#	Requirement	Questions	Answered	Successfully Answered	URL
1	JBNET-2391	Is this something we really want to add to the CSP or simply add to existing functionality within Issue-Tracker?	0	0	https://issues.jboss.org/browse/JBNET-2391
2	JBNET-2376	mike, Got a couple of questions from reading the "Expected tasks". Please read inline: > - create script to set all groups to default "false" preference Should it be setting "Force private comments when escalated form another group" to "true" ? > - create script to set desired JBoss and MetaMatrix groups to "true" preference Should it be setting "Force private comments when escalated form another group" to "false" ? Lisa	1	1	https://issues.jboss.org/browse/JBNET-2376
3	JBNET-2375	Mike, Who in IT/HC I can work with? Thanks, Lisa	1	1	https://issues.jboss.org/browse/JBNET-2375
4	JBNET-2372	Frank, attached, portal_part1.csv are tables (with columns shown) not used/referenced by I-T. Could you please go through them and see if any of the columns are deprecated? Can the deprecated columns be dropped for this iteration? Thanks, Lisa	1	1	https://issues.jboss.org/browse/JBNET-2372
5	JBNET-2372	Frank, One more favor from you please. Can you review the attached, portal_part2.ods file and see if the csp-team references any of the columns that are colored in Red? The colored columns are no longer used by I-T any more. Hopefully dropping them off the schema won't break CSP... thanks, Lisa	1	1	https://issues.jboss.org/browse/JBNET-2372
6	JBNET-2292	The translation fix of the login.html/logout.html, could please do a replacement: s/Please complete a // and this will be the fix. IWO, only the "Please complete a " words are unnecessary. Please note there is a space after "a". Both the pages seems to be non-static page because I can see sessionid are embedded in the pages and I guess we have a template but I'm not sure. Will upload the translated csp help page later.	1	1	https://issues.jboss.org/browse/JBNET-2292
7	JBNET-2273	Jessica, is this a request for SEP or the CSP?	0	0	https://issues.jboss.org/browse/JBNET-2273
8	JBNET-2240	Is there another jira for the logic used to decide who receives each notification?	0	0	https://issues.jboss.org/browse/JBNET-2240
9	JBNET-2123	Lisa, did you just complete the analysis or did we commit these changes at some point?	0	0	https://issues.jboss.org/browse/JBNET-2123
10	JBNET-2099	Does you new table include the label you're using for the drop-down? i.e. "Customer Service Asia/Pacific" label <= "Customer Service APAC" product	1	1	https://issues.jboss.org/browse/JBNET-2099
11	JBNET-2073	results.issues.php line 236 change from if (is_array($opened)) { $search_statement .= "AND ("; $first = TRUE; foreach ($opened as $val) { $search_statement .= $first ? "" : "OR "; $search_statement .= "t.opened_by = '$val' "; $first = FALSE; } To f (is_array($opened)) { $search_statement .= "AND t.opened_by = u.userid AND ("; $first = TRUE; foreach ($opened as $val) { $search_statement .= $first ? "" : "OR "; $search_statement .= "u.username like '$val%' "; $first = FALSE; }	0	0	https://issues.jboss.org/browse/JBNET-2073
12	JBNET-2072	results.issues.php line 253 from if (is_array($assigned)) { $search_statement .= "AND ("; $first = TRUE; foreach ($assigned as $val) { $search_statement .= $first ? "" : "OR "; $search_statement .= "g.assigned_to = '$val' "; $first = FALSE; } To if (is_array($assigned)) { $search_statement .= "AND g.assigned_to = u.userid AND ("; $first = TRUE; foreach ($assigned as $val) { $search_statement .= $first ? "" : "OR "; $search_statement .= "u.username like '$val%' "; $first = FALSE; }	1	1	https://issues.jboss.org/browse/JBNET-2072
13	JBNET-2013	this is for 1.5, not network.jboss.com, as per Mike. -f-	1	1	https://issues.jboss.org/browse/JBNET-2013
14	JBNET-1996	Test Plan: 1. Log in as fmerendagit 2. Go to View Cases 3. Go to case # 127267 4. Go to Attachments 5. Select signature__copy.txt and get the following exception: java.io.FileNotFoundException: /var/files/issues/1/127267/1185224410562-signature__copy_.txt (No such file or directory) java.io.FileInputStream.open(Native Method) java.io.FileInputStream.<init>(FileInputStream.java:106) com.jboss.jbossnetwork.util.FileUtil.getBytesFromFile(FileUtil.java:13) com.jboss.jbossnetwork.ui.actions.ShowAttachmentAction.getStreamInfo(ShowAttachmentAction.java:83) org.apache.struts.actions.DownloadAction.execute(DownloadAction.java:106) org.apache.struts.action.RequestProcessor.processActionPerform(RequestProcessor.java:419) org.apache.struts.action.RequestProcessor.process(RequestProcessor.java:224) org.apache.struts.action.ActionServlet.process(ActionServlet.java:1192) org.apache.struts.action.ActionServlet.doGet(ActionServlet.java:412) javax.servlet.http.HttpServlet.service(HttpServlet.java:697) javax.servlet.http.HttpServlet.service(HttpServlet.java:810) com.jboss.jbossnetwork.util.security.SecurityFilter.doFilter(SecurityFilter.java:125) org.apache.struts2.dispatcher.FilterDispatcher.doFilter(FilterDispatcher.java:414) com.jboss.jbossnetwork.util.EncodingFilter.doFilter(EncodingFilter.java:27) org.jboss.web.tomcat.filters.ReplyHeaderFilter.doFilter(ReplyHeaderFilter.java:96)	0	0	https://issues.jboss.org/browse/JBNET-1996
15	JBNET-1990	The skus were added as part of the work for JBNET-1878. Should this issue remain open until the skus are added in prod?	0	0	https://issues.jboss.org/browse/JBNET-1990
16	JBNET-1945	frank, is this one regarding aliases like matamatrix-support@redhat.com? I discovered a unforseen bug with this one. If we are to use those aliases instead of tao@ for internal users, it breaks the email interface. email udpates bounce if internal users try to update tickets that are opened by customers through those aliases... Things still work if internal users email tao@redhat.com Let me know where we are heading regarding this complication. Lisa	1	1	https://issues.jboss.org/browse/JBNET-1945
17	JBNET-1941	Please ensure that the authentication type can be set by configuration. Should have a constant defined in configuration called "_AUTH_TYPE_" _AUTH_TYPE_ = 1 :: use native authentication, ignore UserService. _AUTH_TYPE_ = 2 :: use new authentication alghorithm (UserService, then native)	0	0	https://issues.jboss.org/browse/JBNET-1941

#	Requirement	Questions	Answered	Successfully Answered	URL
18	JBNET-1932	$where_clause) = $tmpl->tid_gid_filters(true); if ($tmpl->allow_unread_events_sort) { $sql = <<<EOL SELECT t.*, COALESCE(u.eid_count, 0) AS eid_read, t.eid_count - COALESCE(u.eid_count, 0) AS eid_unread FROM $view_name_t LEFT OUTER JOIN $view_name_u USING(tid) WHERE g_gid=gid AND u.uid=${USER['uid']} AND s_uid=${USER['uid']} $where_clause ORDER BY sid != $t_status_reg EOL; } else { if ($tmpl->precache_unread) $tmpl->events_unread_precache($USER['uid']); $sql = <<<EOL SELECT * FROM $view_name_t WHERE g_gid=gid AND s_uid=${USER['uid']}$where_clause ORDER BY sid != $t_status_reg EOL; } $sql .= $tmpl->tid_gid_order_by(); $result = $dbi->query($sql); if ($tmpl->precache_escalate) ticket_escalated_subscribed_to_precache($USER['uid']); $tmpl->tid_gid_output_all("My Subscribed to Issues", $links); Suggested enhancement: list($view_name_t, $view_name_u, $where_clause) = $tmpl->tid_gid_filters_subscribed(true); if ($tmpl->allow_unread_events_sort) { $sql = <<<EOL SELECT t.*, COALESCE(u.eid_count, 0) AS eid_read, t.eid_count - COALESCE(u.eid_count, 0) AS eid_unread FROM $view_name_t LEFT OUTER JOIN $view_name_u USING(tid) WHERE g_gid=gid AND u.uid={$_SESSION['uid']} AND s_uid={$_SESSION['uid']} $where_clause ORDER BY sid != $t_status_reg EOL; } else { if ($tmpl->precache_unread) $tmpl->events_unread_precache($_SESSION['uid']); $sql = <<<EOL SELECT * FROM $view_name_t WHERE g_gid=gid AND s_uid={$_SESSION['uid']}$where_clause ORDER BY sid != $t_status_reg EOL; $sql = $view_name_t." where s.userid=".$_SESSION['uid']." $where_clause ORDER BY sid != $t_status_reg"; } $sql .= $tmpl->tid_gid_order_by(); $result = $dbi->query($sql); if ($tmpl->precache_escalate) ticket_escalated_subscribed_to_precache($_SESSION['uid']); $tmpl->tid_gid_output_all_subscribed("My Subscribed to Issues", $links); portal/modules/issues/tmpl-grouped-issues-pages.php start line 106: Current code: Suggested enhancement: var $sub_tid_pri_all = ''; var $sub_tid_pri_open = ''; var $sub_tid_pub_all = ''; var $sub_tid_pub_open = ''; var $sub_tid_slow_pri_all = ''; var $sub_tid_slow_pri_open = ''; var $tid_gid_pri_all = ''; var $tid_gid_pri_open = ''; var $tid_gid_pub_all = ''; var $tid_gid_pub_open = ''; var $tid_gid_slow_pri_open = ''; var $tid_gid_fast_pri_all = ''; var $tid_gid_fast_pri_open = ''; var $tid_gid_fast_pub_all = ''; var $tid_gid_fast_pub_open = ''; var $sub_tid_pri_all_where = ''; var $sub_tid_pri_open_where = ''; var $sub_tid_pub_all_where = ''; var $sub_tid_pub_open_where = ''; var $sub_tid_slow_pri_all_where = ''; var $sub_tid_slow_pri_open_where = ''; var $tid_gid_pri_all_where = ''; var $tid_gid_pri_open_where = ''; var $tid_gid_pub_all_where = ''; var $tid_gid_pub_open_where = ''; var $tid_gid_slow_pri_open_where = ''; var $tid_gid_fast_pri_all_where = ''; var $tid_gid_fast_pri_open_where = ''; var $tid_gid_fast_pub_all_where = ''; var $tid_gid_fast_pub_open_where = ''; portal/modules/issues/tmpl-grouped-issues-pages.php start line 261: Current code: Suggested enhancement: $s_sub_t_p = " select `t`.`tid` AS `tid`, `t`.`gid` AS `gid`, `t`.`gid` AS `g_gid`, `t`.`opened_by` AS `opened_by_uid`, `t`.`assigned_to` AS `assigned_to_uid`, `t`.`assigned_to` AS `g_assigned_to_uid`, `t`.`opened` AS `opened`, `t`.`opened` AS `escalated`, `t`.`first_response` AS `first_response`, `t`.`first_response` AS `g_first_response`, `t`.`modified` AS `modified`, `t`.`summary` AS `summary`, `t`.`status` AS `sid`, `t`.`severity` AS `severity`, `t`.`etarget` AS `etarget`, `t`.`istatus` AS `isid`, `t`.`priority` AS `priority`, `t`.`score` AS `score`, `t`.`rid` AS `rid`, `t`.`private` AS `t_private`, _utf8't' AS `show_issue`, `t`.`customer_case_id` AS `customer_case_id`, `t`.`cid` AS `cid`, `t`.`product` AS `pid`, `s`.`userid` AS `s_uid`, `g`.`name` AS `group_name`, `g`.`name` AS `group_esc_name`, `g`.`member` AS `group_member`, `g`.`member` AS `group_esc_member`, `g`.`geo` AS `group_geo`, `g`.`geo` AS `group_esc_geo`, `o`.`username` AS `opened_by_username`, `st`.`status` AS `status` from tickets as t join subscriptions as s on (`t`.`tid`=`s`.`tid`) join groups as g on (`g`.`gid` = `t`.`gid`) join users as o on (`o`.`userid` = `t`.`opened_by`) join statuses as st on (`st`.`sid` = `t`.`status`) "; $s_w_open = " and (`t`.`status` not in (4,5,9,78))"; $s_w_pub = " and	0	0	https://issues.jboss.org/browse/JBNET-1932
19	JBNET-1932	For login page (Main page). portal/modules/issues/miniview.issues.php start line 62: Current code: $sql = "SELECT DISTINCT(tid),summary,opened,status "; Suggested enhancement: $sql = "SELECT tid,summary,opened,status "; portal/modules/issues/miniview.issues.php start line 125: Current code: $sql = "SELECT DISTINCT(tid),summary,opened,status,gid "; $sql .= "FROM tickets "; $sql .= "WHERE gid='$val' "; $sql .= "AND status NOT IN ("._CLOSECLIENT_.","._CLOSETECH_.","._AUTOCLOSE_.","._CLOSED_.") "; $sql .= "ORDER BY tid DESC LIMIT $num"; // run query if($result = $dbi->query($sql)){ $issues = TRUE; while(list($tid,$text,$open,$sid,$gid) = $dbi->fetch($result)){ $text = stripslashes($text); // form link to issue $link = "$text"; // get the name of group that owns this ticket $name = group_name($gid); $name = "$name"; // ok write row for ticket $html->html("<tr class=\"colordata\" align=\"center\">"); $html->html("<td width=\"5%\">$tid</td>"); $html->html("<td width=\"20%\">$name</td>"); $html->html("<td align=\"left\">$link</td>"); $html->html("<td width=\"10%\">".usr_date($open)."</td>"); $html->html("<td width=\"10%\">".status($sid)."</td>"); $html->html("</tr>"); } } Suggested enhancement: $sql = "SELECT tid,summary,opened,status,gid "; $sql .= "FROM tickets "; $sql .= "WHERE gid='$val' "; $sql .= "AND status NOT IN ("._CLOSECLIENT_.","._CLOSETECH_.","._AUTOCLOSE_.","._CLOSED_.") "; $sql .= "ORDER BY tid DESC LIMIT $num"; // run query if($result = $dbi->query($sql)) { if($dbi->num_rows($result)>0) { $issues = TRUE; while(list($tid,$text,$open,$sid,$gid) = $dbi->fetch($result)){ $text = stripslashes($text); // form link to issue $link = "$text"; // get the name of group that owns this ticket $name = group_name($gid); $name = "$name"; // ok write row for ticket $html->html("<tr class=\"colordata\" align=\"center\">"); $html->html("<td width=\"5%\">$tid</td>"); $html->html("<td width=\"20%\">$name</td>"); $html->html("<td align=\"left\">$link</td>"); $html->html("<td width=\"10%\">".usr_date($open)."</td>"); $html->html("<td width=\"10%\">".status($sid)."</td>"); $html->html("</tr>"); } } } portal/modules/issues/miniview.issues.php start line 188: Current code: $sql = "SELECT DISTINCT(tid),gid,summary,modified,status "; $sql .= "FROM tickets "; Suggested enhancement: $sql = "SELECT tid,gid,summary,modified,status "; $sql .= "FROM tickets ";	0	0	https://issues.jboss.org/browse/JBNET-1932
20	JBNET-1931	Lisa, was this ever implemented or was this pulled out when we replaced HEAD?	1	0	https://issues.jboss.org/browse/JBNET-1931
21	JBNET-1928	Lisa, was this ever implemented or was this pulled out when we replaced HEAD?	1	0	https://issues.jboss.org/browse/JBNET-1928
22	JBNET-1927	What are you guys using as a driver for your load testing?	0	0	https://issues.jboss.org/browse/JBNET-1927
23	JBNET-1916	This issue has the same situation as JBNET-1915. The label on the list view table is generated within three steps. 1. From $_REQUEST variable. If you altered the view only for this session, the changes will not be stored in database. 2. From database. If you altered the view for not only this session, the changes will be stored in database. The list view on Groups page works this way. 3. From hard codes. If there is no label information available from $_REQUEST and database, it will read from the hard codes. The list views on My Opened, My Assigned, My Subscribed and My Recent pages are using this way. For the third step, I can change it without problem. For the second step, it is not easy to update all the data in database. But it can be easily changed on list view page by replace the string when it needs to be displayed.	1	1	https://issues.jboss.org/browse/JBNET-1916
24	JBNET-1916	These changes only affect the list views on My Opened, My Assigned, My Subscribed and My Recent pages. But not Group page.	1	1	https://issues.jboss.org/browse/JBNET-1916

#	Requirement	Questions	Answered	Successfully Answered	URL
25	JBNET-1915	(from Marvin via email) Hi, Mike: I investigated this issue and JBNET-1916 throughly. I found that the label on the list view table is generated within three steps. 1. From $_REQUEST variable. If you altered the view only for this session, the changes will not be stored in database. 2. From database. If you altered the view for not only this session, the changes will be stored in database. The list view on Groups page works this way. 3. From hard codes. If there is no label information available from $_REQUEST and database, it will read from the hard codes. The list views on My Opened, My Assigned, My Subscribed and My Recent pages are using this way. For the third step, I can change it without problem. For the second step, it is not easy to update all the data in database. But it can be easily changed on list view page by replace the string when it needs to be displayed. Marvin	1	1	https://issues.jboss.org/browse/JBNET-1915
26	JBNET-1896	Don't we already do this? Or is this for Issue-Tracker?	0	0	https://issues.jboss.org/browse/JBNET-1896
27	JBNET-1815	set up group accounts (all noted on the Intranet URL below) where we encourage RH employees to set up their access to our systems and products we offer. This ensures only current employees have access to the free subscriptions (admins deactivate users when they leave the company), and allows us to enter 10 or 20 orders total to give RH employees free subscriptions as opposed to potentially thousands a year. The sku noted below is set up to open up access to all Red Hat/JBoss/RHX products available, and is currently updated by Matt Webbink when new products are set up. Anyone who has asked for an Employee login on one of the RH Employee accounts noted on the Intranet page below will have access to these (all Red Hat) products after they are set up by the admin on the account. I'll leave it up to you guys to decide what's best, but as a fair number of JBoss employees have logins in the Red Hat systems, we thought it might be best to send a note so we can spend less time migrating the free access for internal folks to ensure they are set up the right way in the Red Hat Systems. These users are unique in the respect they will not renew like a customer will, so their account issues won't likely "go away" at renewal time. It almost makes sense to just have these folks request access to the Red Hat accounts set up now rather than migrating. We have tried our best to match current employees to existing Red Hat Employee accounts (thanks Edith) as part of the migration but it is likely there will be others who need/want access, and we wanted to be sure there was a clear way how to do this. That being said, you are the engineer, so please do what will be easiest for your teams during the migration! Jessica and Edith are available anytime if you need additional assistance from Customer Service. Regards, Deborah :) Frank Merenda wrote: > Hi Deborah! Long time no hear! > > We do not recognize this SKU. I just searched email, and I don't see this SKU anywhere. Not an issue, see below. > > I sent Jessica a list of red hat and jboss employees on Friday who currently have access as customers, which she is updating for our data migration with the correct account numbers. I am *assuming* that these accounts will have the magical SKU enabled for it, and all our employees will have access via that mechanism, is that correct? > > Which products does this sky allow access to? Should I assume it's all products? > > Cris, I added this to Jira: http://jira.jboss.com/jira/browse/JBNET-1815 > > -f- > > > Deborah Curtis wrote: >> Hi Chris, Frank! >> >> Jessica was kind enough to draft this email for us to notify former JBoss folks with access to the JBN that they will need to get access on an existing Red Hat Employee Subscription account to access JBoss subscriptions through the Red Hat CSP. Please make sure the CSP is accurately recognizing the Employee Subscription skus (SYS/SVC 0395 Red Hat Employee Subscription), and assist us in getting this communication to the right folks so we can make sure access for JBoss products/support for anyone who needs it is in place before the migration go-live date. Thank you! >> >> Regards, >> Deborah >> >> >>> >>> >>> Hello, >>> >>> You are receiving this email because you are listed as one of the active employees with current access to JBoss subscriptions on the CSP. The JBoss and Red Hat systems are merging to consolidate all accounts into one database. Beginning October 16th, you will be able to access JBoss employee subscriptions using your Red Hat website/RHN login _only_, rather than your current login for https://network.jboss.com . If you wish to continue using your JBoss employee subscription, please verify that you have access to a Red Hat Employee Subscription, which will grant you access to JBoss products. Please check with your manager or visit the following link for more information on accessing employee subscriptions: >>> >>> http://intranet.corp.redhat.com/ic/intranet/wiki.html?wikicat=Customer%20Service%3aEmployee%20Subscriptions&master=00138&toggle=00160 >>> >>> Please email global-customerservice@redhat.com , (INTERNAL ONLY alias), with any questions regarding this process. >>> >>> Have a great day, >>> >>> Jessica >>> Customer Service >>> Red Hat, Inc. >> > -- Deborah Stumpf Curtis Manager, Customer Service Red Hat, Inc. 1801 Varsity Drive Raleigh, NC 27606 888-733-4281, Ext 44225	1	1	https://issues.jboss.org/browse/JBNET-1815
28	JBNET-1812	...this is 1812...did you mean JBNET-1813?	0	0	https://issues.jboss.org/browse/JBNET-1812
29	JBNET-1782	Frank, Do we want to hand-fix those duplicates after they go alive in PROD? Or having the migration script check if an user email exists in portal or not before entering the insertion to the DB? Lisa	1	1	https://issues.jboss.org/browse/JBNET-1782
30	JBNET-1759	To test out permissions for attachments: (11:04:00 AM) ashcrow: you should delete the jboss user and jboss group (11:04:03 AM) ashcrow: leave apache (11:04:11 AM) ashcrow: the rpm only creates jboss:jboss (11:04:16 AM) ashcrow: and puts jboss user in apache group (11:05:16 AM) fjm: so here's how: as root: (11:06:05 AM) fjm: userdel -r jboss && groupdel jboss	0	0	https://issues.jboss.org/browse/JBNET-1759
31	JBNET-1716	Frank, For this preferred_timezone field, what's the default value? NULL or default to "America/New York" as the display Timezone in Issue-Tracker? Lisa	1	1	https://issues.jboss.org/browse/JBNET-1716
32	JBNET-1619	Frank, what's the difference between an org_id and a canonical_id? Lisa	1	1	https://issues.jboss.org/browse/JBNET-1619
33	JBNET-1482	Can we call these URLs final? We need to update our configs to reflect the > final url's. > > https://www.redhat.com/support/esp/APP_HERE > > and it would be: > > ...esp/csp/ ---> jboss:8080/jbossnetwork > ...esp/gatekeeper/ --> :8080/gatekeeper > ...esp/survey/ --> :8080/survey > ...esp/sep/ --> :8080/support-portal > .../esp/downloads --> :8080/slide	0	0	https://issues.jboss.org/browse/JBNET-1482

#	Requirement	Questions	Answered	Successfully Answered	URL
34	JBNET-1481	<gdouglas@redhat.com> Subject: Re: downtime announcement (was [Fwd: Re: ESP 1.3 Status Follow-up, 5/29]) I thought this was going on the login page itself, not on the landing page after login? more below, this all has changed this week because of our two releases, so please bear with me..... Mike Amburn wrote: > Can one of you provide me the data points: > > - what exactly is happening? (i.e. going down, going slow, etc) for 1.3: - Nothing. no users yet. Deploying new web app to production for Command Center users only, but none currently exist. We don't need any changes at this time. I am going to move the jira task to 1.3.1. sorry about this confusion. Gail, Marty needs to know that the front page of the site is final for 1.3. Again, sorry about this confusion. We will send him the 1.3.1 site for approval when it is completed. for 1.3.1: - JBoss users will experience the whole new portal, from the old one of SF to the new one. Support will be down for all *legacy jboss* users for most of the day that saturday when we deploy. network.jboss.com will redirect to our new URLs. > - when is it happening? 1.3: June 9th, all day. no current users. nothing needs to change 1.3.1: TBD. depends on data migration issues. > - when will everything return to normal? as soon as the site comes back up that day. Probably 5 hours for data migration for 1.3.1. > - what other information do you want to communicate? The changes will be two fold: - Command Center users ---- will not notice anything except the outage for the 1.3.1 deployment. -JBoss Users (1.3.1 release - not 1.3) --- will be coming to the new site for the first time. They will see all the new changes for the first time. We should put up a blurb about the new stuff, I listed some of it in the jira task. Sorry about the confusion between these releases... -f- > > > > On May 30, 2007, at 5:55 PM, Frank Merenda wrote: > >> >> Gail Douglas wrote: >>> Mike, >>> I would like for you to read the email string below... it has info >>> for you in two areas: >>> 1) final feedback on UI text changes, and after that change is >>> made we'll have the GSS approval John Arnzen was waiting on >>> 2) if you already have that nifty text that will appear upon >>> logging into CSP, and you can share that with Marty to get his >>> approval, then we can put that issue to rest as well. If you don't >>> have the text yet, we'll have to take another approach to getting you >>> and Marty in sync. And if you don't have that proposed text yet, let >>> me know and I can get people together to decide quickly. >> >> Mike doesn't have this yet. He's getting knowledge transfer from John >> today in ATL. >> >> Mike, can you throw together some new text for the login page. You >> don't have to worry about i18n for this new text for the 1.3 release. >> I have put this in jira for you, >> http://jira.jboss.com/jira/browse/JBNET-1481. >> >> We could get Marty to do this, but it would take a while, so we would >> prefer if you can whip it up. Thanks! >> >> >>> Please let me know your thoughts on these. If you want to chat, I >>> should be in the office and on IM pretty early tomorrow. >>> Thanks! >>> Gail >>> -------- Original Message -------- >>> Subject: Re: ESP 1.3 Status Follow-up, 5/29 >>> Date: Wed, 30 May 2007 14:31:28 -0400 >>> From: Martin C. Messer <marty@redhat.com> >>> Organization: Red Hat, Inc. >>> To: gdouglas@redhat.com >>> CC: Frank Merenda <fmerenda@redhat.com>, Iain Gray >>> <igray@redhat.com>, Lee Mewshaw <lmewshaw@redhat.com> >>> References: <465C521A.5040201@redhat.com> >>> <465DBB35.70408@redhat.com> <465DC02A.3020101@redhat.com> >>> Gail Douglas wrote: >>>> Hi Marty, >>>> >>>> Thanks so much for your quick response to my email below. Just >>>> wanted to >>>> get you feedback on a couple of items... >>>> >>>> For the externally affected users during the outage...because CSP, >>>> Issue >>>> Tracker, their shared db, and Gatekeeper will be down during the >>>> outage, >>>> I assumed we should send external notification of the outage to CSP >>>> and >>>> Issue Tracker users. Also, we will be pushing all of the bug fixes in >>>> Issue Tracker on the 9th, so users may notice a change as a result of >>>> those. I am looking for a list of external users that we can send only >>>> outage notification to...we are looking into creating a short blurb >>>> after CSP login that will notify users that the software has been >>>> updated,	1	1	https://issues.jboss.org/browse/JBNET-1481
35	JBNET-1390	Ahmad, Need some details for this task please: regarding browse the FTP dir part, which FTP service are we talking about here? Lisa	1	1	https://issues.jboss.org/browse/JBNET-1390
36	JBNET-1207	Use this to show how many are RH logins vs how many all unique logins query above Red Hat user count select count(distinct(log_user)) FROM group_users,logs where log_time BETWEEN extract('epoch' FROM '2007-02-07'::date) AND extract('epoch' FROM '2007-02-14'::date) AND log_type='logins' AND group_users.userid=logs.log_user AND group_users.gid = 72;	0	0	https://issues.jboss.org/browse/JBNET-1207
37	JBNET-1198	I'm not sure I understand the problem. If I'm logged in as mamburn and create an issue, then I should show as the owner of the ticket. How can you not know the owner?	1	0	https://issues.jboss.org/browse/JBNET-1198
38	JBNET-1198	Mike, If RH employee X opens an Issue for customer, then yes he is the owner of the ticket. But, if X then needs something from the customer, this will have to be done by person Y from customer staff. So, if there was a customer owner field in addition to the existing RH owner field, then tech Y on the customer side could not hide behind his organisation and claim "oh, I did not know this was for me to do" (happens rather frequently. That is why I made the request, I want to see clearly ownership on the customer side in addition to the ownership we see on the RH side. Does this make the request clearer?	1	1	https://issues.jboss.org/browse/JBNET-1198
39	JBNET-1198	I still don't quite understand. Suppose Mike is a Red Hat employee and Joe is a customer from Company XYZ. Is Mike opening the issue for Joe in the Company XYZ Issue-Tracker group and not under Joe (so that the ticket has no direct customer attachment)? Is that the problem?	1	0	https://issues.jboss.org/browse/JBNET-1198
40	JBNET-954	Is this ticket done? I thought these were created.	1	1	https://issues.jboss.org/browse/JBNET-954
41	JBNET-907	Hua Li began testing using one of our stock users. Here is the feedback I've received thus far: There are some mistakes that I can't find file to modify. 1. the title of http://10.10.29.31:8080/jbossnetwork/restricted/main.html now the title of main.html is "JBoss??????????" , should be modified to "JBoss??????" . 2. the title of http://10.10.29.31:8080/jbossnetwork/help/zh/help-main.htm can't be correctly shown. Please check it. title of help-main.htm 3. please access http://10.10.29.31:8080/jbossnetwork/restricted/knowledgeHome.html I don't know what is "????" ? what is "????" ---- Addressing what I can for these now. Will work with Hua for items needing more translation and clarification.	0	0	https://issues.jboss.org/browse/JBNET-907
42	JBNET-891	Notes on that meeting with Lisa: 1. (Lisa) Create event_types table with the following event types: 0 - Unknown 1 - State Change 2 - Comment 3 - Attachment 4 - Task (for recording time spent on a ticket) 5 - Notification 2. (Lisa) Create event_type field of type integer in the event table 3. (Ahmad) Update CSP schema and mappings 4. (Ahmad) Change all queries for CSP and GK 5. (Ahmad) Change case migration 6. (Lisa) Change IT to store events in correct types 7. (Lisa) Change Notifications to pull from events that are type 5 8. (Lisa) Change migration so that all events are type 5 9. (Lisa & Ahmad) Update dev and wendyc-t	0	0	https://issues.jboss.org/browse/JBNET-891
43	JBNET-873	Lamon, Can you add Active Entitlements to the IssueTracker Batch Process and start testing this? IM me for details. Ahmad	1	1	https://issues.jboss.org/browse/JBNET-873
44	JBNET-831	<llu> what's the possible values for 'Type'? <fmerenda> let me get them one sec <fmerenda> "Question", "Problem" and "Feature Request" <llu> Problem == bug? <fmerenda> could be a problem in - for example - their production environment, and they need help, may be that it is not a bug <fmerenda> one sec looking up details <fmerenda> llu, when we add cases we ususally set to "Problem", unless it is dev assist, or question when they submit the case, then it's a "Question". <fmerenda> your UI just needs to be able to change the type <fmerenda> We only use it for information, not for queueing or anything like that	1	1	https://issues.jboss.org/browse/JBNET-831
45	JBNET-830	Question: How do we want to define and map the filter that will UNBIND all customer URLs when they are accessed? Do we care (i.e. something other than /files/customer/*)?	1	1	https://issues.jboss.org/browse/JBNET-830
46	JBNET-755	frank, what's the difference between software and products? what does repository do? The page for creating/editing products is avai. today in Issue Tracker production. Lisa	1	1	https://issues.jboss.org/browse/JBNET-755
47	JBNET-755	Frank, how high priority is this task? Need it to be done before the integration? Lisa	1	1	https://issues.jboss.org/browse/JBNET-755

#	Requirement	Questions	Answered	Successfully Answered	URL
48	JBNET-309	Question on this: What is happening when you try to update/create an object? Not compatible? Is this a salesforce-connector issue or an actual Axis issue? Has anyone tried to work with the SForce API directly using Axis 1.3?	1	1	https://issues.jboss.org/browse/JBNET-309
49	ENTMQBR-667	This is needed to get Artemis working with Fuse 7. Can you get it into one of the next AMQ patches?	0	0	https://issues.jboss.org/browse/ENTMQBR-667
50	ENTMQBR-644	[~jbertram] thanks for the info, my bad that I considered SSLv2Hello to be SSLv2 use. I need to gain more knowledge of that. I'll have a closer look at your comments after the A-MQ 7.0 GA is out of the door. This was proposed to 7.0.1 and we are currently very busy now. Just one note right now: \\ when I run a client twice, ie: {code} # openssl s_client -connect $(hostname):5671 -cert /opt/A-MQ-certificates/client_cert.pem -key /opt/A-MQ-certificates/client_private_key.pem -pass pass:changeit -tls1_2 # openssl s_client -connect $(hostname):5671 -cert /opt/A-MQ-certificates/client_cert.pem -key /opt/A-MQ-certificates/client_private_key.pem -pass pass:changeit -tls1_2 {code} Broker log contains two INFO messages regarding disallowing SSLv2Hello: {code} 23:32:41,507 INFO [org.apache.activemq.artemis.core.server] AMQ221053: Disallowing use of vulnerable protocol 'SSLv2Hello' on acceptor 'netty-ssl-acceptor'. See http://www.oracle.com/technetwork/topics/security/poodlecve-2014-3566-2339408.html for more details. 23:32:44,517 INFO [org.apache.activemq.artemis.core.server] AMQ221053: Disallowing use of vulnerable protocol 'SSLv2Hello' on acceptor 'netty-ssl-acceptor'. See http://www.oracle.com/technetwork/topics/security/poodlecve-2014-3566-2339408.html for more details. {code} This seems to me in contrast with what you are saying (logging is done when the acceptor is started, not on every connection). Am I missing something? Thanks in advance for explanation.	1	0	https://issues.jboss.org/browse/ENTMQBR-644
51	ENTMQBR-539	I understand why you may want to configure a max size on a per destination basis, however there is currently no scenario where 1 queue paging will stop another. Queues are never starved no matter what state paging is in across the broker. Could you provide more information?	0	0	https://issues.jboss.org/browse/ENTMQBR-539
52	ENTMQBR-526	Hi Karl, What kind of clients are the user having in their test (OpenWire or others) and do you have a test case? Howard	1	1	https://issues.jboss.org/browse/ENTMQBR-526
53	ENTMQBR-526	Ah. I noticed the following in from the customer in the open ticket #01782533: We have clients both using JNDI: Sample property files content: java.naming.factory.initial=org.apache.activemq.artemis.jndi.ActiveMQInitialContextFactory connectionFactory.ConnectionFactory=(tcp://artemis-1.test.iggroup.local:61616,tcp://artemis-2.test.iggroup.local:61616,tcp://artemis-3.test.iggroup.local:61616,tcp://artemis-1.test.iggroup.local:61716,tcp://artemis-2.test.iggroup.local:61716,tcp://artemis-3.test.iggroup.local:61716)?clientID=1234 And code-ally String connectionUrl = "(tcp://artemis-1.test.iggroup.local:61616,tcp://artemis-2.test.iggroup.local:61616,tcp://artemis-3.test.iggroup.local:61616,tcp://artemis-1.test.iggroup.local:61716,tcp://artemis-2.test.iggroup.local:61716,tcp://artemis-3.test.iggroup.local:61716)"; URIFactory parser = new URIFactory(); ActiveMQJMSConnectionFactory cf = new ActiveMQJMSConnectionFactory(parser.expandURI(connectionUrl).toString());	1	1	https://issues.jboss.org/browse/ENTMQBR-526
54	ENTMQBR-512	If [~joef551] will not manage to test this, will you manage this sprint [~dlenoch]?	1	1	https://issues.jboss.org/browse/ENTMQBR-512
55	ENTMQBR-496	[~opiske] what do you mean by purge?	1	0	https://issues.jboss.org/browse/ENTMQBR-496
56	ENTMQBR-496	[~opiske] what would you suggest?	1	0	https://issues.jboss.org/browse/ENTMQBR-496
57	ENTMQBR-486	shouldnt we wait for a QPID-JMS JMS 2.0 to be released before we start raising Jiras?	1	1	https://issues.jboss.org/browse/ENTMQBR-486
58	ENTMQBR-486	Bumping this one before PM call. Any idea what to do with this [~ataylor]? Are we ok for not 100% support JMS 2.0? Release notes for Beta2?	1	1	https://issues.jboss.org/browse/ENTMQBR-486
59	ENTMQBR-422	As a follow-up, I've updated the QE results parser to skip all the issues identified here and in the parent issue ENTMQBR-417. There are still few failing tests that were not investigated yet and thus _may_ indicate a real broker issue: {code} 1) test: N8activemq4test8openwire27OpenwireTempDestinationTestE::testPublishFailsForClosedConnection (F) line: 369 activemq/test/openwire/OpenwireTempDestinationTest.cpp expected exception not thrown - Should throw a InvalidDestinationException since temp destination should not exist anymore. - Expected: cms::InvalidDestinationException 2) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testXAResource_Exception1 (E) uncaught exception of type N3cms11XAExceptionE - 3) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testXAResource_Exception2 (E) uncaught exception of type N3cms11XAExceptionE - 4) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testXAResource_Exception3 (E) uncaught exception of type N3cms11XAExceptionE 5) test: N8activemq4test8openwire18OpenwireSimpleTestE::testWithZeroConsumerPrefetchAndZeroRedelivery (E) uncaught exception of type N8activemq10exceptions17ActiveMQExceptionE - AMQ119025: Cannot delete queue jms.queue.testWithZeroConsumerPrefetchAndZeroRedelivery on binding jms.queue.testWithZeroConsumerPrefetchAndZeroRedelivery - it has consumers = org.apache.activemq.artemis.core.postoffice.impl.LocalQueueBinding 6) test: N8activemq4test8openwire20OpenwireAdvisoryTestE::testConcurrentTempDestCreation (F) line: 177 activemq/test/openwire/OpenwireAdvisoryTest.cpp assertion failed - Expression: thread1.isNoErrors() {code} *^^ These are regression between ER2 and ER3 and still persists on ER12 (except testConcurrentTempDestCreation that is passing again on ER12)* ---------- {code} 7) test: N8activemq4test8openwire23OpenwireTransactionTestE::testSendSessionClose (E) uncaught exception of type N8activemq10exceptions17ActiveMQExceptionE - AMQ119017: Queue jms.queue.testSendSessionClose does not exist {code} *^^ These are regression between ER4 and ER6 and still persists on ER12.* ---------- {code} 8) test: N8activemq4test8openwire18OpenwireSimpleTestE::testDestroyDestination (F) line: 306 activemq/test/openwire/OpenwireSimpleTest.cpp assertion failed - Expression: false - CAUGHT EXCEPTION {code} *^^ These are regression between ER7 and ER8 and still persists on ER12.* ---------- following test are failing occasionally: {code} 9) test: N8activemq4test8openwire24OpenwireQueueBrowserTestE::testReceiveBrowseReceive (F) line: 81 activemq/test/QueueBrowserTest.cpp assertion failed - Expression: enumeration->hasMoreMessages() - should have received the second message (OR sometimes: - should have received the third message) {code} Reproducibility ~25% {code} 10) test: N8activemq4test8openwire24OpenwireQueueBrowserTestE::testBrowseReceive (F) line: 137 activemq/test/QueueBrowserTest.cpp assertion failed - Expression: enumeration->hasMoreMessages() - should have received the first message {code} Reproducibility ~25% Also worth mentioning that following tests were fixed from ER2 (were failing on ER2 but are passing on ER12) * N8activemq4test8openwire18OpenwireSimpleTestE::testMultipleConnections * N8activemq4test8openwire18OpenwireSimpleTestE::testMultipleSessions * N8activemq4test8openwire18OpenwireSimpleTestE::testReceiveAlreadyInQueue * N8activemq4test8openwire27OpenwireTempDestinationTestE::testBasics [~ataylor] I expect that you want to treat these as separate issues? I'll create a separate Jira's for them on Monday if not requested otherwise. Thanks in advance for helping me sort the supported CMS tests out.	1	1	https://issues.jboss.org/browse/ENTMQBR-422
60	ENTMQBR-422	Hi Petr, I followed the steps given by you but after I did 'make check', I couldn't find the executable activemq-test-integration. I do see a lot object files generated thou. What could be the reason for the missing executable? Thanks Howard	0	0	https://issues.jboss.org/browse/ENTMQBR-422
61	ENTMQBR-418	[~ataylor] Can you please add more documentation for user drop down box "in the top right hand corner" and items as "Connect, Console Logging, Core, Editor, Git, Jolokia, Plugins, Reset, Theme". Or only send me more information about this options.	1	0	https://issues.jboss.org/browse/ENTMQBR-418

#	Requirement	Questions	Answered	Successfully Answered	URL
62	ENTMQBR-418	Some typos found: In the section *deploying the plugin* [standalone subsection|https://github.com/rh-messaging/artemis-hawtio/blob/master/docs/user-manual.md#standalone] ??To deploy standalone refer to the HawtIO documentation on installing plugins. for e?? - hanging leftover *"for e"* Adjacent section [Logging in to the Artemis console|https://github.com/rh-messaging/artemis-hawtio/blob/master/docs/user-manual.md#logging-in-to-the-artemis-console] ??If you have installed HawtIO standalone and want to connect to a remote broker you will need to connect to it by clicking on the connect tab and entering the Jolikia path?? - typo in *Jolikia* Subsection *##Sending a message to a Queue* ??Sometimes it is desirable to test an installation by sending a test message or to simply just send a message. To do this choose the Queue from the JMX tab that you want to send the message to and click on the send tab. Fill in the message content and any headers and click on the Send Message` button, like so:?? - lost styling at *Send Message`*?	0	0	https://issues.jboss.org/browse/ENTMQBR-418
63	ENTMQBR-417	the results above (CMS 3.9.1 vs. ER1). Failures CMS 3.9.3 vs A-MQ7 ER3: 1) test: N8activemq4test8openwire20OpenwireAdvisoryTestE::testConnectionAdvisories (F) line: 111 activemq/test/openwire/OpenwireAdvisoryTest.cpp 2) test: N8activemq4test8openwire20OpenwireAdvisoryTestE::testConcurrentTempDestCreation (F) line: 177 activemq/test/openwire/OpenwireAdvisoryTest.cpp 3) test: N8activemq4test8openwire30OpenwireEnhancedConnectionTestE::testDestinationSourceGetters (F) line: 178 activemq/test/openwire/OpenwireEnhancedConnectionTest.cpp 4) test: N8activemq4test8openwire22OpenwireMapMessageTestE::testEmptyMapSendReceive (F) line: 60 activemq/test/MapMessageTest.cpp 5) test: N8activemq4test8openwire30OpenwireMessageCompressionTestE::testTextMessageCompression (F) line: 109 activemq/test/MessageCompressionTest.cpp 6) test: N8activemq4test8openwire30OpenwireMessageCompressionTestE::testStreamMessageCompression (F) line: 221 activemq/test/MessageCompressionTest.cpp 7) test: N8activemq4test8openwire37OpenWireMessageListenerRedeliveryTestE::testQueueSessionListenerExceptionDlq (E) 8) test: N8activemq4test8openwire37OpenWireMessageListenerRedeliveryTestE::testTransactedQueueSessionListenerExceptionDlq (E) 9) test: N8activemq4test8openwire28OpenWireRedeliveryPolicyTestE::testDLQHandling (E) 10) test: N8activemq4test8openwire28OpenWireRedeliveryPolicyTestE::testRepeatedRedeliveryReceiveNoCommit (E) 11) test: N8activemq4test8openwire28OpenWireRedeliveryPolicyTestE::testRepeatedRedeliveryOnMessageNoCommit (E) 12) test: N8activemq4test8openwire18OpenwireSimpleTestE::testWithZeroConsumerPrefetchAndZeroRedelivery (E) 13) test: N8activemq4test8openwire18OpenwireSimpleTestE::testWithZeroConsumerPrefetchWithInFlightExpiration (E) 14) test: N8activemq4test8openwire18OpenwireSimpleTestE::tesstStreamMessage (F) line: 354 activemq/test/openwire/OpenwireSimpleTest.cpp 15) test: N8activemq4test8openwire27OpenwireTempDestinationTestE::testTwoConnections (F) line: 223 activemq/test/openwire/OpenwireTempDestinationTest.cpp 16) test: N8activemq4test8openwire27OpenwireTempDestinationTestE::testPublishFailsForClosedConnection (F) line: 369 activemq/test/openwire/OpenwireTempDestinationTest.cpp 17) test: N8activemq4test8openwire24OpenwireVirtualTopicTestE::testVirtualTopicSyncReceiveAutoAck (F) line: 95 activemq/test/VirtualTopicTest.cpp 18) test: N8activemq4test8openwire24OpenwireVirtualTopicTestE::testVirtualTopicSyncReceiveClinetAck (F) line: 95 activemq/test/VirtualTopicTest.cpp 19) test: N8activemq4test8openwire24OpenwireVirtualTopicTestE::testVirtualTopicSyncReceiveTransacted (F) line: 95 activemq/test/VirtualTopicTest.cpp 20) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testXAResource_Exception1 (E) 21) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testXAResource_Exception2 (E) 22) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testXAResource_Exception3 (E) 23) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testSendReceiveTransactedBatches (F) line: 252 activemq/test/openwire/OpenwireXATransactionsTest.cpp 24) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testSendRollback (E) 25) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testWithTTLSet (E) 26) test: N8activemq4test8openwire26OpenwireXATransactionsTestE::testSendRollbackCommitRollback (E)	1	1	https://issues.jboss.org/browse/ENTMQBR-417
64	ENTMQBR-412	Did we get any feedback on what the preferred directory is here. On Ubuntu based systems there's not always an /etc/sysconfig but I think on RHEL based systems the sysconfig tends to be preferred over /etc/default.	0	0	https://issues.jboss.org/browse/ENTMQBR-412
65	ENTMQBR-411	Seems like [~ataylor] moved this one from test to ready to release. Was that a missclick? If not, we could mark it with label "internal" or close it as a duplicate of ENTMQBR-523. What do you think?	0	0	https://issues.jboss.org/browse/ENTMQBR-411
66	ENTMQBR-379	Thanks, Andy. Is the blue cylinder at the top a message store? Where does the aspect of persistence come into this scenario? The steps w/ descriptions of what's happening are good.	1	0	https://issues.jboss.org/browse/ENTMQBR-379
67	ENTMQBR-379	I agree - the audience will understand your visual. Maybe just a short sentence about the msg store? (I'm a pain, I know.) Do you want to share with Ted and Justin before modifying to reflect the logging with persistence using Topics scenario? Let me know if I can help.	1	0	https://issues.jboss.org/browse/ENTMQBR-379
68	ENTMQBR-379	[~sjayatredhat] the logging scenario will be a different document, I'll get to that soon. [~tedross] [~justi9] guys you can review what Ive done if you like, also I'm not sure If we mean the same thing when we talk about the logging with persistence scenario, what were you thinking of?	1	0	https://issues.jboss.org/browse/ENTMQBR-379
69	ENTMQBR-379	[~rhn-support-pfestoso], I'm not certain if I've shared these mock-ups with Amanda yet? Can you please check what she returned for the Intro recently and share these with her if they're not in that batch? Muchos gracias.	0	0	https://issues.jboss.org/browse/ENTMQBR-379
70	ENTMQBR-378	- If user try extending the "org.apache.activemq.pool.PooledConnectionFactory" as attached and then they with the help of following spring configuration it can be achieved. {code} <!-- The ActiveMQ JMS connection pool The actual pool implementation is 'org.apache.activemq.pool.PooledConnectionFactory', but this doesn't expose any details at runtime, so we use a subclass to provide some information. --> <bean id="activemq-jms-cf-tx" class="custom.activemq.pool.MonitorablePooledConnectionFactory" init-method="start" destroy-method="stop"> <property name="maxConnections" value="${amq.pool.size}" /> <property name="connectionFactory" ref="activemq-jms-connection-factory"/> </bean> <!-- Expose ActiveMQ JMS connection pool statistics via JMS --> <bean id="activemq-pool-monitor" factory-bean="activemq-jms-cf-tx" factory-method="createPoolMonitor"/> <bean id="exporter" class="org.springframework.jmx.export.MBeanExporter" lazy-init="false"> <property name="beans"> <map> <entry key="custom.activemq:name=activemq-jms-pool" value-ref="activemq-pool-monitor"/> </map> </property> </bean> {code} - Can the similar functionality be implemented so that customers need not to write custom code.	0	0	https://issues.jboss.org/browse/ENTMQBR-378
71	ENTMQBR-376	[~ataylor] [~ganandan-Redhat] I have stumble across this issue. Is it still valid? If so, for what version do we target? etc... Thank you	1	0	https://issues.jboss.org/browse/ENTMQBR-376

#	Requirement	Questions	Answered	Successfully Answered	URL
72	ENTMQBR-376	The following is taken from the SFDC case (comment #12): {quote} Created By: Jason Meyer (2/26/2015 10:30 AM) Akhmad, Good morning! This enhancement has been approved and is targeted for the next major release, which is A-MQ 7.0. Schedule-wise it is too soon to tell. We can probably expect A-MQ 7.0 to be released around a year after 6.2. Since this has now been approved we can track it via the enhancement request. Do you have any issues with me closing this out now? We will still be tracking the progress through the enhancement request. Let me know if you have any questions! Thanks. -Jason- {quote} Akhmad is no longer with Nielsen - he's moved on to Raymond James Financial - so I really don't know where this stands wrt the customer.	0	0	https://issues.jboss.org/browse/ENTMQBR-376
73	ENTMQBR-366	QE currently do not have any knowledge of the scheduler, could you possibly point us to a documentation for scheduling messages? Thanks in advance. [~ataylor]	1	0	https://issues.jboss.org/browse/ENTMQBR-366
74	ENTMQBR-362	QE is not sure how to verify this issue, could you possibly provide verification scenario? QE may do some performance testing, but this do not really verify the reuse of the context objects. [~dbosanac]	0	0	https://issues.jboss.org/browse/ENTMQBR-362
75	ENTMQBR-362	[~ataylor], we don't currently re-use the context object for the LDAP connection. [~jcordes], have you run a particular benchmark to determine this is a performance bottleneck or is this simply based on a reading of the code?	0	0	https://issues.jboss.org/browse/ENTMQBR-362
76	ENTMQBR-358	Hello Justin, the configuration reloading is documented in [Apache ActiveMQ Artemis User Manual|https://activemq.apache.org/artemis/docs/1.5.2/book.pdf] (page 99), right? Do we have another subscription of this feature, please?	1	1	https://issues.jboss.org/browse/ENTMQBR-358
77	ENTMQBR-358	According to [Apache ActiveMQ Artemis User Manual|http://activemq.apache.org/artemis/docs/1.5.2/config-reload.html] automatic configuration reloading should work for these modules of "Broker.xml" configuration file: * Address Settings * Security Settings * JMS Queues * JMS Topics Configuration reloading works for "Address settings" and "Security Settings" modules, but the new addressing model do not have a separate JMS configuration and changes in "addresses" module are not reloaded. Is this an expected behavior, please? Note: I already created a Jira to update the Artemis User Manual: [Apache ActiveMQ Artemis User Manual|http://activemq.apache.org/artemis/docs/1.5.2/].	1	1	https://issues.jboss.org/browse/ENTMQBR-358
78	ENTMQBR-355	Wrt to the comment above, shouldn't this be really closed as a duplicate then? Otherwise it will still appear in 7.0.0 queries and will be synced to Polarion. [~ataylor]	1	1	https://issues.jboss.org/browse/ENTMQBR-355
79	ENTMQBR-351	[~grgrzybek] Can you comment on this? Is it one we can do for 6.2.1?	1	0	https://issues.jboss.org/browse/ENTMQBR-351
80	ENTMQBR-351	Are you saying that we already show this in HawtIO?	0	0	https://issues.jboss.org/browse/ENTMQBR-351
81	ENTMQBR-351	Is this issue for AMQ6 or AMQ7? According to "Target release" it is for AMQ7, but all the talk is about AMQ6 if I am correct. [~garytully] [~ataylor] could you please make it clear for us? if it is for AMQ6, could you please move this item to ENTMQ project, where it belongs? Thank you :)	0	0	https://issues.jboss.org/browse/ENTMQBR-351
82	ENTMQBR-348	Hello Justin, could you possibly add any description about this enhancement, how is it implemented or how to test it, please? Thank you in advance.	1	1	https://issues.jboss.org/browse/ENTMQBR-348
83	ENTMQBR-348	Justin, thank you for your clarification of this enhancement. I already can understand how it should work (I hope ;)). I am very sorry, but I have other questions: 1. Is this enhancement specific for durable subscribers only or is it relevant also for non durable ones? 2.How shall I configure the security-setting in broker.xml file? Could you possibly add any example, please? 3. Where can I find any documentation of this enhancement (how it is implemented)?	0	0	https://issues.jboss.org/browse/ENTMQBR-348
84	ENTMQBR-341	What should I pay attention to when testing it? What is the sign of self-discovery?	0	0	https://issues.jboss.org/browse/ENTMQBR-341
85	ENTMQBR-331	Approved - will try to schedule for future release	0	0	https://issues.jboss.org/browse/ENTMQBR-331
86	ENTMQBR-331	Assessment effort started. There are pending questions: * Is certification for both GFS and GFS2 needed? * What platforms are we going to include in certification, all (RHEL5 i686/x86_64, RHEL6 i686/x86_64, RHEL7 x86_64)? * What is desired cluster configuration? (2/3/more nodes?) At the moment it is clear that effort will be manual and that's why those question matter.	1	1	https://issues.jboss.org/browse/ENTMQBR-331
87	ENTMQBR-325	The issue retested with latest jboss-a-mq-6.2.0.redhat-114.zip with same results.	0	0	https://issues.jboss.org/browse/ENTMQBR-325
88	ENTMQBR-325	Additional details: When JBoss A-MQ finds there is a problem to bind to any of ports specified as transport connectors then waits until all ports are available. Once all ports are available to bind then JBoss A-MQ continue to load transport connectors. The issue is still there but probably lower severity, feel free to adjust. The key point of this defect is that no transport connector is loaded until *ALL* transport connector ports specified in broker configuration are ready to bind. Broker is not ready to operate. Possibly two ways to resolve this (sorted by QE preference): * incorporate changes that ALL transport connectors which can be loaded and bind to ports are loaded and can be used, the ones which cannot are waiting and finally failing to initiate after a timeout. This is making sure broker is functional even if one port cannot be bound. * document the current behavior	0	0	https://issues.jboss.org/browse/ENTMQBR-325
89	ENTMQBR-314	[~ataylor] We need more info/description. Do you think "Browse message" [https://github.com/rh-messaging/artemis-hawtio/blob/master/docs/user-manual.md#browsing-a-queue-for-messages] ?	0	0	https://issues.jboss.org/browse/ENTMQBR-314
90	ENTMQBR-309	[~ataylor] Hawtio thing, you mean? Browser just renders what the console plugin and the console itself give it. If this won't be fixed downstream, should I report it to upstream hawtio, as Ernie suggested on the original DISPATCH-422?	0	0	https://issues.jboss.org/browse/ENTMQBR-309
91	ENTMQBR-266	QE assumes that this relates to Hawtio. We expect that Hawtio shall be extended with a browser for management messages. Is that correct? Could you possibly update the components accordingly.	1	1	https://issues.jboss.org/browse/ENTMQBR-266
92	ENTMQBR-262	Could you please add more description? Thank you.	1	1	https://issues.jboss.org/browse/ENTMQBR-262

#	Requirement	Questions	Answered	Successfully Answered	URL
93	ENTMQBR-256	+jms-bridge.message-count=The number of messages successfully committed. +jms-bridge.aborted-message-count=The number of messages aborted/rollback. What I observed is that when a message fails to be committed, it gets counted both to messageCount and to abortedMessageCount. Going by the doc strings in that WildFly UI, I'd say it should be counted only to abortedMessageCount. Reproducer output: {noformat} getMessageCount: 0 getAbortedMessageCount: 0 10:43:12,574 WARN AMQ342009: JMS Bridge N/A failed to send + acknowledge batch, closing JMS objects javax.transaction.RollbackException at com.redhat.mqe.bridge.Main$1$1.commit(Main.java:114) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.sendBatchXA(JMSBridgeImpl.java:1382) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.sendBatch(JMSBridgeImpl.java:1285) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.access$1700(JMSBridgeImpl.java:75) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl$SourceReceiver.run(JMSBridgeImpl.java:1682) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) getMessageCount: 1 getAbortedMessageCount: 1 10:43:14,545 WARN AMQ342009: JMS Bridge N/A failed to send + acknowledge batch, closing JMS objects javax.transaction.RollbackException at com.redhat.mqe.bridge.Main$1$1.commit(Main.java:114) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.sendBatchXA(JMSBridgeImpl.java:1382) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.sendBatch(JMSBridgeImpl.java:1285) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.access$1700(JMSBridgeImpl.java:75) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl$SourceReceiver.run(JMSBridgeImpl.java:1682) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) 10:43:16,048 WARN AMQ342009: JMS Bridge N/A failed to send + acknowledge batch, closing JMS objects javax.transaction.RollbackException at com.redhat.mqe.bridge.Main$1$1.commit(Main.java:114) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.sendBatchXA(JMSBridgeImpl.java:1382) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.sendBatch(JMSBridgeImpl.java:1285) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.access$1700(JMSBridgeImpl.java:75) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl$SourceReceiver.run(JMSBridgeImpl.java:1682) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) 10:43:16,551 WARN AMQ342009: JMS Bridge N/A failed to send + acknowledge batch, closing JMS objects javax.transaction.RollbackException at com.redhat.mqe.bridge.Main$1$1.commit(Main.java:114) at org.apache.activemq.artemis.jms.bridge.impl.JMSBridgeImpl.sendBatchXA(JMSBridgeImpl.java:1382) at	1	1	https://issues.jboss.org/browse/ENTMQBR-256
94	ENTMQBR-248	Lami, would you be able to clarify the details here? Is this from Hawtio? There is a move operation exposed through MBean Operation?	0	0	https://issues.jboss.org/browse/ENTMQBR-248
95	ENTMQBR-244	WE believe this should not be tested by us - QE. Could devs verify this? Adding "internal" label.	0	0	https://issues.jboss.org/browse/ENTMQBR-244
96	ENTMQBR-234	[~chirino_hchirino] can you let us know if this need any engineering work	1	1	https://issues.jboss.org/browse/ENTMQBR-234
97	AMQDOC-2425	What should QE focus on this item? Is it relevant for QE? If not, please close this one.	1	1	https://issues.jboss.org/browse/AMQDOC-2425
98	AMQDOC-2408	Hello Susan, could you possibly send me a link to the document, please?	1	1	https://issues.jboss.org/browse/AMQDOC-2408
99	AMQDOC-2408	Hi [~psvobodo], Most certainly. :-) Because the Portal tools/style sheet specifies all caps for headings in HTML and PDF, it's hard to find "where" one can verify that the capitalization is correct (other than looking at the source)! If you go here: https://access.redhat.com/documentation/en-us/red_hat_jboss_amq/7.0/html/using_amq_broker/ You can browse all of the headings in the grey panel on the left. Thanks for checking it out! Susan	1	1	https://issues.jboss.org/browse/AMQDOC-2408
100	AMQDOC-2399	Hello, Phil, could you possibly send me a link to the document, please?	1	1	https://issues.jboss.org/browse/AMQDOC-2399
101	AMQDOC-2398	After looking through the content in chapters 8.1 and 8.2, this is a first draft of a 'toc' that may suit the new chapter. It seems pretty basic, but it maps to the content currently available when viewed through the lens of 'Users and Roles'. [~nicole.engard] WDYT? = Users and Roles == Overview == Guest Authentication === `login.conf` == User Authentication === `users.properties` == Roles === `roles.properties`	1	1	https://issues.jboss.org/browse/AMQDOC-2398
102	AMQDOC-2398	A few comments: - "Preamble" should not have its own section heading - http://file.rdu.redhat.com/~jross/amq7-documentation-contrib/internal-resources/information-architecture.html#general_guidelines - last bullet - Users wouldn't say "adding a guest" because it's really for a class of users - I recommend "Enabling Guest Access" or "Adding a Guest Account" - The section on configuring the roles property file location - is that necessary? If there is a default, users should stick to that in general. Even if we do retain that part, I don't think it should be primary.	1	1	https://issues.jboss.org/browse/AMQDOC-2398
103	AMQDOC-2398	Thanks for the feedback [~justi9] I've actioned your first two points, but I'm a little unsure about the last one re; the roles.properties file location? While I mention that the 'roles.properties' does not necessarily have to be called that (which I confirmed with an SME at the time of writing the original security content), I don't think that I imply users should do this. On the contrary, a comment left on my initial push of this content called negative attention to how subtly I'd mentioned the option to use a custom filename. Can you elaborate on your take on this? You feel we shouldn't mention this option?	1	1	https://issues.jboss.org/browse/AMQDOC-2398
104	AMQDOC-2384	[~jdanek] can this tool be used with A-MQ 6? Or is it just something for our 7.0 beta and alpha users?	0	0	https://issues.jboss.org/browse/AMQDOC-2384
105	AMQDOC-2224	Where can I check/verify added documentation for this?	1	1	https://issues.jboss.org/browse/AMQDOC-2224
106	AMQDOC-2224	I don't think there is a relevant equivalent part as in EAP doc. In EAP doc, there is expiring of messages from client POV. In broker guide I have found only expiring messages from broker POV in {{16.2.3. Core Queue Management}} Should I consider it as relevant and what this feature jira added?	1	1	https://issues.jboss.org/browse/AMQDOC-2224
107	AMQDOC-2171	[~jdanek], Do you feel that we've worked through the dependencies to ensure that we're tending to ENTESB issues efficiently + reliably? Please mark this Closed or Reopen with other suggestions.	1	0	https://issues.jboss.org/browse/AMQDOC-2171
108	AMQDOC-2171	[~esammons], what do you think? I know that Petr talked to you about whether we are supposed to do QE verification on the ENTESB Jiras on Monday, this is related. It seems to me that the label=artemis may be enough to mark the Jiras for QE purposes as well. The router is using hawt.io too, so we should maybe add the {{dispatch}} label to the filter. We would have to add this to filter so that we don't miss them. Or is it that all issues in Hawtio framework itself get verified by the Fuse QE, so we don't care? I'd like that ;)	0	0	https://issues.jboss.org/browse/AMQDOC-2171
109	AMQDOC-2155	The sudo note should now be included in every A-MQ 7 doc.	1	1	https://issues.jboss.org/browse/AMQDOC-2155

#	Requirement	Questions	Answered	Successfully Answered	URL
110	AMQDOC-2118	Hi - Just curious....who is the customer for this JIRA?. *Dean Richardson* Middleware Account Executive, Strategic Accounts drichardson@redhat.com (513) 382-9606 (mobile) (513) 231-3530 (office) --	1	1	https://issues.jboss.org/browse/AMQDOC-2118
111	AMQDOC-2118	I'm getting this too and don't know why. Might be that Jira is spamming everyone.	0	0	https://issues.jboss.org/browse/AMQDOC-2118
112	AMQDOC-2118	So this is a different version of SPAM. Should we open a JIRA ticket for that?	0	0	https://issues.jboss.org/browse/AMQDOC-2118
113	AMQDOC-2118	The second page of the PDF is empty, and the third just repeats info on the first/front page, perhaps they can be removed? GW: It is automatically rendered by our publishing software. I think these pages are left blank intentionally. In the TOC, Chapter 1 is listed as being about 'installing', which 1.1 and 1.2 then indeed are, but 1.3 is really for configuration. Perhaps it could be split out as a '2' rather than 1.3? GW: The ZIP package cannot work without subsequent configuration. But you comment do make sense. The Zip installation is likely the more universally applicable, maybe those instructions should be first ahead of the RPM ones? GW; I agree. The RPM instructions in "1.1.3. Installing JMS Client" say to install the C++ client. There is a note below it about the instructions being for Alpha, but that still seems off? GW: Yes. These instructions will be removed after Alpha release. In "1.1.7. AMQP JMS A-MQ 1.0 Java Client" it says "Qpid JMS is a complete AMQP JMS 1.0 and JMS 1.1 client". Should be "AMQP 1.0 and JMS 1.1 client"? The Note below this section seems like it needs updated, it describes it as if there are two seperate clients 'both' requiring Java 1.7+, but its just a single client. This section seems like it should perhaps be nearer at the start of the doc rather than mid-installation. GW: Yes. This is confusing. I will figure out better ways to make it more clear. The "1.1.8. AMQP C++ Client" and "1.1.9. Python Language Bindings for AMQP API" sections seem like they should be removed from the document given its JMS focus. GW: I think so. In "1.2. INSTALLING JMS CLIENT (ZIP INSTALLATION)" the initial instructions indicate to download the RPM and extract things from it. I presume the note about alpha means that should be removed from the final, given the instructions in "1.2.1. Downloading the A-MQ Broker ZIP File From the Red Hat Customer Portal" then go on to describe getting the actual zip from the website. The main thing here though was that the title in 1.2.1 should say Client rather than Broker. GW: Yes. Instructions for Alpha release will be removed after the release Downloading packages from Red Hat Customer Portal is not possible and actually irrelevant for the Alpha release. I keep it there as a place holder for future release.	1	1	https://issues.jboss.org/browse/AMQDOC-2118
114	AMQDOC-2117	Hi George, I see that you updated the Status of a handful of doc review sub-task JIRAs to Coding in Progress. I'm not certain if this status is reflective of the activity that's happening in real life. Does it denote you are incorporating SME feedback? Or are you requesting a review to confirm updates you've made? Let's discuss and clarify. I'm sure that our JIRA workflow could use some work.	0	0	https://issues.jboss.org/browse/AMQDOC-2117
115	AMQDOC-53	[~tabish121], I get why we have to add this information, but I can't find any mention of the KahaDB Store version setting. I've been looking in the Configuring Broker Persistence doco [1], but there's no mention. I've also been searching the Apache doco but it doesn't seem to be much different. Can you give me a clue as to where this setting might live? Cheers, Nichola [1] https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_A-MQ/6.2/html-single/Configuring_Broker_Persistence/index.html#FuseMBKahaDB	1	1	https://issues.jboss.org/browse/AMQDOC-53
116	AMQDOC-53	[~tabish121], I've added the setting to the table. Have a look at [1]. I basically used your description because I couldn't explain it any better myself! Can you let me know if it looks OK? Cheers, Nichola [1] http://jenkinscat.gsslab.pnq.redhat.com:8080/job/doc-AMQ-6-2-1/lastStableBuild/artifact/publican/Configuring_Broker_Persistence/tmp/en-US/html-single/index.html#KahaDBConfiguration	1	1	https://issues.jboss.org/browse/AMQDOC-53
117	ENTMQIC-1925	[~tedross] - Please review and update the issue. Should this be a bug instead of enhancement?	1	0	https://issues.jboss.org/browse/ENTMQIC-1925
118	ENTMQIC-74	Thank you for mentioning {{qdmanage}}, I might've forget that it has this functionality too. I noticed that the upstream project has a management web console. Is the console part of Interconnect as well?	1	0	https://issues.jboss.org/browse/ENTMQIC-74
119	ENTMQIC-69	Hi Ted, could you possibly add a description of this issue? Thank you in advance.	0	0	https://issues.jboss.org/browse/ENTMQIC-69
120	ENTMQIC-69	Is this issue in a valid state?	1	1	https://issues.jboss.org/browse/ENTMQIC-69
121	ENTMQIC-69	Is there a problem that Engineering needs to look at?	1	1	https://issues.jboss.org/browse/ENTMQIC-69
122	ENTMQIC-58	Matej, What does "Interior" mode mean? (As opposed to stand-alone mode) I could use a term that better describes this mode. This is in your Mojo doc in the Usage/Differences section.	1	0	https://issues.jboss.org/browse/ENTMQIC-58
123	AEROGEAR-1696	where is the latest info provided ?	0	0	https://issues.jboss.org/browse/AEROGEAR-1696
124	AEROGEAR-1665	Why is that assigned to me? Should I update the plugin names in the XML file? Not sure I understand this ticket	1	1	https://issues.jboss.org/browse/AEROGEAR-1665
125	AEROGEAR-1664	[~lfryc] ALL platforms in one document ? or one per platform (the later I like)	1	1	https://issues.jboss.org/browse/AEROGEAR-1664
126	AEROGEAR-1633	I believe it was partially resolved as part of other issue. Could you take a look and send a PR for few more changes that are required?	1	1	https://issues.jboss.org/browse/AEROGEAR-1633
127	AEROGEAR-1631	[~mwessendorf] Andres refactored the page as I suggested into one long document with TOC: http://staging-aerogearsite.rhcloud.com/docs/unifiedpush/ups_userguide/index/ WDYT?	1	1	https://issues.jboss.org/browse/AEROGEAR-1631
128	AEROGEAR-1616	Hey Andres, could you perhaps look at this as you work on footer changes?	0	0	https://issues.jboss.org/browse/AEROGEAR-1616
129	AEROGEAR-1615	What do you expect other then: "open Visual Studio and press build"?	1	1	https://issues.jboss.org/browse/AEROGEAR-1615
130	AEROGEAR-1607	Could you please do a design work and I will make it work javascript-wise?	1	1	https://issues.jboss.org/browse/AEROGEAR-1607
131	AEROGEAR-1607	Here it is: https://github.com/andresgalante/aerogear.org/tree/new-design We need to save vertical space on the sidebar. Can we use that awesome combo selection box?	1	1	https://issues.jboss.org/browse/AEROGEAR-1607
132	AEROGEAR-1580	It would be interesting to have some styling ready for use in blogs: - icons (such as info, warning, danger), see http://arquillian.org/guides/ - panels (used with icons: http://arquillian.org/guides/getting_started/) Another nice sample of panels with icons is on http://jekyllrb.com/docs/home/ Wdyt [~agalante], could you prepare such styles? A complete list TBD	1	1	https://issues.jboss.org/browse/AEROGEAR-1580
133	AEROGEAR-1580	[~lfryc] How about something like this: https://issues.jboss.org/browse/AEROGEAR-1551 Jekyll samples are great, should I make something fancier?	1	0	https://issues.jboss.org/browse/AEROGEAR-1580
134	AEROGEAR-1557	anymore info? Autoinstall what?	0	0	https://issues.jboss.org/browse/AEROGEAR-1557
135	AEROGEAR-1553	Let's do the HTML validation once again later before going with Site 2.0 to production.	1	1	https://issues.jboss.org/browse/AEROGEAR-1553
136	AEROGEAR-1484	I thought there already was a theme. Is this not what we see when we use the UPS Admin Console to update user information?	1	1	https://issues.jboss.org/browse/AEROGEAR-1484
137	AEROGEAR-1484	Ah gotcha. Thanks Just curious, what is a mobile theme in regards to KC?	1	1	https://issues.jboss.org/browse/AEROGEAR-1484
138	AEROGEAR-1484	Hey [~sebastienblanc], I believe it is sufficient for demoing at this point, we can expand later. WDYT?	1	1	https://issues.jboss.org/browse/AEROGEAR-1484

#	Requirement	Questions	Answered	Successfully Answered	URL
139	AEROGEAR-1464	I have created the project - https://issues.jboss.org/browse/AGCORDOVA/ Matthias, Erik, and Luke admins for now. Erik, could you go in and create the appropriate components, and versions for tracking, and labels as needed. Including marking issues with "Mobile Push 1.0" for issues related to that.	0	0	https://issues.jboss.org/browse/AEROGEAR-1464
140	AEROGEAR-1450	OK, I see I am able to create it. Does one know why to actually import the repo from passos ?	0	0	https://issues.jboss.org/browse/AEROGEAR-1450
141	AEROGEAR-1446	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-1446
142	AEROGEAR-1443	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-1443
143	AEROGEAR-1426	[~supittma] Would it be alright to close if I close this JIRA in favor or the tasks created for the [Revision control|http://aerogear.org/docs/planning/roadmaps/AeroGearConflictResolution] roadmap?	0	0	https://issues.jboss.org/browse/AEROGEAR-1426
144	AEROGEAR-1423	[~supittma] Can I close this in favour of the jiras create in the [Data sync|http://aerogear.org/docs/planning/roadmaps/AeroGearDataSync/] roadmap?	0	0	https://issues.jboss.org/browse/AEROGEAR-1423
145	AEROGEAR-1422	[~supittma] Can I close this in favour of the jiras create in the [Data sync|http://aerogear.org/docs/planning/roadmaps/AeroGearDataSync/] roadmap?	0	0	https://issues.jboss.org/browse/AEROGEAR-1422
146	AEROGEAR-1421	[~supittma] Can I close this in favour of the jiras create in the [Data sync|http://aerogear.org/docs/planning/roadmaps/AeroGearDataSync/] roadmap?	0	0	https://issues.jboss.org/browse/AEROGEAR-1421
147	AEROGEAR-1408	How is purge different from deleting a record from your local store and syncing that?	1	1	https://issues.jboss.org/browse/AEROGEAR-1408
148	AEROGEAR-1408	This is what I am thinking https://docs.google.com/drawings/d/1hYO0P-qJZJvn9sIUN3j_lxxLk4Vm0dWcWqdhJH88cEw/edit?usp=sharing How is `purge` separate from removing a record from your local data and then telling the sync client to sync?	1	1	https://issues.jboss.org/browse/AEROGEAR-1408
149	AEROGEAR-1408	i do see what you are saying though, but in that case, why do we have remove methods on a store, if save/update can do the same thing?	1	1	https://issues.jboss.org/browse/AEROGEAR-1408
150	AEROGEAR-1408	Ahhhh I see now. Basically I wasn't thinking of the sync client spec including data transfer per se when I made this issue. I was more thinking of it as a data inspection API. IE is the data in sync, is there new data? is there data to be pushed, will our syncs conflict, etc. Data transfer is a part of that which I always thought was "move what is new"	1	1	https://issues.jboss.org/browse/AEROGEAR-1408
151	AEROGEAR-1408	I guess we're focusing right now on sync itself rather than the whole picture, storing offline and then sync local storage with remote. Here we're in sending data (no store involved yet). BTW I didn't watch the JIRa issues so was out of sync with you guys. What about using ML thread: http://aerogear-dev.1069024.n5.nabble.com/aerogear-dev-sync-To-pipe-or-not-to-pipe-td6135.html To gist/discuss SyncPipe API (what luke called DataSync). we need to agree on API/vocabulary. wdyt?	1	0	https://issues.jboss.org/browse/AEROGEAR-1408
152	AEROGEAR-1398	This issue doesn't make sense to me inside an integration tests server. How could the app scan the QRCode? If we don't have anything on the client side, I vote to close it.	1	1	https://issues.jboss.org/browse/AEROGEAR-1398
153	AEROGEAR-1394	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-1394
154	AEROGEAR-1372	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-1372
155	AEROGEAR-1347	how about _jaxrs_ (https://jaxrs-aerogear.rhcloud.com) ?	1	1	https://issues.jboss.org/browse/AEROGEAR-1347
156	AEROGEAR-1316	do this still makes sense in face of our recent discussions? (ag-controller is dead)	1	0	https://issues.jboss.org/browse/AEROGEAR-1316
157	AEROGEAR-1297	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-1297
158	AEROGEAR-1279	Does this mean the AGPUSH project would just be UnifiedPush? If so, that means we remove any SimplePush labeled components from AGPUSH, right? I just want this to be clear so I don't have to keep track of JIRAs for a single "project" across multiple JIRA projects. Just like anything push relates for JS goes in AGJS.	0	0	https://issues.jboss.org/browse/AEROGEAR-1279
159	AEROGEAR-1269	Could you provide more info on this? Selecting a component and an affected version plus adding a description referencing some links to your code and steps to reproduce your error would be a good start. Thanks!	0	0	https://issues.jboss.org/browse/AEROGEAR-1269
160	AEROGEAR-1189	what is the desired repo name?	0	0	https://issues.jboss.org/browse/AEROGEAR-1189
161	AEROGEAR-1133	Is there a description / explanation on why we do/want this ?	0	0	https://issues.jboss.org/browse/AEROGEAR-1133
162	AEROGEAR-1133	There's an explanation about why we want to get rid of this. Works for you?	1	1	https://issues.jboss.org/browse/AEROGEAR-1133
163	AEROGEAR-1115	Is this still relevant ?	0	0	https://issues.jboss.org/browse/AEROGEAR-1115
164	AEROGEAR-995	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-995
165	AEROGEAR-970	Push was part of AG 1.1.0 (umbrella) release	0	0	https://issues.jboss.org/browse/AEROGEAR-970
166	AEROGEAR-952	I can't reproduce this issue. Abstractj or matzew, can you help me work out what is breaking between our systems?	0	0	https://issues.jboss.org/browse/AEROGEAR-952
167	AEROGEAR-903	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-903
168	AEROGEAR-888	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-888
169	AEROGEAR-849	Hi guys, may I close it?	1	1	https://issues.jboss.org/browse/AEROGEAR-849
170	AEROGEAR-830	is this still valid as an umbrella ticket? I think AEROGEAR-1374 fits the bill better	0	0	https://issues.jboss.org/browse/AEROGEAR-830
171	AEROGEAR-829	is this still relevant?	1	0	https://issues.jboss.org/browse/AEROGEAR-829
172	AEROGEAR-823	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-823
173	AEROGEAR-815	what name that repo would have? aerogear-examples?	1	0	https://issues.jboss.org/browse/AEROGEAR-815
174	AEROGEAR-791	The current guide has reference to the kitchensink-mobile archetype. If we want to update these guides to be based on the kitchensink-aerogear-js, would we need to also create a new archetype for that, or remove that reference altogether. Also, are the current guides somehow tied to JDF, or can we do anything we want?	0	0	https://issues.jboss.org/browse/AEROGEAR-791
175	AEROGEAR-791	what's the status of this?	0	0	https://issues.jboss.org/browse/AEROGEAR-791
176	AEROGEAR-791	There is no getting started guide for AeroGear Controller, though we do have the User Guide plus the controller demo (which show basic usage). Is that enough or should there be a getting started guide as well?	0	0	https://issues.jboss.org/browse/AEROGEAR-791
177	AEROGEAR-790	is the fixed by correct ?	0	0	https://issues.jboss.org/browse/AEROGEAR-790
178	AEROGEAR-786	So maybe is just a matter of add it at our demos? Let me know, maybe my description was too vague, I can change it.	1	1	https://issues.jboss.org/browse/AEROGEAR-786
179	AEROGEAR-774	We have decided to remove the year from the license. JS is done. Android and iOS, where are you at on this? I'll assign this to me for now and just bug you about it until you're done :)	0	0	https://issues.jboss.org/browse/AEROGEAR-774
180	AEROGEAR-772	Why would we want to create a plugin when we have a javascript implementation for this that works?	1	0	https://issues.jboss.org/browse/AEROGEAR-772

#	Requirement	Questions	Answered	Successfully Answered	URL
181	AEROGEAR-768	Hi Matthias, I'm trying to reproduce the error here. After insert a task with: <script>window.status='Hello World';\"></script> Is still possible to add more tasks. Am I missing something?	1	1	https://issues.jboss.org/browse/AEROGEAR-768
182	AEROGEAR-760	Is this still relevant ?	1	0	https://issues.jboss.org/browse/AEROGEAR-760
183	AEROGEAR-746	Is this iOS specific - or more a general thing for the site? (wondering about the two used components)	1	1	https://issues.jboss.org/browse/AEROGEAR-746
184	AEROGEAR-745	@Kris what did you have in mind that isn't covered under documentation/guides&tutorials/examples?	1	1	https://issues.jboss.org/browse/AEROGEAR-745
185	AEROGEAR-734	Updated http://corscontroller-aerogear.rhcloud.com/aerogear-controller-demo to version 1.0.0.M1.20121217-SNAPSHOT. Was there anything else to be done for this task Bruno?	0	0	https://issues.jboss.org/browse/AEROGEAR-734
186	AEROGEAR-707	shouldn't it be aerogear-store-ios, to avoid confusion?	1	1	https://issues.jboss.org/browse/AEROGEAR-707
187	AEROGEAR-674	what JS authentication demos? can you provide a link?	1	0	https://issues.jboss.org/browse/AEROGEAR-674
188	AEROGEAR-664	Is this task related to AEROGEAR-680 ?	1	0	https://issues.jboss.org/browse/AEROGEAR-664
189	AEROGEAR-629	Is there a 'status' for this?	1	0	https://issues.jboss.org/browse/AEROGEAR-629
190	AEROGEAR-629	regarding hypermedia, LINKing was already implemented (server and client). Is there anything else you had in mind ?	1	0	https://issues.jboss.org/browse/AEROGEAR-629
191	AEROGEAR-628	Is this still on our list ?	0	0	https://issues.jboss.org/browse/AEROGEAR-628
192	AEROGEAR-622	I noticed that PicketBoxLoadUsers has the following comment: {quote} //TODO this entire initialization code will be removed {quote} Can we remove this class since users can register there own accounts now?	1	0	https://issues.jboss.org/browse/AEROGEAR-622
193	AEROGEAR-611	Should the title of the Tutorial be changed also, from JBDS 5 to JDBS 6?	0	0	https://issues.jboss.org/browse/AEROGEAR-611
194	AEROGEAR-607	nice :) shouldn't the container (jax-rs impl) give a warning ?	1	1	https://issues.jboss.org/browse/AEROGEAR-607
195	AEROGEAR-607	Could you please elaborate more? I'm not following you	1	1	https://issues.jboss.org/browse/AEROGEAR-607
196	AEROGEAR-607	Ah..... I think I now understand... So.. This: {code} @ApplicationPath("/") public class Foo extends Application { } {code} will mess around with this: {code} public class Routes extends AbstractRoutingModule { @Override public void configuration() { route() .from("/") .on(RequestMethod.GET) .to(Home.class).index(); } } {code} Or am I wrong ?	1	1	https://issues.jboss.org/browse/AEROGEAR-607
197	AEROGEAR-493	How should the demo app handle being "offline"? It is not one of the current features of the application - at least not at this point.	0	0	https://issues.jboss.org/browse/AEROGEAR-493
198	AEROGEAR-492	Are we going to want the ability for users to build custom builds of past releases? or just the current release?	1	1	https://issues.jboss.org/browse/AEROGEAR-492
199	AEROGEAR-482	[~jbalunas], should we push this to 1.future? I honestly forgot about it and don't think it's a super-high priority if we even do it.	1	1	https://issues.jboss.org/browse/AEROGEAR-482
200	AEROGEAR-475	Hi Daniel, I saw that you've done it here https://github.com/danbev/aerogear-controller/commit/e96423550bfd93f8f158bf77273e3601fa05beb1 Do you want to take a stab at it and send the PR?	1	1	https://issues.jboss.org/browse/AEROGEAR-475
201	AEROGEAR-475	Hi Bruno, I've sent a pull request now. I was not able to assign this jira to myself, do you know if this is excepted?	1	1	https://issues.jboss.org/browse/AEROGEAR-475
202	AEROGEAR-459	Hi Matthias, what do you think about https://community.jboss.org/wiki/MavenSettings ?	0	0	https://issues.jboss.org/browse/AEROGEAR-459
203	AEROGEAR-457	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-457
204	AEROGEAR-447	Is it alright if I take this task?	0	0	https://issues.jboss.org/browse/AEROGEAR-447
205	AEROGEAR-438	Who has the PSD or the ability to modify this image?	1	1	https://issues.jboss.org/browse/AEROGEAR-438
206	AEROGEAR-438	i was able to modify the image http://blog.lholmquist.org/img/app_types_updated.png got the image from a presentation of matthias's, use it?	0	0	https://issues.jboss.org/browse/AEROGEAR-438
207	AEROGEAR-435	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-435
208	AEROGEAR-426	Hi Bruno, would you mind if I take a stab at this task? Thanks, /Daniel	1	1	https://issues.jboss.org/browse/AEROGEAR-426
209	AEROGEAR-412	Do we want an 'index page', located at http://aerogear.github.com/ ? Right now the JS and iOS (sub)projects have some documentation online: - http://aerogear.github.com/aerogear-js/ - http://aerogear.github.com/aerogear-ios/ Would be nice, to link to those, from an index page. What we need for that is just a github repo called 'aerogear.github.com'	1	0	https://issues.jboss.org/browse/AEROGEAR-412
210	AEROGEAR-395	Does this include moving anything that has changed from the quickstart repo to the archetype? Just want to make sure that is covered somewhere as there are changes in the quickstart that are not yet present in the archetype.	0	0	https://issues.jboss.org/browse/AEROGEAR-395
211	AEROGEAR-366	Could you please test it? https://github.com/aerogear/TODO/commit/ecd641e4895592b4126b8f7de844bf6d4221c6e8 I did some tests with: curl -v -H "Accept: application/json" -H "Content-type: application/json" -X POST -d '{"title": "test", "style":"mydescription"}' http://localhost:8080/todo-server/rest/project/ curl -v -H "Accept: application/json" -H "Content-type: application/json" -X GET http://localhost:8080/todo-server/rest/project/ curl -v -H "Accept: application/json" -H "Content-type: application/json" -X PUT -d '{"title":"test","style":"BROdescription"}' http://localhost:8080/todo-server/rest/project/1 for i in {1..10} do curl -v -H "Accept: application/json" -H "Content-type: application/json" -X PUT -d '{"title":"test","style":"$i-description"}' http://localhost:8080/todo-server/rest/project/1 done curl -v -H "Accept: application/json" -H "Content-type: application/json" -X GET http://localhost:8080/todo-server/rest/project/ Optimistic locking is not a requirement for us currently, of course we need to think about it in the future.	0	0	https://issues.jboss.org/browse/AEROGEAR-366
212	AEROGEAR-362	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-362
213	AEROGEAR-360	Could please test it https://github.com/aerogear/TODO/commit/a421a74091158b2c6ba52ab86ff5a035ee6066df and let me know how it goes?	1	1	https://issues.jboss.org/browse/AEROGEAR-360
214	AEROGEAR-342	Is https://issues.jboss.org/browse/AEROGEAR-581 a duplicate of this task perhaps?	0	0	https://issues.jboss.org/browse/AEROGEAR-342
215	AEROGEAR-338	Can you please add some more details for your jira? I'm not exactly what your setup is, or what you are using. I'm assuming this is related to the AeroGear Controller, or one of our examples - right?	1	1	https://issues.jboss.org/browse/AEROGEAR-338
216	AEROGEAR-336	is this out dated since we are deprecating asciidoc on aerogear.org?	0	0	https://issues.jboss.org/browse/AEROGEAR-336
217	AEROGEAR-326	In the Readme for the html5-kitchensink app, it says how to run the tests. Is that what the wiki was suppose to be, or a more in depth on? maybe a more General version? using Arquillian Drone?(haven't actually played with it yet)	1	1	https://issues.jboss.org/browse/AEROGEAR-326
218	AEROGEAR-310	Should these be the same as documented here: https://community.jboss.org/wiki/JBossProjectVersioning The current version of aerogear-controller is '1.0.0.Alpha1-SNAPSHOT', which seems like it fits in with what the above document describes (Alpha for early releases that are not API complete). Is this task only related to versioning in Jira/Roadmap?	1	1	https://issues.jboss.org/browse/AEROGEAR-310
219	AEROGEAR-296	Hi guys, if nobody's against, I'd like to start these requirements to android libraries. Could you please point me out to the specific documentation on iOS and JS?	0	0	https://issues.jboss.org/browse/AEROGEAR-296
220	AEROGEAR-253	What's the status on this ?	0	0	https://issues.jboss.org/browse/AEROGEAR-253

#	Requirement	Questions	Answered	Successfully Answered	URL
221	AEROGEAR-252	This is likely caused because the tests and application use the same datasource. During the M2 release we used deployable datasources and I do not believe we had this issue. However we had to remove that functionality because openshift did not support it at that time. Since then it seems the jboss-as archetypes have switched to using deployable datasources and do not have this issue. If we switch to this configuration we need to make sure that jboss-as archetypes/quickstarts will also be released this way, and that openshift, and the target jboss-as 7.1 release support this.	0	0	https://issues.jboss.org/browse/AEROGEAR-252
222	AEROGEAR-215	Is android/iOS the right component ? I'd assume only Cordova would be appropriate	1	0	https://issues.jboss.org/browse/AEROGEAR-215
223	AEROGEAR-215	Is this now JDF ?	1	0	https://issues.jboss.org/browse/AEROGEAR-215
224	AEROGEAR-174	Started work in https://github.com/mbogoevici/as-quickstarts/tree/upgrade-entities Some notes: - Changing URL scheme from /{entity}/{id}/json to /{entity}/json/{id} (to allow for /{entity}/json scheme prefixing) - Adding Country CRUD services Open points: Are some of the existing resources still needed (e.g. GET or form-based entity create) or shall we just focus on JSON/XML (which is what the actual example is using)?	0	0	https://issues.jboss.org/browse/AEROGEAR-174
225	AEROGEAR-174	Is there any way we can get rid of the XML stuff. We don't use it except to show that it's available do we? It's terrible to work with client side as well so everything is done with JSON now. Just a thought. :)	0	0	https://issues.jboss.org/browse/AEROGEAR-174
226	AEROGEAR-160	I've just tried the kitchensink-html5-mobile quickstart example on android 2.3.6, the member list is vertically scrolled but the one memeber is not displayed on one line. Did you rethink this or am I doing something wrong?	0	0	https://issues.jboss.org/browse/AEROGEAR-160
227	AEROGEAR-157	Kris - this should be two minutes since you already have the env setup. Can you grab this one.	0	0	https://issues.jboss.org/browse/AEROGEAR-157
228	AEROGEAR-130	Isn't this resolved? https://github.com/aerogear/aerogear-parent	0	0	https://issues.jboss.org/browse/AEROGEAR-130
229	AEROGEAR-126	Must now review and update based on latest updates to archetype - namely dir and url updates. Plus screenshots.	0	0	https://issues.jboss.org/browse/AEROGEAR-126
230	AEROGEAR-126	Almost all the updates are already done, but will the latest updates some images need updating (only archetype related ones in tools), and we need to mention the sub-module naming changes in archetype. The best bet is to go through the instructions and just see what is not accurate (should not be much at this point) https://community.jboss.org/wiki/GetStartedWithHTML5MobileWebDevelopmentWithJBoss https://community.jboss.org/wiki/HTML5MobileQuickstartArchetypeDeepDive	0	0	https://issues.jboss.org/browse/AEROGEAR-126
231	AEROGEAR-123	what do you think on having a deployer/handler that do the processing for us on deploy/runtime? (view compilation, minification, etc)	1	1	https://issues.jboss.org/browse/AEROGEAR-123
232	AEROGEAR-123	[~qmx] Any projects in mind? Do you mean a servlet filter with GZIP or a static file preprocessor?	1	1	https://issues.jboss.org/browse/AEROGEAR-123
233	AEROGEAR-121	What are you thinking for this?	0	0	https://issues.jboss.org/browse/AEROGEAR-121
234	AEROGEAR-108	There was no marketing or spins in my previous comment. In fact, your comment is really not appreciated, specially when I offered to help and provide a nice cool demo. Your reply doesn't even contain a thank you by the way. Politeness will take you a long way, you seem to have none. Our definition of purity also seems to be different. JBoss is mostly used by Java devs. A puristic approach would be do have an end-to-end Java solution. So my statement was by no means intended to be marketing related or spin as you put it, but a more elegant solution or showcase where Appcelerator is used at the end of the day. If you took the liberty of editing my previous comment, you can safely remove now both completely. Best of luck. Alfredo	1	1	https://issues.jboss.org/browse/AEROGEAR-108
235	AEROGEAR-107	Do we still showcase PhoneGap build even though we're using a more standard local build setup?	1	0	https://issues.jboss.org/browse/AEROGEAR-107
236	AEROGEAR-107	We should find the JBIDE jira for this and link it. We also need to review phonegap build again and see if this is something we still want to use. I.e. is it keeping up with the core Cordova project.	0	0	https://issues.jboss.org/browse/AEROGEAR-107
237	AEROGEAR-101	Also are you planning on flushing out the "In The Works" down? They could use some attention.	1	1	https://issues.jboss.org/browse/AEROGEAR-101
238	AEROGEAR-97	James - Can we get this video up on .org somewhere for AeroGear launch?	0	0	https://issues.jboss.org/browse/AEROGEAR-97
239	AEROGEAR-97	This is the link correct? Does this replace a current video? http://player.vimeo.com/video/35705170?byline=0&portrait=0	1	1	https://issues.jboss.org/browse/AEROGEAR-97
240	AEROGEAR-96	I've managed to get it onto tools, it will be like this once we manage to deploy https://img.skitch.com/20120206-pftpi1h8wtg1df6q2p21rga88u.jpg what are the next steps? will we wait for M2 to deploy it?	0	0	https://issues.jboss.org/browse/AEROGEAR-96
241	AEROGEAR-95	Comment from the review: The application will be running at the following URL http://localhost:8080/jboss-as-html5-mobile/. Should the context be changed to http://localhost:8080/jboss-as-kitchensink-html5-mobile to match the directory and jar name? Other than that it is good. The helloworld quickstart needs to be updated to match however.	0	0	https://issues.jboss.org/browse/AEROGEAR-95
242	AEROGEAR-86	Archetype notes: ---------- * Archetype versioning * should it sync with quickstart * referring to the /pom.xml not the one in resources * Do we need the .gitignore, or .factorypath? * Java Packages seem messed up ** https://github.com/aerogear/as-archetypes/tree/master/jboss-html5-mobile-archetype/src/main/resources/archetype-resources/src/test/java ** https://github.com/aerogear/as-archetypes/tree/master/jboss-html5-mobile-archetype/src/main/resources/archetype-resources/src/main/java * Still showing older quickstart webapp files ** https://github.com/aerogear/as-archetypes/tree/master/jboss-html5-mobile-archetype/src/main/resources/archetype-resources/src/main/webapp ** Need to get latest from Quickstart here Ideally a project built from the archetype (with the same groupid/artifactid) should be identical to the quickstart.	1	1	https://issues.jboss.org/browse/AEROGEAR-86
243	AEROGEAR-86	Wesley - as I mention in the comments for AEROGEAR-88 [1] PR#72 needs to be merged into aerogear/as-quickstarts, and not into /jbossas/quickstarts. Also can you put this change into as-archetypes? Doug already completed this jira. If needed, post launch we can figure out an automated way of doing this, but we don't have that now :-([1] https://issues.jboss.org/browse/AEROGEAR-88?focusedCommentId=12662058&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12662058	0	0	https://issues.jboss.org/browse/AEROGEAR-86
244	AEROGEAR-78	fixed at https://github.com/aerogear/as-archetypes/compare/c4d09d26a50a9f18b19e85f9c5d947b50bb47b1d^...master and https://github.com/aerogear/as-quickstarts/compare/25b35b0b006f90c85431f9660e91c8c946a3876a...e183e48d64ea6ead61330e412d91c3465e5fbff1 What about the code generated from the archetype? It don't inherits the license file too, right?	1	1	https://issues.jboss.org/browse/AEROGEAR-78
245	AEROGEAR-64	sanity check, did you mean ag-archetype and ag-quickstart?	1	1	https://issues.jboss.org/browse/AEROGEAR-64
246	AEROGEAR-64	i've did one initial round of git wrestling, can anyone review if this is what you all expected? (only quickstarts for now) https://github.com/aerogear/quickstart/tree/new-layout	1	1	https://issues.jboss.org/browse/AEROGEAR-64
247	AEROGEAR-63	Can anyone review the page please? Might need some polish, but I think the basics are there.	1	1	https://issues.jboss.org/browse/AEROGEAR-63

#	Requirement	Questions	Answered	Successfully Answered	URL
248	AEROGEAR-57	{quote} On my Android 2.3.6 it showed the mobile version, and the navigation bar also displayed. This was based on code from yesterday though, not sure if that changed. {quote} I forked the quickstart repo today and applied the patch. {quote} We know Android has some quirks here, any idea on the cause? If not feel free to comment in the PR {quote} I am not aware of the quirks since I do not have a 2.3.6 device to test on. {quote} The content flowing out of bounds is a known issue, not with Kris's This in on the portrait version of the list page - right? You or I should enter a jira for that {quote} It flows out on both orientations in list page. I guess you have a list of known issues? Might be better for you to create that one. {quote} Landscape on iOS This unfortunately is another known issue with no fix Its a problem with iOS and how it handles rotation Again we should probably put in a jira {quote} I did not have this problem with the RichFaces mobile showcase. Is this a JQuery issue?	1	1	https://issues.jboss.org/browse/AEROGEAR-57
249	AEROGEAR-55	can you do a simple google docs diagram showing me your ideas on this? You and Burr are the masterminds of this concept and I want to make sure I understand it fully as it plays out in your mind.	1	1	https://issues.jboss.org/browse/AEROGEAR-55
250	AEROGEAR-54	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AEROGEAR-54
251	AEROGEAR-45	do you think this is still needed?	0	0	https://issues.jboss.org/browse/AEROGEAR-45
252	AEROGEAR-44	I meant to say "What is Aerogear?" section... not the roadmap section. Thanks for clarifying Jay!	0	0	https://issues.jboss.org/browse/AEROGEAR-44
253	AEROGEAR-40	This should likely take the form of a wiki page. wdyt?	1	1	https://issues.jboss.org/browse/AEROGEAR-40
254	AEROGEAR-37	any feedback on this approach? Does it look good enough? https://github.com/qmx/as-archetypes/commit/08487b3bf840d553d4acef9277aafd01d68e55f0	1	1	https://issues.jboss.org/browse/AEROGEAR-37
255	AEROGEAR-33	What do we mean by POH5? Just want to be sure I have all the concepts included.	1	1	https://issues.jboss.org/browse/AEROGEAR-33
256	AEROGEAR-28	Nice! thanks for the warm comments. I just needed a few ideas from you. Thanks!	0	0	https://issues.jboss.org/browse/AEROGEAR-28
257	AEROGEAR-7	Does this need to be more than the archetype and quickstart already created? The archetype and quick start could be updated for full CRUD - at the moment they are CR not UD	1	1	https://issues.jboss.org/browse/AEROGEAR-7
258	AEROGEAR-7	Adding this as a subtask of as-quickstart. Should we add this functionailty to the current as-quickstart if we end up using a JavaScript MVC. We need a decision on that and if this should be a new demo under AeroGear or added to the quickstart.	1	0	https://issues.jboss.org/browse/AEROGEAR-7
259	AEROGEAR-2	Are you talking about the tweets on the right hand column or the RSS feed at the bottom of the main content pages?	1	1	https://issues.jboss.org/browse/AEROGEAR-2
260	AGDROID-599	[~danielpassos] since you removed this, are you ok when i assign this to you?	1	1	https://issues.jboss.org/browse/AGDROID-599
261	AGDROID-587	pipe/src/main/java/org/jboss/aerogear/android/pipe/rest/multipart/MultipartRequestBuilder.java#L155] started read the static fields like _serialVersionUID_ and we don't have get/set for this field to make [_Property_|https://github.com/aerogear/aerogear-android-core/blob/master/src/main/java/org/jboss/aerogear/android/core/reflection/Property.java] happy. Code: https://github.com/danielpassos/aerogear-android-cookbook/tree/shoot-android-n {code} 11-04 17:31:27.754 14859-14859/org.jboss.aerogear.android.cookbook.shootandshare D/CameraActivity: onPictureTaken - wrote bytes: 1066910 to /data/user/0/org.jboss.aerogear.android.cookbook.shootandshare/cache/1478287887748.jpg 11-04 17:32:57.442 14859-15167/org.jboss.aerogear.android.cookbook.shootandshare E/MultipartRequestBuilder: Cannot find get/set to field serialVersionUID (long) on PhotoHolder org.jboss.aerogear.android.core.reflection.PropertyNotFoundException: Cannot find get/set to field serialVersionUID (long) on PhotoHolder at org.jboss.aerogear.android.core.reflection.Property.findAccessorMethods(Property.java:97) at org.jboss.aerogear.android.core.reflection.Property.<init>(Property.java:60) at org.jboss.aerogear.android.pipe.rest.multipart.MultipartRequestBuilder.getProperties(MultipartRequestBuilder.java:158) at org.jboss.aerogear.android.pipe.rest.multipart.MultipartRequestBuilder.getBody(MultipartRequestBuilder.java:64) at org.jboss.aerogear.android.pipe.rest.RestAdapter$4.run(RestAdapter.java:241) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1113) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:588) at java.lang.Thread.run(Thread.java:818) 11-04 17:32:57.444 14859-15167/org.jboss.aerogear.android.cookbook.shootandshare E/UploadService: org.jboss.aerogear.android.core.reflection.PropertyNotFoundException: Cannot find get/set to field serialVersionUID (long) on PhotoHolder java.lang.IllegalStateException: org.jboss.aerogear.android.core.reflection.PropertyNotFoundException: Cannot find get/set to field serialVersionUID (long) on PhotoHolder at org.jboss.aerogear.android.pipe.rest.multipart.MultipartRequestBuilder.getBody(MultipartRequestBuilder.java:146) at org.jboss.aerogear.android.pipe.rest.RestAdapter$4.run(RestAdapter.java:241) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1113) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:588) at java.lang.Thread.run(Thread.java:818) Caused by: org.jboss.aerogear.android.core.reflection.PropertyNotFoundException: Cannot find get/set to field serialVersionUID (long) on PhotoHolder at org.jboss.aerogear.android.core.reflection.Property.findAccessorMethods(Property.java:97) at org.jboss.aerogear.android.core.reflection.Property.<init>(Property.java:60) at org.jboss.aerogear.android.pipe.rest.multipart.MultipartRequestBuilder.getProperties(MultipartRequestBuilder.java:158)	1	1	https://issues.jboss.org/browse/AGDROID-587

#	Requirement	Questions	Answered	Successfully Answered	URL
262	AGDROID-585	After update to Android N the app start to crash. I see it's related to Firebase version in Android N, but I did not had too much time to update the AeroGear Android Push library locally to test with new versions. Code: https://github.com/danielpassos/aerogear-android-cookbook/tree/hellopush-android-n **Error* {code} V/FA: Processing queued up service tasks: 1 E/FA: Task exception on worker thread: java.lang.IllegalAccessError: Method 'void android.support.v4.content.ContextCompat.<init>()' is inaccessible to class 'com.google.firebase.iid.zzg' (declaration of 'com.google.firebase.iid.zzg' appears in /data/app/org.jboss.aerogear.android.cookbook.hellopush-2/base.apk): com.google.android.gms.measurement.internal.zzt.zzEd(Unknown Source) V/FA: Activity paused, time: 27278 V/FA: onActivityCreated --------- beginning of crash E/AndroidRuntime: FATAL EXCEPTION: AsyncTask #1 Process: org.jboss.aerogear.android.cookbook.hellopush, PID: 2690 java.lang.RuntimeException: An error occurred while executing doInBackground() at android.os.AsyncTask$3.done(AsyncTask.java:309) at java.util.concurrent.FutureTask.finishCompletion(FutureTask.java:354) at java.util.concurrent.FutureTask.setException(FutureTask.java:223) at java.util.concurrent.FutureTask.run(FutureTask.java:242) at android.os.AsyncTask$SerialExecutor$1.run(AsyncTask.java:234) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1113) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:588) at java.lang.Thread.run(Thread.java:818) Caused by: java.lang.IllegalAccessError: Method 'void android.support.v4.content.ContextCompat.<init>()' is inaccessible to class 'com.google.firebase.iid.zzg' (declaration of 'com.google.firebase.iid.zzg' appears in /data/app/org.jboss.aerogear.android.cookbook.hellopush-2/base.apk) at com.google.firebase.iid.zzg.zzeC(Unknown Source) at com.google.firebase.iid.zzg.<init>(Unknown Source) at com.google.firebase.iid.zzg.<init>(Unknown Source) at com.google.firebase.iid.zzd.zzb(Unknown Source) at com.google.firebase.iid.FirebaseInstanceId.getInstance(Unknown Source) at com.google.firebase.iid.FirebaseInstanceId.getInstance(Unknown Source) at org.jboss.aerogear.android.unifiedpush.fcm.AeroGearFCMPushRegistrar$2.get(AeroGearFCMPushRegistrar.java:100) at org.jboss.aerogear.android.unifiedpush.fcm.AeroGearFCMPushRegistrar$2.get(AeroGearFCMPushRegistrar.java:96) at org.jboss.aerogear.android.unifiedpush.fcm.AeroGearFCMPushRegistrar$4.doInBackground(AeroGearFCMPushRegistrar.java:145) at org.jboss.aerogear.android.unifiedpush.fcm.AeroGearFCMPushRegistrar$4.doInBackground(AeroGearFCMPushRegistrar.java:135) at android.os.AsyncTask$2.call(AsyncTask.java:295) at java.util.concurrent.FutureTask.run(FutureTask.java:237) at android.os.AsyncTask$SerialExecutor$1.run(AsyncTask.java:234) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1113) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:588) at java.lang.Thread.run(Thread.java:818) V/FA: Activity resumed, time: 27389 E/Surface: getSlotFromBufferLocked: unknown buffer: 0x7fe4e955a240 V/FA: Activity paused, time: 28121 {code}	1	1	https://issues.jboss.org/browse/AGDROID-585
263	AGDROID-574	I have started the work here: https://github.com/danielpassos/aerogear-android-pipe/tree/Android-N What I did: * Bump to Android N * Bump plugins versions (some plugins have Android 23 dependency) * Organize/update travis.yml What I didn't * Bump core version to 3.1.0 (waiting for be on Maven Central) * Update test to the new Android test style	1	1	https://issues.jboss.org/browse/AGDROID-574
264	AGDROID-571	[~danielpassos] Yep, just realised, sorry about that! So can this just be moved to Done?	1	1	https://issues.jboss.org/browse/AGDROID-571
265	AGDROID-551	[~supittma] as discussed in AGPUSH-1636, the entire Google section (including screenshots) need to be updated. Perhaps this is something that [~rsmeral] or [~ssitani] can assist with ?	1	1	https://issues.jboss.org/browse/AGDROID-551
266	AGDROID-551	Which doc is this referring to? https://aerogear.org/docs/unifiedpush/ups_userguide/index/, or https://aerogear.org/docs/unifiedpush/aerogear-push-android/guides/ ? We need changes in both: https://aerogear.org/docs/unifiedpush/ups_userguide/index/ * change mentions of GCM to e.g. "FCM (formerly GCM)" or similar https://aerogear.org/docs/unifiedpush/aerogear-push-android/guides * the whole "Google Related Setup" needs to be redone for FCM - I believe we (Docs team) could handle this by 3.12 * "Your first Android/Push App" and "Send a Push Notification" - these need updates for FCM too, but would require significant on-ramp for us (docs), so I can't commit to this for 3.12 without Engineering help. That being said - [~supittma], which part are you working on here? (you are assigned and Coding in Progress) Also, why was AGPUSH-1636 closed? It covered the Google-related setup, while this issue seems (by the description) to cover some other part. They don't seem to be duplicates. If we re-open AGPUSH-1636, I can take it into this/next sprint.	1	1	https://issues.jboss.org/browse/AGDROID-551
267	AGDROID-551	[~ssitani] can you send a PR with whatever you have? So that [~supittma] and I can continue and solve the doc issue ?	1	1	https://issues.jboss.org/browse/AGDROID-551
268	AGDROID-551	[~mwessendorf] [~supittma] Small update on our (docs team) availability: at this moment, we have to invest 100% of time on RHMAP 4.2 documentation - there's a lot to write/edit: https://issues.jboss.org/browse/RHMAP-9018 I'm unassigning this from [~ssitani] for now. We may be able to do a language review if you can provide a complete draft. WDYT?	1	1	https://issues.jboss.org/browse/AGDROID-551
269	AGDROID-550	I think we could add an empty JSON file to the app folder of the demos, and get some documentation out, what's needed to get the file (content) in place	1	1	https://issues.jboss.org/browse/AGDROID-550
270	AGDROID-546	/cc [~danielpassos] So this has some interesting consequences for the UPS life cycle. Right now we assume that when the success callback is called we are registered with all of the categories, have a token, etc. Now happens inside of a service that we extend. However, what does this mean for the register method? For starters, all of the methods we call are non blocking so we can kill the Callback method (YAY!), but this means that failures to connect to UPS are now handled by a service and not communicated to the client (boo!). There are several solutions to this, but none of them are a clear cut winner. 1) We continue to use the callback, but don't call it until the registration state with UPS is determined. 2) We register a handler that is called with the connection state passed to it (similar to how Google's client SDKs work) 3) UPS state is totally handled by our libraries. Developers will tell us what state they want us to be in and the library will fight the good fight keeping the device in that state. wdyt?	0	0	https://issues.jboss.org/browse/AGDROID-546
271	AGDROID-539	Should we send the global topic to the push service with the categories payload?	1	1	https://issues.jboss.org/browse/AGDROID-539
272	AGDROID-519	[~danielpassos] [~supittma] with Android N coming, will this mean, that there is JDK8 / OpenJDK only ?	0	0	https://issues.jboss.org/browse/AGDROID-519
273	AGDROID-428	yeah, would be nice, no ?	1	0	https://issues.jboss.org/browse/AGDROID-428

#	Requirement	Questions	Answered	Successfully Answered	URL
274	AGDROID-421	[~supittma] Maybe I don't get it but I meant this method: https://github.com/aerogear/aerogear-android-push/blob/45606e795f10e6a6343183981fd2b12dd17f0d49/aerogear-android-push/src/main/java/org/jboss/aerogear/android/unifiedpush/gcm/UnifiedPushConfig.java#L145-L156 There is no way to get the device token from the android libs right? So why can I set it here?	1	1	https://issues.jboss.org/browse/AGDROID-421
275	AGDROID-421	[~danielpassos] I took a quick look but this isn't actually used. In the constructor of GCMPushRegistrar we set the value, but it is always overriden in the .register method. wdyt about removing it?	1	1	https://issues.jboss.org/browse/AGDROID-421
276	AGDROID-376	Did you mean 2015? ;) Much of the benefit of the GCM/XMPP bridge is that it is a two way communication channel that is managed by Android. Device goes down and comes back up? Messages will make it out and the app won't crash. Last I checked the best Java/Android WebSocket client still needed lots of setup around it. Ideally by the time this hits final we will be able to make all of these choices configurable and/or automagic.	1	1	https://issues.jboss.org/browse/AGDROID-376
277	AGDROID-376	|yes, but in a lot of countries delivering data to a 3rd party (which would be the fact here) is not allowed. Have you been spying on me? Because I swear I had this conversation with a different German last night... |Having a solid aerogear-android-websocket repo/library would be benefit for the overall Android community, not just us Writing a solid Android WebSocket library is something I want to do.	0	0	https://issues.jboss.org/browse/AGDROID-376
278	AGDROID-351	Most likely can be resolved. Is there any ETA to release 3.0.0 version of this?	1	1	https://issues.jboss.org/browse/AGDROID-351
279	AGDROID-319	Is there anything blocking progress on this issue? Is it a simple matter of replacing the webview invocation with an appropriate intent? If so I'll gladly provide a PR making this change.	1	1	https://issues.jboss.org/browse/AGDROID-319
280	AGDROID-289	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	1	0	https://issues.jboss.org/browse/AGDROID-289
281	AGDROID-258	That;s actually a new feature on iOS8 [~supittma] If you want to support that on Android (does it make sense?), we can send that to GCM as well	1	1	https://issues.jboss.org/browse/AGDROID-258
282	AGDROID-254	Could you provide a list of which elements you want to test to jumpstart this?	1	1	https://issues.jboss.org/browse/AGDROID-254
283	AGDROID-250	Do we plan to make this gotcha a part of our guides?	1	1	https://issues.jboss.org/browse/AGDROID-250
284	AGDROID-231	Wasn't [~aemmanou] doing similar already for the Java Sender's Http(s)URLConnection ?	0	0	https://issues.jboss.org/browse/AGDROID-231
285	AGDROID-218	We've moved on to 4.3.0 at this point. [~kpiwko] if this is still a problem please reopen with a link to a failing project in github.	0	0	https://issues.jboss.org/browse/AGDROID-218
286	AGDROID-59	Was this for a simple demo (PushEEDemo - used by the tutorial) ? Or was this the ProDoctor/AeroDoc "topic demo" ?	1	0	https://issues.jboss.org/browse/AGDROID-59
287	AGDROID-7	Passos, did you work on this at some point?	1	1	https://issues.jboss.org/browse/AGDROID-7
288	AGCORDOVA-158	I've been doing some testing: If the app is in foreground, it doesn't appear as open, shouldn't it appear as open as the message is received? If the app is in background it works fine in both iOS and Android. If the app is closed: - It doesn't appear as open for Android. - It can appear as 2 messages open for iOS instead of 1	0	0	https://issues.jboss.org/browse/AGCORDOVA-158
289	AGCORDOVA-136	Hi ! all of you, I am tested org.jboss.aerogear.cordova.push plguin . this plugin is deprecated with cordova 6.0.0 version right now. This is working fine with till cordova version 5.x. , I request to aerogear lead developer to update this plugin for further use. Thanks.	1	1	https://issues.jboss.org/browse/AGCORDOVA-136
290	AGCORDOVA-134	[~edewit] [~mwessendorf] [~supittma] adding above code resolves android errors however there are iOS errors as well. Please see https://issues.jboss.org/browse/AGCORDOVA-146 issue for iOS error logs. Is there any way you guys can take a look at this?	0	0	https://issues.jboss.org/browse/AGCORDOVA-134
291	AGCORDOVA-50	3) I'm not sure, but I don't think you must have this, on the ios shoot project it's also not there 4) why should these be remove? Isn't it nice to have an example of how to configure it? 5) when you add android platform you get the information about building via gradle in your terminal I must admit I rather have that we didn't require these additional setup steps, but android and ios libraries are not easy to consume in a cordova project.	1	1	https://issues.jboss.org/browse/AGCORDOVA-50
292	AGCORDOVA-50	3) I am not sure if it is required, but Android version of this application has it. 4) Well I guess that this client Ids belong to somebody on our team. What if somebody just copies the project, build their app on top of it and release it with this Ids? This doesn't concern us? And what about the developer? 5) I know and I think it is not sufficient, because you make the first steps once, but you might want to build the application several times and later forget how to do it. I think what we should strive for is a great developer's experience and easy of use of the examples. Complete and clear documentation is necessary to do that.	1	1	https://issues.jboss.org/browse/AGCORDOVA-50
293	AGCORDOVA-50	I am reopening this, because there are still some problems. Ad 2) Facebook worked fine. But Google & Keycloak buttons still don't work. This time, instead of crashing, their output is alert window with "java.lang.reflect.InvocationTargetException". Does it work for you, am I doing something wrong? With that in mind I just realized that there is no information about the configuration of oauth providers in README. There should probably be a link to Android version which covers setting up the accounts, but it should also mention what changes have to be made in app.js and controllers.js.	1	1	https://issues.jboss.org/browse/AGCORDOVA-50
294	AGCORDOVA-50	have you added the extra dependencies? e.g. compile 'com.google.code.gson:gson:1.7.2' compile 'com.google.guava:guava:18.0' I'll add some more to the readme	1	1	https://issues.jboss.org/browse/AGCORDOVA-50
295	AGCORDOVA-50	[~pstribny] regarding "information about the configuration of oauth providers in README", can you file a new/different JIRA ?	1	1	https://issues.jboss.org/browse/AGCORDOVA-50
296	AGCORDOVA-50	have you updated the plugin with the new version 1.0.1?	1	1	https://issues.jboss.org/browse/AGCORDOVA-50
297	AGCORDOVA-50	[~pstribny] [~edewit] since the error is slightly different (This time, instead of crashing, their output is alert window with "java.lang.reflect.InvocationTargetException".) should we get a new JIRA for this as well ?	1	1	https://issues.jboss.org/browse/AGCORDOVA-50
298	AGCORDOVA-50	I've updated the 1.0.1 release with a fix for this, seems I should have tested all the settings not just google. Could you give it another go?	0	0	https://issues.jboss.org/browse/AGCORDOVA-50
299	AGCORDOVA-50	I deleted the platforms & plugins folder and built the application again and it behaves in the same way. Should I also clear/delete some dependencies?	0	0	https://issues.jboss.org/browse/AGCORDOVA-50
300	AGCORDOVA-35	what's the status here ?	1	1	https://issues.jboss.org/browse/AGCORDOVA-35
301	AGCORDOVA-21	[~edewit] Since the related PR has been merged to a branch, due to the fact that some more work is needed, should the Fix Version be updated to push-1.0.1 ? Let's not make that part of the initial 1.0.0 release	1	1	https://issues.jboss.org/browse/AGCORDOVA-21
302	AGCORDOVA-21	what should be the fix version for this ?	1	1	https://issues.jboss.org/browse/AGCORDOVA-21

#	Requirement	Questions	Answered	Successfully Answered	URL
303	AGDIGGER-113	This is a bit problematic... {{zipalign}} is versioned under {{build-tools}} directory. e.g. on the current Android slave container I checked, it is here: {{/opt/android-sdk-linux/build-tools/25.0.3}} This means, we have to know what version of the build tools are available in the Android SDK container that we mount to Android slave container. [~lrossett] [~weil] what do you think of creating a symlink {{LATEST}} in {{build-tools}} directory in Android SDK container? This should be easy to do with androidctl. However, this brings 2 questions to my mind: 1. Is it ok to use latest build tools all the time, at least in Jenkins Android Signing plugin? 2. Would symlinks in the Android slave container? Note that symlink will be created in Android SDK container, not slave. I cannot think of any other solution except defining the version as an env var in Android slave container's deployment config?	1	1	https://issues.jboss.org/browse/AGDIGGER-113
304	AGDIGGER-113	Shouldn't we just set the "zipalignPath" parameter in the jenkinsfile as the default/recommended version? Users/developers/customers can change it as they want. I think that keeping it as a string with the zipalign path makes it easy and obvious for someone to change it rather than hide it in an env var.	1	1	https://issues.jboss.org/browse/AGDIGGER-113
305	AGDIGGER-108	[~weil] Are you happy to close this one. Is the installation of the instance on OSM4 sufficient to consider this one verified?	1	1	https://issues.jboss.org/browse/AGDIGGER-108
306	AGDIGGER-108	[~lfitzgerald] Yeah, I think so. Does the groovy scripts that updates the templates and setup proxy works fine? I am asking because we are seeing some issues with the jenkins-cli in different versions of Jenkins. To make sure, you need to remove the jenkins-cli.jar from the local dir (like /tmp) before running the playbooks.	1	1	https://issues.jboss.org/browse/AGDIGGER-108
307	AGDIGGER-99	[~wtrocki] thanks for the investigation. really good work! It seems that the integration is easier than I thought. 2 questions: 1. {{## Monitor filesystem to check what files were changed}} how easy is it to do this? This is an awesome idea! We can remove created files afterwards. 2. "slaves on demand" is awesome, but I am worried that the queuing could be problematic. We don't want to have our same machine to have 10s of slaves running at the same time and building things. It might be bad in terms of performance and memory. Any experience with that?	1	0	https://issues.jboss.org/browse/AGDIGGER-99
308	AGDIGGER-96	Template should support one slave. More slaves can be added using gui (which is no more complicated than environment variables in template) We should document this process. Documentation ticket is here: https://issues.jboss.org/browse/AGDIGGER-69 Can we close this as duplicate and provide more useful information to description. ??We reference the node name (slave name) in Jenkinsfile. But ideally, Jenkins should pick any OSX slave which is not busy. Perhaps: https://wiki.jenkins-ci.org/display/JENKINS/Compound+Slaves?? That's really good comment. We should make sure that this is documented and tested. We may just say that users would need to specify slave name in pipeline definition or.. loadballance using mentioned plugin. For now it's best to keep it simple.	0	0	https://issues.jboss.org/browse/AGDIGGER-96
309	AGDIGGER-88	[~luigizuccarelli] - I would say that for this ticket we do not need to provide any runnable scripts. Best to figure out how we can do this operation on one node (for example using labels) to feed android SDK to the build process. Once we would know the best way to do it we can finish this as part of ansible epic. WDYT?	1	1	https://issues.jboss.org/browse/AGDIGGER-88
310	AGDIGGER-87	What's the status here ??	1	1	https://issues.jboss.org/browse/AGDIGGER-87
311	AGDIGGER-80	How to prepare sandbox definiton files: * Create a sandbox file which traces things: {code} ; file logall.sb (version 1) (debug all) (trace "trace.sb") {code} * Use the file {code} sandbox-exec -f logall.sb xcodebuild arg1 arg2 {code} * Open trace.sb, remove duplicate lines in it * Optimize and normalize things in it * It makes sense to change many {{path}} entries with {{literal}} entries * Use the new sandbox file and see if the same command works! The way that tracing works is, * it logs only the things that are not normally allowed. * everything is allowed So, you can enable tracing again in the optimized file and in that case you will only see the requests that are not satisfied.	1	1	https://issues.jboss.org/browse/AGDIGGER-80
312	AGDIGGER-79	[~aliok2] Why don't we make "PLATFORM" parameter type - "Choice parameter" and add just valid choices android/ios ?	1	0	https://issues.jboss.org/browse/AGDIGGER-79
313	AGDIGGER-73	[~wtrocki] Not really sure what you mean. It looks like the DiggerClient is supporting passing the branch info when creating a Jenkins job. Are you saying it is not working? Or did you mean that the we need to integrate the DiggerClient with the platform?	1	1	https://issues.jboss.org/browse/AGDIGGER-73
314	AGDIGGER-71	[~wtrocki] Not really sure what's this ticket is about. What "client" do you mean here? RHMAP or the digkins Java client?	1	1	https://issues.jboss.org/browse/AGDIGGER-71
315	AGDIGGER-70	[~wtrocki] how does this relate to RHMAP-11958 ?	1	1	https://issues.jboss.org/browse/AGDIGGER-70
316	AGDIGGER-70	what permissions would those normal users have ?	1	1	https://issues.jboss.org/browse/AGDIGGER-70
317	AGDIGGER-63	??tell build use "Android 24" image, based on this setting:?? How you see this happening? I would see it as node("Android 24") so you need to specify this as target in your jenkinsfile and do not use actual setting. If we want to extend this idea we can have some function called from Jenkinsfile: {code} var nodeName = "android" + readAndroidVersionFromBuildGradle(); node(nodeName) { ... } {code} ??Perhaps this can be added as a plugin? As a way to tell the kubernetes plugin "launch Android-24 image"?? Kubernetes plugin would have everything for this already. We planning to wrap s2i process into ansible scripts. When images are created then we would need to change kubernetesPodTemplates, but this would happen just once (or better to say every time someone would run playbook). I my missing some context here so correct me if I'm wrong but I think that having plugin here would just complicate things. Let's add this to the topic list for technical discussions and create some additional jiras if needed.	1	1	https://issues.jboss.org/browse/AGDIGGER-63
318	AGDIGGER-63	[~wtrocki] Ok tried the digger-node , the document needs to be updated a bit - It does not show the steps to install node and run it locally - I know what todo but the community will be confused. Also got this error node bin/digkins.js login https://jenkins-jenkins.192.168.2.194.xip.io admin xxxxx info: Cannot connect to jenkins instance Error: self signed certificate in certificate chain at Error (native) at TLSSocket.<anonymous> (_tls_wrap.js:1022:38) at emitNone (events.js:67:13) at TLSSocket.emit (events.js:166:7) at TLSSocket._init.ssl.onclienthello.ssl.oncertcb.TLSSocket._finishInit (_tls_wrap.js:586:8) at TLSWrap.ssl.onclienthello.ssl.oncertcb.ssl.onnewsession.ssl.onhandshakedone (_tls_wrap.js:428:38)	1	1	https://issues.jboss.org/browse/AGDIGGER-63
319	AGDIGGER-58	[~aliok2] https://github.com/aerogear/digger-java/pull/1 PR created. Can you review that? I merged that by mistake, any comments would be resolved as different PR :)	0	0	https://issues.jboss.org/browse/AGDIGGER-58
320	AGDIGGER-41	[~wtrocki] Could you please write a summary of what we've talked during the backlog refinement meeting?	1	0	https://issues.jboss.org/browse/AGDIGGER-41
321	AGDIGGER-33	[~wtrocki] would love to see some instructions with bunch of commands somewhere in our repo? mind preparing one? think about community as well.	1	0	https://issues.jboss.org/browse/AGDIGGER-33
322	AGDIGGER-22	Is there any reason why ant is crossed out? While there is no ant template at this point, can we ensure there are no Ant based apps currently at customers side? AIUI this would be breaking backwards compatibility with 3.15.0.	1	1	https://issues.jboss.org/browse/AGDIGGER-22

#	Requirement	Questions	Answered	Successfully Answered	URL
323	AGDIGGER-22	Q: Where to have the debug keystore? In Openshift? Outside Openshift? In Docker image? etc [~mwessendorf] That's sounds like separate issue for me. Can I create this? A: It is part of the build. We need to provide the debug keystore and users won't provide their own (for MVP) [~mwessendorf] IMHO: That's platform integration part. I have couple ideas for the logic here, but basically we can copy solution we using in old digger.	0	0	https://issues.jboss.org/browse/AGDIGGER-22
324	AGDIGGER-20	[~luigizuccarelli] I now have my Openshift cluster working properly. Gave a try, although we now have other pipelines. Same result with you, all good!	0	0	https://issues.jboss.org/browse/AGDIGGER-20
325	AGIOS-439	is this any longer needed - given that iOS9 is around ?	0	0	https://issues.jboss.org/browse/AGIOS-439
326	AGIOS-427	Why are we removing https://github.com/aerogear/aerogear-ios-sync-client ?	1	1	https://issues.jboss.org/browse/AGIOS-427
327	AGIOS-394	This is the OTP demo, right? https://github.com/aerogear/aerogear-ios-cookbook/tree/master/Two-Factor It's updated to Swift 3 So, can we close this?	1	1	https://issues.jboss.org/browse/AGIOS-394
328	AGIOS-347	how about moiving to future?	0	0	https://issues.jboss.org/browse/AGIOS-347
329	AGIOS-343	Do we really need SockJS ? I'd think that WSS would be good enough here	0	0	https://issues.jboss.org/browse/AGIOS-343
330	AGIOS-342	Should the label _sync-1.0.0.alpha.1_ be removed for this task. I don't think this will be included for 1.0.0.alpha.1?	1	0	https://issues.jboss.org/browse/AGIOS-342
331	AGIOS-277	Good morning [~corinnekrych], I couldn't find any good reference on it. Is the section 4.4 from the RFC what do you have in mind? http://tools.ietf.org/html/rfc6749#section-4.4	1	0	https://issues.jboss.org/browse/AGIOS-277
332	AGIOS-271	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	0	0	https://issues.jboss.org/browse/AGIOS-271
333	AGIOS-255	Once this is there, does that mean the 2.x version of the push/registration SDK can be based on aerogear-ios-http? (since that requires http basic to work)	1	1	https://issues.jboss.org/browse/AGIOS-255
334	AGIOS-231	Initial version is: https://github.com/corinnekrych/aerogear-ios-oauth2 could you move it to aerogear organization? it contains OAuth2: - request authz code - request access token - refresh token - revoke AccountManger for fluid integration and storage of tokens is tracked in this JIRA https://issues.jboss.org/browse/AGIOS-253	1	1	https://issues.jboss.org/browse/AGIOS-231
335	AGIOS-150	Thanks Corinne. From the description i understand that you are looking for an option to allow uploading a new file to Google Drive and also some way to demo modify/update (i.e. could be just as simple as updating the file title) an existing file. Is my understaning correct? Regards, Yagyesh	1	1	https://issues.jboss.org/browse/AGIOS-150
336	AGIOS-150	I have started working on this & hopefully by today i could have a PR ready. In the meantime, would it be possible to get this ticket assigned to me. I tried, but looks like i dont have permissions.	0	0	https://issues.jboss.org/browse/AGIOS-150
337	AGIOS-34	We already have a branch, that does update the requirements to iOS5 and adds AFN 1.2.1 (as a matter of fact) https://github.com/aerogear/aerogear-ios/tree/matzew_afn121	1	1	https://issues.jboss.org/browse/AGIOS-34
338	AGIOS-17	[~corinnekrych] Let's lower the priority for this one. IMO all other 2.2.0 items are way more important. Ok?	0	0	https://issues.jboss.org/browse/AGIOS-17
339	AGIOS-11	we could do it for 1.2.0 as well;	1	1	https://issues.jboss.org/browse/AGIOS-11
340	AGIOS-6	[~abstractj] do you recall why we want/need such a parameter ?	1	0	https://issues.jboss.org/browse/AGIOS-6
341	AGIOS-6	It;'s ok :-) We were not sure as well - let's keep it ... perhaps someday... there is need for it :-)	0	0	https://issues.jboss.org/browse/AGIOS-6
342	AGJS-272	This is now available on master: https://github.com/aerogear/aerogear-sync-server/tree/master/distribution When _mvn package_ is run the produces jar will be in the format: {noformat} target/aerogear-sync-server-version.jar {noformat} How do you want to use this. Do we need to stick the produced jar somewhere? This is actually not part of the build at the moment (is a manual task) but perhaps it should so you can access snapshots. Let me know what you need and we will sort it out.	1	0	https://issues.jboss.org/browse/AGJS-272
343	AGJS-261	the JIRA says PR sent, but the PR was closed. Should it go back to "open" ?	1	1	https://issues.jboss.org/browse/AGJS-261
344	AGJS-260	[~lucas.holmquist] the JIRA says PR sent, but the PR was closed. Should it go back to "open" ?	1	1	https://issues.jboss.org/browse/AGJS-260
345	AGJS-255	[~lucas.holmquist] I was digging into a several repositories and was wondering about where the OAuth2 library is located today. Is here https://github.com/lholmquist/ag-google-drive ?	1	0	https://issues.jboss.org/browse/AGJS-255
346	AGJS-237	Notifier's API's are event based, so Promises would not be recommend here	0	0	https://issues.jboss.org/browse/AGJS-237
347	AGJS-189	I believe this is comment to AGJS-211 ^, am I right?	0	0	https://issues.jboss.org/browse/AGJS-189
348	AGJS-175	By dependencies you mean: {code} Pipeline -depends-on-> Ajax then watch(ajax) --> run(ajax, pipeline, ...)? {code}	1	1	https://issues.jboss.org/browse/AGJS-175
349	AGJS-143	[~lucas.holmquist] do we have any use cases for that? Why current solution isn't sufficient?	0	0	https://issues.jboss.org/browse/AGJS-143
350	AGJS-137	Could this please be unassigned from me? I have no permissions to drop it :)	0	0	https://issues.jboss.org/browse/AGJS-137
351	AGJS-64	Have a suggestion parked in this [branch|https://github.com/danbev/aerogear-js/tree/AGJS-64-subprotocols] until the SockJS [task|https://github.com/sockjs/sockjs-client/issues/130] is completed.	1	0	https://issues.jboss.org/browse/AGJS-64
352	AGJS-47	The QUnit tests reside in the integration repo: https://github.com/tolis-e/aerogear-js-integration Execute ./servers/activemqtest/server.sh and then open the tests/notifier/mqttws.html file in a browser. The problem is that it doesn't work in phantomjs for some reason after updating to activemq-5.9.0 (version which supports MQTT). When using the debug mode I can see: phantomjs console Unexpected response code: 400. Looks like the issue is probably related to https://groups.google.com/forum/#!msg/phantomjs/p4scrV2GwB4/c--SS5owLZIJ	1	1	https://issues.jboss.org/browse/AGJS-47
353	AGJS-41	Hey [~sienaluke], would you mind to review this beast? :-)	0	0	https://issues.jboss.org/browse/AGJS-41
354	AGJS-22	So, is there a generic way to wrap form data as JSON for PUT/POST to a REST endpoint?	1	1	https://issues.jboss.org/browse/AGJS-22
355	AGJS-17	what exactly would it be confused with? If you downloaded the sources and playing around with them, then maybe there would be some confusion since the folder that holds the different types is named adapter. But for a regular developer who is just using the lib, i think type is fine. I think the pay off of this change wouldn't equal the amount of time for deprecation, refactor, etc..	1	1	https://issues.jboss.org/browse/AGJS-17
356	AGJS-8	[~lucas.holmquist] Any update on this? Will it make it in for 1.1 or should we push it?	1	0	https://issues.jboss.org/browse/AGJS-8
357	AGJS-5	[~abstractj] Any objection to pushing this to 1.2.0? I would like to code freeze 1.1.0 so we can complete tests and docs and release 1.1.0 this week as planned.	1	1	https://issues.jboss.org/browse/AGJS-5
358	AGPUSH-2121	[~dimitraz] AGPUSH-2121, AGPUSH-2132 and AGPUSH-2136 look like one task to me. Can you take a look, sum them up in one task and close the others?	0	0	https://issues.jboss.org/browse/AGPUSH-2121

#	Requirement	Questions	Answered	Successfully Answered	URL
359	AGPUSH-2114	While we are still researching the CDI model, does it make sense to have any unit tests other than a simple send and receive? This is what I have currently: https://github.com/dimitraz/aerogear-unifiedpush-server/commit/ff104bb98c37876e877b48f0e86f64ef7d82d91d Based on Matthias' 'SimpleTest' class: https://github.com/matzew/kafka-cdi/blob/master/src/test/java/net/wessendorf/kafka/cdi/SimpleTest.java#L83-L88	1	1	https://issues.jboss.org/browse/AGPUSH-2114
360	AGPUSH-2111	[~mwessendorf], can you say when it is likely that we start working on this topic?	1	0	https://issues.jboss.org/browse/AGPUSH-2111
361	AGPUSH-2108	Hi [~ppatiern], can I message you on IRC if thats ok? What channels are you on?	1	0	https://issues.jboss.org/browse/AGPUSH-2108
362	AGPUSH-2105	[~mwessendorf] What do you think about [spring-kafka|https://projects.spring.io/spring-kafka/] ?	1	1	https://issues.jboss.org/browse/AGPUSH-2105
363	AGPUSH-2105	[~mwessendorf], I have explored the library. Do you have any concrete ideas for extensions you want to add? I think for now and taking into consideration what exactly we have in Kafka-UPS, the library is enough to be used. We can start using the library for integration.	1	1	https://issues.jboss.org/browse/AGPUSH-2105
364	AGPUSH-2103	After the first consumer is ready, what exactly did u mean by this task? I think the replacement of the counter so that it is based only on Kafka is already done in [[AGPUSH-2102]|https://issues.jboss.org/browse/AGPUSH-2102]?	1	1	https://issues.jboss.org/browse/AGPUSH-2103
365	AGPUSH-2102	The task is ready [github branch|https://github.com/polinankoleva/aerogear-unifiedpush-server/tree/installation_metrics_consumer]. I am waiting for your feedback to make a pull request. Is the implementation what did you expect?	1	1	https://issues.jboss.org/browse/AGPUSH-2102
366	AGPUSH-2101	I see that in the updateAnalytics method there's a query to the database and if pushMessageInformation is not null it means the app has been opened. I'm a bit confused as to how to do this without querying the DB. Is it correct to say that if pushMessageId is not null then this is another way to tell if the app has been opened? [~mziccardi] [~mwessendorf]	1	1	https://issues.jboss.org/browse/AGPUSH-2101
367	AGPUSH-2096	Hi [~mwessendorf] I saw this list of bugs and started resolve. I update this file to version 3.2. Among the changes, the main difference is the best way to use the environment variables, since the version 2 of docker-compose it is not a good pratice share the environment variables between services, so I created an .env file that has the variables needed to start the services. I do not know what do you prefer, send the pull request or attach the files here?	1	1	https://issues.jboss.org/browse/AGPUSH-2096
368	AGPUSH-2095	Hi [~mwessendorf], sorry for the time away. I've seen this task and created an endpoint /rest/sys/info/manifest. I resolved it, reading the manifest file and displaying its information. What do you think? The result looks like this: !manifest-details-maven.png|thumbnail!	0	0	https://issues.jboss.org/browse/AGPUSH-2095
369	AGPUSH-2091	Perhaps the new model object needs to be like: {code} FlatPushMessageInformation extends BaseModel { // like currently: @NotNull private String pushApplicationId; private String rawJsonMessage; private String ipAddress; private String clientIdentifier; private Date submitDate = new Date(); private Long appOpenCounter = 0L; private Date firstOpenDate; private Date lastOpenDate; // this is new: // once an error on a variant occurs, a error reason will be added to the list private List<VariantErrorStatus> errors; // this is than "removed" private Long appOpenCounter = 0L; private Date firstOpenDate; private Date lastOpenDate; private Integer servedVariants = 0; private Integer totalVariants = 0; private Set<VariantMetricInformation> variantInformations = new HashSet<>(); // end of not relevant section } package org.jboss.aerogear.unifiedpush.api; import java.util.List; /** * Created by matzew on 5/31/17. */ public class VariantErrorStatus { //@Id private String pushMessageVariantId = "push-job-id" + ":" + "variant-id"; private String errorReason; // the text we receive for the error public VariantErrorStatus() { } public void setPushMessageVariantId(final String bogusID) { pushMessageVariantId = bogusID; } public void applyAllTheThings(final String PMI_ID, final String varID) { // update me } } {code}	0	0	https://issues.jboss.org/browse/AGPUSH-2091
370	AGPUSH-2067	Is this one still in progress [~mwessendorf] ?	0	0	https://issues.jboss.org/browse/AGPUSH-2067
371	AGPUSH-2051	Is this one still in progress [~mziccardi]?	1	0	https://issues.jboss.org/browse/AGPUSH-2051
372	AGPUSH-2049	Is this one still in progress [~mziccardi]?	1	0	https://issues.jboss.org/browse/AGPUSH-2049
373	AGPUSH-2044	[~gunnar.morling] I think it's logical, that we should also replace this sub-select but, looking at the DAO, it basically also runs two query executions - should those be separated ?	0	0	https://issues.jboss.org/browse/AGPUSH-2044

#	Requirement	Questions	Answered	Successfully Answered	URL
374	AGPUSH-2043	IJ000616: Leak detected in pool: KeycloakDS (477ad5a2) (1487162949977): java.lang.Throwable: ALLOCATION LEAK at org.jboss.jca.core.connectionmanager.pool.mcp.LeakDumperManagedConnectionPool.getConnection(LeakDumperManagedConnectionPool.java:96) [ironjacamar-core-impl-1.0.37.Final-redhat-1.jar:1.0.37.Final-redhat-1] at org.jboss.jca.core.connectionmanager.pool.AbstractPool.getSimpleConnection(AbstractPool.java:457) [ironjacamar-core-impl-1.0.37.Final-redhat-1.jar:1.0.37.Final-redhat-1] at org.jboss.jca.core.connectionmanager.pool.AbstractPool.getConnection(AbstractPool.java:429) [ironjacamar-core-impl-1.0.37.Final-redhat-1.jar:1.0.37.Final-redhat-1] at org.jboss.jca.core.connectionmanager.AbstractConnectionManager.getManagedConnection(AbstractConnectionManager.java:344) [ironjacamar-core-impl-1.0.37.Final-redhat-1.jar:1.0.37.Final-redhat-1] at org.jboss.jca.core.connectionmanager.tx.TxConnectionManagerImpl.getManagedConnection(TxConnectionManagerImpl.java:367) [ironjacamar-core-impl-1.0.37.Final-redhat-1.jar:1.0.37.Final-redhat-1] at org.jboss.jca.core.connectionmanager.AbstractConnectionManager.allocateConnection(AbstractConnectionManager.java:499) [ironjacamar-core-impl-1.0.37.Final-redhat-1.jar:1.0.37.Final-redhat-1] at org.jboss.jca.adapters.jdbc.WrapperDataSource.getConnection(WrapperDataSource.java:143) [ironjacamar-jdbc-1.0.37.Final-redhat-1.jar:1.0.37.Final-redhat-1] at org.jboss.as.connector.subsystems.datasources.WildFlyDataSource.getConnection(WildFlyDataSource.java:69) [jboss-as-connector-7.5.11.Final-redhat-1.jar:7.5.11.Final-redhat-1] at org.hibernate.service.jdbc.connections.internal.DatasourceConnectionProviderImpl.getConnection(DatasourceConnectionProviderImpl.java:141) at org.hibernate.internal.AbstractSessionImpl$NonContextualJdbcConnectionAccess.obtainConnection(AbstractSessionImpl.java:301) at org.hibernate.engine.jdbc.internal.LogicalConnectionImpl.obtainConnection(LogicalConnectionImpl.java:214) at org.hibernate.engine.jdbc.internal.LogicalConnectionImpl.getConnection(LogicalConnectionImpl.java:157) at org.hibernate.engine.transaction.internal.jdbc.JdbcTransaction.doBegin(JdbcTransaction.java:67) at org.hibernate.engine.transaction.spi.AbstractTransactionImpl.begin(AbstractTransactionImpl.java:160) at org.hibernate.internal.SessionImpl.beginTransaction(SessionImpl.java:1387) at org.hibernate.ejb.TransactionImpl.begin(TransactionImpl.java:59) at org.keycloak.connections.jpa.JpaKeycloakTransaction.begin(JpaKeycloakTransaction.java:22) at org.keycloak.services.DefaultKeycloakTransactionManager.enlist(DefaultKeycloakTransactionManager.java:22) at org.keycloak.connections.jpa.DefaultJpaConnectionProviderFactory.create(DefaultJpaConnectionProviderFactory.java:39) at org.keycloak.connections.jpa.DefaultJpaConnectionProviderFactory.create(DefaultJpaConnectionProviderFactory.java:25) at org.keycloak.services.DefaultKeycloakSession.getProvider(DefaultKeycloakSession.java:99) at org.keycloak.models.jpa.JpaRealmProviderFactory.create(JpaRealmProviderFactory.java:34) at org.keycloak.models.jpa.JpaRealmProviderFactory.create(JpaRealmProviderFactory.java:16) at org.keycloak.services.DefaultKeycloakSession.getProvider(DefaultKeycloakSession.java:99) at	0	0	https://issues.jboss.org/browse/AGPUSH-2043
375	AGPUSH-2008	Where can the swagger docs be found? What is the url for the generated docs?	1	0	https://issues.jboss.org/browse/AGPUSH-2008
376	AGPUSH-1988	A working installation can be found at http://52.51.162.165:8080/ What we have done so far: * Installed HBASE on 52.51.162.165 * Installed Tomcat on 52.51.162.165 * Installed pinpoint web and pinpoint collector on the newly installed tomcat * Installed the pinpoint agent into the JBoss instance of the eteam cluster * Configured pinpoint * Made some test to check the pinpoint outcomes	0	0	https://issues.jboss.org/browse/AGPUSH-1988
377	AGPUSH-1987	[~mziccardi] found a command tool named [Artillery|https://artillery.io] that read a YAML file with some test configuration, for example "create 50 request every second during 1 hour". It's actually very easy and faster than creating a node/java client (it uses UPS's REST API). It also generates some reports. What do you think [~mwessendorf] EDIT: This was to be an example of 10k requests every second during 1 minute but it has gone too bad, UPS is completely down... EDIT 2: This is a test with 100 request every second during 1 minute: {noformat} Complete report @ 2016-12-20T12:24:13.046Z Scenarios launched: 6000 Scenarios completed: 4765 Requests completed: 4765 RPS sent: 48.03 Request latency: min: 487.8 max: 90958.1 median: 20366.8 p95: 33191 p99: 33426.1 Scenario duration: min: 526.5 max: 90997.2 median: 20385.2 p95: 33205.6 p99: 33459.9 Scenario counts: 0: 6000 (125.918%) Codes: 202: 4764 500: 1 Errors: ECONNRESET: 1235 {noformat}	1	1	https://issues.jboss.org/browse/AGPUSH-1987
378	AGPUSH-1979	An idea from [~jusanche]: {quote} If we switch to FCM for iOS too, we shouldn't need to ask for the certificates ourselves anymore, but document that they have to provide them on firebase console, and the same Sender ID and Server Key/token will work for both iOS and Android, so it might simplify the configuration on our part. {quote} I like this idea, and implementation could be quick, since at the end it's really just an "Android Variant". Perhaps we explicitly call it out FCM/APN mapper ? To avoid confusion w/ the pure Android Variant. Or.... we just give them example doc, how to upload the apple things, to use an Android variant. Heck, we should rename it "FCM variant", no ?	0	0	https://issues.jboss.org/browse/AGPUSH-1979
379	AGPUSH-1979	It seems that the current code can support FCM iOS. Do we have any APN Cert to test it on iOS? This document shows how to push notification to iOS by FCM: https://firebase.google.com/docs/cloud-messaging/ios/certs	0	0	https://issues.jboss.org/browse/AGPUSH-1979
380	AGPUSH-1978	[~jmesnil] Hi Jeff, does the above ring a bell, or is there something we can improve in our HornetQ usage, generally?	1	1	https://issues.jboss.org/browse/AGPUSH-1978
381	AGPUSH-1978	[~kpiwko][~mwessendorf] Looking at PagingManagerImpl and JournalImpl, the objects which got really huge are maps which contain in-flight transactions. Looks like that something started to fail transactions so HornetQ could not finish them. New and new transactions were started but they failed as well (most likely) and this map got really huge causing OOME (OutOfMemoryError). As HornetQ is storing those transactions in file journal as well then when it tries to load journal after restart then new OOME can be thrown as it's putting all "unfinished" transactions to this map again during initialization. By default there is 5 minute timeout for those transactions and then they're removed. But in this case it seems that memory got full before those transactions could be timeouted. At this moment this seems to be a symptom but not the cause. Any idea why the transactions are failing? What are the first error messages in server.log? [~mwessendorf] hornetq-jms.xml just defines which queue/topics should be deployed together with the deployment. I cannot find configuration of the server where this deployment is deployed. Could you help? Is there described what the deployment is doing?	1	1	https://issues.jboss.org/browse/AGPUSH-1978
382	AGPUSH-1978	[~mwessendorf] Thanks! It's pretty default configuration. Nothing interesting there. Could you share with me the heap dump so I can take the look? I would like to take a look at hornetq journal? By default it's located in $JBOSS_HOME/standalone/data/ directory. You can send me the the whole data directory. I'm interested in following directories messagingbindings, messagingjournal, messaginglargemessages, messagingpaging.	1	1	https://issues.jboss.org/browse/AGPUSH-1978
383	AGPUSH-1951	it was resolved as done, is there any produced artifact ?	1	0	https://issues.jboss.org/browse/AGPUSH-1951

#	Requirement	Questions	Answered	Successfully Answered	URL
384	AGPUSH-1951	[~mwessendorf] This was just to track the time spent on looking into using something like seneca or to use docker & openshift. I had a look at both and decided to use openshift. I went through it with my project supervisor and Leigh also put me on to a couple of graduates who used microservices last year in their college projects to discuss my approach. I plan to break up what I have into the following services: push configuration service, lead service & sales agent service. From studying the code and how it is currently it looks like I can separate out the functionality for each one to be its own service, and seems like the most obvious thing to do. I should have really ran it by you also to get your input, apologies. Does this plan sound ok with you? I have extracted the push configuration service and plan to do the others in the next sprint.	0	0	https://issues.jboss.org/browse/AGPUSH-1951
385	AGPUSH-1898	can you do an "ls -la" in the deployments folder and paste the result here ?	0	0	https://issues.jboss.org/browse/AGPUSH-1898
386	AGPUSH-1898	In your Apache log I see "Token audience doesn't match domain. Token issuer is http://URL/auth/realms/aerogear, but URL from configuration is https://URL/auth/realms/aerogear" , have you changed something around httpS ? Basically keycloak does not accept the token because it has not the same domain origin.	1	0	https://issues.jboss.org/browse/AGPUSH-1898
387	AGPUSH-1892	Is it possible, that you use a version that has FCM support already ? Becuase there the Android receivers are not correct :-/	1	0	https://issues.jboss.org/browse/AGPUSH-1892
388	AGPUSH-1605	You mean on the docker images? Or where ?	1	1	https://issues.jboss.org/browse/AGPUSH-1605
389	AGPUSH-1604	That's awesome! Mind giving it a shot? The snippets are located in here: * https://github.com/aerogear/aerogear-unifiedpush-server/tree/master/admin-ui/app/snippets/register-device While the related code for rendering is here: * https://github.com/aerogear/aerogear-unifiedpush-server/blob/master/admin-ui/app/scripts/directives.js * https://github.com/aerogear/aerogear-unifiedpush-server/blob/master/admin-ui/app/directives/ups-client-snippets.html Do you want this is a 1.1.3 release? Or do you think the 1.2.0 series is "good enough" ?	1	1	https://issues.jboss.org/browse/AGPUSH-1604
390	AGPUSH-1603	[~mwessendorf] Looking at the PHP sender, it looks like there aren't any code changes needed. The index.php page can have priority keys added to the message object manually and that will get passed correctly. Can you confirm?	1	0	https://issues.jboss.org/browse/AGPUSH-1603
391	AGPUSH-1600	[~mwessendorf] I just build everything from head and gave it a go. UPS accepted the token and sent a message no problem.	1	1	https://issues.jboss.org/browse/AGPUSH-1600
392	AGPUSH-1596	Matthias - what about OpenShit Cartridges ? Do they get updated with this issue resolved?	1	1	https://issues.jboss.org/browse/AGPUSH-1596
393	AGPUSH-1559	[~mwessendorf] This setup assumes the databases inside kubernets. How could I tweak the setup so that: UPS_DB and KC_DB are for instance inside openstack?	0	0	https://issues.jboss.org/browse/AGPUSH-1559
394	AGPUSH-1541	[INFO] ---------- [INFO] Building UnifiedPush Server (Admin UI) 1.1.1-SNAPSHOT [INFO] ---------- [INFO] ---------- [INFO] [INFO] --- maven-clean-plugin:2.5:clean (default-clean) @ unifiedpush-admin-ui --- [INFO] Deleting /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui/target [INFO] Deleting /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui (includes = [dist/**, .tmp/**, .build-tmp/**], excludes = []) [INFO] [INFO] --- maven-enforcer-plugin:1.3.1:enforce (enforce-java-version) @ unifiedpush-admin-ui --- [INFO] [INFO] --- maven-enforcer-plugin:1.3.1:enforce (enforce-maven-version) @ unifiedpush-admin-ui --- [INFO] [INFO] --- buildnumber-maven-plugin:1.2:create-timestamp (get-build-timestamp) @ unifiedpush-admin-ui --- [INFO] [INFO] --- buildnumber-maven-plugin:1.2:create (get-scm-revision) @ unifiedpush-admin-ui --- [INFO] Checking for local modifications: skipped. [INFO] Updating project files from SCM: skipped. [INFO] Executing: /bin/sh -c cd /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui && git rev-parse --verify HEAD [INFO] Working directory: /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui [INFO] Storing buildNumber: 5f5f167dc16d23e291f5ada71bd3bc7097d30b6d at timestamp: 1448955113979 [INFO] Executing: /bin/sh -c cd /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui && git rev-parse --verify HEAD [INFO] Working directory: /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui [INFO] Storing buildScmBranch: UNKNOWN [INFO] [INFO] --- frontend-maven-plugin:0.0.26:install-node-and-npm (install node and npm) @ unifiedpush-admin-ui --- [INFO] Found proxies: [] [INFO] Node v0.10.31 was installed, but we need version v0.12.8 [INFO] Installing node version v0.12.8 [INFO] Creating temporary directory /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui/node_tmp [INFO] Downloading Node.js from http://nodejs.org/dist/v0.12.8/node-v0.12.8-darwin-x64.tar.gz to /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui/node_tmp/node.tar.gz [INFO] No proxies configured [INFO] No proxy was configured, downloading directly [INFO] Extracting Node.js files in node_tmp [INFO] Unpacking /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui/node_tmp/node.tar.gz into /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui/node_tmp [INFO] Moving node binary to /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui/node/node [INFO] Deleting temporary directory /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui/node_tmp [INFO] Installed node locally. [INFO] Found NPM version 1.4.23 [INFO] [INFO] --- frontend-maven-plugin:0.0.26:npm (npm install) @ unifiedpush-admin-ui --- [INFO] Found proxies: [] [INFO] Running 'npm install --color=false' in /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui [INFO] [INFO] --- frontend-maven-plugin:0.0.26:grunt (grunt build) @ unifiedpush-admin-ui --- [INFO] Running 'grunt dist --no-color' in /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui [INFO] >> Local Npm module "grunt-cli" not found. Is it installed? [INFO] [INFO] Running "bower:install" (bower) task [INFO] >> Installed bower packages [INFO] >> Copied packages to /Users/matzew/Work/JBoss/aerogear-unifiedpush-server/admin-ui/app/bower-components [INFO] [INFO] Running "jshint:all" (jshint) task [INFO] >> 20 files lint free. [INFO] [INFO] Running "clean:server" (clean) task [INFO] [INFO] Running "concurrent:test" (concurrent) task [INFO] [INFO] Running "autoprefixer:dist" (autoprefixer) task [INFO] [INFO] Running "clean:dist" (clean) task [INFO] [INFO] Running "less:main" (less) task [INFO] File .tmp/styles/main.css created: 0 B → 269.96 kB [INFO] [INFO] Running "less:account" (less) task [INFO] File .tmp/styles/account.css created: 0 B → 261.81 kB [INFO] [INFO] Running "less:login" (less) task [INFO] File .tmp/styles/login.css created: 0 B → 259.04 kB [INFO] [INFO] Running "copy:fonts" (copy) task [INFO] Created 3 directories, copied 73 files [INFO] [INFO] Running "copy:misc" (copy) task [INFO] Copied 1 files [INFO] [INFO] Running "useminPrepare:html" (useminPrepare) task	0	0	https://issues.jboss.org/browse/AGPUSH-1541
395	AGPUSH-1538	Alexandre, is this issue also present with the 1.1.0.Final release? We did put a bit of work into the migrator	1	0	https://issues.jboss.org/browse/AGPUSH-1538
396	AGPUSH-1538	Nop, I didn't test it. I found it on production environment (1.0.3) . I did not install 1.1.0 because you found some critical bugs on the release stage. Is it safe to move to 1.1.0 now? Can I migrate from 1.0.3 to 1.1.0?	1	0	https://issues.jboss.org/browse/AGPUSH-1538
397	AGPUSH-1538	we did restage the 1.1.0.Final, after those bug fixes But also now in the middle of 1.1.1 :-) So, that should be out in a few weeks too	0	0	https://issues.jboss.org/browse/AGPUSH-1538
398	AGPUSH-1536	From the mailing list : "Can I suggest that we have a standard name for what this key is? The docs refer to 'Client ID', 'Server Key' and 'Google Cloud Messaging Key' "In the Wizard after you create a PushApplication, click the Add Variant button and fill out the Android options. You will want to use the Client ID and Project Number from the Google API Console in their appropriate fields:" "On the last screen we are finally get to see the actual value of the generated Server Key, which we will use later:" 'Google Cloud Messaging Key' is on the dialogue box that pops up when you create the Android Variant. I think that 'Client ID', 'Server Key' and 'Google Cloud Messaging Key' are actually all the same key, the one created in this attachment Are they all the same k,ey? "	0	0	https://issues.jboss.org/browse/AGPUSH-1536
399	AGPUSH-1536	[~mwessendorf] Makes sense. Does UPS already support FCM? Or is FCM backwards compatible and there was no change necessary?	1	0	https://issues.jboss.org/browse/AGPUSH-1536

#	Requirement	Questions	Answered	Successfully Answered	URL
400	AGPUSH-1528	service jboss.undertow.deployment.default-server.default-host./auth: org.jboss.msc.service.StartException in service jboss.undertow.deployment.default-server.default-host./auth: java.lang.RuntimeException: Failed to construct public org.jboss.aerogear.unifiedpush.keycloak.UpsKeycloakApplication(javax.servlet.ServletContext,org.jboss.resteasy.core.Dispatcher) at org.wildfly.extension.undertow.deployment.UndertowDeploymentService$1.run(UndertowDeploymentService.java:85) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:511) at java.util.concurrent.FutureTask.run(FutureTask.java:266) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) at org.jboss.threads.JBossThread.run(JBossThread.java:320) Caused by: java.lang.RuntimeException: Failed to construct public org.jboss.aerogear.unifiedpush.keycloak.UpsKeycloakApplication(javax.servlet.ServletContext,org.jboss.resteasy.core.Dispatcher) at org.jboss.resteasy.core.ConstructorInjectorImpl.construct(ConstructorInjectorImpl.java:160) at org.jboss.resteasy.spi.ResteasyProviderFactory.createProviderInstance(ResteasyProviderFactory.java:2211) at org.jboss.resteasy.spi.ResteasyDeployment.createApplication(ResteasyDeployment.java:295) at org.jboss.resteasy.spi.ResteasyDeployment.start(ResteasyDeployment.java:236) at org.jboss.resteasy.plugins.server.servlet.ServletContainerDispatcher.init(ServletContainerDispatcher.java:112) at org.jboss.resteasy.plugins.server.servlet.HttpServletDispatcher.init(HttpServletDispatcher.java:36) at io.undertow.servlet.core.LifecyleInterceptorInvocation.proceed(LifecyleInterceptorInvocation.java:117) at org.wildfly.extension.undertow.security.RunAsLifecycleInterceptor.init(RunAsLifecycleInterceptor.java:78) at io.undertow.servlet.core.LifecyleInterceptorInvocation.proceed(LifecyleInterceptorInvocation.java:103) at io.undertow.servlet.core.ManagedServlet$DefaultInstanceStrategy.start(ManagedServlet.java:230) at io.undertow.servlet.core.ManagedServlet.createServlet(ManagedServlet.java:131) at io.undertow.servlet.core.DeploymentManagerImpl.start(DeploymentManagerImpl.java:511) at org.wildfly.extension.undertow.deployment.UndertowDeploymentService.startContext(UndertowDeploymentService.java:101) at org.wildfly.extension.undertow.deployment.UndertowDeploymentService$1.run(UndertowDeploymentService.java:82) ... 6 more Caused by: java.lang.NoClassDefFoundError: org/apache/http/conn/socket/LayeredConnectionSocketFactory at org.keycloak.connections.httpclient.DefaultHttpClientFactory.init(DefaultHttpClientFactory.java:104) at org.keycloak.services.DefaultKeycloakSessionFactory.init(DefaultKeycloakSessionFactory.java:82) at org.keycloak.services.resources.KeycloakApplication.createSessionFactory(KeycloakApplication.java:162) at	1	1	https://issues.jboss.org/browse/AGPUSH-1528
401	AGPUSH-1522	We could do that depending on what the UPS is actually returning for it's response. [~evanshortiss] what version of the UPS are you using. and what node-sender version are you using.	1	1	https://issues.jboss.org/browse/AGPUSH-1522
402	AGPUSH-1522	Hi Guys, I missed this update. Here's the version: AeroGear UnifiedPush Server 1.0.3 and WildFly	1	0	https://issues.jboss.org/browse/AGPUSH-1522
403	AGPUSH-1522	I see this has been resolved. Can you specify the new patched node module version and AeroGear server version that are compatible and will return rich error responses? Sure, will use a mailing list in the future. Can you specify which list specifically?	0	0	https://issues.jboss.org/browse/AGPUSH-1522
404	AGPUSH-1522	Thanks, I still don't feel that the original issue has been resolved though; it can't be without updating the node module to take the AG server response and feeding that as JSON Object or String to the callback along with the status code. Or is there also a separate issue open for that besides this one?	0	0	https://issues.jboss.org/browse/AGPUSH-1522
405	AGPUSH-1496	I'm been looking into this and trying to figure out why this is happening. The version of UPS in [bbe1c0e#diff-23e531d12489411024ea5c379be1d178R70|https://github.com/aerogear/aerogear-unifiedpush-server/commit/bbe1c0e80a4a5295a3c7cefe66f5e8cc07da66be#diff-23e531d12489411024ea5c379be1d178R70] is 1.1.0.Final-SNAPSHOT: {noformat} git show bbe1c0e:pom.xml {noformat} With the addition of the migrator we now have the requirement to create the data base schema using the [migrator|https://aerogear.org/docs/unifiedpush/ups_userguide/index/#schema] and not have it automatically created by hibernate upon application startup. {noformat} ./bin/ups-migrator --logLevel=debug update {noformat} This would populate the database succesfully with a clean database. Now, if we take checkout the commit above: {noformat} git co -b migrator-qe-issue bbe1c0e {noformat} Build and deploy against to the server (with a clean mysql database) will create the data base schema, but an incorrect one. Running the migrator against this database will error with the above reported error. At the time of the bbe1c0e commit ,there was no migrator yet so letting hibernate create the database schema was the way to do it. Now that the migrator was introduced it should be used to generate the schema. One thought that I had is that you mentioned in the thread that you run an integration test with _runMigrator=false_. Depending on the version of UPS used it might be possible that it is creating the incorrect database schema {noformat} $ ./gradlew --stacktrace :spacelift:test -Pwildfly8UpsFromSource -Pdatasource=mysql -PrunMigrator=false {noformat} I tried running this on my local machine but it fail but it looks like this is using _1.1.0-beta.3_ which does not contain the migrator. So this could indeed be the case that the schema is created upon application deployment instead of by the migrator. I'm also wondering about the version used to test. Should we not be using a later version to testa against?	1	0	https://issues.jboss.org/browse/AGPUSH-1496
406	AGPUSH-1493	We could still make topics work like the old categories, when we apply the alias 'filter' in UPS and then remove the devices that don't have specific alias. Then use the topic to send the message, that way we have the same behavior we used to have with the 'old' categories. How does that sound?	0	0	https://issues.jboss.org/browse/AGPUSH-1493
407	AGPUSH-1490	[~fturizo] one potential reason behind this could be if the time on your desktop is out of sync with the Keycloak server. Is the Keycloak server on a separate machine? Can you check the date on your machine and the server and let me know what the difference is?	1	1	https://issues.jboss.org/browse/AGPUSH-1490
408	AGPUSH-1490	Mmm, now that you mention it, the wildfly server and the UPS application are deployed in an AWS server located in the US. West Coat (GMT - 4) and i'm testing in Bogota (GMT - 5). Could this be a reason for this error?	0	0	https://issues.jboss.org/browse/AGPUSH-1490
409	AGPUSH-1490	Yes, I think so - could you try to set your machine to the same time zone and make sure the time is in sync, then confirm if that fixes the issue? If so I'll figure out some fix for keycloak.js to deal with this.	1	1	https://issues.jboss.org/browse/AGPUSH-1490

#	Requirement	Questions	Answered	Successfully Answered	URL
410	AGPUSH-1489	The problem is that the protocolType for the windows variant is never set before we go to editing https://github.com/aerogear/aerogear-unifiedpush-server/blob/6ad62a16bbbb9f838acc85c7c3511a0097c90b9f/admin-ui/app/scripts/services/variantModal.js#L70 Looks like it's set only when we create a variant, but this info is lost after reload. [~lfryc] Could you take a look ?	0	0	https://issues.jboss.org/browse/AGPUSH-1489
411	AGPUSH-1486	[~abstractj] [~sebastienblanc] does this somewhat security related ticket make sense for 1.2.0 ? Or should we do it later ?	0	0	https://issues.jboss.org/browse/AGPUSH-1486
412	AGPUSH-1470	[~abstractj] perhaps this could be something already for 1.2.0? Or would 2.0.0 be a more relaxed time-frame ?	0	0	https://issues.jboss.org/browse/AGPUSH-1470
413	AGPUSH-1458	What is the issue here? We could also use yaml instead of xml.	0	0	https://issues.jboss.org/browse/AGPUSH-1458
414	AGPUSH-1452	Just please consider a need to use alternative delivery network capabilities in underlying impl - I'd much rather implement a http/2 server then http/2 proxy in the future. Seems that pushy support that which is a significant plus from QE perspective! https://github.com/relayrides/pushy/blob/master/src/main/java/com/relayrides/pushy/apns/ApnsEnvironment.java	1	0	https://issues.jboss.org/browse/AGPUSH-1452
415	AGPUSH-1422	should we move to .Final ?	1	1	https://issues.jboss.org/browse/AGPUSH-1422
416	AGPUSH-1422	[~sebastienblanc] do you need some more input from [~kpiwko] or [~smikloso] on this ? Or should we move to 1.1.x ?	1	1	https://issues.jboss.org/browse/AGPUSH-1422
417	AGPUSH-1410	What's the status here?	1	1	https://issues.jboss.org/browse/AGPUSH-1410
418	AGPUSH-1409	What's the status here?	1	1	https://issues.jboss.org/browse/AGPUSH-1409
419	AGPUSH-1393	Should we move to .Final ?	1	1	https://issues.jboss.org/browse/AGPUSH-1393
420	AGPUSH-1392	[~sebastienblanc] [~agalante] Guys, I'd like to know your opinion how this feature should be implemented. In the old UI, the "generate push-config.json" button is little bit independent from the rest of samples. Andres, to get you into a picture, {{push-config.json}} is an alternative mechanism how to configure Cordova based applications. You need to store the JSON file such as [1] side by side your Cordova project sources. It will be loaded automatically when you don't specify configuration in your code. But, the problem is that {{push-config.json}} contains configuration for several platforms, not just one (see the sample [1]). In the old UI we've decided to have a separated button that opened a dialog for variants selection for which you wanted to generate the {{push-config.json}}. Then you was able to download the configuration file or copy its contents. You can see full discussion about this feature here: http://aerogear-dev.1069024.n5.nabble.com/aerogear-dev-Cordova-push-config-json-generated-in-UPS-Console-multiple-variants-tp9799p9824.html My question is whether this is something we want to re-do in a new UI or something that should be re-designed. Could you possibly work this feature in the mockup? [1] https://github.com/jboss-mobile/unified-push-helloworld/blob/master/cordova/www/push-config.json	1	1	https://issues.jboss.org/browse/AGPUSH-1392
421	AGPUSH-1373	Want me to help on the KC related Screenshots ? Once we got 1.3 of KC in ?	1	1	https://issues.jboss.org/browse/AGPUSH-1373
422	AGPUSH-1361	[~lfryc] do we need this? Or can we now close it ?	0	0	https://issues.jboss.org/browse/AGPUSH-1361
423	AGPUSH-1338	[~agalante] we check that on the server once the user hits Create button. Is there strong reason to pre-validate that even before? It's technically feasible, but I would personally avoid that.	1	1	https://issues.jboss.org/browse/AGPUSH-1338
424	AGPUSH-1320	> Just wondering, is that a constant problem, or did it occur just on those two messages? I don't think any new messages could be sent at that time, so I restarted WildFly today. Some messages were sent correctly a few months ago, then some time passed when I don't think any messages were sent, and then two messages were attempted to be sent in the two past days, and failed. So the WildFly server has been running for a very long time. >Another question is, did a reboot of the WildFly server fix the problem? According to the stack, it's some problem related to the transaction execution. I have restarted the WildFly server now, and will see if that fixes the problem. >Would you be willing to test if the problem occurs on 1.0.3 as well? Yes, I can do that. Thanks, Andreas R.	1	1	https://issues.jboss.org/browse/AGPUSH-1320
425	AGPUSH-1320	submitted to PushNetworks for further processing ^[[0m^[[0m15:12:26,776 INFO [JPAPushMessageInformationDao] (EJB default - 2) Processing send request with '[alert=text goes here., data={url=http://www.xkcd.no/nyheter/#push}, criteria=[aliases=null, deviceTypes=null, categories=null, variants=null], sound=default, action-category=null, badge=-1, time-to-live=-1, simplePush=null, content-available=false]' payload ^[[0m^[[33m15:17:26,772 WARN [com.arjuna.ats.arjuna] (Transaction Reaper) ARJUNA012117: TransactionReaper::check timeout for TX 0:ffffc0a8655e:560b48d0:550a97df:6032 in state RUN ^[[0m^[[33m15:17:26,775 WARN [com.arjuna.ats.arjuna] (Transaction Reaper Worker 0) ARJUNA012095: Abort of action id 0:ffffc0a8655e:560b48d0:550a97df:6032 invoked while multiple threads active within it. ^[[0m^[[33m15:17:26,775 WARN [com.arjuna.ats.arjuna] (Transaction Reaper Worker 0) ARJUNA012108: CheckedAction::check - atomic action 0:ffffc0a8655e:560b48d0:550a97df:6032 aborting with 1 threads active! ^[[0m^[[33m15:17:26,788 WARN [org.hibernate.engine.jdbc.spi.SqlExceptionHelper] (EJB default - 2) SQL Error: 0, SQLState: null ^[[0m^[[31m15:17:26,789 ERROR [org.hibernate.engine.jdbc.spi.SqlExceptionHelper] (EJB default - 2) javax.resource.ResourceException: IJ000460: Error checking for a transaction ^[[0m^[[33m15:17:26,791 WARNING [ClientInstallationServiceImpl] (EJB default - 2) Error on 'ios' delivery ^[[0m^[[33m15:17:26,792 WARN [org.hibernate.jpa.spi.AbstractEntityManagerImpl] (EJB default - 2) HHH000326: Cannot join transaction: do not override hibernate.transaction.factory_class ^[[0m^[[33m15:17:27,220 WARN [org.hibernate.engine.transaction.synchronization.internal.SynchronizationCallbackCoordinatorTrackingImpl] (Transaction Reaper Worker 0) HHH000451: Transaction afterCompletion called by a background thread; delaying afterCompletion processing until the original thread can handle it. [status=4] ^[[0m^[[33m15:17:27,220 WARN [com.arjuna.ats.arjuna] (Transaction Reaper Worker 0) ARJUNA012121: TransactionReaper::doCancellations worker Thread[Transaction Reaper Worker 0,5,main] successfully canceled TX 0:ffffc0a8655e:560b48d0:550a97df:6032 ^[[0m^[[31m15:17:27,221 ERROR [org.jboss.as.ejb3.invocation] (EJB default - 2) JBAS014134: EJB Invocation failed on component SenderServiceImpl for method public abstract void org.jboss.aerogear.unifiedpush.message.SenderService.send(org.jboss.aerogear.unifiedpush.api.PushApplication,org.jboss.aerogear.unifiedpush.message.UnifiedPushMessage): javax.ejb.EJBException: JBAS014586: Transaction 'TransactionImple < ac, BasicAction: 0:ffffc0a8655e:560b48d0:550a97df:6032 status: ActionStatus.ABORTED >' is in unexpected state (STATUS_ROLLING_BACK) at org.jboss.as.ejb3.tx.CMTTxInterceptor.endTransaction(CMTTxInterceptor.java:116) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.as.ejb3.tx.CMTTxInterceptor.invokeInOurTx(CMTTxInterceptor.java:277) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.as.ejb3.tx.CMTTxInterceptor.required(CMTTxInterceptor.java:340) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.as.ejb3.tx.CMTTxInterceptor.processInvocation(CMTTxInterceptor.java:239) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:309) at org.jboss.as.ejb3.component.interceptors.CurrentInvocationContextInterceptor.processInvocation(CurrentInvocationContextInterceptor.java:41) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.as.ejb3.component.invocationmetrics.WaitTimeInterceptor.processInvocation(WaitTimeInterceptor.java:43) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:309) at	1	1	https://issues.jboss.org/browse/AGPUSH-1320
426	AGPUSH-1320	Oh, that's bad. Did migrating to 1.0.3 help?	1	1	https://issues.jboss.org/browse/AGPUSH-1320

#	Requirement	Questions	Answered	Successfully Answered	URL
427	AGPUSH-1320	Just to make sure it isn't related to the payload, could you try to send a push message using the Admin console ? (it will just use the 'alert' field)	0	0	https://issues.jboss.org/browse/AGPUSH-1320
428	AGPUSH-1320	the Admin console worked correctly to a test application. I have not tested using the production application, since it is in production. This is a typical log message where the sending fails, which also includes the payload: ^[[0m^[[0m19:53:38,823 INFO [PushNotificationSenderEndpoint] (default task-10) Message submitted to PushNetworks for further processing ^[[0m^[[0m19:53:38,826 INFO [JPAPushMessageInformationDao] (EJB default - 3) Processing send request with '[alert=Rosenborg tapte treningskampen mot Viking. Resultatet ble 4-2., data={url=http://www.foobar-domain.no/100Sport/fotball/forsesong/article507344.snd}, criteria=[aliases=null, deviceTypes=null, categories=null, variants=null], sound=default, action-category=null, badge=-1, time-to-live=-1, simplePush=null, content-available=false]' payload ^[[0m^[[33m19:58:38,825 WARN [com.arjuna.ats.arjuna] (Transaction Reaper) ARJUNA012117: TransactionReaper::check timeout for TX 0:ffffc0a8655e:460e760:550bfe11:53cf7 in state RUN ^[[0m^[[33m19:58:38,827 WARN [com.arjuna.ats.arjuna] (Transaction Reaper Worker 0) ARJUNA012095: Abort of action id 0:ffffc0a8af55e:460e760:550bfe11:53cf7 invoked while multiple threads active within it. ^[[0m^[[33m19:58:38,827 WARN [com.arjuna.ats.arjuna] (Transaction Reaper Worker 0) ARJUNA012108: CheckedAction::check - atomic action 0:ffffc0a8655e:460e760:550bfe11:53cf7 aborting with 1 threads active! ^[[0m^[[33m19:58:38,829 WARNING [ClientInstallationServiceImpl] (EJB default - 3) Error on 'ios' delivery ^[[0m^[[33m19:58:38,832 WARN [org.hibernate.jpa.spi.AbstractEntityManagerImpl] (EJB default - 3) HHH000326: Cannot join transaction: do not override hibernate.transaction.factory_class ^[[0m^[[33m19:58:39,209 WARN [org.hibernate.engine.transaction.synchronization.internal.SynchronizationCallbackCoordinatorTrackingImpl] (Transaction Reaper Worker 0) HHH000451: Transaction afterCompletion called by a background thread; delaying afterCompletion processing until the original thread can handle it. [status=4] ^[[0m^[[33m19:58:39,210 WARN [com.arjuna.ats.arjuna] (Transaction Reaper Worker 0) ARJUNA012121: TransactionReaper::doCancellations worker Thread[Transaction Reaper Worker 0,5,main] successfully canceled TX 0:ffffc0a8655e:460e760:550bfe11:53cf7 ^[[0m^[[31m19:58:39,210 ERROR [org.jboss.as.ejb3.invocation] (EJB default - 3) JBAS014134: EJB Invocation failed on component SenderServiceImpl for method public abstract void org.jboss.aerogear.unifiedpush.message.SenderService.send(org.jboss.aerogear.unifiedpush.api.PushApplication,org.jboss.aerogear.unifiedpush.message.UnifiedPushMessage): javax.ejb.EJBException: JBAS014586: Transaction 'TransactionImple < ac, BasicAction: 0:ffffc0a8655e:460e760:550bfe11:53cf7 status: ActionStatus.ABORTED >' is in unexpected state (STATUS_ROLLING_BACK) at org.jboss.as.ejb3.tx.CMTTxInterceptor.endTransaction(CMTTxInterceptor.java:116) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.as.ejb3.tx.CMTTxInterceptor.invokeInOurTx(CMTTxInterceptor.java:277) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.as.ejb3.tx.CMTTxInterceptor.required(CMTTxInterceptor.java:340) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.as.ejb3.tx.CMTTxInterceptor.processInvocation(CMTTxInterceptor.java:239) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:309) at org.jboss.as.ejb3.component.interceptors.CurrentInvocationContextInterceptor.processInvocation(CurrentInvocationContextInterceptor.java:41) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:309) at org.jboss.as.ejb3.component.invocationmetrics.WaitTimeInterceptor.processInvocation(WaitTimeInterceptor.java:43) [wildfly-ejb3-8.1.0.Final.jar:8.1.0.Final] at	1	1	https://issues.jboss.org/browse/AGPUSH-1320
429	AGPUSH-1320	Andreas, I'm trying to identify which Push Network your application uses, could you list types of variants listed in the application that fails? e.g. GCM, APN.. I guess some sender implementation throws error and does not catch/log it appropriately.	1	1	https://issues.jboss.org/browse/AGPUSH-1320
430	AGPUSH-1320	Andreas, do you have a lot of installations in your production env? I guess the EJB could potentially time out while sending the messages to push networks (depends on your configuration of the wildfly container).	1	1	https://issues.jboss.org/browse/AGPUSH-1320
431	AGPUSH-1320	> Andreas, I'm trying to identify which Push Network your application uses, could you list types of variants listed in the application that fails? e.g. GCM, APN.. The application in Aerogear push has one variant, which is APN (iOS). > You could also try to fetch a debug log so we could get more insight - here is how you can switch wildfly to debug mode for a little while in production and then turn it back while this "defective" push message is being sent: Thanks, will try that. > Andreas, do you have a lot of installations in your production env? I guess the EJB could potentially time out while sending the messages to push networks (depends on your configuration of the wildfly container). There are many thousand installations in the production env.	1	1	https://issues.jboss.org/browse/AGPUSH-1320
432	AGPUSH-1320	> Andreas, do you have a lot of installations in your production env? I guess the EJB could potentially time out while sending the messages to push networks (depends on your configuration of the wildfly container). This is a very likely cause of the problem. Looking at the timestamps in the logs I have attached previously, the sending was aborted after 5 minutes. I logged into the WildFly admin console, and found that in "Container" -> "Transactions" the default timeout was set to 300s (5 minutes). I have now increased this to 7200s, and hope that this will solve the problem the next time a push message is sent. This would also explain why sending worked correctly some months ago, and now the number of device installations has increased significantly. Thanks for the help everyone, I am optimistic now!	1	1	https://issues.jboss.org/browse/AGPUSH-1320
433	AGPUSH-1320	This is a push message service for a iOS news app, so it is important for the users that the push-message is sent as quickly as possible. So the users have said that even five minutes is too long. Is there anything I can do to improve sending performance?	1	1	https://issues.jboss.org/browse/AGPUSH-1320
434	AGPUSH-1317	[~agalante] [~sebastienblanc] what's the status here? are the mocks ready ?	1	1	https://issues.jboss.org/browse/AGPUSH-1317
435	AGPUSH-1311	Here's a quick overview of the cartridge changes that will be needed: # The binaries for the migration tool will need to be added to the cartridge repository (e.g., in versions/1.0.3/usr/lib/migration/). # The binaries for the migration tool will need to be installed on the gear (e.g., in $OPENSHIFT_AEROGEAR_PUSH_DIR/usr/lib/migration). This will need to be done in the cartridge's [install |https://github.com/aerogear/openshift-origin-cartridge-aerogear-push/blob/master/bin/install#L23] script. # The migration tool will need to be invoked from an action hook in the cartridge's [control | https://github.com/aerogear/openshift-origin-cartridge-aerogear-push/blob/master/bin/control] script. Once the above changes have been made, they can be tested using the following commands: {code} rhc snapshot save <APP> rhc app delete <APP> rhc create-app --no-git <APP> <CARTRIDGE_URL> rhc snapshot restore <APP> -f /PATH/TO/APP.tar.gz {code}	1	1	https://issues.jboss.org/browse/AGPUSH-1311
436	AGPUSH-1311	[~fjuma] with the PR for AGPUSH-1310, most of this is already implemented, on version 1.0.3. What is needed to actually give users with previous versions a chance to upgrade? Or is the best option, to offer migration from 1.0.3 and later? (e.g. also enabling git hooks to manually upload WAR files?)	1	1	https://issues.jboss.org/browse/AGPUSH-1311
437	AGPUSH-1306	Looking closer to GCM, looks like they are no "GCM" specific keys for the Push Payload, could [~danielpassos] and [~edewit] confirm this ? If so, I think we can close this one.	1	1	https://issues.jboss.org/browse/AGPUSH-1306
438	AGPUSH-1291	[~qmx] [~sebastienblanc] Not sure I really get this ticket. Isn't applying the changes implicitly part of the merge exercise? If yes, let's close this as out dated	0	0	https://issues.jboss.org/browse/AGPUSH-1291

#	Requirement	Questions	Answered	Successfully Answered	URL
439	AGPUSH-1274	reopening, since it has not been merged, and therefore is not done	1	1	https://issues.jboss.org/browse/AGPUSH-1274
440	AGPUSH-1260	are we going to camelCase those? so titleLocKey, titleLocArgs? or how does the UPS deal with them?	0	0	https://issues.jboss.org/browse/AGPUSH-1260
441	AGPUSH-1259	[~sebastienblanc] oh yeah. Can you do a a sub-task for this ?	0	0	https://issues.jboss.org/browse/AGPUSH-1259
442	AGPUSH-1248	[~agalante] I would like to have your opinion on this one, what is the best regarding user experience ?	1	1	https://issues.jboss.org/browse/AGPUSH-1248
443	AGPUSH-1243	What's the status here? Looks like the PR is no longer open	1	1	https://issues.jboss.org/browse/AGPUSH-1243
444	AGPUSH-1240	[~mwessendorf] on the UI, we no longer expose descriptions, should we remove that also on database layer?	1	0	https://issues.jboss.org/browse/AGPUSH-1240
445	AGPUSH-1223	[~kpiwko] It looks like those tests are for the Unified Push server and not the SimplePush server which is the target component for this particular jira. Was the sub-task meant for a different jira perhaps?	1	0	https://issues.jboss.org/browse/AGPUSH-1223
446	AGPUSH-1222	should we move to 0.11.0 (next release) ?	0	0	https://issues.jboss.org/browse/AGPUSH-1222
447	AGPUSH-1221	Is that still on the bubble? I'd vote nope, to stay in line w/ Mozilla's SimplePush specification	0	0	https://issues.jboss.org/browse/AGPUSH-1221
448	AGPUSH-1220	should we move to 0.10.0 (next release) ?	0	0	https://issues.jboss.org/browse/AGPUSH-1220
449	AGPUSH-1220	should we move to 0.11.0 (next release) ?	0	0	https://issues.jboss.org/browse/AGPUSH-1220
450	AGPUSH-1216	Have this working on this [branch|https://github.com/danbev/aerogear-simplepush-server/tree/new-netty-sockjs]. These is one issue that I discovered regarding [CORS configuration|https://github.com/netty/netty/pull/2427] and details can be found in this [commit|https://github.com/danbev/aerogear-simplepush-server/commit/b23ada1bdccb32f339ff3d544838361f3d71a6b8].	0	0	https://issues.jboss.org/browse/AGPUSH-1216
451	AGPUSH-1210	should we move to 0.10.0 (next release) ?	1	0	https://issues.jboss.org/browse/AGPUSH-1210
452	AGPUSH-1205	The PR for this ticket has been closed; Should we close/reject this ticket as well ?	1	0	https://issues.jboss.org/browse/AGPUSH-1205
453	AGPUSH-1189	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	0	0	https://issues.jboss.org/browse/AGPUSH-1189
454	AGPUSH-1149	[~kpiwko] [~abstractj] for this, the default should stay, as is. and I think it is reasonable to just enable this with a system property, if possible. In case the container is started with -Dorg.aerogear.keycloak.direct.access(=true), the direct access should be there. Using system properties minimizes the risk of users accidently enabling this, by messing around inside of cfg files. What's your thoughts?	1	1	https://issues.jboss.org/browse/AGPUSH-1149
455	AGPUSH-1146	[~edewit] does that mean we need to change the data model ? Or should we just ignore the ERROR?	1	1	https://issues.jboss.org/browse/AGPUSH-1146
456	AGPUSH-1130	I think the questions are: * What's next on this? * Do we support a migration from 1.0.0.final to any new release of the 1.0.x series * Does a fix for this ticket need to be included in the 1.0.3 release ? [~edewit] [~qmx] any comments?	0	0	https://issues.jboss.org/browse/AGPUSH-1130
457	AGPUSH-1129	[~edewit] great work. Can I ask for a small change list screen? - Font size for the text should be 16px, color black 87% (or #212121) and line height 20px - I guess the header is standard windows phone, right? - Dividers lines, if they are not windows phone standard please make them black 12% (or #e0e0e0) - Make padding between the screen and the starting point of the text 16px. - Make padding between the line and the starting point of the text 20px. If they are not easy changes, forget about them :)	1	1	https://issues.jboss.org/browse/AGPUSH-1129
458	AGPUSH-1127	[~mwessendorf], yes, absolutely. What's the easiest way for me to reproduce this and issue these statements? Or should we do a short hangout some time this week to discuss it?	1	1	https://issues.jboss.org/browse/AGPUSH-1127
459	AGPUSH-1126	[~edewit] can we close this? Or is this still something 'reasonable' ?	1	1	https://issues.jboss.org/browse/AGPUSH-1126
460	AGPUSH-1124	is this for both? MPNS and WNS ?	0	0	https://issues.jboss.org/browse/AGPUSH-1124
461	AGPUSH-1111	[~abstractj] hey man, do you have an idea how to get this implemented? We could hangout and discuss it /wrt connection to admin console (AGPUSH-1112).	0	0	https://issues.jboss.org/browse/AGPUSH-1111
462	AGPUSH-1110	[~lfryc] for the all new UI - do we have an epic already for this ?	1	1	https://issues.jboss.org/browse/AGPUSH-1110
463	AGPUSH-1110	[~mwessendorf] [~lfryc] New UI is not on Jira yet. Thanks for reminding me of this screen. I'll add it to the new UI. Is there any other "hidden" screen I should know about?	1	1	https://issues.jboss.org/browse/AGPUSH-1110
464	AGPUSH-1108	You can see it in action here: http://andresgalante.com/ups-console/app-detail-variants.html What do you think?	1	1	https://issues.jboss.org/browse/AGPUSH-1108
465	AGPUSH-1108	[~lfryc] this can be moved to ups-1.1.x ?	1	1	https://issues.jboss.org/browse/AGPUSH-1108
466	AGPUSH-1075	what's the status here?	1	1	https://issues.jboss.org/browse/AGPUSH-1075
467	AGPUSH-1075	[~edewit] ok, do we have a JIRA for the actual migration (https://github.com/aerogear/aerogear-unifiedpush-server/pull/448). If not, can you create one for 1.0.3 ?	0	0	https://issues.jboss.org/browse/AGPUSH-1075
468	AGPUSH-1039	[~sblanc] [~mwessendorf] [~lfryc] I do think there is a need for it. Should this be just a search box? or also sorting options?	1	0	https://issues.jboss.org/browse/AGPUSH-1039
469	AGPUSH-1039	I don't think sorting is required, if yes I would ask which categories it should be based on (username (visible just to admins), # of variants, # of notifications delivered, # of devices registered)?	1	1	https://issues.jboss.org/browse/AGPUSH-1039
470	AGPUSH-1032	Hey [~gcardoso] would you know how to implement this? I can't find any samples in https://www.patternfly.org/wikis/patterns/navigation/	1	1	https://issues.jboss.org/browse/AGPUSH-1032
471	AGPUSH-1032	There is a widget for the horizontal navigation, but no sample for the vertical navigation. What exactly do you need?	1	1	https://issues.jboss.org/browse/AGPUSH-1032
472	AGPUSH-1032	The Admin link opens the User Management / Log Out dropdown directly, but in responsive mode, you see the hamburger icon, which, when you click on, reveals Admin link that opens the User Management / Logout. Would you consider that a critical fix? As I see there are changes in bootstrap/patternfly needed to make this work.	1	1	https://issues.jboss.org/browse/AGPUSH-1032
473	AGPUSH-1000	[~mhusnain] any hint on how this is wrong ? (e.g. what path)	0	0	https://issues.jboss.org/browse/AGPUSH-1000
474	AGPUSH-990	[~lfryc] Done for Keycloak, currently gets the realm name. In our case "aerogear"	1	1	https://issues.jboss.org/browse/AGPUSH-990
475	AGPUSH-986	[~lfryc] do we need this for 1.0.x ? Or would that be just fine for 1.1.x ?	1	1	https://issues.jboss.org/browse/AGPUSH-986
476	AGPUSH-972	[~burrsutter] I am not sure, but should we explain Xcode/IntelliJ here ?	1	1	https://issues.jboss.org/browse/AGPUSH-972
477	AGPUSH-968	[~mwessendorf] Regarding timing, wouldn't it be better to investigate this for 1.1 ?	1	0	https://issues.jboss.org/browse/AGPUSH-968
478	AGPUSH-967	[~sebastienblanc] I wonder, do we expect to get anything with the upgrade? I'd incline to reject this and wait before we do some bigger refactoring. Sounds good?	1	1	https://issues.jboss.org/browse/AGPUSH-967
479	AGPUSH-967	what's the status here? should we do that for alpha.2 ? or not?	1	1	https://issues.jboss.org/browse/AGPUSH-967
480	AGPUSH-965	What's the status here? Should we do that for alpha.2, or later, e.g. beta.1 ?	1	0	https://issues.jboss.org/browse/AGPUSH-965
481	AGPUSH-964	What's the status here? Should we do that for alpha.2, or later, e.g. beta.1 ?	1	0	https://issues.jboss.org/browse/AGPUSH-964
482	AGPUSH-963	What's the status here? Should we do that for alpha.2, or later, e.g. beta.1 ?	1	1	https://issues.jboss.org/browse/AGPUSH-963

#	Requirement	Questions	Answered	Successfully Answered	URL
483	AGPUSH-962	What's the status here? Should we do that for alpha.2, or later, e.g. beta.1 ?	1	1	https://issues.jboss.org/browse/AGPUSH-962
484	AGPUSH-953	should we backport this to 1.0.x as well ?	1	1	https://issues.jboss.org/browse/AGPUSH-953
485	AGPUSH-943	At some point do we want to have this import endpoint exposed also in our UI Console ? Do we want to have the opposite as well , an export function ?	1	1	https://issues.jboss.org/browse/AGPUSH-943
486	AGPUSH-934	[~lfryc] I like what Gabriel did with the styles. Should I make the changes on the less to reflect that?	1	1	https://issues.jboss.org/browse/AGPUSH-934
487	AGPUSH-934	What about preparing one more, "darker" version of pastel colours and opening a poll, where we would let community choose? (can be open to twitter-verse) Also I re-read the github discussion and people reported there that the [colours of the header do not align with colours of buttons and other controls|https://github.com/aerogear/aerogear-unifiedpush-server/pull/384#issuecomment-55136194]. Just to let you know what was the feedback.	1	1	https://issues.jboss.org/browse/AGPUSH-934
488	AGPUSH-934	[~agalante] Andres, could I ask you for your help with redefining the patternfly button colours? I have hard time to apply it properly: https://github.com/aerogear/aerogear-unifiedpush-server/pull/384	0	0	https://issues.jboss.org/browse/AGPUSH-934
489	AGPUSH-928	What's the status here? Should we do that for alpha.2, or later, e.g. beta.1 ?	0	0	https://issues.jboss.org/browse/AGPUSH-928
490	AGPUSH-883	Hey [~abstractj], do you know what these RestEasy WARNINGs are about? {code} 10:27:21,536 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field providers of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,536 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field securityContext of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,537 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field uriInfo of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,537 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field headers of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,538 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field request of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,538 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field response of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,538 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field session of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,539 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field transaction of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,539 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field clientConnection of subresource org.keycloak.services.resources.TokenService will not be injected according to spec 10:27:21,539 WARNING [org.jboss.resteasy.core.ResourceLocator] (http--127.0.0.1-8080-1) Field providerSession of subresource org.keycloak.services.resources.TokenService will not be injected according to spec {code}	1	1	https://issues.jboss.org/browse/AGPUSH-883
491	AGPUSH-864	Ok, nevermind, {{1.4.11}} contains that NPM feature as well, but.... Use of NPM {{1.4.11}} results to {code} [INFO] --- frontend-maven-plugin:0.0.15:npm (npm install) @ unifiedpush-server --- [INFO] Running 'npm install --color='false' in /mnt/workspace/workspaces/aerogear/aerogear-unifiedpush-server/admin-ui [INFO] npm ERR! Error: Cannot find module 'inflight' [INFO] npm ERR! at Function.Module._resolveFilename (module.js:338:15) [INFO] npm ERR! at Function.Module._load (module.js:280:25) [INFO] npm ERR! at Module.require (module.js:364:17) [INFO] npm ERR! at require (module.js:380:17) [INFO] npm ERR! at Object.<anonymous> (/mnt/workspace/workspaces/aerogear/aerogear-unifiedpush-server/admin-ui/node/npm/lib/cache/get-stat.js:6:16) [INFO] npm ERR! at Module._compile (module.js:456:26) [INFO] npm ERR! at Object.Module._extensions..js (module.js:474:10) [INFO] npm ERR! at Module.load (module.js:356:32) [INFO] npm ERR! at Function.Module._load (module.js:312:12) [INFO] npm ERR! at Module.require (module.js:364:17) [INFO] npm ERR! If you need help, you may report this *entire* log, [INFO] npm ERR! including the npm and node versions, at: [INFO] npm ERR! <http://github.com/npm/npm/issues> [INFO] [INFO] npm ERR! System Linux 3.15.6-200.fc20.x86_64 [INFO] npm ERR! command "/mnt/workspace/workspaces/aerogear/aerogear-unifiedpush-server/admin-ui/node/node" "/mnt/workspace/workspaces/aerogear/aerogear-unifiedpush-server/admin-ui/node/npm/bin/npm-cli.js" "install" "--color='false" [INFO] npm ERR! cwd /mnt/workspace/workspaces/aerogear/aerogear-unifiedpush-server/admin-ui [INFO] npm ERR! node -v v0.10.26 [INFO] npm ERR! npm -v 1.4.11 [INFO] npm ERR! code MODULE_NOT_FOUND [INFO] npm ERR! [INFO] npm ERR! Additional logging details can be found in: [INFO] npm ERR! /mnt/workspace/workspaces/aerogear/aerogear-unifiedpush-server/admin-ui/npm-debug.log [INFO] npm ERR! not ok code 0 {code} and use of NPM {{1.4.12}} has this issue: {code} ✓ pre-build test passed successfully npm ERR! cb() never called! npm ERR! not ok code 0 {code}	0	0	https://issues.jboss.org/browse/AGPUSH-864
492	AGPUSH-864	I found another issue which results to build failure: if you kill the build when npm is running, it may result to something like this: {code} [INFO] --- frontend-maven-plugin:0.0.16:grunt (grunt build) @ unifiedpush-admin-ui --- [INFO] Running 'grunt dist --no-color' in /Users/lukasfryc/Documents/workspace/aerogear-unifiedpush-server/admin-ui [INFO] [INFO] module.js:340 [INFO] throw err; [INFO] ^ [INFO] Error: Cannot find module 'findup-sync' [INFO] at Function.Module._resolveFilename (module.js:338:15) [INFO] at Function.Module._load (module.js:280:25) [INFO] at Module.require (module.js:364:17) [INFO] at require (module.js:380:17) [INFO] at Object.<anonymous> (/Users/lukasfryc/Documents/workspace/aerogear-unifiedpush-server/admin-ui/node_modules/grunt-cli/bin/grunt:8:14) [INFO] at Module._compile (module.js:456:26) [INFO] at Object.Module._extensions..js (module.js:474:10) [INFO] at Module.load (module.js:356:32) [INFO] at Function.Module._load (module.js:312:12) [INFO] at Function.Module.runMain (module.js:497:10) {code} the problem here is inconsistent {{node_modules/}}. So cleaning {{node_modules/}} helped. But then, the NPM cache itself was broken: {code} [INFO] --- frontend-maven-plugin:0.0.16:npm (npm install) @ unifiedpush-admin-ui --- [INFO] Running 'npm install --color='false' in /Users/lukasfryc/Documents/workspace/aerogear-unifiedpush-server/admin-ui [INFO] npm WARN package.json ups-admin-ui@0.0.0 No repository field. [INFO] npm ERR! cb() never called! [INFO] npm ERR! not ok code 0 {code} Now we have to refresh the cache (we would normally do that using {{npm clean cache}} but the build has local cache to avoid conflicts in global cache).	1	1	https://issues.jboss.org/browse/AGPUSH-864

#	Requirement	Questions	Answered	Successfully Answered	URL
493	AGPUSH-858	?, ?, 'android', ?) [42121-172] javax.persistence.PersistenceException: org.hibernate.exception.SQLGrammarException: could not prepare statement at org.hibernate.ejb.AbstractEntityManagerImpl.convert(AbstractEntityManagerImpl.java:1387) at org.hibernate.ejb.AbstractEntityManagerImpl.convert(AbstractEntityManagerImpl.java:1310) at org.hibernate.ejb.AbstractEntityManagerImpl.convert(AbstractEntityManagerImpl.java:1316) at org.hibernate.ejb.AbstractEntityManagerImpl.flush(AbstractEntityManagerImpl.java:999) at org.jboss.aerogear.unifiedpush.jpa.VariantDaoTest.createDifferentVariantTypes(VariantDaoTest.java:297) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:17) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:26) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:27) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:271) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:70) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:50) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:238) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:63) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:236) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:53) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:229) at org.junit.runners.ParentRunner.run(ParentRunner.java:309) at org.junit.runner.JUnitCore.run(JUnitCore.java:160) at com.intellij.junit4.JUnit4IdeaTestRunner.startRunnerWithArgs(JUnit4IdeaTestRunner.java:74) at com.intellij.rt.execution.junit.JUnitStarter.prepareStreamsAndStart(JUnitStarter.java:211) at com.intellij.rt.execution.junit.JUnitStarter.main(JUnitStarter.java:67) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at com.intellij.rt.execution.application.AppMain.main(AppMain.java:134) Caused by: org.hibernate.exception.SQLGrammarException: could not prepare statement at org.hibernate.exception.internal.SQLStateConversionDelegate.convert(SQLStateConversionDelegate.java:123) at org.hibernate.exception.internal.StandardSQLExceptionConverter.convert(StandardSQLExceptionConverter.java:49) at org.hibernate.engine.jdbc.spi.SqlExceptionHelper.convert(SqlExceptionHelper.java:124) at org.hibernate.engine.jdbc.internal.StatementPreparerImpl$StatementPreparationTemplate.prepareStatement(StatementPreparerImpl.java:193) at org.hibernate.engine.jdbc.internal.StatementPreparerImpl.prepareStatement(StatementPreparerImpl.java:89) at	1	1	https://issues.jboss.org/browse/AGPUSH-858
494	AGPUSH-853	A few items I noticed, while looking at http://staging.aerogear.org/docs/guides/aerogear-cordova/AerogearCordovaPush/ * cordova plugin add https://github.com/aerogear/aerogear-pushplugin-cordova.git Should it use the latest release, from the plugin repo, instead of the GH repo? * the Helloworld/Quickstart explicitly talk about steps required for Cordova installation: ** "npm install -g Cordova" ** "npm install -g ios-deploy" (for iOS) [~edewit] should we do that here too ? * link to "AeroGear Push plugin API doc" seems to present old docs. At the bottom of http://staging.aerogear.org/docs/specs/aerogear-cordova/Push.html it says _Documentation generated by JSDoc 3.3.0-alpha5 on Sat Jul 26 2014 15:37:23 GMT+0200 (CEST)_ - I think this was generated before our last (0.7.0) release	0	0	https://issues.jboss.org/browse/AGPUSH-853
495	AGPUSH-846	These lines: https://github.com/aerogear/aerogear-unifiedpush-server/blob/master/push/src/main/java/org/jboss/aerogear/unifiedpush/message/sender/APNsPushNotificationSender.java#L116-L117 are invoked in a few cases: * no certificate in the database (e.g. some smart DBA deleted it there) * certificate revoked at the Apple Developer Network site Wondering, if we should replace "No certificate was found" with something like "No (valid) certificate was found. Did you revoke it at the Apple Developer Center" ?	0	0	https://issues.jboss.org/browse/AGPUSH-846
496	AGPUSH-785	Could you [~mwessendorf] find out for which request this is happening. I think there is somehow an invalid json request in there. I cannot reproduce this.	1	1	https://issues.jboss.org/browse/AGPUSH-785
497	AGPUSH-782	I am still trying to figure out how to add a floating action button for "ADD" on ionic. In the mean time here are the screens. Should I create a repository to upload this?	1	1	https://issues.jboss.org/browse/AGPUSH-782
498	AGPUSH-782	Hi Erik Jan de Wit please check out https://github.com/andresgalante/ionicquickstart You'll see the action btw on list.html: ADD And there is css to style and position it at ionic.app.scss If the action btw is a problem, I added a "ADD" btn on the Action Bar. Also let me know if there is any screen or part missing and I'll take care of it. The floating btn is not fully styled yet, I should add an icon, shadows, etc. If you make it work let me know and I'll update the css. Thanks a lot!	1	1	https://issues.jboss.org/browse/AGPUSH-782
499	AGPUSH-781	Should this be a full Responsive website or just mobile?	1	1	https://issues.jboss.org/browse/AGPUSH-781
500	AGPUSH-781	What is the status here ?	0	0	https://issues.jboss.org/browse/AGPUSH-781
501	AGPUSH-780	truly look awesome! nice work [~agalante] mind to share a bit on where can I find the artificacts used on it e.g. font name used icons colors Thanks!	1	1	https://issues.jboss.org/browse/AGPUSH-780
502	AGPUSH-780	Thanks a lot Christos! Primary color #00bcd4 Secondary color #0097a7 Accent color (for login btn) #8bc34a Font color black 87% Font secondary text black 54% Font color white 100% Font secondary text white 70% divider lines color black or white 12% (but you can use the default if there is a ready made solution for the list) Font Family Helvetica Neue (Regular for normal text and Medium for bold text) I'll attach icons soon. Please let me know if there is anything missing. I wish I knew some xcode to help on implementation more.	1	1	https://issues.jboss.org/browse/AGPUSH-780
503	AGPUSH-778	[~agalante] are the sub-tasks still open? or were they all applied to the actual demo apps?	1	1	https://issues.jboss.org/browse/AGPUSH-778
504	AGPUSH-778	What is the status here ?	1	1	https://issues.jboss.org/browse/AGPUSH-778
505	AGPUSH-778	[~mwessendorf] HTML5 app is based on jQuery mobile. I need to learn how to skin jQuery mobile but I always find something to do before. Should I give priority to it? Erik is super busy with FH windows support, I wonder if he can implemente it if I do it.	1	0	https://issues.jboss.org/browse/AGPUSH-778
506	AGPUSH-776	What is the status here ?	1	1	https://issues.jboss.org/browse/AGPUSH-776
507	AGPUSH-775	Hey [~agalante] could you share the icon and colors used on it?	1	1	https://issues.jboss.org/browse/AGPUSH-775
508	AGPUSH-775	Hey [~agalante] what do you think about make splash and registering fullscreen?	1	1	https://issues.jboss.org/browse/AGPUSH-775

#	Requirement	Questions	Answered	Successfully Answered	URL
509	AGPUSH-775	Daniel Passos this is great! very good implementation. If you can make this minor changes to make it look perfect: 1- Background should be white #fff 2- Good idea on the splash page been fullscreen, I'll make one and send it over. 3- I think we should leave the action bar for registering, is that doable? 4- Registering activity bar, can you make it: Background color: #ffcc80 bar color: #ff9800 with: 180dp 5- Font sizes: Action bar: Medium 20sp text in general (all of them): Regular 14sp, Tracking 10, Leading 20 bold text make it Medium weight please. 6- On list, please put a padding top and bottom: 20dp (between the text and the line) 7 - Black screen, please make the icon smaller with 36dp height 40dp and color black 40%. All text on that screen is the same spec as text in general, text but leading is 24. 8- Shadow on action bar is too strong. If that is measure in index, it should be level 1. If any of these are hard to do, forget about them :) Thanks a lot!	1	1	https://issues.jboss.org/browse/AGPUSH-775
510	AGPUSH-775	1) * Done 2) * Don't need update the mockup. Only to clarify I mean without ActionBar not fullscreen, because fullscreen hide NotificationBar 3) * Done 4) * Background color: Done * bar color: Done * with: I have no idea what it means :D 5) * Action bar: Not sure if is a good ideia change the default font size on ActionBar * Text: Added in dimens.xml 6) * Done * I also added a little padding left 7) * Make the icon smaller: Done * Image color: I can't apply alpha on ImageView because alpha was introduced on API leval 11 and the app is for API leval 10. Could you apply that on image and send me? * No sure what you mean by leading on that 8) * Agreed, but it will be a little hard to fix, probably I will need change the ActionBar by toolbar	1	1	https://issues.jboss.org/browse/AGPUSH-775
511	AGPUSH-775	That ActionBar colour is a little different for the Notifications ActionBar colors. That is the ideia? I'm using the same ActionBar colors for registering and notifications for now 6) * Done 7) I have some problem to use *medium* and differents kind of fontFamily on this app because the minSdkVersion is 10. so, I add: * "Nothing here" - 14sp + Bold * "rest of text" - 12sp What do you think? See new screenshots here => https://www.dropbox.com/sh/nvjru5l0e0px3g4/AAApUcF_URMktOalurJ9_t0Va/v3?dl=0	0	0	https://issues.jboss.org/browse/AGPUSH-775
512	AGPUSH-774	[~agalante] are the sub-tasks still open? or were they all applied to the actual demo apps?	1	1	https://issues.jboss.org/browse/AGPUSH-774
513	AGPUSH-774	[~agalante] This is done you made the html yourself	1	1	https://issues.jboss.org/browse/AGPUSH-774
514	AGPUSH-774	[~edewit] yeap, but is it implemented?	0	0	https://issues.jboss.org/browse/AGPUSH-774
515	AGPUSH-774	What is the status here ?	1	1	https://issues.jboss.org/browse/AGPUSH-774
516	AGPUSH-770	What other should Android library do than throwing URISyntaxException if URI contains invalid characters? Still, the example app should handle this exception.	1	1	https://issues.jboss.org/browse/AGPUSH-770
517	AGPUSH-761	The actual issue: {code} TypeError: Cannot read property 'token' of null at request (http://localhost:8080/ag-push/scripts/app.js:178:23) at deferred.promise.then.wrappedCallback (http://localhost:8080/ag-push/bower_components/angular/angular.js:11319:81) at deferred.promise.then.wrappedCallback (http://localhost:8080/ag-push/bower_components/angular/angular.js:11319:81) at http://localhost:8080/ag-push/bower_components/angular/angular.js:11405:26 at Scope.$RootScopeProvider.$get.Scope.$eval (http://localhost:8080/ag-push/bower_components/angular/angular.js:12412:28) at Scope.$RootScopeProvider.$get.Scope.$digest (http://localhost:8080/ag-push/bower_components/angular/angular.js:12224:31) at Scope.$RootScopeProvider.$get.Scope.$apply (http://localhost:8080/ag-push/bower_components/angular/angular.js:12516:24) at HTMLDocument.<anonymous> (http://localhost:8080/ag-push/bower_components/angular/angular.js:9567:22) at HTMLDocument.jQuery.event.dispatch (http://localhost:8080/ag-push/bower_components/jquery/dist/jquery.js:4409:9) at HTMLDocument.elemData.handle (http://localhost:8080/ag-push/bower_components/jquery/dist/jquery.js:4095:28) angular.js:9778 4 Server returned : {code}	0	0	https://issues.jboss.org/browse/AGPUSH-761
518	AGPUSH-746	On the code snippet page I am seeing a different one: {code} ypeError: Cannot read property 'childNodes' of undefined at d (http://judcon-pushee.rhcloud.com/ag-push/scripts/08ae27be.modules.js:37461:264) at p [as highlightBlock] (http://judcon-pushee.rhcloud.com/ag-push/scripts/08ae27be.modules.js:37461:5169) at http://judcon-pushee.rhcloud.com/ag-push/scripts/0b26ed0d.scripts.js:687:10 at http://judcon-pushee.rhcloud.com/ag-push/scripts/08ae27be.modules.js:23205:28 at completeOutstandingRequest (http://judcon-pushee.rhcloud.com/ag-push/scripts/08ae27be.modules.js:13491:10) at http://judcon-pushee.rhcloud.com/ag-push/scripts/08ae27be.modules.js:13792:7 08ae27be.modules.js:18969 {code}	1	1	https://issues.jboss.org/browse/AGPUSH-746
519	AGPUSH-742	where is the URL ?	1	1	https://issues.jboss.org/browse/AGPUSH-742
520	AGPUSH-725	[~hbons] I belive this is related to Keycloak User Management screens. Is this really done?	1	1	https://issues.jboss.org/browse/AGPUSH-725
521	AGPUSH-725	[~bruno_oliveira] How do I access these user management pages to be able to work on this? If we can have a link from the sidebar or if you can give me a a direct path to it I can have a look.	1	1	https://issues.jboss.org/browse/AGPUSH-725
522	AGPUSH-707	Similar to AGPUSH-268 ? If so, please mark AGPUSH-268 as duplicated; Also, since the PR #207 was merged, should this be marked as resolved:done ?	0	0	https://issues.jboss.org/browse/AGPUSH-707
523	AGPUSH-669	Beside org.jboss.resteasy:resteasy-jackson-provider , the other dependencies mentioned are not in the POM of the Java Sender, how can we align these ?	1	1	https://issues.jboss.org/browse/AGPUSH-669
524	AGPUSH-669	You are absolutely correct, those are transitive dependencies. Sorry about that, you may close this issue. What is left is just the use of the aerogear bom (https://issues.jboss.org/browse/AGPUSH-663).Thanks	0	0	https://issues.jboss.org/browse/AGPUSH-669
525	AGPUSH-661	np, so when you are trying to import/update 100k devices. are you doing any batching (eg persisting many entities in one bang instead of one and one)? another thing you could look at if updates is an issue is to enable entity caching, looking at what attributes that can be lazy loaded etc. if you're using the same query over and over again query caching would also speed things up. - as always, caching need to be done right, or it might even slow things down :)	1	1	https://issues.jboss.org/browse/AGPUSH-661
526	AGPUSH-657	[~hbons] any ETA on a new icon ? Or is this something unrelated to the actual icon ?	1	1	https://issues.jboss.org/browse/AGPUSH-657
527	AGPUSH-637	Is that a priority? Or do we have the time to do it for the 1.0.0 ?	1	1	https://issues.jboss.org/browse/AGPUSH-637
528	AGPUSH-637	This is not priority for 0.11.0 and since Erik had problems with adopting aerogear.js in UPS, I wonder whether it is doable in 1.0 at all. ---- Here is a branch where Erik started refactoring: https://github.com/edewit/aerogear-unifiedpush-server/tree/angular-aerogear [~edewit], could you please elaborate little bit on what were blockers here?	1	0	https://issues.jboss.org/browse/AGPUSH-637
529	AGPUSH-637	Should we move this to 'future', [~edewit], [~lfryc], [~lucas.holmquist] ?	1	1	https://issues.jboss.org/browse/AGPUSH-637
530	AGPUSH-629	Hey Erik, is there anything else we should do about this issue?	1	1	https://issues.jboss.org/browse/AGPUSH-629
531	AGPUSH-571	Most of the content has been already moved out, instead of linking to 'real' documentation, we did link to the old 0.10.x branch	0	0	https://issues.jboss.org/browse/AGPUSH-571
532	AGPUSH-567	Is this still relevant ?	0	0	https://issues.jboss.org/browse/AGPUSH-567

#	Requirement	Questions	Answered	Successfully Answered	URL
533	AGPUSH-546	Do you mind to share a link for the main page you are talking about ?	0	0	https://issues.jboss.org/browse/AGPUSH-546
534	AGPUSH-544	Right, REST represents the API. However, in various tests we need to generate applications, variants, etc. Java is the language of choice. We are using REST in the end but we need DTO objects. I was thinking of reusing model jar. Do you plan to have model released in Maven Central? Otherwise me would duplicate and make our own DTO objects.	1	1	https://issues.jboss.org/browse/AGPUSH-544
535	AGPUSH-541	[~supittma] yeah ... I was asking about the registrationID a device/app gets from GCM; Paylload we know: up to 4k - but how long can the actual regID get ? reading some stackoverflow, it seems that 256 is a safe bet (which is what we have, for the registrationId)	1	1	https://issues.jboss.org/browse/AGPUSH-541
536	AGPUSH-530	Are you saying it works on your iOS7 device? That you are seeing a notification on the lockscreen (using the Cordova plugin) ?	1	1	https://issues.jboss.org/browse/AGPUSH-530
537	AGPUSH-530	I think you mean the pull-down/drop-down notification bar, is that what you mean by "touch the notification"?	1	0	https://issues.jboss.org/browse/AGPUSH-530
538	AGPUSH-530	[~edewit] everything works fine w/ iOS6 - did you test w/ iOS7 as well ? The interesting part - the native (non Cordova) bits worked on my iOS 7 (7.0.4) device (meaning I did see an "alert" whlile watching the lockscreen)	1	1	https://issues.jboss.org/browse/AGPUSH-530
539	AGPUSH-528	Should we import this to AeroGear, on Github?	1	1	https://issues.jboss.org/browse/AGPUSH-528
540	AGPUSH-491	What's the status here ?	1	1	https://issues.jboss.org/browse/AGPUSH-491
541	AGPUSH-490	what's the status of this ?	1	0	https://issues.jboss.org/browse/AGPUSH-490
542	AGPUSH-487	[~lfryc] [~sebastienblanc] [~hbons] what is the status ? can we close this out ?	1	1	https://issues.jboss.org/browse/AGPUSH-487
543	AGPUSH-456	What's the status here ?	1	0	https://issues.jboss.org/browse/AGPUSH-456
544	AGPUSH-456	[~lucas.holmquist] what's the status here? Is there anything left?	1	0	https://issues.jboss.org/browse/AGPUSH-456
545	AGPUSH-456	[~lucas.holmquist] what's the status here ? and related tickets? Can we close them ?	1	0	https://issues.jboss.org/browse/AGPUSH-456
546	AGPUSH-446	[~mwessendorf] This ticket made me think of a useful criteria we could have in the future for the client sender API : some kind of alias exclusion filter. For example : { "variants" : ["c3f0a94f-48de-4b77-a08e-68114460857e", "444939cd-ae63-4ce1-96a4-de74b77e3737" ....], "excludeAlias" :["john", "maria"], "categories" : ["someCategory", "otherCategory"], "deviceType" : ["iPad", "AndroidTablet"], "ttl" : 3600, "message": { "alert":"HELLO!", "sound":"default", "badge":7, "someKey":"some value", "anotherCustomKey":"some other value" }, "simple-push": "version=123" }' wdyt ?	0	0	https://issues.jboss.org/browse/AGPUSH-446
547	AGPUSH-400	Which is the best place to add this reset button ? 1. Add an reset action button/link on the pushApps/variants list 2. Add an reset button on the detail screen of the selected pushApp/Variant , in the section "Authentication Info". My preference is #2.	0	0	https://issues.jboss.org/browse/AGPUSH-400
548	AGPUSH-400	i would say, lets put in that modal box, instead of using the confirm, that is the plan for the other ones too https://issues.jboss.org/browse/AGPUSH-315	0	0	https://issues.jboss.org/browse/AGPUSH-400
549	AGPUSH-351	[~sebastienblanc] what is the state of this ticket and the mentioned GIST ?	1	1	https://issues.jboss.org/browse/AGPUSH-351
550	AGPUSH-343	how many seconds do we want to use for the default ? Looks like the default could be pretty high...	1	1	https://issues.jboss.org/browse/AGPUSH-343
551	AGPUSH-343	I think a really high value is fine. How often are you really changing CORS access parameters? If you're developing and making a CORS change, you can clear the browser's cache manually. I think something like 20 or 30 days is fine. If that feels too long, even a week (7 days) would be good.	0	0	https://issues.jboss.org/browse/AGPUSH-343
552	AGPUSH-323	if we are changing things, can we just change the sendTo method to "send" ?	1	1	https://issues.jboss.org/browse/AGPUSH-323
553	AGPUSH-308	A few questions: What information is required to connect to a proxy server? Is it like your regular HTTP proxy or something else? Can we test whether the connection is successful somehow?	0	0	https://issues.jboss.org/browse/AGPUSH-308
554	AGPUSH-307	[~craig.brookes] [~pgough] thanks the PR, and the documentation. Besides that, is there more that's needed to support proxy servers ? If this is really all, I'd like to close this JIRA, as done	0	0	https://issues.jboss.org/browse/AGPUSH-307
555	AGPUSH-287	Mrunal Patel <mpatel@redhat.com> made a comment on [bug 997108|https://bugzilla.redhat.com/show_bug.cgi?id=997108] Matthias, Do you know why aerogear cartridge refuses connection over 8443? Thanks, Mrunal	1	1	https://issues.jboss.org/browse/AGPUSH-287
556	AGPUSH-287	Meng Bo <bmeng@redhat.com> made a comment on [bug 997108|https://bugzilla.redhat.com/show_bug.cgi?id=997108] I have tried create new app with cartridge in comment#7 on INT. Still get the same result. # openssl s_client -connect tls-bmengdev.int.rhcloud.com:8443 connect: Connection refused connect:errno=111 # rhc app show tls tls @ http://tls-bmengdev.int.rhcloud.com/ (uuid: 52130e206cec0e079100095d) ---------- Domain: bmengdev Created: 2:35 AM Gears: 1 (defaults to small) Git URL: ssh://52130e206cec0e079100095d@tls-bmengdev.int.rhcloud.com/~/git/tls.git/ SSH: 52130e206cec0e079100095d@tls-bmengdev.int.rhcloud.com aerogear-aerogear-push-0.8.0 (AeroGear Push Server 0.8.0) ---------- From: https://cartreflect-claytondev.rhcloud.com/reflect?github=aerogear/openshift-origin-cartridge-aerogear-push Gears: Located with mysql-5.1 mysql-5.1 (MySQL Database 5.1) ---------- Gears: Located with aerogear-aerogear-push-0.8.0 Connection URL: mysql://$OPENSHIFT_MYSQL_DB_HOST:$OPENSHIFT_MYSQL_DB_PORT/ Database Name: tls Password: <Hidden> Username: <Hidden> And when I creating same app on stage, the 8443 port can be connected.	1	1	https://issues.jboss.org/browse/AGPUSH-287
557	AGPUSH-282	Can Developers create other Developers (users) or should we create another role ?	1	1	https://issues.jboss.org/browse/AGPUSH-282
558	AGPUSH-263	Sounds good. I fixed the URL in https://github.com/aerogear/openshift-origin-cartridge-aerogear-push/commit/302c4eab79fc50ac92261eb9072b33347d44a6ff and I'm marking this issue as resolved. [~mwessendorf], what would be proper fix version here? Should the issue be moved to a different project if the cartridge has independent versioning? Or do you plan to add versions like *cartridge-0.8.0*?	1	1	https://issues.jboss.org/browse/AGPUSH-263
559	AGPUSH-224	Mrunal Patel <mpatel@redhat.com> made a comment on [bug 997108|https://bugzilla.redhat.com/show_bug.cgi?id=997108] Matthias, Do you know why aerogear cartridge refuses connection over 8443? Thanks, Mrunal	1	1	https://issues.jboss.org/browse/AGPUSH-224
560	AGPUSH-224	Can you take a look, or attach, server.log from the instance that does not work correctly? Perhaps there is an error upon startup. The log can be found in aerogear-push/logs/server.log.	1	0	https://issues.jboss.org/browse/AGPUSH-224
561	AGPUSH-211	AFAIK, for broadcast we don't use already a "message" object. So that is really specific for "selected" messages, right ? https://github.com/aerogear/aerogear-unified-push-java-client/blob/master/src/main/java/org/jboss/aerogear/unifiedpush/SenderClient.java#L89	1	1	https://issues.jboss.org/browse/AGPUSH-211

#	Requirement	Questions	Answered	Successfully Answered	URL
562	AGPUSH-211	[~kborchers] [~mwessendorf] I'm not against that (from the java sender perspective) but that means adding more logic on the Sender impl. And what if we send a "selected" message for both a simplepush network and a "native" network ? We shoudl only choose one of the option, or keep "message" or get rid of it.	1	1	https://issues.jboss.org/browse/AGPUSH-211
563	AGPUSH-195	Do you know the git repo will be 'moved' into this git folder: https://github.com/openshift/origin-server/tree/master/cartridges Or is a different one being used ? Regarding enabling this DB in your MASTER, should we keep this open, or should we close that out ?	1	1	https://issues.jboss.org/browse/AGPUSH-195
564	AGPUSH-162	[~abstractj] "For example here we should add @param annotation of Javadocs doesn't make too much sense:" What do you mean ? We should add the @param ? Or or not (because of the "doesn't make too much sense")	0	0	https://issues.jboss.org/browse/AGPUSH-162
565	AGPUSH-150	I started on this. Just to be clear : we do only paging on the UI side for now ? Nothing to be done the UPS side ? I'm following this approach http://stackoverflow.com/a/13033425 and looks good so far. I'm using a mixin to avoid too much duplication. How many items do we want per page ? 10 ? [~hbons] Any advice about the layout for the pagination ("next", "previous" and page numbers links) ? I will start with a classic / lean design	1	1	https://issues.jboss.org/browse/AGPUSH-150
566	AGPUSH-150	[~hbons] Could you elaborate a bit on patternfly ? Is it a concrete CSS lib or just guidelines ? Samples seems to be based on bootstrap.	1	1	https://issues.jboss.org/browse/AGPUSH-150
567	AGPUSH-116	i'm not sure that plan for this feature, but as we've been building our mobile just a few thoughts on what types of metrics/analytics that could potentially be gathered: Users Date New Returning Total Sessions Date New Returning Total Location Country ex. US, Canada, Germany, etc... City ex. Raleigh, Toronto, Munich, etc... Devices OS ex. Android, iOS, etc... OS Details (Major Version & Minor Version) ex. iOS 7.0.2, iOS 7.1, Android 4.4, etc... Device Maker/Brand ex. iPhone 5 (Global), iPhone 5S (GSM), iPod Touch 5G, SAMSUNG-SGH-I727, HTC One, Full MIPS Android on Emulator, Simulator, etc... Resolution/Width & Height ex. 800x480, 1920x1080, etc... App Version ex. 1.0, 3.1, 2.1.223, etc... Carrier ex. AT&T, T-Mobile, Verizon, etc... Custom Events (Can be used to track a wide variety of custom events that could be used by a funnel, or could be used to drill down into data. Its just a simple server based counter.) ex. Newsletter Signups, Subscriptions, Bought, Video Plays, Articles Read, Button Clicks, Demographics, Behavior flow/how users user the app, etc... Error/Bug Reporting Automatically sends crash report (no sensitive or user data) back to AG Gives user option to give additional error information and perhaps a screenshot Potentially useful links/examples https://github.com/Countly/countly-server http://docs.phonegap.com/en/3.3.0/cordova_device_device.md.html#Device https://github.com/AppBlade/AppBladeSDK/wiki/Feature-Implementation-Guide#iOS-crash-reporting http://devcenter.kinvey.com/phonegap/guides/troubleshooting http://try.crashlytics.com/reports/ https://www.google.com/analytics/mobile/ http://piwik.org/blog/2012/04/how-to-use-piwik-to-track-mobile-apps-activity-clicks-phones-errors-etc/	1	1	https://issues.jboss.org/browse/AGPUSH-116
568	AGPUSH-110	currently SimplePush has _NO_ support for this; but could be good enough for the native plaforms that have support for it	0	0	https://issues.jboss.org/browse/AGPUSH-110
569	AGPUSH-92	I agree this should be configurable, but I'm not sure yet if it should be in present in the UI. OpenShift should change the default path without bothering the user with it. $HOME should normally be writable. If it's not for whatever reason we can ask to change it and show a field in the UI or point to a configuration file. What are potential use cases for changing the default path?	1	0	https://issues.jboss.org/browse/AGPUSH-92
570	AGPUSH-38	How about we show the CURL command for sending, on the Admin UI ?	1	1	https://issues.jboss.org/browse/AGPUSH-38
571	AGPUSH-36	https://developer.android.com/google/gcm/ccs.html#implement The example uses Jive's Smack (an Open Source XMPP (Jabber) library): {code} import org.jivesoftware.smack.ConnectionConfiguration; import org.jivesoftware.smack.ConnectionConfiguration.SecurityMode; import org.jivesoftware.smack.ConnectionListener; import org.jivesoftware.smack.PacketInterceptor; import org.jivesoftware.smack.PacketListener; import org.jivesoftware.smack.XMPPConnection; import org.jivesoftware.smack.XMPPException; import org.jivesoftware.smack.filter.PacketTypeFilter; import org.jivesoftware.smack.packet.DefaultPacketExtension; import org.jivesoftware.smack.packet.Message; import org.jivesoftware.smack.packet.Packet; import org.jivesoftware.smack.packet.PacketExtension; import org.jivesoftware.smack.provider.PacketExtensionProvider; import org.jivesoftware.smack.provider.ProviderManager; import org.jivesoftware.smack.util.StringUtils; import org.json.simple.JSONValue; import org.json.simple.parser.ParseException; import org.xmlpull.v1.XmlPullParser; import java.util.HashMap; import java.util.Map; import java.util.Random; import java.util.logging.Level; import java.util.logging.Logger; import javax.net.ssl.SSLSocketFactory; /** * Sample Smack implementation of a client for GCM Cloud Connection Server. * * <p>For illustration purposes only. */ public class SmackCcsClient { Logger logger = Logger.getLogger("SmackCcsClient"); public static final String GCM_SERVER = "gcm.googleapis.com"; public static final int GCM_PORT = 5235; public static final String GCM_ELEMENT_NAME = "gcm"; public static final String GCM_NAMESPACE = "google:mobile:data"; static Random random = new Random(); XMPPConnection connection; ConnectionConfiguration config; /** * XMPP Packet Extension for GCM Cloud Connection Server. */ class GcmPacketExtension extends DefaultPacketExtension { String json; public GcmPacketExtension(String json) { super(GCM_ELEMENT_NAME, GCM_NAMESPACE); this.json = json; } public String getJson() { return json; } @Override public String toXML() { return String.format("<%s xmlns=\"%s\">%s</%s>", GCM_ELEMENT_NAME, GCM_NAMESPACE, json, GCM_ELEMENT_NAME); } @SuppressWarnings("unused") public Packet toPacket() { return new Message() { // Must override toXML() because it includes a <body> @Override public String toXML() { StringBuilder buf = new StringBuilder(); buf.append("<message"); if (getXmlns() != null) { buf.append(" xmlns=\"").append(getXmlns()).append("\""); } if (getLanguage() != null) { buf.append(" xml:lang=\"").append(getLanguage()).append("\""); } if (getPacketID() != null) { buf.append(" id=\"").append(getPacketID()).append("\""); } if (getTo() != null) { buf.append(" to=\"").append(StringUtils.escapeForXML(getTo())).append("\""); } if (getFrom() != null) { buf.append(" from=\"").append(StringUtils.escapeForXML(getFrom())).append("\""); } buf.append(">"); buf.append(GcmPacketExtension.this.toXML()); buf.append("</message>"); return buf.toString(); } }; } } public SmackCcsClient() { // Add GcmPacketExtension ProviderManager.getInstance().addExtensionProvider(GCM_ELEMENT_NAME, GCM_NAMESPACE, new PacketExtensionProvider() { @Override public PacketExtension parseExtension(XmlPullParser parser) throws Exception { String json = parser.nextText(); GcmPacketExtension packet = new GcmPacketExtension(json); return packet; } }); } /** * Returns a random message id to uniquely identify a message. * * <p>Note: * This is generated by a pseudo random number generator for illustration purpose, * and is not guaranteed to be unique. * */ public String getRandomMessageId() { return "m-" + Long.toString(random.nextLong()); } /** * Sends a downstream GCM message. */ public void send(String jsonRequest) { Packet request = new	0	0	https://issues.jboss.org/browse/AGPUSH-36
572	AGPUSH-20	It looks like there is no way to "GET" a list/one instance other than getting a particular variant and then looking at the instance array. was thinking this could/should go in the corresponding VariantEndpoint.java file ex: /application/{pushAppId}/{variantType}/{variantId}/instance(s) and /application/{pushAppId}/{variantType}/{variantId}/instance/{instanceId} should we reopen?	0	0	https://issues.jboss.org/browse/AGPUSH-20

#	Requirement	Questions	Answered	Successfully Answered	URL
573	AGPUSH-11	those look pretty good. do we need a Push Cordova plugin thing?	0	0	https://issues.jboss.org/browse/AGPUSH-11
574	AGPUSH-11	I've been thinking about this for a while now and I'm honestly not sure. The question is, does Cordova, or one of its plugins bridge JS to the notifications? On the one hand, since it's a native app, it could be covered by the iOS or Android clients but the user would have to know how to incorporate their own or build that bridge. Since it's also using JS, it could just be handled by the JS SimplePush client but then you're only getting in app notifications while it's open. I say we at least create a Cordova component and we'll investigate this more later.	1	1	https://issues.jboss.org/browse/AGPUSH-11
575	AGPUSH-11	Do we need need to duplicate "Unified Push Server" and "SimplePush" so many times? I would recommend # Unified Push # SimplePush # Security # APNs # Server # Client # Admin # UI etc.... Then mix and match as needed - such as "SimplePush, Server, Security". Otherwise you'll end up with tag spagetti imo.	0	0	https://issues.jboss.org/browse/AGPUSH-11
576	AGPUSH-8	Can we mark AEROGEAR-1038 as 'done' ?	0	0	https://issues.jboss.org/browse/AGPUSH-8
577	AGSEC-182	[~kpiwko] reason? Is http://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2013-1624 the reason? We don't make use of CBC	0	0	https://issues.jboss.org/browse/AGSEC-182
578	AGSEC-30	Are you planning to do it? Or should I assign to myself?	0	0	https://issues.jboss.org/browse/AGSEC-30
579	AGSYNC-42	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	0	0	https://issues.jboss.org/browse/AGSYNC-42
580	AGSYNC-41	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	0	0	https://issues.jboss.org/browse/AGSYNC-41
581	AGSYNC-40	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	0	0	https://issues.jboss.org/browse/AGSYNC-40
582	AGSYNC-39	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	0	0	https://issues.jboss.org/browse/AGSYNC-39
583	AGSYNC-38	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	0	0	https://issues.jboss.org/browse/AGSYNC-38
584	AGSYNC-37	Others, like Google (http://blog.chromium.org/2015/02/hello-http2-goodbye-spdy-http-is_9.html) or Jetty are deprecating their SPDY support, in fav. of HTTP/2. Perhaps we should close this one? Or is it really worth it ?	1	0	https://issues.jboss.org/browse/AGSYNC-37
585	AGSYNC-37	this is linked to offline epic - should it be part of AGSYNC ?	1	0	https://issues.jboss.org/browse/AGSYNC-37
586	AGSYNC-30	yes [~supittma] I told [~agalante] that the impl. is _NOT_ scope for alpha.1. However, providing the UX/screenshot/mocks is fine - IMPL on the demo app: _not_ needed alpha.1 :-)	0	0	https://issues.jboss.org/browse/AGSYNC-30
587	AGSYNC-21	[~supittma] [~beve] what should the name of the repo be ? aerogear-android-sync ?	1	1	https://issues.jboss.org/browse/AGSYNC-21
588	AGSYNC-9	[~beve] [~mwessendorf] I wonder, have we discussed which of the two mentioned formats we will use for the Alpha1 impl? While *JSON Patch* is similar to what we have now: {code} [{ "op": "test", "path": "/a/b/c", "value": "foo" }, { "op": "remove", "path": "/a/b/c" }, { "op": "add", "path": "/a/b/c", "value": ["foo", "bar"] }, { "op": "replace", "path": "/a/b/c", "value": 42 }, { "op": "move", "from": "/a/b/c", "path": "/a/b/d" }, { "op": "copy", "from": "/a/b/d", "path": "/a/b/e" }] {code} *JSON Merge Patch* may be actually simpler to adopt and it is IMO also more straightforward to read: {code} { "title": "Hello!", "author" : { "givenName" : "John" }, "tags": ["example"], "content": "This will be unchanged", "phoneNumber": "+01-123-456-7890", "this_will_be_removed": null } {code} The *JSON Patch* is obviously more powerful with its "test", "copy" and "move" operations (the latter two we could potentially use), still they are not necessary.	1	1	https://issues.jboss.org/browse/AGSYNC-9
589	AGSYNC-9	Neither? Is binding our internal representation of the sync data to JSON what we want to do? Also neither of these proposals address the case of small changes to large text values (the diff and merge of the diff sync algorithm). Double also neither of these protocols make any claims for the order of the json fields/values which could also be quite problematic when it comes time to implement shadows and diffing.	1	1	https://issues.jboss.org/browse/AGSYNC-9
590	AGSYNC-9	Sorry, I was incorrect regarding my comment about JSON Merge Patch. It is only the patches that are being sent. My view on this that we could provide different algorithms/storage options for different use cases. The first version, though inefficient, could be used to handle the small changed to large text values like Summers brought up. And a different implementation could use JSON Patch or JSON Merge Patch and store JSON Objects which it then performs the operations on. We could use a content-type property (header depending of transport protocol) to know what type is being targeted. The initial goal of the sync engine was to make this easy to implement, basically you should only have to provide a [ServerSynchronizer|https://github.com/aerogear/aerogear-sync-server/blob/master/api/src/main/java/org/jboss/aerogear/sync/server/ServerSynchronizer.java], and a [ServerDataStore|https://github.com/aerogear/aerogear-sync-server/blob/master/api/src/main/java/org/jboss/aerogear/sync/server/ServerDataStore.java] and you'd be good to go. But it has become painfully obvious that this is not enough since what is called a Diff in the current implementation can vary with the implementation. For example, it might be operations for the current DiffMatchPatch implementation, but for JSON Patch just a patch. I'm currently looking into this refactoring and trying to fix this. @supittma When you say neither. What would you suggest we have here instead or in addition to this?	0	0	https://issues.jboss.org/browse/AGSYNC-9
591	AGSYNC-8	This is about a Java client ? or generally an overall ticket for all different client engines ?	1	1	https://issues.jboss.org/browse/AGSYNC-8
592	AESH-433	is this fixed now with your latest changes [~jdenise]?	1	1	https://issues.jboss.org/browse/AESH-433
593	AESH-399	I think that the appropriate behavior is that if you have something like: "ls "<tab> it should default to call the argument value completor if the command accept arguments, if not it should list options. However if you have something like: "ls -"<tab> or "ls --"<tab> it should list option names and not call the argument value completor. Is this acceptable?	1	1	https://issues.jboss.org/browse/AESH-399
594	AESH-399	If there is a completer on the Argument and the completer has something to show, this could be a valid behavior. But when the completer is not present or there is nothing to show, the user has no clue about what is going on. The argument could be required or not (this is something that should be added, I logged an RFE some time ago). If the argument is not required, do we default to the argument? When a choice is possible between showing argument values or options, something like the following could help: ls <TAB> Command Options: --verbose, -l Command Argument: /subsystem= /core-services= ...	1	1	https://issues.jboss.org/browse/AESH-399
595	AESH-397	the parser currently supports the use of ",' and double quotes "","" using these you can have space in values. is the usage of {,} a standard? if so we could probably extend the parser to support that as well.	0	0	https://issues.jboss.org/browse/AESH-397

#	Requirement	Questions	Answered	Successfully Answered	URL
596	AESH-393	Just immediately I'm thinking that this kind of syntax is very "non-standard". Could we change it to something like:$ ls -l | echo-drm ? If we really have to support this, I think you second suggestion looks like the best. I also would prefer to hide as much as possible.	1	1	https://issues.jboss.org/browse/AESH-393
597	AESH-358	In which context is this needed in the public api?	1	1	https://issues.jboss.org/browse/AESH-358
598	AESH-354	The commits concerning this issue are not merged into the master branch. Is that on purpose?	0	0	https://issues.jboss.org/browse/AESH-354
599	AESH-354	what is the commit not present in master? It seems that it is present but a mismatch in merge message. Stale used "Fix for AESH-354" for something else (interactive mode). In master: https://github.com/aeshell/aesh/commit/cd15eb9780a8d8c356fe5b9d85b710bede9d6c7f In 0.66.+: https://github.com/aeshell/aesh/commit/b1868f4a65c12adde452065f071a1fc50c7cf650 [~stalep] could you confirm?	1	1	https://issues.jboss.org/browse/AESH-354
600	AESH-354	Ok, we'll switch from master to 0.66.+ branch. Although we haven't had any problems using the master branch for our dev in the last 9 month. Thanks for this great project!	1	1	https://issues.jboss.org/browse/AESH-354
601	AESH-352	[~stalep] (CC: [~jdenise]): Any update here? Do you have fixed this issue in Aesh-readline?	1	0	https://issues.jboss.org/browse/AESH-352
602	AESH-352	hi, i have not added functionality to do this yet. I'll try to do it tomorrow, it should not be very difficult. we need æsh to run that command and parse the output. its important that we get the correct width and height. i wonder how infocmp works on hp-ux, do you know [~mkopecky]?	0	0	https://issues.jboss.org/browse/AESH-352
603	AESH-348	I added some unit tests to my maven project to show how to reproduce. There is still an issue with 0.66.7-SNAPSHOT. The "hasOption" is not working. See my README for the full replication steps - just build with the latest snapshot and you'll see the one unit test fails. $ git clone git@github.com:jmazzitelli/aesh-348.git && cd aesh-348 $ mvn clean install -Daesh.version=0.66.7-SNAPSHOT Tests run: 5, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 0.124 sec <<< FAILURE! testMyOptionAtEndWithoutValue(MainTest) Time elapsed: 0.012 sec <<< FAILURE! java.lang.AssertionError: Should have myoption at MainTest.testMyOptionAtEndWithoutValue(MainTest.java:41) This happens when the option is specified last, with no value, such as: mycmd --abc 123 --myoption	1	1	https://issues.jboss.org/browse/AESH-348
604	AESH-348	What is the probability of getting this bug fixed? There is still a problem in the latest 0.66.7.SNAPSHOT release - see my last comment - including the automated test that can be used to demonstrate the problem.	1	1	https://issues.jboss.org/browse/AESH-348
605	AESH-321	with the current impl thats in master child commands of a groupcommans will be autocompleted if the the groupcommand do not contain any options/arguments. atm this also happens when autocompleting on an empty input line: {quote} [aesh@rules]$ alias export group child1 ls prompt test exit foo group child2 man run unalias [aesh@rules]$ {quote} compared to 0.60 behaviour: {quote} [aesh@rules]$ alias exit export foo group ls man prompt run test unalias [aesh@rules]$ {quote} is that ok [~brmeyer]?	1	1	https://issues.jboss.org/browse/AESH-321
606	AESH-306	hi, its been a while since i looked at the code for the groupcommand parser, but i think one of the reasons it only add a space when pressing space for 1. is that the group "main" command also can contain options. the parser could be smarter though and check if there are no options it could instead list all the child commands. does this sound like a ok solution? i dont think i can get this in 0.61 though. im planning to release 0.61 tomorrow or monday and im leaving on vacation after next week. regarding 2. i think that could be done as well. we could add a "separator" char for the group command or something. i need to look into it, but i think its doable. in any case i think it deserves a separate jira :) 3. we have the possibility to specify the color/font/etc for options, but not commands atm. one idea might be to add a feature for commands to have a specified renderer similar to options (OptionRenderer). - as with 2. i also think it deserves a separate jira.	1	1	https://issues.jboss.org/browse/AESH-306
607	AESH-306	{quote} the group "main" command also can contain options. the parser could be smarter though and check if there are no options it could instead list all the child commands. does this sound like a ok solution? {quote} +1 Re: #3, it could be something even simpler. Append ":" to the end of them, assuming you use that as the separator?	1	1	https://issues.jboss.org/browse/AESH-306
608	AESH-300	Im probably a bit slow but i dont see the practical usecase for this, could you give an example? :) - i dont oppose it, i just would like to understand why.	1	1	https://issues.jboss.org/browse/AESH-300
609	AESH-289	hi, sorry for the late reply. do you know which version of æsh you are using? the latest version should work with those keys you describe. if you are using æsh 0.5* then it would be helpful if you tried to run: java -cp aesh-version.jar AeshExample that example should create a aesh.log file in your tmp folder (/tmp on posix) for each keypress there will be a line starting with "GOT" if you could comment here what you see when you press the home/end keys that would be great. æsh have hardcoded some values that might change, but it is also using infocmp on posix systems to try to find the correct key (which works in most cases).	1	1	https://issues.jboss.org/browse/AESH-289
610	AESH-289	I am using whatever is in Forge 2.10.0. I think it's 0.56.1. Could you please give a link to aesh-version.jar? Since we use it as a lib, I only have aesh.jar and aesh-extensions.jar.	1	1	https://issues.jboss.org/browse/AESH-289
611	AESH-281	hi, i did some inital work on this, but didnt get it to work as wanted. could you look at this when you have time [~danielcunha]?	1	1	https://issues.jboss.org/browse/AESH-281
612	AESH-259	Can we close it?	1	1	https://issues.jboss.org/browse/AESH-259
613	AESH-251	What do you prefer? Close the issue or return "" case the method called with parameter without $? (I have PR in gun! Hehehe!)	1	1	https://issues.jboss.org/browse/AESH-251
614	AESH-249	yup, that should work. another thing is in æsh we're using a lot System.getProperty("file.separator") can we still assume that we can use this or do we have to get that from this new api? we also need a method ala: List<AeshResource> listRoots();	1	1	https://issues.jboss.org/browse/AESH-249
615	AESH-223	so, thinking about this a bit further. do you think it would make sense to have an annotation type for pipe/redirection handling? atm the pipe "result" is caught by æsh and stored in a buffer, then the "receiving" command needs to check that buffer to see if there is any input. would it be easier to use the api you suggested and have æsh "automatically" populate the field with data if it receives data through pipe (i guess identical to what you said if im not misunderstanding you but with a specific annotation name) ? ala: {code:java} @Retention(RUNTIME) @Target(FIELD) public @interface PipedList { } public class GrepCommand { @PipedList private List<Object> pipeInput; @Option private String arguments; Result<Set<Object>>, Formatter<Set<Object>> execute(CommandInvocation) { .. grep based on el expression passed in by arguments on each Object in pipeinput and return Set<Object> } } {code}	1	1	https://issues.jboss.org/browse/AESH-223
616	AESH-223	ive started implementing it in the result_and_format branch, but im a but uncertain regarding the Format interface. what is the reason for it, eg what is it meant to do? another thing, lets say we have; command1 | command2 command returns some text that is picked up by æsh and is then injected into the pipelist (or whatever its called) field of command2. the value from command1 is always as text so the question is if we need a "converter" from text to whatever object type is used for the field that is tagged with pipelist?	0	0	https://issues.jboss.org/browse/AESH-223

#	Requirement	Questions	Answered	Successfully Answered	URL
617	AESH-223	after thinking a bit more about this (and implementing a bit), im not sure if this will be very intuitive since it will require a two-step execution of commands. but i would still like to try to get it implemented somehow. what if we do something like: {code:java} public interface Command<T> { List<T> process(CommandInvocation2 ci); CommandResult execute(CommandInvocation ci, List<T> data); } {code} process will not push anything to the shell, just return some data that will be sent as a value to the execute method. CommandInvocation2 (in lack of a better name in this example) is identical to CommandInvocation, but do not have a reference to Shell.out() (not sure if this is required, but it would prohibit commands to push anything to the terminal at this stage). in a pipeline scenario we could either have the values thats passed from the previous command included in the CommandInvocation object or a command might also have a field annotated with {code}@Pipeline{code} where æsh could inject the values. what do you think? [~lincolnthree] and [~gastaldi] any input?	0	0	https://issues.jboss.org/browse/AESH-223
618	AESH-209	AESH-210 should have resolved these afaik. could you please check?	1	1	https://issues.jboss.org/browse/AESH-209
619	AESH-150	I think that a better way to do this would be to implement a ManProvider, which accepts a class, or a command, an returns the URL to a MAN page. Thus, Forge could provide our own ManProvider extension, and do our own resolution. Do you still support AsciiDoc?	1	1	https://issues.jboss.org/browse/AESH-150
620	AESH-108	Or probably we could have a single Completion implementation in Forge 2 and then it could iterate through the available commands. What do you think that would be the best strategy ?	1	1	https://issues.jboss.org/browse/AESH-108
621	AESH-23	yes i agree the way its works now is not good enough. so, do if we change it to: --name "joe bloggs", will set the option name value to "joe bloggs" (without "). it should also accept ' AND change escaped spaces to "normal" spaces does that sound ok?	1	1	https://issues.jboss.org/browse/AESH-23
622	AGOVA-83	Seems resolved. Socializer shows a failure, but it is related to a missing repository. Should that persist I guess I'll create another issue.	0	0	https://issues.jboss.org/browse/AGOVA-83
623	AGOVA-6	Can this be moved to a 0.5.1 target, because existing project that depends on Seam would not benefit from having the DS dependency until DS is "Seam/Solder"-grade?	1	1	https://issues.jboss.org/browse/AGOVA-6
624	AGOVA-6	That's great Ove ! We keep Solder because of Generic Bean mechanism : DS team is not convince to adopt it and I didn't have time to convince them. Do you think we could swith to your ptoject for Generics ?	1	1	https://issues.jboss.org/browse/AGOVA-6
625	AGOVA-3	The obvious way would be to support client side information storage in a cookie, but in some use case we me need to have user information on server side (asynchronous or long batch like crawling user network graph). should we support both scenario or try to have a unique smart one ?	0	0	https://issues.jboss.org/browse/AGOVA-3
626	AJSF-80	Is this going to be fixed with 3.1? Without this fix I will not be able to use RichFaces in my application	0	0	https://issues.jboss.org/browse/AJSF-80
627	AJSF-46	Search components in reRender'ed set changed to same algorithm as in h:outputLabel, h:message components : 0) absolute paths, started from ':' used "as is" 1) search all enclosed NamingContainer components for a given id. 2) descendand search from ViewRoot.	1	1	https://issues.jboss.org/browse/AJSF-46
628	AJSF-33	Can anyone tell us how to do it?	1	0	https://issues.jboss.org/browse/AJSF-33
629	AJSF-7	a4j-trinidad is not yet published to maven repository and not build by default. However, it is possible to build and deploy the package manually. Is there any special reason is package is released to public? Are there any outstanding issue?	1	0	https://issues.jboss.org/browse/AJSF-7
630	ENTMQCL-251	Unclear on scope - wouldn't this require the python-qpid client be ported to python 3.x?	0	0	https://issues.jboss.org/browse/ENTMQCL-251
631	ENTMQCL-164	[~justi9] What is expected to be tested here? I believe that connection failover was supported in python amqp client prior to 1.1.0-alpha. Have there been any changes to this functionality?	1	1	https://issues.jboss.org/browse/ENTMQCL-164
632	ANDIAMO-30	I think the answer here might be to tell Windows users to use Wordpad instead of Notepad. Wordpad, apparently, can handle either LF or CR+LF, just like VIM can on Linux. Any fairly recent Windows version has Wordpad along with Notepad (at least from XP on, which should be more sufficient). Jim, can you test out Wordpad on a file that doesn't work with Notepad. We could possible put a note on this in the knowledge base in the customer portal, and include it the right documentation as a side note.	1	1	https://issues.jboss.org/browse/ANDIAMO-30
633	ANDIAMO-24	I guess the comment was for EPP, Jim, why didn't you report in EPP instead ?	1	0	https://issues.jboss.org/browse/ANDIAMO-24
634	ANDIAMO-20	Jim, can you elaborate on why this is a major problem? I want to understand the priority. Also, if notepad on Windows doesn't respect the carriage return line feeds, what's going to make it easy to edit something in notepad?	1	1	https://issues.jboss.org/browse/ANDIAMO-20
635	ANDIAMO-20	I don't have a Windows box around, but I do have Wine installed, and it has the notepad application. When I open the configuration files in EAP 5.0.1, they all look fine. What's the trick here?	1	1	https://issues.jboss.org/browse/ANDIAMO-20
636	ANDIAMO-20	Jim, Could you also attach the files directly. I want to do a hex diff of the file on your system, and the file on mine.	1	1	https://issues.jboss.org/browse/ANDIAMO-20
637	ANDIAMO-20	Windows XP: Steps: Downloaded EAP 5.0.1 zip file Run from the cmd.exe jar -xf ${name of file} Browse attached file/s via Notepad Just reconfirmed all of the above Not sure how you can not reconfirm? Very interesting. What is your environment? If you are still running Wine, I would guess that the libs are converting these characters correctly for you when you open the file. jar -xf changing the contents of the file on the way out, that sounds pretty crazy to me. Although that is why I call it Computer Art sometimes and not Computer Science or Computer Engineering	1	1	https://issues.jboss.org/browse/ANDIAMO-20
638	ANDIAMO-20	I did attach the two files..below..what files are you seeking..I attached jboss-service.xml and jmx-console-users.properties???? at Aug 16th at 11:18 and 11:17 AM????	1	1	https://issues.jboss.org/browse/ANDIAMO-20
639	ANDIAMO-20	Is a process in place to test this for all products/projects etc so this is never an issue again?	1	0	https://issues.jboss.org/browse/ANDIAMO-20
640	ANDIAMO-7	Does this apply to the AS7 subsystem codebase only or is it relevant to jboss projects that are pulled in as binary dependencies too?	0	0	https://issues.jboss.org/browse/ANDIAMO-7
641	ANDIAMO-2	Discussions on Andiamo mailing list. General consensus seems almost complete at least in terms of the format for messages. No discussion yet on common API though. Outstanding issue around exceptions though: how are they formatted? Jason will complete the format thread and then move to the implementation/API discussion.	0	0	https://issues.jboss.org/browse/ANDIAMO-2
642	APIMAN-1236	Do you want metrics to be disabled entirely or stored in some other system? To disable completely, set and/orchange these properties in apiman.properties: {code} apiman-gateway.metrics=io.apiman.gateway.engine.impl.InMemoryMetrics apiman-manager-ui.metrics.enable=false {code}	1	0	https://issues.jboss.org/browse/APIMAN-1236
643	APIMAN-1234	Thank you Eric. I need to increase buffer size for request as well as response payload. I hope change in these 3 files will increase buffer for both request and response. Will there be any consequences if I make it to 32K ?	1	0	https://issues.jboss.org/browse/APIMAN-1234

#	Requirement	Questions	Answered	Successfully Answered	URL
644	APIMAN-1212	I tried npm install command at location apiman/manager/ui/war on AWS server but getting errors: {color:#d04437}]# npm install npm ERR! Linux 4.4.14-24.50.amzn1.x86_64 npm ERR! argv "/usr/bin/node" "/usr/bin/npm" "install" npm ERR! node v6.3.1 npm ERR! npm v3.6.0 npm ERR! code EINVALIDTYPE npm ERR! typeerror Error: Argument #1: Expected string but got array npm ERR! typeerror at markDeps (/usr/lib/node_modules/npm/lib/install/deps.js:89:5) npm ERR! typeerror at /usr/lib/node_modules/npm/node_modules/slide/lib/async-map.js:52:35 npm ERR! typeerror at Array.forEach (native) npm ERR! typeerror at /usr/lib/node_modules/npm/node_modules/slide/lib/async-map.js:52:11 npm ERR! typeerror at Array.forEach (native) npm ERR! typeerror at asyncMap (/usr/lib/node_modules/npm/node_modules/slide/lib/async-map.js:51:8) npm ERR! typeerror at Array.<anonymous> (/usr/lib/node_modules/npm/node_modules/slide/lib/bind-actor.js:15:8) npm ERR! typeerror at LOOP (/usr/lib/node_modules/npm/node_modules/slide/lib/chain.js:15:14) npm ERR! typeerror at chain (/usr/lib/node_modules/npm/node_modules/slide/lib/chain.js:20:5) npm ERR! typeerror at recalculateMetadata (/usr/lib/node_modules/npm/lib/install/deps.js:114:3) npm ERR! typeerror This is an error with npm itself. Please report this error at: npm ERR! typeerror <http://github.com/npm/npm/issues> {color}	0	0	https://issues.jboss.org/browse/APIMAN-1212
645	APIMAN-1212	You're trying to install npm on an AWS server? What is the goal you are trying to accomplish? Typically gulp would be used on a local developer machine for the purpose of doing iterative development work on the core apiman UI.	1	0	https://issues.jboss.org/browse/APIMAN-1212
646	APIMAN-1157	Can you describe what you would expect to see for the configuration of such a policy?	1	1	https://issues.jboss.org/browse/APIMAN-1157
647	APIMAN-1129	String forwardIdentityHttpHeader; private boolean requireTransportSecurity; private Boolean requireBasicAuth; private MongoIdentitySource mongoIdentitySource; public RMConfig() { } public String getRealm() { return realm; } public String getForwardIdentityHttpHeader() { return forwardIdentityHttpHeader; } public boolean isRequireTransportSecurity() { return requireTransportSecurity; } public Boolean getRequireBasicAuth() { return requireBasicAuth; } public MongoIdentitySource getMongoIdentitySource() { return mongoIdentitySource; } public void setRealm(String realm) { this.realm = realm; } public void setForwardIdentityHttpHeader(String forwardIdentityHttpHeader) { this.forwardIdentityHttpHeader = forwardIdentityHttpHeader; } public void setRequireTransportSecurity(boolean requireTransportSecurity) { this.requireTransportSecurity = requireTransportSecurity; } public void setRequireBasicAuth(Boolean requireBasicAuth) { this.requireBasicAuth = requireBasicAuth; } public void setMongoIdentitySource(MongoIdentitySource rmIdentitySource) { this.mongoIdentitySource = mongoIdentitySource; } } {code} and this is my JSON schema: {code:json} { "title" : "Mongo Authentification Policy Configuration", "description" : "Configure all of the necessary properties used by Mongo policy.", "type" : "object", "properties": { "Authentication Realm": { "title" : "Authentication Realm", "type": "string", "minLength" : 1, "maxLength" : 64 }, "Require Transport Security": { "title": "Require Transport Security", "type": "boolean", "default": false }, "Forward Authenticated Username as HTTP Header": { "title" : "Forward Authenticated Username as HTTP Header", "type": "string", "minLength" : 1, "maxLength" : 64 }, "Require Basic Authentication": { "title" : "Require Basic Authentication", "type": "boolean", "default": false } } } {code} Because until now my MongoIdentitySource is empty. and this is my Policy is the same as Basic authentication policy: {code:java} public class MongoPolicy extends AbstractMappedPolicy<MongoConfig> { private static final MongoValidator MongoValidator = new MongoValidator(); public MongoPolicy() { } @Override protected Class<MongoConfig> getConfigurationClass() { return MongoConfig.class; } @Override protected void doApply(final ApiRequest request, final IPolicyContext context, final MongoConfig config, final IPolicyChain<ApiRequest> chain) { String authHeader = request.getHeaders().get("Authorization"); //$NON-NLS-1$ boolean requireBasic = config.getRequireBasicAuth() == null ? Boolean.TRUE : config.getRequireBasicAuth(); // Handle the case where no authentication credentials are provided if (authHeader == null || authHeader.trim().isEmpty()) { if (requireBasic) { sendAuthFailure(context, chain, config, PolicyFailureCodes.BASIC_AUTH_REQUIRED); return; } else { chain.doApply(request); return; } } // Handle the case where auth credentials are provided but they aren't BASIC // credentials (e.g. BEARER-TOKEN) if (!authHeader.toUpperCase().startsWith("BASIC ")) { //$NON-NLS-1$ if (requireBasic) { sendAuthFailure(context, chain, config, PolicyFailureCodes.BASIC_AUTH_REQUIRED); return; } else { chain.doApply(request); return; } } // Check transport security if (config.isRequireTransportSecurity() && !request.isTransportSecure()) { sendAuthFailure(context, chain, config, PolicyFailureCodes.TRANSPORT_SECURITY_REQUIRED); } // Parse the Authorization http header. String username; String password = null; try { String userpassEncoded = authHeader.substring(6); byte[] decoded = Base64.decodeBase64(userpassEncoded); String data = new String(decoded, "UTF-8"); //$NON-NLS-1$ int sepIdx = data.indexOf(':'); if (sepIdx > 0) { username = data.substring(0, sepIdx); password = data.substring(sepIdx + 1); } else { username = data; } } catch (Throwable t) { // TODO log this error to apiman::logger	1	1	https://issues.jboss.org/browse/APIMAN-1129
648	APIMAN-1129	Since you have a few different issues here, can you ZIP up your entire maven project and attach it to this jira?	1	1	https://issues.jboss.org/browse/APIMAN-1129
649	APIMAN-1129	So where can i find the source code for: {code} import foo.bar.Hello; import tmp.foo.spi.TestHello; {code} ?	1	1	https://issues.jboss.org/browse/APIMAN-1129

#	Requirement	Questions	Answered	Successfully Answered	URL
650	APIMAN-1129	Your code doesn't work either right because I get a failure in the test using Version 1.2.4.Final {code:java} --- T E S T S --- Running com.mongo.example.auth.MongoValidatorTest 10:12:52,132 WARN shiro.ini found at the root of the classpath, but it did not contain any data. Tests run: 1, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 0.25 sec <<< FAILURE! testValidate(com.mongo.example.auth.MongoValidatorTest) Time elapsed: 0.176 sec <<< ERROR! java.lang.ExceptionInInitializerError at com.mongo.example.auth.MongoValidator.validate(MongoValidator.java:54) at com.mongo.example.auth.MongoValidatorTest.testValidate(MongoValidatorTest.java:43) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:17) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:271) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:70) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:50) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:238) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:63) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:236) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:53) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:229) at org.junit.runners.ParentRunner.run(ParentRunner.java:309) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:252) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:141) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:112) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:189) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:165) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:85) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:115) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:75) Caused by: java.lang.RuntimeException: while loading WorkUnit ('class foo.bar.Hello'): Expecting strictly one implementation, but found {[]} at foo.bar.Hello.<clinit>(Hello.java:23) ... 31 more Results : Tests in error: testValidate(com.mongo.example.auth.MongoValidatorTest) Tests run: 1, Failures: 0, Errors: 1, Skipped: 0 {code}	1	1	https://issues.jboss.org/browse/APIMAN-1129
651	APIMAN-1129	Ok I tested My plugin on the new version, it by pass the last error but now I het an error in class Policy where I have the Authorization {code:java} if (result.getResult()) { String forwardIdentityHttpHeader = config.getForwardIdentityHttpHeader(); if (forwardIdentityHttpHeader != null && !forwardIdentityHttpHeader.trim().isEmpty()) { request.getHeaders().put(forwardIdentityHttpHeader, forwardedUsername); } RequestMetric metric = context.getAttribute(PolicyContextKeys.REQUEST_METRIC, (RequestMetric) null); if (metric != null) { metric.setUser(forwardedUsername); } // Remove the authorization header so that it doesn't get passed through to the backend API // TODO: make this optional - perhaps they *want* the auth header passed through? request.getHeaders().remove("Authorization"); //$NON-NLS-1$ chain.doApply(request); } else { sendAuthFailure(context, chain, config, PolicyFailureCodes.BASIC_AUTH_FAILED); } {code} In this part I don't get why? finally I have a 404 not found error trace : {code:java} 14:31:18,857 ERROR [stderr] (default task-43) java.nio.channels.ClosedChannelException 14:31:18,866 ERROR [stderr] (default task-43) at io.undertow.channels.DetachableStreamSinkChannel.write(DetachableStreamSinkChannel.java:238) 14:31:18,879 ERROR [stderr] (default task-43) at io.undertow.server.HttpServerExchange$WriteDispatchChannel.write(HttpServerExchange.java:2000) 14:31:18,885 ERROR [stderr] (default task-43) at io.undertow.servlet.spec.ServletOutputStreamImpl.writeBufferBlocking(ServletOutputStreamImpl.java:567) 14:31:18,892 ERROR [stderr] (default task-43) at io.undertow.servlet.spec.ServletOutputStreamImpl.flushInternal(ServletOutputStreamImpl.java:482) 14:31:18,898 ERROR [stderr] (default task-43) at io.undertow.servlet.spec.ServletOutputStreamImpl.flush(ServletOutputStreamImpl.java:469) 14:31:18,907 ERROR [stderr] (default task-43) at io.apiman.gateway.platforms.servlet.GatewayServlet$3.write(GatewayServlet.java:310) 14:31:18,933 ERROR [stderr] (default task-43) at io.apiman.gateway.platforms.servlet.GatewayServlet$3.write(GatewayServlet.java:298) 14:31:18,942 ERROR [stderr] (default task-43) at io.apiman.gateway.engine.impl.DefaultPolicyFailureWriter.write(DefaultPolicyFailureWriter.java:98) 14:31:18,948 ERROR [stderr] (default task-43) at io.apiman.gateway.platforms.servlet.GatewayServlet.writeFailure(GatewayServlet.java:295) 14:31:18,978 ERROR [stderr] (default task-43) at io.apiman.gateway.platforms.servlet.GatewayServlet$1.handle(GatewayServlet.java:152) 14:31:18,990 ERROR [stderr] (default task-43) at io.apiman.gateway.platforms.servlet.GatewayServlet$1.handle(GatewayServlet.java:103) 14:31:18,995 ERROR [stderr] (default task-43) at io.apiman.gateway.engine.impl.ApiRequestExecutorImpl.lambda$wrapResultHandler$0(ApiRequestExecutorImpl.java:159) 14:31:19,013 ERROR [stderr] (default task-43) at io.apiman.gateway.engine.impl.ApiRequestExecutorImpl.lambda$createPolicyFailureHandler$16(ApiRequestExecutorImpl.java:609) 14:31:19,020 ERROR [stderr] (default task-43) at io.apiman.gateway.engine.policy.Chain.doFailure(Chain.java:230) {code}	1	1	https://issues.jboss.org/browse/APIMAN-1129
652	APIMAN-1122	Sorry but I didn't know about the mailing List or the IRC (which I can't install), anyway do you have an idea how can I get the pluginClassLoader in my plugin? I'm trying this code but it does not work :/ {code:java} PluginClassLoader cl = (PluginClassLoader) Thread.currentThread().getContextClassLoader(); try{ URL toto = cl.getResource("mongoDB.properties"); System.out.println(toto.toURI().getPath()); System.out.println(toto.toString()); } catch (Exception e){ System.out.println("----------"); System.out.println("----------"); System.out.println(e.getMessage()); System.out.println("----------"); System.out.println("----------"); } {code} Knowing that actually there is a class in my JAR that load mongoDB.properties and this file exists in the Jar. I copy it but I have to call the class in the JAR to do the processing. Thank you for your help	1	1	https://issues.jboss.org/browse/APIMAN-1122

#	Requirement	Questions	Answered	Successfully Answered	URL
653	APIMAN-1122	So I have no control on my service loader, but I didn't get your last solution about changing the current thread's context classloader prior to calling whatever library. Can you give me a simple example please?	1	1	https://issues.jboss.org/browse/APIMAN-1122
654	APIMAN-1122	Ok, I will test it once you release it in the website and give you a feedback. When can we have the possibility to extend the policies in APIMAN like BASIC AUTHENTICATION with our needs? It is something necessary I guess. Like in my case I just needed it to use a new identity source but I had to rewrite the doApply function because it called a private function which is validateCrendentials who only let you use the 3 identity source defined by you. I guess it will be interessting to let this function public so I can override it or expand it. What do you think?	1	1	https://issues.jboss.org/browse/APIMAN-1122
655	APIMAN-1122	Can you open up a separate JIRA issue to make the basic auth policy more extensible? :)	1	0	https://issues.jboss.org/browse/APIMAN-1122
656	APIMAN-1122	Would you be able to share your code or create a stripped-down version that reproduces the problem? I'd be happy to give it a try and figure out what's going wrong!	0	0	https://issues.jboss.org/browse/APIMAN-1122
657	APIMAN-1119	Do you want this INI file to be packaged inside the plugin? You should be able to access such a resource by loading it through the plugin's classloader. Something like this: {code} URL configURL = getClass().getClassLoader().getResource("/config/my-properties.ini"); {code} Assuming you put the file into *src/main/resources/config/my-properties.ini*. Note: I can never remember if the leading slash in the string is required. Try both. :)	1	1	https://issues.jboss.org/browse/APIMAN-1119
658	APIMAN-1119	Do you have a full stack trace you can include? I will try to reproduce this locally.	1	1	https://issues.jboss.org/browse/APIMAN-1119
659	APIMAN-1119	Yes I really like too. Especially the Pom.xml!! Where did you put the shiro.ini? In the resources folder?	1	1	https://issues.jboss.org/browse/APIMAN-1119
660	APIMAN-1119	Thank you I have already solved the problem. It was in the pom.xml apiman-plugins parent where the ressource are ignored. I have a question about this part of code: {code:xml} <plugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-enforcer-plugin</artifactId> <version>${version.enforcer.plugin}</version> <dependencies> <dependency> <groupId>de.is24.maven.enforcer.rules</groupId> <artifactId>illegal-transitive-dependency-check</artifactId> <version>1.7.4</version> </dependency> </dependencies> <executions> <execution> <id>enforcer-transitives-check</id> <goals> <goal>enforce</goal> </goals> <phase>process-classes</phase> <configuration> <fail>true</fail> <rules> <illegalTransitiveDependencyCheck implementation="de.is24.maven.enforcer.rules.IllegalTransitiveDependencyCheck"> <reportOnly>false</reportOnly> <regexIgnoredClasses> <regexIgnoredClass>javax\..+</regexIgnoredClass> <regexIgnoredClass>org.w3c.dom\..+</regexIgnoredClass> </regexIgnoredClasses> <useClassesFromLastBuild>true</useClassesFromLastBuild> </illegalTransitiveDependencyCheck> </rules> </configuration> </execution> </executions> </plugin> {code} For what reason do you use this plugin like this? Can you explain please?	1	1	https://issues.jboss.org/browse/APIMAN-1119
661	APIMAN-1093	How to use this new feature? It's easy, just get the current *RequestMetric* object from the policy context and add your own custom attributes to it! {code} RequestMetric metric = context.getAttribute(PolicyContextKeys.REQUEST_METRIC, (RequestMetric) null); metric.getAttributes().put("my-property-1", "value-1"); metric.getAttributes().put("my-property-2", "value-2"); {code}	0	0	https://issues.jboss.org/browse/APIMAN-1093
662	APIMAN-1045	When you say backend services are clustered, wouldn't you be having a cluster IP or basically a load balancer ip such as F5 (Hardware based) or Apache web server (software based) IP? On the APIMAN you could just configure the backend services load balancer IP which in the event of outage for one of the servers could re-direct the traffic to the one that is alive.	1	0	https://issues.jboss.org/browse/APIMAN-1045
663	APIMAN-1027	If [~opontes] and [~eric.wittmann] could comment... > 1) section 4.4.3 is suggesting that redirect URLs should be modified for apimanui only. Other clients (apiman, apiman-geteway) should be mentioned as well Do we need redirect URLs for apiman and apiman-gateway? The user shouldn't be logging directly into either of those endpoints using a browser, AFAIK. Does it not work unless you change it? > 3) The guide seems to lack any information that client configuration in KC needs to be reflected in related section of standalone.xml as well. Additionally the realm file linked in the docs [2] is configured as confidential while default values in standalone.xml are configured for public clients (at least apimanui is). Yes, this is an annoying aberration. For demo and other reasons (it was tricky to get the secrets to work properly when using cURL, it might be fixed now) we turned it off in the example config. We should harmonise it (e.g. provide a separate one with secrets enabled for this guide).	1	1	https://issues.jboss.org/browse/APIMAN-1027
664	APIMAN-1018	How do we envisage this working? It's simple to do given the groundwork already done, but I want some clarification. * What are the preconditions for allowing a plan to be deleted? That it's unused by any ClientApp/Contract? Or we just delete any contracts that use the plan * Are we anticipating that we'd want to do batch deletes of plans in isolation? (seems functionality with low usefulness IMO given delete API, ClientApp & Org)	0	0	https://issues.jboss.org/browse/APIMAN-1018
665	APIMAN-1012	[~eric.wittmann], [~msavy] I've created an initial implementation of this policy here: https://github.com/outofcoffee/apiman-plugins/tree/feature/url-whitelist-plugin The policy normalises the incoming URL, then applies the configured rules, based on URL and HTTP method to determine if the request is permitted to be passed-on to the back-end API. Please see the {{url-whitelist-policy}} directory for implementation and tests. Would you be interested in including this in the main {{apiman-plugins}} repo, or would you prefer I keep it as a separate project?	1	1	https://issues.jboss.org/browse/APIMAN-1012
666	APIMAN-982	Can we rephrase this to read "uninstalling" plugins? Thanks!	0	0	https://issues.jboss.org/browse/APIMAN-982
667	APIMAN-973	We don't use a KC json config. When running in EAP/WF the adapters are configured in the standalone.xml file. Is there a way to configure the tomcat KC adapters in some sort of equivalent tomcat config file? If we can avoid requiring users to fiddle with the WAR that would be ideal.	1	0	https://issues.jboss.org/browse/APIMAN-973
668	APIMAN-961	Wouldn't it make more sense to move as much of apiman into WF module as possible? This would enable us to leverage WF's patching mechanism in the future if we need it.	1	1	https://issues.jboss.org/browse/APIMAN-961
669	APIMAN-942	Great! Would it be possible for you to create a pull request in github? https://help.github.com/articles/using-pull-requests/	1	1	https://issues.jboss.org/browse/APIMAN-942
670	APIMAN-931	Could you clarify how this works. My understanding is that we just need to compare the SOAPAction value to the allowed value(s) and see whether it matches one of those patterns, then go/no-go? One complication is that the SOAP spec seems to indicate the use of multiple headers with identical keys: https://www.w3.org/TR/2000/NOTE-SOAP-20000508/#_Toc478383528 - presumably it's possible to do multiple actions with one message? I'm hesitant to do loads of Map<String, List<String>> due to the overheads, but I guess we should profile that before worrying (allocation is cheap, after all). I can think of a few workarounds.	1	1	https://issues.jboss.org/browse/APIMAN-931

#	Requirement	Questions	Answered	Successfully Answered	URL
671	APIMAN-906	I'm a big +1 to creating new pages with explanations of what certain things mean. For example, creating a page to explain Centralized Security...and another page to explain rate limiting, etc. I've actually been meaning to create a better "How does apiman work?" page or series of pages. Basically something that will explain what the concepts are: API, Client App, Plan. I'm also good with adding a Resources page.	1	1	https://issues.jboss.org/browse/APIMAN-906
672	APIMAN-906	[~eric.wittmann] - Thanks for the quick response. What about breaking up the guides into multiple pages? Are those automatically generated? Can they be broken up? I'm referring to the Installation, User, Developer, and Production guides. Thanks!	0	0	https://issues.jboss.org/browse/APIMAN-906
673	APIMAN-892	Is it something difficult to acheive ? We need apiman to work with a recent version of ES and are considering doing it ourself.	1	0	https://issues.jboss.org/browse/APIMAN-892
674	APIMAN-876	I spent a while on this, and can't reproduce it on 1.2.x at all. I just realised after re-analysing it that this is probably a UI bug, and I think we've since fixed it. I can't find any data flows whereby the apiman-manager-api returns the user's credentials. I think this is the UI's AJAX request breaking somehow, and then displaying it in the error box in a manner that misleads us into thinking it's the backend! Investigating 1.1.9.Final to see if I can reproduce the bug, but seemingly not.	0	0	https://issues.jboss.org/browse/APIMAN-876
675	APIMAN-876	Console output seems to validate my theory {code} Failed to load resource: net::ERR_CONNECTION_REFUSED apiman-manager-ui.js?cid=2016-01-16_16:44:1629 apiman [18:30:33]>> Failed to refresh bearer token: { "readyState": 0, "status": 0, "statusText": "error" }error @ apiman-manager-ui.js?cid=2016-01-16_16:44:1629 apiman-manager-ui.js?cid=2016-01-16_16:44:437 ** Auth type was bearerToken but no bearer-token was found! ** deps.js?cid=2016-01-16_16:44:37465 PUT http://localhost:8080/apiman/gateways // Pressed test again when we no longer have a valid token net::ERR_CONNECTION_REFUSED(anonymous function) @ deps.js?cid=2016-01-16_16:44:37465sendReq @ deps.js?cid=2016-01-16_16:44:37258serverRequest @ deps.js?cid=2016-01-16_16:44:36965processQueue @ deps.js?cid=2016-01-16_16:44:41488(anonymous function) @ deps.js?cid=2016-01-16_16:44:41504Scope.$eval @ deps.js?cid=2016-01-16_16:44:42732Scope.$digest @ deps.js?cid=2016-01-16_16:44:42543Scope.$apply @ deps.js?cid=2016-01-16_16:44:42840(anonymous function) @ deps.js?cid=2016-01-16_16:44:50297jQuery.event.dispatch @ deps.js?cid=2016-01-16_16:44:72334elemData.handle @ deps.js?cid=2016-01-16_16:44:72020 apiman-manager-ui.js?cid=2016-01-16_16:44:437 ** Auth type was bearerToken but no bearer-token was found! ** deps.js?cid=2016-01-16_16:44:37465 PUT http://localhost:8080/apiman/gateways {code}	0	0	https://issues.jboss.org/browse/APIMAN-876
676	APIMAN-876	I've had an idea: How's about we add a filter to the remote call that regexes the returned JSON and replaces anything that looks vaguely like "Password: 1234" with Password: <REMOVED>". If that is too broad (there may be some cases where we actually want to display passwords in the UI that we've downloaded?), we can add an option to turn the filtering off. The main thrust of my point is that it seems sensible _default_ behaviour, as it will avoid future accidental displays of user credentials. It should only be turned off in cases where we are fully aware of the caveats, and hence have _explicitly_ disabled the obfuscation. Seem a reasonable option?	0	0	https://issues.jboss.org/browse/APIMAN-876
677	APIMAN-847	Yeah - I already added initial support for this in the community version. What form would you want the icon to take for f8 services? The community edition supports fa-* (font awesome) icons ATM, but I could support image URLs too I suppose. I'd need to be careful to scale a URL based icon appropriately, of course.	1	1	https://issues.jboss.org/browse/APIMAN-847
678	APIMAN-828	This is something that we have discussed internally and I think is a very good feature to have. Thanks for the feature request. The hardest thing about this feature, of course, is making sure we don't allow the same base URL for multiple APIs. :)	0	0	https://issues.jboss.org/browse/APIMAN-828
679	APIMAN-822	Where can we find this documentation?	1	0	https://issues.jboss.org/browse/APIMAN-822
680	APIMAN-795	[~eric.wittmann]- Do you mean the disabled endpoint field on the Service Implementation tab? Because for the Endpoint tab it looks like that field already has a copy-to-clipboard feature on it.	1	1	https://issues.jboss.org/browse/APIMAN-795
681	APIMAN-795	You get there by "searching for a service to consume" from the Application UI. Or by Find/Consume a service from the Home page.	0	0	https://issues.jboss.org/browse/APIMAN-795
682	APIMAN-790	[~ryordan] - can you add basic validation of the Service Endpoint in the UI? Just start out with ensuring that it starts with "http://" or "https://" . :)	0	0	https://issues.jboss.org/browse/APIMAN-790
683	APIMAN-772	Eric, why are you no longer showing if the serviceType, and the link to the Definition Document?	1	0	https://issues.jboss.org/browse/APIMAN-772
684	APIMAN-750	Can you try something where there is an arrow pointing @ the disabled Publish button with a link/button labeled something like "Why can't I publish?" That link/button could then open a dialog (or something else) with a comprehensive checklist of items required for publishing. That allows us to defer asking the server for the checklist data until the user actually wants to see the checklist. So we won't need to load that data for every tab, instead we do it on demand (on click of the "Why can't I publish?" button). I also like that because advanced users can just easily ignore it. Thoughts?	1	1	https://issues.jboss.org/browse/APIMAN-750
685	APIMAN-722	[~eric.wittmann] What do you think of this idea? Seems a reasonably simple solution, I think.	1	0	https://issues.jboss.org/browse/APIMAN-722
686	APIMAN-688	We can use different plans for different consumers. Eg. A Plan for us consumers and another plan for Japan consumers	0	0	https://issues.jboss.org/browse/APIMAN-688
687	APIMAN-683	What about create a tree with the versions and the correspondent status?	0	0	https://issues.jboss.org/browse/APIMAN-683
688	APIMAN-664	To be clear, the REST api for installing plugins has not changed - you can still use the REST API to install a plugin, but you need the plugin's maven GAV information to do this. The addition of plugin registry support is simply a *list* of possible plugins. You can still install a plugin that is not on the list. You can *retrieve* the list of available plugins via REST: {code} /apiman/plugins/availablePlugins {code} This simply returns a list of all plugins listed by all configured plugin registries. A consumer would need to invoke the above endpoint, then iterate through the results in some way. Once the desired plugin is found in the list, the GAV information for that plugin will be available. What additional registry REST API functionality would be useful?	1	1	https://issues.jboss.org/browse/APIMAN-664
689	APIMAN-652	The "next page" button *should* be adding activity items to the bottom of the list. So if you want to see previous entries you should be able to just scroll back up in your browser. Does that work?	0	0	https://issues.jboss.org/browse/APIMAN-652
690	APIMAN-652	[~ldimaggio], can you check the previous comment? Wondering if it works for you.	0	0	https://issues.jboss.org/browse/APIMAN-652
691	APIMAN-649	What would the level of detail mean? You mean whether to show the details or not?	1	1	https://issues.jboss.org/browse/APIMAN-649
692	APIMAN-647	You're supposed to see a "Load More" button at the end of the activity list if there are more items to see. Is that not working?	1	0	https://issues.jboss.org/browse/APIMAN-647
693	APIMAN-640	I am interested in Russian localization of APIMAN. I have discovered where are few i18n files in project. I have found html-files with "apiman-i18n-key" properties and TemplateScanner.java which generates messages.properties from html. I can't understand how messages.properties backward affect to html. Yes, I have registred at zanata.org. I have entered apiman in search field but found nothing. Is the project registered where. By the way. I'm ready to localize the project for Russian language. How can I start?	1	0	https://issues.jboss.org/browse/APIMAN-640

#	Requirement	Questions	Answered	Successfully Answered	URL
694	APIMAN-588	Is usage by contract different from usage by application? I suppose an application could have multiple contracts. FWIW, the data to do this query is easily available.	1	0	https://issues.jboss.org/browse/APIMAN-588
695	APIMAN-582	[~eric.wittmann]: Hey! For the multiple versions do you mean adding an extra column for the version, and an extra row for each unique version? So, something like this: https://drive.google.com/open?id=0B54p6afFhym-VUcxaU0tcDlLOXc Please excuse the hideous mockup, I don't have any design tools installed on this laptop. :D As for the multiple services an organization has, don't we already have a page/view for that with the path /api-manager/orgs/<org>/services?	1	1	https://issues.jboss.org/browse/APIMAN-582
696	APIMAN-544	{code} TypeError: Cannot read property 'length' of null at Object.compile (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:102059:22) at Object.JSONEditor.compileTemplate (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:95693:19) at Class.JSONEditor.AbstractEditor.Class.extend.getLink (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:96748:32) at Class.JSONEditor.AbstractEditor.Class.extend.addLinks (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:96705:29) at Class.JSONEditor.AbstractEditor.Class.extend.postBuild (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:96640:10) at Class.JSONEditor.defaults.editors.select.JSONEditor.AbstractEditor.extend.postBuild (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:100144:10) at Class.prototype.(anonymous function) [as postBuild] (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:95279:26) at http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:97881:16 at $each (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:95439:10) at Class.JSONEditor.defaults.editors.object.JSONEditor.AbstractEditor.extend.build (http://localhost:8080/apimanui/dist/deps.js?cid=2016-02-11_23:00:97875:7) {code}	1	1	https://issues.jboss.org/browse/APIMAN-544
697	APIMAN-544	[~msavy] - is this regressed? If so please open it again, assign to [~ryordan] and set to Blocker.	1	1	https://issues.jboss.org/browse/APIMAN-544
698	APIMAN-540	That's a really great start! The back-end bits should be easy enough to implement - you are well on your way. Have you thought about how best to display the data in the UI? What kind of chart(s) would you want to see? If you want to take a stab at implementing the UI as well, of course that would be most welcome. If that's not your cup of tea then I'm happy to finish that part off.	1	1	https://issues.jboss.org/browse/APIMAN-540
699	APIMAN-540	It's basically just a vertical bar chart, one resource per bar. I realized just recently that simply aggregating over resource doesn't really cut it because resource is stored with parameters, so if you have 1000000 users searched on a /users endpoint, this means 1000000 entries in the chart. I think there are two things an API owner might want to know: 1. Which parameters are hit the most on a certain endpoint (which users in /users ?) 2. Which endpoints are hit the most? The problem is actually filtering out the endpoint or the set of endpoints (which you get from swagger), especially with endpoints with path parameters. For example: /organizations/:organizationId/applications/:applicationId /organizations/organizationId/plans You'd have to filter by two substring tokens to recognise these two endpoints and aggregate over them. I'm pretty sure ES was never meant for this kind of queries. So at this point I am thinking that you'd need some kind of endpoint identifier field added to the database for easier filtering. This means you'd need to actually recognise the endpoint when it is accessed which further complicates the problem. As far as I know, Apiman uses swagger upload just for the reference and does not actually use that information for anything important.	1	1	https://issues.jboss.org/browse/APIMAN-540
700	APIMAN-538	Out of curiosity question. I vaguely remember [~ibek] telling me about discussion he had with you about the use of JRAPIDoc [1]. Was this idea dismissed or do I remember incorrectly? [1] https://github.com/projectodd/jrapidoc	1	0	https://issues.jboss.org/browse/APIMAN-538
701	APIMAN-516	I'm using AngularJS and calling the application with the simple call : $http.get($scope.service,{'headers': { 'Authorization': getBearerKey() }}); Is there something else i should send ?	1	1	https://issues.jboss.org/browse/APIMAN-516
702	APIMAN-516	You may need to add a CORS policy if you're doing a CORS request. Or do you want the service itself to service the CORS? I suggest you add and configure the CORS plugin policy for best functionality. It should be very easy. There may be an argument for passing CORS requests through to the back end but I'm not convinced that's a good idea.	1	1	https://issues.jboss.org/browse/APIMAN-516
703	APIMAN-516	Do you think this is why the calls not working for me ? I'll try to add that . Are you able to test your fix on http://fadiabdeen.github.io/keycloak-oauth/public_html/ ? This should be pretty easy , just type in the information on the keycloak.json file and and put the same url as your redirect url. you should be able to get a code and then a token , then test your service. When you click "TEST" , it will attach a header and send it , and this is where it sends only the preflight and doesn't finish .	1	1	https://issues.jboss.org/browse/APIMAN-516
704	APIMAN-516	I know i'm a pain :(... I think we are very close. I did not know there was a CORS policy ,, so for the first time i'm using it now. I added the CORS policy on both the 1.1.3 and your version 1.1.7-SNAPSHOT and gave it a shot . i only tried the older version because you said its your fix might be releated , but it didnt work. Your version 1.1.7-SNAPSHOT (before tiring down and rebuilding) gave me problems. XMLHttpRequest cannot load http://localhost:8080/apiman-gateway/express/test/4.0. No 'Access-Control-Allow-Origin' header is present on the requested resource. Origin 'http://fadiabdeen.github.io' is therefore not allowed access. The response had HTTP status code 403. even when i remove the Auth policy completely. Then i decided to change the "Terminate on CORS error" to false . and what i got is : XMLHttpRequest cannot load http://localhost:8080/apiman-gateway/express/test/5.0. Request header field Authorization is not allowed by Access-Control-Allow-Headers. I have the CORS policy setup same as your screen shot but it seems not to be reading the values : Access-Control-Allow-Origin : http://fadiabdeen.github.io, * Access-Control-Allow-Headers : Authorization The Access-Control-Allow-Credentials are false , i'm not passing username/password and i turned off the terminate on cors error That was my cors-fix branch setup from yesterday. should it not work ? When you said update from the latest and rebuild , you mean from master ?	1	1	https://issues.jboss.org/browse/APIMAN-516
705	APIMAN-516	It seems this issue still occurs in APIMan 1.2.3. Does anyone has the same issue?	1	0	https://issues.jboss.org/browse/APIMAN-516
706	APIMAN-516	It is the only policy you have in the whole chain? i.e. client + plan + api? In general you'd probably want CORS on the client (or API if you're using a public API). If it's still not working then you may have your CORS policy setup slightly wrong. Could you share it, please?	1	1	https://issues.jboss.org/browse/APIMAN-516
707	APIMAN-516	Can you paste the precise error you're seeing if that doesn't resolve your situation.	1	1	https://issues.jboss.org/browse/APIMAN-516
708	APIMAN-516	Maybe this sounds like a silly question, but should I have build a new war file and deploy it on my Wildfly? I'm using the standalone version of APIMan. Could I just remove my old version (1.2.3) and then deploy the new one? I'm using the standalone-apiman.xml config. So I'd like to know if when I deploy the new version this going to continue running fine, or should I need to change something?	0	0	https://issues.jboss.org/browse/APIMAN-516
709	APIMAN-503	Are you using Joda time? I've heard that helps. Also must be careful in date ranges as you mention above to be inclusive of all time within a day - a typical daily date range would be something like: 01-01-2015-00-00-00 to 01-01-2015-11-59-59....	1	1	https://issues.jboss.org/browse/APIMAN-503

#	Requirement	Questions	Answered	Successfully Answered	URL
710	APIMAN-470	Can you explain #1 - I don't understand what that means.	1	1	https://issues.jboss.org/browse/APIMAN-470
711	APIMAN-458	Not sure this is an intended use-case, but here's what I did: 1) Create new Apiman/WildFly 8.2 instance 2) Login as admin/admin123! 3) Create organization -> fails, no auto-roles defined 4) Edit Roles/Permissions 5) Create new role called "admin" with some description 6) Select "Auto-Grant Role" 7) Select all available permissions 8) Click "Create Role" Upon doing so, the default admin user has been updated and now has membership association with the admin role to include all admin role permissions. Is this intended? Or is the default "admin" user intended to be like a "super admin" with the expectation that end users of Apiman would be creating new "admin" users for administering organizations? With all that said, I'm trying to put into context the feature description above. Do you mean when creating new users and assigning them the admin role, by default those users should inherit "Publish" and "Register" in addition to "Edit"? It seems as if database pre-population scripts are missing that create default system roles and assign default permissions to those roles, or something along those lines.	1	1	https://issues.jboss.org/browse/APIMAN-458
712	APIMAN-445	Thanks for the pointers! Next question is which version you'd like to use. We strife to be very conservative with APIs, however Lucene doesn't always do the same, so the Hibernate Search version in WildFly 8 turns out to be slightly different than the one in WildFly 10. You can of course ignore the version from WildFly and include any other version.. that's particularly good in case we need to develop a new feature (unlikely but it's good to have the options). Personally I would expect you to move to WildFly 10 since it's stable now, so we could use the version provided by it which is nicely up to date (and will get long term product support), while also avoiding to bloat the download size of ApiMan with new dependencies. So that would be my preference, but I don't know if you can have "master" drop support for WildFly 8 now?	0	0	https://issues.jboss.org/browse/APIMAN-445
713	APIMAN-430	Wow it looks like I never released them! Sorry about that. I'll make sure that when we release 1.1.3.Final (soon) I won't forget the plugins. I am filled with self loathing now.	1	1	https://issues.jboss.org/browse/APIMAN-430
714	APIMAN-426	Hi [~eric.wittmann] [~msavy], I've been working on a CLI for apiman, so you can do things like this: Create a new service: {noformat} $ ./apiman service create \ --server http://localhost:8080/apiman \ --name example \ --endpoint http://example.com \ --initialVersion 1.0 \ --publicService \ --orgName test {noformat} Add a gateway: {noformat} $ ./apiman gateway create \ --server http://localhost:8080/apiman \ --name test-gw \ --endpoint http://localhost:1234 \ --username apimanager \ --password "apiman123!" \ --type REST {noformat} You can find the project at https://github.com/outofcoffee/apiman-cli It uses the management API documented at http://www.apiman.io/latest/api-manager-restdocs.html to perform common operations such as: Manage Gateways: {noformat} apiman gateway test [args...] apiman gateway show [args...] apiman gateway create [args...] apiman gateway list [args...] {noformat} Manage Plugins: {noformat} apiman plugin show [args...] apiman plugin create [args...] apiman plugin list [args...] {noformat} Manage Organisations: {noformat} apiman org show [args...] apiman org create [args...] {noformat} Manage Services: {noformat} apiman service create [args...] apiman service policy [add|remove] [args...] apiman service publish [args...] {noformat} Would you be interested in adopting this as part of the main project at some point?	1	0	https://issues.jboss.org/browse/APIMAN-426
715	APIMAN-426	Thanks [~stalep] - æsh looks interesting as a front-end to the CLI, as it provides completion, which is something that we don't get from Args4j. It also looks like æsh's annotation-based commands are similar to what we do at the moment, so hopefully it wouldn't be too much rework. Do you have an idea about how we can create hierarchical command trees using æsh, like we have right now? For example: {noformat} apiman api list [args] apiman api policy add [args] {noformat} Each subcommand makes things more specific. It looks like {{@GroupCommandDefinition}} might do this - does it work for >1 level deep?	1	0	https://issues.jboss.org/browse/APIMAN-426
716	APIMAN-426	so the gist i created show how to create a groupcommand as you saw. its not possible to go more than one level deep unfortunately, but i would assume that it wont be needed to use the "apiman" name if you start a apiman-cli? im thinking: $ ./apiman-cli.sh (some welcome text here if needed) [apiman] > *user press tab* service gateway org api .... [apiman] > ... not sure if i mange to show my point, but the idea is that we dont need the apiman anymore since we're running in an interactive cli. does that make sense? if so, it should be sufficient to have one level deep group commands as æsh supports i think.	0	0	https://issues.jboss.org/browse/APIMAN-426
717	APIMAN-421	Additional elements - the time period - if "day" can use an ID, why can't they all use id's? Thx! Second: //div[@id='form-page']/div[5]/div/div/div/div/div[2]/div/ul/li[2]/a/span Minute: //div[@id='form-page']/div[5]/div/div/div/div/div[2]/div/ul/li[3]/a/span Hour: //div[@id='form-page']/div[5]/div/div/div/div/div[2]/div/ul/li[4]/a/span Day: link=Day Month: //div[@id='form-page']/div[5]/div/div/div/div/div[2]/div/ul/li[6]/a/span Year: //div[@id='form-page']/div[5]/div/div/div/div/div[2]/div/ul/li[7]/a/span	0	0	https://issues.jboss.org/browse/APIMAN-421
718	APIMAN-408	Eric, I've been trying to find a way to convert the data in chunks, therefore avoiding the storage of the whole data in memory. Unfortunately I haven't found a way to do it. Do you have any suggestions?	1	1	https://issues.jboss.org/browse/APIMAN-408
719	APIMAN-408	In the policy's getResponseDataHandler method, I have the following line of code: final IBufferFactoryComponent bufferFactory = context.getComponent(IBufferFactoryComponent.class); When I run the test, I get the following error: {code} io.apiman.gateway.engine.beans.exceptions.ComponentNotFoundException: Component not found: io.apiman.gateway.engine.components.IBufferFactoryComponent at io.apiman.gateway.engine.impl.DefaultComponentRegistry.getComponent(DefaultComponentRegistry.java:87) at io.apiman.gateway.engine.policy.PolicyContextImpl.getComponent(PolicyContextImpl.java:81) at io.apiman.plugins.transformation_policy.TransformationPolicy.getResponseDataHandler(TransformationPolicy.java:95) {code} How do I register the component in my test set up?	1	1	https://issues.jboss.org/browse/APIMAN-408
720	APIMAN-368	Erm, so :-)... What are we thinking on this issue, Jakub?	0	0	https://issues.jboss.org/browse/APIMAN-368
721	APIMAN-353	The client will send a request with a query string parameter specifying the name of the callback function, e.g., "http://host/service/resource?callback=myCallbackFunctionName". Should the parameter name be parameterized? jQuery provides that flexibility (http://api.jquery.com/jQuery.ajax/ - jsonp). In case jQuery users don't specify a parameter name, it uses 'callback' by default. Should the callback function name parameter be sent to the endpoint? It will be of no use in the endpoint. We could remove it of just leave it there. I'm not sure it could cause issues.	1	1	https://issues.jboss.org/browse/APIMAN-353
722	APIMAN-347	is there any workaround right now ?	1	0	https://issues.jboss.org/browse/APIMAN-347
723	APIMAN-336	I don't see where the CORS policy uses the shared state component... Isn't the IPolicyContext sufficient for this? Specifically getAttribute() and setAttribute() are conversation-scoped (in-memory). The ISharedStateComponent is what should be used if the state needs to be shared across requests.	1	0	https://issues.jboss.org/browse/APIMAN-336

#	Requirement	Questions	Answered	Successfully Answered	URL
724	APIMAN-173	[~jorgemoralespou_2] - did you offer to take this on?	1	1	https://issues.jboss.org/browse/APIMAN-173
725	APIMAN-173	Great, thanks [~jorgemoralespou_2]. Does Keycloak support some sort of command line for creating users? I honestly think it doesn't. Not sure how to implement that last bullet point above.	1	1	https://issues.jboss.org/browse/APIMAN-173
726	APIVIZ-8	Does this _only_ exclude it from the overview summary? Does this also apply to subpackages?	0	0	https://issues.jboss.org/browse/APIVIZ-8
727	APIVIZ-7	This can probably be solved in Edge.java via spec.replaceAll("\\.\\s*\\n\\s+", "") Although this still is a bit of a hack.	0	0	https://issues.jboss.org/browse/APIVIZ-7
728	AF-564	Hi Matt, As we talk before, here are my thought on the subject: 1) Teams are a bad name, I rather prefer to use Organizations. Organizations are how name, it sets basically a namespace. Another thing that contribute to me don't like the Team terminology is that this will limit us in the future to create team in the scope of the Organization. 2) Embrace Organizations are fundamental for multi tenancy (cloud first) 3) Now the technninical `implications`: 3.1) If users can create a project without a `Organization` setup first, where should we store the project? If projects aren't bound to thier organizations, we'll limit ourselves to have one single project name in the whole installation (we'll end up with: hello-world1, hello-world11, hello-world2, hello-world133, hello-world13422, my-hello-world-12,...) 3.2) What exactly projects we'll list under the Teams? 3.3) What about pagination? We'll need paginate everything for sure (multi tenancy in mind)	1	0	https://issues.jboss.org/browse/AF-564
729	AF-561	[~mstevens1] Yes ! I really like option2. If that becomes too big, I would recommend to move some behind a more... link (or something similar) that would then open the wizard with all options listed there. But imho it significantly improves the experience for the most common asset types.	1	1	https://issues.jboss.org/browse/AF-561
730	AF-546	Hello [~mbarkley], did you do some benchmarking before/after these changes? I'd like to have an idea about how fast it was/is now for some example app. Also can we expect some kind of speedup for kie-wb? If so at what point (e.g. the initial loading of the app after we login etc.).	1	0	https://issues.jboss.org/browse/AF-546
731	AF-411	Hi [~ederign] we chatted about two thinigs. 1) the ability to expand multiple categories via API. This is now implemented by the addExpandedCategory(String... categoriesToExpand) method. and 2) the ability to have an event in order to be aware of the categories expand/collapse. Did you add this event too?	0	0	https://issues.jboss.org/browse/AF-411
732	AF-338	Hello [~paulovmr], do I understand it right that this wysiwyg editor is some third party library that is just being integrated at some places in our product? In such case QE would only have to create basic selenium integration test(s) - e.g. set some HTML using editor's source view and verify it's rendered correctly in the places where the edited HTML is used. Or is this editor more like part of the product, with some heavy customizations from our side? In such case we would also need to create some tests like checking that the editor's buttons for inserting images or tables in the visywyg mode are working correctly.	1	0	https://issues.jboss.org/browse/AF-338
733	AF-288	[~porcelli] Thanks. I made a new commit on the PR that uses more adequate images to illustrate the example on uberfire-wires: https://github.com/paulovmr/uberfire-extensions/blob/UF-288-master/uberfire-wires/uberfire-wires-webapp/src/main/resources/org/uberfire/ext/wires/client/resources/images/optaPlannerIconBlue.png https://github.com/paulovmr/uberfire-extensions/blob/UF-288-master/uberfire-wires/uberfire-wires-webapp/src/main/resources/org/uberfire/ext/wires/client/resources/images/optaPlannerIconWhite.png Do you think we could do anything else to try to prevent that to happen? In our case, I think the PR review should catch this kind of thing.	1	1	https://issues.jboss.org/browse/AF-288
734	AF-110	Ok this makes sense. I would anticipate that I will go with PicketLink and a derivative of the Errai's FormAuthenticationScheme. As for tracking this, should I open a ticket in Errai or Picketlink? Also, whats the thought on a release date for UberFire 0.5.0.Final?	1	0	https://issues.jboss.org/browse/AF-110
735	AF-49	Leverage work done by Catherine Robson: Section 508 Compliance - Developer Reference Guide: https://mojo.redhat.com/docs/DOC-175787 1) Prior to Uberfire release, Determine how to test, execute an internal audit of Uberfire. What tools will you use to automate this effort? 2) Documents results of internal audit, 508 non-compliance log bugs 3) Fix 508 non-compliance bugs	1	1	https://issues.jboss.org/browse/AF-49
736	JBAS-9571	Tomaz and Bilge, Should this move to the EAP 7 jira instead or should it also be opened in BZ? I know EAP 6 is basically done, so maybe this belongs somewhere else, where is that somewhere else? Jim Tyrrell Principal JBoss Solutions Architect Public Sector Middleware Advocate 720-839-2251 mobile	1	1	https://issues.jboss.org/browse/JBAS-9571
737	JBAS-9571	Good point on Undertow, do we need to have a test case validating this behavior doesn't occur in Wildly, how do we figure that out? Jim Tyrrell Principal JBoss Solutions Architect Public Sector Middleware Advocate 720-839-2251 mobile	1	1	https://issues.jboss.org/browse/JBAS-9571
738	JBAS-9526	I have very similar issue. +JBoss 5.1.0.+ During the load test with >200 users per instance each instance run into high CPU load >99%. The thread dump analysis and statistics for CPU usage per Application Threads revealed that most of CPU (>75%) is taken by org.apache.tomcat.util.net.JIoEndpoint$Worker threads. During the load test multiple new +JIoEndpoint$Worker threads+ were created, and +each took from 1 to 9% of CPU+. Surprisingly, most of JIoEndpoint$Worker threads which were on high CPU usage were in Waiting state, which doesn't makes much sense to me. To collect CPU per Application thread statistics I used the following linux command: {code:java} ps -mo pid,lwp,stime,time,%cpu,%mem -C java >>/shared/threads_per_cpu.$(date +"%F_%H-%M-%S") {code} Why JIoEndpoint$Worker Threads took so much of CPU? (Especially the Threads which were in Waiting state) Is there any way to workaround this? (e.g. configuration change or Hotfix) Dump example for the heavy CPU threads: {code:java} "http-0.0.0.0-8280-25" daemon prio=10 tid=0x00002b4f7d5ee800 nid=0x6c05 in Object.wait() [0x00002b4f9d64b000] java.lang.Thread.State: WAITING (on object monitor) at java.lang.Object.wait(Native Method) - waiting on <0x000000078f5c3298> (a org.apache.tomcat.util.net.JIoEndpoint$Worker) at java.lang.Object.wait(Object.java:485) at org.apache.tomcat.util.net.JIoEndpoint$Worker.await(JIoEndpoint.java:420) - locked <0x000000078f5c3298> (a org.apache.tomcat.util.net.JIoEndpoint$Worker) at org.apache.tomcat.util.net.JIoEndpoint$Worker.run(JIoEndpoint.java:446) at java.lang.Thread.run(Thread.java:662) {code}	1	0	https://issues.jboss.org/browse/JBAS-9526
739	JBAS-9472	I'm sorry for very slow reaction. We can understand your stance for this issue. But, you revised the implementation of Parameters class for JBoss AS 7.1.1 in AS7-3989. Could you tell us if you have a plan for back-porting it to JBoss AS 5.1.0.GA? Regards,	1	0	https://issues.jboss.org/browse/JBAS-9472
740	JBAS-9472	Hello, We have the same problem in production. Can you gives as a version of jboss modified with a patch ? Thx.	0	0	https://issues.jboss.org/browse/JBAS-9472

#	Requirement	Questions	Answered	Successfully Answered	URL
741	JBAS-9430	Skywalker:testsuite alesj$ ./build.sh one-test -Dtest=org.jboss.test.web.test.ActiveRequestResponseCacheValveUnitTestCase ./build.sh: line 105: ulimit: open files: cannot modify limit: Invalid argument build.sh: Could not set maximum file descriptor limit: unlimited Searching for build.xml ... Buildfile: /Users/alesj/projects/jboss6/trunk/testsuite/build.xml [echo] Using node0=localhost,node0.http.url=http://localhost:8080,node0.jndi.url=jnp://localhost:1099 [echo] Using node1=${env.HOSTNAME},node1.http.url=http://${env.HOSTNAME}:8080,node1.jndi.url=jnp://${env.HOSTNAME}:1099 [echo] Using java.net.preferIPv4Stack=true,java.net.preferIPv6Addresses=true check-env: init-maven-ant-tasks: init-dependencies: init-dependencies: [echo] Loading dependency paths from file: /Users/alesj/projects/jboss6/trunk/testsuite/output/build-dependencies.xml init-paths: server-config: init: one-test: [delete] Deleting: /Users/alesj/projects/jboss6/trunk/testsuite/output/log/test.log [junit] WARNING: multiple versions of ant detected in path for junit [junit] jar:file:/Users/alesj/jboss/tools/ant/lib/ant.jar!/org/apache/tools/ant/Project.class [junit] and jar:file:/Users/alesj/projects/jboss6/trunk/tools/ant/lib/ant.jar!/org/apache/tools/ant/Project.class [junit] Running org.jboss.test.web.test.ActiveRequestResponseCacheValveUnitTestCase [junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 0.536 sec BUILD SUCCESSFUL Total time: 4 seconds	1	1	https://issues.jboss.org/browse/JBAS-9430
742	JBAS-9426	build.sh: Could not set maximum file descriptor limit: unlimited Searching for build.xml ... Buildfile: /Users/alesj/projects/jboss6/trunk/testsuite/build.xml [echo] Using node0=localhost,node0.http.url=http://localhost:8080,node0.jndi.url=jnp://localhost:1099 [echo] Using node1=${env.HOSTNAME},node1.http.url=http://${env.HOSTNAME}:8080,node1.jndi.url=jnp://${env.HOSTNAME}:1099 [echo] Using java.net.preferIPv4Stack=true,java.net.preferIPv6Addresses=true check-env: init-maven-ant-tasks: init-dependencies: init-dependencies: [echo] Loading dependency paths from file: /Users/alesj/projects/jboss6/trunk/testsuite/output/build-dependencies.xml init-paths: server-config: init: test: [delete] Deleting: /Users/alesj/projects/jboss6/trunk/testsuite/output/log/test.log [junit] WARNING: multiple versions of ant detected in path for junit [junit] jar:file:/Users/alesj/jboss/tools/ant/lib/ant.jar!/org/apache/tools/ant/Project.class [junit] and jar:file:/Users/alesj/projects/jboss6/trunk/tools/ant/lib/ant.jar!/org/apache/tools/ant/Project.class [junit] Running org.jboss.test.jmx.test.AnnotatedJMXUnitTestCase [junit] Tests run: 3, Failures: 0, Errors: 0, Time elapsed: 0.575 sec [junit] Running org.jboss.test.jmx.test.AttributesUnitTestCase [junit] Tests run: 5, Failures: 0, Errors: 0, Time elapsed: 0.435 sec [junit] Running org.jboss.test.jmx.test.CPManifestUnitTestCase [junit] Tests run: 4, Failures: 0, Errors: 0, Time elapsed: 4.077 sec [junit] Running org.jboss.test.jmx.test.DeployJarUnitTestCase [junit] Tests run: 3, Failures: 0, Errors: 0, Time elapsed: 0.482 sec [junit] Running org.jboss.test.jmx.test.DeployServiceUnitTestCase [junit] Tests run: 12, Failures: 0, Errors: 0, Time elapsed: 2.167 sec [junit] Running org.jboss.test.jmx.test.DeployXMBeanUnitTestCase [junit] Tests run: 11, Failures: 0, Errors: 2, Time elapsed: 2.102 sec [junit] Test org.jboss.test.jmx.test.DeployXMBeanUnitTestCase FAILED [junit] Running org.jboss.test.jmx.test.EarDeploymentUnitTestCase [junit] Tests run: 2, Failures: 0, Errors: 0, Time elapsed: 1.494 sec [junit] Running org.jboss.test.jmx.test.InterceptableUnitTestCase [junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 0.526 sec [junit] Running org.jboss.test.jmx.test.JMXConnectorUnitTestCase [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=0 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=1 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=2 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=3 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=4 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=5 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=6 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=7 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=8 [junit] handleNotification, event: javax.management.Notification[source=jboss.test:service=InvokerTest][type=InvokerTest.timer][message=], count=9 [junit] Tests run: 8, Failures: 0, Errors: 0, Time elapsed: 9.813 sec [junit] Running org.jboss.test.jmx.test.JMXConsoleUnitTestCase [junit] Tests run: 3, Failures: 0, Errors: 0, Time elapsed: 2.788 sec [junit]	1	1	https://issues.jboss.org/browse/JBAS-9426
743	JBAS-9425	where did you get the download file for hibernate 3.6.6-final. Please comment Thanks Chris	1	0	https://issues.jboss.org/browse/JBAS-9425
744	JBAS-9164	Jozef, what version are you referring to? 2.1.0.GA is the only GA version I see released which is currently in AS 6. Can you be more specific what version you are requesting?	1	1	https://issues.jboss.org/browse/JBAS-9164
745	JBAS-8948	Thanks for the quick reply! Yep, that workaround works for us. The only problem with this was that it requires us to modify all our deployments which would also be required if we went with the jboss-scanning.xml solution. But this is doable. Do you think adding a configuration property is a valid enhancement that you might consider adding to a future version? Thanks, /Daniel	1	1	https://issues.jboss.org/browse/JBAS-8948
746	JBAS-8948	{quote} Do you think adding a configuration property is a valid enhancement that you might consider adding to a future version? {quote} Yeah, that looks reasonable. I'll see what needs to be done to add that support for 6.1.0.	1	0	https://issues.jboss.org/browse/JBAS-8948
747	JBAS-8948	Daniel, I've committed a fix for JBAS-8318. Can you try this application against the latest snapshot containing this fix and see if it helps?	1	0	https://issues.jboss.org/browse/JBAS-8948
748	JBAS-8939	Marius, one obvious thing that comes to mind about this solution is that it will mean that anything that relies on this new feature will not be able to run on earlier versions of AS6, right? In the case of SwitchYard, that's probably going to be an issue (need to check that). I might have a look ar the current Weld war deployer + JndiBinder class and see can we replicate that for SwitchYard. If it works, can we make sure that the resolution of this issue checks that the BeanManger is not already bound in.	0	0	https://issues.jboss.org/browse/JBAS-8939

#	Requirement	Questions	Answered	Successfully Answered	URL
749	JBAS-8745	This is done. metadata-deployer-jboss-beans.xml now has the following config options <bean name="SchemaResolverConfig" class="org.jboss.xb.binding.sunday.unmarshalling.SchemaResolverConfig"> + <property name="schemaInitializerInstances"> + <map keyClass="java.lang.String" valueClass="java.lang.String"> + <entry> + <key>http://java.sun.com/xml/ns/javaee</key> + <value> + <!-- Every specified property in this schema initializer + will override the value of the corresponding SchemaBinding property + which was initialized from binding annotations. + Properties that are commented will not affect the corresponding SchemaBinding properties. --> + <bean class="org.jboss.xb.util.SchemaBindingPropertyInitializer"> + <!-- must all content have a valid binding + <property name="strictSchema">true</property> --> + <!-- Should ${x} references be replaced with x system property + <property name="replacePropertyRefs">true</property> --> + <!-- if all the characters in the mixed content are whitespaces + should they be considered indentation and ignored? + <property name="ignoreWhitespacesInMixedContent">true</property> --> + <!-- whether to trim string values + <property name="normalizeSpace">true</property> --> + </bean> + </value> + </entry> + </map> + </property>	0	0	https://issues.jboss.org/browse/JBAS-8745
750	JBAS-8613	The following tests run in Hudson environment were failing when because the lookup was resolving to localhost:1099 instead of honoring the -Dnode0 IP address passed in to the tests. Specifically, org.jboss.test.cluster.defaultcfg.web.test.PersistentManagerFormAuthTestCase org.jboss.test.cluster.ejb3.clusteredsession.ejbthree1136.unit.ClusteredCacheCleanStartUnitTestCase These tests extend JBossTest* instead of JBossTestCluster* which results in: Stacktrace javax.naming.CommunicationException: Could not obtain connection to any of these urls: localhost:1099 and discovery failed with error: javax.naming.CommunicationException: Receive timed out [Root exception is java.net.SocketTimeoutException: Receive timed out] at org.jnp.interfaces.NamingContext.discoverServer(NamingContext.java:1690) at org.jnp.interfaces.NamingContext.checkRef(NamingContext.java:1761) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:695) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:688) at javax.naming.InitialContext.lookup(InitialContext.java:392) at org.jboss.test.JBossTestServices.getServer(JBossTestServices.java:154) at org.jboss.test.JBossTestServices.invoke(JBossTestServices.java:228) at org.jboss.test.JBossTestCase.invoke(JBossTestCase.java:240) at org.jboss.test.cluster.ejb3.clusteredsession.ejbthree1136.unit.ClusteredCacheCleanStartUnitTestCase.getExtraneousData(ClusteredCacheCleanStartUnitTes return server; }	1	1	https://issues.jboss.org/browse/JBAS-8613
751	JBAS-8610	Due to feedback from the community, I have split the AS7 and AS6 projects and reopened all unscheduled AS6 issues that are a year or less old. This will make it easier community members to find and work on them. Future releases beyond 6.1 can be done provided a community member steps up to coordinate them.	0	0	https://issues.jboss.org/browse/JBAS-8610
752	JBAS-8559	Adding the details on what exactly is required out of this task. The JavaEE 6 spec (section EE.5.17) states that: <quote> A DataSource resource may be defined in a web module, EJB module, application client module, or application deployment descriptor using the datasource element. For example: <data-source> <description>Sample DataSource definition</description> <name>java:app/MyDataSource</name> <class-name>com.foobar.MyDataSource</class-name> <server-name>myserver.com</server-name> <port-number>6689</port-number> <database-name>myDatabase</database-name> <user>lance</user> <password>secret</password> <property> <name>Property1</name> <value>10</value> </property> <property> <name>Property2</name> <value>20</value> </property> <login-timeout>0</login-timeout> <transactional>false</transactional> <isolation-level>TRANSACTION_READ_COMMITTED</isolation-level> <initial-pool-size>0</initial-pool-size> <max-pool-size>30</max-pool-size> <min-pool-size>20</min-pool-size> <max-idle-time>0</max-idle-time> <max-statements>50</max-statements> </data-source> A DataSource resource may also be defined using the DataSourceDefinition annotation on a container-managed class, such as a servlet or enterprise bean class. For example: @DataSourceDefinition(name="java:app/MyDataSource", className="com.foobar.MyDataSource", portNumber=6689, serverName="myserver.com", user="lance", password="secret") (Of course, we do not recommend including passwords to production systems in the code, but it's often useful while testing. Passwords, or other parts of the DataSource definition, can be overridden by a deployment descriptor when the application is deployed.) </quote> So all that needs to be done is to process this Java EE spec defined datasource deployment descriptor element and the corresponding annotation and create a JBoss specific datasource and bind it to JNDI. That's it.	0	0	https://issues.jboss.org/browse/JBAS-8559
753	JBAS-8481	Due to feedback from the community, I have split the AS7 and AS6 projects and reopened all unscheduled AS6 issues that are a year or less old. This will make it easier community members to find and work on them. Future releases beyond 6.1 can be done provided a community member steps up to coordinate them.	1	1	https://issues.jboss.org/browse/JBAS-8481
754	JBAS-8468	Richard, this should be Resolved with this commit: http://fisheye.jboss.com/changelog/JBossAS/?cs=108309 Can you confirm?	1	1	https://issues.jboss.org/browse/JBAS-8468
755	JBAS-8438	There is still more work to do. Can this issue be moved to the next release?	0	0	https://issues.jboss.org/browse/JBAS-8438
756	JBAS-8409	We dropped support of AS6 in jbosgi. Our current implementation targets AS7 M1 and is based on MSC. IMHO, it'd be sufficient to remove the osgi profile and delete the test. Is that acceptable?	1	1	https://issues.jboss.org/browse/JBAS-8409
757	JBAS-8400	There is a separate plugin for jbcache, jopr-jboss-cache-v3-plugin-3.0.0.B06.jar. So it may be that we'd just need to add an Infinispan plugin. Is there one?	1	1	https://issues.jboss.org/browse/JBAS-8400
758	JBAS-8399	This is probably not the best solution to deal with instantiation. The original idea was to have a BeanInstantiator registered per-component (i.e. EJB) in order for the insantiator to know what EJB is instantiated (and retrieve the appropriate AnnotatedType in Weld). But WAR deployments containing EJBs share the same java:comp context (and with java:module) so two EJBs will share the same BeanInstantiator.	0	0	https://issues.jboss.org/browse/JBAS-8399
759	JBAS-8371	Done in weld-snapshot branch, Marius, can you merge to trunk?	1	1	https://issues.jboss.org/browse/JBAS-8371
760	JBAS-8350	Daniel, questions like these are more appropriate in the user forums http://community.jboss.org/. For this specific question, the JCA user forum is appropriate http://community.jboss.org/en/jbossas/jca?view=discussions. Could you please create a topic there?	0	0	https://issues.jboss.org/browse/JBAS-8350
761	JBAS-8315	Stan, what is the deadline for this? Dan, any updates?	1	1	https://issues.jboss.org/browse/JBAS-8315
762	JBAS-8268	I think, this made to 6.0.0.Final, didn't it?	1	1	https://issues.jboss.org/browse/JBAS-8268

#	Requirement	Questions	Answered	Successfully Answered	URL
763	JBAS-8249	classloader of the deployer that loaded it. As explained previously, adding these extensions directly to the Deployment Unit's CL resolves the problem. [1] 12:45:57,716 INFO [org.jboss.profileservice.management.upload.remoting.DeployHandler] invoke, payload: {DeploymentTargetID=names=[org.jboss.weld.tests.producer.method.ManagerProducerTest.war], clientAddress=/127.0.0.1}, parameter: remove 12:45:57,717 WARN [org.jboss.profileservice.management.upload.remoting.DeployHandler] Failed to complete command: [remove] for deployment: names=[org.jboss.weld.tests.producer.method.ManagerProducerTest.war]: org.jboss.profileservice.spi.NoSuchDeploymentException: org.jboss.weld.tests.producer.method.ManagerProducerTest.war at org.jboss.profileservice.plugins.deploy.AbstractDeployHandler.resolveProfileForDeployment(AbstractDeployHandler.java:291) [:0.1.0.Alpha1] at org.jboss.profileservice.plugins.deploy.AbstractDeployHandler.resolveProfiles(AbstractDeployHandler.java:248) [:0.1.0.Alpha1] at org.jboss.profileservice.plugins.deploy.AbstractDeployHandler.removeDeployments(AbstractDeployHandler.java:201) [:0.1.0.Alpha1] at org.jboss.profileservice.management.upload.remoting.DeployHandlerDelegate.removeDeployments(DeployHandlerDelegate.java:84) [:6.0.0-SNAPSHOT] at org.jboss.profileservice.management.upload.remoting.DeployHandler.invoke(DeployHandler.java:156) [:6.0.0-SNAPSHOT] at org.jboss.remoting.ServerInvoker.invoke(ServerInvoker.java:898) [:] at org.jboss.remoting.transport.socket.ServerThread.completeInvocation(ServerThread.java:791) [:] at org.jboss.remoting.transport.socket.ServerThread.processInvocation(ServerThread.java:744) [:] at org.jboss.remoting.transport.socket.ServerThread.dorun(ServerThread.java:548) [:] at org.jboss.remoting.transport.socket.ServerThread.run(ServerThread.java:234) [:] 12:46:16,443 INFO [org.jboss.profileservice.management.upload.remoting.DeployHandler] invoke, payload: {DeploymentTargetID=names=[org.jboss.weld.tests.producer.method.ManagerProducerTest.war], clientAddress=/127.0.0.1}, parameter: start 12:46:19,071 INFO [org.jboss.weld.Version] WELD-000900 1.1.0 (SNAPSHOT) Listening for transport dt_socket at address: 5005 12:46:45,002 ERROR [org.jboss.kernel.plugins.dependency.AbstractKernelController] Error installing to Start: name=vfs:///opt/jboss/jboss-6.0.0-SNAPSHOT/server/default/deploy/org.jboss.weld.tests.producer.method.ManagerProducerTest.war_WeldBootstrapBean state=Create: org.jboss.weld.exceptions.DefinitionException: Exception List with 1 exceptions: Exception 0 : com.google.common.collect.ComputationException: org.jboss.weld.exceptions.WeldException: org.jboss.weld.bean.proxy.BeanInstance at com.google.common.collect.MapMaker$StrategyImpl.compute(MapMaker.java:602) at com.google.common.collect.MapMaker$StrategyImpl.compute(MapMaker.java:462) at com.google.common.collect.CustomConcurrentHashMap$ComputingImpl.get(CustomConcurrentHashMap.java:2045) at org.jboss.weld.bean.proxy.ClientProxyProvider.getClientProxy(ClientProxyProvider.java:104) at org.jboss.weld.manager.BeanManagerImpl.getReference(BeanManagerImpl.java:650) at	0	0	https://issues.jboss.org/browse/JBAS-8249
764	JBAS-8244	org.jboss.test.ws.jaxws.ejb3Integration.WebServiceTestCase Tests run: 5, Failures: 0, Errors: 1, Time elapsed: 6.605 sec Testcase: testRemoteAccess took 0.313 sec Testcase: testWebService took 2.763 sec Testcase: testWebServiceContext took 0.107 sec Testcase: testWebServiceRef took 0.236 sec Caused an ERROR java.lang.RuntimeException: Unable to inject jndi dependency: env/org.jboss.test.ws.jaxws.ejb3Integration.WebServiceRefBean/serviceRef into property org.jboss.test.ws.jaxws.ejb3Integration.WebServiceRefBean.serviceRef: Could not find service named null in wsdl http://localhost:8080/webservices-ejb3/SimpleEndpoint?wsdl javax.ejb.EJBException: java.lang.RuntimeException: Unable to inject jndi dependency: env/org.jboss.test.ws.jaxws.ejb3Integration.WebServiceRefBean/serviceRef into property org.jboss.test.ws.jaxws.ejb3Integration.WebServiceRefBean.serviceRef: Could not find service named null in wsdl http://localhost:8080/webservices-ejb3/SimpleEndpoint?wsdl at org.jboss.ejb3.tx2.impl.CMTTxInterceptor.handleExceptionInOurTx(CMTTxInterceptor.java:183) at org.jboss.ejb3.tx2.impl.CMTTxInterceptor.invokeInOurTx(CMTTxInterceptor.java:251) at org.jboss.ejb3.tx2.impl.CMTTxInterceptor.required(CMTTxInterceptor.java:349) at org.jboss.ejb3.tx2.impl.CMTTxInterceptor.invoke(CMTTxInterceptor.java:209) at org.jboss.ejb3.tx2.aop.CMTTxInterceptorWrapper.invoke(CMTTxInterceptorWrapper.java:52) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.tx.TxPropagationInterceptor.invoke(TxPropagationInterceptor.java:76) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.ejb3.tx.NullInterceptor.invoke(NullInterceptor.java:42) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.ejb3.security.Ejb3AuthenticationInterceptorv2.invoke(Ejb3AuthenticationInterceptorv2.java:186) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.ejb3.ENCPropagationInterceptor.invoke(ENCPropagationInterceptor.java:41) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.ejb3.BlockContainerShutdownInterceptor.invoke(BlockContainerShutdownInterceptor.java:67) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.ejb3.core.context.CurrentInvocationContextInterceptor.invoke(CurrentInvocationContextInterceptor.java:47) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.currentinvocation.CurrentInvocationInterceptor.invoke(CurrentInvocationInterceptor.java:67) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.ejb3.interceptor.EJB3TCCLInterceptor.invoke(EJB3TCCLInterceptor.java:86) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at	1	1	https://issues.jboss.org/browse/JBAS-8244
765	JBAS-8244	Alessio, can you comment on this failure?	0	0	https://issues.jboss.org/browse/JBAS-8244
766	JBAS-8230	Marius, could you try the upgrade that Nik requests in AS? We will update Weld for 1.1, but perhaps we can get the update into AS sooner.	0	0	https://issues.jboss.org/browse/JBAS-8230

#	Requirement	Questions	Answered	Successfully Answered	URL
767	JBAS-8169	If we want to make a service accessible through CORBA 1) We create a standard EJB service using a standard invoker (not IIOP, no IOR Security config). 2) We create the public interfaces accessible through CORBA in an independent IDL file 3) We create the necessary java objects using IDLJ plug-in. 4) We implement the servant by extending the POA object generated by IDLJ (This servant object will then make (relay) a CORBA IIOP call to a standard EJB) 5) We add a listener to our EJB application that will register this CORBA object at deployment using the following schema. // initialisation InitialContext ctx = new InitialContext(); NamingContextExt rootNC = (NamingContextExt)ctx.lookup("java:JBossCorbaNaming"); POA poa = (POA) ctx.lookup("java:JBossCorbaPOA"); ORB orb = (ORB)ctx.lookup("java:JBossCorbaORB"); // using reflexion, we create a IorSecurityConfigMetaData with our custom realm IorSecurityConfigMetaData data = new IorSecurityConfigMetaData(); Class asConfigClass = data.getAsContext().getClass(); Class sasConfigClass = data.getSasContext().getClass(); Field realm = asConfigClass.getDeclaredField("realm"); realm.setAccessible(true); realm.set(data.getAsContext(),"weblogicDEFAULT"); Field propagation = sasConfigClass.getDeclaredField("callerPropagation"); propagation.setAccessible(true); propagation.set(data.getSasContext(),"SUPPORTED"); Any secPolicy = orb.create_any(); secPolicy.insert_Value(data); Policy csiv2Policy = orb.create_policy(CSIv2Policy.TYPE, secPolicy); Policy[] policies = new Policy[]{csiv2Policy}; poa.the_POAManager().activate(); // Create a secure child POA POA childPOA = null; try{ childPOA = poa.find_POA("csiv2POA",true); }catch(org.omg.PortableServer.POAPackage.AdapterNonExistent e){ logger.info("csiv2POA do not exist => Beeing created"); childPOA = poa.create_POA("csiv2POA", poa.the_POAManager(), policies); } // register the servant in the secure POA and activate it childPOA.activate_object(servant); childPOA.the_POAManager().activate(); // add the service in the naming directory org.omg.CORBA.Object o = childPOA.servant_to_reference(servant); createContext(nsPath, rootNC); rootNC.rebind(rootNC.to_name(nsPath), o); As you can see, the way we register CORBA objects do not allow us to use the IOR configuration available in the deployment descriptor. The code presented before allows us to register CORBA services with our custom realm and to change the callerPropagation setting. Nevertheless, we would like to avoid reflexion. We propose you to either give the possibility to set these values through a configuration XML file or by making the constructor of IorSecurityConfigMetaData to public. Furthermore, we remarked that the CSIV2 security settings are not propagated to the EJB in this configuration although the CSIV2 security context is correctly intercepted by org.jboss.iiop.csiv2.SASTargetInterceptor . Why don't you populate the SecurityAssocation variable with these intercepted values? If this is a desired feature, how can we access to the CSIV2 security context in our CORBA servant in order to forward them to the EJB. It is realy important for us to be able to use this architecture to register CORBA services and that the security context is correctly propagated.	1	1	https://issues.jboss.org/browse/JBAS-8169
768	JBAS-8166	The hudson run JBoss-AS-6.0.x-testSuite-sun16 (Build #801 (Jul 15, 2010 2:26:27 AM) ran with Remoting 2.5.3. There are three failures that didn't occur in the previous run that went to completion: * org.jboss.test.cluster.defaultcfg.simpleweb.test.SessionCountUnitTestCase(Default-udp).testMarshalledRedeployWarAndCache * org.jboss.test.profileservice.test.Ejb3MetricsUnitTestCase(profileservice).testSlsbLocal * org.jboss.test.xml.DDValidatorUnitTestCase.testJBossXml They don't appear to have any relationship to Remoting, so I'm closing the case.	0	0	https://issues.jboss.org/browse/JBAS-8166
769	JBAS-8159	Thanks for the 5.x patch, that will be useful to others also using 5.x. The JSR-160 connector can be secured already in 6.x (see umbrella jira JBAS-7818), but if you have improvements (for AS 6) to what we did, that is always welcome. What do think?	0	0	https://issues.jboss.org/browse/JBAS-8159
770	JBAS-8159	Question: How complete is the patch for securing the JSR-160 (example) code in AS 5? I am wondering if its complete or known to need further work (as a solution to patch the current JSR-160 example code in AS 5). In addition to looking at any new proposals for securing JSR-160 support, we should look at merging any parts of this contribution to AS-7 release (that make sense). Marking as fix for AS 7 to support that separate effort.	0	0	https://issues.jboss.org/browse/JBAS-8159
771	JBAS-8148	Just a reminder - M4 is planned to be tagged on the coming Wednesday (July 21). Looking at the AS trunk, these files seem to come from HornetQ distribution (although I don't know why, since these configs should perhaps be in AS). So that means, this will require a HornetQ release before 21st. I believe this is a good improvement to have in terms of boot time, so IMO should make it to M4.	0	0	https://issues.jboss.org/browse/JBAS-8148
772	JBAS-8143	Why there is a JCA task separated from the server? All the security events are done at the Server's Session. (I.e. I'm not sure there's any need for a JCA specific part).	1	0	https://issues.jboss.org/browse/JBAS-8143
773	JBAS-8105	What happened to conf/jbossts-properties.xml in trunk, can't find it?!	0	0	https://issues.jboss.org/browse/JBAS-8105
774	JBAS-8105	Clebert, could you comment on Robert's entry whether that would be the proper way to configure it manually? Thanks in advance - Rico.	1	0	https://issues.jboss.org/browse/JBAS-8105
775	JBAS-8105	Well, for JBoss 6.0.0 this is definitely not the case. So you would say that here, Robert's configuration is correct? Also for 6.1.0 or even 7.x I cannot see that this configuration has made it into the general distribution - can you point me to it in case you found it? At least I think HornetQ's documentation should reflect the necessary configuration for JBoss 6 and later, instead of only for JBoss5 at is atm.	0	0	https://issues.jboss.org/browse/JBAS-8105
776	JBAS-8099	This probably won't happen until AS7, but I don't have a target release version yet. Very minor task.	1	1	https://issues.jboss.org/browse/JBAS-8099
777	JBAS-8093	INFO in that case You might want to have web applications that aren't CDI enabled (and therefore you left out beans.xml). But isn't this a chicken-egg syndrome since if you left out beans.xml, Weld isn't booting (not going the weld-deployers route) and who does the warning in that case?	0	0	https://issues.jboss.org/browse/JBAS-8093
778	JBAS-8093	Underlying issue is that if CDI annotations are present but beans.xml isn't then the AS should complain. However I don't think the AS can do this right now? Maybe in AS7? Jason, could you let me know what the best course of action is.	1	1	https://issues.jboss.org/browse/JBAS-8093
779	JBAS-8078	Tim Fox wrote: The first time you boot HQ it will take longer, since it has to create the journal files. On subsequent boots it will be faster since it just has to open them. I didn't measure first run (heat phase). Tim Fox wrote: You can reduce the number of initial files in the config (see user manual) if you want to speed this up. Why isn't it default in AS?	1	0	https://issues.jboss.org/browse/JBAS-8078
780	JBAS-8061	I'm not sure this is working in M4 -- I'm still getting: 15:11:21,069 SEVERE [javax.enterprise.resource.webcontainer.jsf.application] JSF1005: Cannot instantiate validator of type integerValidator When this class is on the classpath, and I try to reference it in a Facelet via the <f:validator validatorId="integerValidator" /> tag: ---------- ---------- package com.ocpsoft.socialpm.validator; import javax.faces.application.FacesMessage; import javax.faces.component.UIComponent; import javax.faces.context.FacesContext; import javax.faces.validator.FacesValidator; import javax.faces.validator.Validator; import javax.faces.validator.ValidatorException; @FacesValidator("integerValidator") public class IntegerValidator implements Validator { @Override public void validate(final FacesContext context, final UIComponent component, final Object value) throws ValidatorException { //... } }	1	1	https://issues.jboss.org/browse/JBAS-8061

#	Requirement	Questions	Answered	Successfully Answered	URL
781	JBAS-8061	How are you deploying your validator? Is it in a jar or in WEB-INF/classes? If it's a jar, please attach.	1	1	https://issues.jboss.org/browse/JBAS-8061
782	JBAS-8061	Some annotations are found but others are not. Need to investigate why.	1	1	https://issues.jboss.org/browse/JBAS-8061
783	JBAS-8055	Does this happen on AS startup? Or during an actual AS unit test? Which test?	1	1	https://issues.jboss.org/browse/JBAS-8055
784	JBAS-8026	All the nodes that are part of a cluster are supposed to use the same ports. How the service binding manager would help here? Also, I have a question about this phrase: "Currently the discovery group and the broadcast group in hornetq-configuration.xml in AS are used for configuring the group-address and the port: " What you meant by this statement? Are you affirming we are clashing with another port, or that was just informing what the port was for?	0	0	https://issues.jboss.org/browse/JBAS-8026
785	JBAS-8026	>> All the nodes that are part of a cluster are supposed to use the same ports. How the service binding manager would help here? The ServiceBindingManager (SBM) is the central place where all ports and group addresses are centrally configured. Each such address and/or port is part of a group known as ServiceBindingSet. The advantage of this approach is that end user/admins can just use a single resource (the SBM) to manage their ports. If they want to switch their ports, then that just switch the ServiceBindingSet group by passing a JVM system param (like in the example in the wiki). >> Also, I have a question about this phrase: >>"Currently the discovery group and the broadcast group in hornetq-configuration.xml in AS are used for configuring the group-address and the port: " >> What you meant by this statement? Are you affirming we are clashing with another port, or that was just informing what the port was for? Just informing what those ports are for.	0	0	https://issues.jboss.org/browse/JBAS-8026
786	JBAS-8026	on that case.. even with you inform -Dports=, those groups should still be the same... unless they form another group. I'm not sure how the ServiceBinding would help here	0	0	https://issues.jboss.org/browse/JBAS-8026
787	JBAS-8026	"SBM internally adds +100 (just an example) to the default port." This is exactly my point.... if you add another server on the same box... everything should be added +100 except those groups ports. Unless the user wants to make them part of a different cluster. In other words.. most of the same the user will want the same port for those two groups. Notice that the connectors and acceptors are using the binding manager	0	0	https://issues.jboss.org/browse/JBAS-8026
788	JBAS-8020	Dan, can you take care of this please?	0	0	https://issues.jboss.org/browse/JBAS-8020
789	JBAS-8016	Ron, can you work on the remoting docs?	0	0	https://issues.jboss.org/browse/JBAS-8016
790	JBAS-8015	Clebert, can you or someone from hornetQ work on the docs?	0	0	https://issues.jboss.org/browse/JBAS-8015
791	JBAS-8015	Can we have some context about what the purpose of this task is, what's expected of us etc? So far I was under the impression that AS docs would simply reference our existing HornetQ docs.	0	0	https://issues.jboss.org/browse/JBAS-8015
792	JBAS-8012	Anil, can you work on the AS6 docs for security?	0	0	https://issues.jboss.org/browse/JBAS-8012
793	JBAS-8008	Jason, can you take this take this task or re-assign?	0	0	https://issues.jboss.org/browse/JBAS-8008
794	JBAS-8008	I think that the twiddle JIRA (JBAS-8013) could probably be closed as duplicated by JBAS-8008. If no one disagrees, I'll merge my current twiddle changes in to the admin doc. Thinking more about it, we probably need several subtasks for JBAS-8008 (to allow multiple people to work on JBAS-8008). Does Jira allow subtasks for subtasks? We could then have a twiddle subtask.	1	1	https://issues.jboss.org/browse/JBAS-8008
795	JBAS-7988	Can you please check Paul if there is a flag to disable this update check?	1	1	https://issues.jboss.org/browse/JBAS-7988
796	JBAS-7970	Ales, is there a plan that backport this fix into EAP 5.1?	1	0	https://issues.jboss.org/browse/JBAS-7970
797	JBAS-7955	No idea, and as you can see, this was a while ago. I would have to look again. btw: this our jsr77 impl doesn't do much, mostly basic deployment structure What exactly is the issue or you're trying to do?	0	0	https://issues.jboss.org/browse/JBAS-7955
798	JBAS-7889	updated existing wiki pages: Remote bean access http://community.jboss.org/wiki/HowDoIGetRemoteAccessToMyMBean Which interface should I use for remote bean access http://community.jboss.org/wiki/WhichRemoteInterfaceShouldIUseForTheMBeanServer There is some jsr160 support in remoting (see jmx-remoting.sar in http://community.jboss.org/wiki/Remoting) twiddle page http://community.jboss.org/wiki/Twiddle shutdown page http://community.jboss.org/wiki/StartStopJBoss https://community.jboss.org/wiki/SecureTheJmxConsole jconsole page http://community.jboss.org/wiki/UseJDK5JConsole	0	0	https://issues.jboss.org/browse/JBAS-7889
799	JBAS-7857	Sorry for not getting back earlier on this. Does this upgrade still need to happen for AS6? I see both apache-xalan:xalan and xalan:xalan included in the AS6 dep tree, so probably the "apache-xalan" deps should be excluded. And the new patched version should be uploaded to the thirdparty-releases community repo.	1	0	https://issues.jboss.org/browse/JBAS-7857
800	JBAS-7830	Is there anything in jmx-invoker-service.xml that we should keep? :)	1	1	https://issues.jboss.org/browse/JBAS-7830
801	JBAS-7828	Should we do this still? What is the major advantage over jmx-jboss-beans.xml?	0	0	https://issues.jboss.org/browse/JBAS-7828
802	JBAS-7828	I see various examples of other services doing this, such as jboss-threads Should I create a separate subversion project, as you did for jboss-threads? Or should this live in an existing project?	0	0	https://issues.jboss.org/browse/JBAS-7828
803	JBAS-7818	Is this going to be scheduled on the remoting roadmap?	0	0	https://issues.jboss.org/browse/JBAS-7818
804	JBAS-7818	The OSGi remote testsuite relies on the RMIAdaptor. Recent changes made <groupId>org.jboss.jbossas</groupId> <artifactId>jboss-as-client</artifactId> <version>6.0.0.20100216-M2</version> unusable with jbossas/trunk. Could you please provide documentation on the intended migration path? * How should clients remotely connect to JMX in the future? * When is the existing access method intended to go out of service? I assume that access to the RMIAdaptor using the 6.0.0.20100216-M2 client libraries should continue to work for a while and that I am seeing a bug introduced by recent changes.	1	1	https://issues.jboss.org/browse/JBAS-7818
805	JBAS-7818	Scott, It appears that with the recent changes around this, the following lookup: MBeanServerConnection adaptor = (MBeanServerConnection) ctx.lookup("jmx/rmi/RMIAdaptor"); runs into issues (in EJB3 testsuite). This throws a: javax.naming.Reference cannot be cast to javax.management.MBeanServerConnection java.lang.ClassCastException: javax.naming.Reference cannot be cast to javax.management.MBeanServerConnection at org.jboss.test.JBossTestServices.getServer(JBossTestServices.java:154) at org.jboss.test.JBossTestServices.invoke(JBossTestServices.java:228) at org.jboss.test.JBossTestServices.undeploy(JBossTestServices.java:364) at org.jboss.test.JBossTestSetup.undeploy(JBossTestSetup.java:197) at org.jboss.test.JBossTestCase$1.tearDown(JBossTestCase.java:333) at junit.extensions.TestSetup$1.protect(TestSetup.java:24) at junit.extensions.TestSetup.run(TestSetup.java:27) The root cause appears to be the class org.jboss.system.server.jmx.JMXAdapter (present in JBOSS_HOME/lib/jboss-system.jar) missing from the client classpath (which uses the jars in JBOSS_HOME/client folder, especially jbossall-client.jar). Looking at the name of that class, i don't think it would be appropriate to include that in the client jar files.	1	1	https://issues.jboss.org/browse/JBAS-7818

#	Requirement	Questions	Answered	Successfully Answered	URL
806	JBAS-7818	Scott, is this work completed for AS 6.0.0? There is a problem starting AS 6 with IPv6; specifically, the error shown below. This prevents additional testing with IPv6 to proceed. Can you comment on this failure? 4:15:34,374 INFO [JMXConnector] starting JMXConnector on host 3ffe:ffff:100:f101::1:1090 14:15:34,393 ERROR [JMXConnector] Could not start JMXConnector: java.net.MalformedURLException: Bad port number: "": java.lang.NumberFormatException: For input string: "" at javax.management.remote.JMXServiceURL.<init>(JMXServiceURL.java:193) [:1.6.0_13] at org.jboss.system.server.jmx.JMXConnector.start(JMXConnector.java:143) [:6.0.0.20100721-M4 (Build SVNTag:JBoss_6.0.0.20100721-M4 date: 20100723)]	1	1	https://issues.jboss.org/browse/JBAS-7818
807	JBAS-7739	Stan, can you check the versions for jsf.deployer and jsf-impl in jbossas/trunk/component-matrix/pom.xml. There seems to be some discrepancy in the pom. From the JIRA, your comment is to upgrade the jsf.deployer to 1.0.3, but trunk contains 1.0.2: <version.jboss.jsf-deployer>1.0.2</version.jboss.jsf-deployer> http://fisheye.jboss.com/changelog/JBossAS/?cs=108115 There is also a dependency for jsf-impl; <dependency> <groupId>javax.faces</groupId> <artifactId>jsf-impl</artifactId> <version>${version.javax.faces}</version> </dependency> where version.javax.faces is: <version.javax.faces>2.0.2-FCS</version.javax.faces> The 1.0.0.Beta6 of org.jboss.javaee.specs.bom also defines the JSF API version as: <version.jsf.api>2.0.2-FCS</version.jsf.api>	1	1	https://issues.jboss.org/browse/JBAS-7739
808	JBAS-7735	As per http://community.jboss.org/wiki/JBossProjectVersioning 6.0.0.20100215-M2 is the official and correct name of the release. Any formal declaration of the version should use the full and correct name; IMO the logging above is a formal declaration of the version. The one semi-exception to this is the decision not to include the timestamp in the name of the zip files for the dist, as discussed at http://post-office.corp.redhat.com/archives/jboss-tag-public/2010-February/msg00032.html . I say semi-exception, because a name of a zip file isn't really a formal declaration of the version. Changing the directory name seems harmless enough, but doing it means hacking the build scripts to essentially manually set it. Is that worth doing?	0	0	https://issues.jboss.org/browse/JBAS-7735
809	JBAS-7735	Are you thinking more like (build: SVNTag=JBoss_6_0_0_20100215-M2 revision=r100927) ? I'll ping Paul Gier to see what's involved.	0	0	https://issues.jboss.org/browse/JBAS-7735
810	JBAS-7735	What happened to the old codenames (Trinity, Morpheus, Neo). Why they disappeared?	1	1	https://issues.jboss.org/browse/JBAS-7735
811	JBAS-7715	Dimitris - is this planned for 6.0.0.CR1? If javax.script is part of JDK 6 now, what is needed to be done in AS (remove the 2.4.0 dep?)?	1	0	https://issues.jboss.org/browse/JBAS-7715
812	JBAS-7706	On a remote client I get: Failed to initialise session factory javax.jms.JMSException: Failed to initialise session factory at org.hornetq.core.client.impl.ClientSessionFactoryImpl.createSessionInternal(ClientSessionFactoryImpl.java:1043) at org.hornetq.core.client.impl.ClientSessionFactoryImpl.createSession(ClientSessionFactoryImpl.java:790) at org.hornetq.jms.client.HornetQConnection.authorize(HornetQConnection.java:572) at org.hornetq.jms.client.HornetQConnectionFactory.createConnectionInternal(HornetQConnectionFactory.java:602) at org.hornetq.jms.client.HornetQConnectionFactory.createQueueConnection(HornetQConnectionFactory.java:128) at org.hornetq.jms.client.HornetQConnectionFactory.createQueueConnection(HornetQConnectionFactory.java:123) at org.jboss.ejb3.test.dd.mdb.unit.MDBUnitTestCase.getQueueSession(MDBUnitTestCase.java:129) at org.jboss.ejb3.test.dd.mdb.unit.MDBUnitTestCase.testQueue(MDBUnitTestCase.java:155) at org.jboss.ejb3.test.dd.mdb.unit.MDBUnitTestCase.testQueue(MDBUnitTestCase.java:270) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:23) at junit.extensions.TestSetup.run(TestSetup.java:27) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:23) at junit.extensions.TestSetup.run(TestSetup.java:27) Caused by: HornetQException[errorCode=0 message=Failed to initialise session factory] Caused by: java.lang.IllegalArgumentException: Error instantiating connector factory "org.hornetq.integration.transports.netty.NettyConnectorFactory" at org.hornetq.core.client.impl.FailoverManagerImpl.instantiateConnectorFactory(FailoverManagerImpl.java:1016) at org.hornetq.core.client.impl.FailoverManagerImpl.<init>(FailoverManagerImpl.java:228) at org.hornetq.core.client.impl.ClientSessionFactoryImpl.initialise(ClientSessionFactoryImpl.java:237) at org.hornetq.core.client.impl.ClientSessionFactoryImpl.createSessionInternal(ClientSessionFactoryImpl.java:1039) Caused by: java.lang.ClassNotFoundException: org.hornetq.integration.transports.netty.NettyConnectorFactory at java.net.URLClassLoader$1.run(URLClassLoader.java:200) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:188) at java.lang.ClassLoader.loadClass(ClassLoader.java:307) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:252) at org.hornetq.core.client.impl.FailoverManagerImpl.instantiateConnectorFactory(FailoverManagerImpl.java:1011)	0	0	https://issues.jboss.org/browse/JBAS-7706
813	JBAS-7703	Carlo, can you or one of your guys have a look at this? One thing I notice is ContainerDependencyMetaData has an unused jndiPolicy field, which if used would seem to be a hook for making a more sophisticated determination of the jndi name.	1	1	https://issues.jboss.org/browse/JBAS-7703
814	JBAS-7703	Do you have some time to make a servlet that does the same thing as simple.jsp, throw it in the war and invoke on that? If that fails, it's a more significant issue since as you say resource injection in JSP pages is not supported.	1	1	https://issues.jboss.org/browse/JBAS-7703
815	JBAS-7690	The produced Maven artifacts will just use whatever version is in the POM. The same goes for the base directory in the zip. According to the guidelines linked in the summary, the community release should be 6.0.0.Final. Dimitris, is there anything remaining to do on this issue?	1	1	https://issues.jboss.org/browse/JBAS-7690
816	JBAS-7689	Do we want all build artifacts to use that version string? I'm not sure if there is anything to change in the build for this. As long as the version is changed in the pom it will affect all the artifacts and the distribution zip.	0	0	https://issues.jboss.org/browse/JBAS-7689
817	JBAS-7661	Per discussion on jboss-dev, can we move this to a special osgi profile for now?	1	1	https://issues.jboss.org/browse/JBAS-7661

#	Requirement	Questions	Answered	Successfully Answered	URL
818	JBAS-7661	FYI, here a transcript of the chat that I had with Dimitris Friday 09 April 2010 ... [12:38:37] Thomas Diesler: The discussion is however relevant for the upcoming M3 release. ;-) [12:45:25] Thomas Diesler: the beauty about smoke-tests is that folks actually run them often [12:45:45] Dimitris Andreadis: i know, but your test target will also show up in hudson runs [12:45:45] Thomas Diesler: thats why I wanted osgi coverage in there out of the box ... [12:46:28] Thomas Diesler: again, whats your concern against osgi being part of 'all'? is it just boot time? [12:46:42] Dimitris Andreadis: yea, boot time ... [12:51:20] Dimitris Andreadis: I think the 'osgi' config is a good first step to show that "we do OSGi", too. And we can take it from there depending on user interest. [12:53:59] Thomas Diesler: my main concern with that approach is that other services i.e. ejb3, web, hornetq, etc. would not work as expected when a user installs jbosgi using the installer. Because there is no test coverage that tests these services with osgi in place. [12:54:07] Thomas Diesler: Here is a prime example http://community.jboss.org/thread/150453?tstart=0 [12:54:54] Thomas Diesler: Simple EJB3 SLSB breaks when osgi is installed [12:55:13] Thomas Diesler: Issues like this I'd like to see in the AS hudson run [12:55:37] Thomas Diesler: which I would not if its (only) in a separate profile [12:55:53] Dimitris Andreadis: but we could test for that with your custom osgi-test config, no? [12:57:01] Thomas Diesler: only if we run all-tests against an 'all+osgi' configuration [12:57:18] Thomas Diesler: I claim taht it would be better to make osgi part of all [12:57:35] Thomas Diesler: so that all test coverage must work with osgi in place [12:58:01] Dimitris Andreadis: i see your point [13:04:34] Thomas Diesler: Here is another example: https://jira.jboss.org/jira/browse/JBAS-7909 [13:04:51] Thomas Diesler: the 'default' profile would not even boot with osgi in place [13:05:17] Thomas Diesler: because the AOP subsystems makes certain assumptions that osgi bundle deployments do not sattisfy [13:05:59] Thomas Diesler: stuff like this I'd like to catch early before I want to ship osgi as part of AS [13:06:07] Dimitris Andreadis: in that case the only way to satisfy everybody is to copy over osgi for all tests and remove it afterwards. [13:06:20] Dimitris Andreadis: i.e. make all-test a custom config [13:06:38] Thomas Diesler: that would work ... [13:07:47] Dimitris Andreadis: maybe you can throw the idea/reasoning on the dev-list thread [13:08:26] Thomas Diesler: sure, let me add the boot time delta [13:08:38] Dimitris Andreadis: yes, thats imporant [13:19:17] Thomas Diesler: on my machine I get 'all' config boot times of osgi excluded: ~ 1:17min osgi included: ~ 1:37min This is an increase of ca. 29% [13:19:50] Dimitris Andreadis: 1:17 is a lot already ... [13:28:05] Dimitris Andreadis: all+osgi: 1:02 [13:29:52] Dimitris Andreadis: all-osgi: 49 [13:30:33] Dimitris Andreadis: +26.5% [13:31:26] Thomas Diesler: same result [13:31:31] Thomas Diesler: so what do you say? [13:33:13] Dimitris Andreadis: do the isolated testing anyway (custom config with smoke tests + osgi tests anyway) and discuss if we should alter the all testing [13:33:36] Thomas Diesler: ok	1	1	https://issues.jboss.org/browse/JBAS-7661
819	JBAS-7651	Stan, do you still consider this a BLOCKER for 6.0.0.Final?	0	0	https://issues.jboss.org/browse/JBAS-7651
820	JBAS-7612	I see this largely as a usage issue. The work unit stuff should be adapted to the new thread pool code since the old thread pool code predates JDK5 and Executors. Can you describe the requirements (preferably in a forum thread linked into this issue)?	0	0	https://issues.jboss.org/browse/JBAS-7612
821	JBAS-7588	>> The issue here is that org.jboss.system.RMIClassLoader is no longer getting used. >> It was previously configured in conf/jboss-service.xml before any subsystem (e.g. naming) used RMI. If you don't do it that way it won't work. Thanks a lot Adrian, for that help. Without it, this surely would have taken me a lot longer to even get a complete grasp of what's going on :) Just as a quick check, i booted the server with -Djava.rmi.server.RMIClassLoaderSpi=org.jboss.system.JBossRMIClassLoader in the JAVA_OPTS of the run script. The tests started passing consistently. I'll do the necessary changes (which you already explained).	1	1	https://issues.jboss.org/browse/JBAS-7588
822	JBAS-7583	Some time back there was a suggestion to move the SBM configuration to a central location instead of per server config/profile http://community.jboss.org/message/434149. Is that something that could be done?	0	0	https://issues.jboss.org/browse/JBAS-7583
823	JBAS-7563	Do we need to remove JRMPInvokerProxyHA JRMPInvokerProxy? I assume yes but wanted to see if others agree. Class JRMPInvoker has the following method which tells me that we do want to remove the proxies for JRMPInvoker: protected Invoker createDelegateInvoker() { return new JRMPInvokerProxy(this); }	1	1	https://issues.jboss.org/browse/JBAS-7563
824	JBAS-7563	From IRC re: WrapperDataSourceService: <bstansberry> smarlow: is that tied to JRMP? or just to Invoker? <bstansberry> smarlow: i.e. i suspect UnifiedInvoker is fine <bstansberry> yeah, it takes a property named "JMXInvokerName" where "JMXInvoker" means a legacy detached invoker that calls through the JMX bus <bstansberry> UnifiedInvoker meets the requirements <smarlow> well, it uses the Invoker interface but the comment mentions creating a JRMPInvokerProxy <bstansberry> where? <smarlow> around line 273 <smarlow> bstansberry: line 273 of WrapperDataSourceService.java <bstansberry> smarlow: ok, but it actually calls into GenericProxyFactory, which can create a proxy that can deal with any legacy detached invoker <bstansberry> a lot of code comments are from the early days when JRMP was the only detached invoker <smarlow> bstansberry: okay, so we should be good just updating the comment <bstansberry> yes	1	1	https://issues.jboss.org/browse/JBAS-7563
825	JBAS-7542	Shouldn't the solution be to update to use the correct library (or wait for that), rather than reverting to reflection?	1	1	https://issues.jboss.org/browse/JBAS-7542
826	JBAS-7539	Emanuel, if the profileservice-secured.jar is no longer needed, can I also remove ejb-jar.xml, profileservice-users.properties, and profileservice-roles.properties?	0	0	https://issues.jboss.org/browse/JBAS-7539
827	JBAS-7533	Can we make this an alpha3 with JBASM-33 fixed? That will stabilize the testsuite run by preventing the JBM clustering tests screwing it up. I volunteer to do JBASM-33 right away if not throwing the ISE is the right thing to do.	1	1	https://issues.jboss.org/browse/JBAS-7533
828	JBAS-7509	Is the console code going to be removed for M2? If not, please open a separate JIRA for cleaning out the code. The AS dist no longer includes the web-console, so this JIRA should be closed before M2 so that fact is reflected in the release notes.	1	1	https://issues.jboss.org/browse/JBAS-7509
829	JBAS-7509	I might have a go with it. What's the cutoff, Thursday 11th?	0	0	https://issues.jboss.org/browse/JBAS-7509
830	JBAS-7503	Howard, A couple questions: 1) When will the HornetQ upgrade happen? 2) How much effort is it to move trunk to 1.4.3.GA? Is it more than changing component-matrix/pom.xml and running the testsuite? I ask because 6.0.0.M1 includes 1.4.3.GA, and the development branch for the next milestone release does not. That's a regression, plus it makes other efforts to ensure things are working correctly in trunk more suspect. It's not good to leave the AS in a regressed state for a long time, even if we know that in the end JBM is going to be replaced.	1	1	https://issues.jboss.org/browse/JBAS-7503

#	Requirement	Questions	Answered	Successfully Answered	URL
831	JBAS-7457	Ok, the issue is that the NamingService is now deployed as an MCBean: <!-- An obsolete bean only needed if using legacy detached invoker access, or the createAlias/removeAlias operations --> <bean name="NamingService" class="org.jboss.naming.NamingService"> <annotation>@org.jboss.aop.microcontainer.aspects.jmx.JMX(name="jboss:service=NamingService", exposedInterface=org.jboss.naming.NamingServiceMBean, registerDirectly=true)</annotation> <!-- An example of using the unifed invoker as the transport. <depends optional-attribute-name="InvokerProxyFactory" proxy-type="attribute">jboss:service=proxyFactory,type=unified,target=Naming</depends> --> <property name="namingInfo"><inject bean="RemoteNamingBean"/></property> </bean> and this is not reproducing the same attributes as are exposed by default or via the xmbean descriptor. What I see is: {JNPServerSocketFactory=InvocationContext[name=JNPServerSocketFactory, type=null, attributeType=java.lang.String, isReadable=true, isWritable=true, returnType=null, signature=null], BootstrapURL=InvocationContext[name=BootstrapURL, type=null, attributeType=java.lang.String, isReadable=true, isWritable=false, returnType=null, signature=null], ClientSocketFactory=InvocationContext[name=ClientSocketFactory, type=null, attributeType=java.lang.String, isReadable=true, isWritable=true, returnType=null, signature=null], ClientSocketFactoryBean=InvocationContext[name=ClientSocketFactoryBean, type=null, attributeType=java.rmi.server.RMIClientSocketFactory, isReadable=true, isWritable=true, returnType=null, signature=null], InstallGlobalService=InvocationContext[name=InstallGlobalService, type=null, attributeType=boolean, isReadable=true, isWritable=true, returnType=null, signature=null], NamingProxy=InvocationContext[name=NamingProxy, type=null, attributeType=java.lang.Object, isReadable=true, isWritable=true, returnType=null, signature=null], RmiBindAddress=InvocationContext[name=RmiBindAddress, type=null, attributeType=java.lang.String, isReadable=true, isWritable=true, returnType=null, signature=null], Port=InvocationContext[name=Port, type=null, attributeType=int, isReadable=true, isWritable=true, returnType=null, signature=null], NamingInfo=InvocationContext[name=NamingInfo, type=null, attributeType=org.jnp.server.NamingBean, isReadable=true, isWritable=true, returnType=null, signature=null], Backlog=InvocationContext[name=Backlog, type=null, attributeType=int, isReadable=true, isWritable=true, returnType=null, signature=null], JNPServerSocketFactoryBean=InvocationContext[name=JNPServerSocketFactoryBean, type=null, attributeType=javax.net.ServerSocketFactory, isReadable=true, isWritable=true, returnType=null, signature=null], ServerSocketFactory=InvocationContext[name=ServerSocketFactory, type=null, attributeType=java.lang.String, isReadable=true, isWritable=true, returnType=null, signature=null], NamingInstance=InvocationContext[name=NamingInstance, type=null, attributeType=org.jnp.interfaces.Naming, isReadable=true, isWritable=false, returnType=null, signature=null], BindAddress=InvocationContext[name=BindAddress, type=null, attributeType=java.lang.String, isReadable=true, isWritable=true, returnType=null, signature=null], ServerSocketFactoryBean=InvocationContext[name=ServerSocketFactoryBean, type=null, attributeType=java.rmi.server.RMIServerSocketFactory, isReadable=true, isWritable=true, returnType=null, signature=null], RmiPort=InvocationContext[name=RmiPort, type=null, attributeType=int, isReadable=true, isWritable=true, returnType=null, signature=null], UseGlobalService=InvocationContext[name=UseGlobalService, type=null, attributeType=boolean, isReadable=true, isWritable=true, returnType=null, signature=null], CallByValue=InvocationContext[name=CallByValue, type=null, attributeType=boolean, isReadable=true, isWritable=true, returnType=null, signature=null], LookupListenerException=InvocationContext[name=LookupListenerException, type=null, attributeType=java.lang.Exception, isReadable=true, isWritable=false, returnType=null, signature=null]}	0	0	https://issues.jboss.org/browse/JBAS-7457
832	JBAS-7440	I couldn't reproduce locally for some reason (maybe it was fixed by someone already). Could you point me to the failure on hudson? Thanks!	1	1	https://issues.jboss.org/browse/JBAS-7440
833	JBAS-7434	15:44:17,139 DEBUG [org.jboss.ejb3.stateful.StatefulContainer] (WorkerThread#0[127.0.0.1:64635]) Received dynamic invocation for method with hash: -792204721456858463 2009-11-12 15:44:17,142 DEBUG [org.jboss.ejb3.entity.ExtendedPersistenceContextPropagationInterceptor] (WorkerThread#0[127.0.0.1:64635]) ++++ LongLivedSessionPropagationInterceptor 2009-11-12 15:44:17,142 DEBUG [org.hibernate.impl.SessionImpl] (WorkerThread#0[127.0.0.1:64635]) opened session at timestamp: 12580118571 2009-11-12 15:44:17,142 DEBUG [org.hibernate.ejb.AbstractEntityManagerImpl] (WorkerThread#0[127.0.0.1:64635]) Looking for a JTA transaction to join 2009-11-12 15:44:17,143 DEBUG [org.hibernate.jdbc.JDBCContext] (WorkerThread#0[127.0.0.1:64635]) successfully registered Synchronization 2009-11-12 15:44:17,144 DEBUG [org.hibernate.ejb.AbstractEntityManagerImpl] (WorkerThread#0[127.0.0.1:64635]) Looking for a JTA transaction to join 2009-11-12 15:44:17,144 DEBUG [org.hibernate.ejb.AbstractEntityManagerImpl] (WorkerThread#0[127.0.0.1:64635]) Transaction already joined 2009-11-12 15:44:17,162 DEBUG [org.hibernate.event.def.AbstractSaveEventListener] (WorkerThread#0[127.0.0.1:64635]) generated identifier: 1001, using strategy: org.hibernate.id.Assigned 2009-11-12 15:44:17,185 DEBUG [org.hibernate.event.def.AbstractFlushingEventListener] (RequestProcessor-2) processing flush-time cascades 2009-11-12 15:44:17,186 DEBUG [org.hibernate.event.def.AbstractFlushingEventListener] (RequestProcessor-2) dirty checking collections 2009-11-12 15:44:17,189 DEBUG [org.hibernate.event.def.AbstractFlushingEventListener] (RequestProcessor-2) Flushed: 1 insertions, 0 updates, 0 deletions to 1 objects 2009-11-12 15:44:17,189 DEBUG [org.hibernate.event.def.AbstractFlushingEventListener] (RequestProcessor-2) Flushed: 0 (re)creations, 0 updates, 0 removals to 0 collections 2009-11-12 15:44:17,190 DEBUG [org.hibernate.pretty.Printer] (RequestProcessor-2) listing entities: 2009-11-12 15:44:17,190 DEBUG [org.hibernate.pretty.Printer] (RequestProcessor-2) org.jboss.test.cluster.clusteredentity.classloader.Account{id=1001, balance=5, accountHolder=2c6d8085f3f280c2eff2e7aeeae2eff3f3aef4e5f3f4aee3ecf5f3f4e5f2aee3ecf5f3f4e5f2e5e4e5eef4e9f4f9aee3ece1f3f3ecefe1e4e5f2aec1e3e3eff5eef4c8efece4e5f2d0cb8080808080808081828082cc8088ece1f3f4cee1ede5f48092cceae1f6e1afece1eee7afd3f4f2e9eee7bbcc8083f3f3eef180fe8081f8f0f48085d3ede9f4e8f48084b1b0b0b0, branch=94536} 2009-11-12 15:44:17,190 DEBUG [org.hibernate.cache.UpdateTimestampsCache] (RequestProcessor-2) Pre-invalidating space [Account] 2009-11-12 15:44:17,198 DEBUG [org.hibernate.jdbc.AbstractBatcher] (RequestProcessor-2) about to open PreparedStatement (open PreparedStatements: 0, globally: 0) 2009-11-12 15:44:17,198 DEBUG [org.hibernate.jdbc.ConnectionManager] (RequestProcessor-2) opening JDBC connection 2009-11-12 15:44:17,205 DEBUG [org.hibernate.SQL] (RequestProcessor-2) insert into Account (accountHolder, balance, branch, id) values (?, ?, ?, ?) 2009-11-12 15:44:17,209 ERROR [org.hibernate.event.def.AbstractFlushingEventListener] (RequestProcessor-2) Could not synchronize database state with session: org.hibernate.type.SerializationException: could not deserialize at org.hibernate.util.SerializationHelper.deserialize(SerializationHelper.java:185) at org.hibernate.util.SerializationHelper.deserialize(SerializationHelper.java:211) at org.hibernate.type.SerializableType.fromBytes(SerializableType.java:105) at org.hibernate.type.SerializableType.deepCopyNotNull(SerializableType.java:97) at org.hibernate.type.MutableType.deepCopy(MutableType.java:48) at org.hibernate.type.SerializableToBlobType.deepCopy(SerializableToBlobType.java:102) at org.hibernate.type.AbstractType.disassemble(AbstractType.java:78) at org.hibernate.type.TypeFactory.disassemble(TypeFactory.java:474) at org.hibernate.cache.entry.CacheEntry.<init>(CacheEntry.java:69) at org.hibernate.action.EntityInsertAction.execute(EntityInsertAction.java:102) at org.hibernate.engine.ActionQueue.execute(ActionQueue.java:279) at	1	1	https://issues.jboss.org/browse/JBAS-7434
834	JBAS-7434	Open question for anyone following/working this issue. Is there a way to make the com.arjuna.ArjunaOTS.ManagedSynchronizationPOA._invoke run in the user/application thread that commited the transaction. This would be an optimization for the case where the invocation came from a local thread. Assumingly, the user/application thread must be blocked on the completition of the com.arjuna.ArjunaOTS.ManagedSynchronizationPOA._invoke and we probably have a way to make this "local" optimization.	1	0	https://issues.jboss.org/browse/JBAS-7434
835	JBAS-7434	Which version of Hibernate is this? 3.5?	0	0	https://issues.jboss.org/browse/JBAS-7434

#	Requirement	Questions	Answered	Successfully Answered	URL
836	JBAS-7434	AS trunk/component-matrix/pom.xml currently references the following versions: <version.org.hibernate>3.5.0-CR-1</version.org.hibernate> <version.org.hibernate.commons.annotations>3.2.0.Beta1</version.org.hibernate.commons.annotations> <version.org.hibernate.entity.manager>3.5.0-CR-1</version.org.hibernate.entity.manager> <version.org.hibernate.javax.persistence>1.0.0-CR-1</version.org.hibernate.javax.persistence> <version.org.hibernate.validator>4.0.2.GA</version.org.hibernate.validator> Which versions of the above should we be using at this point?	0	0	https://issues.jboss.org/browse/JBAS-7434
837	JBAS-7414	Now results in: java.lang.IllegalStateException: Cannot locate classloader resource: /managers.xml at org.jboss.jmx.adaptor.snmp.agent.TrapEmitter.load(TrapEmitter.java:240) at org.jboss.jmx.adaptor.snmp.agent.TrapEmitter.start(TrapEmitter.java:105)	0	0	https://issues.jboss.org/browse/JBAS-7414
838	JBAS-7384	servlet.war 21:35:22,050 INFO [TomcatDeployment] deploy, ctxPath=/jaxws-endpoint-servlet 21:35:22,288 ERROR [AbstractKernelController] Error installing to PreReal: name=http://jaxws-endpoint-api/jaxws-endpoint state=PostClassLoader mode=Manual requiredState=PreReal org.jboss.deployers.spi.DeploymentException: Cannot bind validator factory. at org.jboss.deployers.spi.DeploymentException.rethrowAsDeploymentException(DeploymentException.java:49) at org.jboss.beanvalidation.deployers.ValidatorFactoryJndiBinderDeployer.deploy(ValidatorFactoryJndiBinderDeployer.java:65) at org.jboss.beanvalidation.deployers.ValidatorFactoryJndiBinderDeployer.deploy(ValidatorFactoryJndiBinderDeployer.java:39) at org.jboss.deployers.spi.deployer.helpers.AbstractSimpleRealDeployer.internalDeploy(AbstractSimpleRealDeployer.java:62) at org.jboss.deployers.spi.deployer.helpers.AbstractRealDeployer.deploy(AbstractRealDeployer.java:50) at org.jboss.deployers.plugins.deployers.DeployerWrapper.deploy(DeployerWrapper.java:171) at org.jboss.deployers.plugins.deployers.DeployersImpl.doDeploy(DeployersImpl.java:1440) at org.jboss.deployers.plugins.deployers.DeployersImpl.doInstallParentFirst(DeployersImpl.java:1158) at org.jboss.deployers.plugins.deployers.DeployersImpl.install(DeployersImpl.java:1099) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:348) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:1633) at org.jboss.dependency.plugins.AbstractController.incrementState(AbstractController.java:935) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:1083) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:985) at org.jboss.dependency.plugins.AbstractController.change(AbstractController.java:823) at org.jboss.dependency.plugins.AbstractController.change(AbstractController.java:553) at org.jboss.deployers.plugins.deployers.DeployersImpl.process(DeployersImpl.java:782) at org.jboss.deployers.plugins.main.MainDeployerImpl.deploy(MainDeployerImpl.java:552) at org.jboss.webservices.integration.endpoint.EndpointAPIHttpServer.publish(EndpointAPIHttpServer.java:121) at org.jboss.ws.core.jaxws.spi.EndpointImpl.publish(EndpointImpl.java:167) at org.jboss.ws.core.jaxws.spi.EndpointImpl.publish(EndpointImpl.java:135) at org.jboss.test.ws.jaxws.endpoint.EndpointServlet.init(EndpointServlet.java:64) at org.apache.catalina.core.StandardWrapper.loadServlet(StandardWrapper.java:1048) at org.apache.catalina.core.StandardWrapper.load(StandardWrapper.java:950) at org.apache.catalina.core.StandardContext.loadOnStartup(StandardContext.java:4122) at org.apache.catalina.core.StandardContext.start(StandardContext.java:4417) at org.jboss.web.tomcat.service.deployers.TomcatDeployment.performDeployInternal(TomcatDeployment.java:333) at org.jboss.web.tomcat.service.deployers.TomcatDeployment.performDeploy(TomcatDeployment.java:150) at	0	0	https://issues.jboss.org/browse/JBAS-7384
839	JBAS-7384	What does beanvalidation deployer do to make this fail?	0	0	https://issues.jboss.org/browse/JBAS-7384
840	JBAS-7376	Can you change the name of the package to org.jboss.test.ejb3.jbas7376.unit (instead of ejbthree7376)? Helps in finding the JIRA :)	0	0	https://issues.jboss.org/browse/JBAS-7376
841	JBAS-7370	So can JBAS-7361 be closed then?	1	1	https://issues.jboss.org/browse/JBAS-7370
842	JBAS-7326	Added code to JaccContextValve that sets the response in the ActiveRequestResponseCache, just as it does with the request object. Test now pass: [sguilhen@BLACKBIRD testsuite]$ ant one-test -Dtest=org.jboss.test.web.test.WebProgrammaticLoginTestCase Buildfile: build.xml one-test: [delete] Deleting: /opt/workspace-jboss/jbossas-b5x/testsuite/output/log/test.log [junit] WARNING: multiple versions of ant detected in path for junit [junit] jar:file:/opt/development-tools/apache-ant-1.7.1/lib/ant.jar!/org/apache/tools/ant/Project.class [junit] and jar:file:/opt/workspace-jboss/jbossas-b5x/tools/lib/ant.jar!/org/apache/tools/ant/Project.class [junit] Running org.jboss.test.web.test.WebProgrammaticLoginTestCase [junit] Tests run: 2, Failures: 0, Errors: 0, Time elapsed: 5.006 sec BUILD SUCCESSFUL Total time: 9 seconds	1	1	https://issues.jboss.org/browse/JBAS-7326
843	JBAS-7277	Would this include all projects pulled into the AS? Or just the modules of AS itself? For a combined AS module javadoc, I opened JBBUILD-526 a while ago, but haven't worked on it yet. Also would this need to be packaged up into the distribution, or deployed to the web somewhere?	0	0	https://issues.jboss.org/browse/JBAS-7277
844	JBAS-7273	Andrew, if these are tests of a non-existing feature, should they be disabled until the feature is there?	0	0	https://issues.jboss.org/browse/JBAS-7273
845	JBAS-7273	Investigate dependency mismatch in AS?	1	1	https://issues.jboss.org/browse/JBAS-7273
846	JBAS-7257	Stan, can you comment on this remaining SEAM test failure?	1	1	https://issues.jboss.org/browse/JBAS-7257
847	JBAS-7257	Is this application using the old version of Facelets? And by old, I mean the non-JSF2 version that you need to bundle with your WAR. I assume that is the case and you just need to put the Facelets jar in the WEB-INF/lib directory. Otherwise, please post web.xml and faces-config.xml.	1	1	https://issues.jboss.org/browse/JBAS-7257
848	JBAS-7249	Ron, where is 2.5.2. Is it in the maven repository?	1	1	https://issues.jboss.org/browse/JBAS-7249
849	JBAS-7244	Is the update-to version now called 'weld'?	1	1	https://issues.jboss.org/browse/JBAS-7244
850	JBAS-7241	Isn't this update critical for this release?	0	0	https://issues.jboss.org/browse/JBAS-7241
851	JBAS-7241	Is this ready to be closed?	1	1	https://issues.jboss.org/browse/JBAS-7241
852	JBAS-7224	Against trunk: [alr@localhost trunk]$ grep -R 'fixRemoteAddressProperty' . Binary file ./main/target/classes/org/jboss/Main$ServerConfigUtil.class matches Binary file ./main/target/classes/org/jboss/Main.class matches ./main/src/main/org/jboss/Main.java: ServerConfigUtil.fixRemoteAddressProperty("java.rmi.server.hostname", defaultAddress); ./main/src/main/org/jboss/Main.java: ServerConfigUtil.fixRemoteAddressProperty("jgroups.bind_addr", defaultAddress); ./main/src/main/org/jboss/Main.java: ServerConfigUtil.fixRemoteAddressProperty("bind.address", null); ./main/src/main/org/jboss/Main.java: public static void fixRemoteAddressProperty(String systemPropertyName,	1	1	https://issues.jboss.org/browse/JBAS-7224

#	Requirement	Questions	Answered	Successfully Answered	URL
853	JBAS-7223	Juergen, could you provide this as a test case. Either VFS, Deployers or JBossAS test. If you need any help, we can continue the discussion on MC or JBossAS dev forums.	1	1	https://issues.jboss.org/browse/JBAS-7223
854	JBAS-7222	I've added hibernate-commons-annotations.jar & hibernate-validator.jar Regarding hibernate-client.jar, I see it produced by the ejb3 module in Branch_4_2: <!-- build hibernate-client.jar --> <!-- <unjar src="${hibernate.hibernate.lib}/hibernate3.jar" dest="${build.classes}" /> --> <jar jarfile="${build.lib}/hibernate-client.jar"> <zipfileset src="${hibernate.hibernate.lib}/hibernate3.jar"> <include name="org/hibernate/AssertionFailure.class"/> <include name="org/hibernate/**/*Exception.class"/> <include name="org/hibernate/collection/**"/> <include name="org/hibernate/exception/Nestable*.class"/> <include name="org/hibernate/exception/ExceptionUtils.class"/> <include name="org/hibernate/engine/SessionImplementor.class"/> <include name="org/hibernate/loader/CollectionAliases.class"/> <include name="org/hibernate/persister/collection/**"/> <include name="org/hibernate/pretty/MessageHelper.class"/> <include name="org/hibernate/type/Type.class"/> <include name="org/hibernate/EntityMode.class"/> <include name="org/hibernate/util/**"/> <include name="org/hibernate/proxy/**"/> </zipfileset> </jar> Do we want to do the same, or just copy the whole hibernate3.jar?	0	0	https://issues.jboss.org/browse/JBAS-7222
855	JBAS-7179	Should be resolved due to sub-task JBAS-7311 Both the upcoming JBAS 5.2 as well as in AS6.	1	1	https://issues.jboss.org/browse/JBAS-7179
856	JBAS-7179	The attached file may fix the issue in JBAS 5.1 but no guarantees.	1	1	https://issues.jboss.org/browse/JBAS-7179
857	JBAS-7179	Can we expect a fix for jboss-4.2.3.GA as well? It's the same behaviour, ExtendedFormAuthenticator.forwardToLoginPage will not invoke SecurityAssociationValve, so calls to WebAutentication will result in java.lang.IllegalStateException: request is null	0	0	https://issues.jboss.org/browse/JBAS-7179
858	JBAS-7085	Can you summarize what the JNDI binding names are for: 1) the remote AOP proxy(s) 2) the local java:* unsecured proxy Also, can you confirm that authentication is enabled by default for the AOP proxy and also describe the configuration steps for turning authentication on and off for it?	1	1	https://issues.jboss.org/browse/JBAS-7085
859	JBAS-7085	> the local unsecured ProfileService is only available via injection from the "ProfileService" MC bean Can you describe how to do this? I'll need to do it from within the jbas5 Jopr plugin, which will need to run both within the AS process (i.e. for the admin console) and external to it (in the JON Agent).	1	1	https://issues.jboss.org/browse/JBAS-7085
860	JBAS-7081	(11:38:59 AM) ccrouch: were you seeing this repeatedly? (11:38:59 AM) ccrouch: JBAS-7081 ConcurrentModificationException when invoking 'stop' managed operation on Queue managed component (11:42:21 AM) ips: no, i only saw that in a multi-threaded PS client i wrote that was sending concurrent requests (11:42:35 AM) ips: never saw it in admin console orJON regular use	1	0	https://issues.jboss.org/browse/JBAS-7081
861	JBAS-7078	Shouldn't this new work be a separate issue? The initial work is already in a released version and included in its release notes	1	1	https://issues.jboss.org/browse/JBAS-7078
862	JBAS-7074	that does not have jboss-dependency.jar in its classpath: java.lang.NoClassDefFoundError: org/jboss/dependency/spi/ControllerState at org.jboss.profileservice.spi.types.ControllerStateMetaType.<init>(ControllerStateMetaType.java:53) at org.jboss.profileservice.spi.types.ControllerStateMetaType.<clinit>(ControllerStateMetaType.java:49) at sun.misc.Unsafe.ensureClassInitialized(Native Method) at sun.reflect.UnsafeFieldAccessorFactory.newFieldAccessor(UnsafeFieldAccessorFactory.java:25) at sun.reflect.ReflectionFactory.newFieldAccessor(ReflectionFactory.java:122) at java.lang.reflect.Field.acquireFieldAccessor(Field.java:918) at java.lang.reflect.Field.getFieldAccessor(Field.java:899) at java.lang.reflect.Field.getLong(Field.java:528) at java.io.ObjectStreamClass.getDeclaredSUID(ObjectStreamClass.java:1614) at java.io.ObjectStreamClass.access$700(ObjectStreamClass.java:52) at java.io.ObjectStreamClass$2.run(ObjectStreamClass.java:425) at java.security.AccessController.doPrivileged(Native Method) at java.io.ObjectStreamClass.<init>(ObjectStreamClass.java:413) at java.io.ObjectStreamClass.lookup(ObjectStreamClass.java:310) at java.io.ObjectStreamClass.initNonProxy(ObjectStreamClass.java:547) at java.io.ObjectInputStream.readNonProxyDesc(ObjectInputStream.java:1583) at java.io.ObjectInputStream.readClassDesc(ObjectInputStream.java:1496) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1732) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:351) at org.jboss.managed.plugins.DefaultFieldsImpl.readObject(DefaultFieldsImpl.java:381) at sun.reflect.GeneratedMethodAccessor9.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at java.io.ObjectStreamClass.invokeReadObject(ObjectStreamClass.java:974) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1849) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:351) at org.jboss.managed.plugins.ManagedPropertyImpl.readVersion1(ManagedPropertyImpl.java:663) at org.jboss.managed.plugins.ManagedPropertyImpl.readObject(ManagedPropertyImpl.java:633) at sun.reflect.GeneratedMethodAccessor25.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at java.io.ObjectStreamClass.invokeReadObject(ObjectStreamClass.java:974) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1849) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:351) at java.util.HashMap.readObject(HashMap.java:1030) at sun.reflect.GeneratedMethodAccessor26.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at java.io.ObjectStreamClass.invokeReadObject(ObjectStreamClass.java:974) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1849) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1947) at	1	0	https://issues.jboss.org/browse/JBAS-7074
863	JBAS-7052	Ian, this jira is not assigned to me but I wanted to try helping with a few questions. I read the above jira information and am left wondering a few things. I assume that you already put in effort to resolve this on your own, so please help me answer these questions: 1. Does the above exception occur after the server side getComponentsForType() has been invoked? I assume yes and that this is caused by the handling of the getComponentsForType() result object being returned. 2. Assuming that this is caused by the result value from the getComponentsForType(), what classes are being returned? 3. If you know, can you tell me why the org.jboss.profileservice.management.ManagedOperationDelegate class is being referenced in the above call stack?	1	1	https://issues.jboss.org/browse/JBAS-7052
864	JBAS-7052	Scott, it looks like you only made this fix in the Branch_5_x branch. I'm using the EAP5 branch. Can you make the change in that branch too?	1	1	https://issues.jboss.org/browse/JBAS-7052

#	Requirement	Questions	Answered	Successfully Answered	URL
865	JBAS-7051	the page and the following exception in the server log: 13:19:35,767 ERROR [PluginContainerResourceManager] Error retrieving configuration for resource: Resource[id=-37, type=Q ueue, key=queue/tutorial/defaultedexample, name=queue/tutorial/defaultedexample, parent=JBoss EAP 5 (default)] org.rhq.core.clientapi.agent.PluginContainerException: Cannot load Resource configuration for [Queue Resource with id [- 37]. at org.rhq.core.pc.configuration.ConfigurationManager.loadResourceConfiguration(ConfigurationManager.java:174) at org.jboss.on.embedded.manager.pc.PluginContainerResourceManager.getResourceConfiguration(PluginContainerResou rceManager.java:235) at org.jboss.on.embedded.ui.configuration.resource.ResourceConfigurationUIBean.resourceConfiguration(ResourceCon figurationUIBean.java:144) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.seam.util.Reflections.invoke(Reflections.java:22) at org.jboss.seam.intercept.RootInvocationContext.proceed(RootInvocationContext.java:31) at org.jboss.seam.intercept.SeamInvocationContext.proceed(SeamInvocationContext.java:56) at org.jboss.seam.transaction.RollbackInterceptor.aroundInvoke(RollbackInterceptor.java:28) at org.jboss.seam.intercept.SeamInvocationContext.proceed(SeamInvocationContext.java:68) at org.jboss.seam.core.BijectionInterceptor.aroundInvoke(BijectionInterceptor.java:77) at org.jboss.seam.intercept.SeamInvocationContext.proceed(SeamInvocationContext.java:68) at org.jboss.seam.core.ConversationInterceptor.aroundInvoke(ConversationInterceptor.java:56) at org.jboss.seam.intercept.SeamInvocationContext.proceed(SeamInvocationContext.java:68) at org.jboss.seam.core.MethodContextInterceptor.aroundInvoke(MethodContextInterceptor.java:44) at org.jboss.seam.intercept.SeamInvocationContext.proceed(SeamInvocationContext.java:68) at org.jboss.seam.intercept.RootInterceptor.invoke(RootInterceptor.java:107) at org.jboss.seam.intercept.JavaBeanInterceptor.interceptInvocation(JavaBeanInterceptor.java:166) at org.jboss.seam.intercept.JavaBeanInterceptor.invoke(JavaBeanInterceptor.java:102) at org.jboss.on.embedded.ui.configuration.resource.ResourceConfigurationUIBean_$$_javassist_12.resourceConfigura tion(ResourceConfigurationUIBean_$$_javassist_12.java) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.el.util.ReflectionUtil.invokeMethod(ReflectionUtil.java:329) at org.jboss.el.util.ReflectionUtil.invokeMethod(ReflectionUtil.java:274) at org.jboss.el.parser.AstMethodSuffix.getValue(AstMethodSuffix.java:59) at org.jboss.el.parser.AstMethodSuffix.invoke(AstMethodSuffix.java:65) at	0	0	https://issues.jboss.org/browse/JBAS-7051
866	JBAS-7051	exceptions are dumped to the server log: 13:25:56,641 ERROR [JmsDestinationComponent] Failed to collect metric for MeasurementScheduleRequest[scheduleId=1, name= deliveringCount, interval=0, enabled=true, dataType=MEASUREMENT, rawNumericType=null] java.lang.reflect.UndeclaredThrowableException: Failed to get property 'deliveringCount' on component 'jboss.mq.destinat ion:name=defaultedtutorial,service=Queue'. at org.jboss.profileservice.management.KernelBusRuntimeComponentDispatcher.get(KernelBusRuntimeComponentDispatch er.java:119) at org.jboss.profileservice.management.DelegatingComponentDispatcherImpl.get(DelegatingComponentDispatcherImpl.j ava:81) at sun.reflect.GeneratedMethodAccessor432.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.aop.Dispatcher.invoke(Dispatcher.java:121) at org.jboss.aspects.remoting.AOPRemotingInvocationHandler.invoke(AOPRemotingInvocationHandler.java:82) at org.jboss.profileservice.remoting.ProfileServiceInvocationHandler.invoke(ProfileServiceInvocationHandler.java :97) at org.jboss.remoting.ServerInvoker.invoke(ServerInvoker.java:891) at org.jboss.remoting.transport.local.LocalClientInvoker.invoke(LocalClientInvoker.java:106) at org.jboss.remoting.Client.invoke(Client.java:1724) at org.jboss.remoting.Client.invoke(Client.java:629) at org.jboss.aspects.remoting.InvokeRemoteInterceptor.invoke(InvokeRemoteInterceptor.java:60) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.MergeMetaDataInterceptor.invoke(MergeMetaDataInterceptor.java:74) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.security.SecurityClientInterceptor.invoke(SecurityClientInterceptor.java:65) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aop.generatedproxies.AOPProxy$1.get(AOPProxy$1.java) at org.jboss.profileservice.management.client.ManagedPropertyDelegate.getValue(ManagedPropertyDelegate.java:151) at org.rhq.plugins.jbossas5.ManagedComponentComponent.getSimpleValue(ManagedComponentComponent.java:278) at org.rhq.plugins.jbossas5.ManagedComponentComponent.getValues(ManagedComponentComponent.java:231) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.rhq.core.pc.inventory.ResourceContainer$ComponentInvocationThread.call(ResourceContainer.java:526) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) at java.util.concurrent.FutureTask.run(FutureTask.java:138) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) at java.lang.Thread.run(Thread.java:619) Caused by: javax.management.AttributeNotFoundException: not found: DeliveringCount at	0	0	https://issues.jboss.org/browse/JBAS-7051

#	Requirement	Questions	Answered	Successfully Answered	URL
867	JBAS-7051	java.lang.reflect.UndeclaredThrowableException: Failed to invoke method 'listAllMessages' on component 'jboss.mq.destination:name=tutorial,service=Queue' with parameters [SimpleMetaType:java.lang.String:null]. at org.jboss.profileservice.management.KernelBusRuntimeComponentDispatcher.invoke(KernelBusRuntimeComponentDispatcher.java:201) at org.jboss.profileservice.management.DelegatingComponentDispatcherImpl.invoke(DelegatingComponentDispatcherImpl.java:95) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.aop.Dispatcher.invoke(Dispatcher.java:121) at org.jboss.aspects.remoting.AOPRemotingInvocationHandler.invoke(AOPRemotingInvocationHandler.java:82) at org.jboss.profileservice.remoting.ProfileServiceInvocationHandler.invoke(ProfileServiceInvocationHandler.java:99) at org.jboss.remoting.ServerInvoker.invoke(ServerInvoker.java:891) at org.jboss.remoting.transport.local.LocalClientInvoker.invoke(LocalClientInvoker.java:106) at org.jboss.remoting.Client.invoke(Client.java:1724) at org.jboss.remoting.Client.invoke(Client.java:629) at org.jboss.aspects.remoting.InvokeRemoteInterceptor.invoke(InvokeRemoteInterceptor.java:60) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.MergeMetaDataInterceptor.invoke(MergeMetaDataInterceptor.java:74) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.security.SecurityClientInterceptor.invoke(SecurityClientInterceptor.java:65) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aop.generatedproxies.AOPProxy$0.invoke(AOPProxy$0.java) at org.jboss.profileservice.management.client.ManagedOperationDelegate.invoke(ManagedOperationDelegate.java:63) at org.rhq.plugins.jbossas5.ManagedComponentComponent.invokeOperation(ManagedComponentComponent.java:218) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.rhq.core.pc.inventory.ResourceContainer$ComponentInvocationThread.call(ResourceContainer.java:525) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) at java.util.concurrent.FutureTask.run(FutureTask.java:138) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) at java.lang.Thread.run(Thread.java:619) Caused by: javax.management.ReflectionException at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:231) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.system.microcontainer.ServiceControllerContext.invoke(ServiceControllerContext.java:190) at	0	0	https://issues.jboss.org/browse/JBAS-7051
868	JBAS-7047	What's the latest on this one?	0	0	https://issues.jboss.org/browse/JBAS-7047
869	JBAS-7047	I'm still not seeing any change of behavior in the latest 5.2 snapshot downloaded from hudson. The exception received client side is an IncompleteDeploymentException which contains a the contextsInError. The context in error that relates to WebBeansBootstrapBean contains the root cause of the exception. See http://anonsvn.jboss.org/repos/webbeans/ri/trunk/jboss-tck-runner/src/test/java/org/jboss/webbeans/tck/jbossas/WebBeansProfileServiceDeploymentExceptionTransformer.java for the way I extract the exception. Is this correct?	1	1	https://issues.jboss.org/browse/JBAS-7047
870	JBAS-7046	Where does this stand?	1	1	https://issues.jboss.org/browse/JBAS-7046
871	JBAS-7013	Hi Stan, I have reproduced this issue on a previously successfully running Seam 2.2.1.EAP5 web application when I upgraded the JBoss AS 5 (EAP) from jbdevstudio 3.0.1 using your cheat sheet (http://community.jboss.org/wiki/upgradejbossastojsf2). This is using Mojarra: 2.0.3 (b05), after the admin-console is successfully loaded but is of no effect - just a nuisance. I, too, am stuck with AS5 rather than deploying to AS5 based on client req't. Any more insights? Regards, Henry Katz	1	1	https://issues.jboss.org/browse/JBAS-7013
872	JBAS-7008	What do we need exposed? Any desired behaviour should be reflected in: http://anonsvn.jboss.org/repos/jbossas/branches/Branch_5_x/testsuite/src/main/org/jboss/test/profileservice/test/Ejb3MetricsUnitTestCase.java	0	0	https://issues.jboss.org/browse/JBAS-7008
873	JBAS-7008	Well the minimum would just be one ManagedObject for each entity bean with a 'name' property whose value is the EJB's name. As for metrics, I'm not sure what's available. If the Hibernate stats are not turned on, are there still any other metrics exposed by the EJB3 container for entity beans?	1	0	https://issues.jboss.org/browse/JBAS-7008
874	JBAS-7008	I've deployed an EJB3 entity bean, and the corresponding MBean in the jmx-console (jboss.j2ee:jar=entity-test.jar,name=EntityTestBean,service=EJB3) contains the following metrics: CreateCount R int MBean Attribute. InvokeStats R org.jboss.ejb3.statistics.InvocationStatistics MBean Attribute. CurrentSize R int MBean Attribute. RemoveCount R int MBean Attribute. MaxSize R int MBean Attribute. AvailableCount R int MBean Attribute. Can we expose these same metrics via the Entity Bean ManagedObject?	1	0	https://issues.jboss.org/browse/JBAS-7008
875	JBAS-7008	Are you sure EntityTestBean is an entity, not some SLSB used to test entities? Because these metrics make no sense for entities; they're not invoked, we don't track create counts, etc.	0	0	https://issues.jboss.org/browse/JBAS-7008
876	JBAS-7008	Yeah, my bad. EntityTestBean is indeed an SLSB. Jason, is it possible for you to use some Hibernate API to obtain the list of deployed entities and their associated deployments?	0	0	https://issues.jboss.org/browse/JBAS-7008
877	JBAS-6982	I'm still getting a "javax.ejb.EJBAccessException: Caller unauthorized" using either set of credentials. It looks like the difference may be in the method I'm using to authenticate. I'm setting the principal and credentials in the env props I pass when creating the InitialContext. My client code is at: https://svn.jboss.org/repos/jopr/trunk/etc/jbas5-ejb-client/src/main/java/test/EjbClient.java Whereas in: https://svn.jboss.org/repos/jbossas/branches/JBPAPP_5_0/testsuite/src/main/org/jboss/test/profileservice/testsecure/ProfileServiceUnitTestCase.java you are instead directly using JAAS to authenticate prior to doing the JNDI lookups. Is the method I'm using supposed to work or do I need to switch to the JAAS method?	1	1	https://issues.jboss.org/browse/JBAS-6982
878	JBAS-6942	Jesper, which MBean exposes the new operation and what is the name of the operation? Also, did you expose the functionality via the Profile Service as well, or only JMX? Thanks, Ian	1	0	https://issues.jboss.org/browse/JBAS-6942

#	Requirement	Questions	Answered	Successfully Answered	URL
879	JBAS-6885	Prompt Input Thread" java.lang.NoClassDefFoundError: Lorg/jboss/config/spi/Configuration; at java.lang.Class.getDeclaredFields0(Native Method) at java.lang.Class.privateGetDeclaredFields(Class.java:2291) at java.lang.Class.getDeclaredField(Class.java:1880) at java.io.ObjectStreamClass.getDeclaredSUID(ObjectStreamClass.java:1610) at java.io.ObjectStreamClass.access$700(ObjectStreamClass.java:52) at java.io.ObjectStreamClass$2.run(ObjectStreamClass.java:425) at java.security.AccessController.doPrivileged(Native Method) at java.io.ObjectStreamClass.<init>(ObjectStreamClass.java:413) at java.io.ObjectStreamClass.lookup(ObjectStreamClass.java:310) at java.io.ObjectStreamClass.initNonProxy(ObjectStreamClass.java:547) at java.io.ObjectInputStream.readNonProxyDesc(ObjectInputStream.java:1583) at java.io.ObjectInputStream.readClassDesc(ObjectInputStream.java:1496) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1732) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1947) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1871) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:351) at java.util.HashMap.readObject(HashMap.java:1030) at sun.reflect.GeneratedMethodAccessor1.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at java.io.ObjectStreamClass.invokeReadObject(ObjectStreamClass.java:974) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1849) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1947) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1871) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1947) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1871) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:351) at java.util.HashMap.readObject(HashMap.java:1030) at sun.reflect.GeneratedMethodAccessor1.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at java.io.ObjectStreamClass.invokeReadObject(ObjectStreamClass.java:974) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1849) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1947) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1871) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1753) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1329) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:351) at org.jboss.aop.joinpoint.InvocationResponse.readExternal(InvocationResponse.java:119) at java.io.ObjectInputStream.readExternalData(ObjectInputStream.java:1792) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1751) at	1	1	https://issues.jboss.org/browse/JBAS-6885
880	JBAS-6885	org.jboss.profileservice.management.TempManagedComponentImpljava.lang.IllegalStateException: java.lang.ClassNotFoundException: org.jboss.profileservice.management.TempManagedComponentImpl at org.jboss.adminclient.command.ListComponentsCommand.execute(ListComponentsCommand.java:66) at org.jboss.adminclient.AdminClientMain.executeCommand(AdminClientMain.java:307) at org.jboss.adminclient.AdminClientMain.executeCommand(AdminClientMain.java:265) at org.jboss.adminclient.AdminClientMain.access$000(AdminClientMain.java:57) at org.jboss.adminclient.AdminClientMain$1.run(AdminClientMain.java:239) at java.lang.Thread.run(Thread.java:613) Caused by: java.lang.ClassNotFoundException: org.jboss.profileservice.management.TempManagedComponentImpl at org.jboss.remoting.serialization.ClassLoaderUtility.loadClass(ClassLoaderUtility.java:103) at org.jboss.remoting.loading.RemotingClassLoader.loadClass(RemotingClassLoader.java:86) at java.lang.ClassLoader.loadClassInternal(ClassLoader.java:374) at java.lang.Class.forName0(Native Method) at java.lang.Class.forName(Class.java:242) at org.jboss.remoting.loading.ObjectInputStreamWithClassLoader.resolveClass(ObjectInputStreamWithClassLoader.java:179) at java.io.ObjectInputStream.readNonProxyDesc(ObjectInputStream.java:1544) at java.io.ObjectInputStream.readClassDesc(ObjectInputStream.java:1466) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1699) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1305) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:348) at java.util.HashSet.readObject(HashSet.java:278) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at java.io.ObjectStreamClass.invokeReadObject(ObjectStreamClass.java:946) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1809) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1719) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1305) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:348) at org.jboss.aop.joinpoint.InvocationResponse.readExternal(InvocationResponse.java:119) at java.io.ObjectInputStream.readExternalData(ObjectInputStream.java:1755) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1717) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1305) at java.io.ObjectInputStream.defaultReadFields(ObjectInputStream.java:1908) at java.io.ObjectInputStream.readSerialData(ObjectInputStream.java:1832) at java.io.ObjectInputStream.readOrdinaryObject(ObjectInputStream.java:1719) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1305) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:348) at org.jboss.remoting.serialization.impl.java.JavaSerializationManager.receiveObjectVersion2_2(JavaSerializationManager.java:238) at org.jboss.remoting.serialization.impl.java.JavaSerializationManager.receiveObject(JavaSerializationManager.java:138) at org.jboss.remoting.marshal.serializable.SerializableUnMarshaller.read(SerializableUnMarshaller.java:123) at org.jboss.invocation.unified.marshall.InvocationUnMarshaller.read(InvocationUnMarshaller.java:59) at org.jboss.remoting.transport.socket.MicroSocketClientInvoker.versionedRead(MicroSocketClientInvoker.java:1220) at	0	0	https://issues.jboss.org/browse/JBAS-6885
881	JBAS-6836	The schema is available via jbosssx.jar and is http://anonsvn.jboss.org/repos/jbossas/projects/security/security-jboss-sx/branches/Branch_2_0/jbosssx/src/resources/schema/security-config_5_0.xsd I updated the schema in the local folder with the latest version. Now I am getting an error as follows: ================================= Failed to parse source: Element {urn:jboss:security-config}policy is not bound as a global element. org.jboss.xb.binding.JBossXBException: Failed to parse source: Element {urn:jboss:security-config}policy is not bound as a global element. at org.jboss.xb.binding.parser.sax.SaxJBossXBParser.parse(SaxJBossXBParser.java:203) at org.jboss.xb.binding.parser.sax.SaxJBossXBParser.parse(SaxJBossXBParser.java:183) at ================================= What does that really mean?	1	1	https://issues.jboss.org/browse/JBAS-6836

#	Requirement	Questions	Answered	Successfully Answered	URL
882	JBAS-6804	Can you comment on these failures?	0	0	https://issues.jboss.org/browse/JBAS-6804
883	JBAS-6796	Emanuel, Did you mark this as resolved, because you thought we could update the Jopr jbas5 plugin's GenericValue adapter code to set the missing field(s) on the wrapped ManagedObject? If so, I'll need a little direction. Is attachmentName the only field I need to set, or are there others? What should I set the attachment name to for the security-domain property? You mentioned that you thought it was a good idea to represent the security-domain prop as a CompositeValue, rather than a GenericValue wrapping an MO. Is that something you plan on doing in the CR1 timeframe? Thanks.	1	1	https://issues.jboss.org/browse/JBAS-6796
884	JBAS-6789	Do you require support specifically for log4j deployments? Or simply the ability for each deployment to provide its own logging configuration?	1	1	https://issues.jboss.org/browse/JBAS-6789
885	JBAS-6787	Comparing JBoss-AS-5.1.x-testSuite-sun15 run #210 (before updata) and run #211, three new failures appeared: org.jboss.test.cluster.defaultcfg.test.ReadWriteClusteredLockManagerUnitTestCase(Default-udp).testLocalLockingStateRejectsSuperiorRemoteCaller org.jboss.test.compatibility.test.SerialVersionUIDUnitTestCase(tests-compatibility).test500Compatibility org.jboss.test.compatibility.test.SerialVersionUIDUnitTestCase(tests-compatibility).test501Compatibility and five failures went away: org.jboss.test.cluster.defaultcfg.web.test.SessionTimestampTestCase(SyncModeNUseJvm-udp).testSessionTimeout org.jboss.test.compatibility.test.SerialVersionUIDUnitTestCase.testJavaEE50Compatibilit org.jboss.test.jbossmessaging.test.SecurityUnitTestCase(tests-jbossmessaging).testAuzInvalidTopicSubscriber org.jboss.test.util.test.TwiddleUnitTestCase.testHelp org.jboss.test.cluster.defaultcfg.test.ReadWriteClusteredLockManagerUnitTestCase(Default-udp).testConcurrentRemoteLock	1	1	https://issues.jboss.org/browse/JBAS-6787
886	JBAS-6787	The following failures appeared betweenJBoss-AS-5.1.x-testSuite-sun15-jts run and 283 JBoss-AS-5.1.x-testSuite-sun15-jts run #284 * diff.org.jboss.test.cluster.defaultcfg.simpleweb.test.BuddyReplicationFailoverUnitTestCase(Default-udp).testFailoverAndFailBack * org.jboss.test.jbossmessaging.clustertest.ClusteredTestCase(tests-jbossmessaging-cluster).testSimpleCluster	1	1	https://issues.jboss.org/browse/JBAS-6787
887	JBAS-6787	The following failures appeared between runs 60 and 61 of JBoss-AS-5.1.x-testSuite-sun15-sun16: * org.jboss.test.cluster.defaultcfg.test.ReadWriteClusteredLockManagerUnitTestCase(Default-udp).testConcurrentRemoteLock * org.jboss.test.jbossmessaging.clustertest.ClusteredTestCase(tests-jbossmessaging-cluster).testSimpleCluster * org.jboss.test.jbossmessaging.clustertest.ClusteredTestCase(tests-jbossmessaging-cluster).testSimpleFailover * org.jboss.test.timer.test.BasicTimerUnitTestCase.testEntityBeanTimer and the following disappeared: * org.jboss.test.timer.test.SecureTimerUnitTestCase(tests-security-basic-unit).testSecuredStatelessSessionBeanTimer	1	1	https://issues.jboss.org/browse/JBAS-6787
888	JBAS-6779	Changing target release to 2.0.0.CR1 as GA depends on JBC 3.1.0.GA which won't be out. I expect no difference between 2.0.0.CR1 and 2.0.0.GA, nor do I expect any changes in JBC between current 3.1.0.CR1 and 3.1.0.GA that will affect ha-server-cache-jbc.	0	0	https://issues.jboss.org/browse/JBAS-6779
889	JBAS-6777	Using the EAP branch r90329, I'm still seeing the interleaving prop always having a value of false, for new ds.xml files that contain the interleaving element, as well as after I've updated the prop's value to true via the Profile Service. Did you forget to check the fix into the EAP branch?	1	1	https://issues.jboss.org/browse/JBAS-6777
890	JBAS-6777	Using EAP branch r90482, the 'interleaving' prop still gets reset to false by the Profile Service after I set it to true via a component update. I tried with both an XA Datasource and a Tx Connection Factory with XA enabled. Did you forget to check the fix into the EAP branch?	1	1	https://issues.jboss.org/browse/JBAS-6777
891	JBAS-6777	Yes, its in the https://svn.jboss.org/repos/jbossas/branches/JBPAPP_5_0 and the test is validating that the interleaving property is true on the underlying component: // Query the interleaving ManagementView mgtView = getManagementView(); ComponentType type = KnownComponentTypes.DataSourceTypes.XA.getType(); ManagedComponent txcf = mgtView.getComponent(jndiName, type); assertNotNull(txcf); ManagedProperty interleaving = txcf.getProperty("interleaving"); assertNotNull("interleaving", interleaving); MetaValue interleavingMV = interleaving.getValue(); assertNotNull("interleaving.value", interleavingMV); assertEquals("interleaving.value is true", SimpleValueSupport.wrap(Boolean.TRUE), interleavingMV); Can you attach the ds.xml file that is showing the false value.	1	1	https://issues.jboss.org/browse/JBAS-6777
892	JBAS-6765	The NPE is now avoided, but a null value is not valid for an ObjectName type property. There is no meaningful default. The stack trace you show does not look like the setting of a template property. Is this the origin of the property?	0	0	https://issues.jboss.org/browse/JBAS-6765
893	JBAS-6750	Where do I get this war from? Is it in the maven repository somewhere?	1	1	https://issues.jboss.org/browse/JBAS-6750
894	JBAS-6750	How does common/deploy work? Do files in there get deployed for all config sets? That directory doesn't even exist in my AS 5.1 build.	1	1	https://issues.jboss.org/browse/JBAS-6750
895	JBAS-6750	Create thread on jboss-development about this - having a common/deploy doesn't make sense if you don't plan to support all configuration under server/ - like minimal. Also remember that the idea of having configurations under server is to enable users to create their own. Are you going to support these also ? Email or forum thread, please.	0	0	https://issues.jboss.org/browse/JBAS-6750
896	JBAS-6750	Hey Paul Ian prepared a non-snapshot version to test the process http://repository.jboss.com/maven2/org/jboss/jopr/jopr-embedded-jbas5/1.2.0.CR1/ Can you make the necessary updates to get this in the JBAS build. We can then just update the version to say 1.2.0.CR2 next week. Thanks Charles	1	1	https://issues.jboss.org/browse/JBAS-6750
897	JBAS-6750	Ian, can you also add the CR1 version of the parent pom? When I tried to update to the new version, maven couldn't find the parent pom.	1	1	https://issues.jboss.org/browse/JBAS-6750
898	JBAS-6731	Did this doc find its way into the AS documentation?	1	0	https://issues.jboss.org/browse/JBAS-6731
899	JBAS-6728	Shouldn't the same be done for the "xa-transaction" property?	1	1	https://issues.jboss.org/browse/JBAS-6728
900	JBAS-6728	That is, should that property be read-only too, or can the transaction type of a datasource be flipped back and forth between XA and Local at runtime? If that is supposed to work, right now it's broken. When I change the value of xa-transaction from false to true and then call updateComponent(), the value gets reset to false.	1	1	https://issues.jboss.org/browse/JBAS-6728
901	JBAS-6721	1) I've updated the Queue and Topic service: the Name is now available and can be set in the template. But is now STATISTIC, as it does not have a setter. 2) done 3) Yes, I've set the messageCounterHistoryDayLimit to STATISTIC, as it does not work - if you specify a value for that the deployment fails, so until this is not fixed it should be ignored. 4) clustered is now readOnly, but still available in the template. Can we close this jira then ?	0	0	https://issues.jboss.org/browse/JBAS-6721
902	JBAS-6716	Emanuel, so the above is the meta data for the securityConfig property generated from the management interface. Can you explain how this maps to the attribute in the xml definition, e.g. <attribute name="SecurityConfig"> <security> <role name="publisher" read="true" write="true" create="false"/> </security> </attribute>	0	0	https://issues.jboss.org/browse/JBAS-6716

#	Requirement	Questions	Answered	Successfully Answered	URL
903	JBAS-6716	Which revision of 5_x are you using, as I cannot reproduce this in r87033?	0	0	https://issues.jboss.org/browse/JBAS-6716
904	JBAS-6716	r87027. Have you tried reproducing using the admin console?	0	0	https://issues.jboss.org/browse/JBAS-6716
905	JBAS-6701	Can you please provide more information? There are some test cases which are removing exploded deployments, so it might be something different.	1	1	https://issues.jboss.org/browse/JBAS-6701
906	JBAS-6701	Ian can this be resolved now?	1	1	https://issues.jboss.org/browse/JBAS-6701
907	JBAS-6697	That sounds close to what we need, but not exactly. In the Jopr jboss-as-4 plugin, when a user creates a new WAR or EAR via Jopr, they are prompted to enter the deploy directory, which is a path relative to the AS configSet dir. They would usually enter "deploy" or "farm", but we had a couple users who deployed to "my.custom.deploy" or something (e.g. when they wanted to keep their deploy dir on a different filesystem). We were hoping to provide this same option at app create time in the jboss-as-5 plugin. Prompting for a deployment profile name seems like it might be too low-level for an end-user; it might be more intuitive to them to just specify the directory they want to deploy to. What do you think? Charles, any thoughts?	0	0	https://issues.jboss.org/browse/JBAS-6697
908	JBAS-6689	Can you verify we are persisting the destination properties marked as DEPLOYMENT_RESTART across restarts in one of the jms destinations testcases Emanuel.	1	1	https://issues.jboss.org/browse/JBAS-6689
909	JBAS-6676	As I mentioned in my forum post, there is a similar issue where the {type=MCBean, subtype=*} ComponentType is used to represent multiple types of underlying managed resources - ServerInfo, etc. From the Jopr perspective, we'd much rather see a different ComponentType used for each different type of underlying entity. Is this feasible?	0	0	https://issues.jboss.org/browse/JBAS-6676
910	JBAS-6669	Can this now be resolved? I've set those config-proerties to readOnly, as config-properties can only be edited for NoTx and TxConnectionFactories.	1	1	https://issues.jboss.org/browse/JBAS-6669
911	JBAS-6664	Files Changed MODIFY /branches/Branch_5_x/connector/src/main/org/jboss/resource/deployers/management/DsDataSourceTemplateInfo.java Are you sure this is going to fix the issue for connection factories too?	1	1	https://issues.jboss.org/browse/JBAS-6664
912	JBAS-6657	Please let me know: was this ever resolved? I am implementing jboss 5.1.0.GA and get this same WARN re: Cannot create String name from non-Simple property: ManagedProperty{JNDIName,JNDIName,metaType=SimpleMetaType:java.lang.String. Is this a warning that can be ignored?	1	0	https://issues.jboss.org/browse/JBAS-6657
913	JBAS-6646	Another missing one is MBean.Platform. Also, for the MCBeans, would MCBean.* (i.e. with subtype=*) encompass all of the other MCBean subtypes (ServiceBindingSet, JTA, etc.), or do separate items need to be added to the MCBean enum for them?	1	0	https://issues.jboss.org/browse/JBAS-6646
914	JBAS-6630	Yep, that's it. See why I volunteered to take on the task? ;-)	0	0	https://issues.jboss.org/browse/JBAS-6630
915	JBAS-6625	This has been fully implemented in the EAP5 branch. See the JMXMappingUnitTestCase: Set<ManagedComponent> comps = mgtView.getComponentsForType(new ComponentType("EJB", "StatefulSession")); for (ManagedComponent comp : comps) { System.out.println(comp.getName()); ManagedProperty property = comp.getProperty("DetypedInvocationStatistics"); if ("jboss.j2ee:jndiName=ejb/StatefulCounterEjb,service=EJB".equals(comp.getName())) { MetaValue value = property.getValue(); System.out.println("Value = " + value); CompositeValue methodStatsMap = (CompositeValue)((CompositeValue) value).get("methodStats"); CompositeValue methodStats = (CompositeValue)methodStatsMap.get("count"); assertEquals(2L, ((SimpleValue)methodStats.get("count")).getValue()); return; } } fail("Could not find EJB!");	0	0	https://issues.jboss.org/browse/JBAS-6625
916	JBAS-6597	Have you tried w/o primkey-field in ejb-jar.xml? The primkey-field element is not used if the primary key maps to multiple container-managed fields (i.e. the key is a compound key). In this case, the fields of the primary key class must be public, and their names must correspond to the field names of the entity bean class that comprise the key.	0	0	https://issues.jboss.org/browse/JBAS-6597
917	JBAS-6587	Hi Adrian, this issue is closed, so I don't know if my comment reaches you. I also used YourKit to profile an app deployed on a JBoss 4.2.3 with an underlying MySQL database. Basically all stack traces where time is spent boil down to the following: <pre> -> our stuff --> some Hibernate method time spent so far: 1% ----> org.jboss.resource.adapter.jdbc.WrappedPreparedStatement.executeQuery() time spent in this method: 95% -------> net.sourceforge.jtds.jdbc.JtdsPreparedStatement.executeQuery() time spent in this method: 4% </pre> The interesting thing now is that 95% of the time is spent in WrappedPreparedStatement, BUT ONLY 4% in JtdsPreparedStatement.executeQuery(). So it looks as if 95% of the time is spent inside Jboss and only 4% in the actual query to the database!!! This also complies with the observation that our application is terribly slow, but the database load is still quite low. I debugged into org.jboss.resource.adapter.jdbc.WrappedPreparedStatement.executeQuery() to see where this time could have been spent and I have the suspicion it is in the synchronized(stateLock) block of the following method (called from WrappedPreparedStatement.executeQuery()): org.jboss.resource.adapter.jdbc.BaseWrapperManagedConnection.checkTransaction() So is this really the desired behaviour? Or are we doing something completely wrong (likely :-)? Any help is highly appreciated! Regards, Christian Naeger	1	0	https://issues.jboss.org/browse/JBAS-6587
918	JBAS-6582	Wasn't agreed that org.jboss.test.ejb3.servlet.unit.ServletUnitTestCase.testEJBServlet is not a valid usecase, and we need to remove from the suite?	1	1	https://issues.jboss.org/browse/JBAS-6582
919	JBAS-6580	I thought that might be the case when I looked closer at the failure. Would the fix also resolve JBAS-6140 that we've seen with OpenJDK?	0	0	https://issues.jboss.org/browse/JBAS-6580
920	JBAS-6568	Wrong component - should have been "JMS (JBossMQ)" - please move, thx.	1	0	https://issues.jboss.org/browse/JBAS-6568
921	JBAS-6554	Scott, is this out of date? The AS trunk component-matrix/pom.xml shows jboss-man at 2.1.1.SP1. A quick look at the subversion history for JBAS-6285 indicates Emanuel ported those changes to AS trunk.	0	0	https://issues.jboss.org/browse/JBAS-6554
922	JBAS-6549	Why do these problems show up in GA versions and not before? All JBoss has to do is load test a few apps and analyze results, no?	0	0	https://issues.jboss.org/browse/JBAS-6549
923	JBAS-6549	We also are suffering from the memory problems in AS5.0.1, so much so we've switched to IBM's JRE (which seems to help slightly with OOME's). Even still we have OOME's at least once a day in our production environment. Has anyone tried AS5.1.0 to see if it's better?	0	0	https://issues.jboss.org/browse/JBAS-6549
924	JBAS-6549	Magnus, thanks for reporting. Is this memory usage after boot up and before accessing your application? Or is this the memory usage of a clean AS 5.1 without your applications deployed?	0	0	https://issues.jboss.org/browse/JBAS-6549

#	Requirement	Questions	Answered	Successfully Answered	URL
925	JBAS-6508	Latest vfs snapshot fixes this. C:\projects\branches\jbossas\Branch_5_0\testsuite>build jboss-minimal-tests Calling ..\tools\bin\ant.bat jboss-minimal-tests Buildfile: build.xml jboss-minimal-tests: [server:start] Starting server "minimal", with command (start timeout is 120 seconds): [server:start] C:\Java\jdk1.5.0_14\bin\java -cp C:\projects\branches\jbossas\Branch_5_0\build\output\jboss-5.0.1.GA\bin\run.jar;C:\Java\jdk1.5.0_14\lib\tools.jar -Djboss.vfs.forceNoReaper=true org.jboss.Main -c minimal -b localhost -g DefaultDomain [copy] Copying 1 file to C:\projects\branches\jbossas\Branch_5_0\build\output\jboss-5.0.1.GA\server\minimal\deploy [echo] Minimal server started, stopping [delete] Deleting: C:\projects\branches\jbossas\Branch_5_0\build\output\jboss-5.0.1.GA\server\minimal\deploy\shutdown.sar BUILD SUCCESSFUL Total time: 23 seconds	0	0	https://issues.jboss.org/browse/JBAS-6508
926	JBAS-6411	Example: $ java -version java version "1.6.0_0" IcedTea6 1.4 Runtime Environment (build 1.6.0_0-b12) OpenJDK Server VM (build 10.0-b19, mixed mode) With the code change, boot.log now shows the same information: 00:40:36,112 INFO [ServerInfo] Java version: 1.6.0_0,Sun Microsystems Inc. 00:40:36,112 INFO [ServerInfo] Java Runtime: IcedTea6 1.4 (fedora-7.b12.fc10-i386) Runtime Environment (build 1.6.0_0-b12) 00:40:36,112 INFO [ServerInfo] Java VM: OpenJDK Server VM 10.0-b19,Sun Microsystems Inc.	1	1	https://issues.jboss.org/browse/JBAS-6411
927	JBAS-6411	Re: Before doing changes, I did the changes last night. Re: isn't boot.log logging *all* system properties already? It is indeed. However, we could use your same argument to remove the following two lines which also come from system properties.: 00:40:36,112 INFO [ServerInfo] Java version: 1.6.0_0,Sun Microsystems Inc. 00:40:36,112 INFO [ServerInfo] Java VM: OpenJDK Server VM 10.0-b19,Sun Microsystems Inc. The idea here is to visually have the same feedback as you get when doing java -version. It's a matter of all or nothing with regards to printing Java/VM version information. I'm happy with either as long as we're consistent. To show information partially, like it was before I did these changes can be confusing.	1	1	https://issues.jboss.org/browse/JBAS-6411
928	JBAS-6402	Is there some other configuration change/deployer update needed to deal with the current VFS issues?	1	1	https://issues.jboss.org/browse/JBAS-6402
929	JBAS-6402	Where is that?	1	1	https://issues.jboss.org/browse/JBAS-6402
930	JBAS-6392	Since this fix was integrated into AS 5 Branch_5_0, the testsuite's web server config fails to shutdown. See: http://hudson.qa.jboss.com/hudson/view/JBoss%20AS/job/JBoss-AS-5.0.x-testSuite-sun15-sun16/415/console The testsuite target is: tests-web-profile. Can you comment on this failure?	1	1	https://issues.jboss.org/browse/JBAS-6392
931	JBAS-6392	Is that a public address? I don't have access to it. I have this problem in 4.2 generally with the HTTP invoker WHENEVER I've used it at least once regardless of this change. It isn't related to initiating a Jaas authentication.	1	1	https://issues.jboss.org/browse/JBAS-6392
932	JBAS-6390	Why adding host@port syntax instead of supporting standard notation for URL containg IPV6 address? Thanks.	1	1	https://issues.jboss.org/browse/JBAS-6390
933	JBAS-6390	Could not we add support for standard URL notation containing IPV6 address? E.g.: jnp://[fec0::f101:21a:a0ff:fea0:d745]:1099 Some test cases in the testsuite rely on HOST:PORT notation. Thanks.	1	1	https://issues.jboss.org/browse/JBAS-6390
934	JBAS-6386	Do I also find that problem in the construction of my applications, can you position us on as we should proceed?	0	0	https://issues.jboss.org/browse/JBAS-6386
935	JBAS-6386	implementation 'org.jboss.web.jsf.integration.injection.JBossDelegatingInjectionProvider' does not implement the InjectionProvider interface. 2009-01-20 16:49:20,884 ERROR [javax.enterprise.resource.webcontainer.jsf] (main) Cant instantiate class: com.sun.rave.web.ui.appbase.faces.ActionListenerImpl. java.lang.InstantiationException: com.sun.rave.web.ui.appbase.faces.ActionListenerImpl at java.lang.Class.newInstance0(Class.java:335) at java.lang.Class.newInstance(Class.java:303) at com.sun.faces.util.Util.createInstance(Util.java:471) at com.sun.faces.config.ConfigureListener.configure(ConfigureListener.java:547) at com.sun.faces.config.ConfigureListener.configure(ConfigureListener.java:487) at com.sun.faces.config.ConfigureListener.contextInitialized(ConfigureListener.java:381) at org.apache.catalina.core.StandardContext.listenerStart(StandardContext.java:3910) at org.apache.catalina.core.StandardContext.start(StandardContext.java:4393) at org.jboss.web.tomcat.service.deployers.TomcatDeployment.performDeployInternal(TomcatDeployment.java:367) at org.jboss.web.tomcat.service.deployers.TomcatDeployment.performDeploy(TomcatDeployment.java:146) at org.jboss.web.deployers.AbstractWarDeployment.start(AbstractWarDeployment.java:460) at org.jboss.web.deployers.WebModule.startModule(WebModule.java:118) at org.jboss.web.deployers.WebModule.start(WebModule.java:96) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.system.microcontainer.ServiceProxy.invoke(ServiceProxy.java:206) at $Proxy36.start(Unknown Source) at org.jboss.system.microcontainer.StartStopLifecycleAction.installAction(StartStopLifecycleAction.java:42) at org.jboss.system.microcontainer.StartStopLifecycleAction.installAction(StartStopLifecycleAction.java:37) at org.jboss.dependency.plugins.action.SimpleControllerContextAction.simpleInstallAction(SimpleControllerContextAction.java:62) at org.jboss.dependency.plugins.action.AccessControllerContextAction.install(AccessControllerContextAction.java:71) at org.jboss.dependency.plugins.AbstractControllerContextActions.install(AbstractControllerContextActions.java:51) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:348) at org.jboss.system.microcontainer.ServiceControllerContext.install(ServiceControllerContext.java:286) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:1598) at org.jboss.dependency.plugins.AbstractController.incrementState(AbstractController.java:934) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:1062) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:984) at	1	0	https://issues.jboss.org/browse/JBAS-6386
936	JBAS-6378	Jason, where does "AS5 Administration And Configuration Guide" live?	1	1	https://issues.jboss.org/browse/JBAS-6378
937	JBAS-6378	Jason, I wrote a new Remoting chapter for EAP 5. Is there any reason I shouldn't commit it for AS 5 as well?	0	0	https://issues.jboss.org/browse/JBAS-6378
938	JBAS-6346	Yeah weird, I didn't realize until your comment that the ClusterChooser had been factored out of the invocation retry (which is how it was done in the EJB2 stack). Is that really wise? It isn't obvious to me that they can really be separate concerns along the chain of command pattern that forms the client interceptor stack.	0	0	https://issues.jboss.org/browse/JBAS-6346
939	JBAS-6334	When do you expect 1.4.3.GA to be ready? In the meantime should we move to 1.4.2.GA.SP1 in AS Branch_5_x?	0	0	https://issues.jboss.org/browse/JBAS-6334

#	Requirement	Questions	Answered	Successfully Answered	URL
940	JBAS-6334	I think we should update jbm 1.4.2.GA.SP1 to Branch_5_x. Because it's the latest release. As for 1.4.3.GA, I'm not sure. Can you make some comment Clebert?	0	0	https://issues.jboss.org/browse/JBAS-6334
941	JBAS-6334	So which release is the target and when to expect it?	1	1	https://issues.jboss.org/browse/JBAS-6334
942	JBAS-6330	I've committed initial work into 5_x branch. Max, can you check if this works for you Tools guys?	1	1	https://issues.jboss.org/browse/JBAS-6330
943	JBAS-6330	What exactly was this list? String, array, ... What did it show? Folders to scan, actual deployments, ... This is probably the old interface: - http://anonsvn.jboss.org/repos/jbossas/branches/Branch_5_x/system-jmx/src/main/org/jboss/deployment/scanner/URLDeploymentScannerMBean.java	1	1	https://issues.jboss.org/browse/JBAS-6330
944	JBAS-6330	This should now be fixed in 5_x branch. Rob, can you try it?	1	1	https://issues.jboss.org/browse/JBAS-6330
945	JBAS-6330	1) Fix removeURL to make sure it undeploys 2) I'm not seeing suspendDeployment or resumeDeployment at all here. Eh?	1	1	https://issues.jboss.org/browse/JBAS-6330
946	JBAS-6330	Cool. Ales - how about 5.0.x ? are those fixes rolled back, the same as in 5.x or some partial/halfbaked state ?	1	1	https://issues.jboss.org/browse/JBAS-6330
947	JBAS-6330	This patch is for unit tests. It has found a bug in listDeployedURLs I am not sure that the unit test works 100% even after listDeployedURLs is fixed so the unit test may need some tweaking. Another issue is that the full test takes 40 seconds to run. This seems a little long for me, but the nature of deployment scanning takes a while. Is there a way to force a scan so that maybe it can go faster?	1	1	https://issues.jboss.org/browse/JBAS-6330
948	JBAS-6330	better unit test. Nothing needs to be changed on the server side... but I did need to increase the delay on the client side. I have a general feeling 5 seconds is too much but what can I do to lower it? If I lower it it fails.	1	1	https://issues.jboss.org/browse/JBAS-6330
949	JBAS-6320	After some more investigating it seems that jboss as jars are now located at http://repository.jboss.org/maven2/org/jboss/jbossas If so, it looks like http://repository.jboss.org/maven2/org/jboss/jbossas/jboss-as-client/5.0.0.GA/jboss-as-client-5.0.0.GA.pom could fit the bill. Could you confirm that this pom is equivalent to the old jbossall-client.jar?	0	0	https://issues.jboss.org/browse/JBAS-6320
950	JBAS-6245	Kabir can you take a look please? http://hudson.qa.jboss.com/hudson/view/JBoss%20AS/job/JBoss-AS-5.0.x-TestSuite-sun15-sun16/lastBuild/testReport/org.jboss.test.asynch/AsynchTestCase/unknown/ I think this started when I synched o.j.invocation.InvokerInterceptor with the one from Branch_4_2 and the new version has a an anonymous inner class /** These objects are safe to reuse in callByValue operations */ static final SafeClone safeToReuse = new SafeClone(){ public boolean isSafeToReuse(Object obj) { if (obj==null) { return false; I can see the class included in the jboss.jar and jboss-client.jar What's the problem with the inner class? Caused by: org.jboss.aop.instrument.TransformationException: Failed to aspectize class org.jboss.invocation.InvokerInterceptor. Could not find class it references org.jboss.invocation.InvokerInterceptor$1 It may not be in your classpath and you may not be getting field and constructor weaving for this class. at org.jboss.aop.instrument.Instrumentor.convertReferences(Instrumentor.java:708) at org.jboss.aop.instrument.Instrumentor.transform(Instrumentor.java:763) at org.jboss.aop.instrument.GeneratedAdvisorInstrumentor.transform(GeneratedAdvisorInstrumentor.java:119) at org.jboss.aop.SuperClassesFirstWeavingStrategy.instrumentClass(SuperClassesFirstWeavingStrategy.java:202)	1	1	https://issues.jboss.org/browse/JBAS-6245
951	JBAS-6244	HI all, when this problem occur , it seems that the exceptions mechanism in Jboss is affected. Another problem is that the affected timer stops. In My opinion the transactions that cannot recover , should be just quit. Can i do this ..somehow? Or is there another way to deal with the situation ? Thanks in advance!	1	1	https://issues.jboss.org/browse/JBAS-6244
952	JBAS-6242	In my opinion, what it is doing is correct. The format of the statistics is determined by the statistics formatter. Which by default returns a String, but it could return anything. A collection of JBossSubPoolStatistics would be the most obvious alternative which would map to a collection of composite types for the same. You need to be able handle the return type dynamically, Its called polymorphism which is an integral part of any object orientated language. :-) i.e. MetaValue returnValue = invoke(...) MetaType type = returnValue.getMetaType(); // etc. If you can't handle that, then I'd suggest we add a new method which hardwires the String formatter with a String return type and only annotate that as a ManagedOperation. But this issue will come up elsewhere where the return type cannot be constrained so easily.	0	0	https://issues.jboss.org/browse/JBAS-6242
953	JBAS-6235	Not joining the cluster until everything else is started isn't really possible, because the proper startup of some of that "everything else" requires that the node has joined the cluster. But, clustered services should still work properly during the startup/shutdown period. To the extent they aren't, that's something we need to address on a case by case basis. Could you open a thread on http://www.jboss.com/index.html?module=bb&op=viewforum&f=155 and give more details about your specific issues?	1	0	https://issues.jboss.org/browse/JBAS-6235
954	JBAS-6226	Needs to be rolled back. Stefan, can u do the rollback?	0	0	https://issues.jboss.org/browse/JBAS-6226
955	JBAS-6221	This "bug report" contains at least two misunderstandings. 1) JMSXDeliveryCount is set by the jms provider. The DLQHandler just references it. (Also, if your jms provider already has DLQ handling on the server - unlike JBossMQ - then you don't need any DLQ handling within JBoss). 2) EJB3 has never referenced standardjboss.xml or had invoker/container-configurations It uses aop and annotations to override/control configurations. For (1) you should talk to your jms provider. For (2) I've discussed this before with the EJB3 team (I won't go into details here). So I'm making this a feature request for them. Discussions around this feature request belong in the forums (which is where also this post should have originated). JIRA is for tracking what gets done, it is the release notes. It should not be used for confusing other users with misguided speculation.	0	0	https://issues.jboss.org/browse/JBAS-6221

#	Requirement	Questions	Answered	Successfully Answered	URL
956	JBAS-6219	client passes back an updated copy that is never written back onto the original template properties. Is there some way your caching the original template info? I have added logging of the DeploymentTemplateInfo as received and after updated, and for this output I have run the testAddDataSource test twice, but I still see null/default values in the template properties that are updated to non-null values: 1552 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: valid-connection-checker-class-name,null 1552 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: valid-connection-checker-class-name,SimpleMetaType:java.lang.String:org.jboss.resource.adapter.jdbc.vendor.DummyValidConnectionChecker 1553 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: jndi-name,SimpleMetaType:java.lang.String: 1553 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: jndi-name,SimpleMetaType:java.lang.String:TestLocalTxDs 1553 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: validate-on-match,SimpleMetaType:java.lang.Boolean:true 1553 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: validate-on-match,SimpleMetaType:java.lang.Boolean:false 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: connection-url,null 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: connection-url,SimpleMetaType:java.lang.String:jdbc:hsqldb:. 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: metadata,null 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: metadata,CompositeValueSupport: metaType=[MutableCompositeMetaType{org.jboss.resource.metadata.mcf.DBMSMetaData items=[name=typeMapping type=java.lang.String]}] items=[typeMapping=null] 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: url-delimiter,null 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: url-delimiter,SimpleMetaType:java.lang.String:+ 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: driver-class,null 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: driver-class,SimpleMetaType:java.lang.String:org.hsqldb.jdbcDriver 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: min-pool-size,SimpleMetaType:java.lang.Integer:0 1554 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: min-pool-size,SimpleMetaType:java.lang.Integer:0 1555 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: prepared-statement-cache-size,SimpleMetaType:java.lang.Integer:0 1555 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: prepared-statement-cache-size,SimpleMetaType:java.lang.Integer:12 1555 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: exception-sorter-class-name,null 1555 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: exception-sorter-class-name,SimpleMetaType:java.lang.String:org.jboss.resource.adapter.jdbc.vendor.DummyExceptionSorter 1555 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property before: max-pool-size,SimpleMetaType:java.lang.Integer:10 1555 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - template property after: max-pool-size,SimpleMetaType:java.lang.Integer:11 1555 [main] DEBUG	1	1	https://issues.jboss.org/browse/JBAS-6219
957	JBAS-6204	In the current trunk I added a test of the security-domain property type and value of the DefaultDS and its returning the expected GenericMetaType/GenericMetaValue: 1056 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - security-domain type: GenericMetaType:org.jboss.managed.api.ManagedObject 1056 [main] DEBUG org.jboss.test.profileservice.test.ProfileServiceUnitTestCase - security-domain value: GenericMetaType:org.jboss.managed.api.ManagedObject:ManagedObject{org.jboss.resource.metadata.mcf.SecurityDomainMetaData} Is there a specific datasource deployment your seeing the incorrect type on?	0	0	https://issues.jboss.org/browse/JBAS-6204
958	JBAS-6200	Shelly, can this be resolved now?	0	0	https://issues.jboss.org/browse/JBAS-6200
959	JBAS-6178	Did you just add the missing org.jboss.invocation.MarshalledValueEX class? In my local copy I've more or less synched the Branch_4_2 changes of org/jboss/invocation/** and there are quite a few. With that I believe I get similar results.	1	1	https://issues.jboss.org/browse/JBAS-6178
960	JBAS-6171	Isn't this moving to the next point release?	0	0	https://issues.jboss.org/browse/JBAS-6171
961	JBAS-6144	Would synchronous replication help? With async there is a small window (or it could be larger depending on the communication channel queue length). If we want to track down the why, maybe some debug logging in key places would help. Certain applications that require absolute data integrity would be happier with sync (IMO). Applications that don't mind a short period of dirty data on the failed over to node, would be happy with async (but the test probably needs enhancing to properly handle async IMO).	1	0	https://issues.jboss.org/browse/JBAS-6144
962	JBAS-6139	Can you coordinate with Rajesh to make the JBOSS-AS-5.0.x build work by accessing the jboss-repository only? http://hudson.qa.jboss.com/hudson/view/JBoss%20AS/job/JBoss-AS-5.0.x/	0	0	https://issues.jboss.org/browse/JBAS-6139
963	JBAS-6133	Shouldn't this issue be moved to JBossAS-6.x?	0	0	https://issues.jboss.org/browse/JBAS-6133
964	JBAS-6133	I updated trunk to use Quartz 1.7.3 in r102577 [1]. I had to add a few methods to JBossQuartzThreadPool to implement the interface. Ales, can you take a look at my changes? I wasn't sure what to put for the new method blockForAvailalbeThreads() so I just return 1 to say that there is always a thread available. [1]http://fisheye.jboss.org/changelog/JBossAS/?cs=102577	1	1	https://issues.jboss.org/browse/JBAS-6133
965	JBAS-6125	Is this setting on the client side?	1	1	https://issues.jboss.org/browse/JBAS-6125
966	JBAS-6125	So, what's then that got to do with tests? Tests == client side. ;-)	1	1	https://issues.jboss.org/browse/JBAS-6125
967	JBAS-6105	Jesper can you profile AS in trunk to find out what's wrong?	0	0	https://issues.jboss.org/browse/JBAS-6105
968	JBAS-6105	Ales, could you take a look at the BaseClassLoader and check what's going on ? Feel free to reassign back to me and I'll take another look once I get back	0	0	https://issues.jboss.org/browse/JBAS-6105
969	JBAS-6105	This was recently fixed: https://jira.jboss.org/jira/browse/JBCL-49 Jasper, can you re-run your profiling, but with jboss-cl build from trunk?	0	0	https://issues.jboss.org/browse/JBAS-6105
970	JBAS-6055	There are two versions of jsp-api and servlet-api currently being used javax.servlet.jsp:jsp-api:jar:2.1 vs. jboss.web:jsp-api:jar:2.1.1.GA javax.servlet:servlet-api:jar:2.5 vs. jboss.web:servlet-api:jar:2.1.1.GA Are these duplicates, or are they both required?	1	1	https://issues.jboss.org/browse/JBAS-6055
971	JBAS-6055	I've cleaned up most of the dependency issues identified by Fernando and Dimitrius. I'm not sure how to resolve the servlet-api and jsp-api issue. Based on the discussion on jboss-development should we use the one from jboss.web? The other remaining issue is the xml-apis duplication. We have: apache-xerces:xml-apis:jar:2.9.1 and this is pulled in transitively. xml-apis:xml-apis:jar:1.0.b2 Which one should we use?	1	1	https://issues.jboss.org/browse/JBAS-6055

#	Requirement	Questions	Answered	Successfully Answered	URL
972	JBAS-6055	Re: FastInfoset.jar, my guess is we should be importing only sun-fi, since this is the one explicitly specified in the component-matrix pom: [INFO] +- sun-fi:FastInfoset:jar:1.2.2:compile [INFO] +- sun-fi:FastInfosetUtilities:jar:1.2.2:compile jbossws uses com.sun.xml.fastinfoset:FastInfoset:jar:1.2.2:compile, but I understand it's essentially the same binary under a different name. The sun-jaxws one, should propably be exluded, too: [INFO] +- sun-jaxws:FastInfoset:jar:2.1.1:compile What's gets deployed with AS is a mystery to me, because of the stupid jbossws-deploy-macros.xml: <patternset id="jbossws.service.lib.patternset"> <include name="**/FastInfoset.jar"/> By comparing the jars, I think the sun-fi is deployed.	1	1	https://issues.jboss.org/browse/JBAS-6055
973	JBAS-6047	Ivo, these tests started failing after your testsuite changes to JBAS-5349. The first issue has been resolved. Can you look at this issue further and determine if your changes affect these remaining two tests. If not, pls reassign back to me.	1	1	https://issues.jboss.org/browse/JBAS-6047
974	JBAS-6013	How do I create such a queue using @MessageDriven annotation only? I tried to build latest jboss where I could find all the new stuff but could not figure out how to do this.	1	0	https://issues.jboss.org/browse/JBAS-6013
975	JBAS-6013	Where can I find documentation on how to use this CreateDestination feature with an annotation based MDB?	0	0	https://issues.jboss.org/browse/JBAS-6013
976	JBAS-6008	What you are doing wrong is posting help requests to JIRA. In what does this issue in your ant task/jndi configuration need any work doing on the application server by any jboss developer? Use the forums to ask for help.	0	0	https://issues.jboss.org/browse/JBAS-6008
977	JBAS-5965	The error-message using 5.0.0-CR2 is as follow: bash-3.2$ ./run.sh ./run.sh: line 89: ulimit: open files: cannot modify limit: Invalid argument run.sh: Could not set maximum file descriptor limit: unlimited ========== ... I hope to find time to fix it this week.	1	1	https://issues.jboss.org/browse/JBAS-5965
978	JBAS-5965	Marc, I having the issue you mentioned with 4.2.3 It looks like this: (I'm using JDK 1.6) jboss/bin/./run.sh: line 89: ulimit: open files: cannot modify limit: Invalid argument run.sh: Could not set maximum file descriptor limit: unlimited ========== JBoss Bootstrap Environment JBOSS_HOME: /Users/Troy/jboss JAVA: java JAVA_OPTS: -Dprogram.name=run.sh -Xms128m -Xmx512m -Dsun.rmi.dgc.client.gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000 CLASSPATH: /Users/Troy/jboss/bin/run.jar jboss starts, but I cannot seem to get a page up, because of this: description The requested resource (/hello) is not available. JBossWeb/2.0.1.GA I would really appreciate any suggestions. I've searched everywhere. the same jboss works fine with jdk 1.6 in windows. Is there something wrong with the environment variables? thanks	1	1	https://issues.jboss.org/browse/JBAS-5965
979	JBAS-5965	I suppose you are talking about OSX Leopard. Did you add the line "ulimit -n 1024" to run.sh? It will set the maximum number of open files. The error message will disappear.	1	1	https://issues.jboss.org/browse/JBAS-5965
980	JBAS-5965	Hi Marc, I didn't see your jboss_init_osx.sh attached. Did you write one? Can you share? Thanks...	1	1	https://issues.jboss.org/browse/JBAS-5965
981	JBAS-5945	Change is a bit more complex than what's in the JIRA description. What's going on in initAfterLoad() is really three things: 1) Storing refs to the session Manager and some of its components. This is eliminated; the session no longer drops its refs via serialization, so it already has them. 2) populateAttributes(). Moved into update(), with real work actually done in the ha-server-cache layer. 3) Activation notifications. This is still called by JBossCacheManager, but we just use the existing notifyDidActivate method.	0	0	https://issues.jboss.org/browse/JBAS-5945
982	JBAS-5930	Is EJBTHREE-1487 suficient for the changes you want to make?	0	0	https://issues.jboss.org/browse/JBAS-5930
983	JBAS-5928	Branch_4_2 haven't IMAPMailFolder/POP3MailFolder too. Seam, they have their branches about mail.rar, right? Will seam depends on Branch_4_2 or trunk after AS 5.0 release?	1	1	https://issues.jboss.org/browse/JBAS-5928
984	JBAS-5927	There are still a couple of dependencies on CR versions of other stuff. jboss-aop-jdk50-client - 2.0.0.CR16 jboss-jca-api - 1.5.0.CR2 jboss-transaction-api - 1.0.1.CR2 Should I wait until these are released before doing a GA release of integration?	0	0	https://issues.jboss.org/browse/JBAS-5927
985	JBAS-5920	The issue is why does it have (JBoss)AppClientMetaData at all. The problem is the AppClientScanningDeployer: String mainClassName = getMainClassName(unit); if(mainClassName == null) return; Class<?> mainClass = unit.getClassLoader().loadClass(mainClassName); log.info("mainClass = " + mainClass); if(hasAnnotations(mainClass)) { // add a dummy application client dd to fire up the ejb3 client deployer dd = new JBossClientMetaData(); unit.addAttachment(JBossClientMetaData.class, dd); } // TODO: should we check for type of annotations? private boolean hasAnnotations(Class<?> cls) { ... } The answer to that question is that it should only process jars with a main-class if it has the EJB3 environment annotations, e.g. @Resource Otherwise it is going to make most jars with a main-class in the manifest appclients (with the likely assumption that any modern class has some annotations). The NPE still needs fixiing: JBossAppMetaData earMD = earUnit.getAttachment(JBossAppMetaData.class); // oops earMD == null when the parent is not an ear. String libDir = earMD.getLibraryDirectory(); The AppClientScanningDeployer should probably include a similar test for the parent being an ear?	1	1	https://issues.jboss.org/browse/JBAS-5920
986	JBAS-5918	Oddly the issue doesn't appear when I execute only one target with few tests between start and stop. Anyways, the bug seems related to https://issues.apache.org/bugzilla/show_bug.cgi?id=41368 Should I care to try out the patch and go for patching the built-in ant, or just explain in the build.xml why the stop command is commented out?	0	0	https://issues.jboss.org/browse/JBAS-5918
987	JBAS-5905	Just in case, opened a forum thread to get input from JBM folks as to why this is this way. But based on following chat w/ Bela, I'll add FD; if there's a reason not to, it can always be removed for 5.0.0.GA. [10:26:22] Brian Stansberry: can you imagine any reason to just have FD_SOCK and not FD in a TCP-based stack? [10:26:46] Bela Ban: No, not really	1	1	https://issues.jboss.org/browse/JBAS-5905
988	JBAS-5900	Hi, Can some one please give me the solution in jboss-structure.xml in JBoss5.1.0? That would be great.	0	0	https://issues.jboss.org/browse/JBAS-5900
989	JBAS-5891	Let me understand: are you fixing those 2 tests as well (or what is needed?)	1	1	https://issues.jboss.org/browse/JBAS-5891
990	JBAS-5872	I agree. Additionally, I have been looking to incorporate the tests-matrix for AS 5.0 considering the client/server combinations: 1) 4.2.x jdk5 ==> 5.0 jdk5 2) 4.2.x jdk6 ==> 5.0 jdk6 3) 4.2.x jdk5 ==> 5.0 jdk6 The tests-matrix test targets and associated sources were removed from jbossas/trunk. Is there some history of why that was deliberately done? With the client/server combination of #1) above, all tests fail due to: 2008-08-21 14:09:51,900 ERROR [org.jboss.remoting.transport.socket.ServerThread] (WorkerThread#0[127.0.0.1:53720]) WorkerThread#0[127.0.0.1:53720] exception occurred during first invocation java.lang.ClassNotFoundException: org.jboss.invocation.MarshalledValueEX at java.net.URLClassLoader$1.run(URLClassLoader.java:200)	1	1	https://issues.jboss.org/browse/JBAS-5872
991	JBAS-5872	If the jdks themselves are inconsistent regarding their SUIDs then it's not our fault :) I mean what can we do to correct this?	1	0	https://issues.jboss.org/browse/JBAS-5872
992	JBAS-5870	I've created this JIRA here because I'm not sure whether such handler would end up going in AS or Common because the handler needs to be available to the System classloader (sun.misc.Launcher$AppClassLoader) so maybe it'd need to live in run.jar?	1	0	https://issues.jboss.org/browse/JBAS-5870

#	Requirement	Questions	Answered	Successfully Answered	URL
993	JBAS-5867	Is this the correct fix? This issue was also seen in the refs tests (JBAS-5806). The refs tests had similar lookups for the remote interface: ic.lookup("BeanImpl3x/remote") but changed the tests to do the lookup using mappedName. So the way to do the lookup is to include the entire interface name in the lookkup? - BeanRemote3x bean = (BeanRemote3x) ic.lookup("BeanImpl3x/remote-org.jboss.test.profileservice.ejb3x.BeanRemote3x"); Or, we could used the bean mappedName too? + BeanRemote3x bean = (BeanRemote3x) ic.lookup("profileservice/ejb3x/BeanRemote3x");	0	0	https://issues.jboss.org/browse/JBAS-5867
994	JBAS-5866	What type of logs do you want. Something that looks weird to me is the testsuite executions do not produce the output/log/test.log anymore.	0	0	https://issues.jboss.org/browse/JBAS-5866
995	JBAS-5866	A test log can be found here: http://hudson.qa.jboss.com/hudson/view/JBoss%20AS/job/JBoss-AS-5.0.x-TestSuite-sun15/869/artifact/trunk/testsuite/output/log/test.log Do you need something more?	0	0	https://issues.jboss.org/browse/JBAS-5866
996	JBAS-5856	This task was about placing properties on the ports.. I was planning to use the old JMX style XML. I thought you were providing that kind of change. (sorry for misunderstanding your comments) We are moving towards JBM 2, and we don't have the MC in place at JBM, so we can't maintain / test this configuration. Is it okay if we keep the old XML style?	0	0	https://issues.jboss.org/browse/JBAS-5856
997	JBAS-5846	java.lang.NullPointerException java.lang.Exception: java.lang.NullPointerException at org.jboss.test.asynch.AsynchTestCase.testLocal(AsynchTestCase.java:93) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) Caused by: java.lang.NullPointerException at org.jboss.aop.Advisor.dynamicInvoke(Advisor.java:1465) at org.jboss.aop.Dispatcher.invoke(Dispatcher.java:101) at org.jboss.aspects.remoting.AOPRemotingInvocationHandler.invoke(AOPRemotingInvocationHandler.java:82) at org.jboss.remoting.ServerInvoker.invoke(ServerInvoker.java:847) at org.jboss.remoting.transport.socket.ServerThread.completeInvocation(ServerThread.java:681) at org.jboss.remoting.transport.socket.ServerThread.processInvocation(ServerThread.java:634) at org.jboss.remoting.transport.socket.ServerThread.dorun(ServerThread.java:459) at org.jboss.remoting.transport.socket.ServerThread.run(ServerThread.java:197) at org.jboss.remoting.MicroRemoteClientInvoker.invoke(MicroRemoteClientInvoker.java:203) at org.jboss.remoting.Client.invoke(Client.java:1685) at org.jboss.remoting.Client.invoke(Client.java:589) at org.jboss.aspects.remoting.InvokeRemoteInterceptor.invoke(InvokeRemoteInterceptor.java:60) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.MergeMetaDataInterceptor.invoke(MergeMetaDataInterceptor.java:74) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.tx.ClientTxPropagationInterceptor.invoke(ClientTxPropagationInterceptor.java:61) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.security.SecurityClientInterceptor.invoke(SecurityClientInterceptor.java:65) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.IsLocalInterceptor.invoke(IsLocalInterceptor.java:48) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.asynch.AsynchProxyInterceptor.invoke(AsynchProxyInterceptor.java:53) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at AOPClassProxy$1.test(AOPClassProxy$1.java) at org.jboss.test.asynch.AsynchTestCase.testLocal(AsynchTestCase.java:88) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23)	0	0	https://issues.jboss.org/browse/JBAS-5846
998	JBAS-5846	Kabir, can you comment on these tests? Do we want to preserve these in AS 5 testsuite?	0	0	https://issues.jboss.org/browse/JBAS-5846
999	JBAS-5846	We can remove the test-aspects target from the suite but want to be sure the failures in the local and collocated tests is not exposing a real issue. Can you confirm that we should exclude the tests for 5.0?	0	0	https://issues.jboss.org/browse/JBAS-5846
1000	JBAS-5846	Is the new release ready?	1	1	https://issues.jboss.org/browse/JBAS-5846
1001	JBAS-5846	The proxy is created and bound to localhost is set by an mbean running on the server. What is the best way to determine the IP address the server is bound to? HostAddress attribute of jboss.system:type=ServerInfo?	1	1	https://issues.jboss.org/browse/JBAS-5846

#	Requirement	Questions	Answered	Successfully Answered	URL
1002	JBAS-5846	I tried using HostAddress attribute of jboss.system:type=ServerInfo, but then I get the following error locally estcase: testRemote took 0.479 sec Caused an ERROR org.jboss.remoting.CannotConnectException: Can not get connection to server. Problem establishing socket connection for InvokerLocator [socket://192.168.1.69:5150/] java.lang.RuntimeException: org.jboss.remoting.CannotConnectException: Can not get connection to server. Problem es tablishing socket connection for InvokerLocator [socket://192.168.1.69:5150/] at org.jboss.test.asynch.AsynchTestCase.testRemote(AsynchTestCase.java:74) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) Caused by: org.jboss.remoting.CannotConnectException: Can not get connection to server. Problem establishing socket connection for InvokerLocator [socket://192.168.1.69:5150/] at org.jboss.remoting.transport.socket.MicroSocketClientInvoker.transport(MicroSocketClientInvoker.java:771) at org.jboss.remoting.MicroRemoteClientInvoker.invoke(MicroRemoteClientInvoker.java:160) at org.jboss.remoting.Client.invoke(Client.java:1708) at org.jboss.remoting.Client.invoke(Client.java:612) at org.jboss.aspects.remoting.InvokeRemoteInterceptor.invoke(InvokeRemoteInterceptor.java:60) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.MergeMetaDataInterceptor.invoke(MergeMetaDataInterceptor.java:74) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.tx.ClientTxPropagationInterceptor.invoke(ClientTxPropagationInterceptor.java:61) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.security.SecurityClientInterceptor.invoke(SecurityClientInterceptor.java:65) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.IsLocalInterceptor.invoke(IsLocalInterceptor.java:48) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.asynch.AsynchProxyInterceptor.invoke(AsynchProxyInterceptor.java:53) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at AOPClassProxy$0.testMethod(AOPClassProxy$0.java) at org.jboss.test.asynch.AsynchTestCase.testRemote(AsynchTestCase.java:60) ... 21 more	1	1	https://issues.jboss.org/browse/JBAS-5846
1003	JBAS-5846	THis is a different error from before. The MBean that sets up the proxy will bind it to the correct IP address. If I start up the server myself and bind it to another IP address, I get the same error as on Hudson, which happens when trying to deploy the test MBean to the server. I think something needs setting in jndi.properties or somewhere to be able to connect to the correct server, since that still uses localhost:1099. How is this normally done? $more ../aspects/output/reports/TEST-org.jboss.test.asynch.AsynchTestCase.txt Testsuite: org.jboss.test.asynch.AsynchTestCase Tests run: 0, Failures: 0, Errors: 1, Time elapsed: 5.126 sec Testcase: unknown took 0 sec Caused an ERROR Could not obtain connection to any of these urls: localhost:1099 and discovery failed with error: javax.naming.CommunicationException: Re ceive timed out [Root exception is java.net.SocketTimeoutException: Receive timed out] javax.naming.CommunicationException: Could not obtain connection to any of these urls: localhost:1099 and discovery failed with error: ja vax.naming.CommunicationException: Receive timed out [Root exception is java.net.SocketTimeoutException: Receive timed out] [Root excepti on is javax.naming.CommunicationException: Failed to connect to server localhost:1099 [Root exception is javax.naming.ServiceUnavailableE xception: Failed to connect to server localhost:1099 [Root exception is java.net.ConnectException: Connection refused]]] at org.jnp.interfaces.NamingContext.checkRef(NamingContext.java:1604) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:636) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:629) at javax.naming.InitialContext.lookup(InitialContext.java:351) at org.jboss.test.JBossTestServices.getServer(JBossTestServices.java:154) at org.jboss.test.JBossTestSetup.getServer(JBossTestSetup.java:108) at org.jboss.test.aop.test.AOPTestSetup.setUp(AOPTestSetup.java:54) at junit.extensions.TestSetup$1.protect(TestSetup.java:18) at junit.extensions.TestSetup.run(TestSetup.java:23) Caused by: javax.naming.CommunicationException: Failed to connect to server localhost:1099 [Root exception is javax.naming.ServiceUnavail ableException: Failed to connect to server localhost:1099 [Root exception is java.net.ConnectException: Connection refused]] at org.jnp.interfaces.NamingContext.getServer(NamingContext.java:276) at org.jnp.interfaces.NamingContext.checkRef(NamingContext.java:1575) ... 12 more Caused by: javax.naming.ServiceUnavailableException: Failed to connect to server localhost:1099 [Root exception is java.net.ConnectExcept ion: Connection refused] at org.jnp.interfaces.NamingContext.getServer(NamingContext.java:250) ... 13 more Caused by: java.net.ConnectException: Connection refused at java.net.PlainSocketImpl.socketConnect(Native Method) at java.net.PlainSocketImpl.doConnect(PlainSocketImpl.java:333) at java.net.PlainSocketImpl.connectToAddress(PlainSocketImpl.java:195) at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:182) at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:432) at java.net.Socket.connect(Socket.java:520) at java.net.Socket.connect(Socket.java:470) at java.net.Socket.<init>(Socket.java:367) at java.net.Socket.<init>(Socket.java:267) at org.jnp.interfaces.TimedSocketFactory.createSocket(TimedSocketFactory.java:84) at org.jnp.interfaces.TimedSocketFactory.createSocket(TimedSocketFactory.java:77) at org.jnp.interfaces.NamingContext.getServer(NamingContext.java:246) ... 13 more	1	1	https://issues.jboss.org/browse/JBAS-5846

#	Requirement	Questions	Answered	Successfully Answered	URL
1004	JBAS-5846	the InvokerLocator is set up. I'll investigate Testcase: testRemote took 0.521 sec Caused an ERROR org.jboss.remoting.CannotConnectException: Can not get connection to server. Problem establishing socket connection for InvokerLocator [s ocket://10.0.0.1:5150/] java.lang.RuntimeException: org.jboss.remoting.CannotConnectException: Can not get connection to server. Problem establishing socket conn ection for InvokerLocator [socket://10.0.0.1:5150/] at org.jboss.test.asynch.AsynchTestCase.testRemote(AsynchTestCase.java:74) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) Caused by: org.jboss.remoting.CannotConnectException: Can not get connection to server. Problem establishing socket connection for Invoke rLocator [socket://10.0.0.1:5150/] at org.jboss.remoting.transport.socket.MicroSocketClientInvoker.transport(MicroSocketClientInvoker.java:771) at org.jboss.remoting.MicroRemoteClientInvoker.invoke(MicroRemoteClientInvoker.java:160) at org.jboss.remoting.Client.invoke(Client.java:1708) at org.jboss.remoting.Client.invoke(Client.java:612) at org.jboss.aspects.remoting.InvokeRemoteInterceptor.invoke(InvokeRemoteInterceptor.java:60) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.MergeMetaDataInterceptor.invoke(MergeMetaDataInterceptor.java:74) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.tx.ClientTxPropagationInterceptor.invoke(ClientTxPropagationInterceptor.java:61) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.security.SecurityClientInterceptor.invoke(SecurityClientInterceptor.java:65) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.IsLocalInterceptor.invoke(IsLocalInterceptor.java:48) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.asynch.AsynchProxyInterceptor.invoke(AsynchProxyInterceptor.java:53) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at AOPClassProxy$0.testMethod(AOPClassProxy$0.java) at org.jboss.test.asynch.AsynchTestCase.testRemote(AsynchTestCase.java:60) ... 21 more Caused by: java.net.ConnectException: Connection refused at java.net.PlainSocketImpl.socketConnect(Native Method) at java.net.PlainSocketImpl.doConnect(PlainSocketImpl.java:333) at java.net.PlainSocketImpl.connectToAddress(PlainSocketImpl.java:195) at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:182) at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:432) at java.net.Socket.connect(Socket.java:520) at org.jboss.remoting.util.SecurityUtility.connect(SecurityUtility.java:969) at org.jboss.remoting.transport.socket.SocketClientInvoker.createSocket(SocketClientInvoker.java:203) at org.jboss.remoting.transport.socket.MicroSocketClientInvoker.getConnection(MicroSocketClientInvoker.java:1084) at	1	1	https://issues.jboss.org/browse/JBAS-5846
1005	JBAS-5836	Sure; we have code freeze for jbossws 3.0.3 at the end of this week. The release date is currently fixed for the 19th of August, but we can try to release earlier if this is required. Does this meet the plans for AS 5.0.0.CR2?	1	1	https://issues.jboss.org/browse/JBAS-5836
1006	JBAS-5764	Updated to JSF Mojarra 1.2_09 Beta. This is only for Branch_4_2. When 1.2_09 goes final, all branches will be updated with the GA version.	0	0	https://issues.jboss.org/browse/JBAS-5764
1007	JBAS-5759	Sorry, I thought we were talking about AS 5. AS 5 uses aop 2. AS 4.2.x uses AOP 1.5.x. I can cut a 1.5.7 release later this week if you want?	0	0	https://issues.jboss.org/browse/JBAS-5759
1008	JBAS-5743	Did you read my post? For 'jboss.management.local:J2EEServer=Local,j2eeType=JVM,name=localhost' the 'name' is hardcoded: JVM jvm = new JVM("localhost", lServer, System.getProperty("java.version"), System.getProperty("java.vendor"), hostName); According to the jsr77 spec the value of name is arbitrary and should identify the resource, in this case the JVM the server is running on. In this regard a more appropriate name would be the value of one (or more) of the properties: java.runtime.name Java(TM) 2 Runtime Environment, Standard Edition java.runtime.version 1.5.0_09-b03 java.specification.name Java Platform API Specification java.specification.vendor Sun Microsystems Inc. java.specification.version 1.5 java.vendor Sun Microsystems Inc. java.vendor.url http://java.sun.com/ java.vendor.url.bug http://java.sun.com/cgi-bin/bugreport.cgi java.version 1.5.0_09 java.vm.info mixed mode java.vm.name Java HotSpot(TM) Server VM java.vm.specification.name Java Virtual Machine Specification java.vm.specification.vendor Sun Microsystems Inc. java.vm.specification.version 1.0 java.vm.vendor Sun Microsystems Inc. java.vm.version 1.5.0_09-b03 For the other mbeans, those are controlled/created by tomcat, so I don't know.	1	1	https://issues.jboss.org/browse/JBAS-5743
1009	JBAS-5733	Tim, can you provide the exact link you are talking about? Is it? http://www.jboss.org/file-access/default/members/jbossas/freezone/docs/Clustering_Guide/beta500/html_single/index.html#clustering-jms	0	0	https://issues.jboss.org/browse/JBAS-5733
1010	JBAS-5733	Does this still need updating? JBoss MQ was superseded by JBoss Messaging. Is HornetQ going to / is it / has it superseded JBoss Messaging in the same way?	0	0	https://issues.jboss.org/browse/JBAS-5733
1011	JBAS-5729	Ron, can you please create a 2.2.2.SP8 and put it in the JBoss repository? I can only see a 2.2.2.SP8-brew version there, and I want to move away from -brew releases for the community project. https://svn.jboss.org/repos/repository.jboss.org/jboss/remoting/	1	1	https://issues.jboss.org/browse/JBAS-5729
1012	JBAS-5727	Is this upgrade fully backwards compatible?	1	1	https://issues.jboss.org/browse/JBAS-5727
1013	JBAS-5726	Yes, Rajesh, that's fine thanks, it uses the right version of jbossws. Might the error be somehow related to file permissions, etc? I'm reassigning this back, should I do something let me know.	0	0	https://issues.jboss.org/browse/JBAS-5726
1014	JBAS-5726	OK, only the jaxr tests are failing now. Should we exclude them since default has no juddi, or what?	0	0	https://issues.jboss.org/browse/JBAS-5726
1015	JBAS-5724	Can we get reliable TCK14 results?	1	1	https://issues.jboss.org/browse/JBAS-5724
1016	JBAS-5703	Galder, what's the status on this one?	0	0	https://issues.jboss.org/browse/JBAS-5703
1017	JBAS-5690	As of CR1, there are two jca tests still failing: org.jboss.test.jcaprops.test.JBossRaXmlOverrideUnitTestCase.testJBossRaXmlOverride AssertionFailedError: Did not expect: [StringRAR=StringFromRARProperties, BooleanRAR=true] expected [StringRAR=XMLOVERRIDE, BooleanRAR=false, ByteRAR=1, CharacterRAR=A, ShortRAR=2, IntegerRAR=3, LongRAR=4, FloatRAR=5000000.0, DoubleRAR=7.0E8] org.jboss.test.jca.test.StatisticsFormatterUnitTestCase.testXmlFormatterStatistics ComparisonFailure: expected:<.....> but was:<... standalone="no"...> See linked issue showing resolution when the changes are committed.	1	1	https://issues.jboss.org/browse/JBAS-5690

#	Requirement	Questions	Answered	Successfully Answered	URL
1018	JBAS-5673	I'm reopening this because the fix can cause a NPE for JSF managed beans. I changed the JSF integration test in the testsuite. If you run org.jboss.test.web.test.JSFIntegrationTestCase you will see the NPE. I used a JButton as a managed bean in faces-config.xml like this: <managed-bean> <description>Use a swing JButton as a simple javabean</description> <managed-bean-name>myJButton</managed-bean-name> <managed-bean-class>javax.swing.JButton</managed-bean-class> <managed-bean-scope>session</managed-bean-scope> <managed-property> <property-name>text</property-name> <value>Initial Value</value> </managed-property> </managed-bean> Note that this is an odd example. I'm not actually using JButton for anything other than the fact that I can use it as a holder to getText()/setText(). But it is perfectly valid to use a JDK class for this purpose. The NPE is in TomcatInjectionContainer.isDymanicBean() where it says: // Check if the object was loaded by a dynamic class loader (e.g. Jasper) String classLoaderName = object.getClass().getClassLoader().getClass().getName(); if(dynamicClassLoaders.contains(classLoaderName)) return true; Since JButton is loaded by the bootstrap classloader, object.getClass().getClassLoader() will return null and you get a NPE. Furthermore, as the code currently stands, it looks like any JSF managed bean that was not loaded by Jasper's classloader will not have annotations processed. Most JSF applications don't use JSP at all, so it looks to me like this code will often fail to process annotations for managed beans. Can you explain exactly what this code is trying to do?	1	1	https://issues.jboss.org/browse/JBAS-5673
1019	JBAS-5672	Do you want this to remain open Scott? Or is it done?	1	1	https://issues.jboss.org/browse/JBAS-5672
1020	JBAS-5672	With the current work that has been done (deployers are now used to parse all metadata, which the web container uses rather than do its own parsing), is there anything more that should be done for AS 6 or can this be considered complete for this release ?	1	0	https://issues.jboss.org/browse/JBAS-5672
1021	JBAS-5640	Can you elaborate on what you mean?	1	1	https://issues.jboss.org/browse/JBAS-5640
1022	JBAS-5636	Yes it looks JBAS-4735. Remy, you said "This cannot be addressed properly in the java.io connectors" in JBAS-4735, I don't get how this comes. Could you advice on it? I did quick hack in JBossWeb to call pause() method on connectors when Server.START_NOTIFICATION_TYPE received. I didn't test it carefully but it works.	1	1	https://issues.jboss.org/browse/JBAS-5636
1023	JBAS-5626	Look inside today CookieUnitTestCase: The failed reason is the cookie include "commented" attribute, (see the comment in CookieUnitTestCase "//RFC2109:Note that there is no Comment attribute"). The strange is I tested both on tomcat 6 and tomcat 5.5, deploy jbosstest-cookie.war, the "commented" attributes are all set .But the test can be passed on Jboss Beta4 and RC1 EncAnnotationsUnitTestCase: The test failed after Aug.14. The exception track: 14:34:22,853 ERROR [[ENCServlet]] Allocate exception for servlet ENCServlet javax.naming.NameNotFoundException: jboss not bound at org.jnp.server.NamingServer.getBinding(NamingServer.java:564) at org.jnp.server.NamingServer.getBinding(NamingServer.java:572) at org.jnp.server.NamingServer.getObject(NamingServer.java:578) at org.jnp.server.NamingServer.lookup(NamingServer.java:317) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:669) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:629) at javax.naming.InitialContext.lookup(InitialContext.java:351) at org.jnp.interfaces.NamingContext.resolveLink(NamingContext.java:1251) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:760) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:629) at org.jboss.ejb3.JndiUtil.lookup(JndiUtil.java:44) at org.jboss.injection.JndiPropertyInjector.lookup(JndiPropertyInjector.java:75) at org.jboss.injection.JndiPropertyInjector.inject(JndiPropertyInjector.java:99) at org.jboss.web.tomcat.service.TomcatInjectionContainer.processInjectors(TomcatInjectionContainer.java:354) at org.jboss.web.tomcat.service.TomcatInjectionContainer.newInstance(TomcatInjectionContainer.java:239) at org.apache.catalina.core.StandardWrapper.loadServlet(StandardWrapper.java:1006) at org.apache.catalina.core.StandardWrapper.allocate(StandardWrapper.java:777) at org.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:129) at org.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:191) at org.jboss.web.tomcat.security.SecurityAssociationValve.invoke(SecurityAssociationValve.java:189) at org.jboss.web.tomcat.security.JaccContextValve.invoke(JaccContextValve.java:90) at org.jboss.web.tomcat.security.SecurityContextEstablishmentValve.invoke(SecurityContextEstablishmentValve.java:96) at org.apache.catalina.core.StandardHostValve.invoke(StandardHostValve.java:127) at org.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:102) at org.jboss.web.tomcat.service.jca.CachedConnectionValve.invoke(CachedConnectionValve.java:157) at org.apache.catalina.core.StandardEngineValve.invoke(StandardEngineValve.java:109) at org.apache.catalina.connector.CoyoteAdapter.service(CoyoteAdapter.java:325) at org.apache.coyote.http11.Http11Processor.process(Http11Processor.java:828) at org.apache.coyote.http11.Http11Protocol$Http11ConnectionHandler.process(Http11Protocol.java:601) at org.apache.tomcat.util.net.JIoEndpoint$Worker.run(JIoEndpoint.java:447) at java.lang.Thread.run(Thread.java:595) It seems some code changed (in org.jboss.web.tomcat.service) caused failing. I will continue looking them.	0	0	https://issues.jboss.org/browse/JBAS-5626
1024	JBAS-5625	What should we do with this JIRA? I can see the tests passing on Hudson, but they fail for me locally in WinXP. Should I close or reschedule?	1	1	https://issues.jboss.org/browse/JBAS-5625
1025	JBAS-5625	The question is do we want to bring in another VFS release in CR1 just for this? I think not.	1	1	https://issues.jboss.org/browse/JBAS-5625

#	Requirement	Questions	Answered	Successfully Answered	URL
1026	JBAS-5625	Ok, looking at the 867 server log I see there is a race between the HDScanner seeing the testEjb3xDeployment.jar and the test removing the jar. Its happening between stop of the deployment, and its removal. During this period it should not be exposed to the HDScanner, so there is a bug in the profile service repository/profile.getModifiedDeployments(). 2008-09-01 01:40:55,515 DEBUG [org.jboss.profileservice.management.upload.remoting.DeployHandler] (WorkerThread#0[10.18.34.175:49212]) invoke, payload: {DeploymentTargetID=names=[testEjb3xDeployment.jar]phase=APPLICATIONcopyContent=truedescription=null, clientAddress=/10.18.34.175}, parameter: stop2008-09-01 01:40:55,515 INFO [org.jboss.profileservice.management.upload.remoting.DeployHandler] (WorkerThread#0[10.18.34.175:49212]) Stop, [testEjb3xDeployment.jar] ... 2008-09-01 01:40:55,667 DEBUG [org.jboss.deployers.plugins.deployers.DeployersImpl] (WorkerThread#0[10.18.34.175:49212]) Fully Undeployed vfszip:/qa/services/hudson/hudson_workspace/workspace/JBoss-AS-5.0.x-TestSuite-sun15/trunk/build/output/jboss-5.0.0.CR2/server/profileservice/deploy/testEjb3xDeployment.jar ... 2008-09-01 01:40:55,701 DEBUG [org.jboss.system.server.profileservice.repository.SerializableDeploymentRepository] (HDScanner) Created deployment: AbstractVFSDeployment(testEjb3xDeployment.jar) ... 2008-09-01 01:40:55,716 DEBUG [org.jboss.profileservice.management.upload.remoting.DeployHandler] (WorkerThread#0[10.18.34.175:49214]) invoke, payload: {DeploymentTargetID=names=[testEjb3xDeployment.jar]phase=APPLICATIONcopyContent=truedescription=null, clientAddress=/10.18.34.175}, parameter: undeploy 2008-09-01 01:40:55,716 INFO [org.jboss.profileservice.management.upload.remoting.DeployHandler] (WorkerThread#0[10.18.34.175:49214]) Undeploy, [testEjb3xDeployment.jar] 2008-09-01 01:40:55,718 INFO [org.jboss.profileservice.management.upload.remoting.DeployHandler] (WorkerThread#0[10.18.34.175:49214]) Undeployed: testEjb3xDeployment.jar ... 2008-09-01 01:40:55,728 DEBUG [org.jboss.deployment.OptAnnotationMetaDataDeployer] (HDScanner) Error during deploy: vfszip:/qa/services/hudson/hudson_workspace/workspace/JBoss-AS-5.0.x-TestSuite-sun15/trunk/build/output/jboss-5.0.0.CR2/server/profileservice/deploy/testEjb3xDeployment.jar org.jboss.deployers.spi.DeploymentException: Cannot process metadata at org.jboss.deployers.spi.DeploymentException.rethrowAsDeploymentException(DeploymentException.java:49) at org.jboss.deployment.AnnotationMetaDataDeployer.deploy(AnnotationMetaDataDeployer.java:181) at org.jboss.deployment.AnnotationMetaDataDeployer.deploy(AnnotationMetaDataDeployer.java:93) at org.jboss.deployers.plugins.deployers.DeployerWrapper.deploy(DeployerWrapper.java:169) ... Caused by: java.util.zip.ZipException: error in opening zip file at java.util.zip.ZipFile.open(Native Method) at java.util.zip.ZipFile.<init>(ZipFile.java:203) at java.util.zip.ZipFile.<init>(ZipFile.java:234) at org.jboss.virtual.plugins.context.zip.ZipFileWrapper.ensureZipFile(ZipFileWrapper.java:173) at org.jboss.virtual.plugins.context.zip.ZipFileWrapper.openStream(ZipFileWrapper.java:207) at org.jboss.virtual.plugins.context.zip.ZipEntryContext.openStream(ZipEntryContext.java:906) ...	1	1	https://issues.jboss.org/browse/JBAS-5625
1027	JBAS-5624	I'm re-opening this issue to discuss it. There are still some cluster-related tests failing every now and then. Lately the JBoss-AS-5.0.x-TestSuite-sun15 run shows 2 tests-jbossmessaging-cluster failures one of which has failed 4 time in a row (so maybe that one is not really transient?) http://hudson.qa.jboss.com/hudson/view/JBoss%20AS/job/JBoss-AS-5.0.x-TestSuite-sun15/888/testReport/ Brian, how you want to deal with those? As things are now, the clustering tests are the last ones that worry me (we have another 3-4 persisting failures to deal with). I know you have created a couple of other JIRAs to improve the predictability of the testsuite runs. Do you want to close this issue and deal with transient failures later as time permits? Do you want to have a look at: http://hudson.qa.jboss.com/hudson/view/JBoss%20AS/job/JBoss-AS-5.0.x-TestSuite-sun15/888/testReport/org.jboss.test.jbossmessaging.clustertest/ClusteredTestCase(tests-jbossmessaging-cluster)/testSimpleCluster/ ?	0	0	https://issues.jboss.org/browse/JBAS-5624
1028	JBAS-5594	The deployment is working now and 7 out of 8 tests are passing. But there is still a CORBA failure: Testcase: testTeller took 0.846 sec Caused an ERROR CORBA BAD_INV_ORDER 4 No; nested exception is: org.omg.CORBA.BAD_INV_ORDER: vmcid: 0x0 minor code: 4 completed: No java.rmi.RemoteException: CORBA BAD_INV_ORDER 4 No; nested exception is: org.omg.CORBA.BAD_INV_ORDER: vmcid: 0x0 minor code: 4 completed: No at com.sun.corba.se.impl.javax.rmi.CORBA.Util.mapSystemException(Util.java:298) at javax.rmi.CORBA.Util.mapSystemException(Util.java:83) at org.jboss.proxy.ejb.DynamicIIOPStub.invoke(DynamicIIOPStub.java:149) at org.jboss.test.bankiiop.interfaces._Teller_Stub.getCustomer(Unknown Source) at org.jboss.test.bankiiop.test.BankStressTestCase.testTeller(BankStressTestCase.java:80) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) Caused by: org.omg.CORBA.BAD_INV_ORDER: vmcid: 0x0 minor code: 4 completed: No at org.jacorb.orb.ORB.work_pending(Unknown Source) at org.jacorb.orb.ORB.perform_work(Unknown Source) at org.jacorb.orb.Delegate.request(Unknown Source) at org.omg.CORBA.portable.ObjectImpl._request(ObjectImpl.java:449) at org.jboss.proxy.ejb.DynamicIIOPStub.invoke(DynamicIIOPStub.java:120) ... 23 more	1	1	https://issues.jboss.org/browse/JBAS-5594
1029	JBAS-5591	I looked at these tests a while ago, they will pass with one-test but not in the entire suite. Did you run only these tests on Windows or the suite?	0	0	https://issues.jboss.org/browse/JBAS-5591
1030	JBAS-5583	Done?	1	1	https://issues.jboss.org/browse/JBAS-5583
1031	JBAS-5580	Does that mean I win a beer?	1	0	https://issues.jboss.org/browse/JBAS-5580
1032	JBAS-5576	Fernando, are you looking after the current console changes?	0	0	https://issues.jboss.org/browse/JBAS-5576
1033	JBAS-5560	The Seam team already has work underway to test Seam in a number of environments before doing a release. However, we need something on the AS side that runs from the AS testsuite so we can always verify that the latest GA version of Seam runs on the latest nightly build of AS. My current thinking is that we would take a few of the Seam sample applications and test them against AS using JSFUnit. So we would simply pull down the release versions of the Seam examples and run them. Currently, only the Booking, DVD Store, and Numberguess examples are being published to the Maven repo. Will these three examples be sufficient for a good test? The only other issue I see with this is that we should go through these examples and make sure that the markup has the ID attribute set for all the JSF components. That will make them testable with JSFUnit.	0	0	https://issues.jboss.org/browse/JBAS-5560
1034	JBAS-5560	"Will these three examples be sufficient for a good test?" Booking tests EJB3, JSF, Facelets, Seam Security DVDStore adds in JBPM Others to consider - seamspace, mail, itext We can easily publish others to maven if you need it, create an issue in JBSEAM	0	0	https://issues.jboss.org/browse/JBAS-5560

#	Requirement	Questions	Answered	Successfully Answered	URL
1035	JBAS-5560	I would add in the JPA example as it will also cover WAR deployment, POJO components, Hibernate JPA, and seam managed transactions. Since it is based on the booking example it should be easy to adjust if needed. There are so many examples based on the booking application that we should try to make all of them testable (in the long run) using the same tests. If we make the IDs match across them it should not be difficult. WDYT? There should probably also be a set of manual tests as part of the AS release process. This could be based on the $SEAM/release-process.txt, although it could be simplified for this purpose.	0	0	https://issues.jboss.org/browse/JBAS-5560
1036	JBAS-5560	> There are so many examples based on the booking application that we should try to make all of them testable (in the long run) using the same tests. If we make the IDs match across them it should not be difficult. WDYT? I think that's a good goal. Creating the JSFUnit test classes that run against different versions of the Booking example should be easy. But keeping the xhtml files in sync will be difficult because there are copies in each different example app. With Maven, you can make a WAR depend on a WAR. This feature is very useful for this sort of thing because you just have to maintain the differences between the main WAR and the derived WAR. So we would only need to keep most of the xhtml files in the main booking example. But for now, I guess we will have to keep them in sync manually.	1	1	https://issues.jboss.org/browse/JBAS-5560
1037	JBAS-5558	Is there more work to it, or can we close this?	0	0	https://issues.jboss.org/browse/JBAS-5558
1038	JBAS-5558	Can we create subtasks to track individual changes needed to the metadata creators? Otherwise this task is too generic to be useful.	0	0	https://issues.jboss.org/browse/JBAS-5558
1039	JBAS-5541	Is there something to do here? The latest jboss-integration artifacts used in AS contain the new jboss-transaction-spi classes (XAResourceRecovery, XAResourceRecoveryRegistry).	0	0	https://issues.jboss.org/browse/JBAS-5541
1040	JBAS-5529	The WebServer, connector and thread pool have been added. I don't see an mbean that provides the vhosts that are available. The vhosts show up as mbeans in a pattern of: host=xxx,type=Host and this has an alias property. Is this what your looking for, all Hosts?	0	0	https://issues.jboss.org/browse/JBAS-5529
1041	JBAS-5529	> The vhosts show up as mbeans in a pattern of:host=xxx,type=Host >and this has an alias property. Is this what your looking for, all Hosts? Yep, exactly.	0	0	https://issues.jboss.org/browse/JBAS-5529
1042	JBAS-5527	There is no plan to release JBoss-4.0.5.SP1. The fix will be in JBoss-5.0.0.CR1 and JBoss-4.2.3.GA. Thanks for the patch!	1	1	https://issues.jboss.org/browse/JBAS-5527
1043	JBAS-5503	Jeff, is this still the case? If so I suppose you can increase the memory in 64m increments until the problem is solved.	1	1	https://issues.jboss.org/browse/JBAS-5503
1044	JBAS-5482	So can we implement the workaround mentioned in the sun bug? Try something like this as a workaround: Context ic = new InitialContext(); ccc = NamingManager.getInitialContext(ResourceManager.getInitialEnvironment(null)); where ccc is a static field Currently it is causing several failures in both AS 5 and EAP 4.2/4.3.	1	0	https://issues.jboss.org/browse/JBAS-5482
1045	JBAS-5467	I don't think it's the SARDeployer to decide to file order. Have you looked at the DeploymentSorter used by the deployment scanner?	0	0	https://issues.jboss.org/browse/JBAS-5467
1046	JBAS-5467	This issue caused us a problem again. We have a deployer script which instead of replacing files, renames them ".backup...". JBoss had been loading both the new file and the .backup file in the lib directory. I.e., lib/hibernate.jar; lib/.backup-hibernate.jar. This problem remained hidden until the JBoss server was copied to another directory. After being moved, it started loading the older files first. Should we have multiple copies of jars in our lib directory? No, obviously this is a problem on our part. Should moving the directory cause a working deployment to stop working? In my opinion, no, the app server should function in a consistant manner regardless of hardware or the physical ordering of the files. Unfortunately, we don't build our JBoss servers, but instead insist on running a version for which a support contact is available. As such, we can't just patch our own copy and ignore the issue. I might be wrong about the class, but I'm pretty sure the micro-kernel uses it when it load the jars in the server's lib directory. Files in the deploy directory do get loaded in a defined order, but not those in the lib directory. Here's the specific method, for which the patch applies. protected File[] listFiles(final String urlspec) throws Exception { URL url = Strings.toURL(urlspec); // url is already canonical thanks to Strings.toURL File dir = new File(url.getFile()); File[] files = dir.listFiles(new java.io.FileFilter() { public boolean accept(File pathname) { String name = pathname.getName().toLowerCase(); return (name.endsWith(".jar") || name.endsWith(".zip")); } }); return files; }	1	0	https://issues.jboss.org/browse/JBAS-5467
1047	JBAS-5422	Ran tests on AS 4.0.4.GA, JBC 1.4.0.SP1, JGroups 2.2.9.2, testing UDP, TCP and TCP_NIO with bundling on and off, REPL_ASYNC and REPL_SYNC. 4 node cluster, no buddy replication. Results are attached. The basic JBC config files are also attached. The versions used for no-bundling tests looked like the attached, except "enabled_bundling" was "false". The versions used for REPL_SYNC tests looked like the attached except for CacheMode being REPL_SYNC and the FC protocol was commented out of the protocol stack. Test app basically puts an 8Kb payload in a single attribute for each request. Both SESSION and ATTRIBUTE granularity were tested. Since the session only has a single attribute, the payload size difference between the two granularities is not great. A big difference is SESSION does not use a TX during replication (it just does a single put, so no need), while ATTRIBUTE does use a TX. Interesting points: 1) TCP_NIO is by far the fastest transport, consistent with what I've seen previously. 2) Disabling bundling actually *hurts* TCP_NIO performance. 3) For UDP and TCP, disabling bundling has only a minor effect (< 1.5%), in most cases negative. This is actually good news, since with the multiplexer a REPL_ASYNC cache may share a channel with a synchrononous HAPartition. It is important that an optimization for one does not harm the other. 4) With SESSION + SYNC, the benefit of turning of bundling is largest. This is interesting, as this combination uses no TX, so the costly 2PC protocol doesn't get involved. 5) Turning off bundling actually hurts a bit in the ATTR + SYNC case. This is odd, as this case uses the full 2PC protocol.	0	0	https://issues.jboss.org/browse/JBAS-5422
1048	JBAS-5408	Who said we need to fix this now?	0	0	https://issues.jboss.org/browse/JBAS-5408
1049	JBAS-5388	Alexey, can you take a look? This testcase must be in the WS testsuite.	1	1	https://issues.jboss.org/browse/JBAS-5388
1050	JBAS-5387	Could you please attach the deployment unit the test is deploying?	1	0	https://issues.jboss.org/browse/JBAS-5387

#	Requirement	Questions	Answered	Successfully Answered	URL
1051	JBAS-5373	console.war 09:26:21,687 ERROR [AbstractKernelController] Error installing to Start: name=jboss.web.deployment:war=/embedded-console state=Create mode=Manual requiredState=Installed org.jboss.deployers.spi.DeploymentException: URL file:/C:/JBossSVN/jboss-head/build/output/jboss-5.0.0.CR1/server/defaul t/tmp/deploy/embedded-console28064-exp.war/ deployment failed at org.jboss.deployers.spi.DeploymentException.rethrowAsDeploymentException(DeploymentException.java:49) at org.jboss.web.tomcat.service.deployers.TomcatDeployment.performDeployInternal(TomcatDeployment.java:350) at org.jboss.web.tomcat.service.deployers.TomcatDeployment.performDeploy(TomcatDeployment.java:139) at org.jboss.web.deployers.AbstractWarDeployment.start(AbstractWarDeployment.java:431) at org.jboss.web.deployers.WebModule.startModule(WebModule.java:112) at org.jboss.web.deployers.WebModule.start(WebModule.java:90) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.system.microcontainer.ServiceProxy.invoke(ServiceProxy.java:206) at $Proxy35.start(Unknown Source) at org.jboss.system.microcontainer.StartStopLifecycleAction.installAction(StartStopLifecycleAction.java:42) at org.jboss.system.microcontainer.StartStopLifecycleAction.installAction(StartStopLifecycleAction.java:37) at org.jboss.dependency.plugins.action.SimpleControllerContextAction.simpleInstallAction(SimpleControllerContext Action.java:62) at org.jboss.dependency.plugins.action.AccessControllerContextAction.install(AccessControllerContextAction.java: 71) at org.jboss.dependency.plugins.AbstractControllerContextActions.install(AbstractControllerContextActions.java:5 1) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:348) at org.jboss.system.microcontainer.ServiceControllerContext.install(ServiceControllerContext.java:271) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:1394) at org.jboss.dependency.plugins.AbstractController.incrementState(AbstractController.java:786) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:914) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:836) at org.jboss.dependency.plugins.AbstractController.change(AbstractController.java:674) at org.jboss.dependency.plugins.AbstractController.change(AbstractController.java:456) at org.jboss.system.ServiceController.doChange(ServiceController.java:664) at org.jboss.system.ServiceController.start(ServiceController.java:436) at org.jboss.system.deployers.ServiceDeployer.start(ServiceDeployer.java:150) at	1	1	https://issues.jboss.org/browse/JBAS-5373
1052	JBAS-5370	Scott, are all the features described at the beginning of that thread currently available? 1) List deployed apps (ears/wars), 2) Deploy apps (ears/wars), 3) Undeploy apps(ears/wars), 4) Get the availability status of the deployed app Where do we stand with the win32 issues? I'm going to close JBAS-5522 so we just have one issue to track. We can spin off any remaining issues as subtasks of JBAS-5246	0	0	https://issues.jboss.org/browse/JBAS-5370
1053	JBAS-5370	When attempting to deploy a war using the following code: DeploymentProgress progress = manager.distribute(tempFile.getName(), ManagedDeployment.DeploymentPhase.APPLICATION, tempFile.toURL()); progress.addProgressListener(this); progress.run(); I get the following NPE: 14:16:29,155 ERROR [ProfileJBossServerComponent] java.lang.RuntimeException: java.lang.NullPointerException at org.jboss.profileservice.management.upload.remoting.StreamingDeploymentTarget.transferDeployment(StreamingDep loymentTarget.java:268) at org.jboss.profileservice.management.upload.remoting.StreamingDeploymentTarget.distribute(StreamingDeploymentT arget.java:108) at org.jboss.profileservice.management.upload.DeploymentProgressImpl.distribute(DeploymentProgressImpl.java:169) at org.jboss.profileservice.management.upload.DeploymentProgressImpl.run(DeploymentProgressImpl.java:82) at org.rhq.plugins.jbossas5.ProfileJBossServerComponent.deployArchive(ProfileJBossServerComponent.java:318) at org.rhq.plugins.jbossas5.ProfileJBossServerComponent.createContentBasedResource(ProfileJBossServerComponent.j ava:283) at org.rhq.plugins.jbossas5.ProfileJBossServerComponent.createResource(ProfileJBossServerComponent.java:138) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:623) at org.rhq.core.pc.inventory.ResourceContainer$ComponentInvocationThread.call(ResourceContainer.java:450) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:334) at java.util.concurrent.FutureTask.run(FutureTask.java:166) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:603) at java.lang.Thread.run(Thread.java:713) Caused by: java.lang.NullPointerException at org.jboss.profileservice.management.upload.remoting.DeployHandler.handleStream(DeployHandler.java:106) at org.jboss.remoting.ServerInvoker.handleInternalInvocation(ServerInvoker.java:1785) at org.jboss.remoting.ServerInvoker.invoke(ServerInvoker.java:895) at org.jboss.remoting.transport.local.LocalClientInvoker.invoke(LocalClientInvoker.java:106) at org.jboss.remoting.Client.invoke(Client.java:1708) at org.jboss.remoting.Client.invoke(Client.java:612) at org.jboss.remoting.Client.invoke(Client.java:1492) at org.jboss.profileservice.management.upload.remoting.StreamingDeploymentTarget.transferDeployment(StreamingDep loymentTarget.java:259) ... 16 more 14:16:29,155 INFO [DeployHandler] Begin start, [valid.war], phase: APPLICATION I'm using an AS5 build from late last week.	1	1	https://issues.jboss.org/browse/JBAS-5370
1054	JBAS-5370	The org.jboss.test.profileservice.test.DeployUnitTestCase are still passing. The line of the NPE in DeployHandler indicates that the DeploymentManager bean in profileservice/deploy/profileservice-jboss-beans.xml does not have a repositoryFactory injected. How is the profileservice config setup?	1	1	https://issues.jboss.org/browse/JBAS-5370

#	Requirement	Questions	Answered	Successfully Answered	URL
1055	JBAS-5370	Thanks, adding the call to DeploymentManager.loadProfile() did the trick. Could you possibly add some more javadoc on that method? If the DeploymentManager is a singleton, then how does one prevent threads from interfering with each other by loading different profiles via loadProfile()?	0	0	https://issues.jboss.org/browse/JBAS-5370
1056	JBAS-5318	Stefan, has you checked if the cts config needs updating?	1	0	https://issues.jboss.org/browse/JBAS-5318
1057	JBAS-5307	We rolledbacked the maven transition for now, so can you put xerces in the legacy repo and switch to it?	1	1	https://issues.jboss.org/browse/JBAS-5307
1058	JBAS-5298	I did not maintain an individual changelog for those CRs (I did not think I would have to produce that much; in the future I'll do a separate changelog entry in future releases). As for AS5 CR1, I was planning to tag a GA this week. It's not a problem as most things are aligned with Tomcat 6, which is stable, Beyond CR1, what is the strategy for including bugfixes in AS 5 ? (given the advancements of the TCK, it looks like GA is not 4 weeks away) I should not have many hugely important fixes to put in, obviously, but most likely some "nice to have" ones. Would starting to work on a JBoss Web 2.1.1 GA (and include it if needed) be the right way ?	1	1	https://issues.jboss.org/browse/JBAS-5298
1059	JBAS-5289	Currently using 0.1.5 (Beta, CR, GA? who knows :-)	0	0	https://issues.jboss.org/browse/JBAS-5289
1060	JBAS-5288	I released the 5.0.0.CR1 version, and updated the app server build. Dimitris, can this issue be closed?	0	0	https://issues.jboss.org/browse/JBAS-5288
1061	JBAS-5286	Who'll be taking over this?	0	0	https://issues.jboss.org/browse/JBAS-5286
1062	JBAS-5269	Why are you raising bug reports as subtasks of closed issues in old releases? 1) Check whether it is fixed in the latest version before raising bug reports 2) Raise a top level bug report against the latest release, not as a subtask of a closed issue. 3) Show relevant information (config, logs, etc.) 4) While an NPE is always a bug (it should give a proper error message telling you what is wrong) it usually means missing configuration. Use the forums to ask for help.	0	0	https://issues.jboss.org/browse/JBAS-5269
1063	JBAS-5265	Is it safe now to assume the upgrade is over?	1	1	https://issues.jboss.org/browse/JBAS-5265
1064	JBAS-5252	Has someone looked more into this? AFAICT, the leaks are coming and going from run to run. E.g. from a relatively recent one: org.jboss.test.classloader.leak.test.Ejb3ClassloaderLeakUnitTestCase.testSimpleEjb org.jboss.test.cluster.classloader.leak.test.FieldGranularityClassloaderLeakUnitTestCase(ClusteredClassLoaderLeak).testSimpleReplicableWar org.jboss.test.cluster.classloader.leak.test.ReplicableClassesClassloaderLeakUnitTestCase(ClusteredClassLoaderLeak).testNoCacheNoReplicableWar org.jboss.test.cluster.classloader.leak.test.ReplicableClassesClassloaderLeakUnitTestCase(ClusteredClassLoaderLeak).testNoCacheReplicableWar org.jboss.test.cluster.classloader.leak.test.ReplicableClassesClassloaderLeakUnitTestCase(ClusteredClassLoaderLeak).unknown	1	1	https://issues.jboss.org/browse/JBAS-5252
1065	JBAS-5249	Why bsf and bsh need to be upgraded together? Does bsf 2.4.0 work with the older bsh 1.3.0?	1	1	https://issues.jboss.org/browse/JBAS-5249
1066	JBAS-5240	I'm not sure I understand what you mean Dimitris. If you see the wiki I pointed out in the description, you'll I've added the following: "By default, command line parameters passed to JBoss Application Server on startup are only logged in the console output. Within the scope of remote support, this leads to users/customers being asked to provide these options explicitly, as console output is often lost. This can be avoided by adding the following line to java execution in run.sh/bat: -Djboss.java.opts="$JAVA_OPTS" \ to make the script look something like this: ... if ["x$LAUNCH_JBOSS_IN_BACKGROUND" = "x"]; then # Execute the JVM in the foreground "$JAVA" $JAVA_OPTS \ -Djava.endorsed.dirs="$JBOSS_ENDORSED_DIRS" \ -Djboss.java.opts="$JAVA_OPTS" \ -classpath "$JBOSS_CLASSPATH" \ org.jboss.Main "$@" JBOSS_STATUS=$? else # Execute the JVM in the background "$JAVA" $JAVA_OPTS \ -Djava.endorsed.dirs="$JBOSS_ENDORSED_DIRS" \ -Djboss.java.opts="$JAVA_OPTS" \ -classpath "$JBOSS_CLASSPATH" \ org.jboss.Main "$@" & JBOSS_PID=$! ... On startup, ServerInfo class prints all system properties set on startup to boot.log and so, by adding this new system property, we're effectively forcing boot.log to print command line parameters passed on startup: 00:18:50,045 DEBUG [ServerInfo] jboss.java.opts: -Dprogram.name=run.sh -server -Xms128m -Xmx512m -Dsun.rmi.dgc.client.gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000 -Xdebug -Xrunjdwp:transport=dt_socket,address=8787,server=y,suspend=y" This solution worked fine for me. Am I missing something?	1	1	https://issues.jboss.org/browse/JBAS-5240
1067	JBAS-5240	I've commited the changes in both run.sh and run.bat. Can you check it works for you and also update the wiki?	1	1	https://issues.jboss.org/browse/JBAS-5240
1068	JBAS-5240	In general both posix and windows scripts are broken if JAVA_OPTS contains quoted param. > export JAVA_OPTS=-Djboss.foo="some param" > ./run.sh Exception in thread "main" java.lang.NoClassDefFoundError: value Caused by: java.lang.ClassNotFoundException: value at java.net.URLClassLoader$1.run(URLClassLoader.java:200) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:188) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:276) at java.lang.ClassLoader.loadClass(ClassLoader.java:251) at java.lang.ClassLoader.loadClassInternal(ClassLoader.java:319) This is caused by -Djboss.java.opts="$(JAVA_OPTS)" so if JAVA_OPTS contains quoted string one has invalid command line cause its evaluated to: -Djboss.java.opts="-server -Djboss.foo="some value"" So you have value" as main class, and that sucks :) The solution would be to escape the $JAVA_OPTS for usage in -Djboss.java.opts.	1	1	https://issues.jboss.org/browse/JBAS-5240
1069	JBAS-5235	Pushing to GA. Jonathan will you do that after CR2?	0	0	https://issues.jboss.org/browse/JBAS-5235
1070	JBAS-5235	Can Jesper or Jeff add to example datasource descriptors an explanatory comment about usage of local datasources in transactions?	0	0	https://issues.jboss.org/browse/JBAS-5235
1071	JBAS-5234	The new params to Main (post 5.0 Beta4) contains the param -Djboss.java.opts="%JAVA_OPTS%" Now this is broken because if any JAVA_OPTS has quoted string there will be double quoting issue resulting in the following error when starting JBoss java.lang.NoClassDefFoundError: Files Caused by: java.lang.ClassNotFoundException: Files at java.net.URLClassLoader$1.run(URLClassLoader.java:200) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:188) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:276) at java.lang.ClassLoader.loadClass(ClassLoader.java:251) at java.lang.ClassLoader.loadClassInternal(ClassLoader.java:319) Exception in thread "main" Shutdown service	1	1	https://issues.jboss.org/browse/JBAS-5234
1072	JBAS-5233	When implementing this, make sure that any potential code for JBAS-5647 does not log "expected" non-received responses at DEBUG.	1	0	https://issues.jboss.org/browse/JBAS-5233
1073	JBAS-5224	Stale, I've created JBVFS-26 to optimize frequently called vfs methods. From what you've seen in your analysis, who are the callers of the vfs methods where we can look to reduce the number of calls to vfs? We need to create JIRA to other projects so their parts can be optimized, too. Thanks	0	0	https://issues.jboss.org/browse/JBAS-5224
1074	JBAS-5221	Is that a profile service issue?	0	0	https://issues.jboss.org/browse/JBAS-5221

#	Requirement	Questions	Answered	Successfully Answered	URL
1075	JBAS-5216	Kabir, is this issue still current? If so, can you revert the fix/hack?	1	1	https://issues.jboss.org/browse/JBAS-5216
1076	JBAS-5203	The necessary code changes were applied. The following wiki pages were modified with the notes about the changes in the configuration. Should the documentation team be notified to propagate the info further? Or how is it supposed to get into the docs?	0	0	https://issues.jboss.org/browse/JBAS-5203
1077	JBAS-5163	FAO Dimitris: Head has been based in Java 5 for quite a while now, so as part of this JIRA, I should be able to convert reflection calls to JDK 5 methods into type safe ones, correct? Obviously, I would use reflection for JDK 6 classes/methods. Thoughts?	1	1	https://issues.jboss.org/browse/JBAS-5163
1078	JBAS-5163	Also, I'd like to see what would the extra jdk6 info be in the output (some example).	1	1	https://issues.jboss.org/browse/JBAS-5163
1079	JBAS-5163	Re: Also, I'd like to see what would the extra jdk6 info be in the output (some example). Yeah, definitely. In facf, not just that, I'm creating 2 deadlock examples (one with synchronised monitors and one with java.util.concurrent Lock class impl) so that you can also see how we can print out deadlocks both JDK5 and JDK6. We could do already print synchronized monitor deadlock with JDK5 already, but was not implemented. One thing I'm changing for all thread dumps is the following: Thread: JBossLifeThread : priority:5, demon:false, threadId:58, threadState:WAITING, lockName:java.l ang.Object@10cff6b java.lang.Object.wait(Native Method) java.lang.Object.wait(Object.java:474) Rather than "lockName: java.lang.Object@10cff6b" in the title, it makes more sense to put it in the stack itself, in the same way that any standard thread dump output does, i.e: Thread: JBossLifeThread : priority:5, demon:false, threadId:58, threadState:WAITING - waiting on java.lang.Object@10cff6b java.lang.Object.wait(Native Method) java.lang.Object.wait(Object.java:474) The closer we are to real non JMX-console thread dumps, the easier for any of us to read them.	1	1	https://issues.jboss.org/browse/JBAS-5163
1080	JBAS-5158	Sorry if I forgot. Is there a corresponding Forum post with this patch? If not, would you mind to start one there? I will also need some additional documentation and also unit test cases so I can check them in properly. Thanks, -Ben	1	0	https://issues.jboss.org/browse/JBAS-5158
1081	JBAS-5154	Brian would you like anything to go to the guides in relation to the examples?	1	0	https://issues.jboss.org/browse/JBAS-5154
1082	JBAS-5147	How those tests are supposed to run?	1	1	https://issues.jboss.org/browse/JBAS-5147
1083	JBAS-5147	down server: scoped-aop-jdk50 <testcase classname="org.jboss.test.aop.test.ScopedUnitTestCase" name="testScoped2" time="0.046"> <error type="javax.management.RuntimeMBeanException">javax.management.RuntimeMBeanException at org.jboss.mx.interceptor.ReflectedDispatcher.handleInvocationExceptions(ReflectedDispatcher.java:180) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:165) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.jmx.connector.invoker.InvokerAdaptorService.invoke(InvokerAdaptorService.java:270) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:138) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:140) at org.jboss.jmx.connector.invoker.SerializableInterceptor.invoke(SerializableInterceptor.java:74) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.invocation.jrmp.server.JRMPProxyFactory.invoke(JRMPProxyFactory.java:179) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:815) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:416) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:294) at sun.rmi.transport.Transport$1.run(Transport.java:153) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:149) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:466) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:707) at java.lang.Thread.run(Thread.java:595) Caused by: java.lang.RuntimeException: Expected ScopedAspect 1 for POJO constructor, was 0 at org.jboss.test.aop.scoped.ScopedTester.testScoped(ScopedTester.java:100) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) ... 41 more </error> <testcase classname="org.jboss.test.aop.test.ScopedUnitTestCase" name="testIntroduction2" time="0.031"> <error type="javax.management.RuntimeMBeanException">javax.management.RuntimeMBeanException at org.jboss.mx.interceptor.ReflectedDispatcher.handleInvocationExceptions(ReflectedDispatcher.java:180) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:165) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at	1	1	https://issues.jboss.org/browse/JBAS-5147
1084	JBAS-5147	I still can't reproduce. Did you update your workspace copy? What OS are you on?	1	1	https://issues.jboss.org/browse/JBAS-5147
1085	JBAS-5139	Heiko a) Is EJB3 team aware of the SPI change? b) AS Branch_4_2 and trunk are now on jboss-native snapshots. When are we moving out of them?	1	1	https://issues.jboss.org/browse/JBAS-5139
1086	JBAS-5135	I've released the GA version http://repository.jboss.org/maven2/org/hibernate/hibernate-annotations/3.3.1.GA/ be careful the ejb3-persistence dependency has changed to 1.0.1.GA as well http://repository.jboss.org/maven2/org/hibernate/ejb3-persistence/1.0.1.GA/ (downgrade in the version, but the right one) can you take it from here and update EJB3 and AS deps?	1	1	https://issues.jboss.org/browse/JBAS-5135
1087	JBAS-5135	I've added HA 3.3.1.GA to the legacy repo, until we switch to maven: https://svn.jboss.org/repos/repository.jboss.org/hibernate-annotations/3.3.1.GA/ What is the story with ejb3-persistence and the legacy repo? Where is it specified? I don't see an entry in the central build/build-thirdparty.xml.	0	0	https://issues.jboss.org/browse/JBAS-5135
1088	JBAS-5134	I've released the GA version http://repository.jboss.org/maven2/org/hibernate/hibernate-entitymanager/3.3.2.GA/ be careful the ejb3-persistence dependency has changed to 1.0.1.GA as well http://repository.jboss.org/maven2/org/hibernate/ejb3-persistence/1.0.1.GA/ (downgrade in the version, but the right one) can you take it from here and update EJB3 and AS deps?	1	1	https://issues.jboss.org/browse/JBAS-5134
1089	JBAS-5133	I am running into issues with the maven release plugin. I'd like to get you jars from a reproducible build. When is the drop dead date for me to get you this release? Otherwise, I'll just manually build you the jars.	1	1	https://issues.jboss.org/browse/JBAS-5133

#	Requirement	Questions	Answered	Successfully Answered	URL
1090	JBAS-5133	two questions dimitris: 1) are y'all using maven now for building? I am trying to figure out where to put these jars for you. Is just in the maven repo good enough? 2) I am past the issues with the release plugin per se. I am still having issues with generating an assembly which is what is used to create a dist bundle for upload to sourceforge. so basically there means that currently there is a 100% chance this CR will not simply clock over to be GA. Is that kosher for you for the time being?	0	0	https://issues.jboss.org/browse/JBAS-5133
1091	JBAS-5133	Assigning to Paul to make the necessary changes to pull that new release into AS. Couple of points: 1) Due to splitting Hibernate out into modules, we will need to account for Hibernate features actually utilized within the AS. AFAIK, we should only need the following 2 artifacts: a) org.hibernate:hibernate-core b) org.hibernate:hibernate-jbosscache2 2) a dependency on the hibernate-core artifact is going to transitively pull in a dep on cglib/asm. should those be excluded?	0	0	https://issues.jboss.org/browse/JBAS-5133
1092	JBAS-5133	I updated the hibernate dependency. The hibernate3 jar seems to have been replaced by hibernate-core and hibernate-jmx. With those two jars, the build seems to run ok. I didn't try a full testsuite, but it seems to compile and start ok. Should any of the other core hibernate components (cache, connection, etc) be included in the build?	0	0	https://issues.jboss.org/browse/JBAS-5133
1093	JBAS-5133	TCK complains about a missing org/slf4j/LoggerFactory class which is a dependency on hibernate-jbosscache. Should this also be installed?	0	0	https://issues.jboss.org/browse/JBAS-5133
1094	JBAS-5120	Am using JBOSS5.0 Beta4 and can see this exception during start up. Is this fixed ? This is what I have done (basically I was trying to achive cluster setup with Beta4) 1. create profile node_1 (copy all directory from all). 2. create profile node_2 (copy all directory from all) 3. deployers/profileservice-beans.xml.bak to deployers/profileservice-beans.xml for node_1 , node_2 and all profile. 4. Start node_1 , node_2 and all I see the above exception in startup of all three profile.	0	0	https://issues.jboss.org/browse/JBAS-5120
1095	JBAS-5117	I've flipped trunk to this new Beta5 version. I had to move the release to the legacy position. Was that on purpose? Added: jboss/metadata/1.0.0.Beta5/ Removed: jboss/jboss-metadata/1.0.0.Beta5/ We don't have a consistent naming convention. Some modules are jboss/xxx others jboss/jboss-xxx, yet others jboss/jbossxxx: <componentref name="jboss/aop" version="2.0.0.CR1"/> <componentref name="jboss/cache" version="2.1.0.CR2"/> <componentref name="jboss/common-core" version="2.2.3.GA"/> <componentref name="jboss/common-logging-jdk" version="2.0.2.GA"/> <componentref name="jboss/common-logging-log4j" version="2.0.2.GA"/> <componentref name="jboss/common-logging-spi" version="2.0.2.GA"/> <componentref name="jboss/integration" version="5.0.0.Beta4"/> <componentref name="jboss/jaxr" version="1.2.0.GA"/> <componentref name="jboss/jboss-ejb3-cache" version="0.11.1"/> <componentref name="jboss/jboss-ejb3-ext-api" version="0.1.2.Beta1"/> <componentref name="jboss/jboss-ejb3-impl" version="0.1.2.Beta1"/> <componentref name="jboss/jboss-ha-client" version="1.0.0.BETA1"/> <componentref name="jboss/jboss-ha-server-api" version="1.0.0.BETA2-SNAPSHOT"/> <componentref name="jboss/jboss-jaspi-api" version="1.0-BETA1"/> <componentref name="jboss/jboss-javaee" version="5.0.0.Beta3Update1"/> <componentref name="jboss/jboss-security-spi" version="2.0.2.Beta3"/> <componentref name="jboss/jbosssx" version="2.0.2.Beta3"/> <componentref name="jboss/jbosssx-client" version="2.0.2.Beta3"/> <componentref name="jboss/jbossts" version="4.3.0.BETA2"/> <componentref name="jboss/jboss-vfs" version="2.0.0.Beta6"/> <componentref name="jboss/jbossws-native50" version="2.0.2.GA"/> <componentref name="jboss/jbossxb" version="2.0.0.CR5"/> <componentref name="jboss/jms-integration-tests" version="1.0.1.GA"/> <componentref name="jboss/messaging" version="1.4.0.SP3"/> <componentref name="jboss/metadata" version="1.0.0.Beta5"/> <componentref name="jboss/microcontainer" version="snapshot"/> <componentref name="jboss/profiler/jvmti" version="1.0.0.CR5"/> <componentref name="jboss/remoting" version="2.2.2.SP3-brew"/> <componentref name="jboss/serialization" version="1.0.3.GA"/> <componentref name="jboss/test" version="1.0.4.GA"/> <componentref name="jboss/web" version="2.1.0.CR11"/>	0	0	https://issues.jboss.org/browse/JBAS-5117
1096	JBAS-5078	With the proper unlimited strength crypto files installes, all SRPLoginModuleUnitTestCase tests pass successfully. I can see only one error in the SRPUnitTestCase's testEchoArgs() method: java.lang.IllegalStateException: Security Context is null" type="java.rmi.ServerException">java.rmi.ServerException: RuntimeException; nested exception is: java.lang.IllegalStateException: Security Context is null at org.jboss.ejb.plugins.LogInterceptor.handleException(LogInterceptor.java:421) at org.jboss.ejb.plugins.LogInterceptor.invokeHome(LogInterceptor.java:136) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invokeHome(ProxyFactoryFinderInterceptor.java:107) at org.jboss.ejb.SessionContainer.internalInvokeHome(SessionContainer.java:640) at org.jboss.ejb.Container.invoke(Container.java:1045) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:815) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:416) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:294) at sun.rmi.transport.Transport$1.run(Transport.java:153) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:149) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:466) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:707) at java.lang.Thread.run(Thread.java:595) Caused by: java.lang.IllegalStateException: Security Context is null at org.jboss.ejb.plugins.SecurityInterceptor.process(SecurityInterceptor.java:229) at org.jboss.ejb.plugins.SecurityInterceptor.invokeHome(SecurityInterceptor.java:199) at org.jboss.ejb.plugins.LogInterceptor.invokeHome(LogInterceptor.java:132) I am now investigating why this is happening.	1	1	https://issues.jboss.org/browse/JBAS-5078
1097	JBAS-5073	How this problem manifestred itself, i.e. what was the problem exactly? :-) BTW, nice to see you back in JIRA Andy!	1	1	https://issues.jboss.org/browse/JBAS-5073
1098	JBAS-5073	the <condition> element was not inside of the target. look at the patch. It meant that if you supplied a user/password ... it was ignored. IT is a one line fix. I don't know why ant lets you randomly stick <condition> tags outside of where they are interpreted but apparently it does. well I have a little consulting side business and do subcontracting for RH among others so it often involves JBoss and building JBoss stuff on frustrating networks with frustrating proxies. You can't imagine the joy of building JBoss behind an NTLM authenticated proxy. I even had to patch ntlmaps which is defunct but has a buffering bug that makes it not work on slow networks... BTW is 4.4 an active branch? Should I have applied it there?	1	1	https://issues.jboss.org/browse/JBAS-5073

#	Requirement	Questions	Answered	Successfully Answered	URL
1099	JBAS-5065	Does this mean converting the .rpm's to .deb?	1	1	https://issues.jboss.org/browse/JBAS-5065
1100	JBAS-5065	Is anyone working on this? Also, I cannot find any information about MEAD that Clive is talking about... I am interested in creating a package for Ubuntu, I will work on doing this, and keep people posted, if no one has allready started this.	0	0	https://issues.jboss.org/browse/JBAS-5065
1101	JBAS-5060	Have you tested that it works? Start jboss, point to http://localhost:8080/web-console/, navigate to 'System', 'JMX MBeans', 'jboss.system', 'jboss.system:type=ServerInfo". Right click on the 'FreeMemory' attribute and choose to create a graph. It hangs, with this StackTrace: Exception in thread "AWT-EventQueue-2" java.lang.NullPointerException at org.jfree.chart.labels.StandardXYSeriesLabelGenerator.createItemArray(StandardXYSeriesLabelGenerator.java:119) at org.jfree.chart.labels.StandardXYSeriesLabelGenerator.generateLabel(StandardXYSeriesLabelGenerator.java:102) at org.jfree.chart.renderer.xy.XYLineAndShapeRenderer.getLegendItem(XYLineAndShapeRenderer.java:1100) at org.jfree.chart.plot.XYPlot.getLegendItems(XYPlot.java:3675) at org.jfree.chart.title.LegendTitle.fetchLegendItems(LegendTitle.java:393) at org.jfree.chart.title.LegendTitle.arrange(LegendTitle.java:464) at org.jfree.chart.JFreeChart.drawTitle(JFreeChart.java:1134) at org.jfree.chart.JFreeChart.draw(JFreeChart.java:1025) at org.jfree.chart.ChartPanel.paintComponent(ChartPanel.java:1275) at javax.swing.JComponent.paint(Unknown Source) at javax.swing.JComponent.paintToOffscreen(Unknown Source) at javax.swing.BufferStrategyPaintManager.paint(Unknown Source) at javax.swing.RepaintManager.paint(Unknown Source) at javax.swing.JComponent._paintImmediately(Unknown Source) at javax.swing.JComponent.paintImmediately(Unknown Source) at javax.swing.RepaintManager.paintDirtyRegions(Unknown Source) at javax.swing.RepaintManager.paintDirtyRegions(Unknown Source) at javax.swing.RepaintManager.seqPaintDirtyRegions(Unknown Source) at javax.swing.SystemEventQueueUtilities$ComponentWorkRequest.run(Unknown Source) at java.awt.event.InvocationEvent.dispatch(Unknown Source) at java.awt.EventQueue.dispatchEvent(Unknown Source) at java.awt.EventDispatchThread.pumpOneEventForFilters(Unknown Source) at java.awt.EventDispatchThread.pumpEventsForFilter(Unknown Source) at java.awt.EventDispatchThread.pumpEventsForHierarchy(Unknown Source) at java.awt.EventDispatchThread.pumpEvents(Unknown Source) at java.awt.EventDispatchThread.pumpEvents(Unknown Source) at java.awt.EventDispatchThread.run(Unknown Source)	1	1	https://issues.jboss.org/browse/JBAS-5060
1102	JBAS-5054	The reason there were so many errors existed is explained here: http://www.jboss.com/index.html?module=bb&op=viewtopic&t=125466 EJBMethodPermission(String EJBName, String methodName, String methodInterface, String[] methodParams) methodParams being an empty array implies method with zero args whereas a null methodParams implies all overloaded method signatures. Who would have thought things would be so subtle in the JACC jungle?	1	1	https://issues.jboss.org/browse/JBAS-5054
1103	JBAS-5054	Scott, what this means is that you are duplicating the hudson situation. Can you please provide me the trace log for the entire JACC run? The security error is because the security checks have not been able to derive the policy context from the policy.	1	1	https://issues.jboss.org/browse/JBAS-5054
1104	JBAS-5054	I had a quick look and the test looks badly written, e.g.: org.jboss.test.security.test.EJBSpecUnitTestCase:614 log.debug("Sent msg to queue/QueueA"); queueConn.start(); QueueReceiver recv = session.createReceiver(queB); msg = recv.receive(5000); log.debug("Recv msg: "+msg); String info = msg.getStringProperty("reply"); recv.close(); session.close(); queueConn.close(); If for some reason recv.receive(5000) times-out, then it returns null and you get an NPE so the session is not closed properly as in: http://hudson.qa.jboss.com/hudson/view/JBoss%20AS/job/JBoss-AS-5.0.x-TestSuite-sun15/837/testReport/org.jboss.test.security.test/EJBSpecUnitTestCase(tests-security-basic-unit)/testMDBRunAs/ Can you improve the test so the timeout is increased and proper cleanup is made in all cases?	0	0	https://issues.jboss.org/browse/JBAS-5054
1105	JBAS-5044	Rescheduling to CR1. Is this issue really critical? It hasn't been touched in Months.	0	0	https://issues.jboss.org/browse/JBAS-5044
1106	JBAS-5041	HI , We have the default Jboss Transaction Timeout set to 3 hours. We have a MDB that has CMT , and want to set the Transaction Timeout for this MDB to 6 hours. How can this be done?? We are using JbossAS-4.2.2GA, and EJB2 When will this issue be fixed? Anuja	1	0	https://issues.jboss.org/browse/JBAS-5041
1107	JBAS-5035	It looks the same, did you attach the correct one? :-) I've deleted both the attachments can you try again.	1	1	https://issues.jboss.org/browse/JBAS-5035

#	Requirement	Questions	Answered	Successfully Answered	URL
1108	JBAS-5031	name=VFSClassLoaderDeployer state=PreInstall java.lang.RuntimeException: java.lang.reflect.UndeclaredThrowableException at org.jboss.aop.util.ClassInfoMethodHashing.methodHash(ClassInfoMethodHashing.java:70) at org.jboss.aop.util.ClassInfoMethodHashing.addDeclaredMethods(ClassInfoMethodHashing.java:169) at org.jboss.aop.util.ClassInfoMethodHashing.populateMethodTables(ClassInfoMethodHashing.java:182) at org.jboss.aop.util.ClassInfoMethodHashing.populateMethodTables(ClassInfoMethodHashing.java:181) at org.jboss.aop.util.ClassInfoMethodHashing.populateMethodTables(ClassInfoMethodHashing.java:181) at org.jboss.aop.util.ClassInfoMethodHashing.populateMethodTables(ClassInfoMethodHashing.java:181) at org.jboss.aop.util.ClassInfoMethodHashing.populateMethodTables(ClassInfoMethodHashing.java:181) at org.jboss.aop.util.ClassInfoMethodHashing.getMethodMap(ClassInfoMethodHashing.java:188) at org.jboss.aop.microcontainer.integration.AOPDependencyBuilder.getMethodAnnotationDependencies(AOPDependencyBuilder.java:197) at org.jboss.aop.microcontainer.integration.AOPDependencyBuilder.getAnnotationDependencies(AOPDependencyBuilder.java:152) at org.jboss.aop.microcontainer.integration.AOPDependencyBuilder.getDependencies(AOPDependencyBuilder.java:130) at org.jboss.classadapter.plugins.BasicClassAdapter.getDependencies(BasicClassAdapter.java:80) at org.jboss.beans.info.plugins.AbstractBeanInfo.getDependencies(AbstractBeanInfo.java:211) at org.jboss.kernel.plugins.dependency.DescribeAction.installActionInternal(DescribeAction.java:53) at org.jboss.kernel.plugins.dependency.KernelControllerContextAction.installAction(KernelControllerContextAction.java:135) at org.jboss.kernel.plugins.dependency.KernelControllerContextAction.installAction(KernelControllerContextAction.java:46) at org.jboss.dependency.plugins.action.SimpleControllerContextAction.simpleInstallAction(SimpleControllerContextAction.java:62) at org.jboss.dependency.plugins.action.AccessControllerContextAction.install(AccessControllerContextAction.java:71) at org.jboss.dependency.plugins.AbstractControllerContextActions.install(AbstractControllerContextActions.java:51) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:327) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:1309) at org.jboss.dependency.plugins.AbstractController.incrementState(AbstractController.java:734) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:862) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:784) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:574) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:398) at	0	0	https://issues.jboss.org/browse/JBAS-5031
1109	JBAS-5031	Can we close it now?	1	1	https://issues.jboss.org/browse/JBAS-5031
1110	JBAS-5023	Jeff, can you take 2 look at those 2 failing tests?	0	0	https://issues.jboss.org/browse/JBAS-5023
1111	JBAS-5023	The patch looks right to me. Do you want to check it in once you have commit rights, Jeff? That way we do justice to you for finding out the location in the code.	0	0	https://issues.jboss.org/browse/JBAS-5023
1112	JBAS-5023	What is left here?	1	1	https://issues.jboss.org/browse/JBAS-5023
1113	JBAS-5022	Remy, can you take a look of this case? I don't think the provided patch is really related.	0	0	https://issues.jboss.org/browse/JBAS-5022
1114	JBAS-5007	Stefan are you looking into this?	1	1	https://issues.jboss.org/browse/JBAS-5007
1115	JBAS-4996	With the "old" jbossws versioning scheme is jbossws-3.0.1-native-2.0.4.GA equivalent to 2.0.4 or 3.0.1?	1	1	https://issues.jboss.org/browse/JBAS-4996
1116	JBAS-4993	Need to investigate this failure too: org.jboss.test.tm.test.TxTimeoutUnitTestCase: Pass: 1 Fail: 0 Error: 1 javax.ejb.EJBException: Should be marked rolled back: STATUS_ACTIVE at org.jboss.test.tm.ejb.TxTimeoutBean.testOverriddenTimeoutExpires(TxTimeoutBean.java:71)	0	0	https://issues.jboss.org/browse/JBAS-4993
1117	JBAS-4992	Stefan, can you comment on this behavior?	1	1	https://issues.jboss.org/browse/JBAS-4992
1118	JBAS-4956	what does "any more" mean in terms of jbossas, tck5 SVN revisions	0	0	https://issues.jboss.org/browse/JBAS-4956
1119	JBAS-4939	The other issue of server start-up being very slow because messages sent are not seen by the sender I think that is somehow related to the above one issue and the change in behavior after RHEL 4U4. So one of the above hopefully fix that: 1. not bind the multicast socket to specific interface or bind to the interface which multicast route is set through 2. set the socket option IP_MULTICAST_LOOP if not already set My feeling is that we should not try to guess correct interface but rather let OS choose the default one or the user use a start-up option to specify one. That will safe us from many issues.	0	0	https://issues.jboss.org/browse/JBAS-4939
1120	JBAS-4939	Dimitris, could you re-enabled the regular noip run for a day? Want to see if my JBAS-4966 fix took care of the problem; expect it will.	1	1	https://issues.jboss.org/browse/JBAS-4939
1121	JBAS-4926	This fails to identify deployments that should fail as failing. What is the real issue with this exception? Why shouldn't the ear deployment be failing?	0	0	https://issues.jboss.org/browse/JBAS-4926
1122	JBAS-4926	If I do not have the juddi-service.sar deployed i.e. there is no failed deployment because of [JBAS-4929] then the ear deployment does not fail either. Also, the junit output contains a list of all failed deployments for every test. This prevents junit-reports because of the xml volume to process. Is it not enough to list the deployment as failed when it actually is deployed for the first time?	0	0	https://issues.jboss.org/browse/JBAS-4926
1123	JBAS-4926	Hi, I allways get this error ("LifecycleException: Pipeline has not been started"), when I deploy my ear derived from Netbeans6.0/JBoss4.2.2. How can I solve this problem?	0	0	https://issues.jboss.org/browse/JBAS-4926

#	Requirement	Questions	Answered	Successfully Answered	URL
1124	JBAS-4916	This is thrown by Xerces. It's validation against the XSD. I tried a small validation test but it passed public void testMain() throws Exception { String xml = "<web-app xmlns='http://java.sun.com/xml/ns/j2ee' " + "xmlns:xsi='http://www.w3.org/2001/XMLSchema-instance' " + "xsi:schemaLocation='http://java.sun.com/xml/ns/j2ee http://java.sun.com/xml/ns/j2ee/web-app_2_4.xsd' " + "version='2.4'>" + "<servlet>" + " <servlet-name>TestService</servlet-name>" + " <servlet-class>org.jboss.test.ws.jaxrpc.enventry.EnvEntryBeanJSE</servlet-class>" + "</servlet>" + "<servlet-mapping>" + " <servlet-name>TestService</servlet-name>" + " <url-pattern>/*</url-pattern>" + "</servlet-mapping>" + "<env-entry>" + " <env-entry-name>jsr109/entry1</env-entry-name>" + " <env-entry-type>java.lang.String</env-entry-type>" + " <env-entry-value>web</env-entry-value>" + "</env-entry>" + "<env-entry>" + " <env-entry-name>jsr109/entry2</env-entry-name>" + " <env-entry-type>java.lang.Integer</env-entry-type>" + " <env-entry-value>8</env-entry-value>" + "</env-entry>" + "</web-app>"; Validator.assertValidXml(xml, new EntityResolver() { public InputSource resolveEntity(String publicId, String systemId) throws SAXException, IOException { String path = "file:///e:/cvsroot/projects/metadata/trunk/src/main/resources/schema/"; if(systemId != null && systemId.startsWith(path)) { path = systemId; } else if("http://java.sun.com/xml/ns/j2ee/web-app_2_4.xsd".equals(systemId)) { path += "web-app_2_4.xsd"; } else if("http://www.w3.org/2001/xml.xsd".equals(systemId)) { path += "xml.xsd"; } else if("http://www.ibm.com/webservices/xsd/j2ee_web_services_client_1_1.xsd".endsWith(systemId)) { path += "j2ee_web_services_client_1_1.xsd"; } else { System.out.println("resolveEntity: '" + publicId + "' '" + systemId + "'"); return null; } return new InputSource(path); }}); } Could you try to come up with a test like this that fails? I tried creating a test in the metadata project but that would also pass. I can commit it if you want.	1	1	https://issues.jboss.org/browse/JBAS-4916
1125	JBAS-4915	org.jboss.test.ws.jaxws.samples.webserviceref.WebServiceRefClientTestCase [junit] javax.naming.NamingException: Could not dereference object [Root exception is org.jboss.ws.WSException: Cannot create service] [junit] at org.jnp.interfaces.NamingContext.getObjectInstanceWrapFailure(NamingContext.java:1337) [junit] at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:765) [junit] at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:629) [junit] at org.jboss.ejb3.JndiUtil.lookup(JndiUtil.java:44) [junit] at org.jboss.injection.JndiFieldInjector.lookup(JndiFieldInjector.java:82) [junit] at org.jboss.injection.JndiFieldInjector.inject(JndiFieldInjector.java:101) [junit] at org.jboss.ejb3.client.ClientContainer.<init>(ClientContainer.java:120) [junit] at org.jboss.ejb3.client.ClientLauncher.launch(ClientLauncher.java:76) [junit] at org.jboss.ejb3.client.ClientLauncher.launch(ClientLauncher.java:179) [junit] at org.jboss.test.ws.jaxws.samples.webserviceref.WebServiceRefClientTestCase.testApplicationClient(WebServiceRefClientTestCase.java:80) [junit] at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) [junit] at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) [junit] at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) [junit] at java.lang.reflect.Method.invoke(Method.java:585) [junit] at junit.framework.TestCase.runTest(TestCase.java:154) [junit] at junit.framework.TestCase.runBare(TestCase.java:127) [junit] at junit.framework.TestResult$1.protect(TestResult.java:106) [junit] at junit.framework.TestResult.runProtected(TestResult.java:124) [junit] at junit.framework.TestResult.run(TestResult.java:109) [junit] at junit.framework.TestCase.run(TestCase.java:118) [junit] at junit.framework.TestSuite.runTest(TestSuite.java:208) [junit] at junit.framework.TestSuite.run(TestSuite.java:203) [junit] at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) [junit] at junit.extensions.TestSetup$1.protect(TestSetup.java:19) [junit] at junit.framework.TestResult.runProtected(TestResult.java:124) [junit] at junit.extensions.TestSetup.run(TestSetup.java:23) [junit] at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.run(JUnitTestRunner.java:297) [junit] at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.launch(JUnitTestRunner.java:672) [junit] at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.main(JUnitTestRunner.java:567) [junit] Caused by: org.jboss.ws.WSException: Cannot create service [junit] at org.jboss.ws.WSException.rethrow(WSException.java:60) [junit] at org.jboss.ws.core.jaxws.client.ServiceObjectFactoryJAXWS.getObjectInstance(ServiceObjectFactoryJAXWS.java:201) [junit] at javax.naming.spi.NamingManager.getObjectInstance(NamingManager.java:304) [junit] at org.jnp.interfaces.NamingContext.getObjectInstance(NamingContext.java:1312) [junit] at org.jnp.interfaces.NamingContext.getObjectInstanceWrapFailure(NamingContext.java:1329) [junit] ... 28 more [junit] Caused by: java.lang.IllegalArgumentException: Cannot create generic javax.xml.ws.Service without wsdlLocation: org.jboss.test.ws.jaxws.samples.webserviceref.TestEndpointClientOne/port4 [junit] at org.jboss.ws.core.jaxws.client.ServiceObjectFactoryJAXWS.getObjectInstance(ServiceObjectFactoryJAXWS.java:139) [junit] ... 31 more [junit] Tests run: 3, Failures: 0, Errors: 1, Time elapsed: 1,399.484 sec [junit] Test org.jboss.test.ws.jaxws.samples.webserviceref.WebServiceRefClientTestCase FAILED	1	1	https://issues.jboss.org/browse/JBAS-4915
1126	JBAS-4914	Ok, I get further, but something is still not built correctly as the reports output shows the test class is not found. This is after I have run ant and ant tests-compile. Testsuite: org.jboss.test.ws.jaxws.samples.context.WebServiceContextJSETestCase Tests run: 1, Failures: 0, Errors: 1, Time elapsed: 0 sec Caused an ERROR org.jboss.test.ws.jaxws.samples.context.WebServiceContextJSETestCase java.lang.ClassNotFoundException: org.jboss.test.ws.jaxws.samples.context.WebServiceContextJSETestCase at java.net.URLClassLoader$1.run(URLClassLoader.java:200) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:188) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:268) at java.lang.ClassLoader.loadClass(ClassLoader.java:251) at java.lang.ClassLoader.loadClassInternal(ClassLoader.java:319) at java.lang.Class.forName0(Native Method) at java.lang.Class.forName(Class.java:164)	1	1	https://issues.jboss.org/browse/JBAS-4914
1127	JBAS-4913	Is it org.jboss.metadata.javaee.spec.ServiceReferenceMetaData that is bound to service-ref? Could you tell which property in that class is supposed to be bound to wsdl-override element?	0	0	https://issues.jboss.org/browse/JBAS-4913
1128	JBAS-4906	How would autodiscovery work if different servers are using different multicast addresses and ports?	1	1	https://issues.jboss.org/browse/JBAS-4906

#	Requirement	Questions	Answered	Successfully Answered	URL
1129	JBAS-4891	There is a usage of such qname type elments in the EjbJar3xUnitTestCase.java, and this is parsing correctly. https://svn.jboss.org/repos/jbossas/projects/metadata/trunk/src/test/java/org/jboss/test/metadata/ejb/EjbJar3xUnitTestCase.java https://svn.jboss.org/repos/jbossas/projects/metadata/trunk/src/test/resources/org/jboss/test/metadata/ejb/EjbJar3x_testServiceRefs.xml <ejb-jar xmlns="http://java.sun.com/xml/ns/javaee" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://java.sun.com/xml/ns/javaee http://java.sun.com/xml/ns/javaee/ejb-jar_3_0.xsd" version="3.0" id="ejb-jar-service-refs"> <enterprise-beans> <session id="session1"> <description>A secured trival echo session bean</description> <ejb-name>StatelessSession1</ejb-name> <home>org.jboss.test.security.interfaces.StatelessSessionHome</home> <remote>org.jboss.test.security.interfaces.StatelessSession</remote> <local-home>org.jboss.test.security.interfaces.StatelessSessionLocalHome</local-home> <local>org.jboss.test.security.interfaces.StatelessSessionLocal</local> <ejb-class>org.jboss.test.security.ejb.StatelessSessionBean</ejb-class> <session-type>Stateless</session-type> <transaction-type>Container</transaction-type> <service-ref id="session1-service-ref"> <!-- JNDI name bound to the service interface in the client's environment context --> <service-ref-name>session1/Hello</service-ref-name> <!-- service interface class --> <service-interface>org.jboss.test.security.interfaces.HelloWorldService</service-interface> <service-ref-type>org.jboss.test.security.interfaces.HelloWorldRef</service-ref-type> <!-- placeholder for published WSDL file --> <wsdl-file>session1-wsdl-file</wsdl-file> <!-- Java<->XML mapping file --> <jaxrpc-mapping-file>META-INF/jaxrpc-mapping.xml</jaxrpc-mapping-file> <service-qname xmlns:session1="http://www.jboss.org">session1:Session1Qname</service-qname> <port-component-ref> <!-- service endpoint interface class --> <service-endpoint-interface>org.jboss.test.security.interfaces.HelloEndPoint</service-endpoint-interface> </port-component-ref> <handler-chains> <handler-chain id="session1-chain1"> <service-name-pattern>*</service-name-pattern> <handler id="sesion1-chain1-handler1"> <handler-name>session1.Handler1</handler-name> <handler-class>session1.Handler1</handler-class> <init-param> <param-name>param1</param-name> <param-value>value1</param-value> </init-param> <soap-header xmlns:session1="http://www.jboss.org">session1:soap-header</soap-header> <soap-role>session1:soap-role</soap-role> <port-name>session1:soap-role</port-name> </handler> </handler-chain> </handler-chains> </service-ref> </session> What is an example descriptor that is failing?	0	0	https://issues.jboss.org/browse/JBAS-4891
1130	JBAS-4891	The problem seems to be related jboss descriptors that use a dtd. Generally, dtd based descriptors don't seem to support namespace attributes. Currently schemas are not available for jboss-client.xml, jboss-web.xml that can also contain <service-qname> elements Another point is backwards compatiblity. Folks might have AS-4.2 descriptors that use dtd based descriptors with qname elements. If we go back to supporting the <service-qname>{http://some-namespace}Foo</service-qname> syntax, we would have different syntaxes in the standard vs. the proprietary descriptors. What do you think?	0	0	https://issues.jboss.org/browse/JBAS-4891
1131	JBAS-4882	Thomas, can you or someone from the ws team review this document? It probably needs updating to reflect the switch to jbossws-2.0.x.	0	0	https://issues.jboss.org/browse/JBAS-4882
1132	JBAS-4877	Are we going for the upgrade or not?	1	1	https://issues.jboss.org/browse/JBAS-4877
1133	JBAS-4869	Is there a link we can point to for supported Operating System/Databases etc?	1	0	https://issues.jboss.org/browse/JBAS-4869
1134	JBAS-4868	Verify the following content in this chapter: 1. Console output on server start up 2. Do we have a production profile? 3. Image snapshot	1	0	https://issues.jboss.org/browse/JBAS-4868
1135	JBAS-4867	Verify the following content in this chapter: 1. Do we have a GUI installer? 2. Installer Filename - jboss-4.2.2.GA.zip 3. Major Installation Steps 4. Shortcuts / XDG Menu applicable?	1	0	https://issues.jboss.org/browse/JBAS-4867
1136	JBAS-4851	Carlo, can you or someone from EJB3 review this document? I made a first pass with some minor changes.	1	1	https://issues.jboss.org/browse/JBAS-4851
1137	JBAS-4850	Should I close the issue now?	1	1	https://issues.jboss.org/browse/JBAS-4850
1138	JBAS-4842	Samson, in the respective subtasks, I see the wrong document attached to the tasks (e.g. the microkernel task has the naming document). Can you please check?	0	0	https://issues.jboss.org/browse/JBAS-4842
1139	JBAS-4834	Brian, Looking at the full AS 4.2.2 beta documentation, it looks as if the HA Singleton topics were included inside HTTP services: 5. HTTP Services 5.1. Configuring load balancing using Apache and mod_jk 5.2. Download the software 5.3. Configure Apache to load mod_jk 5.4. Configure worker nodes in mod_jk 5.5. Configuring JBoss to work with mod_jk 5.6. Configuring HTTP session state replication 5.7. Enabling session replication in your application 5.8. Using FIELD level replication 5.9. Monitoring session replication 5.10. Using Clustered Single Sign On 5.11. Clustered Singleton Services 5.11.1. HASingletonDeployer service 5.11.2. Mbean deployments using HASingletonController 5.11.3. HASingleton deployments using a Barrier 5.11.4. Determining the master node Was this intended? http://labs.jboss.com/file-access/default/members/jbossas/freezone/docs/Clustering_Guide/beta422/html/index.html	0	0	https://issues.jboss.org/browse/JBAS-4834
1140	JBAS-4832	What does "Out of Date" mean for this task? Clearly this document requires some work although it should be included in the installation guide. Installing a datasource in JBossAS is a typical task	0	0	https://issues.jboss.org/browse/JBAS-4832
1141	JBAS-4828	index.xhtml has a header like this: Seam POJO . Should it not be something like JSF-EJB3 example .	0	0	https://issues.jboss.org/browse/JBAS-4828
1142	JBAS-4822	HI, a few point made in this chapter: 1. There are 3 links under the Chapter header are these really necessary? The chapter is only 2 pages long 2. It is mentioned in 4.1 that the ZIP can be downloaded from http://labs.jboss.com/jbossas/downloads/ but the same document in chapter 2 it is talking about installation types for the EAP. Should we explain that this is also available from https://network.jboss.com? 3. 4.3. Post Installation Security Configuration - I found that both JMX & Web Console do not actually, in 4.2.2 ,require authentication when the installation guide says they do - It is mentioned that no user accounts exist, but the JBossAdmin account is present - The section in creating accounts/authentication is not straight forward to a new user (I propose new structure and addition of enabling <security-domain> JAAS, not mentioned) I hope this is of help	0	0	https://issues.jboss.org/browse/JBAS-4822
1143	JBAS-4811	Sure, we should move to SP2. How you propose to incorporate the WS tests in AS releases? Wouldn't the easiest thing be to replicate your hudson configuration in the QA environment?	1	1	https://issues.jboss.org/browse/JBAS-4811
1144	JBAS-4811	The hudson job is configured here. http://hudson.qa.jboss.com/hudson/view/JBoss%20AS/job/JBossWS-2.0.1.SP2-testsuite-AS4.2.x/ There are 3 test failures due to java.lang.ClassNotFoundException: org.apache.ws.scout.registry.ConnectionFactoryImpl Where is this class located?	1	1	https://issues.jboss.org/browse/JBAS-4811

#	Requirement	Questions	Answered	Successfully Answered	URL
1145	JBAS-4811	Shouldn't the tests be such so they run on a out-of-the box AS? Or if they require a different configuration, create it, as part of the test?	1	1	https://issues.jboss.org/browse/JBAS-4811
1146	JBAS-4811	Setting release version back to 2.2.2.GA I've also recreated the tag: https://svn.jboss.org/repos/jbossas/tags/JBoss_4_2_2_GA/ All the testsuites pass, as far as I can see. Rajesh, can we release today (Oct/22nd) ?	1	1	https://issues.jboss.org/browse/JBAS-4811
1147	JBAS-4811	Since you had a look at the hudson script, currently i unzip the juddi-service.sar in the default/deploy directory. unzip thirdparty/juddi-service.sar -d $WORKSPACE/$DIST_NAME/server/default/deploy/juddi-service.sar/ instead if i just copy the sar, the tests fail. cp thirdparty/juddi-service.sar $WORKSPACE/$DIST_NAME/server/default/deploy/ Do you know why that happens?	1	1	https://issues.jboss.org/browse/JBAS-4811
1148	JBAS-4806	Deepak, can you look at this task?	1	1	https://issues.jboss.org/browse/JBAS-4806
1149	JBAS-4795	The JBDOCS-108 task and its subtasks are a dependency of the 4.0.3 guide updates.	1	1	https://issues.jboss.org/browse/JBAS-4795
1150	JBAS-4752	I am working on JBOSS 4.2.3. I am having one issue. I deploy my application on JBOSS. I use spnego protocol for Single Sign On. I am running into an issue where the application is running out of memory. I took heap dumps at regular intervals. I open every dump in Eclipse Memory Analyzer. I open the dominator_tree. I see a big percentage of entries (it grows from small percentage to a bigger one) occupied by class org.jboss.util.TimedCachePolicy. There one entry for each https request to the application. I think they should not get accumulated. My question is if this fix (JBAS-4752) would resolve this issue? Any suggestions / recommendations / questions / details you need? Thanks	0	0	https://issues.jboss.org/browse/JBAS-4752
1151	JBAS-4743	Sorry Eric, can you also add those 3 dtds from here: https://svn.jboss.org/repos/jbossas/trunk/mbeans/src/resources/dtd/	1	0	https://issues.jboss.org/browse/JBAS-4743
1152	JBAS-4743	Eric, there is one more schema file attached (jboss_5_0.xsd) that for some reason is not included with the other dtds/schemas. Can you add this as well, please?	1	0	https://issues.jboss.org/browse/JBAS-4743
1153	JBAS-4706	Scott, you are doing the merging deployer. Can you please align this JIRA issue with whatever issue that you have open for the merging deployer?	0	0	https://issues.jboss.org/browse/JBAS-4706
1154	JBAS-4696	Mark, was there any feedback from Joshua?	1	1	https://issues.jboss.org/browse/JBAS-4696
1155	JBAS-4696	Okay, Mark, I'm on it. Step 1: Move JBEAP4.2 forked content (x6 languages) back into upstream Target is 28 Sept and branch as Enterprise. Step 2: Create test branch and rearchitect to build with Maven2 Step 3: Have AS team examine test branch and if good to go, move to content to head	1	1	https://issues.jboss.org/browse/JBAS-4696
1156	JBAS-4696	Status Update: ========== Step 1: Move JBEAP4.2 forked content (x6 languages) back into upstream Target is 28 Sept and branch as Enterprise. |_DONE: Enterprise branch is JBEAP42CP01. Note will only build on linux with the documentation-devel package Step 2: Create test branch and rearchitect to build with Maven2 |_IN PROCESS: Content in trunk is being rearchitected. Giving myself 2 days to complete this. I have requested some additional help from MarkN. re: using maven. my upstream maven plugin is unstable and I need a baseline to work from. Hopefuly, Mark can give me an example. ETA: COB Oct 2nd, for Installation Guide. Step 3: Have AS team examine test branch and if good to go, move to content to head |_NOT DONE	0	0	https://issues.jboss.org/browse/JBAS-4696
1157	JBAS-4694	Is this already done? I am not sure if this introduced 200+ failing tests in cmp2 in both the AS 4.2 and 5.0 testsuites from around the same time this JIRA was created.	1	1	https://issues.jboss.org/browse/JBAS-4694
1158	JBAS-4694	Either this, or a change by Adrian. Can you look more into that, or simply rollback the hsqldb in a local copy and re-test?	1	1	https://issues.jboss.org/browse/JBAS-4694
1159	JBAS-4688	Where does the transaction-spi live ?	1	1	https://issues.jboss.org/browse/JBAS-4688
1160	JBAS-4688	Pointing trunk to Beta2. I see the integration project currently at 5.0.0.Beta3. Is an update to that necessary, as well?	1	1	https://issues.jboss.org/browse/JBAS-4688
1161	JBAS-4675	It's ok to disallow transaction local delegate construction if the status is not ACTIVE, similar to you cannot create the synchronization. It's NOT ok to disallow retrieval of the delegate if it is was previouslly allowed to be constructed (assuming the transaction hasn't completed - rollback, heuristic or otherwise). It's virtually the same rules as transaction synchronization. i.e. You register the synchronization, setRollbackOnly() and still the afterCompletion() gets invoked. I imagine this also affects the TransactionSynchronizationRegistry implementation of JavaEE5 which has a simpler "TransactionLocal" put/getResource notion?	1	1	https://issues.jboss.org/browse/JBAS-4675
1162	JBAS-4672	Is there any discussion what needs to be done for this task?	1	1	https://issues.jboss.org/browse/JBAS-4672
1163	JBAS-4657	For the AS project we need subtasks (like this) for any jboss/thirdparty component that needs updating, so we know what's left and were we stand. Also, JBossWeb need to be tracked somewhere, no? Looking into it's JIRA there hasn''t been any recorded release in 2007 http://jira.jboss.com/jira/secure/project/ViewProject.jspa?pid=10054 How do we know what feature/fix got into what release, especially now that jboss web is becoming a fork of tomcat? Can this info be found somewhere else?	1	1	https://issues.jboss.org/browse/JBAS-4657
1164	JBAS-4644	3 failures are still there: org.jboss.test.webservice.jbws309.JBWS309TestCase(JACC).testBasicSecuredServiceAccess org.jboss.test.webservice.jbws309.JBWS309TestCase(JACC).testBasicSecuredDIIAccess org.jboss.test.webservice.jbws309.JBWS309TestCase(JACC).testRoleSecuredServiceAccess All fail similarly with: junit.framework.AssertionFailedError: Security exception expected at org.jboss.test.webservice.jbws309.JBWS309TestCase.testRoleSecuredServiceAccess(JBWS309TestCase.java:179) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23)	1	1	https://issues.jboss.org/browse/JBAS-4644
1165	JBAS-4644	Anil is away on vacation, can you take a look Alessio? This has dragged on for a long time and it makes up 3 of the remaining 21 failures of the main testsuite.	1	1	https://issues.jboss.org/browse/JBAS-4644
1166	JBAS-4644	Dimitris, I've looked at this but did not come to a solution. I've just wrote some comments on the linked forum thread. Is Anil still on vacation? I might go on digging this, but don't know how much time this could take me, may be someone with more knowledge on this could speed up the process.	0	0	https://issues.jboss.org/browse/JBAS-4644
1167	JBAS-4644	Anil is coming back on the 27th. Who would be more appropriate to help? If the jbossws endpoint servlet is not called then it's a deployment problem, right?	1	1	https://issues.jboss.org/browse/JBAS-4644
1168	JBAS-4644	> Who would be more appropriate to help? Well, don't know, may be someone from the security project Still looking at this anyway.	0	0	https://issues.jboss.org/browse/JBAS-4644
1169	JBAS-4644	So who's looking at this Stefan or Anil?	0	0	https://issues.jboss.org/browse/JBAS-4644

#	Requirement	Questions	Answered	Successfully Answered	URL
1170	JBAS-4644	Stefan, I am wondering if you can do the following. 1) Create a jacc instance by running the "ant tests-jacc-security" target once. 2) Then start the jacc instance. 3) Now try running this JBWS test as" ant -Dtest=org.jboss.test... one-test" 4) Place break points in the AbstractSecurityDeployer (and sub security deployers), tomcat deployer etc. What I want to know is when does the WS layer attach the web metadata as an attachment into the deployment unit (so you may have to trace log the mc layer etc). The way to fix this is that after the ws layer creates a JBossWebMetaData and attaches it to the deployment unit, we need to send the meta data via the security deployers such that the jacc permissions can be created.	1	1	https://issues.jboss.org/browse/JBAS-4644
1171	JBAS-4643	file:/services/cruisecontrol/checkout/jboss-head-testsuite-sun-1.5/testsuite/output/resources/org/jboss/test/xml/ExtendedByImportUnitTestCase.xml@6,62 org.jboss.xb.binding.JBossXBException: Failed to parse source: file:/services/cruisecontrol/checkout/jboss-head-testsuite-sun-1.5/testsuite/output/resources/org/jboss/test/xml/ExtendedByImportUnitTestCase.xml@6,62 at org.jboss.xb.binding.parser.sax.SaxJBossXBParser.parse(SaxJBossXBParser.java:184) at org.jboss.xb.binding.UnmarshallerImpl.unmarshal(UnmarshallerImpl.java:139) at org.jboss.test.xml.JBossXBTestDelegate.unmarshal(JBossXBTestDelegate.java:118) at org.jboss.test.xml.AbstractJBossXBTest.unmarshal(AbstractJBossXBTest.java:119) at org.jboss.test.xml.AbstractJBossXBTest.unmarshal(AbstractJBossXBTest.java:106) at org.jboss.test.xml.AbstractJBossXBTest.unmarshal(AbstractJBossXBTest.java:72) at org.jboss.test.xml.ExtendedByImportUnitTestCase.testExtended(ExtendedByImportUnitTestCase.java:55) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) Caused by: java.lang.ClassCastException: java.lang.String at org.jboss.kernel.plugins.deployment.xml.BeanPropertyInterceptor.add(BeanPropertyInterceptor.java:48) at org.jboss.xb.binding.sunday.unmarshalling.SundayContentHandler.endElement(SundayContentHandler.java:1197) at org.jboss.xb.binding.sunday.unmarshalling.SundayContentHandler.endElement(SundayContentHandler.java:248) at org.jboss.xb.binding.parser.sax.SaxJBossXBParser$DelegatingContentHandler.endElement(SaxJBossXBParser.java:374) at org.apache.xerces.parsers.AbstractSAXParser.endElement(Unknown Source) at org.apache.xerces.xinclude.XIncludeHandler.endElement(Unknown Source) at org.apache.xerces.impl.XMLNSDocumentScannerImpl.scanEndElement(Unknown Source) at org.apache.xerces.impl.XMLDocumentFragmentScannerImpl$FragmentContentDispatcher.dispatch(Unknown Source) at org.apache.xerces.impl.XMLDocumentFragmentScannerImpl.scanDocument(Unknown Source) at org.apache.xerces.parsers.XML11Configuration.parse(Unknown Source) at org.apache.xerces.parsers.XML11Configuration.parse(Unknown Source) at org.apache.xerces.parsers.XMLParser.parse(Unknown Source) at org.apache.xerces.parsers.AbstractSAXParser.parse(Unknown Source) at org.apache.xerces.jaxp.SAXParserImpl$JAXPSAXParser.parse(Unknown Source) at org.jboss.xb.binding.parser.sax.SaxJBossXBParser.parse(SaxJBossXBParser.java:180) ... 25 more and org.jboss.test.xml.SingletonSchemaBindingResolverUnitTestCase testBean Error Failed to parse source: file:/services/cruisecontrol/checkout/jboss-head-testsuite-sun-1.5/testsuite/output/resources/org/jboss/test/xml/SingletonSchemaBindingResolverUnitTestCaseContainerStrictBean.xml@4,2 org.jboss.xb.binding.JBossXBException: Failed to parse source: file:/services/cruisecontrol/checkout/jboss-head-testsuite-sun-1.5/testsuite/output/resources/org/jboss/test/xml/SingletonSchemaBindingResolverUnitTestCaseContainerStrictBean.xml@4,2 at org.jboss.xb.binding.parser.sax.SaxJBossXBParser.parse(SaxJBossXBParser.java:184) at org.jboss.xb.binding.UnmarshallerImpl.unmarshal(UnmarshallerImpl.java:139) at org.jboss.test.xml.JBossXBTestDelegate.unmarshal(JBossXBTestDelegate.java:118) at org.jboss.test.xml.AbstractJBossXBTest.unmarshal(AbstractJBossXBTest.java:119) at org.jboss.test.xml.AbstractJBossXBTest.unmarshal(AbstractJBossXBTest.java:106) at org.jboss.test.xml.AbstractJBossXBTest.unmarshal(AbstractJBossXBTest.java:72)	0	0	https://issues.jboss.org/browse/JBAS-4643
1172	JBAS-4643	This is last failure. Anil, could you please have a look and either add it to the ignored list or make it valid? Testcase: testJBossWeb took 0.541 sec FAILED 1 from 43 are invalid: [Failed to parse: E:\cvsroot\trunk\testsuite\output\resources\web\programmatic\WEB-INF\jboss-web.xml: Premature end of file. [-1,-1]] expected:<0> but was:<1> junit.framework.AssertionFailedError: 1 from 43 are invalid: [Failed to parse: E:\cvsroot\trunk\testsuite\output\resources\web\programmatic\WEB-INF\jboss-web.xml: Premature end of file. [-1,-1]] expected:<0> but was:<1> at org.jboss.test.xml.DDValidatorUnitTestCase.validate(DDValidatorUnitTestCase.java:148) at org.jboss.test.xml.DDValidatorUnitTestCase.testJBossWeb(DDValidatorUnitTestCase.java:129)	0	0	https://issues.jboss.org/browse/JBAS-4643
1173	JBAS-4642	Alex can you have a quick look at some point, see if the failure rings a bell?	1	1	https://issues.jboss.org/browse/JBAS-4642
1174	JBAS-4637	Clebert, can you help?	0	0	https://issues.jboss.org/browse/JBAS-4637
1175	JBAS-4637	Looking at the test failures in Hudson for org.jboss.test.jbossmessaging.ra. are all passing for me locally. The issue that is seen during the Hudson runs are that deployment fails due to: java.lang.IllegalStateException: Cannot set paging params when active I am unable to get much system information/server.log from Hudson when these errors were encountered but it seems that there may be some issues with paging/swapping during a full test run or insufficiently configured for such? one-test: [junit] Running org.jboss.test.jbossmessaging.ra.RaJMSSessionUnitTestCase [junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 1.418 sec [junit] Running org.jboss.test.jbossmessaging.ra.RaQueueUnitTestCase [junit] Tests run: 4, Failures: 0, Errors: 0, Time elapsed: 24.29 sec [junit] Running org.jboss.test.jbossmessaging.ra.RaJMSSessionUnitTestCase [junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 2.392 sec [junit] Running org.jboss.test.jbossmessaging.ra.RaTopicUnitTestCase [junit] Tests run: 4, Failures: 0, Errors: 0, Time elapsed: 6.861 sec The test failures in: org.jboss.test.jmsra.test.Ra*UnitTestCase are reproducable in my local environment and will investigate further: javax.naming.NoInitialContextException: Need to specify class name in environment or system property, or as an applet parameter, or in an application resource file: java.naming.factory.initial at javax.naming.spi.NamingManager.getInitialContext(NamingManager.java:645) at javax.naming.InitialContext.getDefaultInitCtx(InitialContext.java:247) at javax.naming.InitialContext.getURLOrDefaultInitCtx(InitialContext.java:284) at javax.naming.InitialContext.lookup(InitialContext.java:351)	0	0	https://issues.jboss.org/browse/JBAS-4637

#	Requirement	Questions	Answered	Successfully Answered	URL
1176	JBAS-4636	15:05:26,609 INFO [MainDeployer] deploy, url=file:/X:/cvs/jboss-public/jboss-he ad/testsuite/output/lib/testdeployers-ear1.ear 15:05:27,375 WARN [CollectionPropertyHandler] ClassInfo.getDeclaredConstructor(null) didn't work for org.jboss.metadata.ejb.spec.InterceptorClassesMetaData, fo und the default ctor in ClassInfo.getDeclaredConstructors() 15:05:28,156 INFO [EjbDeployer] installing bean: ejb/testdeployers-bean1ejb.jar #Bean1EJB,uid21412463 15:05:28,156 INFO [EjbDeployer] with dependencies: 15:05:28,156 INFO [EjbDeployer] and supplies: 15:05:28,156 INFO [EjbDeployer] jndi:Bean1EJB 15:05:28,203 INFO [JBossASKernel] Created KernelDeployment for: testdeployers-c lient1.jar 15:05:28,203 INFO [JBossASKernel] installing bean: jboss.j2ee:ear=testdeployers - ear1.ear,jar=testdeployers-client1.jar,name=test-client,service=EJB3 15:05:28,203 INFO [JBossASKernel] with dependencies: 15:05:28,203 INFO [JBossASKernel] and demands: 15:05:28,203 INFO [JBossASKernel] and supplies: 15:05:28,203 INFO [JBossASKernel] Added bean(jboss.j2ee:ear=testdeployers-ear1. ear,jar=testdeployers-client1.jar,name=test-client,service=EJB3) to KernelDeploy ment of: testdeployers-client1.jar 15:05:28,250 INFO [ClientENCInjectionContainer] STARTED CLIENT ENC CONTAINER: t est-client 15:05:28,250 INFO [ClientENCInjectionContainer] STOPPED CLIENT ENC CONTAINER: t est-client 15:05:28,296 ERROR [AbstractKernelController] Error installing to Real: name=vfs zip:/X:/cvs/jboss-public/jboss-head/testsuite/output/lib/testdeployers-ear1.ear state=PreReal mode=Manual requiredState=Real org.jboss.deployers.spi.DeploymentException: Error during deploy: vfszip:/X:/cvs /jboss-public/jboss-head/testsuite/output/lib/testdeployers-ear1.ear/testdeploye rs-mcf1-ds.xml at org.jboss.deployers.spi.DeploymentException.rethrowAsDeploymentExcept ion(DeploymentException.java:49) at org.jboss.deployers.plugins.deployers.DeployerWrapper.deploy(Deployer Wrapper.java:180) at org.jboss.deployers.plugins.deployers.DeployersImpl.doInstallParentFi rst(DeployersImpl.java:970) at org.jboss.deployers.plugins.deployers.DeployersImpl.doInstallParentFi rst(DeployersImpl.java:1023) at org.jboss.deployers.plugins.deployers.DeployersImpl.install(Deployers Impl.java:911) at org.jboss.dependency.plugins.AbstractControllerContext.install(Abstra ctControllerContext.java:348) at org.jboss.dependency.plugins.AbstractController.install(AbstractContr oller.java:1392) at org.jboss.dependency.plugins.AbstractController.incrementState(Abstra ctController.java:784) at org.jboss.dependency.plugins.AbstractController.resolveContexts(Abstr actController.java:912) at org.jboss.dependency.plugins.AbstractController.resolveContexts(Abstr actController.java:834) at org.jboss.dependency.plugins.AbstractController.change(AbstractContro ller.java:672) at org.jboss.dependency.plugins.AbstractController.change(AbstractContro ller.java:455) at org.jboss.deployers.plugins.deployers.DeployersImpl.process(Deployers Impl.java:594) at org.jboss.deployers.plugins.main.MainDeployerImpl.process(MainDeploye rImpl.java:541) at org.jboss.deployment.MainDeployer.deploy(MainDeployer.java:812) at org.jboss.deployment.MainDeployer.redeploy(MainDeployer.java:587) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl. java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAcces sorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatch er.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at	0	0	https://issues.jboss.org/browse/JBAS-4636
1177	JBAS-4636	I still see this issue intermittently in 5.0.0.CR1 when attempting to hot deploy my ear. Ales' last comment is not clear. Is this something that the application implementation needs to do or is the fix embedded in JBAS? Our application deployed fine in JB 4.0.2 and does not use a jboss-app.xml file in its manifest. By 'intermittently' I can get this failure 2, 3 sometimes 6 times in a row, then, for whatever reason, deployment succeeds. Also, when am I going to be able to NOT have to use hot-deployment. There still appears to be deployment ordering issues in CR1. Thanks.	0	0	https://issues.jboss.org/browse/JBAS-4636
1178	JBAS-4635	Shelly, can you make sure we have subtasks to cover all testsuite failures, grouped together as much as possible?	0	0	https://issues.jboss.org/browse/JBAS-4635
1179	JBAS-4591	Were is this the discussion thread for this? If you do it, don't change the MDB to enable this by default. I deprecated this whole feature in JBoss4 and disabled in JBoss5. Not because it is JBossMQ specific but because it leads to broken behaviour. See where this has been discussed it before in the JCA and EJB design forums for why this is a bad idea!	1	0	https://issues.jboss.org/browse/JBAS-4591
1180	JBAS-4586	Hey Troy, are you still out there? I'm sorry this issue's been flying under the radar for so long. I think I know what was happening. The standard EJB3 configuration uses the Remoting "socket" transport, which has a default limit of 300 worker threads on the server side. Once a worker thread is connected to a client, it will keep listening for further invocations from that client until either (1) the socket times out, or (2) the client disconnects. So it's likely that most of the delay experienced by your users was spent waiting for worker thread to free up. The number of worker threads can be increased by changing the EJB3 configuration: <mbean code="org.jboss.remoting.transport.Connector" name="jboss.remoting:type=Connector,name=DefaultEjb3Connector,handler=ejb3"> <depends>jboss.aop:service=AspectDeployer</depends> <attribute name="InvokerLocator">socket://${jboss.bind.address}:3873/?maxPoolSize=600</attribute> <attribute name="Configuration"> <handlers> <handler subsystem="AOP">org.jboss.aspects.remoting.AOPRemotingInvocationHandler</handler> </handlers> </attribute> </mbean> Here I've added the "maxPoolSize" parameter to the "InvokerLocator" attribute. An alternative would be to reduce the timeout value, so worker threads return to the thread pool sooner: <mbean code="org.jboss.remoting.transport.Connector" name="jboss.remoting:type=Connector,name=DefaultEjb3Connector,handler=ejb3"> <depends>jboss.aop:service=AspectDeployer</depends> <attribute name="InvokerLocator">socket://${jboss.bind.address}:3873/?timeout=1000</attribute> <attribute name="Configuration"> <handlers> <handler subsystem="AOP">org.jboss.aspects.remoting.AOPRemotingInvocationHandler</handler> </handlers> </attribute> </mbean> Here I've set the "timeout" parameter to "1000". If you're still interested in pursuing this issue, let me know. Otherwise, I'll assume it's a configuration problem and close the issue.	1	1	https://issues.jboss.org/browse/JBAS-4586

#	Requirement	Questions	Answered	Successfully Answered	URL
1181	JBAS-4572	I created "depends" links to the EJB3 issue that *must* be solved, and "related" links to the known failures, according to Carlo's input: > > org.jboss.ejb3.test.clusteredsession.unit.NestedBeanUnitTestCase(BR-enabled).testStatefulPassivation EJBTHREE-1053 + EJBTHREE-1054 (to some extend) > > org.jboss.ejb3.test.ejbthree832.unit.EntityTestCase.test1 EJBTHREE-832: known issue > > org.jboss.ejb3.test.iiop.unit.IiopRemoteUnitTestCase(jacorb).testTxPropegation EJBTHREE-892: known issue > > org.jboss.ejb3.test.interceptors2.unit.InterceptorsTestCase.testOverrideOrdering EJBTHREE-896: known issue > > org.jboss.ejb3.test.jacc.unit.JaccTestCase.testSomeEntity EJBTHREE-894: known issue > > org.jboss.ejb3.test.jaxws.unit.ContextEJBTestCase.testClientAccessRPC EJBTHREE-757: is sessionContext.getMessageContext supposed to work? It's for JAX-RPC which is obsolete, right? Rephrase: do we do JAX-RPC with EJB 3? > > org.jboss.ejb3.test.webservices.unit.JSR181TestCase.testWebServiceRef3 > > org.jboss.ejb3.test.webservices.unit.JSR181TestCase.testWebServiceRef4 EJBTHREE-899: web service ref injection via dd I've added a test for the post construct issue: EJBTHREE-1055.	0	0	https://issues.jboss.org/browse/JBAS-4572
1182	JBAS-4568	Is 2.2.0.beta still the target AOP release? (currently on snapshot)	0	0	https://issues.jboss.org/browse/JBAS-4568
1183	JBAS-4567	Is the target 2.1.0 or 2.0.1 ?	1	1	https://issues.jboss.org/browse/JBAS-4567
1184	JBAS-4565	Hibernate Validator was implicitly included into Hibernate Annotations 3.2.x, it is now an independent package. Should we include it or not?	1	1	https://issues.jboss.org/browse/JBAS-4565
1185	JBAS-4565	Currently we'll have the following libs on the server/xxx/lib dir: hibernate-annotations.jar hibernate-commons-annotations.jar hibernate-entitymanager.jar hibernate3.jar jboss-hibernate.jar I will be adding hibernate-validator.jar, too. On the client side we have hibernate-annotations.jar hibernate-client.jar Is any other library needed on the client side?	1	1	https://issues.jboss.org/browse/JBAS-4565
1186	JBAS-4564	Yes hibernate-commons-annotations is required. Another problem may arise. Hibernate Validator (the component) is now separate from hibernate-annotations Do we consider it part of the JBoss AS distro, or is it safe to not include it (ie let the user add it if needed)?	1	1	https://issues.jboss.org/browse/JBAS-4564
1187	JBAS-4563	Is this the latest version we can use?	1	1	https://issues.jboss.org/browse/JBAS-4563
1188	JBAS-4562	Ideally, I would like to integrate a CR release of JBoss Web 2.1 in the beta 3 build. This should not be an issue, but there will be a small update needed since this new version slightly updates the annotation integration interface. Would this be acceptable, or would you prefer sticking with a JBoss Web 2.0.1 GA and upgrading post beta 3 ?	1	1	https://issues.jboss.org/browse/JBAS-4562
1189	JBAS-4533	Do we want a different tag than the current "5.0.0.Beta3" one?	0	0	https://issues.jboss.org/browse/JBAS-4533
1190	JBAS-4501	Can we please get 4.2.1.GA building with JDK 1.6?	1	1	https://issues.jboss.org/browse/JBAS-4501
1191	JBAS-4499	Dimitris.. I have assigned this to you.. for your evaluation first.	1	1	https://issues.jboss.org/browse/JBAS-4499
1192	JBAS-4496	Is that ok to reschedule this for 4.2.2 ? I don't want more changes in Branch_4_2 before we release a community 4.2.1.GA	1	0	https://issues.jboss.org/browse/JBAS-4496
1193	JBAS-4438	Please don't shoot the messenger! I filed this in response to Weston Price's suggestion to file a JIRA for it (see http://www.jboss.org/index.html?module=bb&op=viewtopic&t=105141) If the container handles redelivery by NACKing back to the destination, then I can see this doesn't actually make a lot of sense, although could still be useful for JMS providers that don't implement any kind of redelivery delay themselves. Feel free to reject it. (Aside: Is NACKing back to the queue really the correct way to handle rollback? The JMS spec says that session rollback should cause *recovery* of the session (Sec 4.4.7). Session recovery (Sec. 4.4.11) is supposed to stop and restart delivery of the last unacknowledged in that session. If you NACK on rollback then the message will go back to the queue and if there is *another* session consuming from the same queue, then the other session may receive the NACKed message, which would be contrary to the JMS spec which is pretty clear the message should be delivered back to the same session. For this reason, JBoss Messaging (unlike JBossMQ) always triggers session recovery (not NACK) on rollback, to ensure the correct session gets the rolled back messages.)	0	0	https://issues.jboss.org/browse/JBAS-4438
1194	JBAS-4423	Anil, what is the status on this task. The EJB3 team urgently needs this to properly re-factor and stabilize their work. What's your ETA for this to be done. Thanks.	1	1	https://issues.jboss.org/browse/JBAS-4423
1195	JBAS-4401	What release is this reported against?	0	0	https://issues.jboss.org/browse/JBAS-4401
1196	JBAS-4392	Is there more work to it?	0	0	https://issues.jboss.org/browse/JBAS-4392
1197	JBAS-4388	JBPAPP-327 is not related to this issue. JBPAPP-327 concerns security domains. This issue concerns the web context root of current and future administrative services. SecureTheInvokers does not resolve the issue. Placing an authentication requirement on the invokers, console, etc does not prevent dictionary or brute force attacks, or abuse from leaked passwords or ex-employees. In my experience, most major enterprise environments have policies that prohibit public access to administrative login interfaces. All such interfaces are only accessible from the private network, requiring VPN access (usually with two-factor authentication) from remote sources.	1	1	https://issues.jboss.org/browse/JBAS-4388
1198	JBAS-4383	Are there any docs or info on how this works now? Thanks.	0	0	https://issues.jboss.org/browse/JBAS-4383
1199	JBAS-4383	How can I achieve the same with AS 6. I am currently working on deploying an app using JSF 1.2 to AS 6. Thanks	1	1	https://issues.jboss.org/browse/JBAS-4383
1200	JBAS-4378	What's the status of this issue? All the subtasks are not scheduled at all.	1	1	https://issues.jboss.org/browse/JBAS-4378
1201	JBAS-4373	The following permission needs to be added ---------- 12:05:36,760 INFO [TomcatDeployer] deploy, ctxPath=/class-loading, warUrl=.../tmp/deploy/tmp42192class-loading-exp.war/ 12:05:37,081 ERROR [[ClasspathServlet2]] Servlet.service() for servlet ClasspathServlet2 threw exception java.security.AccessControlException: access denied (java.lang.RuntimePermission getClassLoader) at java.security.AccessControlContext.checkPermission(AccessControlContext.java:264) at java.security.AccessController.checkPermission(AccessController.java:427) at java.lang.SecurityManager.checkPermission(SecurityManager.java:532) at java.lang.Class.getClassLoader(Class.java:588) at org.jboss.test.util.Debug.displayClassInfo(Debug.java:43) at org.jboss.test.web.servlets.ClasspathServlet2.processRequest(ClasspathServlet2.java:77) at org.jboss.test.web.servlets.ClasspathServlet2.doGet(ClasspathServlet2.java:50) at javax.servlet.http.HttpServlet.service(HttpServlet.java:690) at javax.servlet.http.HttpServlet.service(HttpServlet.java:803) at sun.reflect.GeneratedMethodAccessor222.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.apache.catalina.security.SecurityUtil$1.run(SecurityUtil.java:244) ---------- The permission should be: permission java.lang.RuntimePermission "getClassLoader"; Try adding that and see if the test cases pass.	1	1	https://issues.jboss.org/browse/JBAS-4373

#	Requirement	Questions	Answered	Successfully Answered	URL
1202	JBAS-4359	Ron... I'm assigning this to you.. You have the fix attached to this issue. All you need to do is to patch jboss-remoting 2.2 (maybe create a SP2?) with the fix. I tried to commit this myself on remoting branch, but I was confused about what branch to use, since I didn't recognize what branch to use based in your naming standard.	1	1	https://issues.jboss.org/browse/JBAS-4359
1203	JBAS-4359	Thank you for catching this, Clebert. I'm running the Remoting test suite with your changes, and as soon as it's done. I wll put a 2.2.0.SP2 jboss-remoting.jar in the repository. Clebert, can you think of a unit test? FYI: - Remoting HEAD is basically Remoting 3 - Remoting_2_x is "HEAD" for Remoting 2.x - Remoting_2_2_0_GA is for patches to Remoting 2.2.0	0	0	https://issues.jboss.org/browse/JBAS-4359
1204	JBAS-4359	Do you guys think this has performance implications? Looks like a local 'flush' after every write?	1	1	https://issues.jboss.org/browse/JBAS-4359
1205	JBAS-4355	Paul, not really a maven issue, but can you look into it?	0	0	https://issues.jboss.org/browse/JBAS-4355
1206	JBAS-4355	I think I"ve fixed it. Can you verify on Unix?	0	0	https://issues.jboss.org/browse/JBAS-4355
1207	JBAS-4320	Hi Dimitris, Thanks for replying. I've put it under feature request, since it works with HSQLDB. Is this a bug or a misconfiguration? Thanks in advance	0	0	https://issues.jboss.org/browse/JBAS-4320
1208	JBAS-4317	With various jbossws tests I was getting IllegalStateException("Security Context in invocation is null"); Please validate the change below. First throwing an ISE and then further down calling seems not logical //Place on the invocation if(mi.getSecurityContext() == null) mi.setSecurityContext(SecurityActions.getSecurityContext()); [tdiesler@tddell trunk]$ svn diff server Index: server/src/main/org/jboss/ejb/SessionContainer.java ==========...=== --- server/src/main/org/jboss/ejb/SessionContainer.java (revision 62601) +++ server/src/main/org/jboss/ejb/SessionContainer.java (working copy) @@ -620,9 +620,9 @@ public Object internalInvokeHome(Invocation mi) throws Exception { - //Validate that there is a security context on the invocation + //Place on the invocation if(mi.getSecurityContext() == null) - throw new IllegalStateException("Security Context in invocation is null"); + mi.setSecurityContext(SecurityActions.getSecurityContext()); String securityDomain = SecurityConstants.DEFAULT_APPLICATION_POLICY; if(sm != null) @@ -633,9 +633,6 @@ mi.getCredential(), securityDomain, null); } SecurityActions.pushCallerRunAsIdentity(mi.getSecurityContext().getRunAs()); - //Place on the invocation - if(mi.getSecurityContext() == null) - mi.setSecurityContext(SecurityActions.getSecurityContext()); Method method = mi.getMethod(); if (method != null && method.getName().equals("remove")) @@ -672,9 +669,9 @@ */ public Object internalInvoke(Invocation mi) throws Exception { - //Validate that there is a security context on the invocation + //Place on the invocation if(mi.getSecurityContext() == null) - throw new IllegalStateException("Security Context in invocation is null"); + mi.setSecurityContext(SecurityActions.getSecurityContext()); String securityDomain = SecurityConstants.DEFAULT_APPLICATION_POLICY; if(sm != null) @@ -685,9 +682,6 @@ mi.getCredential(), securityDomain, null); } SecurityActions.pushCallerRunAsIdentity(mi.getSecurityContext().getRunAs()); - //Place on the invocation - if(mi.getSecurityContext() == null) - mi.setSecurityContext(SecurityActions.getSecurityContext()); try {	1	1	https://issues.jboss.org/browse/JBAS-4317
1209	JBAS-4317	Thomas, the security context either comes over the wire (remote calls) or comes from the thread local (Local EJB invocations). So where-ever the Invocation object is created on the server side, the security context needs to be set on the Invocation object. The IllegalStateException thrown in the containers was one way of validating that whoever was creating the Invocation object has set the security context (just the way they would have done with .setPrincipal, setCredential etc). The primary issue is that there are various integration layers constructing the Invocation object rather than a central place. Some of the examples where the Invocation object is created on the server side include the BaseLocalProxyFactory, ProxyFinderFactory, CMPFieldBridgexxxx. So I will need to revert back the IllegalStateException and need your stack trace so that I can understand where your Invocation is being created. Once the containers have established that the invocation does contain a security context, they set it on the thread local so that the JACC PolicyContext get Subject call always takes care of the RunAsIdentity that came into the specific container.	0	0	https://issues.jboss.org/browse/JBAS-4317
1210	JBAS-4288	build -Dtest=org.jboss.test.jmx.test.RMIAdaptorUnitTestCase one-test Calling ..\tools\bin\ant.bat -Dtest org.jboss.test.jmx.test.RMIAdaptorUnitTestCase one-test Buildfile: build.xml one-test: [delete] Deleting: G:\cygwin\home\Magesh\jbossas\branches\Branch_4_0\testsuite\output\log\test.log [junit] Running org.jboss.test.jmx.test.RMIAdaptorUnitTestCase [junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 4.234 sec BUILD SUCCESSFUL Total time: 8 seconds	1	1	https://issues.jboss.org/browse/JBAS-4288
1211	JBAS-4271	It passes locally. Dimitris, can you test it out locally?	0	0	https://issues.jboss.org/browse/JBAS-4271

#	Requirement	Questions	Answered	Successfully Answered	URL
1212	JBAS-4271	Fails on dev24, pa-risc, HPUX, HPUX JVM, 4.2.0GA Alpha5: SecurityException; nested exception is: javax.security.auth.login.FailedLoginException: Supplied Credential did not match existing credential for null java.rmi.AccessException: SecurityException; nested exception is: javax.security.auth.login.FailedLoginException: Supplied Credential did not match existing credential for null at org.jboss.ejb.plugins.LogInterceptor.handleException(LogInterceptor.java:388) at org.jboss.ejb.plugins.LogInterceptor.invokeHome(LogInterceptor.java:136) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invokeHome(ProxyFactoryFinderInterceptor.java:107) at org.jboss.ejb.SessionContainer.internalInvokeHome(SessionContainer.java:637) at org.jboss.ejb.Container.invoke(Container.java:981) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.server.Invocation.invoke(Invocation.java:86) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) at org.jboss.invocation.pooled.server.PooledInvoker$MBeanServerAction.invoke(PooledInvoker.java:896) at org.jboss.invocation.pooled.server.PooledInvoker.invoke(PooledInvoker.java:383) at org.jboss.invocation.pooled.server.ServerThread.processInvocation(ServerThread.java:233) at org.jboss.invocation.pooled.server.ServerThread.dorun(ServerThread.java:275) at org.jboss.invocation.pooled.server.ServerThread.run(ServerThread.java:156) Caused by: javax.security.auth.login.FailedLoginException: Supplied Credential did not match existing credential for null at org.jboss.security.auth.spi.BaseCertLoginModule.login(BaseCertLoginModule.java:218) at javax.security.auth.login.LoginContext.invoke(LoginContext.java:769) at javax.security.auth.login.LoginContext.access$000(LoginContext.java:186) at javax.security.auth.login.LoginContext$4.run(LoginContext.java:683) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.login.LoginContext.invokePriv(LoginContext.java:680) at javax.security.auth.login.LoginContext.login(LoginContext.java:579) at org.jboss.security.plugins.JaasSecurityManager.defaultLogin(JaasSecurityManager.java:601) at org.jboss.security.plugins.JaasSecurityManager.authenticate(JaasSecurityManager.java:535) at org.jboss.security.plugins.JaasSecurityManager.isValid(JaasSecurityManager.java:344) at org.jboss.ejb.plugins.SecurityInterceptor.checkSecurityAssociation(SecurityInterceptor.java:211) at org.jboss.ejb.plugins.SecurityInterceptor.invoke(SecurityInterceptor.java:158) at org.jboss.ejb.plugins.SSLSessionInterceptor.invokeHome(SSLSessionInterceptor.java:53) at org.jboss.ejb.plugins.LogInterceptor.invokeHome(LogInterceptor.java:132) ... 16 more	1	1	https://issues.jboss.org/browse/JBAS-4271
1213	JBAS-4271	Fails on x86_64, RHEL4, BEA JVM testClientCertSSLAccess, Error SecurityException; nested exception is: javax.security.auth.login.FailedLoginException: Supplied Credential did not match existing credential for null java.rmi.AccessException: SecurityException; nested exception is: javax.security.auth.login.FailedLoginException: Supplied Credential did not match existing credential for null at org.jboss.ejb.plugins.LogInterceptor.handleException(LogInterceptor.java:388) at org.jboss.ejb.plugins.LogInterceptor.invokeHome(LogInterceptor.java:136) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invokeHome(ProxyFactoryFinderInterceptor.java:107) at org.jboss.ejb.SessionContainer.internalInvokeHome(SessionContainer.java:637) at org.jboss.ejb.Container.invoke(Container.java:981) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.server.Invocation.invoke(Invocation.java:86) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) at org.jboss.invocation.pooled.server.PooledInvoker$MBeanServerAction.invoke(PooledInvoker.java:896) at org.jboss.invocation.pooled.server.PooledInvoker.invoke(PooledInvoker.java:383) at org.jboss.invocation.pooled.server.ServerThread.processInvocation(ServerThread.java:233) at org.jboss.invocation.pooled.server.ServerThread.dorun(ServerThread.java:275) at org.jboss.invocation.pooled.server.ServerThread.run(ServerThread.java:156) Caused by: javax.security.auth.login.FailedLoginException: Supplied Credential did not match existing credential for null at org.jboss.security.auth.spi.BaseCertLoginModule.login(BaseCertLoginModule.java:218) at javax.security.auth.login.LoginContext.invoke(LoginContext.java:769) at javax.security.auth.login.LoginContext.access$000(LoginContext.java:186) at javax.security.auth.login.LoginContext$4.run(LoginContext.java:683) at javax.security.auth.login.LoginContext.invokePriv(LoginContext.java:680) at javax.security.auth.login.LoginContext.login(LoginContext.java:579) at org.jboss.security.plugins.JaasSecurityManager.defaultLogin(JaasSecurityManager.java:601) at org.jboss.security.plugins.JaasSecurityManager.authenticate(JaasSecurityManager.java:535) at org.jboss.security.plugins.JaasSecurityManager.isValid(JaasSecurityManager.java:344) at org.jboss.ejb.plugins.SecurityInterceptor.checkSecurityAssociation(SecurityInterceptor.java:211) at org.jboss.ejb.plugins.SecurityInterceptor.invoke(SecurityInterceptor.java:158) at org.jboss.ejb.plugins.SSLSessionInterceptor.invokeHome(SSLSessionInterceptor.java:53) at org.jboss.ejb.plugins.LogInterceptor.invokeHome(LogInterceptor.java:132) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invokeHome(ProxyFactoryFinderInterceptor.java:107) at org.jboss.ejb.SessionContainer.internalInvokeHome(SessionContainer.java:637) at org.jboss.ejb.Container.invoke(Container.java:981) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94)	1	1	https://issues.jboss.org/browse/JBAS-4271
1214	JBAS-4270	This test passes on both jrockit as well as Sun JDK. From cruisecontrol, I see that the test passes for the "tests-security-basic-unit". This test is failing when it is run as part of some other target. We should just be testing it as part of "security basic unit tests". Dimitris, can you figure out which target is this test running that is failing?	1	1	https://issues.jboss.org/browse/JBAS-4270

#	Requirement	Questions	Answered	Successfully Answered	URL
1215	JBAS-4260	here are the proposed changes Index: jdbc2/PersistenceManager.java ==========…========== --- jdbc2/PersistenceManager.java (revision 60904) +++ jdbc2/PersistenceManager.java (working copy) @@ -241,8 +241,6 @@ synchronized protected void createSchema() throws JMSException { - TransactionManagerStrategy tms = new TransactionManagerStrategy(); - tms.startTX(); Connection c = null; PreparedStatement stmt = null; boolean threadWasInterrupted = Thread.interrupted(); @@ -351,7 +349,6 @@ } catch (SQLException e) { - tms.setRollbackOnly(); throw new SpyJMSException("Could not get a connection for jdbc2 table construction ", e); } finally @@ -374,7 +371,6 @@ { } c = null; - tms.endTX(); // Restore the interrupted state of the thread if (threadWasInterrupted) @@ -439,9 +435,7 @@ catch (SQLException e) { tms.setRollbackOnly(); - throw new SpyJMSException(- "Could not resolve uncommited transactions. Message recovery may not be accurate", - e); + throw new SpyJMSException("Could not resolve uncommited transactions. Message recovery may not be accurate", e); } finally { Index: jdbc/JDBCStateManager.java ==========…========== --- jdbc/JDBCStateManager.java (revision 60904) +++ jdbc/JDBCStateManager.java (working copy) @@ -412,13 +412,15 @@ if (createTables) { - JDBCSession session = new JDBCSession(); + Connection connection = null; + connection = dataSource.getConnection(); try { PreparedStatement statement; try { - statement = session.prepareStatement(CREATE_USER_TABLE); + statement = connection.prepareStatement(CREATE_USER_TABLE); statement.executeUpdate(); } catch (SQLException ignored) @@ -427,7 +429,7 @@ } try { - statement = session.prepareStatement(CREATE_ROLE_TABLE); + statement = connection.prepareStatement(CREATE_ROLE_TABLE); statement.executeUpdate(); } catch (SQLException ignored) @@ -436,7 +438,7 @@ } try { - statement = session.prepareStatement(CREATE_SUBSCRIPTION_TABLE); + statement = connection.prepareStatement(CREATE_SUBSCRIPTION_TABLE); statement.executeUpdate(); } catch (SQLException ignored) @@ -451,7 +453,7 @@ { nextQry = (String) iter.next(); - statement = session.prepareStatement(nextQry); + statement = connection.prepareStatement(nextQry); statement.execute(); } catch (SQLException ignored) @@ -462,7 +464,13 @@ } finally { - session.close(); + try { + if(connection != null) + connection.close(); + }catch(SQLException sqle){ + log.trace("Error when closing connection " +sqle); } } } Note that these changes should end up in MS SQL specific subclasses, namely MSSQLPersistenceManager.java and MSSQLJDBCStateManager.java	1	1	https://issues.jboss.org/browse/JBAS-4260
1216	JBAS-4245	Fails on dev24, pa-risc, HPUX, HPUX JVM, 4.2.0GA Alpha5: javax.transaction.xa.XAException at com.arjuna.ats.internal.jta.transaction.arjunacore.jca.XATerminatorImple.commit(XATerminatorImple.java:99) at org.jboss.test.jca.inflow.TestResourceAdapterTxInflow.basicTest(TestResourceAdapterTxInflow.java:93) at org.jboss.test.jca.inflow.TestResourceAdapterTxInflow.run(TestResourceAdapterTxInflow.java:55) at org.jboss.test.jca.inflow.TestResourceAdapter.testTxInflow(TestResourceAdapter.java:78) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.server.Invocation.invoke(Invocation.java:86) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) at org.jboss.jmx.connector.invoker.InvokerAdaptorService.invoke(InvokerAdaptorService.java:266) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:133) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:142) at org.jboss.jmx.connector.invoker.SerializableInterceptor.invoke(SerializableInterceptor.java:74) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) at org.jboss.invocation.jrmp.server.JRMPProxyFactory.invoke(JRMPProxyFactory.java:179) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.server.Invocation.invoke(Invocation.java:86) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:818) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:419) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:294) at sun.rmi.transport.Transport$1.run(Transport.java:153) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:149) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:466) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:707) at java.lang.Thread.run(Thread.java:595)	1	1	https://issues.jboss.org/browse/JBAS-4245

#	Requirement	Questions	Answered	Successfully Answered	URL
1217	JBAS-4244	Account:0 java.rmi.ServerException: EJBException:; nested exception is: javax.ejb.EJBException: Could not transfer 1 from Account:1 to Account:0 at org.jboss.ejb.plugins.LogInterceptor.handleException(LogInterceptor.java:365) at org.jboss.ejb.plugins.LogInterceptor.invoke(LogInterceptor.java:209) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invoke(ProxyFactoryFinderInterceptor.java:138) at org.jboss.ejb.SessionContainer.internalInvoke(SessionContainer.java:648) at org.jboss.ejb.Container.invoke(Container.java:960) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.server.Invocation.invoke(Invocation.java:86) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) at org.jboss.invocation.unified.server.UnifiedInvoker.invoke(UnifiedInvoker.java:231) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.server.Invocation.invoke(Invocation.java:86) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) Caused by: javax.ejb.EJBException: Could not transfer 1 from Account:1 to Account:0 at org.jboss.test.jca.bank.ejb.TellerBean.transfer(TellerBean.java:151) at org.jboss.invocation.Invocation.performCall(Invocation.java:359) at org.jboss.ejb.StatelessSessionContainer$ContainerInterceptor.invoke(StatelessSessionContainer.java:237) at org.jboss.resource.connectionmanager.CachedConnectionInterceptor.invoke(CachedConnectionInterceptor.java:158) at org.jboss.ejb.plugins.StatelessSessionInstanceInterceptor.invoke(StatelessSessionInstanceInterceptor.java:169) at org.jboss.ejb.plugins.CallValidationInterceptor.invoke(CallValidationInterceptor.java:63) at org.jboss.ejb.plugins.AbstractTxInterceptor.invokeNext(AbstractTxInterceptor.java:121) at org.jboss.ejb.plugins.TxInterceptorCMT.runWithTransactions(TxInterceptorCMT.java:350) at org.jboss.ejb.plugins.TxInterceptorCMT.invoke(TxInterceptorCMT.java:181) at org.jboss.ejb.plugins.SecurityInterceptor.invoke(SecurityInterceptor.java:168) at org.jboss.ejb.plugins.LogInterceptor.invoke(LogInterceptor.java:205) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invoke(ProxyFactoryFinderInterceptor.java:138) at org.jboss.ejb.SessionContainer.internalInvoke(SessionContainer.java:648) at org.jboss.ejb.Container.invoke(Container.java:960) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.server.Invocation.invoke(Invocation.java:86) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) at org.jboss.invocation.unified.server.UnifiedInvoker.invoke(UnifiedInvoker.java:231) at org.jboss.remoting.MicroRemoteClientInvoker.invoke(MicroRemoteClientInvoker.java:163) at org.jboss.remoting.Client.invoke(Client.java:612) at org.jboss.remoting.Client.invoke(Client.java:604) at org.jboss.invocation.unified.interfaces.UnifiedInvokerProxy.invoke(UnifiedInvokerProxy.java:175) at org.jboss.invocation.InvokerInterceptor.invokeInvoker(InvokerInterceptor.java:365) at org.jboss.invocation.InvokerInterceptor.invoke(InvokerInterceptor.java:197) at org.jboss.proxy.TransactionInterceptor.invoke(TransactionInterceptor.java:61) at org.jboss.proxy.SecurityInterceptor.invoke(SecurityInterceptor.java:70) at	1	1	https://issues.jboss.org/browse/JBAS-4244
1218	JBAS-4236	Stan, is the update worth it?	1	1	https://issues.jboss.org/browse/JBAS-4236
1219	JBAS-4215	Is this a duplicate of JBAS-5153 ?	0	0	https://issues.jboss.org/browse/JBAS-4215
1220	JBAS-4211	Updated version. Works well! But the solution is a little bit dirty. Does anyone have got an idea how to make it "clean"? The output of the memory checks is in appropriate reports/*.xml file.	1	1	https://issues.jboss.org/browse/JBAS-4211
1221	JBAS-4211	How dirty it is?	0	0	https://issues.jboss.org/browse/JBAS-4211
1222	JBAS-4211	So, it takes the whole junit ant task, changes some classes and put one more class to it. What do you think about that?	1	1	https://issues.jboss.org/browse/JBAS-4211
1223	JBAS-4211	Martin, this issue appears outdated and Ryan no longer here. Does QA plan to implement this feature? Do you need input from dev?	1	0	https://issues.jboss.org/browse/JBAS-4211
1224	JBAS-4209	Shall I supply a patch? If so what should the flag be called? Regards Felix	1	0	https://issues.jboss.org/browse/JBAS-4209
1225	JBAS-4180	Summary of chat with Brian: - A general refactor of stuff that extends HAServiceMBeanSupport such that HA classes can be POJOs, with JMX behavior an added feature. - Similar fashion to the work you done with JBossCache 2.0 where JMX was decoupled. - Examine all the use cases, see what makes sense. E.g. HASingletonController emits JMX notifications, so is a non-JMX-aware version sensible?	0	0	https://issues.jboss.org/browse/JBAS-4180
1226	JBAS-4175	Could you attach the file to this issue so I can check what I have done wrong ?	1	0	https://issues.jboss.org/browse/JBAS-4175
1227	JBAS-4174	Maybe I see what you are saying. Are you talking about the Sun versions of these libraries? Sun forked these libraries and changed the package names to com.sun.org.apache. There is no way we could safely use the standard versions of these libraries. Stan	1	0	https://issues.jboss.org/browse/JBAS-4174
1228	JBAS-4165	This already exists; was added two years ago. Even uses jnp.discoveryTTL as the name. But it, along with several other configurable properties, is not in the documentation. Will open a new JIRA to fix that.	0	0	https://issues.jboss.org/browse/JBAS-4165
1229	JBAS-4152	The last bug I see mentioning a HSQL upgrade was JBAS-1618 which ended up with 1.8 RC9. However, this bug mentions that we were upgrading from 1.8.0.2 to 1.8.0.7. When did 1.8 RC 9 get upgraded and where is it tracked/documented?	0	0	https://issues.jboss.org/browse/JBAS-4152
1230	JBAS-4152	Ok, I don't recall every detail aymore :-) What I did was a <dbindex/> tag in the relationships (as well as for fields) as most DBs don't automatically put indexes on fks (as oposed to what hsql did). Now I put an index on one side of the relation and not on the 2nd one. That 2nd one is now failing, as the test code detects an index, which should not be there. I *guess* that hsqldb is now putting in an index there that does not start with "pk_idx" (those get filtered) which gets detected by the failing test, as it does not expect it. Is there a "preconfigured" jboss as with the new hsqldb and the test(suite) soewhere available?	1	1	https://issues.jboss.org/browse/JBAS-4152
1231	JBAS-4151	Is there *any* integration test for that?	1	1	https://issues.jboss.org/browse/JBAS-4151
1232	JBAS-4148	Can you help us with this?	0	0	https://issues.jboss.org/browse/JBAS-4148
1233	JBAS-4148	Dimitris, who could d this change?	0	0	https://issues.jboss.org/browse/JBAS-4148
1234	JBAS-4148	Can you please commit that to trunk, as well?	1	1	https://issues.jboss.org/browse/JBAS-4148
1235	JBAS-4127	Weston will you include this in CR1 or defer it?	1	1	https://issues.jboss.org/browse/JBAS-4127
1236	JBAS-4100	I've synched the dtds in Branch_4_2 with HEAD, introducing new 4_2 ones in both branches. I did some minor changes where one of them was outdated. The xml schemas were ok, so I just copied over ejb-jar_3_0.xsd & web-app_2_5.xsd from HEAD to Branch_4_2. One discrepency is security-config_4_1.xsd in Branch_4_2 and HEAD. They contain different packages names for the jbxb bindings: org.jboss.security.auth.* (Branch_4_2) org.jboss.security.config (HEAD) Are we supposed to synch them?	0	0	https://issues.jboss.org/browse/JBAS-4100
1237	JBAS-4085	Is this possibly another manifestation of JBAS-4060? Turned out that's what JBAS-4064 was -- similar symptom of getting a 503. Key thing to see is if an earlier test that undeployed the same war shows the JBAS-4060 NPE; if so the first request following deploying the war again gets a 503.	1	1	https://issues.jboss.org/browse/JBAS-4085
1238	JBAS-4085	Brian, can you take this?	1	1	https://issues.jboss.org/browse/JBAS-4085

#	Requirement	Questions	Answered	Successfully Answered	URL
1239	JBAS-4085	It fails again :-(This and the 2 variations of it with JACC and/or security manager enabled: testUnauthorizedAccess Error Expected reply code:403, actual=400 java.io.IOException: Expected reply code:403, actual=400 at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:137) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:108) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:96) at org.jboss.test.web.test.WebIntegrationUnitTestCase.testUnauthorizedAccess(WebIntegrationUnitTestCase.java:445) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23)	1	1	https://issues.jboss.org/browse/JBAS-4085
1240	JBAS-4082	Scott, can you look into this? I've fixed all tests except for ExternalContextUnitTestCase Some change in the packaging causes a ClassCastException here: static Context createProxyContext(Context ctx) { ClassLoader loader = Thread.currentThread().getContextClassLoader(); Class[] interfaces = ctx.getClass().getInterfaces(); InvocationHandler handler = new CachedContext(ctx); Context proxyCtx = (Context) Proxy.newProxyInstance(loader, interfaces, handler); <=== here return proxyCtx; } 20:37:09,907 WARN [ServiceController] Problem starting service jboss.test:servi ce=ExternalContext,jndiName=external/SubMockNamingContext java.lang.ClassCastException: $Proxy67 at org.jboss.naming.ExternalContext$CachedContext.createProxyContext(Ext ernalContext.java:527) at org.jboss.naming.ExternalContext.rebind(ExternalContext.java:307) at org.jboss.naming.ExternalContext.startService(ExternalContext.java:23 0) at org.jboss.system.ServiceMBeanSupport.jbossInternalStart(ServiceMBeanS upport.java:289) at org.jboss.system.ServiceMBeanSupport.jbossInternalLifecycle(ServiceMB eanSupport.java:245) If you run the test a second time against the same server it passes.	1	1	https://issues.jboss.org/browse/JBAS-4082
1241	JBAS-4081	Anil, can you have a look? Only fails with JRockit.	1	1	https://issues.jboss.org/browse/JBAS-4081
1242	JBAS-4076	???? Why are these test using org.jboss.cache.transaction.DummyBaseTransactionManager from JBoss Cache? These tests use: TransactionManagerLocator.getInstance().locate() inside the test client. i.e. they expect to setup a local TM, the old TxManager by default. So it looks to me like the JBossCache tests are doing naughty things?	1	1	https://issues.jboss.org/browse/JBAS-4076
1243	JBAS-4073	8 tests fail with: Failure WEBAPP classloader unregistered junit.framework.AssertionFailedError: WEBAPP classloader unregistered at org.jboss.test.classloader.leak.test.ClassloaderLeakUnitTestCase.warTest(ClassloaderLeakUnitTestCase.java:222) at org.jboss.test.classloader.leak.test.ClassloaderLeakUnitTestCase.testSimpleWar(ClassloaderLeakUnitTestCase.java:144) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) Both sun & jrockit jdks.	0	0	https://issues.jboss.org/browse/JBAS-4073
1244	JBAS-4066	Somewhere in between 1.5.0_04 and 1.5.0_09, the exception thrown when completing class loaders try to load the same class has changed from a ClassCircularityError to the more general LinkageError: 13:42:52,920 INFO [RecursiveCCETests$MyUCL] Unexpected error during load of:org.jboss.test.classloader.circularity.support.HARMIServerImpl_Stub java.lang.LinkageError: org/jboss/test/classloader/circularity/support/HARMIServer at java.lang.ClassLoader.defineClass1(Native Method) at java.lang.ClassLoader.defineClass(ClassLoader.java:620) at java.security.SecureClassLoader.defineClass(SecureClassLoader.java:124) at java.net.URLClassLoader.defineClass(URLClassLoader.java:260) at java.net.URLClassLoader.access$100(URLClassLoader.java:56) at java.net.URLClassLoader$1.run(URLClassLoader.java:195) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:188) at org.jboss.mx.loading.RepositoryClassLoader.findClassLocally(RepositoryClassLoader.java:672) at org.jboss.mx.loading.RepositoryClassLoader.findClass(RepositoryClassLoader.java:652) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at org.jboss.mx.loading.RepositoryClassLoader.loadClassLocally(RepositoryClassLoader.java:190) at org.jboss.mx.loading.ClassLoadingTask$ThreadTask.run(ClassLoadingTask.java:131) at org.jboss.mx.loading.LoadMgr3.nextTask(LoadMgr3.java:399) at org.jboss.test.classloader.circularity.test.RecursiveCCETests$MyUCL.loadClass(RecursiveCCETests.java:199) The logic for retrying the load after competing class loading tasks have been retried depends on the ClassCircularityError being thrown. There are several LinkageError subclasses, and only the ClassCircularityError would seem to have retry semantics. If we are to change to catch a LinkageError, additional logic to avoid repeated errors due to the alternate, non-retryable forms would be needed.	0	0	https://issues.jboss.org/browse/JBAS-4066
1245	JBAS-4064	Stan can you have a look? Could this be due to the http-client update also?	1	1	https://issues.jboss.org/browse/JBAS-4064
1246	JBAS-4042	The timers table name (and optionally schema) can be configured through the deploy/ejb-deployer.xml descriptor: ... <!-- A persistence policy that persists timers to a database. The 2 supported db persistence plugins are: org.jboss.ejb.txtimer.GeneralPurposeDatabasePersistencePlugin org.jboss.ejb.txtimer.OracleDatabasePersistencePlugin The table name defaults to "TIMERS". It can be overriden using the 'TimersTable' attribute if the persistence plugin supports it. When overriding the timers table, an optional schema can be specified using the syntax [schema.]table --> <mbean code="org.jboss.ejb.txtimer.DatabasePersistencePolicy" name="jboss.ejb:service=EJBTimerService,persistencePolicy=database"> <!-- DataSourceBinding ObjectName --> <depends optional-attribute-name="DataSource">jboss.jca:service=DataSourceBinding,name=DefaultDS</depends> <!-- The plugin that handles database persistence --> <attribute name="DatabasePersistencePlugin">org.jboss.ejb.txtimer.GeneralPurposeDatabasePersistencePlugin</attribute> <!-- The timers table name --> <attribute name="TimersTable">TIMERS</attribute> </mbean> ...	0	0	https://issues.jboss.org/browse/JBAS-4042

#	Requirement	Questions	Answered	Successfully Answered	URL
1247	JBAS-4031	One test fails for me (winXp), but seems to pass in cruisecontrol? <testcase classname="org.jboss.test.jacc.test.WebIntegrationUnitTestCase" name="testBadEarRedeploy" time="0.932"> <error message="Expected reply code:200, actual=503" type="java.io.IOException">java.io.IOException: Expected reply code:200, actual=503 at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:137) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:108) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:96) at org.jboss.test.web.test.WebIntegrationUnitTestCase.testBadEarRedeploy(WebIntegrationUnitTestCase.java:422) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) </error>	1	1	https://issues.jboss.org/browse/JBAS-4031
1248	JBAS-4031	The test actually passes with jrockit, but fails with sun jdk5, both on winxp and linux, so I'm re-opening the case. testBadEarRedeploy Error Expected reply code:200, actual=503 java.io.IOException: Expected reply code:200, actual=503 at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:137) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:108) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:96) at org.jboss.test.web.test.WebIntegrationUnitTestCase.testBadEarRedeploy(WebIntegrationUnitTestCase.java:422) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23)	1	1	https://issues.jboss.org/browse/JBAS-4031
1249	JBAS-4025	I only found one other instance of this problem. There is a server.xml file for the tests-webservice-ssl target. However, no tests actually get executed when you run this target. According to build.xml in testsuite, here are the tests it is trying to pull in: <patternset id="webservice-ssl.includes"> <include name="org/jboss/test/webservice/secure/**/*TestCase.class"/> <exclude name="org/jboss/test/webservice/secure/noenv/**/*TestCase.class"/> </patternset> So I guess the tests-webservice-ssl target is obsolete? Can someone from the Web Services team comment? As far as JBoss 5 goes, there are many problems with that right now and I can't run the test. I'm going to open a separate jira task and link to here so it doesn't get lost.	1	1	https://issues.jboss.org/browse/JBAS-4025
1250	JBAS-4025	Thomas. Can someone on your team clean up the tests-webservice-ssl test if it is no longer being used?	1	1	https://issues.jboss.org/browse/JBAS-4025
1251	JBAS-4024	Yes fine, so we can basically drop the test since it doesn't make sense any more. Is there a link to a forum discussion on this? Will you do the code cleanup and track it with another JIRA issue? Thanks.	0	0	https://issues.jboss.org/browse/JBAS-4024
1252	JBAS-4023	Same tests fails here: org.jboss.test.web.test.WebIntegrationUnitTestCase testENCServletViaInvoker Error Expected reply code:200, actual=500 java.io.IOException: Expected reply code:200, actual=500 at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:137) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:108) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:96) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:82) at org.jboss.test.web.test.WebIntegrationUnitTestCase.testENCServletViaInvoker(WebIntegrationUnitTestCase.java:98) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23)	1	1	https://issues.jboss.org/browse/JBAS-4023
1253	JBAS-4023	And here: org.jboss.test.web.test.UserInRoleUnitTestCase(JACC+SecurityMgr) testConflictingUserInRoleJaccVersion Error Expected reply code:200, actual=500 java.io.IOException: Expected reply code:200, actual=500 at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:137) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:108) at org.jboss.test.util.web.HttpUtils.accessURL(HttpUtils.java:96) at org.jboss.test.web.test.UserInRoleUnitTestCase.testConflictingUserInRoleJaccVersion(UserInRoleUnitTestCase.java:144) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23)	1	1	https://issues.jboss.org/browse/JBAS-4023
1254	JBAS-4019	I can't understand why it is not included by default since the law has changed to not restrict crypto (AFAIK). Could someone tell Sun? ;)	0	0	https://issues.jboss.org/browse/JBAS-4019
1255	JBAS-4011	Thanks. If you want to be really picky about this you could throw an InvalidStateException if you attempt to set the attribute twice. And while there, I'm seeing also those failures somewhat related to try to set the tx timeout (but on a DummyTxMgr): JBAS-4079. Is this supported in jbossTS ?	1	1	https://issues.jboss.org/browse/JBAS-4011
1256	JBAS-3993	Did you actually checked your changes work? On Branch_4_0: X:\cvs\jboss-public\jboss-4.0\testsuite>build one-test -Dtest=org.jboss.test.uti l.test.ThreadPoolRunnableUnitTestCase Calling ..\tools\bin\ant.bat one-test -Dtest org.jboss.test.util.test.ThreadPool RunnableUnitTestCase Buildfile: build.xml one-test: [junit] Running org.jboss.test.util.test.ThreadPoolRunnableUnitTestCase BUILD FAILED X:\cvs\jboss-public\jboss-4.0\testsuite\build.xml:3046: X:\cvs\jboss-public\jbos s-4.0\testsuite\${jboss.dist}\client not found. Total time: 3 seconds Press any key to continue . . .	1	1	https://issues.jboss.org/browse/JBAS-3993
1257	JBAS-3993	Did you tried fresh SVN checkout? I tried it right now: [mvecera@dev01 testsuite]$ ./build.sh one-test -Dtest=org.jboss.test.util.test.ThreadPoolRunnableUnitTestCase Searching for build.xml ... Buildfile: /qa/home/mvecera/jboss/JBAS/3993/Branch_4_0/testsuite/build.xml one-test: [mkdir] Created dir: /qa/home/mvecera/jboss/JBAS/3993/Branch_4_0/testsuite/output/log [junit] Running org.jboss.test.util.test.ThreadPoolRunnableUnitTestCase [junit] Tests run: 9, Failures: 0, Errors: 0, Time elapsed: 25.158 sec BUILD SUCCESSFUL Total time: 30 seconds	1	1	https://issues.jboss.org/browse/JBAS-3993
1258	JBAS-3983	It's for 4.0.4.GA, not 4.0.5.GA. Can you re-test?	1	1	https://issues.jboss.org/browse/JBAS-3983
1259	JBAS-3949	Is it possible to have the procedure to upgrade to log4j1.2.14 and commons logging 1.1 for jboss 4.0.5GA ?	0	0	https://issues.jboss.org/browse/JBAS-3949
1260	JBAS-3942	what do you guys think about a "scripts.conf" which is similar to run.conf but centralized for twiddle, shutdown and any scripts that follow.	1	0	https://issues.jboss.org/browse/JBAS-3942
1261	JBAS-3939	Should we try CR8?	1	1	https://issues.jboss.org/browse/JBAS-3939
1262	JBAS-3928	Is this going to happen for AS 5 GA? I'm now having some issues with jmx-invoker-service.xml and wonder if some changes has been made.	1	0	https://issues.jboss.org/browse/JBAS-3928
1263	JBAS-3914	Clive, Fernando mentioned that you and Jon will be working on console related JIRAs. Can you or Jon look into this? Thanks	1	0	https://issues.jboss.org/browse/JBAS-3914
1264	JBAS-3891	I'm updating the thirdparty build to point to JBC 1.4.1.GA. There are 2 binaries in jbc lib. jboss-cache.jar jboss-cache-jdk50.jar What's the difference?	1	1	https://issues.jboss.org/browse/JBAS-3891
1265	JBAS-3888	EJB3 should be refactored to support ejb3 effectively on both platforms (jbossAS 4.2 / 5.0). Link to the corresponding task.	1	1	https://issues.jboss.org/browse/JBAS-3888

#	Requirement	Questions	Answered	Successfully Answered	URL
1266	JBAS-3886	Can we also use jboss-test 1.0.1.GA ? The one we are using does not even have a tag and there are multiple different code versions named "CR1".	1	1	https://issues.jboss.org/browse/JBAS-3886
1267	JBAS-3858	OK, couple things need to happen here: 1. We should deprecate the org.jboss.web.tc5.session.AopMarker (and InstanceOfAopMarker). They are there for JDK1.4 annotation and annoc. And PojOCache 2.0 don't support JDK1.4 anymore. 2. Now that we use JDK5.0, I'd also propose we just use the annotations in PojoCache, e.g., org.jboss.cache.pojo.PojoCacheable (but please also be aware of that I may need to change the annotation name before this Friday, see http://www.jboss.com/index.html?module=bb&op=viewtopic&p=3986033#3986033). 3. Originally, we need Oberserver and Obervable pattern to subscribe to POJO field change event. We won't need that anymore since now PojoCache 2.0 provides PojoCacheListener event that can emit field changes. tc6 can subscribe to the event directly. 4. As a result, we would need tc6-cluster.aop to only consist of jboss-aop.xml. That jboss-aop.xml will be a copy of pojocache-aop.xml in the distro (similar to *-service.xml convention). Another question will then be how can a user run "aopc" to pre-process the POJO? He will need: 1. put the classes under all/lib in the classpath. Or maybe just jboss-cache.jar, pojocache.jar, and jboss-aop-jdk50.jar is enough, I haven't verified. 2. put tc6-cluster.aop (or equivalently, pojocache-aop.xml) in the classpath 3. User class lib (can also include additional 3rd party one as well) to generate the byte code weaving. Maybe we can provide a script to do that.	1	1	https://issues.jboss.org/browse/JBAS-3858
1268	JBAS-3846	This is a 30 second fix. Could someone please take care of it?	1	1	https://issues.jboss.org/browse/JBAS-3846
1269	JBAS-3834	In JBoss 5.0.1 GA, when a plain String with some national characters is the output of a method in the @Management interface of a @Service POJO, then JMX console displays nothing. I mean plain String, no HTML tags in the string. This is even more annoying when the method is a getter for an attribute. Then the web page for the @Service bean shown nothing (HMTL source->nothing) in the JMX console. Try a @Service with method: String getLocalizedString() { return ?š??žýáié; } Also, the java Date.toString() may result sometimes in the non-english characters, and then then again the web page for the @Service bean is nothing.	0	0	https://issues.jboss.org/browse/JBAS-3834
1270	JBAS-3802	Can this be tagged? So i can build and upload to the repository. The build.xml has to be updated to build the zip and the src distributions. The tag can be updated later when releasing on sourceforge.	1	1	https://issues.jboss.org/browse/JBAS-3802
1271	JBAS-3769	Weston, can you please re-apply this in Branch_4_0, and marked done also for 4.0.5.SP?	1	1	https://issues.jboss.org/browse/JBAS-3769
1272	JBAS-3765	Important: Consider the related EJB3 issue as part of it. Be aware Carlo wants to extract EJB3 from the AS code base and use the same code for 4.2.x and 5.x, with pluggable adapters to deal with API differences in the AS integration. Stuff has to be backward compatible. So we'll need to think about that.	1	1	https://issues.jboss.org/browse/JBAS-3765
1273	JBAS-3726	Are you doing this for GA?	1	1	https://issues.jboss.org/browse/JBAS-3726
1274	JBAS-3726	Yes. But unless I can quickly arrange to work on someone's RHEL/Fedora box it might be a no-op on non-Windows environments. What is the absolute deadline on this?	1	1	https://issues.jboss.org/browse/JBAS-3726
1275	JBAS-3726	Wouldn't the next CC build pick it up? We'd like to have work done by Friday.	1	1	https://issues.jboss.org/browse/JBAS-3726
1276	JBAS-3715	Remember the "Unlimited Strength Crypto" that needs to be downloaded for the VM? We did for Sun VM?	1	1	https://issues.jboss.org/browse/JBAS-3715
1277	JBAS-3715	http://java.sun.com/products/jce/index-14.html#UnlimitedDownload This is for JDK4. Should work on JRockit also as per http://dev.project.net/trac/pnet/wiki/InstallationGuides (Q. Where do I get the unlimited strength encryption JCE Policy files for JRockit? A. Use the Sun ones for the appropriate version of JRockit (e.g. 1.4.2 policy files))	1	1	https://issues.jboss.org/browse/JBAS-3715
1278	JBAS-3704	The description is not very clear to me... Lets say there would be ant task <server:addresses prop="jboss.server.iface." env="MYTESTIP_" ifaces="2">. This task would work as follows: x = 1; while (MYTESTIP_x <> "") { jboss.server.iface.(x - 1) = MYTESTIP_x; x++; } while (x <= ${ifaces}) { jboss.server.iface.(x - 10) = fetch_next_interface_from_NetworkInterface(); x++; } Do you imagine it like this?	0	0	https://issues.jboss.org/browse/JBAS-3704
1279	JBAS-3683	Thanks! Why you changed some counters to Gauge32 and left 2 of them as INTEGER?	1	1	https://issues.jboss.org/browse/JBAS-3683
1280	JBAS-3676	What is the SVN branch for this? I need to check if my changes are going to make it in 4.0.5	1	1	https://issues.jboss.org/browse/JBAS-3676
1281	JBAS-3659	Brian, we've started getting OOMEs in the testsuite runs, could this be related?	1	1	https://issues.jboss.org/browse/JBAS-3659
1282	JBAS-3658	Per forum: Seems to me should be handled at the invoker level (i.e. UnifiedInvoker) and use a service like the org.jboss.web.WebService we currently have. The WebService actually looks fairly good to me to what is needed (i.e. downloading classes). Maybe add another configuration attribute for what type of class files can be downloaded (i.e. "com.acme.domain.*") and should be sufficient? Also, where is the client code for downloading classes from the WebService?	1	1	https://issues.jboss.org/browse/JBAS-3658
1283	JBAS-3658	Tom we only have a few days to complete this. What's the status?	1	1	https://issues.jboss.org/browse/JBAS-3658
1284	JBAS-3658	Ok, I"ll defer this task for a next release. Can you fix the test?	1	1	https://issues.jboss.org/browse/JBAS-3658
1285	JBAS-3656	Brian, do we need to re-apply this in Branch_4_0, and tag fixed for 4.0.5.SP1 instead of 4.0.5.GA?	1	1	https://issues.jboss.org/browse/JBAS-3656
1286	JBAS-3652	Just the text file should be added? what about doc and pdf? they are all the same. Should this go in HEAD as well?	1	1	https://issues.jboss.org/browse/JBAS-3652
1287	JBAS-3651	Ryan, could you assign this to somebody that would record testsuite failures on Solaris10?	1	1	https://issues.jboss.org/browse/JBAS-3651
1288	JBAS-3651	What's the latest update?	1	1	https://issues.jboss.org/browse/JBAS-3651
1289	JBAS-3644	Modified version of compiler works as follows: 1. visiting ASTWhere node sets wherePath=true in all ASTWhere's child paths. 2. visiting ASTOr, ASTNullComparison or ASTIsEmpty causes their child paths to have forceLeftJoin=true 3. When visiting ASTPath compiler checks if wherePath is true and adds join based on forceLeftJoin field. If wherePath is false compiler works as before. 4. If some path has forceLeftJoin set, it cannot be reset to innerjoin Otherwise the compiler works as before. Is there something i'm missing or can this be added to default JBoss? This really boosted performance for us.	0	0	https://issues.jboss.org/browse/JBAS-3644
1290	JBAS-3641	Thanks John. I'll commit this to all brances, but can you please verify it works with the latest jboss-4.0.5.CR1 ?	1	1	https://issues.jboss.org/browse/JBAS-3641
1291	JBAS-3615	Sun don't think it is a bug to be fixed because the result meet the JMX spec 1.2, so could you update the test case to the right asseration to avoid misundertsanding?When we run the JBoss AS testsuite 4.0.5,it is reported errors and failed,we must pay much attention to this test case.	1	1	https://issues.jboss.org/browse/JBAS-3615

#	Requirement	Questions	Answered	Successfully Answered	URL
1292	JBAS-3613	Franscisco, the test shows the same error on my machine (WinXP/SunJDK1.5) even after the jacorb 2.3.0.jboss.patch1 update. I even tried 2.2.4.jboss.patch2 without luck. The error still is: test_getException Error Invalid indirection to offset 1704 java.io.IOException: Invalid indirection to offset 1704 at com.sun.corba.se.impl.io.IIOPInputStream$ActiveRecursionManager.getObject(IIOPInputStream.java:2684) at com.sun.corba.se.impl.io.IIOPInputStream.inputClassFields(IIOPInputStream.java:2218) at com.sun.corba.se.impl.io.IIOPInputStream.inputObject(IIOPInputStream.java:1221) at com.sun.corba.se.impl.io.IIOPInputStream.simpleReadObject(IIOPInputStream.java:400) at com.sun.corba.se.impl.io.ValueHandlerImpl.readValueInternal(ValueHandlerImpl.java:327) at com.sun.corba.se.impl.io.ValueHandlerImpl.readValue(ValueHandlerImpl.java:293) at org.jacorb.util.ValueHandler.readValue(Unknown Source) at org.jacorb.orb.CDRInputStream.read_untyped_value(Unknown Source) at org.jacorb.orb.CDRInputStream.read_typed_value(Unknown Source) at org.jacorb.orb.CDRInputStream.read_value(Unknown Source) at com.sun.corba.se.impl.io.ValueHandlerImpl.read_Array(ValueHandlerImpl.java:756) at com.sun.corba.se.impl.io.ValueHandlerImpl.readValueInternal(ValueHandlerImpl.java:325) at com.sun.corba.se.impl.io.ValueHandlerImpl.readValue(ValueHandlerImpl.java:293) at org.jacorb.util.ValueHandler.readValue(Unknown Source) at org.jacorb.orb.CDRInputStream.read_untyped_value(Unknown Source) at org.jacorb.orb.CDRInputStream.read_typed_value(Unknown Source) at org.jacorb.orb.CDRInputStream.read_value(Unknown Source) at com.sun.corba.se.impl.io.IIOPInputStream.inputObjectField(IIOPInputStream.java:1989) at com.sun.corba.se.impl.io.IIOPInputStream.inputClassFields(IIOPInputStream.java:2213) at com.sun.corba.se.impl.io.IIOPInputStream.inputObject(IIOPInputStream.java:1221) at com.sun.corba.se.impl.io.IIOPInputStream.simpleReadObject(IIOPInputStream.java:400) at com.sun.corba.se.impl.io.ValueHandlerImpl.readValueInternal(ValueHandlerImpl.java:327) at com.sun.corba.se.impl.io.ValueHandlerImpl.readValue(ValueHandlerImpl.java:293) at org.jacorb.util.ValueHandler.readValue(Unknown Source) at org.jacorb.orb.CDRInputStream.read_untyped_value(Unknown Source) at org.jacorb.orb.CDRInputStream.read_typed_value(Unknown Source) at org.jacorb.orb.CDRInputStream.read_value(Unknown Source) at org.jacorb.orb.Any.read_value(Unknown Source) at org.jacorb.orb.CDRInputStream.read_any(Unknown Source) at com.sun.corba.se.impl.javax.rmi.CORBA.Util.readAny(Util.java:401) at javax.rmi.CORBA.Util.readAny(Util.java:92) at org.jboss.iiop.rmi.marshal.CDRStream$ObjectReader.read(CDRStream.java:618) at org.jboss.iiop.rmi.marshal.strategy.StubStrategy.readRetval(StubStrategy.java:226) at org.jboss.proxy.ejb.DynamicIIOPStub.invoke(DynamicIIOPStub.java:126) at org.jboss.test.iiop.interfaces._StatelessSession_Stub.getException(Unknown Source) at org.jboss.test.iiop.test.ParameterPassingStressTestCase.test_getException(ParameterPassingStressTestCase.java:355) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25)	1	1	https://issues.jboss.org/browse/JBAS-3613
1293	JBAS-3613	I don't see the error with JacORB 2.3.0.jboss.patch. JacORB 2.2.4jboss.patch1 gives me the error above, but the error goes away when I switch to JacORB 2.3.0.jboss.patch1. I am running Sun's Java 1.5.0_06-b05 on Linux, though, so the VM might behave differently on Windows XP and on Linux. (Possible, but unlikely.) Dimitris: Are you sure you replaced jacorb.jar in all places (build/output/jboss-x.y/client/jacorb.jar, build/output/jboss-x.y/server/all/lib/jacorb.jar and thirdparty/jacorb/lib/jacorb.jar)?	0	0	https://issues.jboss.org/browse/JBAS-3613
1294	JBAS-3613	What changed?	1	1	https://issues.jboss.org/browse/JBAS-3613
1295	JBAS-3613	I don't know. The only thing I've noticed is that this branch is using JacORB 2.3.0.jboss.patch4, so maybe this patched version fixed the problem. Do you want me to try this with HP JDK?	1	1	https://issues.jboss.org/browse/JBAS-3613
1296	JBAS-3603	I am assuing JRockit 1.4. This works fine for me on Windows. What platform are you testing on?	1	1	https://issues.jboss.org/browse/JBAS-3603
1297	JBAS-3580	Am I duplicating JBAS-3221 ?	0	0	https://issues.jboss.org/browse/JBAS-3580
1298	JBAS-3577	We have 2.0.0.Beta - thanks. Are we closing this issue, or moving to Beta2?	1	1	https://issues.jboss.org/browse/JBAS-3577
1299	JBAS-3546	Why can't you just change your appender for this category?	1	1	https://issues.jboss.org/browse/JBAS-3546
1300	JBAS-3524	Is there a way that we can get this to work in the mean time. Or if you can give some pointers on how to implement it that would fantastic.	1	0	https://issues.jboss.org/browse/JBAS-3524
1301	JBAS-3450	How can I reproduce this?	0	0	https://issues.jboss.org/browse/JBAS-3450
1302	JBAS-3434	Closing this in the JBoss JIRA, it is tracked in the installer and qa jira projects. Creating the Installer is on track and will be released immediately after the .zip distro.	0	0	https://issues.jboss.org/browse/JBAS-3434
1303	JBAS-3415	It's not clear from the description what work is left for RAR deployments. Is the missing part the datasource management view? Can Adrian look at it, since Weston is not with us?	1	1	https://issues.jboss.org/browse/JBAS-3415
1304	JBAS-3400	Is Sohil working on this?	1	1	https://issues.jboss.org/browse/JBAS-3400
1305	JBAS-3385	Patched EntityTable in package org.jboss.ejb.plugins.cmp.jdbc2 that used the following for table cache naming: 1. datasource jndi name (':' will be stripped) 2. ejb name 3. abstract schema name The patch is generated by diff -u	0	0	https://issues.jboss.org/browse/JBAS-3385
1306	JBAS-3343	In which release will this issue be fixed? The fix is very important because no JCA RA can be properly configured except with cracking the rar and overwriting the default props.	1	1	https://issues.jboss.org/browse/JBAS-3343
1307	JBAS-3324	Anil, is this done for JB5? If so, please close it, thanks.	1	0	https://issues.jboss.org/browse/JBAS-3324
1308	JBAS-3306	It seems there is a Version 1.1.4-SNAPSHOT version. Should we try this one?	1	1	https://issues.jboss.org/browse/JBAS-3306
1309	JBAS-3280	Is there a common usage of defining serialVersionUID? Your version count or serialver from JDK? If serialvar - what's the easies way to get all the necessary classes into classpath of serialvar tool? I've seen your post (http://www.jboss.com/index.html?module=bb&op=viewtopic&t=76117&postdays=0&postorder=asc&start=0), is this Ant script - build.properties file - available in cvs?	1	1	https://issues.jboss.org/browse/JBAS-3280
1310	JBAS-3272	Tom, will you have a remoting 2 version in the next couple of weeks? (at least a CR)?	1	1	https://issues.jboss.org/browse/JBAS-3272
1311	JBAS-3257	Jerry, can you assign a Fix Version to this? If the answer is "no time in the next 9 months", please set Fix Version to "No Release".	0	0	https://issues.jboss.org/browse/JBAS-3257
1312	JBAS-3235	Rajesh, can you make those tests conditionally excluded from the testsuite?	1	1	https://issues.jboss.org/browse/JBAS-3235
1313	JBAS-3215	Somewhat related: JBAS-3118. (tell the scanner to temporarily ignores a URL while you make modifications to it).	0	0	https://issues.jboss.org/browse/JBAS-3215
1314	JBAS-3215	Will something be done about this? It keeps occuring at the most inconvenient times, mostly when my boss is breathing down my neck.	0	0	https://issues.jboss.org/browse/JBAS-3215
1315	JBAS-3215	Nice idea, we just don't have cycle for it. Do you want to help?	0	0	https://issues.jboss.org/browse/JBAS-3215
1316	JBAS-3193	Where in the codebase this needs fixing?	1	0	https://issues.jboss.org/browse/JBAS-3193

#	Requirement	Questions	Answered	Successfully Answered	URL
1317	JBAS-3192	the following which is not a CMP issue. <error message="SecurityException; nested exception is: javax.security.auth.login.FailedLoginException: Password Incorrect/Password Required" type="java.rmi.AccessException">java.rmi.AccessException: SecurityException; nested exception is: javax.security.auth.login.FailedLoginException: Password Incorrect/Password Required at org.jboss.ejb.plugins.LogInterceptor.handleException(LogInterceptor.java:388) at org.jboss.ejb.plugins.LogInterceptor.invokeHome(LogInterceptor.java:136) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invokeHome(ProxyFactoryFinderInterceptor.java:107) at org.jboss.ejb.SessionContainer.internalInvokeHome(SessionContainer.java:637) at org.jboss.ejb.Container.invoke(Container.java:975) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.server.Invocation.invoke(Invocation.java:86) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:659) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:819) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:420) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:261) at sun.rmi.transport.Transport$1.run(Transport.java:148) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:144) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:460) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:701) at java.lang.Thread.run(Thread.java:534) Caused by: javax.security.auth.login.FailedLoginException: Password Incorrect/Password Required at org.jboss.security.auth.spi.UsernamePasswordLoginModule.login(UsernamePasswordLoginModule.java:213) at org.jboss.security.auth.spi.UsersRolesLoginModule.login(UsersRolesLoginModule.java:152) at javax.security.auth.login.LoginContext.invoke(LoginContext.java:675) at javax.security.auth.login.LoginContext.access$000(LoginContext.java:129) at javax.security.auth.login.LoginContext$4.run(LoginContext.java:610) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.login.LoginContext.invokeModule(LoginContext.java:607) at javax.security.auth.login.LoginContext.login(LoginContext.java:534) at org.jboss.security.plugins.JaasSecurityManager.defaultLogin(JaasSecurityManager.java:601) at org.jboss.security.plugins.JaasSecurityManager.authenticate(JaasSecurityManager.java:535) at org.jboss.security.plugins.JaasSecurityManager.isValid(JaasSecurityManager.java:344) at org.jboss.ejb.plugins.SecurityInterceptor.checkSecurityAssociation(SecurityInterceptor.java:211) at org.jboss.ejb.plugins.SecurityInterceptor.invokeHome(SecurityInterceptor.java:135) at org.jboss.resource.connectionmanager.CachedConnectionInterceptor.invokeHome(CachedConnectionInterceptor.java:189) at org.jboss.ejb.plugins.StatefulSessionInstanceInterceptor.invokeHome(StatefulSessionInstanceInterceptor.java:143) at org.jboss.ejb.plugins.CallValidationInterceptor.invokeHome(CallValidationInterceptor.java:56) at	0	0	https://issues.jboss.org/browse/JBAS-3192
1318	JBAS-3190	I see the same symptom here, as well: Class org.jboss.test.pooled.test.SSLSocketsUnitTestCase testClientCertSSLAccess Error StatelessSessionWithPooledSSL not bound javax.naming.NameNotFoundException: StatelessSessionWithPooledSSL not bound at org.jnp.server.NamingServer.getBinding(NamingServer.java:529) at org.jnp.server.NamingServer.getBinding(NamingServer.java:537) at org.jnp.server.NamingServer.getObject(NamingServer.java:543) at org.jnp.server.NamingServer.lookup(NamingServer.java:296) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:261) at sun.rmi.transport.Transport$1.run(Transport.java:148) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:144) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:460) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:701) at java.lang.Thread.run(Thread.java:534) at sun.rmi.transport.StreamRemoteCall.exceptionReceivedFromServer(StreamRemoteCall.java:247) at sun.rmi.transport.StreamRemoteCall.executeCall(StreamRemoteCall.java:223) at sun.rmi.server.UnicastRef.invoke(UnicastRef.java:133) at org.jnp.server.NamingServer_Stub.lookup(Unknown Source) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:625) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:587) at javax.naming.InitialContext.lookup(InitialContext.java:347) at org.jboss.test.pooled.test.SSLSocketsUnitTestCase.testClientCertSSLAccess(SSLSocketsUnitTestCase.java:85) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23)	0	0	https://issues.jboss.org/browse/JBAS-3190
1319	JBAS-3190	The following 3 tests should pass: org.jboss.test.jrmp.test.SSLSocketsUnitTestCase Reason: The handshakecompleted listener had no time to populate the cipherSuites because the test would proceed to check for the cipherSuite. Added a 2 sec sleep to give the listener time to populate the test instance variable - cipherSuites. org.jboss.test.jrmp.test.SSLFailuresSocketsUnitTestCase Reason: JASSIST-18 did the trick org.jboss.test.pooled.test.SSLSocketsUnitTestCase Reason: was missing a verifier option in the BaseCertLoginModule config	1	1	https://issues.jboss.org/browse/JBAS-3190
1320	JBAS-3190	The following test anyway passes on CC. I need to revert back the following fix made. org.jboss.test.pooled.test.SSLSocketsUnitTestCase Reason: was missing a verifier option in the BaseCertLoginModule config	1	0	https://issues.jboss.org/browse/JBAS-3190
1321	JBAS-3183	We'll try to push more to get the MQ server installed. In the meantime, have you tried installing MQ locally? This is really a blocker for this 4.0.5.CR1 release, so let's say that the end of the next week (Aug/11th) is the absolute deadline.	0	0	https://issues.jboss.org/browse/JBAS-3183
1322	JBAS-3183	Well, if it's that critical, then I suggest we open a support account (unless we already have one) with IBM. I am not sure what type of 'workaround' you expect to be delivered for an IBM related defect especially in the time frame. Can someone tell me why this can't be delivered to this client as a patch after the 4.0.5CR1 release and why an issue with IBM's MQ code has to hold up our release?	1	1	https://issues.jboss.org/browse/JBAS-3183
1323	JBAS-3183	Hi Weston, Thanks for all your efforts to make it happen What we need now is very simple. We need a standalone test case that we can provide IBM support with. Would you be able to produce such code or does it have to be executed within JBossAS?	0	0	https://issues.jboss.org/browse/JBAS-3183
1324	JBAS-3162	This is done in 4.2, but not 5.0. I should be able to finish up today.	0	0	https://issues.jboss.org/browse/JBAS-3162
1325	JBAS-3158	Stan, do you want to schedule this for 4.2.0.CR1 ?	1	1	https://issues.jboss.org/browse/JBAS-3158
1326	JBAS-3156	Could you post instructions on how to run the tests for portal-2.2? A lot of the errors I mentioned above come from deployment errors of forums and missing Database connectors.	0	0	https://issues.jboss.org/browse/JBAS-3156
1327	JBAS-3142	Clebert, will you have this withing the next 2 weeks?	1	1	https://issues.jboss.org/browse/JBAS-3142
1328	JBAS-3142	Clebert, this issues seems out-of-date. Is there something you plan to do?	0	0	https://issues.jboss.org/browse/JBAS-3142
1329	JBAS-3133	What about the other discussion thread - attached? Why isn't this discussed in JBoss/JCA forum rather than the Mail forum?	0	0	https://issues.jboss.org/browse/JBAS-3133

#	Requirement	Questions	Answered	Successfully Answered	URL
1330	JBAS-3124	I see... Isn't it possible to get server address from java.naming.provider.url property on client (to strip port number, of course)? Then some clients may use an IP-address, other clients may use a host name.	1	0	https://issues.jboss.org/browse/JBAS-3124
1331	JBAS-3104	Weston, can you validate this and apply what makes sense?	0	0	https://issues.jboss.org/browse/JBAS-3104
1332	JBAS-3066	Any idea, who can validate the examples?	1	1	https://issues.jboss.org/browse/JBAS-3066
1333	JBAS-3052	The linked issue is fixed, so can we close this?	0	0	https://issues.jboss.org/browse/JBAS-3052
1334	JBAS-3029	This is unnecessary because you can simply change your query: SELECT * FROM (SELECT 'User', 'Roles' FROM USERS UNION SELECT 'Administrator','Roles' FROM ADMINISTRATORS) WHERE username=? The query optimizer of any sane database will notice the where clause and apply it to each select statement so there is no perfomance impact. I vote against this issue and suggest it is closed. Regards Felix	1	0	https://issues.jboss.org/browse/JBAS-3029
1335	JBAS-3012	Scott, could you clarify what you think the cause is here? We are running into this as well. A long running server had over 100K login modules plus 400K SimplePrincipals and 100K SecurityAssociationHandlers, SimpleGroups, etc. which eventually brought it to its knees. I've posted a screen shot from jprofiler at https://trac.openmicroscopy.org.uk/omero/ticket/818 , and confirmed that calling DynamicLoginConfig's flushAuthenticationCaches does free almost everything up. Not, however, the SecurityAssociation$SubjectContext instances. I'm still looking into that. My planned work-around is to periodically call flushAuthenticationCaches. Any thoughts? Many thanks, Josh Moore Glencoe Software, Inc. The Open Microscopy Environment From: jboss-service.xml <mbean code="org.jboss.security.auth.login.DynamicLoginConfig" name="omero:service=LoginConfig"> <attribute name="AuthConfig">META-INF/jboss-login.xml</attribute> <depends optional-attribute-name="LoginConfigService"> jboss.security:service=XMLLoginConfig </depends> <depends optional-attribute-name="SecurityManagerService"> jboss.security:service=JaasSecurityManager </depends> </mbean> From: jboss-login.xml <policy> <application-policy name="OmeroSecurity"> <authentication> <login-module code = "ome.security.JBossLoginModule" flag = "required"> <module-option name = "unauthenticatedIdentity">guest</module-option> <module-option name = "dsJndiName">java:/omeroDS</module-option> <module-option name = "hashAlgorithm">MD5</module-option>	0	0	https://issues.jboss.org/browse/JBAS-3012
1336	JBAS-2975	Jerry, from the forum discussion, etc., I believe this test issue is fixed. If so, can you resolve this issue? Thanks.	1	1	https://issues.jboss.org/browse/JBAS-2975
1337	JBAS-2972	The problem is caused by this change: http://fisheye.jboss.org/changelog/JBoss?cs=Branch_4_0:tdiesler:20060107001625 There is no annotationsCL = null; in cleanup(); This also shows a problem in EJBModuleFactory.destroy() in that deploymentToModuleNameMap() just accumulates DeploymentInfo references. Nothing ever removes things from it.	1	1	https://issues.jboss.org/browse/JBAS-2972
1338	JBAS-2951	I see the commit in HEAD. Are you going to do the same in Branch_4_0?	1	1	https://issues.jboss.org/browse/JBAS-2951
1339	JBAS-2921	Hany, I've forgotten where this stands. Can you remind me?	1	1	https://issues.jboss.org/browse/JBAS-2921
1340	JBAS-2904	Vicky, have you progressed with this?	1	1	https://issues.jboss.org/browse/JBAS-2904
1341	JBAS-2904	Can you port this to Branch_4_0 (again in SVN)? Then I'll have a look on the test undeployment issue	1	1	https://issues.jboss.org/browse/JBAS-2904
1342	JBAS-2899	Does this need to be re-implemented for Tomcat6?	1	1	https://issues.jboss.org/browse/JBAS-2899
1343	JBAS-2872	Down to 1 failing: testSRPLoginWithAuxChallenge Unable to complete login: Failed to complete SRP login, msg=Failed to encrypt aux challenge junit.framework.AssertionFailedError: Unable to complete login: Failed to complete SRP login, msg=Failed to encrypt aux challenge at org.jboss.test.security.test.SRPLoginModuleUnitTestCase.testSRPLoginWithAuxChallenge(SRPLoginModuleUnitTestCase.java:142) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23)	1	1	https://issues.jboss.org/browse/JBAS-2872
1344	JBAS-2861	Brian, can this make it for 4.0.5.CR1?	0	0	https://issues.jboss.org/browse/JBAS-2861
1345	JBAS-2861	I would like to know the plans for delivering this change in terms of time. Could someone post that information here?	0	0	https://issues.jboss.org/browse/JBAS-2861
1346	JBAS-2861	Is the intent still to include in 4.2.0 or will this only be available in the 5 family?	0	0	https://issues.jboss.org/browse/JBAS-2861
1347	JBAS-2861	What are the latest plans for this in terms of planned releases and timeframes? THis issue is of very high importance to our efforts. TIA!	0	0	https://issues.jboss.org/browse/JBAS-2861
1348	JBAS-2861	This issue keeps getting pushed and pushed and pushed. What does it take for this to be resolved?	0	0	https://issues.jboss.org/browse/JBAS-2861
1349	JBAS-2861	I see that this is assigned to the "backlog" for fix version. Is there any intent to implement this functionality?	0	0	https://issues.jboss.org/browse/JBAS-2861
1350	JBAS-2861	Hello, Is there an update on this? is this something that can be done in any way with jboss-5.1.0.GA-jdk6? Thank you for your answer. Luis	1	0	https://issues.jboss.org/browse/JBAS-2861
1351	JBAS-2860	If you want table name to not be checked for reserved words then use quotes in jbosscmp-jdbc.xml. What exactly is wrong with the current approach? The bug you refered to has already been fixed in HEAD.	0	0	https://issues.jboss.org/browse/JBAS-2860
1352	JBAS-2860	>What exactly is wrong with the current approach? The problem is that the quote character is different for different database engines and can even vary depending on configuration settings. For MySQL it could be either a backquote or a double quoute depending on the ANSI setting. Having the quotes in jbosscmp-jdbc.xml forces you to create different EAR's for different databases. >The bug you refered to has already been fixed in HEAD. OK, didn't know since it is still open.	1	0	https://issues.jboss.org/browse/JBAS-2860
1353	JBAS-2853	I see the trunk moved to 2.0.0.GA. Are we updating JGroups, too? (2.5)	0	0	https://issues.jboss.org/browse/JBAS-2853
1354	JBAS-2846	I don't see how this version of tomcat is usable due to the new behavior described in this bug: http://issues.apache.org/bugzilla/show_bug.cgi?id=38579 The tomcat WebappClassLoader is walking over private static and private static final field contents in classes that it has not loaded. Either a tomcat 5.5.16 is needed, or this change needs to be incorporated into a patched tomcat: http://svn.apache.org/viewcvs?rev=378251&view=rev	0	0	https://issues.jboss.org/browse/JBAS-2846
1355	JBAS-2846	Tomcat upgraded to v5.5.17 in Branch 4.0 and HEAD.	0	0	https://issues.jboss.org/browse/JBAS-2846
1356	JBAS-2831	Great! It's a useful feature. Does it mean that the caller identity propagation across servers is already available?	0	0	https://issues.jboss.org/browse/JBAS-2831
1357	JBAS-2737	Was this backported to 4.0.x or 4.2..x? Where is this new parameter documented?	0	0	https://issues.jboss.org/browse/JBAS-2737
1358	JBAS-2699	Is this the same thing as JBAS-2802?	0	0	https://issues.jboss.org/browse/JBAS-2699
1359	JBAS-2699	I just added tests on JBser per JBAS-3102. As 4.0.4 GA is the first release with UnifiedInvokers available, I don't think it would make sense to add these tests into MatrixTestSuite now. Because of that I've created JBAS-3142. I guess I could consider this task solved.	1	1	https://issues.jboss.org/browse/JBAS-2699
1360	JBAS-2696	This looks done. Can we close this?	0	0	https://issues.jboss.org/browse/JBAS-2696

#	Requirement	Questions	Answered	Successfully Answered	URL
1361	JBAS-2655	Are there any other scenarios in which the Clustered SSO valve may have issues when deployed on hosts on multiple services? If you think there are none, please close this JIRA issue. In the future, when customers report problems, this JIRA issue can be reopened for investigation.	0	0	https://issues.jboss.org/browse/JBAS-2655
1362	JBAS-2652	Scott, can you elaborate on this? Is this for checking the jsr77 names follow the spec?	0	0	https://issues.jboss.org/browse/JBAS-2652
1363	JBAS-2652	org.jboss.test.management.test.JSR77SpecUnitTestCase checks for ejb/web/rar names, what else is needed?	1	1	https://issues.jboss.org/browse/JBAS-2652
1364	JBAS-2647	Brian, can you please update this dated task (and the parent task)?	1	1	https://issues.jboss.org/browse/JBAS-2647
1365	JBAS-2634	Anil, is this a requirement for JavaEE5 ?	0	0	https://issues.jboss.org/browse/JBAS-2634
1366	JBAS-2607	Is this work finished?	1	1	https://issues.jboss.org/browse/JBAS-2607
1367	JBAS-2604	4.0.4.CR2 will go out with 1.4.1_final. Do we require the same for jboss-serialization, which is at 1.0.0.CR4, or the issue should be closed?	0	0	https://issues.jboss.org/browse/JBAS-2604
1368	JBAS-2575	Was this issue recreated in messaging? I could not find it int he messaging category. It appears that tim Fox reported JBMESSAGING-1172 http://jira.jboss.org/jira/browse/JBMESSAGING-1172. Is this the same issue or has it been resolved already?	0	0	https://issues.jboss.org/browse/JBAS-2575
1369	JBAS-2574	Hello Dimitris, Thanks for looking into this. We have a few quick questions: 1) Is the jboss_init_hpux.sh added to the 4.0.X branch of the JBoss AS: The message below indicates that the fixed JBoss versions are: Fix Version: JBossAS-3.2.8.SP1 JBossAS-5.0 Alpha However, further below it indicates the fixed versions are: Fix For: JBossAS-5.0 Alpha, JBossAS-4.0.4.GA, JBossAS-3.2.8.SP1 2) I haven't seen the name JBossAS-4.0.4.GA before. What does the "GA" at the end of the release name refer to? Cheers, - John	1	0	https://issues.jboss.org/browse/JBAS-2574
1370	JBAS-2566	Genman, is this still issue pending or can you finish it off? We'd like to be over with the enclosing task, thanks.	1	1	https://issues.jboss.org/browse/JBAS-2566
1371	JBAS-2545	Actually, let's skip 2.2.9 and go to 2.2.9.1 straight. My goal is to have JBoss 5.0 ship with JGroups 2.3 (in beta right now), or even with 2.4 (not likely though)	1	0	https://issues.jboss.org/browse/JBAS-2545
1372	JBAS-2545	What is the status on this task? Are we ready to test JGroups 2.3 yet?	0	0	https://issues.jboss.org/browse/JBAS-2545
1373	JBAS-2545	I upgraded repository.jboss.com to JGroups 2.3 (final). I also updated jboss-head to use 2.3. Can you run the tests one more time and - if they are successful - resolve this task ? Thanks, Bela	1	0	https://issues.jboss.org/browse/JBAS-2545
1374	JBAS-2519	Ok, Chris. Can you attach a test war to this JIRA issue and continue discussion at the forum thread, with some new concerns I have.	1	1	https://issues.jboss.org/browse/JBAS-2519
1375	JBAS-2510	The retrocheck utility I wrote showed that the retrotranslator was incomplete. There are parts of JBossWS where it will throw NoSuchMethodError Besides the fact that retrotranslator is very new and unsupportable, and also potentially conflicts with Hibernate on tested versions of cglib/asm. Why are we releasing an obsolete version in the candidate only to change it for final (which will then not be tested in the wider community)?	0	0	https://issues.jboss.org/browse/JBAS-2510
1376	JBAS-2494	Jerry, we have discussed this before. When you have the bandwidth, can you propose a fix?	0	0	https://issues.jboss.org/browse/JBAS-2494
1377	JBAS-2481	In which version is this solved?	1	1	https://issues.jboss.org/browse/JBAS-2481
1378	JBAS-2481	How is this solved? I want to use it in jboss 5.1.0.GA! Gunnar	1	0	https://issues.jboss.org/browse/JBAS-2481
1379	JBAS-2480	Is this being worked on? On which branch?	0	0	https://issues.jboss.org/browse/JBAS-2480
1380	JBAS-2466	Hi. I just downloaded version 3.2.8 SP1 and all the listed *.java files from the last comment still contain the proprietary license ... can we re-open this bug? The license text in the header is really restrictive. Thanks a lot & best regards, Daniel	1	1	https://issues.jboss.org/browse/JBAS-2466
1381	JBAS-2440	Jerry, do you have time to do this as part of the conversion of DRM to use JBoss Cache?	1	1	https://issues.jboss.org/browse/JBAS-2440
1382	JBAS-2423	Hi, Would it be possible to upgrade HttpClient to its latest stable version (3.0.1) or the current version 3.1-rc1? I am porting over code where I've used standalone with HttpClient and the following call is not supported when using jbossall-client.jar since the HttpClient class it packages is too old: httppost.setRequestEntity (new ByteArrayRequestEntity (ba)); where httppost is of type PostMethod. This is fairly standard code but the API is only available starting from 3.0 on. Even with this bug fix I would still have this problem since it only updates it to version 2.0.2, which is fairly old now. I see this bug was originally filed in November 2005 so a lot of changes have gone into HttpClient since then. It'd be nice to update it to the latest version if possible. Thanks, Jonathan Huang	1	1	https://issues.jboss.org/browse/JBAS-2423
1383	JBAS-2411	I've have some junit tests I written for a xml util library that could easily be modified to support this kind of testing. What is the scope of testing that is required.	1	1	https://issues.jboss.org/browse/JBAS-2411
1384	JBAS-2400	Brian, can you do this ? This should be a script probe.sh/probe.bat in the ./bin directory of JBoss. Thanks !	0	0	https://issues.jboss.org/browse/JBAS-2400
1385	JBAS-2369	Luc, will you be working on this?	1	1	https://issues.jboss.org/browse/JBAS-2369
1386	JBAS-2352	I'd like this functionality as well -- is there no other way to add a "remember me" cookie when using form-based auth? (I'm using JBAS 4.2.3GA).	1	0	https://issues.jboss.org/browse/JBAS-2352
1387	JBAS-2299	Is it possible that the following configuration within jboss-app.xml is the problem? I ask as there are quite a few log entries mentioning the HierarchicalLoaderRepository3 as shown subsequently below <jboss-app> <loader-repository loaderRepositoryClass="org.jboss.mx.loading.HeirarchicalLoaderRepository3"> xMP:type=HeirarchicalClassRepository, service=VXE1_ContentTestService, version=1.0 <loader-repository-config configParserClass="org.jboss.mx.loading.HeirarchicalLoaderRepository3ConfigParser"> java2ParentDelegaton=false </loader-repository-config> </loader-repository> </jboss-app> 2005-09-23 23:58:08,727 DEBUG [org.jboss.mx.loading.RepositoryClassLoader] setRepository, repository=org.jboss.mx.loading.HeirarchicalLoaderRepository3@6ab811, cl=org.jboss.mx.loading.HeirarchicalLoaderRepository3$CacheClassLoader@7daf5b{ url=null ,addedOrder=0} 2005-09-23 23:58:08,727 DEBUG [org.jboss.mx.loading.UnifiedClassLoader] New jmx UCL with url null 2005-09-23 23:58:08,727 DEBUG [org.jboss.mx.loading.RepositoryClassLoader] setRepository, repository=org.jboss.mx.loading.HeirarchicalLoaderRepository3@6ab811, cl=org.jboss.mx.loading.UnifiedClassLoader3@18380a4{ url=null ,addedOrder=0}	0	0	https://issues.jboss.org/browse/JBAS-2299

#	Requirement	Questions	Answered	Successfully Answered	URL
1388	JBAS-2299	Common-BeanUtils is broken: org.jboss.web.tomcat.tc5.WebAppClassLoader@6455597 !--- class org.apache.struts.action.ActionServlet !--- !--- [Ljava.lang.Class;@15889588 !--- !--- !--- java.lang.reflect.Method@2653351 !--- !--- !--- !--- [Ljava.lang.reflect.Method;@14422877 !--- !--- !--- !--- !--- java.util.Hashtable$Entry@23818829 !--- !--- !--- !--- !--- !--- [Ljava.util.Hashtable$Entry;@18693832 !--- !--- !--- !--- !--- !--- !--- java.util.Hashtable@17102985 !---!---!---!---!---!---!---!--- StaticFieldReference private static java.util.Hashtable org.apache.commons.beanutils.MappedPropertyDescriptor.declaredMethodCache Detail There is a list of declaredMethods (Reflection) being referenced inside of MappedPropertyDescriptor. BeanUtils is now being bundled with JBoss/JSF at this point. If you want to avoid this bug for this specific application, having BeanUtils deployed within the application + classLoader isolations will make BeanUtils to keep circular references inside the classLoader what would avoid the contention. We should look for a way to contribute a fix for apache. Besides this reflection cache, it looks also they have a circular reference on their WeakHashMap, having the classLoader on the value. (but I'm not sure about that one). But I know for sure BeanUtils is broken. Dimitris, since this is a Beanutils bug, what should we do here? Close, postpone? It's your call now.	1	1	https://issues.jboss.org/browse/JBAS-2299
1389	JBAS-2299	Clebert, can you please check if this has been fixed in beanutils 1.7.0 we use in Branch_4_2 ? If not we should consider contributing a fix.	1	1	https://issues.jboss.org/browse/JBAS-2299
1390	JBAS-2278	In order to derive the Principal i do this: setIdentity(SimplePrincipal(rs.getString(1))); where rs.getString(1) is the value of company certificate (DN) in users table of JAAS. If fact i mean employee instead of user, i only need use the employee certificate to check which company this employee belongs, after this i need the company identity to inherit the company role(s).	1	0	https://issues.jboss.org/browse/JBAS-2278
1391	JBAS-2278	Doesn't works because i don't have the company certificate to pass into CertificatePrincipal i just have the DN of the company. My problem in other words is, once the user is validated and this user (CEO) belongs to an company that has a DN in the system i need replace the user principal by company principal because this user will only use the company role to perform your tasks in the system. CertificatePrincipal uses and array of X509Certificates as parameter and i only have the DN string of the company at this point because the certificate presented to the JAAS was the user certificate and not the company certificate.	0	0	https://issues.jboss.org/browse/JBAS-2278
1392	JBAS-2251	Can we close this?	1	1	https://issues.jboss.org/browse/JBAS-2251
1393	JBAS-2243	What is an example client context that is exposing this info? The behavior of the login module itself is a feature request to use an alternate msg.	0	0	https://issues.jboss.org/browse/JBAS-2243
1394	JBAS-2235	Ok, I'm hijacking this issue. Did this forced shutdown ever worked? The http://dev45.qa.atl.jboss.com:8585/hudson/view/JBoss%20AS/job/JBoss-AS-5.0.x-TestSuite-sun15-noip run that hungs because of wrong multicast settings on the testing host is never shutdown properly. Who can look into this?	0	0	https://issues.jboss.org/browse/JBAS-2235
1395	JBAS-2235	I have looked at the Ant task source code and related classes in JBossAS test project. What you want is the ServerController throws an Exception (let's say InvalidStateException) when the server proccess must been killed. This exception is handled in StopServerTask and it show appropriate warning. Correct?	1	1	https://issues.jboss.org/browse/JBAS-2235
1396	JBAS-2235	So if we allow test suite to continue execution, then how will we notice that a server failed to shutdown? I think we should do one of: 1. fail the build 2. generate a simple junit xml report file indicating what's wrong	1	1	https://issues.jboss.org/browse/JBAS-2235
1397	JBAS-2235	I've committed the check if any server failed to shutdown in trunk r69559 Now jboss-test needs to be released and included in AS thirdparty. How is that done?	0	0	https://issues.jboss.org/browse/JBAS-2235
1398	JBAS-2233	Can you take a look at this?	1	1	https://issues.jboss.org/browse/JBAS-2233
1399	JBAS-2189	What value do think the default timeout should have? I just used the same one as the PooledInvoker. Also, changed the error message to be: java.rmi.ConnectException: Socket timed out. Waited 3000 milliseconds for response while calling on InvokerLocator [socket://127.0.0.1:8899/?socketTimeout=3000]; nested exception is: java.net.SocketTimeoutException: Read timed out	1	1	https://issues.jboss.org/browse/JBAS-2189
1400	JBAS-2189	I was going to change the default timeout to 30 minutes, but see you beat me to it in the jboss-service.xml and set to 10 min. I did set the default value within remoting to 30 minutes, but the config in jboss-service.xml will over ride it. As for keep alive ping, can't support on all transports, but can have for some (actually some will have this). What exception do you suggest that I throw? I don't want to throw IOException (or anything that inherits from it, such as SocketTimeoutException), since I can not include the cause by, which I want to keep for debugging. If throw original IOException, then can't add on extra information such as the timeout value or the locator was calling on. Why did you change tcp/ip no delay to true. While this will reduce latency, but will cause unneeded network traffic as won't buffer the data. Do you think the default will be that the data is that small as to not need buffering of the packets? Just curious if you already had info on why this was better?	0	0	https://issues.jboss.org/browse/JBAS-2189
1401	JBAS-2156	Would it be okay to back port (for 4.0.4 perhaps) a subset of the functionality that does not alter the public interfaces for the IL/Invoker stuff? The pooling changes (TimeoutFactory) then would only be for supporting the async expiration feature. I'll leave it to somebody else to work out the details for integrating a single pool for the UIL2 stuff, which I don't know much about. It sounds like you're not firm on what you want to do about it either. I don't know what DLQ dependency you're talking about. What was your end-to-end conception of how this feature is deployed?	1	1	https://issues.jboss.org/browse/JBAS-2156
1402	JBAS-2155	There are 2 bundles at http://www.apache.org/dist/xml/xalan-j/ xalan-j_2_7_0-bin.zip xalan-j_2_7_0-bin-2jars.zip Both of the them contain the same big (3Mb) xalan.jar ?	1	1	https://issues.jboss.org/browse/JBAS-2155
1403	JBAS-2155	So the default build does indeed produce the combined: interpretive-xalan + xsltc + xsltc dependencies (bcel, etc.) So what do we want here? a) interpretive-xalan b) interpretive-xalan + xsltc (withouht dependencies) The dependencies are: - java_cup(runtime) 0.10k - apache bcel 5.0rc1 - apache regexp 1.2 I can see jboss 4.0.3 uses: - apache bcel 5.1 - gnu-regexp 1.1.14	1	1	https://issues.jboss.org/browse/JBAS-2155
1404	JBAS-2153	The sprint deployer is released independenly of jboss, so I removed the JB5 release fix. It would be nice to have this done along with the 4.0.5 release though.	1	1	https://issues.jboss.org/browse/JBAS-2153
1405	JBAS-2153	Since Spring deployer is released independenly of JBossAS, should my unit test even be in testsuite? I'll write something similar to ejb3 tests. Ok?	1	1	https://issues.jboss.org/browse/JBAS-2153
1406	JBAS-2148	I don't really see the need for this? If you want programmatic access, why can you just use the platform MXBeans directly?	0	0	https://issues.jboss.org/browse/JBAS-2148
1407	JBAS-2146	Will this issue affect the codebase, or is it simply a doc?	1	1	https://issues.jboss.org/browse/JBAS-2146

#	Requirement	Questions	Answered	Successfully Answered	URL
1408	JBAS-2142	This seems like a special case. I presume also that cache is "optional" depending on whether you run in a clustered config? Since Tomcat5 is the one who controls the lifecycle and is responsible also for starting/stopping the connectors, he is the one who should signal (either explicitly or through a JMX notification) a cache start/stop. So I propose the following: Tomcat5 will have an <depends optional-attribute-name="CacheService>....</depends> on the Cache service mbean. This mbean should have a normal createService()/destroyService(), but an empty startService()/stopService(). Instead you'll implement 2 equivalent startCache()/stopCache() to do the state transfer. Tomcat5, at create() will check if the CacheService attribute is set, in order to locate the cache mbean and do whatever with it later in the start phase. Then just before the connectors are started(or stopped) (during Tomcat5 start) it will also issue a startCache()/stopCache() (again if the cache service is configured). Make sense?	0	0	https://issues.jboss.org/browse/JBAS-2142
1409	JBAS-2132	I think the installer only works with 1.5 to ensure that the user knows that some of the installable features require 1.5. I think the installer should technically be able to run under a 1.4 jre. Scott, can you comment on this? Norman or I could fix this.	1	1	https://issues.jboss.org/browse/JBAS-2132
1410	JBAS-2131	Weston, do you want to have a go with this, for the next release?	1	1	https://issues.jboss.org/browse/JBAS-2131
1411	JBAS-2109	Also, now I had to deploy jbosstestadapter.rar instead of jbosstestadapter.jar to add needed classes to the server's classpath. Local test in the default configuration fails for me with the following exception Testcase: testFailoverLocalMCF took 61.018 sec Caused an ERROR Failed to communicate. Problem during marshalling/unmarshalling; nested exception is: java.net.SocketTimeoutException: Read timed out java.rmi.ConnectException: Failed to communicate. Problem during marshalling/unmarshalling; nested exception is: java.net.SocketTimeoutException: Read timed out at org.jboss.remoting.transport.socket.SocketClientInvoker.transport(SocketClientInvoker.java:264) at org.jboss.remoting.RemoteClientInvoker.invoke(RemoteClientInvoker.java:112) at org.jboss.remoting.Client.invoke(Client.java:226) at org.jboss.remoting.Client.invoke(Client.java:189) at org.jboss.invocation.unified.interfaces.UnifiedInvokerProxy.invoke(UnifiedInvokerProxy.java:111) at org.jboss.invocation.InvokerInterceptor.invokeInvoker(InvokerInterceptor.java:227) at org.jboss.invocation.InvokerInterceptor.invoke(InvokerInterceptor.java:167) at org.jboss.proxy.TransactionInterceptor.invoke(TransactionInterceptor.java:46) at org.jboss.proxy.SecurityInterceptor.invoke(SecurityInterceptor.java:55) at org.jboss.proxy.ejb.StatelessSessionInterceptor.invoke(StatelessSessionInterceptor.java:97) at org.jboss.proxy.ClientContainer.invoke(ClientContainer.java:87) at $Proxy2.testHaLocalConnection(Unknown Source) at org.jboss.test.jca.test.HAConnectionFactoryUnitTestCase.testFailoverLocalMCF(HAConnectionFactoryUnitTestCase.java:55) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) Caused by: java.net.SocketTimeoutException: Read timed out at java.net.SocketInputStream.socketRead0(Native Method) at java.net.SocketInputStream.read(SocketInputStream.java:129) at java.io.BufferedInputStream.fill(BufferedInputStream.java:218) at java.io.BufferedInputStream.read(BufferedInputStream.java:235) at java.io.ObjectInputStream$PeekInputStream.peek(ObjectInputStream.java:2200) at java.io.ObjectInputStream$BlockDataInputStream.peek(ObjectInputStream.java:2490) at java.io.ObjectInputStream$BlockDataInputStream.peekByte(ObjectInputStream.java:2500) at java.io.ObjectInputStream.readObject0(ObjectInputStream.java:1267) at java.io.ObjectInputStream.readObject(ObjectInputStream.java:339) at org.jboss.remoting.marshal.serializable.SerializableUnMarshaller.read(SerializableUnMarshaller.java:73) at org.jboss.invocation.unified.marshall.InvocationUnMarshaller.read(InvocationUnMarshaller.java:32) at	1	1	https://issues.jboss.org/browse/JBAS-2109
1412	JBAS-2107	Scott, can you retest this. If it still fails, can you look at the server.log for the real error message.	1	1	https://issues.jboss.org/browse/JBAS-2107

#	Requirement	Questions	Answered	Successfully Answered	URL
1413	JBAS-2107	15:47:58,656 ERROR [LogInterceptor] EJBException in method: public abstract void org.jboss.test.isolation.interfaces.a.SessionA.invokeSessionB() throws java.rmi .RemoteException, causedBy: javax.naming.NamingException: Could not dereference object [Root exception is ja vax.naming.CommunicationException: Could not obtain connection to any of these u rls: localhost:1099 and discovery failed with error: javax.naming.CommunicationE xception: Receive timed out [Root exception is java.net.SocketTimeoutException: Receive timed out] [Root exception is javax.naming.CommunicationException: Faile d to connect to server localhost:1099 [Root exception is javax.naming.ServiceUna vailableException: Failed to connect to server localhost:1099 [Root exception is java.net.ConnectException: Connection refused: connect]]]] at org.jnp.interfaces.NamingContext.resolveLink(NamingContext.java:1052) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:692) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:701) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:572) at javax.naming.InitialContext.lookup(InitialContext.java:351) at org.jboss.test.isolation.ejb.a.SessionAEJB.invokeSessionB(SessionAEJB .java:36) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl. java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAcces sorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.invocation.Invocation.performCall(Invocation.java:345) at org.jboss.ejb.StatelessSessionContainer$ContainerInterceptor.invoke(S tatelessSessionContainer.java:214) at org.jboss.resource.connectionmanager.CachedConnectionInterceptor.invo ke(CachedConnectionInterceptor.java:149) at org.jboss.ejb.plugins.StatelessSessionInstanceInterceptor.invoke(Stat elessSessionInstanceInterceptor.java:154) at org.jboss.webservice.server.ServiceEndpointInterceptor.invoke(Service EndpointInterceptor.java:54) at org.jboss.ejb.plugins.CallValidationInterceptor.invoke(CallValidation Interceptor.java:48) at org.jboss.ejb.plugins.AbstractTxInterceptor.invokeNext(AbstractTxInte rceptor.java:105) at org.jboss.ejb.plugins.TxInterceptorCMT.runWithTransactions(TxIntercep torCMT.java:335) at org.jboss.ejb.plugins.TxInterceptorCMT.invoke(TxInterceptorCMT.java:1 66) at org.jboss.ejb.plugins.SecurityInterceptor.invoke(SecurityInterceptor. java:139) at org.jboss.ejb.plugins.LogInterceptor.invoke(LogInterceptor.java:192) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invoke(ProxyFacto ryFinderInterceptor.java:122) at org.jboss.ejb.SessionContainer.internalInvoke(SessionContainer.java:6 24) at org.jboss.ejb.Container.invoke(Container.java:873) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl. java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAcces sorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatch er.java:141) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:80) at org.jboss.mx.server.Invocation.invoke(Invocation.java:72) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker. java:245) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:644) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke (JRMPInvoker.java:805) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java: 406) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl. java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAcces sorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at	0	0	https://issues.jboss.org/browse/JBAS-2107
1414	JBAS-2106	Do we use jboss-cache.jar now instead build it from the cache module in 4.x? If we do, we should skip the cache unit testing and focus on the MBean integration testing instead under AS.	0	0	https://issues.jboss.org/browse/JBAS-2106
1415	JBAS-2105	The SecurityUnitTestCase is now passing, but the TxUnitTestCase is still failing on testAnnotated. Is this working for you? 18:56:45,031 ERROR [TxTester] failed java.lang.Exception: failed on mandatory no tx call at org.jboss.test.aop.bean.AnnotatedTxPOJO.callNever(AnnotatedTxPOJO.java:49) at org.jboss.test.aop.bean.TxTester.testAnnotated(TxTester.java:97) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:141) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:80) at org.jboss.mx.server.Invocation.invoke(Invocation.java:72) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:245) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:644) at sun.reflect.GeneratedMethodAccessor76.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.jmx.connector.invoker.InvokerAdaptorService.invoke(InvokerAdaptorService.java:251) at sun.reflect.GeneratedMethodAccessor73.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:141) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:80) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:118) at org.jboss.mx.server.Invocation.invoke(Invocation.java:74) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:127) at org.jboss.mx.server.Invocation.invoke(Invocation.java:74) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:245) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:644) at org.jboss.invocation.jrmp.server.JRMPProxyFactory.invoke(JRMPProxyFactory.java:164) at sun.reflect.GeneratedMethodAccessor72.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:141) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:80) at org.jboss.mx.server.Invocation.invoke(Invocation.java:72) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:245) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:644) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:805) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:406) at sun.reflect.GeneratedMethodAccessor71.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAcces sorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:294) at sun.rmi.transport.Transport$1.run(Transport.java:153) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:149) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:460) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:701) at java.lang.Thread.run(Thread.java:595)	1	1	https://issues.jboss.org/browse/JBAS-2105
1416	JBAS-2102	Will this include JBoss Messaging 1.2?	0	0	https://issues.jboss.org/browse/JBAS-2102

#	Requirement	Questions	Answered	Successfully Answered	URL
1417	JBAS-2090	for some reason. tests-iiop: [junit] Running org.jboss.test.bankiiop.test.BankStressTestCase [junit] Tests run: 8, Failures: 0, Errors: 9, Time elapsed: 76.985 sec [junit] Test org.jboss.test.bankiiop.test.BankStressTestCase FAILED [junit] Running org.jboss.test.excepiiop.test.ExceptionTimingStressTestCase [junit] Tests run: 6, Failures: 0, Errors: 7, Time elapsed: 60.641 sec [junit] Test org.jboss.test.excepiiop.test.ExceptionTimingStressTestCase FAILED [junit] Running org.jboss.test.helloiiop.test.HelloTimingStressTestCase [junit] Tests run: 6, Failures: 0, Errors: 7, Time elapsed: 60.75 sec [junit] Test org.jboss.test.helloiiop.test.HelloTimingStressTestCase FAILED [junit] Running org.jboss.test.hellojrmpiiop.test.HelloTimingStressTestCase [junit] Tests run: 7, Failures: 0, Errors: 7, Time elapsed: 56.328 sec [junit] Test org.jboss.test.hellojrmpiiop.test.HelloTimingStressTestCase FAILED [junit] Running org.jboss.test.iiop.test.NoCodebaseUnitTestCase [junit] Tests run: 1, Failures: 0, Errors: 1, Time elapsed: 6.156 sec ... <testcase classname="org.jboss.test.bankiiop.test.BankStressTestCase" name="testTeller" time="8.484"> <error message="Retries exceeded, couldn't reconnect to 127.0.0.1:3528" type="org.omg.CORBA.TRANSIENT">org.omg.CORBA.TRANSIENT: Retries exceeded, couldn't reconnect to 127.0.0.1:3528 vmcid: 0x0 minor code: 0 completed: No at org.jacorb.orb.iiop.ClientIIOPConnection.connect(Unknown Source) at org.jacorb.orb.giop.GIOPConnection.write(Unknown Source) at org.jacorb.orb.CDROutputStream.write(Unknown Source) at org.jacorb.orb.giop.ServiceContextTransportingOutputStream.write_to(Unknown Source) at org.jacorb.orb.giop.GIOPConnection.sendMessage(Unknown Source) at org.jacorb.orb.giop.GIOPConnection.sendRequest(Unknown Source) at org.jacorb.orb.giop.ClientConnection.sendRequest(Unknown Source) at org.jacorb.orb.giop.ClientConnection.sendRequest(Unknown Source) at org.jacorb.orb.Delegate.invoke_internal(Unknown Source) at org.jacorb.orb.Delegate.invoke(Unknown Source) at org.jacorb.orb.Delegate.is_a(Unknown Source) at org.omg.CORBA.portable.ObjectImpl._is_a(ObjectImpl.java:112) at org.omg.CosNaming.NamingContextHelper.narrow(NamingContextHelper.java:69) at com.sun.jndi.cosnaming.CNCtx.setOrbAndRootContext(CNCtx.java:345) at com.sun.jndi.cosnaming.CNCtx.initOrbAndRootContext(CNCtx.java:205) at com.sun.jndi.cosnaming.CNCtx.<init>(CNCtx.java:69) at com.sun.jndi.cosnaming.CNCtxFactory.getInitialContext(CNCtxFactory.java:32) at javax.naming.spi.NamingManager.getInitialContext(NamingManager.java:667) at javax.naming.InitialContext.getDefaultInitCtx(InitialContext.java:247) at javax.naming.InitialContext.init(InitialContext.java:223) at javax.naming.InitialContext.<init>(InitialContext.java:197) at org.jboss.test.bankiiop.test.BankStressTestCase.getInitialContext(BankStressTestCase.java:58) at org.jboss.test.bankiiop.test.BankStressTestCase.setUp(BankStressTestCase.java:488) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) </error> The cosnaming.jndi.properties being used by the jndi lookup: public BankStressTestCase(String name) throws java.io.IOException { super(name); java.net.URL url = ClassLoader.getSystemResource("cosnaming.jndi.properties"); jndiProps = new java.util.Properties(); jndiProps.load(url.openStream()); } // Override getInitialContext() ---------- protected InitialContext getInitialContext() throws Exception { return new InitialContext(jndiProps); } does still refer to localhost: java.naming.factory.initial=com.sun.jndi.cosnaming.CNCtxFactory java.naming.provider.url=corbaloc::localhost:3528/JBoss/Naming/root java.naming.factory.object=org.jboss.tm.iiop.client.IIOPClientUserTransactionObjectFactory I'm running with jdk5 so perhaps the ORBInitRef.NameService property is not	1	1	https://issues.jboss.org/browse/JBAS-2090
1418	JBAS-2090	With an override of the java.naming.provider.url for the deployment InitialContext, and a JBossIIOPTestCase to configure the cosnaming InitialContext most of the tests are passing: [starksm@banshee9100 testsuite]$ ant -Dnojars=t tests-iiop Buildfile: build.xml Overriding previous definition of reference to xdoclet.task.classpath tests-iiop: [junit] Running org.jboss.test.bankiiop.test.BankStressTestCase [junit] Tests run: 8, Failures: 0, Errors: 0, Time elapsed: 10.625 sec [junit] Running org.jboss.test.excepiiop.test.ExceptionTimingStressTestCase [junit] Tests run: 6, Failures: 0, Errors: 0, Time elapsed: 2.422 sec [junit] Running org.jboss.test.helloiiop.test.HelloTimingStressTestCase [junit] Tests run: 6, Failures: 0, Errors: 0, Time elapsed: 2.719 sec [junit] Running org.jboss.test.hellojrmpiiop.test.HelloTimingStressTestCase [junit] Tests run: 7, Failures: 0, Errors: 0, Time elapsed: 4.844 sec [junit] Running org.jboss.test.iiop.test.NoCodebaseUnitTestCase [junit] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 0.938 sec [junit] Running org.jboss.test.iiop.test.ParameterPassingStressTestCase [junit] Tests run: 20, Failures: 0, Errors: 0, Time elapsed: 3.859 sec [junit] Running org.jboss.test.txiiop.test.IIOPUserTransactionStressTestCase [junit] Tests run: 2, Failures: 0, Errors: 2, Time elapsed: 2.812 sec [junit] Test org.jboss.test.txiiop.test.IIOPUserTransactionStressTestCase FAILED [junit] Running org.jboss.test.txpropiiop.test.TxPropIIOPUnitTestCase [junit] Tests run: 12, Failures: 0, Errors: 0, Time elapsed: 1.985 sec [junit] Test org.jboss.test.txpropiiop.test.TxPropIIOPUnitTestCase The ut test seems to not be setting up the transaction. <testcase classname="org.jboss.test.txiiop.test.IIOPUserTransactionStressTestCase" name="testUserTx" time="1.203"> <error type="java.lang.IllegalStateException">java.lang.IllegalStateException at org.jboss.tm.iiop.client.IIOPClientUserTransaction.commit(IIOPClientUserTransaction.java:303) at org.jboss.test.txiiop.test.IIOPUserTransactionStressTestCase.testUserTx(IIOPUserTransactionStressTestCase.java:57) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) Caused by: org.omg.CosTransactions.NoTransaction: IDL:omg.org/CosTransactions/NoTransaction:1.0 at org.jboss.tm.iiop.client.IIOPClientUserTransaction.getThreadLocalTerminator(IIOPClientUserTransaction.java:152) at org.jboss.tm.iiop.client.IIOPClientUserTransaction.commit(IIOPClientUserTransaction.java:298) ... 22 more	0	0	https://issues.jboss.org/browse/JBAS-2090
1419	JBAS-2085	I changed OTSServant.get_txcontext() to throw CORBA::BAD_INV_ORDER if the transaction is not in the active state. This seems to be the reasonable thing to do, even though the OTS spec does not say a word on it. Other related changes: - TransactionImpl.getPropagationContext() and TransactionImpl.getOTSPropagationContext() now return null if the transaction is not in the active state. - TransactionInactiveException was added to the throws list of Coordinator.getTransactionContext(). - DTMServant.getTransactionContext() now throws that exception if the transaction is not in the active state. - TxServerClientInterceptor now uses !TxUtils.isActive(tx) (instead of TxUtils.isCompleted(tx)) as the condition for not propagating the OTS context. All calls to getTimeLeftBeforeTimeout() are now passing true in the errorRollback parameter. As an aside: the RollbackException currently thrown by this method is a bit misleading, as the transaction might have been committed rather than rolled back. Wouldn't it be better to throw java.lang.IllegalState or org.jboss.tm.remoting.interfaces.TransactionInactiveException?	0	0	https://issues.jboss.org/browse/JBAS-2085

#	Requirement	Questions	Answered	Successfully Answered	URL
1420	JBAS-2073	The entry point is javax.xml.transform.TransformerFactory.newInstance() that tries to instansiate the class stored at property: "javax.xml.transform.TransformerFactory". When the property is not set if fall back to: org.apache.xalan.processor.TransformFactoryImpl (jdk1.4) com.sun.org.apache.xalan.internal.xsltc.trax.TransformerFactoryImpl (jdk5) The problem is org.apache.xalan.* is included with jdk1.4 so unless you use the lib/endorsed overriding mechanism, the jdk1.4 provided xalan 2.4.1 classes will be loaded. They solved the problem in jdk5 by renaming the provided xalan implementation. I don't see how we can go around this, without using the lib/endorsed mechanism, especially if code is trying to access the org.apache.xalan classes directly. As a side note, in the UCL model, how the system classes take precedence in the default UCL repository?	0	0	https://issues.jboss.org/browse/JBAS-2073
1421	JBAS-2050	The first date tested is ambiguous as it has no timezone info. Is that the comparison that is failing? Looking at the problem is the setup of the expected answers: // The expected instance for each inputData value calendar.set(2005, 0, 4, 0, 0, 0); calendar.set(Calendar.MILLISECOND, 0); calendar.setTimeZone(TimeZone.getTimeZone("PST")); Date date1 = calendar.getTime(); calendar.set(2005, 0, 4, 23, 38, 21); Date date2 = calendar.getTime(); calendar.set(2005, 0, 4, 23, 38, 48); Date date3 = calendar.getTime(); The date1 value should not being using the PST timezone since the DateEditor will be using whatever the default is. Try changing this to: // The expected instance for each inputData value calendar.set(2005, 0, 4, 0, 0, 0); calendar.set(Calendar.MILLISECOND, 0); Date date1 = calendar.getTime(); calendar.setTimeZone(TimeZone.getTimeZone("PST")); calendar.set(2005, 0, 4, 23, 38, 21); Date date2 = calendar.getTime(); calendar.set(2005, 0, 4, 23, 38, 48); Date date3 = calendar.getTime();	1	1	https://issues.jboss.org/browse/JBAS-2050
1422	JBAS-2042	>> Use the template files on JIRA for consistency. Where to find them in JIRA? >> Template files for other IDEs(IntelliJ-IDEA, NetBeans should be available here soon. Any progress? I'm an IntelliJ-IDEA user. >> Also, it should not use printStackTrace(), In current version, there is no such usage. >> nor should it use the horrible log.error and throw paradigm Spring deployer followed the idea of AspectDeployer, which includes such stuff. What to do?	1	1	https://issues.jboss.org/browse/JBAS-2042
1423	JBAS-2019	Why doesn't the ClientAddress and various "JMS client properties for the UIL2 transport" address this? See, http://wiki.jboss.org/wiki/Wiki.jsp?page=ConfigUIL2	0	0	https://issues.jboss.org/browse/JBAS-2019
1424	JBAS-2019	is it possible to use a dns name with ConnectAddress ?	0	0	https://issues.jboss.org/browse/JBAS-2019
1425	JBAS-2004	Is the idea here to not throw an exception if the timeout kicks in, e.g. make a sort of REPL_SEMI_SYNC? Otherwise the existing JBC SyncReplTimeout already does this.	1	1	https://issues.jboss.org/browse/JBAS-2004
1426	JBAS-2004	JBC exposes the syncReplTimeout configuration property to control this. With the AS 5 CacheManager different webapps could use different cache configs and thus have separate values, so I don't see much benefit in pushing the config for this up to jboss-web.xml. Is there anything else you were thinking of here, Bela?	1	1	https://issues.jboss.org/browse/JBAS-2004
1427	JBAS-2002	I can't say I'm very familiar with this code, but this look wrong to me. Why have you modified all those constructors to pass the MBeanServer around? Besides breaking the api and lifecycle contract, is should be unnecessary. As far as I can see these are all MBeans and most are extensions to ServiceMBeanSupport. It should be using the MBeanRegistration interface to get the MBeanServer (ServiceMBeanSupport does this for you). And either that interface or the ServiceMBean interface for state changes. There should be no side affects in constructors for MBeans. Can you explain in more detail what you are fixing? There is obviously some problem with stop() and start() not being symmetrical, but that should be fixed by making them symmetrical (internal state machine fixes) rather than changing public api signatures.	1	1	https://issues.jboss.org/browse/JBAS-2002
1428	JBAS-2002	Also, is it possible you can provide patches againt the latest CVS? http://wiki.jboss.org/wiki/Wiki.jsp?page=CVSRepository There have been a number of fixes/changes in the deployers in 4.0.3 so it is pretty difficult for me to apply your patches - or even figure out what you changed in some cases:-(	1	1	https://issues.jboss.org/browse/JBAS-2002
1429	JBAS-2000	Follow-up : should we also remove hibernate-specific jars from the base distro and just dropping jboss-hibernate.jar into the client directory. This would allow each deployed application to more easily use their own copy/version of hibernate3 (provided ejb3 is not deployed).	0	0	https://issues.jboss.org/browse/JBAS-2000
1430	JBAS-1992	So, to summarize: Improvements done (they'll be out with 4.0.3RC2) - org.jboss.util.propertyeditor.StringEditor added, to reduce failed editor lookup in jboss editors package - the ServiceBindingManager is now properly plugable, it won't be called if it hasn't registered with the ServiceController - DeploymentInfo & TomcatDeployer won't try to open .xml files as zip files. Not done - I don't like the changes proposed by ServerConfig/ServerLoader. If I understand correctly it hacks the jboss bootclasspath (libraries in ./lib) by adding commonly used libraries from ./server/xxx/lib. This seems against the current design, so no change there. - I don't see any of the AbstractMBeanInvoker, AbstractWebContainer files attached to evaluate them. Cheers /Dimitris	1	1	https://issues.jboss.org/browse/JBAS-1992
1431	JBAS-1992	Can a blacklist clear mechanism be implemented similar to the existing JBoss blacklist clear algorithms?	0	0	https://issues.jboss.org/browse/JBAS-1992
1432	JBAS-1992	To the submitters: could you please provide more info as to why the classloading patches in case JBAS-2166 are needed? Please add any information on that particular case, this one is now closed.	0	0	https://issues.jboss.org/browse/JBAS-1992
1433	JBAS-1980	I'm wondering if the way we do is correct though, i.e. if invoking on attributes is supported, are we supposed to show in the metadata those attributes appearing both as attributes and operations?	1	1	https://issues.jboss.org/browse/JBAS-1980

#	Requirement	Questions	Answered	Successfully Answered	URL
1434	JBAS-1974	This has been done. All default addresses have been left in place rather than using a single address for backwards compatibility/version interoperability reasons. adrian@adrian-srv01:~/jboss-4.0/workspace/build/output/jboss-4.0.3RC1/bin> ./run.sh --help ========== JBoss Bootstrap Environment JBOSS_HOME: /home/adrian/jboss-4.0/workspace/build/output/jboss-4.0.3RC1 JAVA: /home/adrian/downloads/j2sdk1.4.2_08/bin/java JAVA_OPTS: -server -Xms128m -Xmx128m -Dprogram.name=run.sh CLASSPATH: /home/adrian/jboss-4.0/workspace/build/output/jboss-4.0.3RC1/bin/run.jar:/home/adrian/downloads/j2sdk1.4.2_08/lib/tools.jar ========== usage: run.sh [options] options: -h, --help Show this help message -V, --version Show version information -- Stop processing options -D<name>[=<value>] Set a system property -d, --bootdir=<dir> Set the boot patch directory; Must be absolute or url -p, --patchdir=<dir> Set the patch directory; Must be absolute or url -n, --netboot=<url> Boot from net with the given url as base -c, --configuration=<name> Set the server configuration name -j, --jaxp=<type> Set the JAXP impl type (ie. crimson) -B, --bootlib=<filename> Add an extra library to the front bootclasspth -L, --library=<filename> Add an extra library to the loaders classpath -C, --classpath=<url> Add an extra url to the loaders classpath -P, --properties=<url> Load system properties from the given url -b, --host=<host or ip> Bind address for all JBoss services -g, --partition=<name> HA Partition name (default=DefaultDomain) -u, --udp=<ip> UDP multicast address	1	1	https://issues.jboss.org/browse/JBAS-1974
1435	JBAS-1973	Done, you can now use "-g whatever" or "--partition=whatever" to set the partition name. adrian@adrian-srv01:~/jboss-4.0/workspace/build/output/jboss-4.0.3RC1/bin> ./run.sh --help ========== JBoss Bootstrap Environment JBOSS_HOME: /home/adrian/jboss-4.0/workspace/build/output/jboss-4.0.3RC1 JAVA: /home/adrian/downloads/j2sdk1.4.2_08/bin/java JAVA_OPTS: -server -Xms128m -Xmx128m -Dprogram.name=run.sh CLASSPATH: /home/adrian/jboss-4.0/workspace/build/output/jboss-4.0.3RC1/bin/run.jar:/home/adrian/downloads/j2sdk1.4.2_08/lib/tools.jar ========== usage: run.sh [options] options: -h, --help Show this help message -V, --version Show version information -- Stop processing options -D<name>[=<value>] Set a system property -d, --bootdir=<dir> Set the boot patch directory; Must be absolute or url -p, --patchdir=<dir> Set the patch directory; Must be absolute or url -n, --netboot=<url> Boot from net with the given url as base -c, --configuration=<name> Set the server configuration name -j, --jaxp=<type> Set the JAXP impl type (ie. crimson) -B, --bootlib=<filename> Add an extra library to the front bootclasspth -L, --library=<filename> Add an extra library to the loaders classpath -C, --classpath=<url> Add an extra url to the loaders classpath -P, --properties=<url> Load system properties from the given url -b, --host=<host or ip> Bind address for all JBoss services -g, --partition=<name> HA Partition name (default=DefaultDomain)	1	1	https://issues.jboss.org/browse/JBAS-1973
1436	JBAS-1970	I don't see why we need a new dom4j-jboss thirdparty entry. What does the thirdparty directory name have to do with the dom4j.jar artifact name? This is simply a new version of dom4j that I don't see as any different from a 1.6.1 version. We have no dependency on jaxen. The only reason it was added was because of the binding of theses classes to the root dom4j.jar by the bootstrap classloader, and this resulted in problems with loading jaxen from app level deployments. If people want to use the dom4j/jaxen integration they simply deploy a full dom4j.jar and jaxen.jar as the server config/app deployment level.	0	0	https://issues.jboss.org/browse/JBAS-1970
1437	JBAS-1970	So once this is complete the end result will be that the current dom4j.jar retrieved by checking out the 4.0 branch of JBoss4.0 will not contain the xpath package. If this breaks people's code, then they will need to deploy a full dom4j.jar, but where will they get the complete jar from? Haven't we just overwritten it with an incomplete one? My assumption was that two versions of this jar would need to be present in cvs, an incomplete one, and a complete one. Because both versions need to have the same filename they would have to reside in different directories.	0	0	https://issues.jboss.org/browse/JBAS-1970
1438	JBAS-1970	the new problem is how can I load the full dom4j jar in app(ear) level?I mean the dom4j.jar(jboss) is in the top level class loader,if I use dom4j's xpath function,it must be put in the same folder with jboss version dom4j.jar.	0	0	https://issues.jboss.org/browse/JBAS-1970
1439	JBAS-1965	What is an example test that is failing currently because of this? The jbossxb classes where synched with the latest changes from head, and there is an example of any processing working in the org.jboss.test.naming.test.JNDIBindingUnitTestCase.	1	1	https://issues.jboss.org/browse/JBAS-1965
1440	JBAS-1965	Can we take discussions and random speculations to the forums. This at least is a bug in that it should be giving you a more meaningful error message: java.lang.NullPointerException at org.jboss.beans.metadata.plugins.StringValueMetaData.getValue(StringValueMetaData.java:80) like "I have no idea what types you want me to insert into this list".	1	1	https://issues.jboss.org/browse/JBAS-1965
1441	JBAS-1965	Where does it say you want a started/installed pool in your bean? This says you want a configured pool: <property name="poolService"> <inject bean="PoolService" state="Configured"/> </property> which is not the same thing.	1	1	https://issues.jboss.org/browse/JBAS-1965
1442	JBAS-1955	Just in case Fabiano - can you get me one single .zip containing the latest version of those files? Or do I need to use all 3 zip files to reproduce it?	1	1	https://issues.jboss.org/browse/JBAS-1955
1443	JBAS-1925	I noticed a JMSContainerInvokerMessageDrivenUnitTestCase class in the test suite which just test the JMSContainerInvoker implementation against a Queue using the SimpleMessageDrivenUnitTest. Is it better to implement the testStopDelivery & testStartDelivery in there or Extends this class? How I see the implementation is something like this public void testStopDelivery() throws Exception { Properties props = (Properties) defaultProps.clone(); //set some default props runTest(getStopOperations(), props); } public void testStartDelivery() throws Exception { Properties props = (Properties) defaultProps.clone(); //set some default props runTest(getStartOperations(), props); } with getStartOperations() and getStoptOperations() working the same as SimpleMessageDrivenUnitTest. I guess a few new operations will need to be created to invoke the stopDelivery and startDelivery methods on the JMSContainerInvoker. Also I could be way off the mark here....	1	1	https://issues.jboss.org/browse/JBAS-1925

#	Requirement	Questions	Answered	Successfully Answered	URL
1444	JBAS-1900	Was the customer using FORM authentication? There is any issue with FORM authentication where Tomcat stores login credentials in a transient map in the session. This map is not being replicated (which we need to fix, at least for these values). Without these values, Tomcat can't authenticate the user from the session when redirected. ClusteredSSO fixes this by providing another source for the login credentials, but we need to make them available from the session itself. With BASIC and CERT authentication, the browser presents the credentials with each request, so I expect this to only affect FORM authentication.	0	0	https://issues.jboss.org/browse/JBAS-1900
1445	JBAS-1871	Scott, this issue appears out-of-date. Do you still need these tests written?	1	0	https://issues.jboss.org/browse/JBAS-1871
1446	JBAS-1865	Okay, here it is. Do you want me to add the JIRA list for 2.2.8 as well ? Fast Message marshalling ---------- - Replaced Externalizable for Message with Streamable, resulting in much faster marshalling and reduced size of marshalled messages, allowing for more messages to be sent / second - org.jgroups.tests.MessageSerializationTest2 can be used to compare Externalizable with Streamable: - For 50000 messages, size reduction is almost 50%, marshalling 150% faster and unmarshalling 650% faster flags="-Xmx500M -Xms500M -XX:NewRatio=1 -XX:+AggressiveHeap -verbose:gc -XX:+DisableExplicitGC -XX:ThreadStackSize=32 -XX:CompileThreshold=100" java $flags org.jgroups.tests.MessageSerializationTest2 -num 100000 -add_headers false serialized size=8588935, streamable size=5788899, streamable is 48 percent smaller serialized write=831, streamable write=331, streamable write is 151 percent faster serialized read=1352, streamable read=180, streamable read is 651 percent faster Performance numbers ---------- - Real tests will be produced in 2.2.9, using org.jgroups.tests.perf.Test - With JGroups/conf/fc-fast-minimalthreads.xml, I got ca 5000 1K messages on my laptop: (2 members, 1 sender, 1 receiver, 1 CPU laptop, 20000 1K msgs, 100Mbps switch) - April 20 2005: **5006** msgs/sec (on 192.168.5.1): -Xmx500M -Xms500M -XX:NewRatio=1 -XX:+AggressiveHeap -verbose:gc -XX:+DisableExplicitGC -XX:ThreadStackSize=32 -XX:CompileThreshold=100: combined: num_msgs_expected=20000, num_msgs_received=20000 (loss rate=0%), received=20MB, time=3995ms, msgs/sec=5006.26, throughput=5.01MB/sec UDP ---------- - bind_to_all_interfaces="true" now allows to listen for multicast messages on *all* available interfaces. This requires 1.4, under 1.3 the default interface will be selected MPING ----- - MPING allows for a combination where node discovery in a cluster uses multicast, but the real transport uses TCP. This is an addition to TCP:TCPPING and TCP:TCPGOSSIP. Example (short version of JGroups/conf/mping.xml): <config> <TCP bind_addr="192.168.5.1" start_port="7800" sock_conn_timeout="500" loopback="true"/> <MPING timeout="2000" bind_to_all_interfaces="true" mcast_addr="228.8.8.8" mcast_port="7500" ip_ttl="8" num_initial_members="2" num_ping_requests="1"/> <FD timeout="2000" max_tries="3"/> <pbcast.NAKACK gc_lag="50" retransmit_timeout="600,1200,2400,4800" max_xmit_size="60000"/> <pbcast.STABLE stability_delay="1000" desired_avg_gossip="20000" down_thread="false" max_bytes="0" up_thread="false"/> <pbcast.GMS print_local_addr="true" join_timeout="5000" join_retry_timeout="2000" shun="true"/> </config> Concurrent startup ---------- - When multiple members are started simultaneously, and no other member is running yet, they form singleton groups, and merge after some time. The new version avoids this merge, so merging occurs only after network partitions now, never on concurrent startup of initial members	0	0	https://issues.jboss.org/browse/JBAS-1865
1447	JBAS-1864	I am going to revert this change. I already parameterized this using the spec defined mechanism, i.e. using a system property when ant is invoked. <sysproperty key="java.naming.provider.url" value="${node0.jndi.url}"/> Please, can we avoid doing script/template type rubbish. I am trying my best to get the testsuite into a state where you can run it without having to go through a 50 minute build to run a three line change in a test. Ideally from inside eclipse where it doesn't even need to go through the ant build script if you are not changing the deployment. Where else has this ant template rubbish been done? As a general rule, we should support things using first class notions, not workaround things with code generation, templates, scripting or any other stupid code.	0	0	https://issues.jboss.org/browse/JBAS-1864
1448	JBAS-1839	There should be no whitespace in the hostname or urls so whitespace could simply be another delimiter. What was an example of the whitespace plagued usage?	1	1	https://issues.jboss.org/browse/JBAS-1839
1449	JBAS-1836	Summary of changes in the transaction module ---------- TransactionImpl: added getAssociatedThreadCount(), call out to TransactionSynchronization handler before/after commit/rollback, restore the interrupted status of the thread TransactionManagerService: new SynchronizationService attribute TxManager: holds TransactionSynchronization instance when registered, disassociateThread on tm.suspend() TransactionLocal: Add lock()/unlock() to allow clients to optionally synchronize access to the transaction local TransactionLocalDelegateImpl: se TransactionLocal instead of 'this' as the key to the TransactionLocalSynchronization map TxUtils: added String getStatusAsString(int status) Summary of additions in the transaction module ---------- --- TransactionSynchronization: new interface for a jboss specific singleton synchronization callback with access to TransactionImpl TransactionSynchronizationService: new service encapsulating a TransactionSynchronization implementation of a simple policy that rollback the tx if more than one threads are still associated at commit time. It includes a ForcedRollbackCounter and a setting to enable/disable the policy at runtime. TransactionSynchronizationServiceMBean: the interface of the service SynchronizedTransactionLocal: new Helper for automatic synchronization Summary of addition for transaction tests ---------- MTTest, MTTestMBean: server side driver for multithreaded tests MTOperation: tasks that can be combined in multithread tests MTTransactionManagerUnitTestCase: the test cases	1	1	https://issues.jboss.org/browse/JBAS-1836
1450	JBAS-1824	The reason why we originally wanted "*" to be understood as a "any role mapping" was that the container will only do a (basic) authentication if we have specified an <auth-constraint>, for which we wanted to add "*". The actual problem though is that the container does not perform an authentication on an unprotected resource. So what we really want to have is authentication in the web container, and later an authorization in an ejbs. So it would be really nice if we could configure (on a web app or even resource base), that the web container should do an authentication even if no authorization is required. Shall we add this as new feature request? Regards, Simon	1	1	https://issues.jboss.org/browse/JBAS-1824
1451	JBAS-1824	Anil, will you be working on this?	1	1	https://issues.jboss.org/browse/JBAS-1824
1452	JBAS-1820	<quote> I don't know if this is what A. Brock means by "...this does not give the full semantics of what you want." </quote> The spec says nothing about being able to redeploy the rar without also redeploying the MDB. Which is obviously required because of classloading considerations. Just adding a <depends> does not force a full redeployment of the MDB container. It only controls the ordering of startup and shutdown between the two components Hence this a feature request to get the "full semantics of what you want".	1	1	https://issues.jboss.org/browse/JBAS-1820

#	Requirement	Questions	Answered	Successfully Answered	URL
1453	JBAS-1818	Correct. It's just simpler and more accurate to produce the descriptor from the live metadata. Maybe a flag to control this would do? But this needs to make sure the replaced entries in the descriptor fully match the entity declaration in the DTD, plus it'll have to be changed whenever they change. Do you want to contribute this? :)	0	0	https://issues.jboss.org/browse/JBAS-1818
1454	JBAS-1809	Why isn't this using the default syntax ${jboss.server.lib.url:lib} ?	0	0	https://issues.jboss.org/browse/JBAS-1809
1455	JBAS-1783	Hi Adrian, First of all, thanks for the quick feedback - that was one of the quickiest I ever got on an OSS project :-) Now, back to the new feature, I can provide 2 new patches against the cvs (using jboss-core and Branch_3_2 - please correct me if I'm wrong). But before I do so, I have a few questions: - instead of using AssertionError - which is dependent on JDK 1.4 - we should either choose another JDK RuntimeException or create a new one. I would prefer to create a new one (such as ConnectionNotClosedException), so we can handle it better. - if we opt for creating a new one, should it be connection-specific or resource-generic? In other words, should we call it ConnectionNotClosedException or ResourceNotClosedException? I think it's the latter, but I'm not sure. - also, in which package/module should we add this new Exception? Could it go to the common module and in the org.jboss.util package? Regards, -- Felipe	1	1	https://issues.jboss.org/browse/JBAS-1783
1456	JBAS-1770	What are the immediate requirements for this task ? Is it correct to consider altering the build/packaging process to include the testsuite and its dependencies into the distribution as the primary objective ? Which of the enhancements above are required and which should be added to the wish list ? Do we have to include some kind of documentation about the details on running the different tests ? Those questions might seem stupid to you but, please, bear with me since I am a newcomer and it is not obvious for me.	1	0	https://issues.jboss.org/browse/JBAS-1770
1457	JBAS-1768	Ryan, is there someone from QA that could help with this type of tasks?	0	0	https://issues.jboss.org/browse/JBAS-1768
1458	JBAS-1768	Adrian, do you mean that some of the tests in the JBossAS testsuite are by other projects as well ? If this is true - am I correctly interpreting you comment that you want those common tests moved to a separate project (cvs/svn module) so that everyone who depends on them can use them in a clean way and share the same codebase ?	0	0	https://issues.jboss.org/browse/JBAS-1768
1459	JBAS-1747	The backport as needed by ejb3 is done. There are a couple of changes in head such as the jbxb:addMethod, and a dicussion of how to improve the namespace to xsd resolution. Is there anything missing for the bean deployer?	1	1	https://issues.jboss.org/browse/JBAS-1747
1460	JBAS-1747	Just what needs to be done before we can tag this for incorporation into the 4.0 branch so that the tests can be incorporated and testing can start. The StringBuilder is a jdk5 class. Does the bean deployer also require jdk5 for its base behavior? We can't introduce a jdk5 dependency at the common module level though.	1	1	https://issues.jboss.org/browse/JBAS-1747
1461	JBAS-1713	SourceForge Username: ioparra . Logged In: YES user_id=812998 It appears that the problem is the MANIFEST declaration. You have: application.xml deploying ejb.jar that has 1 classpath entry. JBoss will add that one classpath entry, but will not recursively check each child. Work around: If you want to include lib2.jar in your classpatch, add lib2.jar to MANIFEST.MF for ejb.jar. That'll fix the problem. Solution: Is an EAR container supposed to recursively run through ALL jars that are added to the classpath and add there MANIFEST.MF class-path entries? JBoss is assuming only the top level. The Spec says: If the by-reference approach is used, the ejb-jar file producer does not duplicate the content of the ejb-client JAR file in the ejb-jar file, but instead uses a Manifest Class-Path entry in the ejb-jar file to specify that the ejb-jar file depends on the ejb- client JAR at runtime. The use of the Class-Path entries in the JAR files is explained in the Java 2 Platform, Enterprise Edition specification [9]. Neither the j2EE or the ejb spec talks about recursively solving all jar dependencies.	1	1	https://issues.jboss.org/browse/JBAS-1713
1462	JBAS-1713	SourceForge Username: adrianprice . Logged In: YES user_id=580837 Many thanks for the response. However, I beg to differ: the J2EE spec. does require recursive resolution of bundled jar file Class-Path entries, as does the jar file specification. "J2EE.8.2.1 Assembling a J2EE Application", j2ee-1_3-fr- spec.pdf, page 106: "d. Make sure that each component in the application properly describes any dependencies it may have on other components in the application. A J2EE application should not assume that all components in the application will be available on the ???classpath??? of the application at run time. Each component might be loaded into a separate class loader with a separate namespace. If the classes in a JAR file depend on classes in another JAR file, the first JAR file should reference the second JAR file using the Class-Path mechanism. A notable exception to this rule is JAR files located in the WEB- INF/lib directory of a web application. All such JAR files are included in the class path of the web application at runtime; explicit references to them using the Class-Path mechanism are not needed." Recursive Class-Path resolution is the only approach that makes sense. Unless dependencies are declared locally by Class-Path attributes in the jar files containing the referencing classes and resolved recursively, all such dependencies would have to be understood and surfaced by top-level ejb-jar files. This problem is newly introduced between JBoss-3.0.4 and 3.2.1. My application runs perfectly on the older version, and on other app. servers. HOWEVER, I must have been tired last night and accidentally deployed the ejb.jar instead of the .ear, because when I tried again this morning I discovered that there was an error in the application.xml file and the .ear itself wouldn't deploy. I fixed the error, redeployed, and hey presto - the test case works correctly!!! There's still a problem with my application, but I finally diagnosed the cause as related to jar file indexes. Deleting the META-INF/INDEX.LIST files from all the jars cured the problem and allowed the class loader to use the Class-Path attributes correctly. MORAL: don't use jar file indexes unless there's a 1:1 mapping from package to jar, across all packages.	1	1	https://issues.jboss.org/browse/JBAS-1713
1463	JBAS-1711	18 months since last comment. Can we add this back on the roadmap, please, it's a trivial fix.	0	0	https://issues.jboss.org/browse/JBAS-1711
1464	JBAS-1706	Can you be more specific about the use-case you want to implement? Maybe, provide an example.	0	0	https://issues.jboss.org/browse/JBAS-1706
1465	JBAS-1691	This is your basic deploy the classes in two different classloaders and then do pass by reference problem, which worked fine until we decided to switch to Servlet spec classloading (JBAS-1691) in 4.0.2 rather than the JBoss flat classloading model. e.g. add the following code to your servlet: Object objRef = initial.lookup("gogis/tutorial/slsb/servlet/HelloWorld"); +out.println(HelloWorldHome.class.getClassLoader()); +out.println(objRef.getClass().getInterfaces()[0].getClassLoader()); The other fix is to enable call by value semantics: http://wiki.jboss.org/wiki/Wiki.jsp?page=ClassLoadingConfiguration P.S. It is a good idea to provide an example than actually works out-of-the-box. If we have to fix your example to even compile it (e.g. where is your classpath in build.xml?), it is likely to be ignored/rejected	1	1	https://issues.jboss.org/browse/JBAS-1691
1466	JBAS-1610	Implemented the ability to run jboss using the 'platform MBeanServer' on jdk1.5+ runtimes. Attribute ServerConfig.getPlatformMBeanServer() has been added to indicate if we are using the platform MBeanServer as the main jboss server. Both the 'jboss.platform.mbeanserver' property must be set, and the jvm must be jdk1.5+. to enable this feature. MBeanServerLocator will instead return the platform MBeanServer, in this case. To enable this feature use: './run -Djboss.platform.mbeanserver' and see the 'java.lang' and 'java.util.logging' domains appearing in the jmx-console	0	0	https://issues.jboss.org/browse/JBAS-1610

#	Requirement	Questions	Answered	Successfully Answered	URL
1467	JBAS-1610	This is not working for me with the 4.0.2/jdk1.5.0_03 release when the jmx jconsole management agent has been enabled by setting the com.sun.management.jmxremote system property. When this is set, I'm seeing this startup failure: [starksm@banshee9100 bin]$ run.sh -Djboss.platform.mbeanserver ========== JBoss Bootstrap Environment JBOSS_HOME: C:\cvs\Releases\jboss-4.0.2 JAVA: /usr/java/jdk1.5.0_03/bin/java JAVA_OPTS: -server -Xms128m -Xmx128m -Dcom.sun.management.jmxremote -Dprogram. name=run.sh CLASSPATH: C:\cvs\Releases\jboss-4.0.2\bin\run.jar;C:\usr\java\jdk1.5.0_03\lib \tools.jar ========== 14:31:28,078 INFO [Server] Starting JBoss (MX MicroKernel)... 14:31:28,078 INFO [Server] Release ID: JBoss [Zion] 4.0.2 (build: CVSTag=JBoss_ 4_0_2 date=200505022023) 14:31:28,078 INFO [Server] Home Dir: C:\cvs\Releases\jboss-4.0.2 14:31:28,078 INFO [Server] Home URL: file:/C:/cvs/Releases/jboss-4.0.2/ 14:31:28,078 INFO [Server] Library URL: file:/C:/cvs/Releases/jboss-4.0.2/lib/ 14:31:28,078 INFO [Server] Patch URL: null 14:31:28,078 INFO [Server] Server Name: default 14:31:28,078 INFO [Server] Server Home Dir: C:\cvs\Releases\jboss-4.0.2\server\ default 14:31:28,093 INFO [Server] Server Home URL: file:/C:/cvs/Releases/jboss-4.0.2/s erver/default/ 14:31:28,093 INFO [Server] Server Data Dir: C:\cvs\Releases\jboss-4.0.2\server\ default\data 14:31:28,093 INFO [Server] Server Temp Dir: C:\cvs\Releases\jboss-4.0.2\server\ default\tmp 14:31:28,093 INFO [Server] Server Config URL: file:/C:/cvs/Releases/jboss-4.0.2 /server/default/conf/ 14:31:28,093 INFO [Server] Server Library URL: file:/C:/cvs/Releases/jboss-4.0. 2/server/default/lib/ 14:31:28,093 INFO [Server] Root Deployment Filename: jboss-service.xml 14:31:28,093 INFO [Server] Starting General Purpose Architecture (GPA)... Failed to boot JBoss: javax.management.InstanceNotFoundException: JMImplementation:service=LoaderRepository,name=Default at com.sun.jmx.interceptor.DefaultMBeanServerInterceptor.getMBean(DefaultMBeanServerInterceptor.java:1010) at com.sun.jmx.interceptor.DefaultMBeanServerInterceptor.invoke(DefaultMBeanServerInterceptor.java:804) at com.sun.jmx.mbeanserver.JmxMBeanServer.invoke(JmxMBeanServer.java:784) at org.jboss.system.server.ServerImpl.initBootLibraries(ServerImpl.java:580) at org.jboss.system.server.ServerImpl.doStart(ServerImpl.java:378) at org.jboss.system.server.ServerImpl.start(ServerImpl.java:315) at org.jboss.Main.boot(Main.java:195) at org.jboss.Main$1.run(Main.java:463) at java.lang.Thread.run(Thread.java:595)	1	1	https://issues.jboss.org/browse/JBAS-1610
1468	JBAS-1610	The problem is the agent processing is getting invoked before the system property is set. If you set it directly on the java command, you then have the problem that the classpath is not correct because this is before even main() is invoked. ========== JBoss Bootstrap Environment JBOSS_HOME: /home/adrian/jboss-4.0/workspace/build/output/jboss-4.0.3RC1 JAVA: /home/adrian/downloads/jdk1.5.0_03/bin/java JAVA_OPTS: -server -Xms128m -Xmx128m -Dcom.sun.management.jmxremote -Djavax.management.builder.initial=org.jboss.mx.server.MBeanServerBuilderImpl -Dprogram.name=run.sh CLASSPATH: /home/adrian/jboss-4.0/workspace/build/output/jboss-4.0.3RC1/bin/run.jar:/home/adrian/downloads/jdk1.5.0_03/lib/tools.jar ========== javax.management.JMRuntimeException: Failed to load MBeanServerBuilder class org.jboss.mx.server.MBeanServerBuilderImpl: java.lang.ClassNotFoundException: org.jboss.mx.server.MBeanServerBuilderImpl at javax.management.MBeanServerFactory.checkMBeanServerBuilder(MBeanServerFactory.java:499) at javax.management.MBeanServerFactory.getNewMBeanServerBuilder(MBeanServerFactory.java:530) at javax.management.MBeanServerFactory.newMBeanServer(MBeanServerFactory.java:304) at javax.management.MBeanServerFactory.createMBeanServer(MBeanServerFactory.java:219) at javax.management.MBeanServerFactory.createMBeanServer(MBeanServerFactory.java:180) at sun.management.ManagementFactory.createPlatformMBeanServer(ManagementFactory.java:264) at java.lang.management.ManagementFactory.getPlatformMBeanServer(ManagementFactory.java:512) at sun.management.jmxremote.ConnectorBootstrap.startLocalConnectorServer(ConnectorBootstrap.java:385) at sun.management.Agent.premain(Agent.java:92) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at sun.management.Agent.startAgent(Agent.java:197) Caused by: java.lang.ClassNotFoundException: org.jboss.mx.server.MBeanServerBuilderImpl at java.net.URLClassLoader$1.run(URLClassLoader.java:200) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:188) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:268) at java.lang.ClassLoader.loadClass(ClassLoader.java:251) at javax.management.MBeanServerFactory.loadBuilderClass(MBeanServerFactory.java:441) at javax.management.MBeanServerFactory.checkMBeanServerBuilder(MBeanServerFactory.java:484) ... 13 more Exception thrown by the agent : javax.management.JMRuntimeException: Failed to load MBeanServerBuilder class org.jboss.mx.server.MBeanServerBuilderImpl: java.lang.ClassNotFoundException: org.jboss.mx.server.MBeanServerBuilderImpladrian@adrian-srv01:~/jboss-4.0/workspace/build/output/jboss-4.0.3RC1/bin> adrian@adrian-srv01:~/jboss-4.0/workspace/build/output/jboss-4.0.3RC1/bin>	1	1	https://issues.jboss.org/browse/JBAS-1610

#	Requirement	Questions	Answered	Successfully Answered	URL
1469	JBAS-1598	_____ From: Scott M Stark Sent: Monday, April 11, 2005 1:00 PM To: Clebert Suconic; 'qa@jboss.com' Subject: RE: Create a compatibility matrix which tracks JBoss client vs. server class versions Serialization should be mostly covered by the SerialVersionUIDUnitTestCase in the bundled testsuite. There needs to be basic testing of all j2ee interfaces by running a subset of the existing testuite modules against the current jboss server. This is not hard. 1. For each release of jboss there is a build of the release source. 2. The current jboss server branch is obtained and built. 3. The current jboss server branch binary is run as with the all configuration. 4. For each release of jboss: a. Run a subset of the testsuite unit tests from within the release build so that the release build binaries are used. 5. Merge the testsuite runs into a cross version compatibility matrix. The toplevel view is a simple pass/fail matrix which indicates whether all of the tests run using the x.y.z binaries passed. _____ From: Clebert Suconic Sent: Monday, April 11, 2005 10:52 AM To: Scott M Stark; 'qa@jboss.com' Subject: RE: Create a compatibility matrix which tracks JBoss client vs. server class versions We will need to test Serialization across versions or functionalities across version as well? I'm afraid about the answer you would have, as if we have to test functionality, the amount of work will be huge. What if the Serialization works, but it the interceptor on 4.0.1 puts something different on the Marshalling that the stubs are not prepared to receive. (Just wondering some breaking in the functionality) But If it's serialization only, we can do it fast. Clebert _____ From: Scott M Stark Sent: Monday, April 11, 2005 12:42 PM To: Clebert Suconic; 'qa@jboss.com' Subject: RE: Create a compatibility matrix which tracks JBoss client vs. server class versions No, that is not enough. We need an independent test matrix that is run as part of the nightly qa process that runs the testsuite using the x.y.z release jars against the current 4.0 branch to validate actual functional behavior. Serialization version ids are only one small part of version interop. Addition/removal of classes, and changes in implementation are the other big factors that break interop. The serialization version uid test is the only one that should be part of the normal development testsuite. The cross version matrix needs to be external and not part of the testuite cvs module so as not to bloat the download. _____ From: Clebert Suconic Sent: Monday, April 11, 2005 9:49 AM To: Scott M Stark; 'qa@jboss.com' Subject: RE: Create a compatibility matrix which tracks JBoss client vs. server class versions I can extend what you did into http://jira.jboss.com/jira/browse/JBAS-1584 Do you think that if I compare SerialUID from previous versions against actual version would be enough to keep this compatibility matrix, or should we be doing something beyond that? Also... I created a simple testcase that verifies if classes are defining UIDs. As cruisecontrol is broken for 4.0 now, I think people haven't look at the results. I think we still have classes not defining SerialUUID. My understanding is that this matrix is dependent of clearing SerialUUIDs first. Is that right? Clebert _____ From: Scott M Stark Sent: Monday, April 11, 2005 12:08 AM To: qa@jboss.com Subject: Create a compatibility matrix which tracks JBoss client vs. server class versions What is the status of this issue? http://jira.jboss.com/jira/browse/JBAS-1598	0	0	https://issues.jboss.org/browse/JBAS-1598
1470	JBAS-1598	For web service interoperability. All ws4ee-foo-client.jars must be deployed on a node of the same version as the testsuite running the tests. For that we could modify JBossTestCase to take a node identifier for the deploy(String nodeId, String deployments), undeploy(String nodeId, String deployments) methods. As an optimization for the non-interop scenario the client node does not need to be started and it's Id is simply assigned to that of the server node. I know we have the feature of starting up nodes from the testsuite. Can we do this optionally depending on client/server versions? How do we deploy to a specific node?	1	1	https://issues.jboss.org/browse/JBAS-1598
1471	JBAS-1595	Bill, do you want to reopen this task and talk to Ryan about JDK 5 build requirements? Is retroweaver not an option for JDK 1.4? Ivelin	1	1	https://issues.jboss.org/browse/JBAS-1595
1472	JBAS-1595	This has nothing to do with build requirements. It has to do with release requirements. Do we have a specific JDK 5 release of JB4?	0	0	https://issues.jboss.org/browse/JBAS-1595
1473	JBAS-1595	OK, Ryan is notified. Scott let's clarify the requirements: You want 2 builds for JBoss 4 - one compiling with JDK 1.4 and another one requiring JDK 5. The first one will skip the EJB3 module, the second one will include it. Is this correct? Should we version the one that includes EJB3 JBoss AS 4.5? Ivelin	1	1	https://issues.jboss.org/browse/JBAS-1595
1474	JBAS-1595	Is this really still a "critical" issue?	0	0	https://issues.jboss.org/browse/JBAS-1595
1475	JBAS-1585	Flow control needs a lot more thought than my off the cuff trivial test. e.g. The requests should not just be throttled they should be prioritized so a throttled request can give up its threads to others that can do work. For example a sender that wants to add a message gives up its thread to allow a receiver to clear some memory. This requires co-operation from the invoker framework. i.e. I am about to do perform a send, can I proceed? The link about message serialization is there because ideally during the same process, the sender can write the message body directly into persistent store using buffers, rather than hogging memory with a full read and deserialiation.	0	0	https://issues.jboss.org/browse/JBAS-1585
1476	JBAS-1525	We also see a problem if the userDN cannot directly be used to authenticate and check the pasword using the LDAP bind operation. We have LDAP entries which are aliases and they cannot be used to perform a bind. An extra technical user is needed to perform the bind and that user needs read access on the password field. The login module then can read the password hash the given user password and compare the two hashes. We created an "AdvancedLdapLoginModule" from the 4.0.1 src which is doing the job. This is something we would like to contribute. How can we do that?	1	1	https://issues.jboss.org/browse/JBAS-1525
1477	JBAS-1525	Marcus, why does the bind user option trigger pulling in the user password for validation inside the server rather than simply binding to the ldap server as the as the user with their credentials for authentication as currently done. I'm not inclined to introduce this ldap server specific password processing and userPassword attribute notion. The only point of the bind user for me was to allow for querying of the roles under a context that the user does not have rights to.	1	1	https://issues.jboss.org/browse/JBAS-1525
1478	JBAS-1525	So can you query for the user with referral resolution enabled and bind to the resulting ldap url for the user?	1	1	https://issues.jboss.org/browse/JBAS-1525
1479	JBAS-1475	Can you ping me on Yahoo IM (belaban@yahoo.com) so we can talk ?	1	1	https://issues.jboss.org/browse/JBAS-1475
1480	JBAS-1458	What do you mean by not being able to bind to any port? To bind to an anonymous port use an RMIObjectPort of 0, otherwise specify the desired port.	0	0	https://issues.jboss.org/browse/JBAS-1458
1481	JBAS-1458	We recently have been bitten by this. Our customers run our software (which bundles JBoss) with multiple instances on one *NIX server. They also run other third-party software to which they assign ports explicitly per instance. Because JBoss was picking up a handful random ports from the OS based on the default server config, we hit conflicts where JBoss was holding a port to be used by another service. Since we truly have no need for any remote access to these JBoss instances, we set about trying to disable all port bindings except port 80 for HTTP. We have been able to eliminate all of them except for the RMIObjectPort, and only because the EJB container can't seem to run without it. I see this issue has been in limbo for 3 years... any chance this could be addressed? Is it a trivial change? Would you accept a patch for this?	1	0	https://issues.jboss.org/browse/JBAS-1458

#	Requirement	Questions	Answered	Successfully Answered	URL
1482	JBAS-1458	This issue is out-of-date, so I'm closing it. As of 4.2.x EJB2 doesn't use the JRMP invoker, it uses the UnifiedInvoker (JBAS-3950) which can be changed to use some other protocol e.g. as per the feature request an "in memory transport". The only thing that does use it is the jmx-invoker-service.xml which can be removed or changed to use a different invoker, e.g. the "web" configuration in JBoss5 uses the http invoker instead. What is left of this feature request is now on JBAS-6071which would remove the legacy invokers by default (they are only required for backwards compatibility when talking to older versions of JBoss).	0	0	https://issues.jboss.org/browse/JBAS-1458
1483	JBAS-1443	I think an extra interceptor could work out nicely, any hints on how to implement this interceptor? Is there an interceptor we could use or use as a good base? We would need some kind of waiting/signaling to implement this. I'm not into developing interceptors, so I'm not sure how to do this best?	1	0	https://issues.jboss.org/browse/JBAS-1443
1484	JBAS-1436	While looking for a way to get wait times for connections logged I stumbled across this one. From the activity log it's unclear to me whether there is actually work in progress besides PreparedStatementPool. Is there a chance that we will be able to see logging of how long a thread waited for a connection in JB 4.2.0? I looked at InternalManagedConnectionPool but it seems there is logging of wait time only in case the thread is interrupted (via the exception) but not in the standard case.	0	0	https://issues.jboss.org/browse/JBAS-1436
1485	JBAS-1429	Does "Reauthentication Support" in this feature request mean support by application server or by JCA connection factory? It seems not to be distinguished here. Changes 1-3) are server specific and 4) is JCA provider specific. Changes 1-3 affects more than JDBC only, but any other JCA implementation as well. If you are going to solve the issue by introducing a new interface like ReauthenticationMachanism I doubt that any third-party JCA vendor will implement it; it would be a solution for JBoss' JCA implementation for JDBC only. For us it is important to have the server-side re-authentication support first (1-3), which will allow to use re-authentication provided by the other JCA vendors.	0	0	https://issues.jboss.org/browse/JBAS-1429
1486	JBAS-1426	Actually, after thinking about this is a bigger problem then I thought. Moving checkState() above setAutoCommit() won't help if the underlying connection's autoCommit==true. That would mean the solution on your side would be to: setAutoCommit(true); change readonly status; setAutoCommit(false); which looks more like a special C-JDBC hack. What do you think? In the mean time I'm going to lobby C-JDBC for a better transaction detection system. Mike	0	0	https://issues.jboss.org/browse/JBAS-1426
1487	JBAS-1402	Adrian, can we close this whole task & subtasks out-of-date, or do we keep the JCA tasks as still current?	1	0	https://issues.jboss.org/browse/JBAS-1402
1488	JBAS-1391	I was asked to look into this. Wouldn't it be easier to have a logical DataSource that would maintain a list of target datasources? The logical one would choose a target ds from the list and try to get a connection. If that one fails, it will choose another target, etc. This way we also don't have connections to different databases in the same pool.	1	1	https://issues.jboss.org/browse/JBAS-1391
1489	JBAS-1391	First, it is better to discuss the design in the dev forums. To answer your question Alex, it can be done both ways. The DataSource of DataSources idea is lot more consistent (you don't get a mixture of connections in the pool) in the event of failover, however it has a number of technical problems. 1) It cannot be extended easily to introduce load balancing later 2) Identifying that a DataSource is down is harder, versus the rar that can look at the real connection failures. e.g. maybe you can't get a connection because all the connections are checked out and you hit the timeout or some other temporary failure that should not cause failover of the DataSource? 3) Configuring mulitple datasources is a pain 4) How do you manage the dependencies of the DataSource, i.e. how does it wait for the sub DataSources, does it get undeployed if one of them goes away? 5) How do you manage the dependees of the DataSource? There are a number of services that depend upon the DataSourceBinding getting injected across JMX so they don't try to use the DataSource until it is available.	1	1	https://issues.jboss.org/browse/JBAS-1391
1490	JBAS-1325	I don't see the usecase for allowing the type to be externalized via a system property. Your not using this feature in the example. Is that really a desired feature?	0	0	https://issues.jboss.org/browse/JBAS-1325
1491	JBAS-1313	Hi, I've implemented this change request. A notification is installed into all subdeployers (the listeners are dynamically added/removed from/to the SubDeployers using the notifications from the MainDeployer). It seeks META-INF/log4j.xml / .prop in the classpath (using localCl), and configure the logger for the given deployment. It uses ClassLoaders to track deployments. It needs no code changes in any of the existing JBoss code. (Tested on 4.0.3, with EARs, WARs and spring deployer) One problem though, in an EAR deployment all subdeployments receives the classloader thus, one cannot create seperate logger for sub deployments. (eg. one for the whole EAR and a seperate for the WAR inside) Or a logger in a subdeployment would be used for the whole application. That's why this service is limited for 'root' deployments only at the moment. What's next? Cheers, Tamas	0	0	https://issues.jboss.org/browse/JBAS-1313
1492	JBAS-1310	Michael, I've tried your solution but I've run into some issues. Would you like to discuss this trough direct email or should I use the bug tracker for WfmOpen?	1	0	https://issues.jboss.org/browse/JBAS-1310
1493	JBAS-1292	Is this issue still current, or outdated?	0	0	https://issues.jboss.org/browse/JBAS-1292

#	Requirement	Questions	Answered	Successfully Answered	URL
1494	JBAS-1279	The shutdown command has been updated. [starksm@banshee9100 bin]$ shutdown.sh -S Exception in thread "main" java.lang.SecurityException: Failed to authenticate principal=null, securityDomain=jmx-console at org.jboss.jmx.connector.invoker.AuthenticationInterceptor.invoke(AuthenticationInterceptor.java:72) at org.jboss.mx.server.Invocation.invoke(Invocation.java:74) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:249) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:642) at org.jboss.invocation.jrmp.server.JRMPProxyFactory.invoke(JRMPProxyFactory.java:164) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:324) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:144) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:80) at org.jboss.mx.server.Invocation.invoke(Invocation.java:72) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:249) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:642) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:805) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:406) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:324) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:261) at sun.rmi.transport.Transport$1.run(Transport.java:148) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:144) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:460) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:701) at java.lang.Thread.run(Thread.java:534) at sun.rmi.transport.StreamRemoteCall.exceptionReceivedFromServer(StreamRemoteCall.java:247) at sun.rmi.transport.StreamRemoteCall.executeCall(StreamRemoteCall.java:223) at sun.rmi.server.UnicastRef.invoke(UnicastRef.java:133) at org.jboss.invocation.jrmp.server.JRMPInvoker_Stub.invoke(Unknown Source) at org.jboss.invocation.jrmp.interfaces.JRMPInvokerProxy.invoke(JRMPInvokerProxy.java:118) at org.jboss.invocation.InvokerInterceptor.invokeInvoker(InvokerInterceptor.java:163) at org.jboss.invocation.InvokerInterceptor.invoke(InvokerInterceptor.java:103) at org.jboss.jmx.connector.invoker.client.InvokerAdaptorClientInterceptor.invoke(InvokerAdaptorClientInterceptor.java:51) at org.jboss.proxy.SecurityInterceptor.invoke(SecurityInterceptor.java:55) at org.jboss.proxy.ClientMethodInterceptor.invoke(ClientMethodInterceptor.java:59) at org.jboss.proxy.ClientContainer.invoke(ClientContainer.java:91) at $Proxy0.invoke(Unknown Source) at org.jboss.Shutdown$ServerProxyHandler.invoke(Shutdown.java:240) at $Proxy1.shutdown(Unknown Source) at org.jboss.Shutdown.main(Shutdown.java:211) [starksm@banshee9100 bin]$ shutdown.sh -S -u admin -p admin Shutdown message has been posted to the server. Server shutdown may take a while - check logfiles for completion [starksm@banshee9100 bin]$	0	0	https://issues.jboss.org/browse/JBAS-1279
1495	JBAS-1279	Is this fixed in 4.0.5? I see this issue when JBoss Portal 2.6 is deployed in JBoss and I try running "shutdown.sh -S -u admin". I've used a custom security realm and updated jmx-invoker-service.xml accordingly.	0	0	https://issues.jboss.org/browse/JBAS-1279
1496	JBAS-1107	Any progress on this issue? It's just a matter of adding a few lines to jboss-app.dtd and doesn't hurt anybody. So why not just do it?	1	1	https://issues.jboss.org/browse/JBAS-1107
1497	JBAS-1091	You are kidding, aren't you? If JBoss does not persist timers across shutdowns and you do not consider this a bug, I'll switch to another AS for J2EE 1.4 immediately. From the ejb 2.1 spec: "The timer service is intended for the modelling of long-lived business processes. Timers survive container crashes and the activation/passivation and load/store cycles of the enterprise beans that are registered with them." Now, wouldn't you agree that the shutdown of the AS implies a "passivate/store" of the EJB associated with the timer? And that the restart of the server implies a "activate/load"? This logically implies that the timers must be persisted across shutdowns. Anything else would be ridiculous.	0	0	https://issues.jboss.org/browse/JBAS-1091
1498	JBAS-1091	I took the statement in the tutorial as a reason to write to Sun. I got an answer written by Kenneth Saks (EJB engineering lead) and confirmed by Linda DeMichiel (EJB specification lead). Kenneth says: "... While the spec doesn't explicitly use the phrase "on server shutdown", that was certainly my understanding of the intent. What's the purpose of timers surviving crashes only and not regular shutdowns? If the deployer can't rely on the existing timers being active until they naturally expire or are explicitly removed then they're not useful. ..." Of course, neither the section in the tutorial nor the quoted anser is normative. But the JBoss team seems to allow reason to prevail after all anyway.	0	0	https://issues.jboss.org/browse/JBAS-1091
1499	JBAS-1091	Have you tried with the .zip distro?	1	1	https://issues.jboss.org/browse/JBAS-1091
1500	JBAS-1091	Yes, I have tried full .zip(well, really .tar.bz2) version of RC1 with EJB3-RC8 applied, because it had not ejb3 support by default. But do Timer classes intersect with ejb3 distro jars? Is it important that I am using simple .ejb3 instead of complex .ear packaging? By the way, I have also tried modified version of ejb3_tutorial/timer sample modified to have 5e+6 msec timeout - same result: the item from "timers" table is deleted on server shutdown.	1	1	https://issues.jboss.org/browse/JBAS-1091
1501	JBAS-1091	Well, I'll be thankful if you could give me a link to your code to patch 4.0.4-GA with so I could recompile it. How long could it take to build jboss ejb3 profile on athlon64 3800+ x2 ? Thanks.	1	1	https://issues.jboss.org/browse/JBAS-1091
1502	JBAS-1091	I thought TimerService -related classes themselves are responsible for timer persistence, isn't it so? If not - I'll try to contact ejb3 developers how to sync their code with base profile accordingly. Could you please point out what classes you have worked on to make timer persistent. Thanks.	1	1	https://issues.jboss.org/browse/JBAS-1091
1503	JBAS-1010	SourceForge Username: smapjb . Logged In: YES user_id=559447 I didn't make myself clear.. I meant resource-ref (a URL resource ref) that is not specific to a component.. and the deployment descriptors that I refer to are the jboss ones. I know that you don't define URLs in the J2EE ones. The matter still stands.. All other resource references are defined outside the EAR.. but the URL *resource ref* has to be defined inside the *jboss* deployment descriptors and hence inside the EAR.. this is not consistent and other app servers (specifically Websphere) don't do it that way. Please reread my original post ignoring 'environment'. Apologies for the confusion.	1	1	https://issues.jboss.org/browse/JBAS-1010

#	Requirement	Questions	Answered	Successfully Answered	URL
1504	JBAS-1010	Logged In: YES user_id=175228 I still don't know what your talking about. How would websphere map the following ejb-jar resource-refs to the deployment environment JNDI location? <ejb-jar> <display-name>ENC Tests</display-name> <enterprise-beans> <session> <!-- JDBC DataSources (java:comp/env/jdbc) --> <resource-ref> <description>The default DS</description> <res-ref-name>jdbc/DefaultDS</res-ref-name> <res-type>javax.sql.DataSource</res-type> <res-auth>Container</res-auth> </resource-ref> <!-- JavaMail Connection Factories (java:comp/env/mail) --> <resource-ref> <description>Default Mail</description> <res-ref-name>mail/DefaultMail</res-ref-name> <res-type>javax.mail.Session</res-type> <res-auth>Container</res-auth> </resource-ref> <!-- URL Connection Factories (java:comp/env/url) --> <resource-ref> <description>The JBoss Web Site HomePage</description> <res-ref-name>url/JBossHomePage</res-ref-name> <res-type>java.net.URL</res-type> <res-auth>Container</res-auth> </resource-ref> </session> </enterprise-beans> </ejb-jar>	1	1	https://issues.jboss.org/browse/JBAS-1010
1505	JBAS-1010	SourceForge Username: smapjb . Logged In: YES user_id=559447 Apologies for reopening this but I submitted a patch and noone noticed!! I have written the URLResource service that compliments this feature.. you can now register your URLs in a url-service.xml file and hot deploy them.. finishing off this RFE. If you close this RFE then you will probably also want to close this patch http://sourceforge.net/tracker/index.php?func=detail&aid=1069543&group_id=22866&atid=376687	1	1	https://issues.jboss.org/browse/JBAS-1010
1506	JBAS-855	SourceForge Username: jim_b_o . Logged In: YES user_id=960535 What tools do you use for creating diffs? The reason for setting the identity per statement is that we cannot do it on getConnection(). The third party tools that we use do not request the connection under the identity that is actually going to use the connection. At least that is how it appears from the callkback events which it raises. I'd be happy to try what you're recommending, to check if it's actually do something different under the covers.	1	1	https://issues.jboss.org/browse/JBAS-855
1507	JBAS-855	Logged In: YES user_id=9459 cvs diff -u 5 Will produce a universal diff against the code you checked out with a context of 5 lines either side of the changes. I'm not sure I understand your comment about the thirdparty tool. Can you give an example of the use case? With managed connections, there is a real connection and a handle (the connection object you actually use). The handle is only valid for the time you have a connection checked out from the pool. Doesn't the identity remain the same for the lifetime of that checked out handle, probably the duration of the transaction. Regards, Adrian	1	1	https://issues.jboss.org/browse/JBAS-855
1508	JBAS-855	SourceForge Username: jim_b_o . Logged In: YES user_id=960535 If there's anybody watching... Is the latest attachment OK? Is it likely this will be incorporated in a future release? Can I do anything else to assist?	1	1	https://issues.jboss.org/browse/JBAS-855
1509	JBAS-826	SourceForge Username: pilhuhn . Logged In: YES user_id=217112 The posting of scott were the relese notes btw. See: http://sourceforge.net/docman/display_doc.php?docid=20433&group_id=22866	1	0	https://issues.jboss.org/browse/JBAS-826
1510	JBAS-804	SourceForge Username: drmlipp . Logged In: YES user_id=731113 I'm sorry, but what kind of argument is this? Because some developer might not implement the container/EJB contract correctly, JBoss is entitled to violate the client/container contract? You acknowledge that the client must act differently depending on the kind of problem that has caused the exception. Of course, I have implemented looking at the cause as a workaround. But this is highly JBoss specific (and really, while looking at an exception's cause may help to decide what to do next in some cases, this is clearly a situation where you should be able to distinguish the problem cause "first-hand", i.e. without having to dig down). We maintain an application for both JBoss and Bea WLS here and WLS acts exactly as I had excpected (i.e. throws java.rmi.noSuchObjectException) while JBoss throws an exception that -- from the documentation -- is simply not appropriate. Why should I have to implement workarounds when using JBoss just because some developer cannot implement his finder methods properly?	1	0	https://issues.jboss.org/browse/JBAS-804
1511	JBAS-763	SourceForge Username: pilhuhn . Logged In: YES user_id=217112 This should (at least in theory) be adressed at the Xdoclet guys. Gavin has access to the Xdoclet CVS; so we might try :-)	1	1	https://issues.jboss.org/browse/JBAS-763
1512	JBAS-714	Logged In: YES user_id=175239 Why do you need this? The libraries under jboss/lib should be a fixed set necessary to bootstrap the server. Each configuration can have their own specific set of libraries added to the classloader repository by adding them either to jboss/server/<configname>/lib (where all libraries are added) or jboss/server/<configname>/deploy (in case the library should be hot deployable).	0	0	https://issues.jboss.org/browse/JBAS-714
1513	JBAS-714	SourceForge Username: rmcauble . Logged In: YES user_id=378285 The reason is because in order to workaround the ClassCircularityError which comes up from time to time when using the UnifiedClassLoader, I've had to move everything from jboss/server/<configname>/lib to jboss/lib and move all the jars from all the sars to jboss/lib. (It's safe to leave the jars in the wars since they use WebClassLoader instead of UnifiedClassLoader). Obviously it would be preferable to fix the ClassCircularityError, but there is no way to fix that until Sun fixes the JVM. In the meantime anyone who uses this workaround just won't be able to hot-deploy any sars. Obviously hot-deploy is a nice convenience, but we were seeing the ClassCircularityError pretty frequently. It's just nice to have a simple workaround as an alternative. For more information, see the email thread associated with this message: http://www.mail-archive.com/jboss- user@lists.sourceforge.net/msg32109.html Note that I've also submitted a patch for this: 813664.	1	1	https://issues.jboss.org/browse/JBAS-714
1514	JBAS-664	SourceForge Username: elfring . Logged In: YES user_id=572001 How do you think about a cooperation with the "Open Service Gateway"? http://en.wikipedia.org/wiki/OSGi	1	0	https://issues.jboss.org/browse/JBAS-664
1515	JBAS-590	SourceForge Username: rombergc . Logged In: YES user_id=808269 Hi, the specification of container managed security was done unintentionally, after removal of the security related element your descriptor worked fine. Much thanks! I've attached the final descriptor version. One additional question: Is it possible to include the descriptor into the RAR-file itself? Best regards, Christian Romberg	1	1	https://issues.jboss.org/browse/JBAS-590
1516	JBAS-460	SourceForge Username: d_jencks . Logged In: YES user_id=60525 I don't understand the problem you describe. There is no mixed content in the depends or depends-list-element elements that I can see. depends can have one object name or one embedded mbean config, and so can depends-list-element. If I misunderstand please give a concrete example of the mixed content. The problem I was originally referring to is not related to depends elements. You can have an mbean attribute of type Element (or dom4j element, in jboss 4). Since the structure of the element depends on what the mbean expects, the jboss service.dtd can't validate it. Do you know what happens if an element is wrapped in CDATA? Would this fix the problem?	1	1	https://issues.jboss.org/browse/JBAS-460

#	Requirement	Questions	Answered	Successfully Answered	URL
1517	JBAS-434	SourceForge Username: ikar . Logged In: YES user_id=742459 Basicly what we need is the unpacked SAR+WebDAV deployment, (other types of archives also, but later) so that one could change some attribute/add mbean or whatever and this was understood by several JBosses running on remote boxes. Changing something in the packed SAR is something beyond the level of expertise of our tech support guys. I'll probably take a closer look whether it's a bug or not, but I've only a week of experience with JBoss so don't count on me much. Could you tell me one thing - is SARDeployer capable of deploying the URL like http://whatever/webdav/deploy/my.sar/ ? This means that SARDeployer must be WebDAV-aware, it must know what to do if it realizes that '/' at the end indicates that it's a web folder and it must use WebDAV.. URLDeploymentScanner is made webdav-aware, what about SARDeployer?	1	0	https://issues.jboss.org/browse/JBAS-434
1518	JBAS-351	SourceForge Username: mattmunz . Logged In: YES user_id=340814 Bill, Cool. I'm thinking of using this in the process of development. At the moment, I'm working on MBean Persistence, so I don't really need all the EJB and .Net stuff, just MBeans and the JMX Console. Would your tool work for this purpose? If so, is it only available through the subscription? Also, I imagine that this feature would be useful in an administrative scenario, used as a package management tool. Could it be used for this purpose as well? - Matt	1	1	https://issues.jboss.org/browse/JBAS-351
1519	JBAS-351	SourceForge Username: mattmunz . Logged In: YES user_id=340814 Igor, Thanks for pointing out the lego script for me. It's a great start, and I'm finding it really useful right away. Perhaps it should be used to generate the minimal and default configurations as well? Is there any plan to do this? - Matt	1	1	https://issues.jboss.org/browse/JBAS-351
1520	JBAS-286	SourceForge Username: dsundstrom . Logged In: YES user_id=251431 Ok, so you want to limit the number of guys put in the read-ahead cache. Are you sure you wan to limit this on a pre relationhship basis? I would guess a per bean basis is more desireable. Maybe a configurable (optional) shared cache, so we can say no more then n across all beans readahead per tx. Do you just want to cap it at x or should we use an LRU rule?	1	1	https://issues.jboss.org/browse/JBAS-286
1521	JBAS-182	SourceForge Username: mkotsbak . Logged In: YES user_id=366650 Have you checked jboss.net? It provides support for webservices for beans almost without any extra effort. I think it lacks some security-features, and UDDI, but it works.	1	1	https://issues.jboss.org/browse/JBAS-182
1522	JBAS-156	SourceForge Username: d_jencks . Logged In: YES user_id=60525 Could you go into more detail about what you have in mind? Are you talking about the jboss-jdbc.rar wrapper? XADataSourceImpl?? I think it is unlikely that any changes would be made to the XADataSourceImpl from the 2.4 series since that has been completely eliminated from 3.0. What is LINTER RDMS (link, please)	1	0	https://issues.jboss.org/browse/JBAS-156
1523	JBAS-112	SourceForge Username: nobody . Logged In: NO How does one register for this task - I'd like to have a crack at it!	1	1	https://issues.jboss.org/browse/JBAS-112
1524	JBAS-104	SourceForge Username: khollund . Logged In: YES user_id=316655 For what server .. the jnp - or are you thinking about the webservice/containter you are running ? Knut Erik	1	1	https://issues.jboss.org/browse/JBAS-104
1525	JBAS-98	SourceForge Username: schaefera . Logged In: YES user_id=70434 Where is the problem on the Server Side (where the connector is conntected to) or on the Client Side where you use the connector ? Is the JNDI naming or MBean naming the problem ? Please send your code to me at: aschaefer@ebuilt.net and I will look at your changes and if OK upload to CVS.	1	1	https://issues.jboss.org/browse/JBAS-98
1526	JBAS-70	Is this already part of the security auditing features the security team has added?	1	1	https://issues.jboss.org/browse/JBAS-70
1527	JBAS-59	What other work relates to this task? The only subtask is scheduled for CR1.	0	0	https://issues.jboss.org/browse/JBAS-59
1528	JBAS-29	Perhaps so.. but then why not apply the same logic to mail, datasource and jms resources. I should be able to configure my URL resources outside of my EAR as I can with other app servers.. I'll include a screenshot of websphere. Can you include this patch now and worry about the more general case later. I have followed the same pattern as for the mail resource. At the mo I am having to distribute a patched jboss to our customers which is untidy.	1	1	https://issues.jboss.org/browse/JBAS-29
1529	JBAS-29	Just for clarity... Generalize it to include non trivial urls? or Generalize it as per your comment on 21 Feb? If the latter, then can you elaborate the requirements? they are a little succinct.	1	1	https://issues.jboss.org/browse/JBAS-29
1530	AS7-6776	Maybe the ServletContainerInitializerDeploymentProcessor is not using the service loader the way it should. Can you reassign ?	0	0	https://issues.jboss.org/browse/AS7-6776
1531	AS7-6776	Looks like I have a fix. I just remove the attachment in ServletContainerInitializerDeploymentProcessor and the problem goes away. public void undeploy(final DeploymentUnit context) { context.removeAttachment(ScisMetaData.ATTACHMENT_KEY); } I don't know if this is really fixing the root problem though. What do you think?	0	0	https://issues.jboss.org/browse/AS7-6776
1532	AS7-6755	Hi Stuart, I've submitted the PR here: https://github.com/jbossas/jboss-eap/pull/88/files Could you please help to check if my modifications are enough? Thanks!	0	0	https://issues.jboss.org/browse/AS7-6755
1533	AS7-6548	Can you define what you mean by "unstable" and also post precise steps to reproduce the issue? I don't understand what the problem is.	0	0	https://issues.jboss.org/browse/AS7-6548
1534	AS7-6548	This doesn't seem to be a clustering issue, but rather a domain controller issue. Have you tried this with the latest 7.1.x release?	0	0	https://issues.jboss.org/browse/AS7-6548

#	Requirement	Questions	Answered	Successfully Answered	URL
1535	AS7-6511	Sample stack trace looks like this (the assertion that bundle is ACTIVE fails): --- java.lang.AssertionError: expected:<32> but was:<4> at org.junit.Assert.fail(Assert.java:93) at org.junit.Assert.failNotEquals(Assert.java:647) at org.junit.Assert.assertEquals(Assert.java:128) at org.junit.Assert.assertEquals(Assert.java:472) at org.junit.Assert.assertEquals(Assert.java:456) at org.jboss.as.test.embedded.osgi.bootstrap.StandaloneBootstrapTestCase.testRestartTwoHostsOneFragment(StandaloneBootstrapTestCase.java:78) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:47) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:50) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:222) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:234) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:133) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:114) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:188) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:166) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:86) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:101) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:74)	0	0	https://issues.jboss.org/browse/AS7-6511
1536	AS7-6489	Was the username and password the same from installation to installation?	1	1	https://issues.jboss.org/browse/AS7-6489
1537	AS7-6481	Won't this cover your request https://github.com/aloubyansky/jboss-as/commit/55bd9ea33ae4e8bf7580810ea8baa19ac5a445dc ? It can be configured in jboss-cli.xml with connection-timeout element and/or specified as the launching script argument --timeout=xxx.	1	1	https://issues.jboss.org/browse/AS7-6481
1538	AS7-6481	Ahh yes that does indeed address it. Sorry about that, I was doing find usages on the wrong method in the ModelControllerClient.Factory. What's the policy on closing jiras, is that something that I should do?	0	0	https://issues.jboss.org/browse/AS7-6481
1539	AS7-6481	How about we resolve it as Done, so it shows up as such in the release notes, searches etc? Since there's no other JIRA for the work. The URL for the commit in upstream master is https://github.com/jbossas/jboss-as/commit/55bd9ea33ae4e8bf7580810ea8baa19ac5a445dc, so people interested in the actual patch can find it via the comments in this issue.	1	1	https://issues.jboss.org/browse/AS7-6481
1540	AS7-6423	[~rafaelliu]: Can you open a support case about this so that we can get this fixed in a future release? This JIRA tracker is not the correct location. In the support case, you can tell them that Jimmy Wilson asked you to do this.	1	0	https://issues.jboss.org/browse/AS7-6423
1541	AS7-6386	Unassigned? Blocker for the release?	1	1	https://issues.jboss.org/browse/AS7-6386
1542	AS7-6377	Is this also true for a deployed war on a specific server instance? For the operation "/host=master/server=lunar12/deployment=jboss-as-login.war/subsystem=web:read-resource" the URL should be unambiguous?	0	0	https://issues.jboss.org/browse/AS7-6377
1543	AS7-6377	Ok, I see. How about adding a sub resource with the available URLs? Does that make sense for you? We have already some views in the admin console where the user can follow links and it would definitely make sense if the user can browse the deployed web applications. What do you think?	1	1	https://issues.jboss.org/browse/AS7-6377
1544	AS7-6377	There are so many case where the url could be wrong that I would prefer not to have this feature. Looking to my previous comment how do you provide the url if the user is starting AS with bin/standalone.sh -b myhost? How do I get myhost? In case of an AJP connector how to guess the httpd server name?	1	0	https://issues.jboss.org/browse/AS7-6377
1545	AS7-6350	Paul, what's the scenario? Wouldn't it be easier to rename the jboss-as-* dir at the end of the build and integration test process?	1	1	https://issues.jboss.org/browse/AS7-6350
1546	AS7-6346	Isn't this done?	1	1	https://issues.jboss.org/browse/AS7-6346
1547	AS7-6303	Will this make it until Friday?	1	1	https://issues.jboss.org/browse/AS7-6303
1548	AS7-6264	I guess the fact that i can loads classes from my deployment and i can't load the same classes from the netty handler i would have considered as an issue which could be a bug. Or is this normal that two parts of a deployment are seeing different class pathes?	0	0	https://issues.jboss.org/browse/AS7-6264
1549	AS7-6179	{quote} <hardy> also here https://github.com/jbossas/jboss-as/blob/master/jpa/core/src/main/java/org/jboss/as/jpa/validator/JPALazyValidatorFactory.java#L80 <hardy> the other thing I am wondering is whether it is correct to always initialise Hibernate Validator <smarlow> hardy: should I do that as part of the upgrade? removing the config.addMapping(new ConstraintMapping()); <hardy> this one you can do as part of the upgrade for sure <hardy> this line is no good for anything really {quote}	1	1	https://issues.jboss.org/browse/AS7-6179
1550	AS7-6179	What are the module dependencies for jsoup?	1	1	https://issues.jboss.org/browse/AS7-6179

#	Requirement	Questions	Answered	Successfully Answered	URL
1551	AS7-6152	{noformat} java.lang.NullPointerException at org.jboss.as.console.client.domain.groups.deployment.ContentRepositoryView$3.getValue(ContentRepositoryView.java:134) at org.jboss.as.console.client.domain.groups.deployment.ContentRepositoryView$3.getValue(ContentRepositoryView.java:1) at com.google.gwt.user.cellview.client.Column.render(Column.java:165) at com.google.gwt.user.cellview.client.AbstractCellTableBuilder.renderCell(AbstractCellTableBuilder.java:206) at com.google.gwt.user.cellview.client.DefaultCellTableBuilder.buildRowImpl(DefaultCellTableBuilder.java:129) at com.google.gwt.user.cellview.client.AbstractCellTableBuilder.buildRow(AbstractCellTableBuilder.java:87) at com.google.gwt.user.cellview.client.AbstractCellTable.buildRowValues(AbstractCellTable.java:2347) at com.google.gwt.user.cellview.client.AbstractCellTable.replaceAllChildren(AbstractCellTable.java:2148) at com.google.gwt.user.cellview.client.AbstractHasData$View.replaceAllChildren(AbstractHasData.java:245) at com.google.gwt.user.cellview.client.HasDataPresenter.resolvePendingState(HasDataPresenter.java:1351) at com.google.gwt.user.cellview.client.HasDataPresenter.access$3(HasDataPresenter.java:1062) at com.google.gwt.user.cellview.client.HasDataPresenter$2.execute(HasDataPresenter.java:984) at com.google.gwt.core.client.impl.SchedulerImpl$Task$.executeScheduled$(SchedulerImpl.java:50) at com.google.gwt.core.client.impl.SchedulerImpl.runScheduledTasks(SchedulerImpl.java:228) at com.google.gwt.core.client.impl.SchedulerImpl.flushFinallyCommands(SchedulerImpl.java:327) at com.google.gwt.core.client.impl.Impl.exit(Impl.java:266) at com.google.gwt.core.client.impl.Impl.entry0(Impl.java:257) at sun.reflect.GeneratedMethodAccessor67.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at com.google.gwt.dev.shell.MethodAdaptor.invoke(MethodAdaptor.java:103) at com.google.gwt.dev.shell.MethodDispatch.invoke(MethodDispatch.java:71) at com.google.gwt.dev.shell.OophmSessionHandler.invoke(OophmSessionHandler.java:172) at com.google.gwt.dev.shell.BrowserChannelServer.reactToMessages(BrowserChannelServer.java:293) at com.google.gwt.dev.shell.BrowserChannelServer.processConnection(BrowserChannelServer.java:547) at com.google.gwt.dev.shell.BrowserChannelServer.run(BrowserChannelServer.java:364) at java.lang.Thread.run(Thread.java:680) {noformat}	0	0	https://issues.jboss.org/browse/AS7-6152
1552	AS7-6136	Perfect. Do you know if that will make 7.1.4?	1	0	https://issues.jboss.org/browse/AS7-6136
1553	AS7-6120	Jeff, could you have a look at the various 'entries' attributes in messaging? In a standalone server it looks like those could support expressions (and we allow them from jndi names elsewhere in the AS). Biggest question is how big of a PITA it would be to add the transformation and testing thereof.	0	0	https://issues.jboss.org/browse/AS7-6120
1554	AS7-6084	Why is this jar part of public API javadoc?	0	0	https://issues.jboss.org/browse/AS7-6084
1555	AS7-6084	Have you tracked down how the faulty jar was generated as faulty in the first place? If we have a problem with the HTTP Server build that should be fixed as well otherwise this may come back on the next release.	1	1	https://issues.jboss.org/browse/AS7-6084
1556	AS7-6002	why did you unset the fix version?	1	1	https://issues.jboss.org/browse/AS7-6002
1557	AS7-5980	The current JBoss http connector is not useful for high load environments. Tunable KeepAlive timeout is required, as are all of the other parameters mentioned (and more). This is why all other web servers allow them to be tuned. Did all of these other web server maintainers get it wrong? Apache, nginx, tomcat, glassfish, ... You should send these guys an email to let them know that their implementations are needlessly complex and could be solved with one size fits all defaults like 15s for KeepAlive. I'm sure they would be happy to learn of this wisdom.	1	0	https://issues.jboss.org/browse/AS7-5980
1558	AS7-5960	Hi Pete, I'm not an Eclipse user so I'm not sure what kind of filter you are talking about. Can you post a screen shot that shows the problem in CLI-GUI? Then we can talk about how the filter might work. Thanks, Stan	1	1	https://issues.jboss.org/browse/AS7-5960
1559	AS7-5960	Thanks Pete. One problem we have is that the tree is updated dynamically. So when you first open CLI GUI, we only know about the root and its children. A filter could not do a search beyond what is currently displayed unless we cache the whole tree. I really don't want to make a change like that right now. Do you think a filter that is limited to the root's children would be useful?	1	1	https://issues.jboss.org/browse/AS7-5960
1560	AS7-5960	Here is the simple filter. What do you think? !cli-gui-filter.png! It's only filtering the children of the root. I thought of ways to make it more robust, but it gets too complicated. I think this did turn out to be a big usability improvement. Thanks for the suggestion! Stan	1	1	https://issues.jboss.org/browse/AS7-5960
1561	AS7-5952	Than obvious question is: Why this issue hasnt been assigned to someone?	1	1	https://issues.jboss.org/browse/AS7-5952
1562	AS7-5923	I replaced the JBoss version of Mojarra with the latest regular Mojarra version and I still see the problem. I think it's probably a Weld issue. Jozef, can you take a look?	1	1	https://issues.jboss.org/browse/AS7-5923
1563	AS7-5923	While debugging this I found - https://issues.jboss.org/browse/WELD-909 Stan did you test against 7.1.1 which would have had Weld 1.1.5?	1	1	https://issues.jboss.org/browse/AS7-5923
1564	AS7-5844	Ondra: {quote} When I use JUnit's categories for Arq tests, then the class needs to be added to the deployment. jar.addPackage(CommonCriteria.class.getPackage()); Is there some more convenient solution? E.g. could I somehow instruct arq to wrap certain package to any deployment it does? {quote} Aslak: {quote} You can do this today if you want to, same as the weld core tests do. Create a little AuxiliaryArchiveAppender that packages the categories. As long as that is registered in a LoadableExtension and register as SPI and on Classpath when you run, the produced Archive will be bundled with the deployment. AuxiliaryArchiveAppender example: https://github.com/weld/core/blob/master/tests-arquillian/src/test/java/org/jboss/weld/tests/CategoryArchiveAppender.java LoadableExtension example: (only the service AuxiliaryArchiveAppender is needed in this case) https://github.com/weld/core/blob/master/tests-arquillian/src/test/java/org/jboss/weld/tests/WeldCategoryExtension.java SPI registration: https://github.com/weld/core/blob/master/tests-arquillian/src/test/resources/META-INF/services/org.jboss.arquillian.core.spi.LoadableExtension {quote}	0	0	https://issues.jboss.org/browse/AS7-5844

#	Requirement	Questions	Answered	Successfully Answered	URL
1565	AS7-5844	The idea/requests behind this is that packages only allow to divide tests into a tree. But sometimes, as Stuart wrote, there are some cross-cutting concerns, see the suggested categories. It's not supposed to replace package-based fine-grained division, but to complement it, to allow define groups of tests they need to run, instead of creating and maintaining list of packages or even classes. Another example where current model has some flaws are smoke tests: Separating the tests physically to different module is artificial and not much semantic: If any tests starts being considered as smoke test, what we do? Move it to different module and change package name. Does that make sense? Wouldn't it be better to just add `@Category(Smoke.class)` ? It also collides with the technical solution of separating tests by type of server setup (multinode, ...). If we decided to include a multinode test as smoke, what would we do? Move it to smoke module and create another surefire execution? Doesn't make much sense. Just adding the `Smoke` category and filtering the tests to be run using categories is way better. Last but not least: These categories are totally optional, no one has to go through the tests and add them. It's a tool for use cases when this categorization is necessary, like, security-related tests, it's subset Common Criteria tests, JPA-related tests, "performance smoke tests", etc. Summary: Tests have some cross-cutting concerns, facetes, which can't be covered by a tree structure, by principle. Also, we are bound by technical solutions coming from Maven's and Arquillian's limitations. Having additional way to categorize can only help, without doing any harm.	0	0	https://issues.jboss.org/browse/AS7-5844
1566	AS7-5800	Chris, can you paste the stack trace or log message you are getting now?	0	0	https://issues.jboss.org/browse/AS7-5800
1567	AS7-5800	I agree on the possible leak of information by outputting the value. Thinking more about it, it's also unneeded to have the value. However, having the value *class name* is quite useful, so how about outputting that instead of the the value?	1	1	https://issues.jboss.org/browse/AS7-5800
1568	AS7-5800	How about the attribute name and the type of the value? That doesn't help much if it's a collection and the problem is in an element. Hmm, I have a vague memory of long ago writing some utility to analyze a value and help track this kind of thing down -- walk through a collection or something. Rado -- does that ring a bell? I have a vague feeling it wasn't ever part of any actual release though; just something someone could hack in somehow.	0	0	https://issues.jboss.org/browse/AS7-5800
1569	AS7-5797	Can you have the edit button at the top, and then that goes away once you get into edit mode? Just the save and cancel move to the bottom. Like mentioned, you might want to keep an "easy cancel" towards the top if the user clicked edit by accident.	0	0	https://issues.jboss.org/browse/AS7-5797
1570	AS7-5768	org.junit.Assert.fail(Assert.java:92) at org.junit.Assert.assertTrue(Assert.java:43) at org.junit.Assert.assertNotNull(Assert.java:526) at org.junit.Assert.assertNotNull(Assert.java:537) at org.jboss.as.test.integration.jca.moduledeployment.ModuleIJDeploymentTestCase.testConfiguration(ModuleIJDeploymentTestCase.java:218) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.jboss.arquillian.junit.Arquillian$6$1.invoke(Arquillian.java:270) at org.jboss.arquillian.container.test.impl.execution.LocalTestExecuter.execute(LocalTestExecuter.java:60) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:90) at org.jboss.arquillian.core.impl.EventContextImpl.invokeObservers(EventContextImpl.java:99) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:81) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:135) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:115) at org.jboss.arquillian.core.impl.EventImpl.fire(EventImpl.java:67) at org.jboss.arquillian.container.test.impl.execution.ContainerTestExecuter.execute(ContainerTestExecuter.java:38) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:90) at org.jboss.arquillian.core.impl.EventContextImpl.invokeObservers(EventContextImpl.java:99) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:81) at org.jboss.arquillian.test.impl.TestContextHandler.createTestContext(TestContextHandler.java:89) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:90) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.test.impl.TestContextHandler.createClassContext(TestContextHandler.java:75) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:90) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at	0	0	https://issues.jboss.org/browse/AS7-5768

#	Requirement	Questions	Answered	Successfully Answered	URL
1571	AS7-5758	It seems to follow the same pattern: 401, authentication, GET request, additional 401: {noformat} 13:01:09,790 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 16) GET /console/app/deferredjs/A985E29579FAB3C55245D1782654FB15/25.cache.js HTTP/1.<TRUNCATED> [200 OK] () 13:01:10,119 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 16) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:10,123 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 16) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:10,128 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 17) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:10,213 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 17) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:10,213 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 16) POST /management HTTP/1.1 [200 OK] () 13:01:10,214 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 16) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:10,217 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 16) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:10,321 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 17) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:10,322 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 16) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:10,323 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 18) POST /management HTTP/1.1 [200 OK] () 13:01:10,680 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 18) POST /management HTTP/1.1 [200 OK] () 13:01:13,916 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 18) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:17,616 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 18) POST /management HTTP/1.1 [200 OK] () 13:01:20,222 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 18) POST /management HTTP/1.1 [200 OK] () 13:01:23,941 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 18) POST /management HTTP/1.1 [401 Unauthorized] () 13:01:26,478 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 18) POST /management HTTP/1.1 [200 OK] () 13:01:28,027 FINE [com.sun.net.httpserver] (HttpManagementService-threads - 18) POST /management HTTP/1.1 [200 OK] () {noformat} Does that ring a bell?	1	1	https://issues.jboss.org/browse/AS7-5758
1572	AS7-5758	Home many password prompts are you seeing? I would be expecting additional 401 responses as that is our protection against replay attacks - however the additional 401 responses should be telling the browser that the password is valid but the nonce is no longer valid. It also may be that related to this Safari is not re-using the existing connection which would also contribute to so many challenges in a short space of time.	1	1	https://issues.jboss.org/browse/AS7-5758
1573	AS7-5758	Does this problem only occur when the server is bound to a loopback device?	1	1	https://issues.jboss.org/browse/AS7-5758
1574	AS7-5758	By chance do you have any special safari plugins installed. Does google chrome work for you?	0	0	https://issues.jboss.org/browse/AS7-5758
1575	AS7-5758	Ok the reason I ask is they are both built on the webkit engine. My wire dump looks nothing like yours, which has some really strange things happening. The most interesting is the 40 packets from port 51543. The client repeating the same request, and each time being told it requires auth, and yet it never provides auth info. A later connection tries to send a user of "admin", but that makes no sense unless you have cached credentials somewhere. Here is another experiment; Can you try creating a new clean user on your mac, and login as that clean user, and see if the same thing happens?	0	0	https://issues.jboss.org/browse/AS7-5758
1576	AS7-5729	Do you think this is a question of implementing more logic in mod_cluster itself? In my opinion, the problem of balancer itself being a single point of failure should be addressed on DNS (DNS balancing) level. Hmm... I would advocate for KISS principle whenever possible :-)	0	0	https://issues.jboss.org/browse/AS7-5729
1577	AS7-5578	can you elaborate?	1	0	https://issues.jboss.org/browse/AS7-5578
1578	AS7-5544	What is the bare minimum of bundles that need to be deployed to make the console available? Authentication needs to integrate with the security subsystem. If possible I suggest to leave that out of the picture for now.	0	0	https://issues.jboss.org/browse/AS7-5544
1579	AS7-5544	If opening an SSH socket on the server is the only option, we need to clear that one up first. Do we have alternatives?	1	1	https://issues.jboss.org/browse/AS7-5544
1580	AS7-5544	[org.apache.aries.blueprint.container.BlueprintContainerImpl] (Blueprint Extender: 3) Unable to start blueprint container for bundle org.apache.karaf.jaas.config: org.osgi.service.blueprint.container.ComponentDefinitionException: Unable to intialize bean proxyLoginModuleInitializer at org.apache.aries.blueprint.container.BeanRecipe.runBeanProcInit(BeanRecipe.java:640) at org.apache.aries.blueprint.container.BeanRecipe.internalCreate(BeanRecipe.java:726) at org.apache.aries.blueprint.di.AbstractRecipe.create(AbstractRecipe.java:64) at org.apache.aries.blueprint.container.BlueprintRepository.createInstances(BlueprintRepository.java:219) at org.apache.aries.blueprint.container.BlueprintRepository.createAll(BlueprintRepository.java:147) at org.apache.aries.blueprint.container.BlueprintContainerImpl.instantiateEagerComponents(BlueprintContainerImpl.java:631) at org.apache.aries.blueprint.container.BlueprintContainerImpl.doRun(BlueprintContainerImpl.java:337) at org.apache.aries.blueprint.container.BlueprintContainerImpl.run(BlueprintContainerImpl.java:230) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:439) [rt.jar:1.6.0_39] at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) [rt.jar:1.6.0_39] at java.util.concurrent.FutureTask.run(FutureTask.java:138) [rt.jar:1.6.0_39] at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.access$301(ScheduledThreadPoolExecutor.java:98) [rt.jar:1.6.0_39] at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.run(ScheduledThreadPoolExecutor.java:206) [rt.jar:1.6.0_39] at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:895) [rt.jar:1.6.0_39] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:918) [rt.jar:1.6.0_39] at java.lang.Thread.run(Thread.java:662) [rt.jar:1.6.0_39] Caused by: java.lang.NoClassDefFoundError: org/apache/karaf/jaas/boot/ProxyLoginModule at org.apache.karaf.jaas.config.impl.ProxyLoginModuleInitializer.init(ProxyLoginModuleInitializer.java:32) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) [rt.jar:1.6.0_39] at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) [rt.jar:1.6.0_39] at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) [rt.jar:1.6.0_39] at java.lang.reflect.Method.invoke(Method.java:597) [rt.jar:1.6.0_39] at org.apache.aries.blueprint.utils.ReflectionUtils.invoke(ReflectionUtils.java:225) at org.apache.aries.blueprint.container.BeanRecipe.invoke(BeanRecipe.java:838) at org.apache.aries.blueprint.container.BeanRecipe.runBeanProcInit(BeanRecipe.java:638) ... 15 more Caused by: java.lang.ClassNotFoundException: org.apache.karaf.jaas.boot.ProxyLoginModule from [Module "deployment.org.apache.karaf.jaas.config:2.2.9" from Service Module Loader] at org.jboss.modules.ModuleClassLoader.findClass(ModuleClassLoader.java:190) at org.jboss.modules.ConcurrentClassLoader.performLoadClassUnchecked(ConcurrentClassLoader.java:468) at org.jboss.modules.ConcurrentClassLoader.performLoadClassChecked(ConcurrentClassLoader.java:456) at org.jboss.modules.ConcurrentClassLoader.performLoadClassChecked(ConcurrentClassLoader.java:423) at org.jboss.modules.ConcurrentClassLoader.performLoadClass(ConcurrentClassLoader.java:398) at	0	0	https://issues.jboss.org/browse/AS7-5544

#	Requirement	Questions	Answered	Successfully Answered	URL
1581	AS7-5448	Sorry for the misunderstanding, yes I was referring to the threads subsystem. as defined in the jboss-threads-2.0.0.GA.jar packaged in the jboss 7.1.2.Final distribution org.jboss.threads.JBossThreadPoolExecutor implements org.jboss.threads.management.ThreadExecutorMBean, I think the getQueueSize() method should be added here. org.jboss.threads.JBossThreadPoolExecutor extends java.util.concurrent.ThreadPoolExecutor, so the BlockingQueue I was talking about is available, making the getQueueSize() method a snap to implements. should the issue be filled instead in the jboss-threads jira ?	0	0	https://issues.jboss.org/browse/AS7-5448
1582	AS7-5448	I have done locally all modifications for AS7-5448 and JBTHR-22 though I'm getting the following error inside Java Visual VM: "java.util.MissingResourceException: Cant't find resource for bundle java.util.PropertyResourceBundle, key threadpool.common.queue-size" which is effectively the new attribute that I've added. I guess this is related to a certain .properties file packaged within jboss that I still need to modify. Can someone help me with this one, then I'll be able to send all the code ? Thanks !	0	0	https://issues.jboss.org/browse/AS7-5448
1583	AS7-5405	Does Enterprise OSGi JPA need to be be active at the same time as the EE JPA container/subsystem in AS?	1	1	https://issues.jboss.org/browse/AS7-5405
1584	AS7-5405	{quote} Can OSGi Enterprise JPA use EE JPA (container) support or do we need a new subsystem for OSGi JPA? {quote} I expect the JPA integration to live here: https://github.com/tdiesler/jboss-as/tree/master/osgi/integration/src/main/java/org/jboss/as/osgi {quote} Does OSGi require a dynamic implementation of PersistenceProviderResolver (that could change Hibernate/EclipseLink/OpenJPA/other providers on the fly)? {quote} That won't be necessary.	0	0	https://issues.jboss.org/browse/AS7-5405
1585	AS7-5380	Is Weld 2 an EE7 technology? Which AS release should we target EE7 changes towards?	0	0	https://issues.jboss.org/browse/AS7-5380
1586	AS7-5380	The idea is not to replace Weld 1.x with Weld 2 but rather to host the Weld 2 integration code (which has been hosted in a fork of JBoss AS so far and this represented an additional effort to mantain) in the upstream repository. Ideally, a build profile would be used that would allow AS7 to be built with either Weld 1.x (by default) or Weld 2. We would use this for special AS7 releases that we do along with Weld 2 releases.	1	1	https://issues.jboss.org/browse/AS7-5380
1587	AS7-5315	Why was this issue rejected? Session fixation is possible on AS7 exactly the way Endrigo described it. A new session id needs to be generated just before authentication after the form has been submitted to j_security_check.	1	0	https://issues.jboss.org/browse/AS7-5315
1588	AS7-5315	So you are saying to use the "Disable URL rewriting" option. Is it possible on JBOSS? I made a search on the documentation and didn't found nothing. If you are saying that this is the solution, please tell me where is the documentation or how can i find it.	1	1	https://issues.jboss.org/browse/AS7-5315
1589	AS7-5315	Are you talking about a custom authentication login-module for the security-domain?	1	0	https://issues.jboss.org/browse/AS7-5315
1590	AS7-5306	Is the ApplicationRealm relevant to this test? It's something untidy we should clean up (and we should understand and maybe fix why things behave differently between 7.1 and master as far as configuring it goes) but IIRC the "New missing/unsatisfied dependencies" thing doesn't really prevent tests working. It would only matter if the ApplicationRealm was accessed, and this test shouldn't be doing that.	1	1	https://issues.jboss.org/browse/AS7-5306
1591	AS7-5251	howard, is there anything do for this in AS7 /EAP 6? All HornetQ user credentials attributes allow expressions and can be stored in the vault and are resolved appropriately. Are there other requirements?	1	1	https://issues.jboss.org/browse/AS7-5251
1592	AS7-5251	On what version does this work? It doesn't appear to work for me on 7.1.1.Final (or 7.2.0.Alpha.SNAPSHOT). I started a discussion that includes the exceptions I'm getting at https://community.jboss.org/thread/204759	1	0	https://issues.jboss.org/browse/AS7-5251
1593	AS7-5251	Jeff, thanks clarifying. What workarounds would you suggest in the near-term if we need to mask our SSL keystore passwords on 7.1.1.Final? I'm not sure how long it will be before we can move to 7.2 once the fix become available.	1	0	https://issues.jboss.org/browse/AS7-5251
1594	AS7-5216	Gernot - do you have an option to click on 'Edit' on that page? If so may be quicker to apply the change rather than report it in a Jira.	1	1	https://issues.jboss.org/browse/AS7-5216
1595	AS7-5160	Hello, is there any workaround for this at the moment?	0	0	https://issues.jboss.org/browse/AS7-5160
1596	AS7-5157	In jboss-cli.sh, the following commands seemed to work after deploying 2lc.jar: cd /deployment=2lc.jar/subsystem=jpa/hibernate-persistence-unit=2lc.jar#TEST_PU/ :read-resource(include-runtime=true,recursive=true) Is there DEBUG/TRACE logging that I can enable to see what commands are currently being sent from the admin console? We can probably compare the debug output of what the admin console is trying to do, with the above.	1	1	https://issues.jboss.org/browse/AS7-5157
1597	AS7-5157	scott, can you check why the enabled attribute is set anymore?	1	1	https://issues.jboss.org/browse/AS7-5157
1598	AS7-5145	Can you be more specific? Some attributes already accept expressions - and this is enabled on a per-attribute basic.	0	0	https://issues.jboss.org/browse/AS7-5145
1599	AS7-5140	@Akram Could you post the link to the ResourceURLFactory ? Thanks Marco nevermind. Found Could you explain how to use it ? Thanks Marco	1	1	https://issues.jboss.org/browse/AS7-5140
1600	AS7-5135	Can we make sure that the node identifier is validated to make sure nodeName.getBytes().length is <24 as I don't think that validator is on there atm.	1	1	https://issues.jboss.org/browse/AS7-5135
1601	AS7-5135	what's the status of this patch?	0	0	https://issues.jboss.org/browse/AS7-5135
1602	AS7-5120	I presume that you'd ideally want to run a maven command like: mvn clean package jboss-as:start jboss-as:deploy Where 'jboss-as:start' will look to see if there is a running instance of the AS and then try and deploy your war/ear?	0	0	https://issues.jboss.org/browse/AS7-5120
1603	AS7-5112	I had debate with Emanuel on this topic and we came to conclusion that best way to address this would be not to over complicate things. Aka we do not need to provide "text" info on why this deprecated in DMR as this would just cause us problems with localization. Instead of wording on why and how this was deprecated we could just provide "replacement" property that would tell us what to use instead. Also we do not need info on when this operation was deprecated as current state is deprecated as in "this is the version when it is/was deprecated" more important information would be when this resource would be removed. so my current idea about meta data would be something like this: (example for attribute) {noformat} name: old-attribute type: OBJECT deprecated: true deprecated-replacement: new-attribute deprecated-removed: 1.3.0 (or if we go with version of AS, 7.2.0.Final) {noformat} wdyt?	0	0	https://issues.jboss.org/browse/AS7-5112
1604	AS7-5112	Does it fail if "roles.deprecated" does not exist? It shouldn't. For AS 8 we should change the ResourceDescriptionResolver interface and the standard impl to properly handle this. With the standard impl doing the "roles.deprecated" behavior we'd already introduced.	1	1	https://issues.jboss.org/browse/AS7-5112
1605	AS7-5110	Another easier way to reproduce this bug. Start up with ./domain.sh Create a user. check out the host.xml file and note the following comments: <server name="server-one" group="main-server-group"> <!-- Remote JPDA debugging for a specific server <jvm name="default"> <jvm-options> <option value="-Xrunjdwp:transport=dt_socket,address=8787,server=y,suspend=n"/> </jvm-options> </jvm> --> </server> Change server-three to auto start and then save it. look at the host.xml file. <server name="server-one" group="main-server-group" auto-start="true"/> Now the comments are gone.	0	0	https://issues.jboss.org/browse/AS7-5110

#	Requirement	Questions	Answered	Successfully Answered	URL
1606	AS7-5108	Out of interest how would you suggest we identify all support -D parameters? On the basis that these are just set as system properties not only would we need the list of properties supported by the core of AS we would also need to identify every property defined to be resolved within the configuration, many of these could be user defined - we would also need to identify any properties set in deployment descriptors which would probably be all user defined and then we would need to identify every call to System.getProperty in every single module and in every deployment - some of these items not being available until late in the deployment process. Even if we did that there is always the case that the user wants to set a property that may be used by a future deployment which is not present at boot time.	0	0	https://issues.jboss.org/browse/AS7-5108
1607	AS7-5108	When the object/String that represents the passed in flag is used/injected, can't we increment a counter on that object, and increment it on every use, and then at the end of booting make sure every object at least has a count of 1 or greater. If not then we error out or at least output an error before continuing to boot. That would take care of any random new parameters, our's and any other use case. It would enable people to easily see the custom parameter they thought they were using/injecting is in fact not working as they expected. Of course if they are passing a parameter to their application, then this would generate an error at boot time, hmmm. Thoughts?	1	1	https://issues.jboss.org/browse/AS7-5108
1608	AS7-4965	Your right, I have a old document 'Enterprise JavaBeans 3.1, Proposed Final Draft' between all other final documents by accident. There is an additional paragraph which describe exactly this alternative. I think we should reject the Jira, or do you mean that this should not be deployable?	1	0	https://issues.jboss.org/browse/AS7-4965
1609	AS7-4929	What is the deployed app that is being tested here?	0	0	https://issues.jboss.org/browse/AS7-4929
1610	AS7-4929	Is that the problem? Long story short: It should have been protected.	0	0	https://issues.jboss.org/browse/AS7-4929
1611	AS7-4929	I would suggest if you detect a possible vulnerability in your own site that you should refrain from publicly posting the URL to that site along with any report of the vulnerability. Do you have the details of the full request that is sent to trigger this? Trying the URL in your original comment does not result in an error page with the word DEFACED highlighted however any escaping my have been lost in the conversion to Jira - ideally attach a text file containing the full request that is reported to reproduce this.	0	0	https://issues.jboss.org/browse/AS7-4929
1612	AS7-4929	I will abide by your suggestion. The text in JIRA is the same that we have in the vulnerability report. We have a pdf report tha I can make available to you. Can I email it to you? I looked for an upload button to no avail.	0	0	https://issues.jboss.org/browse/AS7-4929
1613	AS7-4929	Thanks for the attachment - I still can not reproduce anything here. Do you have support for the tool you are using? I think really you are going to need to follow up with the supplier of the scanning utility to verify exactly how it is identifying the failure. The report itself is actually for a specific web server which is not JBoss AS - when I test this myself the contents of the URL are appropriately escaped in the error message so the do not form any type of HTML injection into the error message which is displayed.	0	0	https://issues.jboss.org/browse/AS7-4929
1614	AS7-4929	My Systems Administrator is going to get in touch wiht the makers of the tool. I will post here what they tell us. We appreciate your help. Could this just be a "false positive"?	0	0	https://issues.jboss.org/browse/AS7-4929
1615	AS7-4925	One curiosity of mine is translation - If an application logs with Log4J, is there any translation between that log event and the log4j appender in JBoss (as managed by the log manager after this bug is fixed)? I can understand a translation occurring if the log event comes from JUL, Commons, or SLF - but if Log4J is the logging API and the Handling API - I'm hoping we can get the original event object. Reason I ask is we have a couple of APIs that would like to use Log4J's Object parameter to pump in special objects - and Log4J retains these objects in the LoggingEvent that gets sent to the Appender. However if this is not possible and the message gets converted to a String - we'll understand this as one of the unfortunate side effects of the logmanager subsystem. Our work-around would be to include a special appender in log4J's root logger that would hopefully not get closed or otherwise managed by JBoss - and from that appender we can perform our special logging. Overall this is probably the only EAP 5 to 6 transitional item we have a no-alternative for.	0	0	https://issues.jboss.org/browse/AS7-4925
1616	AS7-4917	IRC log about this issue: [16:49:36] <jmesnil> kkhan, before leaving, brian asked me to look at AS7-4917 [16:49:37] <jbossbot> jira [AS7-4917] paths need to allow expressions in its 'path' attribute [Open (Unresolved) Feature Request, Major, Jeff Mesnil] https://issues.jboss.org/browse/AS7-4917 [16:50:32] <jmesnil> kkhan, afaict, the fix is easy, just allow expressions for the path attribute and I check it works both in standalone and domain mode [16:51:22] <jmesnil> kkhan, but brian told me that there were things to discuss about this. Am I missing something about allowing expression for the paths? [16:51:27] * kkhan looks [16:52:04] <emuckenhuber> perhaps about the implications for backwards compatibility [16:52:42] <jmesnil> emuckenhuber, yeah, we'd need to bump up the management version if we allow expressions for paths [16:53:16] <jmesnil> emuckenhuber, but if it was not allowed before, I don't see how we could break existings paths by allowing it now [16:53:35] <ctomc> jmesnil: maybe it is also how attributes are beeing read, you need to use resolveModelAttribute(context,model)... [16:53:46] <kkhan> ctomc/emuckenhuber: How do expressions work with transformers [16:53:53] <emuckenhuber> kkhan: that is the main problem :) [16:53:55] <kkhan> e.g. DC's version allows expressions [16:54:05] <emuckenhuber> basically they can't be resolved properly [16:54:06] <kkhan> and then the legacy HC does not [16:54:23] <emuckenhuber> i mean we can resolve them against domain properties [16:54:30] <emuckenhuber> however not against host / server properties [16:54:33] <ctomc> kkhan: currently we have no special handling for this, but it will be added as part of "default transformer" [16:55:04] <kkhan> So DC would call resolveExpression on legacy HC and massage the op or something? [16:55:33] <kkhan> jmesnil: ^^ is what I felt it could break :-) [16:55:43] <emuckenhuber> it can't [16:56:00] <emuckenhuber> since there might be multiple server-groups on the legacy HC [16:56:20] <emuckenhuber> so it won't have the information any longer when it starts the server [16:56:36] <emuckenhuber> i mean as it is now... i was thinking adding that information somewhere [16:57:13] <emuckenhuber> we could do it for runtime operations though - although not sure if that makes sense [16:57:53] <kkhan> emuckenhuber: ok, could you not resolve per server group? [16:58:36] <kkhan> I'm just throwing out ideas without looking in case that is not clear :-) [16:58:43] <emuckenhuber> profiles can be shared, so i guess not [17:02:29] <ctomc> jmesnil: from my point of view, fix for this issue should be trivial, except for maybe transformation stuff :| [17:03:56] <kkhan> jmesnil: Nothing springs to mind about what would be bad apart from transformations for me either so go ahead. Surely this will not be the only thing introducing expressions [17:04:44] <jmesnil> kkhan, ok, thanks! :) [17:05:42] <kkhan> jmesnil: Add more stuff to org.jboss.as.controller.services.path.PathsTestCase and if there is stuff we can't see now it will hopefully make it clear [17:06:05] <jmesnil> kkhan, +1	1	1	https://issues.jboss.org/browse/AS7-4917
1617	AS7-4853	can you elaborate on the initial description? do you mean the filesystem location?	1	1	https://issues.jboss.org/browse/AS7-4853

#	Requirement	Questions	Answered	Successfully Answered	URL
1618	AS7-4853	This came out of an issue that I have also filed in that the quickstarts deploy applications via the API, they do not end up in the deployments directory, but as a hashed name in some other folder. The deployments directory has things in it, it would be nice to know where an application is loaded from. Indicating that you will not see it in the deployments directory or something along those lines. I hope that helps. Let me know if it does not?	1	1	https://issues.jboss.org/browse/AS7-4853
1619	AS7-4828	Did 7.1.2 actually not support expressions at all? My understanding is it did, but not properly, i.e. the expressions were resolved by the JGroups system property replacement logic. And there was a bug related to that. That improper support isn't ideal, but it may have worked for a lot of use cases. Rejecting the expression now would break those use cases. In particular it would make a rolling upgrade of a cluster problematic, perhaps impossible. Apologies in advance if my understanding of the situation in 7.1.2 is incorrect.	1	1	https://issues.jboss.org/browse/AS7-4828
1620	AS7-4821	David I agree with you, my idea was just to simplify the way we register operations, so users do not need to write DescriptionProvider's we have done this for resources and attributes with ResourceDefinition & AttributeDefinition but operations ware left out :-(in spirit of converting all subsystems to ResourceDefinition/AttributeDefinition it would be nice to present users friendlier api also for operations. my current idea is implemented here: https://github.com/ctomc/jboss-as/commit/b5a23ab5c70a67f37c644aa5811f6fad00127c74 it adds OperationDefinition (so it is similar to RD and AD) and four methods to ManagementResourceRegistration two for simpler adding of operation by name and attribues and two that enable usage of OperationDefinition {code:java} void registerOperationHandler(String operationName, OperationStepHandler handler, ModelType replyType, ModelType replyValueType, AttributeDefinition ... attributes); void registerOperationHandler(String operationName, OperationStepHandler handler, AttributeDefinition ... attributes); void registerOperationHandler(OperationDefinition definition,OperationStepHandler handler); void registerOperationHandler(OperationDefinition definition, OperationStepHandler handler,boolean inherited); {code} this i think can be added without to much disturbance and will allow to write prettier management code :) wdyt?	1	1	https://issues.jboss.org/browse/AS7-4821
1621	AS7-4800	Are there other references besides https://community.jboss.org/wiki/SubsystemDevelopmentConventions + https://community.jboss.org/message/730384#730384 that I should look at?	1	1	https://issues.jboss.org/browse/AS7-4800
1622	AS7-4720	Brian Should this be removed from AS 7.1.3 (on the 7.1 branch) as well?	1	1	https://issues.jboss.org/browse/AS7-4720
1623	AS7-4713	Randall, Are you able to test against the changes in https://github.com/scottmarlow/jboss-as/tree/AS7-4713_movehibernatecommonjar? Scott	1	1	https://issues.jboss.org/browse/AS7-4713
1624	AS7-4687	Why don't we flip the default ThreadLocal used to be non inheritable? Rather than relying on a system property being set - at the moment I have not seen a valid reason for using an InheritedThreadLocal and this is not the first time it has caused us problems.	1	1	https://issues.jboss.org/browse/AS7-4687
1625	AS7-4646	Why exactly is DIGEST non compliant?	1	1	https://issues.jboss.org/browse/AS7-4646
1626	AS7-4646	At what point do your requirements regarding storage actually apply? Is it that it needs to be irreversible or just a level of encryption that is applied? As I said above we do have an option to use the vault to store the passwords to work with Digest authentication.	1	0	https://issues.jboss.org/browse/AS7-4646
1627	AS7-4580	Your configuration looks correct. The root exception is: Cannot assign requested address. Could you check that your 225.3.2.3 is correctly configured for UDP?	1	1	https://issues.jboss.org/browse/AS7-4580
1628	AS7-4580	Could inserting this netty port directive in the run.sh startup script cause issues with the discovery port not binding properly? -Dhornetq.remoting.netty.port=7777 [rundeck@prdlvm0038 bin]$ ps aux | grep java rundeck 12530 0.0 0.0 61180 736 pts/1 R+ 16:42 0:00 grep java rundeck 23501 0.0 4.8 3632788 393852 ? Sl Apr26 3:37 java -Djnp.port=1099 -Djnp.rmiPort=1098 -Djnp.host=10.40.2.38 -Dhornetq.remoting.netty.host=10.40.2.38 -Dhornetq.remoting.netty.port=7777 -XX:+UseParallelGC -XX:+AggressiveOpts -XX:+UseFastAccessorMethods -Xms3072M -Xmx3072M -Dhornetq.config.dir=../config/stand-alone/clustered -Djava.util.logging.config.file=../config/stand-alone/clustered/logging.properties -Djava.library.path=. -classpath ../lib/twitter4j-core.jar:../lib/netty.jar:../lib/jnpserver.jar:../lib/jnp-client.jar:../lib/jboss-mc.jar:../lib/jboss-jms-api.jar:../lib/hornetq-twitter-integration.jar:../lib/hornetq-spring-integration.jar:../lib/hornetq-logging.jar:../lib/hornetq-jms.jar:../lib/hornetq-jms-client-java5.jar:../lib/hornetq-jms-client.jar:../lib/hornetq-jboss-as-integration.jar:../lib/hornetq-core.jar:../lib/hornetq-core-client-java5.jar:../lib/hornetq-core-client.jar:../lib/hornetq-bootstrap.jar:../config/stand-alone/clustered:../schemas/ -Dcom.sun.management.jmxremote org.hornetq.integration.bootstrap.HornetQBootstrapServer hornetq-beans.xml	1	1	https://issues.jboss.org/browse/AS7-4580
1629	AS7-4580	I don't think setting the netty port is related, Netty is not used at all for UDP communication. Could you try to run HornetQ with -Djava.net.preferIPv4Stack=true and then a 2nd time with -Djava.net.preferIPv6Addresses=true to see if it makes any difference? http://docs.oracle.com/javase/6/docs/technotes/guides/net/properties.html	1	1	https://issues.jboss.org/browse/AS7-4580
1630	AS7-4567	* Where can I find the spec? * What's the correct location to put jars which are shared across the whole AS7 server (accessible for all deployed artifacts)?	1	1	https://issues.jboss.org/browse/AS7-4567
1631	AS7-4567	Gernot, Do you have a sample of your war to attach? Just to try to reproduce the problem, I'm not saying that it will be implemented since there is no plans to this.	0	0	https://issues.jboss.org/browse/AS7-4567
1632	AS7-4567	I am testing this with EAP 6.1 Alpha and apparently it works only partially. The beans packaged in a module can be injected, but it looks like the annotations in those beans are not processed. I have a bean that looks like this inside a module: {code:java} @ApplicationScoped public class MyBean { @PostConstruct public void init() { System.out.println("init"); } } {code} I added this to the {{jboss-deployment-structure.xml}} file of a {{.ear}} where I have EJB that need this bean injected: {code:xml} <dependencies> <module name="org.mymodule" export="true" meta-inf="import"/> </dependencies> {code} Then I have something like this in the EJB: {code:java} @Singleton public class MyEjb ... { @Inject private MyBean myBean; } {code} What happens then is that the injection works, but a new instance of MyBean is created for each call to the EJB. In addition the method marked with {{@PostConstruct}} is never executed. Is this the expected behavior?	1	0	https://issues.jboss.org/browse/AS7-4567
1633	AS7-4566	Remy, would you please provide the name of the system property? I was unable to find it.	1	0	https://issues.jboss.org/browse/AS7-4566

#	Requirement	Questions	Answered	Successfully Answered	URL
1634	AS7-4547	Could the performance loss actually be due to JVM CodeCache exhaustion? Default Java CodeCache size is quite low (48M/96M) and the JIT compiler switches off *permanently* once usage nears the ceiling. We have a large system & needed to up our CodeCache to 256m minimum (now running with 400m as our default recommendation). Our symptoms were subtle & initially misattributed, this is a tricky issue. -XX:ReservedCodeCacheSize=400m Tiered compilation should not cause performance loss, but can contribute to the above exhaustion & subsequent performance loss. See a few references: - http://blogs.atlassian.com/2012/05/codecache-is-full-compiler-has-been-disabled/ -- (don't follow their tip to turn on CodeCache Flushing, there are actually major bugs there that can, on their own, cause severe degraded JVM performance.) - https://blog.codecentric.de/en/2012/07/useful-jvm-flags-part-4-heap-tuning/ - https://www-304.ibm.com/connections/blogs/wcs/entry/performance_roadblock_full_jit_cache?lang=en_us Hope this helps.	0	0	https://issues.jboss.org/browse/AS7-4547
1635	AS7-4508	Yeah, I noticed that part while looking at the code although I didn't mention that as a workaround to the user since it really looked like something that happens to work but probably wasn't intended that way (should we be ignoring duplicate URLs in that comma separated list?). The other thing that I'm working as part of this change is to cleanup HaRemotingNamingStore a bit to make sure that it uses the NamingStoreCache so that appropriate cleanup of channel and connection takes place whenever necessary.	1	1	https://issues.jboss.org/browse/AS7-4508
1636	AS7-4508	with which build is this fix available ??	1	1	https://issues.jboss.org/browse/AS7-4508
1637	AS7-4407	Should this issue be marked out of date? It looks like it was resolved a while ago	1	0	https://issues.jboss.org/browse/AS7-4407
1638	AS7-4390	stan, can you take a look at this?	1	1	https://issues.jboss.org/browse/AS7-4390
1639	AS7-4342	Same applies to maven, if you dont set, you get "build.sh: Could not locate Maven; check $MVN or $MAVEN_HOME." Wasnt this solved back in AS5 days?	1	1	https://issues.jboss.org/browse/AS7-4342
1640	AS7-4342	Are you sure? I don't have $MVN or $MAVEN_HOME set and build.sh works fine for me.	0	0	https://issues.jboss.org/browse/AS7-4342
1641	AS7-4324	Still an issue, if I deploy a bunch of bundles where one dependency is missing, of course the bundles fail to deploy. But when I deploy that one missing dependency the container should detect that the missing dependency is now present and refresh the other bundles which should start as a result? Can we have a feature to turn back on, or make configurable, automatic refresh of bundles as it was in JBoss 7.1?	1	0	https://issues.jboss.org/browse/AS7-4324
1642	AS7-4324	Could you please point me in the direction of the part of the OSGi specification that says bundles must be installed, resolved and started as a set and that automatic refreshing is not allowed when a new bundle is hot deployed? I took a look at v5 and was unable to find anything like that.. although I could have easily missed it.	0	0	https://issues.jboss.org/browse/AS7-4324
1643	AS7-4296	This is a particularly important issue for me in relation to my cloud related projects. We need total automation over the configuration of target containers across development, test, and UAT environments and we also need automation of the overriding of application servers within deployed clusters. Can anyone provide advice on the process to get this issue assigned to a target release?	1	1	https://issues.jboss.org/browse/AS7-4296
1644	AS7-4287	Its network latency, it can not be significantly lower over Europe->USA because electric field travels in wire slightly lower then 300k km/sec and every router on way is adding few milliseconds delay. Look at other web applications. Do they needs minutes to load? Using https://developers.google.com/web-toolkit/doc/latest/DevGuideClientBundle helps	1	0	https://issues.jboss.org/browse/AS7-4287
1645	AS7-4279	Closing, as it contains no useful information. Please ask these sorts of questions on the forums, and actually provide some information that will allow people to answer your questions. We are not familiar with your application and your setup, and we are not mind readers. e.g: Is your server running out of memory? Is your server leaking file handles? Are there any errors in the logs? etc	0	0	https://issues.jboss.org/browse/AS7-4279
1646	AS7-4268	Should this really be its own subsystem? Wouldn't it make more sense to add it to the messaging subsystem, but have it configured outside the <hornetq-server> element?	1	1	https://issues.jboss.org/browse/AS7-4268
1647	AS7-4268	Jeff, just noticed that you began work on AS7-4268. When you complete this, is it something we could just pull the latest and get your update?	0	0	https://issues.jboss.org/browse/AS7-4268
1648	AS7-4268	Is this done since the 2 pull requests have been merged?	1	1	https://issues.jboss.org/browse/AS7-4268
1649	AS7-4268	Hi, I was trying to configure the JMSBridge based on the doc of AS72, but couldn't figure out where would I download the dependency jars from (acmemq-1.2.3.jar and mylogapi-0.0.1.jar) ? A google search didnt help actually. Could anyone help on this ?	1	1	https://issues.jboss.org/browse/AS7-4268
1650	AS7-4268	I understand the documentation refers to non-existent JARs. Can you give a specific example for connecting to prior versions of JBoss, for example, if HornetQ is running in JBoss AS 7.x, how to bridge a topic from a JBoss 4.2 / JBossMQ environment? Which JBossMQ JARs have to be copied, where to put them, etc?	0	0	https://issues.jboss.org/browse/AS7-4268
1651	AS7-4265	Do we have provision to edit an existing module ? For example let's say if i already have a module in place and if i want to add one module dependency to it, can i do so ?	0	0	https://issues.jboss.org/browse/AS7-4265
1652	AS7-4265	This works well for me, however, there's one thing which I had to manually change before it would work: I have a properties file which needs to be outside the WAR file, but yet accessible by the WAR file. I read the instructions on how to do this using the "modules" approach (https://developer.jboss.org/wiki/HowToPutAnExternalFileInTheClasspath), and so using the CLI module "add" command would be perfect. This is what I did: module add --name=com.mycompany.myproject.configuration --resources=myproperties.properties (where the myproperties.properties file is in the same folder where I issue this command) Everything worked well, except that the module.xml file created looked like this: <?xml version="1.0" ?> <module xmlns="urn:jboss:module:1.1" name="com.mycompany.myproject.configuration"> <resources> <resource-root path="myproperties.properties"/> <!-- THIS IS THE LINE WHICH I NEED CHANGED. --> </resources> </module> But what I need is: <?xml version="1.0" ?> <module xmlns="urn:jboss:module:1.1" name="com.mycompany.myproject.configuration"> <resources> <resource-root path="."/> <!-- THIS IS WHAT I NEED. --> </resources> </module> Otherwise the WAR file can't find the properties file. How do I do this? Thanks,	1	1	https://issues.jboss.org/browse/AS7-4265
1653	AS7-4194	Do you know if this be pulled in for 7.1.2.Final-redhat1?	1	1	https://issues.jboss.org/browse/AS7-4194
1654	AS7-4179	How does one go around to adding one category to get only logs from that one category? If you add it as <logger> with a lower level than threshold the logs are ditched, if you lower the threshold you get unwanted logs as well.	0	0	https://issues.jboss.org/browse/AS7-4179
1655	AS7-4151	I would like to see this feature add too, is it on the roadmap?	1	1	https://issues.jboss.org/browse/AS7-4151

#	Requirement	Questions	Answered	Successfully Answered	URL
1656	AS7-4140	I have confirmed with Andy Taylor that for the standalone-full-ha.xml messaging configuration, the broadcast-group and the discovery-group: - should be using the same multicast channel (i.e. same address and port) - work correctly when using the same multicast channel At present, standalone-full-ha.xml is not used for testing in the AS71 testsuite, as far as I can tell. However, in case it is used in furture, we should update the config generation so that it is properly processed. Therefore, we need to add the messaging configuration to the list of XSL transformations in changeIPAddresses.xsl so that the multicast addresses are replaced by -Dmcast, whether -Dipv6 is true or not. Because the ports used are not used by any other multicast channels, this should work for testing. Ondra, can you set this up? Once this is set up and tested, we can close this issue, I believe.	0	0	https://issues.jboss.org/browse/AS7-4140
1657	AS7-4140	I'll get to this at the end of this week, is that ok? I'll include it in a bunch of other updates.	0	0	https://issues.jboss.org/browse/AS7-4140
1658	AS7-4140	Ondra, did you forget about this? I still see hardcoded addresses in standalone-ha-full.xml.	0	0	https://issues.jboss.org/browse/AS7-4140
1659	AS7-4140	Ondro Can you resubmit and link the PR now? Or is there some outstanding issue?	0	0	https://issues.jboss.org/browse/AS7-4140
1660	AS7-4088	Do you mean add expression support to the various attributes in ModeClusterConfigResourceDefinition? If so this is done and this can be resolved against 7.1.2.	1	1	https://issues.jboss.org/browse/AS7-4088
1661	AS7-4084	Implemented the code that Tristan provided into the Infinispan subsystem recently refactored into resources. Encountered the following problem: <quote> Infinispan and JGroups gather statistics in certain associated services (e.g. TransactionManager, RpcManager,...). These services are installed at boot time, but not started until someone deploys a clustering application which uses them. So, if I try to read one of their read-only attributes before a deployment has occurred, I throw an OperationFailedException with a message saying that the required service has not been started. This works fine and seems to be the right way to proceed. On the other hand, there is a smoke test case BasicOperationsUnitTestCase.testRecursiveReadIncludingRuntime() which seems to assume that all reads of runtime-only resources will succeed. </quote> Is it really the case that every read of a read-only attribute should always succeed? If so, what is the best way to work around this? Can I somehow return an informational message instead of throwing an exception? Checking these issues with Emanuel M.	1	1	https://issues.jboss.org/browse/AS7-4084
1662	AS7-4084	In that case, you seem to be activating whole child resources by conditionally returning resources in response to getChild() calls. In what we had set up so far, we need to turn on a subset of attributes for a specific resource; e.g. /susbsystem=infinispan/cache-container=X // always available attributes A B C // attributes available when a particular MSC service is running D E Can the method used in the Messaging case also be used to cater to the above? In other words, what mechanism can be used to "turn-on" and "turn-off" the visibility of the attributes, as opposed to children?	1	1	https://issues.jboss.org/browse/AS7-4084
1663	AS7-4075	Are you certain other components depend on this? I'm not seeing it. In the poms I see some exclusions. But maven dependencies are complex beasts. Luckily, however, we have a modular classloading setup where dependencies have to be explicitly declared. And a "git grep commons.pool" shows me that the only spots in modules where that appears are: build/src/main/resources/modules/org/apache/commons/pool/main/module.xml:<module xmlns="urn:jboss:module:1.1" name="org.apache.commons.pool"> build/src/main/resources/modules/org/infinispan/cachestore/jdbc/main/module.xml: <module name="org.apache.commons.pool"/> build/src/main/resources/modules/org/infinispan/client/hotrod/main/module.xml: <module name="org.apache.commons.pool"/>	1	1	https://issues.jboss.org/browse/AS7-4075
1664	AS7-4075	Right, my bad, I was just only looking at maven -- but looking at JBoss modules is very exact. We should definitely upgrade then. Then, I also think we should upgrade 7.1 branch -- do you agree?	1	1	https://issues.jboss.org/browse/AS7-4075
1665	AS7-4057	Is this worked on? Can we commit for EAP 6.1?	1	1	https://issues.jboss.org/browse/AS7-4057
1666	AS7-4051	{quote}But if this transparent failover fails in our tests unless we change the cookie domain, then why would it pass in a real failover scenario without resetting the cookie domain?{quote} If I understand the question correctly, in real life you store cookies with domain of the load-balancer (its hostname) which makes the failover transparent. Our tests dont use loadbalancer as it would be OS specific.	1	1	https://issues.jboss.org/browse/AS7-4051
1667	AS7-4022	Yuck. Why are these properties and not proper attributes?	0	0	https://issues.jboss.org/browse/AS7-4022
1668	AS7-4022	Also, why are hibernate properties on an infinispan config?	0	0	https://issues.jboss.org/browse/AS7-4022
1669	AS7-4022	And if they are proper attributes, are they stable and supportable for many years? I realize this last question argues in an opposite direction from my previous question. The overall point is how does this feature request fit into an overall long-term supportable management contract for this subsystem?	0	0	https://issues.jboss.org/browse/AS7-4022
1670	AS7-4006	How about using reverse lookup? {code} ::1 {code} becomes {code} 1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.ip6.arpa. {code} So it might work if we set the cookie hostname to "0.0.0.0.ip6.arpa" But it needs a IPv6 capable DNS server...	1	1	https://issues.jboss.org/browse/AS7-4006

#	Requirement	Questions	Answered	Successfully Answered	URL
1671	AS7-4006	# using ipv6 global address (since there are problems with multicast on local ::1) # using any multicast, but ff01:XYZ is enough {noformat} [rhusar@rhusar testsuite]$ time mvn clean install -Dts.noSmoke -Dts.clust -Dipv6 -Dnode0=1:222:68ff:fe0b::1 -Dnode1=1:222:68ff:fe0b::1 -DudpGroup=ff01::1 [INFO] Scanning for projects... [INFO] -- [INFO] Reactor Build Order: [INFO] [INFO] JBoss Application Server Test Suite: Aggregator [INFO] JBoss Application Server Test Suite: Integration [INFO] JBoss Application Server Test Suite: Integration - Clustering [INFO] [INFO] -- [INFO] Building JBoss Application Server Test Suite: Aggregator 7.1.1.Final-SNAPSHOT [INFO] -- [INFO] [INFO] --- maven-clean-plugin:2.4.1:clean (default-clean) @ jboss-as-testsuite --- [INFO] Deleting /home/rhusar/git/jboss-as/testsuite/target [INFO] [INFO] --- maven-enforcer-plugin:1.0:enforce (ban-bad-dependencies) @ jboss-as-testsuite --- [INFO] [INFO] --- maven-enforcer-plugin:1.0:enforce (enforce-java-version) @ jboss-as-testsuite --- [INFO] [INFO] --- maven-enforcer-plugin:1.0:enforce (enforce-maven-version) @ jboss-as-testsuite --- [INFO] [INFO] --- buildnumber-maven-plugin:1.0:create (default) @ jboss-as-testsuite --- [INFO] Checking for local modifications: skipped. [INFO] Updating project files from SCM: skipped. [INFO] Executing: /bin/sh -c cd /home/rhusar/git/jboss-as/testsuite && git show [INFO] Working directory: /home/rhusar/git/jboss-as/testsuite [INFO] Storing buildNumber: cc7bd73d4c16acfcb225e1161fbfa3fc8272556a at timestamp: 1331041067133 [INFO] Executing: /bin/sh -c cd /home/rhusar/git/jboss-as/testsuite && git show [INFO] Working directory: /home/rhusar/git/jboss-as/testsuite [INFO] Storing buildScmBranch: UNKNOWN [INFO] [INFO] --- maven-help-plugin:2.1.1:active-profiles (help.active-profiles) @ jboss-as-testsuite --- [INFO] Active profile report written to: /home/rhusar/git/jboss-as/testsuite/target/help.active-profiles.txt [INFO] [INFO] --- maven-help-plugin:2.1.1:effective-pom (help.effective-pom) @ jboss-as-testsuite --- [INFO] Effective-POM written to: /home/rhusar/git/jboss-as/testsuite/target/help.effective-pom.txt [INFO] [INFO] --- maven-help-plugin:2.1.1:effective-settings (help.effective-settings) @ jboss-as-testsuite --- [INFO] Effective-settings written to: /home/rhusar/git/jboss-as/testsuite/target/help.effective-settings.txt [INFO] [INFO] --- maven-help-plugin:2.1.1:system (help.system) @ jboss-as-testsuite --- [INFO] System report written to: /home/rhusar/git/jboss-as/testsuite/target/help.system.txt [INFO] [INFO] --- properties-maven-plugin:1.0-alpha-2:write-project-properties (default) @ jboss-as-testsuite --- [INFO] [INFO] --- maven-antrun-plugin:1.6:run (banner) @ jboss-as-testsuite --- [INFO] Executing tasks main: [echo] ********** [echo] [echo] JBoss Application Server Testsuite [echo] [echo] Usage: [echo] [echo] mvn [clean] install [-DallTests | (-Dts.<group>)+ | -Dtest=...] [<params>] [echo] [echo] integration-tests.sh [-DallTests | (-Dts.<group>)+ | -Dtest=...] [<params>] [echo] [echo] Where: [echo] [echo] <group> = one of smoke, basic, clust, iiop. [echo] [echo] <params> = parameters changing the configuration like database, IP stack, etc. [echo] [echo] Params: [echo] -Dds.db=<database> Configures the AS to use given database. [echo] Additional properties to be provided: [echo] -Dds.jdbc.url, -Dds.jdbc.user, -Dds.jdbc.user [echo] [echo] -Dipv6 Configures the testsuite to run with IPv6. [echo] [echo] -Dtest=<test classes mask> Runs matching tests. [echo] The group needs to be specified to let it prepare proper AS config. [echo] ********** [INFO] Executed tasks [INFO] [INFO] --- maven-dependency-plugin:2.3:copy (copy-annotations-endorsed) @ jboss-as-testsuite --- [INFO] Configured Artifact: org.jboss.spec.javax.annotation:jboss-annotations-	1	1	https://issues.jboss.org/browse/AS7-4006
1672	AS7-3992	Should this streamlining also affect legacy-to-new command line parameters? Current output: {noformat} -b=<value> Set system property jboss.bind.address to the given value -b <value> Set system property jboss.bind.address to the given value {noformat} Suggested: {noformat} -b=<value>, -b <value> Set system property jboss.bind.address to the given value {noformat}	0	0	https://issues.jboss.org/browse/AS7-3992
1673	AS7-3992	{quote}I don't see any reason to mark them as deprecated.{quote} Isn't one syntax preferred over the other? It seems like we should just pick either with or without the "=" as the recommended syntax, and only list this preferred syntax in the help to make it simpler. The other syntax would still be supported for compatibility, but just not listed in the standard help message if there is no reason to use it other than backwards compatibility.	0	0	https://issues.jboss.org/browse/AS7-3992
1674	AS7-3987	The reason it's needed is because currently it requires setting both properties to run the arquillian modules and the testsuite against an AS zip. I'm not sure I understand why jboss.home can't be used for both. Can you describe a use case where (1) is not the same as (2) in your comment?	1	1	https://issues.jboss.org/browse/AS7-3987
1675	AS7-3965	The mod_cluster does support expressions for certain properties. Which ones are missing?	0	0	https://issues.jboss.org/browse/AS7-3965
1676	AS7-3965	At least "balancer". But can you verify for others metrics? This is key for profile definition. Then a "ha" profile could be used by several server-groups to define different clusters.	0	0	https://issues.jboss.org/browse/AS7-3965
1677	AS7-3965	Is there any reason not to support an expression for the ssl config's ca-revocation-url? The other attributes that point to external resources support expressions.	1	1	https://issues.jboss.org/browse/AS7-3965
1678	AS7-3950	I'm working on it but the --silent option seems broken. When the user provide the minimal information (username, password) without --silent it's fails. Can I make by default the console silent when the user provide the necessary information ? I would like to do something like : {code} ./add-user.sh -u ggrossetie -p abcd # Add ggrossetie with password abcd to realm ManagementRealm (Management user) in silent mode ./add-user.sh --user ggrossetie --password abcd -r MyRealm -a # Add ggrossetie with password abcd to realm MyRealm (Application user) in silent mode ./add-user.sh --user ggrossetie --password abcd --role Master # Add ggrossetie with password abcd and role Master to realm ManagementRealm (Management user) in silent mode ./add-user.sh --user ggrossetie # Default mode ask for additional information (password) {code} @see https://issues.jboss.org/browse/AS7-2965	1	1	https://issues.jboss.org/browse/AS7-3950
1679	AS7-3950	Can you please clarify what you mean by 'it fails'? I am not sure silent by default would be a good change, if a user has not requested it to be silent they should see the relevant output messages indicating the success.	0	0	https://issues.jboss.org/browse/AS7-3950
1680	AS7-3950	I'm working on master : {code} ./add-user.sh ggrossetie abcd Exception in thread "main" java.lang.IllegalStateException: JBAS015232: Aucune java.io.Console disponible pour interagir avec l'utilisateur. {code} When the user provide two or three args the console is null and the code check : {code:java} if ((theConsole == null) && (stateValues.isSilent() == false)) { throw MESSAGES.noConsoleAvailable(); } {code} The --silent option seems to be ignored : {code} ./add-user.sh ggrossetie abcd --silent Exception in thread "main" java.lang.IllegalStateException: JBAS015232: Aucune java.io.Console disponible pour interagir avec l'utilisateur. {code} {quote} I am not sure silent by default would be a good change {quote} Me neither but can we instanciate a JavaConsole ?	1	1	https://issues.jboss.org/browse/AS7-3950
1681	AS7-3950	Do you have a stack trace for those messages?	0	0	https://issues.jboss.org/browse/AS7-3950
1682	AS7-3891	John is on leave for a couple of weeks, who else can take this?	0	0	https://issues.jboss.org/browse/AS7-3891
1683	AS7-3874	Jaikiran, How much is left on this?	1	1	https://issues.jboss.org/browse/AS7-3874

#	Requirement	Questions	Answered	Successfully Answered	URL
1684	AS7-3850	This issue is being relocated to 7.1.2 so that 7.1.1 only contains criticals, blockers, and EAP LA issues (which are probably not yet at the right status, or need an extra triage pass). If these issues are completed in a 7.1.1 timeframe please change the fix version. Although if you forget I will bulk correct.	1	1	https://issues.jboss.org/browse/AS7-3850
1685	AS7-3848	Could you please let us know how did you verify with some sample program..thanks	1	1	https://issues.jboss.org/browse/AS7-3848
1686	AS7-3848	{quote} Could you please let us know how did you verify with some sample program..thanks {quote} You don't have to do anything. Just deploy the @Clustered beans on the server side and invoke them from a client. Here's one example of the client invoking on remote beans https://docs.jboss.org/author/display/AS71/EJB+invocations+from+a+remote+client+using+JNDI. That example isn't specific to @Clustered beans though.	1	0	https://issues.jboss.org/browse/AS7-3848
1687	AS7-3842	It doesn't look like this project has been progressing recently. Is there any chance of an official version?	1	1	https://issues.jboss.org/browse/AS7-3842
1688	AS7-3842	I agree this would be nice to have, but does this issue belong in the AS7 project JIRA? Seems like this is an external effort.	1	1	https://issues.jboss.org/browse/AS7-3842
1689	AS7-3816	Expansion of variables in web.xml and jboss-web.xml do not appear to be supported in JBoss 7.1 either (and were in JBoss 5.x). Will expansion of variables in web.xml and jboss-web.xml be supported too in 7.1.2? For details please see: http://stackoverflow.com/questions/10032095/references-to-systemproperties-in-web-xml-and-jboss-web-xml-works-in-jboss5-but	0	0	https://issues.jboss.org/browse/AS7-3816
1690	AS7-3816	Expansion wasn't working for jboss-service.xml for a SAR deployment (Jboss AS 7.1.1). Would this be supported with this fix?	1	0	https://issues.jboss.org/browse/AS7-3816
1691	AS7-3809	Hi, can someone explain why was this commit reverted?	1	1	https://issues.jboss.org/browse/AS7-3809
1692	AS7-3807	Can this be resolved? The dependent issue is fixed and the fixed JBoss TS is incorporated in AS 7 master.	1	1	https://issues.jboss.org/browse/AS7-3807
1693	AS7-3793	What is the use case for this?	1	0	https://issues.jboss.org/browse/AS7-3793
1694	AS7-3793	<stuartdouglas> opalka: That is not a real world requirement <stuartdouglas> appclient is not supposed to act as a server <opalka> stuartdouglas, of course it isn't. It's just for our testing purposes. <stuartdouglas> opalka: can you redesign the test so it actually runs inside the appclient JVM? <opalka> stuartdouglas, I think yes. But this is exactly I wanna avoid (it's about 200+ tests) <stuartdouglas> If you really want to use the appclient as a server can you just use a custom appclient.xml in your tests? <opalka> stuartdouglas, yes, this is workable for us too <stuartdouglas> opalka: does having the remoting subsystem in appclient actually allow remote naming to work? <opalka> stuartdouglas, yes, it works for me <opalka> stuartdouglas, I verified the proof of concept last week <opalka> stuartdouglas, If you will say no remoting protocol in appclient I'm fine with it. In that case I'll provide either custom appclient.xml for our tests or XSLT <opalka> stuartdouglas, I just wanna know the decision :) <opalka> stuartdouglas, to open or not to open remoting protocol on appclient, that's the question :) <stuartdouglas> opalka: I'm not going to add remoting to appclient just for your test suite <opalka> stuartdouglas, ok, NP. The decision have been taken. I'm fine with it. <stuartdouglas> as it is something that most real world users would then want to remove <opalka> stuartdouglas, +1 <opalka> stuartdouglas, I'm going to remove it from my pull request <stuartdouglas> ok <opalka> stuartdouglas, removed - https://github.com/jbossas/jboss-as/pull/1548	0	0	https://issues.jboss.org/browse/AS7-3793
1695	AS7-3789	Hi Jeff, Could you review this HQ patch?	0	0	https://issues.jboss.org/browse/AS7-3789
1696	AS7-3789	tomaz, did you attach a patch? I can only see your compiled hornetq-core.jar in the attachments...	1	1	https://issues.jboss.org/browse/AS7-3789
1697	AS7-3789	Are we waiting for a new HornetQ release with a fix?	1	1	https://issues.jboss.org/browse/AS7-3789
1698	AS7-3789	How is the testing going Tomaz?	1	1	https://issues.jboss.org/browse/AS7-3789
1699	AS7-3789	Thanks Tomaz. BTW, is this fixed now? I don't see the delay in booting standalone-full.xml anymore?! (win7/32b)	1	1	https://issues.jboss.org/browse/AS7-3789
1700	AS7-3789	dimitris, yes the fix has been merged into HornetQ 2.2.15.Final. That's why you don't see the delay anymore. There is still room for performance improvements (do not use reflection, etc.) but the main bottleneck has been fixed. Does the current boot time is fast enough to downgrade the priority of this issue (or close it)?	0	0	https://issues.jboss.org/browse/AS7-3789
1701	AS7-3789	Just tried standalone-full.xml at 2855ms, so that looks pretty competitive to me. Maybe leave some optimizations for later releases? :)	1	1	https://issues.jboss.org/browse/AS7-3789
1702	AS7-3770	In order to cherrypick this in the EAP ER2 build, we would need a resolution by Feb/27th. Is this doable?	1	1	https://issues.jboss.org/browse/AS7-3770
1703	AS7-3739	Do we want to preload the cache on server startup? Since file store is not the fastest persistence mechanism in the world, IMO it will be sufficient to load the sessions when they are requested.	0	0	https://issues.jboss.org/browse/AS7-3739
1704	AS7-3715	Actually, it is indeed byte array that i used internally. Could this even be represented in a different way? If no then just reject.	1	1	https://issues.jboss.org/browse/AS7-3715
1705	AS7-3696	Can this JIRA be resolved with today's testsuite commit?	0	0	https://issues.jboss.org/browse/AS7-3696
1706	AS7-3692	This could be made better, perhaps, but IMHO it's not a critical issue. Better == optional params to the top-level resource :add operation that provide the info needed to build the child resources. Reduces things to one op. In this specific case though that would be fairly useless, since a human would never type such a command in the interactive CLI -- just too much data is needed. They would end up using a script. In which case, why not use the batch commands?	1	1	https://issues.jboss.org/browse/AS7-3692
1707	AS7-3659	Will this only be in EAP, as a manual test?	1	1	https://issues.jboss.org/browse/AS7-3659
1708	AS7-3659	Vladimir, Ironjacamar 1.0.8 is in: https://github.com/jbossas/jboss-as/commit/9872ebdf9c6a7cfff14c1d7f01510fcaeba5a844#L9R150 Can you please rebase and send a pull request?	1	1	https://issues.jboss.org/browse/AS7-3659
1709	AS7-3633	Should this be a Blocker? Major issues for 7.1.0.Final are in danger of being pushed out at any time, so bump the priority if that shouldn't happen.	0	0	https://issues.jboss.org/browse/AS7-3633
1710	AS7-3633	Is this still an issue?	1	0	https://issues.jboss.org/browse/AS7-3633
1711	AS7-3607	Can someone from the JBoss Web team take look please?	0	0	https://issues.jboss.org/browse/AS7-3607
1712	AS7-3607	This handles the opposite case - if I prefer an IPv6 stack and i'm given an IPv4 address, convert it to an IPv6 mapped address. But even with this fix, if I do prefer an IPv6 stack and I do pass in an IPv4 address, it still comes out as an IPv4 socket. The problem is further down. I tried locating the code for AprEndpoint to insert some debugging statements to confirm my suspicions. J-F, where is it located (i.e. which source repo and which jar?) I see JBoss Web 7.0.10.Final included in the pom; when I go to JBoss Web page, no hint of any version called 7.0.10.Final, as far as I could tell.	1	1	https://issues.jboss.org/browse/AS7-3607

#	Requirement	Questions	Answered	Successfully Answered	URL
1713	AS7-3561	Hello, This change broke IDEA support of JBoss 7 (http://youtrack.jetbrains.com/issue/IDEA-81489) In IDEA, we have used 'validate-address' operation in 2 cases: 1) check whether server is ready for remote administration - call to 'validate-address' without address 2) check whether deployment is registered on the server - call to 'validate-address' with 'deployment=<artifact name>' as the address As I can see from the AS7-3561 and from the actual CLI of the JBoss 7.1.0 final, now, validate-address: - could be executed without address only (execution with 'deployment=<artifact name>' fails with JBAS014739: No handler for validate-address at address [("deployment" => "=<artifact name>")] - requires the 'value' attribute But, 'value' attribute seems not to expect address or just broken itself, since every value I've passed as the 'value' fails with JBAS014749: Operation handler failed: null (I've tried empty string, "/", "/deployment=<deployed artifact name>/"...) Could you please describe what is expected as 'value' or confirm 'validate-address' is broken now? Regards, Michael	1	0	https://issues.jboss.org/browse/AS7-3561
1714	AS7-3556	Thomas, could you please list the test classes involved so I can be sure we're talking about the same thing?	1	1	https://issues.jboss.org/browse/AS7-3556
1715	AS7-3556	These unit tests are actually quite valuable. They test the inner workings of these classes and therefore mock out their environment. However I do see value in providing additional system tests for the OSGi subsystem management. Shall we rename this bug to 'Add system tests for OSGi subsystem model'?	1	1	https://issues.jboss.org/browse/AS7-3556
1716	AS7-3510	Have you tried any 7.1.0 release candidate?	1	0	https://issues.jboss.org/browse/AS7-3510
1717	AS7-3489	I don't think we have the best coverage of the HTTP interface at the moment - doesn't this set of ignores completely remove all testing of the HTTP interface?	1	1	https://issues.jboss.org/browse/AS7-3489
1718	AS7-3488	That doesn't make sense to me. At this point "wilma" is not a valid cache name. Why is that better than leaving default-cache undefined?	1	1	https://issues.jboss.org/browse/AS7-3488
1719	AS7-3488	Also, relying on batch mode is awkward from a usability standpoint. Can't we solve the chicken and egg problem? What if you add a required "default-cache-type" parameter to the add operation? Then the add operation would just create a cache for the default-cache. The user would probably need to go back and configure the new cache. But at least you would always maintain a valid configuration. So this would execute successfully: {noformat}/subsystem=infinispan/cache-container=fred:add(default-cache=wilma,default-cache-type=LOCAL){noformat} And read-resource would immediately look like this: {noformat}/subsystem=infinispan/cache-container=fred/:read-resource(recursive=true,proxies=false,include-runtime=false,include-defaults=true) { "outcome" => "success", "result" => { "aliases" => undefined, "default-cache" => "wilma", "distributed-cache" => undefined, "eviction-executor" => undefined, "invalidation-cache" => undefined, "jndi-name" => undefined, "listener-executor" => undefined, "replicated-cache" => undefined, "replication-queue-executor" => undefined, "start" => undefined, "transport" => undefined, "local-cache" => {"wilma" => { "batching" => false, "cache-mode" => "LOCAL", "eviction" => undefined, "expiration" => undefined, "file-store" => undefined, "jdbc-store" => undefined, "locking" => undefined, "name" => "wilma", "remote-store" => undefined, "store" => undefined, "transaction" => undefined }} }, "response-headers" => {"process-state" => "reload-required"} }{noformat}	1	1	https://issues.jboss.org/browse/AS7-3488
1720	AS7-3488	It's impossible for the cache to be present. You can't create the cache until you create the cache container. Is my proposed solution workable?	1	1	https://issues.jboss.org/browse/AS7-3488
1721	AS7-3488	Was this completed?	0	0	https://issues.jboss.org/browse/AS7-3488
1722	AS7-3419	Tomaz, could you look at this? Basically this means providing proper expression support for this attribute. This subsystem isn't using AttributeDefinition. For 7.1.0 I absolutely don't want to try and change it to that (too late). If changing this one element to use AttributeDefinition is the efficient way to fix this, it's ok by me. It will look a bit odd having just one AttributeDefinition, but c'est la vie. AttributeDefinition isn't required for proper expression support though; it just has most of the code already there.	1	1	https://issues.jboss.org/browse/AS7-3419
1723	AS7-3419	What does not work? I have tested it and it work properly... i have tested with this configuration and all expressions resolved properly: {code:xml} <connector name="https" protocol="HTTP/1.1" scheme="https" socket-binding="https" secure="true"> <ssl name="ssl.name" key-alias="alias" password="${ssl.keystore.password:defaultPassword}" certificate-key-file="${jboss.server.config.dir}/keystore.p12" verify-client="true" ca-certificate-file="${jboss.server.config.dir}/megapos.jks" ca-certificate-password="${my.custom.sysprop:password}" keystore-type="PKCS12" truststore-type="JKS"/> </connector> {code} I didn't test with vault expression but other then that all expressions resolved properly	1	1	https://issues.jboss.org/browse/AS7-3419
1724	AS7-3371	What exactly is that stress test, testing? I haven't taken a look. But without that detail, I don't think this will be assigned to the right person.	0	0	https://issues.jboss.org/browse/AS7-3371
1725	AS7-3371	Tomas, Can you try setting native="false" in the web config?	1	1	https://issues.jboss.org/browse/AS7-3371
1726	AS7-3371	I don't know much about thread pool size settings. I tried to follow http://www.masterthebos.com/jboss-application-server/341-jboss-as-7-performance-tuning.html?start=3, but I am not really sure about that numbers. I tried to double values for core-threads and max-threads, but didn't see any change. Can somebody tell me more or point me to some doc?	0	0	https://issues.jboss.org/browse/AS7-3371
1727	AS7-3363	Jan,what are allowed socket bindings ? Is this a fix set of values or do they change in runtime ?	1	1	https://issues.jboss.org/browse/AS7-3363
1728	AS7-3359	Mirek, can you create a pull request for this? You might have to wait until a ha profile is available upstream (JBCTS-1230)	1	1	https://issues.jboss.org/browse/AS7-3359
1729	AS7-3357	Is this planned to be exposed in metadata?	0	0	https://issues.jboss.org/browse/AS7-3357
1730	AS7-3347	So was the driver list command added?	1	0	https://issues.jboss.org/browse/AS7-3347
1731	AS7-3322	Darran: I was going to explain some details for Tomaz, but now I'm confused. The schema for the <management-interfaces><http-management-interface> type is: <xs:complexType name="http-management-socket-binding-refType"> <xs:annotation> <xs:documentation> Reference to the configurations of the sockets to be used by a standalone server's exposed HTTP and HTTPS management interface. </xs:documentation> </xs:annotation> <xs:attribute name="http" type="xs:string" use="optional"> <xs:annotation> <xs:documentation> Name of a socket-binding configuration declared in the server's socket-binding-group to use for a HTTP socket. </xs:documentation> </xs:annotation> </xs:attribute> <xs:attribute name="https" type="xs:string" use="optional"> <xs:annotation> <xs:documentation> Name of a socket-binding configuration declared in the server's socket-binding-group to use for a HTTPS socket. </xs:documentation> </xs:annotation> </xs:attribute> </xs:complexType> There are 2 attributes, one for HTTP one for HTTPS. What's missing?	1	1	https://issues.jboss.org/browse/AS7-3322

#	Requirement	Questions	Answered	Successfully Answered	URL
1732	AS7-3322	I really don't think this is still the issues. as configuration that I use for quite a while looks like this: {code:xml}<connector name="https" protocol="HTTP/1.1" scheme="https" socket-binding="https" secure="true">{code} and under socket-binding I have: {code:xml}<socket-binding name="https" port="443"/>{code} Am I missing something?	1	1	https://issues.jboss.org/browse/AS7-3322
1733	AS7-3309	This issue is being relocated to 7.1.2 so that 7.1.1 only contains criticals, blockers, and EAP LA issues (which are probably not yet at the right status, or need an extra triage pass). If these issues are completed in a 7.1.1 timeframe please change the fix version. Although if you forget I will bulk correct.	1	1	https://issues.jboss.org/browse/AS7-3309
1734	AS7-3290	<smarlow> galderz: which code module contains the code for constructing the global component registry (that gets Infinispan core to load the Hibernate Infinispan service entries)? <ttarrant> smarlow, the module that constructs that is org.jboss.as.clustering.infinispan <ttarrant> smarlow, the gcr is constructed with that module's classloader <ttarrant> smarlow, you probably need to add hibernate-infinispan to that module <ttarrant> with services exported <ttarrant> smarlow, can you try ? <smarlow> ttarrant, galderz: org.jboss.as.clustering.infinispan doesn't really know about the Hibernate Infinispan module, but even if org.jboss.as.clustering.infinispan had a dependency on Hibernate Infinispan, I'm not sure what that would do but I can try it as a hacj <ttarrant> smarlow, you are consuming caches created by the infinispan clustering subsystem are you not ? <smarlow> ttarrant: yes <smarlow> I believe that is the case <ttarrant> smarlow, the infinispan GlobalConfigurationRegistry is initialized by that subsystem with its module's classloader <ttarrant> smarlow, therefore the ServiceLoader will find only the services "exported" at that level <ttarrant> and visible to that classloader * smarlow is adding <module name="org.hibernate.infinispan" services="export"/> to org.jboss.as.clustering.infinispan <smarlow> ttarrant, galderz: ^ worked! :)	1	1	https://issues.jboss.org/browse/AS7-3290
1735	AS7-3262	[Host Controller] 12:13:47,878 INFO [org.jboss.as.server.deployment.repository.impl] (management-handler-threads - 4) JBAS014900: Content added at location /home/jenkins/jenkins-work/workspace/as7-param-pull/testsuite/domain/target/domains/CLITestSuite/master/content/da/39a3ee5e6b4b0d3255bfef95601890afd80709/content [Server:main-one] 12:13:48,309 INFO [org.jboss.as.server.deployment] (MSC service thread 1-1) Starting deployment of "SimpleServlet.war" [Server:main-three] 12:13:48,316 INFO [org.jboss.as.server.deployment] (MSC service thread 1-3) Starting deployment of "SimpleServlet.war" [Server:other-two] 12:13:48,319 INFO [org.jboss.as.server.deployment] (MSC service thread 1-5) Starting deployment of "SimpleServlet.war" [Server:main-one] 12:13:48,349 ERROR [org.jboss.msc.service.fail] (MSC service thread 1-4) MSC00001: Failed to start service jboss.deployment.unit."SimpleServlet.war".STRUCTURE: org.jboss.msc.service.StartException in service jboss.deployment.unit."SimpleServlet.war".STRUCTURE: Failed to process phase STRUCTURE of deployment "SimpleServlet.war" [Server:main-one] at org.jboss.as.server.deployment.DeploymentUnitPhaseService.start(DeploymentUnitPhaseService.java:121) [jboss-as-server-7.1.0.Final-SNAPSHOT.jar:7.1.0.Final-SNAPSHOT] [Server:main-one] at org.jboss.msc.service.ServiceControllerImpl$StartTask.startService(ServiceControllerImpl.java:1824) [jboss-msc-1.0.1.GA.jar:1.0.1.GA] [Server:main-one] at org.jboss.msc.service.ServiceControllerImpl$StartTask.run(ServiceControllerImpl.java:1759) [jboss-msc-1.0.1.GA.jar:1.0.1.GA] [Server:main-one] at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) [:1.6.0_29] [Server:main-one] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) [:1.6.0_29] [Server:main-one] at java.lang.Thread.run(Thread.java:662) [:1.6.0_29] [Server:main-one] Caused by: org.jboss.as.server.deployment.DeploymentUnitProcessingException: Failed to mount deployment content [Server:main-one] at org.jboss.as.server.deployment.module.DeploymentRootMountProcessor.deploy(DeploymentRootMountProcessor.java:82) [jboss-as-server-7.1.0.Final-SNAPSHOT.jar:7.1.0.Final-SNAPSHOT] [Server:main-one] at org.jboss.as.server.deployment.DeploymentUnitPhaseService.start(DeploymentUnitPhaseService.java:115) [jboss-as-server-7.1.0.Final-SNAPSHOT.jar:7.1.0.Final-SNAPSHOT] [Server:main-one] ... 5 more [Server:main-one] Caused by: java.util.zip.ZipException: error in opening zip file [Server:main-one] at java.util.zip.ZipFile.open(Native Method) [:1.6.0_29] [Server:main-one] at java.util.zip.ZipFile.<init>(ZipFile.java:127) [:1.6.0_29] [Server:main-one] at java.util.zip.ZipFile.<init>(ZipFile.java:143) [:1.6.0_29] [Server:main-one] at org.jboss.vfs.VFSUtils.unzip(VFSUtils.java:845) [jboss-vfs-3.1.0.CR1.jar:3.1.0.CR1] [Server:main-one] at org.jboss.vfs.VFS.mountZipExpanded(VFS.java:536) [jboss-vfs-3.1.0.CR1.jar:3.1.0.CR1] [Server:main-one] at org.jboss.vfs.VFS.mountZipExpanded(VFS.java:567) [jboss-vfs-3.1.0.CR1.jar:3.1.0.CR1] [Server:main-one] at org.jboss.as.server.deployment.repository.impl.ServerDeploymentRepositoryImpl.mountDeploymentContent(ServerDeploymentRepositoryImpl.java:73) [jboss-as-deployment-repository-7.1.0.Final-SNAPSHOT.jar:7.1.0.Final-SNAPSHOT] [Server:main-one] at org.jboss.as.server.deployment.module.DeploymentRootMountProcessor.deploy(DeploymentRootMountProcessor.java:78) [jboss-as-server-7.1.0.Final-SNAPSHOT.jar:7.1.0.Final-SNAPSHOT] [Server:main-one] ... 6 more [Server:main-one] [Server:other-two] 12:13:48,345 ERROR [org.jboss.msc.service.fail] (MSC service thread 1-2) MSC00001: Failed to start service jboss.deployment.unit."SimpleServlet.war".STRUCTURE: org.jboss.msc.service.StartException in service jboss.deployment.unit."SimpleServlet.war".STRUCTURE: Failed to process phase STRUCTURE of deployment "SimpleServlet.war" [Server:other-two] at	1	1	https://issues.jboss.org/browse/AS7-3262
1736	AS7-3251	Even though ConfigAdmin is now an OSGi-independent subsystem, will we keep the console GUI part under OSGi? Or should we introduce an independent ConfigAdmin subsystem in the Console too?	1	1	https://issues.jboss.org/browse/AS7-3251
1737	AS7-3251	The latter. The reasoning is that the OSGi subsystem may get removed completely, which should then also happen (by itself) in the console. Could you please check that the console can react adequately to subsystem availability?	1	1	https://issues.jboss.org/browse/AS7-3251
1738	AS7-3241	Hi Kurt, Do you plan to do this before Feb 7? (My guess is no)	0	0	https://issues.jboss.org/browse/AS7-3241
1739	AS7-3229	Scott, are you still getting this error? SFSB serialization has changed a bit since this jira was opened - and I wonder if this is jira isn't obsolete.	0	0	https://issues.jboss.org/browse/AS7-3229
1740	AS7-3229	Can this be closed, or is there still something outstanding?	0	0	https://issues.jboss.org/browse/AS7-3229
1741	AS7-3203	Hi Stuart, I can't find anywhere in the class loading guide (https://docs.jboss.org/author/display/AS71/Class+Loading+in+AS7) that mentions the <module .. export="true" />. Can you quote in detail where I can find it? Thanks. John Wu	1	0	https://issues.jboss.org/browse/AS7-3203
1742	AS7-3195	Is this fix expected to ship with 7.1.2, or is there more work required?	1	1	https://issues.jboss.org/browse/AS7-3195
1743	AS7-3195	My app was defining properties in an external .properties file, then loading them into system properties for later use. They seemeed to be available in time for expansion into persistence.xml with 7.0.2, but that behavior seems to be different for 7.1.1. Defining the property directly in standalone.xml seems to work, and make the values available for substitution during deployment. Properties were externalized so that deployables could ship with their config files in a 'hot' directory, and deployed to different as7 servers as needed. This let us define one hot directory in standalone.xml, and drop deployable .properties files there, instead of having to mangle standalone.xml for each deployable. For properties that need to be available when a deployable is being deployed, is standalone.xml the right place to put those? Is there an approved way to externalize those? The admin guide (https://docs.jboss.org/author/display/AS7/General+configuration+concepts) doesn't mention any alternatives. Thanks	1	1	https://issues.jboss.org/browse/AS7-3195
1744	AS7-3195	This bug is now present in EAP 6.1 on OpenShift. Is it slated for fix there?	0	0	https://issues.jboss.org/browse/AS7-3195

#	Requirement	Questions	Answered	Successfully Answered	URL
1745	AS7-3174	Would it make sense to make this even more generic? AFAICT only top-level stuff can be excluded.	1	1	https://issues.jboss.org/browse/AS7-3174
1746	AS7-3166	Rado, Is this reproducible on master? This issue was opened when Infinispan was at 5.1.0.CR1 - Infinispan is now at 5.1.0.CR4.	0	0	https://issues.jboss.org/browse/AS7-3166
1747	AS7-3159	Thanks Stuart, but what about Hibernate Search? does it have the same reason to reject?	0	0	https://issues.jboss.org/browse/AS7-3159
1748	AS7-3148	"If D has a dependency on B and they are deployed at the same time I'll get:" .... Do you mean, if they're *not* deployed at the same time? If the scanner picks up both in the same run you shouldn't be getting this error.	1	1	https://issues.jboss.org/browse/AS7-3148
1749	AS7-3148	At the time of writing I was receiving the error in the description when A and D were picked up in the same scan. I hadn't tried this in a while as my deployment script currently waits on the creation of E.ear.deployed. I just tried it again and now get a different error that looks more like a race condition. This might not be the right JIRA for it though if what I am requesting is already the expected behaviour. In my case B.jar that contains entities and my persistence unit which provides an EntityManagerFactory via JNDI. D.war contains various seam2 components and its ManagedPersistenceContext tries to look up the EMF when the seam application starts. If both A and D are deployed in the same run the EMF is not available when D is starting. Is this intended or a bug? I will attach a log showing a deployment run where both are picked up at the same time and fail and then when they are deployed sequentially and succeed. The error when they are deployed concurrently is below: Caused by: java.lang.IllegalArgumentException: EntityManagerFactory not found in JNDI : java:/PlutoMainEntityManagerFactory at com.britensw.pluto.web.seam.ManagedPersistenceContext.getEntityManagerFactoryFromJndiOrValueBinding(ManagedPersistenceContext.java:188) [classes:] at com.britensw.pluto.web.seam.ManagedPersistenceContext.createContext(ManagedPersistenceContext.java:91) [classes:] at com.britensw.pluto.web.seam.ManagedPersistenceContext.create(ManagedPersistenceContext.java:84) [classes:] at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) [:1.6.0_30] at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) [:1.6.0_30] at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) [:1.6.0_30] at java.lang.reflect.Method.invoke(Method.java:597) [:1.6.0_30] at org.jboss.seam.util.Reflections.invoke(Reflections.java:22) [jboss-seam-2.2.2.Final.jar:2.2.2.Final] at org.jboss.seam.util.Reflections.invokeAndWrap(Reflections.java:144) [jboss-seam-2.2.2.Final.jar:2.2.2.Final] at org.jboss.seam.Component.callComponentMethod(Component.java:2283) [jboss-seam-2.2.2.Final.jar:2.2.2.Final] at org.jboss.seam.Component.callCreateMethod(Component.java:2198) [jboss-seam-2.2.2.Final.jar:2.2.2.Final] at org.jboss.seam.Component.newInstance(Component.java:2158) [jboss-seam-2.2.2.Final.jar:2.2.2.Final] ... 37 more Caused by: javax.naming.NameNotFoundException: Error looking up PlutoMainEntityManagerFactory, service service jboss.naming.context.java.PlutoMainEntityManagerFactory is not started at org.jboss.as.naming.ServiceBasedNamingStore.lookup(ServiceBasedNamingStore.java:126) at org.jboss.as.naming.ServiceBasedNamingStore.lookup(ServiceBasedNamingStore.java:74) at org.jboss.as.naming.NamingContext.lookup(NamingContext.java:177) at org.jboss.as.naming.InitialContext.lookup(InitialContext.java:113) at org.jboss.as.naming.NamingContext.lookup(NamingContext.java:213) at javax.naming.InitialContext.lookup(InitialContext.java:392) [:1.6.0_30] at com.britensw.pluto.web.seam.ManagedPersistenceContext.getEntityManagerFactoryFromJndiOrValueBinding(ManagedPersistenceContext.java:186) [classes:] ... 48 more	0	0	https://issues.jboss.org/browse/AS7-3148
1750	AS7-3147	This issue is being relocated to 7.1.2 so that 7.1.1 only contains criticals, blockers, and EAP LA issues (which are probably not yet at the right status, or need an extra triage pass). If these issues are completed in a 7.1.1 timeframe please change the fix version. Although if you forget I will bulk correct.	1	1	https://issues.jboss.org/browse/AS7-3147
1751	AS7-3142	Same problem on our side (see stacktrace)! Can someone help us ? {code} 10:54:05,428 ERROR [org.jboss.msc.service.fail] (MSC service thread 1-3) MSC00001: Failed to start service jboss.deployment.unit."XXX.war".STRUCTURE: org.jboss.msc.service.StartException in service jboss.deployment.unit."XXX.war".STRUCTURE: Failed to process phase STRUCTURE of deployment "XXX.war" at org.jboss.as.server.deployment.DeploymentUnitPhaseService.start(DeploymentUnitPhaseService.java:121) [jboss-as-server-7.1.0.CR1b.jar:7.1.0.CR1b] at org.jboss.msc.service.ServiceControllerImpl$StartTask.startService(ServiceControllerImpl.java:1824) [jboss-msc-1.0.1.GA.jar:1.0.1.GA] at org.jboss.msc.service.ServiceControllerImpl$StartTask.run(ServiceControllerImpl.java:1759) [jboss-msc-1.0.1.GA.jar:1.0.1.GA] at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) [:1.6.0_29] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) [:1.6.0_29] at java.lang.Thread.run(Thread.java:680) [:1.6.0_29] Caused by: java.lang.IllegalArgumentException: Given parent is not an ancestor of this virtual file at org.jboss.vfs.VirtualFile.getPathNameRelativeTo(VirtualFile.java:116) at org.jboss.vfs.VirtualFile.getPathNameRelativeTo(VirtualFile.java:122) at org.jboss.vfs.VirtualFile.getPathNameRelativeTo(VirtualFile.java:110) at org.jboss.as.server.deployment.module.ManifestClassPathProcessor.createAdditionalModule(ManifestClassPathProcessor.java:193) [jboss-as-server-7.1.0.CR1b.jar:7.1.0.CR1b] at org.jboss.as.server.deployment.module.ManifestClassPathProcessor.handlingExistingClassPathEntry(ManifestClassPathProcessor.java:185) [jboss-as-server-7.1.0.CR1b.jar:7.1.0.CR1b] at org.jboss.as.server.deployment.module.ManifestClassPathProcessor.deploy(ManifestClassPathProcessor.java:162) [jboss-as-server-7.1.0.CR1b.jar:7.1.0.CR1b] at org.jboss.as.server.deployment.DeploymentUnitPhaseService.start(DeploymentUnitPhaseService.java:115) [jboss-as-server-7.1.0.CR1b.jar:7.1.0.CR1b] ... 5 more {code}	0	0	https://issues.jboss.org/browse/AS7-3142
1752	AS7-3116	Can you open a new JIRA for this as it is a different issue.	1	1	https://issues.jboss.org/browse/AS7-3116
1753	AS7-3114	Another factor in the equation is that whether or not to have the console needs to be a configurable attribute, probably on the http management interface element. It's quite reasonable that people using AS 7 / EAP will want to use the HTTP management connector but not run a console. So you have a matrix: One access: Is the console available? (yes/no) Other access: Does user want console? (yes/no/didn't say)	1	1	https://issues.jboss.org/browse/AS7-3114

#	Requirement	Questions	Answered	Successfully Answered	URL
1754	AS7-3114	Tomaz, can you look into this one? There are two aspects: 1) Add a non-required boolean config attribute the the <management-interfaces><http-interface> element to allow users to disable the console. Default is true. 2) Detect the absence of the module that provides the console, and deal with that absence appropriately. How to deal with the absence of the console is an open question. One thought I had was to throw an exception if the config attribute was actually set to "true". If the config attribute is undefined or false then deal with it using ConsoleMode.NO_CONSOLE.	1	1	https://issues.jboss.org/browse/AS7-3114
1755	AS7-3036	Thanks for reply! But, where can I find the bundle files? Can I get from the original project repos and then deploy into AS7?	1	1	https://issues.jboss.org/browse/AS7-3036
1756	AS7-3030	I see the jacorb build (2.3.1.jboss-1) in the JBoss repository: https://repository.jboss.org/nexus/content/repositories/thirdparty-releases/org/jacorb/jacorb/ Should there also be a matching version of jacorb-idl-compiler?	1	1	https://issues.jboss.org/browse/AS7-3030
1757	AS7-2999	Hi, what is the current status of this bug? I was told that it should be ready in the beginning of this week ... and that my pull request will be hold off until then. Thanks for the info.	1	1	https://issues.jboss.org/browse/AS7-2999
1758	AS7-2999	If I remove @Ignore(AS7-2999) for example in InjectionAnnSFSBtoSFSBTestCase, I get an exception like this: javax.ejb.NoSuchEJBException: Could not find SFSB InjectionSFSBtoSFSB with {[-32, 80, 121, 29, -10, 122, 68, 78, -87, 114, -8, 77, -101, -123, -17, -11]} at org.jboss.as.ejb3.component.stateful.StatefulComponentInstanceInterceptor.processInvocation(StatefulComponentInstanceInterceptor.java:62) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.as.ejb3.tx.CMTTxInterceptor.invokeInOurTx(CMTTxInterceptor.java:228) at org.jboss.as.ejb3.tx.CMTTxInterceptor.required(CMTTxInterceptor.java:304)	1	1	https://issues.jboss.org/browse/AS7-2999
1759	AS7-2991	Is there any documentation explaining how to enable gzip compression in jboss as 7.1.1 (or maybe before the official release?)?	1	1	https://issues.jboss.org/browse/AS7-2991
1760	AS7-2991	Any chance this regression is related to your change? AS7-4124 Which repo is the change you made at?	0	0	https://issues.jboss.org/browse/AS7-2991
1761	AS7-2991	So is this actually resolved? I added -Dorg.apache.coyote.http11.Http11Protocol.COMPRESSION=true to the vm properties but my pages are not being compressed.	0	0	https://issues.jboss.org/browse/AS7-2991
1762	AS7-2991	How to enable this ? I'm using JBoss AS 7.1.1 and compression is not enabled by default.	1	1	https://issues.jboss.org/browse/AS7-2991
1763	AS7-2991	Can GZIP compression be enable on the HTTP connector in standalone.xml?	1	0	https://issues.jboss.org/browse/AS7-2991
1764	AS7-2991	It seems that JBoss AS 7.1.1 HTTP compression only compress the document itself and not static files (e.g., *.js, *.css, ...). Is there a way to define which the compressible MimeTypes are?	1	0	https://issues.jboss.org/browse/AS7-2991
1765	AS7-2984	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 765759|https://bugzilla.redhat.com/show_bug.cgi?id=765759] Prabhat Jha <prabhat.jha@jboss.com> made a comment on jira ISPN-1799 Can this be fixed so that it gets fixed in next ER build for JDG?	0	0	https://issues.jboss.org/browse/AS7-2984
1766	AS7-2984	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 765759|https://bugzilla.redhat.com/show_bug.cgi?id=765759] Tristan Tarrant <ttarrant@redhat.com> made a comment on jira ISPN-1799 Dan could you please try and get it in 5.1.4.CR1 ?	0	0	https://issues.jboss.org/browse/AS7-2984
1767	AS7-2984	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 765759|https://bugzilla.redhat.com/show_bug.cgi?id=765759] Tristan Tarrant <ttarrant@redhat.com> made a comment on jira ISPN-1799 Dan could you please try and get it in 5.1.4.CR1 or at least a patch for JDG	0	0	https://issues.jboss.org/browse/AS7-2984
1768	AS7-2984	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 765759|https://bugzilla.redhat.com/show_bug.cgi?id=765759] dex chen <dex80526@yahoo.com> made a comment on jira ISPN-1799 we saw this error in our 3 node cluster in which the network connectivity to one of nodes is not reliable and has high latency. Does this exception have anything to do with the network connectivity and high latency?	0	0	https://issues.jboss.org/browse/AS7-2984
1769	AS7-2958	Misty Stanley-Jones <misty@redhat.com> made a comment on [bug 831442|https://bugzilla.redhat.com/show_bug.cgi?id=831442] Sande, I think you might be a good fit for this topic. Maybe something like the following: About HA Singleton Differences from the EAP 5 Implementation Configure the HA Singleton Dummy Interface Use HA Singleton Can you give the above links a quick read through within the next week or so, and give Rich an estimate of how long it will take to write these topics? You can work with Tom to get an asynchronous addition to the Development Guide. In fact this might be a good candidate for a rough draft in Confluence for the developers and community to help with.	0	0	https://issues.jboss.org/browse/AS7-2958
1770	AS7-2954	Is there a reason why this is a bad idea? I mean all assert does is throw an exception? Are you trying to limit the usage of junit to allow usage of testng?	0	0	https://issues.jboss.org/browse/AS7-2954
1771	AS7-2954	What is the reasoning behind this?	0	0	https://issues.jboss.org/browse/AS7-2954
1772	AS7-2954	That was my question as well :) Ondrej, could you explain this?	0	0	https://issues.jboss.org/browse/AS7-2954
1773	AS7-2891	I don't know where it was configured in AS 6, also. Was it using a realm element in server.xml (or context.xml) ?	0	0	https://issues.jboss.org/browse/AS7-2891
1774	AS7-2881	Shouldn't this JIRA be assigned to youself Max?	0	0	https://issues.jboss.org/browse/AS7-2881
1775	AS7-2881	Thanks Max (y) When you say codefreeze next week you mean of 7.1.0.final? Should we expect to have the release that was planned to Tuesday, Feb7th, to be next week?	0	0	https://issues.jboss.org/browse/AS7-2881
1776	AS7-2871	Won't Fix {code} (07:42:38 PM) tdiesler: Nihility, I guess this attribute is special in the sense that no other subsystem is activated lazily. The attr pertains to subsystem startup and is not a configuration of the subsystem itself. WDYT? (07:45:01 PM) Nihility: tdiesler: yeah i think you have a good point that this could be a generic notion, i really need bstansberry's thoughts on this though (07:45:21 PM) Nihility: bstansberry: hey what do you think about ^ (07:46:00 PM) ***bstansberry looks (07:49:08 PM) bstansberry: Nihility: I don't particularly object to attributes on the subsystem element (07:49:19 PM) bstansberry: and I now see 5 open JIRAs to fix them (07:49:59 PM) Nihility: bstansberry: yeah i was never really against them. I know David had some reasons, but I dont remeber what they were or if I ever agreed completely :) (07:50:00 PM) bstansberry: I don't really believe those are going to be fixed (07:50:23 PM) bstansberry: Nihility: the basic reason is readability (07:50:32 PM) bstansberry: AIUI (07:51:06 PM) bstansberry: tdiesler: I'm fine with leaving that attribute (07:51:14 PM) Nihility: yeah i guess in some cases we had some bad attributes on the susbsytem level, activation does seem to belong there though, i mean i buy tdiesler's argument on this one {code}	0	0	https://issues.jboss.org/browse/AS7-2871
1777	AS7-2848	Tomaz, can you take a look at this? I think this should be a matter of registering ReloadRequiredWriteAttributeHandler for all the attributes. This subsystem has 1 resource with a big long list of attributes. Potentially some attributes could be immediately changed in the runtime service (I have no idea which, and quite possibly none could be) but looking into that is a post-7.1.0.Final task If possible it would be nice to convert this subsystem to use ResourceDefinition, but that's better done after 7.1.0.Final since we have much more critical things to do.	1	1	https://issues.jboss.org/browse/AS7-2848

#	Requirement	Questions	Answered	Successfully Answered	URL
1778	AS7-2792	Rosta This question is best placed for Brian, but i'll present my understanding of it as it applies to the JGroups subsystem. If I understand correctly, this issue is not particular to socket-binding definitions, but to any model attribute whose value is an expression. When an attribute whose value is an expression is viewed from the management interface, the unresolved expression is displayed, not the resolved value, as it reflects the value of the attribute in the persistent model, and the persistent model needs to be server-instance neutral. Such expressions are resolved when services are started in a particular server context, but those resolved values are generally not available for viewing. The option include-runtime does not cause the resolved value of the attribute to be displayed, but rather displays attributes which are defined as runtime attributes (i.e. are not part of the persistent configuration, which is server instance neutral) - not the same thing. It would be possible to define a runtime attribute which read and made observable the resolved multicast address, but that would be one particular solution to a more general problem. I expect that you would like to have an option where you may connect to a server and see the model with all expressions resolved in the context of that server's environment; like: {noformat} :read-resource(resolve-expressions=true) {noformat} so that the resolved model can be viewed. Is that correct?	0	0	https://issues.jboss.org/browse/AS7-2792
1779	AS7-2787	What would you prefer the flow to be? I can add support for 'y' or do you think I should just keep repeating "Is this correct yes/no?" until they say "no"?	0	0	https://issues.jboss.org/browse/AS7-2787
1780	AS7-2770	In the past there were the same issue with GWT opened check here: http://code.google.com/p/google-web-toolkit/issues/detail?id=1898 and since GWT 1.5 there has been added a method ensureDebugId()for UIObject We need to change current implementation somewhat in following way... 1. Create new button ToolButton button = new ToolButton("Remove", new ClickHandler() { @Override public void onClick(ClickEvent event) { Feedback.confirm("Remove Module", "Really remove module?", new Feedback.ConfirmationHandler() { @Override public void onConfirmation(boolean isConfirmed) { if(isConfirmed) { Module module = ((SingleSelectionModel<Module>) modules.getSelectionModel()).getSelectedObject(); presenter.onRemoveModule(form.getEditedEntity(), module); } } }); } }); 2.Set Debug Id button.ensureDebugId("RemoveButtonEEView"); 3. Add to Panel moduleTools.addToolButtonRight(button);	0	0	https://issues.jboss.org/browse/AS7-2770
1781	AS7-2763	can anybody please have an update on this ? Since Hibernate Multi-tenant is not fully supportive due to the conflict with second level caching we can't think about Hibernate Multi-tenant. So the only way in front of us is to use one SessionFactory per tenant by registering SessionFactory using "hibernate.session_factory_name" configuration property. Since JBoss AS 7 is not supporting JNDI binding, we are really in trouble.	1	0	https://issues.jboss.org/browse/AS7-2763
1782	AS7-2763	Hi Stuart, Can you please tell the implementation details ?	0	0	https://issues.jboss.org/browse/AS7-2763
1783	AS7-2756	Is the same strength check applied to management user? I would assume management user passwords have more stringent requirement, at least should be the same as app user. But I can create a management user using password that consists of all letters (more than 8 letters). Can we validate the password immediately after the password is entered, before asking the user to enter a potentially useless password a second time, and before asking the user to assign a role. Currently the validation kicks in after adding role: Realm (ApplicationRealm) : Username : user1 Password : Re-enter Password : What roles do you want this user to belong to? (Please enter a comma separated list, or leave blank for none)[ ]: * Error * JBAS015266: Password must not have atleast one digit. The above error message is confusing, and that issue is tracked in AS7-5631	1	1	https://issues.jboss.org/browse/AS7-2756
1784	AS7-2751	we do basically need all of it? but why is your resource description different from mine? are you using the latest AS7 master?	1	1	https://issues.jboss.org/browse/AS7-2751
1785	AS7-2751	I used [domain@localhost:9999 /] /profile=full-ha/subsystem=modcluster:read-resource-description(recursive=true) vs. [domain@localhost:9999 /] /profile=full-ha/subsystem=modcluster:read-resource(recursive=true) The read-resource operation gives me the same list as you posted but read-resource-description operation gives a bit more attributes, which looks strange. Shouldn't it be only detailed description of the same sets of attributes ? Not sure if we need all of them, it depends what we want to manage through UI config IMO. Yes I do use the latest master jboss-as SNAPSHOT, have updated yesterday	1	1	https://issues.jboss.org/browse/AS7-2751
1786	AS7-2746	yes, the "wrong credentials" approach is what I had in mind. do you think that's feasible? does it works across browsers?	0	0	https://issues.jboss.org/browse/AS7-2746
1787	AS7-2729	Stuart, is this issue still valid?	0	0	https://issues.jboss.org/browse/AS7-2729
1788	AS7-2703	Pull request?	0	0	https://issues.jboss.org/browse/AS7-2703
1789	AS7-2703	Whether we create new interceptors or reassemble them, we need to replicate enough information for that to happen. For example, the SFSBDestroyInterceptor constructor is passed the SFSBXPCMap to be used. Currently, I'm serializing the applicationDeploymentBagKeyName so that the (per deployment) SFSBXPCMap can be looked up (https://github.com/scottmarlow/jboss-as/commit/bce74f9775efe37977e1ad2be8aac0ff3530cfcc#L2R93). Which class would we recreate the interceptors from? This also needs to work for passivation/activation as well. Do we have a passivation/activation test case already? I just pushed an update to https://github.com/scottmarlow/jboss-as/tree/cluster7 that includes a basic cluster failover test change to include an extended persistence context. Its not really ready to be shared yet though as I'm using a Hibernate 4.0.1 snapshot build.	1	1	https://issues.jboss.org/browse/AS7-2703
1790	AS7-2695	where does the lib/name1-name2/libxxx.so, name1? name2? module which module? Did you look to http://downloads.jboss.org/jbossnative//2.0.10.GA/jbossweb-native-2.0.10.jar at least the jbossweb code could load libraries from those location.	1	1	https://issues.jboss.org/browse/AS7-2695
1791	AS7-2695	"System.loadLibrary() will automatically look there first." What? How? and there is where?	1	1	https://issues.jboss.org/browse/AS7-2695
1792	AS7-2695	{quote} "System.loadLibrary() will automatically look there first." What? How? and there is where? {quote} System.loadLibrary() always uses the caller's class loader first to locate native libs. "There" is the "lib" directory in the module which executes that method. {quote} There are few technical reasons why having product/lib/${platform}-${arch}/native.so would create some issues. {quote} For #1, correct, you should either link statically or rely upon the operating system to provide dependencies. For #2, you have it backwards - by having arch-specific directories we allow a single distribution which works on any operating system. For #3, I don't see this as any different than an OS distributing Java JARs. We don't use the OS for this case, so we shouldn't use the OS for the JNI case either, especially given that a JNI library is so tightly bound to its JAR counterpart. If you want to do it this way then you certainly can, but we don't normally support the distribution of some components separately from the AS because it's a bit of a pain from the user's perspective. If we can find a way to get around that then fine.	1	1	https://issues.jboss.org/browse/AS7-2695

#	Requirement	Questions	Answered	Successfully Answered	URL
1793	AS7-2695	What does ${os}-${arch} expand into?	1	1	https://issues.jboss.org/browse/AS7-2695
1794	AS7-2695	I put the latest jboss-native into the jboss web module, using the naming convention outline here, but I don't see anything in the log telling its using the native connector? Is there anything that is supposed to spit out?	1	1	https://issues.jboss.org/browse/AS7-2695
1795	AS7-2695	What if you have multiple modules using the same native component? Since the .os/.dll can be loaded only once per VM instance second invocation will always fail because JNI methods have already been bound to a different class loader. Next what if multiple native components have the same dependency. Static linkage of dependencies means that dependency fix would cause fixing all components using it. Anyhow, setting LD_LIBRARY_PATH to a whatever location you have installed jboss-native/native before starting AS works equally well, so I see no reason why would this be something mandatory.	1	1	https://issues.jboss.org/browse/AS7-2695
1796	AS7-2695	Jean-Frederic - you said you tested this successfully on linux. I get the following after putting the native libraries in: <jboss-as>/modules/org/jboss/as/web/main/lib/$<os>-$<arch> 07:24:39,236 INFO [org.apache.catalina.core.AprLifecycleListener] (MSC service thread 1-3) The Apache Tomcat Native library which allows optimal performance in production environments was not found on the java.library.path: /usr/lib/jvm/java-1.6.0-openjdk-1.6.0.0.x86_64/jre/lib/amd64/server:/usr/lib/jvm/java-1.6.0-openjdk-1.6.0.0.x86_64/jre/lib/amd64:/usr/lib/jvm/java-1.6.0-openjdk-1.6.0.0.x86_64/jre/../lib/amd64:/usr/java/packages/lib/amd64:/usr/lib64:/lib64:/lib:/usr/lib I also noticed in my boot.log that os was Linux, and arch was amd64, so I tried linux-x86_64, Linux-amd64, and linux-amd64, but none of them worked for me. I did notice that linux-x86_64 is working for HornetQ though, but so far no luck in this getting picked up. I also looked at the schema for jboss web, and native-true is the default for the configuration, so I shouldn't need to do anything in the configuration xml. What was your setup that worked?	0	0	https://issues.jboss.org/browse/AS7-2695
1797	AS7-2695	{quote} What if you have multiple modules using the same native component? Since the .os/.dll can be loaded only once per VM instance second invocation will always fail because JNI methods have already been bound to a different class loader. Next what if multiple native components have the same dependency. Static linkage of dependencies means that dependency fix would cause fixing all components using it. Anyhow, setting LD_LIBRARY_PATH to a whatever location you have installed jboss-native/native before starting AS works equally well, so I see no reason why would this be something mandatory. Well, since everybody seems to be 'module' oriented why not creating a separate module for all native components a. Native component can be loaded only once per VM b. In that case they could be shared among different modules c. It would solve module redeployment cause they would be loaded by parent class loader. d. User could safely deploy his native component into that module and be sure it'll be loaded by a correct classloader {quote} What are you talking about? Modules can't be redeployed. Modules are loaded once per VM, that's it. Anyway, there are no parent class loaders in AS 7. None of this makes any sense... I think you must have some wrong assumption.	1	1	https://issues.jboss.org/browse/AS7-2695
1798	AS7-2695	{quote} What are you talking about? Modules can't be redeployed. Modules are loaded once per VM, that's it. Anyway, there are no parent class loaders in AS 7. None of this makes any sense... I think you must have some wrong assumption. {quote} OK, my bad, wrong assumption. So if they are fixed why is so important to have them inside module/lib directory? This is not technical reason, its just convention you are imposing.	1	1	https://issues.jboss.org/browse/AS7-2695
1799	AS7-2695	What do you meany by: "Either the library is there and it's loaded, or it's not and it's skipped" Loaded by whom and when? I still do not understand what are you talking about. AS works out of box on my Fedora using tomcat-native rpm without any setup. Are you saying we should refuse to load anything outside module/lib?	1	1	https://issues.jboss.org/browse/AS7-2695
1800	AS7-2695	Heretics should be burnt at the stake, don't you know that ? :)	1	1	https://issues.jboss.org/browse/AS7-2695
1801	AS7-2695	Mladen, Why should we have inconsistent mechanisms for jni library loading?	1	1	https://issues.jboss.org/browse/AS7-2695
1802	AS7-2695	David, Where is the code that intercepts System.loadLibrary	1	1	https://issues.jboss.org/browse/AS7-2695
1803	AS7-2695	Jean-Frederic forwarded me your explanation for findLibrary Thanks. Now, how can we make that not hard coded? It would be nice to have some configuration file which would be used to construct different layouts.	0	0	https://issues.jboss.org/browse/AS7-2695
1804	AS7-2676	Hello. Looks like it is not working as expected. Example: script.properties: test=aaa my-script.cli: /subsystem=security/security-domain=SomeDomain:add() /subsystem=security/security-domain=SomeDomain/authentication=classic:add(login-modules=[{"code"=>"org.picketbox.datasource.security.SecureIdentityLoginModule","flag"=>"required","module-options"=>[("username"=>"someUsername"),("password"=>"${test}")]}]) And the result is: <security-domain name="SomeDomain"> <authentication> <login-module code="org.picketbox.datasource.security.SecureIdentityLoginModule" flag="required"> <module-option name="username" value="someUsername"/> <module-option name="password" value="${test}"/> </login-module> </authentication> </security-domain> Example 2: my-script.cli: /subsystem=security/security-domain=${test}:add() Result is: <security-domain name="aaa" /> What am I missing here? It is working only when I change <resolve-parameter-values> to true in jboss-cli.xml.	1	0	https://issues.jboss.org/browse/AS7-2676
1805	AS7-2676	You are commenting on a resolved issue, I would suggest first a forum post, your question is "What am I missing here?" so haven't quite reached the point to confirm as a bug. If there is a bug here it will need raising under the WFCORE project.	0	0	https://issues.jboss.org/browse/AS7-2676
1806	AS7-2667	Scott, what's the type of the context on which you're invoking getChildTarget() ? Could you please post a larger code snippet ? Thank you.	0	0	https://issues.jboss.org/browse/AS7-2667
1807	AS7-2667	I would like to bind my remote object from external JVM to Jboss AS 7 and use its JNDI as central Naming Directory. How is that possible to create/get org.jboss.msc.service.StartContext from a standalone application?	0	0	https://issues.jboss.org/browse/AS7-2667
1808	AS7-2664	I need more convincing arguments here. '=' are appended after request parameters and node types. Is it missing somewhere else? ':' is a separator between the node path (address part) and the operation name. How would the CLI know when you finished the node path and want to start an operation? I mean all these are valid ':read-resource', '/subsystem=datasources:read-resource', /subsystem=datasources/data-source=myds:read-resource'. '(' is a start of the parameter list. It's not required if the operation doesn't have parameters or does have but can be executed w/o them. So, it's questionable. Adding it automatically could be annoying.	1	1	https://issues.jboss.org/browse/AS7-2664

#	Requirement	Questions	Answered	Successfully Answered	URL
1809	AS7-2664	There would be a difficulty here in using the Tab completion to open the optional arguments, such as adding a "(" character. For many of the existing operations, the operation name can be filled with Tab completion, but another operation contains the same string of words. For example: bq. [standalone@localhost:9999 subsystem=datasources] :read-resource If you then hit Tab, you get the following: bq. read-resource read-resource-description The read-resource operation, however, also has some optional arguments. If you open the "(" bracket: bq. [standalone@localhost:9999 subsystem=datasources] :read-resource(And then hit Tab to show the options: bq. include-runtime= proxies= recursive= This would make it difficult to differentiate between the user's intention to fill the name of the operation, or add arguments. Jim, what would you think about having a {{--help}} wrapper, as raised in [AS-1758|https://issues.jboss.org/browse/AS7-1758]? This would allow the user to use the existing tab compeltion to fill out the operation name, and then use the {{operation-name --help}} option to effectively shortcut to running the {{read-operation-description(name=NameOfOperation)}}. The advantage I see here is a nice compromise between the existing Tab completion, and being able to query the use of optional arguments without interupting the workflow of the existing operation (in other words, not having to undo the Tab completion of the operation name, and retype the read-operation-description, read it, and then type out the original operation again). Thoughts?	1	1	https://issues.jboss.org/browse/AS7-2664
1810	AS7-2663	Can you update this ticket with the duplicate issue.	1	0	https://issues.jboss.org/browse/AS7-2663
1811	AS7-2653	Emanuel, there's still an @Ignore("AS7-4653") test in this class. Was that an oversight?	1	0	https://issues.jboss.org/browse/AS7-2653
1812	AS7-2649	I'm reluctant to require the client of this API to add a jboss-modules dependency in the manifest. It would be equally wrong to generate a hard dependency on apache scout in a DUP. A possible solution would be to generate a dependency on an intermediary module (i.e. org.jboss.jaxr.provider) which reexports the impl and such provides a configurable indirection layer. Another option would be to add the value of the generated dependency as a config option to the jaxr subsystem.	0	0	https://issues.jboss.org/browse/AS7-2649
1813	AS7-2640	What if I make a typo can I up arrow and make a fix and "replay" the command like I can in a Linux/Dos Shell? Are these '\' added automatically on long lines? Did you look at the MySQL client which I think has behavior that is close to this?	1	0	https://issues.jboss.org/browse/AS7-2640
1814	AS7-2624	Is "bin/standalone.sh -v" sufficient for this?	1	1	https://issues.jboss.org/browse/AS7-2624
1815	AS7-2536	Another vote. Saving to alternative locations and changing the length of history were two in my topic spec for CLI docs. If you have a separate config file for CLI, what else might it cover?	1	1	https://issues.jboss.org/browse/AS7-2536
1816	AS7-2531	Rado, can you suggest a better wording for this error message?	1	1	https://issues.jboss.org/browse/AS7-2531
1817	AS7-2530	When this test case is launched, has the AS instance already been started with Byteman? http://community.jboss.org/wiki/ABytemanTutorial	1	1	https://issues.jboss.org/browse/AS7-2530
1818	AS7-2480	brian, you commented the @Ignore some times ago (https://github.com/jbossas/jboss-as/commit/53c5582650180fa1c0d167c3b07afc0ba9a90118). Have you seen these failures again?	1	1	https://issues.jboss.org/browse/AS7-2480
1819	AS7-2469	Hi all, does this change mean the end for sub-profiling idea. What a pity! A profile extending by means of sub-profiles would be a very useful approach. Imagine hundred applications running in the same domain, which would be using the same default profile. But each application would like to have e.g. different logging configuration: levels, appenders etc. Without sub-profiling we need to create a separate profile for all applications requiring even small differences to the default profile and these profile will have up to 99% the same content as default. Assume we did so. On a profile update we have to change all hundreds profiles. That's why we are unhappy to resign sub-profiling idea. Regards	1	1	https://issues.jboss.org/browse/AS7-2469
1820	AS7-2437	I want to lock down requirements for this. Can we confirm that a maven command that removes the code is ok? You would call it like: {code}mvn clean -Pblank{code} and we could offer this as an option when generating: {code} mvn archetype:generate \ -DarchetypeArtifactId=jboss-javaee6-webapp \ -DarchetypeGroupId=org.jboss.spec.archetypes \ -DarchetypeVersion=7.0.2.Final -Dblank=true {code} (see the last line). Burr, please approve this, and I will add it.	0	0	https://issues.jboss.org/browse/AS7-2437
1821	AS7-2419	can you add the sample deployment?	1	1	https://issues.jboss.org/browse/AS7-2419
1822	AS7-2415	Jakiran any thoughts on this?	0	0	https://issues.jboss.org/browse/AS7-2415
1823	AS7-2392	What happens is that the JSR88 layer generates a tmp file and passes the URL to which to the AS7 management API. During the building of the ADD step an InputStream for that URL is opened and stored as content to the ADD step. When deployment is done, the JSR88 layer attempts to delete the tmp file, which fails on w2k8 presumably because the associated input stream is not closed. See the client [test log file|https://hudson.qa.jboss.com/hudson/view/JBossOSGi/job/jboss-as7-as2392/ws/testsuite/integration/smoke/target/workdir/target/test.log] and search for deployment-app.ear. {code} 03:19:18,396 INFO [org.jboss.as.ee.deployment.spi] (Thread-179) JBAS016101: End deploy: [target=http://org.jboss.as.ee.deployment/jsr88?targetType=as7,parent=null,type=ear,id=file:/C:/DOCUME~1/hudson/LOCALS~1/Temp/deployment-app.ear] 03:19:18,397 WARN [org.jboss.as.ee.deployment.spi] (Thread-179) JBAS016103: Cannot delete deployment file C:\DOCUME~1\hudson\LOCALS~1\Temp\deployment-app.ear, will be deleted on exit {code} This particular test fails because it tries to deploy the same EAR twice, but cannot use the same tmp file name. In branch [as2392|https://github.com/tdiesler/jboss-as/tree/as2392] I implemented a workaround that assigns a different tmp file name when the tmp file already exists and cannot be deleted. The underlying issue IMHO is however in the AS7 deployment layer and its handling of input streams associated with user provided URLs. This should be a general issue for all management clients running on w2k8	1	1	https://issues.jboss.org/browse/AS7-2392
1824	AS7-2380	How do we distribute jars that are not intended to be used by the server? Should they still go into the modules hierarchy?	1	1	https://issues.jboss.org/browse/AS7-2380
1825	AS7-2361	Also, how do you do external JNDI lookups? Or what is the recommended way?	1	1	https://issues.jboss.org/browse/AS7-2361
1826	AS7-2360	Shouldn't this happen by default?	1	1	https://issues.jboss.org/browse/AS7-2360
1827	AS7-2360	"The command line switches do not appear to be going into the server or boot logs at all. It does scroll by at startup, but is not durably stored for support to have a chance to trouble shoot." Isn't this clear....if not.. Then currently in EAP 5/AS 5 the command line switches are written to the log, this seems like a regression and hurts supports ability to quickly ascertain what is going on in an environment. Do you not agree this is a regression?	0	0	https://issues.jboss.org/browse/AS7-2360
1828	AS7-2360	I am not even sure what you are asking..by default out of the box..java_opts should to into the boot server log....is that clear enough?	0	0	https://issues.jboss.org/browse/AS7-2360
1829	AS7-2360	"By default out of the box" really needs to be defined. The option is already there as a TRACE level message. We can add the JAVA_OPTS to print, as David suggests, as a DEBUG message like the system properties do. We would likely not want it at any lower level than that. Is that what you are looking for?	0	0	https://issues.jboss.org/browse/AS7-2360

#	Requirement	Questions	Answered	Successfully Answered	URL
1830	AS7-2353	Is this available in the console? I did not see anything that mentioned https in the drop downs around connectors....so I am confused...	1	0	https://issues.jboss.org/browse/AS7-2353
1831	AS7-2350	But shouldn't there be ways to edit view the settings in it, or are the "hidden" in some of the other metaphors that are in the current console? I would think the values of the host.xml need to be exposed somehow, or am I just missing something. Please help me understand why host settings should not be exposed in the console... For example the hosts, IP address, management port, creating serves, ie related to the cloning comment earlier....I do not think we can discount host file "editing" in spirit, but I of course do agree we should be editing XML in the console, but values need to be exposed.	0	0	https://issues.jboss.org/browse/AS7-2350
1832	AS7-2345	Do you mean they should be within the table?	1	1	https://issues.jboss.org/browse/AS7-2345
1833	AS7-2344	What is the branch/tag for this issue?	0	0	https://issues.jboss.org/browse/AS7-2344
1834	AS7-2339	Do you mean two instances of the same server configuration? or simply duplicating a given server configuration under a different name?	1	1	https://issues.jboss.org/browse/AS7-2339
1835	AS7-2336	I personally like that we have these services marked as lazy. After all why enable those services if nothing uses it? What's the advantage of making these eager?	0	0	https://issues.jboss.org/browse/AS7-2336
1836	AS7-2336	I am a server admin, I do not have a war file, how do I test out...see a cluster....in the past I just did about 10 mins of stuff and then I had a cluster. With AS 7.X I would have to find some docs, find a setting that is not documented, is not in the console, and then I am kinda lost. From a end user perspective out of the box, changing behavior that is several years old is challenge for people being great quickly.	1	1	https://issues.jboss.org/browse/AS7-2336
1837	AS7-2334	What does this relate to? Can you elaborate?	1	1	https://issues.jboss.org/browse/AS7-2334
1838	AS7-2327	This really should be a separate config though. The HTTP API != management console. It should be possible to have the former without the latter. For example if you want to ensure that your admins use JON.	1	1	https://issues.jboss.org/browse/AS7-2327
1839	AS7-2324	Can you provide a list of contents that should into the report?	0	0	https://issues.jboss.org/browse/AS7-2324
1840	AS7-2320	What do you mean interim values?	0	0	https://issues.jboss.org/browse/AS7-2320
1841	AS7-2318	That sounds good to me. Does that would mean "Log Level" combo box would have an empty or none entry?	1	1	https://issues.jboss.org/browse/AS7-2318
1842	AS7-2318	{quote} Does that would mean "Log Level" combo box would have an empty or none entry? {quote} Doug, I haven't checked how the admin console displays that information. But it would be the same as what it would display for any other handler, when the threshold (which is an optional field) is not set.	1	1	https://issues.jboss.org/browse/AS7-2318
1843	AS7-2317	Curent problem or challenge is the embedded logger editor only allows for the creation/managing of one logging results file. How do I create multiple log files that are specific to needs. One request from the field is how do I log security related events in one specific file. This is a big deal for security conscious use cases ie like in government agencies. Hope that makes more sense.	1	0	https://issues.jboss.org/browse/AS7-2317
1844	AS7-2280	Is this still relevant? Cache configuration can be modified by the management API, is there anything else that is required to resolve this issue?	1	1	https://issues.jboss.org/browse/AS7-2280
1845	AS7-2272	Isn't this better explained in a README? Where do you draw the line? Should we have 100 XML files that existed in AS6 with a message that says "use this standalone.xml instead"? Should we have twiddle.sh launch jboss-admin.sh? So this will either be inconsistent, or we will have a ridiculous amount of clutter	1	1	https://issues.jboss.org/browse/AS7-2272
1846	AS7-2272	What came out of a quick but pretty thorough discussion of this at the usability workshop is: I'm an AS7 newbie, perhaps a JBoss AS newbie, but not at all a "try out new stuff" newbie. I download, unzip and do one of two things: 1) Read README.txt in the root 2) or, since it's a server I downloaded, it should start and do something. It's a web server so it probably has an http accessible welcome page with helpful information on it. So I look in bin and try to run the start script. A start script usually has a name that indicates it starts or runs something, so that's what I look for. The goal is to immediately get such people on the right path, without forcing them to right off the bat go download docs to figure out what to do. (IMO any output from the script should reference the README.txt, not "See the documentation...") The README.txt should include a brief explanation of running a standalone server versus running a managed domain. Any run.sh/bat should not actually attempt to launch a server, because the configuration and behavior of the server will be different from previous AS releases.	1	1	https://issues.jboss.org/browse/AS7-2272
1847	AS7-2239	{noformat} java.lang.IllegalArgumentException: Can not set java.net.URL field com.sun.ts.tests.ejb30.bb.session.stateless.annotation.resource.ResourceFieldBean.url to java.lang.String {noformat}	1	1	https://issues.jboss.org/browse/AS7-2239
1848	AS7-2239	Will this be fixed today? If not, please push to CR1.	1	1	https://issues.jboss.org/browse/AS7-2239
1849	AS7-2222	What is this actually used for? The issue is, we don't have References for these objects. We may need to create references for this, but I am curious how the reference is used by the Referencable objects.	1	1	https://issues.jboss.org/browse/AS7-2222
1850	AS7-2179	Maybe hide inherited operations by default? Actually I find it counter-intuitive that I have to specify operations=true to find out about the local operations. Maybe read-resource-description should include local operations by default and then you also have :read-resource-description(inherited-operations=true) or something like that...	0	0	https://issues.jboss.org/browse/AS7-2179
1851	AS7-2172	Heiko, can this be scheduled for 7.1.1? Thanks.	1	1	https://issues.jboss.org/browse/AS7-2172
1852	AS7-2169	{quote} 1. Call System.getenv("X") and if it exists, use it as expression value 2. else, call System.getProperty("env.X") and if it exists, use it as expression value {quote} Isn't it usually the other way around? Developers usually use the system properties to override the environment entries (for example by passing -Denv.X). Would it be more consistent if we did: 1. Call System.getProperty("env.X") and if it exists, use it as expression value 2. else, call System.getenv("X") and if it exists, use it as expression value	0	0	https://issues.jboss.org/browse/AS7-2169
1853	AS7-2164	Stuart, I'm creating CdiAnnotationProcessor, CdiAnnotations and CdiDeploymentMarker classes that will act as Jaxrs equivalent classes: JaxrsAnnotationProcessor, JaxrsAnnotations and JaxrsDeploymentMarker respectively. The differnce will be of course in the CdiAnnotations that will contain the CDI annotations of JSR 299. Is that right?	1	1	https://issues.jboss.org/browse/AS7-2164
1854	AS7-2127	RESTEasy 2.2.3.GA is now released. Does stuart want to bump the version?	0	0	https://issues.jboss.org/browse/AS7-2127
1855	AS7-2115	This issue is being relocated to 7.1.2 so that 7.1.1 only contains criticals, blockers, and EAP LA issues (which are probably not yet at the right status, or need an extra triage pass). If these issues are completed in a 7.1.1 timeframe please change the fix version. Although if you forget I will bulk correct.	1	1	https://issues.jboss.org/browse/AS7-2115
1856	AS7-2113	This issue is being relocated to 7.1.2 so that 7.1.1 only contains criticals, blockers, and EAP LA issues (which are probably not yet at the right status, or need an extra triage pass). If these issues are completed in a 7.1.1 timeframe please change the fix version. Although if you forget I will bulk correct.	1	1	https://issues.jboss.org/browse/AS7-2113
1857	AS7-2101	Can this be resolved?	1	1	https://issues.jboss.org/browse/AS7-2101

#	Requirement	Questions	Answered	Successfully Answered	URL
1858	AS7-2096	I have a hard time getting explanations and details ... So it was the opposite setting then ;) I am not sure about the usefulness about 100 as the default. It has been used because it is the httpd default, but I don't really see how it could help prevent bad things from happening. I'll review that, because I find it odd. On the other hand, the default is relatively big, which implies a low impact. Generally, if this is about adding a setting which would only be useful for configs which are bad in production, then why should it be put in ? You have a means to configure it as is.	1	1	https://issues.jboss.org/browse/AS7-2096
1859	AS7-2074	[--controller=host:port] [--connect] [--file=file_path] [--commands=command_or_operation1,command_or_operation2...] [--command=command_or_operation] [--user=username --password=password] --help (-h) - prints (this) basic description of the command line utility. --version - prints the version info of the JBoss AS release, JVM and system environment. --controller - the default controller host and port to connect to when --connect option. (described below) is specified or when the connect command is issued w/o the arguments. The default controller host is localhost and the port is 9999. --connect (-c) - instructs the CLI to connect to the controller on start-up (to avoid issuing a separate connect command later). --file - specifies a path to a file which contains commands and operations (one per line) that should be executed (in a non-interactive mode). The CLI will terminate the session immediately after the last command has been executed or if some command or operation failed. --command - specifies a single command or an operation that should be executed in the CLI session. The CLI will terminate the session immediately after the command or the operation has been executed. Note: --command argument is optional in a sense that any word (or phrase without whitespaces in it) will be assumed to be a command (or an operation). --commands - specifies a comma-separated list (the list must not contain whitespaces) of commands and operations that should be executed in the CLI session. The CLI session will be automatically terminated as soon as the last command or operation has been executed or after the first error. Note: --commands argument is optional in a sense that any comma-separated list at the end of the argument list will be assumed to be the list of commands and operations. --user - if the controller requires user authentication, this argument can be used to specify the user name as a command line argument. If the argument isn't specified and the authentication is required the user will be prompted to enter the user name when the connect command is issued. --password - specifies the password for authentication while connecting to the controller as a command line argument. If the argument isn't specified and the authentication is required the user will be prompted to enter the password when the connect command is issued. During the command line session for a list of commands available in the current context execute help --commands The resulting listing may depend on the current context. E.g. some of the commands require an established connection to the controller (standalone or domain). These commands won't appear in the listing unless the connection has been established. Other commands may depend on the availability of specific subsystems. E.g. if the messaging subsystem is not available, messaging-related commands will not be listed. Here are some of the most basic supported commands: cn (or cd) - change the current node path to the argument; connect - connect to the specified host and port; deploy - deploy an application; help (or h) - print this message; history - print or disable/enable/clear the history expansion. ls - list the contents of the node path; pwn (or pwd) - prints the current working node; quit (or q) - quit the command line interface; undeploy - undeploy an application; version - prints the version and environment information. For a more detailed description of a specific command, execute the command with '--help' as the argument. Tab-completion is supported for the commands, just press the tab key to start. Operation requests start with './' or '/' and follow the format: [node-type=node-name (, node-type=node-name)*] : operation-name ['('[name=value [, name=value]*]')'] e.g. /subsystem=web/connector=http:read-attribute(name=protocol) If the operation request doesn't require a node path then the request can start with ':' followed by an operation name and the property list if	0	0	https://issues.jboss.org/browse/AS7-2074
1860	AS7-2053	Could we at least update to 1.7.3? Version which we distribute is not compatible with the one Hibernate Search 4 uses, leading to problems with applications wishing to use JAX-RS and HSEARCH together.	1	1	https://issues.jboss.org/browse/AS7-2053
1861	AS7-2052	What do you think about this John?	1	1	https://issues.jboss.org/browse/AS7-2052

#	Requirement	Questions	Answered	Successfully Answered	URL
1862	AS7-1994	discussion, and our needs. I'm sending this email mainly as reminder for me to discuss these points with Brian after his and mine vacations regard S. * pferraro_ (~pferraro@adsl-99-97-30-7.dsl.wlfrct.sbcglobal.net) has joined #jboss-as7 <emuckenhuber_> maeste: pong * emuckenhuber_ is now known as emuckenhuber <maeste> emuckenhuber_: Hi, I'm trying to implement our metrics. Just to recall we need this issue Brian has resolved to implement them: AS7-772 <jbossbot> jira [AS7-772] Allow operation handlers to register child ModelNodeRegistrations [Resolved (Done) Feature Request, Critical, Brian Stansberry] https://issues.jboss.org/browse/AS7-772 <maeste> emuckenhuber: what I'm trying to do is: * bbrowning (~bbrowning@redhat/jboss/bbrowning) has joined #jboss-as7 * ChanServ gives voice to bbrowning <maeste> emuckenhuber: https://github.com/maeste/jboss-as/commit/d08f758b39d6a83257745cf86fe9f7ced481c4cb#L1R78 <jbossbot> git [jboss-as] d08f758.. Stefano Maestri stats prototype <maeste> emuckenhuber: but I'm getting this exception <maeste> emuckenhuber: https://gist.github.com/1111413 <maeste> emuckenhuber: have you an idea what I'm doing wrong? <emuckenhuber> maeste: yes <maeste> emuckenhuber: sorry to ask you for a Brian's issue, but AFAIK Brian is away until 11 August <maeste> emuckenhuber: cool <emuckenhuber> maeste: basically you would need to register the handlers inside that operation <maeste> emuckenhuber: yup, moreover after a service activated by this operation has been started <maeste> emuckenhuber: it's the reason because I've added a listener <maeste> emuckenhuber: line 91 <emuckenhuber> maeste: no, i'm saying operations have to be registered as part of operation handler itself <emuckenhuber> which is basically the exception <maeste> emuckenhuber: you have lost me <emuckenhuber> actually all operation handlers should be registered as part of the Extension <maeste> emuckenhuber: oki, but it's what AS7-772 try to solve <jbossbot> jira [AS7-772] Allow operation handlers to register child ModelNodeRegistrations [Resolved (Done) Feature Request, Critical, Brian Stansberry] https://issues.jboss.org/browse/AS7-772 <maeste> emuckenhuber: at least AFAI understood <maeste> emuckenhuber: this is our original problem <emuckenhuber> hmm, imho that does not make any sense <maeste> emuckenhuber: we have to register operation (metics in this case) after a datasource have been created <emuckenhuber> well perhaps a bit <emuckenhuber> anyway - you would have to register that as part of populateModel() * asoldano is now known as asoldano_lunch <maeste> emuckenhuber: because metrics are dynamic for us and need a concrete instance of ResourceAdapter/DS <emuckenhuber> yeah, but it basically does not really comply with how the management model works <maeste> emuckenhuber: there was a discussion on ML in May about that, and then Brian created 772 <emuckenhuber> you remember the subject? <maeste> emuckenhuber: discussion subject was "access to the ModelNodeRegistration during handling of an operation and during deployment" * kevinpollet (~kevinpoll@217.112.54.72) has joined #jboss-as7 * abstractj (~abstractj@186.202.49.66) has joined #jboss-as7 * hbraun (~hbraun@ppp-93-104-112-69.dynamic.mnet-online.de) has joined #jboss-as7 * hbraun has quit (Client Quit) * miclorb_ has quit (Remote host closed the connection) <emuckenhuber> maeste: hmm, ok that's an issue then <maeste> emuckenhuber: so, basically you are saying it's not implented <maeste> emuckenhuber: oki, I'll put down in my todo list <maeste> emuckenhuber: better to wait for Brian, since I've discussed about that a lot with him <emuckenhuber> maeste: well i would not say 'not implemented' - we actually have a check to prevent exactly this <maeste> emuckenhuber: oki much stronger <maeste> emuckenhuber: <emuckenhuber> maeste: well perhaps you want to do a single operation handler which gives you all the metrics for now? <emuckenhuber> not sure how critical that is for the console work <maeste> emuckenhuber: not possible, it was already done, but it doesn't work <emuckenhuber> why? <maeste> emuckenhuber: or better it works, but not as expected <maeste> emuckenhuber: because sometimes some metrics are not available <maeste> emuckenhuber: to be	1	1	https://issues.jboss.org/browse/AS7-1994
1863	AS7-1994	For the Datasource case, does StatisticsPlugin have a field that forms a natural key? Looking at https://github.com/maeste/jboss-as/commit/d08f758b39d6a83257745cf86fe9f7ced481c4cb#L1R78, the concern I have is it appears you are modifying a generic resource registration by adding instance-specific metrics. So, the management interface for /subsystem=datasources/datasource=* has one set of attributes and ops, /subsystem=datasources/datasource=myDS has another set (a superset if we could make that work) and /subsystem=datasources/datasource=anotherDS has a 3rd set. It's almost like an abstract superclass and concrete subclasses. That might be possible, but I don't much like the idea that a resource's definition changes. Something like /subsystem=datasources/datasource=myDS/statistics-provider=oracle and /subsystem=datasources/datasource=myDS/statistics-provider=postgres sounds nicer. The statistics-provider level works like "subsystem" works.	1	1	https://issues.jboss.org/browse/AS7-1994
1864	AS7-1963	Can anyone comment on how jboss (or the webapp in question) knows to use the valve defined in jboss-web.xml? how does it know it is any different from any non-security related valve?	0	0	https://issues.jboss.org/browse/AS7-1963

#	Requirement	Questions	Answered	Successfully Answered	URL
1865	AS7-1932	IST) Jaikiran: ? (09:47:50 IST) jbossbot: jira [AS7-1934] EJB client context configuration and management ops [Open (Unresolved) Sub-task, Blocker, jaikiran pai] https://issues.jboss.org/browse/AS7-1934 (09:48:32 IST) kpiwko_afk is now known as kpiwko (09:49:13 IST) dmlloyd: sure (09:49:47 IST) Jaikiran: a part of it ties into AS7-1932 (09:49:49 IST) jbossbot: jira [AS7-1932] Remoting outbound connection service [Open (Unresolved) Sub-task, Blocker, Unassigned] https://issues.jboss.org/browse/AS7-1932 (09:50:03 IST) dmlloyd: yeah there's two pieces (09:50:19 IST) dmlloyd: basically we first need a way to configure outbound connections in the remoting subsystem configuration (09:50:53 IST) dmlloyd: they need to be named, and have some base set of configuration parameters to facilitate authentication (think data sources) as well as configuring the socket itself (09:51:14 IST) dmlloyd: the connections would be established in the context of a source endpoint so they would reside within the endpoint configuration (09:51:41 IST) dmlloyd: then our EJB client configuration in the EJB subsystem would declare client contexts and have a listing of connections to associate with them (09:52:36 IST) dmlloyd: an important note though - the connection service lifecycle cannot be tied to the connection lifecycle (09:52:38 IST) Jaikiran: so ultimately the deployment descriptor would just be a named ref to that connections (09:52:45 IST) dmlloyd: right (09:53:02 IST) Jaikiran: isn't it the other way round? (09:53:16 IST) dmlloyd: it's both ways around (09:53:23 IST) Jaikiran: i mean the connection is destroyed on service destroy? (09:53:25 IST) dmlloyd: a service lifecycle cannot be controlled by external factors (09:53:41 IST) dmlloyd: so it can't wait for the connection to be up before starting for example (09:53:47 IST) dmlloyd: and killing the connection cannot bring down the service (09:54:07 IST) Jaikiran: oh ok i see what you mean (09:54:25 IST) Jaikiran: so typically, the named ref dependency doesn't guarantee a connection that is UP (09:54:39 IST) dmlloyd: right (09:54:55 IST) dmlloyd: we need a way to get a connection from the ref if it's up - probably best to return an IoFuture<Connection> always (09:55:05 IST) dmlloyd: the connection is either up or in the process of being created/restored (09:55:21 IST) dmlloyd: and anything which uses it has to be able to go back for a new connection in the event of a failure (09:55:38 IST) dmlloyd: this will take some thinking when it comes to EJBReceiver since we don't want the EJBReceiver disappearing (09:55:47 IST) dmlloyd: also, we should have a way to designate a client context as the "default" which is applied to deployments if none is specified for that deployment (09:55:55 IST) dmlloyd: (once we're at that part) (09:56:08 IST) lfryc [~lfryc@cst-prg-121-2.cust.vodafone.cz] entered the room. (09:56:35 IST) Jaikiran: although we don't "configure" the default context we do have the necessary setup ready to allow that default to be made available to the deployments currently (09:56:47 IST) stliu left the room (quit: Quit: Leaving). (09:57:11 IST) Jaikiran: btw, i was looking at the remoting subsystem xsd and it appears that the source endpoint is just 1? (09:57:24 IST) Jaikiran: i mean all connections would thus be in the scope of the endpoint? (09:58:00 IST) dmlloyd: yeah I'm not sure how it's set up, but one endpoint per VM is a reasonable policy (09:58:23 IST) Jaikiran: ok (09:58:41 IST) Jaikiran: in terms of specifying the destination for those connections (09:58:56 IST) Jaikiran: i believe we can rely on the outbound-socket-binding support that we added recently (09:59:04 IST) Jaikiran: which allows for specifying a host and a port (09:59:16 IST) Jaikiran: is there more that needs to be considered? (09:59:17 IST) dmlloyd: yes and no - remoting accepts a URI for the destination (09:59:34 IST) dmlloyd: which includes protocol as well, and basically supports anything a URI supports (09:59:57 IST) dmlloyd: for example, "local:" is a valid destination (10:00:05 IST) Jaikiran: so the outbound-client-socket would just be one type of the destination then (10:00:12 IST) dmlloyd: at some point "http:" may be as well (10:00:46 IST) dmlloyd: yeah we'd need some way of referencing it (10:01:25 IST) dmlloyd: e.g. "theproto://#{host}:#{port}/blah?foo" (10:01:31 IST) dmlloyd: which is really ugly :) (10:01:36 IST) dmlloyd: but I can't think of anything better (10:01:43 IST) dmlloyd:	1	1	https://issues.jboss.org/browse/AS7-1932
1866	AS7-1928	Hi David, I think you might misunderstand my intention. What I'm trying to do is to *access the resources in META-INF/ in a JAR (defined as a module) from a class in that same JAR*. I tried tweaking the module.xml and jboss-deployment-structure.xml in various ways but no luck. Did I miss anything?	1	1	https://issues.jboss.org/browse/AS7-1928
1867	AS7-1928	Are you using the current classes class loader to try and access META-INF or are you trying to use the TCCL?	1	0	https://issues.jboss.org/browse/AS7-1928
1868	AS7-1920	So how did this end up working ? I would expect myfile.xml myfile.xml.dodeploy to work or how ?	1	1	https://issues.jboss.org/browse/AS7-1920
1869	AS7-1920	okey - is there a jira for Part 2 ? btw. does this mean .xml is special and .dodeploy will just be ignored ? i.e. we currently apply this to all files we deploy (.jar, .war, .ear, .xml etc.) and just wondering if that will be hurtful for this kind of deploys ? (I assumed these would be symmetric, but sounds like they aren't)	1	1	https://issues.jboss.org/browse/AS7-1920
1870	AS7-1917	Wouldn't it be much better solution to make it possible to set URIEncoding and useBodyEncodingForURI on the connector element instead of using system properties? Connector element: <Connector name="http" ... URIEncoding="UTF-8" useBodyEncodingForURI="true"/> Like it's done for JBossWeb 3.0.x http://docs.jboss.org/jbossweb/3.0.x/config/http.html http://docs.jboss.org/jbossweb/7.0.x/config/http.html System properties: org.apache.catalina.connector.URI_ENCODING org.apache.catalina.connector.USE_BODY_ENCODING_FOR_QUERY_STRING	1	0	https://issues.jboss.org/browse/AS7-1917
1871	AS7-1917	Okay. So is it official to use these properties I've found in the code? Are they documented anywhere?	1	0	https://issues.jboss.org/browse/AS7-1917
1872	AS7-1908	Planning to make a management function available: {noformat} /subsystem=jgroups/stack=udp:export-native-configuration {noformat} which returns a String containing the XML JGroups configuration. Would that satisfy this request?	1	1	https://issues.jboss.org/browse/AS7-1908
1873	AS7-1893	This issue is being relocated to 7.1.2 so that 7.1.1 only contains criticals, blockers, and EAP LA issues (which are probably not yet at the right status, or need an extra triage pass). If these issues are completed in a 7.1.1 timeframe please change the fix version. Although if you forget I will bulk correct.	0	0	https://issues.jboss.org/browse/AS7-1893
1874	AS7-1893	Richard, there are integration clustering tests which use jdbc cache store and their config (https://github.com/jbossas/jboss-as/blob/master/testsuite/integration/clust/src/test/config/ispn/cache-container-with-jdbc-store.xml) is currently out of date. Can you please fix this?	1	1	https://issues.jboss.org/browse/AS7-1893
1875	AS7-1893	One more thing.. I'm looking at https://github.com/rachmatowicz/jboss-as/commit/9fed5aad5faacde351966d15f117b04ebf16b896#L0R589 and https://github.com/rachmatowicz/jboss-as/commit/9fed5aad5faacde351966d15f117b04ebf16b896#L1R539 and I see a problem. When we have XSD types for all three jdbc cache stores (string-keyed-jdbc-store, binary-keyed-jdbc-store and mixed-keyed-jdbc-store), I would suppose that when I use mixed-keyed-jdbc-store, I can really only create mixed jdbc cache store from it. However, from the code logic it seems that when a user specifies mixed-keyed-jdbc-store and inside that only binary-keyed-table, they will actually get binary jdbc cache store (the same for string-based one when they specify only string-keyed-table). I would suppose here that the XSD will say that binary-keyed-table has minOccurs=1 and maxOccurs=1 as well as string-keyed-table (they have "0" now): https://github.com/rachmatowicz/jboss-as/blob/9fed5aad5faacde351966d15f117b04ebf16b896/build/src/main/resources/docs/schema/jboss-as-infinispan_1_2.xsd#L589 WDYT?	1	1	https://issues.jboss.org/browse/AS7-1893
1876	AS7-1833	Hi All Does this bug really fix in EAP 6.1?I also had the bug on JBoss EAP6.1 final version. My environment is :Linux 5.3 + JBoss EAP 6.1.0 + JDK 1.6.0_26 Here are the log : [standalone@localhost:9999 /] :reload Communication error: java.util.concurrent.ExecutionException: Operation failed: Operation failed: Channel closed [disconnected /] Please correct me if I have any misunderstanding. Regards Xin	1	1	https://issues.jboss.org/browse/AS7-1833
1877	AS7-1830	Emanuel, this ties in with the "flattening" thread we had with Heiko Rupp. Where do things stand with that? Basically a VALUE_TYPE for an interface criteria attribute is impossible if we allow arbitrary nesting.	0	0	https://issues.jboss.org/browse/AS7-1830

#	Requirement	Questions	Answered	Successfully Answered	URL
1878	AS7-1809	Just curious - Can't this library jar (of any version of your choice) be just packaged within the application and start using it? Or is JBoss AS adding this as an implicit dependency to your deployments (which would require some additional jboss-deployment-structure.xml configurations, to use a different version of that module)?	0	0	https://issues.jboss.org/browse/AS7-1809
1879	AS7-1807	The parts that we customized for EAP 6 so far include: (1) The log messages (JBPAPP-6831) (2) The welcome pages (JBPAPP-6838) (3) The admin console - logos, colors, etc (JBPAPP-7111, JBPAPP-7031) (4) The config profiles (JBPAPP-7032) A single jar containing this content/configuration sounds like a good idea. Could this be configured by changing standalone.conf and domain.conf?	0	0	https://issues.jboss.org/browse/AS7-1807
1880	AS7-1807	David - does your suggestion avoid the .bat/.sh scripts being the place this is defined ? See above why we should not make .bat/.sh's the place to control product disto settings. basically any settings controlled by .bat/.sh is ignored by any launches done via IDE's, arquillian and other "jboss launchers" because of the non-portability of control over processes started via .bat/.sh.	1	1	https://issues.jboss.org/browse/AS7-1807
1881	AS7-1796	Is this happening by Final? If not please punt to 7.1.1	1	1	https://issues.jboss.org/browse/AS7-1796
1882	AS7-1786	Hi Kabir, What is the current status of this on master? Thanks, David	1	1	https://issues.jboss.org/browse/AS7-1786
1883	AS7-1784	Any update? Should this be reassigned?	0	0	https://issues.jboss.org/browse/AS7-1784
1884	AS7-1782	Isn't more clear with subresources? See AS7-923, and let me know what other informations you need	0	0	https://issues.jboss.org/browse/AS7-1782
1885	AS7-1781	Marcus, can you provide the relevant information.	0	0	https://issues.jboss.org/browse/AS7-1781
1886	AS7-1781	Thanks Marcus, that's already a good starting point. Can you provide some examples for " /subsystem=security/security-domain=other:add()"?	1	1	https://issues.jboss.org/browse/AS7-1781
1887	AS7-1781	Can someone update the resource description? it's still empty and practically not usable for us to move forward. although some examples for the child resources would be good.	0	0	https://issues.jboss.org/browse/AS7-1781
1888	AS7-1780	Hi Paul, What is the status on this one?	1	1	https://issues.jboss.org/browse/AS7-1780
1889	AS7-1780	Richard, will there be a possibility to change the infinispan subsystem for the whole domain or server group so that all servers in the domain/server group are affected? I'm just trying to find out what we'll be able to test around EDG running in the domain mode. Thanks	1	1	https://issues.jboss.org/browse/AS7-1780
1890	AS7-1776	I think this is already done for the English language - the only thing that needs to be done is the German translation (the message bundle contains 'TRANSLATE ME' messages). Or did I miss anything?	1	1	https://issues.jboss.org/browse/AS7-1776
1891	AS7-1772	Where does this stand? Perhaps it's done? Andy did some commits a week or so back that may have wrapped it up. Either way, please let QE (Rajesh) know where it stands, if you haven't already.	0	0	https://issues.jboss.org/browse/AS7-1772
1892	AS7-1772	Is the configuration that is needed for this in the XML schema? I opened another JIRA against the XML schema because I couldn't find it there? Isn't that missing too, which would be critical, since they wouldn't be able to configure it.	0	0	https://issues.jboss.org/browse/AS7-1772
1893	AS7-1772	Hi Andy, seems like this functionality is done and AS7 is able to handle HQ in collocated topology. Is it truth? Thank you, Mirek	1	1	https://issues.jboss.org/browse/AS7-1772
1894	AS7-1772	Hi, I've managed to start AS7/HQ in collocated topology. Great work. :-) There is only question whether jms client is able to failover to backup when remote jndi is not ready yet. Do we have any update? Thank you, Mirek	0	0	https://issues.jboss.org/browse/AS7-1772
1895	AS7-1772	Hi Andy, I've noticed that '<ha>true</ha>' cannot be parsed in configuration. Could you add it please? Thank you, Mirek	0	0	https://issues.jboss.org/browse/AS7-1772
1896	AS7-1772	Can this be closed? I believe the HA issue is fixed.	0	0	https://issues.jboss.org/browse/AS7-1772
1897	AS7-1752	The second issue... {quote} 2) The launch script directly references the standalone/configuration/logging.properties and standalone/log/boot.log files. So, if those files aren't present the server doesn't boot properly, and both, e.g. serverA and serverB will use the standalone/log/boot.log as their boot log file (where logging goes until the boot process starts the logging subsystem.) {quote} ...is still present in 7.1.2.Final when using Mac OS X. Should i create a new Jira Issue on this or can someone reopen this one?	0	0	https://issues.jboss.org/browse/AS7-1752
1898	AS7-1751	The idea would be the following to have the standalone.xml having ${VAR} for every variable that makes sense to be defined at run-time. standalone.conf - is the configuration file not only for standalone.sh BUT can also for the standalone.xml via environment variables. Then one would have to take care to configure one place and it get's pushed to the xml file. AS7 already overwrites the standalone.xml file every time it starts at min. It also provides a history. So it can be a 'nice' fit - in my opinion. Think a bash shell managing an xml file	0	0	https://issues.jboss.org/browse/AS7-1751
1899	AS7-1735	I tried to set the cause of the StartException but it is not taken into account anyway in its toString() method. I checked the code and if I understand it correctly, it is up to the messaging system to provide a description with the error inside the StartException. Wouldn't it be simpler to let StartException print out the inner cause l10ned message (if there is one) rather than do this many times in the messaging subsystem (and likely other systems)?	1	1	https://issues.jboss.org/browse/AS7-1735
1900	AS7-1730	Pavel, can you check AS7-2616 and see if it satisfies your requirements?	0	0	https://issues.jboss.org/browse/AS7-1730
1901	AS7-1730	I'm closing it, it is verified and resolved. Currently available impl. satisfies my requirements.	1	1	https://issues.jboss.org/browse/AS7-1730
1902	AS7-1724	Is this a good candidate for "Assign to Community?"	1	1	https://issues.jboss.org/browse/AS7-1724
1903	AS7-1709	Does this issue include remote JNDI look up for JMS objects (queues, factories ...)?	1	0	https://issues.jboss.org/browse/AS7-1709
1904	AS7-1701	I don't understand the description. What would you like to see as the result of jboss-admin.sh --help? Would what is printed as the result of "jboss-admin.sh help" do?	0	0	https://issues.jboss.org/browse/AS7-1701
1905	AS7-1674	According to the JAX-RS spec there are three ways to activate a restful endpoint. These three methods are specified in the JAX-RS 1.1 specification in section 2.3.2. 1) Not including an application class and using web.xml as described above 2) subclassing javax.ws.rs.core.Application and using web.xml 3) subclassing javax.ws.rs.core.Application and using the ApplicationPath annotation. See https://docs.jboss.org/author/display/AS7/JAX-RS+Reference+Guide The idea that resteasy should be activated if there happens to be a JAX-RS annotation in the deployment does not really work. For one thing, where would you map it to? For another, this means that if a 3rd party app library includes a class with a JAX-RS annotation then resteasy would be activated wether you wanted it or not.	0	0	https://issues.jboss.org/browse/AS7-1674
1906	AS7-1666	What are you interested in seeing? That was a feature request to list it like that. BTW, before opening new issue, consider discussing them first.	1	1	https://issues.jboss.org/browse/AS7-1666
1907	AS7-1652	Brian, can you comment on this?	0	0	https://issues.jboss.org/browse/AS7-1652
1908	AS7-1652	Are you looking for a date? Let me know if it's blocking you.	1	1	https://issues.jboss.org/browse/AS7-1652
1909	AS7-1640	Clebert can you comment on this one? Brian spent a lot of effort on getting everything mapped, so most of the work is done, but we still need a hornetq expert to scan for major problems.	0	0	https://issues.jboss.org/browse/AS7-1640

#	Requirement	Questions	Answered	Successfully Answered	URL
1910	AS7-1633	Was this already completed?	1	1	https://issues.jboss.org/browse/AS7-1633
1911	AS7-1614	I believe with the commits that we made as part of AS7-2134 this issue should now be covered. Brian, do we want to keep this JIRA open to track the system property substitution (if at all we plan to do it)?	1	1	https://issues.jboss.org/browse/AS7-1614
1912	AS7-1590	Needs work on "What are RESTful Web Services"	0	0	https://issues.jboss.org/browse/AS7-1590
1913	AS7-1554	You can check the jiras documented here https://github.com/tdiesler/org.apache.felix.webconsole/commits/3.1.6.SP1/ Why is this still relevant? I thought we were removing that console.	1	0	https://issues.jboss.org/browse/AS7-1554
1914	AS7-1545	Does this issue still occur? I tried "./build.sh -X clean install" and "mvn -X clean install" from the root directory, and both comleted without error.	1	1	https://issues.jboss.org/browse/AS7-1545
1915	AS7-1531	Hi Anil, What's left to do on this (uf anything at all)?	1	1	https://issues.jboss.org/browse/AS7-1531
1916	AS7-1500	Hey all, There are several other paths that's needed for my application. <path name="sun/security"/> <path name="sun/security/krb5"/> <path name="com/sun/jndi/url"/> <path name="com/sun/jndi/url/dns"/> Can these be added as well? Thank you, Oved	1	1	https://issues.jboss.org/browse/AS7-1500
1917	AS7-1489	Why is this needed?	1	1	https://issues.jboss.org/browse/AS7-1489
1918	AS7-1484	not working in 7.1.0.Beta1 anymore :(. Can you reopen this issue ?	0	0	https://issues.jboss.org/browse/AS7-1484
1919	AS7-1459	The changes to standalone.sh is not working for me. + (-d|--debug) ... + (--) ... + (*) What's this opening bracket in lines 13, 20 and 23? Never saw this syntax in shell-scripts. Bourne- and Korn-shell on Solaris do not like them.	1	1	https://issues.jboss.org/browse/AS7-1459
1920	AS7-1459	Hi Frank, Is it working whithout the opening brackets ? Guillaume.	1	1	https://issues.jboss.org/browse/AS7-1459
1921	AS7-1426	What do I need to do to get {code} @Stateless @LocalBean public class SimpleStatelessSessionBean implements Echo { @Resource private BundleContext context; {code} to work (i.e. lookup by type).	0	0	https://issues.jboss.org/browse/AS7-1426
1922	AS7-1411	Could you give a little more information than this? What are the errors? Do you have a branch I can look at?	1	1	https://issues.jboss.org/browse/AS7-1411
1923	AS7-1407	There's another issue with our Naming store :(I have the following piece of code in our web.xml DD: --- <env-entry> <env-entry-name>jaxws/entry1</env-entry-name> <env-entry-type>java.lang.String</env-entry-type> <env-entry-value>web</env-entry-value> <injection-target> <injection-target-class>org.jboss.test.ws.jaxws.enventry.EnvEntryBeanJSE</injection-target-class> <injection-target-name>strValue</injection-target-name> </injection-target> </env-entry> <env-entry> <env-entry-name>jaxws/entry2</env-entry-name> <env-entry-type>java.lang.Integer</env-entry-type> <env-entry-value>8</env-entry-value> </env-entry> --- This results in the following output in the CLI: --- ... "jaxws-enventry-jse.war" => { "modules" => {"jaxws-enventry-jse" => {"java:module" => { ... "env" => { "class-name" => "javax.naming.Context", "children" => { "entry1" => { "class-name" => "org.jboss.as.naming.context.ModularReference", "value" => "web" }, "entry2" => { "class-name" => "org.jboss.as.naming.context.ModularReference", "value" => "8" } } }, }}} } ... --- Note the missing 'jaxws' prefix? And at the end when I was debugging our JNDI naming store, I noticed that missing 'jaxws' context is present in the debugger. It's just not displayed. Another thing that I noticed is when I'm doing the following lookup: --- new InitialContext().lookup("java:comp/env/"); --- then for above web.xml DD configuration "env/jaxws" is returned instead of expected "env/" context :(	1	1	https://issues.jboss.org/browse/AS7-1407
1924	AS7-1393	Can you be more specific as to what problem this is causing? Currently, if the user sets property "hibernate.cache.infinispan.cachemanager", we add a dependency on the specified value (in which case the cache manager will be started). If an application isn't using a 2lc, they wouldn't set this property (which means the cache isn't started). I don't think that we want the Infinispan cache to be started by default.	1	1	https://issues.jboss.org/browse/AS7-1393
1925	AS7-1381	For the records: "what about resource-adapter/connection-definition/config-property or resource-adapter/admin-object/config-property?" How can these be made accessible?	1	1	https://issues.jboss.org/browse/AS7-1381
1926	AS7-1379	What "attributes" ? If you mean the resource adapter config-property's, then they are only writeable if the resource adapter is version 1.6, and supports dynamic updates. That should already be supported.	1	1	https://issues.jboss.org/browse/AS7-1379
1927	AS7-1379	Hi Jesper, Currently transaction-support is readonly. [standalone@localhost:9999 resource-adapter=xmpp-ra.rar] :write-attribute(name=transaction-support,value=Notransaction) { "outcome" => "failed", "failure-description" => "Attribute transaction-support is not writeable", "rolled-back" => true } This should be modifiable from cli or console? So far, we have some RWHandler to support read-write properties about pool. Should we make every properties modifiable?	1	1	https://issues.jboss.org/browse/AS7-1379
1928	AS7-1365	If this one is rejected, what still needs to happen for mod_proxy and mod_jk support?	1	1	https://issues.jboss.org/browse/AS7-1365
1929	AS7-1365	Remy, Are you talking about the automagic logic in the mod_cluster subsystem? There are a couple issues with this: * The autogenerated jvm-routes are UUIDs. While the generated values are indeed consistent (i.e. not random), they aren't really known until the server starts. Since mod_jk/mod_proxy needs to define these explicitly (e.g. in workers.properties), this becomes a little awkward to setup. * mod_jk/mod_proxy users may not realize that they need the mod_cluster subsystem for jvm-route autogeneration (it's not obvious). This is more of a documentation issue, really.	1	1	https://issues.jboss.org/browse/AS7-1365
1930	AS7-1365	Is there an expectation that the jvmRoute won't be configurable at all in the future? If so, I understand not adding a config option for it to our management API. If not, IMHO forcing people to use a system property to make a common configuration change is poor. However it gets sets, we should definitely at least expose a read-only attribute for it.	0	0	https://issues.jboss.org/browse/AS7-1365
1931	AS7-1365	Wouldn't the mod_cluster service in AS 7 display that actual jvmroute ? I don't quite know.	0	0	https://issues.jboss.org/browse/AS7-1365
1932	AS7-1365	A user asked about this on the AS7 forum http://community.jboss.org/message/632900#632900. I see that this issue is marked as resolved. Just to be sure, the system property is currently available and the open item is to find a suitable name for this? Should this JIRA be marked as re-opened to reflect this pending task?	1	0	https://issues.jboss.org/browse/AS7-1365
1933	AS7-1358	Chaps, this change doesn't appear to have made it into jboss-as master yet, as of now. {{git blame transactions/src/main/java/org/jboss/as/txn/TransactionSubsystemAdd.java}} reveals Chris's commit, but the line to add the alias "java:/jboss/UserTransaction" is missing (despite showing up when you check {{git show 919c16e2}}). Is this by design? If not, can we get it in?	1	1	https://issues.jboss.org/browse/AS7-1358
1934	AS7-1358	I checked the jboss admin console. The value of jndi java:jboss/UserTransaction is listed as unknown. Maybe this is the problem? I read Seam transaction 3.1.0's code, it does look up this jndi. What's the value for that jndi in your jboss admin console?	0	0	https://issues.jboss.org/browse/AS7-1358
1935	AS7-1356	ISnt this done?	0	0	https://issues.jboss.org/browse/AS7-1356

#	Requirement	Questions	Answered	Successfully Answered	URL
1936	AS7-1338	How come this issue got rejected? Can't find another issue regarding accessing the JNDI from a client (or other server) - or am I missing something?	0	0	https://issues.jboss.org/browse/AS7-1338
1937	AS7-1338	Thanks, that's what I'm looking for. Guess you can't say when in time the preview build will contain it?	0	0	https://issues.jboss.org/browse/AS7-1338
1938	AS7-1338	Was previously rejected: <opalka> wolfc, dmlloyd I'm not able to connect to JNDI on AS7. in the AS7 config I see it should be running on 1099, but I'm getting http://fpaste.org/wgy8/ :(<opalka> wolfc, dmlloyd in the config I see the following ports defined <opalka> --- <opalka> <socket-binding name="jndi" port="1099"/> <opalka> <socket-binding name="jmx-connector-registry" port="1090"/> <opalka> <socket-binding name="jmx-connector-server" port="1091"/> <opalka> --- <dmlloyd> we don't do JNDI over RMI anymore <dmlloyd> that listing is an error <opalka> dmlloyd, So we're not able to do JNDI lookups to AS7 anymore and don't plan to? <opalka> dmlloyd, or we'll be doing JNDI over different binary protocol? <dmlloyd> no, we will be able to, it'll just work somewhat differently <opalka> dmlloyd, ok, I'm excluding the tests then ... <opalka> dmlloyd, Is there a JIRA covering this issue so I can use it in my excludes comment? <dmlloyd> not that I'm aware of, it's just an outstanding task <dmlloyd> remote invocation is not planned to be covered in next months' beta release <opalka> dmlloyd, ok, thx for info ;) But now it's back :)	1	1	https://issues.jboss.org/browse/AS7-1338
1939	AS7-1338	Hi all, I'm just adding some brief info from John about newly implemented remote jndi support and how to use it. As I understood everything what should be reachable in jndi from outside the server must be registered under "java:jboss/exported/" context. As an example I used it in standalone jms client to get connection factory and queue from jndi: {code} // jms client code // initialize context final Properties env = new Properties(); env.put(Context.INITIAL_CONTEXT_FACTORY, "org.jboss.naming.remote.client.InitialContextFactory"); env.put(Context.PROVIDER_URL, "remote://localhost:4447"); Context context = new InitialContext(env); // do lookups ConnectionFactory cf = (ConnectionFactory) context.lookup("jms/RemoteConnectionFactory"); Queue queue = (Queue) context.lookup("jms/queue/test"); // do whatever you want {code} How are jms objects registered to jndi in messaging subsystem in standalone-ha.xml: {code} <jms-connection-factories> ... <!-- configuration for connection factory --> <connection-factory name="RemoteConnectionFactory"> <connectors> <connector-ref connector-name="netty"/> </connectors> <entries> <entry name="java:jboss/exported/jms/RemoteConnectionFactory"/> <!- this is the important part --> </entries> </connection-factory> ... </jms-connection-factories> <jms-destinations> ... <jms-queue name="testQueue"> <entry name="java:jboss/exported/jms/queue/test"/> <!- this is the important part --> </jms-queue> ... </jms-destinations> {code} Used libraries (in my simple client): hornetq-core-2.2.11.Final.jar hornetq-jms-2.2.11.Final.jar jboss-as-naming-7.1.0.Final-SNAPSHOT.jar jboss-jms-api_1.1_spec-1.0.0.Final.jar jboss-logging-3.1.0.GA.jar jboss-marshalling-1.3.6.GA.jar jboss-marshalling-river-1.3.6.GA.jar jboss-modules.jar jboss-msc-1.0.2.GA.jar jboss-remote-naming-1.0.0.Beta4.jar jboss-remoting-3.2.0.CR9.jar jboss-sasl-1.0.0.Beta9.jar netty-3.2.6.Final.jar xnio-api-3.0.0.GA.jar xnio-nio-3.0.0.GA.jar Note: I had to add static block to my class so I could connect to server through remoting: {code} static { Security.addProvider(new JBossSaslProvider()); } {code} Hopefully I didn't forget anything and this info will help. Thanks for this cool feature, Mirek	1	1	https://issues.jboss.org/browse/AS7-1338
1940	AS7-1338	I have followed these instructions and have successfully sent a remote message. However, after each message, I see the following logs on the server. Is there something else I need to do? ERROR [org.jboss.remoting.remote] JBREM000200: Remote connection failed: java.io.IOException: An existing connection was forcibly closed by the remote host INFO [org.jboss.as.naming] JBAS011806: Channel end notification received, closing channel Channel ID 58ff39c8 (inbound) of Remoting connection 01889fce to null	1	1	https://issues.jboss.org/browse/AS7-1338
1941	AS7-1338	Ed what does your code look like? And when you say after each "message" what kind of message do you mean? Also is this an Windows OS? And do those log messages affect any functionality? Just trying to understand what the problem might be.	1	1	https://issues.jboss.org/browse/AS7-1338
1942	AS7-1338	@Ed when you say "after each message". Is this a java test class, which you are running each time, so creating a new process/VM for each message? We see this on VM shutdown, but not per request.	1	1	https://issues.jboss.org/browse/AS7-1338
1943	AS7-1338	@Jason, you had posted some example code on this thread, to illustrate how to connect to a remote queue. Where did that code go? I am having problems with this after downloading 7.1.0.Final, and I just want to make sure I am following the steps correctly. Can you please point me to a link where those directions are?	1	1	https://issues.jboss.org/browse/AS7-1338
1944	AS7-1338	Hello. I have the experience of 11 years with JBOSS. Recently we wanna to migrate from JBOSS 6 to JBOSS 7.1.0 final Here is a part of my code final Hashtable jndiProperties = new Hashtable(); jndiProperties.put(Context.PROVIDER_URL, "remote://localhost:4447"); jndiProperties.put(Context.INITIAL_CONTEXT_FACTORY, "org.jboss.naming.remote.client.InitialContextFactory"); jndiProperties.put("remote.connection.default.username","usera"); jndiProperties.put("remote.connection.default.password","passa"); Context context; try { context = new InitialContext(jndiProperties); I may put username and password, I may not put it - all the time I got ERROR: JBREM000200: Remote connection failed: javax.security.sasl.SaslException: Authentication failed: all available authentication mechanisms failed javax.naming.NamingException: Failed to create remoting connection [Root exception is java.lang.RuntimeException: javax.security.sasl.SaslException: Authentication failed: all available authentication mechanisms failed] at org.jboss.naming.remote.client.ClientUtil.namingException(ClientUtil.java:36) at org.jboss.naming.remote.client.InitialContextFactory.getInitialContext(InitialContextFactory.java:117) at javax.naming.spi.NamingManager.getInitialContext(NamingManager.java:684) at javax.naming.InitialContext.getDefaultInitCtx(InitialContext.java:307) at javax.naming.InitialContext.init(InitialContext.java:242) at javax.naming.InitialContext.<init>(InitialContext.java:216) JBoss is of course running and the user added to application users. Any help? Is JBoss 7.1.0 robust enough to port the system from Jboss 6 to it?	1	1	https://issues.jboss.org/browse/AS7-1338
1945	AS7-1338	Hi, Can someone please summarize what is the final resolution for this bug? As I understand from the bug, "JBAS011806: Channel end notification received, closing channel Channel ID 72d0f3d7 (inbound) of Remoting connection 67ffa9c7 to null" - this message in logs is fine and will come on Windows and currently no way to suppress it for this OS. We are testing with JBoss 7.1.1.Final and see this message at INFO level in JBoss console. Also it seems that this statememnt is logged via System.err.println also from code. We just need to know what is the final stance on this one. Thanks, Himanshu	1	0	https://issues.jboss.org/browse/AS7-1338
1946	AS7-1301	How is the client supposed to calculate the system loader paths? Internally I see JDKPaths being used. public static DependencySpec createSystemDependencySpec(final PathFilter importFilter, final PathFilter exportFilter, final Set<String> loaderPaths)	1	1	https://issues.jboss.org/browse/AS7-1301

#	Requirement	Questions	Answered	Successfully Answered	URL
1947	AS7-1301	The client won't normally want to do this as it incurs an expensive computation at boot. What is the use case here? Normally in AS you'd simply add a dependency on javax.api, and maybe sun.jdk and org.jboss.modules.	1	1	https://issues.jboss.org/browse/AS7-1301
1948	AS7-1301	Thinking about it ... I always have an explicit import filter, so the set of paths does not need to be greater than what the filter accepts. So I could derive the paths from the filter. WDYT?	1	1	https://issues.jboss.org/browse/AS7-1301
1949	AS7-1293	Is there anyway to get documentation on how to use this? Specifically on what NameSpace these entries are bound into? I am trying to use this, but can't as I can't get a read on how to reference them from my code and/or jboss-web.xml for resource references.	0	0	https://issues.jboss.org/browse/AS7-1293
1950	AS7-1236	What about the marking of the jar files as *.failed? Is that correct behavior?	0	0	https://issues.jboss.org/browse/AS7-1236
1951	AS7-1236	Okay I guess on the .failed markers, but I have to ask a question, so if we are deploying just things likes .war, .ear, .sar, and a host of other well known extensions, why would this even happen. In the past I thought there was a "deployable" list of extensions JBoss AS/EAP would look at what changed requirement/need wise, or was this just some oversight? Also curious does stuff like this end up in unit tests? How would I look at those if they exist?	1	1	https://issues.jboss.org/browse/AS7-1236
1952	AS7-1236	Yeah I get that a jar file is a standard deployment, but it is usually in the context of a war/ear/etc....just found it odd that something in a directory that did not end in the well known types (war, ear, sar, etc) would be scanned, and then that it would be crawled it, and .failed file would be created in the exploded directory..just odd to me. Are there cases where other "things" could be done to an exploded directory?	1	1	https://issues.jboss.org/browse/AS7-1236
1953	AS7-1190	Hi Jim, What kind of class loader isolation are you looking for? Typically, in AS7 users can package the libraries of their choice within their application and have it picked up by the server. Some documentation is here https://docs.jboss.org/author/display/AS7/Class+Loading+in+AS7 and also in this forum thread http://community.jboss.org/message/608887#608887. Is that what you are looking for?	1	0	https://issues.jboss.org/browse/AS7-1190
1954	AS7-1177	Should a new subsystem be created ?	1	1	https://issues.jboss.org/browse/AS7-1177
1955	AS7-1177	Hi, i am already working on mail subsystem https://github.com/ctomc/jboss-as/tree/mail but I am not done yet... As for your configuration goes, what would be the purpose of multiple mail-sessions? my idea for config was something similar to what was in as4/5/6 like this: {code:xml} <subsystem xmlns="urn:jboss:domain:mail:1.0"> <jndi-name>java:/Mail</jndi-name> <properties> <property name="mail.store.protocol" value="pop3"/> <property name="mail.transport.protocol" value="smtp"/> </properties> </subsystem> {code} wdyt?	1	1	https://issues.jboss.org/browse/AS7-1177
1956	AS7-1177	Tomaz, I would like to help you on implementing this. Is that possible ?	1	1	https://issues.jboss.org/browse/AS7-1177
1957	AS7-1177	See my comments on AS7-1375 - once you have this solved then why don't you go ahead with a pull request and we'll start the first round of code review. Thanks!	1	1	https://issues.jboss.org/browse/AS7-1177
1958	AS7-1163	This top-level license.txt file is showing LGPL v3. Is that deliberate or should it be v2.1?	1	0	https://issues.jboss.org/browse/AS7-1163
1959	AS7-1136	Last time I booted up AS it seemed to pick up the host name and node name properly. Can I close this?	1	1	https://issues.jboss.org/browse/AS7-1136
1960	AS7-1095	Does this JIRA cover restart of the host controller without having to restart the servers? This is an important use case when we want a host controller to take over the responsibilities of a failed domain controller from another host (same domain).	1	1	https://issues.jboss.org/browse/AS7-1095
1961	AS7-984	Is this issue going to be assigned to an AS7 version and component?	0	0	https://issues.jboss.org/browse/AS7-984
1962	AS7-967	From email discussion and AS7-1465, it sounds like we need to add joda time module and have HV 4.2 depend on it. See email discussion on http://lists.jboss.org/pipermail/hibernate-dev/2011-August/006863.html Also, should this jira be closed, since we upgraded to HV 4.2 already?	1	1	https://issues.jboss.org/browse/AS7-967
1963	AS7-936	<bstansberry> darranl: when you do, can you do a general cleanup on the interface configs? get down to 2 interface definitions, management and ???public??? and no example configs etc, just boilerplate	1	1	https://issues.jboss.org/browse/AS7-936
1964	AS7-931	Could you please add some coverage to the wiki on the set of supported operations/attibutes that manage the osgi subsystem? There is an important distinction between the OSGi subsystem already being activated or not. Some of the configuration affect boot time only (i.e. activation, modules). How do changes of these properties affect the subsystem? Generally, should boot time config even be writeable? How is adding/removing a module that effects system classpath packages supposed to work? I believe that framework properties are read-only by spec. Is this reflected in the management operations? I realize and assume that much of the existing code is part of a learning exercise. Following this I suggest to go back to the drawing board and define a concise set of mgmnt operations that make sense for the OSGi subsytem. A discussion on the ml (referenced from here) would also be cool.	0	0	https://issues.jboss.org/browse/AS7-931
1965	AS7-928	{quote} and model the new module after one of the current modules listed there (perhaps org.javassist). org.javassist is in as7/build/src/main/resources/modules/org/javassist/main). {quote} i don't get this.... {quote} 5. Next, add the "org.hibernate.envers" dependency to the "org.hibernate" module (in as7/build/src/main/resources/modules/org/hibernate/main.xml) {quote} iirc, "org.hibernate" should be a dependency of "org.hibernate.envers" and, envers also has a hibernate-tools dependency, should it be a module too?	1	1	https://issues.jboss.org/browse/AS7-928
1966	AS7-925	This issue is being relocated to 7.1.2 so that 7.1.1 only contains criticals, blockers, and EAP LA issues (which are probably not yet at the right status, or need an extra triage pass). If these issues are completed in a 7.1.1 timeframe please change the fix version. Although if you forget I will bulk correct.	1	1	https://issues.jboss.org/browse/AS7-925
1967	AS7-911	Is there any more work that needs to be done on this?	1	1	https://issues.jboss.org/browse/AS7-911
1968	AS7-909	Hi Clebert, I see there is a 2.2.3-SNAPSHOT at the repo site https://repository.jboss.org/nexus/content/groups/public/org/hornetq/hornetq-core/ There is no 2.2.3-Final. Should the snapshot be used? Howard	0	0	https://issues.jboss.org/browse/AS7-909
1969	AS7-909	Can we validate this?	1	0	https://issues.jboss.org/browse/AS7-909
1970	AS7-887	We are using Jboss EAP6.0.1 and used foo-commonj workmanager.sar that implemented on ServiceMBean and ServiceMBeanSupport. It received same exception as reported here and included previous version supported jars. Can I use your fix that is 6.1.0 jars to resolve this issue. Please help me in this. -Rajesh	1	1	https://issues.jboss.org/browse/AS7-887
1971	AS7-887	Pl kindly tell me which corrected jars i can use for Jboss EAP 6.0.1. Can you pl send any nightly build link if you have. Thanks Rajesh	1	1	https://issues.jboss.org/browse/AS7-887

#	Requirement	Questions	Answered	Successfully Answered	URL
1972	AS7-887	>> Sorry I am bothing you as I dont have support and I have to stick to EAP 6.0.1. EAP 6.0.1 binary is a paid product and you are entitled to support. Where did you get it from? You might want to check with your team to see who has access to the customer service portal. >> We are migrating app from Websphere to Jboss 6.0.1 EAP. We have websphere's workmanager feature which needs to be migrate using commonj. I tried to foo-commonj.sar to make use of workmanager features from ejb in Jboss 6.0.1 EAP. I used previous versions support jars to register commonj jndi in Jboss 6.0.1 EAP. You already created 2 other JIRAs to ask that question. Please do use the JIRAs for such questions. The user forums are the place where this should be discussed.	0	0	https://issues.jboss.org/browse/AS7-887
1973	AS7-838	I'm trying to stand up a SecurityDomainService with a custom LoginModule during app deployment on AS7. Have a look at the code excerpt here: https://gist.github.com/991956 I've pored over it a few times and it looks like it should do what I want, but when I try to authenticate against it, I get errors about no users.properties files as though it doesn't load my custom login module and instead defaults to UsersRolesLoginModule. 17:05 lanceball: asaldhan: hey anil - I have a question about standing up a SecurityDomainService with a custom LoginModule during app deployment on AS7 17:05 lanceball: if you have a sec, can you take a look at this gist? https://gist.github.com/991956 17:06 lanceball: I've pored over it a few times and it looks like it should do what I want, but when I try to authenticate against it, I get errors about no users.properties files 17:06 lanceball: as though it doesn't load my custom login module and instead defaults to UsersRolesLoginModule 17:07 asaldhan: lanceball: not sure we have actually figured it out on AS7 17:07 asaldhan: lanceball: is this as7 or as6 17:07 lanceball: as7 17:08 lanceball: asaldhan: I've tried to walk the jaas code but had a hard time figuring out where you load up login modules 17:09 asaldhan: lanceball: doubt it works yet. We still figuring out stuff. https://issues.jboss.org/browse/AS7-838 17:09 jbossbot: ┐jira┐☆ [└3AS7-838☆] Allow individual security domains to be deployed [└10Open (Unresolved) Feature Request☆, └7 Major☆, └6 Marcus Moyses☆] https://issues.jboss.org/browse/AS7-838 17:09 asaldhan: lanceball: marcus and I were just broaching this topic. Plan is to work on it either tomorrow (Marcus) or Monday(I) depending on who finishes the current item 17:09 lanceball: asaldhan: thanks - I'll add my gist and comments to the Jira! 17:10 asaldhan: lanceball: there is a method to create security domains tho 17:10 asaldhan: lanceball: let me get it for u 17:10 lanceball: asaldhan: I'm successfully creating security domains on the fly with existing jboss LoginModules 17:11 lanceball: it's just a problem when I want to use a custom login module 17:11 asaldhan: lanceball: not sure if we have tested against custom LM yet. I will defer it to Marcus tomorrow morn. I will ask him once he is online	1	1	https://issues.jboss.org/browse/AS7-838
1974	AS7-838	We do have a unit test for custom login modules (CustomLoginModuleTestCase currently in teststuite2/internals) that was created for this purpose. It tests for the most basic scenario which is including the custom login module in a web app and trying to use it. Where is this TorqueBoxLoginModule class included?	1	1	https://issues.jboss.org/browse/AS7-838
1975	AS7-834	maybe we should introduce something like Phase.java for configurations, so we have one place that we set the Configurator order [09:56:06] <Jaikiran> it would help with readability but other than that probably won't make much of a difference [09:56:07] <Jaikiran> would it/ [09:56:08] <Jaikiran> ? [09:56:12] <stuartdouglas> because I think as we start adding more interceptors etc figuring out the order is just going to turn into a mess [09:56:26] <Jaikiran> yeah, I agree with that [09:59:09] <stuartdouglas> how hard is it to add the transaction interceptor back in? because I think that would fix all the EPCPropagationTestCase failures [09:59:10] *** jamezp_afk has quit IRC [10:00:06] <Jaikiran> the tx interceptors are already in [10:00:12] <stuartdouglas> and other than that there is only really the weld failures I mentioned before, and the multiple bindings failures, which I think dmlloyd is working on [10:00:24] <Jaikiran> the EPCpropagation tests are failing due to this EJBContext issue that i'm fixing [10:00:43] <Jaikiran> yeah, the multiple bindings thing is affect the persistencecontext/unit test [10:00:58] <stuartdouglas> I see those tests failing with javax.persistence.TransactionRequiredException [10:01:09] <stuartdouglas> I don't think the transaction interceptor is actually enabled [10:01:10] * Jaikiran looks again [10:01:38] <stuartdouglas> that is after my changes, so your problem is probably hiding the transaction one [10:02:23] <Jaikiran> oh might be [10:02:27] <Jaikiran> because mine is failing with http://pastebin.com/k4BDxyhY [10:02:39] <Jaikiran> i'll take a look at those tests while i work on this fix [10:02:42] *** rmaucher has joined #jboss-as7 [10:02:51] <Jaikiran> can you pastebin your log [10:03:03] <Jaikiran> just to see what i might have missed with tx interceptors [10:04:18] *** rawbdor has quit IRC [10:04:21] <stuartdouglas> http://pastebin.com/jqp9cMXA [10:05:23] <stuartdouglas> actually it could be something else [10:05:28] * stuartdouglas looks at test again [10:06:18] *** AndyTaylor has joined #jboss-as7 [10:06:18] *** ChanServ sets mode: +v AndyTaylor [10:09:09] <stuartdouglas> I'm not really sure what is going on, I will try and debug it [10:09:54] *** wolfc has joined #jboss-as7 [10:09:54] *** ChanServ sets mode: +v wolfc [10:10:14] *** davidbos has joined #jboss-as7 [10:11:09] <Jaikiran> brb after lunch [10:11:13] *** Jaikiran is now known as Jaikiran|Lunch [10:23:42] *** alesj has joined #jboss-as7 [10:26:06] <stuartdouglas> Jaikiran|Lunch: It seems to be an ordering issue the transaction interceptor and the SessionInvocationContextInterceptor appear after the ComponentDispatcherInterceptor, and so never get invoked [10:26:51] *** aslak has quit IRC [10:27:22] *** aslak has joined #jboss-as7 [10:27:22] *** ChanServ sets mode: +v aslak [10:27:40] *** aslak has quit IRC [10:28:02] *** stalep has joined #jboss-as7 [10:31:30] *** hardy has joined #jboss-as7 [10:31:46] *** Nihility has quit IRC [10:41:04] *** galderz has joined #jboss-as7 [10:41:04] *** ChanServ sets mode: +v galderz [10:41:45] *** zroubali has joined #jboss-as7 [10:41:45] *** ChanServ sets mode: +v zroubali [10:54:47] *** kkhan has joined #jboss-as7 [10:54:47] *** ChanServ sets mode: +v kkhan [11:01:45] *** davidbos1 has joined #jboss-as7 [11:02:24] *** galderz has quit IRC [11:02:35] *** davidbos has quit IRC [11:05:20] <stuartdouglas> Jaikiran|Lunch: You there? [11:06:15] *** davidbos1 has quit IRC [11:25:14] *** Jaikiran|Lunch is now known as Jaikiran [11:25:16] <Jaikiran> stuartdouglas: yep [11:26:00] <stuartdouglas> I changed the order round so the interceptor runs, but now I am getting a NPE as EjbComponent.txAttrs is null [11:26:19] <stuartdouglas> and I am not 100% sure how to build that up from the info that is in the component description [11:26:46] <Jaikiran> i'll take that up, since it all requires rewiring the components like we had in the previous ee framework [11:31:45] *** aslak has joined #jboss-as7 [11:31:45] *** ChanServ sets mode: +v aslak [11:34:23] <stuartdouglas> Jaikiran: I have pushed my (minor) changes to my ejb branch [11:34:43] <Jaikiran> ok, i'll get them into my workspace [11:34:50] *** sannegrinovero has joined #jboss-as7 [11:34:50] *** sannegrinovero has quit IRC [11:34:50] *** sannegrinovero has joined #jboss-as7 [11:34:50] *** ChanServ sets mode: +v sannegrinovero [11:36:08] *** miclorb_ has joined #jboss-as7 [11:37:24] <stuartdouglas> I	1	1	https://issues.jboss.org/browse/AS7-834
1976	AS7-832	What is the reason for not implementing the feature to be able to deploy EJBs in OSGi bundles? Looks like a very good match.	0	0	https://issues.jboss.org/browse/AS7-832
1977	AS7-832	Is it still open or resolved ?	1	0	https://issues.jboss.org/browse/AS7-832
1978	AS7-826	Reopen. Why does this needs to be ignored?	0	0	https://issues.jboss.org/browse/AS7-826
1979	AS7-783	I'd just make this "Open to Community" but my inclination would be to reject such a change anyway, so I don't want to mislead people. If deployments are not done in a batch, and 2.jar depends on 1.jar, 2.jar will fail if the scanner deploys it first. And the user will have no way out, unless we start adding in all that terrible pluggable deployment sorting logic we had in AS < 7.	1	1	https://issues.jboss.org/browse/AS7-783

#	Requirement	Questions	Answered	Successfully Answered	URL
1980	AS7-772	task needs to be opened. Stefano Maestri also posted about the missing aspects on the dev list in July. Here's an IRC snippet with Scott Marlow that illustrates the issue: [1:08pm] smarlow: if the /deployment=foo.ear/subsystem-jpa/... structure is going to change, as I evolve the support, that would probably break the cli users, I would think [1:08pm] bstansberry: correct; the structure should be nailed down [1:09pm] smarlow: so, might need to build more long term support now but I'm conflicting with doing more than is needed [1:09pm] smarlow: okay, good to know [1:09pm] bstansberry: but I don't think you can really do that across JPA providers [1:09pm] bstansberry: there's no contract of what people are supposed to expose [1:10pm] bstansberry: I could see /deployment=foo.ear/subsystem-jpa/provider=hibernate/session-factory=xxx [1:10pm] smarlow: I was thinking that it would be nice to allow people to do hibernate 4 stuff, eventually, someone will want to see hibernate 3 stuff possibly also. And OGM [1:11pm] bstansberry: so /deployment=foo.ear/subsystem-jpa/provider=hibernate -- anything under that should be stable [1:12pm] bstansberry: but /deployment=foo.ear/subsystem-jpa/provider=ogm -- whatever under that is provided by the ogm guys [1:12pm] bstansberry: and whether /deployment=foo.ear/subsystem-jpa/provider=hibernate exists or not depends on whether hibernate is the provider [1:12pm] bstansberry: which is fine [1:13pm] smarlow: If I start out with only support for /deployment=foo.ear/subsystem-jpa/provider=hibernate and later have other things, I think that is locked down enough [1:13pm] bstansberry: BTW, see org.jboss.as.webservices.dmr.WSExtension for how they register their stuff under /deployment=foo.ear/subsystem=webservices [1:13pm] bstansberry: right [1:15pm] bstansberry: https://github.com/jbossas/jboss-as/blob/master/webservices/server-integration/src/main/java/org/jboss/as/webservices/dmr/WSExtension.java#L72 [1:17pm] smarlow: hmm [1:17pm] smarlow: final ManagementResourceRegistration deployments = subsystem.registerDeploymentModel(new DescriptionProvider() { [1:18pm] smarlow: this is subsystem start time code? [1:18pm] smarlow: or app deployment time? [1:18pm] smarlow: I guess the method name tells registerDeploymentModel [1:19pm] bstansberry: this is Extension initialization code; so normally boot time [1:19pm] bstansberry: it's the same as your JPAExtension [1:20pm] bstansberry: but instead of just registering resources for the /subsystem=webservices part of the tree, the subsystem.registerDeploymentModel will make them available for use use /deployment=*/subsystem=webservices/ [1:21pm] bstansberry: the tricky part for you is you don't know what providers to register [1:22pm] bstansberry: how are providers integrated? [1:23pm] smarlow: sound like I need to register this management code statically for all known providers that I want to manage [1:23pm] bstansberry: you could do a ServiceLoader thing [1:24pm] bstansberry: if the modules for the providers are going to be visible to the JPA extension module when JPAExtension is called [1:24pm] smarlow: currently, the provider can be packaged with the app (picked up via ServiceLoader thing) or can be an AS7 module (org.hibernate:3 or org.hibernate:main currently) [1:24pm] bstansberry: ah, ok [1:24pm] smarlow: or org.hibernate:ogm [1:25pm] smarlow: there could be a eclipsement provider someday [1:25pm] smarlow: there could be a EclipseLink provider someday [1:25pm] • bstansberry looks at modules.xml for org.jboss.as.jpa [1:25pm] smarlow: but I'm not worried about that this year [1:26pm] bstansberry: ok, org.jboss.as.jpa has no dependency on org.jboss.as.jpa.hibernate.3 or hibernate.4 [1:27pm] bstansberry: so you couldn't use a ServiceLoader to pull in an SPI impl those modules would provide [1:27pm] smarlow: right, its all dynamically done at deployment time now [1:29pm] smarlow: we could include it in the org.jboss.as.jpa.hibernate:3 or org.jboss.as.jpa.hibernate:4 integration classes [1:29pm] smarlow: that is where the integration spi implementation classes are [1:29pm] bstansberry: https://issues.jboss.org/browse/AS7-772 [1:29pm] bstansberry: which only partly solves the problem [1:30pm] bstansberry: https://github.com/jbossas/jboss-	0	0	https://issues.jboss.org/browse/AS7-772
1981	AS7-771	My question about going back and forth between ASYNC and SYNC logging. This is really a number of steps. If there is no ASYNC logger defined, we would need to add one and attach all the handlers to it and update the root logger to use the ASYNC. Should this really be a single op, or should we just make sure we have all the supporting ops available to perform the task?	0	0	https://issues.jboss.org/browse/AS7-771
1982	AS7-753	Jaikiran, yes correct. Is that functionality available through the detyped API?	0	0	https://issues.jboss.org/browse/AS7-753
1983	AS7-738	Is this just a console issue now? I doubt anything more will be done on this for AS 7.1 on the server side, other than fixing any bugs that come up.	1	1	https://issues.jboss.org/browse/AS7-738
1984	AS7-737	/subsystem=datasources/data-source="java:/H2DS":test-connection-in-pool { "outcome" => "success", "result" => [true], "compensating-operation" => undefined } means that we have sucessefully opened a connection with H2DS. Isn't it enough? This operation is already there.	1	1	https://issues.jboss.org/browse/AS7-737
1985	AS7-712	Should we raise a helpdesk ticket for this one, so that Mark Newton and his team and take a look at this?	1	1	https://issues.jboss.org/browse/AS7-712
1986	AS7-705	Can we use i18n resource bundles to customise the look and feel of gwt consoles for the platforms? The projects could then provide a default bundle and the platforms could replace it with another.	0	0	https://issues.jboss.org/browse/AS7-705
1987	AS7-705	Can you elaborate on that? The look and feel consists of - css files - images resources I don't really see how this relates to i18n bundles.	0	0	https://issues.jboss.org/browse/AS7-705
1988	AS7-680	How come this is not targeted for Beta 4? Waiting until the CR seems like a big risk for such a highly used feature as this.	1	1	https://issues.jboss.org/browse/AS7-680
1989	AS7-672	What is the action item for this issue?	0	0	https://issues.jboss.org/browse/AS7-672
1990	AS7-659	Just got hit by the same issue while testing changes on JSPs. As I can understand you don't want to enable dev mode by default, how about triggering that dev mode by an environment variable on startup? We could then adjust the default environment settings directly in JBoss Tools' runtime launch configuration. WDYT?	1	0	https://issues.jboss.org/browse/AS7-659
1991	AS7-650	Does straight dependency mean a jar which users can add to their classpath (even where maven isn't involved)? That would be pretty useful, like the 4.x versions of jbossall-client.jar (ofcourse a better assembled version of it).	0	0	https://issues.jboss.org/browse/AS7-650
1992	AS7-650	Where should hybrid (testing spec features mixed with public API) tests go? org.jboss.as.testsuite.integration.jpa.hibernate.SessionFactoryTestCase is a hybrid test that consists of "spec" API use and non-spec API use also. SessionFactoryTestCase checks that legacy applications can still inject a PersistenceContext into a Hibernate session and a PersistenceUnit into a Hibernate session factory.	1	1	https://issues.jboss.org/browse/AS7-650
1993	AS7-643	So is this done?	0	0	https://issues.jboss.org/browse/AS7-643
1994	AS7-642	Due to feedback from the community, I have split the AS7 and AS6 projects and reopened all unscheduled AS6 issues that are a year or less old. This will make it easier community members to find and work on them. Future releases beyond 6.1 can be done provided a community member steps up to coordinate them.	1	1	https://issues.jboss.org/browse/AS7-642
1995	AS7-623	This is not an issue with respect to CDI 1.0: in that version of the specification, extensions are per-container and not per-BeanManager (cf. 11.5 of the spec). Leaving this open in case a later version of a CDI spec will provide a more granular extension model.	1	1	https://issues.jboss.org/browse/AS7-623
1996	AS7-623	Mark, are you saying my initial assumption is correct? Should it be detected and scoped to the webapp inside the EAR?	0	0	https://issues.jboss.org/browse/AS7-623
1997	AS7-623	See here for people with related issues all over Weld / Seam 3: SEAMFACES-76 SEAMFACES-114 WELD-888 WELD-889 This is a pretty serious issue. It affects more or less each non-trivial EAR-deployment, and certainly almost everyone using Seam 3 (e.g. how should one use Seam Solder / I18N / Config ... in WAR _and_ EJB-Modules?). In my opinion it's not an option to wait for JBoss 7.	1	1	https://issues.jboss.org/browse/AS7-623
1998	AS7-623	Marius, please take care of this issue. It's really a blocker. Do you see a realistic chance to get a 6.1 release containing this bug fix in the near future?	0	0	https://issues.jboss.org/browse/AS7-623

#	Requirement	Questions	Answered	Successfully Answered	URL
1999	AS7-620	What do you think about this Alexey?	1	1	https://issues.jboss.org/browse/AS7-620
2000	AS7-612	Is this really still open? I can build (and run) master (7.1.0.Alpha2-SNAPSHOT) on F-15 with the shipped Maven 3 and OpenJDK. {code}[cbredesen@monza jboss-as]$ java -version java version "1.6.0_22" OpenJDK Runtime Environment (IcedTea6 1.10.4) (fedora-60.1.10.4.fc15-i386) OpenJDK Server VM (build 20.0-b11, mixed mode) [cbredesen@monza jboss-as]$ mvn3 -version /usr/lib/jvm/java Apache Maven 3.0.3 (rNON-CANONICAL_2011-06-28_15-20_mockbuild; 2011-06-28 20:50:31+0530) Maven home: /usr/share/maven Java version: 1.6.0_22, vendor: Sun Microsystems Inc. Java home: /usr/lib/jvm/java-1.6.0-openjdk-1.6.0.0/jre Default locale: en_US, platform encoding: UTF-8 OS name: "linux", version: "2.6.40.6-0.fc15.i686", arch: "i386", family: "unix"{code}	1	0	https://issues.jboss.org/browse/AS7-612
2001	AS7-608	That'll do. But why the discrepancy? How come the QueuelessThreadPoolService and UnboundedThreadPoolService both provide an ExecutorService while the BoundedQueueThreadPoolService only provides an Executor?	1	1	https://issues.jboss.org/browse/AS7-608
2002	AS7-582	Done. https://github.com/jbossas/jboss-as/pull/1985 I'm not sure about the list of artifacts: 1) Does the module name always match the groupId? I counted on that. If not, we need some additional effort. 2) The list I got is simply a list of modules which do not contain jboss.api value="private". Is that enough?	1	1	https://issues.jboss.org/browse/AS7-582
2003	AS7-582	artifact="jbossxts"/> <maven-resource group="org.jboss.jbossts" artifact="jbossxts-api"/> <maven-resource group="org.jboss.jbossts" artifact="jbosstxbridge"/> </module-def> {code} {code} ##### With exported dependencies, converted from module names to groupIDs: echo '' > packages.tmp.txt for i in `find build/src/main/resources/modules/ -name module.xml` ; do FILE=`grep 'value="private"' --files-without-match $i`; if ["$FILE" == ""] ; then continue; fi; echo " Public module: $i" ## Extract module name. PKG=`grep '<module .* name="' $FILE | head -1 | sed 's#<module .* name="\([^"]*\).*"#\1#' | sed 's#/\?>##'`; echo " == $PKG" echo $PKG >> packages.tmp.txt ## Exported dependencies. #grep '<module name="' $FILE | grep 'export="true"' | sed 's#<module name="\([^"]*\).*"#\1#' | sed 's#\s*\(.*\)\s*#\1#' | sed 's#.*# Exported dep: \0#' #grep '<module name="' $FILE | grep 'export="true"' | sed 's#<module name="\([^"]*\).*"#\1#' | sed 's#/\?>##' | sed 's#\s*\(.*\)\s*#\1#' >> packages.tmp.txt grep '<module name="' $FILE | grep 'export="true"' | sed 's#<module name="\([^"]*\).*"#\1#' | sed 's#/\?>##' | sed 's#\s*\(.*\)\s*#\1#' | tee --append packages.tmp.txt | sed 's#.*# Exported dep: \0#' done sort packages.tmp.txt | uniq > modules.tmp2.txt #cat packages.tmp2.txt | sed 's#.*#<include>\0:*</include>#' ### Now we have a list of public API modules, e.g. javax.management.j2ee.api ### Let's convert it into a list of groupIDs. echo > groupIDs.tmp.txt while read -r MODULE ; do echo "Artefacts for module: $MODULE" GROUP_ID=`xsltproc --stringparam moduleName "$MODULE" convertModuleNameToGroupID.xsl build/build.xml` echo " GroupID: $GROUP_ID" echo $GROUP_ID >> groupIDs.tmp.txt done < modules.tmp2.txt cat groupIDs.tmp.txt | sort | uniq > groupIDs.tmp-sorted.txt ### Wrap it as includes for pom.xml. cat groupIDs.tmp-sorted.txt | sed 's#.*#<include>\0</include>#' {code} And the complete output is: {code} com.h2database:h2 commons-configuration:commons-configuration dom4j:dom4j javax.activation:activation javax.enterprise:cdi-api javax.faces:jsf-api javax.inject:javax.inject javax.jws:jsr181-api javax.mail:mail javax.validation:validation-api joda-time:joda-time org.apache.juddi:juddi-client org.apache.juddi.scout:scout org.apache.juddi:uddi-ws org.hibernate.common:hibernate-commons-annotations org.hibernate:hibernate-core org.hibernate:hibernate-entitymanager org.hibernate:hibernate-envers org.hibernate:hibernate-infinispan org.hibernate:hibernate-validator org.hibernate.javax.persistence:hibernate-jpa-2.0-api org.hornetq:hornetq-core org.hornetq:hornetq-jms org.infinispan:infinispan-core org.jboss.as:jboss-as-controller-client org.jboss.com.sun.httpserver:httpserver org.jboss.invocation:jboss-invocation org.jboss:jboss-dmr org.jboss:jboss-ejb-client org.jboss:jboss-remote-naming org.jboss.jbossts:jbosstxbridge org.jboss.jbossts:jbossxts org.jboss.jbossts:jbossxts-api org.jboss:jboss-vfs org.jboss.logging:jboss-logging org.jboss.logging:jul-to-slf4j-stub org.jboss.logmanager:jboss-logmanager org.jboss.marshalling:jboss-marshalling org.jboss.msc:jboss-msc org.jboss.remotingjmx:remoting-jmx org.jboss.remoting3:jboss-remoting org.jboss.sasl:jboss-sasl org.jboss.spec.javax.annotation:jboss-annotations-api_1.1_spec org.jboss.spec.javax.ejb:jboss-ejb-api_3.1_spec org.jboss.spec.javax.el:jboss-el-api_2.2_spec org.jboss.spec.javax.enterprise.deploy:jboss-jad-api_1.2_spec org.jboss.spec.javax.faces:jboss-jsf-api_2.1_spec org.jboss.spec.javax.interceptor:jboss-interceptors-api_1.1_spec org.jboss.spec.javax.jms:jboss-jms-api_1.1_spec org.jboss.spec.javax.management.j2ee:jboss-j2eemgmt-api_1.1_spec org.jboss.spec.javax.resource:jboss-connector-api_1.6_spec org.jboss.spec.javax.rmi:jboss-rmi-api_1.0_spec org.jboss.spec.javax.security.auth.message:jboss-jaspi-api_1.0_spec org.jboss.spec.javax.security.jacc:jboss-jacc-api_1.4_spec org.jboss.spec.javax.servlet:jboss-servlet-api_3.0_spec org.jboss.spec.javax.servlet.jsp:jboss-jsp-api_2.2_spec org.jboss.spec.javax.servlet.jstl:jboss-jstl-api_1.2_spec org.jboss.spec.javax.transaction:jboss-transaction-api_1.1_spec org.jboss.spec.javax.ws.rs:jboss-jaxrs-	1	1	https://issues.jboss.org/browse/AS7-582
2004	AS7-582	Happy to help :) Perhaps we could add some CSS. Do we have some for Javadoc? I've also asked James Cobb.	0	0	https://issues.jboss.org/browse/AS7-582
2005	AS7-580	is there a reason we can't house those javadoc resources inside the AS javadoc aggregate project, rather than attempting to deal with file collisions during dependency-source aggregation? At least then we could be more explicit WRT which resources are used, and which are not. I would think file collisions would be common in this case.	1	1	https://issues.jboss.org/browse/AS7-580
2006	AS7-580	How close does this get you: http://maven.apache.org/plugins/maven-javadoc-plugin/examples/aggregate-dependency-sources.html Pay special attention to the bottom example, which is focused on making javadoc resources available from dependency projects for use in the aggregation step...	0	0	https://issues.jboss.org/browse/AS7-580
2007	AS7-560	Did some of the prep work (isolated Hibernate specific details into org.jboss.as.jpa.hibernate package.	0	0	https://issues.jboss.org/browse/AS7-560
2008	AS7-548	@javax.ejb.Remote annotation scanning support is available since Beta2 http://community.jboss.org/wiki/EJB3supportinJBossAS700Beta2 : {quote} @Remote view: Currently, the @Remote annotation is processed and a view is created out of it. However, remote view semantics don't apply and it works the same as a @Local. {quote} Is this JIRA still valid?	1	1	https://issues.jboss.org/browse/AS7-548
2009	AS7-524	What's left on this task is to: 1) Add in handling of any schema elements that are ommitted in what I've done. 2) Get proper testing. 3) Update the ServerMaker class as mentioned in previous comment.	1	1	https://issues.jboss.org/browse/AS7-524
2010	AS7-514	What's the status on this? It says "Pull Request Sent" but there's no link.	0	0	https://issues.jboss.org/browse/AS7-514
2011	AS7-514	Is there any docs for this? I can't find any in confluence: https://docs.jboss.org/author/display/AS71/Logging+Configuration Here are few bits I have extracted from few discussions - no guarantees: Configured in your app in WEB-INF/classes/logging.properties or log4j.properties or log4j.xml You can turn it off with sysprop {{org.jboss.as.logging.per-deployment = false}} - see http://middlewaremagic.com/jboss/?p=1458 See the commit diff, https://209.132.182.49/message/749989 and http://jboss-as7-development.1055759.n5.nabble.com/logger-config-per-deployment-app-td5709805.html	0	0	https://issues.jboss.org/browse/AS7-514
2012	AS7-514	i created a logging.properties and put it in META-INF folder for the .jar file and deployed, the specific log file is being created ,but app is not logging anything into it..can you point out what am i missing please?	1	0	https://issues.jboss.org/browse/AS7-514
2013	AS7-504	Can we get a yes/no on this issue? If it's a "no" at this point, let's just close it out.	1	0	https://issues.jboss.org/browse/AS7-504

#	Requirement	Questions	Answered	Successfully Answered	URL
2014	AS7-501	time. Following is the discussion we had where we discussed an approach, and outlined its benefits. Perhaps could be useful if someday we decided to pursue this: <maeste2> bstansberry: I wonder if these undocumented generic properties (that JBossTS map to system propeerties) have to be specified in tx subsystem or if we can make anew subsystem for generic properties (not only for tx) <maeste2> bstansberry: in second case we have 2 plus IMHO <maeste2> bstansberry: an xsd documenting the xml sturcture and a subsystem to manage via console <maeste2> bstansberry: #2 the ability to add here some special properties also during EAP phase <maeste2> bstansberry: sincee we are going to use a generic API for this properties/values <maeste2> bstansberry: JBossTS (or other services needing special props) can just depends on this subsystem's service <maeste2> bstansberry: opinion? * bstansberry is pondering this a bit <bstansberry> maeste2: the entire use case goes beyond subsystems though, so this would almost be a core item <bstansberry> maeste2: so to register a handler, that would really be another high-level part of what an Extension does bstansberry> maeste2: so we'd add something in SubsystemRegistration? <bstansberry> maeste2: I'd been thinking to just use system properties, but your ideas have gotten me intrigued <bstansberry> maeste: looking at <property-group subsystem="blah"><property name="blah">value</property><property-group> ... <bstansberry> I see two theoretically benefits over simple system properties <bstansberry> one is if "value" had some more structure, wasn't just a string <bstansberry> but TBH I doubt that's practical, since this needs to be generic <maeste> bstansberry: yup <bstansberry> the other is the subsystem="blah" bit <bstansberry> which is an example of providing some scope to the property <maeste> bstansberry: even if we can use an attribute called type and we can validate against it <maeste> bstansberry: it's not more structure, but it still a plus against system property <maeste> bstansberry: another benefit is that they are manageable (at least theoretically) <bstansberry> maeste: well, system properties are manageable too <maeste> bstansberry: but they haven't custom handlers <bstansberry> maeste: right <bstansberry> that was the intriguing part <maeste> bstansberry: theoretically they could be runtime manageble <maeste> bstansberry: why can't be an OperationStepHandler? <bstansberry> maeste: I guess it can be <bstansberry> maeste: there's an ordering issue <bstansberry> any OSH will only be registered after the relevant extension is added <bstansberry> but all the extension properties have to be evaluated before any subsystem is added <bstansberry> otherwise the subsystem can start without knowing the value <maeste> sure <bstansberry> maeste: so that just means 1) this would have to be a core thing, not a subsystem and 2) in the boot handling its operations have to come before the subsystem ones <bstansberry> maeste: so that's solvable <bstansberry> maeste: we'd also need to add that outside subsystems <maeste> it's a different approach in model, but we are still saying that there are properties scoped at subsystem level <bstansberry> e.g. for core stuff bstansberry> TBH I really dislike adding that to every subsystem though <maeste> oki <bstansberry> I've had way too much experience with what it takes to add some new capability to every subsystem <maeste> returning to ordering issue <bstansberry> that's solvable pretty easily <maeste> oki <bstansberry> just put this new element in the xsd after extension and before subsystem <bstansberry> before profile <maeste> gotcha <bstansberry> maeste: it just took me a bit to realize that <maeste> a bit more for me <bstansberry> maeste: so the advantages of this are 1) runtime handling and 2) scoping <bstansberry> i.e. a finer scope than the entire server <bstansberry> which is the scope of a system property <bstansberry> how about <property-group scope="subsystem=datasource/datasource=blah"><property name="blah">value</property><property-group> ? <bstansberry> scope or "address" or something <maeste> bstansberry: cool <bstansberry> maeste: that starts to make scoping a real benefit <maeste> bstansberry: where type reflect DMR ones and make possible to add a validator on the property? <bstansberry> maeste: so it	0	0	https://issues.jboss.org/browse/AS7-501
2015	AS7-500	What's the status on this one?	1	1	https://issues.jboss.org/browse/AS7-500
2016	AS7-494	Hi Brian, Is there anything wrong with HARMIClient and HARMIServer, other than the fact that they depend on RMI? -Ron	1	1	https://issues.jboss.org/browse/AS7-494
2017	AS7-493	Carlo, is this an issue in the HornetQ integration or the Jacamar integration or elsewhere?	1	1	https://issues.jboss.org/browse/AS7-493
2018	AS7-492	Is this done? We have a contract; seems any remaining task would just be to review it and check if it's complete and well polished.	0	0	https://issues.jboss.org/browse/AS7-492
2019	AS7-485	Is this still valid? I don't see where any integration tests are using JSFUnit right now. As far as I know JSFUnit is only in the smoke tests and that uses a WAR.	0	0	https://issues.jboss.org/browse/AS7-485
2020	AS7-476	Thomas, this issue does not have seem to have a specific corresponding deliverable. If you do have a specific deliverable, can you note it and assign this issue to a release? Otherwise it should be closed or moved to another project. Thanks.	0	0	https://issues.jboss.org/browse/AS7-476
2021	AS7-474	Stuart, can you tell be the current state of affairs? Especially wrt to your last comment?	1	0	https://issues.jboss.org/browse/AS7-474
2022	AS7-470	Brian, can you assign this as you see fit, or simply close it out? Thanks.	0	0	https://issues.jboss.org/browse/AS7-470
2023	AS7-450	Can we get a status of this, please?	0	0	https://issues.jboss.org/browse/AS7-450
2024	AS7-450	Does this task include full support for EJB 2.1 Entity, Session, and Message Driven Beans?	1	1	https://issues.jboss.org/browse/AS7-450
2025	AS7-436	Is this handled by Stefano?	0	0	https://issues.jboss.org/browse/AS7-436
2026	AS7-431	Is the intent here that zipped archives are going to be served directly from the deployments/ dir?	0	0	https://issues.jboss.org/browse/AS7-431

#	Requirement	Questions	Answered	Successfully Answered	URL
2027	AS7-426	Saw this when using interactive demo with both a master and slave DC. Had deployed a war to all servers in all groups, undeployed from main-server-group, then did "undeploy and remove" from other-server-group. This was reported on the master: [Host Controller] 10:59:39,615 ERROR [org.jboss.as.protocol.connection] (pool-1-thread-13) Failed to read a message: java.io.IOException: Failed to handle management operation [Host Controller] at org.jboss.as.protocol.mgmt.AbstractMessageHandler.handleMessage(AbstractMessageHandler.java:63) [Host Controller] at org.jboss.as.protocol.ConnectionImpl.safeHandleMessage(ConnectionImpl.java:254) [Host Controller] at org.jboss.as.protocol.ConnectionImpl$1$1.run(ConnectionImpl.java:213) [Host Controller] at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:441) [:1.6.0_24] [Host Controller] at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) [:1.6.0_24] [Host Controller] at java.util.concurrent.FutureTask.run(FutureTask.java:138) [:1.6.0_24] [Host Controller] at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.access$301(ScheduledThreadPoolExecutor.java:98) [:1.6.0_24] [Host Controller] at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.run(ScheduledThreadPoolExecutor.java:206) [:1.6.0_24] [Host Controller] at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) [:1.6.0_24] [Host Controller] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) [:1.6.0_24] [Host Controller] at java.lang.Thread.run(Thread.java:680) [:1.6.0_24] [Host Controller] at org.jboss.threads.JBossThread.run(JBossThread.java:122) [Host Controller] Caused by: java.lang.IllegalArgumentException [Host Controller] at org.jboss.dmr.ModelValue.asPropertyList(ModelValue.java:100) [Host Controller] at org.jboss.dmr.ModelNode.asPropertyList(ModelNode.java:251) [Host Controller] at org.jboss.as.domain.controller.DomainControllerImpl.getOpsByGroup(DomainControllerImpl.java:627) [Host Controller] at org.jboss.as.domain.controller.DomainControllerImpl.execute(DomainControllerImpl.java:322) [Host Controller] at org.jboss.as.domain.controller.DomainControllerImpl.execute(DomainControllerImpl.java:110) [Host Controller] at org.jboss.as.controller.AbstractModelController.execute(AbstractModelController.java:46) [Host Controller] at org.jboss.as.controller.remote.ModelControllerOperationHandlerImpl$ExecuteAsynchronousOperation.sendResponse(ModelControllerOperationHandlerImpl.java:209) [Host Controller] at org.jboss.as.protocol.mgmt.ManagementResponse$2.handle(ManagementResponse.java:119) [Host Controller] at org.jboss.as.protocol.mgmt.AbstractMessageHandler.handleMessage(AbstractMessageHandler.java:59) [Host Controller] ... 11 more	0	0	https://issues.jboss.org/browse/AS7-426
2028	AS7-424	Hi, Can you please provide example of using EJB3 with IIOP. How to access the EJB3 with IIOP/Corba.? Thanks	0	0	https://issues.jboss.org/browse/AS7-424
2029	AS7-423	Hi Do you work also on HAJNDI lookup ? see last post in: http://community.jboss.org/message/622860#622860	0	0	https://issues.jboss.org/browse/AS7-423
2030	AS7-421	For Hibernate, as a "pesky legacy subsystems" that uses a standard Java-defined API {{Context#bind}} to attempt to make users lives easier, what do you suggest we do?	1	1	https://issues.jboss.org/browse/AS7-421
2031	AS7-403	Will you upgrade Mojarra to 2.0.6? http://download.java.net/maven/2/com/sun/faces/jsf-api/2.0.6/ http://download.java.net/maven/2/com/sun/faces/jsf-impl/2.0.6/	1	1	https://issues.jboss.org/browse/AS7-403
2032	AS7-403	I don't think it would be wise to upgrade to 2.0.6 with CR and final just a few days away. Is there something critical for you in that release?	1	1	https://issues.jboss.org/browse/AS7-403
2033	AS7-402	Brian, is this Jira still relevant? If yes, can you update the description please? I tried to shutdown and restart domain controller, the remote host controller reconnected correctly: {code} [Host Controller] 14:19:51,773 WARN [org.jboss.as.domain.controller] (Thread-20) Could not connect to remote domain controller 127.0.0.1:9999 [Host Controller] 14:19:59,776 WARN [org.jboss.as.domain.controller] (Thread-20) Could not connect to remote domain controller 127.0.0.1:9999 [Host Controller] 14:20:03,466 INFO [org.jboss.as.domain.controller] (Thread-20) Reconnected to master {code}	1	1	https://issues.jboss.org/browse/AS7-402
2034	AS7-395	Alexey, can you drive this issue to completion?	1	0	https://issues.jboss.org/browse/AS7-395
2035	AS7-393	Darran, can you drive this to completion or assign it to someone who can?	1	1	https://issues.jboss.org/browse/AS7-393
2036	AS7-392	Was this finished?	1	1	https://issues.jboss.org/browse/AS7-392
2037	AS7-389	Is this available in AS7 upstream?	0	0	https://issues.jboss.org/browse/AS7-389
2038	AS7-389	Is this complete John?	1	1	https://issues.jboss.org/browse/AS7-389
2039	AS7-388	Any update, or can we drive this to conclusion?	1	1	https://issues.jboss.org/browse/AS7-388
2040	AS7-385	Scott, Brian indicated that you were looking into this. Can you give an update?	1	1	https://issues.jboss.org/browse/AS7-385
2041	AS7-383	Is this something you want to do? Assign back if not. Should be a very minor change.	0	0	https://issues.jboss.org/browse/AS7-383
2042	AS7-367	From tooling I'm more interested to knowing *actual* status of deployments such as: Is it running ? Did it deploy - any warnings? any errors ? stacktraces ? Inner/nested details that arquillian is asking for is nice but i'm much more interested in getting status for deployment overall first with access to error info.	1	1	https://issues.jboss.org/browse/AS7-367
2043	AS7-364	Brian, did you have a specific approach in mind for this?	1	1	https://issues.jboss.org/browse/AS7-364
2044	AS7-364	Brian, is the EAP distinction done by appending -redhat? Or could the information that shows up in boot log at 22:13:52,235 INFO [org.jboss.as] JBoss EAP 6.0.0.Alpha2 (AS 7.1.0.Beta1-redhat-1) starting be the version string returned -- meaning "EAP 6... (....)" ?	0	0	https://issues.jboss.org/browse/AS7-364
2045	AS7-354	What does this mean?	0	0	https://issues.jboss.org/browse/AS7-354
2046	AS7-328	has this been merged?	1	0	https://issues.jboss.org/browse/AS7-328
2047	AS7-318	Aren't we done here? I'd like to close all of these generic holder tasks. If there is any commands we need for usability reasons we should just have a specific issues for them.	1	1	https://issues.jboss.org/browse/AS7-318
2048	AS7-311	Emanuel, could you have a look and see if this is still valid?	0	0	https://issues.jboss.org/browse/AS7-311

#	Requirement	Questions	Answered	Successfully Answered	URL
2049	AS7-308	Hi, I tried to use this fix to solve issue JBAS-7882 on Jboss AS 7.1.0 Final. But it does not seem to work. I put all my security jars in ear/lib and mentioned in ear/META-INF/jboss-deployment-structure.xml like below: <jboss-deployment-structure xmlns="urn:jboss:deployment-structure:1.1"> <ear-subdeployments-isolated>true</ear-subdeployments-isolated> <deployment> <resources> <resource-root path="entbase.jar" use-physical-code-source="true"/> <resource-root path="entarchive.jar" use-physical-code-source="true"/> ... </resources> </deployment> ... But still i am getting Self-Integrity Check FAILED: java.lang.SecurityException: java.security.PrivilegedActionException: java.util.zip.ZipException: zip file is empty: java.lang.SecurityException: Self-Integrity Check FAILED: java.lang.SecurityException: java.security.PrivilegedActionException: java.util.zip.ZipException: zip file is empty at com.entrust.toolkit.security.provider.JCEVerifierImpl.checkIntegrity(Unknown Source) [entbase.jar:7.1.0 build 941 "Jul 25 2005" "03:44:15"] at com.entrust.toolkit.security.provider.Entrust.<init>(Unknown Source) [entbase.jar:7.1.0 build 941 "Jul 25 2005" "03:44:15"] Am i missing something? Please guide. Thanks	0	0	https://issues.jboss.org/browse/AS7-308
2050	AS7-308	I confirm that use-physical-code-source is working with a recent version (7.2) I added it in the jboss-deployment-structure.xml of a war to solve the zip file is empty error about a bouncycastle jar It's working if I deploy the war directly in the application server deployments directory but sadly it doesn't work if I put it in an EAR. I tried to define it as subdeployment of the ear but it didn't help. Do you think that it is normal or I should report a new issue ?	0	0	https://issues.jboss.org/browse/AS7-308
2051	AS7-306	IRC discussion for reference: {quote} (06:56:39 IST) dmlloyd: EE.8.3.2 (06:57:00 IST) dmlloyd: portable applications must not depend on having or not having access to these classes or resources: The content of any EJB jar files included in the same ear file. dmison_a` dmlloyd (06:58:09 IST) Jaikiran: dmlloyd: indeed! (06:58:16 IST) Jaikiran: i wasn't aware that wasn't portable (06:58:24 IST) Jaikiran: but would be really useful to enable that in AS& (06:58:26 IST) Jaikiran: AS7 (06:58:33 IST) Jaikiran: assuming there's no technical issue (06:58:34 IST) dmlloyd: well if you weren't aware I imagine most of our users aren't either (06:58:41 IST) dmlloyd: so perhaps it should be a configurable default (06:59:23 IST) Jaikiran: from what i have seen, users (and even maven?) just creates the top level ejb jar and expects the classes in ejb deployment to be visible in web (06:59:29 IST) fbricon: last time I checked, my sample project ran on glassfish 3.x (06:59:50 IST) hardy [~hardy@81-235-238-123-no92.tbcn.telia.com] entered the room. (07:00:02 IST) dmlloyd: Jaikiran: you mean, in an EAR right? dmison_a` dmlloyd (07:00:07 IST) Jaikiran: dmlloyd: yes (07:00:18 IST) Jaikiran: both war in EJB jar in the .ear (07:00:46 IST) fbricon: dmlloyd: Jaikiran: ^that's my setup (07:00:58 IST) dmlloyd: sounds like a config param for the EE subsystem: <isolate-ear-contents value="true"/> (07:01:12 IST) Jaikiran: that "sounds" different (07:01:18 IST) dmlloyd: can one of you guys open a JIRA for that? assign it to CR1 or something? (07:01:27 IST) Jaikiran: given our usage of ear isolation in previous AS version (07:01:39 IST) dmlloyd: I think the more strict behavior should be the default though. That will encourage people to move towards better spec compliance in their deployments (07:01:50 IST) fbricon: dmlloyd: ok I can create the issue and attach my sample project (07:01:53 IST) Jaikiran: basically, isolate-ear in previous releases meant that each ear would have a separate cl and wouldn't see each other (07:02:03 IST) Jaikiran: so it might be better to use a different term for this attribute (07:02:06 IST) dmlloyd: true, we may want to choose better wording (07:02:09 IST) Jaikiran: yeah (07:02:14 IST) Jaikiran: but yes, a JIRA would be a starting point (07:02:30 IST) Jaikiran: fbricon: yes, please do attach a app for easy testing {quote}	1	1	https://issues.jboss.org/browse/AS7-306
2052	AS7-65	C:\Users\Sacha\workspace\jboss-as>java -version java version "1.6.0_18" Java(TM) SE Runtime Environment (build 1.6.0_18-b07) Java HotSpot(TM) 64-Bit Server VM (build 16.0-b13, mixed mode) Typical failure: [INFO] [INFO] ---------- [INFO] Building JBoss Application Server: OSGi Subsystem 7.0.0.Beta1-SNAPSHOT [INFO] ---------- [INFO] [INFO] --- maven-enforcer-plugin:1.0-beta-1:enforce (enforce-java-version) @ jboss-as-osgi --- [INFO] [INFO] --- maven-enforcer-plugin:1.0-beta-1:enforce (enforce-maven-version) @ jboss-as-osgi --- [INFO] [INFO] --- maven-resources-plugin:2.4.3:resources (default-resources) @ jboss-as-osgi --- [INFO] Using 'UTF-8' encoding to copy filtered resources. [INFO] Copying 2 resources [INFO] [INFO] --- maven-compiler-plugin:2.1:compile (default-compile) @ jboss-as-osgi --- [INFO] Nothing to compile - all classes are up to date [INFO] [INFO] --- maven-checkstyle-plugin:2.5:checkstyle (check-style) @ jboss-as-osgi --- [INFO] Starting audit... Audit done. [WARNING] Unable to locate Source XRef to link to - DISABLED [INFO] [INFO] --- maven-resources-plugin:2.4.3:testResources (default-testResources) @ jboss-as-osgi --- [INFO] Using 'UTF-8' encoding to copy filtered resources. [INFO] Copying 13 resources [INFO] [INFO] --- maven-compiler-plugin:2.1:testCompile (default-testCompile) @ jboss-as-osgi --- [INFO] Nothing to compile - all classes are up to date [INFO] [INFO] --- maven-antrun-plugin:1.4:run (build-dist) @ jboss-as-osgi --- project.artifactId [INFO] Executing tasks modules: [echo] Initializing module -> javax.transaction.api [INFO] Executed tasks [INFO] [INFO] --- maven-surefire-plugin:2.5:test (default-test) @ jboss-as-osgi --- [INFO] Surefire report directory: C:\Users\Sacha\workspace\jboss-as\osgi\target\surefire-reports ---------- T E S T S ---------- Running org.jboss.as.osgi.xservice.BundleDependsOnLocalModuleTestCase 12:09:41,359 INFO [org.jboss.osgi.framework.bundle.BundleManager] (Thread-2) Install bundle: system.bundle:0.0.0 12:09:41,382 INFO [org.jboss.as.osgi] (Thread-2) Starting OSGi Framework 12:09:41,383 INFO [org.jboss.osgi.framework.bundle.FrameworkState] (Thread-2) JBossOSGi Framework Core - 1.0.0.Alpha11 12:09:41,471 INFO [org.jboss.osgi.framework.plugin.internal.StartLevelPluginImpl] (Thread-2) Starting bundles for start level: 1 12:09:41,471 INFO [org.jboss.osgi.framework.bundle.FrameworkState] (Thread-2) Framework started 12:09:41,583 INFO [org.jboss.as.deployment] (main) Activating deployment: bundle1 12:09:41,605 WARN [org.jboss.vfs] (main) VFS was unable to set the URLStreamHandlerFactory. This will have unpredictable results 12:09:41,704 INFO [org.jboss.as.osgi] (Thread-2) Installing deployment: [bundle1:1.0.0,location=vfs:/C:/Users/Sacha/workspace/jboss-as/osgi/deploymen ts/bundle1/] 12:09:41,723 INFO [org.jboss.osgi.framework.bundle.BundleManager] (Thread-2) Install bundle: bundle1:1.0.0 12:09:41,770 INFO [org.jboss.as.osgi] (Thread-2) Uninstalling deployment: [bundle1:1.0.0,location=vfs:/C:/Users/Sacha/workspace/jboss-as/osgi/deploym ents/bundle1/] 12:09:41,772 WARN [org.jboss.osgi.deployment.deployer.SystemDeployerService] (Thread-2) Cannot uninstall: [bundle1:1.0.0,location=vfs:/C:/Users/Sacha /workspace/jboss-as/osgi/deployments/bundle1/]: java.lang.IllegalStateException: Bundle uninstalled: bundle1:1.0.0 at org.jboss.osgi.framework.bundle.AbstractBundle.assertNotUninstalled(AbstractBundle.java:764) at org.jboss.osgi.framework.bundle.AbstractBundle.uninstall(AbstractBundle.java:484) at org.jboss.osgi.framework.bundle.AbstractUserBundle.uninstall(AbstractUserBundle.java:381) at org.jboss.osgi.framework.bundle.BundleWrapper.uninstall(BundleWrapper.java:186) at org.jboss.osgi.deployment.deployer.SystemDeployerService.undeployInternal(SystemDeployerService.java:142) at	1	1	https://issues.jboss.org/browse/AS7-65
2053	ARQ-2121	[~chco] Hi Christian, thanks for opening this issue! Do you think you could create a pull request on GitHub for this at https://github.com/arquillian/arquillian-container-was/ ?	1	1	https://issues.jboss.org/browse/ARQ-2121
2054	ARQ-2082	Could you share reproducer on GitHub? That would tremendously help us analyzing the problem.	0	0	https://issues.jboss.org/browse/ARQ-2082

#	Requirement	Questions	Answered	Successfully Answered	URL
2055	ARQ-2051	How are you downloading the liberty instance from Maven? If you're using the liberty:create-server goal it already provides an option for a configFile like we packaged it in our wlp-managed-85 test. Why is this more useful than that? I would think that if we start adding this, we will need further changes to support other features as well and at that point it just might be easier to include the server.xml as part of the test in the first place anyway; What do you think?	1	1	https://issues.jboss.org/browse/ARQ-2051
2056	ARQ-2051	I've now created a new docker container that works with the current wercker environment and I also moved to the 1.2 version of the liberty-maven-plugin. The bad news is that at least for me the ApplicationMBean is never shown over the localConnector with the versions I tested with from the public maven repositories. I see a few other WebSphere:* MBeans, but not the ApplicationMBean. [~atijms] are you seeing this MBean? Can you tell me which exact JDK and liberty version you're using? tnx!	1	1	https://issues.jboss.org/browse/ARQ-2051
2057	ARQ-2047	[~cthorburn8], I have tried to reproduce this issue as per steps mentioned by you. But I couldn't reproduce your mentioned behavior for lambda expression. You can look for my test project [here|https://github.com/dipak-pawar/ARQ-2047-test] Can you please share your project where this issue is reproducing? Maybe something is missing in my project.	1	0	https://issues.jboss.org/browse/ARQ-2047
2058	ARQ-2040	[~stefanutti] Hi Antonin, if you really mean JTA then this has been already fixed in https://github.com/arquillian/arquillian-container-weld/issues/25 [~mjobanek] [~bmajsak] I can't assing issue to myself. Can you please update my permissions?	1	1	https://issues.jboss.org/browse/ARQ-2040
2059	ARQ-2032	[~thomas.diesler], I do not understand how is this issue connected with Wildfly... I just wanted to track issue somewhere. Shall I rather use github issues?	0	0	https://issues.jboss.org/browse/ARQ-2032
2060	ARQ-2017	[~thomas.diesler] : I notice that the "org.jboss.arquillian.osgi" package is listed both under <Private-Package> and also under <_exportcontents>. I assume that you would prefer one or the other. Let me know which one is appropriate and I will remove the duplicate from the other element. I am also curious why you have declined to inline the embedded dependencies, but to embed them as jars instead. There may be classloader improvements, among other things, if they are inlined instead. Are there licensing issues that require you to embed the jars? Otherwise, I would like to try inlining them instead. Please let me know. Thanks!	0	0	https://issues.jboss.org/browse/ARQ-2017
2061	ARQ-2016	I don't see the "Assign to me" option in your JIRA. Is there a way that I can assign this story, and sub-tasks that I create, to myself?	1	1	https://issues.jboss.org/browse/ARQ-2016
2062	ARQ-2016	I'm asking specifically about the AuxiliaryArchiveAppender. When the auxiliary archive is created, what's the process by which it is deployed? And is it a bundle or is it only intended as a library to go into a web archive? I'm using a jar archive for my test, so I'm trying to understand the best way to make Spock available within the container.	0	0	https://issues.jboss.org/browse/ARQ-2016
2063	ARQ-2016	Hello!, We have adapt the Arquillian OSGi Extension so it can use custom extensions, now we are using it sucessfully in the Arquillian Liferay Extension. The idea that we have implemented is: 1. The Arquillian OSGi Bundle is automatically generated (with all the extensions that are declared in the classpath). (This is done only once) 2. For every test: 2.1. The Arquillian OSGi Bundle is installed. 2.2. The test deployment is installed without any enhacement. 2.3. A fragment for the Arquillian OSGI Bundle that contains only the test class is generated and installed. 2.4 The tests are executed 2.5 All the previously installed bundle are uninstalled. We have created a pull (https://github.com/arquillian/arquillian-container-osgi/pull/51) were you can check the changes. We think that this maybe can be usefull for someone. Maybe it would be interesting try to merge it with the current Arquillian OSGi Container. What do you think? If you need more information, please, let us know. Thanks!	1	1	https://issues.jboss.org/browse/ARQ-2016
2064	ARQ-2008	java.lang.RuntimeException: Could not create new instance of class org.jboss.arquillian.test.impl.EventTestRunnerAdaptor at org.jboss.arquillian.test.spi.SecurityActions.newInstance(SecurityActions.java:165) at org.jboss.arquillian.test.spi.SecurityActions.newInstance(SecurityActions.java:102) at org.jboss.arquillian.test.spi.TestRunnerAdaptorBuilder.build(TestRunnerAdaptorBuilder.java:52) at org.jboss.arquillian.junit.Arquillian.run(Arquillian.java:113) at org.eclipse.jdt.internal.junit4.runner.JUnit4TestReference.run(JUnit4TestReference.java:86) at org.eclipse.jdt.internal.junit.runner.TestExecution.run(TestExecution.java:38) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.runTests(RemoteTestRunner.java:459) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.runTests(RemoteTestRunner.java:675) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.run(RemoteTestRunner.java:382) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.main(RemoteTestRunner.java:192) Caused by: java.lang.reflect.InvocationTargetException at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeConstructorAccessorImpl.newInstance(Unknown Source) at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(Unknown Source) at java.lang.reflect.Constructor.newInstance(Unknown Source) at org.jboss.arquillian.test.spi.SecurityActions.newInstance(SecurityActions.java:161) ... 9 more Caused by: java.lang.RuntimeException: Could not create new instance of class org.jboss.arquillian.core.impl.ManagerImpl at org.jboss.arquillian.core.spi.SecurityActions.newInstance(SecurityActions.java:165) at org.jboss.arquillian.core.spi.SecurityActions.newInstance(SecurityActions.java:102) at org.jboss.arquillian.core.spi.ManagerBuilder.create(ManagerBuilder.java:77) at org.jboss.arquillian.test.impl.EventTestRunnerAdaptor.<init>(EventTestRunnerAdaptor.java:62) ... 14 more Caused by: java.lang.reflect.InvocationTargetException at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeConstructorAccessorImpl.newInstance(Unknown Source) at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(Unknown Source) at java.lang.reflect.Constructor.newInstance(Unknown Source) at org.jboss.arquillian.core.spi.SecurityActions.newInstance(SecurityActions.java:161) ... 17 more Caused by: java.lang.NoClassDefFoundError: org/jboss/arquillian/container/test/impl/enricher/resource/OperatesOnDeploymentAwareProvider at java.lang.ClassLoader.defineClass1(Native Method) at java.lang.ClassLoader.defineClass(Unknown Source) at java.security.SecureClassLoader.defineClass(Unknown Source) at java.net.URLClassLoader.defineClass(Unknown Source) at java.net.URLClassLoader.access$100(Unknown Source) at java.net.URLClassLoader$1.run(Unknown Source) at java.net.URLClassLoader$1.run(Unknown Source) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(Unknown Source) at java.lang.ClassLoader.loadClass(Unknown Source) at sun.misc.Launcher$AppClassLoader.loadClass(Unknown Source) at java.lang.ClassLoader.loadClass(Unknown Source) at org.jboss.arquillian.protocol.servlet.ServletExtension.register(ServletExtension.java:40) at org.jboss.arquillian.core.impl.loadable.LoadableExtensionLoader.load(LoadableExtensionLoader.java:79) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(Unknown Source) at java.lang.reflect.Method.invoke(Unknown Source) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at	1	1	https://issues.jboss.org/browse/ARQ-2008
2065	ARQ-2008	What does your POM file look like now?	1	1	https://issues.jboss.org/browse/ARQ-2008
2066	ARQ-2008	What happens when you run "mvn clean install" from command line?	1	1	https://issues.jboss.org/browse/ARQ-2008
2067	ARQ-1975	Hi, is this the same feature as this one? https://developer.jboss.org/message/935605 If so, there is already a jira for that: ARQ-540	1	1	https://issues.jboss.org/browse/ARQ-1975

#	Requirement	Questions	Answered	Successfully Answered	URL
2068	ARQ-1968	So I don't quite understand. Are you saying that Arquillian already works for Fuse and it's just a matter of RH saying it's a tested framework like Arquillian on EAP (https://access.redhat.com/solutions/275433)?	1	0	https://issues.jboss.org/browse/ARQ-1968
2069	ARQ-1954	How can this be used? I'm running Arquillian Tests inside a remote WildFly instance and I cannot get it to work. :-(	0	0	https://issues.jboss.org/browse/ARQ-1954
2070	ARQ-1921	[~aslak] you mean like you get it in lookup method in "Annotation... qualifiers" parameter? How these annotations would be called? I can look at it.	1	1	https://issues.jboss.org/browse/ARQ-1921
2071	ARQ-1902	I assume you would want to skip the whole Before|Test|After phase? Not just Test? Do you want it 'marked' as Skipped in report?	1	1	https://issues.jboss.org/browse/ARQ-1902
2072	ARQ-1887	I have run into another issue that is related to how the scripts are split before execution. This time I'm trying to insert a simple JavaScript script into my table and all script's content between curly braces is cut out, which causes the test failure, since the script cannot be parsed by the Nashorn engine - example below. {code:title=script.sql|borderStyle=solid} INSERT INTO SCRIPT(ENGINE, BODY) VALUES('JavaScript', 'function execute(context) {var Result = Java.type("xx.xxxxxx.xxxx.xxxxx.Result"); var result = new Result(); if (context.service === undefined) {result.code = 100;} else { result.code = 0;} return result;}'); {code} Result: {code:title=output.log|borderStyle=solid} INSERT INTO SCRIPT(ENGINE, BODY) VALUES('JavaScript', 'function execute(context) else return result;}'); {code} The question is: shall I make a separate issue for this or maybe this comment will be enough? Thanks!	1	1	https://issues.jboss.org/browse/ARQ-1887
2073	ARQ-1887	Please go ahead with dedicated issue. Would you be interested in contributing the fix for this issue or the other one? It's rather straightforward and you can count on my full support :)	1	1	https://issues.jboss.org/browse/ARQ-1887
2074	ARQ-1883	tests/src/test/java/org/jboss/arquillian/integration/persistence/cucumber/SeedingStrategyTest.java]) but its failing (i think its a detail) because persistence unit is not found. I've changed the deployment from *jar* to *war* cause i think cukespace doesn't work in jar project (in fact I didnt tried to flat the jar so changed to war) but the [@PersistenceContext inside test|https://github.com/rmpestano/arquillian-extension-persistence/blob/master/int-tests/src/test/java/org/jboss/arquillian/integration/persistence/cucumber/SeedingStrategyTest.java#L52] isn't working, here is test output: {noformat} org.jboss.arquillian.container.spi.client.container.DeploymentException: Cannot deploy: test.war at org.jboss.as.arquillian.container.ArchiveDeployer.deployInternal(ArchiveDeployer.java:83) at org.jboss.as.arquillian.container.ArchiveDeployer.deployInternal(ArchiveDeployer.java:64) at org.jboss.as.arquillian.container.ArchiveDeployer.deploy(ArchiveDeployer.java:46) at org.jboss.as.arquillian.container.CommonDeployableContainer.deploy(CommonDeployableContainer.java:146) at org.jboss.arquillian.container.impl.client.container.ContainerDeployController$3.call(ContainerDeployController.java:161) at org.jboss.arquillian.container.impl.client.container.ContainerDeployController$3.call(ContainerDeployController.java:128) at org.jboss.arquillian.container.impl.client.container.ContainerDeployController.executeOperation(ContainerDeployController.java:271) at org.jboss.arquillian.container.impl.client.container.ContainerDeployController.deploy(ContainerDeployController.java:127) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.invokeObservers(EventContextImpl.java:99) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:81) at org.jboss.arquillian.container.impl.client.container.DeploymentExceptionHandler.verifyExpectedExceptionDuringDeploy(DeploymentExceptionHandler.java:50) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.container.impl.client.ContainerDeploymentContextHandler.createDeploymentContext(ContainerDeploymentContextHandler.java:78) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.container.impl.client.ContainerDeploymentContextHandler.createContainerContext(ContainerDeploymentContextHandler.java:57) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at	1	1	https://issues.jboss.org/browse/ARQ-1883
2075	ARQ-1852	CommandService is part of Core already. Every Protocol has it's own implementation of it to allow InContainer to communicate with Client side: https://github.com/arquillian/arquillian-core/blob/master/container/test-spi/src/main/java/org/jboss/arquillian/container/test/spi/command/CommandService.java E.g. the ServletCommandService from the Servlet Protocol: https://github.com/arquillian/arquillian-core/blob/master/protocols/servlet/src/main/java/org/jboss/arquillian/protocol/servlet/runner/ServletCommandService.java After re-reading, you want the opposite direction? Client to send commands to Container at any point?	1	0	https://issues.jboss.org/browse/ARQ-1852
2076	ARQ-1852	At which points do the Client know? The normal execution flow in Arquillian is to for each @Test on client side execute the full In Container cycle(BeforeSuite/BeforeClass/Before/Test...), then take the result from container and move on. This is per say single threaded, so the Client can't really send anything at 'random' points to the container. It only sends one 'command'; run the complete test. It's from within the container side you can send fine grained commands back to the client within any of the events(BeforeSuite, BeforeClass..).. If you can explain in more detail what you're trying to do.	1	0	https://issues.jboss.org/browse/ARQ-1852
2077	ARQ-1842	I think this is related to https://issues.jboss.org/browse/ARQ-1807 Could you verify with arquillian-core upstream master branch?	1	1	https://issues.jboss.org/browse/ARQ-1842
2078	ARQ-1842	@Aslak Knutsen: Could you please where can we find the last 2.0.0.CR1 version? Is there any maven repository to download it from?	1	1	https://issues.jboss.org/browse/ARQ-1842
2079	ARQ-1842	Please, is there any estimate on when it is going to be available on maven repos? I would like to use this extension on my project but I keep getting this error. Btw: I am using 1.1.10.Final	1	1	https://issues.jboss.org/browse/ARQ-1842
2080	ARQ-1842	Partial fix based on [~dashorst] work has been already rolled out with version 1.0.2.Final. Is that the one you are using? There is one additional improvement awaiting with this [pull request|https://github.com/arquillian/arquillian-extension-transaction/pull/12]. I will wrap it up this week.	1	1	https://issues.jboss.org/browse/ARQ-1842
2081	ARQ-1842	[~bmajsak] Sorry to ask: is there anywhere I can check if this was released? Is there anything I can do to help this matter?	1	1	https://issues.jboss.org/browse/ARQ-1842

#	Requirement	Questions	Answered	Successfully Answered	URL
2082	ARQ-1809	I cannot reproduce this behaviour, but I have improved sorted and filtered sets so that they have now proper metadata. As for the suggested fix : both YAML and JSON implementation are based on FlatXML data sets, where datatype is also set to unknown. If I understand DBUnit right, the real types are then extracted from underlying db and conversion is applied. But this might be wrong reasoning :) I will take a closer look. In the meantime, could you try to run your test against [latest master|https://github.com/arquillian/arquillian-extension-persistence] and let me know? Does this error occurs each and every time you run the test?	0	0	https://issues.jboss.org/browse/ARQ-1809
2083	ARQ-1800	Thanks for the feature request. Of course this will be handy feature. Shall we consider UTF-8 as a default? Moreover would you be interested in contributing this improvement? That would be simply awesome.	0	0	https://issues.jboss.org/browse/ARQ-1800
2084	ARQ-1797	What kind of strategy do you want to use here when you put @Screenshot annotation on a method? That screenshot annotation has to have some default values. It is clear that specified parameter to that annotation overrides settings in arquillian.xml. In case that property is not specified there, what setting should "win"? Annotation one which is not specified explicitly or property in arq.xml ?	1	1	https://issues.jboss.org/browse/ARQ-1797
2085	ARQ-1789	Have you tried running the Managed/Remote 7 against a 8 Container?	1	1	https://issues.jboss.org/browse/ARQ-1789
2086	ARQ-1756	Thanks for opening this feature request! What would this feature do exactly? Do other arquillian adapters provide this function already? Please also look at the section at http://arquillian.org/guides/getting_started/#test_across_containers and see how this is documented to work for jboss-as artifacts. I think this should also work with Liberty. What do you think?	1	1	https://issues.jboss.org/browse/ARQ-1756
2087	ARQ-1751	This is a _almost_ working prototype from DevNation hackfest; https://github.com/aslakknutsen/arquillian-recorder/tree/ARQ-1751 The missing parts; * probably integrate bowler with Maven to build the UI JS parts automatically * some missmatch in URL used by the UI and backend atm * possible some missmtach in the format used General improvements require some more changes in recorder core. e.g.; * Should the same JAX-B model be used on Client and Container side? * JAX-B in general generates XML, XML is hell'ish in a JS client ** Currently worked around by XML from Test client to Server, and Gson converted jax-b model from Server to JS client * The full report is pushed to server for each @Test * Missing image upload.	1	0	https://issues.jboss.org/browse/ARQ-1751
2088	ARQ-1750	I would've assumed that we use the configuration mechanisms used by liberty itself such as the ones documented in http://pic.dhe.ibm.com/infocenter/wasinfo/v8r5/index.jsp?topic=%2Fcom.ibm.websphere.wlp.nd.doc%2Fae%2Ftwlp_tun.html to set mx and debug parameters; what is the added value of having this as part of the arquillian.xml configuration? Do other container implementations already support this?	1	1	https://issues.jboss.org/browse/ARQ-1750
2089	ARQ-1744	[~kpiwko] [~tkriz] Is it possible to have this in 2.0.0.Alpha1? Having this particular JIRA done, it would simplify Droidium code tremendously and I would move with other things as well for which this one being resolved is very needed prerequisity. Waiting for Alpha2 is just time blocker for me.	1	1	https://issues.jboss.org/browse/ARQ-1744
2090	ARQ-1741	Are you trying to handle a TypeNotPresentException?	1	0	https://issues.jboss.org/browse/ARQ-1741
2091	ARQ-1716	Would it be possible to just bind to Arquillian Events? This is a long term maintenance nightmare - JUnit, TestNG, Cukespace, Spock, etc...	1	1	https://issues.jboss.org/browse/ARQ-1716
2092	ARQ-1572	This will work nicely for JARs and WARs, but not for EARs, as there is no 'global classloader' for the ear package as far as I can say. I will keep it as jar in the lib/ folder with the fixed name, eg. 'arquillian-persistence-datasets.jar'. Would it work for you?	1	1	https://issues.jboss.org/browse/ARQ-1572
2093	ARQ-1572	So the conclusion will be: * JAR and WAR will have it flattened (included in the test archive), so that JRebel works out of the box * For EAR we have {{arquillian-persistence-datasets.jar}} as regular (not exploded) archive in the {{lib/}} folder in the root of the archive. Does it sound good enough?	1	1	https://issues.jboss.org/browse/ARQ-1572
2094	ARQ-1567	[~smikloso], would it be possible to check system property java.home as well? While this might point to JRE in some cases, it might make setup easier.	0	0	https://issues.jboss.org/browse/ARQ-1567
2095	ARQ-1555	Actually, think I found a way to leverage @ArquillianResource. Since it uses a class, the class provided could be the @ApplicationPath annotated class. If you did @ArquillianResource(JaxRSActivator.class) MyInterface restClient; REST extension could look for that application path (as found on that class) and leverage that as the base URI. If your JAXRSActivator was annotated @ApplicationPath("/rest") it would use rest. Otherwise it should default to /, since this is the default behavior of JAX-RS. WDYT?	0	0	https://issues.jboss.org/browse/ARQ-1555
2096	ARQ-1555	Are we still thinking this is not easily fixable to use ArquillianResource? Seems like it shouldn't be that difficult. What am I missing?	0	0	https://issues.jboss.org/browse/ARQ-1555
2097	ARQ-1461	How does this effect arquillian-extension-ios?	1	1	https://issues.jboss.org/browse/ARQ-1461
2098	ARQ-1453	Could you define 'Drone-like'?	0	0	https://issues.jboss.org/browse/ARQ-1453
2099	ARQ-1444	I have created a pull request for this issue [1]. Can anyone review it? Thanks in advance. [1] https://github.com/arquillian/arquillian-extension-persistence/pull/38	1	1	https://issues.jboss.org/browse/ARQ-1444
2100	ARQ-1444	Hi guys, is this issue really resolved? im trying to clean a oracle db using @CleanupUsingScript or @ApplyScriptAfter/Before and still having ORA-00911: invalid character in Alpha. My scripts are simple DELETE from TABLE; my arquillian.xml is as follows: <extension qualifier="persistence-dbunit"> <property name="schema">TR_DETRAN_RED_AD</property> <!-- <property name="datatypeFactory">org.dbunit.ext.oracle.Oracle10DataTypeFactory</property>--> <property name="filterTables">true</property> <!--<property name="filterForeignKeysEnabled">false</property>--> <!--<property name="customTableFilter">org.jboss.arquillian.persistence.dbunit.filter.OracleDatabaseSequenceFilterProvider</property>--> </extension> <extension qualifier="persistence"> <property name="defaultDataSource">java:jboss/datasources/RedTesteDS</property> <!--<property name="defaultDataSeedStrategy">CLEAN_INSERT</property> --> <!-- <property name="defaultCleanupPhase">NONE</property> --> <property name="dumpData">false</property> </extension> <!-- <extension qualifier="persistence-script"> <property name="sqlStatementDelimiter">;</property> </extension> --> thanks in advance.	1	1	https://issues.jboss.org/browse/ARQ-1444
2101	ARQ-1419	Yes, it's HTTP only. With regard to the descriptor deployment - the only relevant documentation I found is this: https://docs.jboss.org/author/display/ARQ/Descriptor+deployment. However I checked the AS7/WildFly container adapter and it seems it does not support this either. So I think we can live without this feature for now and add the support later. WDYT?	1	1	https://issues.jboss.org/browse/ARQ-1419
2102	ARQ-1388	Do you get the same even if you don't cache the ear and the jar between tests?	1	0	https://issues.jboss.org/browse/ARQ-1388

#	Requirement	Questions	Answered	Successfully Answered	URL
2103	ARQ-1387	<weld.extensions.version>1.1.Final-redhat-1</weld.extensions.version> <hibernate.version>4.1.3.Final-redhat-1</hibernate.version> <jboss.bom.version>1.0.0.M12-redhat-1</jboss.bom.version> <seam3.version>3.1.0.Beta3</seam3.version> <jsf.version>2.1.7-redhat-1</jsf.version> <arquillian.version>1.0.2.Final</arquillian.version> <selenium.version>2.24.1</selenium.version> <mongo.driver.version>2.9.0</mongo.driver.version> <morphia.version>0.99</morphia.version> <sonarVersion>2.2</sonarVersion> <uel.version>2.1.2-b04</uel.version> <resteasy.version>2.3.3.Final</resteasy.version> <infinispan.version>5.1.4.FINAL-redhat-1</infinispan.version> <maven.build.timestamp.format>MMMM dd, yyyy 'at' hh:mm aa</maven.build.timestamp.format> <buildNumber>${maven.build.timestamp}</buildNumber> <deployFarmName>dev</deployFarmName> <version.org.jboss.arquillian.graphene>2.0.0.Alpha1</version.org.jboss.arquillian.graphene> <metrics.version>3.0.0-BETA2</metrics.version> <pentaho.version>4.4.0-stable</pentaho.version> <jacoco.version>0.6.4.201312101107</jacoco.version> Above is all the versions for different artifact but I don't see one defined for ASM. What version we should use for that?	1	0	https://issues.jboss.org/browse/ARQ-1387
2104	ARQ-1387	We actually exclude the asm in the pom. Should we include our own?	1	0	https://issues.jboss.org/browse/ARQ-1387
2105	ARQ-1385	Is this a default behavior you need? Or you want to have some value being the default and configurate via webdriver extension?	1	1	https://issues.jboss.org/browse/ARQ-1385
2106	ARQ-1385	[~lfryc] What's the proper SPI lifecycle event in your opinion?	0	0	https://issues.jboss.org/browse/ARQ-1385
2107	ARQ-1384	Thank you very much for your effort Tamás. Would you be so kind and create pull request out of it?	0	0	https://issues.jboss.org/browse/ARQ-1384
2108	ARQ-1370	Sure, you're welcome. Regarding your proposed solution (referenced in the description): without admittedly knowing Warp's internals, this looks more like a workaround than a solution. The root cause of this problem seems to be that {{Alpha3}} requires a dedicated custom SSL context (or TrustManager, respectively) while {{Alpha2}} was able to make use of the container's existing SSL context as described in ARQ-1365 and the forum thread (referenced above). The question is: why is that and would it be possible to let {{CommandService}} reuse the existing SSL context as well? Just ignore my comment if I'm totally mistaken ;-)	1	0	https://issues.jboss.org/browse/ARQ-1370
2109	ARQ-1370	Great, thanks for the doc pointers! The Arquillian GitHub wiki at [2] would have been the last place for me to look for such Warp details ;-) {quote} I actually think that any SSL communication wouldn't work in Warp in Alpha2, beucase Warp is a Man in the Middle from the SSL handshake perspective {quote} Well, as described earlier, it does - at least for client-side tests in a test class annotated with {{@WarpTest}}. That's why this issue seemed like a Alpha2->Alpha3 regression from a user's (my) perspective: for some reason Warp-enriched _client-side_ tests do fail with Alpha3 while they work with Alpha2 (via SSL). You know better than me whether that kind of client-side tests is also affected by the new {{CommandService}} infrastructure/workflow. I guess the major difference compared to Alpha2 is the _additional_ request in Alpha3 you describe: while Alpha2 (enriched) client-side tests, issuing requests via Drone/WebDriver, simply run in JBoss's configured SSL context, Alpha3 (enriched) client-side tests involve {{CommandService}} creating a new request _itself_, somehow outside the container's SSL context...? If my assessment is correct the question should be, can you tap into the container's SSL context or adopt it somehow? If not, you most likely need to establish your own, as suggested by your proposed workaround, if Warp's architecture itself can't be changed/improved in that regard. IMHO, such a workaround would be less optimal, however, because a functional JSF test naturally involves the web app's container, including potential security measures in place (like SSL or authentication). Thus it would be nice to run such tests in their actual environment, without the need for a parallel security setup that (potentially) needs to be kept in sync with the container's setup in order to produce meaningful test results (without superfluous artefacts). After all, this is what makes Arquillian such a great tool, right? What do you reckon? Seems like an interesting problem :-)	1	0	https://issues.jboss.org/browse/ARQ-1370
2110	ARQ-1370	Hey Oliver, I agree that Warp should be able to work with SSL. {quote} Well, as described earlier, it does - at least for client-side tests in a test class annotated with @WarpTest. That's why this issue seemed like a Alpha2->Alpha3 regression from a user's (my) perspective: for some reason Warp-enriched client-side tests do fail with Alpha3 while they work with Alpha2 (via SSL). {quote} The thing is that Warp needs to be able to intercept the request flowing through proxy from client to server - it basically needs to know request/response headers (in order to be able select particular request which should be intercepted). I'm not sure whether LittleProxy current supports SSL handshake - it seems so, however as I said, we are not interested just in transporation level (HTTP-over-SSL), but rather in presentation layer (details of HTTP request). It seems that LittleProxy team recently added support for SSL interception (Man in the Middle - MITM) [1] to 0.6.0-SNAPSHOT (Warp currently uses 0.5.3). {quote} I guess the major difference compared to Alpha2 is the additional request in Alpha3 you describe: while Alpha2 (enriched) client-side tests, issuing requests via Drone/WebDriver, simply run in JBoss's configured SSL context, Alpha3 (enriched) client-side tests involve CommandService creating a new request itself, somehow outside the container's SSL context...? If my assessment is correct the question should be, can you tap into the container's SSL context or adopt it somehow? If not, you most likely need to establish your own, as suggested by your proposed workaround, if Warp's architecture itself can't be changed/improved in that regard. {quote} That's interesting, what have happened in Alpha2: {code} client --HTTPS--> proxy --HTTPS with payload--> server {code} * on the proxy, we had to have access to HTTP headers * in order to add additional request header (which contains payload) How it works in Alpha3: {code} client --HTTPS--> proxy --HTTPS with flag--> server proxy --CommandService--> server {code} * we still need to have access to request headers (adding a flag (serialId of the message) * but proxy needs to be able to connect independently to server You are talking about SSL contexts, but I think it actually stands for SSL certificates setup for {{client}} (browser?) and {{proxy}} (test JVM) and in Alpha3, the {{test}} itself (CommandService, test JVM). Any chance you could look at this issue more closely? I can provide you with guidance. [1] https://github.com/adamfisk/LittleProxy/commit/8cb5b470ad3368badc15d6da67f8c8ffe6fe94f6	1	1	https://issues.jboss.org/browse/ARQ-1370
2111	ARQ-1370	Hi Lukáš, Thanks for getting back to this! How do you want me to help? Would it be sufficient to somehow replace the current version of LittleProxy with 0.6.0-SNAPSHOT and see if it works. I hope I can find some time over the weekend but it's somewhat hard to come by right now.	1	0	https://issues.jboss.org/browse/ARQ-1370
2112	ARQ-1364	Karel, could you please review this pull request?	0	0	https://issues.jboss.org/browse/ARQ-1364
2113	ARQ-1363	Karel. could you please review this pull request?	0	0	https://issues.jboss.org/browse/ARQ-1363
2114	ARQ-1362	Which Arquillian Protocol are you using? Servlet or JMX(default) ?	1	0	https://issues.jboss.org/browse/ARQ-1362

#	Requirement	Questions	Answered	Successfully Answered	URL
2115	ARQ-1341	As Aslak [pointed out|https://community.jboss.org/message/809204#809204], the REST extension would be better fit for separate module: {quote} We could put this in the arquillian-extension-rest. What's there now is just an early prototype with very little movement. We could 'move' the master branch to a branch, replace master with warp-rest, then retro fit the old master branch if needed. {quote} What do you think, [~jmnarloch]?	1	1	https://issues.jboss.org/browse/ARQ-1341
2116	ARQ-1341	Ok, thanks. Quick question since we are moving the extension to repository *arquillian-extension-rest* how do you want me to name the packages? At the moment I had used *arquillian-warp-jaxrs* prefix.	1	1	https://issues.jboss.org/browse/ARQ-1341
2117	ARQ-1341	org.jboss.arquillian.extension.rest.warp ? org.jboss.arquillian.extension.rest.warp.jaxrs ?	1	1	https://issues.jboss.org/browse/ARQ-1341
2118	ARQ-1330	I'll admit that there might be a few edge cases where it makes sense to use a behavior-driven, mock-style approach. But I still hesitate to add a mock framework to JSFUnit. This is partially because I don't often see those theoretical edge cases. But mostly it's because a mock approach is a completely different style of testing. I'm not saying it's invalid, but it causes you to think about your tests in a completely different way. The classic paper on this is by Martin Fowler. If you haven't read it, I'd highly recommend it. It looks like he updated it since I last studied it so maybe I need to re-read it as well: http://martinfowler.com/articles/mocksArentStubs.html To address the use cases the two of you have presented thus far: * For the one about analyzing generated html, I think this is solved with the speed and ease of Arquillian. If I understand your requirements correctly, I'd argue that it's easier to set up a JSFUnit test for this than to set up a mock environment. * For testing if an attribute is set only once per phase, I'm not sure I understand why you need that. What is the actual use case? Can you use a java assert statement to achieve the same thing? * For testing exceptions, that's pretty straightforward. You just create the conditions that throw the exception and see what happens. If that can't be done with a live application then you need to question your design because perhaps you are testing something that can't happen in the real world. Some of the above use cases might stem from the fact that JSFUnit tests aren't fine-grained enough. You are required to wait until the lifecycle completes and there is no opportunity to peek at the live object until rendering is done. To make fine-grained tests easier, I've thought of implementing JSFUNIT-59. This feature would call back into your tests before and after each phase. Then you could examine the live object at different points during the JSF lifecycle. You could even call its methods from your test, but I'd discourage that most of the time. The only tricky part to implementing JSFUNIT-59 would be to figure out how to register the callback with a PhaseListener. But it's quite doable. It would be great if one of you wants to take that on. I think it would make some of the above use cases even easier.	1	1	https://issues.jboss.org/browse/ARQ-1330
2119	ARQ-1319	Hey Aris, great job with investigations! The initial intention was to verify that Warp-based tests coverage is reported correctly. Configuring Warp's code coverage is actually nice consequence and I will merge it once it will work correctly. The problem here is that even in the ant-generated reporting, I can't see coverage of e.g. TestingServlet ({{ftest}} project, {{org.jboss.arquillian.warp.ftest}} package), which should be actually hit several times - this is what fundamentally Warp users wants. I guess there is some problem with reporting from Warp-based tests, e.g. does the command service does not work correctly? Does it provide correct data?	1	1	https://issues.jboss.org/browse/ARQ-1319
2120	ARQ-1319	It is possible to decouple response's writing and closing from {{WarpFilter}} {{doFilterWarp}} method, but this means that I should have to do the following changes: * Make {{RequestContextHandler}} observe {{BeforeSuite}} & {{AfterSuite}} instead of {{BeforeRequest}} & {{AfterRequest}} events. This way requestContext will be active throughout Arquillian's event lifecycle. * Add two observer methods on {{HttpRequestProcessor}}. The first one will observe {{BeforeSuite}} event (with precedence 100) and replace {{HttpServletResponse}} with a {{NonWritingResponse}}. The second one will observe {{AfterSuite}} event (with precedence -100) and will write and close the {{ServletOutputStream}}. Do you think that this is a good approach?	1	1	https://issues.jboss.org/browse/ARQ-1319
2121	ARQ-1319	Interesting approach. Do you plan on starting the implementation soon, or would you like me to give it a try? I believe it might take me some time, since I cannot commit to it full-time. ARQ-934 seems like a good starting point	1	1	https://issues.jboss.org/browse/ARQ-1319
2122	ARQ-1315	Fabian, do you have a proposal how it should look like in the configuration file?	1	1	https://issues.jboss.org/browse/ARQ-1315
2123	ARQ-1313	[~lfryc] [~jpapouse] isn't this done via http://search.maven.org/#artifactdetails%7Corg.jboss.arquillian.extension%7Carquillian-phantom-driver%7C1.1.0.Beta1%7Cjar ? And who actually manages arquillian-phantom-driver now?	1	0	https://issues.jboss.org/browse/ARQ-1313
2124	ARQ-1307	Okay, I will release them. [~jpapouse], what do you think about moving those projects to https://github.com/qa/ repo?	1	1	https://issues.jboss.org/browse/ARQ-1307
2125	ARQ-1307	Jan, I find the configuration quite cumbersome for end-user (all the profiles which need to setup). ---- I did a proof of concept that we could resolve the artifact with ShrinkWrap Resolver for Maven: {code:java} @Test public void test() { Maven.resolver().resolve("org.jboss.arquillian.phantomjs:phantomjs-binary:jar:unix_amd64:1.8.1") .withoutTransitivity().as(JavaArchive.class); } {code} With {{org.jboss.shrinkwrap.resolver:shrinkwrap-resolver-depchain:2.0.0-beta-2}} I was able to get local resolution in ~0.5sec. I consider it quite fast, but we should still leverage local binary cache as you did already. ---- Could you re-implement it using the technique mentioned above?	0	0	https://issues.jboss.org/browse/ARQ-1307
2126	ARQ-1303	Ok, But in such case, how can I use that in my test? I would have to 'lookup' the bean that is being test each time from the application context. You won't be able to inject you configured beans into the the tests directly. I was thinking about different aproach. How about resigning from class level annotations in favor of 'factory method'. Something like: {code} public class TestCase { @SpringApplicationContext public static ApplicationContext createApplicationContext() { return ... // the application context could be created fully programticly } public void test() { } } {code} In such way we could just @Autowire the beans in to the test class afterwards. Also we would have access to the ApplicationContext by simply adding @Autowired ApplicationContext field. It would work out of the box (likely this feature has been added somewhere in Spring 2.5.x version). The other thing I was thinking was using the same convention as in Spring Test 3.1. You are able to define static inner class annotated with @Configuration and define your beans there. Something like: {code} @SpringAnnotationConfiguration public class TestCase { @Configuration static class ContextConfiguration { @Bean public SomeBean create() { } } public void test() { } } {code}	1	1	https://issues.jboss.org/browse/ARQ-1303
2127	ARQ-1303	Does this make sense?	1	1	https://issues.jboss.org/browse/ARQ-1303

#	Requirement	Questions	Answered	Successfully Answered	URL
2128	ARQ-1303	You don't need to lookup the bean, we can provide a feature to autowire beans in a test class by using reflection. please take a look at AbstractSpringInjectionEnricher.java registerBean method. in my point of view, it might be better to keep away the developer from complexity of creating spring context and let Arquillian take care of it so having a prepared context and controlling it by annotation is more convenient than creating it. second approach sounds cool but you only will be able to use it in Spring 3.x and it does not work with static application context. another problem that thanks to Arquillian it is solved now is to use your test context in server. it is really cumbersome specially if you are using Grizzly (jersey) but it is not our concern in this case. WDYT ?	1	0	https://issues.jboss.org/browse/ARQ-1303
2129	ARQ-1292	Hi Aslak and Andrew This looks very interesting. Shall I do some investigations?	0	0	https://issues.jboss.org/browse/ARQ-1292
2130	ARQ-1269	Hey, can I take this over? I have a pull request ready to go.	1	1	https://issues.jboss.org/browse/ARQ-1269
2131	ARQ-1262	Hey Pavol, could you please setup browser-compatibility tests for Warp? They could use CloudBees/OpenShift and SauceLabs (I can provide all required accounts). This issue does not have high priority for Beta1.	0	0	https://issues.jboss.org/browse/ARQ-1262
2132	ARQ-1261	Hey Pavol, could you please setup container-compatibility tests for Warp? They could use CloudBees and mock-browser (htmlUnit now, phantomjs if available). This issue has high priority for Beta1.	0	0	https://issues.jboss.org/browse/ARQ-1261
2133	ARQ-1260	[~lfryc] how you want to figure out what is the last deployment? Should you mark it somehow on the filesystem?	1	1	https://issues.jboss.org/browse/ARQ-1260
2134	ARQ-1247	Hey John, could you please update to latest master and try full build? Tests seem to work fine on continuous integration: https://buildhive.cloudbees.com/view/My%20Repositories/job/arquillian/job/arquillian-extension-warp/ The latest Warp builds were published to Maven repo today.	0	0	https://issues.jboss.org/browse/ARQ-1247
2135	ARQ-1237	Hey Bartosz, do you have difficulties with implementing it? Any insight you would like to share?	1	1	https://issues.jboss.org/browse/ARQ-1237
2136	ARQ-1216	Most important - you should avoid breaking what currently works on developers/testers. This will ease migration to the new way. Everything old can be marked as deprecated, and only drive new features in to the new setup. *Do not break the old setup.* The first issue that needs to be addressed is why you can only have one container on the classpath at a time. I'm assuming there is some marker interface that Arquillian uses to determine this is a container (probably the local implementation of LoadableExtension that's in every container). Once that's done, couldn't you start by adding some new tags to arquillian.xml, allowing the selection of which container to run and set each of these with an ID? Then setup some kind of JVM argument or something indicating which arquillian test set to run based on these IDs? Then make sure it allows you to run multiple (either serial or parallel). Maybe.. mvn test -DarquillianContainer=weld_ee -DarquillianContainer=jbossas_managed_7 I would almost argue that these need to be two different issues, one covering containers and the other covering extensions (though they are essentially both extensions...)	1	0	https://issues.jboss.org/browse/ARQ-1216
2137	ARQ-1213	Do we still consider this issue relevant?	1	0	https://issues.jboss.org/browse/ARQ-1213
2138	ARQ-1212	Is this issue still relevant?	0	0	https://issues.jboss.org/browse/ARQ-1212
2139	ARQ-1176	Karel, would you prefer to have this implemented in Arquillian Drone?	0	0	https://issues.jboss.org/browse/ARQ-1176
2140	ARQ-1176	+1 reject it wdyt [~kpiwko]?	1	0	https://issues.jboss.org/browse/ARQ-1176
2141	ARQ-1169	is this related to ARQ-483 ?	0	0	https://issues.jboss.org/browse/ARQ-1169
2142	ARQ-1169	Seems related. But Lincoln's POC there handles simply the embedded case. Sounds like 483 would really be incorporated by this, which would aim to provide a target runtime. And could have a remote, managed, and embedded container to accompany it. What do you say we link this as "incorporates" 483?	0	0	https://issues.jboss.org/browse/ARQ-1169
2143	ARQ-1166	[10:53] <aslak> http://pastebin.com/bQP8rJiS line:50 [10:53] <aslak> the original file name is linked in. the Asset might have a different name in the Archive then on the Filesystem [10:54] <aslak> in this case we're setting up rebel.xml to monitor test-web.xml, but does rebel know that that is our web.xml ? [10:55] == jharting [jharting@redhat/jboss/jharting] has joined #jbosstesting [10:56] <aslak> most of those shouldn't be linked under "web". [10:57] <blabno> aslak, sorry, but i don't understand [10:57] == kpiwko [kpiwko@redhat/jboss/kpiwko] has joined #jbosstesting [10:57] <blabno> i think jrebel will not know that test-web.xml is web.xml in archive [10:58] <aslak> but it should ? [10:58] <blabno> well, would be nice for use [10:58] <blabno> for us, but I don't think they have forseen this [10:58] <blabno> and I would not expet them to support this [10:59] <aslak> blabno, mm, I can't seem to find any 'rename/redirect' type element [11:00] <aslak> we basically want to say: as / in web, link from src/test/resources/ file test-web.xml as web.xml [11:00] <blabno> i think that we will need to copy all resources that have different name in archive then in filesystem to jrebel temp [11:01] <aslak> mm yea, rebel should still be able to detect the changes when you run so. the user shouldn't notice [11:02] <aslak> i'll add it to my 'jrebel-feature-request' list [11:02] <blabno> :) yes and let's gather bunch of people to vote for it :D	1	0	https://issues.jboss.org/browse/ARQ-1166
2144	ARQ-1143	Hi Bartosz, I tried to update the runner to use Groovy 2.1 and Spock 0.7, but for some reason I haven't figured out yet AuxiliaryArchiveAppender is not picking groovy classes into arquillian-spock.jar even if groovy-all is on the classpath no matter of the scope. One of the workarounds I came with was to include groovy-all and 0.7-spock-${groovy-version} fetched by ShrinkWrapMavenResolver based on what groovy/spock version user specifier/transitively depends on in his pom.xml. How does it sound to you?	1	1	https://issues.jboss.org/browse/ARQ-1143
2145	ARQ-1096	Do you suggest supporting both these variants or full migration to <default-container>foo</default-container>?	0	0	https://issues.jboss.org/browse/ARQ-1096
2146	ARQ-1096	Using default=true + arquillian.launch with Maven Resource Filtering works. Does not seem nice but still works with current infrastructure. Do we want to change the behavior somehow in future?	0	0	https://issues.jboss.org/browse/ARQ-1096
2147	ARQ-1089	I don't think this is actually critical and I would leave the fix for future. I can assess it again if you will provide more info about use case? Is that integration issue?	0	0	https://issues.jboss.org/browse/ARQ-1089
2148	ARQ-1082	This is awesome! Does that mean that we could use persistence on the client side?	0	0	https://issues.jboss.org/browse/ARQ-1082
2149	ARQ-1077	Guys, I think this issue can be fixed if: * Drone extension could be enabled in IN_CONTAINER mode or * Persistence extension could be enabled in AS_CLIENT mode or * Test run in IN_CONTAINER mode, however Drone tests are forced to run AS_CLIENT So, I think following scenario should work without any code modifications required: {code} @RunWith(Arquillian.class) @DataSource("java:jboss/datasources/ExampleDS") @UsingDataSet("OneUserAccount-Admin.yml") public void CombinedTest { // notice testable=true, hence run in container by default @Deployment(testable=true) public static Archive<?> deployment() { // return full archive for Drone here } @Test @RunAsClient public void droneTest(@Drone WebDriver driver) { driver.XYZ(); } {code} Bryan, can you try this approach and let us know?	0	0	https://issues.jboss.org/browse/ARQ-1077

#	Requirement	Questions	Answered	Successfully Answered	URL
2150	ARQ-1077	Bryan, At least in my case I"m getting failures due to graphene not being in the archive. Did you have to add anything to your WAR file for that?	0	0	https://issues.jboss.org/browse/ARQ-1077
2151	ARQ-985	Now corect me if I'm wrong, but will this have any chance to work with embedded container? Assuming that the XmlClientApplicationContextProducer will create the instance of TestScopeApplicationContext and then the XmlRemoteApplicationContextProducer will be run and create in CDI context new TestScopeApplicationContext, will that completly overide the previouse one?	1	1	https://issues.jboss.org/browse/ARQ-985
2152	ARQ-985	Could I try to define different CDI scopes for the created TestScopeApplicationContext to make that working with the embedded containers?	1	1	https://issues.jboss.org/browse/ARQ-985
2153	ARQ-970	Where can these tests be found?	0	0	https://issues.jboss.org/browse/ARQ-970
2154	ARQ-967	images, scripts, css, etc. <lfryc> that's the time when we expect AJAX request, but request for those evil resources come :-) <lfryc> so it's time to define smart filters which decides what request to enrich <aslak> client multi threads the resource request ? <lfryc> but it can't be written generally enough, since you can't simply map request URL to content-type or so <lfryc> yeah, browsers usually loads up to 5 resources at time <lfryc> it is not issue with htmlunit (that's why unit tests pass just fine) <lfryc> but real browsers have cause issues <aslak> lfryc, you can't look at the accept header ? <lfryc> aslak: it's not silver-bullet <lfryc> browser may not need to send that header <aslak> true <lfryc> I think we have two options: <lfryc> 1) smart heuristics <lfryc> 2) framework-dependent filters <aslak> you have control over the proxy, so you should be able to sync up at that level how ever you want it <lfryc> you mean I can force client to send accept-header? :-) <aslak> lfryc, hehe no, as in race conditions i mean <lfryc> race condition - I mean, request fighting in what will be enriched :-) <aslak> lfryc, you could in theory throttle the request <aslak> aa ok <aslak> there should be multiple types of filters.. accept headers, url pattern etc <lfryc> hmm but browsers usually sents correct accept-headers for images/icons/etc. <lfryc> for resource types they know what to expect <aslak> depending on the use case, not all will work. if browsers don't sent accept headers, that filter migh tbe more useful in a REST request <lfryc> yes, we could provide multiple types of filters <lfryc> thinkinga bout framework-specific and smart ones <lfryc> user may decide which to use for given use case <aslak> yea <lfryc> use the most appropriate one for the use case <lfryc> and choose/configure default one in arquillian.xml <aslak> i think you can provide in warp core http based filters * jharting has quit (Quit: Leaving.) <aslak> i guess that would be, request content, url and headers * alesj has quit (Quit: Leaving.) <lfryc> that's going to be huge :-) <lfryc> and then framework specific filters, which can be easier to use <aslak> where does arq.xml come in? * rmannibucau has quit (Remote host closed the connection) * ge0ffrey has quit (Ping timeout: 240 seconds) <lfryc> configure defaults <lfryc> like, I want to use JSF filter with FacesServlet mapping on /xyz <aslak> default what? fitler ? <lfryc> yep <lfryc> when no filter specified, default can be used <aslak> arn't filters on a .execute() level ? <lfryc> yep, they are <lfryc> oh, no, they are on the proxy level * DavideD has quit (Ping timeout: 245 seconds) <aslak> ok, so basically via arq.xml you can say that only ever enrich jsf requests in this suite, what ever i might be calling <lfryc> which is on client-side <lfryc> but evaluated after client action <aslak> lfryc, sure, but i mean you don't need global filters. the individual .execute methods would define them <aslak> (not saying you can't have global) <lfryc> yes, but main use case for Warp is verifying request pages <aslak> eh, to rephrase.. a Filter is active from .execute start til end <lfryc> not resources <lfryc> so global might be handy <lfryc> but still you can define per Warp execution * #jbosstesting :End of Channel Ban List <aslak> and it work now based on just selecting the first request within the execute block <lfryc> yep <lfryc> first come, first served <aslak> so unless you wanted to enrich resources, you could really just split up the requests in two executes <lfryc> yep * qmx is now known as qmx|away <aslak> HTTP filters should work with Resources and AJax request, by either headers or url <aslak> a JSF filter would basically be a URL filter no? <lfryc> yes, it will be URL pattern filter <aslak> yea <aslak> what was the issue ? <aslak> :) <aslak> known how to filter? <aslak> or the browser doing multi threaded requests that potentially create multiple enriched requests ? <lfryc> know what reuqest to filter :-) <lfryc> know what to filter because browser may request pages/resources in parallel <aslak> does the parallel matter ? <lfryc> it matters, because without filtering, you can't choose what request to enrich/verify <lfryc> so you can either enrich first or enrich all * ge0ffrey (~ge0ffrey@174.195-241-81.adsl-dyn.isp.belgacom.be) has joined #jbosstesting * ge0ffrey has quit (Changing host) * ge0ffrey (~ge0ffrey@redhat/jboss/ge0ffrey) has joined #jbosstesting * balunasj is now known as balunasj|lunch <lfryc> when enriching all requests, you can always wait	1	1	https://issues.jboss.org/browse/ARQ-967
2155	ARQ-967	What's missing: * user-centric API * relation between filters and requests * integrity checks ** a filter is used ** how many requests user expects (at overall / per each filter) * standard set of common filters (TBD) * framework-specific filters * default filters configuration / SPI * groups of filters	0	0	https://issues.jboss.org/browse/ARQ-967
2156	ARQ-957	Guys, could you please review and comment what needs to be improved to stay aligned with other demos?	0	0	https://issues.jboss.org/browse/ARQ-957
2157	ARQ-932	We could solve the issue by forcing client to use only APIs of client tools (like selenium-api) and package them to the WAR. Is it possible to make anonymous classes static? -> No http://stackoverflow.com/questions/758570/is-it-possible-to-make-anonymous-inner-classes-in-java-static	1	1	https://issues.jboss.org/browse/ARQ-932
2158	ARQ-932	Would it be possible to rewrite the Anonymous Inner Class to a Static class on 'deploy', assuming of course the Inner class doesn't reference any object instances from the parent. Using e.g. javasist or similar ?	1	1	https://issues.jboss.org/browse/ARQ-932
2159	ARQ-928	What's this all about? :)	0	0	https://issues.jboss.org/browse/ARQ-928
2160	ARQ-924	What did you have in mind? I'm thinking of something similar, allowing you to build tests that run against servicemix/jboss fuse using the OSGi runtime of karaf.	1	1	https://issues.jboss.org/browse/ARQ-924
2161	ARQ-924	I'm curious how this works, or maybe we just need to wait for some docs to come out. Can I point to KARAF_HOME to run tests, and the container will start up for me?	0	0	https://issues.jboss.org/browse/ARQ-924

#	Requirement	Questions	Answered	Successfully Answered	URL
2162	ARQ-918	The following suggestions indicate to me I need to clarify what behavior we're attempting to control. * testMode = TestMode.CLIENT etc... * runMode = ArquillianRunMode.IN_CONTAINER|CLIENT|CLIENT_ENRICHED * Do the 'supports' and 'testing' bookends reinforce the meaning of 'inContainer' that much? A deployment is "in container" by definition. We aren't trying to control whether or not it's deployed to the container. Otherwise, we wouldn't be defining a deployment in the first place. What we are trying to control is whether Arquillian adds its own infrastructure to the deployment (effectively repackaging it). That infrastructure adds extra services to the deployment that become available once the deployment is deployed. Those services are (as of now): * a service to run tests in-container * a service to monitor, analyze or execute callbacks on the server-side code We are not trying to control how the tests are run. That behavior is controlled on the test class or test method using the @RunAsClient annotation. When absent, the default is to run the test "in-container" if the deployment supports in-container testing or "as client" if the deployment does not. (If there are multiple deployments, @OperateOnDeployment tells Arquillian which deployment to use for that test). I'll summarize the three ways a deployment is packaged and the behavior it provides: * Add the test infrastructure to the deployment - capable of running tests in-container (white box testing) * Add instrumentation to the deployment - supports client testing w/ knowledge of internal state (gray box testing) * Use the deployment exactly as it's defined - supports purely client testing (black box testing) With that in mind, I have another proposal for how to control the packaging of these services: {code} @Deployment(testingServices = ALL | INSTRUMENTATION_ONLY | NONE) {code} or as a separate annotation {code} @Deployment @TestingServices(NONE) {code} ALL could be IN_CONTAINER instead.	1	1	https://issues.jboss.org/browse/ARQ-918
2163	ARQ-918	Do we want to continue supporting "mixed-mode"? Currently, this is possible doing something like this: {code} @RunWith(Arquillian.class) public class ExampleIT { @Deployment public static WebArchive createDeployment() { // ... } @Test public void inContainerTest() { // ... } @Test @RunAsClient public void clientTest(@ArquillianResource URL baseURL) { // ... } } {code}	1	0	https://issues.jboss.org/browse/ARQ-918
2164	ARQ-917	In what container have you tested this? It being a field, won't it be initialized when accessing/loading the class and fail to execute shrinkwrap in a managed/remote container?	0	0	https://issues.jboss.org/browse/ARQ-917
2165	ARQ-912	Done -> Create a JIRA component to track issues and releases Component: "Extension - REST" Version: rest_1.0.0.next Done -> Create a new repository: arquillian-extension-rest https://github.com/arquillian/arquillian-extension-rest Done -> Add the repository to the enlistments on http://ohloh.net/p/arquillian	0	0	https://issues.jboss.org/browse/ARQ-912
2166	ARQ-898	Aha, now I see where the disconnect in my logic is. I suspected that the name attribute on the @Deployment annotation determined the middle segment (the deployment name), but I didn't realize that the first segment came from the @TargetsContainer annotation (the target name). I had thought that the target name came from the qualifier in arquillian.xml. If there is no @TargetsContainer annotation, is it possible to use the qualifier of the active configuration, if any? Or are these orthogonal things?	1	0	https://issues.jboss.org/browse/ARQ-898
2167	ARQ-892	Aslak, last time we've talked about this you mentioned that jrebel doesn't pick up some resources. Could you give me more detailed description? I've retested the pull request and everything seems to work.	0	0	https://issues.jboss.org/browse/ARQ-892
2168	ARQ-846	Hi, The status is resolved in 1.1.0.Beta1, where can I find this version please? Regards	1	0	https://issues.jboss.org/browse/ARQ-846
2169	ARQ-835	reset history and push with force ?	1	1	https://issues.jboss.org/browse/ARQ-835
2170	ARQ-809	Which source are you using? arquillian-testenricher-cdi should depend on Weld Core 1.1.5 which is in central	0	0	https://issues.jboss.org/browse/ARQ-809
2171	ARQ-798	Lol, duplicating my own. How come goggle didn't find it for 'Arquillian "+httpd"' ? :)	0	0	https://issues.jboss.org/browse/ARQ-798
2172	ARQ-796	Deleted the {{GlassFishRestDeploymentAppender}}. Aslak, I've not deleted the older SPI file referenced in the comment - the {{glassfish-embedded-3}} artifact is no longer built, since it is not registered as a Maven module in the project. Would it be wise to delete the {{glassfish-embedded-3}} and {{glassfish-remote-3}} directories, seeing that they're not built anymore ? I think we should test the current adapter against GF 3.0 (we test against 3.1 currently) and then proceed to delete the older unused ones. NB - The modules built in CI currently are {{glassfish-managed-3.1}}, {{glassfish-remote-3.1}}, {{glassfish-embedded-3.1}} and {{glassfish-common}}.	0	0	https://issues.jboss.org/browse/ARQ-796
2173	ARQ-782	Is this supported? I'm looking to use this feature and it doesn't look like my profile properties are being picked up, instead it reads ${var} as a string literal.	1	0	https://issues.jboss.org/browse/ARQ-782
2174	ARQ-733	Does it currently not do that? Or.. If so that's a bug	1	0	https://issues.jboss.org/browse/ARQ-733
2175	ARQ-710	Aslak: {quote} Not sure how to attack this.. in theory the stream that is wrapped in the managed container should be able to tell you which server it's from, so hench could be redirected somewhere, but not sure how to get 'client'. In AS4-6 CONSOLE was just another logger, i'm sure AS7 is similar. Wouldn't setting server1, server2 to log to specific file get you what you want? If no server logs to CONSOLE, then the assumption is the rest is 'client' ? {quote} I thought about that, but that would only separate by server. I believe it would be very convenient if separated by server _and_ test.	1	1	https://issues.jboss.org/browse/ARQ-710
2176	ARQ-710	David Lloyd about implementing Appender-like logging: {quote} (17:16:59) OndrejZizka: Hi David, where can I find some Appender class in JBoss logging? (17:17:09) OndrejZizka: in sources (17:19:27) dmlloyd: JBoss Logging has no appenders, it's just a log frontend (17:19:39) dmlloyd: unless you're talking about the older stuff from AS 5 (17:22:24) OndrejZizka: No, AS 7 (17:22:37) OndrejZizka: So, if I want some rolling appender used in AS 7, (17:22:51) OndrejZizka: how would I do it - like in log4j.properties? (17:24:43) OndrejZizka: logging.properties goes directly to the underlying framework or is somehow processed by jboss logging? (17:29:35) dmlloyd: the logging backend in AS7 is jboss-logmanager (17:29:44) dmlloyd: it's a java.util.logging.LogManager implementation (17:29:54) dmlloyd: but it also has extensive compatibility APIs (17:30:13) dmlloyd: jboss-logmanager is what processes logging.properties (17:43:30) OndrejZizka: I haven't worked with JUL much; if I wanted to implement an appender, and set it to be used in AS7, where would I start? (17:45:30) dmlloyd: if it's specifically for AS7, the best thing to do is implement org.jboss.logmanager.ExtHandler or one of its subclasses (17:48:19) OndrejZizka: Yes, it's for AS 7 + arquillian - to separate log lines per test case. {quote}	1	1	https://issues.jboss.org/browse/ARQ-710
2177	ARQ-690	What's left to do: Add {{-Djava.security.debug}} See http://docs.oracle.com/javase/1.5.0/docs/guide/plugin/developer_guide/debugger.html {code} -Djava.security.manager -Djava.security.policy==xxx.policy -Djava.security.debug=access,failure {code} To ignore the JDK's default policy which merges with the specified as default behavior, the policy path is prepended with '='.	0	0	https://issues.jboss.org/browse/ARQ-690

#	Requirement	Questions	Answered	Successfully Answered	URL
2178	ARQ-666	Implementation note: No warning are issue. Configuration is hidden from tests, so user will see no warnings based on implementationClass value. This leads to following situation: {code} @Drone HtmlUnitDriver driver; {code} {code:xml} <property name="implementationClass">...FirefoxDriver</property> {code} This will use HtmlUnitDriverFactory to get instance of FirefoxDriver. This is incorrect and will fail (because HtmlUnitDriverFactory has specific instantiation). Following must be specified in order to pass, however than implementation specific details are no triggered, e.g. no ChromeSwitches will be executed. {code} @Drone WebDriver driver; {code} Is that generic approach enough?	1	1	https://issues.jboss.org/browse/ARQ-666
2179	ARQ-653	Hey Aslak, would it be possible for you to review this pull request? Thanks!	0	0	https://issues.jboss.org/browse/ARQ-653
2180	ARQ-639	Which information are you looking for specifically ?	1	1	https://issues.jboss.org/browse/ARQ-639
2181	ARQ-638	I see there are some bits in WildFly about this functionality, but mainly around embedded container. Shall we keep the issue open? /cc [~aslak]	0	0	https://issues.jboss.org/browse/ARQ-638
2182	ARQ-633	The problem is, which module should it be in? This started as a archive.toURL extension that I wanted in ShrinkWrap, but it was refused at the time. We should probably revisit what it does and move it to ShrinkWrap.	1	0	https://issues.jboss.org/browse/ARQ-633
2183	ARQ-600	{code:java}org.jboss.arquillian.container.tomcat.managed_6.TomcatManagedConfiguration{code} contains a serverName attribute that does not appear in the other container configurations. Should I remove this or propagate it to the other managed container configurations?	1	1	https://issues.jboss.org/browse/ARQ-600
2184	ARQ-600	I've just had a go at adding tomcat-managed-8 to this source base. And now I can see that the remaining code in tomcat-managed-5.5 and tomcat-managed-6 is effectively identical - they both use "Servlet 2.5" protocol. tomcat-managed-7 and tomcat-managed-8 both use "Servlet 3.0" protocol and use the same (overridden) deploy/undeploy/list commands. ie. There's only two different flavours of tomcat-managed-N arquillian container integrations: * 5.5 and 6 * 7 and 8 I'm thinking of moving all the java code into tomcat-managed-common and just leaving the unit tests in the tomcat-managed-N modules and building tomcat-managed-N.jars that just contain the appropriate org.jboss.arquillian.core.spi.LoadableExtension file. What do you think?	1	1	https://issues.jboss.org/browse/ARQ-600
2185	ARQ-597	Is this container still needed? I'm hoping to drop support for it going forward per ARQ-1853. If there are any objections please comment on that issue.	0	0	https://issues.jboss.org/browse/ARQ-597
2186	ARQ-567	Hello Bernard, thank you for pinpointing the suite extension and your and aslaks work on that. I've tested extension and the suite alone is working properly, but we are using persistence extension (1.0.0.Alpha5) and then following exception occurs: java.lang.RuntimeException: Could not create new instance of class org.jboss.arquillian.test.impl.EventTestRunnerAdaptor at org.jboss.arquillian.persistence.core.deployment.PersistenceExtensionArchiveAppender.requiredLibraries(PersistenceExtensionArchiveAppender.java:81) Unfortunately the real exception is eaten and i can't debug within persistence extension. Thanks, Peter	1	0	https://issues.jboss.org/browse/ARQ-567
2187	ARQ-567	Hi, Is there any news on this issue? I find it pretty interesting as we have more than 60 test cases per app and a part from build time we need big infrastructure to handle them. How do you do that in other projects? Thanks in advance. Regards. Esteve	1	0	https://issues.jboss.org/browse/ARQ-567
2188	ARQ-567	Hi, Is there a way to run multiple tests with a single Deployment?	0	0	https://issues.jboss.org/browse/ARQ-567
2189	ARQ-561	Aslak, this workaround is useful only in a limited set of use cases. Here are the problems I've encountered: 1/ @Rule enricher is triggered before the Arquillian one...I was not able for example to do the following: {code} @Inject HelloService service1; @Inject HelloService service2; HelloService ruleService; @Rule public MethodRule dataProvider = new ParameterRule("ruleService", service1, service2); {code} It is clear the only way how to change that is to have a @Rule annotated method, which is not possible in JUnit currently 2/ Test method is reported as a single execution although it actually contains multiple tests 3/ MethodRule is deprecated Open question: Why do the sample implementation skip parameter enrichement while not running in container?	1	1	https://issues.jboss.org/browse/ARQ-561
2190	ARQ-561	> Open question: Why do the sample implementation skip parameter enrichement while not running in container? It would be executed both incontainer and on the client side if not one of them were skipped. We could have enhanced that with support for e.g. @RunAsClient. If we had executed the Client side multiple times for a Remote call, the Remote side would be invoked multiple times as well	1	1	https://issues.jboss.org/browse/ARQ-561
2191	ARQ-556	Hi guys. Is there any update to this feature request?	0	0	https://issues.jboss.org/browse/ARQ-556
2192	ARQ-521	Karel, could you use the right URL if it's available (configured or obtained from Container based events) and if not, use the default?	1	0	https://issues.jboss.org/browse/ARQ-521
2193	ARQ-517	[~alrubinger] [~aslak] Shall we pay this technical debt? Or is it already sorted out? :)	0	0	https://issues.jboss.org/browse/ARQ-517
2194	ARQ-474	I was thinking more in the line of @DependsOn("otherMethodName"). What do you mean by does not mean anything? Are you talking the annotation name "InOrder"?	0	0	https://issues.jboss.org/browse/ARQ-474
2195	ARQ-474	How about providing an annotation, lets say, @Ordered with which a whole test class could be annotated and then the test methods would be ordered according to the order in which they are written in the test class (from top->down). Do you think it would be possible?	0	0	https://issues.jboss.org/browse/ARQ-474
2196	ARQ-474	What about the non ordered ones? executed before or after? Default order is 0, then you can order as you like + or - ?	0	0	https://issues.jboss.org/browse/ARQ-474
2197	ARQ-467	Got any details here? pom, stack trace etc..	1	0	https://issues.jboss.org/browse/ARQ-467
2198	ARQ-448	I've made the changes locally to implement ARQ-448. I'm not so familiar with git. Can you please advise me on what I need to do to submit them? Thanks	1	1	https://issues.jboss.org/browse/ARQ-448
2199	ARQ-437	Lukas, If was thinking about this: You can't configure this via an annotation as during the BeforeSuite event you have no idea what files are in the suite and thus you can't go through source code creating appropriate browser instances. It is difficult to get this added to Arquillian runners, as JUnit has no {After,Before}Suite events. So only thing actually available in BeforeSuite event is configuration from arquillian.xml. So I can create a Drone Extension, which if activated, creates all instance in BeforeSuite and TestEnricher will later inject all field instances from SuiteScope context. If no browser instance is found, it will safely fail. This means that you will have possibility to use either BEFORE_SUITE + BEFORE_METHOD or BEFORE_CLASS + BEFORE_METHOD browsers, latter being the default. Do you find this usable? Proposed extension: {code:xml} <extension qualifier="drone-lifecycle"> <property name="browser-lifecycle">suite|class</property> </extension> {code} WDYT?	1	1	https://issues.jboss.org/browse/ARQ-437
2200	ARQ-437	One more important note: With SuiteScoped browsers you become dependent on class ordering. You can add a TestClass, thus changing order leading to unexpected failures. I'm not sure if class ordering is preserved on different platforms, afaik Surefire uses filesystem ordering as default. Are the advantages of this future bigger than drawbacks?	1	1	https://issues.jboss.org/browse/ARQ-437
2201	ARQ-436	Seems to me that new Arquillian Core is logging testng to console. Lukas, can you check Ajocado ftest example, which uses TestNG module if there is enough logging?	1	1	https://issues.jboss.org/browse/ARQ-436

#	Requirement	Questions	Answered	Successfully Answered	URL
2202	ARQ-435	Did I got right that idea of this is to make a screenshot and log actual state of the application (e.g. generated HTML) if test fails? Can you please elaborate on how this should be used? Do you want something like following? {code} @FailureLogging @RunWith(Arquillian.class) public class Test { ... } {code} and {code} @RunWith(Arquillian.class) public class Test { @Test @FailureLogging public void test() { ... } } {code}	0	0	https://issues.jboss.org/browse/ARQ-435
2203	ARQ-433	The tests pass from the old https://github.com/arquillian/arquillian.git repo. A notable difference is in the old repo the Servlet 2.5 protocol is being used, whereas in the new repo it's the Local protocol. The {{org.jboss.arquillian.protocol.servlet.v_2_5.ServletProtocolDeploymentPackager}} adds auxiliary archives such as {{arquillian-tomcat-testenrichers.jar}} [to the web application archive|http://community.jboss.org/message/551100]: {noformat} . ./META-INF ./META-INF/context.xml ./WEB-INF ./WEB-INF/beans.xml ./WEB-INF/classes ./WEB-INF/classes/org ./WEB-INF/classes/org/jboss ./WEB-INF/classes/org/jboss/arquillian ./WEB-INF/classes/org/jboss/arquillian/container ./WEB-INF/classes/org/jboss/arquillian/container/tomcat ./WEB-INF/classes/org/jboss/arquillian/container/tomcat/embedded_6 ./WEB-INF/classes/org/jboss/arquillian/container/tomcat/embedded_6/TestBean.class ./WEB-INF/classes/org/jboss/arquillian/container/tomcat/embedded_6/TestServlet.class ./WEB-INF/classes/org/jboss/arquillian/container/tomcat/embedded_6/TomcatEmbeddedInContainerTestCase.class ./WEB-INF/classes/org/jboss/arquillian/protocol ./WEB-INF/classes/org/jboss/arquillian/protocol/servlet ./WEB-INF/classes/org/jboss/arquillian/protocol/servlet/runner ./WEB-INF/classes/org/jboss/arquillian/protocol/servlet/runner/SecurityActions$1.class ./WEB-INF/classes/org/jboss/arquillian/protocol/servlet/runner/SecurityActions$GetTcclAction.class ./WEB-INF/classes/org/jboss/arquillian/protocol/servlet/runner/SecurityActions.class ./WEB-INF/classes/org/jboss/arquillian/protocol/servlet/runner/ServletTestRunner.class ./WEB-INF/lib ./WEB-INF/lib/arquillian-core.jar ./WEB-INF/lib/arquillian-junit.jar ./WEB-INF/lib/arquillian-tomcat-testenrichers.jar ./WEB-INF/lib/weld-servlet-1.0.1-Final.jar ./WEB-INF/web.xml {noformat} The {{org.jboss.arquillian.container.test.impl.client.protocol.local.LocalDeploymentPackager}} does not: {noformat} ./META-INF ./META-INF/context.xml ./WEB-INF ./WEB-INF/beans.xml ./WEB-INF/classes ./WEB-INF/classes/org ./WEB-INF/classes/org/jboss ./WEB-INF/classes/org/jboss/arquillian ./WEB-INF/classes/org/jboss/arquillian/container ./WEB-INF/classes/org/jboss/arquillian/container/tomcat ./WEB-INF/classes/org/jboss/arquillian/container/tomcat/embedded_6 ./WEB-INF/classes/org/jboss/arquillian/container/tomcat/embedded_6/TestBean.class ./WEB-INF/classes/org/jboss/arquillian/container/tomcat/embedded_6/TestServlet.class ./WEB-INF/classes/org/jboss/arquillian/container/tomcat/embedded_6/TomcatEmbeddedInContainerTestCase.class ./WEB-INF/lib ./WEB-INF/lib/weld-servlet-1.0.1-Final.jar ./WEB-INF/web.xml {noformat} Was the change to the Local protocol intentional, and if so any tips on how test enrichment should work with it? Also are there any immediate plans to switch Hudson to the new repo?	1	1	https://issues.jboss.org/browse/ARQ-433
2204	ARQ-433	Can we make this for 1.0.0.Final? this/next week? :)	0	0	https://issues.jboss.org/browse/ARQ-433
2205	ARQ-417	Are we still interested in this feature?	0	0	https://issues.jboss.org/browse/ARQ-417
2206	ARQ-414	Thanks for your input guys! Shall we move issues to GitHub so that they are closer to the code and use this one as an umbrella?	1	0	https://issues.jboss.org/browse/ARQ-414
2207	ARQ-397	I was interrested in testing it so I checkouted your branch and tryed to install it localy. First I had the following error : {quote}{color:red}'dependencies.dependency.version' for org.jboss.shrinkwrap:shrinkwrap-extension-glassfish:jar is missing. @ line 43, column 19{color}{quote} I just added the version to pom.xml (1.0.0-alpha-12) and now I cannot install because the "IntegrationWarTestCase" fails on : {quote} Failed tests: shouldBeAbleToInjectEJBAsInstanceVariable(org.jboss.arquillian.container.glassfish.embedded_3_1.app.IntegrationWarTestCase) {quote} It seems that the test case in error was not able to inject the referenced EJB : {code:java} @EJB private NoInterfaceEJB bean; @Test public void shouldBeAbleToInjectEJBAsInstanceVariable() throws Exception { Assert.assertNotNull("Verify that the Bean has been injected", bean); Assert.assertEquals("Arquillian", bean.getName()); } {code} Having no clue why this could happend, I just installed it skipping the tests, and tried the installed snapshot on my project. The sun-resources.xml is deployed correctly : {quote} INFO: command add-resources result: PlainTextActionReporterSUCCESSDescription: add-resources AdminCommandnull JDBC connection pool jdbc/wakwi-pool created successfully. JDBC resource jdbc/wakwi-ds created successfully. {quote} However ... I have the problem that the ejb is not injected in my test. My test code : {code:java} @RunWith(Arquillian.class) public class TestFileManagementServiceTest { @Deployment public static WebArchive createDeployment() { return ShrinkWrap.create(WebArchive.class, "testFileManagementService.war").addClasses(FileManagementService.class, File.class, Category.class, FileStore.class, CurrentFileStore.class, DefaultFileStore.class). addAsResource("test-persistence.xml", ArchivePaths.create("META-INF/persistence.xml")). addAsResource(EmptyAsset.INSTANCE, "WEB-INF/beans.xml"); } @EJB private FileManagementService service; @Test public void testSaveFile() { assertNotNull("Service should have been injected.", service); } } {code} Do you have the same problem injecting an EJB using the branch ARQ-397 ? Is it related ?	1	1	https://issues.jboss.org/browse/ARQ-397
2208	ARQ-395	I got this message. What does this mean? How to resolve it?	0	0	https://issues.jboss.org/browse/ARQ-395
2209	ARQ-395	Hi, I also get this error-msg: "DeploymentScenario contains targets not maching any defined Container in the registry". @Aslak --- re. "no container integration module on classpath"... I have this in my pom-file: <dependency> <groupId>org.jboss.arquillian.container</groupId> <artifactId>arquillian-jbossas-embedded-6</artifactId> <version>1.0.0.Alpha5</version> <scope>test</scope> </dependency> <dependency> <groupId>org.jboss.jbossas</groupId> <artifactId>jboss-as-depchain</artifactId> <version>6.0.0.Final</version> <type>pom</type> </dependency> Isnt that enough? Anything else I need to do to launch JBoss from Arquillian? (I do not have any @TargetsContainer annotation in my code)	1	1	https://issues.jboss.org/browse/ARQ-395
2210	ARQ-395	The dependencies looks right, it's basically the arquillian-jbossas-embedded-6 dep it's complaining missing. Do you have it setup in a Maven Profile? And if so, are you sure it's activated?	0	0	https://issues.jboss.org/browse/ARQ-395
2211	ARQ-359	"envy" is the server hostname? Can the server resolve "envy" to it self ? What is your configuration ? this sounds similar to the localhost ubuntu resolve issue..	0	0	https://issues.jboss.org/browse/ARQ-359

#	Requirement	Questions	Answered	Successfully Answered	URL
2212	ARQ-358	Is it still something we would like to implement [~aslak]?	1	0	https://issues.jboss.org/browse/ARQ-358
2213	ARQ-351	How is this related to ARQ-456? Does it mean that @Before annotated method can be used if whole test is marked as @RunAsClient?	0	0	https://issues.jboss.org/browse/ARQ-351
2214	ARQ-351	This issue just bit me now when trying to restructure part of WFLY testsuite. Aslak, where is this on the roadmap? How hard would it be for a contributor to fix this?	0	0	https://issues.jboss.org/browse/ARQ-351
2215	ARQ-348	Hi Aslak, is it possible to get byteman support to Arquilian for AS7. I was just informed in AS7-2530 that current implementation is in experimental phase. This thing seems to be needed for us. Thank you, Mirek	1	1	https://issues.jboss.org/browse/ARQ-348
2216	ARQ-347	Is this issue "dead"?	0	0	https://issues.jboss.org/browse/ARQ-347
2217	ARQ-337	<aslak> pbakker, we will create a common arq 'ContextController' you can inject, since all of this is non standard between cdi impls and containers.. with a common api exposed by Arq your tests can be portable between impls pbakker, if you want to help prototype that your more then welcome.. :) <pbakker> what else should the contextcontroller be able to do? because request and session are part of the test already <aslak> pbakker, session/req shouldn't be auto in the test case(maybe with some annotations but..). it just works by chance currently by chance between embedded vs remote container that is <pbakker> hehe ok i wasn't aware of that <aslak> it works basically because the remote test execution currently use HTTP to execute, so when in cdi, you will have req and session there.. but when you move to e.g. ejb protocol for execution then you only have request the scopes should be controlled by the test, not 'randoms' in the environment.. ;) <pbakker> ok than a ContextController makes a lot of sense so i guess step one would be to get it working consistently over types of containers step 2 to look at other cdi impl <aslak> pbakker, yea, start with weld-ee weld-se and jboss remote probably the easies let CDI manage the ContextController as a bean, and only use CDi to inject it as well <pbakker> i agree <aslak> possible we could do something like @Scopes(Request, Session) @Test as a 'around invoke style, but that is a abstraction over the ContextController that would be step 3/4 or so <pbakker> yeah that would be good i guess we just have to write different implementations for each container because different way of executing the tests if there is a real http request you probably want to run the test in that same request, not in a simulated one <aslak> i think one for weld based containers should be ok, since weld works the same in all, but need different impl between openweb beans / reisin etc resin <pbakker> but in a remote test it now uses a http request to execute the test right, and not in an embedded container (lots of assumptions here, haven't really looked at the code yet...) <aslak> hmm.. not quite sure. you might be able to reuse the one there, depends a bit on how strict the container is.. weld allow you to stop and start new requests within a reques etc.. as pete commented, don't do that if you want it to work.. hehe pbakker, correct <pbakker> lol, that sounds extremely hacky, but it might be a way around it and work glassfish-remote and jboss-remote in a similar way? <aslak> yes	1	0	https://issues.jboss.org/browse/ARQ-337
2218	ARQ-336	how to hard kill a embedded server ?	1	1	https://issues.jboss.org/browse/ARQ-336
2219	ARQ-336	All the functionality is implemented already but not yet pushed. A user will be able to inject a ContainerController: {code:java} @ArquillianResource ContainerController controller; {code} and start/stop/kill server instances like this: {code:java} @Test public void killTest() throws Exception { //CONTAINER2 variable corresponds to container's name specified via qualifier arquillian.xml // ... <container qualifier="container2" managed="false"> controller.start(CONTAINER2); //do some test controller.stop(CONTAINER2); //do some test controller.start(CONTAINER2); //do some test controller.kill(CONTAINER2); //do some test } {code} The kill method has a default implementation called DefaultServerKillProcessor (implements org.jboss.arquillian.container.spi.ServerKillProcessor) which is used if no other impls are available. These can be shipped with DeployableContainer implementations or directly within a test. The default implementation does the same as stop() so when there is a need to hard kill a server instance, the concrete implementation for the target container has to be supplied. There are certain limitations, though. When running in a Client mode (@RunAsClient), all servers can be started/stopped/killed. To run in InContainer mode only containers other than the one running the test can be stopped/killed.	1	1	https://issues.jboss.org/browse/ARQ-336
2220	ARQ-336	This is starting to tie into this issue a bit: ARQ-236 Maybe we simple should call it "mode"/"controlmode". Where possible values are: suite|class|manual	1	1	https://issues.jboss.org/browse/ARQ-336
2221	ARQ-336	We already had a two state logic in Arquillian - RunAs(CLIENT, INCONTAINER) and it was replaced in Alpha5. Couldn't we move this to the class directly? Something like: {code} @ContainerSuiteLifecycle class FooTest { ... } {code} or {code} @ContainerManualLifecycle("jboss1") @ContainerClassLifecycle("jboss2", "jboss3") class Foo2Test { @Deployment @TargetsContainer("jboss1") ... } {code} In XML/XSD, this mapping would correspond rather to an element than attribute, so user can easily provide his own implementation and not depend on suite/class/manual possibilities, similar to protocols.	0	0	https://issues.jboss.org/browse/ARQ-336
2222	ARQ-336	What happens when: @ContainerSuiteLifecycle public class Foo1Test {} and @ContainerClassLifecycle public class Foo2Test {} exist in the same Suite ?	1	1	https://issues.jboss.org/browse/ARQ-336
2223	ARQ-336	Could this be pushed to Arquillian and then to AS7 soon? We (QE) really need such functionality for our test development. Thanks.	0	0	https://issues.jboss.org/browse/ARQ-336
2224	ARQ-336	Yep, Martin what's the status of this?	0	0	https://issues.jboss.org/browse/ARQ-336
2225	ARQ-336	This feature looks like exactly what we need to test crash recovery in XTS. However, I have one question... We need to be very precise about when the server crashes. To achieve this we use Byteman to inject some crash code into the server at an exact time and place. Will this feature work when the server has killed itself? This is roughly what our usage would look like: {code} @ArquillianResource ContainerController controller; @Test public void killTest() throws Exception { controller.start(CONTAINER1); //start a test scenario //Wait for byteman to kill the server controller.start(CONTAINER1); //Wait for recovery to complete //Verify correct recovery outcome } {code}	1	1	https://issues.jboss.org/browse/ARQ-336
2226	ARQ-336	This won't work because the "start" method will start the container only if it's not started yet. When you don't call stop() or kill(), the container will be still marked as "started": https://github.com/mgencur/arquillian-core/blob/12d710e821e51e9231114cdb3417858fd0f8fdbf/container/impl-base/src/main/java/org/jboss/arquillian/container/impl/client/container/ContainerLifecycleController.java . We would need to reconsider the approach again. However, you can provide your own implementation of a https://github.com/mgencur/arquillian-core/blob/12d710e821e51e9231114cdb3417858fd0f8fdbf/container/spi/src/main/java/org/jboss/arquillian/container/spi/ServerKillProcessor.java and do the waiting for byteman to kill the server in its kill() method (an example can be found here: https://github.com/mgencur/edg-tests/blob/master/failover/src/main/java/org/arquillian/edg/failover/extension/JBossAS7ServerKillProcessor.java). Then you would normally call controller.kill(CONTAINER1). Is that clear?	1	1	https://issues.jboss.org/browse/ARQ-336

#	Requirement	Questions	Answered	Successfully Answered	URL
2227	ARQ-336	Martin, Thanks for the info. Looks like writing our own impl of ServerKillProcesor shouldn't be too difficult. We should be able to use the example in edg-tests to figure out how to specify our own impl of ServerKillProcesor. Does this make sense?	1	1	https://issues.jboss.org/browse/ARQ-336
2228	ARQ-332	i do see your point. At least in my mind, Arquillian should be able to provide the Switzerland of configuration injection though, at least from my perspective. how would the testng stuff work with @Deployment in an Arquillian test? i never got in to testng, so i'm not sure at what point the config parameters are available; but they look like they're available in the equivalent of @Before.	0	0	https://issues.jboss.org/browse/ARQ-332
2229	ARQ-328	org.jboss.arquillian.container.weld.ee.embedded_1_1.mock.AbstractDeployment.configureServices(AbstractDeployment.java:67) at org.jboss.arquillian.container.weld.ee.embedded_1_1.mock.FlatDeployment.configureServices(FlatDeployment.java:32) at org.jboss.arquillian.container.weld.ee.embedded_1_1.mock.AbstractDeployment.<init>(AbstractDeployment.java:50) at org.jboss.arquillian.container.weld.ee.embedded_1_1.mock.AbstractDeployment.<init>(AbstractDeployment.java:55) at org.jboss.arquillian.container.weld.ee.embedded_1_1.mock.FlatDeployment.<init>(FlatDeployment.java:26) at org.jboss.arquillian.container.weld.ee.embedded_1_1.mock.TestContainer.<init>(TestContainer.java:212) at org.jboss.arquillian.container.weld.ee.embedded_1_1.WeldEEMockContainer.deploy(WeldEEMockContainer.java:93) at org.jboss.arquillian.container.impl.client.container.ContainerDeployController$3.call(ContainerDeployController.java:148) at org.jboss.arquillian.container.impl.client.container.ContainerDeployController$3.call(ContainerDeployController.java:115) at org.jboss.arquillian.container.impl.client.container.ContainerDeployController.executeOperation(ContainerDeployController.java:258) at org.jboss.arquillian.container.impl.client.container.ContainerDeployController.deploy(ContainerDeployController.java:114) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:90) at org.jboss.arquillian.core.impl.EventContextImpl.invokeObservers(EventContextImpl.java:99) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:81) at org.jboss.arquillian.container.impl.client.ContainerDeploymentContextHandler.createDeploymentContext(ContainerDeploymentContextHandler.java:78) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:90) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.container.impl.client.ContainerDeploymentContextHandler.createContainerContext(ContainerDeploymentContextHandler.java:57) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:90) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.container.impl.client.container.DeploymentExceptionHandler.verifyExpectedExceptionDuringDeploy(DeploymentExceptionHandler.java:50) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:90) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:134) at	1	1	https://issues.jboss.org/browse/ARQ-328
2230	ARQ-322	In follow-up to your other commands: ---- *deploying glassfish-resources.xml* I agree that ideally we should wait on the fix in GlassFish. If that stalls, I think we could probably hit most of the use cases by parsing the XML file ourselves and issuing the individual delete-* commands in reverse. Let's open a separate tracking issue for that. *setting the debug port* I also agree that the start-domain command should accept a debug port. I did, however, look into what it would take to change the port in the container adapter. In a nutshell, you can only change the debug port when the DAS is running (or you have to edit the domain.xml file manually). Here's the command: {code} asadmin set server.java-config.debug-options=-Xdebug -Xrunjdwp:transport=dt_socket,server=y,suspend=y,address=9009 {code} There is another issue at hand here. In order for the developer to make use of debugging w/ the managed container, it must be possible to set the suspend flag on the jdwp option. So --debug alone isn't sufficient (again, probably something start-domain should support). Here's the code that would be required to do this today: {code:title=GlassFishServerControl#start()} if (config.isDebug()) { result = executeAdminCommand("Setting debug options", "set", Arrays.asList(config.getTarget() + ".java-config.debug-options=" + "-Xdebug -Xrunjdwp:transport=dt_socket,server=y,suspend=y,address=" + config.getDebugSocketPort())); if (result > 0) { executeAdminDomainCommand("Stopping container", "stop-domain", new ArrayList<String>()); throw new LifecycleException("Could not set debug options"); } // NOTE we can't use restart-domain with suspend=y because connecting the debugger // does not resume startup result = executeAdminDomainCommand("Stopping container", "stop-domain", new ArrayList<String>()); if (result > 0) { throw new LifecycleException("Could not stop container to restart in debug mode"); } result = executeAdminDomainCommand("Starting container in debug mode", "start-domain", Arrays.asList("--debug")); if (result > 0) { throw new LifecycleException("Could not restart container in debug mode"); } } {code} {code:title=GlassFishServerControl#stop()} if (config.isDebug()) { result = executeAdminCommand("Restoring debug options", "set", Arrays.asList(config.getTarget() + ".java-config.debug-options=" + "-Xdebug -Xrunjdwp:transport=dt_socket,server=y,suspend=n,address=" + config.getDebugSocketPort())); if (result > 0) { logger.warning("Could not reset debug options, you may need to change them manually"); } } {code} and of course the configuration property: {code:title=GlassFishManagedContainerConfiguration} /** * Socket port to use for debugging */ private int debugSocketPort = 9009; public int getDebugSocketPort() { return debugSocketPort; } public void setDebugSocketPort(int debugSocketPort) { this.debugSocketPort = debugSocketPort; } {code} *add support for JVM flags* If setting the debug port wasn't complex enough, the JVM arguments are just crazy. Again, JVM options can be set with the DAS is running. However, it only supports deleting or appending, not replacing. To do this reliably, you have to read in the JVM options, figure out if the ones you want to add overlap, etc, etc. Here's the asadmin command set to set an option: {code} asadmin --echo create-jvm-options -Xmx512m {code} This is really where I think manually editing the domain.xml would make sense, because working with the available commands is like having your hands tied. ---- There is one other option to all of this. We could support creating a domain configuration to set custom properties such as the debug port and jvm options, then restart into that configuration. The configuration is created using the copy-config command [1]. The benefit to this approach is that it would also side step the issue of having to revert any changes we make for the test run (such as the debug port). [1] http://docs.oracle.com/cd/E18930_01/html/821-2433/copy-config-1.html#scrolltoc	0	0	https://issues.jboss.org/browse/ARQ-322

#	Requirement	Questions	Answered	Successfully Answered	URL
2231	ARQ-322	Aha!!! The REST API no longer allows you to request XML as a response type. If you change the "accepts" from MediaType.APPLICATION_XML to MediaType.APPLICATION_JSON, you'll see that it's trying to return this string: {"message":"","command":"server","exit_code":"SUCCESS","extraProperties":{"commands":[],"methods":[{"name":"GET"},{}],"childResources":{"server":"http:\/\/localhost:4848\/management\/domain\/servers\/server\/server"}}} Can you gut out the XML parsing and replace with JSON parsing? Then we'll be in good shape. Btw, I'm pretty sure the admin REST API has always support JSON responses. For some reason, Jason picked XML when we wrote the first prototype.	0	0	https://issues.jboss.org/browse/ARQ-322
2232	ARQ-322	Ah, right, I should have guessed there was another issue based on the fact that we've discussed it in the past. Thanks for the pointer. As for XML vs JSON, I wonder which would be faster. I would imagine that the developer would prefer speed over one less dependency, since parse speed affects every test execution. I could see the extra dep as a problem if it had to be packaged, but in this case it's purely for client-side usage. Should we at least do an analysis of the speed? If so, I can open a separate issue.	1	1	https://issues.jboss.org/browse/ARQ-322
2233	ARQ-315	Is there a way to do this with JBoss 4.2?	1	0	https://issues.jboss.org/browse/ARQ-315
2234	ARQ-315	Hi Dan, any updates on this? Will this be working for all JBoss AS containers?	0	0	https://issues.jboss.org/browse/ARQ-315
2235	ARQ-311	where is this commit?	1	1	https://issues.jboss.org/browse/ARQ-311
2236	ARQ-297	What Thomas did could be replaced with: mvn clean install -pl container/osgi-embedded/,containers/osgi-remote -am The same goes for the ppl that are interested in only one container: mvn clean install -pl containers/jbossas-remote-6,junit -am	0	0	https://issues.jboss.org/browse/ARQ-297
2237	ARQ-297	I've updated the build page accordingly: https://www.jboss.org/arquillian/build.html We may need a docs update too (I can't remember if we have build instructions there). Should we still move forward with creating some basic profiles? Do we want to get rid of the OSGi profile?	1	0	https://issues.jboss.org/browse/ARQ-297
2238	ARQ-286	Hi, I would like to take a look at this issue. In the end the behavior for the @Rule annotation should be the same as for @Before/@After (execute the rule only on the container if we are in Client mode). Is that correct?	0	0	https://issues.jboss.org/browse/ARQ-286
2239	ARQ-286	It seems to me that it is still failing in RunAsClient scenario. Consider this: {code} @RunWith(Arquillian.class) @RunAsClient public class TestCase { @Rule public ExpectedException e = ExpectedException.none(); @Test public void test() { e.expect(RuntimeException.class) throw new RuntimeException(); } } {code} Observing org.jboss.arquillian.test.spi.event.suite.After and org.jboss.arquillian.test.spi.TestResult still tells me that result.getStatus() == Status.FAILED is true. I am using {code} <dependency> <groupId>org.jboss.arquillian.junit</groupId> <artifactId>arquillian-junit-container</artifactId> <scope>test</scope> </dependency> {code} in the test project. Should I file related jira or client mode was not treated yet? Thanks	0	0	https://issues.jboss.org/browse/ARQ-286
2240	ARQ-286	I'm triing arquillian 1.1.7 with the same code in mode 'in container' so: {code:java} @RunWith(Arquillian.class) public class TestCase { @Rule public ExpectedException e = ExpectedException.none(); @Test public void test() { e.expect(RuntimeException.class) throw new RuntimeException(); } } {code} I see that the rule is executed anyhow in the client side. Should it run in the container side?	1	0	https://issues.jboss.org/browse/ARQ-286
2241	ARQ-286	Do you have the code/pom example somewhere?	1	0	https://issues.jboss.org/browse/ARQ-286
2242	ARQ-286	Ok I reproduced the problem I met. In the RuleTest there is a runtime exception that is throwed if the test is not in the server side, so the rules.Container class starts only once and my test fail. Only a question. Can I choose if run the test only in a mode instead of both leaving the exception inside the instantiation of the rule? I updated the git project with the issue	0	0	https://issues.jboss.org/browse/ARQ-286
2243	ARQ-271	Hi Thomas, are you still working on this? What is the current status? Thanks	0	0	https://issues.jboss.org/browse/ARQ-271
2244	ARQ-271	Any update on this? It looks like it didn't make it into 1.0.0-Final. Is there a work around? Can I force the enricher to execute manually from within my test's @BeforeClass method?	0	0	https://issues.jboss.org/browse/ARQ-271
2245	ARQ-269	Do we know 100% that this will always be the case? GlassFish has a notion of DomainAdminServer, which is responsible for Deploying to the Nodes. Will the Deployer IP and the 'Testing' Node IP always be the same ? ARQ-153 will move the removeServerAddress out of the GlassFish container configuration as a Protocol Configuration as well..	1	1	https://issues.jboss.org/browse/ARQ-269
2246	ARQ-261	(01:08:32) ozizka-here: jweiss: So, did you work with CLI? (01:08:43) jweiss: ozizka-here: a tiny bit (01:08:51) jweiss: cswiii is more experienced with it (01:09:27) ozizka-here: FWIK, it's for a client side, (01:09:27) ozizka-here: so my idea is to get agent to think it's registered, so it doesn't hang on start, (01:09:42) ozizka-here: and then connect through CLI and just operate the objects available (01:10:03) jweiss: that may work, like i said i was guessing	1	1	https://issues.jboss.org/browse/ARQ-261
2247	ARQ-248	Are there API changes that need to be covered by the upgrade? or is this just a examples/docs/default compile v update?	1	1	https://issues.jboss.org/browse/ARQ-248
2248	ARQ-244	Spring cleaning - how relevant is this issue after 6y? Can we close it?	0	0	https://issues.jboss.org/browse/ARQ-244
2249	ARQ-238	do you have a jndi.properties file on your classpath? alla: http://github.com/aslakknutsen/arquillian/blob/ARQ-238/containers/jbossas-remote-4.2/src/test/resources/jndi.properties	1	1	https://issues.jboss.org/browse/ARQ-238
2250	ARQ-238	I cloned the branch but it seems that now there are some compilation errors because of recente changes on upstream. Are you planning to work on this one any time soon? I think I can help on this since it's needed for the project I'm working right now. Thanks for the great work, Davide	1	1	https://issues.jboss.org/browse/ARQ-238
2251	ARQ-238	Brilliant! Here's my prototype: https://github.com/DavideD/arquillian/commits/ARQ-238 I've created the managed and the remote container following the wiki about JSR88 at http://community.jboss.org/wiki/JSR88Client I saw that you were using a different approach...is there any particular reason? I have to clean it but it seems to work. Cheers	1	1	https://issues.jboss.org/browse/ARQ-238
2252	ARQ-197	Hi, Is there a way to run multiple tests with a single Deployment?	1	0	https://issues.jboss.org/browse/ARQ-197
2253	ARQ-197	Hi [~huangdihu], Thank you so much for taking your time to reply to my questions patiently. I currently have a project that has a deployment class Deployments with the extension @ArquillianSuiteDeployment added and two test classes Greeter and Greeter2 both JUnit Tests Running the Tests gives an exception and tests fail. Could you let me know what I am missing in the Test Classes? Or is it possible to provide ur mail ID so that I could send a snippet of my classes?	1	0	https://issues.jboss.org/browse/ARQ-197
2254	ARQ-197	Hi [~huangdihu] That was really helpful. Thank you for the link. Is it possible to use ArquillianSuiteExtension in Arquillian with JUnit tests running in Apache Tomcat Container?	0	0	https://issues.jboss.org/browse/ARQ-197

#	Requirement	Questions	Answered	Successfully Answered	URL
2255	ARQ-197	Hi [~aslak] I've been following this for well over half a year now. I was also trying out the [arquillian-suite-extension|https://www.google.ch/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&cad=rja&uact=8&ved=0ahUKEwjv1oS21NTJAhXImw4KHcfvBxgQFggdMAA&url=https%3A%2F%2Fgithub.com%2Fingwarsw%2Farquillian-suite-extension&usg=AFQjCNFGF2kTBjM9wZt16f2NuDZHCT15zQ] from ingwarsw, but it wouldn't fit my needs. Any update on this? Read something on Twitter that it might be ready in January? Thanks Martin	1	0	https://issues.jboss.org/browse/ARQ-197
2256	ARQ-197	A very quick update on this would be really appreciated. Is there hope that this will be released or should we assume it has been dropped? After all, this is the second most voted for issue in this JIRA... (!)	0	0	https://issues.jboss.org/browse/ARQ-197
2257	ARQ-197	Hi and thank's to all who are working on Arquillian and especially on this issue. I really appreciate this great work. Is there a planned release date for Version 1.2.0?	1	0	https://issues.jboss.org/browse/ARQ-197
2258	ARQ-194	where is the code for this?	1	1	https://issues.jboss.org/browse/ARQ-194
2259	ARQ-187	To be 'true' to what we show, we need to be able to put Arquillian into a dry-run mode, a Profile where the defined container is overridden with a ExportContainer that can export this format on deploy. We should reuse as much of the same Arq core / setup as possible but without touching a live container. We of course also need to exclude the @Test from execution. (note to self, if we use use the DeployableContainer SPI for exporting, how do we deal with Manual deployment which are only happening under @test)	0	0	https://issues.jboss.org/browse/ARQ-187
2260	ARQ-166	Two open questions remain: 1. Should XML be all uppercase in the configuration properties domainXml and resourcesXml? 2. Should we provide the configuration defaultDataSourceName that sets up an org.apache.derby.jdbc.EmbeddedDataSource? The reason for #2 is that the default data source in GlassFish (named jdbc/__default) connects to a remote Derby instance, which will not likely be available to an embedded GlassFish instance.	1	1	https://issues.jboss.org/browse/ARQ-166
2261	ARQ-155	On the other hand, isn't it more intuitive to bring up GF on the standard port, which the user would be expecting? Seems this is more common than the need to avoid conflicts with a running instance, which of course a user shouldn't do anyway. Think of it like this: if they were bringing up two instances on their own, they'd have to configure the 2nd to avoid a port conflict, not expect that as a default.	1	0	https://issues.jboss.org/browse/ARQ-155
2262	ARQ-148	If done on a container level, groups, protocols and extensions won't profit of this, is that a correct assumption?	0	0	https://issues.jboss.org/browse/ARQ-148
2263	ARQ-148	Where can I find this change? I'm pulling in Arquillian 1.0.0.Final-SNAPSHOT and the change is not getting picked up: > mvn dependency:tree [INFO] org.jboss.as.test.clustering:my-arquillian-testsuite:jar:1.0-SNAPSHOT [INFO] +- org.jboss.spec.javax.servlet:jboss-servlet-api_3.0_spec:jar:1.0.0.Final:provided [INFO] +- org.jboss.arquillian.junit:arquillian-junit-container:jar:1.0.0.Final-SNAPSHOT:compile [INFO] | +- org.jboss.arquillian.junit:arquillian-junit-core:jar:1.0.0.Final-SNAPSHOT:compile [INFO] | +- org.jboss.arquillian.test:arquillian-test-api:jar:1.0.0.Final-SNAPSHOT:compile [INFO] | +- org.jboss.arquillian.test:arquillian-test-spi:jar:1.0.0.Final-SNAPSHOT:compile [INFO] | | +- org.jboss.arquillian.core:arquillian-core-spi:jar:1.0.0.Final-SNAPSHOT:compile [INFO] | | \- org.jboss.shrinkwrap:shrinkwrap-api:jar:1.0.0-beta-5:compile ... <container qualifier="jboss" default="false"> <configuration> <property name="jbossHome">target/clustering-udp-0</property> <property name="javaVmArguments">${ip.server.stack}</property> <property name="bindAddress">192.168.0.103</property> <property name="serverConfig">standalone.xml</property> <property name="startupTimeoutInSeconds">30</property> <property name="outputToConsole">true</properties> </configuration> </container> and i'm passing in <systemPropertyvariables><ip.server.stack>-Djava.net.prefer.IPv4Stack=true</ip.server.stack></systemPropertyVariables> Running TestSuite Sep 12, 2011 10:33:07 AM org.jboss.as.arquillian.container.managed.ManagedDeployableContainer startInternal INFO: Starting container with: [java, ${ip.server.stack}, -Djboss.home.dir=target/clustering-udp-1, -Dorg.jboss.boot.log.file=target/clustering-udp-1/standalone/log/boot.log, -Dlogging.configuration=file:target/clustering-udp-1/standalone/configuration/logging.properties, -Djboss.modules.dir=target/clustering-udp-1/modules, -jar, /home/nrla/maven/my-arquillian-testsuite/target/clustering-udp-1/jboss-modules.jar, -mp, target/clustering-udp-1/modules, -logmodule, org.jboss.logmanager, -jaxpmodule, javax.xml.jaxp-provider, org.jboss.as.standalone, -server-config, standalone.xml] Sep 12, 2011 10:33:07 AM org.jboss.remoting3.EndpointImpl <clinit> INFO: JBoss Remoting version 3.2.0.Beta2 Exception in thread "main" java.lang.NoClassDefFoundError: ${ip/server/stack} Caused by: java.lang.ClassNotFoundException: ${ip.server.stack} at java.net.URLClassLoader$1.run(URLClassLoader.java:202) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:247) Could not find the main class: ${ip.server.stack}. Program will exit.	1	0	https://issues.jboss.org/browse/ARQ-148
2264	ARQ-141	Is this documented anywhere? I don't see it here: https://docs.jboss.org/author/display/ARQ/Reference+Guide especially here where I'd expect to: https://docs.jboss.org/author/display/ARQ/Servlet+3.0 . A Google search for "deploymentExportPath site:docs.jboss.org/author/display/ARQ/" doesn't find it.	1	1	https://issues.jboss.org/browse/ARQ-141
2265	ARQ-141	When tested with Arquillian 1.0.0.Final as per: https://github.com/arquillian/arquillian-examples/blob/master/jpa2/src/test/resources/arquillian.xml and: https://community.jboss.org/thread/155043 no deployment archive appears to be produced anywhere within the target/ tree. Is this still supported? Has usage changed?	1	1	https://issues.jboss.org/browse/ARQ-141
2266	ARQ-141	Does it find your arquillian.xml at all ? It's normally located in src/test/resources, but anywhere on the classpath should do. Depends on your IDE I guess if non .java files are actually moved to classpath from src/test/java	0	0	https://issues.jboss.org/browse/ARQ-141
2267	ARQ-138	This looks alpha/beta/release quality now. Submitting the pull request to make an initial release. Tests have been written for * EAR and WAR deployments, * Injection of dependencies into tests using @Resource, @Ejb and @Inject annotations, What might need review - * The project POMs. Since this is a Java EE 5 container, the artifacts from the JBoss repo have been used. * Use of the WebLogic JMX classloader within the project. What is not yet supported (or might not work yet): * Deployment and execution of tests against WLS clusters	0	0	https://issues.jboss.org/browse/ARQ-138

#	Requirement	Questions	Answered	Successfully Answered	URL
2268	ARQ-136	(02:25:06 PM) ALR@Freenode: It cannot get this transitively through the jbossas-embedded-60 container (02:25:11 PM) ALR@Freenode: As import is not transitive. (02:25:16 PM) aslak: aa.. right (02:25:23 PM) ALR@Freenode: But the more overarching problem: (02:25:34 PM) ALR@Freenode: Is that we shouldn't be leaking out everything in compile scope. (02:25:37 PM) ALR@Freenode: Because we did. (02:25:56 PM) aslak: is test transitive ? (02:25:57 PM) ALR@Freenode: And since import scope wasn't transitive, we missed out on the exclusions defined by depchain. (02:26:01 PM) ALR@Freenode: aslak: No. (02:26:10 PM) ALR@Freenode: aslak: It's not our job to define transitive deps. (02:26:13 PM) aslak: that means we have to push it to the user (02:26:21 PM) ALR@Freenode: Once APIs are locked, everything should play nice. (02:26:23 PM) ALR@Freenode: ie: (02:26:39 PM) ALR@Freenode: Using a specific version of Arquillian should not dictate a specific version of a target container. (02:26:52 PM) ALR@Freenode: If that's the case, we're not helping users. (02:26:59 PM) ALR@Freenode: We're actually limiting their ability to choose. (02:27:06 PM) aslak: no, but a specific v. of a DeployableContainer will (02:27:08 PM) ALR@Freenode: And leaking stuff onto their classpath in the process (02:27:14 PM) ALR@Freenode: aslak: Nope. (02:27:21 PM) ALR@Freenode: aslak: We can't do that either. (02:27:29 PM) jharting [~jharting@r11de73.net.upc.cz] entered the room. (02:27:35 PM) jharting left the room (quit: Client Quit). (02:27:45 PM) ALR@Freenode: aslak: The client must make the dependencies upon their container. (02:27:48 PM) aslak: i mean v. as in DeployableContainer impl (02:27:56 PM) ALR@Freenode: aslak: I know. We can't do that. (02:28:50 PM) ALR@Freenode: The user must declare the dep on their target container. (02:29:00 PM) aslak: ok..... maybe i'm not following but.. how do you propose you run a DeployableContainer that is build for Jboss 4 on a jboss 6 ? (02:29:18 PM) ALR@Freenode: 4 and 6 are different APIs, so we should have different impls sure. (02:29:34 PM) ALR@Freenode: Let's take AS 6.0.0.M2 vs. M3. (02:29:39 PM) ALR@Freenode: They have similar APIs (02:29:47 PM) ALR@Freenode: But very different runtime components. (02:29:57 PM) ALR@Freenode: Or a better example is GFv3 vs. 3.0.1 (02:29:59 PM) ALR@Freenode: Same API. (02:30:13 PM) aslak: that's what i'm saying. a deployable is limited to a specific v. of a target container. whether you want to make that implicit or not is different (02:30:30 PM) ALR@Freenode: aslak: No, it's can't be limited. (02:30:41 PM) ALR@Freenode: We abide by a container's deployment API. (02:30:47 PM) aslak: yea, jboss m2 and m3 are different enought to need impl changes.. same with 5.0 to 5.1 (02:31:24 PM) ALR@Freenode: If the deployment API of the backing container changes, yes, we need to change too. (02:31:26 PM) ALR@Freenode: BUT (02:31:39 PM) ALR@Freenode: M2 and M3 aren't locked APIs. (02:31:43 PM) ALR@Freenode: They're milestone releases. (02:31:51 PM) ALR@Freenode: GFv3 vs. 3.0.1 (02:32:00 PM) ALR@Freenode: Should be the same Arquillian DeployableContainer impl (02:32:09 PM) aslak: your not limiting the user by implicid defining the v. your saying, this is tested and build against this v. if you want to try it with a newer v. you may do so, but keep in mind that it might not work. use exclude and include your own (02:32:33 PM) ALR@Freenode: aslak: I know you're trying to help the user. (02:32:44 PM) ALR@Freenode: But you're stepping into their responsibility/ (02:32:53 PM) ALR@Freenode: They must define their runtime CP. (02:32:54 PM) ALR@Freenode: Not us. (02:33:09 PM) ALR@Freenode: This is a symptom of Maven compile scope by default. (02:33:30 PM) ALR@Freenode: They don't need the backing container for compilation, therefore we should not be exporting compile scope. (02:33:38 PM) ALR@Freenode: If we do, then we leak out that version. (02:33:39 PM) aslak: whats the point of a build system if you have to define all your dependencies ? :) (02:33:49 PM) ALR@Freenode: aslak: To build. (02:34:01 PM) ALR@Freenode: aslak: Don't mistake convenience for doing too much. (02:34:08 PM) ALR@Freenode: Try to follow me here. (02:34:18 PM) ALR@Freenode: It took me a long time to come around to understand this. (02:34:31 PM) ALR@Freenode: We export GFv3. (02:34:38 PM) ALR@Freenode: Say the user wants 3.0.1 (02:34:44	1	1	https://issues.jboss.org/browse/ARQ-136
2269	ARQ-127	Hi, can you provide some update on this? It would be a big good-to-have for my usage of arquillian-testng-standalone + graphene2. Right now I'm getting this error: {noformat} java.lang.IllegalStateException: No TestRunnerAdaptor found, @BeforeSuite has not been called at org.jboss.arquillian.testng.Arquillian.verifyTestRunnerAdaptorHasBeenSet(Arquillian.java:246) at org.jboss.arquillian.testng.Arquillian.arquillianBeforeClass(Arquillian.java:101) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.testng.internal.MethodInvocationHelper.invokeMethod(MethodInvocationHelper.java:84) at org.testng.internal.Invoker.invokeConfigurationMethod(Invoker.java:564) at org.testng.internal.Invoker.invokeConfigurations(Invoker.java:213) at org.testng.internal.Invoker.invokeConfigurations(Invoker.java:138) at org.testng.internal.TestMethodWorker.invokeBeforeClassMethods(TestMethodWorker.java:175) at org.testng.internal.TestMethodWorker.run(TestMethodWorker.java:107) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at java.lang.Thread.run(Thread.java:722) ... {noformat}	1	0	https://issues.jboss.org/browse/ARQ-127
2270	ARQ-127	Do you have a test project to point to?	1	0	https://issues.jboss.org/browse/ARQ-127
2271	ARQ-127	Hi Aslak (and others), today I came back to this issue after some time as it's still seems to be a critical thing which we would appreciate with arquillian-testng-stangalone library. I still used the same example just with updated dependencies. I tried some hacking at Arquillian.class and it seems like the usage of ThreadLocal is breaking this TestNG parallelism feature... {noformat} ... Running TestSuite BEFORE SUITE, runner:null Going to set test runner:org.jboss.arquillian.test.impl.EventTestRunnerAdaptor@266c6ce1, threadLocal:java.lang.ThreadLocal@18603b7c BEFORE CLASS, runner:null BEFORE CLASS, runner:null AFTER TEST, runner:null AFTER TEST, runner:null AFTER TEST, runner:null AFTER TEST, runner:null AFTER TEST, runner:null AFTER TEST, runner:null AFTER CLASS, runner:null AFTER CLASS, runner:null AFTER SUITE, runner:org.jboss.arquillian.test.impl.EventTestRunnerAdaptor@266c6ce1 {noformat} It's initialized once at the suite level, but none of the class/test methods are able to access it and it ends up with same message: {noformat} java.lang.IllegalStateException: No TestRunnerAdaptor found, @BeforeSuite has not been called at org.jboss.arquillian.testng.Arquillian.verifyTestRunnerAdaptorHasBeenSet(Arquillian.java:252) at org.jboss.arquillian.testng.Arquillian.arquillianBeforeClass(Arquillian.java:104) {noformat} [~aslak] - can you please confirm that there is any chance that Arquillian will be sometime able to handle this kind of parallelism (at least on class level)? Or maybe there is some quick hint you can provide... or should I just forget this requirement or workaround it at "ugly way"? Thanks a lot for your help! And I also thanks [~lfryc] for his help and support :)	0	0	https://issues.jboss.org/browse/ARQ-127
2272	ARQ-127	[~tremes] adding arquillian to the test groups is required, but does NOT resolve the TestRunner & Nullpointer exceptions. link to my post [https://developer.jboss.org/thread/266910]	0	0	https://issues.jboss.org/browse/ARQ-127
2273	ARQ-123	Don't we have one, or at least the info is there in the docs. Do we still need this issue?	1	1	https://issues.jboss.org/browse/ARQ-123
2274	ARQ-103	Should we name this something more then just reloaded ? - jbossas-reloaded - mc-reloaded ?	1	1	https://issues.jboss.org/browse/ARQ-103
2275	ARQ-40	Do we have enough for Alpha1?	0	0	https://issues.jboss.org/browse/ARQ-40

#	Requirement	Questions	Answered	Successfully Answered	URL
2276	ARQ-25	via DeployableAppender SPI any module can provide a Archive<?> at runtime, but we're missing a 'smart' way to merge/decide what archive type to deploy.. hardcoded as EnterpriseArchive atm.	1	1	https://issues.jboss.org/browse/ARQ-25
2277	ARQ-13	Isn't this done already? Or we have it in the docs?	0	0	https://issues.jboss.org/browse/ARQ-13
2278	ARQGRA-500	Why won't I go via normal means like a user would, arrive to this parameterized page through other pages? Sometimes this path is too complex, the pages involved are too heavy - they have their own tests, other teams support those pages. Sometimes there may not be such a path at all, because some pages are linked directly with parameters by 3rd party applications.	1	0	https://issues.jboss.org/browse/ARQGRA-500
2279	ARQGRA-500	Then again, some parameter provider class wouldn't know the context to decide on parameters. Only the test does. Some location decider probably can get the current test instance and call some callback on it. The test just needs to implement the interface ParameterProvider - this would be the trigger. No need to alter the \@Location annotation. Upd: it's probably not possible to "get the current test instance" without hardcoding some specific test framework like JUnit. Upd2: it should be possible! Otherwise how would Graphene's annotations work. They know the test object, see org.jboss.arquillian.test.spi.TestEnricher.	1	0	https://issues.jboss.org/browse/ARQGRA-500
2280	ARQGRA-500	Dumping some additional concerns that bother me. -- Generalize to dynamic path manipulation (DPM)? "Parameters" only make sense for the HTTP scheme currently. The linked SO's OP also wanted more than that. But generalize too much and it looks like another LocationDecider. Might as well check if the test object implememts the LocationDecider interface. No, I think the generalized DPM should be an implementation helper for Deciders. E.g. the HTTP decider would use it to implement the parameter provider functionality and path substitution functionality ("http://static/\{dynamic\}") - 2 different interfaces. But how to expose this functionality in Graphene.goTo? Either be content with some generic DPM interface argument, or refactor goTo to return a specific builder somehow. Make goTo return a builder? Nasty breaking change - goTo now doesn't trigger navigation. Some new method (arguments like goTo) returns a builder, and an argumentless goTo triggers navigation? This could work. -- Need to descriminate the pages. Problem: LocationDecider only knows a String location. Needs to know pageClass too. --- Annotated parameters could accept it as {{\@LocationParameters(for = MyPage.class)}}.	1	0	https://issues.jboss.org/browse/ARQGRA-500
2281	ARQGRA-485	Interesting reading [~asotobu]. Do you have some API in mind? It would be worth to start a discussion with API snippets on http://discuss.arquillian.org/.	1	0	https://issues.jboss.org/browse/ARQGRA-485
2282	ARQGRA-446	Hey Chris, nice catch. Any suggestion how to avoid this issue? Is that just metaclass property that needs to be avoided in Groovy ecosystem or are we facing bigger issue?	1	1	https://issues.jboss.org/browse/ARQGRA-446
2283	ARQGRA-440	We can upgrade to Core 1.1.4 in 2.0 branch, but imho Drone 2.x upgrade should happen in 2.1 branch.	0	0	https://issues.jboss.org/browse/ARQGRA-440
2284	ARQGRA-440	[~jhuska] does it make sense to upgrade to Drone 2.0 from RichFaces tests perspective? We should really help Drone test the compatibility. At least we could enable a profile that would help test Drone 2.0 compatibility.	0	0	https://issues.jboss.org/browse/ARQGRA-440
2285	ARQGRA-439	What do you think about prototyping this in 2.1 branch and if the complexity of the fix will be low, we can include it in 2.0 as well?	0	0	https://issues.jboss.org/browse/ARQGRA-439
2286	ARQGRA-431	+1 adding .equalsIgnoreCase(String), .matches(regexp) Karel, could you look into adding this capabilities?	1	1	https://issues.jboss.org/browse/ARQGRA-431
2287	ARQGRA-425	[~lfryc] Please, what do you think about this addition ?	1	1	https://issues.jboss.org/browse/ARQGRA-425
2288	ARQGRA-420	Does it makes sense to place it into {{impl}} under {{org.jboss.arquillian.graphene.integration}} package? It would need to make sure that extension won't try to register itself unless appropriate API artifacts are on classpath.	0	0	https://issues.jboss.org/browse/ARQGRA-420
2289	ARQGRA-411	Emil, afaik GrapheneElement should wait for an element to be present, could you please provide more details, i.e. stacktrace?	0	0	https://issues.jboss.org/browse/ARQGRA-411
2290	ARQGRA-409	+1 we can do that for both of frameworks, I just don't know Ember.js internals. Angular wraps all XHR objects and timeouts in order to be able run its model/view synchronization process. Ember doesn't have to allow similar mechanism at all. So let's go ahead and research whether there is a way. :-) ---- On side note, for AngularJS, it is very handy to have special selectors based on ng- attributes such as ng-click, ng-model, etc, such as: {code} @FindByNg(model = "username") WebElement usernameInput; @FindByNg(action = "search()") WebElement searchButton; {code} Is there anything similar for Ember? ---- I would love to have extensions for Graphene and certain frameworks in place. so you could just add artifact on classpath and you will be all setup: {code} graphene-ember graphene-angular {code}	1	0	https://issues.jboss.org/browse/ARQGRA-409
2291	ARQGRA-409	[~smikloso] [~kpiwko] I believe this extension is not requires at this point? Can I resolve this and see whether there will be a need for that in the future?	1	0	https://issues.jboss.org/browse/ARQGRA-409
2292	ARQGRA-408	[~jhuska] [~lfryc] Do you have any particular requirements how this should be done? I would go for similar style as implemented in that smart url mini extension (1) which was heavily inspired from (2). Talking concretely, it could be like {code} @Location(scheme = Scheme.HTTP) // http:// @Location(scheme = Scheme.HTTPS) // https:// @Location(scheme = Scheme.FILE) // file:// @Location(scheme = Scheme.RESOURCE) // resource:// @Location(scheme = Scheme.DROIDIUM) // and-activity:// {code} where HTTP would be default. (1) https://github.com/smiklosovic/arquillian-smart-url (2) https://issues.jboss.org/browse/ARQ-1321	1	1	https://issues.jboss.org/browse/ARQGRA-408
2293	ARQGRA-408	Sounds good. 1) Wouldn't Droidium users appreciate that they won't need to configure scheme on each location? I.e. they should be able to configure scheme globally: {code} <extension name="graphene"> <!-- use HTTPS everywhere --> <property name="scheme">org.jboss.arquillian.graphene.location.Scheme.HTTPS</property> </extension> {code} 2) Shouldn't we make Scheme an extensible concept? As you said, having extensible set of schemes {HTTP,HTTPS,FILE,RESOURCE} that could e.g. Droidium extend with Scheme.DROIDIUM. There should be SPI for registration new schemes. Wdyt?	1	1	https://issues.jboss.org/browse/ARQGRA-408
2294	ARQGRA-406	Hey Juraj, could you please take care of this update?	0	0	https://issues.jboss.org/browse/ARQGRA-406
2295	ARQGRA-392	Linking a pull request where I have submitted the changes. Ready for review.	0	0	https://issues.jboss.org/browse/ARQGRA-392
2296	ARQGRA-382	Yes, it was intentional. I just forget to remove {{Field}} literal from annotation declaration. But now I am thinking of following enhancement: One test class is often used for testing one page, or often we need to load the same page in each test, and then redirect to another. What if it would be possible to have {{@InitialPage}} over field, which would result in automatic loading of the page URL before each test execution? Usage of {{@InitialPage}} in test parameter would override this global setting. Can I change this to Enhancement ?	1	1	https://issues.jboss.org/browse/ARQGRA-382
2297	ARQGRA-382	Given that *initial ~ existing or occurring at the beginning*, this can be both understood as BeforeClass (my understanding) and Before (your understanding). However, all other fields in tests preserve the state among tests, so Before would be an exception. That said, I would either prefer to remove the possibility of annotating a field or making a paramater to @InitialPage, which would additionally alter it's behavior - e.g. where do you want URL to be fired, which some reasonable default identified by broader audience. Nice Enhancement for post 2.0.0.GA release?	0	0	https://issues.jboss.org/browse/ARQGRA-382

#	Requirement	Questions	Answered	Successfully Answered	URL
2298	ARQGRA-382	Agree, let's make this issue opened for post-2.0 development and make the annotation work only for --fields-- method parameters. Juraj, could you please take this over? ^	0	0	https://issues.jboss.org/browse/ARQGRA-382
2299	ARQGRA-381	Juraj, could you please explain what's your use case of browser detection? This issue might be relevant to ARQGRA-135.	1	0	https://issues.jboss.org/browse/ARQGRA-381
2300	ARQGRA-380	[~jhuska], could you please provide a stack trace so it will be evident where is the issue?	0	0	https://issues.jboss.org/browse/ARQGRA-380
2301	ARQGRA-377	Few points: 1. for determining whether an element is enabled, you should use the simplest API possible: {code} WebElement#isEnabled() {code} http://selenium.googlecode.com/git/docs/api/java/org/openqa/selenium/WebElement.html#isEnabled() 2. (y) +1 for being able to build assertions and waiting conditions in unified manner (if statements doesn't apply afaik) ---- 3. we could leverage interceptors to wrap instances and wait on conditions specified in closing word (method), e.g.: {code} element.isEnabled(); waitGui().until().element(element).isEnabled(); assertThat(element, isEnabled()); {code} the problem is that {{WebElement#isEnabled}} returns false, but it should return {{FluentBase<?>}} otherwise. One of the solutions comes from extending capabilities of {{WebElement}} interface and transform it to fluent API, consumed later by e.g. {{GrapheneElement}}. Say we have interface: {code} public interface ElementOperation<B, T> { B isEnabled(); T getText(); ... } {code} then we can: {code} public class GrapheneElement implements WebElement, ElementOperation<Boolean, String> { Boolean isEnabled(); String getText(); ... } {code} Based on a returned type, you could leverage that same API in assetions/waiting conditions: {code} public class FluentElement implements ElementOperation<FluentBoolean, FluentString> { FluentBoolean isEnabled(); FluentString getText(); } {code} ---- Do you have another ideas? I think re-implementation is not possible at this stage, as we are approaching 2.0 Final release. I would rather take it as suggestions for improvement in 3.0 or alternative API in 2.Future.	1	0	https://issues.jboss.org/browse/ARQGRA-377
2302	ARQGRA-353	org.jboss.arquillian.graphene.javascript.DefaultExecutionResolver.execute(DefaultExecutionResolver.java:84) at org.jboss.arquillian.graphene.javascript.JSInterfaceHandler.intercept(JSInterfaceHandler.java:58) at org.jboss.arquillian.graphene.cglib.ClassImposterizer$ClassWithSuperclassToWorkAroundCglibBug$$EnhancerByCGLIB$$4f42a4eb.getHeight(<generated>) at org.jboss.arquillian.graphene.ftest.javascript.EnrichmentTestCase.testScreenMethods(EnrichmentTestCase.java:88) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.jboss.arquillian.junit.Arquillian$6$1.invoke(Arquillian.java:270) at org.jboss.arquillian.container.test.impl.execution.LocalTestExecuter.execute(LocalTestExecuter.java:60) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.invokeObservers(EventContextImpl.java:99) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:81) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:135) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:115) at org.jboss.arquillian.core.impl.EventImpl.fire(EventImpl.java:67) at org.jboss.arquillian.container.test.impl.execution.ClientTestExecuter.execute(ClientTestExecuter.java:53) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.invokeObservers(EventContextImpl.java:99) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:81) at org.jboss.arquillian.graphene.ftest.intercept.InterceptorRegistrationExtension.register_interceptor(InterceptorRegistrationExtension.java:59) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.container.test.impl.client.ContainerEventController.createContext(ContainerEventController.java:142) at org.jboss.arquillian.container.test.impl.client.ContainerEventController.createTestContext(ContainerEventController.java:129) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at	0	0	https://issues.jboss.org/browse/ARQGRA-353
2303	ARQGRA-329	Hey Oliver, that's interesting idea. How would you imagine a usage? Could you share a snippet of the test code which shows how validation should be used?	1	1	https://issues.jboss.org/browse/ARQGRA-329
2304	ARQGRA-329	Oliver, could you prototype this feature on an own branch first?	0	0	https://issues.jboss.org/browse/ARQGRA-329
2305	ARQGRA-326	Hey Juraj, do you have any requirements for fix deadline?	0	0	https://issues.jboss.org/browse/ARQGRA-326
2306	ARQGRA-325	What kind of settings you would like to set?	1	1	https://issues.jboss.org/browse/ARQGRA-325
2307	ARQGRA-325	We could use Drone's configuration service which loads configuration from arquillian.xml into custom class representing configuration. This service does automatic type conversion. I think this would be simpler to implement than page-fragment specific configuration. wdyt?	1	1	https://issues.jboss.org/browse/ARQGRA-325

#	Requirement	Questions	Answered	Successfully Answered	URL
2308	ARQGRA-299	browsers/) using SauceLabs Unfortunately a lot of tests has to be refactored because of errors: {code} Results : Failed tests: testWithTimeout(org.jboss.arquillian.graphene.ftest.wait.WebDriverWaitTest): The waiting time shouldn't be much bigger than 500 MILLISECONDS, but was 770 ms. Tests in error: testGuardHttp(org.jboss.arquillian.graphene.ftest.parallel.TestGraphene): ERROR Job ab918384be4f4630b2d6a2799912c72c is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/ab918384be4f4630b2d6a2799912c72c testWaitWithElements(org.jboss.arquillian.graphene.ftest.parallel.TestGraphene): ERROR Job ab918384be4f4630b2d6a2799912c72c is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/ab918384be4f4630b2d6a2799912c72c testWaitWithBys(org.jboss.arquillian.graphene.ftest.parallel.TestGraphene): ERROR Job ab918384be4f4630b2d6a2799912c72c is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/ab918384be4f4630b2d6a2799912c72c testGuardXhr(org.jboss.arquillian.graphene.ftest.parallel.TestGraphene): ERROR Job ab918384be4f4630b2d6a2799912c72c is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/ab918384be4f4630b2d6a2799912c72c org.jboss.arquillian.graphene.ftest.parallel.TestGraphene: ERROR Job ab918384be4f4630b2d6a2799912c72c is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/ab918384be4f4630b2d6a2799912c72c testContextVsDrone(org.jboss.arquillian.graphene.ftest.parallel.TestInterceptors): ERROR Job 125a372c4d3947ab90ee033b266b1f34 is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/125a372c4d3947ab90ee033b266b1f34 org.jboss.arquillian.graphene.ftest.parallel.TestInterceptors: ERROR Job 125a372c4d3947ab90ee033b266b1f34 is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/125a372c4d3947ab90ee033b266b1f34 testHeadersViaAttributes(org.jboss.arquillian.graphene.ftest.parallel.TestPageFragments): ERROR Job a3f876ff4e6b4542817db07aacd94ebd is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/a3f876ff4e6b4542817db07aacd94ebd testNotNull(org.jboss.arquillian.graphene.ftest.parallel.TestPageFragments): ERROR Job a3f876ff4e6b4542817db07aacd94ebd is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/a3f876ff4e6b4542817db07aacd94ebd testHeadersViaMethod(org.jboss.arquillian.graphene.ftest.parallel.TestPageFragments): ERROR Job a3f876ff4e6b4542817db07aacd94ebd is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/a3f876ff4e6b4542817db07aacd94ebd testDroneInPageFragments(org.jboss.arquillian.graphene.ftest.parallel.TestPageFragments): ERROR Job a3f876ff4e6b4542817db07aacd94ebd is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/a3f876ff4e6b4542817db07aacd94ebd testJavaScriptInterfaceInPageFragments(org.jboss.arquillian.graphene.ftest.parallel.TestPageFragments): ERROR Job a3f876ff4e6b4542817db07aacd94ebd is not in progress. It may have recently finished, or experienced an error. You can learn more at https://saucelabs.com/jobs/a3f876ff4e6b4542817db07aacd94ebd testJavaScriptExecutorInPageFragments(org.jboss.arquillian.graphene.ftest.parallel.TestPageFragments): ERROR Job a3f876ff4e6b4542817db07aacd94ebd is not in	1	1	https://issues.jboss.org/browse/ARQGRA-299
2309	ARQGRA-299	There was a problem with parallel browsers. When I fixed it, the number of failed/error tests reduced to 6: {code} Results : Failed tests: testWithTimeout(org.jboss.arquillian.graphene.ftest.wait.WebDriverWaitTest): The waiting time shouldn't be much bigger than 500 MILLISECONDS, but was 1050 ms. Tests in error: testDroneInPageFragments(org.jboss.arquillian.graphene.ftest.parallel.TestPageFragments): unexpected invocation exception during invocation of org.openqa.selenium.remote.RemoteWebDriver#getCurrentUrl(), on target 'org.openqa.selenium.htmlunit.HtmlUnitDriver@3943f337': object is not an instance of declaring class testJavaScriptExecutorInPageFragments(org.jboss.arquillian.graphene.ftest.parallel.TestPageFragments): unexpected invocation exception during invocation of org.openqa.selenium.remote.RemoteWebDriver#executeScript(), on target 'org.openqa.selenium.htmlunit.HtmlUnitDriver@3943f337': java.lang.ClassCastException@47c87a55 testDroneInPageObjects(org.jboss.arquillian.graphene.ftest.parallel.TestPageObjects): unexpected invocation exception during invocation of org.openqa.selenium.remote.RemoteWebDriver#getCurrentUrl(), on target 'org.openqa.selenium.htmlunit.HtmlUnitDriver@117a5f43': object is not an instance of declaring class testJavaScriptExecutorInPageObjects(org.jboss.arquillian.graphene.ftest.parallel.TestPageObjects): unexpected invocation exception during invocation of org.openqa.selenium.remote.RemoteWebDriver#executeScript(), on target 'org.openqa.selenium.htmlunit.HtmlUnitDriver@117a5f43': java.lang.ClassCastException@5cd974e0 test(org.jboss.arquillian.graphene.ftest.javascript.TestCustomJSInterface): Can't invoke the javacript org.jboss.arquillian.graphene.ftest.javascript.TestCustomJSInterface$Document#getElementsByTagName() Tests run: 245, Failures: 1, Errors: 5, Skipped: 0 {code}	1	1	https://issues.jboss.org/browse/ARQGRA-299
2310	ARQGRA-296	Jan, what does this issue mean /wrt implementation?	1	1	https://issues.jboss.org/browse/ARQGRA-296
2311	ARQGRA-294	Is there any reason why Selenium WebElement can't be injected by web element wrappers?	0	0	https://issues.jboss.org/browse/ARQGRA-294
2312	ARQGRA-294	Couldn't we rather get those changes upstream?	0	0	https://issues.jboss.org/browse/ARQGRA-294
2313	ARQGRA-294	[~jhuska], could you please open a new issue addressing the mentioned requirements ^?	1	1	https://issues.jboss.org/browse/ARQGRA-294
2314	ARQGRA-293	Jan, could you please list some cases where it is necessary to fire a DOM event manually instead of using WebDriver?	1	1	https://issues.jboss.org/browse/ARQGRA-293
2315	ARQGRA-283	Could you start with an outline which we can brain-storm about?	1	0	https://issues.jboss.org/browse/ARQGRA-283
2316	ARQGRA-283	First proposal: # Introduction ## Testing difficulties in general: _unreliable waits_, ## _testing of ajax behavior_, ## _coupling of HTML and test code_ # Introducing Graphene ## Coolest features: _request guards_, ## _Page Objects_ - support for {{@Page}}, nested pages, {{@FindBy}}, _JQuery_ locators, locators from field names, enriching of other injection points like {{@Drone}}, {{@Location}} mechanism ## _Page Fragments_ - rooting mechanism, reference to RichFaces page fragments, dynamically creatable ## _fluent API_ for waits - example ## _JS interfaces_ - example ## _Interceptors_ - example for resolving {{Stale..}} exceptions ## highly extendable and usable with {{Arquillian}} life cycle # Wrap-up ## nice table which summarizes all the features in comparison to {{WebDriver}} General notes: * Where it is possible I will reference to {{WebDriver}} to showcase the difference. So for example when I will be speaking about waits I will also show how it is done in {{WebDriver}}. What do you think Lukas ?	1	0	https://issues.jboss.org/browse/ARQGRA-283
2317	ARQGRA-283	Could you please update a proposal with the current thoughts on the subject? On a sprint planning meeting you have mentioned some FAQs: * Page Abstractions initialization * Select injection	0	0	https://issues.jboss.org/browse/ARQGRA-283
2318	ARQGRA-282	Could you start with an outline so we could discuss publicly what should be included?	1	1	https://issues.jboss.org/browse/ARQGRA-282

#	Requirement	Questions	Answered	Successfully Answered	URL
2319	ARQGRA-282	IMHO it can be done in this way: * change the feature reference _Drive the Browser_ from [here|http://arquillian.org/features/], to contain _Graphene_ reference. It would be nice to have there also Graphene logo (ARQGRA-69)(together with Drone logo) * reuse and alter the already existing [functional testing using drone|http://arquillian.org/guides/functional_testing_using_drone/], because IMHO there are valuable parts, which are needed (_Assumptions_, _Introducing Client Mode_, _Enabling Client Mode_, _Drone Overview_). So dropping the whole guide is not a good option for me. * it can be done in a similar way as the [guide for JPA testing|http://arquillian.org/guides/testing_java_persistence/] is done. So _Introducing the application we are going to write tests for_, _Preparation for testing with Graphene_, _Developing tests_, _Running the tests_ * basically I have not recognized other candidates from Lukas's ones, I just ordered them in the outline. The guide outline itself can look like: # Assumptions # Introducing Client Mode # Enabling Client Mode # Drone Overview ## There are mentioned all the browsers drivers Drone currently has support for, we can emphasize _Graphene_ here and say that we are going to setup a test suite project with use of it from now on. # Introducing the application we are going to write tests for ## We can show how to develop a simple _HTML 5_ page, which will use _AJAX_ calls to retrieve backend data from _JAX-RS_ endpoints. ## *Or* we can reuse again the _JSF_ sample which is already there, logging fragment, and to add some parts to easily showcase _Graphene_ features on it. My concert here is the _JSF_. I would like to showcase that _Graphene_ can be used with all front end frameworks and most importantly with now most coolest POH5 applications. By this IMHO we can get much bigger audience. ## It does not have to be all in this guideline. There can be just snippets from whole Graphene showcase, which will be developed as well. (ARQGRA-174) # _Graphene_ scaffolding # Senium server setup, remote reusable, chrome browser -> fast development turnaround # Developing tests ## Showcasing some of the _Graphene_ features: _request guards_, _fluent waitings_, _Page Ojects_, _Page Fragments_ ## This should be probably the biggest part of the guideline. My concern here is that it will be too big. On the other side I do not have a major problem with it. IMHO the navigation which is used in arquillian.org is quite clear and that the platform is suitable for long posts as well. # Running and debugging the test in IDE # Best practices from developing and running tests What do you think guys ?	1	1	https://issues.jboss.org/browse/ARQGRA-282
2320	ARQGRA-282	+(really big number). Existing guide should be rewritten to show Graphene2 primarily. Other Drone supported Selenium/WebDriver could be handled by infobox saying something like: _Not starting a new test suite from scratch? Do you already have Selenium 1 and WebDriver (and Graphene1 :-P) tests? See Drone reference documentation how to run those instead of Graphene 2_ No reason why to keep existing content.	0	0	https://issues.jboss.org/browse/ARQGRA-282
2321	ARQGRA-282	Jura, can you please sent this as PR against arquillian.org repository? Much easier to comment that way.	0	0	https://issues.jboss.org/browse/ARQGRA-282
2322	ARQGRA-282	Hey Juraj, could you please review the guide from the grammar perspective?	1	1	https://issues.jboss.org/browse/ARQGRA-282
2323	ARQGRA-282	[~bleathem], could you please check a grammar of the guide?	0	0	https://issues.jboss.org/browse/ARQGRA-282
2324	ARQGRA-280	How would you solve a problem of cache invalidation in case of stale reference?	1	1	https://issues.jboss.org/browse/ARQGRA-280
2325	ARQGRA-279	I don't get why we have this {{element(By)}} version. Couldn't you use {code} WebElement el = driver.findElement(By) Graphene.waitAjax().until().element(el).is().present(); {code} [~jpapouse], doesn't the Alpha4 feature ensure that the element does not have to be present during invocation of {{#findElement}}?	1	1	https://issues.jboss.org/browse/ARQGRA-279
2326	ARQGRA-267	Jan, do we still need to upgrade HtmlUnit manually? It's version is specified by version of Selenium.	1	0	https://issues.jboss.org/browse/ARQGRA-267
2327	ARQGRA-265	I don't like to introduce own solution - could we discuss if it is really necessary? https://issues.jboss.org/browse/ARQGRA-258?focusedCommentId=12753402&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12753402	1	1	https://issues.jboss.org/browse/ARQGRA-265
2328	ARQGRA-264	I looked through the pull request and I have 2 questions: * Is it possible to switch to default frame? It seems that either {{index}} or {{nameOrId}} should be specified. Is it needed at all, or elements without {{@InFrame}} are always looked up in default frame? * What about adding possibility to specify {{@InFrame}} on class for the whole fragments which live in a frame?	1	1	https://issues.jboss.org/browse/ARQGRA-264
2329	ARQGRA-264	You are right Aliaksei, it would not work for all cases. It is not supporting nested structure of iframes, only one level of iframes. The nested functionality if needed can be IMHO implemented as I described [here|https://github.com/arquillian/arquillian-graphene/pull/81#discussion_r4743171]. Is there a real life scenario where this multiple iframe levels are needed ? Otherwise, well documented feature would be sufficient IMHO. The {{@InFrame}} annotation can not be used to switch to the default iframe. IMHO the legacy tests should be rewritten to use this annotation. Another option is to implement such {{@InFrame}} which would be capable doing so. What do you think Lukas ?	1	1	https://issues.jboss.org/browse/ARQGRA-264
2330	ARQGRA-261	Hey Jan, I believe this is already possible, have you tried that?	1	0	https://issues.jboss.org/browse/ARQGRA-261
2331	ARQGRA-258	Hey Jan, would it be possible to inject a proxy of WebElement, which will be initialized later?	1	1	https://issues.jboss.org/browse/ARQGRA-258
2332	ARQGRA-258	Have you associated any commits with this issue since it is Done? Should we resolve it with other Resolution?	1	1	https://issues.jboss.org/browse/ARQGRA-258
2333	ARQGRA-258	We doesn't have to inject *a proxy for the Select*, we can rather inject *the Select with a proxy to the WebElement.* Could we?	1	1	https://issues.jboss.org/browse/ARQGRA-258
2334	ARQGRA-255	Hey Jura, could you please drive an incorporation and improvements of documentation for new features in these releases?	1	1	https://issues.jboss.org/browse/ARQGRA-255
2335	ARQGRA-230	Hey Jan, if we don't have time to improve it in scope of Alpha3, could you please create request for improvement and resolve this? Note that Alpha3 code-freeze is Fri Jan 4th.	1	1	https://issues.jboss.org/browse/ARQGRA-230
2336	ARQGRA-228	Hey Jan, could you please comment what do you think about my proposal?	1	1	https://issues.jboss.org/browse/ARQGRA-228
2337	ARQGRA-219	Hey Lukas, we scheduled first meeting for this on Monday (October 29, 2012). Should we reschedule it so you can be there Lukas ? Or it is enough to send pull request and you can review it there ? Thanks.	1	1	https://issues.jboss.org/browse/ARQGRA-219
2338	ARQGRA-200	Jiri, could you please clarify where guards aren't enough in your use case?	1	1	https://issues.jboss.org/browse/ARQGRA-200
2339	ARQGRA-200	Thanks for the informations. So, should I (always) use some explicit waiting after guarded request? What are the use-cases of the guards now?	1	1	https://issues.jboss.org/browse/ARQGRA-200
2340	ARQGRA-199	Hey Juraj, could you please provide more details? I consider this report quite incomplete.	1	1	https://issues.jboss.org/browse/ARQGRA-199
2341	ARQGRA-195	Juraj, could you please propagate exception and add information about field which is being enriched?	1	1	https://issues.jboss.org/browse/ARQGRA-195
2342	ARQGRA-193	Karel, what is prefered way to configure via system properties then?	0	0	https://issues.jboss.org/browse/ARQGRA-193
2343	ARQGRA-193	_arq.extension.name...._ system properties are already available from descriptor, am I right?	0	0	https://issues.jboss.org/browse/ARQGRA-193
2344	ARQGRA-191	Jan, could you please file an issue for Drone that it could be initialized without having {{@Drone}} directly in test case?	1	1	https://issues.jboss.org/browse/ARQGRA-191
2345	ARQGRA-190	Hey Juraj, we would like to support inner classes with Page Abstractions. Could you please look at it and tell us if it's feasible?	1	1	https://issues.jboss.org/browse/ARQGRA-190
2346	ARQGRA-186	Hey Juraj, I think this was rather misuderstanding. I don't think it';s necessary to rename those interfaces to *PageFragment. I would leave them as *Component or remove the suffix completely. WDYT?	1	1	https://issues.jboss.org/browse/ARQGRA-186

#	Requirement	Questions	Answered	Successfully Answered	URL
2347	ARQGRA-186	Removing the suffix completely can lead to some conflicts with class names, e.g. {{Calendar}} vs. {{java.util.Calendar}}. Suffix also sounds good for me because it tells me what that classes really are, and what they are good for. Page Fragment is little bit long but personally I do not mind. Do you find this justifying reasonable ? Let me know if I should let it with Component suffix. Thanks.	1	1	https://issues.jboss.org/browse/ARQGRA-186
2348	ARQGRA-170	Hey Jan, could we use following entry points? * {{GrapheneContext.getDriverProxy();}} * {{GrapheneContext.getDriverProxyWithInterfaces(Class<?>[] interfaces);}} * {{GrapheneContext.getDriverProxyByImplementation(Class<? extends WebDriver> implClass);}} * {{GrapheneContext.getConfiguration();}} * {{GrapheneContext.getProviders();}} ** returns {{ProviderHolder}}	1	1	https://issues.jboss.org/browse/ARQGRA-170
2349	ARQGRA-160	Great idea, could you please elaborate which configuration options we should start with? For now, I can think just about wait timeouts, right?	1	1	https://issues.jboss.org/browse/ARQGRA-160
2350	ARQGRA-156	Hi Jan, could you please look at this issue? It should be simple one.	0	0	https://issues.jboss.org/browse/ARQGRA-156
2351	ARQGRA-155	[~okiss], could you please share the code what you have so we can take a look? Let's open a branch on your fork of Graphene, please.	0	0	https://issues.jboss.org/browse/ARQGRA-155
2352	ARQGRA-152	Karel, how critical is this issue? Does it mean instantiation of browser makes serious delays in tests, so thus it is blocking the machine where the test runs?	1	1	https://issues.jboss.org/browse/ARQGRA-152
2353	ARQGRA-132	Karel, could you please clarify what is the reference usage of BOM? Since {{arquillian-graphene}} depchain brings you all the dependencies, there is no need for BOM, or is it?	1	1	https://issues.jboss.org/browse/ARQGRA-132
2354	ARQGRA-131	If I read this correctly, Arquillian Drone can handle configuration, creation and destruction of multiple AjaxSelenium instances. However, it would be up to user to switch to correct context manually by his own. Suppose following: {code:java} @Drone AjaxSelenium b1; @Test public void test(@Drone AjaxSelenium b2) { // we can't switch to b2, as the first method might be b1.click(...); // no way how to check method body } {code} The only other idea I got so far is that AjaxSelenium will set correct context in each call, e.g. {code:java} public class AjaxSelenium { public void click(...) { AjaxSeleniumContext.setContext(this); // do the click() } } {code} Lukas, what do you think?	1	0	https://issues.jboss.org/browse/ARQGRA-131
2355	ARQGRA-114	Should we create a sub space on Arquillian community site ? or just reuse the Arquillian one? I guess we can reuse for now, non of them is flooded with blogs so.. ;)	1	0	https://issues.jboss.org/browse/ARQGRA-114
2356	ARQGRA-84	Hi Juraj, could you please look at it once Page Extensions will be implemented (ARQGRA-71)? We can discuss the implementation if that would not be familiar - e.g. JavaScript part.	1	1	https://issues.jboss.org/browse/ARQGRA-84
2357	ARQGRA-84	The pull request is not ready for merging, since tests does not pass on buildhive. Juraj, could you please take a look into failing tests?	1	1	https://issues.jboss.org/browse/ARQGRA-84
2358	ARQGRA-76	I have moved configuration of plugins needed for Jenkins profile into parent pom.xml. However, there is need for support for two different paths, from where the arquillian.xml is copied. One for parent pom: ${basedir}/../qa/richfaces-arquillian-xml-property-resolver/target/conf/arquillian.xml and second for all child poms: ${basedir}/../../qa/richfaces-arquillian-xml-property-resolver/target/conf/arquillian.xml I have solve this with defining property in parent pom, with the value of path which satisfy parent pom, and also overwritten this property with value which satisfy child poms. Can this be right solution ? Another solution which works is to define only path which satisfy child poms and to disable enforcer plugin, which just not copy arquillian.xml for parent. Which one will be more suitable please ?	0	0	https://issues.jboss.org/browse/ARQGRA-76
2359	ARQGRA-76	Karel, isn't there any way to point to the top-level pom such as use parent.relativePath? Jura, you can simulate this for you and add property topLevelProject.relativePath to each module and use it after basedir path: ${basedir}/${topLevelProject.relativePath}/qa/richfaces-arquillian-xml-property-resolver/target/conf/arquillian.xml	1	0	https://issues.jboss.org/browse/ARQGRA-76
2360	ARQGRA-76	Great, Karel, thanks for your insight. Jura, could you please adapt this property for your needs?	0	0	https://issues.jboss.org/browse/ARQGRA-76
2361	ARQGRA-76	Pavol, could you look at this issue?	0	0	https://issues.jboss.org/browse/ARQGRA-76
2362	ARQGRA-74	Perfect, could you create some functional tests for {{element(...)}} and {{attribute(...)}}? I assume methods {{initElement/s}} should be aligned with Juraj's Component Objects ARQGRA-40.	1	1	https://issues.jboss.org/browse/ARQGRA-74
2363	ARQGRA-67	(1) As commented on pull request - could you change that to {{window.Graphene}}? (2) could we just extend DefaultExecutionResolver instead of replacing? (3) why it should be exposed publicly? I don't have strong objections, just GrapheneProxy was intented to be used by internal implementation. How does interception matches with ARQGRA-79?	0	0	https://issues.jboss.org/browse/ARQGRA-67
2364	ARQGRA-45	Karel, I suppose there is no need for object encapsulation of request headers? Do you want to contribute some code around authorization to Ajocado utilities? I can imagine to have encodeCredentials there.	1	1	https://issues.jboss.org/browse/ARQGRA-45
2365	ARQGRA-45	Can you please elaborate and API you have suggested? You propose to use following to add authorization? {code:Java} String credentials = URLUtils.encodeBase64Credentials(username, password); Header authorizationHeader = new Header("Authorization", "Basic " + credentials); selenium.addCustomRequestHeader(authorizationHeader); {code} I propose to introduce utility class for manipulation with common request like: {code:Java} Header authorizationHeader = HeaderUtils.createBasicAuthorization(username, password); {code}	0	0	https://issues.jboss.org/browse/ARQGRA-45
2366	ARQGRA-40	Juraj, could we work on incorporating the Components Objects into Graphene?	1	1	https://issues.jboss.org/browse/ARQGRA-40
2367	ARQGRA-37	Jura, could you please create separated JIRAs for issues you have mentioned? Add them to doc component.	0	0	https://issues.jboss.org/browse/ARQGRA-37
2368	ARQGRA-33	Juraj, it's good point, could you please provide pull request for this feature?	1	1	https://issues.jboss.org/browse/ARQGRA-33
2369	ARQGRA-28	Karel, could you please elaborate more on this issue - it is affecting compatibility between Se1 and Se2?	0	0	https://issues.jboss.org/browse/ARQGRA-28
2370	ARQRUSH-38	Could you move those checkboxes to main panel of ProjectManager? So it will be immediately clear what filters are applied? Additionally, could you add Show Differs quick filter similar to Show All?	1	1	https://issues.jboss.org/browse/ARQRUSH-38
2371	ARQRUSH-28	Hey Jan, could you please send pull request for those changes? Right now, we are basing development on Jakub's repository: https://github.com/cube88/arquillian-rusheye	1	1	https://issues.jboss.org/browse/ARQRUSH-28
2372	ARTIF-748	[~eric.wittmann], thanks again for the help. Am I missing anything here?	1	1	https://issues.jboss.org/browse/ARTIF-748
2373	ARTIF-748	Thanks [~eric.wittmann]. [~csa], any thoughts on that? See my comment on KEYCLOAK-1539. Essentially, a call to JSON REST using Errai's Caller may be unauthorized. But instead of returning a 401 (or anything helpful), KeyCloak returns a login page's HTML. GWT's JSON marshallers, of course, fall over themselves. However, I was a little surprised that the error didn't hit the Caller's ErrorHandler *at all*. Should it? I'll need to play with it again to come up with a stacktrace. But regardless, if a GWT marshaller throws an exception on a REST call, could that bubble up to the ErrorHandler somehow? Or, is there a way to deal with this with some sort of Interceptor, like Eric suggested?	1	1	https://issues.jboss.org/browse/ARTIF-748

#	Requirement	Questions	Answered	Successfully Answered	URL
2374	ARTIF-748	[~brmeyer], I agree with your summary. Since the response code is 200 OK, Errai will try to demarshall the response which it expects to be JSON-based, as specified in @Consumes and @Produces. So, the server is really returning an invalid response and you will see an unhandled exception. Interceptors are probably not an ideal solution here as you don't know which REST call will fail after the token expired, so you'd have to intercept them all. I could easily add error handling for this case to Errai 3.2, if needed. BTW, did you already upgrade to 3.2 and GWT 2.7?	1	0	https://issues.jboss.org/browse/ARTIF-748
2375	ARTIF-730	Completed in https://github.com/brmeyer/s-ramp/tree/wildfly9. But, for now, holding. The server overlay from ARTIF-739's Keycloak upgrade no longer supports WF 8, so we'd have to follow suit and drop it. Also, WF 9 upgraded to Hibernate Search 5.2, which is really different than the 4.x in EAP 6 and WF 8. Supporting both with the distro *might* be possible, but it'll be horrifically complex to do so in the actual repo module (some packages are different, etc.). So, this might imply dropping EAP 6.4 as well. Not sure if 1.1.0.Final is the place to do all that -- how soon is too soon?	1	1	https://issues.jboss.org/browse/ARTIF-730
2376	ARTIF-674	Can't the extension support be added without breaking compliance?	1	1	https://issues.jboss.org/browse/ARTIF-674
2377	ARTIF-659	On hold. At least guava-gwt, and therefore Errai 2, require Java 6-7 and fail with Java 8.	1	1	https://issues.jboss.org/browse/ARTIF-659
2378	ARTIF-659	Hey [~csa]! Quick question: Has anyone had success building against Errai 2 with Java 8? guava-gwt:14.0.1 appears to require 6 or 7, so at least errai-data-binding is an issue. Can I safely upgrade guava-gwt, etc., or is stepping up to Errai 3 the only option?	1	1	https://issues.jboss.org/browse/ARTIF-659
2379	ARTIF-659	Where did you get that info from? I think that's outdated. To my knowledge 2.8 will have support for Java 8. In fact, you can already use a JDK8 now as long as you don't use features that won't translate to JS.	1	0	https://issues.jboss.org/browse/ARTIF-659
2380	ARTIF-645	Like the idea of a pluggable architecture - especially in respect of the UI. The current UI is very much a generic browser to navigate around a wide range of artifacts and view their metadata and relationships. However the information held within the repository could be idea for business specific views - so this pluggable approach could be one way in which organisations could define their build specific views, presenting the key information they are interested in, and allowing users to navigate through the relevant subsets of information they are interested in. For example, one usecase I heard related to the deployment lifecycle information being stored in SRAMP, was having service reports that showed how long it took to progress though the various stages of the lifecycle, and who (i.e. which teams) delayed the deployment. Although with custom views, it may also be good to look at this capability in the context of fine grained access control, to potential enable only some of these custom views to be offered/presented to specific users - e.g. a business user may not be presented with the generic browser and instead may only see a specific set of the custom views, to restrict (a) the information they see (presented at a suitable high level) and (b) the way in which they can navigate around the model.	1	0	https://issues.jboss.org/browse/ARTIF-645
2381	ARTIF-601	[~eric.wittmann], this will be similar to the updateContent discussion. And really, same question for deleting an artifact itself. Presumably, updateContent and deleteContent will both delete the derived artifacts in some fashion. What should happen if relationships *target* something that's deleted or swapped out? Should the relationships automatically get deleted? Should we throw some sort of constraint exception and block the delete? I'd vote for the latter -- less risky and *far* less complicated. Thoughts?	1	1	https://issues.jboss.org/browse/ARTIF-601
2382	ARTIF-583	I think is good to have a maven tree view depending on trhe complex situations we can reach to. I thought the current implementation was complete and it displayed the artifacts included in a group and listed correctly the groups. What are the key areas that are being missed? And what are the complex situations maybe we have to deal in the future? Just for my understanding. I imagine the artifact grouping could be one situation. Just as the user view, I think having a view like this is useful. As this url will be added as maven repository, I like the idea of having access to the artifacts. What would be good also, is to have a kind of tree organized view in s-ramp.	0	0	https://issues.jboss.org/browse/ARTIF-583
2383	ARTIF-578	Hey [~virchete], did you have anything specific in mind for this?	1	1	https://issues.jboss.org/browse/ARTIF-578
2384	ARTIF-578	You were so confident that there'd be issues, that you opened a ticket in advance?! Just kidding ;)	1	1	https://issues.jboss.org/browse/ARTIF-578
2385	ARTIF-573	Shouldn't Ctrl-U and Ctrl-K clear the text (from cursor to the start and end of the line, respectively)? Ctrl-C is SIGINT - shouldn't that exit the shell?	1	1	https://issues.jboss.org/browse/ARTIF-573
2386	ARTIF-506	Having spent ages having to meddle with this metadata on a previous project - I should mention that it can be extremely tricky to get the results you would expect (unless maven metadata is available). For instance, in my experience the following aren't uncommon in externally provided JARs: - `META-INF` directory to be missing entirely (essentially not a valid JAR) - Manifest data often doesn't correspond with what you would expect, because the user decided can some arbitrary value in there (e.g. Bundle-Version => 2.1.0.201209190230-r instead of 2.1.0, presumably someone decided it'd be nice to have date and time in there). [1] - There may be correct data available, but it might not be stored in the key you'd expect (sometimes people put them in Bundle-Version other times in Specification-Version) - If there is a 'meta jar' (i.e. Jar wrapping a whole load of other Jars), it may not even have a `properties.pom`/maven metadata which is useful. Ultimately for the worst cases I ended up having to use a heuristic based approach, with regexp for the ultimate fallback if all other methods failed (would we just reject completely malformed ones?). While I was developing the aforementioned tooling, I wrote a little program which will scan through a directory tree, cracks open all of the jars and builds a set of keys (and values) from all the manifest files. This gives you a very good idea of which fields you may want to look at to derive the data you need, in addition to the sort of rubbish data it may contain you'll have to try to avoid! Sample output (huge set): https://gist.github.com/msavy/cefc9c9b01416176128f Note format is: key => { v0, v1, ..., Vn } I'll post a source link as a priv comment as I haven't put any copyright headers and whatnot on it.	1	0	https://issues.jboss.org/browse/ARTIF-506
2387	ARTIF-506	For now, rejecting this. However, note that some semblance of this is now available in the Maven facade. When you configure Artificer as a repo in your pom, you can add "?artifactType=FooType" to the URL. This sets the type for the *primary* artifact only. Since this entire discussion revolves around primary types, I'm calling that "good enough" for now. I'd rather keep it to something simple, rather than require users to add custom metadata (and have Artificer parse it).	0	0	https://issues.jboss.org/browse/ARTIF-506
2388	ARTIF-490	Thanks [~ncross] -- that's in SRAMP-489. Am I correct that it's limited to the WAR projects?	1	1	https://issues.jboss.org/browse/ARTIF-490
2389	ARTIF-488	Does this duplicate SRAMP-412?	1	1	https://issues.jboss.org/browse/ARTIF-488
2390	ARTIF-487	From [~eric.wittmann]: {quote} Consider: don't allow updates for non-SNAPSHOT artifacts. For SNAPSHOT artifacts add a new one (do not update). So in other words, *never* update the s-ramp artifact. Instead, either fail (non-SNAPSHOT) or create new (SNAPSHOT). The latter should never be done, but can be supported anyway (users probably shouldn't be putting snapshot artifacts in s-ramp). {quote}	0	0	https://issues.jboss.org/browse/ARTIF-487
2391	ARTIF-450	Another question - is the commons installer going to setup the 'admin' user, with the password entered by the user?	0	0	https://issues.jboss.org/browse/ARTIF-450

#	Requirement	Questions	Answered	Successfully Answered	URL
2392	ARTIF-448	Futhermore, the attribute nameInSource is applicable for the relational tables, not for the whole TeiidModel (*.xmi), or am I wrong?	1	0	https://issues.jboss.org/browse/ARTIF-448
2393	ARTIF-447	Instead of splitting the installers back up, I'm wondering if we could do something simpler/smarter. At least with respect to S-RAMP, really the only *runtime* difference between EAP 6.1 and 6.3 was the ModeShape distro that the installer pulls down. Everything else, including compile time, IDP, etc. can stay the same. Knowing that, could we attempt to simply decide, during the installation, what ModeShape distro to grab, based on the EAP instance's full path (assuming it's not renamed and still contains the version within the dir name)? If that fails, fall back on the latest? Theoretically, that *should* allow S-RAMP to simultaneously support both, but not splitting up the installers in overlord, overlord-commons, and s-ramp. [~eric.wittmann]/[~objectiser], thoughts?	1	1	https://issues.jboss.org/browse/ARTIF-447
2394	ARTIF-447	Can you both review https://github.com/Governance/s-ramp/pull/430? Comments appreciated. The only thing I don't like about it is that the s-ramp-distro-assembly now has to include both versions of ModeShape, but not sure if there's a way around that.	0	0	https://issues.jboss.org/browse/ARTIF-447
2395	ARTIF-434	Is there any example of this in our current code? I tried to search about this in all the java files and no result.	1	0	https://issues.jboss.org/browse/ARTIF-434
2396	ARTIF-433	I would defer to those with more eventing/JMS experience than me. However, it does seem like a serialized java object isn't ideal. I would also think that each event would at least need a small set of standard properties such as 'type' - no? I'm wondering what a typical CEP system would need to see. Does anyone have experience with Drools CEP? I did some Esper a long time ago. I think those are the types of things we'd want to leverage to react in interesting ways to the events.	1	1	https://issues.jboss.org/browse/ARTIF-433
2397	ARTIF-433	Gary, thanks for all that -- makes sense. > have a default queue name configured in properties (not with a dtgov specific name, so could potentially be used by any consumer) My only concern with that is shouldn't it be reserved solely for dtgov? What happens if dtgov is running, but a user attempts to pop messages off the "default queue" at the same time? Unless the other clients use a JMS session rollback (as I understand it), dtgov won't receive it, right? That was the intention behind using a dtgov-specific queue. Or am I missing something there? I hadn't intended for the topic(s) to be configurable at all, but I suppose that might make sense in certain situations. > if you use destinations, then you won't need to worry about whether the jndi name refers to a topic or queue Great point -- will do.	1	1	https://issues.jboss.org/browse/ARTIF-433
2398	ARTIF-433	What are we trying to avoid by installing S-RAMP with a default "dtgov" queue? Is it hard to configure queues? Will the DTGov installer not be able to set up its own queue in S-RAMP? Just trying to understand. Logically I would expect S-RAMP to be configured with a default notification topic and a simple way to support multiple named queues (for each of the consumers, of which dtgov is one). A named queue should be relatively easy to create, such that the DTGov installer would add it to the S-RAMP config during DTGov install time.	1	1	https://issues.jboss.org/browse/ARTIF-433
2399	ARTIF-431	Can't auto-download EAP though, right? Do we use wildfly instead? Also, why not use the embedded containers? (question from an Arquillian n00b)	1	1	https://issues.jboss.org/browse/ARTIF-431
2400	ARTIF-431	"As the selenium tests would be just junit tests, they would be located in the src/test folder of every ui project." As they need to be run within the scope of a running target server, won't it just be easier to include these as part of the integration test infrastructure? QE also use selenium, so may be easier to exchange tests.	1	1	https://issues.jboss.org/browse/ARTIF-431
2401	ARTIF-431	Re-wrote the description, after scraping off my Arquillian rust ;) All, take a look at https://github.com/brmeyer/s-ramp/compare/integration-tests and let me know your thoughts. So far, it downloads Tomcat, unpacks it, and runs the full S-RAMP installer on it. Arquillian runs Tomcat in a managed-mode, then runs a simple client test (passes). EAP setup will be identical, minus the downloading (unpack the distro zip from a known location or support $JBOSS_HOME). Arquillian does not currently have a Jetty 8 container. Since there is Jetty 7 & 9, I may look into contributing one for 8. However, note that *none* of them support managed-mode -- all require embedded. Do we want to try that for Jetty, or look into writing a managed one?	1	1	https://issues.jboss.org/browse/ARTIF-431
2402	ARTIF-382	[~objectiser], could we simply rename (or deprecate) s-ramp-server-eap61 and create s-ramp-server-eap6? Looking through it, that should work for 6.1.0 through 6.3.0 (and possibly 6.4.0, if there is one). Unless I'm missing something, I don't see anything in s-ramp-server-eap61 that would need to be changed for 6.3. Thoughts?	1	1	https://issues.jboss.org/browse/ARTIF-382
2403	ARTIF-382	[~objectiser], stupid question: where do the docs live? For the life of me, can't find the src...	1	1	https://issues.jboss.org/browse/ARTIF-382
2404	ARTIF-309	NOTE: This network/graph based representation could also form the basis of the static representation that underpins the 'service dependency' graph used in the runtime governance project, which currently is dynamic, based upon dynamic service interactions that have occurred within a configurable timeframe. Therefore when implementing this task, co-ordination with the runtime governance project should occur to ensure that the relevant common functionality can be factored out into share components - while at the same time not making either project dependent upon the other (e.g. at build, deployment or runtime).	0	0	https://issues.jboss.org/browse/ARTIF-309
2405	ARTIF-303	After discussing with the UX team, creating a new unique URL for every change made to the filters would result in too many items added to the browser history. Instead, we decided to solve this issue by doing the following: 1) Remember the "last" value of the filter settings, so that when the user navigates back to the Artifacts page from the Artifact Details page (for example) we can show the old filter settings (note: this includes the page #) 2) Offer a "share this" feature on every page in the UI - this allows the user to bookmark the page. Ideally the "share this" link generated by this feature will be to a permalink servlet which contains all of the query parameters needed by the page. The permalink URL will not contain a # - the permalink servlet will process the URL and redirect the user to an "application" url that *does* include the #. This also solves the problem of SSO redirects breaking application URLs.	1	1	https://issues.jboss.org/browse/ARTIF-303
2406	ARTIF-303	[~eric.wittmann], should this have been resolved?	1	0	https://issues.jboss.org/browse/ARTIF-303
2407	ARTIF-260	[~rhauch], any thoughts/caveats on this? We're interested in full-text search, both for node properties as well as file content. But, concerned about scalability, as we have a few Artificer uses on the table that involve 10+ million primary artifacts, the bulk of which would have binary content. 1.) Am I correct that MS 4 does not currently support a query index on the file content itself? 2.) Would we be better off using a combination of MS full-text queries on metadata, then an external Lucene/ElasticSearch instance on the file contents (we use the filesystem for binary storage)? Would we be able to use the file's binary storage pathname (looks to be a UUID) to then lookup the associated nodes that reference it?	1	1	https://issues.jboss.org/browse/ARTIF-260

#	Requirement	Questions	Answered	Successfully Answered	URL
2408	ARTIF-250	Andrej Podhradsky <apodhrad@redhat.com> made a comment on [bug 1014005|https://bugzilla.redhat.com/show_bug.cgi?id=1014005] Unfortunately, pastebin at http://pastebin.test.redhat.com/167242 is empty :(So, I add errors which I see. After importing the quickstart 's-ramp-demos-switchyard-multiapp' there are the following errors: Required Capability Missing: The "Bean" capability should be configured on the project when using this type of implementation. Required Capability Missing: The "Bean" capability should be configured on the project when using this type of implementation. Required Capability Missing: The "Camel Core (SEDA/Timer/URI)" capability should be configured on the project when using this type of binding. Required Capability Missing: The "Camel Route" capability should be configured on the project when using this type of implementation. Required Capability Missing: The "SOAP" capability should be configured on the project when using this type of binding. Required Capability Missing: The "SOAP" capability should be configured on the project when using this type of binding. Unresolvable Bean Class: "org.switchyard.quickstarts.demo.multiapp.service.InventoryServiceBean" cannot be resolved for "InventoryService". Unresolvable Bean Class: "org.switchyard.quickstarts.demo.multiapp.service.OrderServiceBean" cannot be resolved for "OrderService". Is this fixed or is it a new bug?	0	0	https://issues.jboss.org/browse/ARTIF-250
2409	ARTIF-209	Stefan Bunciak <sbunciak@redhat.com> made a comment on [bug 979958|https://bugzilla.redhat.com/show_bug.cgi?id=979958] I've installed FSW 6.0.0.ER4 (using the installer). The configuration/standalone.xml still doesn't contain srampDS neither is sramp cache configured to use jdbc-store instead of file-store. Do I miss something or why is it still not fixed?	0	0	https://issues.jboss.org/browse/ARTIF-209
2410	ARTIF-196	I think ModeShape already has the contentHash, would it make sense you use that?	0	0	https://issues.jboss.org/browse/ARTIF-196
2411	ARTIF-189	Hey [~kcbabo], any info available on this? I'm not having any luck finding details. Thanks!	0	0	https://issues.jboss.org/browse/ARTIF-189
2412	ARTIF-159	Can this task be extended to enable auditing data to be recorded through the REST API. Usecase: the governance workflow defines multiple approval steps before an artifact can be promoted into production. The actual audit change to the artifact, once that decision has been made, will potentially only be associated with the last person to approve the artifact, whereas for compliance purposes, it may be necessary for all people who approved the decision to be recorded in the audit logs. Therefore in straightforward cases we can rely on the actual modifications made to the artifacts in the repository to be enough of an audit. However in cases where this is not sufficient, we need the workflow to be able to explicitly add extra audit entries.	0	0	https://issues.jboss.org/browse/ARTIF-159
2413	ARTIF-117	[~eric.wittmann], do you still think this is necessary, since we're using ModeShape's query indexing?	0	0	https://issues.jboss.org/browse/ARTIF-117
2414	ARTIF-113	We should probable default to looking for "artifact in the repository" by checking the name, file size, version (if available) and hash key. If all three match, then link to that artifact rather then adding a new one. Do we need to propose hash key to be on the baseArtifactType in the spec?	1	1	https://issues.jboss.org/browse/ARTIF-113
2415	ARTIF-113	Yes DocumentType: http://governance.github.com/s-ramp/spec/s-ramp-v1.0/documentation/coremodel.html#Link11 The same place where we store fileSize and contentType. Does that mean we need an extendDocumentType? Or else we won't store with fileSize and fileHash and contentType for extended Artifacts.	1	1	https://issues.jboss.org/browse/ARTIF-113
2416	ARTIF-86	Are the derived artifact handlers pluggable?	1	1	https://issues.jboss.org/browse/ARTIF-86
2417	ARTIF-47	The name is left blank now if the slug is not provided. Should we default it to 'newfile'? We can guess an extension based on the mimeType and s-ramp model used for the upload?	1	1	https://issues.jboss.org/browse/ARTIF-47
2418	ARTIF-47	If the Slug is not present, a reasonable name (e.g. newartifact.wsdl) will be generated for it.	0	0	https://issues.jboss.org/browse/ARTIF-47
2419	ARTIF-40	Can this be closed? I think this all works no?	1	1	https://issues.jboss.org/browse/ARTIF-40
2420	ARTIF-16	Hi [~brmeyer]! Could you please have a look and provide a feed-back on the UI Mockup? Thank you!	1	1	https://issues.jboss.org/browse/ARTIF-16
2421	ARTIF-16	(CC [~eric.wittmann] & [~objectiser]) [~bouskaj], hello! Glad to have your help! The approach looks like a great start! You've encompassed the 2 main use cases: uploading an artifact to S-RAMP, and downloading artifacts from S-RAMP and incorporating them into your project. Here's a few comments: - When you upload an artifact, it would be useful to use a popup screen where the user could provide artifact metadata (name, artifact type, properties, etc.). - Your S-RAMP repo screen currently has 1 "Query" box. Definitely keep that -- that will allow developers to use the S-RAMP query language. However, we might also need a sidebar with an advanced search form, really similar to what the web UI has. That way, users can still filter the artifacts, even if they don't know the full query language. Other than that, you're on the right track! Thanks again for the help!	1	1	https://issues.jboss.org/browse/ARTIF-16
2422	ARTIF-16	Hello [~brmeyer]. I made changes according to your advice. Could you have a look on the new UI Mockup? Thank you.	1	1	https://issues.jboss.org/browse/ARTIF-16
2423	ARTIF-16	Hi [~brmeyer]! The 2 main use cases for the Eclipse plugins are basically finished - you can give it a try https://github.com/bouskaJ/-S-RAMP_repository_browser/tree/master/PoC. However, we are struggling with dependencies inclusion. As you will notice, the only dependency is s-ramp-client, but it's transitive deps. makes the final plugin jar quite big (including almost 40 other jars). My suggestion is to create a s-ramp-client plugin that could be installed into the target platform/eclipse and S-RAMP IDE plugin could just reference it in its MANIFEST as required bundle. Any hints on this? Thanks!	1	1	https://issues.jboss.org/browse/ARTIF-16
2424	ARTIF-16	Apologies for the delayed response! {quote} My suggestion is to create a s-ramp-client plugin that could be installed into the target platform/eclipse and S-RAMP IDE plugin could just reference it in its MANIFEST as required bundle. {quote} I guess I'm not quite following as to why that would be an improvement. Doesn't that simply introduce another installation step and maintain the same overall download size (between the 2 plugins)? That being said, I'm sure I could trim down the size of s-ramp-client. It currently pulls in s-ramp-common and s-ramp-atom in their entirety -- both have fairly large transitive trees. I'll take a look under SRAMP-626.	1	1	https://issues.jboss.org/browse/ARTIF-16
2425	ARTIF-16	Can you bump sramp.version up to 0.6.0.Final or, even better, 0.7.0-SNAPSHOT?	1	1	https://issues.jboss.org/browse/ARTIF-16
2426	ARTIF-16	I did finally have a chance to build and install the plugin. Very, very impressive for a proof of concept! Once polished, this is going to be a huge win for the project and a really useful tool. Good work! Are you guys interested in any feedback, yet? Really, it's just a few nitpicky improvements for the look and feel. Functionality wise, this is great. Once you're satistifed with where it's at, I'd say let's go ahead and integrate it as a module in the actual project, but within a branch until it's "finished".	1	1	https://issues.jboss.org/browse/ARTIF-16

#	Requirement	Questions	Answered	Successfully Answered	URL
2427	ARTIF-16	SRAMP-626 is going to be a big deal of help. Thanks! {quote} Can you bump sramp.version up to 0.6.0.Final or, even better, 0.7.0-SNAPSHOT? {quote} Sure. {quote} Are you guys interested in any feedback, yet? Really, it's just a few nitpicky improvements for the look and feel. Functionality wise, this is great. {quote} Of course we are ;-) {quote} Once you're satistifed with where it's at, I'd say let's go ahead and integrate it as a module in the actual project, but within a branch until it's "finished". {quote} I suppose we could consider it finished later this month. It depends on how much feedback will you give us, but basically from functionality perspective I think we are done for now. So, once we incorporate your feedback we can send a pull request.	1	1	https://issues.jboss.org/browse/ARTIF-16
2428	ARTIF-16	Sounds good. Can you guys bump the version up to 0.7.0-SNAPSHOT and re-test? I hit at least one error -- attempting to download a WSDL and import it into a project resulted in an exception on the server. When you're satistifed, I'll give it another in-depth go. Thanks!	1	1	https://issues.jboss.org/browse/ARTIF-16
2429	ARTIF-16	Version updated, wsdl issue fixed (problem was when importing derived artifacts). You can give it another try ;-) Do you have some estimate when SRAMP-626 could be fixed? I suppose you target it to 0.7.0, right?	1	1	https://issues.jboss.org/browse/ARTIF-16
2430	ARTIF-16	[~sbunciak] / [~bouskaj], do you guys have any desire to do a demo webcast with me and other members of the team? Would love to give you the chance to show off your hard work. In the meantime, in the middle of reviewing the pull request. This looks great!	1	1	https://issues.jboss.org/browse/ARTIF-16
2431	BLACKTIE-440	This issue happens on reading disconnect receipt when service dispatcher close receiver socket {code} 2013-03-16 10:01:39,175 [main] TRACE (server.ServiceDispatcher :116) - closing: ConvService-Dispatcher-0 2013-03-16 10:01:39,175 [main] TRACE (server.ServiceDispatcher :118) - Closing receiver 2013-03-16 10:01:39,175 [main] DEBUG (stomp.StompReceiverImpl :131) - close 2013-03-16 10:01:39,175 [main] DEBUG (stomp.StompReceiverImpl :136) - closing socket: Socket[addr=localhost/127.0.0.1,port=61613,localport=39235] 2013-03-16 10:01:39,176 [main] DEBUG (stomp.StompManagement :28) - close 2013-03-16 10:01:39,176 [main] TRACE (stomp.StompManagement :70) - Writing on: java.net.SocketOutputStream@3851af4e 2013-03-16 10:01:39,176 [main] TRACE (stomp.StompManagement :91) - Wrote on: java.net.SocketOutputStream@3851af4e 2013-03-16 10:01:39,176 [main] DEBUG (stomp.StompManagement :35) - Sent disconnect 2013-03-16 10:01:39,176 [main] TRACE (stomp.StompManagement :96) - Reading from: java.net.SocketInputStream@7da22e4a {code} hanging on read socket input stream as it looks like stomp-connect dose not response the disconnect command. Tom, can you config the job to archive the jboss-as/standalone/logs/*.log ? I just close the issue as there are two success builds with no changes. feel free to reopen it if this happens again and hope to find more details with the jboss-as log.	1	0	https://issues.jboss.org/browse/BLACKTIE-440
2432	BLACKTIE-434	Hi Amos, Please can you take a look at this, for some reason nothing is working any more! Thanks, Tom	1	1	https://issues.jboss.org/browse/BLACKTIE-434
2433	BLACKTIE-434	Tom, can you dropbox the last failing jobs logs ? Amos	1	1	https://issues.jboss.org/browse/BLACKTIE-434
2434	BLACKTIE-434	[~zhfeng], {quote} can you dropbox the last failing jobs logs ? {quote} Do you still need the logs? I have put them on dropbox for you anyway. Thanks, Tom	1	1	https://issues.jboss.org/browse/BLACKTIE-434
2435	BLACKTIE-434	Thanks Tom, It looks like linux32 build is still failing, http://172.17.131.2/job/blacktie-linux32/898/ can you dropbox it ? Amos	1	1	https://issues.jboss.org/browse/BLACKTIE-434
2436	BLACKTIE-414	Great work Amos, are you going to try a pull request and check the output and see what we get? It might be worth rdesktoping onto taco and launching an old JBoss first to make sure it gets killed?	1	1	https://issues.jboss.org/browse/BLACKTIE-414
2437	BLACKTIE-414	Tom, can we add the build parameter "pullNumber" both on btny-pulls-blacktie and btny-pulls-blacktie-windows2008 ? also btny-pulls-blacktie-poller need to "wget btny-pulls-blacktie and wget btny-pulls-blacktie-windows2008" ?	0	0	https://issues.jboss.org/browse/BLACKTIE-414
2438	BLACKTIE-402	Hi Mike, Please can you comment on whether you think this will be quicker than REST-TX? I am trying to reduce the complexity of the client library. Wondering about the 5kloc of mongoose.c we have in our repo and maintainance of it. Assuming you think its viable, please can you hand it over to Amos? Tom	1	1	https://issues.jboss.org/browse/BLACKTIE-402
2439	BLACKTIE-402	Hi Tom, what is a binary protocol your thinking ? does it use between client and server or between blacktie and TM ? Amos	1	0	https://issues.jboss.org/browse/BLACKTIE-402
2440	BLACKTIE-399	Thanks Amos, much appreciated. Can you link to a git hub branch so we can see this work evolve?	1	1	https://issues.jboss.org/browse/BLACKTIE-399
2441	BLACKTIE-399	Hi Amos, How is this coming along? Do you have a feel for how much longer it may take, if so, please can you update the Jira with an estimate? Thanks! Tom	1	1	https://issues.jboss.org/browse/BLACKTIE-399
2442	BLACKTIE-399	Tom, 1. finish the work on SocketServer 2. working on the socket endpoint to send/recv messages both on client and server 3. working on the socket session to alert corba endpoint with socket endpoint 4. alert HybridConnectionImpl to use socket session 5. test with xatmi I think 1w on (2) and 1w on (3,4) but not sure on (5). I notice your comment on BLACKTIE-402 with the binary protocol. does it use in the socket endpoint ? currently, I just simply learned some pack/unpack codes from Mike's transport/src/main/cpp/HttpSessionImpl.cxx Amos	1	1	https://issues.jboss.org/browse/BLACKTIE-399
2443	BLACKTIE-399	1.update session to use socket endpoint to send/recv messages. 2.fight with memory issues and try to pass hybrid tests 3.test with xatmi 4.I'm thinking about the similar alternative implementation in jatmibroker-xatmi to remove the use of corba Tom, do you have any idea about point 4 ?	1	1	https://issues.jboss.org/browse/BLACKTIE-399
2444	BLACKTIE-399	Great! Please can you squash the commits into a single commit? If you don't feel comfortable with doing it I can take a look but it is usually easier for you to do it! Thanks again!	1	1	https://issues.jboss.org/browse/BLACKTIE-399
2445	BLACKTIE-360	In 64bit system, how can we deal with long ? man 3 endian #include "endian.h" ... uint64_t htobe64(uint64_t host_64bits); uint64_t htole64(uint64_t host_64bits); uint64_t be64toh(uint64_t big_endian_64bits); uint64_t le64toh(uint64_t little_endian_64bits); but these functions are not available on all UNIX systems even on WINDOWS.	1	1	https://issues.jboss.org/browse/BLACKTIE-360
2446	BLACKTIE-294	Hi Mike, Would it be possible for you to take a look at this for me please? Tom PS It could be the case that this is something we don't need to support but can you update the Jira with details why?	1	1	https://issues.jboss.org/browse/BLACKTIE-294
2447	BLACKTIE-255	I have added valgrind to our core build on Linux (unfortunately it is not available on windows so I had to leave that). Now, when you run a standard build (e.g. mvn install) the tests are executed under valgrind. If the tests fail, then as before the mvn process will also fail. Unfortunately this has had the not unexpected result of increasing the time of the build, on my machine it has gone up from approx 25 minutes to 50 minutes now! - I did have to increase the timeout value of a couple of tests to cope with the slow down. The other "disadvantage" is that now memory certain errors do not cause the build to fail as valgrind can cope with certain types of memory mismanagement. In order to run the build without valgrind, all you do is pass -Duse.valgrind=false (or anything other than True, true, TRUE or 1) What this does mean is that we are now in a position whereby we can easily keep track of the memory issues in the product so long as after your fixes you run the build with valgrind enabled (the default) and take a look at the outputted valgrind log (e.g. <BLACKTIE_SRC_DIR>/<PROJECT_NAME>/target/cpp-test-classes/<PROJECT_NAME>-valgrind.log to check for any new issues.	0	0	https://issues.jboss.org/browse/BLACKTIE-255

#	Requirement	Questions	Answered	Successfully Answered	URL
2448	BGBUILD-375	The systemd/systemctl issue is a much more difficult problem than it first appeared. After investigating and speaking with rwmj, due to the way libguestfs boots the images none of the system's services and/or daemons are running. This renders systemctl useless and it doesn't seem to be possible to start systemd in any useful way after booting (it generally expects to be PID#1). So, for now I shall continue to simply do the task manually by symlinking the relevant directories, but hopefully a solution might be possible (or emulable via something akin to) virt-sysprep {quote} 22:22 msavy: yeah, haven't had any success with that. i suppose the best solution for boxgrinder would be fully booting the VM independently and somehow piping the commands through to it ala libguestfs' sh/command. i realise that's outside of guestfs' purpose though. 22:22 msavy: so that "everything is there", so to speak 22:24 rwmjones: msavy: firstboot right? we added that to virt-sysprep recently 22:24 rwmjones: (although with the big caveat that it doesn't yet work for windows) 22:24 rwmjones: https://rwmj.wordpress.com/2012/08/18/new-in-virt-sysprep-add-firstboot-scripts/ 22:24 rwmjones: http://libguestfs.org/virt-sysprep.1.html#firstboot- 22:25 msavy: well, presently BG only does Fedora, RHEL and various derivatives 22:25 msavy: so the windows element wouldn't matter 22:25 rwmjones: should work for that subset then 22:25 msavy: although i am presently looking for a way to use something such as lorax instead of venerable (and badly breaking) appliance-creator 22:26 msavy: and cursing the frequent language mismatches that make life so difficult :) 22:26 rwmjones: btw I'm not necessarily suggesting using virt-sysprep; you can add the equivalent functionality to bg after examining what virt-sysprep does 22:27 msavy: indeed, thanks. {quote}	1	1	https://issues.jboss.org/browse/BGBUILD-375
2449	BGBUILD-364	redacted appliance-creator output: Installing: apr-util-devel ##################### [265/266] Installing: httpd-devel ##################### [266/266] Found GRUB Legacy package. Writing grub /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/boot/grub/device.map ls: cannot access /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/boot/initramfs*: No such file or directory Writing grub config /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/boot/grub/grub.conf Copying /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/usr/share/grub/x86_64-redhat/e2fs_stage1_5 to /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/boot/grub/e2fs_stage1_5 Copying /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/usr/share/grub/x86_64-redhat/stage1 to /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/boot/grub/stage1 Copying /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/usr/share/grub/x86_64-redhat/stage2 to /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/boot/grub/stage2 Installing grub to /dev/loop0 Probing devices to guess BIOS drives. This may take a long time. GNU GRUB version 0.97 (640K lower / 3072K upper memory) [Minimal BASH-like line editing is supported. For the first word, TAB lists possible command completions. Anywhere else TAB lists the possible completions of a device/filename.] grub> device (hd0) /dev/loop0 grub> root (hd0,1) Filesystem type is ext2fs, partition type 0x83 grub> setup --stage2=/path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/boot/grub/stage2 --prefix=/grub (hd0) Checking if "/grub/stage1" exists... yes Checking if "/grub/stage2" exists... yes Checking if "/grub/e2fs_stage1_5" exists... yes Running "embed /grub/e2fs_stage1_5 (hd0)"... 15 sectors are embedded. succeeded Unable to create appliance : Unable to install grub bootloader Unmounting directory /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/boot Unmounting directory /path/build/appliances/x86_64/centos/5/some_long_appliance_name_many_chars1/1.0/centos-plugin/tmp/imgcreate-FDcjk3/install_root/ Removing compat symlinks Unmapping /dev/loop0 Losetup remove /dev/loop0 [root@appdev04 boxgrinder]# echo $? 1 [root@appdev04 boxgrinder]#	1	1	https://issues.jboss.org/browse/BGBUILD-364
2450	BGBUILD-347	What's the current state? Can we build Fedora 17 appliances? And another thing: can we use BG on a Fedora 17 host?	1	1	https://issues.jboss.org/browse/BGBUILD-347
2451	BGBUILD-346	I'll take care about the spec file, it needs to be rewritten a bit, have a proposal locally, but waiting for two deps (httparty and crack) to push it. Which spec fails?	1	1	https://issues.jboss.org/browse/BGBUILD-346
2452	BGBUILD-272	Caveat: according to the documentation, the aws-sdk gem it only works on Ruby 1.8.7. This means that it will not work on Fedora 13, which only has 1.8.6. I don't really think this should block us, as we technically dropped Fedora 13 support for the 0.9.3 release due to EOL status.	1	1	https://issues.jboss.org/browse/BGBUILD-272
2453	BGBUILD-271	I can't find the host parameter in the history anywhere? Am I missing something here?	1	1	https://issues.jboss.org/browse/BGBUILD-271
2454	BGBUILD-267	Please use the branch BGBUILD-267 to replicate, it was trivially easy to add the basics (yay BoxGrinder!) I can't find the CentOS GRUB SRPM on their mirrors to take a look :(http://mirror.centos.org/centos/6/. The error seems to emanate from the GRUB RPM itself when installing, so I don't see how we can easily interfere with appliance creator to solve it (e.g. if we created and symlinked the mismatched directories it might work, but can we do that _before_ appliance-creator tries to install the RPMs?). Edit: Just to confirm, the shipped CentOS ISO has the directory structures as described above, and they work fine. So need to look elsewhere.	1	1	https://issues.jboss.org/browse/BGBUILD-267
2455	BGBUILD-267	Good job! Do you need intervention from my side?	1	1	https://issues.jboss.org/browse/BGBUILD-267
2456	BGBUILD-257	Do we see the time difference also for appliances built for other than EC2 formats? If yes - we need to think about more general solution, like NTP for all platforms.	0	0	https://issues.jboss.org/browse/BGBUILD-257
2457	BGBUILD-246	Could we also improve performance by allocating more memory for bigger instances?	1	1	https://issues.jboss.org/browse/BGBUILD-246
2458	BGBUILD-224	Could you please attach at least the hardware section of your appliance definition file? Thanks!	1	1	https://issues.jboss.org/browse/BGBUILD-224

#	Requirement	Questions	Answered	Successfully Answered	URL
2459	BGBUILD-222	Done including tests, just needs to be merged in from the 'force' branch. Simply use 'overwrite' to force an overwrite (latest build for a snapshot is the one overwritten) in EBS or S3 plugins. For EBS there are some special considerations and an additional option: - The plugin will terminate (fail) if you have any instances still running, in order to allow you to rescue any important transient data before clobbering the EBS volume. - The associated snapshot *will be deleted* unless you enable 'preserve_snapshots: true' in the EBS plugin config - We simply look for an EBS device on the registered EBS image we wish to overwrite called /dev/sda1 (this is where the EBS plugin created and mounted the volume initially). Do *NOT* manually replace /dev/sda1 with your own important data volume and then use overwrite. It *WILL* delete it as BoxGrinder has no way of knowing the UUID of the device between sessions, it simply uses the known device path. - The image will be deregistered for you, and the new one created with an identical name.	1	1	https://issues.jboss.org/browse/BGBUILD-222
2460	BGBUILD-221	We should allow to specify 32 bit packages. The question is - how will be the dependency resolution look like? If some package requires glibc 32 bit and we'll have 64 bit - dependency resolution will pass, but the app will not work. This needs some testing.	0	0	https://issues.jboss.org/browse/BGBUILD-221
2461	BGBUILD-211	My first idea is to upload image to selected host using SFTP plugin to /var/lib/libvirt/images/{appliance_path} and then use virt-install command to actually register the image. Is this what you guys want or you think about different process?	1	1	https://issues.jboss.org/browse/BGBUILD-211
2462	BGBUILD-211	Yeah, that would definitely satisfy what I was thinking. My goal would be to deploy appliance to a KVM/Xen box and be able to start it with either Virtual Machine Manager or virsh. As I'm looking more closely at the documentation, did I completely misunderstand how much support there is for stuff like this in VMWare/VirtualBox?	1	1	https://issues.jboss.org/browse/BGBUILD-211
2463	BGBUILD-201	After all of this effort I think I might have screwed up :-(. @mgoldmann should I have set the delete_on_termination: true in delivery options?	1	1	https://issues.jboss.org/browse/BGBUILD-201
2464	BGBUILD-166	We can set a mandatory requirement for all users like this: {code} gconftool-2 --direct --config-source xml:readwrite:/etc/gconf/gconf.xml.mandatory -s /apps/nautilus/preferences/media_automount --type bool false {code} But it'll not affect current logged in users, so it's useless. I think the right solution for this would be add the solution to [FAQ|http://boxgrinder.org/faq/] and do not modify the code because it would be one big hack. Do you agree?	0	0	https://issues.jboss.org/browse/BGBUILD-166
2465	BGBUILD-148	14:50 skvidal: gholms: okay wanna try something fun? 14:50 mgoldmann: created a placeholder for BG: https://issues.jboss.org/browse/BGBUILD-148 14:50 gholms: skvidal: What is it? 14:50 skvidal: gholms: on the Fedora system run 14:51 skvidal: db_dump /var/lib/rpm/Packages > myrpmdb.out 14:51 skvidal: that will be a BIG file 14:51 skvidal: then inside the chroot, running on el5's rpm do this 14:51 skvidal: cd /var/lib/rpm 14:51 skvidal: mv Packages Packages.bak 14:52 skvidal: cat myrpmdb.out | db_load Packages 14:52 skvidal: rpm --rebuilddb 14:52 skvidal: rpm -qa 14:52 skvidal: see if that works 14:52 gholms needs to install db4-utils in the chroot 14:55 gholms: skvidal: Yeah, that works. Do you suppose I have to do the same thing for the rest of the DBs there as well? 14:55 skvidal: no 14:55 skvidal: packages is the only db 14:55 skvidal: everything else is more or less an inde 14:55 skvidal: x 14:56 gholms: Can I just nuke the rest before I db_load it? 14:56 skvidal: rpm --rebuilddb recreates the rest of them 14:56 gholms: Ah, ok. 14:56 skvidal: so it matters not 14:56 gholms: It apparently doesn't recreate them when they already exist since it complained about bad hash versions in the Pubkeys db. 14:56 skvidal: oh 14:56 skvidal: pubkeys 14:57 skvidal: is different 14:57 skvidal: :) 14:57 skvidal: pubkeys is the gpg pub keys database 14:57 gholms: Completely separate from the rpmdb? 14:57 skvidal: sort of 14:57 skvidal: you can just do the same thing with that db, too 14:57 skvidal: but that really should be it 14:57 jhford has left IRC (Excess Flood) 14:58 gholms: If I just nuke that db then it just makes me re-import the keys later, right? 14:58 skvidal: should, yep 14:58 zooz has joined (~zooz@host86-163-127-233.range86-163.btcentralplus.com) 14:58 zooz: hi people 14:58 nirik sighs at his 6rd tunnel. POS. I guess I will stop trying to get it working again for now. 14:58 skvidal: but I've not actually tested that	1	1	https://issues.jboss.org/browse/BGBUILD-148
2466	BGBUILD-143	h4. Delivery Plugins * (i) In BG<=0.8.0 environment variables cannot be used in paths for delivery plugins in ~/.boxgrinder/config [BGBUILD-160] * Hash checks are performed to determine whether the file arrived intact or not (or some equivalent check, depending on the delivery method) h5. Local File h6. Packaged (Compressed) {code} plugins: local: path: /root/appliances overwrite: false package: true {code} ||OS||Version||Delivery||Hash Check||Notes|| |Fedora|13|SUCCESS|PASS| https://gist.github.com/833382 | |Fedora|14|SUCCESS|PASS| https://gist.github.com/833317 | |CentOS|5.5|SUCCESS|PASS| https://gist.github.com/833157 | h6. Unpackaged (Uncompressed) {code} plugins: local: path: /root/appliances overwrite: false package: false {code} ||OS||Version||Delivery||Hash Check||Notes|| |Fedora|13|SUCCESS|PASS| https://gist.github.com/842397 | |Fedora|14|SUCCESS|PASS| https://gist.github.com/843036 | |CentOS|5.5|SUCCESS|PASS| https://gist.github.com/842625 | (i) All working correctly as of 0.8.1 h5. SFTP {code} plugins: sftp: host: 192.168.0.40 username: boxgrinder path: /home/boxgrinder/appliances {code} BG only supports key authenticated access. ||OS||Version||Delivery||Hash Check||Notes|| |Fedora|13|SUCCESS|PASS| https://gist.github.com/842401 | |Fedora|14|SUCCESS|PASS| https://gist.github.com/833316 | |CentOS|5.5|FAIL (!) |FAIL (!) Remote file created, but is zero bytes| See [BGBUILD-162] | h5. S3 {code} plugins: s3: bucket: msavy-test-20101101-a access_key: XXXXXXXXXXXXXXX secret_access_key: XXXXXXXXXXXXXXX account_number: XXXXXXXXXXXXXXX cert_file: ~/.ec2/cert-XXXXXXXXXXXXXXX.pem key_file: ~/.ec2/pk-XXXXXXXXXXXXXXX.pem {code} ||OS||Version||Delivery||Hash Check||Notes|| |Fedora|13|SUCCESS|PASS| https://gist.github.com/842593 | |Fedora|14|SUCCESS|PASS| https://gist.github.com/833321 | |CentOS|5.5|SUCCESS|PASS| https://gist.github.com/833363 | h5. AMI ||OS||Version||Delivery||Hash Check||Notes|| |Fedora|13|SUCCESS|PASS| https://gist.github.com/842828 | |Fedora|14|SUCCESS|PASS| https://gist.github.com/842995 | |CentOS|5.5|SUCCESS|FAIL(!) ^1| [BGBUILD-169] | ^1 May be a new problem since 0.8.1, or problem specific to testing virtual image /hardware (needs verification) h5. EBS EBS testing pending (requires image running in EC2 to test, will be completed with AMI tests).	1	1	https://issues.jboss.org/browse/BGBUILD-143
2467	BGBUILD-77	@Bright - I'm aware of the post section capabilities. We use it pretty extensively for our SteamCannon appliances (http://github.com/steamcannon/steamcannon-appliances). However, when the work that needs to be accomplished in the post section is a bit more complex, the post section is a little limiting. Also possible that you can do more with YAML block literals than I am aware of. I'd like to be able to use conditionals, set variables, etc. Is it possible for example to have something like this? {code} post: base: -| if [-e /path/to/file] ; then $RUBY_CMD $RUBY_ARGS else $OTHER_CMD $OTHER_ARGS fi {code}	0	0	https://issues.jboss.org/browse/BGBUILD-77

#	Requirement	Questions	Answered	Successfully Answered	URL
2468	BGBUILD-73	Defaults can be set for many things but there is still problem with OS name and version. There is no section in kickstart files that describes the operating system we're installing. Or at least I'm not familiar with such one. I'm using in every plugin OS name/version combination for creating paths, etc. What options we have? 1. Comments in .ks files, like this: # bg_os_name: fedora # bg_os_version: 14 Seth - take a look at every init.d script - it has chkconfig control structures in comments. 2. Try to parse repo sections and read the OS version and name. But this approach may be fragile, because sometimes we'll not have the OS name included in repo path and sometimes we'll have more than one. 3. Other?	1	1	https://issues.jboss.org/browse/BGBUILD-73
2469	BGBUILD-69	Is this a dupe of BGBUILD-60?	1	0	https://issues.jboss.org/browse/BGBUILD-69
2470	BGBUILD-47	Sanne, Yah, I'm considering adding selection for filesystem type. What do you mean by mount options? You want to add custom entries to /etc/fstab? How it should look? What is the use case?	1	1	https://issues.jboss.org/browse/BGBUILD-47
2471	BGBUILD-47	Thanks for clearing this.The problem with the full line is that you cannot be sure what will be the device name, for one platform it'll be /dev/sda1, for other /dev/xvda1. Another approach could be that you can specify values for partitions, like this for example: "/": size: 10 type: "ext4" options: "barrier=0,nodelalloc,nobh,noatime,data=writeback,commit=60,nouser_xattr,noacl,max_batch_time=30000" If you want add additional entries (mount external filesystems) manually you can use post section: post: base: - "echo '/dev/something /mnt ext4 defaults,ro 0 0' >> /etc/fstab" WDYT?	1	1	https://issues.jboss.org/browse/BGBUILD-47
2472	BGBUILD-46	Does this mean that the resulting image will be the same arch as the host operating system atm? This seems to be implied, but is not mentioned in the existing documentation.	0	0	https://issues.jboss.org/browse/BGBUILD-46
2473	BGBUILD-28	20:40 bobmcw: generally speaking, I think I'd like the example to pull from the s3-delivery or fedora-os plugins that exist 20:40 bobmcw: reading the page, I see I need execute() but don't quite grok what I'd put there 20:40 bobmcw: then again, I'm code ignorant of BG 20:40 bobmcw: but even if psuedo-code 20:40 bobmcw: replace_kernel_with_xen() 20:41 bobmcw: or s3_upload(appliance.artifact) 20:41 bobmcw: or whatever, just to give me an idea 20:41 bobmcw: the page tells me "how", but examples of what and why I'd access @config and @appliance_config would be good 20:41 bobmcw: would light my imagination of what could be done 20:42 bobmcw: (I know we have code in github, so links might suffice)	0	0	https://issues.jboss.org/browse/BGBUILD-28
2474	BPMC-83	Kanchan Desai <kadesai@redhat.com> made a comment on [bug 958393|https://bugzilla.redhat.com/show_bug.cgi?id=958393] Hi Ivo, Doc Text has been added. Can you please review the CCFR.	0	0	https://issues.jboss.org/browse/BPMC-83
2475	BPMC-70	After working with suspend/resume for a while during testing, I found that I had to always think about what the words "true/false" in the "Suspended" column mean. Is it possible to use the words "active" and "retired" instead to avoid confusion?	0	0	https://issues.jboss.org/browse/BPMC-70
2476	BPMC-63	as far as I can see in the code it is not yet fixed. Perhaps good idea would be to make it even more flexible - for instance allow to set it as JVM parameter (-Dorg.jbpm.report.data.dir) instead of relying on catalina.home. That could make it a bit more server independent. What do you think about it?	0	0	https://issues.jboss.org/browse/BPMC-63
2477	BPMC-63	yes, you're right, it's a good idea to use a JVM parameter. I think a good name could be : -Dorg.jbpm.report.engine.dir. Is there currently a workaround ? it's quite annoying not having reports on my Linux install.	0	0	https://issues.jboss.org/browse/BPMC-63
2478	BPMC-63	I am fine with these changes. Can somebody provide a patch? That would greatly simplify my life.	0	0	https://issues.jboss.org/browse/BPMC-63
2479	BPMC-58	Which versions are the patches supposed to work with?	1	0	https://issues.jboss.org/browse/BPMC-58
2480	BPMC-37	Thanks, I've applied your patches to both codebases. Can you verify everything works as expected? If that's successful you can close this issue.	1	1	https://issues.jboss.org/browse/BPMC-37
2481	BPMC-8	Hi Heiko, Could you please help me review apply if it looks good. ;-) Thanks Jeff	1	1	https://issues.jboss.org/browse/BPMC-8
2482	BUILDTRK-230	If we have multiple jars to test, can we have a batch process for this as well? For example, upload text file with name and checksum of jars, return list of jars with their releases. The text file could be (e.g) XML, or could be sent via REST, and the return could be a web page, or REST data.	1	1	https://issues.jboss.org/browse/BUILDTRK-230
2483	BUILDTRK-204	Why should we do it? Fill the description, please.	0	0	https://issues.jboss.org/browse/BUILDTRK-204
2484	BUILDTRK-178	I think, this method could be only overloaded. Is there any reason to not let there this method?	0	0	https://issues.jboss.org/browse/BUILDTRK-178
2485	BUILDTRK-159	Couldn't Deltaspike @ConfigProperty be used?	0	0	https://issues.jboss.org/browse/BUILDTRK-159
2486	BUILDTRK-91	Can you describe this issue more?	1	1	https://issues.jboss.org/browse/BUILDTRK-91
2487	BUILDTRK-81	Can somebody reproduce it?	0	0	https://issues.jboss.org/browse/BUILDTRK-81
2488	BUILDTRK-40	Could you describe it more?	0	0	https://issues.jboss.org/browse/BUILDTRK-40
2489	BUILDTRK-35	Which tab in UI do you have in mind? Project->Release->Dependent Releases? I saw there some code, which looks good to me. Maybe some bug in backend? I remeber, that it worked - in past there were 2 dependent releases for Hibernate 4.2.3, but now when I run it in debugger the DAO method (ReleaseDAOImpl: findDependentReleases) returned 0 instances. I don't understand the DAO method... Why are there these conditions d1.releaseArtifact = true, d2.releaseArtifact = false?	1	1	https://issues.jboss.org/browse/BUILDTRK-35
2490	BUILDTRK-29	This methods are used just for projects, statuses and licenses. getAllProjects() is used for title page - Me and jsenko thinks that it shouldn't be paginated. Currently it looks like for licenses and statuses pagination is not needed, but can be added in future Pagination was creasted for list of releases (in projects page)	0	0	https://issues.jboss.org/browse/BUILDTRK-29
2491	BUILDTRK-12	Kuba, can you please comment on this?	0	0	https://issues.jboss.org/browse/BUILDTRK-12
2492	BUILDTRK-8	Would it possible to open a REST API to list all components, for each component, with the information: {name, version, groupId}, groupId should be available if the component is NOT a native component? Thank you. :-)	1	1	https://issues.jboss.org/browse/BUILDTRK-8
2493	BUILDTRK-8	All data in BuildTracker can be accessible via REST API. We just need to gather requirements for it to design it. => Yes, we can provide you any information you want. To your requirement (we need more precise specification) : Would you like to list all projects with groupId and artifactId? To every project get all releases and to every release get all artifacts?	1	1	https://issues.jboss.org/browse/BUILDTRK-8
2494	BXMSDOC-1613	Is it ok to close this or will there be an update with the examples that I should review?	1	1	https://issues.jboss.org/browse/BXMSDOC-1613

#	Requirement	Questions	Answered	Successfully Answered	URL
2495	BXMSDOC-1613	Ah, very helpful, [~Rikkola]! Thanks. I think these examples are actually worth including in the doc. I'll work on that and send back to you one last time for final OK. However, meanwhile, I'm not quite clear how Resolved Hit works in reality. You say it's the same as First Hit, but you can mix the order of rows. And in the description in the doc, we say for example you can give priority to row 10 over row 5, or whatever. But how do you actually give that priority? Just with a separate rule somewhere that you can create to determine that? Or is it a setting with the table itself? Thanks, Stetson	1	1	https://issues.jboss.org/browse/BXMSDOC-1613
2496	BXMSDOC-1613	Okay [~Rikkola], so I've updated the doc with a new "Hit Policy Examples" section based on what I've understood from your excellent examples. I went a simpler route that I hopefully is still accurate. Here's the link: http://file.rdu.redhat.com/~sterobin/BXMSDOC-1366/#hit_policies_concept Could you please review that Examples section and let me know two things in particular: - Is it all technically accurate? - How does the user specify row/hit priority with the Resolved Hit policy? What I've written about this is still incomplete because I myself don't know. See my comment above. Here's the specific section related to my Resolved Hit question: --- First Hit/Resolved Hit: With both First Hit and Resolved Hit policies, only one of the discounts can be applied. For First Hit, the discount that is satisfied first in the list is applied and the others are ignored. Example: A senior citizen who is also a student and a military veteran will receive only the senior citizen discount of 10%, since that is listed first in the table. For Resolved Hit, the discount that you give priority to, regardless of listed order, will be applied first. *+<How do you set this priority? A separate rule somewhere? In the table itself somehow?>+* Example: If military discounts are prioritized higher than age or student discounts, despite the listed order, then a senior citizen who is also a student and a military veteran will receive only the military discount of 10%, regardless of age or student status. Key difference: With First Hit, the listed order of rules is significant, while with Resolved Hit, listed order is irrelevant. ---	1	1	https://issues.jboss.org/browse/BXMSDOC-1613
2497	BXMSDOC-1579	[~rbarrius] [~nmirasch] Could you please review the changes I have incorporated from Maciej's community docs? You can find the diff here: https://gitlab.cee.redhat.com/red-hat-jboss-bxms-documentation/bxms-docs/commit/ef595325e6310341d2a6c413c3b3b3dc9e02bf08 and view in jenkins here: http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BPMS/job/doc-BPMS-Administration_And_Configuration_Guide-branch-BXMSDOC-1430Reopened/lastSuccessfulBuild/artifact/index.html#chap-execution-error-management Thanks, Gemma.	1	1	https://issues.jboss.org/browse/BXMSDOC-1579
2498	BXMSDOC-1579	Hi Gemma, if you are done incorporating the review comments, can we close this sub-task? I see that the main Jira BXMSDOC-1430 is marked to resolved.	1	1	https://issues.jboss.org/browse/BXMSDOC-1579
2499	BXMSDOC-1485	[~ksuta] thank you. Is it simply a matter of replacing executeCommands() with executeCommandsWithResults() or do the examples need further changes?	1	1	https://issues.jboss.org/browse/BXMSDOC-1485
2500	BXMSDOC-1485	[~ksuta] I've updated the example. Would you please confirm. This command is still in the second example: public void executeCommands() { Is that right?	1	1	https://issues.jboss.org/browse/BXMSDOC-1485
2501	BXMSDOC-1462	I guess we can mention also about multi-module projects [1] point 3. What do you think [~emmurphy], [~tkobayashi]? [1] https://github.com/kiegroup/kie-wb-common/pull/836#issuecomment-299555115	1	1	https://issues.jboss.org/browse/BXMSDOC-1462
2502	BXMSDOC-1452	[~marek.winkler]- I think, and maybe [~cwatkins] can correct me if I am wrong, that Prakasha had mentioned that BBVA, for whom this LA is planned at, is a fresh BxMS customer and won't be using 6.x and hence won't need migration. Did we miss something?	1	0	https://issues.jboss.org/browse/BXMSDOC-1452
2503	BXMSDOC-1447	Thank you [~ksuta], 1. Incorporated (together with other comments in the MR - you can re-check the MR and the 7.0 backport too) 2. The payload has the following line: {{<scanner poll-interval="5000" status="CREATED"/>}}. Do I need to add further ruleconfig config? (see the docs for the full XML payload) I just feel strange putting scanner with status "CREATED" into the payload and then seeing {{<scanner poll-interval="0" status="STOPPED"/>}} in the response in the official documentation...	1	1	https://issues.jboss.org/browse/BXMSDOC-1447
2504	BXMSDOC-1379	Hello [~vidya_iyengar], regarding your question about BPMS and BRMS distinction: in BPMS we need all points 1-5 mentioned in the description. in BRMS we only need points 1, 3 and 5 (NOT 2 and NOT 4), because there are no "jBPM Runtime"-related perspectives in BRMS (Process instances, tasks, dashboard). Moreover the login-module configuration *must not* be included in the standalone.xml because it would be impossible to even login to BRMS business-central (because the implementation for the login-module is only included in BPMS). Why is login-module needed in BPMS? Because those runtime perspectives are using the *currently logged in user* to load data from kie-server (it must be checked, that this user actually has access to see process definitions/instances/tasks from given container on kie-server etc). It ensures that currently logged in user in business central gets automatically logged into kie-server as well (just by virtue of being logged into business-central), so he can retrieve data from it.	1	1	https://issues.jboss.org/browse/BXMSDOC-1379
2505	BXMSDOC-1379	Hi [~jhrcek], Can you tell me if this fix if for 7.0-LA or 7.0-GA?	1	1	https://issues.jboss.org/browse/BXMSDOC-1379
2506	BXMSDOC-1346	Hi [~ksuta], did you mean the first endpoint in the [Unmanaged Intelligent Process Server Environment|http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BPMS/job/doc-BPMS-Development_Guide-branch-BXMSDOC-1297/lastSuccessfulBuild/artifact/index.html#unmanaged_intelligent_process_server_environment] (updated). [~anton.giertli], Check out [Unmanaged Intelligent Process Server Environment|http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BPMS/job/doc-BPMS-Development_Guide-branch-BXMSDOC-1297/lastSuccessfulBuild/artifact/index.html#unmanaged_intelligent_process_server_environment]. I have: * Linked to Chapter 16. The REST API for Intelligent Process Server Execution, where we explain using curl. * Updated Chapter 16 so as to mention the X-KIE-ContentType header. * In [POST] /config, provided examles of curl commands that work with each marshaller type. WDYT?	1	1	https://issues.jboss.org/browse/BXMSDOC-1346
2507	BXMSDOC-1346	[~ksuta], my apologies, should have added it before. How about [now|http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BPMS/job/doc-BPMS-Development_Guide-branch-BXMSDOC-1297/lastSuccessfulBuild/artifact/index.html#unmanaged_intelligent_process_server_environment]?	1	1	https://issues.jboss.org/browse/BXMSDOC-1346
2508	BXMSDOC-1238	[~ksuta], please check the updated docs; I added the information into the section _Persistence Setups - Troubleshooting_ of the Installation Guide: http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BPMS/job/doc-BPMS-Installation_Guide-branch-BXMSDOC-1238/lastSuccessfulBuild/artifact/index.html#special_setup_for_ibm_db2_database Let me know if everything is correct. Also, should this be available in the BRMS version of the Installation Guide, or is this applicable only to BPMS?	1	1	https://issues.jboss.org/browse/BXMSDOC-1238
2509	BXMSDOC-1238	BPMS is enough, this is a jBPM issue. Just a question about point 4 - why do you specifically mention there those insert statements? They are part of postgresql-jbpm-lo-trigger-clob.sql script.	1	1	https://issues.jboss.org/browse/BXMSDOC-1238

#	Requirement	Questions	Answered	Successfully Answered	URL
2510	BXMSDOC-1214	[~ksuta]- Will fix version for this be 6.4.3?	1	1	https://issues.jboss.org/browse/BXMSDOC-1214
2511	BXMSDOC-1183	[~vidya_iyengar]- Could you please add all deets here? TIA!	1	1	https://issues.jboss.org/browse/BXMSDOC-1183
2512	BXMSDOC-1113	Vidya, thank you for trusting me with the review. Here are my suggestions, and again, feel free to ignore them if you do not find them helpful: * _KIE Server and Container Migration_, _the KIE server templates_ -> are you sure that it should be "KIE server"? Shouldn't it be Intelligent Process Server or Realtime Decision Server? Please let me know, I am very curious about this since I though we should not use "KIE server" at all (?) * _system git repository_ -> capitalization: +Git+ repository * _6.1 to 6.3 (or above)_ -> avoid using brackets as much as possible. I think saying _6.1 to 6.3 or above_ without them is fine. * _So to retain the container data while migrating from Red Hat JBoss BPM Suite 6.1 to 6.3 (or above) , you must first_ -> I have more suggestion to this sentence. I would omit the word "so" (it sounds too informal to me in this context) and I think it is superfluous to repeat again the phrase "migrating from Red Hat JBoss BPM Suite 6.1 to 6.3 (or above)", because the exactly same wording is there before. Also, I think it is superfluous to say "you must", instead of that, I would just say "upgrade to". There is a spare space before one comma too. So, the sentence would be: _To retain the container data, upgrade to Red Hat JBoss BPM Suite 6.2 and then to Red Hat JBoss BPM Suite 6.3._ It is much shorter and it says the exactly same thing. That's all, I hope it is helpful! Klara	1	1	https://issues.jboss.org/browse/BXMSDOC-1113
2513	BXMSDOC-1075	Hi [~mczernek], I have a few minor comments: 1)I'd suggest not to use WIH abbreviation at all for WorkItemHandler. Also, it is not the first occurrence of "WorkItemHandler" in this topic. If u meant Work Item Handler, then it is ok. 2)With the following pom.xml: It is a split sentence but acting like an independent sentence. Can we say "Use the following pom.xml" instead? 3)Also, if the context of 'Asynchronous' is not command, then it should be "+a+synchronous" and not "+A+synchronous". For example in the first sentence, When the workflow reaches an Asynchronous WorkItemHandler (WIH). Here, it should be "asynchronous work item handler" -Same for Work Item Definition. If you mean it as the name of the editor, and it it appears as "Work Item Definition" in the UI, it is fine. Otherwise, it should be "work item definition". 4)" In the same project, create a Business Process. Click New Item → Business Process to do so. " Can simply say Click New Item → Business Process to create a create a Business Process. 5)Step 8: b) is not something that you are entering. It is a save button. You can get rid of a) and b) and just say "Enter the value of CommandClass as ..... and click Save". 6)Step 9 : Save your process in the process designer. Click Save. SImply say Click Save to save your process.	1	1	https://issues.jboss.org/browse/BXMSDOC-1075
2514	BXMSDOC-1073	Also is "Authorization: Bearer 123" literal value of the property to be set, or should 123 be replaced by some placeholder? How should user know what to use instead of that "123"?	1	1	https://issues.jboss.org/browse/BXMSDOC-1073
2515	BXMSDOC-1073	Hi Jan, could you re-check the link above? I have re-written the note (in fact, it is not a note any longer). Do you think it's clearer now?	1	1	https://issues.jboss.org/browse/BXMSDOC-1073
2516	BXMSDOC-1057	Hi, I deleted the REST API from [User Guide|http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BRMS/job/doc-BRMS-User_Guide-branch-BXMSDOC-1057/lastSuccessfulBuild/artifact/index.html]. Can we merge and close this issue?	1	1	https://issues.jboss.org/browse/BXMSDOC-1057
2517	BXMSDOC-1038	bq. Chapter 14. Integration with Aries Blueprint bq. bq. This section explains the integration elements of Red Hat JBoss BPM Suite specific to Apache Aries Blueprint. section --> chapter bq. The <kie:kmodule> element defines a collection of KieBase and associated KieSessions. The id parameter is required. and associated KieSessions --> and it's associated KieSessions bq. Set scope to singleton to use the same bean instnace every time the bean is called. instnace --> instance bq. Specifies the reference to the event listeners group. For more information, see the section called “Defining a Group of Listeners”. see the section called "Defining.." --> see the "Defining.." section, or even just "see .." bq. Scanning Interval in milli seconds. milli seconds --> milliseconds bq. ...only KIE objects identified by referenced releaseId element are inicialized and imported into the Blueprint context. inicialized --> initialized referenced releaseId element --> a referenced releaseId element bq. You can enable the KIE scanning feature, that is enableScanner,... bq. The tags have identical names to the listener interfaces, that is: you can afford to drop "that is" bq. The kie-aries-blueprint module allows you to configure the loggers using XML tags using identical names, that is: replace second "using" with "with"? bq. Anonymous bean. :)	0	0	https://issues.jboss.org/browse/BXMSDOC-1038
2518	BXMSDOC-1026	Assigning to [~marek.winkler] Hi Marek, could you please confirm links to the correct version of Integration Stack for BxMS 6.3: http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BPMS/job/doc-BPMS-Getting_Started_Guide-branch-BXMSDOC-987/lastSuccessfulBuild/artifact/index.html#installing_the_jboss_developer_studio_plug_ins http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BRMS/job/doc-BRMS-Getting_Started_Guide-branch-BXMSDOC-987/lastSuccessfulBuild/artifact/index.html#installing_the_jboss_developer_studio_plug_ins And 6.4: http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BPMS/job/doc-BPMS-Getting_Started_Guide-branch-BXMSDOC-987-6.4/lastSuccessfulBuild/artifact/index.html#installing_the_jboss_developer_studio_plug_ins http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BRMS/job/doc-BRMS-Getting_Started_Guide-branch-BXMSDOC-987-6.4/lastSuccessfulBuild/artifact/index.html#installing_the_jboss_developer_studio_plug_ins Thanks, Gemma.	1	1	https://issues.jboss.org/browse/BXMSDOC-1026
2519	BXMSDOC-1013	Hi [~jschwan], Can you please take a look at the changes? Thanks again!	1	1	https://issues.jboss.org/browse/BXMSDOC-1013
2520	BXMSDOC-1013	hi [~jschwan], I have fixed the first two feedback that you gave, however I am not sure what you meant by "If scanner is started for a container and is deployed the new version of your deployment, then your container will be automatic update by scanner." Can you please explain/elaborate a bit more on this? Thanks!	1	1	https://issues.jboss.org/browse/BXMSDOC-1013
2521	BXMSDOC-986	Hi, I have added the following section into the docs: Asynchronous WorkItemHandler Hello World Implementation. Before we make any further changes, could you please check and go through the new section?	1	1	https://issues.jboss.org/browse/BXMSDOC-986
2522	BXMSDOC-986	Hi [~mczernek], great job! Just some minor things: {noformat} 2. Package your implementation into a JAR file, upload it in the Business Central, and add into your project dependencies. {noformat} Maybe something like "Build your maven project, upload resulted JAR to Business Central and add it into your project dependencies." will be better. {noformat} 8. Enter the following values: a. Task name: async. b. CommandClass: docs.command.HelloWorldCommand. {noformat} "TaskName" - is property name, the same as "CommandClass" in point b. so in my opinion it's better to write it in consistent way. And I guess periods are little bit dangerous in this place (some one may use copy-paste method to test how it's work). {noformat} 10. Register AsyncWorkItemHandler in Business Central: b. Click Add under the Work Item handlers category and add the following: async: new org.jbpm.executor.impl.wih.AsyncWorkItemHandler(org.jbpm.executor.ExecutorServiceFactory.newExecutorService(null)) {noformat} I guess it's will be easier to understand when this part will be like "set Name to 'async' and identifier to 'new org.jbpm.executor.impl.wih.AsyncWorkItemHandler(org.jbpm.executor.ExecutorServiceFactory.newExecutorService(null))'"	1	1	https://issues.jboss.org/browse/BXMSDOC-986

#	Requirement	Questions	Answered	Successfully Answered	URL
2523	BXMSDOC-981	Hi [~psiroky], we have a clustering section in our Install Guide. Now, because BxMS for EAP 7.0 won't have a JAR installer, afaik, I have done the following change--see the docs [here|http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20BPMS/job/doc-BPMS-Installation_Guide-branch-BXMSDOC-981/lastSuccessfulBuild/artifact/index.html#clustering_custom_configuration], ctrl+f -> "Furthermore, when deploying on EAP 7.0," That's the small addition I made. Is this correct and sufficient in your opinion?	1	1	https://issues.jboss.org/browse/BXMSDOC-981
2524	BXMSDOC-981	Could you re-check the link above to see the changes? It should now be more explicit. I didn't think it that important before so I didn't include the original value but as per request, it's now in the draft.	1	1	https://issues.jboss.org/browse/BXMSDOC-981
2525	BXMSDOC-870	Hello [~mmacik], could you please doublecheck this? The doc text added seems OK, but I'm not really sure if this shouldn't go to WLS guide instead of WAS guide. Searching in the community sources I only found a mention of "org.uberfire.start.method" in the [README of weblogic assembly| https://github.com/droolsjbpm/kie-wb-distributions/blob/6.5.x/kie-wb/kie-wb-distribution-wars/src/main/assembly/weblogic12/README.md#jvm-custom-properties] and nothing for WAS.	1	1	https://issues.jboss.org/browse/BXMSDOC-870
2526	BXMSDOC-849	I have tweaked the table formatting. Does it look better now? I'm afraid I can't make it much better without a major overhaul.	0	0	https://issues.jboss.org/browse/BXMSDOC-849
2527	BXMSDOC-822	Marek, very good job! My suggestions are as follows, feel free to ignore them if you don't find them helpful: * which is called kie-spring-VERSION-redhat-MINORVERSION.jar. --> should be _VERSION_ and _MINOR_VERSION_ (italics and with underscore) * "How you intend to use the Spring modules in your application affects how you configure them." --> for some reason, I don't like this sentence. I would rewrite it to: "You can configure the Spring modules in the following ways:" or something similar. * "single runtime manager instance use" --> comma is missing? --> "single runtime manager instance, use" * "the jBPM services." --> is that correct? Can we use the term jBPM? * "include following factory beans" --> "include _the_ following factory beans"? * "TaskServiceFactoryBean is required only for Shared Task Service." --> I would put this note on the same line as the org.kie.spring.factorybeans.TaskServiceFactoryBean bean, or into a separate NOTE below. Now it seems like it is some part of the org.kie.spring.factorybeans.TaskServiceFactoryBean bean. * "Shared Task Service" --> I would write "shared task service", because you are not capitalizing task service anywhere in the section. * "suitable for manual population." --> is really "population" a correct term? If so, an article is missing. * "taken from KJAR identified by releaseID or GAV (Group, Artifact, Version)." --> KJAR is "tagged" here, but I don't see why (and also, I would say kJAR, but I don't know if KJAR is a literal or not). I would write "taken from kJAR +and+ identified by +`releaseID`+ or GAV (Group, Artifact, Version)." Also, what is releaseId? I would definitely tag that and maybe specify. * "Built from classpath that consists of kmodule.xml descriptor." --> class path is correct (see IBM Style Guide). And missing articles? * "TaskServiceFactoryBean creates TaskService instances based on given properties. Creates a single instance only." I am not completely sure about this one, but I feel like these two sentences conflicts. In the first one, you are saying that it creates instances, but in the second that it creates a single instance. Maybe I just don't understand the topic, but I recommend rewriting this. * "List of TaskLifeCycleEventListener." --> what does that mean? List of TaskLifeCycleEventListener instances? I would try to be more specific. * "Setup entity manager factory and transaction manager" --> set up + missing articles? * "process is included from the classpath." --> class path * "jBPM services" --> again, is that correct? Maybe it is, let me know. * "spring application:" --> Spring application, missing article * "Configure security provides:" --> providers? * "Configure runtime manager factory:" --> missing article? * "with Spring container." --> missing article? I am not very sure about the articles because that is something I am not very good at, so please make sure that none are missing in the text and that they are correct. :)	1	1	https://issues.jboss.org/browse/BXMSDOC-822
2528	BXMSDOC-818	Vidya, could you please list all the changes that were made? I see only the ones in Table 2. Custom JVM Properties, but nothing new in the section Creating JMS Connection Factories.	1	1	https://issues.jboss.org/browse/BXMSDOC-818
2529	BXMSDOC-810	Hi [~tlivora], Just noticed that the section "3.4. Setup JMS Resources" in the IBM Websphere guide has the following note: If you will be installing the Intelligent Process Server later, add the KIE.SERVER.REQUEST and KIE.SERVER.RESPONSE with the corresponding JNDI names: jms/conn/KIE.SERVER.REQUEST and jms/conn/KIE.SERVER.RESPONSE respectively. Does this also need correction? jms/conn/KIE.RESPONSE instead of jms/conn/KIE.SERVER.RESPONSE AND jms/conn/KIE.REQUEST instead of jms/conn/KIE.SERVER.REQUEST? Regards, Vidya	1	1	https://issues.jboss.org/browse/BXMSDOC-810
2530	BXMSDOC-809	Hi [~vidya_iyengar], thanks, it looks good! I have a couple of comments: - I am not sure if a developer who wants to find some information about embedding Drools/jBPM engine into their own application would check the User Guide - this was the reason why I thought that Dev guide might be more appropriate. On the other hand, the information also can be useful in User Guide, would it make sense to add a reference to the Dev Guide (alternatively, maybe better, the other way around, i.e. add the content to Dev Guide and reference it from User guide)? - the part {{Only if you are using an already compiled kjar}} might be a bit misleading, since you need to add the dependencies even in case when you compile the kjar in your application and then you are using it. My understanding of the sentence is that I need to add the dependencies only when I use a kjar which had already been compiled (and the reader might think that it applies only for a previous kjar compilation by {{kie-maven-plugin}}). Maybe it is just my understanding, what do you think? - looking at the User Guide, I realized that the same applies for other asset types, namely {{Guided Scorecards}} and {{Guided Rule Templates}}, the Maven dependencies for these are as follows: {code} <dependency> <groupId>org.drools</groupId> <artifactId>drools-workbench-models-guided-scorecard</artifactId> </dependency> <dependency> <groupId>org.drools</groupId> <artifactId>drools-workbench-models-guided-template</artifactId> </dependency> {code}	1	1	https://issues.jboss.org/browse/BXMSDOC-809
2531	BXMSDOC-801	Hi Vidya, I can't seem to open this link for some reason. Is there somewhere else I can review the changes? Thanks, Gemma.	1	1	https://issues.jboss.org/browse/BXMSDOC-801
2532	BXMSDOC-748	Hi, is there only one merge request?	1	1	https://issues.jboss.org/browse/BXMSDOC-748
2533	BXMSDOC-728	After a discussion with Maciej, I have changed the matrix. I'm not sure about the stars, could you please check whether everything is ok? Edit: re-check the link above to see the changes.	1	1	https://issues.jboss.org/browse/BXMSDOC-728

#	Requirement	Questions	Answered	Successfully Answered	URL
2534	BXMSDOC-728	# Why do we distringuish between {{\-}} and {{_}}? It is really hard to see the difference since these characters are very similar. I would suggest to unify them under one character {{\-}} with the description: * _indicates that the user role may not do the specified operation or it is not an operation that matches the user's role ("not applicable")._ # Why is *_task stakeholder_* only allowed to execute *_start_* operation? # How can actual owner activate a task? Task can be activated when it is in Created state (see [task lifecycle|http://docs.jboss.org/jbpm/release/6.4.0.Final/jbpm-docs/html/ch07.html#jBPMTaskLifecycle]). This means it cannot have any actual owners. It can have potential owners but that does not necessarily mean they can execute this operation. # I would suggest to remove all those {{*}} because all the operations depend on the state in which the task is. This just makes the table more complicated and does not add any value. Other parts of the table seems to be correct. You can compare it with *_section 7.1.5_* of [WS-HumanTask Specification|http://docs.oasis-open.org/bpel4people/ws-humantask-1.1-spec-cs-01.html]. The only thing that is different is *_MAY_* value from the specification which depends on our implementation. Note that [community documentation|http://docs.jboss.org/jbpm/release/6.4.0.Final/jbpm-docs/html/ch07.html#d0e4794] which was probably the basis for the product one is all wrong when it comes to task permissions. It should also be fixed.	1	1	https://issues.jboss.org/browse/BXMSDOC-728
2535	BXMSDOC-723	The system property org.jbpm.designer.perspective=ruleflow has been added to the example box, but in opposite to other system properties mentioned there, this one lacks any description above the example box. Could you please describe it briefly? Thanks.	1	1	https://issues.jboss.org/browse/BXMSDOC-723
2536	BXMSDOC-723	Hi Radovan, I have added the description for this property based on one that was previously mentioned in Bugzilla 1130466 from a couple of years ago. Could you please confirm this is still the correct description? Jenkinscat link: http://jenkinscat.gsslab.pnq.redhat.com:8080/job/doc-BRMS-Oracle_Weblogic_Installation_and_Configuration_Guide-6.3%20(html-single)/lastSuccessfulBuild/artifact/tmp/en-US/html-single/index.html#Setting_Environment_Variables	1	1	https://issues.jboss.org/browse/BXMSDOC-723
2537	BXMSDOC-714	[~mczernek], looks good! A couple of comments: - Example 3.2. Red Hat JBoss BPM Suite Project Dependencies - {{jbpm-human-task-jpa}} is not mentioned there, although {{jbpm-human-task-core}} and some other *-jpa artifacts are. If a user wants to use jBPM with human task with persistence, this artifact should be there. - Example 3.5 JUnit dependencies - I would mention that the user does not need all of them. It depends on what they want to test (e.g. do they want to test persistence? If not, hibernate, btm and h2 would be redundant.). From my point of view, the current description creates impression that you need all of them for testing with Drools/jBPM.	1	1	https://issues.jboss.org/browse/BXMSDOC-714
2538	BXMSDOC-714	[~ksuta], would you mind taking a look at {{Example 3.3. Intelligent Process Server Dependencies}}? Thanks!	1	1	https://issues.jboss.org/browse/BXMSDOC-714
2539	BXMSDOC-714	sry for barging in, but here are my few notes: Example 3.2 - doesn't mention jbpm Service API at all - it still imports the old API, i.e. jbpm-runtime-manager. I think engineering recommends to use Service API, instead of old RuntimeEnvironment/RuntimeManager API Why there is this note for slf4j dependency? "<!-- slf4j is not suitable for production purposes --> " Who are we to claim this, and why are we claiming this? Isn't it up to customer which logging provider he uses? Also note that if you do not import some implementation (beside importing API) you will see this in the logs: SLF4J: Failed to load class "org.slf4j.impl.StaticLoggerBinder". SLF4J: Defaulting to no-operation (NOP) logger implementation SLF4J: See http://www.slf4j.org/codes.html#StaticLoggerBinder for further details. That's why I also import slf4j-simple to my projects. Example 3.3 - Intelligent Process Server Dependencies - these are dependencies for client app interacting with Intelligent Process Server - I think the name of this section should reflects this. Example 3.5 - why is jbpm-shared-services included? Isn't it more convenient to just import whole jbpm-test? It usually has everything I need to perform my tests. I still believe, that we should have these poms suggested by engineering.. because we may have quite some mistakes in these pom examples, even with SME reviews. Wdyt [~alessandrolt] ? Just to mention one example of what has been overlooked: There is no mention of jbpm-workitems anywhere. I have checked all the dependencies from Example 3.2 and none of those artifacts brings transitively jbpm-workitems - which means your RestWorkItemHandler et all would fail with CNFE. I think jbpm-workitems is a dependency present in most projects for example.	1	1	https://issues.jboss.org/browse/BXMSDOC-714
2540	BXMSDOC-714	I have made the following changes: * Example 3.1 was containing one example dependency. This was removed since it did not make sense in view of the example dependencies we provide below it. * _jbpm-human-task-jpa_ was added. * Example 3.3 - Title was changed. * Example 3.5 - comments were added so that it's not clear users may not need all the dependencies * Mention of _jbpm-test_ was added. I need clarification on: * _slf4j_ - should we remove the production suitability comment? * Should we change _jbpm-runtime-manager_ to the new api? (that is _jbpm-services-api_ if I understand it correctly) * Should we add _jbpm-workitems_ into example 3.2? Any other dependencies missing? Once this Jira is done, I'll make a PR for upstream to update upstream pom.xml examples, if necessary and desirable.	1	1	https://issues.jboss.org/browse/BXMSDOC-714
2541	BXMSDOC-714	[~tlivora], would you mind taking a look at jBPM part? Thanks!	1	1	https://issues.jboss.org/browse/BXMSDOC-714
2542	BXMSDOC-714	Thanks [~tlivora]! [~mczernek], do we have something similar for BRMS? There could be *Application with embedded Drools engine* containing the following dependencies: - {{org.drools:drools-compiler}} - it will bring in other transitive dependencies, such as {{org.drools:drools-core}}, {{org.kie:kie-api}}, etc. - {{org.drools:drools-persistence-jpa}} - if you want to have persistence support (to persist KieSession using JPA). - {{org.kie:kie-ci}} - if your application needs to load kJARs from a Maven repository using KieScanner. - {{org.kie:kie-ci-osgi}} - if your application needs to load kJARs from a Maven repository using KieScanner *in an OSGi environment* (e.g. Fuse). - {{org.drools:drools-decisiontables}} - if you want to use decision tables (xls or csv) (does not cover Guided Decision Tables - gdst). - {{org.drools:drools-templates}} - if you want to use Drools templates. - {{org.drools:drools-scorecards}} - if you want to use Drools scorecards. - for using other asset types at runtime, import appropriate {{org.drools:drools-workbench-models-*}} artifacts (e.g. {{org.drools:drools-workbench-models-guided-dtable}}). - for building a kJAR as a part of your application, consider using {{kie-maven-plugin}}.	1	1	https://issues.jboss.org/browse/BXMSDOC-714
2543	BXMSDOC-714	Thank you all! [~marek.winkler], I have made desired changes. Can you review the doc again?	1	1	https://issues.jboss.org/browse/BXMSDOC-714
2544	BXMSDOC-1613	Ah, I see. Okay, that makes sense. Thanks [~Rikkola]! One last question: For the "Has priority over row" column, to specify two rows, can I say "1 and 2" or "1, 2"? In my example, I want row 3 to take priority over rows 1 and 2, and I wasn't sure if there's a strict formatting protocol for that. I'll make some adjustments and send you one last ping when it's up to just make sure finally that all's well in your eyes. Thank you!	1	1	https://issues.jboss.org/browse/BXMSDOC-1613

#	Requirement	Questions	Answered	Successfully Answered	URL
2545	BYTEMAN-281	Hi Radim, I think it's a very good idea to be able to add rules at points where objects are allocated.I'll leave this issue open while I push release 2.2.2 and the try to address it ASAP by confirming a syntax for the location (AT) specifier and special variables and then either implementing it myself or try to get someone else to implement it (you are welcome to post your own patch :-). I am surprised that you were not able to inject into java.lang.Object METHOD <init> AT ENTRY. There is no reason why injection into java.lang.Object should not work. Are you sure you remembered to set -Dorg.jboss.byteman.transform.all when you tested it? (if you install the agent using bmsubmit.sh then this is set automatically). I will set up my own test to see if I can get this to work.	1	0	https://issues.jboss.org/browse/BYTEMAN-281
2546	BYTEMAN-281	Ah, ok. thanks for spotting this. I think I know what this is now! I believe it's to do with ensuring that injection into a constructor does not place any code in front of a call of to a parallel constructor or super constructor (that's disallowed by the verifier). Of course the one place you don't ever see another constructor call is in Object.<init>. AT RETURN works because it si the only location-specific checker/transformer which does not bother to look for a prior constructor call (it just injects when it see a return bytecode). Could you please raise this as a separate JIRA (can only inject into Object.<init> AT RETURN). I'll pick it up from there.	0	0	https://issues.jboss.org/browse/BYTEMAN-281
2547	BYTEMAN-273	"Is this an OSGi related limit? Am I doing something wrong?" Yes, I am fairly sure that you are doing something wrong. Specifically, the following step almost certainly accounts for the behaviour you are seeing: "I have deployed Byteman jar as a wrapped OSGi bundle, so that a proper class loader for that .jar was created and made available to all the other bundles." Unfortunately, the JVM does not know or care about your Byteman bundle. When the JVM loads any agent, Byteman included, the jar which contains its implementation (i.e. byteman-2.1.0.jar or whatever version you ave used) is added to either i) the system classpath or ii) the bootstrap classpath and the JVM then loads the agent code via the associated loader. So, the agent initialization is all done using code which is not located in your bundle. Just as an aside, option i is the default. If you install the agent from the command line then you can select option ii by adding boot:/path/to/byteman-X.Y.Z.jar as an option to the -javaagent argument to the java command. If you use bminstall to install the agent then flag -b will select option ii. Whatever way you do the agent load your Byteman rules get loaded into a rule database which is indexed from a class (static) variable of class Transformer. That's the version of class Transformer as loaded into the system or bootstrap loader. This class will be the one which injects rules into your application code. Whether they are provided at boot time or uploaded via the Transformer's agent listener it is the version of the Tranformer class loaded by the JVM that gets asked if it would like to transform your application class's bytecode. If it finds rules in its rule database whose name matches your app class's name it transforms them so that they call Rule.execute. The class in your bundle never gets asked to install rules or transform bytecode. However, when the transformed bytecode is linked that is done using the OSGi linkage of the transformed class. Now, you added the byteman agent jar as a module and made this module visible to your app. So, the transformed bytecode will be linked against the copy of of Byteman classes in your bundle. The call to Rule.execute will use the version of class Rule in this bundle and it, in turn, will search the rule database associated with the copy of class Transformer in this bundle. Which will be *empty*. Rule.execute gets called, finds no rule definition and so completes without doing anything. Why is this not flagged as an error! I hear you ask? Well, that's because triggering a rule which is no longer present is a perfectly legitimate outcome. A rule may be in the process of being deleted by the agent while another thread is triggering it concurrently. The agent has to remove the rule from the database to disable it atomically then schedule retransformation of the bytecode to remove the call to Rule.execute. It cannot do it the other way around and maintain consistency. Ok, so how *do* you get your application classes to be able to see the correct version of class Rule -- the one the JVM has placed in either the system or bootstrap classpath. What you need to do is tell OSGi that any class whose name starts with prefix "org.jboss.byteman." should be looked up by delegating to the JVM's system loader. I am not an OSGi expert but I believe you need to add this prefix to the package list specified by OSGi property org.osgi.framework.bootdelegation (n.b. you may or may not want to include a final trailing "." in the prefix). See http://wiki.osgi.org/wiki/Boot_Delegation for more details of how to override delegation.	1	1	https://issues.jboss.org/browse/BYTEMAN-273
2548	BYTEMAN-268	Seems this issue is caused by targeting hadoop bytecode version to 1.7, following change involved in pom.xml of hadoop project: {code} <!-- ProtocolBuffer version, used to verify the protoc version and --> <!-- define the protobuf JAR version --> <protobuf.version>2.5.0</protobuf.version> <protoc.path>${env.HADOOP_PROTOC_PATH}</protoc.path> <zookeeper.version>3.4.5-gphd-2.1.0.0-SNAPSHOT</zookeeper.version> +<javac.source>1.6</javac.source> This line is added. +<javac.target>1.7</javac.target> This line is added <plugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-compiler-plugin</artifactId> <version>2.5.1</version> <configuration> -<source>1.6</source> This line is deleted -<target>1.6</target> This line is deleted +<source>${javac.source}</source> This line is added +<target>${javac.target}</target> This line is added </configuration> </plugin> {code} So previously hdfs can be started with jdk1.6.0_21 and byteman works well with this version of HDFS. After the change involved, hdfs can't be started with jdk1.6 with the following error info. and byteman also doesn't work any more: {code} Exception in thread "main" java.lang.UnsupportedClassVersionError: org/apache/hadoop/hdfs/server/namenode/NameNode : Unsupported major.minor version 51.0 at java.lang.ClassLoader.defineClass1(Native Method) at java.lang.ClassLoader.defineClassCond(ClassLoader.java:632) at java.lang.ClassLoader.defineClass(ClassLoader.java:616) at java.security.SecureClassLoader.defineClass(SecureClassLoader.java:141) at java.net.URLClassLoader.defineClass(URLClassLoader.java:283) at java.net.URLClassLoader.access$000(URLClassLoader.java:58) at java.net.URLClassLoader$1.run(URLClassLoader.java:197) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:307) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:248) Could not find the main class: org.apache.hadoop.hdfs.server.namenode.NameNode. Program will exit. {code}	1	0	https://issues.jboss.org/browse/BYTEMAN-268
2549	BYTEMAN-268	"So previously hdfs can be started with jdk1.6.0_21 and byteman works well with this version of HDFS. After the change involved, hdfs can't be started with jdk1.6 with the following error info. and byteman also doesn't work any more" I don't understand what you are saying here. Your pom is targeting 1.7. So, why are you trying to start hdfs with jdk1.6? The error above says that your compiled class file for class NameNode has a 1.7 major/minor version. That is what would be expected if your pom sets target 1.7. It is also expected that jdk1.6 will not run code compiled with target 1.7. Can you explain why you think Byteman is responsible for this problem rather than the specification in your pom?	1	1	https://issues.jboss.org/browse/BYTEMAN-268

#	Requirement	Questions	Answered	Successfully Answered	URL
2550	BYTEMAN-268	Hi Andrew, Thanks for the response and sorry for the misunderstanding. I will make it more clear about this issue: Let's say there are two HDFS version A and B. The compiler target of A is 1.6, while the compiler target of B is 1.7. For HDFS version A, I have been using Byteman to do fault inject testing well on it. However for HDFS version B, since the target is 1.7, I have to start the namenode with jdk1.7, otherwise it will hit the error listed above, which is also explained by you as the expected error. My point is that currently byteman doesn't work with HDFS version B, which will hit the VerifyError during rule bmsubmit. The reason why I listed the namenode start error is that I searched one issue https://issues.jboss.org/browse/BYTEMAN-211 and thought that this issue occurred again. Please ignore namenode start error, which do confused you. Sorry about that. Now I suspect that it's the compiler target change from 1.6 to 1.7, which blocks byteman to be used in HDFS. Could you please help take a look at this scenario? Any additional debug information you need, please just let me know. Thanks again. Thanks, Zhenglin	1	0	https://issues.jboss.org/browse/BYTEMAN-268
2551	BYTEMAN-268	Hi Zhenglin Thanks for the explanation. I now understand the situation you are describing. Can you provide me with more information about the byteman tests which are failing? Maybe you can provide me wit the rules you are loading. Is there a repo I can check out and use to build HDFS and instructions on how to run the Byteman tests? regards, Andrew Dinn	1	1	https://issues.jboss.org/browse/BYTEMAN-268
2552	BYTEMAN-268	Hi Zhenglin, "1.Go to http://pan.baidu.com/disk/home and download HDFS file hadoop-2.2.0.tar.gz Login Username/Password: sztao1987@163.com/welcome1 Hopefully it works in US." Well, I am not actually in the US but in the UK. Anyway, I am afraid I cannot access your code via this link. Is there a Hadoop project repository where I can access the Hadoop 2.2.0 code you are using? Have you made changes to Hadoop 2.2.0 which are not available in that repository?	1	1	https://issues.jboss.org/browse/BYTEMAN-268
2553	BYTEMAN-268	Hi Andrew, Do you have a chance to download the binary attached in this jira? Thanks, Zhenglin	1	0	https://issues.jboss.org/browse/BYTEMAN-268
2554	BYTEMAN-261	I am not exactly clear about the problem scenario. Could you give a concrete example of something you are tryng to do that is broken? Without fully understanding your requirements I don't know whether this will help but perhaps you need to attach your Byteman rule to the class initializer for your Main class (or some other class, perhaps). To do that use a rule like the following: RULE trigger side effect at Main class init CLASS org.my.Main METHOD <clinit> IF TRUE DO trace("id", filename") ENDRULE	0	0	https://issues.jboss.org/browse/BYTEMAN-261
2555	BYTEMAN-260	Could you explain what you want in more detail? In particular, how does what you need involve some form of nesting? Can you give a sketch of how a nested set of rules for the use case you are proposing would look?	1	1	https://issues.jboss.org/browse/BYTEMAN-260
2556	BYTEMAN-249	Hi Andrew, sorry to answer so late. And excuse me for missing the getObjectSize of Helper class. My custom Helper (MyHelper) extends Helper, so I can use the Helper.getObjectSize method. So this request can be closed.Thanks to point it. This method gives the shallow size of the object, so could you add a method getDeepSize or getRetainedHeap in Helper class that add the size of the objects tree referenced by an object? At github site there are 2 examples of implementation : - https://github.com/afeinberg/sizeof this class uses Instrumentation to get the size of each objet referenced, but you can replace this call by Helper.getObjectSize - https://github.com/dweiss/java-sizeof more sophisticated and more recent and no need of Implementation. But javac 7 complains because of using sun.** packages The retained heap is very usefull, when running stress tests to know the retained heap at Http Session Level, at JPA EntityManager Session (for example with Hibernate). What do you think about it? If you agree, I can add this new request.	1	1	https://issues.jboss.org/browse/BYTEMAN-249
2557	BYTEMAN-248	Hi Tom, Thanks for reporting this issue. Could you also please post the rule which is being injected -- the contents of the script file if you are using a BMScript annotation or the annotation field values if you are using a BMRule annotation. Thanks. regards, Andrew Dinn	1	1	https://issues.jboss.org/browse/BYTEMAN-248
2558	BYTEMAN-248	Hmm, that doesn't seem to tally with the error message: "org.jboss.byteman.agent.Transformer : possible trigger for rule Throw SQLException in WCMHelper on 2nd write in class com.maginatics.raptor.metadata.dbfs.DatabaseFs" Byteman says it is injecting code into a class called DatabaseFs. Can you explain to me what this class is and how it relates to the classes mentioned in your annotations, Class1 and Class2?	1	1	https://issues.jboss.org/browse/BYTEMAN-248
2559	BYTEMAN-245	I tested this code and runs as expected. Sorry i don't know how to correctly indent ;-\. Edit it to cut/paste. Could it be committed on master-branch at Github repo ? for (String scriptPath : scriptPaths) { try { FileInputStream fis = new FileInputStream(scriptPath); byte[] bytes = new byte[fis.available()]; fis.read(bytes); String ruleScript = new String(bytes); scripts.add(ruleScript); } catch (IOException ioe) { // take a new chance with system ClassLoader => ClassLoader.getSystemResourceAsStream(scriptPath); try { InputStream in = ClassLoader.getSystemResourceAsStream(scriptPath); byte[] bytes = new byte[in.available()]; in.read(bytes); String ruleScript = new String(bytes); scripts.add(ruleScript); } catch (IOException e) { System.err.println("org.jboss.byteman.agent.Main: unable to read rule script file : " + scriptPath); // throws original ioe throw ioe; } } }	1	1	https://issues.jboss.org/browse/BYTEMAN-245
2560	BYTEMAN-245	I am not really convinced this is a good idea (not that I am thereby convinced it is a bad idea but the onus is on you to sway me to put this into the release version). First, your use of getSystemResourceAsStream means that the scriptPath can be resolved against any jar in the system classpath, not just the Byteman jar. That's a much wider remit than the functionality you have stated you want/need and probably has a lot more room for security issues to arise (although, to be fair, once the java command line is prejudiced I guess all bets are off anyway :-) Secondly, I am not sure about the full implications of trying to load a resource from the system classpath. I try to do as little work as possible in Main and you may notice that I never let any other Byteman classes be resolved in that class. That's because the "boot:xxx" option in the -javaagent arguments list may be used to relocate the agent jar into the bootstrap classpath. Main is doubly special because nothing in the Byteman code references it. That's why this relocation is safe. So, before making any changes here I'd want to be sure that any accesses to the agent jar from Main don't prejudice resolution against the jar in the bootstrap loader. They probably won't but this needs checking. Finally, let's cut to the chase and ask the $64,000,00 question as I am not sure I really appreciate this idea of bundling everything into the agent jar itself. What is your motive for doing this? Indeed, for another $64,000,000: what benefits do you think this provides? Is there a deployment management issue here that it is really important to allow for? Or something else?	1	1	https://issues.jboss.org/browse/BYTEMAN-245
2561	BYTEMAN-239	Is the ThreadID not mandatory also ?	1	1	https://issues.jboss.org/browse/BYTEMAN-239
2562	BYTEMAN-237	Great news! Btw., is there any way to influence how the rule bindings are nested (for a multiple matching rules)?	1	1	https://issues.jboss.org/browse/BYTEMAN-237
2563	BYTEMAN-236	I have tried to bind the variable using class name of the interface (java.lang.Runnable) but the problem from BYTEMAN-235 was still there. Do you expect to fix it in order to let it work at least with the specific binding?	1	0	https://issues.jboss.org/browse/BYTEMAN-236

#	Requirement	Questions	Answered	Successfully Answered	URL
2564	BYTEMAN-233	I'm not certain exactly what to do here. I'm working on the assumption that the agent startup code needs to i) employ a clone of class ModulesPolicy renamed as, say, BytemanPolicy ii) during agent installation execute Policy.setPolicy(new BytemanPolicy(Policy.getPolicy())); Is that all that is needed to ensure that Byteman code accesses are validated using BytemanPolicy? If so then I propose to add a -javaagent option setPolicy:true which will cause the agent to call this code. A few observations which might qualify any answer to that last question: a) The JVM installs the agent Main class in the system classloader. b) The agent entry routines Main.premain and Main.agentmain() may, depending upon configuration, install the byteman agent jar (which contains Main and all other agent classes) into the bootstrap class path, thereby allowing other agent classes to be loaded via the bootstrap loader. c) These two agent main routines use reflection to create a Transformer instance which does all the rest of the startup using reflection, ensuring that all subsequent load resolution can go via the bootstrap loader if so configured. This is fine because no agent classes refer to Main. So, I am not certain what the implications of this are for i) when to execute the getPolicy call ii) when to create the BytemanPolicy instance iii) when to execute the setPolicy call I suspect the BytemanPolicy instance needs to be created by the Transformer object so it can be loaded via the bootstrap loader (where bootstrap loading is configured). If not then we risk seeing two versions of the class, one in the system loader and one in the bootstrap loader. That implies the Transformer will also have to make the call to setPolicy. I could call getPolicy in the main routine and pass it to the Transformer in the reflective call or create it in the Transformer code. Any idea whether this matters in terms of what permissions are granted? I would guess that creating it in a bootstrap-loaded class is better but I don't know for sure.	1	1	https://issues.jboss.org/browse/BYTEMAN-233
2565	BYTEMAN-217	Hi Andrew, thanks for the detailed explanation. I always thought, the compiler already handled these details when compiling the code to be instrumented. {quote} it can be hard to decide which monitorexit a rule is supposed to be injected before or after. This problem arises because when you consider both normal and exception control flow there is more than one way for control to flow out of a synchronized block. Each exit needs to go through a monitorexit. Exceptions which cross several monitor regions give rise to multiple monitorexits, one per region crossing. {quote} Would it not be possible, to inject code for each of these monitorexists? I would expect my injected code to be called whenever the synchronized block is exited, no matter whether an exception was thrown or the block finished normally.	1	0	https://issues.jboss.org/browse/BYTEMAN-217
2566	BYTEMAN-190	Andrew, This issue seems rather general. Can it be focused on just the Maven plugin aspect? Alternatively, I could create a sub-issue for the Maven plugin bit.	1	1	https://issues.jboss.org/browse/BYTEMAN-190
2567	BYTEMAN-189	Hi James, Thanks for the suggestion and the patch. However, I don't think this is going to be quite as easy as it looks. My first concern is that the model you have implemented is counting catch regions as they appear in bytecode. That's not necessarily going to work because bytecode catch regions are not necessarily going to tally 1:1 with catch blocks as they appear in source code. In particular, there is a very complicated interaction between catch regions and synchronization regions which needs to be taken into account. The situation which complicates things occurs when a (or possibly several) thrown exception(s) propagates through a synchronized region. In this case the bytecode compiler has to insert a (or possibly several) 'synthetic' catch region(s) at the end of the synchronized region (i.e. just after the primary monitorexit), catch the thrown exception(s), execute a corresponding monitorexit in the synthetic catch handler and then rethrow the exception. This is complicated enough to spot and decipher when there are no try catch regions either inside or wrapped around the synchronization region. In the presence of such try catch regions it becomes very difficult to decide whether a catch region is synthetic, original to the source code or some combination of the two (the bytecode compiler may merge other handler code in). As if that is not complicated enough there is also the problem of finally code to consider. Finally blocks are normally implemented using catch-all regions (i.e. a catch region with a null type). However, these regions may also occur because the source code includes a catch block for a Throwable (also, if these are used in combination -- bizarre but possible -- the bytecode compiler might legitimately merge them). So, it is not clear how to classify these (do they count as a catch block?). Finally (no pun intended), there is another issue with finally. Prior to JDK6 the code in the finally block would usually be executed via a JSR in order to allow it to be entered from parallel catch regions as well as from a catch-all region. Since JDK6 JSR has been banned so the finally code is generated multiple times (n.b. this actually means that injection into finally blocks is currently a little iffy as there may be multiple candidate trigger points even if the source only has one apparent trigger point). This does not necessarily add any further grief to the above but it may do if the finally block code itself introduces the possibility of extra exception control flow. I would need to think about this very hard before I can be sure that all these problems can be resolved. Also, it's not just making sense of the semantics of the count argument to AT CATCH type count. There is also the meaning of the type argument. I am not sure whether the trigger injection code can actually identify the necessary type relationships in the exception superclass hierarchy to handle every possible choice of value for the type argument to the AT CATCH location. At injection time it is not possible to load and check the superclass hierarchy for the resolved target class (recursive entry into the transformer code is prohibited so recursive loads would mean no transformation of the relevant superclass). I have worked round this in a few places but there are limits to how I can use this capability. Analysing the exception flow to handle the problems above may require knowing about which exceptions shadow other exceptions so this also needs investigation. If you want to understand more about this second problem I already have a long-standing issue wrt it -- see BYTEMAN-113. But be warned, the relevant code (in and under class CFG) is the scariest piece of the Byteman puzzle. So, expect to take a long time to understand it -- I usually spend a week or two each time I have to fix a bug in it :-) Anyway, nice try and most definitely nothing stupid about your idea or your patch. However, much as I would like to add this capability to Byteman I think it is going to need a lot more thinking and probably some very careful programming before it can be made safe for all cases.	1	1	https://issues.jboss.org/browse/BYTEMAN-189
2568	BYTEMAN-141	Script files now use .btm as the extension. All scripts in the release use this convention and the docs also use it for example script names.	0	0	https://issues.jboss.org/browse/BYTEMAN-141
2569	BYTEMAN-114	Hi Andrew, good reasoning. Still, there is a valid use case: instrumenting many or even unknown classes. So how about clearly documenting that using wildcards is slow and going for the combined way of first performing a hash lookup or two and then checking expressions? If there are no wildcard patterns for CLASS and/or METHOD, this should be as fast as before, but if I need this feature, it would be available, although slower? Regarding implementing this as a contribution, do you have any pointers, where to put this? Cheers, Wolfgang	1	1	https://issues.jboss.org/browse/BYTEMAN-114
2570	BYTEMAN-114	Hello, I would also like this functionality in Byteman, especially to quickly add traces in some unknown code. Was any work on this done? If not I might have a look...	1	0	https://issues.jboss.org/browse/BYTEMAN-114
2571	BYTEMAN-103	The syntax for referring to the INVOKEd method's arguments needs to be unambiguous wrt local and parameter variables of the triggering method which propels us into the realm of punctuation characters. I propose $@ which will only be available in an AT INVOKE method and will bind the arguments passed to the invoked method along similar lines to $* i.e. if the invoked method has N >= 0 parameters $* will be bound to an Object array of size N + 1. Entries 1 .. N will be the N argument values on the stack. Entry 0 will be the recipient for the method call if it is an instance method/null if the invoked method is static.	0	0	https://issues.jboss.org/browse/BYTEMAN-103

#	Requirement	Questions	Answered	Successfully Answered	URL
2572	BYTEMAN-70	uh-oh: } else if (binding.isReturn()) { // at some point we will allow reference to the current return value so we need // to be sure that the method has a non-void return type and that the location // specifier is AT RETURN but for now we do nothing } Looks like $! isn't implemented? I tried to use it and got some parsing errors. Perhaps this JIRA issue is more than just for documentation - I would like $! to be an implemented feature.	1	1	https://issues.jboss.org/browse/BYTEMAN-70
2573	BYTEMAN-29	Andrew, isn't such a plugin partially done? - I've seen somewhere on the internet something what was an eclipse plugin for byteman. Do you have any project(some diagrams or sth) for such a plugin - what describes more precisely how it should exactly work? Kind regards, Lukasz	0	0	https://issues.jboss.org/browse/BYTEMAN-29
2574	CAPEDWARF-168	How does "com.pagecrumb.client.local.Footer" look like? (and its whole super class hierarchy)	1	1	https://issues.jboss.org/browse/CAPEDWARF-168
2575	CAPEDWARF-168	You clearly have what the error states -- 2 scope annotations. I suggest in the future you first open forum post with the issue, and only *after* we determined this is a CD bug, should you open a CD JIRA, OK?	1	1	https://issues.jboss.org/browse/CAPEDWARF-168
2576	CAPEDWARF-166	Kerby, can you attach a .war with the issue?	1	1	https://issues.jboss.org/browse/CAPEDWARF-166
2577	CAPEDWARF-166	I don't get this stack-trace: https://gist.github.com/alesj/5527634 How did you start app. server? Either: ./capedwarf.sh or ./standalone.sh -c standalone-capedwarf.xml	1	1	https://issues.jboss.org/browse/CAPEDWARF-166
2578	CAPEDWARF-75	o) If you login and logout. Then the login screen should display the email you used in the last login session but it uses the default. o) authDomain and userId should be set Using the dev login, calling User currentUser = UserServiceFactory.getUserService().getCurrentUser(); returns {code} currentUser = {com.google.appengine.api.users.User@9969}"foo@bar.com" email = {java.lang.String@9974}"foo@bar.com" authDomain = "" userId = null federatedIdentity = null {code} where in the case of Google's Dev server it returns {code} currentUser = {com.google.appengine.api.users.User@4849}"foo@bar.com" email = {java.lang.String@5526}"foo@bar.com" authDomain = {java.lang.String@5527}"gmail.com" userId = {java.lang.String@5528}"12431731640520613711" federatedIdentity = null {code} Note the authDomain and userId are set. The userId is preserved over logins.	1	1	https://issues.jboss.org/browse/CAPEDWARF-75
2579	CAPEDWARF-31	Is this done?	0	0	https://issues.jboss.org/browse/CAPEDWARF-31
2580	CAPEDWARF-23	Is there a list of supported attributes? (not just this single one ...)	1	0	https://issues.jboss.org/browse/CAPEDWARF-23
2581	CDI-703	What part exactly would you like to clarify? As a part of CDI-441 (2.0) we enhanced the wording making it clear that the relation is *transitive*. And that, in my opinion, should have quite clear implications in this regard.	1	1	https://issues.jboss.org/browse/CDI-703
2582	CDI-702	The description of this issue is unclear. Are you seeing all events being fired or only certain events? Where in the archive is the extension registered, a JAR within the EAR's lib folder? Within the WAR file? Within a library within the WAR file?	1	1	https://issues.jboss.org/browse/CDI-702
2583	CDI-702	My view is that the observable should follow the classloading visibility rules. Why extensions are special? This violent the classloading rules.	1	0	https://issues.jboss.org/browse/CDI-702
2584	CDI-702	[~meetoblivion] The issue is that the CDI extension does something different from other classes in the ear. e.g. {noformat} myApp.ear lib\myLib.jar (LibExtensionA.class, LibOne.class) myWarA.war (WarAExtension.class, myWarAServlet.class) myWarB.war (WarBExtension.class, myWarBServlet.class) {noformat} The myWarAServlet will not be able to see classes in myWarB.war. Why are the CDI extensions special? How come all extension classes can see everything in the ear. Basically, my argument is that CDI extensions should follow the same classloading criteria instead of defining itself visibility rules with even being documented.	0	0	https://issues.jboss.org/browse/CDI-702
2585	CDI-701	Thanks [~mkouba] for pointing me to that issue. The interesting fact is that it didn't work on TomEE either (with OpenWebBeans). That what led me to think it was not supported by the spec. I was aware of the trimmed discovery mode, but I'm still relying on Java EE 7 application servers and for people like me, there's unfortunately still a journey before we can use CDI 2. But anyway, what I'm trying to do is a library that would be put into the classpath of a war, so I wouldn't be in control of the discovery mode of the war itself. Since this is required to work by the spec, could a test be added to the TCK? I understand that would put us in a complicated situation, where the RI wouldn't pass it, but the fact is that the absence of that test makes it difficult to know whether it's expected to work or no (to the point that the other certified implementation doesn't work either). Would a disabled test do the trick? It wouldn't hurt certifications, but users looking at it (like me in this case) would be alerted that there's a bug.	1	1	https://issues.jboss.org/browse/CDI-701
2586	CDI-701	Thanks everyone for so fast respones! [~manovotn] my original conclusion was exactly that (feature not supported by spec). Maybe the spec could incorporate a note regarding this, saying "Implementations are NOT MANDATED to treat stereotypes added by this method as bean defining annotations. Implementations are ALLOWED to do it in a non portable manner". That would clarify the situation and avoid confusions. A way to add annotations to the set of "bean defining annotations" was what I was really looking for. I just wanted to clarify that the addStereotype() was not the way to go before creating a feature request. Has that been already discussed in the past?	1	1	https://issues.jboss.org/browse/CDI-701
2587	CDI-701	bq. That would clarify the situation and avoid confusions. Perhaps, but that would have to be added in 2.1 I think. You cannot expect a change to 1.2 spec at this point. bq. Has that been already discussed in the past? I don't recall a CDI issue regarding that. If you wish to see a non-portable Weld option working in CDI 1.2 (Weld 2.x), I think it would be proper to file a Weld feature request issue.	0	0	https://issues.jboss.org/browse/CDI-701
2588	CDI-700	The title of this ticket mentions "AfterDeploymentValidation" - this should be allowable. The body of this ticket mentions "AfterBeanDiscovery" - this should not be allowable. Agree?	0	0	https://issues.jboss.org/browse/CDI-700
2589	CDI-697	Good question. We don't call it out, but I would assume it should. How do others feel? Martin? Mark? Antoine?	0	0	https://issues.jboss.org/browse/CDI-697
2590	CDI-695	IMO, I think it would bring more confusion than useful: * A Stereotype is usually created to be used on many elements. What would be the interest to apply same priority in different places? * {{@Priority}} is used and activate Interceptors and Alternatives and not having this info directly on the element can be source of confusion IMO. Could you give a use case where it could be interesting to have this?	0	0	https://issues.jboss.org/browse/CDI-695
2591	CDI-688	Right now, this is undefined, i.e. it's completely up to the implementations (assuming the requirements for normal scopes are met: "6.3. Normal scopes and pseudo-scopes"). What kind of pluggability do you have in mind?	1	0	https://issues.jboss.org/browse/CDI-688
2592	CDI-686	Hmm, isnt one of the goals of this factory to make custom beans (Bean<>) easily integrated with CDI interceptors? If so it is common to have java.lang.reflect based proxies there (see deltaspike ones for instance) and it is very nice and smooth to be able to reuse cdi there. Here a trivial but obvious example: take a repository like proxy based implementation, don't you want to integrate it with jcache which has its own interceptor without coding more than the factory itself? Today it requires to either be bound to a jcache implementation or reimplement the interceptor chain logic from the container which is a pain.	1	1	https://issues.jboss.org/browse/CDI-686
2593	CDI-686	[~antoinesabot-durand] how is it supposed to work with an API which doesn't use @Produces and doesn't have a concrete implementation but just a Proxy instance?	1	1	https://issues.jboss.org/browse/CDI-686

#	Requirement	Questions	Answered	Successfully Answered	URL
2594	CDI-686	is that so hard? I mean, I don't think we need a CDI rule but just a standard java one where you can "cast" the instance since it is out of CDI matching rules world, no? Or do you want to limit to the "injected" generic type?	1	1	https://issues.jboss.org/browse/CDI-686
2595	CDI-686	What's the _"injected" generic type_?	1	1	https://issues.jboss.org/browse/CDI-686
2596	CDI-682	Hi [~saasira] Could you be please more specific?	1	1	https://issues.jboss.org/browse/CDI-682
2597	CDI-682	"CDI was designed to create the bean graph once at boot time giving us performance and limiting runtime exception on DI to the maximum." -- i agree that was primary purpose of CDI, and that is achieved with excellence. now, having reached a stable state of the stated goal, can't we scale to the next step? protective defaults are good for most of the cases, but it is those few cases which require adaptive flexibility to build "extra-mile" features. EJB provides UserTransaction besides Container Managed Transactions -- to let the developers handle specialized situations that cannot used CMT (of course, EJBContainer is also monolithic). for a start, you can leave the responsibility of validating bean-injection-point-interceptor-decorator-qualifier graph to the application developer; after all, it is his/her code is triggering registration/unregistration of beans at runtime, so he/she should be having all the details while processing bean registration requests (on installation of a module, for example). let the application developer face the consequences of using this API wrongly with well defined exceptions. or if it is possible to validate the bean graph in CDI container itself (if no time or resource constraints), that would be a great thing to happen to Java EE -- a modular web application platform that was missing for several years. at least we should learn from the success of PHP CMS like Drupal which were modular for more than a decade now. Apache Deltaspike (a CDI implementation) provides some convenient API, though not to the extent of individual beans. ContextControl contextControl = BeanProvider.getContextualReference(ContextControl.class); // stoping the request context will dispose all the @RequestScoped beans contextControl.stopContext(RequestScoped.class); // and now restarting the request context, request scoped beans get reloaded by CDI container contextControl.startContext(RequestScoped.class); possibly some time in future they may also implement BeanManager.register(XYZBean.class), and BeanManager.unregister(XYZBean.class) on popular demand. why lose the opportunity (before the next CDI spec) to make this a standard feature (that may be made optional for CDI containers to implement or ignore)? as a side note : even the API for starting & stopping CDI container is non-standard and left to implementations.	1	1	https://issues.jboss.org/browse/CDI-682
2598	CDI-682	another argument in support of registering & unregistering beans, interceptors, decorators, producers etc at runtime could be "micro-services" - each module (with cdi beans & interceptors, injection-points, etc) being deployed & undeployed (or installed & uninstalled) at run time - each module providing a subset of the functionality that a monolithic web application provides in entirety. isn't microservices the hot topic that is being targeted for javaee 8?	0	0	https://issues.jboss.org/browse/CDI-682
2599	CDI-659	Would it make sense to leverage this ticket as the intro to specifying servlet container support?	1	0	https://issues.jboss.org/browse/CDI-659
2600	CDI-650	+1 for the idea. How user would access to an implementation of this {{NotificationOption}} class?	0	0	https://issues.jboss.org/browse/CDI-650
2601	CDI-650	On event receiving side, how an user can get hold of this NotificationOption, which is an interface? Will this be part of observation such as EventMetaData? After getting hold of it, what the receiver is supposed to do?	1	1	https://issues.jboss.org/browse/CDI-650
2602	CDI-632	I am not a fan of implicit bean archives, but wouldn't it be better if SE worked the same?	0	0	https://issues.jboss.org/browse/CDI-632
2603	CDI-626	well a bit more than taste cause it implies how an application can test if CDI is there: is getCdi() enough or should it tries to call one CDI method too? Also not sure what "container" means in your comment, I see CDI as the CDI container handle so anything before is not really CDI (= the CDIProvider is a bridge to CDI but doesn't require CDI). That said agree we just need to make it clear. One point can maybe be to check the api jars out there and see if changing the impl isn't easier since people will maybe not upgrade their API :s. Also it viciously brings the question of the contextuality of CDIProvider which is maybe not something we want to "spec" yet (ie is the provider bound to an impl and one context once the lookup is done or is CDI only contextual?).	1	1	https://issues.jboss.org/browse/CDI-626
2604	CDI-625	Firing such an explicit pre-destroy event for the application context is probably ok. But please lets not introduce this in a generic way ;) What speaks for not being a problem to move it generally to 'When the Application Context is about to be destroyed": Where can yo observe any @Destroyed(AppllicationContext.class) event if it would only get fired after the ApplicationContext gets shut down? For sure not in an @ApplicationContext scoped bean. Nor in any other CDI bean ;)	1	1	https://issues.jboss.org/browse/CDI-625
2605	CDI-625	This is interesting idea but I am not sure this "interceptor" concept fits well to this situation. What would happen in the case you have multiple observers of this kind? I guess you cannot call {{ctx.proceed()}} multiple times. I am also not sure whether this is doable with this new payload type. It would mean to fire the event twice to stay at least little bit compatible.	1	1	https://issues.jboss.org/browse/CDI-625
2606	CDI-625	How about extending {{@Destroyed}}? {code:java} public @interface Destroyed { Class<? extends Annotation> value(); boolean beforePhase() default false; } class Test { void observeAfter(@Observes @Destroyed(ApplicationScoped.class) Object object) { } void observeBefore(@Observes @Destroyed(value = ApplicationScoped.class, beforePhase = true) Object object) { } } {code}	1	1	https://issues.jboss.org/browse/CDI-625
2607	CDI-625	I checked OWB and we always fire it after destroying the beans right now. What does Weld do?	1	1	https://issues.jboss.org/browse/CDI-625
2608	CDI-625	[~mkouba] being said today the behavior is uniform and the predestroy issue you describe actual it is always possible to use predestroy merging them in a bean (or more commonly having an orchestrator bean intended for the destruction) is beforeshutdown defined in this shutdown lifecycle? if not shouldnt it be the solution? I know it has been prevented but supporting it in any bean would work I think	1	1	https://issues.jboss.org/browse/CDI-625
2609	CDI-625	bq. ...it is always possible to use predestroy merging them in a bean... Not really. If you have application scoped Bean1 and Bean2 you cannot simply create Bean3 and destoy them there... in other words, this "merging" is not always possible. bq. is beforeshutdown defined in this shutdown lifecycle? {{BeforeShutdown}} is a container lifecycle event and is fired after all contexts are destroyed.	0	0	https://issues.jboss.org/browse/CDI-625
2610	CDI-625	You can always remove @PreDestroy in an extension and call it in Bean3 but was more thinking this way: since the user hits this issue the workaround is to not use @PreDestroy when it is an issue or order is important (which is not that common compared to @PreDestroy cases). Right about BeforeShutdown, know early OWB versions needed to move it before the actual destruction in some cases (could check the details and current state if needed) cause BeforeShutdown extensions are often accessing app scoped beans which was not possible in the mentionned order, maybe no more an issue with this new @BeforeDestroyed. About the semantic if we introduce @BeforeDestroyed we should likely introduce @AfterDestroyed for the consistency and do the same for @Initialized which would let us with two events doing the same in term or API (but easily collapsable in the implementation so no perf issue, just an API issue). Is that possible to deprecate not prefixed events? Or maybe @AfterXX should be a @Stereotype which would auto document the new API?	0	0	https://issues.jboss.org/browse/CDI-625

#	Requirement	Questions	Answered	Successfully Answered	URL
2611	CDI-623	Are you talking about the spec doc or javadoc? It seems to me that the [javadoc|https://github.com/cdi-spec/cdi/blob/2.0-EDR2/api/src/main/java/javax/enterprise/inject/spi/AfterBeanDiscovery.java#L82] is pretty clear on this. The configuration is not reusable and that was talked through a thousand times. The idea is that every call of {{abd.add*()}} gives you a new configurator and once you exit the observer method call, it will create a respective bean or observer method. I cannot really see any 'dangling' configuration in this flow. Could you elaborate a bit please?	1	1	https://issues.jboss.org/browse/CDI-623
2612	CDI-621	I am also unsure of what the ambiguity is in this case. Do you mean you would rather rename the method to something like {{removeAllAnnotations}}? In that case you might say that all the methods are ambiguous...	1	1	https://issues.jboss.org/browse/CDI-621
2613	CDI-621	[~meetoblivion] as Martin and Matej already said, your ambiguity is... ambiguous ;). Could you detail please.	1	0	https://issues.jboss.org/browse/CDI-621
2614	CDI-621	Well, the point is the {{AnnotatedTypeConfigurator}} and friends ({{AnnotatedMethodConfigurator}}, etc.) can ONLY be used to modify/configure annotations. So it's implicit. Also the javadoc is clear.	1	1	https://issues.jboss.org/browse/CDI-621
2615	CDI-621	Today the only thing you can remove is annotations. Maybe at a later point you can change other things. When you can, what does {{removeAll}} mean?	1	0	https://issues.jboss.org/browse/CDI-621
2616	CDI-616	Are you referencing: {code} @Inject private transient Foo foo; {code} Or the {{@TransientReference}} from Mark?	0	0	https://issues.jboss.org/browse/CDI-616
2617	CDI-616	Are we talking about a real java 'transient' field? Of course adding transient to an injected CDI field is utter rubbish ^^. Or are we talking about InjectionPoint#isTransient() and the CDI @TransientReference annotation? That is the one I was basing my explanation on.	1	1	https://issues.jboss.org/browse/CDI-616
2618	CDI-616	> Do nothing but add clarification in the spec : Injection in transient field is not supported. We probably are talking about thesame thing, but the wording is not quite as discussed on the ML as far as I interpret it. Why should the injection itself not be supported? I think what we agreed on is that we just don't do anything for RE-injecting it in case of deserialisation. At least that was I understood during our discussion. Not sure we need to clarify anything as this is standard Java behaviour: the developer is responsible to get back any transient fields himself by default.	0	0	https://issues.jboss.org/browse/CDI-616
2619	CDI-610	bq. For the Bean<T> case this would surely work as well... Which other cases do you have in mind? I wouldn't add this to the {{BeanManager}}. I think it's bound to the lifecycle of a {{@Dependent}} bean. bq. It should be as much aligned with injecting the InjectionPoint in a producer method. Not only producer methods but any {{@Dependent}} bean is allowed to obtain the injection point. bq. I don't know from the top of my head... Yes, it's a definition error.	1	0	https://issues.jboss.org/browse/CDI-610
2620	CDI-608	What's the type of {{cdiEventEndpoints}}?	1	1	https://issues.jboss.org/browse/CDI-608
2621	CDI-608	Not sure, what type does it need to be to support this?	1	0	https://issues.jboss.org/browse/CDI-608
2622	CDI-601	If I'm already subscribed via a `@PreDestroy` method, shouldn't that continue to work?	1	0	https://issues.jboss.org/browse/CDI-601
2623	CDI-601	From your point of view, does observer method ordering also take into account dependencies of beans that the observers will use?	0	0	https://issues.jboss.org/browse/CDI-601
2624	CDI-601	But doesn't this imply that a new {{@Dependent Worker}} bean would then be _created_ in order to handle this {{ApplicationStopped}} event? This seems like "non-obvious" behaviour when I'm trying to destroy existing beans and shut the application down.	0	0	https://issues.jboss.org/browse/CDI-601
2625	CDI-590	I'm not sure I understand the use case. Do you want to declare a producer method whose set of bean types will not be derived from the return type? A more general description would be great.	1	0	https://issues.jboss.org/browse/CDI-590
2626	CDI-590	I don't really like this idea. It feels 'hacky' and it could also contradict the type-safe approach. If I got it right, the idea is to have this producer: {code} @Produces @MyProducerQualifier // this Qualifier has some annotation saying it is 'generic' public Object create() { // do some work and return Object } {code} and then for this one producer to have the following injections: {code} @Inject @MyProducerQualifier String string; @Inject @MyProducerQualifier Integer int; {code} Could you give a more specific use-case showing where this would make sense (or even better, be neccessary)?	1	1	https://issues.jboss.org/browse/CDI-590
2627	CDI-590	bq. it doesn't support a lot of types I thought a custom producer for any type [can be implemented|https://github.com/apache/deltaspike/blob/master/deltaspike/core/api/src/main/java/org/apache/deltaspike/core/api/config/ConfigProperty.java#L78]. bq. Just use PropertyEditor API... I don't know what do you mean. Probably {{java.beans.PropertyEditor}}? bq. Tamaya has its own coercing system which wouldnt fit CDI very well... For Tamaya I think it would be more meaningful to create a producer for {{org.apache.tamaya.Configuration}}. But I understand that it would be nice to leverage {{PropertyConverter}} and produce configuration properties directly. Look, I'm not against this feature! I'm just trying to identify all pros and cons ;-)	0	0	https://issues.jboss.org/browse/CDI-590

#	Requirement	Questions	Answered	Successfully Answered	URL
2628	CDI-584	To that I could answer, "what's the use of ProcessBean if you can't do anything about the bean tied to the event" ;) ? But I have an example. I was asked to build a prototype extension to create an injection point when a given annotation ({{@CacheContext}}) is found on a field. Here's the portable code I had to wrote to create the feature : * An AnnotatedType [implementation|https://github.com/antoinesd/CDI-Sandbox/blob/fb18eb718ab93f68dcef6b934d70589b789d251b/auto-inject/src/main/java/org/cdisandbox/autoinject/Extension/AutoInjectingAnnotatedType.java] * An AnnotatedField [implementation|https://github.com/antoinesd/CDI-Sandbox/blob/fb18eb718ab93f68dcef6b934d70589b789d251b/auto-inject/src/main/java/org/cdisandbox/autoinject/Extension/AutoInjectingAnnotatedField.java] * and the [Extension|https://github.com/antoinesd/CDI-Sandbox/blob/fb18eb718ab93f68dcef6b934d70589b789d251b/auto-inject/src/main/java/org/cdisandbox/autoinject/Extension/AutoInjectExtension.java] Extension code is small but {{AnnotatedType}} and {{AnnotatedField}} impls are verbose and the whole is cumbersome to read and maintain (3 classes and 178 lines of code) If we only target OWB, the extension code is: {code} package org.cdisandbox.autoinject.Extension; import org.cdisandbox.autoinject.CacheContext; import javax.enterprise.event.Observes; import javax.enterprise.inject.spi.*; import java.util.HashSet; import java.util.Set; /** * Extension looking for the @CacheContext annotation to transform it in an InjectionPoint */ public class AutoInjectExtension implements Extension { private Set<AnnotatedType<?>> atWithCacheContext = new HashSet<>(); /** * This Observer captures all AnnotatedType containing @CacheContext */ public void CreateInjectionPoint(@Observes @WithAnnotations(CacheContext.class) ProcessAnnotatedType<?> pat) { atWithCacheContext.add(pat.getAnnotatedType()); } /** * This observer add an injection point in all bean having @Cachecontext annotaion on some field */ public <T> void CheckBeansForCacheContext(@Observes ProcessManagedBean<T> pmb, BeanManager bm) { AnnotatedType<T> beanAT = pmb.getAnnotatedBeanClass(); if(atWithCacheContext.contains(beanAT)) { Bean<T> bean = pmb.getBean(); Set<InjectionPoint> ips = new HashSet(bean.getInjectionPoints()); for (AnnotatedField<? super T> f : beanAT.getFields()) { if(f.isAnnotationPresent(CacheContext.class)) { ips.add(bm.createInjectionPoint(f)); } } } } } {code} One class, 42 lines.	1	1	https://issues.jboss.org/browse/CDI-584
2629	CDI-580	[~mkouba] Missing default constructor or why? Wouldn't that be the same issue as when applying a normal scope to the product.... WELD-001435: Normal scoped bean class com.knitelius.NonDefaultConBean is not proxyable because it has no no-args constructor	1	1	https://issues.jboss.org/browse/CDI-580
2630	CDI-580	That would be solved by defining the interceptors for the product as an attribute of the produces annotation. {code} @Produces(interceptors={MyInterceptor.class}) {code} Or am I missing something else?	0	0	https://issues.jboss.org/browse/CDI-580
2631	CDI-580	We discussed this feature during last F2F meeting. My suggestion was to introduce a {{BeanInstanceBuilder}} interface to generate an unmanaged instance for a given {{AnnotatedType}}. This is inspired by the existing {{UnManaged}} helper class for injecting and calling lifecycle callbacks unmanaged instances. {{UnManaged}} only uses a class (containing @Inject, @PostConstruct and @Presdetroy annotations) to create instances, {{BeanInstanceBuilder}} would use an {{AnnotatedType}} allowing the use of a synthetic {{AnnotatedType}} containing annotations added by the user like interceptor bindings or {{@Inject}} mentions not present in the original class. We should forbid or ignore other annotations (scope, qualifier on class This would allow applying interceptor on a produced class but also interceptor on method in the produced class. And finally it could also be used to have interceptor on custom beans. Code would look like this {code} public class MyAdvancedProducerBean { public BeanInstanceProducer<MyClass> bip = Builders.beanInstanceProducer(); @Produces @RequestScoped public MyClass produceTransactionalMyClass() { AnnotatedTypeBuilder<MyClass> atb = Builders.annotatedType(Myclass.class) .read(MyClass.class) .addToType(new TransactionalLiteral()); return bip.read(atb.build()).inject().postConstruct().intercept().decorate().get(); // injection, callbacks and proxies creation } public void disposeMyClass (@Disposes Myclass td) { bip.dispose(td); } } {code} The important point is that BeanInstanceProducer is an interface so implementation will use specific proxy feature of Weld or OWB. wdyt?	0	0	https://issues.jboss.org/browse/CDI-580
2632	CDI-580	Can't the disposal part be implicit? The container knows it destroys the produced instance and that it is an intercepted one so it should be able to destroy it alone, no?	0	0	https://issues.jboss.org/browse/CDI-580
2633	CDI-580	What about extensions adding {{Bean<T>}} instances via {{AfterBeanDiscovery.addBean}}? The above proposal looks like it could work for them, but not sure. Is there any existing way really to add interceptors to manually created {{Bean<T>}} instances?	1	1	https://issues.jboss.org/browse/CDI-580
2634	CDI-574	Sorry Emily. I understand roughly your question, but could you detail your use case. You're talking about producer, there is not code for them. I don't understand where the CounterProducerConsumerModified2 bean is defined. From here it looks like a corner case, but I may be wrong. Could you format your code and add your whole example. Thx.	1	1	https://issues.jboss.org/browse/CDI-574
2635	CDI-574	Imo OWB is spec conform as well. It's all a matter of how you interpret the term BDA in the CDI spec. The inconsistency in Weld is a direct consequence of Weld interpreting the term 'BDA' as per-jar (section 12 interpretation option) and not as per-module (section 5 interpretation option). See CDI-18 and other issues for the discussion. In OWB the alternative is 'per ee module' by default. So the whole EE module (e.g. a single WAR) either has the alternative enabled or not. And thus the whole specialized bean is enabled for the whole ee module or not. That way we are a.) much faster (as we don't need to check the invoker chain each proxy invocation) and b.) more consistent. EE-modules usually are 1:1 to ClassLoaders and BeanManagers. Thus for each BeanManager a bean is clearly either enabled or not. It's an old story that the bda-per-jar interpretation is inconsistent within itself. How should that be solved? Will the observer method in the specialized away bean get invoked or not? What if I just add a SINGLE jar which don't have this Alternative enabled? Will the observer method then get invoked again? What about implicit bean archives? Those are almost always present, Just add a single slf4j api or whatever other util jar. In practice this will render disabling beans via specialisation useless. etc, .	1	1	https://issues.jboss.org/browse/CDI-574
2636	CDI-574	I agree that the spec is not perfectly accurate. And yes, OWB solves this black spot by merging beans.xml files internally. I am working with big CDI projects since 6 years and I _never_ had the urge to have a bean enabled in one jar and disabled in another. Nor does any of the 5k++ conference listeners I asked in the many talks I gave around the globe. Actually the per-jar enablement of alternatives and interceptors is the No1 most-hated behaviour of CDI ;) Anyway, lets get back to the original question with @Alternative @Specializes (which is also a prominent example in the spec since CDI-1.0) and 2 jars where in 1 the alternative is enabled and in the other it isn't. What is the expected outcome?	0	0	https://issues.jboss.org/browse/CDI-574
2637	CDI-574	bq. What is the expected outcome? The specializing bean is enabled in one bean archive and disabled in the second bean archive. And you will probably get unsatisfied dependency because the specialized bean is completely ignored.	1	1	https://issues.jboss.org/browse/CDI-574

#	Requirement	Questions	Answered	Successfully Answered	URL
2638	CDI-574	You both ignored my question: what happens with an observer in there? How to disable them effectively? In Weld it is simply not possible it seems.	0	0	https://issues.jboss.org/browse/CDI-574
2639	CDI-574	bq. what happens with an observer in there? Sorry, I'm not sure I understand your question. However, if there is an observer declared on the specialized bean, then this observer is ignored. See also 4.3. Specialization: _"Even if the second bean defines a producer or observer method, the method will never be called."_	0	0	https://issues.jboss.org/browse/CDI-574
2640	CDI-573	PR sent. Code was simplified relying on JDK {{ServiceLoader}} class. We could simplify more since the only provider used is the first found (see {{configuredProvider}} initialization in {{getCDIProvider()}} method). Question is, as we didn't had tests for this code, how can we make sure that the behavior stays the same ?	1	1	https://issues.jboss.org/browse/CDI-573
2641	CDI-573	Hi Antoine. If I read these comments correctly, will CDI.getCDIProvider() visibility be back to public in Weld 3.0.0.Alpha17? Because I just tried updating to Weld 3.0.0.Alpha16 and ran into getCDIProvider() no longer being visible. Thanks!	0	0	https://issues.jboss.org/browse/CDI-573
2642	CDI-572	Also what should happen if a validation fails?	1	0	https://issues.jboss.org/browse/CDI-572
2643	CDI-572	[~mkouba] yes, @Inject might be needed. But one could also use @PostConstruct to init this field. Or use Method injection. All that doesn't matter. Point is that _after_ all the injection was done and the instance is fully initialized it could be useful to invoke JSR-303 on it. > Also what should happen if a validation fails? Well that's a good question. What happens if there is some Exception in @PostConstruct? Guess the same should happen here as well. Or even throw a nice validation Exception. I agree with [[~Sven Linstaedt] that those checks only make sense at runtime. It might as well be that it fails for one request but passes for another one... If you think about business data producers this also makes perfect sense.	1	0	https://issues.jboss.org/browse/CDI-572
2644	CDI-570	This has not much to do with the CDI spec. It's imo a container implementation detail. How long should the shutdown wait on an async event? What if an async event hangs? Should the server still shut down?	1	1	https://issues.jboss.org/browse/CDI-570
2645	CDI-567	Sorry but I don't understand much. I am not really sure what you would like to achieve. Can you elaborate please (some code example might be helpful)? Can you point me to some extension (as an example) which uses ThreadLocal?	1	0	https://issues.jboss.org/browse/CDI-567
2646	CDI-565	erm, now that I think of it. Changing the behavior such that {{null}} is injected instead of having the deployment crash isn't breaking backward compatibility. All old applications deployed has obviously resolved all their unsatisfied dependencies, right? Question is what type of behavior do we want to see for future applications (deployments)? I have talked to several Java developers at this point. They all agree with my comments made about {{null}} earlier. The most intuitive thing is to inject {{null}}, leaving the application a fair chance to actually manage this situation which provingly, is a good feature to have from time to time without complex workarounds using "guru" API:s like {{Instance}}. To turn things around, which issue are we trying to solve, or how exactly are we trying to "help" developers by failing the deployment for unsatisfied dependencies? I can only see the down sides of this behavior. If we know what's right, then as men (and women for some of us!), we are obliged to do the right thing without paying attention to the "fear of change". I cannot stress enough that Java EE is a complex technology stack and if we are to stand a chance in a competitive environment, we have to stop looking for quick fixes and workarounds. Isn't it possible to cast a vote on the JBoss forum or something, asking people that don't already know the answer, what they think should happen to an unsatisfied dependency? That is what I asked my developers. All but one new guy said {{null}}.	1	1	https://issues.jboss.org/browse/CDI-565
2647	CDI-565	What do we expect to happen here: {code} Person person = new Person(); String name = person.getName(); {code} Yeah, only a problem if we try to dereference the name! =) CDI is one huge "getter". No difference. Most if not all getter methods we write every day return {{null}} and let the client deal with it. If the client require a value and rather have his application blow up with a more specialized exception, then he probably call a more niche method: {code} String name = person.requireName(); {code} I.e, here's what I believe we should do if and only if we want the deployment to fail: {code} @Inject @Required String name; {code} Or maybe even more straight forward: {code} @Inject @DeploymentRequired String name; {code} This way, we can even configure a bit how we want the deployment to fail: {code} @Inject @Required(throwClass=MySuperAwesomeException.class) String name; {code} This "issue" really comes down to what we think is the most intuitive behavior and, whether or not it really brake backward compatbility and if so, how bad would that actually be. I believe 1) {{null}} is the most intuitive thing to do here and it is a good thing leaving the application to deal with it. 2) We don't brake backward compatibility at all. Let me know what you think.	1	1	https://issues.jboss.org/browse/CDI-565
2648	CDI-565	Or {{java.util.Optional}} since JDK 8. Yeah, I think we all agree that the ideal world make a minimal use of {{null}}. But I sure don't think we can ever manage without {{null}}. In fact, {{null}} has some really good use cases. What would you think if {{Map.get(Object)}} threw an exception instead of returning {{null}} when the mapping doesn't exist? Well, the CDI container is one huge map. I ask for a mapped value given a set of types and/or qualifiers - the keys. What is the difference? Even if you do think {{null}} is absolute evil, then it is with us already for better or worse. Convenience, or what we together believe is the most intuitive approach, has to take precedence. Fail-fast is a good practice, but having it implemented just because one can may be a brutal way of enforcing dogma upon others. I see this "fail-fast" approach as both braking intuition and reducing the service offer. I'm not saying I am right. I am trying to find out why some of us developers has to pay the price, for what benefit, for what gain? You say it "helps developers" by "avoiding surprising NPE:s at runtime". All exceptions not handled is a surprise, whether that happens during deployment or runtime. But let me point out that if anything is a surprise, then it is the exception types CDI define and all vendor-specific messages the application servers output in the log file, and you know I am right on this point. New developers in particular will have to walk though a lot of struggles until they have a solid understanding of the CDI vocabulary like "ambiguous" and "unsatisfied". Compare that with a a {{java.lang.NullPointerException}} in the log file with a clear stack trace to the exact source code line where the reference was dereferenced. What a NPE is and how to deal with it is probably what all new developers learn on their very first day. Non-developers like to praise programmers as being some kind of super humans. What I often find myself saying in response to my friends, and experiences I think you share with me, is that I am lazy. I am not really a problem solver. I am not really smart. I'm trying to find the best architecture such that the design can do the thinking for me. I am trying to reduce complexity, trying to simplify and automate everything. Most importantly, I will not try to fix a problem before I have one. And that is super important. By trying to fix a problem before I have it, I might end up introducing one instead. Been there done that many times over. Recognize the situation? That's why we are having this discussion. "Premature optimization is the root of all evil" and ya'll know it =) Least you can do is {{@Optional}}.	1	1	https://issues.jboss.org/browse/CDI-565

#	Requirement	Questions	Answered	Successfully Answered	URL
2649	CDI-565	I'm not sure I fully understood you on that "dropping behavior"-thing. NPE:s is kind of what I try to advertise here =) {{null}} is and will always be a valid candidate to represent the absence, or lack of, a value. *YES*, we all agree that it is easy to misuse {{null}} and {{null}} should be avoided if we so can. My point is definitely not that {{null}} is a god sent gift. Please reread my comments and see if you can answer the questions I asked, for example, what is the difference between {{Map}} and the CDI container (when it comes to dependency lookup)? Also compare {{EntityManager.find()}} from the JPA specification which I believe is another comparable interface. Everything is an API. I don't understand this notion really or how it relates to the topic of discussion. I googled the CDI specification, {{null}} return values are used everywhere. Would you like to advocate that we setup tens of new exception types into our type system instead? Speaking of API:s, I would appreciate an {{Optional}} return type from a service only if it is conceivable and often enough the case that no return value is present. For example: {{tryCatchAFish()}}, {{getLotteryWinner()}}. Or, if it is conceivably so that the client doesn't bother about the presence of the value. For example: {{myDataSource.getSomeData().ifPresent(someValueContainer::setThatField)}}. I myself overused {{Optional}} once it was out in JDK 8, singing praise to Jesus every minute of the day. Only to discover that my overuse of it introduced complex boilerplate code. Eventually I got so lazy that I dereferenced those optionals anyway ({{Optional.get().doSomething()}} and so it was only a matter of time before I found myself back in square one. But this time, not only did I have NPE:s, all my clients had to make an extra method call all over the place for no apparent gain. More than anything else, I learnt by my own mistake that if a parameter is optional for the client, then accept {{null}}. Don't force the clients to type {{Optional.empty()}} instead of the oh so simple blueish keyword {{null}}. Null as parameters make the life easier for the implementations too. Also, if you know a "modern language" that had any success avoiding {{null}}, please let me know. That would be interesting. For example, how would I create a new 10 element array? Am I required to have all ten elements right of the bat? Guys, what is the most intuitive thing to do here? Why do you believe that the dependency lookup should have different semantics than {{Map.get()}}?	1	1	https://issues.jboss.org/browse/CDI-565
2650	CDI-563	bq. I'm not sure to follow him there, especially when CompletionStage has a dependency on CompletableFuture with the toCompletableFuture() method It's kind of a strange dependency, intended probably as a compatibility/interoperation feature. However, implementations are not required to support this conversion. bq. Martin Kouba, you're probably the best person to know if this default Executor is the best for us or if we should use our own Executor... I don't think {{ForkJoinPool.commonPool()}} is a good fit for Java EE. But it's ok for Java SE. I'm for: * having {{CompletionStage}} in the CDI API * declaring explicitly that *the stage's default asynchronous execution facility is not defined*, in other words it is implementor and environment specific * declaring explicitly that *the context propagation is not guaranteed for ANY execution* (default execution, default async execution and custom async via a supplied Executor) What do you think?	1	1	https://issues.jboss.org/browse/CDI-563
2651	CDI-563	{quote}bq. I'm not sure to follow him there, especially when CompletionStage has a dependency on CompletableFuture with the toCompletableFuture() method It's kind of a strange dependency, intended probably as a compatibility/interoperation feature. However, implementations are not required to support this conversion.{quote} I'm not sure, but I think we should, since it's the only way to access the {{get()}} method to allow user to wait for the outcome... {quote} I don't think {{ForkJoinPool.commonPool()}} is a good fit for Java EE. But it's ok for Java SE. {quote} +1 {quote}I'm for: * having {{CompletionStage}} in the CDI API{quote} Yes we also agree here. {quote}* declaring explicitly that *the stage's default asynchronous execution facility is not defined*, in other words it is implementor and environment specific {quote} My favorite option as well ;). I'm just not sure where to add it in the chapter {quote}* declaring explicitly that *the context propagation is not guaranteed for ANY execution* (default execution, default async execution and custom async via a supplied Executor) {quote} Regarding propagation, I thought that the mention in [10.5.3. Observer method invocation context|http://docs.jboss.org/cdi/spec/2.0.EDR1/cdi-spec.html#observer_method_invocation_context], was explicit enough: bq. If the observer method is asynchronous, it is called in a new lifecycle contexts, a new transaction context but the same security context as the invocation of Event.fireAsync(). What do you think we should add?	1	1	https://issues.jboss.org/browse/CDI-563
2652	CDI-563	bq. My favorite option as well . I'm just not sure where to add it in the chapter I would add a new paragraph in 10.2.3. The Event interface... something like "The fireAsync() methods return a CompletionStage whose default asynchronous execution facility is specific to the container implementation." bq. What do you think we should add? I don't like the wording "called in a new lifecycle contexts". What exactly does it mean? Maybe we should just rephrase this.	1	1	https://issues.jboss.org/browse/CDI-563
2653	CDI-563	{quote}bq. What do you think we should add? I don't like the wording "called in a new lifecycle contexts". What exactly does it mean? Maybe we should just rephrase this.{quote} I'm not a big fan of this wording either. But it came from discussion with Jozef, and the fact that for the sake of non-denormalize the spec we shouldn't repeat what is said in [6.7|http://docs.jboss.org/cdi/spec/2.0.EDR1/cdi-spec.html#builtin_contexts]: bq. The context associated with a built-in normal scope propagates across local, synchronous Java method calls. The context does not propagate across remote method invocations or to asynchronous processes. Perhaps we could find a way to have a clearer phrase pointing to this assertion. It lets the possibility for a third party to create a normal scope supporting propagation.	1	1	https://issues.jboss.org/browse/CDI-563
2654	CDI-562	I think there are two similar open issues which might offer a better solution for this use case: # CDI-515 - allow to obtain Bean metadata from {{javax.enterprise.inject.Instance}}, so that if a user wants to get the metadata the programmatic lookup would help # CDI-10 - if we introduce a marking interface for client proxies, e.g. {{BeanProxy}}, we could also probably add {{getBean()}} method I'm not so sure it's a good idea to add a method to the BeanManager. Let's have the following producer: {code:java} clas FooProducer { @Alpha // qualifier @Produces Foo produceAlphaFoo() { ... } @Bravo // qualifier @Produces Foo produceBravoFoo() { ... } } {code} How to identify a bean for any {{Foo}} instance? We would have to use some internal construct even for contextual instances of dependent beans with no interceptors/decorators. And that's not possible - such a bean does not have to be proxyable. Or am I missing something?	0	0	https://issues.jboss.org/browse/CDI-562
2655	CDI-558	We should do this for all SPI metadata in the spec. That would bring us: * {{AnnotatedTypeBuilder}} * {{BeanAttributesBuilder}} (Ancestor of BeanBuilder to be used in ProcessBeanAttributes) * {{InjectionPointBuilder}} * {{ObserverMethodBuilder}} Less sure about: * {{InjectionTargetBuilder}} (we already have a factory) * {{ProducerBuilder}} (Factory here as well) The question is how should we retrieve these builders. I don't think it's a good idea to restrict their access to lifecyle event since we could need them outside an extension. {{BeanManager}} seems the obvious place but there are already too much method in it. Perhaps it's time to create an interface for such service method, that the BeanManager would implement and that could be injected as well to provide access to these builders. WDYT ?	0	0	https://issues.jboss.org/browse/CDI-558

#	Requirement	Questions	Answered	Successfully Answered	URL
2656	CDI-558	bq. The question is how should we retrieve these builders For {{BeanBuilder}}, {{BeanAttributesBuilder}} and {{ObserverMethodBuilder}} I think it is a good idea to restrict the access to container lifecycle events (or [CDI SE builder|https://github.com/weld/core/blob/master/environments/se/core/src/main/java/org/jboss/weld/environment/se/Weld.java#L488] respectivelly). Maybe we should create a table and for each SPI add some description, purpose and/or use-case, possible implementation problems, etc.	0	0	https://issues.jboss.org/browse/CDI-558
2657	CDI-558	{quote}CDI SE builder may allow to access these builders - see my previous comment.{quote} Sorry, I missed that. {quote} bq. or if we decide later to allow addition of some meta data at runtime Well, if we decide later we could add this later. I think it's a good start to have it only accessible where it's needed ATM. {quote} Ok, but that doesn't solve the place for {{AnnotatedTypeBuilder}}. Should it be one of the few helpers classes of our API in {{javax.enterprise.util}} package beside {{AnnotationLiteral}}?	1	1	https://issues.jboss.org/browse/CDI-558
2658	CDI-558	Wrt {{AnnotatedTypeBuilder}} what do you mean by "from outside CDI"? What about to introduce {{AnnotatedTypeBuilder}} as a built-in bean?	1	1	https://issues.jboss.org/browse/CDI-558
2659	CDI-552	Isn't this what DeltaSpike provides already with {{BeanProvider.injectFields(myObject)}}? https://deltaspike.apache.org/documentation/core.html#_beanprovider This method clearly only injects the dependencies, it does not invoke any lifecycle methods and does not add the object to any scope. It would be useful to make this into the CDI spec, such as {{CDI.current().injectDependencies(Object o)}}. It might also make sense to execute the lifecycle methods by {{CDI.current().invokePostConstruct(Object o)}} and {{invokePreDestroy(Object o)}, although they may be executed as directly on the object.	0	0	https://issues.jboss.org/browse/CDI-552
2660	CDI-547	Yes maybe the third option looks as the best one to me. Because when you introduce new method then there is a question: What all should it return in case of ASYNC version? It would make sense to return all ASYNC + SYNC observers but I think this will bring only more confusion.	0	0	https://issues.jboss.org/browse/CDI-547
2661	CDI-540	I fully agree with this. Why developer needs to ask {{CDIProvider}} whether {{CDI}} is initialized? And what does it mean actually? I think {{CDIProvider.isInitialized()}} method would make more sense as method on {{CDI}} and named like {{isRunning}} (since we have shutdown method).	0	0	https://issues.jboss.org/browse/CDI-540
2662	CDI-535	This is really interesting and useful. Can this make it to CDI 2.0 ?	1	1	https://issues.jboss.org/browse/CDI-535
2663	CDI-535	Good point. Can we enhance the Priority annotation or would that take ages to be accepted?	1	1	https://issues.jboss.org/browse/CDI-535
2664	CDI-535	Why? Isn't that what priority is for? Or are you referring to implementing Comparable in this object?	1	1	https://issues.jboss.org/browse/CDI-535
2665	CDI-535	I think {{@Priority}} must be extended to support producer methods and fields, but for now I believe it´s fine to support it in a type-level usage only, WDYT?. For example, how different would a bean be from an interceptor (or an Alternative) that uses this same approach?	1	1	https://issues.jboss.org/browse/CDI-535
2666	CDI-535	[~marembo2008], I agree, so in order to avoid instantiating the objects and comparing against each other, I think using {{@Priority}} or another kind of comparison (eg. annotation) in the type declaration would be a better choice.	1	1	https://issues.jboss.org/browse/CDI-535
2667	CDI-535	I would say more than anything, my concern is on the reuse/misuse of {{Priority}} and using something new like {{OrderBy}} would not be an issue. Priority introduces misuse in my opinion. Even then, if you have multiple beans, all from the same field (all dependent scoped for example), how would you order them?	1	0	https://issues.jboss.org/browse/CDI-535
2668	CDI-535	[~meetoblivion] [~gastaldi] I don't think {{@Priority}} would introduce a misuse here: {quote} The Priority annotation can be applied to classes to indicate in what order the classes should be used. The effect of using the Priority annotation in any particular instance is defined by other specifications that define the use of a specific class. {quote} But as I said we would not be able to add priority to producer methods and fields. So it's -1. A new annotation might work as well. {{OrderBy}} would be mixed up with {{javax.persistence.OrderBy}}. bq. Even then, if you have multiple beans, all from the same field... How could that happen?	0	0	https://issues.jboss.org/browse/CDI-535
2669	CDI-527	+1. This would simplify CDI integration in a lot of cases. Do we have the same issue with / should talk about interceptors too?	1	1	https://issues.jboss.org/browse/CDI-527
2670	CDI-527	{quote}EJB does allow this.{quote} Can you elaborate on how EJB allows that? Based on the following: {quote}Only private methods of the bean class and any superclasses except java.lang.Object may be declared final{quote} I interpret the EJB spec as explicitly prohibiting this.	1	1	https://issues.jboss.org/browse/CDI-527
2671	CDI-527	[~jharting] I digged a bit deeper. Seems EJB-3.2 introduced this (non backward compat) restriction. But *only* for NIVs! Previously this restriction was only for 'Business Methods'. Which always have been only public methods (this is what is different to CDI).	1	1	https://issues.jboss.org/browse/CDI-527
2672	CDI-527	I still fail to see how EJB "allows this". Can you provide a code sample of a session bean with non-private final method that is valid according to EJB 3.2?	1	1	https://issues.jboss.org/browse/CDI-527
2673	CDI-527	@Jozef: just use a parent class with a protected final method. This is not forbidden. That said 3.2 should remove this constraint since that's a regression from 3.1. last point: why do we care about EJB at all? CDi just allow an integration with EJB but EJB shouldn't be the first model (in particular because it defines this contract definition which is not the case in CDI where internal methods are seen as well).	1	1	https://issues.jboss.org/browse/CDI-527
2674	CDI-527	I am not sure I follow on this. On the NIV's it's explicitly stated in {{4.9.8 Session Bean's No-Interface View}} that: {quote} Only private methods of the bean class and any superclasses except java.lang.Object may be declared final. {quote} There is no problem from CDI point of view AFAIK. In the case of interface view types I can't see any problem as well because the proxy instance is created for the interface type (not impl) which has all methods public and not final. What's the reason for CDI to behave differently? In fact I think we can't supress EJB validations anyway.	1	1	https://issues.jboss.org/browse/CDI-527

#	Requirement	Questions	Answered	Successfully Answered	URL
2675	CDI-527	I wrote down a summary of the problem and possible solutions for proxying with final fields: Java8 introduced a few final methods to existing classes. This now leads to a javax.enterprise.inject.UnproxyableResolutionException at bootstrap if you have a producer method or field for those classes. EJB also seemingly allows this. But it's specced on a completely different level. EJB disallows final 'business methods'. That means that whenever you have an interface only those methods will get proxied. The methods of the implementation do not matter at all. With NIV all methods are 'business methods' and must be proxyable imo (can someone verify this?). I personally like the non-proxyable rule. It really prevents tons o user errors. As you all know CDI heavily relies on subclassing proxies. Means we override all methods from the proxied type. But we obviously cannot override (and thus proxy) static, private and final methods. While static and private methods do not really matter, all non-private non-static final methods do. The problem appears when people invoke such a final method on the proxy. In that case they will just hit the proxy and not the Contextual Instance. To prevent the user from making this error we throw the UnproxyableResolutionException at bootstrap. We now have a few options to deal with that: A.) no it's not a problem at all. - We often got the report for ConcurrentHashMap. But people can already just use the ConcurrentMap interface as return type in their producers. - Also the report that this only makes a problem since Java8 is plain wrong imo. ConcurrentHashMap has the following method since at least Java7: final V put(K key, int hash, V value, boolean onlyIfAbsent) - There are a few situations where it hurts though. E.g. extending ConcurrentHashMap with own functionality. This would require introducing an own interface for this new functionality. B.) Generally allow proxying of beans with final methods. - I'm rather -1 for this. The current behaviour was introduced exactly because people made too many errors and blew up at runtime. Removing this rule would really make things worse. Because people might then add new final methods. And if you call those then you will only hit the proxy class and not the Contextual Instance. C.) Introduce an annotation to explicitly declare Beans to be proxyable - Something like @AllowProxying Can be applied to producer methods, producer fields and managed beans. - That would make it explicit for programmers that they leave the sunshine path. People will blow up with the UnproxyableResolutionException and can then decide whether they want to provide that bean anyway or not. - Con: Only the author of the Producer or managed bean knows that he added it. Thus other users will first need to look at the producer to become aware. But that might not be so much of a problem in most cases as it's only a matter of JavaDoc. D.) Introduce an exclude list for knowingly 'bad' classes. - We could e.g. introduce that all classes from java.* and javax.* are proxyable by definition. - Con: hard to know what's going on and why the code doesn't work if a user calls such a final method.	1	1	https://issues.jboss.org/browse/CDI-527
2676	CDI-527	[~emilyj] never has really been supported on CDI side, JDK broke some API (in a more vicious way than this jira) but doesnt mean CDi has to do anything. Do you have an example of a broken application cause of *application* code (and not JVM)?	1	1	https://issues.jboss.org/browse/CDI-527
2677	CDI-527	[~emilyj] Also remember that CDI-1.2 is only specified for Java7. Also note that EJB doesn't support this neither for NIVs... How would D look like in practice? System properties and excludes can be done in a non-portable way anyway. But adding system properties to any spec is in 99% a sign of a broken design. -1 for B. Just added it for completeness ;) We also could have another option E.) introduce some class excludes via beans.xml But that would have the same issues like D imo. Imo C is the most explicit and easiest for users to grasp.	1	1	https://issues.jboss.org/browse/CDI-527
2678	CDI-527	Option A is easy but it does not help the customers with legacy applications. It is not their fault in the first place. The JDK change is beyond their control. If we want to make CDI popular, we should help the customers who are stuck if they are unable to update their legacy applications. Option D: Support proxy on all java.x or javax... classes. Since the system property does not sell well, how about introduce a method setAllowProxyingOnUnProxyable(boolean allowProxying) on CDI.java? If not set, false is the default.	1	1	https://issues.jboss.org/browse/CDI-527
2679	CDI-527	[~french_c] Imo a perfectly valid use case. In the case you described only the option with the excludes or with the special Annotation would be possible. [~emilyj]] Not sure about this. Most of this classes already have final protected methods as of Java7. Which class do you have in mind?	1	1	https://issues.jboss.org/browse/CDI-527
2680	CDI-527	>>Not sure about this. Most of this classes already have final protected methods as of Java7. Which class do you have in mind? HashMap#initHashSeedAsNeeded, as mentioned in your description.	1	1	https://issues.jboss.org/browse/CDI-527
2681	CDI-519	what do you suggest? # to forbid this or, # to explicitly state that {{Instance<T>}} can be used to destroy any dependent scoped bean instance of T. In your example if we go for 1, there will be no solution to destroy the instance you requested thru {{CDI}}	0	0	https://issues.jboss.org/browse/CDI-519
2682	CDI-518	Do you mean "8.3 Class Loading Requirements"? If so then don't forget that ServletContext is not an application class. Therefore, this needs to be defined explicitly.	1	0	https://issues.jboss.org/browse/CDI-518
2683	CDI-518	Oki, getting tricky now. I think we should take a step back and finally do our homework regarding EAR handling. * What classes are seen where. * How do we handle/deal with isolation and visibility constraints of CDI running in EARs. * How does Extension visibility look like * Which classes should get delivered to which Extension * getBeans in EARs * @Named in EARs (thats a TOTAL mess right now) * getInterceptors, decorators etc in EARs ... After that the observer resolution will solve itself pretty quickly.	0	0	https://issues.jboss.org/browse/CDI-518
2684	CDI-515	That would be way more complicated if possible at all. Let's assume you have two producers for dependent {{Foo}} and each specifies a different qualifer. Then you inject {{Instance<Foo>}} and select Instance for qualifier {{@A}} and get instance of Foo, and select Instance for qualifier {{@B}} and get instance of Foo. How would you differentiate these instances? What happens if you pass Foo {{@B}} reference to the Instance for Foo {{@A}}? You would have to track all the instances provied by each Instance.	1	1	https://issues.jboss.org/browse/CDI-515
2685	CDI-515	I don't see that limited to instances provided by a given Instance<>. Like you can use Instance.destroy() to destroy any instance today. A use case could be: obtaining metadata for a bean defined by a class you don't own. Today we store all bean instances to manage their lifecycle, so how difficult would it be to compare a given instance with these stored instances and retrieving the bean used to generated it. It's probably costly but could be very useful.	1	0	https://issues.jboss.org/browse/CDI-515
2686	CDI-514	Right, for CDI 2.0 this will need to be revisited. As for the performance impact I don't really buy this. How much time does it take to do a instanceof or class identity check? Especially given that method is usually called with 0, 1 or 2 qualifiers? We should prefer correctness over dubious performance optimizations.	1	1	https://issues.jboss.org/browse/CDI-514
2687	CDI-513	{quote} Passivating scopes must be explicitly declared @NormalScope(passivating=true). {quote} This would effectively render the whole addScope Extension method useless... How would you add a scope for the JSF @ViewScoped annotation in that case? (This annotation is neither @NormalScoped nor @Scope). I suggest we remove this whole sentence as it contradicts other spec parts.	1	0	https://issues.jboss.org/browse/CDI-513

#	Requirement	Questions	Answered	Successfully Answered	URL
2688	CDI-513	Let's approach it from a different angle: does the combination non-normalscope AND passivating make sense? Imo it does. This are all the use cases where CDI does not handle the lifecycle (nor applies a proxy) but we would like to get the Serializable check. If anyone comes up with a blocker for this case then I'm perfectly happy to add a check that addScope(Xxx.class, false, true) is an illegal situation. If not then it must get supported.	1	0	https://issues.jboss.org/browse/CDI-513
2689	CDI-504	What's wrong with {{java.lang.System.getProperty()}} and {{java.lang.System.getenv()}}? I mean the main advantage of DS Configuration is IMO the resolution mechanism and the ability to extend the set of {{ConfigSource}} s. If you only need system properties and env variables a simple producer method and qualifier (or even a simple util method) would be sufficient. In other words I see no reason to standardize this.	1	1	https://issues.jboss.org/browse/CDI-504
2690	CDI-504	+1 with few minors notes: - @Env and @SystemProperty maybe (instead of EnvProperty) to match JVM naming (@Property would be less meaningful IMO) - conversion should be defined on a close set of types (primitives, wrapper + few java.lang/java.util + Collection + Map only) to ensure portability - not sure it is matre enough to support a good SPI ATM and if then it shouldnt be in config spec tamaya tries to prepare. wdyt?	1	1	https://issues.jboss.org/browse/CDI-504
2691	CDI-504	{quote}So we open the door to include a converter SPI in the spec. I'm worried about the out of scope and anti-diet aspect of the spec here...{quote} Maybe it's [high time to standardize Java EE converters?|http://arjan-tijms.omnifaces.org/2014/08/high-time-to-standardize-java-ee.html] ;) Not particularly looking forward to add CDI to the list I jotted down above.	1	0	https://issues.jboss.org/browse/CDI-504
2692	CDI-502	{quote} An actual type is a type that is not a wildcard nor an unresolved type variable. There is therefore no rule in 5.2.4 where a wildcard would be considered on the bean type side. {quote} Thanks, now I understand. It would be nice if legal bean types would be the types you can instantiate in Java (That doesn't covers all types a variable can have(eg. `List<? extends Optional<?>>`), but covers types, like `List<Optional<?>>`). And assignability to be exactly the same what in Java is with the exception for raw generic types. I guess I could report feature request. {quote} Correct. Not sure why it is so. {quote} Should I report it as a clarification or bug?	1	0	https://issues.jboss.org/browse/CDI-502
2693	CDI-499	{quote}For backward compatibility reason the possibility to invoke an observer asynchronously should be let to the observer{quote} What exactly are the reasons for this? I can think of one: - observer method injecting a @RequestScoped (or similar) bean Are there other arguments for this? {quote}3. Observer bound to a transaction phase these observer will be invoked in the right transaction phase but asynchronously{quote} How's that different from Event.fire()?	1	1	https://issues.jboss.org/browse/CDI-499
2694	CDI-499	{quote} - So to be notified asynchronously an observer should have {{asyncSupported}} member to true. otherwise it will be called synchronously. I think a call to fireAsync() should be processed asynchronously by all observers. But I have the feeling that it is in contradiction with this statement {quote} ... {quote} {color:blue} For backward compatibility reason the possibility to invoke an observer asynchronously should be let to the observer {color} What exactly are the reasons for this? I can think of one: - observer method injecting a @RequestScoped (or similar) bean Are there other arguments for this? {quote} I'm not very happy with this approach as well : It could bring confusion for end user. This possible backward compatibility issue was raised by [~meetoblivion] or [~struberg] during last IRC meeting. I'd be glad to have a better solution here. {quote} {color:blue} 3. Observer bound to a transaction phase these observer will be invoked in the right transaction phase but asynchronously {color} How's that different from Event.fire()? {quote} probably none ;). How did you handle that in Weld Alpha [~jharting]?	1	1	https://issues.jboss.org/browse/CDI-499
2695	CDI-499	Easy to add, huge impact on the world. If you want "basic tasks", then why do we want to add an asynchronous feature to begin with? Josés point about being true to the nature of asynchronous Java API:s are undoubtedly a good one. Either don't implement the feature, or go all the way and deliver a great product. Time to step up our game if we want to compete with other technologies available on the market. I don't think {{asyncSupported()}} should default to {{false}}. It should default to {{true}}. Firstly, backward compatibility must not stand in the way of technology advances. That's kind of why the major release version become a number "2". And if I'm thinking clearly here, then {{true}}, which is without doubt (?) the most natural fit, does not break backward compatibility. Legacy code already call {{Event.fire()}}. If I write new and modern code and call {{Event.fireAsync()}}, then isn't it it my responsibility to make sure that if I use old event container types and thus target old observer methods, that the observer can handle asynchronicity? To put things in another perspective: An observer method that does *not* support asynchronicity is knowledgeable and dependent on who fired the event, how and where. He is dependent on a particular context. That sounds like a design smell to me. But, the one firing the event should have a saying if he want the call to be blocking or not. But imagine the confusion a default false would produce: calling {{fireAsyn()}} is most likely to actually block. Depends on the observer methods! Simply put it, the developer must scan through the entire code base every time to make sure his asynchronous request really will be honored. Even worse, he can never be sure that will always be the case. 5 minutes later, another developer add an observer method and will *by default* make all unknowingly asynchronous events of this type in the system turn synchronous. At best, this will have such an enormous impact on user experience and create large enough bottlenecks to be discovered. At worst, the application will slowly die a painful death as the number of parallell tasks going down the serialized lane increase. Simply put it, don't fix what's not broken. Letting the observer methods have a saying smells bad. Having a method be called {{fireAsync()}} but who in fact work nondeterministically smells even worse and can create all kinds of drama. I think {{asyncSupported}} should be dropped. If it must be in place, then let it by default be {{true}}. If so, it will be the one implementing the observer method that might destroy the application, not the unknowing poor guy who wrote the world's most beautiful code, thinking a method called {{fireAsync()}} actually work as advertised. As an example, I don't use {{@Asynchronous}} from EJB because of the nondeterministic behavior. I am not guaranteed that the container actually run my task. He might reject the task. And the way exception handling work in EJB (which is contrary to the spec's own design goal), I can never be sure that my portable application code can determine why the EJB call crashed. So basically, I must use {{ManagedExecutorService}} and be prepared to catch a {{RejectedExecutionException}}. So, EJB asynchronicity is dead and of no use for serious application code that want to be portable. Please don't make the same mistake with CDI. Oh you should also specify that fireAsync() might throw RejectedExecutionException? h2. {{asyncSupported}} To continue my "ranting" - or improvement ideas =) .. I see what might be a problem if asynchronous observer methods are prioritized, submitted before synchronous observers are called. Today, if a non-transactional observer method crash, then the event processing is aborted - no more observers called. But what if we already submitted asynchronous tasks to a queue somewhere, are they aborted? Are they guaranteed to be aborted if execution of the tasks has not begun when a synchronous observer crash? If we want to keep the "feature" that an observer method should be able to stop propagation, then asynchronous observers should be called ??after?? all synchronous observers has been called. OR be very clear in JavaDoc and in the new specification that it is *not defined* whether or not synchronous observers that crash stop event propagation to asynchronous observers. I understand that it is more pleasing to have the asynchronous tasks be submitted as fast as	0	0	https://issues.jboss.org/browse/CDI-499
2696	CDI-497	I'm not sure to understand your issue. Can you give a use case?	1	1	https://issues.jboss.org/browse/CDI-497
2697	CDI-496	Isnt the issue just the fact "contract" of Beans is not exlpicit? contract of an EJB is its API (Local, Local Bean or Remote), for a ManagedBean it is its explicit or implicit @Typed, etc...In summary it is bean.getTypes() no?	1	0	https://issues.jboss.org/browse/CDI-496

#	Requirement	Questions	Answered	Successfully Answered	URL
2698	CDI-496	excepted the CDI container doesnt have to know the EJB actual impl type so how can it validate it? I think it should be done in EJB container - anyway it should be done in EJB container anyway and lifecycle between EJB and CDI container doesnt say today which one would fail first	0	0	https://issues.jboss.org/browse/CDI-496
2699	CDI-495	I am just wondering if it's worth of specifying. Wasn't there any special case when it is problematic?	1	1	https://issues.jboss.org/browse/CDI-495
2700	CDI-492	[~edburns] I understand those concerns, but I think there are solution for most of them. The points given don't take into account evolutions that could be done in CDI and JSR 330 to make it possible. To take all the points and give first ideas to solve them : *backward compatibility* : introducing a qualifier for these beans (like it's done in Deltaspike to work with CDI 1.0 and 1.1) is a first lead like {{@Servlet}} to tell that these beans are owned by Servlet spec. *Dependency on CDI* ; as you know we're going to work on a new AtInject version. We could imagine having the servlet impl only be dependent on AtInject (which will stay very small : around 10 interfaces / annotations). The only thing that is missing in AtInject today is the concept of {{@Producer}} that could be added to this new version (dagger want to move its {{@Provides}} annotation to AtInject. The integration can off course be done at spec level (the worse case would be to have a dependency on CDI API). So I really don't understand Greg point. For CDI 2.0 our approach is to talk with other specs to see what they need to improve their CDI integration. So let's talk about these needs before deciding it's not doable.	1	1	https://issues.jboss.org/browse/CDI-492
2701	CDI-486	Hi, This JEP(http://openjdk.java.net/jeps/176) proposes a new way of enforcing the security permissions. Does this mean that the SPI methods will be annotated with @ CallerSensitive How does one update the issues list with these types of queries assuming it is valid ? I am still learning the ropes. Thanks, Mohan On Thu, Oct 2, 2014 at 12:44 PM, Jozef Hartinger (JIRA) <issues@jboss.org>	1	0	https://issues.jboss.org/browse/CDI-486
2702	CDI-485	And what about {{INSTANCE}} for annotations having members like {{@New}}, {{@Typed}}, {{@Destroyed}} or {{@Initialized}} ?	1	1	https://issues.jboss.org/browse/CDI-485
2703	CDI-480	Can we mark this as resolved now?	0	0	https://issues.jboss.org/browse/CDI-480
2704	CDI-479	I've found this. # We have {{org.jboss.cdi.tck.tests.event.broken.observer.beanNotManaged.ObserverMethodOnIncorrectBeanTest}} where we check that observer is disabled (not resolved) at non-managed bean and at abstract class with abstract method observer which is overriden in extending class. # We have {{org.jboss.cdi.tck.tests.event.observer.ObserverNotificationTest#testObserversNotified}} where we check that observer is disabled at bean which is not selected alternative. I think we can expand both cases in TCK and probably close this issue. What do you think [~antoinesabot-durand]?	0	0	https://issues.jboss.org/browse/CDI-479
2705	CDI-475	How is this different from CDI-472?	1	0	https://issues.jboss.org/browse/CDI-475
2706	CDI-473	Indeed, {{@Startup}} may be a candidate. I see two questions in order to determine the best way to capture that: * Does that apply to {{ApplicationScoped}} bean only? * It that worth enough to depend on an other specification for that purpose only?	1	1	https://issues.jboss.org/browse/CDI-473
2707	CDI-473	There are some downsides to using Initialized, if not done safely. For one, you don't want to do the initialization logic in there unless you're ok with it being done multiple times (since anyone can fire this event). It's better to still put it in PostConstruct, and leave this method as is. Which begs the question, is there a benefit to having empty methods just to declare observers?	1	1	https://issues.jboss.org/browse/CDI-473
2708	CDI-471	[~antoinesabot-durand] I think this should be one of the issues with highest priority for upcoming CDI F2F meeting. WDYT?	1	1	https://issues.jboss.org/browse/CDI-471
2709	CDI-470	Hum... not that clear to me, but well. Shall it stay like that ? I tend to add a package-level {{@Veto}} on my entities ({{org.mystuff.model}})... but in big projects my model is usually spllitted into sub packages ({{org.mystuff.model.customer}}, {{org.mystuff.model.order}}, ...). Recursive veto would be interesting. Any thoughts ?	0	0	https://issues.jboss.org/browse/CDI-470
2710	CDI-469	[~meetoblivion] I'm sorry but I don't follow you. Could you be more specific about what "nonbinding producer method" means? Simple example would be helpful as well ;-)	1	1	https://issues.jboss.org/browse/CDI-469
2711	CDI-457	BTW, @TransientReference would work in the producer method solution, if it were allowed on methods/fields, rather than params only. WDYT about allowing @TransientReference on producer methods? Allowing the producer to control the behavior of the injected reference, rather than the injection point controlling it?	1	1	https://issues.jboss.org/browse/CDI-457
2712	CDI-457	Please let me ask one question: why should the container know about dependent scoped beans which are _not_ injected into some normalscoped bean? We do not spare any resources for them (at least in OWB). Even if they have interceptors or decorators applied. Btw, there is kind off such a mechanism already. All you need to do is to keep the Bean<T> and the CreationalContext you used to create the contextual instance. Then just use Bean.destroy(T, cc);	1	1	https://issues.jboss.org/browse/CDI-457
2713	CDI-457	[~meetoblivion] Well, I still don't understand why the new interface is needed. {quote} @UnManaged would be a new annotation indicating that the returned object is not contextual {quote} IMO does not make sense, if you inject dependent SomeClass into some normal-scoped bean, it's not non-contextual. {quote} Disposable object goes out of scope. {quote} What does it mean? How would this work? {quote} I suspected that OWB didn't wrap them into contextuals, whereas weld seems to (at least in the 1.1.x line). {quote} Weld returns a subclass if an interceptor or decorator is associated. Did you mean this by "wrapping into contextuals"?	1	1	https://issues.jboss.org/browse/CDI-457
2714	CDI-452	[~edburns] I just want to confirm. Your expectation here is that you can bridge the existing request context/session context from the servlet request in to the new async thread, allowing the same instances of the beans (and in turn, their data) to be accessed in both threads. If so, what synchronization would you expect, if any, to be provided by the CDI container? Or would you prefer to leave it to the developer to handle synchronization concerns?	0	0	https://issues.jboss.org/browse/CDI-452
2715	CDI-447	Thanks for the clarification. Do you agree that it's not clear in the spec that this only applies to embeddable EJB containers?	1	0	https://issues.jboss.org/browse/CDI-447
2716	CDI-427	To be consistent with the decision that's been chosen in CDI-408, I think we should stick to what was decided in CDI-377, CDI-404 and 406 (Normal Scope type, Interceptor, Decorator and Stereotype). Anyway, here's my POV on class level annotation not in Bean defining annotation set. * Qualifiers (built-in or user defined). As @Qualifier are defined in JSR 330, we shouldn't take them to avoid the same issue than in CDI-377 * @Alternative : could bring a lot of issue regarding if the bean this alternative replace has been detected as a bean. And since activating an alternative requiers using beans.xml, it's not very far to switch the bean-discovery-mode to all * @Specialize : same remark about the specialized bean being considered as a bean * @Typed : the only class level annotation that would be easily added to the set. But is it necessary ?	0	0	https://issues.jboss.org/browse/CDI-427
2717	CDI-425	This means we should enrich Javadoc for all initialization event. Do you propose a modification in the spec as well ?	1	1	https://issues.jboss.org/browse/CDI-425

#	Requirement	Questions	Answered	Successfully Answered	URL
2718	CDI-420	bq. The only difference in practice is to create the AnnotatedType + fire the PAT events... Well, if a bytecode scanning lib such as [jandex|https://github.com/wildfly/jandex] is used then a user may supply a pre-built index and you only need to load the index from the jar to be able to skip irrelevant classes. But you're right that the difference probably will not be significant. bq. But the huge benefit is that most CDI Extensions would work... Maybe we should build a list of extensions not working with {{annotated}} and add a FAQ entry on [http://cdi-spec.org/]? bq. Regarding using @WithAnnotations as input for the scanner. Imo that would be a chicken-egg problem. Well, even now an extension is allowed to add a bean defining annotation, see for example {{BeforeBeanDiscovery.addStereotype()}}. But you're right that it's problematic (Weld does not support this in all environments atm). bq. How about bean-discovery-mode = 'scanning'... Sounds better. Let's continue with ideas. What about {{activated}} in the sense that a type is activated if it has a bean defining annotation after the type discovery finished?	1	1	https://issues.jboss.org/browse/CDI-420
2719	CDI-420	you were faster than me [~mkouba] ;). So behavior of type discovery in annotated mode is well defined. The question is: is it better to introduce a new mode (and confusion) or change the behavior of {{annotated}} mode (what would be the real side effect).	1	1	https://issues.jboss.org/browse/CDI-420
2720	CDI-420	I am a bit confused. I thought this discussion is to include some classes that are not beans in the implicit bean archives. In Martin's suggestion, it is to exclude classes. Do you mean "include" or I might have missed something?	1	1	https://issues.jboss.org/browse/CDI-420
2721	CDI-420	Martin, can you elaborate a bit more on how to achieve the requirement of processing all classes with the exclude filter?	1	1	https://issues.jboss.org/browse/CDI-420
2722	CDI-420	I fear it's confusing for the user. Please remember that the whole ticket is mostly for fixing a behaviour which is imo broken since the very early days: to prevent automatically pick up all classes as @Dependent scoped beans, even if they have no single CDI annotation at all. Those are in 99.9999% cases just pure waste in our memory. Otoh we cannot enable this for all classes as there might be some jars which depend on it. Thus we need to activate this per classpath entry. There was a first attempt to fix this behaviour with the bean-discovery-mode="anotated". And it didn't work out because of the aforementioned problems. I thought about Martins proposal for two days now. It is certainly interesting but I don't yet see any benefit in having a second configurable exclude mechanism which exceeds this basic rules. Are there any other use cases you have in mind? Because if the above reason is the only one then I find it overreaching.	1	1	https://issues.jboss.org/browse/CDI-420
2723	CDI-420	I had a further thought on this: We could introduce 'include'. {code:java} <beans xmlns="http://xmlns.jcp.org/xml/ns/javaee" bean-discovery-mode="annotated"> <!-- These filters are applied before type discovery --> <scan> <include name="my.packages.no.beans.*" /> </scan> </beans> {code} This will cause the classes to be scanned and opted in for type discovery. It is similar to Martin's approach. Is this easier to be understood? just my 2cents.	1	1	https://issues.jboss.org/browse/CDI-420
2724	CDI-420	Why not {{<turn-off-beans-if-they-dont-have-bean-defining-annotations />}} (I'm kidding)? It needs a good term. I like the idea, but the naming needs to be better for it to be usable I think.	0	0	https://issues.jboss.org/browse/CDI-420
2725	CDI-409	Antonio, how would the underlying runtime environment know which DataSource to load? And why would only one alternative be enough? Certainly, developers will want to have more options (integration, qa, pre-production).	0	0	https://issues.jboss.org/browse/CDI-409
2726	CDI-408	Thanks [~jason.greene] and [~pmuir] . That means that we should enhanced the list of _bean defining annotations_ by adding {{@Observes}} and {{@Produces}} to them and introduce the fact that _bean defining annotations_ can be *inside* a class. Is it ok for everyone?	1	1	https://issues.jboss.org/browse/CDI-408
2727	CDI-408	extending the bean defining annotations with annotations from the class itself is no problem. But having to go into the class and define producer methods or fields, etc as bean defining would be much more work. You should also think about what happens with a producer method which is contained in a super class. Does that mean we should pick up the child class as dependent scoped? Imo we should have a few very clear rules. Not too many rules and not too complex ones. If a user expects a class to get picked up is only a matter the learning curve ;)	1	1	https://issues.jboss.org/browse/CDI-408
2728	CDI-408	Honestly I don't understand [~struberg]. Today in CDI 1.0 and in CDI 1.1 with {{bean-discovery-mode=all}} mode we do these deep class analysis. So the code is already there in implementations. Ok, it probably has to be adapted but I don't see where is the "much more work" in term of dev. Solution at the impl level are multiple : * Start as if we where in {{bean-discovery-mode=all}} and use the same mechanism that reject non eligible class (after finding no default constructor for instance) or that {{@Vetoed}} it, if it doesn't contain bean defining annotation * Use annotation processing to decide at compile time. My understanding of the {{annotated}} bean discovery mode is not about optimization, it's about CDI automatic activation. If it's optimized it's better but the first goal is to provide CDI implicitly. For me it's not important if the {{annotated}} bean discovery mode take the same amount of time at boot time than {{all}} bean-discovery-mode and that's the only risk we take here in term of "much more work". IMO The real question is "does this feature is interesting for the end user ?" and then "how much does it costs in terms of dev and perf ?". In this particular case I won't say the feature is awesome but the lack of it is very not user friendly. We have zillion forum post and stack exchange "bug" out there about CDI 1.0 users forgetting {{beans.xml}} in their app. Now that we provide the mean in CDI 1.1+ to get rid of this file I really think we should do it as complete as possible. Regarding the producer method in superclass, I'm not sure to totally understand your use case. Can you give an example ?	1	1	https://issues.jboss.org/browse/CDI-408
2729	CDI-407	{quote} @Named public class A @Named @Alternative public class B extends A What should acutally happen if B is activated via beans.xml? {quote} I think this is already covered in [4.1. Inheritance of type-level metadata|http://docs.jboss.org/cdi/spec/1.2/cdi-spec.html#type_level_inheritance], [5.3.1. Ambiguous EL names|http://docs.jboss.org/cdi/spec/1.2/cdi-spec.html#ambig_names] and other related sections. Note that bean names are not inherited unless specialization is used! Regarding your example: * bean A has a bean name "a" * bean B has a bean name "b" * so the names do not collide and even if you declare B as a selected alternative for a bean archive, EL expressions like #{a.foo} will not be affected, "a" will be resolved to bean A * bean name "b" will only be resolved to bean B if it's selected If we modify the example to: {code:java} @Named("a") public class A @Named("a") @Alternative public class B extends A {code} * both beans have the same bean name * if you declare B as a selected alternative, the name "a" will be resolved to bean B	1	1	https://issues.jboss.org/browse/CDI-407
2730	CDI-407	Yep, thats what i expected. In this case, is there any reason to still install/configure A as active bean? As there is no qualifier or stereotype available, it would be possible to automatically install B as @Named("a") without explicity giving it a name. Or do i miss something here? @Named("a") is a workaround but feels "ugly" and you always have to synchronize the names.	1	1	https://issues.jboss.org/browse/CDI-407
2731	CDI-407	[~tandraschko] {quote} In this case, is there any reason to still install/configure A as active bean? {quote} I'm not sure. Probably not in this very special scenario. But this is how it is defined right now and changing the behaviour wouldn't be backwards compatible. {quote} @Named("a") is a workaround but feels "ugly" and you always have to synchronize the names. {quote} You can still use String constants to avoid error-prone sync :-)	1	1	https://issues.jboss.org/browse/CDI-407

#	Requirement	Questions	Answered	Successfully Answered	URL
2732	CDI-407	[~struberg] Mark, I'm not following the example with interfaces. Could you try to describe it in another way? :-)	1	1	https://issues.jboss.org/browse/CDI-407
2733	CDI-392	Maybe a magic formula "non-portable behavior results" would be more consistent with the spec. than "in order to prevent random behaviour" as it is used throughout the whole text. WDYT? Otherwise good for me as well.	1	1	https://issues.jboss.org/browse/CDI-392
2734	CDI-392	If you know upfront that the container throws an exception, how is that random?	1	1	https://issues.jboss.org/browse/CDI-392
2735	CDI-392	Thanks for the input [~jharting]. What do you think of this wording guys ? {quote} Most operations of BeanManager may be called at any time during the execution of the application. However calling some BeanManager operations before the AfterDeploymentValidation event is fired will cause the container to throw an IllegalStateException (The container is permitted to define a non-portable mode in which these operation may be called from an observer of the AfterBeanDiscovery event) {quote}	1	1	https://issues.jboss.org/browse/CDI-392
2736	CDI-392	Even if it looks like a common statement after negotiations between 2 political party ;-) What do you think [~struberg] ?	1	1	https://issues.jboss.org/browse/CDI-392
2737	CDI-392	{quote} Well, there is nothing in the API which says that this is optional for the BeanManager. {quote} If I look at the javadoc of {{getBeans(Type beanType, Annotation... qualifiers)}} at the end I read {quote} *Throws :* *IllegalStateException* - if called during application initialization, before the AfterDeploymentValidation event is fired The container is permitted to define a non-portable mode in which getBeans(Type, Annotation...) may be called from an observer of the AfterBeanDiscovery event. {quote} My understanding was that it's not allowed but that an implementation is allowed to define a non-portable behavior by allowing the call after AfterBeanDiscovery event is fired. From what you say [~struberg], I understand that I misunderstood that point. So why don't we forbid definitively call to these operations before AfterDeploymentValidation event to avoid junior developers to create futur bugs ?	1	1	https://issues.jboss.org/browse/CDI-392
2738	CDI-392	Do we have an updated version of the change?	1	1	https://issues.jboss.org/browse/CDI-392
2739	CDI-392	I can't honestly remember why it is there. Have you checked the commit log?	1	1	https://issues.jboss.org/browse/CDI-392
2740	CDI-392	Since the behavior of these extension is not predictable on Weld 2.x, wouldn't be a better idea to remove the mention and the support (don't even know if Weld 2.x implement this non portable behavior) to make things clear ? Correct me if I'm wrong but Seam Solder is deprecated ?	1	1	https://issues.jboss.org/browse/CDI-392
2741	CDI-392	As we agree that those sentences make no sense: should we remove those?	1	1	https://issues.jboss.org/browse/CDI-392
2742	CDI-382	Thx. Yes the new workflow is only applied on "bug" and "feature request". I'm going to investigate why.	0	0	https://issues.jboss.org/browse/CDI-382
2743	CDI-370	Shouldn't the web socket initiate an HTTP request to emulate this? or is this a sign that the CDI spec is too closely tied to the Servlet spec WRT scope maintenance.	0	0	https://issues.jboss.org/browse/CDI-370
2744	CDI-370	What is the status on this? Will there be an errata or something?	0	0	https://issues.jboss.org/browse/CDI-370
2745	CDI-370	Is there any alternate solution till this get's fixed?	1	0	https://issues.jboss.org/browse/CDI-370
2746	CDI-370	Just wondering, could CDI-30 be used to solve cases like this?	0	0	https://issues.jboss.org/browse/CDI-370
2747	CDI-352	I've tried to make it clearer that this is wrong. Is there a CDI TCK issue for this?	0	0	https://issues.jboss.org/browse/CDI-352
2748	CDI-348	Please clarify "inherits a scope from a stereotype". Does this mean that if there's a scope defined as part of the stereotype that it is not applied to the bean that is annotated with the stereotype?	0	0	https://issues.jboss.org/browse/CDI-348
2749	CDI-323	Isn't it treated by 6.5.3? {quote} The container must ensure that every injection point of type InjectionPoint and qualifier @Default of any dependent object instantiated during this process receives: * an instance of InjectionPoint representing the injection point into which the dependent object will be injected, or * a null value if it is not being injected into any injection point. {quote}	1	0	https://issues.jboss.org/browse/CDI-323
2750	CDI-320	Your quote of the spec seems a bit old. The current spec says {quote} The container must fire an event, before it processes a type, for: • every Java class, interface or enum discovered • each Java class, interface or enum that must be delivered to a ProcessAnnotatedType observer, where the event parameter is annotated with @WithAnnotations. {quote} I think the use case of adding an annotation thru an {{AnnotatedType}} in {{BeforeBeanDiscovery.addAnnotatedType()}} or {{AferTypeDiscovery.addAnnotatedType()}} is very unlikely, but we could change the previous wording by {quote} The container must fire an event, before it processes a type, for: • every Java class, interface or enum discovered • *every {{AnnotatedType}} representing a Java class, interface or enum added by {{BeforeBeanDiscovery.addAnnotatedType()}} or {{AferTypeDiscovery.addAnnotatedType()}}* • each Java class, interface or enum that must be delivered to a ProcessAnnotatedType observer, where the event parameter is annotated with @WithAnnotations. {quote} Wdyt ?	1	1	https://issues.jboss.org/browse/CDI-320
2751	CDI-319	Is splitting a package between the WEB-INF/classes directory of a war and a directory in the JVM classpath portable, then?	0	0	https://issues.jboss.org/browse/CDI-319
2752	CDI-304	Per JLS 5.1.9 "Use of an unchecked conversion generates a mandatory compile-time warning" (this is a unchecked conversion), therefore, I think CDI should do just the same. It should work, but generate a warning. WDYT?	1	1	https://issues.jboss.org/browse/CDI-304
2753	CDI-304	Mark, I don't follow your comment. What do you mean with "ambiguous situations"? If you have multiple producer methods "which map to different generic types of the same base type" each PM has a different bean type (e.g. Foo<String> and Foo<Integer>), or not? I also found this sentence in 3.3 Producer methods: "If a producer method return type contains a wildcard type parameter the container automatically detects the problem and treats it as a definition error." This seems a little bit inconsistent to me. We would allow raw return type and disallow type with wildcard... but there is probably some reason I don't see.	1	1	https://issues.jboss.org/browse/CDI-304
2754	CDI-304	Just for the record, there are other combinations of required/bean type where CDI assignability rules are different from Java assignability rules. However, I do not have a strong argument against the proposed solution. Pete, you originally mentioned that a warning should be generated in such a case. Will this be a spec requirement?	1	1	https://issues.jboss.org/browse/CDI-304
2755	CDI-304	5.2.4. "Assignability of raw and parameterized types" now says {quote}A raw bean type is considered assignable to a parameterized required type if the raw types are identical.{quote} 8.3.1. "Assignability of raw and parameterized types for delegate injection points" says: {quote}A raw bean type is considered assignable to a parameterized delegate type if the raw types are identical *and all type parameters of the delegate type are either unbounded type variables or java.lang.Object.*{quote} Wouldn't it make sense to make these two assertions coherent? Note that the fact that the spec mentions this specifically in 8.3.1 indicates that the current wording of 5.2.4 (changed as a result of this issue) had not been originally intended in CDI 1.0!	1	1	https://issues.jboss.org/browse/CDI-304
2756	CDI-294	What do people think should happen instead if a Servlet class is injected? Options I can see: (a) UnsatisfiedResolutionException (b) return the instance the Servlet container is currently using I prefer (a). Also note that the Java EE EG is discussing this problem.	0	0	https://issues.jboss.org/browse/CDI-294
2757	CDI-289	Where is the FAQ? I would like to know how to proceed in such cases	1	0	https://issues.jboss.org/browse/CDI-289

#	Requirement	Questions	Answered	Successfully Answered	URL
2758	CDI-282	In fact I was thinking of something like this: {code} @Vetoed @Target({ TYPE, METHOD, PARAMETER, FIELD }) @Retention(RUNTIME) @Documented @Qualifier public @interface Predator { } /** * Predator annotation is vetoed. Does it mean the Tiger bean has Any and Default qualifiers only? */ @Predator public class Tiger { } {code}	1	1	https://issues.jboss.org/browse/CDI-282
2759	CDI-282	Martin, should we support @Vetoed on annotations at all? We don't fire ProcessAnnotatedType for them neither, isn't? It just sounds crazy to use a veto annotation on a CDI annotation. You could simply comment out the @Qualifier and get the same effect ;)	0	0	https://issues.jboss.org/browse/CDI-282
2760	CDI-282	package annotations are broken by design anyway. Didn't we drop that already?	0	0	https://issues.jboss.org/browse/CDI-282
2761	CDI-280	I won't go back in terminology war. I'm the newbie here and understand my user / usage oriented pov is not widely shared ;). Anyway, I think we're too late to produce a quality work on this. I propose we postpone this lexical dispute for 2.0 where most of the spec will more deeply reviewed. So I think we should remove this ticket from 1.2. WDYT?	0	0	https://issues.jboss.org/browse/CDI-280
2762	CDI-280	bq. The proxy for a Contextual Instance should be referred to as 'Contextual Reference' Or why not just call the proxy a proxy? If the origin of the proxy is not clear, call it a CDI proxy. In either way, I find that "contextual reference" is confusing.	1	0	https://issues.jboss.org/browse/CDI-280
2763	CDI-280	What is a "tenant proxy"? You mean a proxy written by the application developer himself? If so, then I'd not worry about it at all. Half of the rest is related to EJB and JPA. If CDI enforce different proxy types, then it should be clarified (but it doesn't right?). "CDI proxy" would be a valid term that encapsulate all types (but again, there are no different CDI proxy types?), or just "proxy" for short. The bottom line is that if we deal with a proxy, why not just call it that? Everything in the spec is "contextual" already. The word feels a bit bloated. Actually, I think that the word "proxy" is one of those things that is quite clear without semantic noise. Now you want to replace all references made to a proxy with "contextual reference"? I just don't get it why.	0	0	https://issues.jboss.org/browse/CDI-280
2764	CDI-277	On event1 I would expect to be called: observer1 of BeanA on an instance of BeanB observer1 of BeanA on an instance of BeanC observer1 of BeanB on an instance of BeanB observer1 of BeanC on an instance of BeanC observer1 of BeanA is NOT called on an instance of BeanA because it is specialized On event2 I would expect to be called: observer2 of BeanB on an instance of BeanB observer2 of BeanC on an instance of BeanC observer2 of BeanA is NOT called because BeanB and BeanC override it and no instance of BeanA exists, because BeanC specializes it. On event3 I would expect to be called: observer3 of BeanB observer3 of BeanC observer3 of BeanA is NOT called because BeanA is specialized On event4 I would expect to be called: observer4 of BeanB observer4 of BeanC observer4 of BeanA is NOT called because BeanA is specialized wdyt?	1	1	https://issues.jboss.org/browse/CDI-277
2765	CDI-277	~jharting ok, let me rephrase: A non-activated @Alternative is _never_ target of a CDI event. The @Alternative must be enabled to get a CDI event. What I meant is that just having another @Alternative class B implements X for class A implements X will NOT disable any observer method in class A!	0	0	https://issues.jboss.org/browse/CDI-277
2766	CDI-277	What's X?	0	0	https://issues.jboss.org/browse/CDI-277
2767	CDI-276	That brings another point [~petemuir], In the spec we define a very binary behavior : either the container accept the defined state or the container "treats this as definition error". Shouldn't we also have "the container should raise a definition warning" for case like this one?	1	1	https://issues.jboss.org/browse/CDI-276
2768	CDI-276	Java EE specifications rarely contain as many error conditions as CDI, nor do they specify what happens in these situations as much as CDI does. So we could discuss with the EG introducing this warning requirement, but it's diverging a lot from other specs. Also, it would be very hard to test that it actually happens AFAICT.	0	0	https://issues.jboss.org/browse/CDI-276
2769	CDI-275	oki, in other words if I have a {code} @RequestScoped @Stateful public class MyBean { ... } {code} Then what happens with the lifecycle? Will it get destroyed at the end of the request or not?	0	0	https://issues.jboss.org/browse/CDI-275
2770	CDI-275	yes this part explains it but something is still not very clear IMO: should a remove method be mandatory or does it work without any @Remove method? there is nothing blocking to make it optional technically since it was needed only to be able to run the remove hook but since we manage the lifecycle now we can do it internally so i think a sentence should be added to clarify it wdyt?	0	0	https://issues.jboss.org/browse/CDI-275
2771	CDI-271	Arne: What do you think about my proposal for CDI-20? Maybe this can be used here too.	1	1	https://issues.jboss.org/browse/CDI-271
2772	CDI-271	look like good additions. Could you make it to the next weekly EG meeting on monday 18:00 CEST ?	0	0	https://issues.jboss.org/browse/CDI-271
2773	CDI-271	I'm not sure that it's right to reuse InjectionPoint, as it has other stuff on it that perhaps doesn't make sense (e.g. what do getMember() return, getAnnotated() , isDelegate() and isTransient() return)? And also, the case of doing the fire via BeanManager would not fit well. How about introducing {code} public interface EventCaller { public Set<Annotation> getEventQualifiers(); public Type getEventType(); } Which could be injected into an event observer?	1	1	https://issues.jboss.org/browse/CDI-271
2774	CDI-271	Just saw, that section 5.5.7 already is changed so that Instance.get() is an injection point. For this issue it would suffice to add Event.fire() here and to clarify that for Instance.select(...).get() and Event.select(...).fire(...) the InjectionPoint.getType() returns the selected type and InjectionPoint.getQualifiers() contains the selected qualifiers. We don't have to touch the Annotated then. WDYT? Shall I write a proposal?	1	1	https://issues.jboss.org/browse/CDI-271
2775	CDI-271	Arne, can you write a patch on the spec? I can review and apply then...	1	1	https://issues.jboss.org/browse/CDI-271
2776	CDI-271	Where's the link to the pull request?	1	1	https://issues.jboss.org/browse/CDI-271
2777	CDI-271	What happens when an observer defines an InjectionPoint parameter but the event was fired using BeanManager instead of Event? The spec does not currently seem to define this.	0	0	https://issues.jboss.org/browse/CDI-271
2778	CDI-271	I like CallerAttributes, but would this be just for Events? If so wouldn't EventAttributes make sense? @Arne - why would getMember and getAnnotated return null? Wouldn't getMember be the BeanManager and getAnnotated also be the BeanManager?	1	1	https://issues.jboss.org/browse/CDI-271
2779	CDI-271	At this point, I don't want to add new features like supporting BeanManager here. So Arne, can you create a new issue for that?	0	0	https://issues.jboss.org/browse/CDI-271
2780	CDI-265	I'm not sure I'm following. Provided we have an EAR with two web modules. What is the payload of an @Initialized(ApplicationScoped.class) event? Object?	0	0	https://issues.jboss.org/browse/CDI-265
2781	CDI-259	Pete, do you know if this is still needed? I think we have the same by simply observing the @Initialized(ConversationScoped.class) and injecting a Context which can then be asked for .isTransient(); So you can tell apart whether someone really did hit the Conversation#begin() or if it will end at the Request boundary	1	0	https://issues.jboss.org/browse/CDI-259

#	Requirement	Questions	Answered	Successfully Answered	URL
2782	CDI-255	h3. Proposal The Java EE container should create a producer for every resource available to the application, following these rules: * If there is more than one resource of that type available to the container, a producer is not created, and should be aliased using a resource producer * If the application already defines an enabled bean of that type, the container should not create the producer. OPEN QUESTION: Should the container create a qualified producer in this case? * The qualifiers should be {{@Default}} and {{@Any}} * The bean types should be the transitive closure of the type hierarchy of the type of the resource * The scope should be {{Dependent.class}} h3. Implementation Notes * The Java EE container can register each producer using the {{AfterBeanDiscovery.addBean()}} event * The Java EE container can create the bean by: *# Implementing the {{Producer}} interface, simply returning the relevant resource (e.g. using a JNDI lookup) in {{produce()}} and doing nothing in {{dispose()}}. The set of injection points can be empty. *# Calling BeanManager.createBean(), passing in an implementation of {{BeanAttributes}} that implements the requirements described above (with null for name, an empty set for stereotypes, and false for alternative and nullable), the producer created in (1) and the type of the resource for the bean class *# Calling {{AfterBeanDiscovery.addBean()}} passing in the bean created in (2)	0	0	https://issues.jboss.org/browse/CDI-255
2783	CDI-255	Yes you're right, sorry for the mistake. Anyway don't you think this should be in EE 8 spec not CDI ?	1	0	https://issues.jboss.org/browse/CDI-255
2784	CDI-248	do you fire customers for any other reason? I'm thinking you can use [EventMetadata|http://docs.jboss.org/cdi/api/1.1/javax/enterprise/inject/spi/EventMetadata.html] to find out if an appropriate qualifier is on there.	0	0	https://issues.jboss.org/browse/CDI-248
2785	CDI-245	I would prefer to explicitly allow extensions to use BeanManager.getExtension(). WDYT?	0	0	https://issues.jboss.org/browse/CDI-245
2786	CDI-245	> Today, the only thing you can inject in a portable fashion is the BeanManager. In CDI-1.0 it's just undefined what you can inject at bootup. We can only exclude what explicitely is forbidden or defined to get started afterwards. But there is currently no wording that specifies that e.g. @Stateless beans cannot be @Injected. For Extensions I looked at 11.5: {quote} The container instantiates a single instance of each extension at the beginning of the application initialization process and maintains a reference to it until the application shuts down. The container delivers event notifications to this instance by calling its observer methods. For each service provider, the container must provide a bean of scope @ApplicationScoped and qualifier @Default, sup- porting injection of a reference to the service provider instance. The bean types of this bean include the class of the service provider and all superclasses and interfaces. {quote} Thus a.) all Extensions get loaded before the first container lifecycle event gets fired b.) all Extensions are beans which are available for injection c.) all Extensions are @ApplicationScoped. I guess the problem is that the ApplicationContext only gets started after AfterDeploymentValidation, right? In DeltaSpike we currently use the Extension injection for runtime communication between an Interceptor and it's Extension and not at boot time.	1	1	https://issues.jboss.org/browse/CDI-245
2787	CDI-245	I pondered for a while over that same paragraph, when I came to my conclusion that extensions aren't defined as beans. Typically, when the spec wants something to be a bean, it says "is a ... bean" e.g. 3.1.1 (Managed beans) {quote} A top-level Java class is a managed bean if it is defined to be a managed bean {quote} or 3.2 (Session beans) {quote} A session bean is a bean that is implemented by a session bean with an EJB 3.x client view. {quote} On the other hand, when the spec wants the container to provide a bean which provides access to an externally managed resource it says "must provide a bean" e.g. 3.7 (Additional built-in beans) {quote} A Java EE or embeddable EJB container must provide the following built-in beans, all of which have qualifier @Default: a bean with bean type javax.transaction.UserTransaction, allowing injection of a reference to the JTA UserTransaction, (etc) {quote} To me, it's fairly clear that the spec isn't saying "UserTransaction is bean", but is saying "There must be a bean that represents the UserTransactions". Here, the spec clearly uses the second form, so my reading is that "the container must provide a bean that represents the extension". I think this reading is backed by the spec saying "An extension is a service provider..." and "For each service provider,...", and then going on to refer to the "extension", and not the "extension bean". Firstly, because it clearly says "an extension *is a* service provider" and not "an extension is a bean implemented by a service provider" (or better written words to that effect). Secondly, because the spec is quite clear about when something is a bean or not elsewhere, and here it makes no effort to say "an extension is a service provider". WDYT? Do you see why I came to this conclusion? Regarding field injection into a service provider, I don't believe that has ever been allowed specifically by the specification (it always talked about injection into the observer method only), and I think that this makes sense. The "If other beans are injected into a service provider" was something I added to EDR1, that I think was a mistake now. I also choose not to address defining anything other than the container firing lifecycle events as non-portable, as I don't believe I have a good answer to that question right now. I would actually much prefer we restrict it by defining something on observer methods such as "an extension observer method need only to respond to container lifecycle events fired by the container and need not respond to events fired by an application" or something like that. WDYT?	0	0	https://issues.jboss.org/browse/CDI-245
2788	CDI-233	Sorry, I meant bundled library, not shared library. However it looks we're on roughly on the same level here. Mark, do you have another proposal?	0	0	https://issues.jboss.org/browse/CDI-233
2789	CDI-232	I doubt there is an implementation that would be following this. What can be implied from the spec is that if you obtain the Bean<?> object for the built-in Instance bean and call Bean.getTypes() you should get a huge set of Instance<X> and Provider<X> where X is every possible Java type that matches "2.2.1. Legal bean types". I think that this would not even be a finite set most of the time so typesafe resolution cannot clearly work this way for the Instance bean.	1	1	https://issues.jboss.org/browse/CDI-232
2790	CDI-232	yes [~jharting], that's why I was starting to propose a new way to define the behaviour. Saying that the container doesn't provide all the Instance for all bean type but only check that the Instance you injected is related to a valid bean type. Do you mean that the current behaviour allow me to put a non valid bean type (inject Instance<Foo> where no beans exist with Foo type) ?	1	1	https://issues.jboss.org/browse/CDI-232
2791	CDI-232	So now what doe getBeanTypes return?	1	1	https://issues.jboss.org/browse/CDI-232
2792	CDI-229	Wdyt about naming it @Propagate? Because that's what it does basically...	0	0	https://issues.jboss.org/browse/CDI-229
2793	CDI-229	Definitely useful. However, note that this does not match the existing SPIs as nicely as it could because CDI currently does not expect stereotypes to have members and thus works with annotation classes instead of annotation instances. For example: BeanManager: Set<Annotation> getStereotypeDefinition(Class<? extends Annotation> stereotype) BeanAttributes: Set<Class<? extends Annotation>> getStereotypes()	0	0	https://issues.jboss.org/browse/CDI-229
2794	CDI-228	folks, this is still not clarified but an important point for CDI-1.1 Are method parameters for producer methods now assumed 'Transient' or 'Keep' ? The CDI-1.1 draft completely ignores this new 'Transient' stuff by now, thus all our serialisation rules are a bit broken atm as well...	1	1	https://issues.jboss.org/browse/CDI-228

#	Requirement	Questions	Answered	Successfully Answered	URL
2795	CDI-228	Can someone please change the title of this ticket.	1	1	https://issues.jboss.org/browse/CDI-228
2796	CDI-228	Would it make sense to expose the fact that @DestroyedAfterBlock is present (or whatever the annotation name is) through the InjectionPoint interface?	0	0	https://issues.jboss.org/browse/CDI-228
2797	CDI-228	Hmm, still not sure what to call this. @Independent is just too weird for me. It sounds like the opposite of @Dependent, so therefore it's a new pseudo scope, but of course it's not. I understand that it doesn't feel like your FacesContext is destroyed after the block, but your reference to it would be (it won't be disposed, it will be destroyed). A better version might be @ReleaseAfterBlock? Which is still indicating that you are changing the scope of the reference from the class to the block. WDYT?	0	0	https://issues.jboss.org/browse/CDI-228
2798	CDI-228	Hmm, I agree with your objections about @Independent, but annotations like @DestroyedAfterBlock or @ReleaseAfterBlock really don't fit into the spirit of CDI annotations. CDI annotations describe attributes (i.e. @Red as qualifier, @Transactional as an interceptor binding). The annotations proposed here describe behavior, which is something completely new for the CDI spec What about @IndependentReference or @TransientReference ?	0	0	https://issues.jboss.org/browse/CDI-228
2799	CDI-226	Thanks for picking this up from the discussion on the list! It's really better to have this here in the tracker. Pete, is there a way to link to a jboss mail archive? Or should we just copy the gist of the discussion over here?	0	0	https://issues.jboss.org/browse/CDI-226
2800	CDI-224	[~antoinesabot-durand] That does make life more difficult, but would it really hurt if the Decorator would also be applied in case of "helper class extension". Since those are highly unlikely to be injected? [~thenovot] Yes, however I don't see how the user error risk outways the benefit, it would also be true if a different annotation would be introduced or adding a new element to @Decorator. Alternatively, the type could be inferred from the @Delegate injection in case of extension. Another option which pops to mind, not my favorite, is declaring the to be decorated type in the annotation @Decorator(MyRealBean.class). This could then override the interface type inference and be mandatory for extension.	0	0	https://issues.jboss.org/browse/CDI-224
2801	CDI-223	@Stateful + @EJB User may turn the use-case into something rigid/theoretical, don't you think? Back to to the example, @Remote can't marry with CDI's jurisdiction boundaries. It's obvious and at the same time I feel there's going to be a lot of questions on user forums. It kind of reminds me of the lazy loading issue on detached JPA entities. Well, so, logically (though it will require careful documenting), the only option is to pass a reference to the content you want to use within your EJB. Either you ask developpers to create an intermediate cdi-free "copy" of User(sounds oldschool to me), or take the road Mark outlined. Unless there are so few Remote-EJB/CDI users that this issue is pointless (for this CDI revision at least), I think we have no choice.	0	0	https://issues.jboss.org/browse/CDI-223
2802	CDI-220	Well Martin, I fear the CDI spec is not 100% clear as it says that session beans also must be treated as managed beans... It would be more clear if it would state 'Managed CDI bean' or anything like that. Look at "3.1.1. Which Java classes are managed beans?" > A top-level Java class is a managed bean if it is defined to be a managed bean by any other Java EE specification, Which a session bean certainly is, isn't? hmm but it also states: > It is not annotated with an EJB component-defining annotation or declared as an EJB bean class in ejb-jar.xml. The confusion gets deeper because EJBs must also be held as Bean<T> in the BeanManager as required by the spec. JSR-316 (EE umbrella) btw seems to catched that as well and explicitly ranted about it ;) > "CDI-style managed beans" > We use this term to refer to classes that become managed beans per the rules > in the CDI specification (JSR-299), thus excluding managed beans declared using > the @ManagedBean annotation as well as EJB session beans, both of which would > be managed beans even in the absence of CDI. *scratches his head*	0	0	https://issues.jboss.org/browse/CDI-220
2803	CDI-219	I actually don't have strong arguments for @PreDestroy, the main one is symmetry, which is not completely fool since it's common to use the same resources in @PreDestroy as in @PostConstruct. Joe, I'm afraid I don't understand 2), (and I coudn't find Martin's comment): my suggestion is that in @Singleton @PostConstruct (and eventually @PreDestroy) the current RequestContext, if any, is "suspended" and "resumed" upon exit of the callback. Does this contradict 6.3? This would be coherent with transactional behaviour; otherwise we end up with a new transaction opened on the @Singleton @PostConstruct invocation, suspending the current one, all within the same RequestContext which could lead to weird side effects I think.	0	0	https://issues.jboss.org/browse/CDI-219
2804	CDI-219	Doesn't the TransactionContext (see CDI-121) already violate this principle? I suppose that, if a transaction is suspended and a new one is begun, the corresponding context is suspended as well and a new one is created/activated. I'm just guessing here about TransactionContext so maybe I'm wrong, but it seems to me that the threading semantics of CDI 1.0 is a bit limiting.	0	0	https://issues.jboss.org/browse/CDI-219
2805	CDI-219	Saying (a) "Contexts must always be joined if they are alive" and (b) "There aren't two alive at the same point in time on the same thread" are different things, right? My proposal for @Singleton @PostConstruct (and @PreDestroy) is that the RequestContext should behave exactly like the TransactionContext, that is, to suspend the current context and activate a new one; this does violate (a) but not (b), just like TransactionContext. Reading 6.3, it seems to me that (b) is clearly implied, but I cannot find (a), both in CDI 1.0 and in the current draft of 1.1. Am I missing something?	0	0	https://issues.jboss.org/browse/CDI-219
2806	CDI-213	Tom Jenkinson says {quote} Would it make sense to maintain the same restrictions as JMS? In JMS you can initiate a new transaction from onMessage, analogous to a method decorated with @Observes. If the method returns without completing the transaction then an error is logged and the transaction rolled back. Flowing a transaction from an event producer to an event consumer isn't a great idea (it doesn't work in JMS either). But allowing a consumer to control their own transaction does seem to make sense to me. Admittedly this is a gut reaction, I read through the Jira and the doc you linked to though, and used my JEE experience to draw analogies, do let me know if I got the wrong end of the stick please! {quote} So it sounds like it's sensible to specify as non-portable, as people might want to experiment with this. I'll raise a new issue for TBD to consider specifying how TX control works in observer methods.	1	1	https://issues.jboss.org/browse/CDI-213

#	Requirement	Questions	Answered	Successfully Answered	URL
2807	CDI-213	As far I can tell we can remove the current non portable statement and refer to the JTA spec instead. IN_PROGRESS observers are covered by http://java.net/jira/browse/JTA_SPEC-5, BEFORE_COMPLETION/AFTER_COMPLETION observers are covered by http://java.net/jira/browse/JTA_SPEC-6. How do you refer to a possibly clarified behavior in a future spec? ;) It looks like JTA_SPEC-5 is on it's way, JTA_SPEC-6 (that I created) shouldn't be to difficult after reading Marina's and Paul's comments on that topic (at least it could be declared non portable). ---------- Paul Parkison wrote: I would agree with the general approach/alignment that "CDI should rely on what's already specified for the platform especially for container and bean managed transactions and JTA synchronization callbacks" and that any ambiguous, etc. factors should be addressed in the JTA group/spec. [...] The CMT/EJB/CDI @Transactional/TransactionScopeaspect is being covered here: http://java.net/jira/browse/JTA_SPEC-5 I think that may still leave the finer clarification of what is legal in beforeCompletion. The introduction of the TransactionSynchronizationRegistry resulted in elaborating on the specification of Synchronization behavior (and of course beforeCompletion) as so the chapter 5 API ref is a better place to look, but indeed things like the ability to suspend/resume in a beforeCompletion method are not specified (the spec would actually seem to suggest that would not be legal/portable but at the same time I know it is possible, and likely used, in most if not all appservers).	1	0	https://issues.jboss.org/browse/CDI-213
2808	CDI-213	Jens. Can you tell me exactly what changes I need to make to the CDI spec, if any? Thanks!	0	0	https://issues.jboss.org/browse/CDI-213
2809	CDI-212	I'm somewhat confused. @Decorator cannot be applied to a sterotype. Wouldn't it be better to have an explicit createDecorator method?	1	1	https://issues.jboss.org/browse/CDI-212
2810	CDI-207	Any ideas what the payload should be? Object?	1	1	https://issues.jboss.org/browse/CDI-207
2811	CDI-207	Interesting. Should we actually use the Conversation id as the payload always?	1	1	https://issues.jboss.org/browse/CDI-207
2812	CDI-203	What are the bean types and qualifiers of the MDB?	1	1	https://issues.jboss.org/browse/CDI-203
2813	CDI-197	Jozef, what section of the spec does this relate to?	1	1	https://issues.jboss.org/browse/CDI-197
2814	CDI-196	erm, what do you like to tell us? This issue doesn't contain any issue report nor change request... Please insert .) what is fishy currently .) how do you like it to get changed.	1	0	https://issues.jboss.org/browse/CDI-196
2815	CDI-195	Does the spec state explicitely that instances can hold DIRECT references to the metadata. The container could inject serializable proxies of the metadata and restore the underlying metadata on deserialization.	0	0	https://issues.jboss.org/browse/CDI-195
2816	CDI-194	Do you have a good use case for this? I tempted to defer it.	1	1	https://issues.jboss.org/browse/CDI-194
2817	CDI-194	I worked around the issue Stuart described by observing the ProcessBean event and building up my own registry of annotated types. What do you think about this approach?	1	1	https://issues.jboss.org/browse/CDI-194
2818	CDI-194	Sorry, I'm not sure about this. I'm just utilizing ProcessBean events for "classpath scanning" (well correctly said "bean archive scanning", right;). In fact, as you mentioned, if there's a concrete class subclassing the abstract type you can traverse up the type hierarchy and find out the abstract one. However, to get the AnnotatedType representation of the identified abstract type, I'm utilizing Solder's AnnotatedTypeBuilder which of course isn't aware of any XML configuration. So if one could retrieve a Bean for the abstract type and if there was a Bean.getAnnotatedElement() method, this issue could be fixed nicely. Does my explanation make sense to you?	1	1	https://issues.jboss.org/browse/CDI-194
2819	CDI-194	My use case is the following: my framework operates on @Component annotated classes (@Component is a custom annotation). I want to rely on AnnotatedType infrastructure for two reasons: * CDI/the container does classpath scanning, why should I scan it with some other library again (additional bootstrap time/additional JAR libraries required for scanning etc.)? * I want to gain XML configurability through Solder XML without re-implementing it {code} public <X> void processBean(final @Observes ProcessBean<X> event, final BeanManager beanManager) { final Annotated annotated = event.getAnnotated(); // detect @Component annotated types as managed types if (annotated instanceof AnnotatedType && annotated.isAnnotationPresent(Component.class)) { log.debug("Processing {}", annotated); final Type type = annotated.getBaseType(); // remember annotated type's super types (e.g., for the detection of abstract super classes, which are not // detected by CDI) typeClosureLoop: for (final Type superType : annotated.getTypeClosure()) { if (type.equals(superType)) continue typeClosureLoop; Class<?> superClass = null; if (superType instanceof Class) superClass = (Class<?>) superType; else if (superType instanceof ParameterizedType) superClass = (Class<?>) ((ParameterizedType) superType).getRawType(); // only classes and parameterized types, Object.class is ignored if (superClass != null && !superClass.equals(Object.class) && !discoveredComponentTypes.containsKey(superType)) { discoveredComponentTypes.put(superType, asAnnotated(superClass)); } } // remember annotated type itself discoveredComponentTypes.put(type, (AnnotatedType<?>) annotated); } } private <X> AnnotatedType<X> asAnnotated(Class<X> type) { return new AnnotatedTypeBuilder<X>() .readFromType(type) .create(); // how to do this with BeanManager? } {code} So, all @Component types discovered through @Observes ProcessBean<X> event are enriched with XML config annotations. However, as ProcessBean isn't fired for abstract types, and as there isn't another way to obtain an annotated type from BeanManager, I have to use AnnotatedTypeBuilder which doesn't honor XML configurations. When I was able to retrieve a properly Xml enriched annotated type from Bean or BeanManager, this issue was solved. PS: I use ProcessBean to ensure that Solder's XML config already wrapped the AnnotatedType with XML specified annotations. I cannot specify a dependency on Solder's extension being processed before mine, right?	1	1	https://issues.jboss.org/browse/CDI-194
2820	CDI-194	As Stuart points out, we need to return {{AnnotatedField}} for a producer field, {{AnnotatedMethod}} for a producer method, and {{AnnotatedType}} for managed and session beans. We also want to provide requirements for custom beans. {{Annotated}} is the only supertype of all these types, so should we add {{Annotated<?> getAnnotated()}} to {{BeanAttributes}}?	1	1	https://issues.jboss.org/browse/CDI-194
2821	CDI-184	Mathieu, would you be able to work up a proposal for CDI?	0	0	https://issues.jboss.org/browse/CDI-184
2822	CDI-180	>> However it doesn't cover various invocation scenarios. For example it's >> possible to do in-VM invocation (via remote interface) from the same >> deployment or from a different deployment. > >I don't see why the spec implies this is treated any differently. I think the real question here is that if you have 2 CDI applications, and they communicate via in-vm method calls, does the bean in the second application live for a single invocation, or does it last for the life of the calling request scope. I would assume the former.	0	0	https://issues.jboss.org/browse/CDI-180
2823	CDI-180	I think we need the exact definition of a remote method invocation lifespan from EJB to clear this up properly. Does anyone have this?	1	1	https://issues.jboss.org/browse/CDI-180

#	Requirement	Questions	Answered	Successfully Answered	URL
2824	CDI-180	I'm not sure if this falls under this category, but do we generally like to define what happens with remote and also with _async_ parts? E.g. if I have a bean which starts an async EJB task. At first it sounds very tempting to just promote the current RequestContext and SessionContext also over to this async task. But now think a bit about it. The user starts an async EJB method which needs 30 minutes (or even longer) to succeed. Immediately after kicking off the task (long before it finishes), the user logs out and the logOutHandler therefor destroys the Session. What will happen? Also please be aware that the methods to programmatically access a Session with a given Id got removed from the Servlet spec long time ago. Plus there are serialization issues. Back to your original problem (which is different, but the effect is pretty much the same): whenever we cross boundaries (threads, servers), this might work in some situations, but we will fail to support it in a generally valid manner.	1	1	https://issues.jboss.org/browse/CDI-180
2825	CDI-177	Would it not be cleaner to generally discuss how @Named qualifiers are defaulted, and then make the EL Name use that?	0	0	https://issues.jboss.org/browse/CDI-177
2826	CDI-169	I still don't understand where the bug is. Does spec forbid use of parametrized type as an event payload or not? Right now Weld detects problems on user-defined events only (not container lifecycle events). It's confusing but I'm not sure it is bug.	0	0	https://issues.jboss.org/browse/CDI-169
2827	CDI-164	Multiple reasons for reopening: h3. Bean Manager The spec says: Any decorators associated with BeanManager will not be applied. It is not clear whether this means that decorators are not applied to BeanManager objects injected into Extensions during bootstrap only or whether it means that injected BeanManager objects are never ever decorated. h3. The extension problem For each extension there is a bean that allows the extension to be injected into other beans. Should the injected extension support decoration? We can only support this at runtime (no decorators would be applied to extension instances at bootstrap). If we decide to support decoration of injected extension we should think about whether extension instances obtained via BeanManager.getExtension() should support decoration or not. h3. Consistency The spec says: "Decorators may be associated with any managed bean that is not itself an interceptor or decorator, with any EJB session bean or with any *built-in* bean." The spec currently defines a set of container-provided beans out of which a subset is referred to as "built-in" beans. For example, the injected InjectionPoint is not explicitly labeled as "built-in bean". The same applies for injected Bean, Interceptor, Decorator (5.5.8), Validator, ValidatorFactory, ConstraintValidator and extension beans. We should either say that decorators may be associated with any container-provided bean or we should explicitly label all the container-provided beans as "built-in". h3. Passivation Most of the built-in beans are passivation capable dependencies. We should required decorators of built-in beans to be passivation capable. Otherwise the container cannot store a decorator's state when it passivates e.g. an Event instance.	0	0	https://issues.jboss.org/browse/CDI-164
2828	CDI-162	Would the extended persistence context be bound to a scope (like the seam managed PC) or to a component (like a SFSB)? Personally I think a Extended PC bound to a scope would be much more useful than an extended PC bound to a component.	0	0	https://issues.jboss.org/browse/CDI-162
2829	CDI-162	What would the obstacles be from a CDI point of view? I think that handling extended persistence context inheritance across different component types would be possible (in theory). Since clustering an extended persistence context, isn't defined by any specifications, it probably isn't a blocking requirement to figure out how that would be done. However, serializing/deserializing the component chain when a managed bean is the top level component, would be good to consider (with regard to how that could work for implementations that support clustering of extended persistence contexts).	1	1	https://issues.jboss.org/browse/CDI-162
2830	CDI-162	Also would be good to consider what happens when a managed bean inherits the extended persistence context from a stateful bean and the stateful bean is passivated. * Does the managed bean also get passivated? * Or does the managed bean stay in memory and what about its reference to the extended persistence context (does the xpc stay in memory or is the reference cleared?)	1	1	https://issues.jboss.org/browse/CDI-162
2831	CDI-162	I think it would need to be bound to a scope, not to a component. I don't particularly see any obstacles from a CDI perspective. As Mark points out, trying to serialize the EntityManager isn't possible, so defining this for non-passivating scopes only is really our only option, in a clustered environment (or any other environment which tries to serialize the session). Scott, can you elaborate on what you mean about "serializing/deserializing the component chain when a managed bean is the top level component" - what do you need concretely?	1	1	https://issues.jboss.org/browse/CDI-162
2832	CDI-162	The extended persistence context (EntityManager) could be serializable. If the extended persistence context is not serializable the stateful bean will not be serializable as well. I'm not sure if this helps [with the concern raised here|https://struberg.wordpress.com/2012/04/25/is-there-a-way-to-fix-the-jpa-entitymanager]. From EJB3 core: {quote} 4.2 Conversational State of a Stateful Session Bean The conversational state of a stateful session object is defined as the session bean instance's field values, its associated interceptors and their instance field values, plus the transitive closure of the objects from these instances' fields reached by following Java object references. To efficiently manage the size of its working set, a session bean container may need to temporarily transfer the state of an idle stateful session bean instance to some form of secondary storage. The transfer from the working set to secondary storage is called instance passivation. The transfer back is called activation. In advanced cases, a session object's conversational state may contain open resources, such as open sockets and open database cursors. A container cannot retain such open resources when a session bean instance is passivated. A developer of a stateful session bean must close and open the resources in the A container may only passivate a stateful session bean instance when the instance is not in a transaction. A container must not passivate a stateful session bean with an extended persistence context unless the following conditions are met:[9] * All the entities in the persistence context are serializable. * The EntityManager is serializable. A stateless session bean is never passivated. [8] PrePassivate and PostActivate lifecycle callback interceptor methods. [9] The container is not permitted to destroy a stateful session bean instance because it does not meet these requirements. {quote} What I mean concretely about the top level component, is as follows. The extended persistence context is created at the time that a stateful bean is created (injection or jndi lookup) and that stateful bean is the "top most" stateful bean. When the "top most" stateful bean, injects or lookups another stateful bean that also has the same named extended persistence context, the additional bean inherits the extended persistence context. For the concrete example, please consider two cases, where the managed bean is the top most bean that has looked up a stateful session bean. I'm not quite sure what it would mean for the stateful session bean instance to be passivated (since the top most managed bean will still reference the persistence context (keeping the xpc pinned in memory). Even worse, would be if we are clustering and we failover the stateful session bean to another cluster node. I think for both of these respective cases, the managed bean would need to also be passivated or failed over to the same cluster node as the stateful bean. Does this make any sense?	1	1	https://issues.jboss.org/browse/CDI-162

#	Requirement	Questions	Answered	Successfully Answered	URL
2833	CDI-162	Actually, the EntityManager bean that is produced by Arne's code sample is passivation capable. As defined by 3.6.1, the bean is a "resource bean". And from 6.6.2 "all resources are passivation capable dependencies". So the CDI container will not choke on this definition. However, the real problem is the one Arne outlines in the last paragraph. Scott, the clustering team @ JBoss always planned to do some work on this. Do you know the status?	1	1	https://issues.jboss.org/browse/CDI-162
2834	CDI-162	I agree on the CDI problems and I state that this is even broken in EJB. Yes, I know the old answer that EJB container vendors might implement this as Serializable. But what about the users? How can a user ever implement a portable feature based on this non-reliable functionality? IF the JPA entities used are only using optimistic locking and have a way to transport id+version+state info _then_ it would be possible to auto-reattach via a Serializable EM. But that is beyond the business of this very JSR.	1	1	https://issues.jboss.org/browse/CDI-162
2835	CDI-160	Sure it would be a great solution but is it realistic to have this in a future CDI 1.1 MR or should we aim next release	0	0	https://issues.jboss.org/browse/CDI-160
2836	CDI-158	Are you talking about 'beans' in the sense of Bean<T> or in the sense of Contextual Instances?	0	0	https://issues.jboss.org/browse/CDI-158
2837	CDI-158	Are you taking about modifying all EJB and EE metadata via PAT events? This would be a very major change. From an AS7 perspective this would require a fairly big re-work of all affected subsystems, and would also have a negative performance impact as all services would end up with a dependency on weld bootstrap. I imagine other app servers would also require fairly major changes to support this.	0	0	https://issues.jboss.org/browse/CDI-158
2838	CDI-157	I know we talked about this on IRC, but where would an ordinal come into play for an extension? You won't have the same extension loaded twice. We're talking about executing extensions with a higher ordinal first, correct?	1	0	https://issues.jboss.org/browse/CDI-157
2839	CDI-154	{quote} we currently only check for Serializable and completely ignore Externalizable in the spec {quote} Isn't that because Serializable is a superinterface of Externalizable and therefore every class implementing Externalizable also (indirectly) implements Serializable?	1	0	https://issues.jboss.org/browse/CDI-154
2840	CDI-151	That's a pitty because I already spent some time creating TCK tests - e.g. {{org.jboss.cdi.tck.tests.context.jms.MessageListenerContextTest}} - and even filed an issue - see AS7-3107. Anyway where is this defined, JMS spec? It would be nice to have it referred...	0	0	https://issues.jboss.org/browse/CDI-151
2841	CDI-151	I talked to Nigel Deakin, JMS spec lead about this. Here is our correspondence: {quote} On 24/08/2012 22:51, Pete Muir wrote: > Thanks. > > On 24 Aug 2012, at 18:27, Nigel Deakin wrote: > >> Pete, >> >> On 24/08/2012 15:12, Pete Muir wrote: >>> Nigel, >>> >>> Could you cast your eye over https://issues.jboss.org/browse/CDI-151 and check that it does make sense? >>> >>> Related spec section is at http://docs.jboss.org/cdi/spec/1.0/html/contexts.html#d0e5377 >> >> I spotted this the other day and wondered about it. The references to a javax.jms.MessageListener are a little >> ambiguous. >> >> There are two ways in which messages can be delivered to a MessageListener in Java EE: >> >> 1. In the Java EE EJB container, applications can define a message-driven bean which implements >> javax.jms.MessageListener. >> >> 2. In the Java EE application client container (or Java SE), applications can define any class they like which >> implements javax.jms.MessageListener, create an instance, and register it using the JMS API. However applications >> running in the Java EE web or EJB containers are forbidden from doing this (this prohibition is in the platform >> spec). >> >> When the CDI spec 6.7.3 refers to "message delivery to a MessageListener", which of the above is it referring to? >> >> If it is referring to (1) then this is already covered in the preceding bullet point which mentions "message >> delivery to any EJB message-driven bean". >> >> If it is referring to (2) then this really only makes sense if the application context is active in other places in >> the application client as well. Since there is no other mention of the application client I suspect this was not >> intended. > > CDI doesn't require support for application clients, so I guess it must be referring to (1). Are you ok with me > pasting this conversation into the issue? Yes, that's fine. > I'll resolve it by just deleting that text :-) Yes, I agree that's the most appropriate thing to do. Nigel {quote}	1	1	https://issues.jboss.org/browse/CDI-151
2842	CDI-151	So you're not allowed to consume JMS messages in EE container unless you use MDBs? Sounds a little bit limiting. {quote} However applications running in the Java EE web or EJB containers are forbidden from doing this (this prohibition is in the platform spec). {quote} Where is this defined? Cannot not find it in the EE 6 platform spec...	0	0	https://issues.jboss.org/browse/CDI-151
2843	CDI-144	Since I have been running into this issue lately: JPA does not recommend EntityManager interactions within JPA lifecycle callbacks*. Shouldn't we add a note somewhere that injected resources require special attention regarding JPA compliance? And because of that firing CDI events from a JPA entity listener is particularly dangerous? In my case I tried to implement global available created/updated/deleted entity events which turned out to be an bad idea if implemented on EntityListener level. * JPA Spec 3.5.2: In general, the lifecycle method of a portable application should not invoke EntityManager or Query operations, access other entity instances, or modify relationships within the same persistence context.[45] A lifecycle callback method may modify the non-relationship state of the entity on which it is invoked.	1	1	https://issues.jboss.org/browse/CDI-144
2844	CDI-142	here is an example of such a problematic scenario * EAR with webapp1 and webapp2 * EAR contains @ApplicationScoped MailService (in a shared ejb-jar) * webapp2 contains @Specializes MyOtherMailService extends MailService * we have proper ClassLoader isolation between webapp1 and webapp2 now we use @Inject MailService in webapp1 and webapp2, * what do they get? * what if we shutdown webapp2 programmatically? No impl specifics around here, just pure specc issues imo...	1	1	https://issues.jboss.org/browse/CDI-142
2845	CDI-139	I think we don't need anything like that. It happens that all (non EG) people I talked about this very problem (using Instance<T> for creating @Dependent beans) did not even understand that they created a memory leak! What can we do against it? We could of course add a destroy(Object) instance, but this might be hard to use. Of course, users would really need to use it and old code would still be broken. My solution: the injected Instance<T> itself is a managed contextual instance as well, and it has Scope @Dependent. Thus it is possible to add a @PreDestroy to it which will automatically get called by the container whenever the containing contextual instance will get destroyed. If we now in Instance#get() check if the created contextual instance is @Dependent and then store it in a List inside the respective Instance<T> instance, then we can also easily cleanup them in the @PreDestroy of Instance<T> which got used to create all the Dependent contextual instances. Thus we would not create any mem leaks anymore.	0	0	https://issues.jboss.org/browse/CDI-139

#	Requirement	Questions	Answered	Successfully Answered	URL
2846	CDI-139	> I think we don't need anything like that. It happens that all (non EG) > people I talked about this very problem (using Instance<T> for creating > @Dependent beans) did not even understand that they created a memory leak! > > What can we do against it? We could of course add a destroy(Object) instance, > but this might be hard to use. Of course, users would really need to use it > and old code would still be broken. > My solution: the injected Instance<T> itself is a managed contextual instance > as well, and it has Scope @Dependent. Thus it is possible to add a @PreDestroy > to it which will automatically get called by the container whenever the containing > contextual instance will get destroyed. > > If we now in Instance#get() check if the created contextual instance is @Dependent > and then store it in a List inside the respective Instance<T> instance, then we > can also easily cleanup them in the @PreDestroy of Instance<T> which got used to > create all the Dependent contextual instances. Yes, we should certainly clarify that this is the behavior required, and this is actually what creates the "memory leak" in the first place. The issue is most apparent you associate the Instance with an application scoped bean, it means the memory won't be released until the application is undeployed. This is correct, but has unfortunate side effects. Weld implements this today, not via the impl you propose, but it has the same effect. There are a few optimisations we can do - if there is no @PreDestroy callback on the dependent bean, there is no need to hold a reference to it at all. We could also offer a mode where only a weak reference is held by Weld, so that when the app drops a reference, it can be gc'd. We could of course suggest people don't use Instance with application scoped beans, but this seems a bad restriction. I think offering a mode where people can manually clean up their Instances is useful. It won't affect the functionality of an app if you don't use it, but it can offer an optimisation people can use. > Additionally we need to specify, when @Dependent instances that are obtained via CDI.current() or > BeanManager.getReference(...) are destroyed. At least this should be happen on container-shutdown. Beans obtained via BeanManager must be manually cleaned up using Contextual.destroy(). We should add that instances obtained via CDI.current() will be treated as belonging to the lifecycle of the application, and so automatically cleaned up when the application is stopped. Can you file an issue for that? IMO this latter point also emphasises the need for a Instance.destroy() method, most of the time people will want their beans cleaned before the application is stopped.	0	0	https://issues.jboss.org/browse/CDI-139
2847	CDI-139	The proposal has been updated to account for CDI-89. Open issue is OPEN ISSUE: What name should the AlterableContext interface be given?	1	1	https://issues.jboss.org/browse/CDI-139
2848	CDI-132	things get more complicated. I think the current AnnotatedType maintains no 'relation' to it's parent class. Instead it contains the 'effective' model only. But this is imo not enough. It also doesn't say anything about modifiers public/protected/private. What should the AnnotatedType of the following class provide?: @ApplicationScoped public class MyBaseSvc { private @Inject EntityManager; ... } @ApplicationScoped public class DongService extends MyBaseSvc { ... }	1	1	https://issues.jboss.org/browse/CDI-132
2849	CDI-129	It might well be that the spec is contradicting. But however we turn this, the fact remains that there are 2 different scopes needed. 1 Webapplication scope and 1 ear scope. I doubt that the original intention of @ApplicationScoped was for EAR usage! It got added way later, but the historic evolution of the spec, and the fact that it heavily and easily integrates with EJBs allows that we just use EJBs for those cases you mentioned. We can still use EJBs for those things. If we would interpret @ApplicationScoped your way, then it would behave *heavily* different depending if the war gets used/tested standalone or as part of an EAR. That sounds just plain wrong imo. Btw, it IS possible to fulfil the the spec. In OWB we handle JMS, and quartzed jobs, etc as own 'applications' inside the EAR. So they get an own BeanManager and can even have their own Extensions. See what I mean?	1	1	https://issues.jboss.org/browse/CDI-129
2850	CDI-129	{quote} It might well be that the spec is contradicting. But however we turn this, the fact remains that there are 2 different scopes needed. 1 Webapplication scope and 1 ear scope. {quote} I think that there's a general agreement on that ;). It just remains to decide which direction should we head into. {quote} I doubt that the original intention of @ApplicationScoped was for EAR usage! It got added way later, but the historic evolution of the spec, and the fact that it heavily and easily integrates with EJBs allows that we just use EJBs for those cases you mentioned. We can still use EJBs for those things. {quote} Unless we have a clear record that @AppScoped was intended to be used in a way or another, any discussion about the original intent is purely conjectural. An @ApplicationScoped singleton and a @Singleton EJB are not even remotely the same thing. I can only read it clearly from the spec that the intent was to create a scope shared by all the components of an application. We are discussing now what exactly "application" means, and I cannot see any reason why would it mean something different than the Java EE application/module/component breakdown, especially since components outside a "web application" are called out explicitly. {quote} If we would interpret @ApplicationScoped your way, then it would behave heavily different depending if the war gets used/tested standalone or as part of an EAR. That sounds just plain wrong imo. {quote} How different exactly? If you are deploying two WARs as part of the same EAR or independently you are expected to see differences anyway, starting with the behaviour of the "java:app" namespace, CDI or not. {quote} Btw, it IS possible to fulfil the the spec. In OWB we handle JMS, and quartzed jobs, etc as own 'applications' inside the EAR. So they get an own BeanManager and can even have their own Extensions. See what I mean? {quote} So each EJB is an application by itself, or all EJBs share the same 'application' context? Because the latter is what the specification requires - see above.	1	1	https://issues.jboss.org/browse/CDI-129
2851	CDI-129	I can say with 99.9% certainty that Gavin (at least) intended the @ApplicationScoped to be "ear scoped" when deploying inside an EAR. Unfortunately we can't easily go back and check. Mark, can you explain what you mean by it behaving differently in this situation. As I see it, it actually is more consistent to make it work like this. Regarding {quote} Btw, it IS possible to fulfil the the spec. In OWB we handle JMS, and quartzed jobs, etc as own 'applications' inside the EAR. So they get an own BeanManager and can even have their own Extensions. See what I mean? {quote} this does seem somewhat orthogonal to the question. In Weld an EJB jar would also get it's own BeanManager (and it should get it's own extensions, this is a bug in Weld right now) but it shares the Application context with all other modules in the application. So unless you can clarify what you mean here, I think we can ignore this. Marius, thanks for digging up those references, I think with those we can be quite confident that that the spec intends @ApplicationScoped to be "ear scoped" and that the reference to the Servlet spec was a mistake. In particular the comment "in an application deployed as an ear" indicates that the CDI spec regards an EAR as an application, and that coupled with the comment "the application context is shared ... within the same application" does indicate that the context spans the entire ear. To me, this is a stronger statement than "represent the standard scopes defined by the Java Servlets specification" which seems somewhat more descriptive. Mark, happy to consider adding an @WebApplicationScoped, and we should definitely clarify what @ApplicationScoped means.	1	1	https://issues.jboss.org/browse/CDI-129

#	Requirement	Questions	Answered	Successfully Answered	URL
2852	CDI-129	Mark, it's *not* "absolutely clear" from the CDI specification that this was the intent, if it were we wouldn't even be having this discussion ;-) We also can't base decisions on you saying "it's absolutely clear to me", we need to base them on facts ;-). So we have to leave aside such subjective statements. I understand the concerns about the behaviour changing between war and ears for some who have interpreted the spec one way, however this runs both ways. So this argument is essentially nullified in this discussion (an app which assumed one will stop working when we require the other). So far, I believe the case for it being interpreted as war scoped are a) 2.4.1 "The @RequestScoped, @ApplicationScoped and @SessionScoped annotations defined in Section 6.7, "Context management for built-in scopes" represent the standard scopes defined by the Java Servlets specification." b) That other specs in Java EE (namely JSF, JAX-RS, JAX-WS) refer to the "application scope" as scoped to the war (could you provide references to their specs for this) and as ear scoped a) that the spec describes an ear as an application, and says that the application context is shared across an application So to summarize: * the spec is clear that the application context should be shared across the "application" * the confusion is around what "application" means in the context of an ear * in one place the CDI spec mentions that the scopes "represent" those from the servlet spec, implying that the application means war scoped * in another place the CDI spec mentions that the ear represents the application. I think we need to gather two more opinions: * what was the original intent (I'll ping Gavin) * how does Java EE define "application" in an ear (I'll ping Linda and Bill) Mark, is this a fair summary from your side of the argument? I want us to agree on a summary so that we can move forward with a discussion about which course to take, rather than get stuck in debating the spec ;-)	1	1	https://issues.jboss.org/browse/CDI-129
2853	CDI-129	Over all checking for discrepances, I completely skipped the points we already seem to agree: a) the current @ApplicationScoped is not enough (no matter how we interpret it) b) we need at least one more Scope for EAR scenarios. This could either be b.1) @EnterpriseScoped (if we interpret the current @ApplicationScoped as Webapp singleton) b.2) @WebApplicationScoped (if we interpret the current @ApplicationScoped as EAR-singleton) agree?	1	1	https://issues.jboss.org/browse/CDI-129
2854	CDI-129	Bill Shannon points out that we also need to consider visibility of the bean when clarifying this. For example, if you have the same bean class specified in two separate wars, both of which are application scoped, you can't share the state between the two (as you would get a CCE). In fact, this makes me wonder if there isn't an "obvious" solution staring us in the face? If we clarify that @ApplicationScoped means "shared across the ear" and also that the same bean class packaged in two locations which are not visible to one another (e.g. warA and war B) and cannot share state. Now we provide both options of ear scoped state and war scoped state, and the state is only shared if the bean class is packaged in a shared library or shared module (e.g. ejb module). This seems like a good solution: a) it means we don't need to proliferate the number of "application scoped" annotations b) it makes logical sense (if a user wants per-war state, package the bean in a war, if the user wants shared-ear state, then package it as a shared lib). WDYT?	1	1	https://issues.jboss.org/browse/CDI-129
2855	CDI-129	Can you point me at where @Singleton is defined like this? Or help me understand how you are interpreting the definition like this?	1	1	https://issues.jboss.org/browse/CDI-129
2856	CDI-129	Probably because of the heavily underspecified classloader behaviour in EE in this area? I was really confused about this too also. Because the Servlet spec (which the EE umbrella spec references prominently) imo _exactly_ defines that WebApplications must not interfere with other WebApplications but must be isolated well from each other. But most old EE servers (JBoss4 and 5 are perfect examples for this, as are the old WebLogics) ship with almost no isolation at all it seems! I actually never understood why they behave this way as _default_! @Marius: A Bean can only be used in a shared library if it's type is available on the shared classpath too. E.g. in a ejb-jars lib. And even if one webapp uses an @Alternative which is only available in this very one webapp, then there is no problem. Because injecting an @ApplicationScoped bean will still give you the contextual reference (the proxy). They simply refer to different contextual instances. There is really not much difference to a @SessionScoped bean of which it's class is in a shared ejb-jars lib. Or am I overlooking something?	1	1	https://issues.jboss.org/browse/CDI-129
2857	CDI-129	You are right, if the EE container does not implement the suggestions of the first paragraph of JSR-316 EE.8.3.1 but decides to use reduced or no classloader isolation at all then 'non-portable behaviour occurs'. I personally consider this an insane scenario, and the EE-7 spec should find very clear words that this is crap and should only be used in exceptional cases, e.g. for backward compatibility with old projects ;) For the (web) @ApplicationScoped only applying to one webapp - this is also possible if there is no classloader isolation at all. For handlig this case in OpenWebBeans, we are just setting the 'one' @ApplicationScoped for a certain request to a ThreadLocal (See OWBs WebContextsService - we are doing this via a RequestListener subsequently via initRequestContext()). Remember that we have plugable lifecycle handlers in OWB. This would not be necessary if there is a well defined ClassLoader isolation in place (as suggested - we could then just use the ContextClassLoader in this case) but this works really well with very old WebLogic servers in a big swiss bank already ;)	1	1	https://issues.jboss.org/browse/CDI-129
2858	CDI-129	{quote} But most old EE servers (JBoss4 and 5 are perfect examples for this, as are the old WebLogics) ship with almost no isolation at all it seems! I actually never understood why they behave this way as default! {quote} Believe it or not, older versions of JBoss AS (than 4) did provide more isolation by default, but by user request the isolation was downgraded. Also worth pointing out that JBoss AS 4 and 5 did support much stronger isolation if you enabled it (we always recommended you did with Seam, and enabled it automatically with Weld so I always forget this was the default!). {quote} @Marius: A Bean can only be used in a shared library if it's type is available on the shared classpath too. E.g. in a ejb-jars lib. And even if one webapp uses an @Alternative which is only available in this very one webapp, then there is no problem. Because injecting an @ApplicationScoped bean will still give you the contextual reference (the proxy). They simply refer to different contextual instances. There is really not much difference to a @SessionScoped bean of which it's class is in a shared ejb-jars lib. {quote} I'm a little lost in this sentence... What are you trying to say? Jens/Marius, I think your point is a very reasonable one for the spec. Of course, we would need an FAQ explaining the details of this ;-) Mark, please take this comment to the EE expert group, as you know we can only decide on CDI here ;-) This proposal is certainly implementable, no concerns there. Gerhard, we all agree with this rhetoric here (no real need to repeat it), we just need to find a solution that (a) makes sense technically and (b) works in the environment we are in. What is proposed does make sense (to me) and does work technically except in the case that people use non-isolated classloading, which EE keeps as an option for backwards compatibility. IOW what we have is a solution that makes sense with the recommended deployment model (isolated classloading), but we need to include, in the spec, a note about non-portability for the backwards compatible non-isolated classloading case.	1	1	https://issues.jboss.org/browse/CDI-129
2859	CDI-129	Gerhard, please can you clarify how we can change the current proposal in order to remove the non-portability clause. We can't make changes to the EE spec here.	1	1	https://issues.jboss.org/browse/CDI-129

#	Requirement	Questions	Answered	Successfully Answered	URL
2860	CDI-129	There is a similar issue related to events. I am not sure if the proposed solution for the @ApplicationScoped issues fixes the events issue implicitly or a clarification of the Events chapter is needed. The issue follows: Should a web application bundled within an ear should be notified of an event fired by another web application bundled within the same ear? The observer resolution algorithm should take bean visibility into account and notify only those observer methods from whom the event firer is visible.	1	1	https://issues.jboss.org/browse/CDI-129
2861	CDI-129	Oki to pick up that stuff again @StuartDouglas: > ... beans that execute within the same application. Too bad it's neither in the CDI nor the EE spec defined what they mean with application. Both sometimes refer to web-applications and sometimes to enterprise-applications. It's just not clear. @PMuir From the historical perspective. I read through the very early mails in the webbeans dev mailing list and the original intention seems to have been to glue together JSF and EJB. And as such the CDI scopes were meant as replacement for JSF scopes. All the EAR stuff finds a first mentioning in spec drafts in middle 2008. Now guess how JSF ApplicationScoped behaves? Exactly, 1 per JSF Application == 1 per webapp. With other words. There exists a defined behaviour for ApplicationScoped beans in Java EE since 10 years, and it is 1 per webapp! Oki, lets take aside how we name that baby. The real problem is not how to cope with @ApplicationScoped, but how injecting into anything which has a scope bigger than some class visibility would work. The biggest unique visibility area is a single web-application. Thus we are talking about any scope 'bigger' than 1 per webapp. It doesn't matter if this is @EnterpriseApplicationScoped or a @Statefuf EJB which is shared across the whole EAR - the underlying problem is the same! And just to be clear: the containers I tested so far just blow up with the above mentioned scenario. Oki here it goes. The scenario is as following * A shared jar with interface X and @EnterpriseApplicationScoped public class A implements X and an @EnterpriseApplicationScoped public class MailService { @Inject X;. * webapp1 with @SessionScoped public class BackingBean1 { @Inject X x; * webapp2 with @Specializes public class B extends A and @SessionScoped public class BackingBean2 { @Inject X x; a.) I've now talked with lots of different folks, including JSR spec leads, and NO ONE did expect that behaviour. b.) what proxy will be used for the injection point in MailService#x ? Most containers atm just inject a proxy with the Bean of the first request. If webapp2 gets accessed first, then it will create a contextual instance of MailService and the injection point will be filled with the proxy for B. If now a request from webapp1 arrives then the proxy will return an instance of B which will cause a ClassNotFound exception or similar. Because webapp1 cannot see class B. In the best case you will 'only' randomly end up with the wrong class... This is btw the same if you inject a @SessionScoped or @RequestScoped bean into a bean with a larger scope than 1 per webapp. @Specializes and @Alternative beans will be randomly broken for any injection point of a bean with a broader scope than 1 per webapp. So there are 2 broken situations with most implemented proxies atm 1. some webapps might get classes/impls which dont belong to them 2. some webapps might get NO @Alternative or @Specializes despite being configured This could perfectly get fixed by such proxies by just implementing the injection point type and do getBeans(injectionPointType) + resolve the instance. But this is not defined that way in CDI as Jozef pointed out. Currently the proxies contain the Beans, so we have no chance to do this properly in a spec conform way for CDI beans. The easiest solution would be to just interpret @ApplicationScoped as 1 per webapp and we are done. Jozef, could you please also describe your preferred way? Afair it was 2 options: Option a.) just disallow @Specializes and @Alternatives for all beans defined in a shared jar. Regardless of it's scope. Option b.) inject only beans defined in the shared jar into another bean of a shared jar. E.g. MailService would always get an instance of A injected. Imo this is pretty crucial as @Inject X x; in a JSF backing bean would give you a different contextual instance (even of a totally different type) than the same code @Inject X x; in the MailService.	1	1	https://issues.jboss.org/browse/CDI-129
2862	CDI-129	> Firstly, here is how the modularity work in CDI: Let's say this is your view and is one of n possible interpretations of the spec wording. Due to the well defined TCCL behaviour in the EE umbrella spec it is perfectly viable that A might see B IF it is invoked in a request of the webapp containing A. So technically from the EE spec there is NO such restriction you mentioned. It's just easier for a container vendor to implement. And your example is also _perfectly_ broken from a user experience pov! If you inject X in a shared jar you would get another contextual instance than injecting X in a JSF backing bean that way! Hell, this is SICK SICK SICK! No user would be able to use that stuff! I have a ┼@Inject User usr; in my MailFormular backing bean and the same code in my MailService and would get 2 DIFFERENT users? You're kidding me, right? The are only 2 practical ways to do this properly for users: 1.) restrict any @Specializes or @Alternative of an interface to the visibility area it got defined in. In other words if you have a User interface in a shared jar in your EAR we disallow @Specializes or @Alterative in any WAR. But again, this is NOT defined in CDI-1.0 yet! 2.) use a proper proxy which only proxies the injection point type (in our case X) and not only resolve the contextual instance but also the proper Bean<X> dynamically on each invocation IF it got injected into a 'broader scoped bean'. And here we come back to my first argument: by keeping the specified behaviour of CDI-1.0 as 1 per webapp we would not have any 'broader' scope in CDI and thus we dont need to deal! We could introduce all the necessary handling in CDI-1.1 together with the broader @EnterpriseApplicationScoped. Another point: who likes to go to all the other spec leads like JPA and most important JSF and confess that we messed up? Especially who gonna tell Ed that he cannot drop JSF ManageBeans because we have no replacement for his ┼┼┼JSF @ApplicationScoped which exists since 10 years and he finally agreed to get rid of?	1	1	https://issues.jboss.org/browse/CDI-129
2863	CDI-129	> In other words, CDI contexts are never isolated between bean archives. I don't get that. Where do you read that I like to do that? What I was talking about was that you always see only exactly 1 'slice' of a context for each Thread. All users share the same SessionContext, but still each request can only see the contextual instances of the single currently active session of each request. The same would happen for webapplications. Each request knows to which webapp it belongs. And a single request can only belong to exactly one single webapp! So he will only see the @ApplicationScoped contextual instances of this very webapp the request is for.	1	1	https://issues.jboss.org/browse/CDI-129
2864	CDI-129	Sorry Jozef, you are a pretty clever dude in general and you did hit quite some nails. But in this case you started to drop "it is that way" sentences but missed the most important part of it: the actual argument! I had expected something like "I think it is that way BECAUSE ${argument}, ${quote}, etc" ;) And no, it is not unrelated to cdi-1.0. Another argument why we need further clarification. If you get the BeanManager from JNDI then according to your interpretation you would like to get different results for getBeans() depending on what? How would that ever work? Please shed some light, I don't get it. I think we would need to add some BeanManager#getBeans(injectionclass, type, qualifiers) + resolve with the same + getReference ... We do not have all that currently.	1	1	https://issues.jboss.org/browse/CDI-129
2865	CDI-129	Pete, afair we just agreed that most containers implement it that way currently and that this is basically broken. Please think about the consequences! That would render CDI almost useless! Also please comment on my BeanManager clarification questions. What will a BeanManager#getBeans() return you? On that very meeting I told you that I will write up concrete examples and here we go.	1	1	https://issues.jboss.org/browse/CDI-129

#	Requirement	Questions	Answered	Successfully Answered	URL
2866	CDI-129	Mark says: {quote} Due to the well defined TCCL behaviour in the EE umbrella spec it is perfectly viable that A might see B IF it is invoked in a request of the webapp containing A. So technically from the EE spec there is NO such restriction you mentioned. It's just easier for a container vendor to implement. {quote} I don't believe the TCCL is relevant here, it's only ever used explicitly to load classes, not implicitly if you instantiate via the new keyword, which is the model we are trying to follow. {quote} Another point: who likes to go to all the other spec leads like JPA and most important JSF and confess that we messed up? Especially who gonna tell Ed that he cannot drop JSF ManageBeans because we have no replacement for his ┼┼┼JSF @ApplicationScoped which exists since 10 years and he finally agreed to get rid of? {quote} I'm quite happy to take on any tasks that involve discussions with other spec leads. As JSF is prepared to deprecate scopes, and move people over to standard CDI scopes, there is a code change required anyway. Therefore, if we were to introduce a @ModuleScoped or similar, I think JSF apps could move from JSF application scope to this.	1	1	https://issues.jboss.org/browse/CDI-129
2867	CDI-129	Proposed question for the poll is {quote} Should @ApplicationScoped bean instances be shared between all modules (web apps, ejb modules etc), or be per web application? {quote} Mark, wdyt?	1	1	https://issues.jboss.org/browse/CDI-129
2868	CDI-129	Do you have the code somewhere? What I observed in JBossAS and Glassfish is that the WebApp which touches the shared MailService the first time will do all the injection based on the active beans in this very webapp. Regardless if the resolved Bean<X> is visible in the other webapps or not. This is not only wrong but might in some cases even create the obvious NoClassDefFound and similar (great candidate: serialisation time). Of course one could argue this is an impl bug, but otoh how to implement that correctly? And now we are back to the 3 listed ways to do it: 1.) injecting only other beans of the shaed-lib and completely ignore anything defined in the webapps. Jozefs way is technically valid, but a no-go from a user perspective imo as it would lead to having different contextual instances for the same injection points across your application. 2.) restricting @Specializes and @Alternatives to only shared libs. This is technically valid as well but this restriction is certainly not covered by the spec. 3.) make @ApplicationScoped 1 per WAR and thus remove the problem with scopes over multiple isolation levels. All our classloader slices would always see exactly 1 well defined scenario. We should nonetheless think about any scope broader than 1 per war in CDI-1.1, but for CDI-1.0 it would perfectly work (And to repeat again: this is also what is explicitly stated in the spec 2.4.1).	1	1	https://issues.jboss.org/browse/CDI-129
2869	CDI-129	{quote} Should @ApplicationScoped bean instances be shared between all modules (web apps, ejb modules etc), or be per web application? {quote} again the term 'bean' which is so highly ambiguous I guess you fear that the public audience is not yet fully aware of 'Contextual Instance'? :) I'd rather try it a bit more sluggish: {quote}What would you expect to get from an @ApplicationScoped CDI bean in an EAR? 1 instance per each Web-Application or 1 instance for the full EAR?{quote} Regarding the Bill quote: I've now read through the EE6 spec and it appears to me that they mostly refer to EARs as "Java EE Application" (in older spec versions it was "J2EE Application"), see e.g. EE.2.11.3, EE.2.4, EE.2.12.2, EE.5.5.2 and 185 more references. There is other prominent mentioning of "JSF Application", "Web Application", even "Desktop Application" and "Java Application" amongst a dozen others.	1	1	https://issues.jboss.org/browse/CDI-129
2870	CDI-129	> The same thing would happen for a @Dependent scoped bean > or a @Request scoped bean. Although Mark thinks that restricting > @ApplicationScoped per war fixes this it actually does not. Again Jozef, this is JUST NOT TRUE for CDI! The problem is not the scope of the bean which gets injected, but the scope of the bean you inject into. Are we d' accord on this part? It just doesn't friggin depend whether the injected bean is @Dependent, @SessionScoped or any other scope. So we are just talkin about the scope of the bean which helds the injection point. And here comes the clue: By not providing a scope bigger than 1 per WAR, we DO NOT HAVE THOSE AFOREMENTIONED PROBLEMS IN CDI. We could go on and define the behaviour for scopes larger than 1 per WAR in CDI 1.1. The EJB EG would probably need to clarify the restrictions for injecting into shared stateful EJBs, but indeed that is not so much a problem of CDI. > I disagree. on which part? On the 'unusable' or the 'multiple contextual instances per Injection Point for the same call stack'?	1	1	https://issues.jboss.org/browse/CDI-129
2871	CDI-129	{quote} By not providing a scope bigger than 1 per WAR, we DO NOT HAVE THOSE AFOREMENTIONED PROBLEMS IN CDI. {quote} OK, assume there is a shared @RequestScoped bean R in the shared library injected both into a Servlet in your web app and into the MailService in the shared lib. No alternatives nor specialization this time. What you say is that assuming the Servlet invokes MailService, the Servlet would operate on a different instance of the same R Bean (Bean<R>) than the MailService? In the same call stack during processing of the single same request? How is that consistent? Maybe we could do a poll if this is what people want ;-) This is not CDI. {quote} on which part? On the 'unusable' or the 'multiple contextual instances per Injection Point for the same call stack'? {quote} I disagree that it is unusable.	1	1	https://issues.jboss.org/browse/CDI-129
2872	CDI-129	what is 'shared mail service' ? If we assume {code} @ApplicationScoped public class MailService { @Inject R r; ... } {code} then each webapp would get it's own contextual instance. Thus the MailService instance of your very webapp inoking r.someMethod() would work on an instance of S. For another webapp you might end up with an instance of R if this webapp doesn't specialize and not provide an alternative for S. > S, however S is not accessible from the MailService so this is wrong Dude, who is creating the instance of MailService? It is the BeanManager which uses the TCCL of your webapp. Also: this BeanManager of your webapp knows all enabled Bean<?> of your webapp, so it will perfectly be able to resolve Bean<S>.	1	1	https://issues.jboss.org/browse/CDI-129
2873	CDI-129	here is the original quote from pete {quote}I don't believe the TCCL is relevant here, it's only ever used explicitly to load classes, not implicitly if you instantiate via the new keyword, which is the model we are trying to follow.{quote} That doesn't sound like 'it is definitely that way'. Indeed you almost NEVER do do 'new SomeClass()' in dependency injection frameworks at all! That's the whole idea of DI containers to always do this dynamically. Imagine we do all this by hand: step 1: create a MailService mailService = new MailService() step 2: create a S s = new S; // S extends R step 3: mailService.setR(s); What is the problem with that? I cannot see any classloading issue! Btw: Weld 1.x and 2.x implement it EXACTLY THAT WAY ALREADY! Of course, the MailService cannot have a @Inject S s; but only a @Inject R r; But that doesn't restrict us to only inject instances of type R into it if we ONLY go on using those beans in our very own webapp!	1	1	https://issues.jboss.org/browse/CDI-129
2874	CDI-129	Can you tell me which part of Section 5.1, which describes how modularity works in CDI, allows you to inject S into MailService?	1	1	https://issues.jboss.org/browse/CDI-129
2875	CDI-129	Jozef, thanks for your input so far. There are some valid points in your argumentation and we need to clarify those. I guess we now have a pretty clear picture that EAR scenario is broken in CDI 1.0 if we use scopes broader than 1 contextual instance per war. We also outlined 3 different approaches to treat such scopes. We will schedule an EG meeting to discuss this in the EG. Pete, do you have the example app somewhere so we can look at the sources?	1	1	https://issues.jboss.org/browse/CDI-129

#	Requirement	Questions	Answered	Successfully Answered	URL
2876	CDI-129	Gosh Jozef, please don't drive me nuts! just read the thread above to find it! a.) It is not clear what the term 'Bean' used in 5.1 means b.) Pete was not sure about the TCCL, go on read the TCCL paragraph again please. Imo it is relevant. c.) your interpretation is contradicted by the example shown in "2.2. Bean types" d.) the whole BeanManager API design contradicts your arguments what else? The spec is just not clear. I f***n don't care on which end. If A contradicts B then there are 3 ways to fix that: fix A, fix B or fix both. To me the only sure thing is that we need to fix something as no possible interpretation closes all the gaps. Although the interpretation with the 1 per WAR would imo be the closest one and is also the only interpretation EXPLICITLY defined in the spec. All other interpretations are just indirect consequences of other definitions.	1	1	https://issues.jboss.org/browse/CDI-129
2877	CDI-129	There is a lot of stuff here, but firstly: a) Bean is quite well defined in CDI. If it gets a ProcessBean event thrown then it is a bean. b) You are correct and TTCL is irrelevant here. Class loading is actually pretty irrelevant here, the EE platform spec defines class loading in terms of classes that must and must not be visible, and allows vendors to implement that with whatever class loading arrangement works So from what I see the problem is that there is a conflict in the spec to do with how injection point resolution works with scoping. 5.2.1 makes it pretty clean that every injection point must resolve to one and only 1 bean. In Mark's example, it actually resolves to different beans, due to scoping issues, and this is a problem that we need to sort out in the spec. Making ApplicationScoped be a per web application thing however is not a solution to this problem. This would mean that different beans resolve to the same injection point depending on the current context, which completely destroys static wiring, and we get back to a Seam 2 style resolution at injection time. It also makes the CDI deployment dependent on the current web app. What happens when there is no web app in an ear? What happens with a standalone jar? what happens with a SE app? Mark, from what I am reading it sounds like you are trying to bring back the worst of Seam 2 (I know you probably did not use Seam 2, but the approach was pretty much exactly what you are describing, and it had a lot of problems). IMHO we should just add @WebApplicationScoped and tighten up the spec to remove the scope problems, so if a bean ends up resolving injection points differently depending on the current archive it is a deployment error.	1	0	https://issues.jboss.org/browse/CDI-129
2878	CDI-129	that's also not "more objective", because it isn't only mark who has this opinion. since openwebbeans handles it that way, the owb team agreed on it as well. i know a lot of big projects which (have to) use EARs as "deployment-bundles" for huge applications, but every web-app has to be independent, but they share e.g. portable cdi extensions like deltaspike, codi, seam3, ... if you don't follow his recommendation in this case, you break them all (for weld based servers).	1	1	https://issues.jboss.org/browse/CDI-129
2879	CDI-129	Yesterday Gerhard, David and I had a really good late night discussion. I think I now understand a bit better where you guys are coming from conceptually. The isolation rules Jozef proposed are the ones from EJB. But CDI beans are not EJBs! CDI introduced the notion of Contexts and proxying. With EJBs you don't have cross-scope injection abilities, etc. A lot concepts are fundamentally more open and different in CDI. Thus I fear we also cannot apply EJB rules 1:1 to CDI beans. Now think about it again: you basically like to define @ApplicationScoped as alias for (EJB) @Singleton, right? That's what it boils down to if I read your proposal. Now the question is: why the hack do we need that? Just use an @Singleton EJB if you like to share state between different WAR files, MDB and stuff! So we do not even need to introduce any new scope because a thinkg like "@EnterpriseApplicationScoped" already exists: @Singleton. Maybe that was the reason why Gavin explicitly specified @ApplicationScoped web-app/servlet centric as 1-per-WAR by writing "The @RequestScoped, @ApplicationScoped and @SessionScoped annotations defined in Section 6.7, “Context management for built-in scopes” represent the standard scopes defined by the Java Servlets specification." The argument with the TCK is also not the strongest one. The JCP defines the following rule when resolving ambiguous definitions: 1.) JavaDoc 2.) Spec Wording 3.) TCK We (OWB) already reported 30++ bugs in the TCK and I guess others did the same. If you consider how many tests there are then this is only a small part, but still. Btw, it *really* rocks that the TCK is open sourced! This allows us to easily fix shortcomings, add more use cases and most importantly: talk about scenarios! With closed TCKs we would not even be allowed to talk about TCK use cases in a public bug tracker! @Jozef your argumentation with semi-static wiring is not quite right. This is perfectly allowed and even necessary in CDI 1.0. Please look at Bean<T> and InjectionPoint. Then look at Bean#getInjectionPoints(). This doesn't return a list of Bean<?> to be injected but only a list of InjectionPoints. And those just contain the Type and not the resolved Bean<?>. So upon creating the Contextual Instance via T Bean<T>#create(); a CDI container has to first do getInjectionPoints() and then for each InjectionPoint RESOLVE the correct Bean<?>. This is *not* defined in a static way in CDI 1.0 but fully dynamic!	1	1	https://issues.jboss.org/browse/CDI-129
2880	CDI-129	"Now the question is: why the hack do we need that? Just use an @Singleton EJB if you like to share state between different WAR files, MDB and stuff!" The problem with this is that this only works for EE applications. So if you want a singleton bean, you have to use @javax.ejb.Singleton for EE apps, and @ApplicationScoped for non EE apps, which I don't think is that great.	0	0	https://issues.jboss.org/browse/CDI-129
2881	CDI-129	I am not sure which part of my comment does your argumentation relates to. I assume you accept that there are the two possible approaches. I assume you accept that the first approach does not conflict with CDI 1.0 (Talking about the wiring only. Nothing about the @ApplicationScoped yet). Can we both agree up to this point? What I guess you are arguing about is whether the semi-static approach is valid or not.	1	1	https://issues.jboss.org/browse/CDI-129
2882	CDI-129	Note that this approach actually solves some other problems as well, such as what happens when an extension in a war modifies a shared bean? At the moment this change will be visible everywhere, which violates the 'wars are isolated' principle. If you set the war to be an isolated deployment then this will not longer be an issue, as the war will only modify its own copy of the bean.	0	0	https://issues.jboss.org/browse/CDI-129
2883	CDI-129	I would not say "vertical approach is implemented". It is probably a side effect of specialization never being properly implemented in Weld 1.x. This has other unexpected side effects. Can you try your test with alternatives?	0	0	https://issues.jboss.org/browse/CDI-129
2884	CDI-129	+1 for @ApplicationScoped == "@EARScoped" Q: If no common agreement arises, could this be configured via a property? By default @ApplicationScoped == @EARScoped but you can configure it to be @WARScoped via a property?	0	0	https://issues.jboss.org/browse/CDI-129

#	Requirement	Questions	Answered	Successfully Answered	URL
2885	CDI-129	I have quite a few questions / remarks about it (some of which are already mentioned in CDI-129, but don't have definite answers). # both @ApplicationScoped and @ModuleScoped visibility are needed (see [DELTASPIKE-335|https://issues.apache.org/jira/browse/DELTASPIKE-215]) Unlike https://issues.jboss.org/browse/CDI-129?focusedCommentId=12602519&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12602519, I don't think we should automatically share @ApplicationScoped depending on whether the bean archive is packaged in EAR or in WAR (i.e. some cdi extension could use @ModuleScoped for JSF phase listeners and @ApplicationScoped for another purpose, so automatically sharing @ApplicationScoped would result in an artificial jar split for such extensions and packaging difficulties for the end-user). # What happens if a @ApplicationScoped bean is packaged in WEB-INF/lib for EAR applications ? IMO it should result in an exception, otherwise for EAR containing multiple WARs it will lead to inconsistencies. # What happens if a @ApplicationScoped bean is packaged in WEB-INF/lib for WAR application (no EAR) ? IMO it should be fine otherwise, we won't be able to use CDI in tomcat or with simple wars on JBoss otherwise. # What happens for EAR with multiple WARs if a CDI bean with same name is packaged in WEB-INF/lib of each wars ? With JBoss 7.1.x this leads to DeploymentException: WELD-001414 Bean name is ambiguous. This means that CDI bean archives must absolutely be packaged at the EAR level on JBoss 7.1.x (and for JSF webapp, JSF libs must be packaged in WEB-INF/lib in order for JSF to scan faces-config.xml - this leads to some difficulties if I package JSF artifacts and CDI beans in the same archive). This leads to the question : can CDI bean archives be packaged in WEB-INF/lib and still be portable (not in Java EE 6) ? IMO, it shouldn't result in an exception. # Same question if the a bean of the same class is packaged in WEB-INF/lib ? IMO, it shouldn't result in an exception. # What happens if a @ApplicationScoped bean of an EAR is Specialized or has an Alternative in a WEB-INF/lib ? Discussed previously : https://issues.jboss.org/browse/CDI-129?focusedCommentId=12602718&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12602718 IMO, CDI impl should raise an exception (if we agree for 2. that @ApplicationScoped beans can only be defined at the EAR level for an EAR).	1	0	https://issues.jboss.org/browse/CDI-129
2886	CDI-128	{quote} In CDI, the meta-data will likely exist in an annotation, but it could exist in an XML file (Candi, and Seam XML Extension for CDI). {quote} The same applies for EJB3, by the way. EJB3 can be completely driven by deployment descriptor (ejb-jar.xml) or annotations or a combination of both. {quote} For example, I am working on creating a standard interceptor for JCache 107. I can read the annotation from the getMethod of the InvocationContext, {quote} Note that the getMethod will return null for lifecycle interceptors. So if the annotation is on a lifecycle method of the bean, then this won't work either. {quote} I propose we have an extension interface that extends InvocationContext called CDIInvocationContext that has a getAnnotated. This way if someone annotates in an XML file, then it is available to implementors for interceptors. {quote} Sorry, I don't fully understand that last part. Could you please elaborate on what annotation in a XML file means? Maybe I am misunderstanding this feature request, but, I am not too sure if we need this CDI specific invocation context. For example, in EJB3 land too, the intercepted method can have various spec specified annotations on it and at the same time each of those annotations has a corresponding deployment descriptor based XML equivalent.	1	0	https://issues.jboss.org/browse/CDI-128
2887	CDI-127	h1. Notes Enums are somewhat different to normal classes, as Java guarantees that any instances are singletons. Therefore we cannot treat them as classes. Furthermore, do we want to support them as beans, allowing them to be injected? Differences are: * No callback on instantiation, when do we inject? * Shared between deployments as singletons, how do we handle dependent scoped injections? * We do not control instantiation, cannot perform constructor injection h1. Proposal We should support enums as non-contextual objects, performing field and method injection only. Enums should be assumed to be injected only after the application is fully initialized for portability. If an enum is on a shared classloader, any injections should correctly resolve to the deployment of the calling context. In other words, if a shared library or installed library has an enum A: {code} enum A { FOO, BAR; @Inject B b; String getName() { return b.name; } } {code} where B is a class defined in the shared library {code} @ApplicationScoped class B { String name; } {code} Then assuming that the name has been set to Pete in deployment 1, and Marius in deployment 2, when the enum A is used in deployment 1, the method getName() should return Pete, and in deployment 2 it should return Marius.	1	1	https://issues.jboss.org/browse/CDI-127
2888	CDI-123	What's the status? I'm tempted to set this to "won't fix". I've now used CDI in lots of real world projects and customers for more than 2 years and I know no single user who actually used it at the end of the day. There were 2 customers who moved from Spring to CDI and originally intended to use it. But both ended up either using the Spring-CDI bridge or migrating parts to native annotation based CDI because re-doing all the configuration as CDI-xml was way too much work. Much more than just moving the Services to plain CDI.	1	0	https://issues.jboss.org/browse/CDI-123
2889	CDI-121	I'm not sure how you can handle the "downgrade" to "per-bean-method-invocation" scope without requiring that the bean have an interceptor. Without an interceptor to let you know when the invocation is finished, how will you know when to destroy the bean?	0	0	https://issues.jboss.org/browse/CDI-121
2890	CDI-121	What do we think the status of this is? I am interested in moving this forward. I think this is particularly important now with JMS 2 in early draft and Java EE 7 introducing @javax.transaction.Transactional separate and apart from EJB (applaud!).	1	1	https://issues.jboss.org/browse/CDI-121
2891	CDI-121	We should make a detailed list of which Annotations/Interfaces we need for implementing the support in CDI. Imo those are * @Transactional * @TransactionAttribute * TransactionAttributeType * @TransactionScoped anything missing? Would those all fit to javax.transaction? Is there any use-case where we like to use @TransactionScoped etc _without_ the need for the rest of the javax.transaction stuff? In this case we would need an own package/spec for it, isn't? Linda, I guess you can answer this best I fear that @TransactionScoped would clash because it requires the @NormalScope meta-annotation from the CDI spec api :/ xCan we make it just a JSR-330 @Scope and define in the CDI-1.1 spec that it must be handled like a non-passivating NormalScope?	0	0	https://issues.jboss.org/browse/CDI-121
2892	CDI-120	Mark, would you be able to take a stab at this issue?	0	0	https://issues.jboss.org/browse/CDI-120
2893	CDI-120	Do we have an update for this? I'm sure Ken would also be interested in this.	0	0	https://issues.jboss.org/browse/CDI-120
2894	CDI-117	This is very useful for non-servlet environment (plain JavaSE, or frameworks like Play! that does not use Servlet). Wouldn't it be better if the spec specified a "BeanManagerFactory" factory class that could return the proper BeanManager regardless of environment? Implementation would check JNDI, then Servlet Context, then fallback on singleton or custom lookup service. CDI is just too useful to be constrained on JEE/Servlet only environments.	1	0	https://issues.jboss.org/browse/CDI-117
2895	CDI-114	Do you mean generally adding support for dynamic languages such as Groovy?	0	0	https://issues.jboss.org/browse/CDI-114

#	Requirement	Questions	Answered	Successfully Answered	URL
2896	CDI-114	Hi Vincent! Maybe we need to clarify further what we both understand under the term 'dynamic Bean' What I meant is: dynamically adding Bean<T> metainformation to the container at runtime. E.g. disabling a class or adding a Bean<T> under load which would completely change the behaviour of the container. I think this is very dangerous. What will perfectly work already is the following: Have a Bean<T> which gets parsed and scanned on startup and which behaves differently in Contextual<T>#create() depending on the situation and it might even change it's behaviour at runtime (to some degree). Currently it will be a bad idea to e.g. change the EL name of a bean at runtime. This will produce crap if the EL-Resolver caches the name-> bean mapping for example. We could introduce a Bean<T>#isDynamic() to prevent the container from applying caching in CDI-1.1. But a fundamental design decision as of today is that the rules which lead to finding that Bean<T> do not change. Maybe you could explain what kind of 'dynamics' you have in XWiki?	1	0	https://issues.jboss.org/browse/CDI-114
2897	CDI-114	Btw, OSGi is a perfect sample why to NOT allow _all_ things to get changed dynamically. How often did you manage to install/update a new plugin without having to restart Eclipse? Most of the times you update Eclipse 10 times and it even doesn't start anymore...	1	0	https://issues.jboss.org/browse/CDI-114
2898	CDI-114	Hi Vince! Will need to sleep over it, but currently this looks more like a 'hot redeploy' than a real 'dynamic Bean' change. The problem with allowing to add dynamic Beans is that they e.g. can also 'disable' other Beans, change Interceptors, Decorators, etc. So by adding one Bean, you might in the worst case change the behaviour of all the application. Another area which we need to clarify: what happens if a definition error occurs as result of adding such a new bean? It would also be very hard (or close to almost impossible) to guarantee the behaviour of previously serialized beans. I'm not saying that it's impossible, but we must really think about the consequences and how to cope with them.	1	0	https://issues.jboss.org/browse/CDI-114
2899	CDI-114	Hi guys. I was wondering if there was now a recommended way of supporting the main use case I listed above back in 2011, namely: If you have a CDI application, how can you add support for installing extensions at runtime for this application (for example by adding JARs containing beans to the classloader), and tell the Bean Manager to take into account those news beans? I suspect this is a common need for apps as they grow (i.e. the ability to dynamically add extensions). As I mentioned back then, we've been supporting this in XWiki for about 10 years now with our custom handmade solution but I hate using custom solutions when there exists standard solutions so I'm always on the lookout for such solutions. I've been wanting to try to use CDI for years now but have dropped it in 2011 because of the lack of a solution. The only option I've heard so far would be to tell the Bean Manager to "restart" and thus re-run the process of creating/resolving the Bean graph. Is that still the only solution? However, I'm worried that this would take too long for our needs since XWiki has thousands (possibly tens of thousands) of Beans running, coming from hundreds of JAR files. Is this "restart"/"reload" already implemented and optimized so that there's no rescanning of JAR files for example? Does anyone have any experience with this? Thanks a lot!	1	0	https://issues.jboss.org/browse/CDI-114
2900	CDI-114	What is your exact use case? Usually the bean graph and Classpath per application is pretty much static once booted. Except during development. If it's just for development then you might look at https://github.com/HotswapProjects	1	0	https://issues.jboss.org/browse/CDI-114
2901	CDI-114	[~atijms] That is totally up to the container. It's simply not defined, so you cannot be sure. [~vmassol] I see. So you allow that your users upload some jar which you then dynamically add to you Classpath? Or how does the XWiki container become aware of that new functionality? Does XWiki use OSGi (which would be suited for such a dynamic reload)? Or do you roll your own ClassLoader system?	1	0	https://issues.jboss.org/browse/CDI-114
2902	CDI-114	{quote}So you allow that your users upload some jar which you then dynamically add to you Classpath? Or how does the XWiki container become aware of that new functionality? Does XWiki use OSGi (which would be suited for such a dynamic reload)? Or do you roll your own ClassLoader system?{quote} Yes admin users (or any user having the right permission) can install JAR extensions. XWiki doesn't use OSGI. We asked ourselves the question initially but it was way too compex and heavy so we decided that we wouldn't be able to have several versions of the same class with different versions in XWiki. We have a simple custom Component system to manage this (the main class is here if you're curious: https://github.com/xwiki/xwiki-commons/blob/master/xwiki-commons-core/xwiki-commons-component/xwiki-commons-component-default/src/main/java/org/xwiki/component/embed/EmbeddableComponentManager.java). We do have our own classloader (we have one per namespace - XWiki supports multiple wikis and a namespace is for example one wiki, to provide isolation per wiki). When a JAR extension is uninstalled, we recreate the classloader, removing the JAR URL. FYI: https://github.com/xwiki/xwiki-platform/blob/master/xwiki-platform-core/xwiki-platform-classloader/xwiki-platform-classloader-xwiki/src/main/java/org/xwiki/classloader/xwiki/internal/ContextNamespaceURLClassLoader.java and and https://github.com/xwiki/xwiki-commons/blob/master/xwiki-commons-core/xwiki-commons-classloader/xwiki-commons-classloader-api/src/main/java/org/xwiki/classloader/URIClassLoader.java (this later one is just there to overcome JDK bugs/limitations).	1	0	https://issues.jboss.org/browse/CDI-114
2903	CDI-114	I see. So you have multiple 'apps' already running as each of your Wiki-namespace is essentially an own 'app' from a CDI pov, right? Would it be possible to bootstrap a new namespace instance and then just swapping out the original one with the freshly booted one? Looks much cleaner to me.	0	0	https://issues.jboss.org/browse/CDI-114
2904	CDI-114	I do that in an app (except I have a single plugin but design is reusable for that case) starting N children CDI context inheriting from the parent one (which would be xwiki webapp). Child having access to the parent they can register services etc and still use CDI. Limitation is children are isolated but it can be a nice thing too for plugins. Would it be an option? CDI 2 with its se API should enable that since the managed instance (EE/Servlet) will start several SE contexts.	1	0	https://issues.jboss.org/browse/CDI-114
2905	CDI-114	[~vmassol] don't you mix xwiki component and cdi bean? a macro is not a Bean generally until you generate a custom class at runtime which in any case doesn't need CDI no? The producer option would still work but is still overkill to handle wiki macros since there is a manager in xwiki, no?	1	0	https://issues.jboss.org/browse/CDI-114
2906	CDI-114	[~rmannibucau] If we need to keep our XWiki Component Manager then there's little interest in moving to CDI since (if I understand correctly) we would be able to use the CDI features only for core components that are registered statically at application boot time (and thus have two distinct Component models depending on whether you're coding a core extension or a non-core one). We're also reducing our number of core components to the maximum as time progresses to have the smallest possible core and all the rest are extracted into extensions. We currently have hundreds of extensions and this will grow even more in a very near future. So having 2 models is not a good idea IMO :) {quote}but producers generally means 'do it yourself'{quote} I don't understand this part. AFAIK you can use CDI constructs in producers, using injections, qualifiers, etc. Aren't there plenty of CDI apps that use producer constructs too without requiring to implement a Component Manager that would handle instantiation/injection? If the object created in the producers requires to be injected other beans, CDI doesn't provide any helper for that? Thanks	1	0	https://issues.jboss.org/browse/CDI-114

#	Requirement	Questions	Answered	Successfully Answered	URL
2907	CDI-114	Hmm, does it means xwiki is leaking cause previously injected values can be unregistered ("remove extension")? Or did you mean you never inject the component by type but always Provider<X> everywhere? If so you also do lazy validation which means you accept to start a not functional app cause of a missing component - CDI prevents it. Is that right? If so I fear you need something to "mock" until you can inject an actual component. Means you can refactor your code but not change it drastically cause of these rules.	0	0	https://issues.jboss.org/browse/CDI-114
2908	CDI-114	{quote}Hmm, does it means xwiki is leaking cause previously injected values can be unregistered ("remove extension")?{quote} Actually I don't know how we would behave if you remove at runtime a component that's been previously injected in another component. In theory yes there should be a problem.In practice in 12 years of XWiki I've never seen this happen. {quote}Or did you mean you never inject the component by type but always Provider<X> everywhere? {quote} No we don't do that systematically, although we recommend it. It's done mostly at places where it's important to be able to use the latest registered component for a given role. {quote}If so you also do lazy validation which means you accept to start a not functional app cause of a missing component - CDI prevents it. Is that right?{quote} Yes we do allow this. Thanks	1	0	https://issues.jboss.org/browse/CDI-114
2909	CDI-114	let me put this issue in a different context: (1) let us say an application is deployed with three CORE modules that are essentials and cannot be disabled or removed, each having their own named CDI beans, which may depend on other such beans from the other two core modules. this setup works perfectly as all the modules (and CDI beans, interceptors, qualifiers, decorators, etc) are always loaded , and at deployment (2) a new module is installed into this system at runtime - with the following CDI usage: inject existing CDI beans, supplied by one or more of the always present core modules, into its own injectable classes like Servlets, JSF backing beans, transaction Services, etc. i suppose this should not be an issue at all, since the new module is only trying to use existing CDI components. if for example for Events that need to be fired, the newly registered CDI components from the installed module (plugin, for clarity) can be added as EventListeners or interceptors during bean registration. the core modules shall never depend on the additional modules. it is only the CDI components from additional modules that may listen to or intercept the events and method invocations of CDI components from core modules. ---- what is the problem in supporting this feature? (3) let us say module-x and module-y are deployed as additional modules, and module-y specified dependency on module-x because one or more of the CDI component in module-y is using or listening or intercepting one or more of the CDI components from module-x. in this case - the module framework will not let the installation of module-y if module-x is not installed; similarly, the framework will not let the uninstallation of module-x without the uninstallation of module-y. so, we can always be sure that the CDI components required by module-y are always present in the application because when module-y is there, the module-x shall always be there. at no point in time, a CDI component in module-x can refer to (depend on) a CDI component of module-y. that dependency management is the responsibility of the module framework, if it fails there, it is not the problem of CDI. why can't this be supported? (4) i guess you could be thinking about a scenario where one CDI component from module-x depends on other CDI compoent of module-y while another CDI component of module-y depends on yet another CDI component of module-x.... i.e. cyclical dependency of modules... i agree that makes it much more complex, and require checking boundaries too often-- but why can't a well-defined dependency graph be supported? all that is needed is registering the CDI components (by component, i mean bean/interceptor/qualifier/producer/etc) at runtime, and broadcasting the events (including method invocation - for interceptors) to runtime registered CDI components. what is so dangerous in this? as a realistic use case : suppose a web app is deployed with some core set of JPA Entities, JSF backing beans, JSF pages, CDI Transaction beans that may offer REST services too (because EJB container do not allow dynamic adding of beans, we rely solely on CDI) . now into this system, a new module gets installed that brings in new JPA entities, JSF backing beans, JSF pages, its own CDI transactional components, etc. now, any of the CDI components (including backing beans) may like to listen for or depend or intercept any of CDI components deployed at startup (or installed before). why can't this be achieved since the previously installed CDI components can never depend on the newly installed CDI components? what problems do you see in this? in fact, we can easily uninstall this newly installed module and remove all of its CDI components with out any impact to the core web app system that was deployed initially. — all that is needed is to track what came in new, and which new ones are listening to or intercepting around which old ones. refusing to implement this feature because improper usage may cause unpredictable errors is like saying synchronized key word cannot be supported because its improper use leads to deadlocks in multi-threaded systems.	0	0	https://issues.jboss.org/browse/CDI-114
2910	CDI-113	>> However, these cases are well specified. They are well specified for runtime but not for tools. How should tools recognize BeanManager implementations? What if there are a few classes that implement this interface? Only particular CDI container knows what to load to provide spec required services. And one more problem. CDI container don't have to look in jars without beans.xml at all. So no beans.xml no any need to scan such jars. But for tools it's not the case. Tools have to look in all the jars in class path if they want to find all BeanManager implementations. Otherwise they need to know all the available CDI-implementations to work around this problem using hard-coded names/paths.	1	1	https://issues.jboss.org/browse/CDI-113
2911	CDI-113	bq. Nothing limits CDI container to have a few implementation in class path but use the only one (at least at once). bq. For runtime it's not a problem. You know that you must provide a BeanManager and you provide it. But what if the developer added the following class in the classpath (I don't know why he might do it): bq. public class MyNotUsedBeanManagerImplimentation implements BeanManager { ... } bq. Runtime still works since this dead code is not used by CDI container but how tools can know that there is some difference between those BeanManager implementations? bq. Agree that may be this example is too tricky. And for 99.9% there will be the only implementation for every required service. But I just wanted to show that what is obvious for runtime is not so for design time. Alexey, can you give example of why tools needs to know this - how would it represent this info to the user? bq. Second problem may be more important. Even if we know for sure that there is the only one BeanManager implementation, the only one UserTransaction implementation, etc. we should scan every jar in the class path since we don't have any marker like beans.xml to skip not-CDI jars. It may be too long I don't understand this problem, why do you need to do this?	1	1	https://issues.jboss.org/browse/CDI-113
2912	CDI-113	bq. Alexey, can you give example of why tools needs to know this - how would it represent this info to the user? @Inject BeanManager manager; JBoss Tools used to complain that no bean is eligible for "@Inject BeanManager manager;" until we hardcoded BeanManagerImpl (for Weld) as a special bean. bq. I don't understand this problem, why do you need to do this? How we can find actual bean classes without scanning the entire class path? So far we just hardcoded weld-*.jars name template and scan such jars along with jars which contain beans.xml.	1	1	https://issues.jboss.org/browse/CDI-113

#	Requirement	Questions	Answered	Successfully Answered	URL
2913	CDI-113	{quote} {quote}Alexey, can you give example of why tools needs to know this - how would it represent this info to the user?{quote} @Inject BeanManager manager; JBoss Tools used to complain that no bean is eligible for "@Inject BeanManager manager;" until we hardcoded BeanManagerImpl (for Weld) as a special bean. {quote} I don't think you should be making BeanManagerImpl available for injection, just BeanManager (which all impls *must* provide). So hardcode BeanManager, not BeanManagerImpl. {quote}{quote}I don't understand this problem, why do you need to do this?{quote} How we can find actual bean classes without scanning the entire class path? So far we just hardcoded weld-*.jars name template and scan such jars along with jars which contain beans.xml.{quote} Why do you need to find actual bean classes? All that you should be making available is the interfaces provided by the spec. I think we need to talk about this offline as there is clearly a misunderstanding here, which I think may stem from the way we did things in Seam 2.	1	1	https://issues.jboss.org/browse/CDI-113
2914	CDI-112	Not sure what ordering has to do with this? What would ordering enablements achieve? I think we should just fold this into CDI-18.	1	1	https://issues.jboss.org/browse/CDI-112
2915	CDI-112	RE: The <first> <others> turned out to be not flexible enough for lots of situations. Thus I prefer an ordinal with a float value. Hmmm.. Ok.	1	0	https://issues.jboss.org/browse/CDI-112
2916	CDI-110	Then what would be the difference between a ServiceHandler and an Interceptor ? As stated before, the goal of a ServiceHandler is to provide an implementation to an abstract class or an interface. I believe that allowing multiple ServiceHandler(s) would lead to confusion, as when to create an Interceptor and when to create a ServiceHandler. Thoughts welcome	1	1	https://issues.jboss.org/browse/CDI-110
2917	CDI-110	Maybe we should require Service Handler implementations to implement the InvocationHandler interface. WDYT ? See: http://docs.oracle.com/javase/7/docs/api/java/lang/reflect/InvocationHandler.html Or even better: Rename ServiceHandler to InvocationHandler feature	0	0	https://issues.jboss.org/browse/CDI-110
2918	CDI-110	Could work, do you want to explore that in the spec changes?	0	0	https://issues.jboss.org/browse/CDI-110
2919	CDI-109	Can we expand on this? I am not sure I understand. Is there some supporting document blog, email thread? This is not clear.	0	0	https://issues.jboss.org/browse/CDI-109
2920	CDI-108	same here, no clue, the JSR-330 EG list is quiet. Where is this info from?	0	0	https://issues.jboss.org/browse/CDI-108
2921	CDI-106	More details. I like the ideas of JMS support for CDI. But where should it live. Can we release extensions like JMS in a separate jar file? Or somehow make it modular so as not to preclude CDI from being used in non-Java EE environments. Can we spin off other specifications for JMS CDI Extension? Should a CDI JMS extention be part of the JMS spec.?	0	0	https://issues.jboss.org/browse/CDI-106
2922	CDI-105	Why can't the extension just load the xml file from classpath? Can you state a use case for this?	0	0	https://issues.jboss.org/browse/CDI-105
2923	CDI-105	Will this allow access to the beans.xml at a time when it is possible to add beans and annotated types CDI?	1	1	https://issues.jboss.org/browse/CDI-105
2924	CDI-104	From my understanding, this need has been addressed with [{{UnmanagedInstance}}|http://docs.jboss.org/cdi/api/1.1/javax/enterprise/inject/spi/Unmanaged.UnmanagedInstance.html] as documented in [Obtaining non-contextual instance|http://docs.jboss.org/cdi/spec/1.2/cdi-spec.html#bm_obtain_unmanaged_instance] in CDI 1.1. Isn't it?	1	1	https://issues.jboss.org/browse/CDI-104
2925	CDI-102	Is this a Spec feature or just an extra Weld feature? I think this is bad either way. We are introducing more coupling between Weld and Servlet. The relationship here should be inverted IMO. This is causing bugs with apps that already wrote producers for Http**** artifacts.	0	0	https://issues.jboss.org/browse/CDI-102
2926	CDI-102	This is a spec feature. Weld is already coupled to Servlet, so I'm not concerned about that. But the point about conflicting with existing extensions is relevant. Do you want to reopen?	0	0	https://issues.jboss.org/browse/CDI-102
2927	CDI-102	+1 for conflict, couldn't it be a kind of alternative to avoid to activate it when an extension is already here? (the feature is good btw but out of the box will break too much IMO)	0	0	https://issues.jboss.org/browse/CDI-102
2928	CDI-102	org.jboss.weld.exceptions.DeploymentException: WELD-001409 Ambiguous dependencies for type [HttpServletRequest] with qualifiers [@Default] at injection point [[UnbackedAnnotatedField] @Inject private org.jboss.solder.servlet.http.HeaderParamProducer.request]. Possible dependencies [[org.jboss.weld.bean-test.jar-Built-in-org.jboss.weld.bean.builtin.ee.HttpServletRequestBean, Producer Method [HttpServletRequest] with qualifiers [@Any @Default] declared as [[UnbackedAnnotatedMethod] @Produces @Typed @RequestScoped protected org.jboss.solder.servlet.http.ImplicitHttpServletObjectsProducer.getHttpServletRequest()]]] at org.jboss.weld.bootstrap.Validator.validateInjectionPointForDeploymentProblems(Validator.java:368) at org.jboss.weld.bootstrap.Validator.validateInjectionPoint(Validator.java:297) at org.jboss.weld.bootstrap.Validator.validateGeneralBean(Validator.java:157) at org.jboss.weld.bootstrap.Validator.validateRIBean(Validator.java:184) at org.jboss.weld.bootstrap.Validator.validateBean(Validator.java:470) at org.jboss.weld.bootstrap.ConcurrentValidator$1.doWork(ConcurrentValidator.java:74) at org.jboss.weld.bootstrap.ConcurrentValidator$1.doWork(ConcurrentValidator.java:72) at org.jboss.weld.executor.IterativeWorkerTaskFactory$1.call(IterativeWorkerTaskFactory.java:60) at org.jboss.weld.executor.IterativeWorkerTaskFactory$1.call(IterativeWorkerTaskFactory.java:53) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:334) at java.util.concurrent.FutureTask.run(FutureTask.java:166) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:603) at java.lang.Thread.run(Thread.java:722) Exception 0 : org.jboss.weld.exceptions.DeploymentException: WELD-001409 Ambiguous dependencies for type [HttpServletRequest] with qualifiers [@Default] at injection point [[UnbackedAnnotatedField] @Inject private org.jboss.solder.servlet.http.RedirectBuilderImpl.request]. Possible dependencies [[org.jboss.weld.bean-test.jar-Built-in-org.jboss.weld.bean.builtin.ee.HttpServletRequestBean, Producer Method [HttpServletRequest] with qualifiers [@Any @Default] declared as [[UnbackedAnnotatedMethod] @Produces @Typed @RequestScoped protected org.jboss.solder.servlet.http.ImplicitHttpServletObjectsProducer.getHttpServletRequest()]]] at org.jboss.weld.bootstrap.Validator.validateInjectionPointForDeploymentProblems(Validator.java:368) at org.jboss.weld.bootstrap.Validator.validateInjectionPoint(Validator.java:297) at org.jboss.weld.bootstrap.Validator.validateGeneralBean(Validator.java:157) at org.jboss.weld.bootstrap.Validator.validateRIBean(Validator.java:184) at org.jboss.weld.bootstrap.Validator.validateBean(Validator.java:470) at org.jboss.weld.bootstrap.ConcurrentValidator$1.doWork(ConcurrentValidator.java:74) at org.jboss.weld.bootstrap.ConcurrentValidator$1.doWork(ConcurrentValidator.java:72) at org.jboss.weld.executor.IterativeWorkerTaskFactory$1.call(IterativeWorkerTaskFactory.java:60) at org.jboss.weld.executor.IterativeWorkerTaskFactory$1.call(IterativeWorkerTaskFactory.java:53) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:334) at java.util.concurrent.FutureTask.run(FutureTask.java:166) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) at	0	0	https://issues.jboss.org/browse/CDI-102

#	Requirement	Questions	Answered	Successfully Answered	URL
2929	CDI-102	Any solution to this issue? I got error when deploy application with Solder-impl.3.2.0.Final: Exception 0 : org.jboss.weld.exceptions.DeploymentException: WELD-001409: Ambiguous dependencies for type HttpServletRequest with qualifiers @Default at injection point [BackedAnnotatedField] @Inject private org.jboss.solder.servlet.http.CookieParamProducer.request at org.jboss.solder.servlet.http.CookieParamProducer.request(CookieParamProducer.java:0) Possible dependencies: - WELD|AbstractSyntheticBean|epen-login-1.4.30-SNAPSHOT.war/WEB-INF/lib/solder-impl-3.2.0.Final.jar|HttpServletRequest, - Producer Method [HttpServletRequest] with qualifiers [@Any @Default] declared as [[BackedAnnotatedMethod] @Produces @Typed @RequestScoped protected org.jboss.solder.servlet.http.ImplicitHttpServletObjectsProducer.getHttpServletRequest()] at org.jboss.weld.bootstrap.Validator.validateInjectionPointForDeploymentProblems(Validator.java:369) at org.jboss.weld.bootstrap.Validator.validateInjectionPoint(Validator.java:282) at org.jboss.weld.bootstrap.Validator.validateGeneralBean(Validator.java:133) at org.jboss.weld.bootstrap.Validator.validateRIBean(Validator.java:164) at org.jboss.weld.bootstrap.Validator.validateBean(Validator.java:507) at org.jboss.weld.bootstrap.ConcurrentValidator$1.doWork(ConcurrentValidator.java:68) at org.jboss.weld.bootstrap.ConcurrentValidator$1.doWork(ConcurrentValidator.java:66) at org.jboss.weld.executor.IterativeWorkerTaskFactory$1.call(IterativeWorkerTaskFactory.java:60) at org.jboss.weld.executor.IterativeWorkerTaskFactory$1.call(IterativeWorkerTaskFactory.java:53) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:334) at java.util.concurrent.FutureTask.run(FutureTask.java:166) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:603) at java.lang.Thread.run(Thread.java:722)	0	0	https://issues.jboss.org/browse/CDI-102
2930	CDI-99	The spec covers what happens if there is no instance for a given class but the result may as well be ambiguous. Should we specify what happens in that case?	1	1	https://issues.jboss.org/browse/CDI-99
2931	CDI-99	Don't we fear about security issues? Giving an end-user access to all runtime Extensions might be a bit too much...	1	1	https://issues.jboss.org/browse/CDI-99
2932	CDI-99	while chatting with jozef and ales, we had the following idea: What about just allowing to inject Extensions at bootup (as we already do with BeanManager)? public doThaTrick(@Observes BeforeBeanDiscovery bbd, MyOtherExtension otherExtension) {...}; or even: public doThaTrick(@Observes BeforeBeanDiscovery bbd, Instance<Extension> allExtensions) {...}; We would need to clarify 6.7.3 that Extensions, BeanManager, Event, ...(what else being @ApplicationScoped and available at bootup?) must be available for injection into SystemEvent methods.	1	1	https://issues.jboss.org/browse/CDI-99
2933	CDI-99	Jozef, how can this be ambiguous? There is one extension per extension class.	1	1	https://issues.jboss.org/browse/CDI-99
2934	CDI-92	It is hard to even consider this without some further explanation. Can you clarify? I went to the link. This needs more discussion.	0	0	https://issues.jboss.org/browse/CDI-92
2935	CDI-92	Reopening, since multiple implications are not handled by this change. 1.) The qualifiers (@Intercepted, @Decorated) are missing from the API and from the specification document. 2.) Similarly to the Injection point metadata section, the Bean metadata section should name possible invalid scenarios and required the container to treat them as a definition error. These include: * Interceptor<?> injection point in a non-interceptor bean * @Intercepted Bean<?> injection point in a non-interceptor bean * Decorator<?> injection point in a non-decorator bean * @Decorated Bean<?> injection point in a non-decorator bean 3.) The spec should define what happens in the following scenatio: {code:JAVA} class Foo { @Inject Bean<Bar> metadata; } {code} Should the container validate the generic type of a bean metadata injection point? Should the abovementioned example be treated as a definiton error? 4.) The spec should perhaps explicitly state that it is possible to inject bean metadata into a disposer method and that this metadata always refers to the producer field/method that created the instance being destroyed.	0	0	https://issues.jboss.org/browse/CDI-92
2936	CDI-86	Jozef, could you review the changes I've made and check you feel it sufficiently specifies your scenario?	1	1	https://issues.jboss.org/browse/CDI-86
2937	CDI-86	John, take a look at section 6.7.1 in the 1.0 spec and you will see that CDI requires the request scope is active for JMS message listener message delivery. If the request context is initialized specifically for this, then we will will fire a CDI event with qualifiers @Request @Initialized, and the event should have a payload (in servlet it's a ServletRequestEvent which allows the observer access to the context of the request). What should it be for JMS message listeners, if anything?	1	1	https://issues.jboss.org/browse/CDI-86
2938	CDI-86	Hmmm I missed that in 1.0. Do you mean MDB's or any message listener within a CDI enabled JMS object? E.g. a message listener created that is bound using addMessageListener on a consumer? Also, would you be referring to a CDI enabled message listener (e.g. @Inject @MessageType("A") MessageListener ml;) ? This one I know doesn't work in AS7.	0	0	https://issues.jboss.org/browse/CDI-86
2939	CDI-86	I have updated the pull req with Jozef's suggestion. Jozef, can you spin the idea about extending this to non-built-in scopes to a new issue please?	0	0	https://issues.jboss.org/browse/CDI-86
2940	CDI-86	{quote}Jozef, can you spin the idea about extending this to non-built-in scopes to a new issue please?{quote} CDI-165	1	1	https://issues.jboss.org/browse/CDI-86
2941	CDI-85	What happened to this issue? The pull request is not available any more and the wording in the spec still seems to be wrong.	1	1	https://issues.jboss.org/browse/CDI-85
2942	CDI-82	Marius, where is this defined as an error?	1	1	https://issues.jboss.org/browse/CDI-82
2943	CDI-80	Why is the conversation propagation limited to *GET* request parameter only? If I got it right this rule prevents form-processing servlet-only applications from using the conversation scope.	1	1	https://issues.jboss.org/browse/CDI-80
2944	CDI-78	[~antoinesabot-durand] This is still unclear to me what should happen when I inject {{InjectionPoint}} in decorator. Section {{6.5.4. Contextual reference for a bean}} states: {quote} The container must ensure that every injection point of type InjectionPoint and qualifier @Default of any dependent object instantiated during this process receives: • an instance of InjectionPoint representing the injection point into which the dependent object will be injected, or • a null value if it is not being injected into any injection point. {quote} Will I get null or instance of {{InjectionPoint}} for the IP injected into decorator?	1	1	https://issues.jboss.org/browse/CDI-78

#	Requirement	Questions	Answered	Successfully Answered	URL
2945	CDI-78	2 scenarios here. h2. InjectionPoint relates to the decorator bean itself As a decorator can't be injected directly since it's not available for injection as stated in 5.1.4, so {{InjectionPoint}} of a decorator should always return null. h2. InjectionPoint relates to the decorated bean It could be nice to have information about {{InjectionPoint}} of the decorated bean in the decorator, but it can be confusing to deal with it: decorator can decorate dependent and non-dependent beans at the same time. So, if we don't want to break the spec, we should forbid it if at least one decorated bean is not dependent. h2. My proposal {{InjectionPoint}} with @Default qualifier injected in decorator (or interceptor BTW), should be considered as scenario 1 and always return null. To support injection of decorated bean's {{InjectionPoint}} we should use the {{@Decorated}} qualifier ({{@Inject @Decorated InjectionPoint}}) allowing to have the IP instance if decorated bean is dependent or null if it is not. It won't violate the spec and be less confusing for users (we can assume that a user adding such a qualifier probably knows what he's doing). Again this proposal is also interesting for Interceptors. WDYT ?	0	0	https://issues.jboss.org/browse/CDI-78
2946	CDI-77	Hmm, if we don't talk about constructor injection, this is not a problem at all. So I don't think, we should forbid this. And for constructor injection we could create a proxy. We already create proxies for Dependent beans (i.e. when they have interceptor bindings). So why should we not create a proxy in this situation?	1	0	https://issues.jboss.org/browse/CDI-77
2947	CDI-77	Hmm, but why (but of a technical point of view) should the spec differentiate between normal scoped beans and dependent beans? Either we feel that circular injection is not a good way to design applications and want to forbit it in either situation or we allow it and maybe document behavior where it simply does not work. By the way, how does the spec treat situations where a dependent bean has an interceptor-binding or decorator and is not proxiable?	1	0	https://issues.jboss.org/browse/CDI-77
2948	CDI-75	From my mail to cdi-dev: Initially I thought it would be a good idea to ignore an observer method if a scope is not active and notifyObserver is IF_EXISTS. But now I believe this is partially wrong. There is a difference between A) the observer bean scope is active without an active bean in scope and B) the observer bean scope is not active IF_EXISTS refers to A) only, I believe. Because if IF_EXISTS applies to A) and B) I am unable to support ALWAYS for active scopes while being ignored if a scope is not active. Why would someone need that? Very often we initialize our application during startup. During startup request and application scope are active. Usually a session scoped observer should be ignored during startup. But at runtime - with an active session scope - I need ALWAYS semantics for those observers. So I would vote for: - ALWAYS observer methods will throw an ContextNotActiveException(?) if scope is not active. - IF_EXISTS observer methods will throw an ContextNotActiveException(?) if scope is not active. - IF_EXISTS observer methods will be ignored if scope is active but no active bean is in scope. - (The newly created) IF_ACTIVE observer method (naming TBD) will be ignored if scope is inactive. If scope is active it will be treated as ALWAYS. - There might be a need for IF_ACTIVE_AND_EXISTS;). This might break backwards compatibility though.	1	1	https://issues.jboss.org/browse/CDI-75
2949	CDI-75	From the CDI 1.1 EG meeting, 19th Sept {quote} Mark argued that this isn't 100% clear either way - it's arguable both ways. All agreed that it shouldn't be thrown by default. Will clarify that it isn't thrown, and possibly add a new state in which it is thrown (e.g. Reception.Strict or something somewhat less horribly named ;-) {quote} Jens, WDYT? Do you need this new state?	1	1	https://issues.jboss.org/browse/CDI-75
2950	CDI-74	Marius, I think we need to specify effects, not impl. Can you do this?	0	0	https://issues.jboss.org/browse/CDI-74
2951	CDI-70	Where is your problem here? Should we add a second part to the second sentence like "The container must not directly call the Java Reflection API to discover them." We are just talking of discovery here and that perfectly fits with the javadoc. We are able to discover that the developer has not added any constructor to a class by retrieving an empty set from javax.enterprise.inject.spi.AnnotatedType.getConstructors().	0	0	https://issues.jboss.org/browse/CDI-70
2952	CDI-58	Basically ok, but don't we need an information about it's ID in the AnnotatedType itself? That could be useful for Extensions, isn't? Otherwise an Extension like Seam-XML could not differentiate between the built-in and the external AnnotatedType.	0	0	https://issues.jboss.org/browse/CDI-58
2953	CDI-51	Do we want to do this? What are the use cases?	0	0	https://issues.jboss.org/browse/CDI-51
2954	CDI-50	To be honest: I for one am not so fond about @Veto and @Requires yet ... @Veto: completely underspecified! * How can it be modified via Extensions? * In which phase does it get applied? Directly at startup? Or only after ProcessAnnotatedType? * all corner cases and exceptional behaviour, e.g. how does it fit together with @New? @Requires: not user friendly * against JSR-250 * adds tons of NoClassDefFound pitfalls * how to prevent @Requires("wrongString") People should rather modularize their apps a bit better.	0	0	https://issues.jboss.org/browse/CDI-50
2955	CDI-50	Our preface still mentions @Requires. Plus it's not clear in 11.5.6 what happens if there are package-info.java @Vetoed for the same package in different JARs. Do we need to scan all jars for those first and merge the packages upfront? That could slow down scanning. It' also not clear from the Java Lang Spec what happens if you don't use package-info.java but annotate the package in a single class file. package-info.java is recommended, but the spec allows other ways as well: http://docs.oracle.com/javase/specs/jls/se5.0/html/packages.html#7.4.1.1	0	0	https://issues.jboss.org/browse/CDI-50
2956	CDI-45	From my understanding, that need is met with [{{javax.enterprise.inject.Instance}}|http://docs.oracle.com/javaee/7/api/javax/enterprise/inject/Instance.html]: {code} @Inject Instance<MyBean> bean; if (!bean.isUnsatisfied()) bean.get(); {code} What would be the added value of the {{@Optional}} approach compared to the {{Instance}} concept? Besides, that approach would lead to having potential NPE or to check against nullity which is not very elegant. That being said, with the CDI 2.0 / Java 8 enhancement stream, [{{java.util.Optional}}|http://docs.oracle.com/javase/8/docs/api/java/util/Optional.html] could be automatically supported: {code} @Inject Optional<MyBean> bean; bean.isPresent(); bean.orElse(...); {code}	0	0	https://issues.jboss.org/browse/CDI-45
2957	CDI-45	I think there could be an interface between Provider<T> and Instance<T>, which will have method isUnsatisfied(). So that way backward compatibility will be preserved. The question is not only of how we declare, but also if the dependecy is ambiguous, should resolution fail for "optional" dependecy?	0	0	https://issues.jboss.org/browse/CDI-45
2958	CDI-45	Then one potential compromise would be to enrich the existing abstractions, either {{Instance}} or {{Provider}} with the equivalent semantic that of {{Optional}}. For example: {code} @Inject Provider<MyBean> bean; if (bean.isPresent()) bean.get(); {code} Or if we think {{Optional}} is a convenient abstraction: {code} @Inject Provider<MyBean> bean; bean.getOptional(); {code} In any case, I think that a {{isSatisfied()}} / {{isResolvable()}} method is missing on the {{Instance}} interface when alternatives are enabled for example and that would avoid writing {{!isUnsatisfied() && !isAmbiguous()}}. As per backward compatibility for these interfaces, are they really meant to be implemented aside from the implementations? That would avoid introducing another interface. If not maybe default method in Java 8 could help.	0	0	https://issues.jboss.org/browse/CDI-45

#	Requirement	Questions	Answered	Successfully Answered	URL
2959	CDI-45	I guess there will be a way to toggle this property using a CDI-Extension? How would this look like? {code} void processInjectionPoint(@Observes ProcessInjectionPoint pit) { ... pit.getInjectionPoint().setOptional() } {code}	0	0	https://issues.jboss.org/browse/CDI-45
2960	CDI-45	How about we implement this by allowing Java 8 Optional? {code} @Inject private Optional<MyBean> myBean; {code} Would also be a nice Java 8 alignment.	1	0	https://issues.jboss.org/browse/CDI-45
2961	CDI-45	[~sknitelius] What's wrong with {{Instance}}? We're going to introduce {{Instance.isResolvable()}} within CDI-589.	1	0	https://issues.jboss.org/browse/CDI-45
2962	CDI-45	-1 for adding those methods on {{Instance}} and -1 for introducing {{Optional}} built-in bean (mainly due to serialization issues and the need for special handling, e.g. disallow programatic lookup). {{Instance.getOptional().ifPresent(...)}} etc. is imho easier to understand.	1	0	https://issues.jboss.org/browse/CDI-45
2963	CDI-45	[~meetoblivion] On F2F we decided to add {{isResolvable()}} (is satisfied and not ambiguous), it's already part of the [Weld API|https://github.com/weld/api/blob/2.4/weld/src/main/java/org/jboss/weld/inject/WeldInstance.java#L53]... I suppose this is basically what {{isPresent()}} does. Do you find the name acceptable? Also I quite like {{asOptional()}}.	0	0	https://issues.jboss.org/browse/CDI-45
2964	CDI-43	{quote} I think that should should also apply to ProcessBean, ProcessProducer etc. In this case the even is only delivered if the declaring class (i.e. Bean.getBeanClass()) meets these requirements. {quote} Can you elaborate on why? Also, if we specify this as a filter, I'm concerned that operating on the bean class is inconsistent, as a user would properly expect it to take into account modified type information from AnnotatedType.	0	0	https://issues.jboss.org/browse/CDI-43
2965	CDI-37	OK, I understand, what he means, but I don't understand why it is needed. In what way does an @ApplicationScoped bean not fulfill the mentioned requirements	1	0	https://issues.jboss.org/browse/CDI-37
2966	CDI-30	[~meetoblivion] [First PR|https://github.com/cdi-spec/cdi/pull/291/files] looks good! Would it be an idea to set a wrapped (or decorated) context controller? What I'm looking for is a way to start contexts possibly earlier than a specific container implementation normally starts them, and stop them possibly later. If the context is started earlier in say a request, then eventually the container will also start that same context later in the request. If that will throw an exception the request will not proceed, which is of course not intended. Likewise, when the container stops the context this will thus be too early. With a wrapped context controller, the {{startContexts}} and {{stopContexts}} calls can be ignored for those contexts the application (extension/library) wanted to start earlier and stop later. This does of course assume that the container itself will also use this context controller instead of the current proprietary code. Additionally, the PR does not contain the check anymore to see if the context is already active or already stopped. Is that still intended to be added? This would be particularly useful for code that needs to be portable between different containers, where one container may already have started the context at a given point, where another container has not (catching the exception would be an alternative here but is obviously not so nice). Another point, what about letting the caller of the {{startContexts}} method optionally pass in the objects (or a lambda to obtain them) on which that context primarily depends (if any). E.g. for the request scope, pass in an {{HttpServletRequest}} instance.	1	1	https://issues.jboss.org/browse/CDI-30
2967	CDI-27	Pete, does this issue mean to support EJBs @TransactionAttribute and stuff out of the box? This would require a very deep integration with the EJB container and the support of various transaction models. Not sure if we really like to dig that deep. This is workwise basically a duplication of the EJB spec in this area. But technically it would be different because in EJB transactions are auto-on whereas in CDI they would still need to get applied automatically.	1	1	https://issues.jboss.org/browse/CDI-27
2968	CDI-27	> This was not my intention. I'd rather see the CDI spec just providing the basic mechanisms to make such an Interceptor possible (which it already does). But I'd not specify such complex and history-bloated things in the CDI spec itself, but maybe in the EJB spec? Could you please clarify? If you do not intend to introduce an additional annotation, then I'm o.k. with that. You could easily re-implement EJB 3.1 transactions with CDI 1.0.	0	0	https://issues.jboss.org/browse/CDI-27
2969	CDI-26	The more I think about it, the more I'm sure that we need to push this to the EE7 umbrella spec. Booting the CDI container alone is nice, but by far not sufficient. EJB has already an API to boot the container (javax.ejb.embeddable.EJBContainer), CDI might get one too. But still the user will end up not being able to start his whole environment! Imagine a situation where a user likes to use openejb or JBossAS7-embedded in a SE environment. Starting the container alone would not be enough. Otoh one could argue that in those cases a user might easily use the containers native implementation specific classes directly. Wdyt?	0	0	https://issues.jboss.org/browse/CDI-26
2970	CDI-21	I would also favor this. Have a use case where I observe "ProcessAnnotatedType<T>" events and I don´t want to reinit a variable each time.	0	0	https://issues.jboss.org/browse/CDI-21
2971	CDI-20	Good point, Antoine. And even with ordered event observers: What would be the meaning of "propagate = false" for transactional event observers?	0	0	https://issues.jboss.org/browse/CDI-20
2972	CDI-18	Can you be more specifc Rick as to what elements you are thinking of?	1	1	https://issues.jboss.org/browse/CDI-18
2973	CDI-18	> RE: The <first> <others> turned out to be not flexible enough for lots of situations. Thus I prefer > an ordinal with a float value. Hmmm.. Ok. My comment was targeted at the mechanism used in the faces-config.xml which I assume has the same mechanism (need to re-read the servlet 3.0 spec though). Problem there is that the ordering is for the whole config file, isn't? So you could only define the order of all the stuff in the beans.xml. Not sure if we need that, but in case of the faces-config.xml I definitely regretted it already: I did want to have the EL-order in a different way than the facelets resource order which doesn't work.	1	1	https://issues.jboss.org/browse/CDI-18
2974	CDI-18	>> Folks, introducing any hierarchy is exactly the same way broken as BDA itself. Let's just get rid of that piep! I would say having that *piep* is not so *piep* at all, as it allows you to abstract away diff CL rules. Who's to say which CL rules are right. ;-) e.g. JBossAS pre-7 was doing well with that big-ball-o-mud, until that bbom become unmanageable (via JEE expansion ...). And even if we went directly to modularity / CL, you still need a mechanism to "visit" things API way. Since CL doesn't expose that, you need a custom layer -- here enter BDA. ;-)	1	1	https://issues.jboss.org/browse/CDI-18
2975	CDI-18	Ales, fact is that lots of applications I know which must get deployed to Glassfish or JBossAS6 (not sure yet about AS7) currently just *unpack* all their dependencies to WEB-INF/classes, BECAUSE BDA IS COMPLETELY FUUUU UP.... Have you ever worked with big customers to get their _big_ CDI apps working? I did, and trust me, BDA is really a nightmare! This is really the worst part of CDI-1.0, and it's _very_ easy to get rid of it! Just look at what seam-solder does with @DefaultBean to make parts of this undone with some dirty tricks... @DefaultBean disables the BDA by emulating other mechanism - but at a pretty high performance and complexity cost. And then look at how many parts of Seam are using @DefaultBean? 80%? 90%? Even more?	1	1	https://issues.jboss.org/browse/CDI-18
2976	CDI-18	Do you see problems with resolution around BDAs as well?	1	1	https://issues.jboss.org/browse/CDI-18

#	Requirement	Questions	Answered	Successfully Answered	URL
2977	CDI-18	I don't know BDA. What is it? I do like JBoss modules (but it can't change the language like Jigsaw will). Mark Reinhault's devoxx presentation about Jigsaw is really worth seeing.	0	0	https://issues.jboss.org/browse/CDI-18
2978	CDI-18	I'm not saying BDA impls are perfect :-), but I do think we need some abstraction like BDA. It's easy to say we don't need it, and let's then make everything flat. Which is - as you say - probably how big JEE apps look like, but not everything is JEE. ;-) e.g. we had a big customer with MicroContainer (AS5/6 kernel), which used custom CL rules, and having BDA would make it possible for them to use CDI, if they wanted (this was pre-CDI). btw: how do you then follow bean visibility rules in OWB in hierarchy .ears?	1	1	https://issues.jboss.org/browse/CDI-18
2979	CDI-18	>> @Ales: In fact, JBossAS7 already killed the BDA partly! You already contradict the spec by allowing definitions in WEB-INF/beans.xml to count for all jars in WEB-INF/lib. Please go back and read the spec - this is clearly defined different! As per the CDI-1.0 spec, any definitions in WEB-INF/beans.xml must ONLY affect WEB-INF/classes/** and MUST NOT have any effect on jars in WEB-INF/lib! I don't think that's true for AS7. Are you 100% sure? It was true for AS6, but that was a mistake / bug. And what do you mean by "affect"?	1	1	https://issues.jboss.org/browse/CDI-18
2980	CDI-18	The whole BDA discussion is also important for Beans added via Extensions. Such Beans do not have any BDA. Same question applies if AnnotatedType gets modified or added freshly via Extensions. What BDA are they assigned to?	1	1	https://issues.jboss.org/browse/CDI-18
2981	CDI-18	issue to allowing default/global enablement of interceptors, decorators and alternatives was to provide a global ordering solution. We discussed the three proposals: (1) Relative ordering of individual interceptors, decorators and alternatives using either a name or an ordering qualifier (2) Ordinal/magic number based ordering of individual interceptors, decorators and alternatives (3) Ordinal/magic number based ordering of lists of interceptors, decorators and alternatives (i.e. the ability to specify an order of beans.xml, which themselves specify a partial ordering of interceptors, decorators and alternatives) We discussed the merits of each. In general, we didn't feel any solution is excellent, however this is a problem we must solve, so we must pick the best, even if it isn't perfect. Pete expressed concerns over the understandability of the merge of partial orders done in (3). Consider an application which has 10 libraries, each of which specify 10 interceptors in an order. There is some overlap between the lists. For a user to comprehend the overall order is not trivial in this case, it requires a tool to tell you the order. We discussed (1) and Mark expressed concerns that you can't easily express some well know use cases using it, without explicitly specifying the interceptor by name. For example, consider a TransactionInterceptor, which needs to be "outside all". This must specify "outside all". However, any other interceptor which wants to be outside everything should specify "outside all, inside TransactionInterceptor". This gets complex quite quickly. Its a less complex algorithm to comprehend than (3), but still requires quite a bit of thought to understand the overall order. Finally, we discussed (2). We agreed that it has some good precedents (e.g. httpd, SYS-V init) that have worked well for a long time. It's trivial to comprehend, and also quick to compute. It does require a knowledge of the position of other algorithms, but so do all other approaches considered. We agreed to recommend approach (2). A lower number indicates the interceptor or decorator should appear earlier in the call stack, a higher number indicates it should appear later in the call stack. A higher number alternative takes precedence over a lower number alternative (OPEN ISSUE: how does this interact with class hierarchies?). Two interceptors, decorators and alternatives with the same number have the same precedence, and their order is not defined. We discussed that we should use recommended ranges for the magic numbers. For example: 1-100: System interceptors, decorators and alternatives (interceptors specified by CDI or the Java EE platform specs) 100 - 1000: Extension interceptors, decorators and alternatives (interceptors provided by extension libraries such as Deltaspike) 1000 - 2000: Application interceptors, decorators and alternatives (interceptors provided by an application or internal extension) 2000 - 3000: Extension interceptors, decorators and alternatives 3000 - 3100: System interceptors, decorators and alternatives Note that we need to allow system interceptors to go before (e.g transaction) or after (e.g. validation). Using this "pyramid" of ranges allows for this. Next, we discussed that this needs to be overridable by the SPI, and that the SPI should present a merged view of all interceptors, decorators and alternatives. This should be delivered once, for the whole application. The extension can mutate the list, to change which interceptors, decorators and alternatives are enabled, and the ordering. We looked at the events, and agreed that deferring it as late as possible is good, to allow extensions to have as much info as possible. AfterBeanDiscovery is as late as it can go, as we need this info to perform validation. Next, we discussed where the ordinal number can be placed, and said that: * A bean archive which declares an interceptor (alternative or decorator) can set the ordinal, by adding an entry in beans.xml like: <beans> <interceptors> <class priority="1001">com.acme.FooInterceptor</class> </interceptors> </beans> OPEN ISSUE: Is priority the right attribute? * The application can override the ordering, by specifying the interceptor class, and providing another ordinal. It can able disable an interceptor, decorator or alternative, by adding the enabled="false" attribute. OPEN ISSUE: Can other libraries disable interceptors and change the precedence? If so, which one wins? Finally, we decided to remove the	1	1	https://issues.jboss.org/browse/CDI-18
2982	CDI-18	Maybe we define that you can use one style or the other?	1	1	https://issues.jboss.org/browse/CDI-18
2983	CDI-18	"a default priority for interceptors. I proposed 1000, the lowest priority for application interceptors." If negative numbers are allowed too, is there a specific reason not the default to 0? High priority => use positive number (for example +10). Low priority => use a negative number (for example -10). Normal priority => leave it at 0.	0	0	https://issues.jboss.org/browse/CDI-18
2984	CDI-18	Hi together, as the follow-up to the thread https://community.jboss.org/thread/200853. INTERCEPTORS The specification says that the "beans.xml" is valid per beans archive. But as a specification user I would prefer to have the validity of registered interceptors hierarchically. That means following: IN CONTAINER 1. The first level where the beans.xml file must appear is the bean archive (jar file) if and only if the classes in jar utilize CDI. Already there the first interceptors can be defined. In other words - I CAN say which interceptors I want to use per default for my interception bindings. Defined interceptors will be active only if this definition is not overridden on upper level. AND - if within the jar I declare interceptors I MUST create META-INF/beans.xml where I CAN register my interceptors. My registered interceptors will be active for all binding and not only in this particular jar only in case my jar is together with other jars in class path and my interceptors are not overridden in upper layers or in jar which use interceptor bindings itself. 2. The second and third level are the web/service/resource and enterprise applications respectively. On these levels the application assembly must be able to override the interceptors defined on lower level. Means: if I have registered interceptor for @Foo @Default in jar, the application which have this jar in class path can override registered interceptor by registering another interceptor for the same binding and qualifier in own beans.xml file. DESKTOP 1. The first level as in container. 3. The desktop applications which utilize CDI instantiating implementation directly must be able to provide top-level beans.xml to the implementation. Ddefinitions in this beans.xml will override existing interceptor registrations on class path. ---------- PROVIDERS I think it is necessary to clarify the situation with providers also. what happens so far if I have two providers of the same type in class path????	0	0	https://issues.jboss.org/browse/CDI-18
2985	CDI-18	Jozef, can you file new issues for errors in the PRDas it is now released	1	1	https://issues.jboss.org/browse/CDI-18
2986	CDI-18	{quote}Even JBossAS nor GlassFish are NOT following the strict CDI-1.1 rules right now because they are unusable!{quote} Alternatives are enabled as required by the spec (per bean archive) in JBoss AS (not sure about GF). Could you provide evidence supporting your statement?	0	0	https://issues.jboss.org/browse/CDI-18

#	Requirement	Questions	Answered	Successfully Answered	URL
2987	CDI-18	{quote} Alternatives are enabled as required by the spec (per bean archive) in JBoss AS {quote} Bean Archive is defined as each JAR being an own Bean Archive in CDI-1.0. As far as my tests showed an enabled alternative in WEB-INF/beans.xml is available in all jars in WEB-INF/lib/*.jar. I agree that this makes more sense than having each jar being an own BDA - but it's clearly against CDI-1.0 if implemented that way. Can you verify this please? I tested this with JBossAS-7.0.1 and 7.1.0 back then afair.	1	1	https://issues.jboss.org/browse/CDI-18
2988	CDI-16	Pete: do you know what happened to http://java.net/projects/javaee-spec/lists/jsr342-experts/archive/2012-03/message/74 ? At some point I stopped following. I personally would hope that @ManagedBean would be deprecated, and CDI would ignore this spec.	1	0	https://issues.jboss.org/browse/CDI-16
2989	CDI-16	Currently I dislike CDI treat all classes(which has a constructor with no arguments) as CDI managed bean, I know Jboss Solder provides a @Veto to resolve this issue. I think the matrix 2 is clear. I like the two classes solution. 1. the first level, the bean is injectable, has own scope. such EJB, basic Pojo(annotated with @ManagedBean), @Named etc. And other bean such as EJB can improve the lifecyle support transaction, security etc. Compare to spring framework, it provides @Component, and @Controller, @Service etc. but the later are refined from the @Component. Why the @Named, @EJB can not be refined to use @ManagedBean(JEE)? ... @ManagedBean ..... interface @Named{ } ... @ManagedBean ..... interface @Stateful{ } 2. such as entity listener, faces converter, faces validator, bean validator, servlet & servlet listener, filter, etc, can inject other beans. For example, it should allow us inject EmtityManager in faces conventer. But it can NOT be injected in other bean, and can NOT be defined in its own scope. In fact, as a developer I only want to simplify programming, but do not want to change their lifecyle. But besides the CDI interceptors, decorators, I think it should be optional, developer can select if use CDI to manage it. For the 2, is there a better solution to make it managed? Spring provides another @Configurable. In JEE 7, the jee7 expert group can consider a similar solution, for example, provides a @Managed annotation to improve the lifecyle. @WebServlet @Managed class MyServlet extends HttpServlet{ //you can use @Inject now, } if it is not annotated with @Managed...it is no relation to CDI, only a generic servlet. @WebServlet class MyServlet extends HttpServlet{ //@Inject is NOT supported, } This maybe force the different spec providers leave some space for the CDI hooks in their lifecycle. In summary, I hope all JEE specs looks smooth and seamless to developers. Regards Hantsy	1	1	https://issues.jboss.org/browse/CDI-16
2990	CDI-15	Also, the spec currently contradicts itself when it says: {quote} A qualifier type is a Java annotation defined as @Target(\{METHOD, FIELD, PARAMETER, TYPE\}) and @Retention(RUNTIME). {quote} and then uses {code:JAVA} @Qualifier @Target(PARAMETER) @Retention(RUNTIME) public @interface Role { String value(); } {code} in 10.2.2 of CDI 1.0	0	0	https://issues.jboss.org/browse/CDI-15
2991	CDI-14	To clarify this a bit: Currently we need * BeanManager#getBeans * BeanManager#resolve * BeanManager#getReference I don't think Instance<Object> instance(); is enough. It should be more like Instance<T> instance(T, Annotation<? extends Qualifier>); right? But to be honest: imo we should not make too much functionality available in the BeanManager and bloat our interface if it is easily doable already. This is really only used in situations where you cannot @Inject your Instance object (means in unmanaged instances like a JPA EntityListener or similar). But in those rare cases, it's easily doable already. Maybe we just add an own section to the spec and explain how this 3 methods work together for getting the contextual reference?	1	1	https://issues.jboss.org/browse/CDI-14
2992	CDI-14	At first I roared on the fact a convenience method wasn't available. As everyone else, I wrote a utility class with the 3 lines needed (resolve/createCreationalContext/getReference). It was actually longer to understand the difference between these and InjectionTarget#produce() because the documentation chapter 16.4 was confusing me. Then I realized you can almost always find a way to work properly, so getting bean refs by hand becomes very, very rare (eg: phaselisteners, components you really have to new() yourself). What bothered me much more was the lifecycle of such beans. I aksed myself questions like: "If the bean instance was actually instantiated via getReference, will its lifecycle be still respected?" Eventually, I thought the only thing needed would be extra documentation. I actually shared Mark's point sometimes. Now I'd say it's useless to let every EE developper write the exact same 3 lines of code, so it's better to have a couple of convenience methods (Rick's proposal sounds neat)...BUT I already see deviant behaviors coming amongst the developpers, so I'd be extremely clear on the spec, in the javadoc and in the reference doc about what getting bean refs by hand implies.	1	1	https://issues.jboss.org/browse/CDI-14
2993	CDI-14	Rick, George, why do you want to create a new API for programmatic lookup when we already have Instance? What is your motivation for adding methods to BeanManager over exposing/enhancing the existing API?	1	1	https://issues.jboss.org/browse/CDI-14
2994	CDI-14	Yes as previously commented I was saying simply that if we made it available outside of a managed object it would be via the same mechanisms that we use for all other artifacts from CDI you can look up like this (i.e. currently BM). I'm not in love with the idea that we expose Instance directly, as I think requiring extension authors to write: {code} Foo foo = beanManager.dynamic().select(Foo.class).get(); {code} (or whatever else this method get's called) is not that onerous given this is extension authors only. With hindsight I would have preferred to split up BeanManager more than was done in 1.0 but we can't got back and do that now!	1	1	https://issues.jboss.org/browse/CDI-14
2995	CDI-10	Stuart, can you clarify (perhaps with some code) what you are asking for here (I suspect your terminology may be wrong?).	1	1	https://issues.jboss.org/browse/CDI-10
2996	CDI-10	{quote}If an Extension hacker really needs such a hardcore hack, then he might as well take the direct route:{quote} Given an unknown proxy, how do you know the scope, type and qualifiers? I think the original request is stall valid, for instance for those case where you need to take action based on a bean's class (it's class name may be a key into some map for instance) but all you got is the proxy. The workaround seems a bit circular, but maybe I misunderstood. If you want to find out the original bean type of a proxy, you need the type to obtain a Bean instance via which you can obtain the original type. But if I had the type already, I wouldn't need to obtain a Bean instance in the first place. See also http://stackoverflow.com/questions/14511433/original-class-name-of-a-proxy-without-manual-string-manipulation	1	0	https://issues.jboss.org/browse/CDI-10
2997	CDI-4	One thing we could pretty easily do is to add a Float value representing an ordinal. But that really doesn't fit well to the general observer/observable pattern. What is your exact use case? From a quick think-through I don't see anything which cannot be done in a different way also.	1	1	https://issues.jboss.org/browse/CDI-4

#	Requirement	Questions	Answered	Successfully Answered	URL
2998	CDI-4	A few points of contention here: _"the @Before doesn't scale out in practice. You cannot always change your event if someone adds a jar into the classpath which needs that event"_ Can you give a solid example of this? I mean, one the reasons why I am proposing this, is that the sorting _can_ deal with the addition of new beans which have ordered observers. The solution that you're proposing -- which is to use a sort of "marathon pattern" -- seems to suffer greatly from a coupling problem, because you're relying there being a discretely defined event type for each step. That, in itself, almost entirely precludes the adding of new beans to the container that change the ordering -- at all. I'm struggling to see how my proposal introduces coupling. Having both {{@Before}} and {{@After}} is, for the very reason of dealing with potentially unseen model, very necessary. Say I have an event {{BusEvent}}. And it has qualifiers like {{@Initialize}} and {{@PostInitialize}}. If I add a jar in the future which contains a bean which has say, something like this: {code} public void beforeInit(@Observes @Before(Initialize.class) BusEvent event) { ... } {code} ... I would get _exactly_ the kind of behavior I would expect. So I fail to see the problem you're referring to. There'd be no guarantee of ordering for all the @Before(Initialize.class) events, but I somehow doubt that would be any trouble in such a use case. --- _"The service could be something like javax.enterprise.inject.spi.ObserverComparator, applied after the ProcessObserverMethod event."_ I think this is actually a really good idea. But it's really orthogonal to the problem I'm trying to solve. For the very example I point out above. It seems to me that event ordering is not something which violates decoupling of events, and is a concept which is purely a component-container contract issue. Simply having say, a framework like Forge implement comparators, doesn't help readability of the model at all. In fact, it sort of suffers from the same problem that putting the configuration in an XML file suffers from: you can't get any sense of the ordering just by looking at any given observer.	1	1	https://issues.jboss.org/browse/CDI-4
2999	CDI-4	> I'm struggling to see how my proposal introduces coupling. If I understood correctly, then your proposal relies that you have to know the other/all the other observer(s) in your 'before' Observers. If you don't call this coupling, well... *edit*: ah I see now, your @After and @Before references a qualifier and not the other obserer class. Well then there might be no direct coupling, but this most probably wont work with the currently defined qualifier pattern (depending how we fix CDI-7). > If I add a jar in the future which and what if you add another jar with another bean having another @Observes BusEvent? How to make sure your beforeInit also runs _before_ this new observer? Let's go back to the Portland Pattern Repository or the GoF Design Patterns book and check the Observer/Observable pattern. The whole idea/mechanism is built based on the paradigma that all the involved parties know NOTHNING about each other.	1	1	https://issues.jboss.org/browse/CDI-4
3000	CDI-4	This problem is very acute for a few people, but doesn't affect a majority. Jason's solution may well be the sweetspot as it allows custom solutions to be implemented. Mike, can you impl your solution on top of Jason's proposal?	0	0	https://issues.jboss.org/browse/CDI-4
3001	CDI-4	_Mike, can you impl your solution on top of Jason's proposal?_ Most likely. However, it's probably not something we could target in our own client CDI stuff. So if the {{@Before}} {{@After}} stuff was not standardized, we'd be forced to couple our stuff to a specific extension. That's the only downside for us. It seems to me that, while not necessarily a problem for the majority, we should strongly consider if this is a minority we should ignore. My opinion is that, while at any given time this is a minority problem, the majority will probably run into it at some point or another. My personal preference would be for _both_ an SPI and an standard for ordering in the programming model.	1	1	https://issues.jboss.org/browse/CDI-4
3002	CDI-4	I have a comment about the syntax.I quite like the idea of using @Before and @After coupled with qualifiers for ordering event handlers and IMO partial ordering can go a pretty long way in solving the issues that we have. But, rather than using @Before and @After with the event qualifiers, I think that qualifiers should be placed on event handlers themselves. So, this may work: {code} public void myFirstObserver(@Observes @After(Initialization.class) MyEvent event) { ... } public void mySecondObserver(@Observes @Initialization MyEvent event) { ... } {code} but what if I want to order two observers for @Initialization MyEvent ? I can do: {code} public void myFirstObserver(@Observes @Initialization @After(Initialization.class) MyEvent event) { ... } public void mySecondObserver(@Observes @Initialization MyEvent event) { ... } {code} but that does not do a great favour to readability. On the other hand, what do we do with observers that do not have any qualifier at all? Mike suggested an @OrderingQualifier, but mixing it up with the rest of the qualifiers does not improve readability. On the other hand, ordering is relative to the method not to the event being observed, so why not qualify observer methods themselves (note how I stole the idea of placing the @Before annotation from Christian)? {code} @Initialization public void myFirstObserver(@Observes MyEvent event) { ... } @Before(Initialization.class) public void mySecondObserver(@Observes MyEvent event) { ... } {code} @Initialization can be a special type of annotation, or can be a simple @Qualifer. I also think that the idea can be reused in scenarios such as interceptor/decorator ordering, if we ever consider automatic interceptor/decorator enablement.	0	0	https://issues.jboss.org/browse/CDI-4
3003	CDI-4	And now that we have {{@Priority}}, can't we just use it ? {code} public class MyBean public void myFirstObserverInPriority(@Observes @Priority(1000) MyEvent event) { ... } } public class MyBean public void mySecondObserverInPriority(@Observes @Priority(42) MyEvent event) { ... } } {code}	1	1	https://issues.jboss.org/browse/CDI-4
3004	CDI-4	If we have something at the same priority level, is it just unspecified behavior?	1	1	https://issues.jboss.org/browse/CDI-4
3005	CDI-3	see my comment in CDI-58. All this is depending on our understanding of AnnotatedType: Is there 1 per Class or do we allow 'artificial' AnnotatedTypes?	1	1	https://issues.jboss.org/browse/CDI-3
3006	CDI-1	Add a note to clarify that the PC's should propagate as defined by JPA specs. Stuart if I do this, is it sufficient from your understanding?	1	1	https://issues.jboss.org/browse/CDI-1
3007	CDITCK-592	[~sameer-pandit] I think it should be available [here|https://repository.jboss.org/nexus/content/repositories/snapshots]. Can you please check it?	1	1	https://issues.jboss.org/browse/CDITCK-592
3008	CDITCK-486	Thank you, will you be excluding for the 1.2.4 as that is what the licensee is currently using. The description, that we provided you with is what we received from the Java EE licensee. Certainly happy to get more info as needed and please do not hesitate to ask if you need more clarity and we can follow up with the licensee.	1	0	https://issues.jboss.org/browse/CDITCK-486
3009	CDITCK-467	Which another tests do you mean? If you are talking about VerifyValuesTest#testManagedBeanAttributes there is the @Alternative selected in the beans.xml of test deployment archive.	0	0	https://issues.jboss.org/browse/CDITCK-467
3010	CDITCK-467	{quote}The posted algorithm will create random results depending on the processing order{quote} I cannot see how it creates random results. Can you explain?	1	0	https://issues.jboss.org/browse/CDITCK-467
3011	CDITCK-463	Hi Romain, Can you please elaborate why do you think it is not portable? The test appears to me ok.	1	1	https://issues.jboss.org/browse/CDITCK-463
3012	CDITCK-463	[~tremes] yes, sorry of the brief description. 1) I didnt find any reference to this field in the spec - did I miss it? 2) geronimo spec jar doesnt use it (https://svn.apache.org/repos/asf/geronimo/specs/trunk/geronimo-jcdi_1.1_spec/src/main/java/javax/enterprise/inject/spi/CDI.java)	1	1	https://issues.jboss.org/browse/CDITCK-463

#	Requirement	Questions	Answered	Successfully Answered	URL
3013	CDITCK-436	Mark, could you please describe the problem in more detail? {{WithAnnotationsAppliedToIllegalContainerLifecycleEventParameterTest}} does not try to use the BeanManager. I suppose the problem is with the injection point on the {{org.jboss.cdi.tck.AbstractTest}} (each TCK test class extends this class) - but AFAIK {{CDIInjectionEnricher}} always tries to lookup the {{BeanManager}} and inject the test class, no matter whether it has some injection points or not. Maybe there's a problem in a way Arquillian/{{BeanManagerProducer}} looks up the BM. If it doesn't find the BM via JNDI, it's using a {{CDIExtension}} and this class holds a BM reference in a static field. In this particular case weld-embedded is not affected because Weld is validating extensions before any container lifecycle event is delivered. Also this shouldn't be a problem for managed and remote containers, because {{@ShouldThrowException}} will force the client test run mode (the test archive is not enhanced).	1	1	https://issues.jboss.org/browse/CDITCK-436
3014	CDITCK-436	inject members at org.jboss.arquillian.testenricher.cdi.CDIInjectionEnricher.injectClass(CDIInjectionEnricher.java:135) at org.jboss.arquillian.testenricher.cdi.CDIInjectionEnricher.enrich(CDIInjectionEnricher.java:78) at org.jboss.arquillian.test.impl.TestInstanceEnricher.enrich(TestInstanceEnricher.java:52) at org.jboss.arquillian.container.test.impl.ClientTestInstanceEnricher.enrich(ClientTestInstanceEnricher.java:51) at sun.reflect.GeneratedMethodAccessor33.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.invokeObservers(EventContextImpl.java:99) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:81) at org.jboss.arquillian.container.test.impl.client.ContainerEventController.createContext(ContainerEventController.java:142) at org.jboss.arquillian.container.test.impl.client.ContainerEventController.createBeforeContext(ContainerEventController.java:124) at sun.reflect.GeneratedMethodAccessor32.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.test.impl.TestContextHandler.createTestContext(TestContextHandler.java:102) at sun.reflect.GeneratedMethodAccessor9.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.test.impl.TestContextHandler.createClassContext(TestContextHandler.java:84) at sun.reflect.GeneratedMethodAccessor8.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.test.impl.TestContextHandler.createSuiteContext(TestContextHandler.java:65) at sun.reflect.GeneratedMethodAccessor7.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:135) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:115) at org.jboss.arquillian.test.impl.EventTestRunnerAdaptor.before(EventTestRunnerAdaptor.java:95) at org.jboss.arquillian.testng.Arquillian.arquillianBeforeTest(Arquillian.java:130) at sun.reflect.GeneratedMethodAccessor31.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.testng.internal.MethodInvocationHelper.invokeMethod(MethodInvocationHelper.java:84) at	1	1	https://issues.jboss.org/browse/CDITCK-436
3015	CDITCK-427	Where in the spec we can read, that the incomplete instance should be used?	0	0	https://issues.jboss.org/browse/CDITCK-427
3016	CDITCK-416	The tests in PR should pass now. Can someone verify that and merge the PR?	1	1	https://issues.jboss.org/browse/CDITCK-416
3017	CDITCK-359	Should it implement support for patterns (e.g TestCircle#test*) or multiple methods (e.g TestCircle#testOne+testTwo) also?	1	1	https://issues.jboss.org/browse/CDITCK-359
3018	CDITCK-277	Could you expand on why EarSpecializationTest is not needed? It this covered by an existing CDI TCK test already?	1	1	https://issues.jboss.org/browse/CDITCK-277
3019	CDITCK-277	EarSpecializationTest differs only by using the @WebListener. Is that a special case that needs to be tested? I didn't find anything covering this concern in the specification.	1	1	https://issues.jboss.org/browse/CDITCK-277
3020	CDITCK-209	Ales, does this apply to the Application context only? In other words, is it safe to destroy other contexts besides Application (Request, Session) explicitly in the test?	1	1	https://issues.jboss.org/browse/CDITCK-209
3021	CDITCK-118	The maven repositories are declared since that check-in. If artifacts are missing, that is a different problem. Is this still a problem? I have succeeded in running the TCK on a platform using 1.0.1 final without any problems.	1	0	https://issues.jboss.org/browse/CDITCK-118
3022	CDITCK-114	Jozef, could you take a look at this?	1	1	https://issues.jboss.org/browse/CDITCK-114
3023	CDITCK-81	Pete, could you provide more information pls?	1	0	https://issues.jboss.org/browse/CDITCK-81
3024	CDITCK-41	Pete, could you provide more info pls?	1	1	https://issues.jboss.org/browse/CDITCK-41
3025	CDITCK-29	Do we need tests for this? The isn't that clear on how this is treated...	0	0	https://issues.jboss.org/browse/CDITCK-29
3026	CDITCK-20	Jozef, would you be able to write a quick test for this on Monday? If you can't find an assertion, mark it implied.	1	1	https://issues.jboss.org/browse/CDITCK-20
3027	CIRRAS-46	Another alternative, perhaps not exactly what's being asked for, but.. I think we're already passing YAML. We could add a field which points to a URL (or S3 bucket) containing the on-boot script(s). We wouldn't face the 8k limit, and URL/bucket referencing might be easier to handle than passing full shell scripts around. Just a thought.	1	0	https://issues.jboss.org/browse/CIRRAS-46
3028	CIRRAS-46	Hello Bob, what do you think about the possible solution I tried to briefly explain in the description? I believe that will be most straightforward for the users to use and understand. Without any need to encode the data in some complicated way. It should be something like this: 1. user wants to only set JBoss instance configuration: param1=1 param2=2 ... 2. user wants to set JBoss configuration but also wants a custom script: #!/bin/bash # JBoss instance specific config cat > /etc/jboss.conf << EOF param1=1 param2=2 ... EOF # User custom script yum install net-snmp passwd --stdin root <<< mypassword cat > /tmp/my.storage.example.com.crt << EOF <CA certificate data here> EOF curl --cacert /tmp/my.storage.example.com.crt https://my.storage.example.com/my/long/cool/setup/script.ps | perl	0	0	https://issues.jboss.org/browse/CIRRAS-46
3029	CIRRAS-34	Not sure I understand this issue. Brett, could you please tell me more? What components you mean?	1	1	https://issues.jboss.org/browse/CIRRAS-34
3030	CLOUD-1844	[~goldmann]: [~porcelli] wanted to be able to track this with a guvnor sprint too... Can't we have both sprints included here?	1	1	https://issues.jboss.org/browse/CLOUD-1844
3031	CLOUD-1616	I see INFO in black, WARN in yellow, and ERROR in red. Could this be a difference in our environments?	1	1	https://issues.jboss.org/browse/CLOUD-1616

#	Requirement	Questions	Answered	Successfully Answered	URL
3032	CLOUD-1616	Yeah, I see product messages also colorized in ocp web console, but I meant eg. message: "WARNING: Environment variable OPENSHIFT_KUBE_PING_NAMESPACE undefined. Clustering will be unavailable. Please refer to the documentation for configuration." Do you see that message also in yellow?	1	0	https://issues.jboss.org/browse/CLOUD-1616
3033	CLOUD-1592	I propose we add a parameter so the user can specify which tag they want to use for the DB (e.g.5.5 for mysql). I think we should also decide which versions we should use as the default (e.g. 5.5 or 5.6 for mysql). [~rettori], do you have any opinions here? Also, there's a related issue around the client versions included in the eap images. I'll link that to this issue.	1	1	https://issues.jboss.org/browse/CLOUD-1592
3034	CLOUD-1592	How about we default to latest but add the param so the user can lock-in a major.minor version?	0	0	https://issues.jboss.org/browse/CLOUD-1592
3035	CLOUD-1536	[~jpechanec], could you please provide more details? I'm not quite sure what you're suggesting.	1	0	https://issues.jboss.org/browse/CLOUD-1536
3036	CLOUD-1534	The Hawkular agent is configured through the following environment variables: * AB_HAWKULAR_REST_URL - the URL of the Hawkular service (required, if not set, Hawkular agent will not be configured) * AB_HAWKULAR_REST_USER - the user for authenticating with the Hawkular service * AB_HAWKULAR_REST_PASSWORD - the password for authenticating with the Hawkular service * AB_HAWKULAR_REST_FEED_ID - the feed ID, defaults to autogenerate, must be globally unique to Hawkular service * AB_HAWKULAR_REST_TENANT_ID - the tenant ID, defaults to hawkular * AB_HAWKULAR_REST_KEYSTORE - the keystore file name, used to configure truststore for the Hawkular service (required if Hawkular service is using https) * AB_HAWKULAR_REST_KEYSTORE_DIR - the directory in which the keystore file is located * AB_HAWKULAR_REST_KEYSTORE_PASSWORD - the password for the keystore file * AB_HAWKULAR_AGENT_SUBSYSTEM_SNIPPET - the full path to a subsystem xml snippet which should be used to configure the Hawkular subsystem; defaults to JBOSS_HOME/modules/system/add-ons/hawkular-agent/org/hawkular/agent/main/subsystem-snippet.xml	1	1	https://issues.jboss.org/browse/CLOUD-1534
3037	CLOUD-1534	[~mazz], thanks for the explanation. For eap, we're currently configuring the agent through the subsystem and simply mimicking the logic of the eap installer (i.e. we add the modules, add the extension, subsystem and security realm to standalone.xml). By default, we use the subsystem-snippet.xml that's included with the module. That said, are you saying that none of this is necessary and that all we need to do is add something like -javaagent=hawkular.jar to the command line and supply a yaml config?	1	1	https://issues.jboss.org/browse/CLOUD-1534
3038	CLOUD-1534	What is the difference between using the java agent and the eap subsystem? Which mechanism are we going to recommend to customers? It also seems like there are some tweaks that need to be made to eap to make the java agent work (adding modules to system packages), which I'm guessing are not necessary if going the subsystem route.	1	1	https://issues.jboss.org/browse/CLOUD-1534
3039	CLOUD-1534	Last I heard, [~kconner] wants all our product images (including EAP) to try to use the same agent mechanism - which in this case means the Java Agent. I would talk to [~kconner] and [~pgier] about this as I'm not the right person to talk to about productizing and packaging. But to answer your questions: Question 1. "are you saying that none of this is necessary and that all we need to do is add something like -javaagent=hawkular.jar to the command line and supply a yaml config?" Answer 1: If you are using the Java Agent, yes. You just do that and no need to mess with subsystem stuff because the agent in that mode doesn't run "inside" the wildfly container. Question 2. "What is the difference between using the java agent and the eap subsystem?" Answer 2: One runs inside the WildFly container as a subsystem extension and thus follows its lifecycle (its part of the WildFly server). The java agent is simply attached to the WildFly JVM as a java agent and isn't directly affected by it (e.g. when you "reload" the WildFly server, the java agent doesn't get reloaded along with it since it isn't a subsystem that is affected by things like that). The other difference is the subsystem agent is configured directly in standalone.xml and the java agent is configured in its own yaml config. Question 3: "Which mechanism are we going to recommend to customers?" Answer 3: Not really a question for me as I'm responsible for supporting both :) [~jdoyle] can give you the official answer. But as I say [~kconner] wants all our images to support a single way of attaching the agent (this includes things like Fuse that aren't necessarily EAP-based) so that would probably mean we are going to recommend using the Java Agent. Question 4: "also seems like there are some tweaks that need to be made to eap to make the java agent work (adding modules to system packages), which I'm guessing are not necessary if going the subsystem route." Answer 4: To get the Java Agent to run properly in EAP, you have to add jboss logging to the EAP base classloader due to some known limitations that EAP/WildFly imposes on java agents (I don't have links handy to explain that - but suffice it to say the WildFly dev team knows about this but there isn't much that can be done about it at this time short of what I document in the README found at https://github.com/hawkular/hawkular-agent/tree/master/hawkular-javaagent) Question 5: " if subsystem is the way to go, can you configure the subsystem to read from a yaml config too?" Answer 5: No. If the agent is run as a subsystem, you have to configure it like all the other WildFly subsystems - via standalone.xml. It cannot read an external yaml file. The good thing about this is when running as a subsystem, you can read its config and change its config via the Jboss CLI (which you cannot do if using the Java Agent because, as I say, the agent in that case is not running inside WildFly as a subsystem extension).	0	0	https://issues.jboss.org/browse/CLOUD-1534
3040	CLOUD-1534	[~rcernich] BTW. Why the "AB_" prefix?	1	1	https://issues.jboss.org/browse/CLOUD-1534
3041	CLOUD-1482	Do we also need to modify assembly to use scl enable?	0	0	https://issues.jboss.org/browse/CLOUD-1482
3042	CLOUD-1482	Just quick heads up, maven 3.5 was released last week. We could even consider moving to 3.5 right away, as otherwise it might take another 5 years to update maven in builder images again.... Btw, couldn't maven be installed simply by unpacking tar.gz file in right location of image?	1	1	https://issues.jboss.org/browse/CLOUD-1482
3043	CLOUD-1482	{quote}We could even consider moving to 3.5 right away, as otherwise it might take another 5 years to update maven in builder images again....{quote} With SCL we will be moving faster than we do currently. That's the promise os SCL, in fact :) {quote}Btw, couldn't maven be installed simply by unpacking tar.gz file in right location of image?{quote} No, everything needs to come out of Brew.	0	0	https://issues.jboss.org/browse/CLOUD-1482
3044	CLOUD-1476	* Added, processed, and tagged PR in openshift-kieserver. * Removed WIP prefix from jboss-dockerfiles and application-templates PRs. * Once we have a brew build of openshift-kieserver 1.2.0.Final-redhat-1, the jboss-dockerfiles and application-templates PRs can be processed. * Added a PR in openshift-quickstarts, which updates them to latest KIE version. Although the current 1.3 quickstarts work fine (I regression tested), I would like to process the quickstarts upgrade and tag them (1.4?). Then another PR for application-templates using the new tag should be created. (I have already also tested with the new version, so this will be easy.) * ce-testsuites is still WIP, however.	0	0	https://issues.jboss.org/browse/CLOUD-1476
3045	CLOUD-1389	{quote}The assumption has been that the majority of the artifacts should be coming from our production repositories, i.e. built and shipped by ourselves, rather than from maven central.{quote} So, we (did) assume most of artifacts from users' apps running on top of our (java, eap) images are present on our mirrors? Is it safe to make such (audacious) assumption? Or did I get it wrong?	1	1	https://issues.jboss.org/browse/CLOUD-1389

#	Requirement	Questions	Answered	Successfully Answered	URL
3046	CLOUD-1389	Does the order of redhat repos affect performance (e.g. ga -> tech preview -> early access)? What is the percentage of "red hat" artefacts pulled from each of those repos?	0	0	https://issues.jboss.org/browse/CLOUD-1389
3047	CLOUD-1389	So, should we close this as resolved?	0	0	https://issues.jboss.org/browse/CLOUD-1389
3048	CLOUD-1386	Work with OpenShift Online team to identify appropriate tags for other images (e.g. what tags identify process server?).	0	0	https://issues.jboss.org/browse/CLOUD-1386
3049	CLOUD-1384	Forgive my ignorance, the additions would need to happen to all the templates here: https://github.com/jboss-openshift/application-templates correct? Also, if that is true, would we add a parameter in a similar fashion to the following: {code:JSON} { "displayName": <>, "description": <>, "name": "OPENSHIFT.IO", "value": <>, "required": true }, {code} from this example https://github.com/jboss-openshift/application-templates/blob/master/amq/amq62-basic.json#L18-L24 Is this what needs to be done?	0	0	https://issues.jboss.org/browse/CLOUD-1384
3050	CLOUD-1384	Do I need special permission to access the link above? I get a 404 when I try to access it	0	0	https://issues.jboss.org/browse/CLOUD-1384
3051	CLOUD-1383	Should we move this to Sprint 7?	1	1	https://issues.jboss.org/browse/CLOUD-1383
3052	CLOUD-1382	Should we move this to Sprint 7?	0	0	https://issues.jboss.org/browse/CLOUD-1382
3053	CLOUD-1355	[~jpechanec]How did you configure the resource adapter? Are you specifying RESOURCE_ADAPTERS?	0	0	https://issues.jboss.org/browse/CLOUD-1355
3054	CLOUD-1354	[~jpechanec] How did you configure the resource adapter? Are you specifying RESOURCE_ADAPTERS?	0	0	https://issues.jboss.org/browse/CLOUD-1354
3055	CLOUD-1327	[~filippe.spolti] Is there any reason why the value is hardcoded to 600000? (Why exactly 10 minutes and why is it not configurable?)	1	1	https://issues.jboss.org/browse/CLOUD-1327
3056	CLOUD-1327	[~rcernich] [~maschmid] Agreed, I'll make it configurable so user can define which configuration type to use. [~maschmid] Did you get this issue only with postgresql? Can you share the steps to reproduce it? Thanks.	0	0	https://issues.jboss.org/browse/CLOUD-1327
3057	CLOUD-1292	Thanks Jiri! [~rcernich] [~bdecoste] Can you guys take a look at it? Please schedule, approve an link MR if it's OK.	1	1	https://issues.jboss.org/browse/CLOUD-1292
3058	CLOUD-1233	[~kconner] this was related to the EAP 64 image. Does this mean the EAP 64 image will start using the Java S2I image?	0	0	https://issues.jboss.org/browse/CLOUD-1233
3059	CLOUD-1180	So, I went through all our images and created list of files added to these images to the bin folder (wherever it is located). I do not see anything unusual here. In almost all images the only thing living in root bin/ directory is the launch script itself and readiness/liveness probes. The only difference is the kieserver image which has 4 files added additionally besides this, but no launch directory at all. [~kconner] How would you like to proceed me here? I do not see anything wrong at the first glance. Did I miss something?	1	1	https://issues.jboss.org/browse/CLOUD-1180
3060	CLOUD-1170	Shouldn't this be closed as duplicated and JBIDE-23628 pointed to original issue? CLOUD-940 was reopened and reopening this issue with it just doubles the work with tracking same issue. Another solution could be adding EAP component and then this issue could be verified against only EAP templates	1	0	https://issues.jboss.org/browse/CLOUD-1170
3061	CLOUD-1146	Is this an RH-SSO issue or an EAP issue?	1	1	https://issues.jboss.org/browse/CLOUD-1146
3062	CLOUD-946	I tested it with jboss-eap-6/eap64-openshift:1.5-5 and it seems these fixes are not included because i don't see <!-- ##EXTENSION_RESOURCE_ADAPTERS## --> in /opt/eap/standalone/configuration/standalone-openshift.xml. Can you check [~goldmann] if these fixes are included?	0	0	https://issues.jboss.org/browse/CLOUD-946
3063	CLOUD-944	I have some questions about the resource adapter configuration: # It seems that module slot / module id is required, how can I configure resource adapter not from a module, but deployed in deployments/foo.rar ? # according to https://access.redhat.com/documentation/en-US/JBoss_Enterprise_Application_Platform/6/html/Administration_and_Configuration_Guide/Configure_a_Deployed_Resource_Adapter1.html#Configure_a_Deployed_Resource_Adapter1 there are two separate kinds of <config-property/> elements (either for configuring resource adapter, and separate for the connection) .. According to the scripts/os-eap-launch/added/launch/resource-adapters-common.sh we are using FOO_PROPERTY_ as a connection config property, with no way to specify resource adapter config properties... (it should perhaps be FOO_PROPERTY_... and FOO_CONNECTION_PROPERTY_ as separate kinds of env...)	0	0	https://issues.jboss.org/browse/CLOUD-944
3064	CLOUD-927	Closing again because it was done in CLOUD-994. My fault.	1	1	https://issues.jboss.org/browse/CLOUD-927
3065	CLOUD-906	I merged your MR. Should this change be also applied to EAP 7? I think we want to keep them configured the same way. If yes, please clone this ticket. Yes, these tests should be moved to EAP 7 in case we do not want to update EAP7 image with this change. Otherwise these should be removed. Since this change introduces new way of adding datasources we need to add new tests that use new environment variables, please add them.	0	0	https://issues.jboss.org/browse/CLOUD-906
3066	CLOUD-893	The products already have the version information available, this can be inspected by docker. Can you please clarify what you are asking for here?	0	0	https://issues.jboss.org/browse/CLOUD-893
3067	CLOUD-877	[~jonh.wendell], can you please add release notes to this describing the new variables?	1	0	https://issues.jboss.org/browse/CLOUD-877

#	Requirement	Questions	Answered	Successfully Answered	URL
3068	CLOUD-865	Here are the published documentation references related to proxy environment variables and s2i builds. The gist is that it looks like the xPaas images bring in 5 additional env vars not used in other non-xpaas images in Openshift. The additional env vars are: HTTP_PROXY_HOST HTTP_PROXY_PORT HTTP_PROXY_USERNAME HTTP_PROXY_PASSWORD HTTP_PROXY_NONPROXYHOSTS The common Openshift proxy env vars are: HTTP_PROXY HTTPS_PROXY NO_PROXY These difference may lead to customers encountering failed builds of xpaas image related applications. We can KCS this and clearly explain what is required. Should we look to make any changes around these proxy env vars? Documentation references: Openshift Installation and Configuration (https://access.redhat.com/documentation/en/openshift-enterprise/3.2/single/installation-and-configuration/#s2i-builds) 18.5. CONFIGURING S2I BUILDS FOR PROXIES The following are the supported proxy environment variables with example values: HTTP_PROXY=http://USERNAME:PASSWORD@10.0.1.1:8080/ HTTPS_PROXY=https://USERNAME:PASSWORD@10.0.0.1:8080/ NO_PROXY=master.hostname.example.com ---------- Openshift Developers Guide (https://access.redhat.com/documentation/en/openshift-enterprise/3.2/paged/developer-guide/chapter-10-builds) 10.4.5. Environment Variables For example, defining a custom HTTP proxy to be used during build and runtime: dockerStrategy: ... env: - name: "HTTP_PROXY" value: "http://myproxy.net:5187/" ---------- xPaas EAP image doc (https://access.redhat.com/documentation/en/red-hat-xpaas/0/paged/red-hat-xpaas-eap-image/chapter-3-reference-information) 3.6.4. Environment variables HTTP_PROXY_HOST Hostname or IP address of a HTTP proxy for Maven to use. Example: 192.168.1.1 HTTP_PROXY_PORT TCP Port of a HTTP proxy for Maven to use. Example: 8080 HTTP_PROXY_USERNAME If supplied with HTTP_PROXY_PASSWORD, use credentials for HTTP proxy. Example: myusername HTTP_PROXY_PASSWORD If supplied with HTTP_PROXY_USERNAME, use credentials for HTTP proxy. Example: mypassword HTTP_PROXY_NONPROXYHOSTS If supplied, a configured HTTP proxy will ignore these hosts. Example: some.example.org|*.example.net ---------- KCS article offering solution to ENV var differences (https://access.redhat.com/solutions/2194751) You can set HTTP_PROXY_HOST and HTTP_PROXY_PORT environment variables under the sourceStrategy in BuildConfig. Raw spec: strategy: sourceStrategy: env: - name: HTTP_PROXY_HOST value: <host> - name: HTTP_PROXY_PORT value: <port>	1	1	https://issues.jboss.org/browse/CLOUD-865
3069	CLOUD-865	At this point, I don't think there is anything to change outside of maybe documentation. [~edrich] - Do we replicate this openshift.com documentation in Red Hat docs anywhere? If not, do you have a preference on whether this needs to be reflected in any other docs?	1	1	https://issues.jboss.org/browse/CLOUD-865
3070	CLOUD-865	Hey [~trogers-rh] - if these are xPaaS specific environment variables, then they should be updated in xPaaS docs only. If you want them updated, could you please file a bug report? https://bugzilla.redhat.com/enter_bug.cgi?product=Red%20Hat%20xPaaS (you can just link to this JIRA). Eric is right: OSE docs are published here: https://docs.openshift.com/enterprise/latest/welcome/index.html and here: https://access.redhat.com/documentation/en/openshift-enterprise/ docs.openshift.com gets updated nightly, while the portal gets updated every 12 hours. xPaaS docs are only published here: https://access.redhat.com/documentation/en/red-hat-xpaas/	1	1	https://issues.jboss.org/browse/CLOUD-865
3071	CLOUD-865	[~trogers-rh], [~jdowland] The redirect was set up when the new name was made live. It is redirecting for me from https://access.redhat.com/documentation/en/red-hat-xpaas/ to https://access.redhat.com/documentation/en/red-hat-jboss-middleware-for-openshift/ correctly. Is it not working for you guys?	1	1	https://issues.jboss.org/browse/CLOUD-865
3072	CLOUD-865	[~trogers-rh] [~jdowland] The docs tooling team pushed the Akamai rules that control the redirects again. They have mentioned that it should be working correctly now, as per their tests. Can you please verify if the redirects are working correctly now?	1	1	https://issues.jboss.org/browse/CLOUD-865
3073	CLOUD-818	Not sure if we should be writing to SSO_TRUSTSTORE_DIR (would it work if the secret was mounted as readOnly?)	1	0	https://issues.jboss.org/browse/CLOUD-818
3074	CLOUD-794	Could you please be more specific? What template do you mean? is this template a part of application being deployed?	0	0	https://issues.jboss.org/browse/CLOUD-794
3075	CLOUD-794	What to communicate to the customer? Will this be implemented?	0	0	https://issues.jboss.org/browse/CLOUD-794
3076	CLOUD-792	Can you please point out where did you find that this option is disabled? If this option is not on the configuration file means that it is enabled, since its default value is true. You can check in the container: http://<TOMCAT_CONTAINER_ADDRESS/docs/config/host.html	0	0	https://issues.jboss.org/browse/CLOUD-792
3077	CLOUD-792	Any updates for this request? Do you require any further information at this time?	0	0	https://issues.jboss.org/browse/CLOUD-792
3078	CLOUD-792	Any updates for this request? Do you require any further information at this time?	0	0	https://issues.jboss.org/browse/CLOUD-792
3079	CLOUD-788	The current image restricts access to "127\.\d+\.\d+\.\d+|::1|0:0:0:0:0:0:0:1" which if I understand correctly boils down to localhost | localhost | localhost. So I don't think there is anything to do there. What is the issue with 404.jsp? This is part of the standard tomcat distro	0	0	https://issues.jboss.org/browse/CLOUD-788
3080	CLOUD-786	I am not sure about this one either. when container is created it runs some command inside and when this command exits container itself is destroyed. There cannot be a container which does not run anything. What are these shutdown.sh and catalina.sh? Are these scripts usually triggered by administrator to stop tomcat or they are called by the tomcat itself when it received a terminating signal?	1	1	https://issues.jboss.org/browse/CLOUD-786
3081	CLOUD-786	[~trogers-rh], [~rhn-support-akokshar], should the implementation be changed to remove the shutdown port completely?	0	0	https://issues.jboss.org/browse/CLOUD-786
3082	CLOUD-785	Could you please provide me with an example how we can add necessary property in production environment? This all is about docker image, so we cant modify its content. we can pass parameters to the application through environment variables and possibly through command line parameters, but "command line" is not always the option in case we deploy application on OpenShift.	1	1	https://issues.jboss.org/browse/CLOUD-785
3083	CLOUD-785	Do we not support customizing the docker image to change configuration (either by using Dockerfile to replace the default configuration, or using openshift docekrbuild or s2i build to replace configuration files)?	0	0	https://issues.jboss.org/browse/CLOUD-785
3084	CLOUD-784	Are you suggesting to give the permissions only for the jboss user? The description is a little bit confusing.	0	0	https://issues.jboss.org/browse/CLOUD-784
3085	CLOUD-784	There are write privileges for group, is it really ok?	0	0	https://issues.jboss.org/browse/CLOUD-784
3086	CLOUD-717	Is this a dupe of CLOUD-743?	0	0	https://issues.jboss.org/browse/CLOUD-717
3087	CLOUD-693	Doesn't seem to work with jboss-webserver-3/webserver30-tomcat8-openshift:1.3-2 {code} oc start-build ews-sti-build --from-file=ROOT.war {code} Note the build logs about the paths... copying from /home/jboss and /home/jboss/source/deployments .. .should it be from /home/jboss/source instead of /home/jboss ? {noformat} Receiving source from STDIN as file ROOT.war Copying all war artifacts from /home/jboss directory into /opt/webserver/webapps for later deployment... Copying all war artifacts from /home/jboss/source/deployments directory into /opt/webserver/webapps for later deployment... Pushing image 172.30.45.210:5000/maschmid/ews-sti-image:latest ... Pushed 6/7 layers, 86% complete Pushed 7/7 layers, 100% complete Push successful {noformat}	1	1	https://issues.jboss.org/browse/CLOUD-693

#	Requirement	Questions	Answered	Successfully Answered	URL
3088	CLOUD-693	Does it work with the tomcat7 image (code is the same)? Looks like the copy_artifacts command is being invoked with $(pwd) as the source, with the assumption that the assemble script is being invoked from the source directory (/tmp/...). I wonder if there is a difference of behavior between a regular source build and a binary build?	1	1	https://issues.jboss.org/browse/CLOUD-693
3089	CLOUD-528	[~bdecoste] I'm pretty sure we're adding this for jdv. If so, could you please link those issues to this one and close this one out?	1	1	https://issues.jboss.org/browse/CLOUD-528
3090	CLOUD-528	Was this included in tomcat as well in the end? (is there any related JWS JIRA for this issue?)	0	0	https://issues.jboss.org/browse/CLOUD-528
3091	CLOUD-528	OK... can we remove it from the Components list from this JIRA then>	1	1	https://issues.jboss.org/browse/CLOUD-528
3092	CLOUD-498	We already configure the following on our datasources, what else do you think is missing? {quote} <validate-on-match>true</validate-on-match> <valid-connection-checker class-name=\"$8\"></valid-connection-checker> <exception-sorter class-name=\"$9\"></exception-sorter> {quote}	0	0	https://issues.jboss.org/browse/CLOUD-498
3093	CLOUD-430	[~rcernich] or [~kconner], could you approve this so that I can proper tag it as 'pull request sent'? Meanwhile, here are the links: https://ce-gitlab.usersys.redhat.com/ce/jboss-dockerfiles/merge_requests/1127 https://github.com/jboss-openshift/application-templates/pull/122	1	1	https://issues.jboss.org/browse/CLOUD-430
3094	CLOUD-106	I just tried and didn't work. I opened [this issue|https://issues.jboss.org/browse/ENTMQ-2115] to track it. IMHO, we could use the JVM wide configuration, [~rcernich], what do you think?	1	1	https://issues.jboss.org/browse/CLOUD-106
3095	DASHBUILDE-237	Regarding individual chart filters (see example => !chartfilter.png|thumbnail!) I think both "master" toolbar and individual chart filters can co-exist. A toolbar at the top of the dashboard screen could be implemented as another displayer component which allows for filtering on any of the properties of a given data set. Dashboard designer users can decide which approach to take. For toolbar, a toolbar filter component is dropped at the top of the page and filtering is disabled for every single chart. [~uxdlc] [~mstevens1] wdyt?	1	1	https://issues.jboss.org/browse/DASHBUILDE-237
3096	DASHBUILDE-237	[~davidredhat] I see these filter tags used mostly with big data, faceted/complex filtering. I believe the use case there is that there is so much criteria (and objects) that you're filtering for you really need this kind of visual summary. But for individual chart filters I'm not sure the use case exists, perhaps a slightly more complex filter control would be enough, like a multi-select dropdown: http://dotansimha.github.io/angularjs-dropdown-multiselect/docs/#/main Wdyt?	1	1	https://issues.jboss.org/browse/DASHBUILDE-237
3097	DASHBUILDE-237	There are some cases where individual chart filters control might make sense. See attached exampleB & exampleC. These are simple scenarios where we do not require a full featured filter toolbar as they are quite simple uses cases. A filter toolbar is more suitable for full dashboard scenarios, see exampleA. Anyway, users can always choose whether to use individual chart filters or a global toolbar. [~uxdlc] We do already have selector dropdowns for single column filtering, unlike the filter toolbar which is multi-column. [~mstevens1] 1) As per each icon's tooltip: "Export to CSV", "Export to Excel", "Show/hide columns" 2) +1. The string is a hard coded literal 3) What are you trying to build? Can you elaborate on this please?	1	1	https://issues.jboss.org/browse/DASHBUILDE-237
3098	DASHBUILDE-237	I seem to be losing the thread here a bit. A quick conversation might help. Some stray questions (none of which are directly related to the project Metrics page anymore)... 1) A filter at the top of the screen (a global filter) - does that data column need to be present in every chart on the screen? I can only globally filter by _time_, for example, if every chart has a time value. 2) When filtering one chart on the screen (local filter), should you also filter other charts that have the same data column? I filter one chart by "David" and the other charts with people update? 3) If answering yes to #2, should you then show the "tags" on the other charts to make it clear that the chart is filtered? 4) Have I gathered correctly that Dashbuilder uses NVD3? 5) Does the user get to choose whether the chart displays a legend? And, further, whether that legend is clickable / filters the chart? 6) Are all of these things customizable by the user anyway and my questions are moot?	1	0	https://issues.jboss.org/browse/DASHBUILDE-237
3099	DASHBUILDE-237	Hello [~davidredhat] do you have time to chat on Thursday or Friday? [~uxdlc] Thank you, Matthew	1	0	https://issues.jboss.org/browse/DASHBUILDE-237
3100	DASHBUILDE-237	The following screenshot shows the toolbar implemented in RedHat's BPMS => !bpms_toolbar.png|thumbnail! The first dropdown ("Id") shows a set of predefined columns that support text based search. The rest of the dropdowns ("State", "Errors", "Name" and "Last update") are extra filters for specific columns. They have been expanded in the toolbar to make it easier to operate with them as well as to provide custom filter controls depending on the column type. Notice, for example, the "Last update" column is an advanced control for selecting a date range. In order to support the features above the following requirements must be also considered: 1.- The toolbar requires a list of columns to display in the generic column dropdown 2.- It also requires a sublist of the columns to be expanded in the toolbar 3.- Each expanded column may optionally have a divider element. This enables the concept of column groups in the expanded section The main goal is to provide a generic toolbar able to support different toolbar configurations, including the one used in BPMS. Regarding the discussion in previous comments about having individual chart filters. The best option is to remove them and rely only in either the toolbar component or the existing selectors components (dropdown, label set and slider). Whenever a filter is required in a dashboard, a filtering component needs to be dropped into the page. Charts will no longer provide filtering functionality. That will not only reduce the codebase but it will also keeps the filtering approach, Patternfly and the UXD on the same page. [~uxdlc] [~mstevens1] wdyt?	0	0	https://issues.jboss.org/browse/DASHBUILDE-237
3101	DASHBUILDE-232	Hi Team, It would be a great help if a workaround is provided to have default ascending sort of category. Currently. there is no sorting and category axis appears in random order,	1	1	https://issues.jboss.org/browse/DASHBUILDE-232
3102	DASHBUILDE-230	This is the old implementation => [^dashbuilder_filter_labels.mp4] This is how it looks the new version => [^dashbuilder_filter_labels_new.mp4] [~uxdlc] wdyt? :-)	1	1	https://issues.jboss.org/browse/DASHBUILDE-230
3103	DASHBUILDE-230	[~davidredhat] I sent my comments to a design related jira, next time will just post them here directly. The visual styling of the labels themselves looks more like PF, but some of the interaction aspects aren't aligned. * For future iterations/enhancements: The PF filters are typically used in conjunction with a global [toolbar control|http://www.patternfly.org/pattern-library/forms-and-controls/toolbar/]. It's presented at the top of the screen (or content view) to indicate that everything on that screen is being filtered. The filter labels are shown there to help users to know at a glance which filters have been applied. I haven't seen an example where they are applied out of that context. * Usability concerns with the current dashboard presentation are: _Because the labels are visually associated with one widget, an assumption will be made that the filters have only be applied to that widget. _Users need to hunt around the screen to find the labels to figure out which filters have been applied, whereas in the PF toolbar design they are all presented in one location at the top of the screen. To align the interaction with PF (in addition to the styling) could we look into the options for picking up the toolbar? Of it that can't be used for some reason, perhaps we could look into some alternative design solutions. Let me know what you think, thanks!	0	0	https://issues.jboss.org/browse/DASHBUILDE-230

#	Requirement	Questions	Answered	Successfully Answered	URL
3104	DASHBUILDE-229	Initial implementation done. Preview video => [^dashbuilder_perspective_component.mp4] As the video shows, there are issues with the layout rows. The logic behind the calculation of the row height seems not to be working very well. [~ederign] any ideas?	0	0	https://issues.jboss.org/browse/DASHBUILDE-229
3105	DASHBUILDE-228	[~davidredhat] or [~porcelli], can you guys please put this on Sprint5?	0	0	https://issues.jboss.org/browse/DASHBUILDE-228
3106	DASHBUILDE-226	This issue is a bit tricky. The purpose of having a fixed with for a label set is mostly to expand its size (see !selectorlabelswidth.png|thumbnail!). Even in this example there might be cases where one item's name is still too large. Also, it does not make sense to set a very small width since the item's name won't fit in the button. So, what's the solution? I think a good one could be as follows: 1.- Expand If the total width is larger than the selector width then all the items shall be displayed without cuts and the remaining width should be split and added up evenly to every item. The total width is the sum of all items width whereas an item's width is calculated as a function based on its number of chars. For example, if we have the following items: "Red", "Green" and "Blue", the total length is 12 chars and the total width 120px (provided we set 1 char=10px). If, for instance, the selector width is set to 210 then the width of each item will be: 210-120 = 90px excess width => we split evenly 90/3 = 30px => Red=3chars*10+30=60px, Green=5chars*10+30=80px, Blue=4chars*10+30=70px As a result, every button is expanded 30px. 2.- Reduce If the width set is smaller than the total width then we need to reduce each item so that they fit into the button reduced space. Given the example above and a selector width=90. then 90-120= -30px to reduce overall, -10px each item => Red=30-10=20px, Green=50-10=40px, Blue=40-10=30px. The buttons will show as "Re", "Gree", "Blu" with their corresponding tooltip showing the full name. [~jhrcek] wdyt?	1	1	https://issues.jboss.org/browse/DASHBUILDE-226
3107	DASHBUILDE-218	[~uxdlc], [~jhrcek], the current implementation is ready to push. I've implemented two alternatives: 1.- Item's menu attached to the type icon on the left => [^dashbuilder_navitem_actions2.mp4] 2.- Item's menu separated into a new icon on the right => [^dashbuilder_navitem_actions.mp4] Option #2 makes more explicit the presence of a menu since the icon shows up only when the mouse is over the item. Which do you like the most?	1	1	https://issues.jboss.org/browse/DASHBUILDE-218
3108	DASHBUILDE-218	[~jhrcek][~davidredhat] Cool! And I like Option #2 better as well. Couple of ideas: - I noticed that the cog icon turns blue on hover. For consistency, could we have all of the action icons (on the right) turn blue on hover? - Except, wdyt of having the delete (trash can) icon turn red on hover to indicate that it's a "destructive" action? For example: !Screen Shot 2017-05-24 at 10.40.20 AM.png|thumbnail!	1	1	https://issues.jboss.org/browse/DASHBUILDE-218
3109	DASHBUILDE-93	Hello [~davidredhat], I'm looking at latest kie-wb and dashbuilder showcase webapps, but I don't see any effect of this setting (nor in the chart preview in dataset editor, nor in the created perspectives). For which charts is this setting expected to work? For every that has it in the Display tab? I tried with Bar/Area/Line charts and I didn't see any effect on the axis labels.. Could you please verify that it works?	0	0	https://issues.jboss.org/browse/DASHBUILDE-93
3110	DASHBUILDE-70	I cannot reproduce it yet. I have tried it using the dashbuilder's showcase and the whole kie-wb, both webapps tested in Chrome, Firefox and Safari (on mac), and the text alignment is fine... any ideas?	1	1	https://issues.jboss.org/browse/DASHBUILDE-70
3111	DASHBUILDE-28	Ok Roger. Talking about Patternfly , could you please share a short info on this initiative? I saw the webpage [https://www.patternfly.org/], but I have no idea what's the connection to our projects. Does it mean using some additional library of widgets or something, or is it just a set of design recommendations to make the apps more user friendly / unified?	1	1	https://issues.jboss.org/browse/DASHBUILDE-28
3112	DASHBUILDE-10	Questions about this issue: - The project already depends on resteasy (2.3.7). It is a dependency of Uberfire. Should I keep this version? - Is there a code formatter, to decrease merge needs? - I think that the rest-api should reside on module 'dashbuilder-webapp'. Is it ok? - Any consideration about packages structure/names?	1	0	https://issues.jboss.org/browse/DASHBUILDE-10
3113	DBZ-304	{noformat} for i in `egrep -io 'create database [^;]+' setup.sql | awk '{print $3}'`; do echo "====================== $i ================="; grep -l -r $i ../../java/; done ====================== emptydb ================= ../../java/io/debezium/connector/mysql/ConnectionIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorIT.java ====================== readbinlog_test ================= ../../java/io/debezium/connector/mysql/ReadBinLogIT.java ../../java/io/debezium/connector/mysql/MySqlUnsignedIntegerIT.java ../../java/io/debezium/connector/mysql/ConnectionIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorJsonIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorIT.java ../../java/io/debezium/connector/mysql/MetadataIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorRegressionIT.java ../../java/io/debezium/connector/mysql/MySqlGeometryIT.java ====================== connector_test ================= ../../java/io/debezium/connector/mysql/SnapshotReaderIT.java ../../java/io/debezium/connector/mysql/MySqlUnsignedIntegerIT.java ../../java/io/debezium/connector/mysql/BinlogReaderIT.java ../../java/io/debezium/connector/mysql/MySqlDdlParserTest.java ../../java/io/debezium/connector/mysql/MySqlConnectorJsonIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorRegressionIT.java ../../java/io/debezium/connector/mysql/FiltersTest.java ../../java/io/debezium/connector/mysql/MySqlTaskContextTest.java ../../java/io/debezium/connector/mysql/MySqlGeometryIT.java ../../java/io/debezium/connector/mysql/MySqlSchemaTest.java ====================== connector_test_ro ================= ../../java/io/debezium/connector/mysql/SnapshotReaderIT.java ../../java/io/debezium/connector/mysql/BinlogReaderIT.java ../../java/io/debezium/connector/mysql/MySqlDdlParserTest.java ../../java/io/debezium/connector/mysql/MySqlConnectorIT.java ../../java/io/debezium/connector/mysql/MySqlTaskContextTest.java ====================== regression_test ================= ../../java/io/debezium/connector/mysql/MySqlUnsignedIntegerIT.java ../../java/io/debezium/connector/mysql/BinlogReaderIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorJsonIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorRegressionIT.java ../../java/io/debezium/connector/mysql/MySqlGeometryIT.java ====================== json_test ================= ../../java/io/debezium/connector/mysql/MySqlUnsignedIntegerIT.java ../../java/io/debezium/connector/mysql/MySqlConnectorJsonIT.java ../../java/io/debezium/connector/mysql/MySqlGeometryIT.java ====================== geometry_test ================= ../../java/io/debezium/connector/mysql/MySqlUnsignedIntegerIT.java ../../java/io/debezium/connector/mysql/MySqlGeometryIT.java ====================== table_maintenance_test ================= ../../java/io/debezium/connector/mysql/MySqlTableMaintenanceStatementsIT.java ====================== binary_column_test ================= ../../java/io/debezium/connector/mysql/MySqlFixedLengthBinaryColumnIT.java ====================== unsigned_integer_test ================= ../../java/io/debezium/connector/mysql/MySqlUnsignedIntegerIT.java {noformat} * Can we split the monster file to file per database? * Can we create/drop databases during test runtime as is done with schemas in case of PostgreSQL connector? * Can we minimize resuing of databases between tests to only logically grouped cases?	0	0	https://issues.jboss.org/browse/DBZ-304
3114	DBZ-252	Hi [~anton_nazaruk], I think we already largely ignore those definitions (at least on master). Could you provide an example for the problematic definitions you encounter? Thanks!	1	1	https://issues.jboss.org/browse/DBZ-252

#	Requirement	Questions	Answered	Successfully Answered	URL
3115	DBZ-248	Taking a step back, couldn't this also be handled at the level of a management platform such as Kubernetes? Which would monitor the container running Connect and restart it if a task failed. I.e. when running within such environment, we'd get that functionality for free. I haven't formed an opinion really, I can see advantages to both approaches, solving restarts within the connector or pushing it to some outer layer.	0	0	https://issues.jboss.org/browse/DBZ-248
3116	DBZ-232	Added a [pull request|https://github.com/debezium/debezium/pull/229] that removes the recommenders. It's not clear how valuable these recommenders actually are. First, it's not clear about the expected semantics: can the user use values that don't appear in the recommended values? Second, the recommenders that return large numbers of values can be slow and can result in very large REST API responses. Debezium was using recommenders to return the database and table/collection names, but these lists can be very large for large databases. Rather than cap the number of recommended values and have the recommender return a subset of all potential values, we will instead remove the recommenders altogether.	1	1	https://issues.jboss.org/browse/DBZ-232
3117	DBZ-222	[~gunnar.morling], shouldn't this be targeted to 0.5.1 if it is already merged?	1	1	https://issues.jboss.org/browse/DBZ-222
3118	DBZ-211	On the decoderbufs end adding filtering is easy. In fact I already hacked this in for my database, as 99.9% of all updates are from unmonitored tables, and it improves performance substantially. The main question there is how the filter should be specified, eg should the client provide a list of the tables it is interested in, or should it pass in some whitelist/blacklist regexp filter expressions, eg the same as those used to configure the connector in the first place? I think I prefer the explicit list of tables, but the regexp expressions would likely be easier to implement.	1	0	https://issues.jboss.org/browse/DBZ-211
3119	DBZ-211	There may be a flaw in this scheme of only outputting data for rows in the relevant table: if the output plugin doesn't output anything, will the client still get the LSN? If it doesn't, it can't tell the server that it reached that LSN, so the server won't understand that WAL can be discarded, and WAL will build up until one of the monitored tables gets an update. In a busy database where the monitored tables are not updated much this could lead to a huge WAL buildup (which is what I'm seeing in my testing - not 100% sure that this is the cause yet though).	1	0	https://issues.jboss.org/browse/DBZ-211
3120	DBZ-209	[~rahulravindran0108], thanks for logging this. Perhaps I'm missing something, but does the Debezium PostgreSQL connector need to do something different to handle this case? And I thought Amazon RDS doesn't currently allow you to add your own logical decoding plugin, which means that the Debezium PostgreSQL connector can't currently be used with Amazon RDS running PostgreSQL.	0	0	https://issues.jboss.org/browse/DBZ-209
3121	DBZ-199	I'm not ready to say "no", but I think I need more convincing before "yes". :-) Specifically, I'm not clear about the advantages of a _separate_ repository, probably because we're using it very simply and within the same repo. I understand you'd like to reuse it, but do you need to support multiple Kafka versions, and you'd like to have a separate artifact per Kafka version?	1	0	https://issues.jboss.org/browse/DBZ-199
3122	DBZ-199	I guess I still don't understand the benefits of a separate repo over the current approach. What would having it in a separate repo buy you?	1	0	https://issues.jboss.org/browse/DBZ-199
3123	DBZ-196	[~p5k6], I've marked this issue on the MySQL connector. Just to clarify, can you confirm that you've not run into something similar with the PostgreSQL connector?	1	1	https://issues.jboss.org/browse/DBZ-196
3124	DBZ-177	[~hchiorean], can you provide a PR that changes the PostgreSQL connector to do what this PR did for the MySQL connector?	1	1	https://issues.jboss.org/browse/DBZ-177
3125	DBZ-175	[~anton_nazaruk], have you tried updating any other properties? The problem with updating {{table.whitelist}} is that currently the connector does not support performing anything but a snapshot upon connector creation with the _initial_ list of filtered tables. IOW, if you change the filters (either via whitelist of reducing scope of the blacklist) to add more tables than the initial set, the connector will _not perform additional snapshots_. Additionally, whether or not the connector "sees" the table(s) newly-included by the filter changes depends a lot on when the tables' {{CREATE TABLE}} statements were in the database history. If not, then they're essentially missing from the history and no subsequent changes on the tables will be processed. If the connector configuration update is specific to the above limitation not handling the newly-included tables, then we can rescope this issue to enhance the connectors to support secondary snapshots. However, if there are still connector configuration update issues, then we can keep this issue and probably have a different issue to track the new feature for performing additional snapshots when needed.	1	1	https://issues.jboss.org/browse/DBZ-175
3126	DBZ-175	I didn't try to touch other props, *white* list tables manipulation was the main target. I should probably have named this issue differently, exactly like you explained with "inability to perform secondary snapshots". Thanks for pointing to DBZ-165, looks like we ended up with similar workarounds)) I agree this functionality shall be treated as an "improvement". Anyway, it will be very useful one. We have thousands of tables (in dozens of databases), and only small fraction of them will go through debezium (20-25%). Anyway, when we move forward with materialized views usage (built from CDC events) more and more use cases appear and new tables are needed there. And we can't load all of them with initial snapshot because of many restrictions. Can you roughly estimate time efforts you need to get it done? Also, what is a policy for code contribution? Thanks in advance.	1	1	https://issues.jboss.org/browse/DBZ-175
3127	DBZ-175	[~anton_nazaruk], thanks for responding. I concede that this would be a useful feature and that it should be possible to do. Perhaps the biggest challenge is going to be coming up with the desired functional behavior. For example, is it acceptable/sufficient to perform a complete snapshot whenever new tables are needed, or is it necessary that the connector only perform incremental snapshots on the newly-visible tables? The latter is more attractive and makes more sense, but it's harder to implement such that the output change events are entirely consistent. {quote} Also, what is a policy for code contribution? {quote} We accept contributions from anyone, so any help with this would be really appreciated. See http://debezium.io/community/ for details.	1	1	https://issues.jboss.org/browse/DBZ-175
3128	DBZ-175	I think the described algorithm is sound, i've got one question, though. Do we actually have to converge to one joint binlog reader in steps 4 and 5? Or would it be possible to simply let the binlog reader set up in step 2 run for the time the connector is running and let it continue to process the tables from S1? And after step 3, another, separate binlog reader would be set up for processing S2. When the connector is restarted eventually, and there's no further change to the set of tables, a single binlog reader would be started for S3. This could be a bit less complex. But I'm not sure whether it's possible (or perhaps possible but resulting bad performance) to have multiple binlog readers?	1	1	https://issues.jboss.org/browse/DBZ-175
3129	DBZ-175	One question - cannot we have a BinlogReader running an conssuming all elevents but the events that belongs to S2 would be directed into a temporary topic. When the snapshot is completed events taken between snapshot and current time would be moved from temporary topic into the final destination and then thhe BinLogReader would be redirected to the final destination again?	1	0	https://issues.jboss.org/browse/DBZ-175

#	Requirement	Questions	Answered	Successfully Answered	URL
3130	DBZ-175	I have a slightly modified version of the proposal in [randall's comment|https://issues.jboss.org/browse/DBZ-175?focusedCommentId=13418488&page=com.atlassian.jira.plugin.system.issuetabpanels%3Acomment-tabpanel#comment-13418488] for how to handle the new table snapshots. Instead of attempting to immediately calculate S2 on startup, which is complicated[1], I propose we in-line snapshot generation during binlog reading. When we encounter a (non-create) event for some table A that we should deal with that we don't have schema history for, we can run a snapshot of the table. This snapshot will occur in a new thread. This snapshot will get a transaction on the table it is snapshotting which will be at some offset O1. -The binlog reader will continue, ignoring any A events until it reaches or exceeds O1. If the binlog reader reaches O1, it will begin storing any events for A in a queue. Once the snapshot for A completes, the binlog reader will run through all the events in the queue before processing events normally again, but now including events for A normally.- Edit: I realize that this approach would force the binlog reader to block until the snapshot is able to start a transaction, which is unideal. Sticking with the idea of re-processing some events as needed once the snapshot completes is better. It is possible that multiple tables will be added at the same time. In this case, we will only begin snapshotting the second table once the first is completed. I can think of 2 ways to do this: 1. while a snapshot for table A is running, ignore any events of any other new table B. Once snapshot A is completed, the next event for table B we see will kick off a snapshot for table B. 2. while a snapshot for table A is running, if we see an event for a new table B, we can put table B in a queue. Once table A's snapshot is completed, we can look at the queue to see if any tables are due to be snapshotted. [1]Why is calculating S2 on startup complicated? It isn't always complicated, but there are some edge cases that would force a heavyweight process to try to calculate S2 on each startup. I'll describe one such edge case. Lets say we have a config that uses blacklists to enumerate the tables we want to capture in DBZ. Now imagine the connector goes down for some long period of time. During this time, a new table A is added to the database, but the CREATE event for it falls off the binlog. When the connector starts back up again, no config changes have occurred, but we will need to snapshot the new table A. The only way to discover this is by querying the database and re-calculating the set of tables we are expected to cover. Since we want to be able to snapshot at the same time as we read the binlog, there does not seem like a strong reason to make sure the snapshots occur at startup, since they will be occurring in parallel to binlog reading anyway. The primary risk I can see from this approach is that if some table is added that has no or very few events related to it, it may take a long time (or, possibly never) for a snapshot to get triggered. However, I feel like this is a very unlikely event. Additionally, if this is a well-documented issue, it would be very easy to trigger a snapshot simply by adding a dummy row (and then deleting it) or "modifying" an existing one. Are there any issues that this proposal overlooks? Feedback would be welcome. Edit: Looking more into it, I'm more confused about, what, exactly needs to be done on startup (it seems like something needs to be done on startup...?)	1	0	https://issues.jboss.org/browse/DBZ-175
3131	DBZ-175	Thanks, [~mtagle], for the proposal! I have a few comments/questions: Schema history includes tables that don't satisfy the filter. Why? It's complicated, but for the most part boils down to the fact that we might need that schema history later. For example, if table {{bar}} doesn't satisfy the filters, we might come across {{CREATE TABLE foo AS SELECT FROM bar}}, at which point we need the schema for {{bar}} just to know what table {{foo}} looks like. One of the challenge as I see it is how a connector can determine when it restarts that its filters now include tables that weren't included the last time it was running. I think we may be able to solve this by simply adding the 2 database and 2 table filter properties to the offset. I was just looking through the Kafka Connect code to verify that it does very little with the offset map until it comes time to flush the most recent offset per partition. So adding more fields to that offset map is really not much overhead whatsoever, but it does give the connector the ability when starting up to know what the filters were the previous startup. With that info, we can compare the previous and current filters, and: * continue as-is if they are identical, or * use the previous filters to calculate S1, the new filters to calculate S3, and then calculate S2. Now, you make a good point about edge cases making the computation of S2 difficult. Rather than store the 4 filter properties, would could instead store in the offsets one additional property that contains the names of the actual tables. When the connector starts, it could calculate the set of tables that satisfies the current filter (S1) and generate a single string containing all of the fully-qualified table names. It could then pass that string in all of the offsets, and as mentioned before there is little overhead to this. Upon restart, the connector can read that string and parse into fully-qualified table names for S1, use the new filters to calculate S3, and then determine S2 and snapshot as needed. The only real disadvantage of this approach is that this string could get long, though that's not really that big a deal. One significant advantage, though, is that this is very backward compatible. Upon restart, if the connector doesn't see this property it simply skips the check to determine if snapshots are needed, but then always generates this new property in the offsets. The *next* time the connector restarts, any changes in the sets of tables will be discovered correctly. This does require that people upgrade before changing any filters, but from then on they could change filters at any time. Now, I do like your idea about watching during binlog reading for "new" tables that satisfy the filter and then snapshotting them. This could get fairly complicated, so we'll have to watch that. Hopefully we can use the same snapshot mechanism to do this. What do you think?	1	1	https://issues.jboss.org/browse/DBZ-175
3132	DBZ-175	[~rhauch] that general approach for calculating S2 sounds good to me. When you talk about the "4 filter properties" do you mean database whitelist/blacklist and table whitelist/blacklist?	1	0	https://issues.jboss.org/browse/DBZ-175
3133	DBZ-164	[~tanooren], what do you think about just catching and handling the exception for a database and skipping it? I'm hesitant to exclude tables from the DDL step because the connector currently always processes all DDL statements, even if the databases and tables are excluded.	1	1	https://issues.jboss.org/browse/DBZ-164
3134	DBZ-164	Here's a bit more detail. Although databases are excluded as a whole, it is still possible for some tables in one database to be defined using tables in another database, such as: {code} CREATE TABLE a LIKE b; {code} where tables {{a}} and {{b}} can be in different databases. Right now, the connector processes all database- and table-related DDL statements (e.g., view and procedure definition statements are skipped), and so when this situation occurs the connector still knows what the definition of of table {{b}} might be, even if it is in a database that has been filtered out. I'd rather not add yet-another-configuration-property that controls the behavior, so my proposal is to continue look for the tables in all databases but to instead better handle any error with each database and simply skip everything to do with that database in the rest of the snapshot. We'd still log that we're skipping the database name and include the SQL exception message (but not stack trace). [~tanooren], thoughts?	1	1	https://issues.jboss.org/browse/DBZ-164

#	Requirement	Questions	Answered	Successfully Answered	URL
3135	DBZ-151	h3. 1. Execution of the tests We have to determine how to run individual tests locally (from the Maven command line and from an IDE) and how to run all of the integration tests using Maven's Failsafe. I expect that the best/easiest way to do this is to require a separate test method in a JUnit test class that specifies the directory and runs the connector test. This should require very little code, but it would work very nicely with the Maven build and IDEs using current approaches. As long as we keep each method implementation to at most a few lines, this does not seem to be much of a burden. h3. 2. Organization of the test classes So if we do this, then where should this test class (or test classes) be placed? If we follow current conventions, we'd simply place them inside the connector Maven module. However, this makes it much harder when wanting to run the integration tests against multiple DBMS containers. For example, when running the MySQL build with the {{-Passembly}} argument, the build starts the integration test phase of the build by starting _all Docker containers_ that will be required by the tests, and then runs the integration test suite multiple times (once for each DBMS topology). Overall this works pretty well, and despite the {{pom.xml}} complexity all of the container's are built using shared artifacts (e.g., SQL files, config files, etc). A major disadvantage, however, is that as the number of containers increases this becomes harder to maintain and requires more local resources during the build to run all of the containers at once. An alternative is to create one or more test modules for each connector, where each test module would apply to a single DBMS topology. So a connector with tests for 5 DBMS topologies would mean 5 separate test modules for the connector. Each {{pom.xml}} file becomes simpler, but each test module for a connector has a lot of duplicate code: the {{pom.xml}} would largely be the same, the DBMS config files, the connector configuration files, the expected results. However, this may actually be a good thing, since different DBMS topologies may have subtly different config files, connector config files, and expected results. h3. 3. Output Our current integration build results in significant output, even at the INFO level, simply because of the number of test runs of the connectors. Normal builds shouldn't need any of this output, but developers should be able to easily enable logging when running the tests individually. h3. 4. Configuration and expected results If we use a separate directory for each test, then we could place in the directory: * a properties file (with a well-known name) holding connector's configuration; and * a file containing the expected output of the connector, perhaps in JSON format. It is likely the test framework can read in the expected events and compare sequence of output and expected events. We may have to provide a way to ignore certain event fields (e.g., the timestamp that the event was generated or that the row was changed), but the goal would be that this could be automated completely and require very little customization for each connector. Another benefit of having expected results stored in a file is that Git can manage the history of the output, so any changes to the connector output is captured. A nice feature would be for the test framework to _output_ the expected results file if it does not exist. This will save the developer from having to hand-code the expected results. This would also be great if coupled with the ability to *generate* larger volumes of data in the source databases, as the expected outputs can be generated and manually inspected. It is often useful to test the stopping and restarting of the connector. Could this be accomplished by having special _control statements_ in the expected output file? For example, this line in the expected output file: {code} { "connector" : "restart" } {code} might instruct the connector to stop and restart after the change event on the previous line were output, after which the output should continue with the change events that follow. This might make it theoretically very easy to verify that connector restarts do not lose/miss/duplicate any output events.	1	1	https://issues.jboss.org/browse/DBZ-151
3136	DBZ-151	[~hchiorean], are you far enough with the PostgreSQL connector to have any feedback or input?	1	1	https://issues.jboss.org/browse/DBZ-151
3137	DBZ-151	Hopefully PG will remain that straightforward. Might there be a need to run tests against multiple PG topologies, for example to test failing over a connector from one PG server to another and verifying no data is lost? This is the case for MySQL (e.g., single w/o GTIDs, single w/ GTIDs, master & slave, multi-master slave, etc.) and MongoDB (e.g., simple replica set, sharded cluster).	1	1	https://issues.jboss.org/browse/DBZ-151
3138	DBZ-147	Awesome! Do the docs get updated automatically?	1	0	https://issues.jboss.org/browse/DBZ-147
3139	DBZ-143	Note: this was tested with MySQL 5.7.12 community edition	1	1	https://issues.jboss.org/browse/DBZ-143
3140	DBZ-140	[~rhauch] I'm running into this issue as well, but I was wondering why acquiring a global lock with {{FLUSH TABLES WITH READ LOCK}} is needed here. Would it not be sufficient to acquire locks on only the tables we're interested in by first looking at the whitelist/blacklist configs, regardless of whether or not RDS is being used? This way we could avoid the weirdness with checking {{@@basedir}} altogether.	1	1	https://issues.jboss.org/browse/DBZ-140
3141	DBZ-140	[~konstantin.zadorozhny], do you have any thoughts about this? In particular, what might we change in the connector to keep the strict consistency logic when talking to MySQL but allow it to work for RDS?	1	1	https://issues.jboss.org/browse/DBZ-140
3142	DBZ-140	[~richoconnell], thanks for using Debezium and for reporting the connector's problem using RDS as well as providing a patch. Your fix should absolutely work and looks great. Do you have any interest in creating a pull request with the change? While a patch is perfectly acceptable, the history in Git doesn't show you as contributing that fix. I do think it's great if contributor's changes _are_ reflected in the history, and creating a PR would do that. But it's up to you.	1	1	https://issues.jboss.org/browse/DBZ-140
3143	DBZ-137	Hi Randall, [~everton.portela] : Nice. I am very interested in Oracle CDC and have been long awaiting the successor for DataBus from LinkedIn, I understand LinkedIn is working on an update to Databus. Glad to finally see an Open-Source solution for Oracle CDC. It may be a while before I get a chance to get my hands dirty with this. Can you keep us updated as you progress. I started looking into the StreamSets' Oracle CDC implementation on GitHub : [OracleCDCSource.java|http://bit.ly/2oeMZXp]. Up until now, I thought Redo Logs mining was difficult and due to complexity it wasn't easy to even get a handle on a mechanism to read it. Thanks to [DBZ-137|https://issues.jboss.org/browse/DBZ-137] I got introduced to LogMiner. Is it something that will support Oracle RAC?, wonder what performance overhead of enabling LogMiner on the database would be. From reviewing Oracle 12c LogMiner [documentation |https://docs.oracle.com/cd/B28359_01/server.111/b28319/logminer.htm#i1014380], there are special considerations for Oracle RAC enabled databases and some [concerns |http://www.freelists.org/post/oracle-l/CDC-using-LogMiner-in-RAC-environment,1]: # [Explicitly specifying log files|http://docs.oracle.com/database/122/SUTIL/oracle-logminer-utility.htm#SUTIL1605] for each redo thread. There does appear to be queries that can help query the log files. I'm not an Oracle Expert as you can probably tell, but I do appreciate if someone with Oracle expertise can help clarify these options for RAC and how [V$LOGMNR_LOGS|https://docs.oracle.com/html/E10643_07/rcmsubcl002.htm] works in RAC. # Continuous Mine option to avoid explicitly specifying log files. Option is allowed for Oracle RAC only if redo threads are not enabled. : See [Example LogMiner session without specifying logfiles explicitly|https://docs.oracle.com/cd/B28359_01/server.111/b28319/logminer.htm#i1027809] and [Future query values|https://docs.oracle.com/cd/B28359_01/server.111/b28319/logminer.htm#i1028809] From just a glance I have been able to take of the [OracleCDCSource.java|http://bit.ly/2oeMZXp] I see they use the Continuous Mine option. Yet to determine if that is sufficient for Oracle RAC. Thanks, Pavan	1	0	https://issues.jboss.org/browse/DBZ-137

#	Requirement	Questions	Answered	Successfully Answered	URL
3144	DBZ-137	Hey [~jjose2004], interesting, hadn't heard about dbms_change_notification before. Do you have a pointer to what the "CQ" is? On dbms_change_notification, do you know whether there's any performance implication of this? I.e. would it'd be suitable for implementing CDC on large number of tables? Also would we be able to capture old and new state of a modified row?	1	0	https://issues.jboss.org/browse/DBZ-137
3145	DBZ-134	[~vhorodchuk], do you have any specific metrics in mind?	1	1	https://issues.jboss.org/browse/DBZ-134
3146	DBZ-133	There can be scenario where previous snapshot failed during taking data snapshot and connector was restarted with schema snapshot mode. What should be the behaviour ? Just throw an error message and stop or continue ? and this should include check for row level bin logs too ?? } else if (taskContext.isSchemaSnapshotOnly() && rowBinlogEnabled) { this.snapshotReader.includeData(false); }	1	1	https://issues.jboss.org/browse/DBZ-133
3147	DBZ-133	{quote} and this should include check for row level bin logs too ?? {code} } else if (taskContext.isSchemaSnapshotOnly() && rowBinlogEnabled) { this.snapshotReader.includeData(false); } {code} {quote} Okay, good point. This logic needs a bit of refactoring to handle the various cases more clearly. How about changing this entire if-block to: {code:java} if (taskContext.isInitialSnapshotOnly()) { logger.warn("This connector will only perform a snapshot, and will stop after that completes."); this.snapshotReader.onSuccessfulCompletion(this::skipReadBinlog); } else { // We're expecting to read the binlog ... this.snapshotReader.onSuccessfulCompletion(this::transitionToReadBinlog); if (!rowBinlogEnabled) { throw new ConnectException("The MySQL server is not configured to use a row-level binlog, which is " + "required for this connector to work properly. Change the MySQL configuration to use a " + "row-level binlog and restart the connector."); } } {code} IOW, if we're performing an initial snapshot only, then stop after the snapshot; otherwise, always try to follow up by reading the binlog and fail if that is not the case. Then rather than {{SnapshotReader}} exposing the boolean flag, change the {{SnapshotReader}}'s constructor to then do this: {code:java} this.includeData = !taskContext.isSchemaSnapshotOnly(); {code} The benefit is that the {{includeData}} boolean field can then be {{final}}, which is a huge plus since immutable is much better.	1	1	https://issues.jboss.org/browse/DBZ-133
3148	DBZ-133	@rhauch ..for schema based snapshot, what i wanted to know was that suppose i start connector in snapshot.mode = initial while taking data snapshot it went off or restarted or config was modified..now it was restarted with snapshot.mode = schema, it will again add create tables and all this is fine, but this time it will not take data snapshot , so my kafka queues will have some old data snapshot (incomplete) and now will start reading from current binlog position. Is this behaviour fine ?	1	1	https://issues.jboss.org/browse/DBZ-133
3149	DBZ-133	[~prannoymittal], any progress on this? Do you want to create a pull request for review? You can always update the PR if needed.	1	1	https://issues.jboss.org/browse/DBZ-133
3150	DBZ-126	Will this require that [com.github.shyiko:mysql-binlog-connector-java|https://github.com/shyiko/mysql-binlog-connector-java] need to be updated to support this?	1	1	https://issues.jboss.org/browse/DBZ-126
3151	DBZ-126	{quote} Will this require that com.github.shyiko:mysql-binlog-connector-java need to be updated to support this? {quote} No, that library allows for some customizations that we're already using. If that means we need to duplicate some logic, then we may want to contribute an enhancement to the library. Either way, we can make it work.	1	1	https://issues.jboss.org/browse/DBZ-126
3152	DBZ-121	[~renwu58], can you share an example of a regex that you're using. The above samples for the source tables are great, but it's not clear how they map to topic names. For example, I can imagine some mappings between source tables and topics to be like these: ||Source Table Pattern||Topic Name Pattern|| | {{\{base_table\}_\{tenant\}_\{YYYYMM\}}} | {{\{base_table\}_\{tenant\}}} | | {{\{base_table\}_\{tenant\}_\{YYYYMM\}}} | {{\{base_table\}}} | | {{\{base_table\}_\{seqno\}}} | {{\{base_table\}}} | | {{\{base_table\}_\{tenant\}_\{seqno\}}} | {{\{base_table\}_\{tenant\}}} | Is this close? Can you give a few more specific examples? Would a single regex-based mapping be able to identify all of the source tables that need to be merged, or do we need to allow multiple regex-based mappings?	1	1	https://issues.jboss.org/browse/DBZ-121
3153	DBZ-121	[~renwu58], do the different shards have overlapping primary keys? If so, then the events key (which currently consists only of the table's primary key) may need to be changed to include a "shard ID" or "shard name". BTW, properly merging sharded tables affects a few more things than just using the same topic for a set of matching (sharded) tables, and we may need a bit more time to make sure it's done properly and that it will be useful. It's currently scheduled for 0.3.2, but I was hoping to get 0.3.2 out this week so I may push this into 0.3.3 so that I can release 0.3.2 on Monday. There is not currently a timeframe for 0.3.3, so we could do it relatively quickly after this feature is completed (and perhaps after DBZ-126 with JSON data type).Thoughts?	1	1	https://issues.jboss.org/browse/DBZ-121
3154	DBZ-121	[~renwu58], with this new approach, it should be possible to achieve your goal by specifying one or more regular expressions and replacement patterns. Note that the regular expressions must match the _entire fully-qualified table name_, not just the table name. For example, you could specify the new configuration property like this: {code} topic.naming.strategy=^(\w*?)[.](\w*?)(?=(?:_(?:\d+)|$)+$)(_(?:21|22|51|52|91))?.* = $1$2, table {code} Note this defines two strategies: # The first uses the regular expression {{^(\w*?)\[.\](\w*?)(?=(?:_(?:\d+)\|$)+$)(_(?:21|22|51|52|91))?.*}} to match _fully-qualified table names_, and when they match capture groups 1 and 2 are used to construct the topic name; and # If no previous strategy matches for a given table, then the {{table}} strategy is used as a default Or, if it's easier you could write and specify a class that implements the {{TopicNamingStrategy}} interface with custom logic. What do you think?	1	1	https://issues.jboss.org/browse/DBZ-121
3155	DBZ-115	[~uchouhan], the "Marking coordinator ... dead" is a signal that the Kafka Connect process has lost (or never had) a connection to the Kafka broker. Are there other messages in either the Kafka broker or Kafka Connect processes?	1	1	https://issues.jboss.org/browse/DBZ-115

#	Requirement	Questions	Answered	Successfully Answered	URL
3156	DBZ-115	[~uchouhan], can you please confirm that you're using a single broker and a single Kafka Connect process with a single connector? The log messages you mention are similar to [https://issues.apache.org/jira/browse/KAFKA-2985|KAFKA-2985], so you might try looking at some of the [optional Kafka Connect worker settings|http://docs.confluent.io/3.0.0/connect/userguide.html#configuring-workers] or [overriding the Kafka producer and consumer settings|http://docs.confluent.io/3.0.0/connect/userguide.html#overriding-producer-consumer-settings], both of which can be set on the Kafka Connect docker image via {{-e CONNECT_{property_name_using_underscores}}}. For example, do any tables have large numbers of columns and/or columns with very large values? The value of the {{max.partition.fetch.bytes}} Kafka consumer setting controls the maximum size of each response, which may contain one or more messages. However, if one event message is large enough, it might be greater than this maximum size. The [http://kafka.apache.org/documentation.html#newconsumerconfigs|Kafka documentation] says of this property: {quote} The maximum amount of data per-partition the server will return. The maximum total memory used for a request will be {{#partitions * max.partition.fetch.bytes}}. This size must be at least as large as the maximum message size the server allows or else it is possible for the producer to send messages larger than the consumer can fetch. If that happens, the consumer can get stuck trying to fetch a large message on a certain partition. {quote} The default value for this property is 1048576 bytes, or 1Mi. So, you may try to bump up this value if any of your rows are large enough to be serialized (along with the schema) to be larger than 1Mi. I'm guessing that you didn't change the converter, so this serialized format would be in terms of JSON (which can get pretty verbose). There is a similar {{max.request.size}} configuration property for the Kafka producer, and it also defaults to 1Mi.	1	1	https://issues.jboss.org/browse/DBZ-115
3157	DBZ-115	[~uchouhan] and [~rjharmon0316], can both of you explain the process with every step? For example, when are you stopping the connector, when are you restarting the connector, is the MySQL server stopped or restarted, etc. And, it'd be helpful to attach the log output for each time the connector (obviously with any sensitive information redacted/removed) is started; I'm particularly interested in the number of tables, rows found in each table, and the binlog positions at various points as captured during the snapshot phase and binlog reading phases.	1	1	https://issues.jboss.org/browse/DBZ-115
3158	DBZ-115	[~uchouhan], thanks for the info. The {{Error registering AppInfo mbean}} is actually a WARNING, so it can be ignored. I'm not sure why Kafka Connect is doing this, but I'll see if that's something we can fix (see DBZ-124). {quote} Also, notice that the binlog filename and position are blank, whereas I suppose that should be valid. I have run into the same issue with no events on the kafka topic even when there was a binlog filename/position in this log ... {quote} Do you have the row-level binlog enabled on MySQL? The connector will output some related MySQL server properties, like: {code} ... 2016-09-16 09:24:55,786 INFO MySQL|myServer|snapshot MySQL server variables related to change data capture: [io.debezium.connector.mysql.SnapshotReader] 2016-09-16 09:24:55,795 INFO MySQL|myServer|snapshot binlog_cache_size = 32768 [io.debezium.connector.mysql.SnapshotReader] 2016-09-16 09:24:55,796 INFO MySQL|myServer|snapshot binlog_checksum = CRC32 [io.debezium.connector.mysql.SnapshotReader] 2016-09-16 09:24:55,796 INFO MySQL|myServer|snapshot binlog_direct_non_transactional_updates = OFF [io.debezium.connector.mysql.SnapshotReader] 2016-09-16 09:24:55,796 INFO MySQL|myServer|snapshot binlog_error_action = ABORT_SERVER [io.debezium.connector.mysql.SnapshotReader] 2016-09-16 09:24:55,796 INFO MySQL|myServer|snapshot binlog_format = ROW [io.debezium.connector.mysql.SnapshotReader] ... {code} This will indicate whether you're using the binlog format. Please verify that it is {{ROW}} and that you've properly set up permissions according to http://debezium.io/docs/connectors/mysql/	1	1	https://issues.jboss.org/browse/DBZ-115
3159	DBZ-115	Yes, it is row and I'm not sure if it does not show that information because I'm using snapshot.mode 'never'. I'm using 0.3.1, should I be trying out the latest from git? If so, is there a README I can follow to run the latest code?	1	1	https://issues.jboss.org/browse/DBZ-115
3160	DBZ-115	I think that might be the problem. I purged the mysql logs before a certain date and so its likely that the create table statements do not exist. Do you have any suggestions on how I can go about fixing this?	1	1	https://issues.jboss.org/browse/DBZ-115
3161	DBZ-115	{quote} I think that might be the problem. I purged the mysql logs before a certain date and so its likely that the create table statements do not exist. Do you have any suggestions on how I can go about fixing this? {quote} It is very likely that _any_ MySQL server that's been running a while will have some part of its binlog purged, and this is exactly why the connector supports taking a snapshot and _then_ reading the binlog from the time at which the snapshot was started. Getting a connector to perform a snapshot after it's already completed a snapshot is a bit challenging, but IIRC you're using Kafka in a Docker container, so this will easier for you to try. Or, just start a new connector with a different logical name and snapshot mode. BTW, [our tutorial|http://debezium.io/docs/tutorial/] goes through this exact process, talks about the log output and what each step is doing, and even walks you through stopping and starting the MySQL connector. Be careful trying to stop and restart the MySQL server running in a Docker container, though -- our tutorial doesn't persist the data outside the container, so once the MySQL container stops all data is lost.	1	1	https://issues.jboss.org/browse/DBZ-115
3162	DBZ-115	[~rjharmon0316], what version of MySQL are you using? What version of Debezium? Several things could be causing the problem. First, please check whether the MySQL user is properly configured to be able to read all of the binlog information; see the [documentation|http://debezium.io/docs/connectors/mysql/] for specifics. The MySQL connector outputs information about the related system variables and the user (with their grants) used for snapshot and binlog reading, so that should help. Have you verified that the offsets reported by the connector are valid offsets in the binlog? If not, that would take this in a very different direction. Did you look to see the changes at or near that point in the binlog to see if there's anything interesting? Also, you might want to double check that changes are being made to your table after that point. Third, have you tried running the connector without the database filter? The filtering logic should be the same for both snapshot and binlog modes, but there may be some subtle differences such that a bug in the filtering is causing the binlog mode to ignore events from all the tables. Fourth, you did get it to work using the Docker images? If so, I presume that was using the JSON converter, or did you rebuild the Docker image to support the Avro Converter? Fifth, the Debezium community to my knowledge doesn't test with the Confluent Platform. You may talk to them to find out what if any differences their are between their Platform and the Kafka Connect that's released with Kafka. For example, have they patched anything? Our 0.3.1 release uses Kafka Connect 0.10.0.1, but IIUC only Confluent Platform 3.0.1 uses this same Kafka Connect version.	1	1	https://issues.jboss.org/browse/DBZ-115
3163	DBZ-115	[~rjharmon0316], one more thing. Can you provide the properties that specify your database and table filters? (Change the names if you want to sanitize them.) The more I think about it, the more I wonder if the filtering is a problem.	1	1	https://issues.jboss.org/browse/DBZ-115

#	Requirement	Questions	Answered	Successfully Answered	URL
3164	DBZ-115	{quote} FYI, I used these dbz settings when creating the Connect service: {code} database.server.name=db database.whitelist=inventory database.history.kafka.bootstrap.servers=kafka:9092 database.history.kafka.topic=dbh.inventory {quote} I thought you said earlier that you're setting a table filter. If so, did you take that out? (Or perhaps I was mistaken.) {quote} The "my_connect_offsets" topic reflects new binlog offsets (with a ~30s delay) after ALTER TABLE in the single whitelisted database, but is never updated on INSERT, UPDATE or DELETE. {quote} Okay, so the {{my_connect_offsets}} topic is where Kafka Connect is _periodically_ recording the MySQL connector's "source offsets", which contains the most recently recorded binlog filename and position (among other things). You can control the frequency at which Kafka Connect records these with the {{offset.flush.interval.ms}} worker configuration property (see [Confluent's documentation for Kafka Connect|http://docs.confluent.io/3.0.0/connect/userguide.html#configuring-workers]). There are a couple of possible causes for not seeing any INSERT, UPDATE or DELETE events for your table: # The connector may be incorrectly skipping the binlog events associated with that row if the table filter or database filter are not specified correctly or if there's a bug in the binlog reading logic. Debug logging will output more information if this is the case. You can turn on debug logging on the Kafka Connect process in Docker by setting the {{LOG_LEVEL}} environment variable to {{DEBUG}}; if you're using Confluent then just crank up the debug logging for {{io.debezium.connector.mysql}} or if that's too verbose then debug for just {{io.debezium.connector.mysql.BinlogReader}}. You can see the [{{handleInsert}} method|https://github.com/debezium/debezium/blob/v0.3.1/debezium-connector-mysql/src/main/java/io/debezium/connector/mysql/BinlogReader.java#L342] outputs either "Recorded {n} insert records for event ..." or "Skipping insert row event ..." messages for INSERTs, and similar messages for DELETE and UPDATE. # The connector may be getting confused about the structure of the table. Presumably this is not the case since you get events from the snapshot, meaning the connector has an internal model of the table and its columns. It is possible that an ALTER screwed up this internal model, however, though in the past that has always resulted in an exception about failing to parse DDL. Regardless, if this is the case you would see "Skipping ..." debug log messages as mentioned above. # The connector is not able to actually parse the binlog events for MySQL 5.5. We've done testing with 5.6 and 5.7, but not with 5.5. If this is the case, then if you use TRACE logging then the connector is calling the [{{ignoreEvent}} method|https://github.com/debezium/debezium/blob/v0.3.1/debezium-connector-mysql/src/main/java/io/debezium/connector/mysql/BinlogReader.java#L186] and you would see messages like "Ignoring event due to missing handler...". Either way, it seems like the next step is to turn up logging to DEBUG to see if 1 or 2 are more likely, or if that's not clear turn up logging to TRACE to see if 3 is more likely. BTW, thanks a lot for your patience trying to figure this out and working with me to hopefully resolve the situation!	1	1	https://issues.jboss.org/browse/DBZ-115
3165	DBZ-115	{quote} {code} 2016-09-20 23:16:37,767 INFO MySQL|db|binlog Connected to MySQL binlog at 192.168.1.132:3306, starting at binlog file 'mysql-bin.000208', pos=107, row=0 [io.debezium.connector.mysql.BinlogReader] 2016-09-20 23:16:37,768 DEBUG || Nothing to send to Kafka. Polling source for additional records [org.apache.kafka.connect.runtime.WorkerSourceTask] 2016-09-20 23:16:37,786 DEBUG MySQL|db|binlog Rotating logs: Event{header=EventHeaderV4{timestamp=0, eventType=ROTATE, serverId=1, headerLength=19, dataLength=24, nextPosition=0, flags=32}, data=RotateEventData{binlogFilename='mysql-bin.000208', binlogPosition=107}} [io.debezium.connector.mysql.BinlogReader] 2016-09-20 23:16:37,788 DEBUG MySQL|db|binlog Clearing table converters [io.debezium.connector.mysql.RecordMakers] 2016-09-20 23:16:43,119 DEBUG || Received successful heartbeat response for group 1 [org.apache.kafka.clients.consumer.internals.AbstractCoordinator] {code} {quote} This is interesting. The {{Rotating logs...}} message signals that a [ROTATE event|https://dev.mysql.com/doc/internals/en/rotate-event.html] was found, which is usually the last event written to a binlog file before switching to a new one (e.g., {{mysql-bin.000208}}). Considering that this is the last event mentioned in the log, my guess is that the connector is not able to find this log and for some reason we're either spinning or we're blocked waiting for something to happen. I'd love to know where it is stuck so we can prevent this from happening again. Would you mind checking the binlog files to see if {{mysql-bin.000208}} is still there? Then, if you're still willing to help track this down, could you get a [stack dump|http://stackoverflow.com/questions/10756105/how-to-get-a-complete-stack-trace-of-a-running-java-program-that-is-taking-100] of the Kafka Connect process? Unfortunately, if Kafka Connect is running in a Docker container, you'll have to connect to it with a shell first: {code} docker exec -it connect /bin/bash {code} where {{connect}} is the Kafka Connect container name. Much appreciated!	1	1	https://issues.jboss.org/browse/DBZ-115

#	Requirement	Questions	Answered	Successfully Answered	URL
3166	DBZ-115	restarting connect container (see second stack). Finally, I used a shell loop to capture a series of 60 traces during connect startup, and another one to grep -A 20 -B 20 'ebezium' from those stacks; results in 3rd code block. Please let me know if you want any specifically-numbered trace from that series. kafka@35eea237734d:~$ jstack 1 {code} 2016-09-21 15:19:48 Full thread dump OpenJDK 64-Bit Server VM (25.92-b15 mixed mode): "Attach Listener" #47 daemon prio=9 os_prio=0 tid=0x00007fdfc80b3000 nid=0x246 runnable [0x0000000000000000] java.lang.Thread.State: RUNNABLE "blc-keepalive-192.168.1.132:3306" #46 prio=5 os_prio=0 tid=0x00007fdfd4da7000 nid=0x229 waiting on condition [0x00007fdfc1ce6000] java.lang.Thread.State: TIMED_WAITING (sleeping) at java.lang.Thread.sleep(Native Method) at com.github.shyiko.mysql.binlog.BinaryLogClient$3.run(BinaryLogClient.java:555) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:511) at java.util.concurrent.FutureTask.run(FutureTask.java:266) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) "blc-192.168.1.132:3306" #45 prio=5 os_prio=0 tid=0x00007fdfe00f3000 nid=0x228 runnable [0x00007fdfc1fe9000] java.lang.Thread.State: RUNNABLE at java.net.SocketInputStream.socketRead0(Native Method) at java.net.SocketInputStream.socketRead(SocketInputStream.java:116) at java.net.SocketInputStream.read(SocketInputStream.java:170) at java.net.SocketInputStream.read(SocketInputStream.java:141) at com.github.shyiko.mysql.binlog.io.BufferedSocketInputStream.read(BufferedSocketInputStream.java:51) at com.github.shyiko.mysql.binlog.io.ByteArrayInputStream.readWithinBlockBoundaries(ByteArrayInputStream.java:202) at com.github.shyiko.mysql.binlog.io.ByteArrayInputStream.peek(ByteArrayInputStream.java:175) at com.github.shyiko.mysql.binlog.BinaryLogClient.listenForEventPackets(BinaryLogClient.java:692) at com.github.shyiko.mysql.binlog.BinaryLogClient.connect(BinaryLogClient.java:462) at com.github.shyiko.mysql.binlog.BinaryLogClient$5.run(BinaryLogClient.java:623) at java.lang.Thread.run(Thread.java:745) "kafka-producer-network-thread | 45604805-72b7-49c7-a60f-3035074d2d3a" #44 daemon prio=5 os_prio=0 tid=0x00007fdfe0262800 nid=0x227 runnable [0x00007fdfc1de7000] java.lang.Thread.State: RUNNABLE at sun.nio.ch.EPollArrayWrapper.epollWait(Native Method) at sun.nio.ch.EPollArrayWrapper.poll(EPollArrayWrapper.java:269) at sun.nio.ch.EPollSelectorImpl.doSelect(EPollSelectorImpl.java:93) at sun.nio.ch.SelectorImpl.lockAndDoSelect(SelectorImpl.java:86) - locked <0x00000000fcb069b0> (a sun.nio.ch.Util$2) - locked <0x00000000fcb069a0> (a java.util.Collections$UnmodifiableSet) - locked <0x00000000fcb06888> (a sun.nio.ch.EPollSelectorImpl) at sun.nio.ch.SelectorImpl.select(SelectorImpl.java:97) at org.apache.kafka.common.network.Selector.select(Selector.java:454) at org.apache.kafka.common.network.Selector.poll(Selector.java:277) at org.apache.kafka.clients.NetworkClient.poll(NetworkClient.java:260) at org.apache.kafka.clients.producer.internals.Sender.run(Sender.java:229) at org.apache.kafka.clients.producer.internals.Sender.run(Sender.java:134) at java.lang.Thread.run(Thread.java:745) "kafka-producer-network-thread | 45604805-72b7-49c7-a60f-3035074d2d3a" #43 daemon prio=5 os_prio=0 tid=0x00007fdfe00f6000 nid=0x226 runnable [0x00007fdfc1ee8000] java.lang.Thread.State: RUNNABLE at sun.nio.ch.EPollArrayWrapper.epollWait(Native	1	1	https://issues.jboss.org/browse/DBZ-115
3167	DBZ-115	[~rjharmon0316], thanks for those stack traces. When reading the binlog, the connector uses a 3rd party library to talk to MySQL and read/fetch the binlog events, and this {{com.github.shyiko.mysql.binlog.BinaryLogClient}} class uses two separate threads. One of these is a thread that keeps the connection alive even if there is infrequent activity: {code:java} "blc-keepalive-192.168.1.132:3306" #46 prio=5 os_prio=0 tid=0x00007fdfd4da7000 nid=0x229 waiting on condition [0x00007fdfc1ce6000] java.lang.Thread.State: TIMED_WAITING (sleeping) at java.lang.Thread.sleep(Native Method) at com.github.shyiko.mysql.binlog.BinaryLogClient$3.run(BinaryLogClient.java:555) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:511) at java.util.concurrent.FutureTask.run(FutureTask.java:266) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) {code} The second thread is actually polling MySQL for the next set of MySQL events and handing them off to Debezium logic for converting them into Debezium events and enqueuing them in an in-memory queue: {code:java} "blc-192.168.1.132:3306" #45 prio=5 os_prio=0 tid=0x00007fdfe00f3000 nid=0x228 runnable [0x00007fdfc1fe9000] java.lang.Thread.State: RUNNABLE at java.net.SocketInputStream.socketRead0(Native Method) at java.net.SocketInputStream.socketRead(SocketInputStream.java:116) at java.net.SocketInputStream.read(SocketInputStream.java:170) at java.net.SocketInputStream.read(SocketInputStream.java:141) at com.github.shyiko.mysql.binlog.io.BufferedSocketInputStream.read(BufferedSocketInputStream.java:51) at com.github.shyiko.mysql.binlog.io.ByteArrayInputStream.readWithinBlockBoundaries(ByteArrayInputStream.java:202) at com.github.shyiko.mysql.binlog.io.ByteArrayInputStream.peek(ByteArrayInputStream.java:175) at com.github.shyiko.mysql.binlog.BinaryLogClient.listenForEventPackets(BinaryLogClient.java:692) at com.github.shyiko.mysql.binlog.BinaryLogClient.connect(BinaryLogClient.java:462) at com.github.shyiko.mysql.binlog.BinaryLogClient$5.run(BinaryLogClient.java:623) at java.lang.Thread.run(Thread.java:745) {code} The Kafka Connect thread that is running the task polls the connector for the latest events (which are dequeued from the in-memory queue) and writes them to Kafka and (periodically) updates the offsets. (Note that since the queue is bounded so that any latency or problems writing the events to Kafka creates back pressure on the MySQL binlog client thread by preventing additions to the queue. Also, Kafka Connect records the offsets of those events actually written to Kafka, so if anything bad happens upon restart the connector will start reading the binlog from the last offset recorded by Kafka Connect.) Anyway, it's interesting that this {{blc-192.168.1.132:3306}} thread is at the same position in all of the thread dumps (after restarting the connector). That may just be a coincidence, but it may also be stuck trying to read the next byte from the socket (to MySQL). Have you seen any other stack for this thread in any of the stack traces (perhaps any that you didn't upload)?	1	1	https://issues.jboss.org/browse/DBZ-115
3168	DBZ-105	{quote} I don't think the connector has access to the producer configs {quote} I was assuming that you were creating an out of band KafkaProducer for the history topic. Is that not the case?	1	1	https://issues.jboss.org/browse/DBZ-105

#	Requirement	Questions	Answered	Successfully Answered	URL
3169	DBZ-99	Version 0.4 of the MySQL Binlog client library [is available|https://gitter.im/debezium/dev?at=57b2370c1a7d02075681c97d] with SSL support. Thanks, Stanley! Per the [documentation|https://github.com/shyiko/mysql-binlog-connector-java#using-ssl], it is configured exactly like the [MySQL Connector/J TLS support|https://dev.mysql.com/doc/connector-j/5.1/en/connector-j-reference-using-ssl.html] by either setting system properties in the JVM process: {code} -Djavax.net.ssl.keyStore=<path_to_keystore_file> -Djavax.net.ssl.keyStorePassword=<password> -Djavax.net.ssl.trustStore=<path_to_truststore_file> -Djavax.net.ssl.trustStorePassword=<password> {code} or set via System properties within the code: {code:java} System.setProperty("javax.net.ssl.keyStore","path_to_keystore_file"); System.setProperty("javax.net.ssl.keyStorePassword","password"); System.setProperty("javax.net.ssl.trustStore","path_to_truststore_file"); System.setProperty("javax.net.ssl.trustStorePassword","password"); {code} I think it's better if we use the first approach where the system properties are to be set on the JVM: they're generic and set up the JVM for accessing the keystore & trust store, plus I think it's a bad idea to deal with the passwords via configuration properties. So when it comes to the MySQL connector using SSL, we probably still need to have an explicit {{use.ssl}} configuration property that tells the MySQL JDBC driver and Binlog client library to use the keystore & truststore and connect using TLS. Auto detecting is possible, but it doesn't distinguish the case when a _different_ connector is set up to use TLS but one of the MySQL connectors should not use it. [~criccomini], what do you think?	1	1	https://issues.jboss.org/browse/DBZ-99
3170	DBZ-99	Given DBZ-105 needs configuration properties to set the MySQL connector's database history producer and consumer, I wonder if we should define a single set of these four SSL configuration properties and then use the same values for MySQL connections and the Kafka producer & consumer. We'd still need two additional properties: * {{database.ssl.mode}} for the MySQL connections, with values {{disabled}}, {{preferred}}, {{required}}, {{verify_ca}}, and {{verify_identity}}; and * {{database.history.kafka.ssl.mode}} for whether to use these for the database history producer & consumer, with values {{required}} and {{disabled}}. [~criccomini], thoughts? It doesn't make sense to use different keystore & truststore properties for MySQL and Kafka, does it?	1	1	https://issues.jboss.org/browse/DBZ-99
3171	DBZ-98	[~hchiorean], isn't this very similar to the [pull request|https://github.com/debezium/debezium/pull/84] you created for DBZ-97? [~mvj3], it seems like you approve of that pull request. Is there something else this issue addresses that DBZ-97 does not?	1	1	https://issues.jboss.org/browse/DBZ-98
3172	DBZ-95	[~mvj3], the above error seems to be in {{SourceInfo.java}}, which in the [current {{master}}|https://github.com/debezium/debezium/blob/a589d9ea8468ceb338c6b365c71845c1125d3679/debezium-connector-mysql/src/main/java/io/debezium/connector/mysql/SourceInfo.java#L230] is: {code:java} result.put(BINLOG_FILENAME_OFFSET_KEY, binlogFilename); {code} That error implies that the {{binlogFilename}} is null, and that it actually has nothing to do with the schema in your test database. Can you describe the configuration of the connector and MySQL server? The {{binlogFilename}} should not be null, and I'd like to figure out how it became null.	1	1	https://issues.jboss.org/browse/DBZ-95
3173	DBZ-95	But I think dbz should detects and points out which keys that passed in are invalid? would it be better? Usually, it is done by `MySqlConnectorConfig`.	1	1	https://issues.jboss.org/browse/DBZ-95
3174	DBZ-92	[~mvj3], can you provide more information about your environment? What version of MySQL are you using? Is row-level binary logging enabled? Does the user provided to the connector have the [proper privileges|http://debezium.io/docs/connectors/mysql/#create-a-mysql-user-for-the-connector], namely {{GRANT SELECT, RELOAD, SHOW DATABASES, REPLICATION SLAVE, REPLICATION CLIENT ON *.* TO ...}}? The reason I ask about the latter is the connector was unable to even see any tables as evidenced by this output in your log: {code} 2016-08-02 09:27:15,543 INFO MySQL|mysql_host_192_168_99_100_port_3307|snapshot Step 8: scanning contents of 0 tables [io.debezium.connector.mysql.SnapshotReader] 2016-08-02 09:27:15,543 INFO MySQL|mysql_host_192_168_99_100_port_3307|snapshot Step 8: scanned contents of 0 tables in 00:00:00.0 [io.debezium.connector.mysql.SnapshotReader] {code} Although there may be other causes for this, in the past this has happened to me when I neglected to add the proper {{SELECT}} and/or {{SHOW DATABASES}} privileges for the user. However, this and other error messages in the log your provided above is also a problem, signaling that either Kafka Connect or the connector's own consumer is not able to communicate properly with Kafka: {code} Error while fetching metadata with correlation id 0 : {schema-changes.dp-test=LEADER_NOT_AVAILABLE {code} BTW, the reason that 0.3.0 has not been released is that I've also been having problems with Kafka connectivity performing tests of the whole system. That may be a bug in our Docker images, but as of right now I'm not solved it. (Our unit and integration tests work fine, but don't involve Kafka.)	1	1	https://issues.jboss.org/browse/DBZ-92
3175	DBZ-92	[~mvj3], glad to hear it works now. In the interest of trying to make sure that others don't run into the same problems, would you care to elaborate? Since the connector doesn't use a URL for connection, was your connector configured to use database and/or table filters?	1	1	https://issues.jboss.org/browse/DBZ-92
3176	DBZ-52	Some possible fields for the top-level container structure might include: * {{op}} - Required? Defines the type of event via constants for the operation (e.g., insert, update, delete, read (?), and other future event types) * {{before}} - Optional. Holds a {{Struct}} with the state of the row _before_ the change was made. The {{Schema}} of this structure is connector-specific. * {{after}} - Optional. Holds a {{Struct}} with the state of the row _after_ the change was made. The {{Schema}} of this structure is connector-specific. * {{patch}} - Optional. Holds a description of the individual changes to the record. The format of this description is TBD. * {{source}} - Optional. Holds a {{Struct}} with the connector-supplied source information. The {{Schema}} of this structure is connector-specific. * {{offset}} - Optional. Holds a {{Struct}} with the connector-supplied offset information. The {{Schema}} of this structure is connector-specific. * {{gen-id}} - Optional. Holds a monotonically-increasing integer value representing the snapshot generation in which the event was produced. For example, when a connector needs to generate a consistent snapshot of the contents of a database, it determines the next generation id, and produces a _read_(?) event for each record (including the generation id in each event). Much of these were derived from the discussions on DBZ-31 and DBZ-45, although wWe don't have to lock many of these down to complete this issue. Of course, feedback is welcome.	1	1	https://issues.jboss.org/browse/DBZ-52
3177	DBZ-52	[~xmlking5], it is worth considering. What in particular are you wanting to use that value for? The database's notion of its _logical timestamp_ is already represented in the offset (e.g., for MySQL this would be the binlog filename, position in the binlog, and row within each binlog event).	1	1	https://issues.jboss.org/browse/DBZ-52
3178	DBZ-52	[~xmlking5], wouldn't it be better to filter based upon something in the original database, such as the binlog position in the case of MySQL? Clocks are unreliable, inaccurate, and not even monotonically increasing. And, if problems occurred in the connector, upon restart the same event might be rewritten with a different timestamp.	1	1	https://issues.jboss.org/browse/DBZ-52

#	Requirement	Questions	Answered	Successfully Answered	URL
3179	DBZ-52	Few thoughts: # A version field would be useful. Not sure if it should be at the message level, or the op level (i.e. can operators evolve protocol versions independently?). I lean toward protocol-level version. # Should think about how backwards incompatible changes are to be made. Should we always allow for emission of various version of the protocol, and expose a config that allows end-users to define which version they receive? # Another version of a timestamp that might be useful is simply the timestamp that the connector wrote the message to Kafka. This would be useful for debugging purposes. It's a slightly different ask from [~xmlking5]'s, though. # Adding documentation for these should be exposed in the Debezium docs. (might go without saying, but just saying :)) In general, I agree that these fields can be formalized and added over time. The main thing is getting the barebones stuff set up, and an agreed-upon container structure.	1	1	https://issues.jboss.org/browse/DBZ-52
3180	DBZ-52	{quote} 1. A version field would be useful. Not sure if it should be at the message level, or the op level (i.e. can operators evolve protocol versions independently?). I lean toward protocol-level version. 2. Should think about how backwards incompatible changes are to be made. Should we always allow for emission of various version of the protocol, and expose a config that allows end-users to define which version they receive? {quote} I don't think we need a schema version, since most fields will be optional anyway as few will appear in all messages. That means that we can add other optional fields without consequence; any consumers that don't know about a recently-added optional field will never use it. And I don't think we should reinvent the wheel: simply using the Confluent Avro converters (serialization/deserialization) means we get much of the rest for free, including having consumers be able to choose which "version" of the schema they use. The requirement, though, is that evolution needs to proceed in a manner compatible with [Avro schema evolution|https://martin.kleppmann.com/2012/12/05/schema-evolution-in-avro-protocol-buffers-thrift.html] (namely prevent backward-incompatible changes). {quote} 3. Another version of a timestamp that might be useful is simply the timestamp that the connector wrote the message to Kafka. This would be useful for debugging purposes. It's a slightly different ask from Sumanth Chinthagunta's, though. {quote} Unfortunately, we have no insight or access to when the message is written to Kafka by Kafka Connect. {quote} 4. Adding documentation for these should be exposed in the Debezium docs. (might go without saying, but just saying :-)) {quote} Absolutely. {quote} In general, I agree that these fields can be formalized and added over time. The main thing is getting the barebones stuff set up, and an agreed-upon container structure. {quote} +1. The minimum structure IMO is a single-level envelope with the following fields: {{op}}, {{before}}, {{after}}, {{source}}, and {{offset}}. Other things can be added over time as needed. For example, the {{gen-id}} might be added as part of DBZ-31. Or we might decide that a general purpose {{metadata}} field containing a map of name-value pairs is sufficient for customization.	0	0	https://issues.jboss.org/browse/DBZ-52
3181	DBZ-52	{quote} Unfortunately, we have no insight or access to when the message is written to Kafka by Kafka Connect. {quote} I just meant System.currentTimeMillis() in the Kafka connector. Something approximately correct, so I know if things go wrong, when a message was sent to Kafka by the connector. {quote} The minimum structure IMO is a single-level envelope with the following fields: op, before, after, source, and offset {quote} +1 One nit: could we name it something other than {{offset}}, given that "offset" means something very specific in Kafka-land?	0	0	https://issues.jboss.org/browse/DBZ-52
3182	DBZ-52	This looks awesome. What will the before/after look like for reads? Both set to the same thing? {quote} Notice how the Schema is significantly larger {quote} Personally, not concerned about this, since we will be running with the Confluent schema registry. The schema won't be embedded in the message for us.	1	1	https://issues.jboss.org/browse/DBZ-52
3183	DBZ-50	Your second example has maxLength and masked. Does this make sense? Also, what about {{truncateLength}} just to keep it from getting confused with the length of the VARCHAR?	1	1	https://issues.jboss.org/browse/DBZ-50
3184	DBZ-50	{quote} Your second example has maxLength and masked. Does this make sense? {quote} Yes. The {{maxLength}} parameter comes from the fact that the column's type is {{VARCHAR(255)}}, and thus the column length is 255. Of course, in this case the masking value might be 12 asterisks, and so the {{maxLength}} of the {{email}} in any change events will always be {{************}} and will never have a length approaching 255. We could easily have the {{MaskStrings}} column mapper reset the {{maxLength}} to 12 (for example), but since the masking value is a constant its length really should be {{length}}. Unfortunately, once a parameter is set, it cannot be unset. {quote} Also, what about {{truncateLength}} just to keep it from getting confused with the length of the VARCHAR {quote} This gets to the heart of the issue: what exactly are the semantics of such Schema parameters? Now that I think more about this, it seems to me that _*the Schema parameters should pertain what the Schema describes: the values of the fields in the Struct objects*_. In this case, I think {{length}} and {{maxLength}} are no longer appropriate, since the length can obviously be determined directly from each field value, while the possible lengths of values is dictated and described by the database schema. On the other hand, because the _column mappers_ are altering the field values (relative to the values in the upstream source), it is very appropriate for them to add metadata parameters in the Schema. So, I might propose {{masked=true}} for masked fields, and something like {{truncateLength=N}} for fields that are truncated to specific lengths (only truncation is even possible, namely when _N_ is smaller than the length of the original column).	1	1	https://issues.jboss.org/browse/DBZ-50
3185	DBZ-45	[~criccomini], I think you're right that since most of the popular DBMSes have the ability to include differences, the connector should be able to include the old values in the update event. Should we always do this, or should it be a configuration option? Adding a configuration option adds complexity, but including old+new values in the updates will produce larger messages and may require placing more load on the source DBMS. We do need to think about how to incorporate the old values into the update event. A simple approach is to "reserve" a field name that can hold the old values (in a Kafka Connect {{Struct}}), though this field name (e.g., "{{$old-values}}") can't conflict with any of the field names in the column. {code} { "customerId" : 1234, "customerName" : "Janie Anne Doe", "$old-values" : { "customerId" : 1234, "customerName" : "Jane Doe", } } {code} Alternatively, we rethink the whole structure of the event so that the top level of the event is a "container" struct with fields for new values, old values, type of event/operation, and other metadata not included in the record per se: {code} { "before" : { "customerId" : 1234, "customerName" : "Janie Anne Doe", }, "after" : { "customerId" : 1234, "customerName" : "Jane Doe", }, "op" : "u" } {code} This approach is a bit more like the MongoDB events, although rather than "before" and "after" fields MongoDB events might have a "patch" field containing a description of what changed. I currently believe this is the most flexible and will actually allow a common top-level structure for nearly all connectors, and is worth the additional overhead. Thoughts? What do we do about delete events? Should they contain the old values as well?	1	1	https://issues.jboss.org/browse/DBZ-45

#	Requirement	Questions	Answered	Successfully Answered	URL
3186	DBZ-45	{quote} Should we always do this, or should it be a configuration option? {quote} Config option, but I think the schemas should always look the same. In cases where it's disabled, the "before" should be left empty (see below). {quote} Alternatively, we rethink the whole structure of the event so that the top level of the event is a "container" struct with fields for new values, old values, type of event/operation, and other metadata not included in the record {quote} +1 {quote} What do we do about delete events? Should they contain the old values as well? {quote} This is a little tricky. On the one hand, it's very useful to have the "before" in some cases. On the other hand, if the value of the message is non-null, then Kafka log compaction won't delete it. Need to think about this a bit.	1	1	https://issues.jboss.org/browse/DBZ-45
3187	DBZ-45	I agree that it is promising. Plus I like that the top-level structure for all message is very similar, except for the hard delete. Are these options orthogonal and thus require multiple configuration options, or are they highly-related such that they should be controlled via a single configuration option?	1	1	https://issues.jboss.org/browse/DBZ-45
3188	DBZ-45	{quote} Are these options orthogonal and thus require multiple configuration options, or are they highly-related such that they should be controlled via a single configuration option? {quote} Can you be specific about which options you're referring to?	1	1	https://issues.jboss.org/browse/DBZ-45
3189	DBZ-45	[~konstantin.zadorozhny] wrote: {quote} Also switching between {{minimum}} and {{normal}} or {{full}} will might result in non compatible Avro schema change. {quote} This is a great point: {{minimum}} would change the schema, so it's probably not worth it. Let's drop the {{minimum}} idea. {quote} Based on our experience evolving our CDC schema, it will be very useful to have freedom to extend metadata information sent with every message without breaking backward compatibility. ... I don't think that omitting any metadata is viable option. I have trouble to come up with the scenario when consumers will not need source (binlog in MySQL case) timestamp. {quote} Can you clarify the last sentence? Are you suggesting we include the binlog filename _and_ position (what the connector uses as Kafka Connect offsets) be included as metadata?	1	1	https://issues.jboss.org/browse/DBZ-45
3190	DBZ-45	{quote} Can you clarify the last sentence? Are you suggesting we include the binlog filename and position (what the connector uses as Kafka Connect offsets) be included as metadata? {quote} What we found that, is that having binlog metadata is very useful. This is our metadata field: {code:c} /** Cache record meta data information */ struct MetaData { /** The timestamp of cache update event emitted by binlog replicator */ 1: i64 timestamp; /** The server id of binlog event */ 2: optional i64 binlog_server_id; /** The binlog file name */ 3: optional string binlog_file_name; /** The position in the binlog file */ 4: optional i64 binlog_position; /** The timestamp binlog event */ 5: optional i64 binlog_timestamp; } /* The cache record that represents user table update */ struct User { 1: ... 2: ... ... 9999: optional MetaData meta_data; } {code} Our cache consumers use {{metadata.timestamp}} and {{metadata.binlog_timestamp}} field to tell how far they are behind with the binlog and Kafka topic. (I'm not suggesting Debezium support following use case). In case of the "widespread" client regression, we have an option to reset binlog reader to the last known good position and replay. The point I was trying to make is that, once Debezium user base will grow, the need to extend record metadata will become more apparent. Different databases have their own quirks. Having Kafka Connect record model that could accommodate db specific metadata will be beneficial.	1	1	https://issues.jboss.org/browse/DBZ-45
3191	DBZ-45	This is really interesting. What version of MySQL is this?	1	0	https://issues.jboss.org/browse/DBZ-45
3192	DBZ-45	[~criccomini], I guess I only tested MySQL 5.7. If you're interested in an earlier version of MySQL, would you mind investigating the behavior and reporting back?	0	0	https://issues.jboss.org/browse/DBZ-45
3193	DBZ-41	Is something like this sufficent - https://github.com/jpechane/mysql-for-debezium ?	0	0	https://issues.jboss.org/browse/DBZ-41
3194	DBZ-38	{quote} Sorry, what does "even if they are applied to the definitions inside `Tables`" mean? It's a little technical that related to Debezium's class type system. {quote} Well, I corrected and elaborated that comment. {quote} Thanks, and support rename table would be helpful. {quote} I was afraid of this. The current listener design isn't really able to handle multiple targets (e.g., the old table and the new table). Actually, this helps show that the current design, while being simple, is actually fairly limiting and provides the bare minimum of information with each statement. There are a couple of ways to go about addressing this: 1. Allow multiple TableId objects to be passed to the listener's method. The benefit is that the listeners still can be fairly simple, but unfortunately the listener won't have much detail (e.g., you will not really know which table was old and which was new, unless we use a convention and that becomes fairly brittle). 2. Expand the listener interface to have more specific methods, like createTable, renameTable, alterTable, etc. The benefit is that this is much more explicit and much easier to understand what is going on. The disadvantage is that it is verbose. 3. Change the listener to use an event model, where the listener basically just has a single method and the type of action is described by various event classes. This keeps the API simple, is able to easily handle new types of events, and each event can have different data associated with it. The downside is that it adds complexity to the framework, and it complicates the listener implementation a bit. Thoughts? I'm tempted to use option 3 since IMO it offers the best combination of pros & cons.	1	1	https://issues.jboss.org/browse/DBZ-38
3195	DBZ-38	[~mvj3], can you take a look at the pull request? If it makes sense, I'd like you to try it out before I merge it. Thanks!	1	1	https://issues.jboss.org/browse/DBZ-38
3196	DBZ-37	GTID is only available after MySQL5.6, is there a plan for previous versions of MySQL?	0	0	https://issues.jboss.org/browse/DBZ-37
3197	DBZ-37	{quote} GTID is only available after MySQL5.6, is there a plan for previous versions of MySQL? {quote} Is it possible to replicate from a MySQL <5.6 server into a MySQL 5.6 slave? If so, you could replicate into a 5.6 slave, and use Debezium against the slave.	1	1	https://issues.jboss.org/browse/DBZ-37
3198	DBZ-37	[~criccomini], is there anything else we can/need to do to support HA MySQL clusters other than what's in this issue's description? Do we also need to set a list of host:port pairs and use this ordered list to connect?	1	1	https://issues.jboss.org/browse/DBZ-37
3199	DBZ-37	{quote} is there anything else we can/need to do to support HA MySQL clusters {quote} As far as I know, no. I can't really tell you for sure, until we try it for real, but I believe GTIDs should get us there. {quote} Do we also need to set a list of host:port pairs and use this ordered list to connect? {quote} This would probably be useful for some folks. In our case, we run behind a VIP (or HA Proxy), so we always talk to the same IP/DNS entry. Kafka uses the host:port,host:port model with some success for their bootstrap configs--seems like a reasonable pattern to copy. Like us, at LinkedIn, they use a VIP with a single bootstrap host, though (for legacy reasons).	1	1	https://issues.jboss.org/browse/DBZ-37

#	Requirement	Questions	Answered	Successfully Answered	URL
3200	DBZ-37	[~criccomini], I do have one question, though. The BinlogClient library requires that we either turn on GTIDs if it's going to be aware of them. That's a bit unfortunate, since it would have been nice if it figured it out from the server. But perhaps I'm missing something, and it is useful to be able to run the connector without GTIDs even when the server has them enabled. Does this make sense? If so, then we might need to explicitly enable GTID-awareness on the connector. OTOH, if it only makes sense to be aware of GTIDs when the server uses them, then the connector can figure this out automatically. We actually get the GTIDs via {{SHOW MASTER STATUS}} when running an initial snapshot, but don't yet do anything like this when (re)starting and just reading the binlog. Thoughts? [~konstantin.zadorozhny], do you have any input?	1	1	https://issues.jboss.org/browse/DBZ-37
3201	DBZ-37	IMO, we should default to GTID if it's available. Don't see why not.	1	1	https://issues.jboss.org/browse/DBZ-37
3202	DBZ-37	[~rhauch], thanks for the work on this. It's really appreciated. One question: do these tests include a test that actually fails over from a GTID master to a GTID slave?	1	0	https://issues.jboss.org/browse/DBZ-37
3203	DBZ-31	Looks like you can get the binlog location: http://stackoverflow.com/questions/12025193/mysql-is-there-a-way-to-include-the-bin-log-position-in-the-dump-file-without Based on some cursory browsing, it doesn't sound like mysqldump can produce files that adhere to the binlog protocol. It sounds like it does CSV or SQL DDL/DML. Do you have the ability to parse those commands into Structures?	1	1	https://issues.jboss.org/browse/DBZ-31
3204	DBZ-31	{quote} Based on some cursory browsing, it doesn't sound like mysqldump can produce files that adhere to the binlog protocol. It sounds like it does CSV or SQL DDL/DML. Do you have the ability to parse those commands into Structures? {quote} We should be able to easily parse the DDL/DML form -- we already parse/consume the DDL statements, and the DML (mostly INSERTs?) should be pretty straightforward. And when we do, we can just create the corresponding binlog events and pass them to the same handler methods we already have.	1	1	https://issues.jboss.org/browse/DBZ-31
3205	DBZ-31	If you did (1), would Debezium always read from the replica? Or would Debezium switch back to the master after bulk loading the table? And on a related note, regardless of master vs replica, have you thought about what you would do if the MySQL server you are listening explodes and gets replaced with a new instance? The binlog coordinates that Debezium remembers would no longer apply to the new MySQL instance. Are you using binlog coordinates or GTIDs?	1	1	https://issues.jboss.org/browse/DBZ-31
3206	DBZ-31	{quote} If you did (1), would Debezium always read from the replica? Or would Debezium switch back to the master after bulk loading the table? {quote} As of right now, it would continue using the replica. IMO, that's another reason why #1 is not preferred. {quote} And on a related note, regardless of master vs replica, have you thought about what you would do if the MySQL server you are listening explodes and gets replaced with a new instance? The binlog coordinates that Debezium remembers would no longer apply to the new MySQL instance. Are you using binlog coordinates or GTIDs? {quote} This was just brought up in the mailing list on [this topic|https://groups.google.com/forum/#!topic/debezium/9OWmxcr-Kk4] about HA clusters. It does need to be addressed, but probably via another issue (or several).	1	1	https://issues.jboss.org/browse/DBZ-31
3207	DBZ-31	[~konstantin.zadorozhny], I have a question about your description of what `mysqldump` does. You suggest that the read lock is released before the data from the tables are dumped. Because the read lock is kept for a shorter period of time, the database use by other clients can resume more quickly. On the other hand, keeping the read lock until the dump is complete will give you a perfectly clean snapshot, but at the cost of preventing writes for a longer period of time. I'm trying to work out in my head if releasing the lock before dumping the table data could ever be a problem for downstream consumers. Once the read lock is removed, then the data in a table (or even the table itself) can be changed while the table's data is being dumped. Depending upon where the data dump process is, it may be past the modified rows (in which case the change will not be seen in the snapshot, though the change will still be recorded correctly in the binlog) OR otherwise the snapshot will include the newly committed changes. So while this definitely won't *miss* data changes, it might cause downstream consumers to see change events that were already effectively applied to the data. I can imagine most scenarios would find this perfectly acceptable (the end state of the data will be the same either way), I could imagine a few where this is not ideal. Am I making this more complex than it needs to be?	1	1	https://issues.jboss.org/browse/DBZ-31
3208	DBZ-31	[~rhauch], for [~konstantin.zadorozhny]'s approach, would we *stop* tailing the bin log until the load has completed, or would be reading the binlog WHILE the load is happening?	1	1	https://issues.jboss.org/browse/DBZ-31
3209	DBZ-31	{quote} ... for [~konstantin.zadorozhny]'s approach, would we *stop* tailing the bin log until the load has completed, or would be reading the binlog WHILE the load is happening? {quote} [~criccomini], if snapshotting were enabled, I don't think the connector would start reading the binlog until after the snapshot was completed. When configuring the connector you'd have the option of either: # starting from the beginning of the binlog (possible in Debezium 0.1); or # starting from a specific point in the binlog (possible in Debezium 0.1); or # creating a snapshot (using [~konstantin.zadorozhny]'s approach, including getting the position in the binlog where the snapshot occurred) that would probably (?) appear as INSERT/UPDATE events in the log, and then once that is complete starting reading from the binlog at the measured point. (If the connector had been running, but then was down for longer than the binlog is kept, the only way to recover might be to start it with option 3 so that the snapshot results in new INSERT/UPDATE events.) From a consumer's perspective, a Kafka topic for a particular table would contain an INSERT/UPDATE for each row as read during the snapshot, followed by any INSERT, UPDATE, and DELETE events describing changes committed _after_ the snapshot was started (after the read lock is released). Recall that events from an INSERT and UPDATE look the same, since the event includes a representation of the row *after* the change took effect. If we go with this approach, then the one connector task would do everything, including starting up multiple threads to read the various tables (which are independent from each other). This is simple and would work correctly. (Ideally, the connector could reconfigure itself to start up additional tasks to do that work, but I don't think that will work because of the need to do the work within a single transaction to ensure a consistent snapshot. This could be attempted at a later time.) Does that help?	1	1	https://issues.jboss.org/browse/DBZ-31
3210	DBZ-31	Agree that it's complex. One thing that I've never been able to figure out - and Chris you worked at LinkedIn so maybe you have more experience with this - if you delete the topic, then don't you need to tell all your downstream consumers to dump their datasets and reconsume the topic from scratch? Otherwise, their data may diverge from the stream (like I described above). Or is the small amount of divergence acceptable in some cases? That would be use-case dependent, so if it can happen, it should probably pointed out in the docs somehow.	1	1	https://issues.jboss.org/browse/DBZ-31

#	Requirement	Questions	Answered	Successfully Answered	URL
3211	DBZ-31	{quote} If we going to have generation ID in each record, emitting deletes will be quite easy. Once we have finished publishing snapshot we can read the topic from the beginning, zip the keys, emit tombstones for all the keys with max(GenID) = Current GenID - 1. After that we can safely resume reading from the binlog. {quote} Are you suggesting the connector consume the topics (for each table) and generate the tombstones after it completes a snapshot and before it starts processing the binlog? Where does the list of tables come from, and what if some tables that once existed don't appear in the snapshot?	1	1	https://issues.jboss.org/browse/DBZ-31
3212	DBZ-31	I don't recall us covering how we might trigger these backfills. For use case number (2) in my original list (trigger bootstrap after corruption), how might we trigger Debezium to re-bootstrap?	1	1	https://issues.jboss.org/browse/DBZ-31
3213	DBZ-31	[~konstantin.zadorozhny] wrote: {quote} The [task list endpoint|http://docs.confluent.io/2.0.0/connect/userguide.html#get--connectors-(string-name)-tasks] could be used to monitor the state of the refresh. Once there are no active tasks, the connector config could be changed back to the normal mode. {quote} Are you suggesting the transition from refresh/snapshot mode to normal mode be _manual_? Also, Kafka Connect does not appear to know when a task has _completed_. Even in Kafka 0.10, the task will always remain running (even if it does not do anything) unless the task throws an {{InterruptedException}} (in which case the task will stop running but will commit any remaining offsets) or any other Exception (in which case the exception is logged as an error and the task will commit any offsets). [~criccomini], I think we need an enhancement to Kafka Connect. I think that if the {{SourceConnector}} context provided the {{OffsetStorageReader}}, then we could do everything we need. Luckily that's a non-breaking change, and connector implementations wouldn't have to use it. I'll probably submit an issue and create a PR, but since Kafka 0.10 just came out it will have to wait for the next release. (I don't think they were taking many changes during the past 3 weeks anyway.)	1	1	https://issues.jboss.org/browse/DBZ-31
3214	DBZ-31	[~wushujames] wrote: {quote} We need a way to indicate when a task is "done". Or alternatively, a way for a Task to communicate state back to the Connector. {quote} James, what did you come up with?	1	1	https://issues.jboss.org/browse/DBZ-31
3215	DBZ-31	[~ewencp], first, thanks for taking the time to respond here! This discussion in this issue touches on a number of subjects, but only today have we discussed exactly how to map this to Kafka Connect's concepts. I'm not sure I agree that the offset is simply the state of the snapshot. What happens when you're 99.9% complete on a snapshot when a rebalance or failure occurs? Upon restart, you have to disregard the previously-incomplete snapshot and instead start all over with a new snapshot. To get around that problem, we can use the fact that obtaining a consistent snapshot in MySQL requires first obtaining a write lock. As long as you have the lock, then each table can be read independently -- IOW, once you've completed reading a table and recorded this in an offset, that table doesn't have to be read again. This provides much more favorable recovery behavior, but now the challenge becomes how to start a potentially large number of table names in the offsets for that one partition -- for lots of database tables, this becomes very inefficient. An alternative we talked about above was to instead consider each "independent table" to be a separate snapshot partition, distinct from the original partition that was put on hold. When we completed reading a table, we'd record with the last SourceRecord an offset that denoted that this table was complete. We could then move to the next table. Once all tables are completely read, then the snapshot process terminates, and we record this in the normal partition offset and continue the "normal" reading of the binlog. Essentially it comes down to whether or not these "snapshot source partitions" are a valid concept. IMO, they are since they actually would exist forever and would be reused with future snapshots, should they every be needed (see above for discussion of this). Of course the only way to get this to work (whether one source partition or with "snapshot source partitions") is to still use a single task that handles all of these "source partitions", simply because there's no way for connector instances to monitor or communicate with tasks. That's acceptable, though as mentioned above we would very likely have to manage our own threads rather than rely upon the parallelization of Kafka Connect tasks. See [my email on the Kafka Dev list|http://mail-archives.apache.org/mod_mbox/kafka-dev/201605.mbox/%3cetPan.5744ba0d.6d87b0fa.162@Randalls-MacBook-Pro.local%3e] about a proposal to enable connectors to be able to read the offset storage so they can monitor the progress of tasks (even if they don't ever stop) and to be smarter about the task configurations they create.	1	1	https://issues.jboss.org/browse/DBZ-31
3216	DBZ-31	[~ewencp], [~rhauch], I can talk more about the Kinesis scenario and why it may need "tasks that end". Should I move that discussion to the dev mailing list?	1	1	https://issues.jboss.org/browse/DBZ-31
3217	DBZ-31	Okay, so here's a minor setback. The only way to give Kafka Connect an offset is via {{SourceRecord}} objects produced by the connector. This means that the only way to record that the snapshot has completed is to either: # Change the offset after the snapshot is complete so that the first non-snapshot related {{SourceRecord}} persists this offset information. The disadvantage of this is it requires an update to have occurred in the database since the snapshot started, and while probably common it cannot be guaranteed. *This is not really an option.* # Change the offset on the very last {{SourceRecord}} associated with a table row read during the snapshot. Of course, this will not work if there are no rows read, which can happen if database & table filters are set up and those tables happen to be empty. *This is not really an option.* # Change the offset after the snapshot is complete and _always_ write an additional {{SourceRecord}} to some topic. Database history is out, since we're using our own producer and thus never writing to it via Kafka Connect. We also don't always write to the schema change topic (see above), so that's out, too. However, *this is an option if we always write to a topic*, regardless of whether we find table rows. So the question is, to what topic should we write? One possibility is to have an additional topic that the connector always writes some kind of status and/or state transitions for the connectors. In fact, this could be the same topic for all Debezium connectors, in which case we might name it something like {{debezium.status}}. If multiple Debezium connectors are writing to this topic, then we'd have to be careful about its messages by probably using something very generic like flat JSON messages. However, I'd rather use as few meta-topics as possible, and one for all Debezium has some advantages and potential. This is the best idea I have. Anybody have any thoughts or other ideas? Again, I am going to request that the Kafka Connect API be changed to allow explicit recording of offsets as well as flushing, but that may or may not be desired by others. And in the meantime, we have to work around it as discussed above.	1	1	https://issues.jboss.org/browse/DBZ-31

#	Requirement	Questions	Answered	Successfully Answered	URL
3218	DBZ-31	{quote} the connector always begins by taking a consistent snapshot and then continues reading the binlog from the point in time of the snapshot. {quote} One concern that I do have about this approach is that the bootstrap could take quite a long time. In such a case, if something goes wrong, and a bootstrap is (improperly) triggered (perhaps due to a temporary blip), I'm not sure you'd be able to stop the bootstrap once it's started. Put another way: if a bootstrap starts, must you complete it, if a prior bootstrap has run? If so, perhaps some form of a deadman switch is in order to prevent an unintentional bootstrap from triggering, and causing hours (or days) of "downtime" where you have to rebootstrap a large table (or tables) for no real reason. I don't really have a proposal on what the deadman switch might look like, but it does seem like making the bootstrap a manual process (or require a manual switch flip) might be desirable. {quote} There is a previously-recorded offset, but the binlog coordinates in that offset are no longer available in the MySQL server/cluster's binlog. {quote} Can an offset translate between MySQL machines by accident? This gets into the LSN vs GTID discussion a bit, but is it possible that a fail over to another MySQL machine might just magically line up with the file name/LSN, and Debezium wouldn't do a full bootstrap, but might lose/duplicate data if it ends up in another part of the log? {quote} SET AUTOCOMMIT = 0 {quote} Why set auto commit to false? {quote} tables that exist in the connector's schema model but not in the current databases are dropped using DROP TABLE IF EXISTS <table>; {quote} What's the connector's schema model? Kind of curious about this. {quote} Therefore, if the connector is stopped after before Step 11 completes and the offsets are flushed {quote} Pretty sure this is a typo. Should just be _before_. {quote} When the connector starts up for the first time against a specific logically-named MySQL server/cluster {quote} Can you decsribe how the logical naming will work (or does work)? I'm thinking about this from a fail over perspective. We have some SQL machines behind a load balancer. The fail overs happen fairly transparently (i.e. we'd point Debezium at the load balancer). In such a case, does Debezium think it's talking to one server? This is related to my question above about LSNs translating between machines by accident. {quote} Therefore, this change will remove the ability to specify the binlog coordinates. {quote} +1 {quote} Chris Riccomini, how about changing the focus of this issue to "Connector should start by making snapshot" and then logging a different issue for "Add support for ad hoc snapshots"? {quote} +1 Perhaps we can come up with a work-around for how to handle ad hoc rebootstrapping in the short-term, though? Does renaming the MySQL server/cluster solve the problem (albeit very heavy handedly)? I think it might, though it would require a full re-bootstrap of every table in ever DB in the cluster. {quote} Anybody have any thoughts or other ideas? {quote} Can we just double write the last row? The second time (when we know there are no more records), remove the snapshot flag? Results in a duplicate in the topic. Not great, but an idea nonetheless. {quote} in which case we might name it something like debezium.status {quote} We have a very similar topic in Samza, called the control stream. We used it to store offsets, store config, and send signals to the Samza jobs (e.g. restart, change config, etc). It worked pretty well, but you basically have to treat it like a full-blown protocol. It can get very complicated and messy, so documenting every message, the semantics, etc is really important.	1	1	https://issues.jboss.org/browse/DBZ-31
3219	DBZ-31	perform snapshot | perform snapshot | read from beginning of binlog | | Offsets specify valid binlog position | read binlog from position | read binlog from position | read binlog from position | | Offsets specify binlog position MySQL no longer has | error | perform snapshot | error | | Previous snapshot incomplete | perform snapshot | perform snapshot | error | I think this provides the most control over how the connector starts up, especially in non-nominal situations. {quote} One concern that I do have about this approach is that the bootstrap could take quite a long time. In such a case, if something goes wrong, and a bootstrap is (improperly) triggered (perhaps due to a temporary blip), I'm not sure you'd be able to stop the bootstrap once it's started. Put another way: if a bootstrap starts, must you complete it, if a prior bootstrap has run? {quote} If a bootstrap has started but did not complete, then the connector had already written events to the Kafka topics but those events allow consumers to reconstruct only some of the data. Either a new snapshot is started, or you have to clean up those consumers, topics, and offsets manually -- and you would have that option with the {{never}} snapshot mode described above. If you need or want to stop an ongoing snapshot, simply stop the connector and deal with the Kafka topics, consumers, and offsets. BTW, when a new snapshot is started, the connector writes messages to the schema change topic that contain DDL statements that drop all tables and databases (including any that the connector thought existed prior to the snapshot but that are no longer in the database) and then create the databases and tables that will be in the snapshot. These messages will have the same {{source}} structure as the CREATE events for each row in the tables, giving a consumer everything required to coordinate reading from the normal table streams and the schema change stream. You can read more about the schema change topic [here|http://debezium.io/docs/connectors/mysql/index.html#schema-change-topic]. {quote} Can an offset translate between MySQL machines by accident? This gets into the LSN vs GTID discussion a bit, but is it possible that a fail over to another MySQL machine might just magically line up with the file name/LSN, and Debezium wouldn't do a full bootstrap, but might lose/duplicate data if it ends up in another part of the log? {quote} The purpose of DBZ-37 is to add support for GTIDs so that the connector can recover by using the last GTID to obtain the correct binlog coordinates for the current server. To achieve this, the connector will write in each offset the GTID, the binlog filename, and binlog position; then, upon restart, if a GTID is in the most recently-recorded offset, the connector will use this to start reading the binlog from the correct location. (The MySQL binlog client library we use already supports this, so the connector simply needs to track GTID to the offsets and pass it to the client library when reconnecting.) So, this should also help when connecting to MySQL via a single hostname for the entire cluster resolved via DNS. {quote} Why set auto commit to false? {quote} Because the connector performs multiple operations during the one transaction used in a snapshot. We don't want each of those operations to be in a different transaction, which is what auto commit would do. {quote} Can you decsribe how the logical naming will work (or does work)? I'm thinking about this from a fail over perspective. We have some SQL machines behind a load balancer. The fail overs happen fairly transparently (i.e. we'd point Debezium at the load balancer). In such a case, does Debezium think it's talking to one server? This is related to my question above about LSNs translating between machines by accident. {quote} The 0.1 MySQL connector also used logical name, which is specified with the {{database.server.name}} configuration property and described in the [connector documentation|http://debezium.io/docs/connectors/mysql/index.html#topic-names]. Basically, it allows you go assign a meaningful name to the connector and use that as a prefix for all Kafka topics. Since this does not correspond to anything physical, you have full control over what this is. IIUC, when you're using a load balancer to front the whole MySQL cluster, you'd configure the connector to use the hostname that would resolve to the load balancer. It's pretty clear how a load balancer will help	1	1	https://issues.jboss.org/browse/DBZ-31
3220	DBZ-31	{quote} IOW, this should work as expected as long as the connector is able to establish and use a connection to a specific MySQL server instance. {quote} But didn't we establish that binlog coordinates don't translate between MySQL servers? This was is why we're moving to GTIDs. If the binlog coord don't translate, then don't you run the risk of duplicating/losing data when the failover occurs behind the load balancer, as described above, since the binlog coordinates might go to a different spot in the log? Or would it just fail outright because binlog coords never translate between machines (even inaccurately). [This|https://issues.jboss.org/browse/DBZ-31?focusedCommentId=13247169&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13247169] comment is really fantastic. Is it documented on the website somewhere? Lastly, I'm really curious how this is going to perform on large tables, or many tables. Since it's single threaded, I could imagine bootstrapping taking many days.	1	1	https://issues.jboss.org/browse/DBZ-31

#	Requirement	Questions	Answered	Successfully Answered	URL
3221	DBZ-31	[~criccomini], sorry I misunderstood you point. Yes, I think you're right -- _automatic failover_ to a different MySQL instance pretty much requires a HA MySQL cluster configuration where GTIDs are enabled and some sort of load balancer/proxy on top of the MySQL cluster. But the connector can still _recover_ from failures without a MySQL cluster using GTIDs. In these cases, MySQL replication is used for [backup reasons|https://dev.mysql.com/doc/refman/5.7/en/replication-solutions-backups.html] or [scale out|https://dev.mysql.com/doc/refman/5.7/en/replication-solutions-scaleout.html], and can even use [deeper replication topologies|https://dev.mysql.com/doc/refman/5.7/en/replication-solutions-performance.html] or [multi-source replication|https://dev.mysql.com/doc/refman/5.7/en/replication-multi-source.html], where a single follower tracks multiple distinct masters and is useful for merging multiple shards, backing up multiple servers to a single server, or to consolidating data from multiple servers. In all these cases, if GTIDs are not use then the connector should be configured to use a single specific instance in that cluster. If that instance goes down, the connector becomes unavailable; to recover, either bring up the target MySQL instance (in which case the connector automatically recovers) or explicitly change the connector configuration to point to a different machine. This is not _automatic_ failover, but the connector still recovers quite easily and the resulting event streams remain consistent. Does that help?	1	1	https://issues.jboss.org/browse/DBZ-31
3222	DBZ-31	{quote} Does that help? {quote} Yea, I think it does. Thanks! To paraphrase, I think the guidance is, if you aren't using GTID, then point Debezium at a *specific* machine, rather than the load balancer. {quote} or explicitly change the connector configuration to point to a different machine {quote} When this is done, we'll have to trigger a full snapshot again (presumably by using a different logical name), correct?	0	0	https://issues.jboss.org/browse/DBZ-31
3223	DBZ-29	I don't understand option (C). Can you provide an example?	1	1	https://issues.jboss.org/browse/DBZ-29
3224	DBZ-29	I'm guessing that you can't have: {noformat} column.truncate.server.users.dsecription=100 {noformat} Because it can be a regex, right? IMO, this would be the most intuitive. If we can't do it that way, then I agree, using slightly more words (as in your .chars example) seems prudent. > such as column.truncate.to.100.chars. What about byte arrays? Can we truncate them as well? Not quite sure if that's useful.	1	1	https://issues.jboss.org/browse/DBZ-29
3225	DBZ-29	{quote} I'm guessing that you can't have: {code} column.truncate.server.users.dsecription=100 {code} Because it can be a regex, right? IMO, this would be the most intuitive. If we can't do it that way, then I agree, using slightly more words (as in your .chars example) seems prudent. {quote} I think including the name pattern in the property _name_ is less useful, and it's harder to see what is part of the name vs part of the prefix. And, yes, they could be regex, so that can't be done in a property name. {quote} > such as column.truncate.to.100.chars. What about byte arrays? Can we truncate them as well? Not quite sure if that's useful. {quote} We can truncate byte arrays, but I don't see how that truncated value could be used.	1	1	https://issues.jboss.org/browse/DBZ-29
3226	DBZ-29	One more thing to consider: should anything be included in the generated record written to Kafka when column values are truncated? In other words, is it important for downstream consumers that information has been truncated? For truncated values we could append the ellipsis character (&\#8230; or '…') to signify that the value was truncated. Consider a column truncated to 100 characters, in which the resulting record might contain a value for the column with an ellipsis at 1-based position 101. This approach would allow clients to determine whether a value's length is smaller than or equal to the limit (e.g., 100 characters in our example), versus a value that has been truncated to the limit. As for blacklisted columns, I would suggest we generate records as if those blacklisted column values were not there.	1	1	https://issues.jboss.org/browse/DBZ-29
3227	DBZ-29	> We can truncate byte arrays, but I don't see how that truncated value could be used. Yea, pretty much the only thing I can think of off the top of my head is someone using a BINARY field, but stuffing chars into it. Even then, it'd have to be ASCII so truncating doesn't cause malformed UTF8 (or whatever). > One more thing to consider: should anything be included in the generated record written to Kafka when column values are truncated? In other words, is it important for downstream consumers that information has been truncated? This seems interesting, but I can't think of a concrete use case off the top of my head. > For truncated values we could append the ellipsis character (… or '...') Not so crazy about this idea. Would prefer a flag in the message payload. I suppose string parsing isn't that expensive if you're just doing an indexed lookup, but it seems a bit clunky. Just a personal taste, though. :) > As for blacklisted columns, I would suggest we generate records as if those blacklisted column values were not there. Yea, we faced this issue recently. There are some use cases to keep the columns, but mask the values (e.g. so people can see what the source schema looks like, even if they can't see the data). The three we came up with were: # Mask the values, but keep the column. # Hash the values (SHA3, or something), and keep the column. # Remove the column entirely. (2) is the most useful, but opens you up to dictionary/rainbow attacks. We opted not to do that. (1) is useful for the reason that I mentioned above--people can write queries that will work against "production" (real source) data, even if they can't see the fields. The trade off that we're worried about with this approach is filling masked integer/boolean/whatever with random or zero'd values is a little non-obvious to users; they have to be told that the data is masked, or risk assuming that it's real when it's not. (3) is the safest/cleanest, though.	1	1	https://issues.jboss.org/browse/DBZ-29
3228	DBZ-29	{quote} Perhaps this is acceptable behavior after all, since the user can always properly escape the . characters if the "literal" matches more than expected. {quote} I agree with this statement. {quote} Added a pull request that changes the MySQL connector to use regular expressions in the database {quote} As you move forward with Postgres and others, do you have any thoughts on how to share this generically? It seems like it should be shared amongst all MySQL connectors. Can the regexes be applied across multiple DBs/tables? (e.g. .*email) Overall, this sounds pretty useful to me. Thanks for adding it.	1	1	https://issues.jboss.org/browse/DBZ-29

#	Requirement	Questions	Answered	Successfully Answered	URL
3229	DBZ-29	{quote} As you move forward with Postgres and others, do you have any thoughts on how to share this generically? It seems like it should be shared amongst all MySQL connectors. {quote} The code that does all the work is in {{debezium-core}} in the {{io.debezium.relational}} package. The {{Selectors}} class makes it easy to build a {{Predicate<TableId>}} based upon the table and database whitelists, and a {{Predicate<ColumnId>}} based upon the column blacklist. Creating them is quite easily done in the connector's task initialization code: {code} Predicate<TableId> tableFilter = Selectors.tableSelector() .includeDatabases(config.getString(MySqlConnectorConfig.DATABASE_WHITELIST)) .excludeDatabases(config.getString(MySqlConnectorConfig.DATABASE_BLACKLIST)) .includeTables(config.getString(MySqlConnectorConfig.TABLE_WHITELIST)) .excludeTables(config.getString(MySqlConnectorConfig.TABLE_BLACKLIST)) .build(); {code} and {code} Predicate<ColumnId> columnFilter = Selectors.excludeColumns(config.getString(MySqlConnectorConfig.COLUMN_BLACKLIST)); {code} Likewise, the `ColumnMappers` class makes it easy to register various mappers when the connector's task initializes. Here's what the MySQL connector does, relying upon some utility lambda logic inside the core `Configuration` object to find matching property names using a regular expression, extract the integer value using the regex's group, and then for each matching property name/value call a `ColumnMapper`'s builder method to register the column mapper based upon the regular expression list in the property value: {code} ColumnMappers.Builder columnMapperBuilder = ColumnMappers.create(); config.forEachMatchingFieldNameWithInteger("column\\.truncate\\.to\\.(\\d+)\\.chars", columnMapperBuilder::truncateStrings); config.forEachMatchingFieldNameWithInteger("column\\.mask\\.with\\.(\\d+)\\.chars", columnMapperBuilder::maskStrings); ColumnMappers columnMappers = columnMapperBuilder.build(); {code} {quote} Can the regexes be applied across multiple DBs/tables? (e.g. .*email) {quote} Yes.	1	1	https://issues.jboss.org/browse/DBZ-29
3230	DBZ-20	[~rhauch] Using Oracle CDC in any way is a bad idea... it's completely deprecated and really shouldn't be considered. However... what CDC and OGG have in common is their use of the Log Miner package... it's a built in PL/SQL package that can be used to programmatically mine changes from an Oracle Database. Log Miner pulls changes out of Oracle REDO logs... the logs that are there for recoverability. This is the route most CDC-like approaches take. http://docs.oracle.com/database/121/ARPLS/d_logmnr.htm#ARPLS022 The other possibility is using Oracle Flashback functionality. Oracle is at it's core a multi-version database... it manages completely different versions of the database for different sessions due to transaction isolation modes. It does this using what's called the UNDO space... it's the "Time Travel" capabilities that exist in the database for managing what my version of the DB looks like, what your version of the DB looks like, etc. Flashback provides the ability to query the DB in "meta" mode... and ask questions like "What are all the rows and all the different versions of those rows that existed between these two timestamps?" Flashback is rarely used as a source for CDC, but in my mind it should be. Most users are not aware of the VERSIONS BETWEEN syntax, which gives us access to the complete list of changes between to timestamps. We would need to manage the current window when querying, and managing what window to use to select the next time, etc. Similar to how Kafka manages offsets. http://docs.oracle.com/database/121/ADFNS/adfns_flashback.htm#ADFNS01004	1	1	https://issues.jboss.org/browse/DBZ-20
3231	DBZ-3	> If we add support for [{{BIT(N)}} and {{BIT VARYING(N)}}], they should be represented as Schema.STRING according to the Postgres docs for these types. Wouldn't it be better to match the behavior of the existing MySQL connector? Seems like the best representation in change events is as a {{byte[]}} value with our {{io.debezium.data.Bits}} semantic schema type.	1	1	https://issues.jboss.org/browse/DBZ-3
3232	DEVEX-92	{quote} Introspecting the box itself seems also pretty difficult {quote} Have you tried? I am wondering whether there is already something one could inspect. If not something should be added and probably reported via the {{adbinfo}} plugin.	1	1	https://issues.jboss.org/browse/DEVEX-92
3233	DEVEX-92	Thanks for the info. That the VM is not running at the time of runtime detection is of course a bummer. Not sure what do say about this. Something else I was wondering is the .cdk file itself. IMO it does not really belong to the openshift-vagrant repo. Instead the installer should create the file and put it at the right location. Would you agree?	1	1	https://issues.jboss.org/browse/DEVEX-92
3234	DEVEX-82	The design was provided on this wireframe. Is there something else that you needed from UX for this? https://projects.invisionapp.com/d/main#/console/4496779/129434259/preview	1	1	https://issues.jboss.org/browse/DEVEX-82

#	Requirement	Questions	Answered	Successfully Answered	URL
3235	DEVEX-57	[~gercan] and I met earlier this morning to discuss how to start spiking this issue. We broke off for lunch/other meetings, and plan to continue. As we were exploring the concepts, there were some things that I didn't get down. I think it would serve us well if I asked a few questions to refine my understanding of the task at hand. I'll do my best to seek out the answers, and when I find answers, I'll update this JIRA comment: 1. In more detail, what are the basic use cases of this story? * Want to support `rsync` behaviour, such that updating deployed JavaScript files is seen by Node.js. Developer should be able to use `oc rsync` updated files from local directory to deployed OpenShift instance without restarting corresponding containers. Developer should also not be required to push local changes to GitHub and rebuild image in order to see changes. * Developers should be able to use grunt and gulp to watch specific directories, which execute `oc rsync` as part of the `watch` task. Once the files are copied into OpenShift, the developer should be able to see their changes immediately. 2. What sorts of boundary cases are we aware of? (This seems ambiguous right now. As I start to gather more answers, I'll have a better picture of the answer to this question.) * Developers will naturally want to deploy their code for production. We should see if OpenShift has a 'development mode' that we can hook into, and only provide debugging support if the developer is in 'development mode'. * 3. More explicitly, what is meant by 'existing command line tools'? Simply grunt and gulp, which execute `oc rsync` as a `watch` task. 4. Concretely, what does the final product look like? What repositories does it affect? * sti-nodejs. There will be a PR such that nodemon (or the like) will be installed on the corresponding docker file, the `run` script will npm install production if the user is not currently in development mode, the `run` script will also ensure that nodemon works correctly. * Red Hat Developer blog explaining how to do code sync to OpenShift v3 using grunt and gulp. 5. Can we expand the definition of both acceptance criteria? At this point, I think I've whittled them down far enough. 6. What tool should we use to let developers see their local changes immediately using `oc rsync`? Nodemon has been suggested by [~gercan]. 7. Under the assumption that Nodemon is the used tool, should we allow Nodemon to be deployed along with a developers application? Two questions to keep in mind are: what overhead does Nodemon incur? What security issues exist with doing this? ... 8. There already seems to be a hot deploy scenario for sti-nodejs [here](https://github.com/openshift/sti-nodejs/tree/master/0.10#hot-deploy). Is this exactly the development scenario we're trying to avoid? Should the PR also modify this doc, and perhaps link to the developer blog? Yes, the developer shouldn't need to put in this much work to get 'hot deploy' functionality. This issue is all about making it easier for the developer. I think that the doc should be modified, and should in some way link to the blog. 9. The sti-nodejs repository doesn't appear to have any _real_ tests. Is it possible to add in tests? This may be a bigger change to the repository, but better now or never rather than later. 10. To keep in mind: We may need/want to throttle how often Nodemon reloads on OS. === Current todo: -Get familiar with Grunt-: http://gruntjs.com/getting-started -Get familiar with Gulp-: https://github.com/gulpjs/gulp/blob/master/docs/getting-started.md -Get familiar with rsync-: http://linux.die.net/man/1/rsync -Get familiar with Nodemon-: https://github.com/remy/nodemon	0	0	https://issues.jboss.org/browse/DEVEX-57
3236	DEVEX-56	In a meeting with [~gercan] the other day, we mentioned that we want to be able to pick up the debug port in the sti `run` script, and run the corresponding node scripts with the specified debug port. The constraint is the the debug port that the docker container knows about (to run node-inspector/a node debugger) must be the same as the debug port that node is running with. As of npm 2.0.0, we can pass arguments using this: https://github.com/npm/npm/pull/5518. We can let the user specify the debug port environment variable in the Dockerfile, and choose 5858 if the user doesn't do this. We can then pick up the environment variable in the `run` script. One caveat is that the user needs to modify the Dockerfile if they wish to use a different debug port. Is this something that we're OK with? Another caveat is that this may conflict with the debug port specified in the `package.json`. I'm not certain what the best thing to do in this situation is, but we can address this when we come to it. Perhaps the best solution is to dissuade developers from specifying the debug port in the `package.json` file. Really, I guess the first thing is to decide where the truth lies - an environment variable, or the package.json. The former is much easier implementation wise. Any thoughts on this [~gercan]?	1	1	https://issues.jboss.org/browse/DEVEX-56
3237	DEVEX-56	[~gercan] It looks like I may not be able to setup node in debug mode in the run script. The run script invokes "npm start", and it doesn't look like the "--debug" option can be passed through the npm command. I could write something that accesses the package.json file and inserts the "--debug" argument and executes the new command, but that may be more convoluted than necessary. The alternative is to instruct the user to add the --debug option in their start command in their package.json. So, to rephrase in a more inquisitive manner: * Do you know of any way to pass in the --debug option using npm start? * Should I access the package.json and insert the argument into the command manually, or instruct the user to do so before building/deploying their image? Thoughts?	0	0	https://issues.jboss.org/browse/DEVEX-56
3238	DEVEX-49	why was the acceptance criterias removed or rather reduced ?	1	1	https://issues.jboss.org/browse/DEVEX-49
3239	DEVEX-22	[~crobson], can you comment on a recommended place to put a cancel button based on the current mocks?	0	0	https://issues.jboss.org/browse/DEVEX-22
3240	DMR-9	How is the test buggy? It's checking how the CLI handles escaped text and represents that escaped text correctly per the CLI syntax. Even if it is, if ModelNode.fromString() is returning a different answer now than it did before, that's a problem. [~aloubyansky] fyi.	1	1	https://issues.jboss.org/browse/DMR-9
3241	DMR-9	{quote}Unfortunately the test unintentionally doesn't test this bug,{quote} Why do you conclude it's unintentionally? {quote}because it contains invalid DMR string in second argument,{quote} Strangely, that's the point. {quote}The problem in this case is also how StringValueOperand behaves in its constructor. It ignores IllegalArgumentExceptions thrown by DMR parsers and sets the value to the provided one if DMR parser chokes on broken input. This design decision is bad one IMO.{quote} Perhaps it's a bad design. This is not the only place for this kind of parsing though. Any command and operation argument is parsed this way. The CLI supports two syntaxes, i.e. it accepts values in a DMR format (which is not really a user-friendly one to type in values manually) and a simplified one (which is supposed to remedy help). We can't make DMR syntax not supported. And we can't support only the DMR. So there needs to be a way to determine the format of the value while handling it. The way it's done (it's officially documented too) the parser will first try the DMR syntax. If it worked, the value is assumed to be in a DMR format. If the DMR parser failed, the simpler format parser will be chosen instead. Which also may fail. And if it fails, the user will get an error. If it passed, it passed. {quote}Old DMR parser (based on CookCC) reported it properly as error. New DMR parser (based on new streaming API) doesn't detect this problem because it will finish string processing once first quote is found close to the end.{quote} Why the new doesn't then? What happens with the ending quote?	1	1	https://issues.jboss.org/browse/DMR-9

#	Requirement	Questions	Answered	Successfully Answered	URL
3242	RHDEVDOCS-520	{quote} For our internal purposes, using Bundler and Rake would be overkill. {quote} Why, I would argue repeatable builds are always a good thing. How do you ensure the version of asciidoctor? Setting up a proper build is not only useful for "external" contributions, but it makes internal hand-overs easier as well. Setting up a simple build with Bundler and Rake is not really hard and has a lot of benefits. {quote} Anyway, up until now, all people who accessed the repo a cloned it to their local computers used Linux workstations where all that was required to build the docs was to install AsciiDoctor. {quote} How do you version asciidoctor or all of its dependencies? {quote} And, more importantly, the docs are usually built and tested using an internal tool called 'ccutil' -- that's a specialized tool that builds docs for the Customer Portal. {quote} What does it do? Should that be the build tool of choice? Can this run on any machine? If not maybe one needs a container build? {quote} On top of that, we have set up automated builds through Jenkins. It rebuilds the docs whenever new stuff is pushed to any branch. New build jobs for new branches are created automatically. You can access all those builds at http://ccs-jenkins.gsslab.brq.redhat.com:8080 {quote} I's argue that this feedback cycle is way to slow and remote. Having to push to a remote and then having to wait until I see the changes does not seem to favor quick developer/author feedback. Reallyo one wants to see changes locally the moment the file is saved. I think docs are no different here than actual code.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-520
3243	RHDEVDOCS-520	bq. Why, I would argue repeatable builds are always a good thing. We're not missing out on repeatable local builds. Building and viewing the docs is as simple as executing a single command. In fact, on my machine, I have a hook set up that automatically rebuilds the docs for me whenever I do a local commit. And because no external people have ever been interested in working directly with the repo, there was no pressure to do anything more in this regard. That said, I'm happy to accommodate your requests if it makes your work with the repo easier. I just wanted to explain that for our internal purposes, we have had no need for anything of the sort you suggest (so far). bq. How do you version asciidoctor or all of its dependencies? We don't because we have no need to (AFAICS). AsciiDoctor has only been used for review builds, and the way it renders the docs has not been changing to any appreciable degree in between versions. bq. What does it do? Should that be the build tool of choice? Can this run on any machine? If not maybe one needs a container build? ccutil builds the docs the same way they're built for publishing on the Customer Portal. So yes, for the docs team, it must be the build tool of choice because it's the only way to locally simulate the exact build conditions of the Customer Portal. AFAIK, ccutil only runs on Linux. It shouldn't be difficult to containerize it -- there has just never been any need for that. Only documentation writers need to run it. For content review purposes, AsciiDoctor builds are sufficient. bq. Having to push to a remote and then having to wait until I see the changes does not seem to favor quick developer/author feedback. Reallyo one wants to see changes locally the moment the file is saved. As I mentioned in the previous comment and again now, we _do_ have the possibility of doing local builds, and we use it for local viewing. The Jenkins set-up I described is for sharing the built docs with colleagues, reviewers, and basically anyone interested in seeing the (WIP) docs.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-520
3244	RHDEVDOCS-511	Work on this will begin Monday with research into the correct hardware prerequisites. I'll be checking past mailing list discussions and pinging developers for input if I can't find them quickly. Should we default to reusing the 2.4 prereqs? That seems like it might lead to them being too low or outdated by a few months.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-511
3245	RHDEVDOCS-511	What are the current requirements? I don't really make any statements about this for Minishift, but I guess in terms of the CDK we can go with what we had for CDK 2.x.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-511
3246	RHDEVDOCS-510	I've removed references to vagrant from this section. I then created a merge request, altho I think there may be some problems. * Is the 90-day entitlement still correct? * Do Docker images use the credentials given as `minishift start --username <rhn user> --password <rhn passwd>` ?	0	0	https://issues.jboss.org/browse/RHDEVDOCS-510
3247	RHDEVDOCS-508	In the pull request you seem to only include {code} include::content/module-using-docker-service.adoc[] {code} what are the actual changes you want us to review. Also, would the pull request not include the changes needed for _module-using-docker-service.adoc_. It would be much easier to give feedback, if you would get the ability to review the pull request. On the JIRA issues there is a lot of context missing and looking at the HTML makes it hard to provide comments. IMO we should try to get to a model, where a change to the docs is made, a pull request created and people requested for the review are cc'ed in the pull request. Then we can go through several review cycles on the pull request and in the end merge it.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-508
3248	RHDEVDOCS-507	[~hardy.ferentschik] [~kumarpraveen] [~lalatendum] [~arout] [~agajdosi] I've updated the GSG's introductory section to include a description with as much depth as the 2.4 documentation. Can you please check this out for technical review to be certain that it's accurate? I'm particularly concerned about outdated content. [~psargent] I'm looking for peer review for this section. Do you have a moment? I've asked in IRC but no response yet. http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Getting_Started_Guide-branch-3.1-master-topic_RHDEVDOCS-507/lastBuild/artifact/index.html#introducing_red_hat_container_development_kit	1	1	https://issues.jboss.org/browse/RHDEVDOCS-507
3249	RHDEVDOCS-507	{quote} The Container Development Environment is provided in the form of a Red Hat Enterprise Linux virtual machine. The Container Development Environment itself can also be installed in a virtual machine. {quote} What does that mean, especially "The Container Development Environment itself can also be installed in a virtual machine. Makes not a lot of sense to me. Are you saying I can install the CDK on a VM itself? First of all this would require nested virtualization and secondly, you would have such a bad performance that it would be impossible to work with.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-507
3250	RHDEVDOCS-507	bq. Are you saying I can install the CDK on a VM itself? First of all this would require nested virtualization and secondly, you would have such a bad performance that it would be impossible to work with. Off topic: Just stumbled upon this. [~hardy.ferentschik], I used CDK 2.x all the time in a nested virt. environment, and the performance was quite good -- I mean, I didn't notice any performance problems at all. Has the situation changed with the switch to Minishift? I know a number of people (internal and external) also asked about using CDK within a VM. Thanks.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-507
3251	RHDEVDOCS-507	{quote} Much of the content that's not in the minishift docs is from 2.4 and was requested, {quote} From? {quote} I'll have more time to address the general problems with the docs next sprint when I've closed most of the open Jira. {quote} So here you loose me. You rather just put patches on things and close JIRAs, instead of tackling the root cause of the problem? Is this not wasted effort?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-507

#	Requirement	Questions	Answered	Successfully Answered	URL
3252	RHDEVDOCS-505	[~budhrg] I've added an explanation of the flags. http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Installation_Guide-branch-3.0-master/lastSuccessfulBuild/artifact/index.html#skip-registration The section needs to be reviewed for technical accuracy before being published. Would you mind doing that and letting me know if it appears correct? Thanks in advance for your time.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-505
3253	RHDEVDOCS-505	[~psargent] I reworded it to say "disconnected" instead of offline. Would you mind having another look when you get a change, to see it this way is clearer?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-505
3254	RHDEVDOCS-498	Need to do draft diagrams next, and then setup meeting with Jess and Emma.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-498
3255	RHDEVDOCS-493	Looks like relevant docs for OpenShift containing the same / similar content was released yesterday (July 12). I need to add links to that content (or an earlier version). [~lalatendum][~rkratky] Also, we need scoping for this. Where should I draw the line, and refer users to the OpenShift Online docs rather than covering it inside this document? I'm leaning toward doing so with the least content possible, to prevent outdated content like Hardy spoke of earlier today in a meeting we had with him. When I've got that question answered, I can begin writing this content. I don't see anything in minishift docs about it (yet).	1	0	https://issues.jboss.org/browse/RHDEVDOCS-493
3256	RHDEVDOCS-493	bq. refer users to the OpenShift Online docs OpenShift Container Platform bq. Where should I draw the line For starters, let's stick only with the get started that's in upstream, i.e. no additions. Then the additional get starteds can be added (for example the GSS-requested article on Maven+Wildfly).	0	0	https://issues.jboss.org/browse/RHDEVDOCS-493
3257	RHDEVDOCS-492	Need more information. How are clusters controlled (oc commands)?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-492
3258	RHDEVDOCS-492	[~lalatendum][~hardy.ferentschik][~kumarpraveen][~arout][~agajdosi] We still need information on technical details, we have none in this Jira. The best that I can find is an email thread from March, describing what the feature might look like. Are the details for this ready (has the devel work been completed well-enough to document)? What's the scope of this work? Is there a clear line where content for this should end, or should I document everything because it's minishift? If and when related content lands in the minishift docs repo, I will rebase again and pull in that content. If we know that's going to happen early enough in the current or next sprint, then I'll wait for it.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-492
3259	RHDEVDOCS-489	[~hardy.ferentschik][~lalatendum][~kumarpraveen][~arout][~agajdosi] How is the UX in minishift changed? Are the current minishift docs (or similar) correct wrt those changes, or are there still more changes coming?	0	0	https://issues.jboss.org/browse/RHDEVDOCS-489
3260	RHDEVDOCS-462	What are the conflicts that users have seen? I need a little more information to know how to set requirements for this and make the docs correct.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-462
3261	RHDEVDOCS-462	[~rkratky] Do we need to document how to avoid overwriting the `oc` and `minishift` executables installed to ~/.minishift/cache? What's the recommended way to handle that? My first thought is something like move the executables to ~/bin/minishift-${version} and ~/bin/oc-${version}, with a symlink pointing to the most recent version from ~/bin/minishift and ~/bin/oc. Does that sound reasonable, or am I misunderstanding?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-462
3262	RHDEVDOCS-462	This predates the new workflow using merge requests. The content is not in a topic branch on gitlab. I have a topic branch locally on my desktop, but it's got extra bits in it. I can create a good topic branch with cherry-picked commits or a fast-forward of the split from HEAD if that's useful, but I can see that might potentially cause problems. [~rkratky][~hardy.ferentschik] Would you like me to clean up the topic branch, push it to gitlab, and create a new merge request?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-462
3263	RHDEVDOCS-460	So here is the thing I am wondering regarding the CDK docs. Why don't you not literally take the Minishift docs and modify/extend from there? I think they are in a pretty good shape by now and cover really most of the relevant aspects. I thought that was the plan from the beginning. It would then also be easier to move shared content upstream to Minishift.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-460
3264	RHDEVDOCS-460	{quote}There's not a direct mapping, as there are even different documents and sections to address the different user needs. {quote} Why are there different user needs. I think the whole of the existing Minishift documentation makes sense for CDK users as well. Biggest difference is probably the installation. Also what are these different documents and sections? Why not trying to align? {quote} There were significant modifications required within the sections copied from minishift docs, {quote} Could you provide some concrete examples. This would help to understand better why the docs (have to) differ so much.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-460
3265	RHDEVDOCS-460	The documentation plan has been reviewed and I don't know of anything else that needs to be completed. I set the state to "pull request sent" earlier to reflect that it had been submitted for review. [~rkratky] Can this be resolved as complete?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-460
3266	RHDEVDOCS-455	https://developers.redhat.com/products/cdk/overview/ *Get started with containers on macOS, Microsoft Windows, or Linux* * talks about Vagrant -- not used in 3.0 *Want to try OpenShift 3 Platform-as-a-Service?* * talks about Vagrantfiles *New easy setup for Windows and macOS users* * fails to mention devsuite for RHEL https://developers.redhat.com/products/cdk/hello-world/ *Linux* * talks about and links to CDK 2.4 * says that devsuite isn't available for Linux *macOS* and *Windows* * links to devsuite1.3 * links to old OpenJDK (shouldn't link at all; devsuite install OpenJDK) * contains outdated screenshots *Linux*, *macOS*, and *Windows* * contains instructions for Vagrant-based installation * links to CDK 2.4 docs * links to outdated "Build your first..." tutorials https://developers.redhat.com/products/cdk/help/ * links to outdated discussions https://developers.redhat.com/products/cdk/community/ * talks about Project Atomic -- wouldn't Minishift and OpenShift Origin be more appropriate? https://developers.redhat.com/products/cdk/updates/ * "interactying"	0	0	https://issues.jboss.org/browse/RHDEVDOCS-455

#	Requirement	Questions	Answered	Successfully Answered	URL
3267	RHDEVDOCS-453	The Docs team has the following comments on dev.rh.com: *Overview:* * Add that devSuite is available on RHEL7 using the standard package management system as a Technology Preview feature *Download* * Group the downloads together by operating systems: Online installer Windows Installer Windows Online installer macOS Installer macOS *Hello World* * In the Windows tab, update to +Red Hat Development Suite 1.4+ instead of Red Hat Development Suite 1.3; also state if this is the online installer or the installer (bundled) * update screenshots for step 5 and step 6 for Windows; vagrant is not available in devSuite 1.4; Kompose is available * In the macOS tab, update to +Red Hat Development Suite 1.4+ instead of Red Hat Development Suite 1.3, also state if this is the online installer or the installer (bundled) * update screenshots for step 5 and step 6 for MacOS; vagrant is not available in devSuite 1.4; Kompose is available * include a tab for RHEL * include the installation steps for RHEL *Help* * This tab says "See our most popular forum discussions related to Red Hat Development Suite". I don't see a link to the discussions. Should there be a link to the discussions under Community Q&A? *Update* * under May 2017, include additional components: OpenJDK 1.8.0_131. Kompose 0.6.0; see Table 3.1 Component Details for version updates given by operating system in devSuite 1.4 Installation Guide.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-453
3268	RHDEVDOCS-452	Hi [~rick_wagner], We would like to clarify a couple of items in the devSuite IG and have a couple of questions: What is needed other than a Red Hat Subscription for all of the different components in devSuite? Does CDK require a separate Red Hat subscription? Thank you !	0	0	https://issues.jboss.org/browse/RHDEVDOCS-452
3269	RHDEVDOCS-444	Hi [~mhusnain] where can I find the updated doc? The one at http://ccs-jenkins.gsslab.brq.redhat.com:8080/job/doc-Red_Hat_JBoss_Developer_Studio_Integration_Stack-Installation_Guide-branch-master/lastSuccessfulBuild/artifact/index.html still contains the warning and wrong links in chapters 6 and 8.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-444
3270	RHDEVDOCS-442	Hi [~mhusnain], please fix the component versions for 10.3 as follows | BPEL | 1.1.0.Final_1.3.200.Final | | BPMN2 | 1.3.3.Final_1.3.1.Final | Then, the the following statement in Chapter 2 is not true _"Red Hat JBoss Developer Studio Integration Stack 10.3.0 does not include any bug fixes or updates to the product. This goal for this release is to synchronize with the Red Hat JBoss Developer Studio 10.3.0 release. "_ There are bug fixes for Fuse and SwitchYard tooling. Section 3.1 - some issues are twice, such as FUSETOOLS-2318, FUSETOOLS-2319, FUSETOOLS-2320, FUSETOOLS-2325 - two issues are missing - FUSETOOLS-2377 and FUSETOOLS-2378 Section 3.2 - please add the following bug fix https://issues.jboss.org/browse/SWITCHYARD-2999 Section 4.2 - issue RHBRMS-1416 should be ncluded in section 4.4 since it is Drools related issue	1	1	https://issues.jboss.org/browse/RHDEVDOCS-442
3271	RHDEVDOCS-442	Hi [~mhusnain] it is ok now. Btw it would be good to mention in the chapter 4 that there is a list of most important issues (critical and blockers issues). Or maybe it is not necessary and it is obvious since there is no reason to list all issues (minor, trivial,...). What do you think?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-442
3272	RHDEVDOCS-439	[~psargent], do you already know what this involves? What will need to be documented about Kompose from the devsuite perspective? Let me know if you need anything in this regard.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-439
3273	RHDEVDOCS-404	Sure, Petra. My comments: * 3.1. VERSION UPDATES - This part should IMO be renamed Component versions or something like that. Updates suggests a comparison between the old and new versions but the actual contents appear to just be a list of components. Additionally, the table could just be a list of components with a note about the one component that is only in the one installer. The table seems overkill to present mostly the same information, but that's my opinion. * 3.2. INSTALLER UPDATES - minor changes: - 'Detection is in progress' - if this is the literal label for the indicator, it should be marked up as something from the GUI. - "Cygwin installs with hidden window." - needs and article. "... with a hidden window"? - "The login-in window" - I assume this an accidental hybrid when looking up whether to use login or log in? :) * 4.2. CLOSED ISSUES - just a link to a filter. Shouldn't this be a list of the JIRAs? I wondered if it was because there were too many issues but it shows just 17 (to me). HTH :)	1	1	https://issues.jboss.org/browse/RHDEVDOCS-404
3274	RHDEVDOCS-398	I have tested update from 10.2. We propably have to update all rh-* packages, I got bundle exception because package rh-java-common-google-gson wasn't updated automatically. {code} !ENTRY org.eclipse.jgit.pgm 4 0 2017-05-12 14:03:44.808 !MESSAGE FrameworkEvent ERROR !STACK 0 org.osgi.framework.BundleException: Could not resolve module: org.eclipse.jgit.pgm [231] Unresolved requirement: Import-Package: org.eclipse.jgit.lfs.server.fs; version="[4.6.1,4.7.0)" -> Export-Package: org.eclipse.jgit.lfs.server.fs; bundle-symbolic-name="org.eclipse.jgit.lfs.server"; bundle-version="4.6.1.201703071140-r"; version="4.6.1"; uses:="javax.servlet, javax.servlet.http, org.eclipse.jgit.lfs.server, org.eclipse.jgit.lfs.lib" org.eclipse.jgit.lfs.server [230] Unresolved requirement: Import-Package: com.google.gson; version="[2.2.4,3.0.0)" Unresolved requirement: Import-Package: org.eclipse.jgit.lfs.server; version="[4.6.1,4.7.0)" -> Export-Package: org.eclipse.jgit.lfs.server; bundle-symbolic-name="org.eclipse.jgit.lfs.server"; bundle-version="4.6.1.201703071140-r"; version="4.6.1"; uses:="javax.servlet.http, org.eclipse.jgit.lfs.lib" ... ... !ENTRY org.eclipse.wst.jsdt.js.node 4 0 2017-05-12 14:03:46.661 !MESSAGE FrameworkEvent ERROR !STACK 0 org.osgi.framework.BundleException: Could not resolve module: org.eclipse.wst.jsdt.js.node [653] Unresolved requirement: Require-Bundle: com.google.gson; bundle-version="[2.2.4,3.0.0)" {code} Steps to reproduce: # Disable RHSCL 2.4 and Devstudio snapshot repo # Add RHSCL 2.3 and Devstudio stable repo # yum remove rh-* # yum install rh-eclipse46-devstudio (should install 10.2.0.GA) # check if eclipse works (I had no errors in the log) # Disable RHSCL 2.3 and Devstudio stable repo # Enable RHSCL 2.4 and Devstudio snapshot repo # rm -rf ~/.eclipse # Start the Eclipse with -clean # There should be error I mentioned above. I didn't see any impact to users, egit worked fine. # sudo yum update "rh-*" {code} Updating: rh-java-common-google-gson noarch 2.2.4-1.el7 rh-eclipse46-build 177 k rh-java-common-javapackages-tools noarch 4.3.2-1.14.el7 rh-eclipse46-build 95 k rh-java-common-lucene5 noarch 5.4.1-2.3.el7 rh-eclipse46-build 2.2 M rh-java-common-lucene5-analysis noarch 5.4.1-2.3.el7 rh-eclipse46-build 1.3 M rh-java-common-lucene5-analyzers-smartcn noarch 5.4.1-2.3.el7 rh-eclipse46-build 3.4 M rh-java-common-lucene5-queries noarch 5.4.1-2.3.el7 rh-eclipse46-build 211 k rh-java-common-lucene5-queryparser noarch 5.4.1-2.3.el7 rh-eclipse46-build 339 k rh-java-common-lucene5-sandbox noarch 5.4.1-2.3.el7 rh-eclipse46-build 246 k rh-java-common-python-javapackages noarch 4.3.2-1.14.el7 rh-eclipse46-build 87 k rh-maven33-runtime x86_64 1-18.el7 rh-eclipse46-build 34 k {code} #Restart Eclipse, it will start without errors. I thought that yum/RPM should resolve all dependencies automatically, shouldn't it? Are all dependencies defined correctly? (I don't know how RPM internally works.) CC: [~nickboldt]	1	1	https://issues.jboss.org/browse/RHDEVDOCS-398

#	Requirement	Questions	Answered	Successfully Answered	URL
3275	RHDEVDOCS-398	rpm install/update SHOULD pull in all the compatible dependencies as defined in the spec file. HOWEVER, if we don't state minimum versions, or we don't state enough items, then the update won't necessarily pull in everything that's needed. Here's the spec file [1]. This might be something we should fix for 11.0. Can you open a JIRA requesting that we add more Requires to the devstudio rpm spec file for plugin dependencies (eg., jgit 4.6.1, google gson [2.2.4,3.0.0)). [1] https://github.com/jbdevstudio/jbdevstudio-product/blob/master/rpm/devstudio.spec.template#L48 ---------- Another known issue is that if a user installs EVERYTHING from central (including both the Tycho m2e connector 0.8.1 and Fuse Tools 9.1, remediation will be called because Fuse Tools 9.1 requires Tycho m2e 0.9, not 0.8.1. Once Fuse Tools 9.2 is released in a couple week, this problem will go away as 9.2 depends on 0.8.1 OR 0.9.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-398
3276	RHDEVDOCS-395	[~rkratky] This isn't getting attention, after tagging stakeholders. Can this safely be closed, given its age?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-395
3277	RHDEVDOCS-385	This was fixed in the patches: https://gitlab.cee.redhat.com/red-hat-developers-documentation/red-hat-container-development-kit/commit/5d9f72ec9934ef8685dd937fbad1779ad1476b09 https://gitlab.cee.redhat.com/red-hat-developers-documentation/red-hat-container-development-kit/commit/e4b3969758fc29d7d4d49b344f83ff8eb4a7c685 A recent jenkins build of the document containing the changes can be found at: http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Installation_Guide-branch-2.4-master/lastSuccessfulBuild/artifact/index.html [~arout] [~agajdosi] Can one of you check out the versions in the jenkins build above, to be certain they're correct?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-385
3278	RHDEVDOCS-384	[~lalatendum] I know that you're busy, but could you look at the document again, to see if the templates section is correct? http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Getting_Started_Guide-branch-2.4-master/lastSuccessfulBuild/artifact/index.html#_using_openshift_from_the_command_line Skip to section 8.2.6 for the relevant content. I've only added templates that are shown in `oc get templates -n openshift`. I also looked under /opt/adb/openshift/templates/common, and saw those templates. Two templates listed are in /opt/adb/openshift/templates/cdk. But I don't see /opt/adb/openshift/templates/cdk/image-streams-rhel7.json in the output of `oc get templates -n openshift`. Should that template (image-streams-rhel7) be included in the list of templates? Do I need to perform some task to import that template, to have an accurate description in the document? I'm not sure how to compare the templates in /opt/adb/openshift/templates with output from `oc get templates -n openshift`, or if one should be a reflection of the other. Can you clear up the relationship between `oc` output and the .json files? Also, is there any other content that needs updating around that section on templates?	0	0	https://issues.jboss.org/browse/RHDEVDOCS-384
3279	RHDEVDOCS-383	[~bbrock], did you ask (for peer review) someone specifically? I suggest tagging that person from this issue, so that it's clear even to other people.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-383
3280	RHDEVDOCS-383	Hi [~bbrock]. I'd be happy to, but am not sure where I can view the changes. I did a quick search for that commit hash in various branches (2.4 master, 2.4 stage, master) but can't find it. Could you provide a link to either the commit diff or a rendered version with specifics of which parts have changed for a focused review?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-383
3281	RHDEVDOCS-381	Thanks, [~apodhrad] - is this new version better? --- From Red Hat JBoss Developer Studio Integration Stack 10.1.0.GA onward, SwitchYard 1.x is deprecated in the Fuse Tooling component and SwitchYard 2.x should be used. ---	1	1	https://issues.jboss.org/browse/RHDEVDOCS-381
3282	RHDEVDOCS-376	[~bbrock], how's it looking with this? In case you haven't received any feedback, please, ping Lala again and add also Praveen and Kumar.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-376
3283	RHDEVDOCS-376	[~bbrock], have you contacted the SMEs again? If not, let's close this.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-376
3284	RHDEVDOCS-373	[~rkratky] [~budhrg] I've updated the source for the document with a comment containing the fix (2.4-master). I can update it quickly and push live, if that's the plan. Should this be fixed, or dropped because of age?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-373
3285	RHDEVDOCS-370	Relevant? RHDEVDOCS-373 (Update install-cli '--path' option value to use folder).	1	1	https://issues.jboss.org/browse/RHDEVDOCS-370
3286	RHDEVDOCS-370	[~bbrock] Sorry for late reply, missed it last week in mails. I am not sure if I understand your comment well: "I need to verify the version of docker that CDK installs when a user specifies one version." Could you give me some hints or steps I can follow? Will work on it later today and give you needed information.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-370
3287	RHDEVDOCS-369	Is there a description of the changes from cdk 2.3 -> 2.4 ?	0	0	https://issues.jboss.org/browse/RHDEVDOCS-369
3288	RHDEVDOCS-369	There will most likely be only the following changes between 2.3 and 2.4: * version number * component versions (OCP, Docker, Kubernetes) * supported version of Vagrant * bugfixes (probably only for some vagrant-service-manager issues) [~hardy.ferentschik], could you confirm or provide details, please?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-369
3289	RHDEVDOCS-369	{quote} could you confirm or provide details, please? {quote} sounds reasonable. At least this is the plan afaiu. However, I am not doing the actual release engineering. [~kumarpraveen] would be in a better position to confirm this.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-369
3290	RHDEVDOCS-368	The doc needs to include an explanation of the 2.4 vs 3.0 dichotomy -- what's what, what's the support status for both, etc.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-368
3291	RHDEVDOCS-368	I don't see any information on the etherpad linked in the description. Is there somewhere that I can get the background, myself?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-368
3292	RHDEVDOCS-368	[~arout] [~agajdosi] 2.4 release notes and known issues documents have all new content. Could you review the documents for technical accuracy? Are there any 2.3 issues which were not resolved in 2.4? From the looks of the etherpad, there are none. But that's not canonical.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-368
3293	RHDEVDOCS-368	Quick review notes: 1.1: * "Guide guide" 2.4.2: * Is the v in "OpenShift Container Platform v3.4.0.40" required for the product? I haven't seen this version scheme before so calling attention to it. 3.1.1: * "Per-user" seems an awkward term to use here. I would suggest changing the focus to the file, so "Vagrantfile for Individual Users..." or some such instead of per-user. * "does not use [recommended: the] supplied username" * Minor quibble, avoid latin root words, problematic for potential translation. Instead of via, I'd recommend "using the" or something of that sort. * "In use without a proxy, the username and password supplied in ~/.vagrant.d/Vagrantfile is used, and the user is not prompted to enter a username or password." Would recommend changing the phrasing in the first few words and removing the redundant information at the end, also active phrasing would be better. Example: "If no proxy is set, <CDK> uses the username and password in <filepath and name>." 3.2: * "users using" sounds awkward. I'd recommend something like s/"that only affect users"/"relevant to". 3.2.1: * the markup on "Cygwin" confuses me. If this is the name of a component, it should probably be "cygwin" since we rarely find uppercase letters in naming components. Alternatively, this would be the name of the product and probably doesn't need marking up, but would indeed be Cygwin, not cygwin. * Should C:\tmp be C:\temp? Just checking. 3.3: * 3.3. Issues on Red Hat Enterprise Linux appears to not have any issues in it. 4: * Additional resources seems like a repeat of the information in the understanding the documentation chapter. Is it necessary? If so, would it be easier to just link to the cdk docs splash page instead of the many duplicated links that must be updated every release?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-368

#	Requirement	Questions	Answered	Successfully Answered	URL
3294	RHDEVDOCS-368	[~mhusnain] I've edited the documents to reflect your review. A document built in jenkins with those changes is at: http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Release_Notes-branch-2.4-master/13/artifact/index.html A couple of things I want to clarify, please let me know if these make sense: "Cygwin" is the name of a project and application, so I originally capitalized it. Also, in at least one case "Cygwin" is capitalized, even on the command line: https://cygwin.com/cygwin-ug-net/ov-ex-win.html - Is that correct? C:\tmp is not a preexisting directory in windows, so using "tmp" instead of "temp" won't break anything. Windows OS tends to use "temp" internally, and windows apps (including cygwin) frequently use "tmp" directories. I clarified that by making the directory C:\tmp as part of the instructions.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-368
3295	RHDEVDOCS-368	[~mhusnain], thanks for the thorough review. One comment: bq. Additional resources seems like a repeat of the information in the understanding the documentation chapter. Is it necessary? If so, would it be easier to just link to the cdk docs splash page instead of the many duplicated links that must be updated every release? I believe the links are good to have there, especially in a RN doc. We want to make access to more specific info as straightforward as possible for users. Also, if the user opens the doc in the 'paged' format, they won't have immediate access to the 'Understanding' section from the KI chapter. As far as keeping the links up-to-date, that's not a problem because they're included as attributes, so the maintenance is all done in the attributes.adoc file -- in fact, we only need to update the product-version attribute, and the links will be automatically updated.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-368
3296	RHDEVDOCS-362	Use http://awestruct.org?	0	0	https://issues.jboss.org/browse/RHDEVDOCS-362
3297	RHDEVDOCS-359	[~mhusnain] [~supriya.bharadwaj], 1. N&N - if you look here, there is also an item for Hibernate and Central: http://tools.jboss.org/documentation/whatsnew/jbosstools/4.4.3.Final.html#central Can you add it? 2. Wayland workaround - it seems a bit strange to just describe a workaround without describing an existing problem. Workaround to what? So I would prefer if the title was something like "Running JBoss Developer Studio on Fedora 25 with Wayland" and then starting the chapter by stating the problem: There is a problem when running JBoss Developer Studio with Wayland on Fedora 25. For details, see JIRA_URL... However, there is a workaround for this: ... The rest of the chapter is ok except for this subtle detail: "$DEVSTUDIO_INSTALL_DIR/devstudio-wayland.sh where DEVSTUDIO_INSTALL_DIR" - the dollar sign should either be in both places or none. I would go with no dollar sign personally. You may argue that the wayland issue is listed in the known issues section, but if we're already dedicating an article to this problem, I think it's a good idea to briefly describe the problem here. In fact, I would be in favor or removing the JIRA from that list of known issues, because we would already have it here and the title says "customer-reported known issues" which is not really entirely true AFAIK. (But if you want to keep it there I'm fine with that.) Other than that, all else seems ok.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-359
3298	RHDEVDOCS-359	No problem, @mmalina. 1. Added items for Hibernate and Central and also some items missing from the OpenShift part. 2. Agreed. Added context and link to JIRA and removed the dollar sign. Could you have a look quickly and confirm this looks okay now, Martin? If so, I'll push this to readiness for release right away.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-359
3299	RHDEVDOCS-352	The 2.4 content of course uses install-cli. There's not an equivalent minishift command that I can find (although I'll check again, maybe I'm missing something). The link to the mailing list in the description points to one post asking to define our stance on how to do this, but offers no solutions. This is different from using different versions of the executables after installation. This Jira is about how to install other versions of those binaries in the first place. I'm still looking how to install a different version of `oc` and `docker` without use the vagrant install-cli command, but it's difficult to be efficient about it. [~hardy.ferentschik] you posed the original question on the mailing list. Do you remember how or when / where that question was resolved? If you can give me a pointer, that would be very helpful.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-352
3300	RHDEVDOCS-352	[~hardy.ferentschik] I've got a version of the Getting Started Guide that adds the info for using `minishift start --ocp-tag`. Can you review the section contain the new info, 3.2.2 (Choosing the Version of OpenShift Container Platform Tools to Install with minishift)? http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Getting_Started_Guide-branch-3.1-master/lastSuccessfulBuild/artifact/index.html#flags If I'm incorrect, please point me in the right direction. I was basing this off information I found on various mailing lists, so I may simply have been looking in the wrong direction or misunderstanding the intent of this issue.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-352
3301	RHDEVDOCS-352	I was afraid of that. What's the actual procedure I should use? If this were RHEL, I'd simply do `yum install @virtualization` against the RHEL repo, or reuse the repo listed in the install guide for `oc`. Is that correct?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-352
3302	RHDEVDOCS-352	{quote} I've got a version of the Getting Started Guide that adds the info for using `minishift start --ocp-tag`. Can you review the section contain the new info, 3.2.2 (Choosing the Version of OpenShift Container Platform Tools to Install with minishift)? {quote} Hmm, I think there are several issues here: {quote} The command minishift start automatically installs supported versions of utilities, including docker and oc. Users may which to install custom versions of those utilities. To specify a version of OpenShift Container Platform tools to install during initial startup, use the --ocp-tag option like this: {quote} * We are neither "installing" oc nor docker. We download/cache _oc_ on the host, but in this context "installing" is not the right term. Regarding _docker_, the statement is plain wrong. We are not installing any Docker client. * --ocp-tag is not about the version of the client tools to install, it specified the OCP version (tag) to provision. It is the OpenShift version you want to have running in the end.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-352
3303	RHDEVDOCS-352	Thanks, Hardy. When you say this: When it comes to the oc binary the version matching your requested OpenShift version will be downloaded and cached. Do you mean the version given with --ocp-tag? (sorry, replied to soon... catching up)	1	1	https://issues.jboss.org/browse/RHDEVDOCS-352
3304	RHDEVDOCS-352	bq. Regarding docker, the statement is plain wrong. We are not installing any Docker client. [~hardy.ferentschik], one question about the 'docker' client binary: Do you mean that it's not possible to use a 'docker' client on the host system to interact with the Docker daemon running within CDK like it was in CDK 2.x? Thanks.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-352
3305	RHDEVDOCS-352	OK, I'm unclear about the use case and it sounds like I've been getting it wrong. I'd like to clear that up. My idea was that users may have (or want to install and then have) the docker client on their host desktop communicate with the docker containers inside the rhel VM. To do that, we want to explain installation of docker on their host desktop, from the same environment where they user would also run `minishift start`. That seems like the point of `minishift docker-env`, similar to the reason for `minishift oc-env`. Is that everyone else's understanding, or do I have something wrong? This is completely separate from running `docker` from the shell a user gets with `minishift ssh`.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-352

#	Requirement	Questions	Answered	Successfully Answered	URL
3306	RHDEVDOCS-352	bq. Why? This is not really a requirement for working with the CDK. Just trying to understand where this is coming from? Hi [~hardy.ferentschik], it's not a requirement, but it's one of the main use cases that have been promoted since the beginning of CDK -- the ability to use its Docker engine/daemon directly from your host system (that and the same thing with 'oc'). In fact, it was always preferred to SSH'ing into the VM. I have not heard about any change in this regard. Is there any reason we wouldn't want users doing that?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-352
3307	RHDEVDOCS-352	Yes, that section is already included in the CDK docs. It sounds like you're saying that for docker (docker-env) you want only that section on reusing the docker daemon and no other usage for docker-env. Is that correct?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-352
3308	RHDEVDOCS-352	{quote} Yes, that section is already included in the CDK docs. It sounds like you're saying that for docker (docker-env) you want only that section on reusing the docker daemon and no other usage for docker-env. Is that correct? {quote} Yes.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-352
3309	RHDEVDOCS-350	[~bbrock] I don't see this in the [Installation guide|http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Installation_Guide-branch-3.0-master/lastSuccessfulBuild/artifact/index.html#hypervisor-requirement], where is it?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-350
3310	RHDEVDOCS-345	[~bbrock], CDK 3.0 will use CIFS for shared folders. Have you discussed it with SMEs, so that you know how to document it? If not, please, reach out to the developers (Hardy, Praveen) to get their input.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-345
3311	RHDEVDOCS-344	I might have missed it, but I could not find the _setup-cdk_ command explained. The CDK stays unusable until this command was run. It does not exist in Minishift and hence has no counterpart in the documentation there. Also, the devsuite installer makes sure to call it, for the case that the installer is used. However, for manual install, _setup-cdk_ is required. [~bbrock], is this on your radar?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-344
3312	RHDEVDOCS-343	[~mhusnain], tagging you as the reviewer for this. Have I missed any occurrence of "2.0"?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-343
3313	RHDEVDOCS-339	[~psargent], could you please specify which parts of the doc changed? Or is really a review of the entire guide called for?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-339
3314	RHDEVDOCS-339	[~psargent], here it is: ch. 1: * "One or more of the products that Red Hat Development Suite installs may have a separate Red Hat subscription." // this is vague and should be clarified * "Development support for any of the installed products requires an active Red Hat subscription for that product." // what's "Development support"? * info about support for RHEL install. needs to be added to this chapter ch. 2: * info about RHEL needs to be added * "OpenJDK, Cygwin, VirtualBox and Kompose" // use serial comma * "If some components in the list are not detected, they are installed." // this is ambiguous; should clarify what 'detected' means in this context (it can be read as 'if you have it installed but the installer is unable to detect it') * "macOS-only Installers (Alpha)" // it's not alpha any more * "The default path for Windows is c:\DevelopmentSuite and the default path for macOS is /Applications/DevelopmentSuite." // comma before 'and' * "Red Hat Development Suite allows the user to enable Hyper-V without exiting the installer." // untrue, unless I'm mistaken ch. 6: * link should be provided to ch. on uninstallation ch. 8: * info about RHEL needs to be added	1	1	https://issues.jboss.org/browse/RHDEVDOCS-339
3315	RHDEVDOCS-336	[~dgolovin] I have a question on one of your comments: 8. Also, this issue has been closed in JIRA. Should we still include in the troubleshooting section?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-336
3316	RHDEVDOCS-335	Do you still need review on this, [~psargent]? I can have a look and add any minor comments tomorrow if you like (and that's not too late). Also, I'd recommend making any comments with internal links restricted to Red Hat Employees. I tried to edit your comment to set that (better immediately with security stuff) so that should be okay for now. For future reference, you should see a little lock icon when writing a comment next to "Viewable by All Users". You can use this to change it to Red Hat Employees for sensitive information (GA dates, etc.) and internal links.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-335
3317	RHDEVDOCS-334	[~psargent] In our [Carnation Release Train Deliverables Deck | https://docs.google.com/a/redhat.com/presentation/d/17LcyBOiRL3nYHRWrM3O_jEBI2HiGlTlgF4r9V_3GKXk/edit?usp=sharing] slide #11 comments [~mikgue] asked for 'What's New' section in release-notes. That sounds like a good Idea if we can add it and include in it that paragraph with version updates and add couple more explaining UI changes. Then mention in the end that more details are available in 'New and Noteworthy Issues' section below. WDYT?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-334
3318	RHDEVDOCS-334	[~dgolovin] It sounds like a good idea. The draft has a section called Updates. I could change the name from Updates to What's New. Here is a possible outline of the What's New section: 3. What's New in DevSuite 1.3? 3.1 Version updates Explain the version updates here Cygwin 2.6.0 -> 2.6.1 DevStudio - 10.2.0.GA -> 10.3.0.GA CDK 2.3 GA to CDK 2.4 GA OpenJDK 1.8_111 to 1.8.0_121 3.2 Installer Changes Explain the Installer changes here An Installer that includes most of the required components and their dependencies for installation. Vagrant and Virtualbox are not included in the installer and have to be downloaded during the installation. How does that sound?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-334
3319	RHDEVDOCS-334	Sure [~psargent]. Notes below: 2. Installation Options * I have a nitpick with calling the bundled installed "An Installer" while we have a name for the "Online Installer". I'd say that giving them separate names helps distinguish elsewhere in the text which installer is being referred to, instead of just generic installer = bundled and Online = the online one. But that's my little nitpick. * The warning about support status seems repeated and redundant. Specifically this text: "The Online and Bundled Installers for macOS are Alpha Technology Preview features. For more information on Red Hat Technology Preview features support scope, see http://access.redhat.com/support/offerings/techpreview/." WARNING The macOS online and bundled installers are Alpha Technology Preview features and are not supported by Red Hat. As a result, they must not be used in production environments. I would combine the text into one IMPORTANT or WARNING admonition and keep it succinct. This seems like we're saying the same thing twice, once in an admonition, but once it is not important enough to be in an admonition, which may be confusing. 3. What's new in Red Hat Development Suite? * Not sure we want question marks in the titles at all. Any thoughts, [~rkratky]? * "The following components and their dependencies have been updated to the latest versions:" - I'd remove the latest versions bit because CDK 3.0 is technically the "latest" version, which is not included here. I'd recommend something like "Red Hat Development Suite includes the following components:" 3.2. Installer Updates * Another nitpick: s/The installer/The installers, since there are multiple installers. * The last two points here seem very similar. I'd combine them to just one point about whether the component is selected or deselected instead of a separate point about selected and a separate one about deselected. 4.2. Noteworthy Issue * Not sure what a Noteworthy issue would be. I'd prefer just calling it resolved if it's closed or known is it's open. The terminology of noteworthy may be unfamiliar and confusing. Hope that helps!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-334
3320	RHDEVDOCS-329	Do I need to wait for some time to see the changes? I'm asking because I don't see the changes in the doc.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-329
3321	RHDEVDOCS-325	Hi [~mhusnain]. Thanks for this. Just FYI, the content is now at: * https://developers.redhat.com/products/cdk/get-started/#fndtn-macos * https://developers.redhat.com/products/devsuite/get-started/#fndtn-macos	0	0	https://issues.jboss.org/browse/RHDEVDOCS-325

#	Requirement	Questions	Answered	Successfully Answered	URL
3322	RHDEVDOCS-316	On a related note, do you think your devsuite IG content about the mac installer could be repurposed to create a page similar to this one: https://developers.redhat.com/products/cdk/get-started/#fndtn-windows -- except for mac?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-316
3323	RHDEVDOCS-316	Thanks, Robert. I noticed this is reopened. Are we still tracking this issue for our dept?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-316
3324	RHDEVDOCS-314	[~mhusnain], how's this coming along? Can I help?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-314
3325	RHDEVDOCS-314	I should be able to get Mike Guerette's attention if needed, but I believe we don't need him for this. I have an account in the Developers Blog, and I can publish stuff, so we're OK in that regard. Do you have an idea about the expected content of the blog?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-314
3326	RHDEVDOCS-314	I was going to blog a bit about the new mac installers (and reiterate that the windows ones are available & fully supported, etc.) and point users to our docs for the full procedure. The real objective is to get exposure for the access.redhat.com docs and hopefully feedback. WDYT?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-314
3327	RHDEVDOCS-314	[~mhusnain], HTML would be ideal, but not necessary. I'm not sure what you mean by the 'blog part of the repo'. What repo?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-314
3328	RHDEVDOCS-314	Hi [~mhusnain], I noticed your note about this -- deciding what should or shouldn't be included. Do you have a draft somewhere? I'd be happy to offer some input. (If there's no draft, I'll just add it here.)	1	0	https://issues.jboss.org/browse/RHDEVDOCS-314
3329	RHDEVDOCS-314	Hi [~mhusnain], why did you close this? Are you not going ahead with the blog and the survey?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-314
3330	RHDEVDOCS-312	[~rkratky] What style conventions need to be applied? Should we be doing this in 2.4 or in 3.x instead?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-312
3331	RHDEVDOCS-304	Fixed in http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Installation_Guide-branch-master/lastSuccessfulBuild/artifact/index.html#_setting_up_container_development_kit_for_use_behind_an_http_proxy [~budhrg], could you please have a look?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-304
3332	RHDEVDOCS-303	Documented in CDK Installation Guide. See, for example: http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Installation_Guide-branch-master/lastSuccessfulBuild/artifact/index.html#starting_cdk_vagrant_box_on_rhel [~lalatendum], could you please have a quick look?	0	0	https://issues.jboss.org/browse/RHDEVDOCS-303
3333	RHDEVDOCS-302	Hi [~budhrg], bq. '--cli-version' : 'Specify the version of the binary as value to be installed' (These are the valid origin oc binary versions like 1.2.0, 1.2.1, 1.3.0, 1.3.1 etc) Two questions about this option: # Will all of these `oc` versions work correctly with the OS server in CDK? # What versions of the `docker` binary are supported? bq. More in this link (Note: here --path description mentions relative path which is an issue being tracked here) Will the problem with relative path be fixed in the VSM that will ship with CDK 2.3?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-302
3334	RHDEVDOCS-302	[~rkratky], > 1. Will all of these `oc` versions work correctly with the OS server in CDK? Yes, as far we come to know after having discussion in mailing list. Valid versions are listed [here|https://github.com/openshift/origin/tags]. But recommand to use latest one like 1.2.1, 1.3.0, 1.3.1. However, default is 1.2.0 at the moment. > 2. What versions of the `docker` binary are supported? Any versions listed [here|https://github.com/docker/docker/tags] but recommended one will be installed by default without specifying '--cli-version' > Will the problem with relative path be fixed in the VSM that will ship with CDK 2.3? Unfortunately, it may not be part of CDK 2.3.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-302
3335	RHDEVDOCS-302	[~budhrg], in view of [1] (_"With my new weekly cdk, I found strange errors testing openshift tools. I was using oc 1.2.x, and got random errors. I was told to upgrade to newest. I switched to 1.4.1. Version 1.4.1 also failed against recent weekly cdk build."_), should the docs mention something to keep users from installing a wrong version? In other words, is there a list of versions that don't work? [1] http://post-office.corp.redhat.com/archives/devtools-devsuite/2016-December/msg00000.html	1	1	https://issues.jboss.org/browse/RHDEVDOCS-302
3336	RHDEVDOCS-302	[~budhrg], could you please review the new docs? Added to sections about Docker and OpenShift: http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Getting_Started_Guide-branch-master/lastSuccessfulBuild/artifact/index.html#_preparing_host_system_for_using_docker_from_the_command_line http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Getting_Started_Guide-branch-master/lastSuccessfulBuild/artifact/index.html#_using_openshift_from_the_command_line	1	1	https://issues.jboss.org/browse/RHDEVDOCS-302
3337	RHDEVDOCS-298	I agree. I suspect the current descriptions are just some "can't leave it empty, but don't have time to really think about it now" placeholders I put there when I was creating the links -- and then forgot about it. As for *including* vs *with*, I find both a bit hard to understand in this context. How about we just spell it out outright? Something like: _Starting development of container and cloud-based applications using Red Hat JBoss Developer Studio with the help of Red Hat CDK, OpenShift, and Docker._	1	1	https://issues.jboss.org/browse/RHDEVDOCS-298
3338	RHDEVDOCS-298	[~mmalina], if you have a few minutes today, could you have a look at the comment earlier where I've tagged you and give us your thoughts on which is more technically accurate? Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-298
3339	RHDEVDOCS-292	[~mhusnain] Could you please add these two TD issues to known issues? https://issues.jboss.org/browse/TEIIDDES-2940 https://issues.jboss.org/browse/TEIIDDES-2880 Thank you	0	0	https://issues.jboss.org/browse/RHDEVDOCS-292
3340	RHDEVDOCS-290	Hi [~mhusnain], just making sure: is this supposed to be version "10.1", or "10.1.0"? I suppose the 'x.y.z' format would be better to keep it consistent...	1	1	https://issues.jboss.org/browse/RHDEVDOCS-290
3341	RHDEVDOCS-290	Thanks. Created new branch "10.1.0-stage" and set up corresponding jobs in Pantheon. Get it on local by running: {code} git checkout --track origin/10.1.0-stage {code} Btw, isn't it weird that there won't be any 10.0.1 version?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-290
3342	RHDEVDOCS-290	Hi [~rkratky], I noticed you didn't make a 10.1.0 master this time, unlike last release. Did this change and I just missed it? Also, re: 10.0.1, there may well be, it usually ties in with the tooling product releases and the plan is usually to look at the upcoming release plus the next one.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-290
3343	RHDEVDOCS-290	bq. I noticed you didn't make a 10.1.0 master this time, unlike last release. Did this change and I just missed it? That's intentional. 10.1.0 is -- at this moment -- the latest we've got, so 'master' is the place to work on 10.1.0 docs. For staging, merge/cherrypick from 'master' to '10.1.0-stage'. As soon as one of the following two things happens, we'll create '10.1.0-master': * docs for an even newer version need to be started * 10.1.0 docs get published to prod.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-290
3344	RHDEVDOCS-283	[~budhrg], [~agajdosi], I modified the instructions to include RHEL Workstation specifics. Could you please have a quick look? file:///home/rkratky/Documents/rh/work/books/cdk/installation_guide/master.html#_installing_software_and_configuring_the_host_system_2	1	1	https://issues.jboss.org/browse/RHDEVDOCS-283
3345	RHDEVDOCS-277	[~apodhrad], can you confirm that the SOA 5.x information should be removed for 10.0, since this should now be deprecated?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-277
3346	RHDEVDOCS-277	[~apodhrad], same as other doc, could you have a look or tag someone from your team if you're busy today? Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-277

#	Requirement	Questions	Answered	Successfully Answered	URL
3347	RHDEVDOCS-267	I thought the main issue with Release Notes was an objection to the current practice of just adding a query in there instead of a list of some sort of of the known and resolved issues. Is that included in this JIRA?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-267
3348	RHDEVDOCS-259	[~eng.praveenkum], is this happening for CDK 2.3?	0	0	https://issues.jboss.org/browse/RHDEVDOCS-259
3349	RHDEVDOCS-254	[~agajdosi], [~arout], could you please review Known Issues for CDK *2.2* [1] and identify the ones that are still relevant for CDK *2.3*? Thanks! [1] https://access.redhat.com/documentation/en/red-hat-container-development-kit/2.2/paged/release-notes-and-known-issues/chapter-3-known-issues	1	1	https://issues.jboss.org/browse/RHDEVDOCS-254
3350	RHDEVDOCS-254	[~agajdosi], [~arout], could you please review the statement about VirtualBox Compatibility [1] and confirm that it's valid for CDK 2.3? Likewise, could you please review the Vagrant Compatibility Matrix [2] and confirm that it's correct for CDK 2.3? [1] http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Release_Notes-branch-master/lastSuccessfulBuild/artifact/index.html#_virtualbox_compatibility [2] http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Release_Notes-branch-master/lastSuccessfulBuild/artifact/index.html#_vagrant_compatibility_matrix	1	1	https://issues.jboss.org/browse/RHDEVDOCS-254
3351	RHDEVDOCS-254	Hello [~rkratky], I updated some known issues in ehterpad. Results of 2.2 KI on 2.3, mostly tested by [~arout], are lower: 3.1.1. DNS service provided by xip.io only works with an Internet connection - Found in 2.3 3.1.2. Red Hat Subscription Delay - succeeded in 5 minutes after registration, however with less luck or faster vagrant up it can be still there, good to mention this general problem in KI I think 3.1.3. The nodejs-4-rhel7 Container Image Does Not Work Behind Proxy - Fixed in 2.3 3.1.4. Credentials Defined in Global Vagrantfile Ignored by vagrant-registration - Fixed in 2.3 3.1.5. Certificate Error upon Log in to OpenShift - Fixed in 2.3 3.1.6. Vagrant Reload Does Not Work behind Proxy - Found in 2.3 3.1.7. Invalid URL Offered upon Registration Fail - Fixed in 2.3 3.1.8. Specifying Path for Installation of Client Binaries Does Not Work as Expected - Fixed in 2.3 3.1.9. The vagrant-registration Plugin Does Not Work with Proxy and Self-Signed TSL Certificate - Needs to be tested 3.2.1. vagrant-registration Plugin on Cygwin Bash Does Not Hide Password - Found in 2.3 3.3.1. The docker command Cannot Communicate with Container Development Kit - Fixed in 2.3 3.3.2. libvirt Fails to Unlink Virtual Machine Image - Fixed in 2.3 Talking about support matrix, I can confirm that CDK 2.3 was tested with Vagrant 1.8.1. on all platforms. We will do more about 1.7.4. tomorrow. About VB - Win7 and Win10 were tested with VirtualBox 5.0.26. On Win7 with VB 5.0.30 I tested vagrant up + OpenShift sanity only now - is ok.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-254
3352	RHDEVDOCS-254	Excellent, thank you [~agajdosi]! One question: bq. 3.1.8. Specifying Path for Installation of Client Binaries Does Not Work as Expected - Fixed in 2.3 In my testing of the latest builds, I was still running into this issue. Could you please confirm or offer more details about the scope of the fix?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-254
3353	RHDEVDOCS-251	[~dgolovin], a few questions about the Mac OS X installer: * does the mac installer also have the Microsoft Visual C++ 2010 SP1 Redistributable Package (x86) as a prerequisite? * does a mac also require the BIOS virtualization option enabled? * where can i download the latest working version of the mac installer at the moment (I'll test this tomorrow on a local mac)	1	1	https://issues.jboss.org/browse/RHDEVDOCS-251
3354	RHDEVDOCS-251	Thanks, [~dgolovin]. Do we support Sierra (10.12) yet, or plan to?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-251
3355	RHDEVDOCS-251	Thanks, [~dgolovin]. I'm adding a draft for you to have a look at if you have time shortly, but have a question about uninstalling on the mac. The current instructions are: * use the command prompt (*cmd*) to run the file. * use the Powershell with the following command: + ---- powershell -ExecutionPolicy bypass -File ${target-folder}\uninstaller\uninstall.ps1 ---- What are the ways to run the file, or powershell (or equivalent) command to run the script) on max os x?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-251
3356	RHDEVDOCS-250	[~dgolovin], can you confirm if the online installer for win (tech preview for 1.1) is now fully supported, or if it remains TP? Same for the Mac OS X version of the online installer?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-250
3357	RHDEVDOCS-250	Thank you, [~dgolovin]. I've fixed it up and the new version without that one JIRA is staged at the same link as before. Are there any other changes or are you happy to ACK this for release?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-250
3358	RHDEVDOCS-249	[~supriya.bharadwaj], I noticed that there is no Mac listed in 8. JBoss Developer Studio 9.0. What happened there? Otherwise it looks good.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-249
3359	RHDEVDOCS-248	Hi [~mhusnain], will it be possible for you to take this up? If I am done with the other release tasks earlier than expected, I'll take care of this. Please let me know. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-248
3360	RHDEVDOCS-248	Hi [~supriya.bharadwaj], 1.3. About this Release JBoss Developer Studio 10.2 is an upgrade of JBoss Developer Studio 10.1. The upgrade has the following purposes: - I know that we've been using this phrasing for a long time now. But I'd suggest to change this. It sounds really a bit strange to say that the purpose of the upgrade is to include Eclipse Neon or EAP 7, when those are not new things in this upgrade. Also, I would suggest saying update instead of upgrade in the first sentence. So how about this? JBoss Developer Studio 10.2 is an update of JBoss Developer Studio 10.1. It has the following features: 1.4. Eclipse Automated Reporting Interface (AERI) - it talks about JBoss Tools CAT being suspended. Maybe it's time to drop it and simply talk about AERI? What do you think? 1.6. Issues - the queries look ok. But what about the issues you added: JBDS-4179 - this was reported by Marian (i.e. not a customer), but there was a duplicate reported by a user, so I guess that's ok - assuming this is what we want. [~mhusnain], where did you get the list of "customer reported issues" in the past? I would expect this to be issues that are reported via our customer portal and not just any user, no? JBDS-4181 - reported by Marian, this shouldn't be there JBDS-4180 - again, reported by Marian JBDS-4204 - reported by a user, but again, it's not clear if this qualifies as a "reported by a customer" Maybe the key thing is whether or not there is a Help Desk Ticket Reference in the JIRA? If that's the case, none of the new issues qualify. But at the same time, there may actually be new JIRAs to add.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-248
3361	RHDEVDOCS-248	[~rkratky] thanks for adding the info! [~mhusnain] and [~mmalina], Robert and I discussed about adding this info to the 1.8 Apply this Release section and decided NOT to add it because that section talks about installing devstudio through 1) Eclipse Marketplace and 2) universal installer. It does not mention anything about other ways to install. So, there is no point adding this to a list that could have also had at least three more ways of installation. Would you two agree? (Also, I think the info is apt where Robert placed it.) Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-248
3362	RHDEVDOCS-247	Supriya, could you find text like "10.1" or "10.1 GA" and replace it with 10.2, please? Please review also chapter 7.2. there is wrong path to DevStudio jar. Build produces file named like this "*devstudio-10.2.0.GA-v20161128-1220-installer-standalone.jar*" and you can find it in this folder */path/to/jbdevstudio-product/installer/target*.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-247

#	Requirement	Questions	Answered	Successfully Answered	URL
3363	RHDEVDOCS-247	When will the new 10.2 docs be live? I only see the 10.1 docs here [0]. [0] https://access.redhat.com/documentation/en/red-hat-jboss-developer-studio/10.2/ Marketing announcement of devstudio 10.2 is tomorrow so it'd be good to get this page updated [1] today. [1] https://developers.redhat.com/products/devstudio/docs-and-apis/ See also https://issues.jboss.org/browse/RHD-1808	1	1	https://issues.jboss.org/browse/RHDEVDOCS-247
3364	RHDEVDOCS-247	[~supriya.bharadwaj], can we close this, or is there still some work to do?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-247
3365	RHDEVDOCS-234	gitlab repo created and populated with skeleton.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-234
3366	RHDEVDOCS-234	Ah, got it. I misunderstood when you said you created a new repo. Actually, making it a subdirectory like you did, instead of a whole new repo, makes a lot of sense. Only thing is, this isn't happening for 2.4, so, could you please move it from '2.4-master' to 'master'? (No rush though.)	1	1	https://issues.jboss.org/browse/RHDEVDOCS-234
3367	RHDEVDOCS-234	[~rkratky] This was created a while ago. We have branches for 3.0- and 3.1- master and stage. For 3.1-master, I have topic branches for each jira. I'm going to be importing them into 3.0, wherever I haven't already. Can this be closed? It seems outdated and complete.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-234
3368	RHDEVDOCS-233	Questions for [~lalatendum]: # Can we set up a repo for this doc alongside other docs repos in internal RH Gitlab? (I.e. will devels be able and willing to contribute in that way?) Or is there a more suitable place? # Where would we want to publish this doc? The RH Customer Portal? The RH Developer Portal? (If so, where? How would it fit into the current structure?) Elsewhere?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-233
3369	RHDEVDOCS-233	[~rkratky] +1 to the internal gitlab repo for collaboration. I am sure everyone in CDK team will be willing to participate. Yes, it should somehow be part of landing page of CDK in developers.redhat.com . Can we also release it as a blog? I think blog will get immediate attention.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-233
3370	RHDEVDOCS-233	what information can be removed from the guide? Should it be entire sections, or just reduce the verbiage but leave the sections intact?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-233
3371	RHDEVDOCS-233	Should this be in 3.1, or pushed off further?	0	0	https://issues.jboss.org/browse/RHDEVDOCS-233
3372	RHDEVDOCS-233	[~lalatendum] Do we need another guide, that would be a concise version of the CDK Getting Started Guide? Today in the minishift/CDK docs sync meeting, we briefly discussed how there's a difference between the CDK docs and the minishift docs, but we'd like to move to a common format as much as possible. Is it time to revisit what documents are required, for the 3.2 time-frame? If we simply don't need this kind of document, I'd like to close this request. Thanks.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-233
3373	RHDEVDOCS-233	I think we can close for this. However I have a question, do we have any plan to add manual installation steps for CDK without Devsuite?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-233
3374	RHDEVDOCS-224	[~supriya.bharadwaj], could you please bump this up in your priority list? There is a lot of conversation about this issue on the mailing lists, so we should try to have this documented for users as soon as we can.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-224
3375	RHDEVDOCS-218	[~jrichter1], can you review, please? Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-218
3376	RHDEVDOCS-218	"This article is only valid for Red Hat Development Suite 1.0 and 1.1." Does that mean all the info should be valid for both versions? (even something specific, like the component versions)	1	1	https://issues.jboss.org/browse/RHDEVDOCS-218
3377	RHDEVDOCS-217	[~dgolovin1], would you mind reviewing the doc ? I don't have a Windows machine so I can't verify the steps that Misha documented.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-217
3378	RHDEVDOCS-217	Hi [~dgolovin], Did you mean to add a link to the docker tools in devstudio where we say users should install docker separately? I feel adding a link to our docker tooling may be confusing, as if we are implying a connection between the docker tooling in devstudio and installing docker outside of devstudio. I don't see docker mentioned anywhere else. What do you think?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-217
3379	RHDEVDOCS-216	[~lalatendum], will this be in Banksia? I expect it won't, but just checking.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-216
3380	RHDEVDOCS-214	[~mhusnain], I wonder, you point me to a stage server for the release notes, but release notes is already up live here: https://access.redhat.com/documentation/en/red-hat-jboss-developer-studio/?version=10.1/ Are you aware?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-214
3381	RHDEVDOCS-214	I briefly tried to compare http://tools.jboss.org/documentation/whatsnew/jbosstools/4.4.1.Final.html to https://doc-stage.usersys.redhat.com/documentation/en/red-hat-jboss-developer-studio/10.1/single/release-notes/ And I noticed that Hibernate N&N were missing in your document (the release notes). Can you check?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-214
3382	RHDEVDOCS-202	Thanks for the clarification, [~mmalina]. There was some confusion because we had a separate thread about this and then [~rkratky] logged RHDEVDOCS-208 for this. Apologies for the confusion. I'll change the Fedora versions back, and everything else is the same as for 10.0. Does that seem correct to you? (same link as before for review after a few minutes).	1	1	https://issues.jboss.org/browse/RHDEVDOCS-202
3383	RHDEVDOCS-200	[~mhusnain], is the version referred to as 10.0 or 10.0.0?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-200
3384	RHDEVDOCS-200	[~rkratky] do you mean within the docs or on the splash page? We have the full x.y.z format in things like Release Notes titles and within the docs where necessary (such as the support doc). On the splash page itself, the versions rarely go beyond x.y since maintenance releases are included within the y-stream version rather than having their own separate docs version on the splash page. Does that answer the question or did you mean some other context?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-200
3385	RHDEVDOCS-199	[~pleacu], can you confirm which versions (if any) of DV, Fuse and BPM Suite/BRMS are supported for 10.0.0?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-199
3386	RHDEVDOCS-199	[~apodhrad], could you have a quick look at the above today please, or flag someone from your team to have a look if you're busy please?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-199
3387	RHDEVDOCS-196	[~mhusnain] now I noticed that the link refers to Release Notes. Could you provide a link which refers to the Configuration Support, please?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-196
3388	RHDEVDOCS-195	[~apodhrad], does this looks okay to you now, or should we wait for more issues for the release notes?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-195
3389	RHDEVDOCS-182	Hi [~jrichter1], can you please confirm that there are no changes to the Install instructions at [1] other than the URLs from where the users will download the product (the URLs should now show 10.1 instead of 10)? Thanks! Hi [~akashche], can you please confirm that there are no chnages to be made to the OpenJDK section in [1]? Thanks! [1]: https://access.redhat.com/articles/2358611	1	1	https://issues.jboss.org/browse/RHDEVDOCS-182
3390	RHDEVDOCS-179	Thanks [~mhusnain]. There's one discrepancy. I believe that according to the gdoc, "configuring_devstudio_for_use_with_jboss_eap_and_jboss_web_framework_kit.adoc" should be in "Getting Started with Devstudio Tools": {quote}Chapter 2b: Start Developing Your First Applications with JBoss Developer Studio Tools EAP + WFK (pointer back to the Maven intro for introductory instructions for setting up Maven, if needed){quote} Am I understanding it correctly? To avoid more misunderstandings, could you please sort the topics as well? I think I got it right, but to make sure. You can also see my latest version here: https://gitlab.cee.redhat.com/red-hat-developers-documentation/red-hat-devstudio/tree/topic_fix-get-started-org	1	1	https://issues.jboss.org/browse/RHDEVDOCS-179
3391	RHDEVDOCS-176	[~rkratky], since you have a clone of this one for CDK, do you know who, if anyone, will be giving us a rough timeline to provide to users for new accounts to be picked up by the sub management system?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-176
3392	RHDEVDOCS-176	[~newtonm], pinging again. Could you please provide us with a number that can be used in docs?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-176
3393	RHDEVDOCS-176	WFM, [~rkratky]. Do you want me to keep this issue open while we wait for a definitive answer (and set it to 10.x for post-GA) or are you logging a new JIRA to track the actual answer later?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-176

#	Requirement	Questions	Answered	Successfully Answered	URL
3394	RHDEVDOCS-169	[~mhusnain], do we actually have instructions on Uninstallation? Can I see them? Because you refer to uninstall here and I'm not sure what that involves. And I'm not sure what exactly the installer is capable of detecting now. {quote} If a Vagrant version older than 1.8.1 is installed, uninstall this component before you upgrade. {quote} - This sounds like vagrant is detected during installation. But what if I have vagrant 1.8.1? {quote} If a VirtualBox version older than 5.0.26 is installed, uninstall this component before you upgrade. {quote} - This sounds like VirtualBox is detected during installation. But what if I have version 5.0.26 (or newer)? [~jrichter1], can you review this, please?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-169
3395	RHDEVDOCS-169	[~jrichter1], you say: {quote}The supported versions are: * 1.8.1 and newer for Vagrant{quote} I thought QE tests _only_ 1.8.1, not higher (I believe there was a problem with 1.8.5). [~optak], could you comment, please?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-169
3396	RHDEVDOCS-169	[~jrichter1], what should we be recommending to users, that specific version and nothing older or newer, or that version plus newer with a cautionary note?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-169
3397	RHDEVDOCS-168	[13:46] <amit_> rkratky, Devsuite installer uses virtualbox 5.0.26 for CDK RC3 [13:54] <agajdosi> rkratky, I used VB 5.0.26 for testing at win7 as dev suite 1.1 recomends it [13:55] <rkratky> amit_, agajdosi, thanks - but is this the only supported version, or .26 and higher? [14:00] <amit_> rkratky, we have no idea if devsuite is going to use the higher version. would you please confirm with devsuite team denis ("Denis Golovin" <degolovi@redhat.com>) [14:01] <rkratky> amit_, yes, I'll do that. But are higher version supported for standalone CDK? [14:14] <lalatenduM> rkratky: we will support the version Devsuite is using , we had some discussion with Denis and team about it in th epast	1	1	https://issues.jboss.org/browse/RHDEVDOCS-168
3398	RHDEVDOCS-168	[~dgolovin1], could you please confirm the supported version(s) of VirtualBox for devsuite 1.1? Is it only 5.0.26 (higher?), or other versions as well?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-168
3399	RHDEVDOCS-168	{quote}Denis Golovin: Vagrant 1.8.1 cannot detect latest VirtualBox 5.1.X installed during vagrant up and starts download for VirtualBox 5.0.10 in order to proceed. If you let it finish it will downgrade installed VirtualBox 5.1.X to 5.0.10 and continue with 'vagrant up'. Should we mention that somewhere in release notes?{quote} [~dgolovin1], I assume this is only true on Windows. Is that correct?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-168
3400	RHDEVDOCS-167	[~optak], could you please confirm (or deny) that QE tested starting CDK with the following option set in the Vagrantfile? config.registration.skip = true Thank you.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-167
3401	RHDEVDOCS-164	Can you [~pavanbhat] check also fedora? I expect there will be ruby-devel from default repo, but it need to be checked for compatibility (and s/yum/dnf for fedora).	1	0	https://issues.jboss.org/browse/RHDEVDOCS-164
3402	RHDEVDOCS-162	Are we good to resolve this issue, [~rkratky]?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-162
3403	RHDEVDOCS-162	Hi [~rkratky], thanks for the fix! I had a look at the changes you made here: https://gitlab.cee.redhat.com/red-hat-developers-documentation/red-hat-devstudio/commit/cab282f980982068cc5cf8b09eda82169a723167. I understood most of it but there are a few that I did not understand may be because I am not used to looking at these files for changes. Will it be too much for you to run me through the changes you made so that, going forward, I can keep these in mind? May be a quick bjn whenever you have time?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-162
3404	RHDEVDOCS-155	Thanks, [~pleacu]. Fixed the formatting issue. I'll also be requesting a redirect from the old support page to this one since we've just moved to this new toolchain. I anticipate no issues, but just FYI. Could you have a look at that and let me know if you're happy for me to close this issue off now?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-155
3405	RHDEVDOCS-155	Thanks, [~tsedmik]! Could you have a look (same URL) now and let me know if this is okay/if there are any other changes required?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-155
3406	RHDEVDOCS-130	[~supriya.bharadwaj] - how is this looking for release?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-130
3407	RHDEVDOCS-122	Hi [~budhrg], I added a section [1] about using CDK behind a proxy to the CDK Installation Guide. Could you please review it? I'm unsure about the configuration using the VSM plugin. I used the Vagrantfile syntax from [2]: {code:ruby} config.servicemanager.http_proxy = '<proxy-server>:<port>' config.servicemanager.http_proxy_user = '<username>' config.servicemanager.http_proxy_password = '<password>' config.servicemanager.services = 'docker, openshift' {code} but when I look at the rhel-ose Vagrantfile supplied in the cdk.zip file, I see the following: {code:ruby} ENV['PROXY'] ? (config.registration.proxy = PROXY = ENV['PROXY']) : PROXY = '' ENV['PROXY_USER'] ? (config.registration.proxyUser = PROXY_USER = ENV['PROXY_USER']) : PROXY_USER = '' ENV['PROXY_PASSWORD'] ? (config.registration.proxyPassword = PROXY_PASSWORD = ENV['PROXY_PASSWORD']) : PROXY_PASSWORD = '' {code} This looks as if the vagrant-registration plugin is used to set the proxy environment variables within CDK. So, which is the preferred way? ---- Also, there's the following in the default Vagrant file: {code:ruby} # prevent the automatic start of openshift via service-manager by just enabling Docker config.servicemanager.services = "docker" {code} However, the example you included in the PR comment on GitHub starts both Docker and OpenShift. Does this mean that the fourth line (config.servicemanager.services = 'docker, openshift') should be removed from the VSM proxy instructions? Thanks! [1] http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Installation_Guide-branch-master/lastSuccessfulBuild/artifact/index.html#_setting_up_container_development_kit_for_use_behind_an_http_proxy_3 [2] https://github.com/projectatomic/vagrant-service-manager/pull/400#issuecomment-251665902	1	1	https://issues.jboss.org/browse/RHDEVDOCS-122
3408	RHDEVDOCS-122	Hi [~budhrg]. Thank you for the detailed explanation -- much appreciated. I modified the doc to reflect this information. Could you please check it again at [1]? I have one follow-up question: Do I understand it correctly that in order for the the _registration process_ to work behind a proxy, the proxy-related environment variables need to be set manually before 'vagrant up'? I.e. the VSM proxy settings don't work for the reg. process. [1] http://ccs-jenkins.gsslab.brq.redhat.com:8080/view/Topic%20branches%20-%20CDK/job/doc-Container_Development_Kit-Installation_Guide-branch-master/lastSuccessfulBuild/artifact/index.html#_setting_up_container_development_kit_for_use_behind_an_http_proxy_3	1	1	https://issues.jboss.org/browse/RHDEVDOCS-122
3409	RHDEVDOCS-122	[~budhrg], again, thanks for the detailed explanation. I have one additional follow-up question: Based on what you wrote in the gist, I expect that if the user wants to use VSM to set up proxy and, at the same time, wants to use the k8s Vagrantfile (i.e. not the rhel-ose Vagrantfile you described in the gist), then the following lines need to be removed from that Vagrantfile: {code:ruby} # prevent the automatic start of openshift via service-manager by just enabling Docker config.servicemanager.services = "docker" # Explicitly enable and start kubernetes config.vm.provision "shell", run: "always", inline: <<-SHELL PROXY=#{PROXY} PROXY_USER=#{PROXY_USER} PROXY_PASSWORD=#{PROXY_PASSWORD} /usr/bin/sccli kubernetes echo "kubernetes single node cluster setup successfully" SHELL {code} Is that correct? Or is the VSM proxy solution only valid for the rhel-ose Vagrantfile?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-122

#	Requirement	Questions	Answered	Successfully Answered	URL
3410	RHDEVDOCS-120	Two questions: a) we have some known issues with installation of 3rd party software INTO devstudio when it's been installed via RPM. These are NOT a problem when installed into devstudio installed via installer/updatesite/Marketplace, only when using the devstudio rpm. From devstudio integration stack: * Teiid Designer 10.0.x (fixed in 11.0) * SwitchYard 2.2.0 Other 3rd party tools: * Spring IDE 3.7.3 to 3.8.2 (fixed in 3.8.3.CI) * ZeroTurnaround JRebel 7.0.0 (fixed in 7.0.1.CI) * Atlassian connector for Eclipse (fix possible if I can fork their code and run their Ant build successfully) How should this be documented? I can help with screenshots and steps to reproduce the problem, or URLs to provide workarounds (nightly sites that work rather than release sites that don't). b) How should I label the rpm download page so it's clear the rpm is technology preview? Should this label also be in the rpm's metadata so a yum search will state that information too? Currently: {code} Name : rh-eclipse46-devstudio Arch : x86_64 Epoch : 0 Version : 10.2 Release : 0.20161122.1325.el7 Size : 373 M Repo : @System From repo : rh-eclipse46-devstudio-snapshot-10.2 Summary : Red Hat Developer Studio URL : http://developers.redhat.com/products/devstudio/overview/ License : EPL Description : Red Hat Developer Studio. {code} I can change the Summary or Description [2] fields as you'd like. [1] https://github.com/jbdevstudio/jbdevstudio-product/blob/master/rpm/devstudio.spec.template#L11 [2] https://github.com/jbdevstudio/jbdevstudio-product/blob/master/rpm/devstudio.spec.template#L53-L54	1	1	https://issues.jboss.org/browse/RHDEVDOCS-120
3411	RHDEVDOCS-120	re: labelling the rpm, I've been told: {quote}(2016-11-24 12:39:49) mbenitez: nboldt, fwiw I don't remember any rpms being built with tech preview metadata when I used to do productization work, everything was explained in rel notes / mktg announcements (2016-11-24 12:45:07) mbenitez: nboldt, yes, I just talked to a couple of the top prod people in EAP prod team and they confirmed{quote} and{quote}(2016-11-24 12:56:39) nboldt: mbenitez: so I should stick "Technology Preview" into the HTML on a /stable/ URL? (2016-11-24 12:59:19) mbenitez: nboldt, I don't think it's necessary, unless you want to add an extra layer of information{quote} re: how to *unbreak your install* if you accidentally do so via installation of a 3rd party feature/plugin, I put comments in the google doc. re: workarounds for how to install stuff previously available in Central ... I can provide links for installation alternatives if we want to suggest people use nightly builds in place of GA releases. Do we want to suggest that? Or is it better to tell people, "Sorry, there's no supported version of Teiid Designer 10.x for Devstudio 10.2 rpm; if you want Teiid Designer 10.x in Eclipse on RHEL, use the installer jar [1] instead of the rpm." ?? [1] http://developers.redhat.com/products/devstudio/download/	1	1	https://issues.jboss.org/browse/RHDEVDOCS-120
3412	RHDEVDOCS-120	bq. I can provide links for installation alternatives if we want to suggest people use nightly builds in place of GA releases. Do we want to suggest that? No, we definitely don't want to suggest anything of the sort. bq. Or is it better to tell people, "Sorry, there's no supported version of Teiid Designer 10.x for Devstudio 10.2 rpm; if you want Teiid Designer 10.x in Eclipse on RHEL, use the installer jar [1] instead of the rpm." ?? Exactly -- that's what we need to do. GSS would have our hides if we did otherwise.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-120
3413	RHDEVDOCS-116	Hi [~mhusnain], could you please contact the people from "your" docs plans and request their final approval? I did so for CDK, and I indicated that unless they voice any complaints by the end of this week, I'll consider the plan approved. Note: All four plans (CDK, devsuite, devstudio, devstudio-is) were formally presented to the whole group last week during the program call.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-116
3414	RHDEVDOCS-112	[~rkratky], if you have time, could you upload this adoc file to the devstudio-is repo so I can get back to work directly in there on Tue?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-112
3415	RHDEVDOCS-112	[~rkratky], can you ping [~apodhrad] here when you have merged the changes and this is ready to review?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-112
3416	RHDEVDOCS-112	[~apodhrad], Misha has updated the following two docs for devstudio-is 9.0.1: * Components and Supported Configurations - http://jenkinscat.gsslab.pnq.redhat.com:8080/view/Devtools/job/doc-devstudio-is_Components_and_Supported_Configurations%20(html-single)/lastSuccessfulBuild/artifact/index.html * Release Notes - http://jenkinscat.gsslab.pnq.redhat.com:8080/view/Devtools/job/doc-devstudio-is_Release_Notes%20(html-single)/lastSuccessfulBuild/artifact/index.html Could you please review them and let us know whether any other changes are required?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-112
3417	RHDEVDOCS-112	[~mhusnain], could you update the versions to 9.0.1.GA, please?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-112
3418	RHDEVDOCS-112	Thanks, Robert. [~apodhrad], are you happy to consider these two documents ok for release at this point?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-112
3419	RHDEVDOCS-111	[~dgolovin1], during today's program call, you mentioned that the devsuite installer detects Cygwin. Does this mean that the issue mentioned in JBDS-3862 has been resolved? The JIRA doesn't seem to state that clearly.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-111
3420	RHDEVDOCS-109	[~mhusnain], this seems kinda stale... Is it still relevant? Also, can you offload it on Supriya?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-109
3421	RHDEVDOCS-102	[~mmalina], [~ldimaggio], do you remember this? Did we actually publish an installer for this?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-102
3422	RHDEVDOCS-101	[~gercan], I'm clearing out some of the lower priority issues we've had open for some time. Could you comment on whether this issue is still relevant, and if it is, I'll document it, otherwise we'll just close it as outdated. Thank you in advance!	1	0	https://issues.jboss.org/browse/RHDEVDOCS-101
3423	RHDEVDOCS-101	Hey [~gercan], the link to the docs is included above. Could you have a look and let me know if this is still required, please? Thanks!	1	0	https://issues.jboss.org/browse/RHDEVDOCS-101
3424	RHDEVDOCS-100	[~mmurray] should this JIRA be moved to TOOLSDOC? Is there a section in the install guide or GSG about multi-user windows installs ?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-100
3425	RHDEVDOCS-100	This is potentially worth documenting in the new Forge doc if it is still valid. [~supriya.bharadwaj], since you recently worked on Forge, can you touch base with the relevant SMEs about whether this is still valid for devstudio or if we can just close this as outdated?	1	0	https://issues.jboss.org/browse/RHDEVDOCS-100
3426	RHDEVDOCS-100	Hi [~nickboldt], is this still valid for Devstudio 10? Thanks!	1	0	https://issues.jboss.org/browse/RHDEVDOCS-100
3427	RHDEVDOCS-84	Hi [~ldimaggio], [~mmalina], [~omajid], [~akashche] I have incorporated your feedback in the doc [1]. Can you all please go through it once before I can publish it on the Portal. Thanks for your review on this! [1]: https://docs.google.com/document/d/1fpz8zvDQAh8YMobyKvJVmP8Vk58B_69keYa-2IJK0BM/edit#	1	1	https://issues.jboss.org/browse/RHDEVDOCS-84

#	Requirement	Questions	Answered	Successfully Answered	URL
3428	RHDEVDOCS-84	Hi Martin, _At the top, you start with Install OpenJDK 8. I understand that we want to promote our own OpenJDK for Windows, but we still support OracleJDK just the same. We don't want users that already have (or prefer) OracleJDK 8 installed to think that they need OpenJDK 8._ I was told that OpenJDK should be the preferred option. Is that not the case? Also, the section about installing OpenJDK starts with: _"Red Hat JBoss Developer Studio 10 requires Java 8 to run. Before commencing the installation of devstudio, make sure that Java 8 is installed on your system. Unless you already have Java 8 installed on your system, use the OpenJDK distribution provided by Red Hat."_ I would say that this makes it clear that OpenJDK isn't required -- but it's there for those who don't have any Java yet. _I actually think that we shouldn't have a specific Install OpenJDK/Java chapter - we should just start with the requirements. Most installation guides start with requirements. The requirements chapter link the support matrix document and it does mention java also, so we could just expand this to say that if you wish to install OpenJDK 8, you can check http://developers.redhat.com/products/openjdk/overview/ - that includes all the information._ Historically, it hasn't been possible to rely on info on dev.rh.com. Also, I'd argue that "official" docs should contain _all_ instructions. It's nice to have some whittled down version for quick start on dev.rh.com, but it's not a full substitute. So, I think we should keep the instructions (though we might want to further improve and reorganize the doc as you suggested).	0	0	https://issues.jboss.org/browse/RHDEVDOCS-84
3429	RHDEVDOCS-84	_I was told that OpenJDK should be the preferred option. Is that not the case?_ I'm not sure. Who told you that? :) I assume that Red Hat want to promote open source technologies, so in that sense we probably want to encourage users to use OpenJDK. But from the point of view of testing and what we support, there shouldn't be any difference for the user - it shouldn't really matter if a user chooses to use Oracle JDK or OpenJDK, both are equally supported. I agree now that it's clear that OpenJDK is not required from this paragraph: {quote} Red Hat JBoss Developer Studio 10 requires Java 8 to run. Before commencing the installation of devstudio, make sure that Java 8 is installed on your system. {quote} But perhaps that text should be above Install OpenJDK 8? Because it's not really part of OpenJDK 8 installation. I get your argument for including all instructions in this doc, so that's fine. I wanted to check what other products do and these are my findings: It varies from product to product. For example the EAP installation guide only specifies Java as a requirement, but doesn't say how to install it: https://access.redhat.com/documentation/en/red-hat-jboss-enterprise-application-platform/version-7.0/installation-guide/ "Ensure that a supported Java Development Kit (JDK) has been installed." On the other hand, JBoss Web Server install guide does include that: https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Web_Server/3/html-single/Installation_Guide/index.html#chap-Installing_JBoss_Web_Server_on_Microsoft_Windows JBoss ON only points to their support matrix document: https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Operations_Network/3.3/html-single/Installation_Guide/index.html {quote} The list of supported platforms, databases, and other requirements such as Java, are listed at https://access.redhat.com/knowledge/articles/112523. {quote} So both approaches are probably fine. But in any case the information has to be correct, which it currently is not :)	1	1	https://issues.jboss.org/browse/RHDEVDOCS-84
3430	RHDEVDOCS-84	bq. I'm not sure. Who told you that? I asked about it specifically during one of the program calls before the Summit ("So, when the official OpenJDK builds become available, should the docs say that it's the preferred way to run devstudio?" -- "Yes, that's correct."). I don't remember who supplied the answer though. bq. But perhaps that text should be above Install OpenJDK 8? Because it's not really part of OpenJDK 8 installation. Agreed. bq. I wanted to check what other products do and these are my findings: It varies from product to product. Ideally, all our docs should be consistent, but we're mostly trying to hang in there... bq. But in any case the information has to be correct, which it currently is not Right.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-84
3431	RHDEVDOCS-84	Summary ([~supriya.bharadwaj], could you please look at this when working on RHDEVDOCS-98?): * Move the following text above "Install OpenJDK 8" because it's not really part of OpenJDK 8 installation. "Red Hat JBoss Developer Studio 10 requires Java 8 to run. Before commencing the installation of devstudio, make sure that Java 8 is installed on your system." * Instructions about openjdk installation need fixing in a similar way that RHDEVDOCS-97 suggests. For example you cannot have a paragraph that says "to install openjdk, install devsuite" and then right below say "install the msi installer". * Information that this is all 64-bit Windows only is missing from the text.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-84
3432	RHDEVDOCS-84	Hi [~mmalina], I have incorporated your feedback. Can you please have a look at https://docs.google.com/document/d/1fpz8zvDQAh8YMobyKvJVmP8Vk58B_69keYa-2IJK0BM/edit?pli=1#heading=h.utrct4nxovxn? Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-84
3433	RHDEVDOCS-84	[~supriya.bharadwaj], I looked at the current version and it looks good. Good job! I have one small suggestion. In the note about devsuite ("Note: You can also install Devstudio as part of Red Hat Development Suite.") - perhaps we should also mention up front that this is 64-bit Windows 7,8 and 10 only? You can read that information from the linked devsuite install guide, but I think it would be nice to mention it here also. WDYT?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-84
3434	RHDEVDOCS-84	Hi [~mmalina], I have added the 64-bit Windows sentence to the note. Can you please have a look so that we can wrap this up. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-84
3435	RHDEVDOCS-82	Attaching a screenshot in case someone is unable to view the draft, but this is where I think it makes sense to add a note about making sure the env variable is correctly set. The change is also visible here: https://access.redhat.com/node/2388671/draft [~mmalina], would you mind reviewing this?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-82
3436	RHDEVDOCS-82	I'm afraid it's not correct: 2. Press the Ctrl and 3 keys and in the Quick Access bar, type Launch CDE (for Container Development Environment). 3. From the results, click Launch CDK Server. - it's inconsistent now. Step 3 needs to change also. Also, are you sure about the need for $VAGRANT_HOME? [~rob.stryker], can you confirm?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-82
3437	RHDEVDOCS-82	_I'm afraid it's not correct:_ _2. Press the Ctrl and 3 keys and in the Quick Access bar, type Launch CDE (for Container Development Environment)._ _3. From the results, click Launch CDK Server._ _it's inconsistent now. Step 3 needs to change also._ [~mmalina], I know it appears inconsistent, but does the UI really say "Launch CDE Server"? I cannot test this at the moment.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-82
3438	RHDEVDOCS-82	[~rkratky], are you implementing the changes yourself? Feel free to assign this issue to yourself if you are, it gets it off my to-do list. :)	1	1	https://issues.jboss.org/browse/RHDEVDOCS-82
3439	RHDEVDOCS-82	Thanks, [~mmalina], that's really good feedback! I've changed the text to this: _____ If an error appears about the {rhcdk} adapter not being able to find the Vagrant metadata and plugins, ensure that the VAGRANT_HOME variable is pointing to a non-default location (the default location is ~/.vagrant.d). To change the environment variable value, use the CDK Launch Configuration > Environment Variables tab. _____ Does this seem okay to you now?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-82

#	Requirement	Questions	Answered	Successfully Answered	URL
3440	RHDEVDOCS-82	[~mhusnain], I'm sorry, but "ensure that the VAGRANT_HOME variable is pointing to a non-default location" sounds to me like asking users to point it to a non-default location. Is my understanding of English wrong here? I thought that's how the word "ensure" is usually used, no? And we want the exact opposite - make sure the variable is either not set at all or set to the default location. Instead of "ensure that" I would put "check if". And add: "If set, it needs to be pointed to the default location."	1	1	https://issues.jboss.org/browse/RHDEVDOCS-82
3441	RHDEVDOCS-82	Sorry, [~mmalina], I misunderstood your comment from before and thought that is what we wanted to advise. How is this: _____ If an error appears about the Red Hat Container Development Kit adapter not being able to find the Vagrant metadata and plugins, check whether the VAGRANT_HOME variable is pointing to the default location (~/.vagrant.d). To change the environment variable value, use the CDK Launch Configuration > Environment Variables tab. _____	1	1	https://issues.jboss.org/browse/RHDEVDOCS-82
3442	RHDEVDOCS-75	[~rkratky], can you provide a quick note on where you got to with this issue so that I know where to pick it up again? At the moment, it looks like a lot of discussion across article comments, emails, and possibly in IRC.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-75
3443	RHDEVDOCS-75	Thanks, [~dgolovin]. I've fixed the social accounts thing and the recurring link. The registration delay issue is an open issue and being tracked for devsuite here: https://issues.jboss.org/browse/RHDEVDOCS-176 As you can see in the JIRA, we're waiting to hear about a definitive interval to recommend to users, so I'll add that when we get a response on that JIRA. For the rest, does the text now look okay to you (same link)? If so, I'll set this issue to resolved. Thanks again!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-75
3444	RHDEVDOCS-75	[~dgolovin], [~mmalina], can one of you review this please?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-75
3445	RHDEVDOCS-73	Hi [~mlabuda] and [~fbricon], can you please review the article at: https://docs.google.com/document/d/15D17sS76YyS59_tIovgYI5YP0Iyw0iMGq3ri2r7Ift4/edit#? I have followed the video that Fred suggested (his Microservices demo) and have not yet tested the steps (shall do that tomorrow when I procure a Windows machine). So, it'll be great if you can let me know of any detail that I may have missed out on. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-73
3446	RHDEVDOCS-73	Hi [~fbricon], I have incorporated your feedback and have also made some minor changes in the document based on your feedback. Can you please have a look at it [1] and let em know if the doc is okay now? Thanks! [1]: https://docs.google.com/document/d/15D17sS76YyS59_tIovgYI5YP0Iyw0iMGq3ri2r7Ift4/edit#	1	1	https://issues.jboss.org/browse/RHDEVDOCS-73
3447	RHDEVDOCS-73	Hi [~mhusnain], can you please peer review this doc [1]? Please note, that after trying my best, I could not test the entire procedure due to a CDK issue that I am facing. But, the article is reviewed by Fred, Ondrej, Radim, and Marian and they have all okay'd it for accuracy of content. Thanks! [1]: https://docs.google.com/document/d/15D17sS76YyS59_tIovgYI5YP0Iyw0iMGq3ri2r7Ift4/edit#	1	1	https://issues.jboss.org/browse/RHDEVDOCS-73
3448	RHDEVDOCS-72	[~supriya.bharadwaj], today is release day, can you let me know the status and if the changes are updated in the repo?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-72
3449	RHDEVDOCS-72	Yes, Misha. One of the articles (RHDEVDOCS-73) is with the QE for review and that's the same article that I am testing. I am waiting for the QE (Marian) to get back on this (I pinged Marian y'day to check if he has gotten to review it, but it seems he was busy with something). The other two articles: - RHDEVDOCS-74: Has one section complete; I and Fred have been working on what more needs to be added to the article - RHDEVDOCS-72: I had an IRC chat with Marian about this article because Fred, who suggested for this use case being documented, couldn't provide any video/doc that I can have a look at to begin work on this. Marian will get back about this after GA. Does that clarify things? Please let me know if it doesn't.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-72
3450	RHDEVDOCS-65	[~mhusnain][~supriya.bharadwaj] Scope is a tricky word, I noticed that one of our key/focal mobile docs has no mention of "Red Hat Mobile (aka Feed Henry)". https://access.redhat.com/articles/1295913 Can you provide me with URLs to the docs you mention in the comment? It could be that all we need to do is link one doc to the other.	1	0	https://issues.jboss.org/browse/RHDEVDOCS-65
3451	RHDEVDOCS-65	What is the status on this? Please follow up with Burr.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-65
3452	RHDEVDOCS-65	Hi Supriya, can you follow up on this with Burr or Alexey one last time to see if you can determine the required use case for this. If not, we'll just close it because we don't have enough information to do anything else with it.	0	0	https://issues.jboss.org/browse/RHDEVDOCS-65
3453	RHDEVDOCS-65	Hi [~akazakov], I see that you had asked [~gercan] for his input on this. If this is no longer relevant, we may want to close this issue. WDYT? -Thanks	1	0	https://issues.jboss.org/browse/RHDEVDOCS-65
3454	RHDEVDOCS-55	I reviewed the section and commented in the doc. BTW, do we have any documentation of the different vagrantfiles? I can see that there is a README file in components/rhel/misc/shared_folder/rhel-ose/README.rst, but that doesn't say anything about the shared folders / sshfs. Do we have anything in the overall cdk documentation online?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-55
3455	RHDEVDOCS-55	BTW, [~rkratky], I cannot access https://access.redhat.com/articles/2361631 ATM - access denied. Is it not yet published?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-55
3456	RHDEVDOCS-55	One more thing, [~mhusnain], I don't quite understand this at the beginning: {quote} There are two ways to install the Red Hat Container Development Kit with JBoss Developer Studio: Installation and configuration in one of two ways: 1. a. Automatically install and configure the Red Hat Container Development Kit and Red Hat JBoss Developer Studio using the Red Hat JBoss Developer Studio Platform Installer. b. Manually download, install, and configure the Red Hat Container Development Kit and Red Hat JBoss Developer Studio. 2. Use the installed components for Container-Based Development. {quote} Which are the two ways? I would think that the two ways are automatic or manual installation, i.e. 1.a. and 1.b. So what is 2. then? If it's about using docker and openshift, then that's surely not one of the ways to install it, no?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-55
3457	RHDEVDOCS-52	[~mlabuda], can we call this doc complete or are we looking at more edits please. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-52
3458	RHDEVDOCS-52	Hi [~mhusnain], can you please publish these upstream. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-52
3459	RHDEVDOCS-49	[~optak], is sshfs replacing all sync methods (NFS, VirtualBox shared folders), or only rsync? Even if the other methods remain available, do we want docs to mention them, or only mention sshfs?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-49
3460	RHDEVDOCS-49	Also, rsync was used during the provisioning phase in 2.0, i.e. rsync needed to be installed (Vagrant wouldn't start on Win without rsync.exe in PATH). Is this still the case, or is rsync now not needed at all?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-49
3461	RHDEVDOCS-44	Hi [~mlabuda], can you please review the artilce at: https://docs.google.com/document/d/1Nq7DogBtqVhPqxYFwER-CfzU2uK0dZHBFeNAO2Nf_xw/edit#heading=h.7o8zpk6u9vhw. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-44
3462	RHDEVDOCS-44	> WIP: Incorporating Marian's feedback	1	1	https://issues.jboss.org/browse/RHDEVDOCS-44
3463	RHDEVDOCS-44	Hi [~mlabuda], do you think this article looks good now? Or are we looking at more edits? Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-44
3464	RHDEVDOCS-44	Hi [~mhusnain], can you please publish these in the community. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-44

#	Requirement	Questions	Answered	Successfully Answered	URL
3465	RHDEVDOCS-41	[~ldimaggio] good question and I'm not sure. Can you tag the person who would have the information for Hibernate versions for the 9.x and 10 release?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-41
3466	RHDEVDOCS-41	Wouldn't this be Nick? [~nickboldt]	0	0	https://issues.jboss.org/browse/RHDEVDOCS-41
3467	RHDEVDOCS-41	[~nickboldt], can you let me know the Hibernate versions please, or we'll have to release as-is without this information. Thanks! :)	1	1	https://issues.jboss.org/browse/RHDEVDOCS-41
3468	RHDEVDOCS-41	[~mhusnain], right now we cannot view neither https://access.redhat.com/node/427493/ nor https://access.redhat.com/node/427493/draft. Both give access denied. This is a big issue, because the former is live document that is linked from many places. Have you accidentally made the document inaccessible? Update: This is now fixed and the old version is now published again and viewable without logging in.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-41
3469	RHDEVDOCS-41	Regarding the hibernate versions. [~mhusnain], can you check the previous versions to see why/when were the hibernate versions removed? Or did we never have them listed for devstudio 9?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-41
3470	RHDEVDOCS-41	Thanks Martin and Rastislav! [~kaerts], could you confirm the above please?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-41
3471	RHDEVDOCS-41	Thanks, Koen. Since we had some unexpected behaviour with customer portal articles not displaying at all when in draft state, I've saved the 10 version here instead: https://access.redhat.com/articles/2377301 I'll copy the contents from this article to the original one for GA to prevent breaking any links to the old doc. [~mmalina], is the content at the new link looking ok?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-41
3472	RHDEVDOCS-40	Copied over and updated document here: https://access.redhat.com/articles/2360061 [~mmalina], could I get you to review the updates, which are as follows: * the "Update a JBoss Developer Studio 10.x Installation" section is entirely new * the table in the "Update the Version of Eclipse Used with JBoss Developer Studio BYOE" section has a new row for 10.0	1	1	https://issues.jboss.org/browse/RHDEVDOCS-40
3473	RHDEVDOCS-40	Can we add a note to the document informing the users that while they cannot update from DevStudio 9 to 10, they are able to use workspaces created with DevStudio 9 with DevStudio 10? Thx!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-40
3474	RHDEVDOCS-40	[~ldimaggio] no problem. Added the bold line to the paragraph: Each major version of JBoss Developer Studio is based on a different version of Eclipse. As a result, it is not possible to update an older major version of JBoss Developer Studio to a newer major version. For example, it is not possible to update a devstudio 8.x installation to 9.x, nor is it possible to update a 9.x installation to 10.x. *However, it is possible to use the workspaces created using JBoss Developer Studio 9.x with JBoss Developer Studio 10.x.* Does this look OK to you as well now? Edit: Link to verify - https://access.redhat.com/articles/2360061	1	1	https://issues.jboss.org/browse/RHDEVDOCS-40
3475	RHDEVDOCS-39	Is this just for CDK 2.1, or does it apply to 2.0 as well?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-39
3476	RHDEVDOCS-38	I was given this bit of (unverified) info by one of ADB devs: "Apparently docker-1.10 and docker-1.9 can both pull containers without regard to which version built the container as long as the pull is done by tag. If you pull by SHA-256 digest, images built with docker-1.10 cannot be pulled by docker-1.9. I have not independently verified this information yet." Could you confirm that this is the case, please? [~optak]	1	1	https://issues.jboss.org/browse/RHDEVDOCS-38
3477	RHDEVDOCS-33	[~mmalina], could you or someone on your team have a look at this and review please. I have a couple of things in particular to confirm with QE: - In the **About This Release** section: a. "It includes Eclipse Neon" - Formerly, this was Eclipse Mars.1 for 9.1. Do we need to state anything other than Eclipse Neon for 10? b. "t requires a minimum of Java 8 to run." - Is this still true? I'll fill in the N&N content once that is available, but I'd like to get review on the rest earlier, if possible.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-33
3478	RHDEVDOCS-33	I have another comment: While it's true that we still require Java 8 to run devstudio, I wonder if we need this mentioned in the About This Release article. Because it says: {quote}The upgrade has the following purposes: ... It requires a minimum of Java 8 to run.{quote} I don't think it makes sense to say that the java 8 requirement is one of the purposes of the new version. Another issue is here: Apply This Release To install JBoss Developer Studio BYOE 10.0 in Eclipse Mars ... - this needs to be updated to say "Eclipse Neon" Another thing I was wondering about: Do we need to keep repeating all that stuff in About Red Hat JBoss Developer Studio and Use Cases for JBoss Developer Studio ? Especially since it's just copy&paste from the previous release. Maybe the first short overview would be enough? [~ldimaggio], what do you think? And now the usual thing - JIRA queries :) Resolved Issues Right on the first sight I can see that it's wrong - the versions simply don't match: {code} (project in (JBDS) AND affectedVersion < 10.0.0.GA AND fixVersion in (10.0.0, 10.0.0, 10.0.0.Alpha2, 10.0.0.Alpha3, 10.0.0.Alpha4, 10.0.0.Alpha5, 10.0.0.Alpha6, 10.0.0.Beta1, 10.0.0.Beta2, 10.0.0.CR1, 10.0.0.CR2, 10.0.0.CR3, 10.0.0.CR4, 10.0.0.CR5, 10.0.0.Final, 10.0.0.GA) OR project in (JBIDE) AND affectedVersion < 4.4.0.GA AND fixVersion in (4.4.0.Alpha1, 4.4.0.Alpha1-TP, 4.4.0.Final, 4.4.0.GA)) AND type in (Bug) AND resolution in (Done) {code} Can you fix that? Thanks! Likewise for the Known Issues query - please fix the versions, there is no 4.4.0.GA.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-33
3479	RHDEVDOCS-33	Thanks for the feedback, [~mmalina] 1. "This upgrade has the following purposes" - You're right, this sounds odd. I've changed it to features. 2. s/Mars/Neon in one instance 3. For the introductory matter, we include it because we have to provide such context for releases. I've removed the extra matter such as style and legal information, if you compare to the previous ones. 4. Sorry, I'm not sure I understood the problem with the query. Is it just that there isn't a 4.4.0.GA version, or is there something else wrong in there?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-33
3480	RHDEVDOCS-33	[~mmalina], could you let me know what the issue is with the query please? And any other issues, please let me know so I can lock this down for GA.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-33
3481	RHDEVDOCS-26	[~rkratky], I can't seem to access the devtools-team list archives. Can you paste the subject of the thread in question so I can look it up, since I am subscribed to the list in my inbox anyway.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-26
3482	RHDEVDOCS-26	[~rkratky], ping - have you heard any more about this?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-26
3483	RHDEVDOCS-18	Hi [~supriya.bharadwaj], How does it look with the images? Anything I can help with?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-18
3484	RHDEVDOCS-18	Hi [~rkratky], I am afraid I have not been able to work on those. Was busy filling my Compass on Friday and on Monday was trying to get the OpenShift 3 things sorted for the devstudio 10 release. I'll work on the images today and tomorrow, so that I can send them to you by tomorrow EOD (June 08). Is that okay?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-18
3485	RHDEVDOCS-16	> Completed conversion to adoc [~mhusnain], can you please post these upstream. Thanks!	0	0	https://issues.jboss.org/browse/RHDEVDOCS-16
3486	RHDEVDOCS-16	Hi [~mhusnain], I don't think you created a PR for this article for upstream. Can you please do that so that I can resolve this issue (or we ca do this when you walk me through creating a PR as we discussed y'day). The doc's published downstream for the 10.0 release. Thanks!	0	0	https://issues.jboss.org/browse/RHDEVDOCS-16

#	Requirement	Questions	Answered	Successfully Answered	URL
3487	RHDEVDOCS-14	Found the following issues that may affect documentation: JBIDE-22146 - OpenShift - Create Resource wizard = automated way to "oc create" JBIDE-21806 - OpenShift - Create OS service adapted from service/route context menu JBIDE-21524 - OpenShift - deploymet view added JBIDE-21519 - OpenShift (related to above) - builds tab for deployment view JBIDE-21328 - OpenShift - Create Server Adapter for OpenShift 3 applications JBIDE-20990 - CDK - Add "CDK Server" for start/stop for the CDK JBDS-3846 - CDK - Create container oriented perspective (not done yet, but a good one to plan on for later) Added for Platform Installer: JBDS-3787 Installer - Installation of each component should be optional JBDS-3666 - Back off and retry if download fails	0	0	https://issues.jboss.org/browse/RHDEVDOCS-14
3488	RHDEVDOCS-13	Hi [~supriya.bharadwaj], can you add the status for this here and whether it will be ready for GA.	1	1	https://issues.jboss.org/browse/RHDEVDOCS-13
3489	RHDEVDOCS-13	Hi [~mhusnain], can you please post the article upstream. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-13
3490	RHDEVDOCS-13	Hi [~mhusnain], can you please create a PR for this to be published upstream? The doc's published downstream for the 10.0 release. Thanks!	1	1	https://issues.jboss.org/browse/RHDEVDOCS-13
3491	RHDEVDOCS-10	[~rkratky], what is the scope for these changes? * Just for devsuite docs * devsuite + devstudio 10 * devsuite + devstudio <currently released version, possibly 9.1> onwards	1	1	https://issues.jboss.org/browse/RHDEVDOCS-10
3492	RHDEVDOCS-10	[~rkratky], I was looking into the devstudio docs and it looks like all instances of naming should be covered by the common content definitions now. Do you think this is still worth combing through the docs?	1	1	https://issues.jboss.org/browse/RHDEVDOCS-10
3493	RHDEVDOCS-2	Thanks Robert for link .. yes it fulfills the requirement There is one thing to update in doc vagrant is showing as "vafrant" starting from : 2.4.1. Initializing Vafrant boxes with Hyper-V 2.4.2. Initializing Vafrant boxes with VirtualBox 2.4.3. Registering and Checking the Vafrant Box	1	0	https://issues.jboss.org/browse/RHDEVDOCS-2
3494	JBDOCS-323	Hi Mark Should we move this over to JDOCBOOKSTYLE? I know it doesn't relate specifically to the XSLT for the community base styles, but it might be something we can manage as part of the offering under PressGang? Cheers J	1	0	https://issues.jboss.org/browse/JBDOCS-323
3495	JBDOCS-303	Why is this assigned to me?	0	0	https://issues.jboss.org/browse/JBDOCS-303
3496	JBDOCS-279	Mark, I have a feeling you use Windows? If so, can you take a look? As I don't have Windows. Jay can provide more details I think - he says it is bad enough we should fix it for the next release.	0	0	https://issues.jboss.org/browse/JBDOCS-279
3497	JBDOCS-177	Added docbook-xsd to docbook-support. You can link to the schema like this: (obviously using the correct relative path) <article xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:noNamespaceSchemaLocation="../../docbook-support/support/docbook-xsd/docbook.xsd"> <!-- ... --> </article> This should allow full validation by your editor while being compatible with the existing docbook-support tools. I haven't thoroughly tested it, but it does appear to be working. I'll probably migrate the admin guide over as a test. Until we get more experience, this should be considered purely experimental.	0	0	https://issues.jboss.org/browse/JBDOCS-177
3498	JBDOCS-38	These are supposed to be updates to the admin-devel/adminguide in cvs right? I could see only the following missing: throw-runtime-exceptions <!-- Should runtime exceptions be thrown when unexpected exception occurs ? The throw-runtime-exceptions element must be one of the two following: <throw-runtime-exceptions>true</throw-runtime-exceptions> <throw-runtime-exceptions>false</throw-runtime-exceptions> --> <!ELEMENT throw-runtime-exceptions (#PCDATA)> This is used when getting a connection from the datasource to create or find EJB instances. If DataSource.getConnection() failed and throw-runtime-exceptions is true then EJBException will be thrown. If throw-runtime-exceptions is false then CreateException will be thrown in case of instance creation and FinderException will be thrown in case finder.	1	1	https://issues.jboss.org/browse/JBDOCS-38
3499	JBDOCS-30	SourceForge Username: ejain . Logged In: YES user_id=474829 +1 This would seem trivial to set up, is there any specific reason why this has not been done yet?	1	1	https://issues.jboss.org/browse/JBDOCS-30
3500	JBDOCS-28	To be clear what this asks for is a simple page explaining: 1. What is an MBean notification 2. How do you create a listener for one 3. How do you get a notification that the JBAS server is done starting up. With liberal code samples. The referenced page does not really answer that.	0	0	https://issues.jboss.org/browse/JBDOCS-28
3501	JBDOCS-22	(sorry about the truncated note) The above link contains docs that were last maintained August 9 2009. Is it possible to use them as a starting point?	0	0	https://issues.jboss.org/browse/JBDOCS-22
3502	TOOLSDOC-734	Hi Nick, If I've understood the requirement, what normally happens with a branding change is that the old versions of the products docs remain under the old name (https://access.redhat.com/documentation/en/red-hat-jboss-developer-studio/) and the new versions of the combined product will be a brand new product called Red Hat Developer Studio and would be under the url https://access.redhat.com/documentation/en/red-hat-developer-studio/ I'm adding [~rkratky] to this so that he has visibility on this request and because he's working on merging the old docs for the new product, etc.	0	0	https://issues.jboss.org/browse/TOOLSDOC-734
3503	TOOLSDOC-726	Hi [~ldimaggio], from the email, I see that the requested information was this: I see the following use cases: 1. Installing CDK on Win/Mac/Linux manually without our all-in-one installer. As far as I understand it will be covered by the CDK get started guide. 2. Using manually installed CDK in JBDS. Users have to do a few additional but not obvious steps in JBDS to set up a CDK (create a CDK server adapter, set up oc cli tool, etc). 3. Using our all-in-one windows installer. Some info about how to use the installer. 4. Getting started with CDK in JBDS after using our all-in-one win installer. In that case the installer does set up a CDK in JBDS for you and you don't need to do steps from #2. So, this case is the next step for #2 and #3. I think we've covered all of this in the CDK tooling article here: https://access.redhat.com/articles/2388671 What do you think, Len? I think we can probably close this issue.	0	0	https://issues.jboss.org/browse/TOOLSDOC-726
3504	TOOLSDOC-720	[~ldimaggio] I'm not familiar with how to do this for Windows. Is it sufficient to just tell users to enable virt in the BIOS or provide the steps?	1	1	https://issues.jboss.org/browse/TOOLSDOC-720
3505	TOOLSDOC-718	[~fbricon] - I sent a following up question to others forgetting that you responded to this for me, so maybe you're the one to ask: Is the server adapter the only real use case we want to talk about for CDK? Basically: does it do anything else we want to document?	1	1	https://issues.jboss.org/browse/TOOLSDOC-718
3506	TOOLSDOC-702	Hi [~mhusnain], I'm sorry I made one mistake in the list. In JBDS-IS 8.0.4.GA the Drools plugin was updated to 6.3.0.Final (the other components are correct). Could you update the version of Drools plugin to 6.3.0.Final, please? Thank you!	1	1	https://issues.jboss.org/browse/TOOLSDOC-702
3507	TOOLSDOC-697	[~psrna], can you please check if the version of Apache Cordova in the list are correct for Beta1? Or possibly other versions where you are familiar with what's included. Thanks!	1	1	https://issues.jboss.org/browse/TOOLSDOC-697
3508	TOOLSDOC-697	[~supriya.bharadwaj], can you order the numbers by their value? It's strange to have 3.9.2 between 4.0.1 and 3.5.0.	1	1	https://issues.jboss.org/browse/TOOLSDOC-697
3509	TOOLSDOC-694	[~ldimaggio], what do you think about the idea of not having an article at all? The page on CSP has enough details already. And now there is the blog post at http://tools.jboss.org/blog/ga-cve-for-mars.html so we can just link to it if necessary.	1	1	https://issues.jboss.org/browse/TOOLSDOC-694
3510	TOOLSDOC-694	+1 Yes - let's keep this simple - can we simply refer to the blog post? Thx!	1	0	https://issues.jboss.org/browse/TOOLSDOC-694

#	Requirement	Questions	Answered	Successfully Answered	URL
3511	TOOLSDOC-687	Hi [~mhusnain], I started looking at the release notes and I have one question. Where are the RNs from 9.0.0.CR1? Looking at https://access.redhat.com/documentation/en/red-hat-jboss-developer-studio/ the CR1 version is there, but no docs. Shouldn't it be there? (For the record, I find it easier to have a previous version of RNs open when reviewing a new version, so that I can easily see what is actually changed and what is just carried over from the past.)	0	0	https://issues.jboss.org/browse/TOOLSDOC-687
3512	TOOLSDOC-684	Downloaded Eclipse Mars We must change Chapter 1. Beta Release (@Misha, where can I get this info from?) Install JBDS BYOE when Online (https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Developer_Studio/9.0/html/Install_Red_Hat_JBoss_Developer_Studio/Install_JBoss_Developer_Studio_BYOE_when_Online.html): Fig. 2 must be changed; but, we wouldn't have this SS ready until GA is released on Tuesday. ^ step 1: Changed Beta 1 to GA	1	1	https://issues.jboss.org/browse/TOOLSDOC-684
3513	TOOLSDOC-684	> Spoke to Jan and he isn't sure if they will release GA or CR2 next week. [~mhusnain], wherever necessary, wdyt should the version of JBDS be in the Install guide, CR2 or GA? > I have taken the SS to be included in the Install JBDS BYOE Online chapter with "CR1". I will not be able to take a CR2/GA SS until the product is released next week and is available to be installed through the eclipse marketplace (right now since it isn't released it is not available through marketplace)	1	1	https://issues.jboss.org/browse/TOOLSDOC-684
3514	TOOLSDOC-679	Hi [~ldimaggio], in case you want to have a look at the instructions, here's the link to it: https://docs.google.com/document/d/16Glf94lJVrdZbV5NUEBw8AN0DgHoP7i48_3J9UA_NTQ/edit#heading=h.z46fmakmdrhm Thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-679
3515	TOOLSDOC-679	Hi [~ldimaggio], the link that you have mentioned in your comment speaks of the same steps that I have documented in the Prerequisites section of the doc. Do you mean that I must let the existing pre-req steps be as-is and also explicitly mention the link in the introduction to the section? Thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-679
3516	TOOLSDOC-679	I see that now - thanks - but the actual steps on how to do that - where are they?	0	0	https://issues.jboss.org/browse/TOOLSDOC-679
3517	TOOLSDOC-679	[~mlabuda], I have made changes to the doc as discussed with you. Can you please have a look at it an dlet me know if it needs anymore edits? Thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-679
3518	TOOLSDOC-677	Hi [~ldimaggio], tagging you in case you wish to have a look at the doc. Thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-677
3519	TOOLSDOC-677	Hi [~rawagner], can you please review this?	0	0	https://issues.jboss.org/browse/TOOLSDOC-677
3520	TOOLSDOC-677	Hi [~rawagner] and [~ldimaggio], I have incorporated the feedback and also added a few sentences and a screenshot for webhooks in the article. Can you please have a look at it and let me know of any edits? Location: https://docs.google.com/document/d/1mevNqMkYVe41Lc_-NcfvYC8lcbD-FFJqoSZaxUzr7-s/edit#heading=h.lqc95xrzxc31 Thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-677
3521	TOOLSDOC-677	[~mlabuda], thanks for the detailed review! I have incorporated your feedback in the article. Can you please have a look at it and let me know of any more edits? - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-677
3522	TOOLSDOC-676	Update: added this to add a procedure about adding OSv3 tooling to JBDS. Ready for review.	1	1	https://issues.jboss.org/browse/TOOLSDOC-676
3523	TOOLSDOC-676	Few comments: - instead of "OpenShift 3 is not included in JBoss Developer Studio but is instead available as a Technology Preview feature." I would use "OpenShift 3 is not installed by default in JBoss Developer Studio but is instead available as a Technology Preview feature.", because saying it is not included may make felling that it is not available - i don't see there "JBoss Tools EGit Integration", I think it is bundled together with JBoss OpenShift 3 Tools, my guess. Do you really see also EGit integration feature in software update tab? - connect to an OpenShift 3 instance has step "3 a. Ensure that is selected from the Connection drop down menu." I think there is missing something. Shouldn't it be "a. Ensure that <New Connection> is selected from the Connection drop down menu."? - how is it with terminology? There are on several places used term "field", but the server type and connection are combo boxes. Is it desired? Term field is more likely to be used for text fields instead, I think - step "e. Click Finish to complete the connection configuration." - completing configuration of connection? I would say we are created a new connection to an OpenShift 3 instance or establishing a new connection, but not configuring connection, what do you think? :)	1	1	https://issues.jboss.org/browse/TOOLSDOC-676
3524	TOOLSDOC-676	Re: the field question - I've used field to talk about text boxes where the user must type in a value and combo boxes and other items where users need to select one value out of several included, etc. Re: Completing connection configuration - Agreed, thank you! I've made this change as well. [~mlabuda], could you have a look now and let me know if you have any further feedback?	1	1	https://issues.jboss.org/browse/TOOLSDOC-676
3525	TOOLSDOC-676	I still see "e. Click Finish to complete the connection configuration." in there. And there are 2 comments from Rastislav Wagner, could you please incorporated them as well? Thanks.	1	1	https://issues.jboss.org/browse/TOOLSDOC-676
3526	TOOLSDOC-676	Ok. The article looks fine. Last comment: the screenshot of a New Connection wizard - wouldn't it be better to have a valid credentials for localhost instead of invalid URL on a screenshot? E.g. using https://localhost:8443. To show users how it really could look like.	1	1	https://issues.jboss.org/browse/TOOLSDOC-676
3527	TOOLSDOC-676	Updated here: https://access.redhat.com/articles/1609003 And created a PR for this change in JBT docs here: https://github.com/jbosstools/jbosstools-website/pull/490 If these are now OK, we can set this to resolved. What do you think, [~mlabuda]?	1	1	https://issues.jboss.org/browse/TOOLSDOC-676
3528	TOOLSDOC-674	Is this ready to close off, [~supriya.bharadwaj]?	1	1	https://issues.jboss.org/browse/TOOLSDOC-674
3529	TOOLSDOC-673	Hi [~mhusnain], the QE review for this how-to (https://docs.google.com/document/d/1W24dw5V3FReluKwT03286CmFqnMcsIvYkC08HBGLiP0/edit#) is complete (JBQA-12366). Can you please have a look at it so that I can then convert it to MarkDown and ASCIIDOC. Thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-673
3530	TOOLSDOC-673	[~mhusnain], can you please post these in the community. Thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-673
3531	TOOLSDOC-670	Added/changed the info. Sent .adoc file and screenshot to [~mhusnain] as part of another JIRA (TOOLSDOC-608)	0	0	https://issues.jboss.org/browse/TOOLSDOC-670
3532	TOOLSDOC-669	Hi [~mmalina], could you review the added note here and let me know if you have any feedback? Once the text is finalized, I'll add it to the JBT version of this article as well: https://access.redhat.com/articles/1515503	1	1	https://issues.jboss.org/browse/TOOLSDOC-669
3533	TOOLSDOC-650	[~supriya.bharadwaj], it occurs to me that in the future it may reduce the margin for error by having you locate the appropriate adoc file in the jboss tools docs git repo and copy it locally, then make the changes/insertions required, test it locally by viewing it and then adding that file as an attachment. The benefits of this approach are: * You have greater control over the placement and formatting of the changes/insertions * You test it as you write the content, so this means I can just paste the exact changes you require in and let you know they are available. This occurs to me because while you got the image::<name of image> part correct this time, the formatting of the image is not correct: Select the Android SDK Location image::first_hybrid_app_android_preferences.png[Select the Android SDK Location] . Click `Apply` and then click `OK`. This would not render corrected upon testing. This should be: .Select the Android SDK Location image::first_hybrid_app_android_preferences.png[Select the Android SDK Location] What do you think about the new approach?	1	1	https://issues.jboss.org/browse/TOOLSDOC-650

#	Requirement	Questions	Answered	Successfully Answered	URL
3534	TOOLSDOC-650	[~mhusnain], I agree with you on this approach. It would be great if you could help me, if the need be, in understanding how to do this the first time around. I have worked in Git just once so I am not very sure how to go about with this. Is this doable?	1	1	https://issues.jboss.org/browse/TOOLSDOC-650
3535	TOOLSDOC-647	[~mhusnain], can you please set the metadata for the issue? Thanks - Supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-647
3536	TOOLSDOC-647	> Nick has confirmed that both the zips are required for the offline install (JBDS-3491). The section "http://docbuilder.usersys.redhat.com/23020/#Install_JBoss_Developer_Studio_BYOE_when_Offline" states this already. [~mhusnain], I think we are good to publish the guide now. WDYT? thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-647
3537	TOOLSDOC-646	This information does not replace the important information, which warns users about a separate issue. Prerequisites sounds good, but this information was in two articles. Which one did you mean, or is it both?	1	1	https://issues.jboss.org/browse/TOOLSDOC-646
3538	TOOLSDOC-645	Hi [~rawagner], can you please review the maven how to? I briefly looked at it, but it will be better if someone with better knowledge of the maven tooling takes a look. Regarding the note from Misha: {quote} <!-- Note for Martin: If I understand it correctly, m2 is now included in JBDS and therefore there is no extra step to configure this beyond this unlike in JBDS 7.1 (https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Developer_Studio/7.1/html/Maven_Tools_Reference_Guide/adding_maven_support.html) -> {quote} m2 is included, but I think it was included even in JBDS 7.1. What the linked article is about is how to add maven support to an old-style faceted project - when you use the new maven project wizard, you don't have to do this. Am I correct, [~rawagner] ?	1	1	https://issues.jboss.org/browse/TOOLSDOC-645
3539	TOOLSDOC-645	Changes: * P1: Removed "ensure" from the SNAPSHOT instructions to add instructions about using a unique value that may or may not be the same as the example. * P2: Updated the text so that it is less restrictive about the value of *SNAPSHOT * s/Project Explorer tab/Project Explorer view in the entire document * P3: Updated name and explanation for Procedure 3 to hopefully address the fact that the previous procedures tell us how to create a maven supported project and module from scratch while the third is to add maven support to an existing application without maven support. * P3: Regarding the right click on a new project, click configure and in the menu, select "convert to maven project", the problem is that I can't spot any option for maven projects (in 8.1). in 8.1, would there be an alternate way to do this configuration through the right click project name > Maven> <some option>? And in 9.0, is it going to be right-click project name>configure>convert to maven project? <-- [~rawagner]	1	1	https://issues.jboss.org/browse/TOOLSDOC-645
3540	TOOLSDOC-645	[~rawagner] - I also wrote up a small procedure at the end of this page, could you please review that as well if you have the time: https://docs.google.com/a/redhat.com/document/d/13L561ZsVvhqMyVM4AJHgGZ8xgR2IqIjz4KX-b2_mlvQ/edit?usp=sharing	1	1	https://issues.jboss.org/browse/TOOLSDOC-645
3541	TOOLSDOC-643	Hi [~burrsutter] - I would actually like to get a basic How To for Maven done for 8.1.0. i was looking at JBDS-3279 and trying to get some ideas. For a first basic How To, what kind of tasks are users looking to do? Create a Maven project + ... For a more advanced How To, what kind of tasks? Import an existing Maven project + ... I definitely need some help with ideas here! I'm about as good as your average JBDS user who can't figure out what to do with Maven in JBDS :)	1	1	https://issues.jboss.org/browse/TOOLSDOC-643
3542	TOOLSDOC-643	[~fbricon] Fred, I would be interested in your comments as well my comments 2. "Install Maven in JBDS" - it should already be pre-installed b. Create a maven module - does this mean a multi-module project? how to enable auto download of source & javadoc - this is a great idea explain autoconfiguration of plugins - this is a great idea How about "how to recover when things get screwed up"? when "clean", "build", "compile", "install" are simply not good enough the magic of "Update Project" and "Select Maven Profiles"	1	1	https://issues.jboss.org/browse/TOOLSDOC-643
3543	TOOLSDOC-641	Thanks, [~apodhrad] - I did have a couple of questions: 1. I'm trying to develop information about why a user would set up a remote server as an intro for this article. My general impression is that it is to have the web application and its data stored in a location that is not corrupted if the local data is corrupted or lost, and it allows users to interact remotely with applications and change and update things as required. Can you tell me if I'm on the right track or completely off? 2. For the limited functionalities in step 4 above, I have included them because that seemed to be the direction we pointed users to in the 7.1 edition of these instructions. If these are not necessary, I am happy to replace it with something more open-ended, such as for example, telling users to select either "Filesystem and shell operations" or "Management operations" depending on their remote server requirements. Would this be better? I'll implement the input above tomorrow, and if you can find the time to add some quick answers to the above, I'd appreciate it, Andrej. :)	1	1	https://issues.jboss.org/browse/TOOLSDOC-641
3544	TOOLSDOC-641	[~mmalina], could you help at all with this? Thank you for your help, Andrej! :)	1	1	https://issues.jboss.org/browse/TOOLSDOC-641
3545	TOOLSDOC-641	Hi [~mhusnain], I reviewed the current draft and have a few comments: 4.b Note: If you select Management Operations for the Controlled by field, you must set up an admin user on the server by editing the $SERVER_HOME/bin/add-users.sh file and enter the same credentials in the server editor or during the server start. - this is not correct. You're not supposed to edit the shell script. The shell script is used to add the user. (And also, it's add-users.bat on windows, but I'm not sure how much detailed you want to be.) So instead of "by editing" I would say "by using" 4.c The Server is externally managed. Assume server is started field is used when the user wants to deploy the server but does not want the IDE to stop or start the server for them. - "deploy the server" -> "deploy to the server" 4.d - it's good except for the end: "or leave the box unselected as a default" - this is not true - default is that it is checked (and a new runtime will be created in the next step if there is none existing) 5.a - I suggested that we could say something like "You may need to create a new host." Do you not agree? 5.c.d This location is within the directory specified in the Remote Server Home field. - while technically it is true that it is within the directory, it's not directly there, it's deeper - it's in SERVER_HOME/BASE_DIRECTORY/configuration/ - maybe it would be better explicitly say that?	1	1	https://issues.jboss.org/browse/TOOLSDOC-641
3546	TOOLSDOC-641	All of the above seem reasonable, Martin. I've implemented them in suggestion mode here: https://docs.google.com/document/d/1sd5_RaThCjbRxeA4D-YcYjXLYXl3U3zEzpxyjARXk18/edit Suggestion mode provides an inline diff in green so hopefully this makes it easier to review the changes. If you are OK with all of the changes, I'll apply them to the enterprise and community docs and push them public. Could you have a look at your earliest convenience, please [~mmalina]?	1	1	https://issues.jboss.org/browse/TOOLSDOC-641

#	Requirement	Questions	Answered	Successfully Answered	URL
3547	TOOLSDOC-641	Hi [~mhusnain], Thanks for the update, but it's still not complete ;) In 4.b you changed "by editing" to "by using", but it would be good to also change "file" to "script" - "by using the $SERVER_HOME/bin/add-users.sh script" - this is a minor thing, but since there are still more changes necessary, I thought I'd include it as well. But this is the only thing that is changed. What about all the rest? None of them are fixed: 4.c The Server is externally managed. Assume server is started field is used when the user wants to deploy the server but does not want the IDE to stop or start the server for them. "deploy the server" -> "deploy to the server" 4.d it's good except for the end: "or leave the box unselected as a default" this is not true - default is that it is checked (and a new runtime will be created in the next step if there is none existing) 5.a I suggested that we could say something like "You may need to create a new host." Do you not agree? 5.c.d This location is within the directory specified in the Remote Server Home field. while technically it is true that it is within the directory, it's not directly there, it's deeper - it's in SERVER_HOME/BASE_DIRECTORY/configuration/ - maybe it would be better explicitly say that? BTW, how does the suggestion mode work? I can only see View -> Show Spelling Suggestions, but that doesn't seem to do anything.	1	1	https://issues.jboss.org/browse/TOOLSDOC-641
3548	TOOLSDOC-641	[~mhusnain], these two are still not done: 4.d it's good except for the end: "or leave the box unselected as a default" - this is not true - default is that it is checked (and a new runtime will be created in the next step if there is none existing) 5.a I suggested that we could say something like "You may need to create a new host." Do you not agree?	1	1	https://issues.jboss.org/browse/TOOLSDOC-641
3549	TOOLSDOC-641	Thanks, [~mmalina]. Comments: 4.d: The default part is removed, so I'm not sure what you mean here. Did you mean 4.c.? What I see in 4d at the moment: A remote server can now be created without assigning a runtime to it. Depending on the requirements, select the Assign a runtime to this server check box (and select an existing runtime or create a new one) or leave the box unselected. 5.a: I added the new host line in 5 (intro) rather than 5a, but I've fixed it now.	1	1	https://issues.jboss.org/browse/TOOLSDOC-641
3550	TOOLSDOC-628	Hi [~mhusnain], could you a link to the approprite doc which should be reviewed? Btw, the following issues from [1] should not be included in release notes (they were reopened) RH Bugzilla * 454749 * 459209 Eclipse Bugzilla * 1178839 * 1184422	1	1	https://issues.jboss.org/browse/TOOLSDOC-628
3551	TOOLSDOC-628	Hey [~mhusnain] - Was there any closed Jira info from Drools/jBPM6 ? Looks fine otherwise. --paull	1	1	https://issues.jboss.org/browse/TOOLSDOC-628
3552	TOOLSDOC-622	Hi Misha, I looked at http://docbuilder.usersys.redhat.com/23017/#Known_Issues and tried to compare it to CR1 and there is no difference in the JIRAs listed under Known Issues - it's exactly the same as in CR1. Can you specify what you added? I have replied to your email about the queries - let's wait for Michelle's feedback on the Known Issues query. But for Resolved Issues, I believe we can use this: {code} (project in (JBDS) AND affectedVersion < "8.1.0.Beta1" AND fixVersion in (8.1.0.Beta1, 8.1.0.CR1, 8.1.0.CR2, 8.1.0.GA) OR project in (JBIDE) AND affectedVersion < "4.2.3.Beta1" AND fixVersion in (4.2.3.Beta1, 4.2.3.CR1, 4.2.3.Final)) AND type in (Bug) AND resolution in (Done) {code} Also, I have a question. I wanted to view the 8.1.0.CR1 Release Notes, but they are nowhere to be found. I looked here: https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Developer_Studio/ I'm pretty sure they were there under Beta1 last week when I reported TOOLSDOC-618. I understand that we might want to remove them once GA is released, but that hasn't happened yet. Do you know what happened? I used this for the comparison: http://docbuilder.usersys.redhat.com/23055/#Known_Issues	1	1	https://issues.jboss.org/browse/TOOLSDOC-622
3553	TOOLSDOC-612	Thanks, Andrej. In order: Section 2: Amended list to *updated* and removed early access text. Section 3: Amended and it should now have 7.1 issues only. (Update: Paul responded as well while I was writing this, so I have incorporated only the three he pointed out) Section 5: Removed former Section 5 about Drools. (New) Section 5: Added Fuse section with the one issue specified above. Section 6: Waiting to hear. Section 7: Former section 7 is excluded. For BPMN2 issues: *Added section for BPMN2 Modeler * Added issues from 1.0.3, 10.4 (nothing for 1.0.5). Did not separate these by version, but can do if needed. [~apodhrad], could you have a look at the fixed issues above while we wait for teiid feedback please?	0	0	https://issues.jboss.org/browse/TOOLSDOC-612
3554	TOOLSDOC-610	Hi [~mhusnain], can you please post these upstream? Thanks - Supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-610
3555	TOOLSDOC-609	What is the status for this and is there a link to the draft?	1	1	https://issues.jboss.org/browse/TOOLSDOC-609
3556	TOOLSDOC-609	Hi [~supriya.bharadwaj] - what's the status for this issue? This needs to be done for GA so please update here.	1	1	https://issues.jboss.org/browse/TOOLSDOC-609
3557	TOOLSDOC-609	Hi [~mhusnain], this issue is similar to the icon change one (JBDS-3270). I have documented the steps and got them reviewed from Pavol (instructions at: https://docs.google.com/document/d/124ptroykZasftYjSQTGCd3rSYRibWABUfqqDlVTwV30/edit). But, since we have been unable to get the app to show on the emulator (JBDS-3454), I have not tested the steps. Thanks - Supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-609
3558	TOOLSDOC-608	[~mhusnain], can you please post these upstream? thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-608
3559	TOOLSDOC-595	[~supriya.bharadwaj], could you comment with an update on this issue please?	1	1	https://issues.jboss.org/browse/TOOLSDOC-595
3560	TOOLSDOC-595	Updated info for JBDS 9.0Beta2 in the install guide and published the guide. Location: https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Developer_Studio/9.0/html/Install_Red_Hat_JBoss_Developer_Studio/Install_JBoss_Developer_Studio_BYOE_when_Offline.html Info yet to be edited for JBDS IS; shall check when can this be completed and will add a comment as needed.	1	1	https://issues.jboss.org/browse/TOOLSDOC-595
3561	TOOLSDOC-583	we don't really have docs grouped by components. maybe we should "tag" the docs and generate a page with those related to openshift ? it won't be a "get from A to B" kind of page though ...is that usable ?	1	1	https://issues.jboss.org/browse/TOOLSDOC-583
3562	TOOLSDOC-576	what does DYK mean? :) Btw I'd rather add info that livereload will not work with java 7 if web engine is javaFx. Why - see [FAQ|http://tools.jboss.org/documentation/faq/browsersim.html#main-features]	1	1	https://issues.jboss.org/browse/TOOLSDOC-576
3563	TOOLSDOC-521	[~mmurray], sorry for the delay. As for the "Resolved" query. I can see the new "fixVersion not" in there - what's the purpose? Is that to avoid listing JIRAs that have been fixed in the maintenance (4.1.x) as well as in 4.2? Otherwise, I still don't like the "affectedVersion" there - if I remove it, the list gets from 180 JIRAs to more than 1000. If I search only for JIRAs which have affectedVersion empty, I get 128 JIRAs - these are JIRAs that could have been reported for any version at all, we don't know. But ok, the majority of JIRAs has this set. Still, the fact that a JIRA has "affects 4.2.0.Beta3" does not necessarily mean that 4.1 is not affected - I would estimate that more often than not, 4.1 will be affected as well. But I'm not sure what the right query should be then. Personally I would like it better if we just removed the "affectedVersion" from it. [~ldimaggio], [~maxandersen], what do you think about this? What kind of query should we have in the release notes? The Known issues filter looks ok to me.	1	0	https://issues.jboss.org/browse/TOOLSDOC-521
3564	TOOLSDOC-521	Is there a way in JIRA to save a static set of query results?	1	0	https://issues.jboss.org/browse/TOOLSDOC-521
3565	TOOLSDOC-521	A few comments: affectedVersions is not reliably set - I would not rely on it. < and > in jira queries does not work reliably - thus I would not use that either. One should use fixversion and even if it is fixed in two versions at the same time that just mean we back or forward ported a fix. Is the resolved query only about CR2/GA and not about all the fixes done from Alpha to GA for JBDS 8 ? Then the resolved one is fine.	1	0	https://issues.jboss.org/browse/TOOLSDOC-521

#	Requirement	Questions	Answered	Successfully Answered	URL
3566	TOOLSDOC-516	Hi [~bobjboss], [~anmahaja], [~nids1409], [~dmichael], [~nshendye], [~aakanksha_writer], [~mandarjoshi], [~Nidhis], can one of you please pick up the following article for a peer review? Article name: Establish an IDE Connection to OpenShift Online Link: https://docs.google.com/a/redhat.com/document/d/1a2tXP8JKWbhoesw7rKzkBROFWWyadjWRHz_geBbiRbA/edit#) Thanks, Supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-516
3567	TOOLSDOC-515	[~nids1409] and [~nidis], can you two please review the article at: https://access.redhat.com/articles/1160993 As part of the peer review please also do the following: * Perform the steps listed in the article so as to ratify them * Let me know if the article would be more useful for the end user with/without the screen captures Thanks!	1	1	https://issues.jboss.org/browse/TOOLSDOC-515
3568	TOOLSDOC-498	{quote} True we don't audit our GA for security vulnerabilities but saying we don't audit for Beta isn't actually inaccurate. {quote} When a user reads this information, I think it's pretty easy for them to imply that JBDS GA is audited. I sure would think that if I didn't know better. Sure, mathematically speaking it is a true statement, but so would be saying "This beta release does not allow you to make unicorns" ;) But overall I don't think it's a big deal and agree that for simplicity it's probably better to leave the same message here as everywhere else. So I'm fine with leaving it as it is.	0	0	https://issues.jboss.org/browse/TOOLSDOC-498
3569	TOOLSDOC-494	See TOOLSDOC-576: "Btw I'd rather add info that livereload will not work with java 7 if web engine is javaFx. Why - see FAQ"	0	0	https://issues.jboss.org/browse/TOOLSDOC-494
3570	TOOLSDOC-493	[~maxandersen], can you provide an update on this? Is GTK3 still an issue for JBDS 8 and does it need this work around?	1	1	https://issues.jboss.org/browse/TOOLSDOC-493
3571	TOOLSDOC-485	New text layer was hidden when I exported the PNG. My fault. Have fixed this. Version 1.4 is listed on the screencap.	1	1	https://issues.jboss.org/browse/TOOLSDOC-485
3572	TOOLSDOC-482	Should I mention in DId You Know Section or we should include it only when we productize the How to Article for Enterprise? * You can download EAP server through JBDS/JBoss Tools, Click Windows > Preferences > Server > Runtine Environments. Click Add and search for EAP.	0	0	https://issues.jboss.org/browse/TOOLSDOC-482
3573	TOOLSDOC-477	[~ldimaggio], [~mmalina]: is this still an issue for the latest Luna release?	1	1	https://issues.jboss.org/browse/TOOLSDOC-477
3574	TOOLSDOC-477	[~ldimaggio], can you check? I never got this because I never run Eclipse directly from command line. And I tried now and didn't see it on OS X. But I see this regularly when running e.g. WildFly 8.1 with java 8. I think there was another JIRA for this and the conclusion was that it will have to be left as is for now. So yeah, we could document this. But first Len will have to confirm that it still happens for him.	1	1	https://issues.jboss.org/browse/TOOLSDOC-477
3575	TOOLSDOC-477	[~ldimaggio] can you comment on this? Is this issue even relevant anymore?	1	1	https://issues.jboss.org/browse/TOOLSDOC-477
3576	TOOLSDOC-460	"All the project wizards" -> it is not all (i.e. openshift won't create a maven project if you create php app type and once mobile tools get there it most likely won't have maven out of box either). maybe "Many of the project wizards" ? or most or some. Besides that if you have jboss central installed you automatically get these two. What is more interesting is that each project can state it works best with a specific plugin (for example GWT example linked to the GWT plugins) and that gets to be an option to download when running the example. That workflow is worth documenting.	0	0	https://issues.jboss.org/browse/TOOLSDOC-460
3577	TOOLSDOC-446	[~maxandersen], is there a reason you picked out the 6.0 OpenShift reference guide and not the 7.1 User Guide - OpenShift Tools chapter [1]? [1] https://access.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Developer_Studio/7.1/html-single/User_Guide/index.html#chap-OpenShift_Tools	1	1	https://issues.jboss.org/browse/TOOLSDOC-446
3578	TOOLSDOC-446	[~maxandersen], [~adietish], [~burrsutter]: Is 8787 specifically a debug port of EAP/AS or OpenShift or both?	1	1	https://issues.jboss.org/browse/TOOLSDOC-446
3579	TOOLSDOC-446	[~maxandersen], sometimes it's good to be pedantic ;) So how would a user identify which port is the one being used for debugging? Is there a way to tell from within the IDE? To me the port forward window just show a list of port numbers and I have no idea how to find out which ones do what. I know that there is info in the server editor stating ports for port offset, web and management.	1	1	https://issues.jboss.org/browse/TOOLSDOC-446
3580	TOOLSDOC-446	Thanks for the clarification [~maxandersen]. I need some further info, please. When making a debug configuration (in the video), in the sources tab you added the JBoss Maven Source Container. Why is this necessary? I've looked at the explanation in [1] but I need it in plainer language please. [1] https://community.jboss.org/en/tools/blog/2012/05/21/jboss-tools-and-developer-studio-beta-3	1	1	https://issues.jboss.org/browse/TOOLSDOC-446
3581	TOOLSDOC-446	[~maxandersen], need one tiny bit of further clarification. "That is where jboss maven source container helps since it will use the various indexes available from maven repositories to try and locate the proper source jar." What is the container? * Is it a list of references to online JBoss Maven repositories? e.g., Red Hat Maven Online Repository (tech preview) * Is it using information from the users' settings.xml file? * Why is it called 'JBoss'? Is this because the JBoss Maven Source Container is just looking up EAP and JBoss AS Maven repositories?	0	0	https://issues.jboss.org/browse/TOOLSDOC-446
3582	TOOLSDOC-445	We can explain the project structure so that users can be more informed when making decisions about which files/dirs to place under version control. Since JBDS does not have it's own version controls actions we can't document exactly what to put under version control. As part of JBQA-8565, I stated: {quote} Gorkem Ercan, about "need to explain the different directories and what they do". This would be best included in the context of a description of standard Cordova projects and how JBDS produced Cordova projects vary. Standard/basis of Cordova projects are given in '8.2.1. About Apache Cordova'. If you think there should be more details please provide me with some pointers of what's missing. JBDS Cordova project distinctions can be given in '8.2.2. About Hybrid Mobile Tools', where the project wizard is detailed. Again, please provide any specifics of how JBDS Cordova projects differ. From the info above one distinction seems to be the absence of a platforms directory. {quote} Let's enact this for JBDS 8.0.0 (JBT 4.2.1). Information I have is as follows: 1. Cordova (CLI) project structure [1] = config.xml, hooks/, merges/, www/ with index.html, platforms/, plugins/ 2. Hybrid Mobile Tools Cordova project structure = .cordova/ with config.json, merges/, plugins/, .project, .settings/, www/ with config.xml and index.html *What info do I have wrong or am I missing?* [1] https://github.com/apache/cordova-cli#project-directory-structure	1	1	https://issues.jboss.org/browse/TOOLSDOC-445
3583	TOOLSDOC-440	You seem to have been looking at two sets of docs there. The first link was for 7.1 version but second link is for 7.0. This has probably happened because you are viewing in html format and each page has to be Ctrl+F5'ed to clear the cache. Alternatively you could change to the html_single version and then you just need to do Ctrl+F5 once. As to defining BYOE, * 1.1 About JBDS states 'For developers already running Eclipse, JBoss Developer Studio can also be installed through Eclipse Marketplace. For distinction, this latter JBoss Developer Studio installation is called JBoss Developer Studio BYOE (Bring Your Own Eclipse).' * 3.1 Install Options states 'Alternatively, if you already have Eclipse on your system you may prefer to install JBoss Developer Studio BYOE. This option installs JBoss Developer Studio in Eclipse.' I don't want to have to explain at every instance of mentioning JBDS BYOE what it is. But I could expand the info in section 1.1 with another sentence to add more BYOE detail. What info specifically do you think is missing or it should say? [Any fix would only be released as part of 7.1.1 or 8.0 - the 7.1.0 ship has sailed.]	1	1	https://issues.jboss.org/browse/TOOLSDOC-440
3584	TOOLSDOC-426	[~nickboldt], can you answer the description Qs please.	1	1	https://issues.jboss.org/browse/TOOLSDOC-426

#	Requirement	Questions	Answered	Successfully Answered	URL
3585	TOOLSDOC-426	Latest version is here: http://download.jboss.org/jbosstools/targetplatforms/jbtcentraltarget/4.31.0.CR1-SNAPSHOT/ But we'll have a new version soon as there's more stuff going into JBT 4.1.1 Central soon. See JBDS-2710 for the overall issue listing all the changes / updates / additions. I guess really we should pull the latest zip off there and push it up onto devstudio.jboss.com for archival and ease-of-use purposes. WDYT?	1	1	https://issues.jboss.org/browse/TOOLSDOC-426
3586	TOOLSDOC-426	For JBDS, it needs to have a JBDS URL, like the update zip. So yes, a zip should be pulled from jbt site and pushed elsewhere. Info I have about the update site is: "You can download the update .zip from https://access.redhat.com on the Customer Portal website or https://devstudio.jboss.com/updates/7.0/ on the Red Hat JBoss Developer update website." Can we duplicate this for central zip? so have it on access.redhat and devstudio.blah?	1	1	https://issues.jboss.org/browse/TOOLSDOC-426
3587	TOOLSDOC-425	Hi [~xcoulon], I am adding info about AsciiDoc ability to LiveReload chapter in JBDS User Guide. 2 questions: 1) I haven't used LiveReload for a while so am trying to remind myself how it works. Couldn't get it to work with xhtml files but is that because it only works with html files (and now asciidoc)? 2) I think the AsciiDoctor.js is just rendering the asciidoc file in the browser. And LiveReload (and/or LiveReload browser extensions) does all of the LiveReload work. Is that right? If so and LiveReload Server in IDE is doing all of the work, why is the AsciiDoctor.js necessary?	0	0	https://issues.jboss.org/browse/TOOLSDOC-425
3588	TOOLSDOC-411	{quote} Is there a silent install for JBoss Developer Studio (JBDS)? User needs to roll out the JBDS installation across multiple developer workstations. {quote} Yes - from the JBT/JBDS update site zip (or the JBDS installer jar, which can be treated as an update site zip too): https://www.jboss.org/tools/download/installation/update_4_1#fromscript See also JBDS-2768, which involved me writing scripts for headlessly performing the izpack install, but hasn't yet been +1'd by [~maxandersen] as being "ready for consumption" yet. Code lives here: https://github.com/jbdevstudio/jbdevstudio-website/tree/master/content/updates/scripted-installation But is not live on devstudio.jboss.com yet.	1	1	https://issues.jboss.org/browse/TOOLSDOC-411
3589	TOOLSDOC-409	Additional features as per Gorkem's email: {quote} * Export wizard for native project (Xcode and/or Android projects) * Export wizard for "ready to sign" applications (Xcode and Android) What will be released to Beta1 today * Cordova Plugin Discovery wizard * Extensions to Project explorer regarding Cordova plugins * JS library container for Cordova {quote}	0	0	https://issues.jboss.org/browse/TOOLSDOC-409
3590	TOOLSDOC-402	Generated a skeleton doc in CCMS but still need to populate new features, fixed and known issues. Also, should it refer to 7.1.0.Beta rather than simply 7.1.0? Something to think on.	1	1	https://issues.jboss.org/browse/TOOLSDOC-402
3591	TOOLSDOC-396	[~nickboldt], need your input. In the Release Notes for 7.0.1 I've written: {quote} Chapter 2. About this Release JBoss Developer Studio 7.0.1 is an update to JBoss Developer Studio 7.0.0. The update has two purposes: * It resolves issues identified in JBoss Developer Studio 7.0.0. * It contains an Eclipse Kepler update, SR1, for JBoss Developer Studio 7.0.0 packaged with Eclipse. Users of JBoss Developer Studio BYOE must apply the Eclipse Kepler SR1 update independently. {quote} Is the bit about JBDS BYOE needing to apply Kepler SR1 independently correct? Or is is part of the JBDS update they will get too? If the latter, what happens if they have already applied SR1?	1	1	https://issues.jboss.org/browse/TOOLSDOC-396
3592	TOOLSDOC-390	Made PRs based on 4.1.1.x branches thinking they could be applied to 4.1.1.x and 4.1.x. Alexy found conflicts with applying to 4.1.x (no surprising since I built off 4.1.1.x! lesson learnt). Lead to this discussion: {quote} (07:41:01 PM) maxandersen: akazakov: whats the question ? :) (07:41:57 PM) akazakov: maxandersen, https://github.com/jbosstools/jbosstools-javaee/pull/177 (07:41:57 PM) jbossbot: git pull req [jbosstools-javaee] (open) michellemurray TOOLSDOC-390: Doc versions updated for 4.1 and 7.1 https://github.com/jbosstools/jbosstools-javaee/pull/177 (07:41:58 PM) jbossbot: jira [TOOLSDOC-390] Update book versions for 7.1.0/4.1.1 [Open (Unresolved) Task, Major, Michelle Murray] https://issues.jboss.org/browse/TOOLSDOC-390 (07:42:06 PM) akazakov: is it ok to push to 4.1.1.x? (07:42:25 PM) maxandersen: akazakov: doc changes into 4.1.1.x ? hmm.... murraym remind me if did we do that in past ? I *thought* it was just built from master.....in any case, probably simpler if go into 4.1.x instead imo. (07:44:03 PM) murraym: 4.1.1.x will have wrong version numbers in docs (07:44:50 PM) murraym: maxandersen: I don't think we did anything before, hence why the doc version numbers are so outdated (07:45:16 PM) murraym: master will be for 4.2 and 8.0 - right? (07:45:51 PM) murraym: i was making PRs to update version numbers in docs in 4.1.1.x and 4.1.x branches (07:48:18 PM) jbossbot: git [jbosstools-base] push master 4e09f7e.. Fred Bricon JBIDE-16176 : use proper stacks archetype id during lookup... (07:48:18 PM) jbossbot: jira [JBIDE-16176] Wrong archetype id used to lookup stacks archetypes [Resolved (Done) Bug, Critical, Rob Stryker] https://issues.jboss.org/browse/JBIDE-16176 (07:48:18 PM) jbossbot: git [jbosstools-base] push master URL: http://github.com/jbosstools/jbosstools-base/commit/4e09f7e85 (07:50:13 PM) murraym: maxandersen: if you want me to make changes in 4.1.x only then say so, I think I can cherry-pick the commits i have onto 4.1.x branch and make new PRs for 4.1.x branches (07:50:29 PM) murraym: would like to get this done tonight (07:55:33 PM) maxandersen: murraym: just do it for 4.1.x. 4.1.1 train has left the building for noncritical fixes ;) (07:56:06 PM) murraym: maxandersen: ok, will do - will go close my existing PRs for 4.1.1.x so please disregard {quote} So original PRs requesting merging with 4.1.1.x and 4.1.x closed with being applied by me. Made the following new PRs just for 4.1.x branch: https://github.com/jbosstools/jbosstools-birt/pull/16 https://github.com/jbosstools/jbosstools-bpel/pull/20 https://github.com/jbosstools/jbosstools-central/pull/192 https://github.com/jbosstools/jbosstools-esb/pull/46 https://github.com/jbosstools/jbosstools-hibernate/pull/34 https://github.com/jbosstools/jbosstools-javaee/pull/178 https://github.com/jbosstools/jbosstools-portlet/pull/23 https://github.com/jbosstools/jbosstools-server/pull/182 https://github.com/jbosstools/jbosstools-webservices/pull/102 These set the JBDS doc version to 7.1 and JBT doc version to 4.1.x. Going with 4.1.x rather than 4.1.1 so won't need to be updated if there is a 4.1.blah. This brings the docs numbering in line with JBDS - so new set of docs for each minor version only (x.y) and not maintenance (x.y.z). Need to update Jenkins job to reflect this.	1	0	https://issues.jboss.org/browse/TOOLSDOC-390
3593	TOOLSDOC-388	Source info: (12:34:11 PM) mmurray: so JBDS 7.1, the universal installer will be shipped with Eclipse Kepler SR1 (12:34:43 PM) mmurray: could a user use JBDS BYOE 7.1 with Eclipse Kepler without the SR1 update? (12:35:06 PM) mmurray: or do we only support JBDS BYOE 7.1 with SR1 applied? (01:35:46 PM) nboldt: pretty sure the 4 options will work: Kepler R + JBDS 7.0, Kepler SR1 + JBDS 7.1, Kepler R + JBDS 7.1, Kepler SR1 + JBDS 7.0 (01:35:49 PM) nboldt: mmurray: ^ (01:36:27 PM) nboldt: traditionally we support the release on all three equivalent Eclipse platforms: Kepler R, SR1, SR2 (Eclipse 4.3.0, 4.3.1, 4.3.2) (01:37:25 PM) mmurray: nboldt: great, that's what I wanted to know (01:37:28 PM) nboldt: and since 7.0 and 7.1 should be upgrade compatible, the 2x2 matrix (3x3 matrix in March 2014) should all be supported: {4.3.0, 4.3.1, 4.3.2} x {7.0.0, 7.1.0, 7.x} (01:40:03 PM) nboldt: usually with JBDS we "strongly recommend" users grab the latest the latest SR2 because of the fixes and updates	1	1	https://issues.jboss.org/browse/TOOLSDOC-388

#	Requirement	Questions	Answered	Successfully Answered	URL
3594	TOOLSDOC-388	Also from email conv with nick: 7.1 will be released to the same Marketplace entry (and backing URL) so I'll only need to update the entry to say "7.1.0.GA" instead of "7.0.0.GA". As it should work with Kepler as well as SR1 and the forthcoming SR2 (Feb 2014), it will still have "(Kepler)" on it, but we might update the copy [1] to state that SR1 (and later SR2) is recommended due to performance/bug fixes, like we did for 6.0.1 [2]. [1] https://marketplace.eclipse.org/content/red-hat-jboss-developer-studio-kepler [2] https://marketplace.eclipse.org/content/red-hat-jboss-developer-studio-juno On 09/10/2013 09:21 PM, Michelle Murray wrote: > Hi Nick, > > Trying to plan all the changes that I will need to make for the docs for > 7.1 and I have a question. > > Currently on marketplace (screen cap at end), the two JBDS versions are > clearly distinguishable by '(Kelper)' or '(Juno)' after the product > name. In smaller text it also specifies the exact JBDS version (7.0.0.GA > or 6.0). > > How will the 7.1 release on Kepler be distinguished for users viewing > the marketplace to avoid confusion with the existing JBDS Kepler release? > * Will it be titled 'Red Hat JBoss Developer Studio (Kepler SR1)'? > * Will the only difference be in the smaller text where the specific > JBDS version is stated? > * Will this not matter as the 7.0.0 version will be replaced on the > marketplace by the 7.1.0 version? > > Thanks, > Michelle >	1	1	https://issues.jboss.org/browse/TOOLSDOC-388
3595	TOOLSDOC-388	I've already made changes in the 4.1.x/7.1.x branch so that references within JBDS to its update site (located at https://devstudio.jboss.com/updates/7.0/) refer to it as the 7.x site, not the 7.0 or 7.1 site. That way I won't have to update it again from 7.1 to 7.[whatever is decided*] later, when we do a maintenance release to pick up Kepler SR2. You might want to consider updating docs similarly to "7.x" so that you too won't need to change explicit refs in docs to "7.1" to "7.1.1" or "7.2" in 6 months when Kepler SR2 is out.	1	1	https://issues.jboss.org/browse/TOOLSDOC-388
3596	TOOLSDOC-367	Question for us all (especially Burr and Jiri P.) - which set of resolved bugs do we want to highlight in the 7.0.GA release notes? Right now - we're only showing the (82) bugs with the affected version equal to the final builds: (project in (JBIDE) AND affectedVersion in ("4.0.0.Final", "4.0.1.Final") AND fixVersion >= "4.1.0.Alpha1" OR project = JBDS AND fixversion in ("6.0.0.GA", "6.0.1.GA") AND fixVersion >= "7.0.0.Alpha1") AND type in (Bug) AND resolution in (Done) Don't we want to show all the 700+ bugs fixed since the 6.0 release? (project = JBIDE AND fixVersion >= "4.1.0.Alpha1" OR project = JBDS AND fixVersion >= "7.0.0.Alpha1") AND type in (Bug) AND resolution in (Done)	0	0	https://issues.jboss.org/browse/TOOLSDOC-367
3597	TOOLSDOC-359	Thanks [~nickboldt]. I think it is all pretty much as you said. Regards Jenkins jobs, as I see it: * Not sure how often 4.0 will need to be rebuilt. The docs are so flawed that I am just leaving them alone. But having a jenkins job won't do any harm for the eventuality that something does need fixing or someone else is wanting to take on the job of fixing (as in TOOLSDOC-316). * Will need a job for developing docs (from master). As soon as 4.1 is released then this branch will be developing towards 4.2. * Once 4.1 is released and branch 4.1 created, will need a job for this so we can fix bugs in docs and applying changes that I didn't get time for prior to GA. Creating the 4.1.x branch should coincide with GA rather than be done in advance, since as far as I can see docs doesn't need stabilized docs for Betas and leaving until GA negates the need to apply changes to master and 4.1 branch continually.	0	0	https://issues.jboss.org/browse/TOOLSDOC-359
3598	TOOLSDOC-345	[~rob.stryker], more questions: * In the server editor, what's the difference between 'Publishing' and 'Application reload behavior'? Does one outrank the other or are they doing different things? I guess the answer is the latter but I don't understand why. * Also, why is this .. {quote}The "Process Terminated" polling mechanism should not be used for servers operating with a Remote System Explorer behavior.{quote} * {quote}JBoss Server Tools provides the ability to customize default filesets added to new servers based on the application server type.{quote} I get part of this, for example it could be a config file. But what happens to the fileset when it is add to the server? Does it get copied to a folder of the local application server? What happens to the files when the application server is remote? * {quote}JBoss Server Tools provides the ability to customize classpath filesets based on the application server type.{quote} What are classpath files? What are they used for? What happens after you specify the classpath filesets? Are the class files physically added somewhere? Or are we just stating class files that are available to the server?	1	1	https://issues.jboss.org/browse/TOOLSDOC-345
3599	TOOLSDOC-345	1) Publishing is moving an app into a location for the app server to read. "Application reload behavior" is on what patterns we FORCE the server to redeploy the app. This is because our incremental publisher only copies changed files. If you change a .html, it already will be read live from your deployment. But if you change a .class, or if change a .jar file inside your app, the server will not recognize it needs to redeploy your application. So, we let you set a pattern where, if any file that matches it is changed, we force the server to 'redeploy' the app. 2) Filesets are a completely local construct and are simply meant to give you quick access to local files. It is a very old part of astools and has never been updated for remote servers. The files are not copied anywhere. When you create a fileset, those files are just shown in the view, and, if you double-click on one, it will open the file in the editor so you can modify it or look through it. Since most people do not keep their actual app-server inside eclipse, this is juts a handy mechanism to quickly browse your server's config files. 3) When you create a project, like a dynamic web project, there is a classpath container added to your project. The files aren't copied anywhree, but rather, eclipse adds these jars to your project's classpath so that you can properly code your application. Different users want different things on their classpath. We automatically add about 10 default filesets for things like ejb-api and other javaee api jars, but we always get new requests by a user. Can you add the mail jar? can you add this? Can you add that? After a while, which jars to add became a big debate. We only wanted to automatically add whatever everyone would use, but some users wanted pretty obscure jars added to their project's classpath. We couldn't in good conscience do this, so we added this mechanism for users to customize paths that should be added. Hope this helps!	1	1	https://issues.jboss.org/browse/TOOLSDOC-345

#	Requirement	Questions	Answered	Successfully Answered	URL
3600	TOOLSDOC-345	During a publish, changed files in the workspace will be moved to your deployment. So for example, {workspace}/project1/webcontent/myhtml.html will be copied over to {jbossasroot}/standalone/deployments/project1.war/myhtml.html. Any changes will now be visible. The server itself, the running JBoss server, will not 'redeploy' the module. This is because html files are served directly without changes by the app server. HTML files aren't run or interpreted or executed on teh server at all. They're simply served as regular files. So an updated html file will be completely ignored by the app server, and the server will serve out the new updated one the same as it served out the old one. If you change a .class file somewhere though, the server ALSO will not 'redeploy' the application. This is because the server is not scanning EVERY file in your module for updates. Class files *are* used by the server's logic, but, the server team has decided that checking the timestamps on all files in a deployment is a big waste of time for the server. If a user wants a module to be re-deployed, they should use one of the re-deploy mechanisms. The server will not note that the .class file has been changed and respond to that with a redeploy of the application. SO... when you modify a .html file and republish the app, the server doesn't redeploy, but the new .html file is able to be served out to clients. If you modify a .class file and republish, though, the server ALSO doesn't redeploy, but the new .class file will NOT take effect in teh server. The server will only SEE the new .class file if the module is properly undeployed and redeployed. Users will complain about this behavior. I changed an html file and can see the results immediately, but I changed a System.out.println line in my .java file and published the changes and the server is not printing anything. Yes, this is intended. The server did not scan for timestamp changes on the .class file and did not redeploy the application. On disk, the new .class file is inside your deployments folder, but, the server only loaded those classes when it initially deployed the application. The updated .class file you published has so far been completely ignored by the running server. To fix this, when modifying .java (and thus .class) files, you need to use FULL PUBLISH or RESTART on your module. This ensures the server actually deletes what's in running memory, and reloads your module from disk. Users would then complain, I don't WANT to have to undeploy and redeploy my app every time I make a change to a .java / .class file. Why doesn't the running server notice the change? OR at least why doesn't JBT notice the change and tell the server to redeploy FOR me. I'm waaaay too lazy to click full publish, or undeploy / redeploy. In response to this, we added application reload behavior. So now, when a user changes a .java file, and it gets changed to a .class file, and copied into a standalone/deployments folder, JBT will then ALSO check if ANY changed files match the 'application reload behavior' pattern. If one does, we will ALSO tell the server to redeploy the application, not just copy the changed files over. As I said, in the past, if you changed a .java file, you get a new .class file, the .class file gets copied to standalone/deployments/MyProject.war/some/file.class, and the running server completely ignores it. NOW, if your application reload behavior includes .class, JBT will ALSO tell your server "hey, user wants this application re-deployed from disk, so, undeploy and redeploy it for him". Of course users don't wnat this behavior on every change. Since html files are served directly, they never need a redeploy. So if we simply forced a redeploy after every single change, users who just change a line in a html file will have to wait while the server undeploys and redeploys the application. This could take any number of seconds dependign on the size of the deployment. So it makes sense as a customizable setting for the users.	1	1	https://issues.jboss.org/browse/TOOLSDOC-345
3601	TOOLSDOC-345	Thanks [~rob.stryker], great explanation. I feel like I'm up to speed now :) I think some of the things I was seeing in JBDS were adding to my confusion but *I'm not suggesting anything be changed*. So when I changed a .html file, there was action in the servers tab showing republishing but then there was also action in the console tab that appeared to state the app had been 'redeployed'. And then when I changed a .java file, I got the same two messages - servers tab showed 'republish' and the console tab stated 'redeployed'. Both changes looked like they had the same effect. I have one more question about application reload behavior. I think this works as follows, am I right? * Check box not selected, with the default '\.jar$' regex, results in automatically checking for .jar changes. * The user selects the check box if they want to further customize the reload behavior. * And they could delete '\.jar$' if they didn't want changes to .jar to result in an automatic redeloy.	1	1	https://issues.jboss.org/browse/TOOLSDOC-345
3602	TOOLSDOC-343	[~akazakov], sorry another Q relating to this. I get that the @FacesConverter validation is resolved. But are there any validation issues with JSF in general that report false warnings/errors? Or is this just a CDI issue now?	1	1	https://issues.jboss.org/browse/TOOLSDOC-343
3603	TOOLSDOC-341	Section: 1.1. About JBoss Developer Studio Text: Additionally, each of the tools comprising JBoss Developer Studio has great features to assist your application development: Comment: I think it would be better to refer to JBDS as a whole, as opposed a set of tools. Maybe this sentence can be reworded as: Additionally, JBoss Developer Studio has great features to assist your application development: ------ Section: 1.2. Use Cases of JBoss Developer Studio Text: JBoss Developer Studio assists JEE developers by integrating JBoss technology and APIs in a single development environment. Here are a few ways that JBoss Developer Studio makes the work of a developer easier. Comment: Very small point - end with a ":" not a "." ------ Section: 1.2. Use Cases of JBoss Developer Studio Text: BrowserSim allows you to view your web applications on a variety of simulated mobile devices so that you can ensure they will be correct formatted. Comment: This is a small point. Can we change this to read: "correctly formatted"? ------ Chapter 2. New Features You asked for input on this subject: JBoss Central NeedInfo: Cheat sheets for 'Start from a sample' project wizards are now automatically opened when projects are created and the cheat sheets can be subsequently accessed via the Show In option in project context menus. See the attached screenshot for an illustration on what this looks like. and: JBoss Server Tools NeedInfo: Additionally, you can choose to either stop or undeploy an application from the server via the administrative console. I'm not sure what your question is here. Users can deploy/undeploy applications to servers from either inside JBDS, or from the Admin Console provided by EAP or SOA-P.	1	1	https://issues.jboss.org/browse/TOOLSDOC-341
3604	TOOLSDOC-338	[~fbricon] and [~akazakov] can you take a look ?	1	1	https://issues.jboss.org/browse/TOOLSDOC-338
3605	TOOLSDOC-337	[~koen.aers], just need a bit of clarification. https://issues.jboss.org/browse/JBIDE-11217 release notes text states that 'A new feature has been added that allows for a user to run the Forge Verbose command from within the Forge console, and any exceptions are recorded in the Eclipse error log.' From your comment above, I guess that bit about recording exceptions in the error log is wrong. Or is it the case that exceptions are recorded in the error log but not shown in the error log tab even it the tab is opened manually?	1	1	https://issues.jboss.org/browse/TOOLSDOC-337
3606	TOOLSDOC-336	{quote} If this feature is working, can someone please give me a workflow so that I can have a play? I have no idea how to get started with this. {quote} A cheatsheet is the standard Eclipse feature described at http://help.eclipse.org/juno/topic/org.eclipse.platform.doc.isv/guide/ua_cheatsheet.htm We have added some functionalities described in JBIDE-14331. {quote} What is the current situation with JBDS/JBT cheat sheets? {quote} There are cheatsheets for the helloworld, login, numberguess, kitchensink quickstarts and Project Examples Tutorial. There should be created cheatsheets for archetypes, other quickstarts ...	1	1	https://issues.jboss.org/browse/TOOLSDOC-336

#	Requirement	Questions	Answered	Successfully Answered	URL
3607	TOOLSDOC-336	[~snjeza] Testing with JBDS Beta1-v20130526-0839-B241. I followed the 'Start from a sample' wizards for the numberguess, kitchensink, helloworld projects. The cheat sheet did not automatically open. In the project explorer tab, I right-clicked on the project and clicked Show In > Cheat Sheets - this opened the Cheat Sheets tab but only showed a default message and not the cheat sheet. Is JBIDE-14332 and JBIDE-14334 broken or have I done something wrong?	1	1	https://issues.jboss.org/browse/TOOLSDOC-336
3608	TOOLSDOC-336	{quote} I followed the 'Start from a sample' wizards for the numberguess, kitchensink, helloworld projects. The cheat sheet did not automatically open. {quote} This is correct behaviour because no one of the quickstarts contains either cheatsheet.xml or .cheatsheet.xml. {quote} In the project explorer tab, I right-clicked on the project and clicked Show In > Cheat Sheets - this opened the Cheat Sheets tab but only showed a default message and not the cheat sheet. Is JBIDE-14332 and JBIDE-14334 broken or have I done something wrong? {quote} If you right-click the project and call the Show In > Cheat Sheets action, the Eclipse platform automatically opens the Cheat Sheets view and restores the last used cheatsheet. JBoss Tools doesn't do anything. I think we can't change this behaviour unless we create our own action. If you right-click on a cheatsheet file, JBoos Tools will show the file in the Cheat Sheet view.	1	1	https://issues.jboss.org/browse/TOOLSDOC-336
3609	TOOLSDOC-336	{quote}This is correct behaviour because no one of the quickstarts contains either cheatsheet.xml or .cheatsheet.xml.{quote} Oh? but in a earlier reply you wrote: {quote}There are cheatsheets for the helloworld, login, numberguess, kitchensink quickstarts and Project Examples Tutorial.{quote} Sorry, I'm confused :(To me these replies seem to contradict each other. Please can you clear things up for me. Let's start again. * Do the JBDS/JBT quickstarts have cheat sheets? * If yes, where do you find or how do you open the cheat sheets?	1	1	https://issues.jboss.org/browse/TOOLSDOC-336
3610	TOOLSDOC-336	{quote} There are cheatsheets for the helloworld, login, numberguess, kitchensink quickstarts and Project Examples Tutorial. {quote} Sorry for the confusion. There are cheatsheets for the helloworld, login, numberguess and kitchensink quickstarts. They are placed here: https://anonsvn.jboss.org/repos/jbosstools/workspace/snjeza/jbossas7-examples/helloworld.zip https://anonsvn.jboss.org/repos/jbosstools/workspace/snjeza/jbossas7-examples/kitchensink.zip https://anonsvn.jboss.org/repos/jbosstools/workspace/snjeza/jbossas7-examples/login.zip https://anonsvn.jboss.org/repos/jbosstools/workspace/snjeza/jbossas7-examples/numberguess.zip These quickstarts have been moved to JDF and I am not sure why the cheatsheets aren't included in it. Suppose they will be added after fixing JBIDE-14331. {quote} Do the JBDS/JBT quickstarts have cheat sheets? {quote} Not currently. {quote} If yes, where do you find or how do you open the cheat sheets? {quote} You can download the mentioned archives and import them. The helloworld, kitchensink, and numberguess cheatsheet should be opened automatically. The login cheatsheet won't be opened automatically because it is on a separate location: cheatsheet/login.xml. You can open it by right-clicking on login.xml and selecting Show in>Cheat Sheets. You have to enable opening a cheatsheet automatically when importing a project using the Window>Preferences>JBoss Tools>Project Examples preferences page.	1	1	https://issues.jboss.org/browse/TOOLSDOC-336
3611	TOOLSDOC-335	[~maxandersen], [~nickboldt], [~dgolovin] and [~fbricon]: can one or more of you please comment on this JIRA.	1	1	https://issues.jboss.org/browse/TOOLSDOC-335
3612	TOOLSDOC-333	Thanks [~xcoulon], great explanation. I had seen the video from JBoss Tools blog but never with the sound. I've got another question about the proxy server configuration. This aspect looks great - far better than having to install the LiveReload browser extension. I like that we can simplify things for users by keeping it all in the IDE with minimal installing/configuration. About proxy and web browsers ... so a user selects a file and clicks 'Show With' > 'Web Browser via LiveReload Proxy' and then the IDE browser preference is used in deciding which type of browser window to open. * With the browser extension, users were limited to chrome and firefox because of using file URLs. Is this restriction lifted with the proxy server because it is using HTTP? So if a user's IDE browser preference was 'system default' and their system default was 'Safari' then LiveReload would now work? * I typically have my IDE browser preference set to 'internal web browser' but LiveReload doesn't seem to work with this browser. Why is that?	1	1	https://issues.jboss.org/browse/TOOLSDOC-333
3613	TOOLSDOC-333	Another question: * In Servers tab, user starts LiveReload server. User double-clicks LiveReload server and opens configuration info in new tab. User selects 'Enable Proxy Server' check box and perhaps changes some other settings. User presses Ctrl+S to save the changes. Does the user then need to stop and restart the LiveReload server for changes to take effect or does the server stop&start automatically? Or is stop&start not necessary?	1	1	https://issues.jboss.org/browse/TOOLSDOC-333
3614	TOOLSDOC-317	Ah. I can perhaps see part of the issue now. The N&N says 'There is now a keyboard shortcut that brings up a pop-up with a list of commands that can be applied in the current context. By default this command is 'Ctrl+4' ' But Ctrl+4 is also used to start Forge. So once Forge is started, will Ctrl+4 bring up the list of commands?	1	1	https://issues.jboss.org/browse/TOOLSDOC-317
3615	TOOLSDOC-311	Somehow getting the manual in pdf and commenting in jira feels like splitting things that should stick together. Shouldn't we at least attach the pdf to this jira or find a way to even comment inline?	1	1	https://issues.jboss.org/browse/TOOLSDOC-311
3616	TOOLSDOC-311	*Procedure 1.2. Generate and Upload SSH Keys to OpenShift* {quote} 6. In the Private Key field, type a name for the private key file. 7. In the Private Key Passphrase field, type a passphrase for use in accessing the private key. 8. In the Public Key field, type a name for the public key file. T ypically the file name of the public key is that of the private key with .pub appended. {quote} *Private Key* and *Public Key* were relabeled to *Private Key File Name* and *Public Key File Name* in JBIDE-13447	1	1	https://issues.jboss.org/browse/TOOLSDOC-311
3617	TOOLSDOC-310	im not following you guys... why do you want users and even us self having to read a page of multiple versions of possible links to link to it ? what am I missing ? and not grokking how sftp symbolic links are a problem - everyone else seem to be using them just fine.	1	1	https://issues.jboss.org/browse/TOOLSDOC-310
3618	TOOLSDOC-310	i.e. how do I get to the latest 3.3.x docs ? which *stable* url do I use ? how about 4.0.x, 4.1.x etc. ?	1	1	https://issues.jboss.org/browse/TOOLSDOC-310

#	Requirement	Questions	Answered	Successfully Answered	URL
3619	TOOLSDOC-310	New kick at the can for https://www.jboss.org/tools/docs/reference/ : !http://snag.gy/XioYn.jpg! Changes include: * remove all links to content which is actually located on docs.jboss.org * add links to Additional documentation and Archived documentation --- And for http://docs.jboss.org/tools/ : !http://snag.gy/rVRnb.jpg! Changes include: * stuff in OLD/ moved to archive/, then symlinked back * milestone docs in OLD/ moved to archive/milestones/ * links to 3.x docs (except for 3.3.0, the last one) removed from index page * 3.3.0.Final doc zip moved into archive/ * index page in resources/ removed as it was inaccurate (wrong dates, wrong file sizes, missing links); generated index is 100% accurate * experimental vpe_xulrunner zip renamed and moved into resources/; README updated to clarify that this is from JBT 2.0, from 2007, and is UNSUPPORTED. * link for movies relabelled as "Movies, Screencasts, and Video Walkthroughs" * fixed link to archived movies in http://docs.jboss.org/tools/archive/movies/ * fixed updated date to be ISO compliant * "nightly" watermark removed from index page * added new links to http://www.jboss.org/tools/docs/ and http://www.jboss.org/tools/ * latest/ is now a symlink to 4.0.1.Final instead of an index.html redirect to https://www.jboss.org/tools/docs/reference/ IMHO, this is much better. Does everyone agree? Here's the PR for the above changes: https://github.com/jbosstools/jbosstools-documentation/pull/51	1	1	https://issues.jboss.org/browse/TOOLSDOC-310
3620	TOOLSDOC-310	The links on https://www.jboss.org/tools/docs/reference/ still seem to use exact specific versions (i.e. http://docs.jboss.org/tools/4.0.1.Final/) instead of a "latest" link to http://docs.jboss.org/tools/4.0.x/ ? Meaning if/when we do an update of the docs we still need to update the magnolia page. Thus i'm missing where the symbolic links are setup. And wouldn't it be better to create a structure like: 3.3/latest 3.3/3.3.0.Final 4.0/4.0.1.Final 4.0/4.0.0.Final 4.0/latest -> 4.0.0.Final which removes all need for moving things to archive ? (could also just have it in one big root directory....my main concern is just not having to move things around over time; once url is used it stays. btw. I noticed you said symbolic links result in absolute paths - that should not happen if done via sftp like seam and richfaces guys do it: http://www.seamframework.org/Seam3/ModuleReleaseProcedure sftp seam@filemgmt.jboss.org sftp> cd docs_htdocs/seam/3/{module} sftp> rm latest sftp> ln 3.0.m.Q latest	1	1	https://issues.jboss.org/browse/TOOLSDOC-310
3621	TOOLSDOC-310	Other than the fact that the 4.0.0.Final folder no longer has an index page, and needs to contain one, I see your suggestions [~maxandersen] as being more effort than is justified. Here are my concerns: a) even if https://www.jboss.org/tools/docs/reference/ pointed at http://docs.jboss.org/tools/4.0.x/ instead of http://docs.jboss.org/tools/4.0.1.Final/, we'd still have to update the page when 4.1 is released. And as there's no plan for a 4.0.2 at this time, the only next time we'd have to update the page is when 4.1 is released. Your argument that there's more maintenance to be done here is only valid IFF a 4.0.2 or 3.3.3 happens for which the docs need to be rebuilt. Plus, I don't really mind using Magnolia as I have to use it once every 6 weeks to do milestone releases anyway. Changing a line of text in a field in Magnolia it is actually less work IMHO than creating a topic branch, pulling updates from origin, committing a code change, pushing to my fork, generating a PR, waiting for the PR to be approved, pushing the PR, and deleting my topic branch. b) moving things into subfolders gains exactly nothing except that we then have subfolders. What problem are you trying to solve by nesting? We can still have a symlink to "latest". Is there a usecase for 3.3/latest and 4.0/latest ? Because I can't see the need. Has anyone other than yourself requested this? (Is there a customer or community requirement?) c) First you asked for 4.0.x, then you asked for 4.0/latest. Which is your preferred approach? d) Regarding symlinks, you are welcome to do it your way (which is the way I used to until eng-ops changed stuff on the server) but when that doesn't work, read this and do it my way instead: https://engineering.redhat.com/rt/Ticket/Display.html?id=191655 :D	1	1	https://issues.jboss.org/browse/TOOLSDOC-310
3622	TOOLSDOC-310	>a) even if https://www.jboss.org/tools/docs/reference/ pointed at http://docs.jboss.org/tools/4.0.x/ >instead of http://docs.jboss.org/tools/4.0.1.Final/, we'd still have to update the page when 4.1 is >released. That is *one* time a year. >And as there's no plan for a 4.0.2 at this time, the only next time we'd have to update the page is when >4.1 is released. Your argument that there's more maintenance to be done here is only valid IFF a 4.0.2 or >3.3.3 happens for which the docs need to be rebuilt. I'm not following you. I'm talking about getting a structure that works for the future where there are plenty of releases and having one structure that works for all is to me better than a mix. >Plus, I don't really mind using Magnolia as I have to use it once every 6 weeks to do milestone releases >anyway. Magnolia is going away, it is not versionable, it is not sharable, permissions are a bitch to manage... >Changing a line of text in a field in Magnolia it is actually less work IMHO than creating a topic >branch, >pulling updates from origin, committing a code change, pushing to my fork, generating a PR, waiting >for >the PR to be approved, pushing the PR, and deleting my topic branch. No idea why this is relevant. My suggestion is about creating a symbolic link. No need for PR's every 6 weeks. (but I find it super valuable it is actually possible to do a PR so others can help and improve, with magnolia no such thing) >b) moving things into subfolders gains exactly nothing except that we then have subfolders. What problem >are you trying to solve by nesting? That we stop having to move things into "archived". i.e. less work, less changing urls. >We can still have a symlink to "latest". Is there a usecase for 3.3/latest and 4.0/latest ? Because I can't >see the need. Has anyone other than yourself requested this? (Is there a customer or community >requirement?) Yes, as mentioned before other docs link to /latest assuming it is 3.3.x docs. You can search this jira for the /latest url and find an issue like https://issues.jboss.org/browse/JBIDE-12243 Beyond that it is similar to what other projects do (seam, richfaces, infinispan and others), gives *less* work and easier to manage and use for users (stable urls) so I don't grok why you do not see the need/advantage. >c) First you asked for 4.0.x, then you asked for 4.0/latest. Which is your preferred approach? I don't care. >d) Regarding symlinks, you are welcome to do it your way (which is the way I used to until eng-ops changed >stuff on the server) but when that doesn't work, read this and do it my way instead: >https://engineering.redhat.com/rt/Ticket/Display.html?id=191655 No idea why you are having problems the existing links works fine and I just created one inside archives - it works perfectly. What I did find is that things I symbolic links in /tools directly gets a permission denied when fetching from browser, so yes something is wrong but all existing links works fine thus i'll ping helpdesk to see what is wrong.	1	1	https://issues.jboss.org/browse/TOOLSDOC-310
3623	TOOLSDOC-310	hmm - looking more into this....why is jbosstools-documentation now suddenly a partial mirror of what is on docs.jboss.org/tools ? Thats not how it was in the past (for good reasons since this repo != what is on the website) I must be missing something. In the past we published controlled parts of this to the site - what is published today ? i.e. where is the rsync scripts for publishing ?	1	1	https://issues.jboss.org/browse/TOOLSDOC-310
3624	TOOLSDOC-310	How about we create a docs.jboss.org folder in the root of jbosstools-documentation/ and move all the web content into there? Then I can put all the symlinks, HTML content, etc. in there and not pollute the root folder. [~mmurray] [~maxandersen] WDYT?	1	1	https://issues.jboss.org/browse/TOOLSDOC-310

#	Requirement	Questions	Answered	Successfully Answered	URL
3625	TOOLSDOC-310	[~maxandersen] If you don't like having a backup of the server's HTML content via https://github.com/jbosstools/jbosstools-download.jboss.org then I challenge you to propose a better alternative. Meanwhile, for docs.jboss.org, we ought to have a similar solution to back up the contents on the server so that a) the content can be developed offline and pushed to the server, and b) in case of accidental overwrite / server death, it can be recovered. If not a github repo, then what?	1	1	https://issues.jboss.org/browse/TOOLSDOC-310
3626	TOOLSDOC-307	Hi [~mhusnain], can you please create a PR for this? Thanks - supriya	1	1	https://issues.jboss.org/browse/TOOLSDOC-307
3627	TOOLSDOC-306	Hope this helps? Assigning back.	1	1	https://issues.jboss.org/browse/TOOLSDOC-306
3628	TOOLSDOC-303	Email from Yahor: Great job, you found and described nearly everything! Some comments regarding regarding to this doc: 1. (1.1.2 System requirements) I would add that BrowserSim work on Linux and Mac too, not Win32 only. 2. Could we use words Simulated Devices instead of Replicated Devices? For me it sounds more clear, because it is mobile simulator :-) 3. (1.2.1) There is one more way to open .html files: right click on a file and select Open With>BrowserSim 4. (1.2.3) There is also Weinre debugger in the latest builds: https://issues.jboss.org/browse/JBIDE-13338	1	1	https://issues.jboss.org/browse/TOOLSDOC-303
3629	TOOLSDOC-303	Email from Michelle to Yahor: Thanks for the feedback Yahor - very speedy. I am implementing it all now. Regarding Firebug Lite and Weinre, what's the difference? I think it is this, but am I right? Firebug Lite = local filesystem apps only, for standard web browsers and not mobile devices Weinre = deployed/remote apps, specifically for mobile devices Also, I think you guys have already implemented another function: * take a screenshot https://issues.jboss.org/browse/JBIDE-13434 Any info on this so I can add it to the documentation? Screenshots and Weinre will need news and noteworthy entries. I guess under alpha 2?	1	1	https://issues.jboss.org/browse/TOOLSDOC-303
3630	TOOLSDOC-303	Email from Yahor: >Firebug Lite = local filesystem apps only, for standard web browsers and not mobile devices Any web or local applications, but it can work only on the same device where the app is run. This means that Firebug does not support remote debugging, i.e. if an applications is run on Android device you cannot debug it from your desktop with Firebug. >Weinre = deployed/remote apps, specifically for mobile devices Weinre supports remote debugging. You can debug your applications which are run on Android or iPhone devices from your desktop browser. But from the BrowserSim's users perspective the major difference between these two debuggers is that Firebug Lite is slightly faster and its UI is similar to UI of Firebug in Firefox, and Weinre is slightly solver and the UI is similar to Chrome/Safari Developer Tools. >Also, I think you guys have already implemented another function: >* take a screenshot https://issues.jboss.org/browse/JBIDE-13434 >Any info on this so I can add it to the documentation? We just added a screenshot to this JIRA. It should be pretty self-explanatory. If you want to try it now, you may use nightly JBoss Tools build or browsersim.jar from https://hudson.jboss.org/hudson/view/JBossTools/view/JBossTools_Trunk/job/jbosstools-vpe-browsersim-standalone_master/ And then run it with the command: java -jar browsersim.jar > Screenshots and Weinre will need news and noteworthy entries. I guess under alpha 2? Yes, we are going to create these entries as usual, when a corresponding branch is created.	1	1	https://issues.jboss.org/browse/TOOLSDOC-303
3631	TOOLSDOC-303	Email Michelle to Yahor: I have downloaded the latest universal installer nightly build and had a play with the new BrowserSim functions. For Screenshot, there are three options: save, save as and copy to clipboard. I can see how the last two work but not the first one. What does 'save' do? No dialogue window appears and I cannot find the saved file. I assume it is also saving as .png but where? Is there a default save location? For Open Synchronized Window, a note in the 7.0 requirements document for this feature states that clicking on a browser will make all other browsers go to the top too. Is this something you are looking to implement or not for 7.0.0?	1	1	https://issues.jboss.org/browse/TOOLSDOC-303
3632	TOOLSDOC-303	Email from Yahor: > For Screenshot, there are three options: save, save as and copy to clipboard. I can see how the last two work but not the first one. What does 'save' do? No dialogue window appears and > I cannot find the saved file. I assume it is also saving as .png but where? Is there a default save location? Yes, it saves to a default location. You can see/change this location in preferences. > For Open Synchronized Window, a note in the 7.0 requirements document for this feature states that clicking on a browser will make all other browsers go to the top too. Is this >something you are looking to implement or not for 7.0.0? We completely forgot about this requirement :-) Thank you for reminding us!	0	0	https://issues.jboss.org/browse/TOOLSDOC-303
3633	TOOLSDOC-281	Hey Rastislav Wagner, Thanks for this info :) What's the "Convert to Maven Project" option for then? Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-281
3634	TOOLSDOC-281	Hey Rastislav, Thanks for that info :) I was wondering, when adding maven support via Facets you cannot set the packaging to 'pom', which seems weird. Is that a bug? Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-281
3635	TOOLSDOC-281	Alrighty, thanks for the extra info Fred. Rastislav and Fred, so I get this right, would it be true that the 'Facet' approach is currently they best way for JBT 3.3.x and JBDS 5.0.x streams; however in the future the "Convert to Maven Project" method will supersede the 'Facet' way?	1	1	https://issues.jboss.org/browse/TOOLSDOC-281
3636	TOOLSDOC-280	The linked issue has been resolved, verified and closed so does that make this JIRA outdated now? Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-280
3637	TOOLSDOC-280	Hey Rastlislav, Alrighty, makes sense ;) Can you please link me to the JIRA for the workspace crash issue then? If there isn't one can you please create one and link me to that so I can document it in the release notes. Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-280
3638	TOOLSDOC-274	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 811977|https://bugzilla.redhat.com/show_bug.cgi?id=811977] Rick Wagner <rwagner@redhat.com> made a comment on jira JBDS-2046 I find it a little confusing-- is that dividing line after the first 2 items meant to divide by 'Enterprise vs. Community' or 'Get the package vs. Ala-carte'? Both seem to be true. The top 2 categories (SOA and BRMS) are our 'Enterprise' offerings. The bottom categories are 'Community' or ambiguous. But the top 2 could also be taken as aggregations, while the bottom ones are singletons you can pick from.... I'd think this could be confusing. Could we have some description 'for Enterprise users', 'for Community users', or something similar? Thanks, Rick	1	0	https://issues.jboss.org/browse/TOOLSDOC-274
3639	TOOLSDOC-274	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 811977|https://bugzilla.redhat.com/show_bug.cgi?id=811977] Len DiMaggio <ldimaggi@redhat.com> made a comment on jira JBDS-2046 Aren't we also shipping jBPM5 with SOA-P 5.3?	1	0	https://issues.jboss.org/browse/TOOLSDOC-274
3640	TOOLSDOC-274	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 811977|https://bugzilla.redhat.com/show_bug.cgi?id=811977] Julian Coleman <jcoleman@redhat.com> made a comment on jira JBDS-2046 I agree - that last set of screenshots above is good. What do we need to do to make this happen?	1	0	https://issues.jboss.org/browse/TOOLSDOC-274

#	Requirement	Questions	Answered	Successfully Answered	URL
3641	TOOLSDOC-274	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 811977|https://bugzilla.redhat.com/show_bug.cgi?id=811977] Ken Johnson <kejohnso@redhat.com> made a comment on jira JBDS-2046 Comments on a few of the above entries: First on the JBDS SOA Update site above. 1. I don't think "SOA-P" means anything to customers. It's a very internal term. Suggest simply "SOA" since we've done away with the old collection named with "SOA" 2. The ModeShape tools go under JBoss Data Servcies rather than SOA Second on the most recent screenshots above. They look good and are less confusing. Couple points: 1. Do you think users get the difference between "Project Tooling" and "Product Tooling". Assuming your eyes pick up the subtle wording, does that carry significant meaning to users? Should it be "Community Project Tooling" and "Product Tooling" or a similar distinction? Or maybe it's "JBoss Tooling by Project" vs "JBoss Tools by Product"? Maybe splitting hairs a bit. 2. Under Product we have two products SOA-P and BRMS. There are actually 3 products: SOA Platform, Data Services Platform and BRMS. Rather than creating a new product category, perhaps rename the "SOA-P Tooling" to "SOA and Data Services Tooling". Since Data Services is a superset this makes more sense than breaking them out. Here the Teiid and ModeShape tools would be under "SOA and Data Services Tooling" with ESB, jBPM 3, BPEL, etc. Thoughts?	1	0	https://issues.jboss.org/browse/TOOLSDOC-274
3642	TOOLSDOC-274	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 811977|https://bugzilla.redhat.com/show_bug.cgi?id=811977] Len DiMaggio <ldimaggi@redhat.com> made a comment on jira JBDS-2046 Just to re-re-re-cap 1) We will end up with these (2) choices: * SOA and Data Services Tooling * BRMS Tooling 2) And - these same titles will be used in both JBoss Central and the update site. Correct?	0	0	https://issues.jboss.org/browse/TOOLSDOC-274
3643	TOOLSDOC-274	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 811977|https://bugzilla.redhat.com/show_bug.cgi?id=811977] Max Rydahl Andersen <max.andersen@redhat.com> made a comment on jira JBDS-2046 Len - BRMS is a product name - so no, that is not what it would be. "Business Rules" was my suggestion for general category name for it. Nick, and yes there are two one for JBT and JBDS and the suggestion is we use the same category names to avoid having confusion. About jbpm5 I think this issue has outgrown its ability to handle actual content questions. Open a new for that and ask Doug to work on that ?	0	0	https://issues.jboss.org/browse/TOOLSDOC-274
3644	TOOLSDOC-274	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 811977|https://bugzilla.redhat.com/show_bug.cgi?id=811977] Nick Boldt <nboldt@redhat.com> made a comment on jira JBDS-2046 To avoid further confusion, I'm going to close this and have opened to subtasks: https://issues.jboss.org/browse/JBDS-2112: should jbpm5 be included in both SOA and Business Rules connectors / categories? https://issues.jboss.org/browse/JBDS-2113: relabel categories / connectors in Central and Update sites Let's move the more focussed discussions there.	0	0	https://issues.jboss.org/browse/TOOLSDOC-274
3645	TOOLSDOC-271	BTW, another thing in the server editor should get documented: JBIDE-10319. I'm not sure what is the process here - should I add a documentation label on that JIRA?	1	1	https://issues.jboss.org/browse/TOOLSDOC-271
3646	TOOLSDOC-261	Isaac, maybe I am missing something, but where could I find this information? In official Seam3 servlet module documentation? Or we have documentation for seam3 tooling? Because as I know we haven't had it.	1	1	https://issues.jboss.org/browse/TOOLSDOC-261
3647	TOOLSDOC-261	Jaroslav: I am assuming this wasn't done in the 6.0 documentation. Please can you confirm that this is still a relevant issue? If it is then I will push it to JBT4.1/JBDS7.0 and make sure something is included.	1	0	https://issues.jboss.org/browse/TOOLSDOC-261
3648	TOOLSDOC-259	Hey Peter, I understand your sentiment here and you are correct, but there is only one of me for all of JBDS documentation. Currently the JBDS and JBT docs are built from the same source XMLs. There is also the possibility that JBDS could be used with BRMS and community bits so perhaps we can mention both - with an emphasis on BRMS? Are we talking about just mentioning BRMS whenever Drools Guvnor is mentioned, or are you looking for whole new sections or chapters on BRMS and JBDS? Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-259
3649	TOOLSDOC-259	Hey Peter, I am happy to add in information concerning the tooling working with BRMS, however considering there is about 10 days until release I am not in a position to learn BRMS. Can you provide the content for the guide at all? If not, then this will miss this release. Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-259
3650	TOOLSDOC-250	That's different. Here we're talking about project examples files - and project examples XML files for both JBoss Tools and JBDS appear at that same location - http://download.jboss.org/jbosstools/examples/ (you can see some explicitly have "jbds" in the title). We're simply moving these files to a common location so folks can find and add them quickly to the tooling. Does that help? Do you need steps to document for the user here 'cause I can do that... (added below) Let me know if you need anything else. Happy to help! ---------- To add the examples in JBDS 4.1.0 or JBDS 4.1.1, do the following: 1. Go to Windows->Preferences and look for the JBoss Tools->Project Examples page. 2. Click Add 3. Give the new example site a name and this url: http://download.jboss.org/jbosstools/examples/soa-p-5.x-samples.xml 4. Click OK to close the dialog, then apply and close the preference change. 5. Go to Help->Project Examples 6. Click "Show Experimental Sites" at the top of the dialog. 7. Select your new examples site, select one of the newly available examples: aggregator business_service custom_action fun_cbr helloworld_action helloworld simple_cbr static_router transform_XML2POJO Each of these is suffixed with "_eclipse" when imported to differentiate them from the quickstarts included with the runtime itself. Each example includes one or more deployable ESB projects and a test project. Once you have the project(s) imported, if your server was created ahead of time, the example projects should be ready to go with no additional changes. You can then deploy the non "_test" projects and run the ESB or JMS message testers from the "_test" project after a successful deployment. ----------	1	1	https://issues.jboss.org/browse/TOOLSDOC-250
3651	TOOLSDOC-250	Brian, hopefully just one final question: you mentioned this: "Once you have the project(s) imported, if your server was created ahead of time..." What happens if a user creates the server they wish to deploy the examples to after importing them?	1	1	https://issues.jboss.org/browse/TOOLSDOC-250
3652	TOOLSDOC-250	If you do it backwards (i.e. import the example and then create the server), you'd have to fix the project build path by hand (i.e. tweak the ESB and server library entries to point to the new server)... 1. Right-click on the project, select Build Path->Configure Build Path 2. Select the Libraries tab on the "Java Build Path" property page. 3. Select the JBoss Enterprise Application Platform 5.x Runtime library. 4. Click Edit... 5. Select the runtime you want from the list. 6. Click Finish. 7. Select the JBoss ESB Runtime library. 8. Click Edit... 9. Select the "Server Contained" entry in the list that matches the server picked in step 5. 10. Click Finish. 11. Click OK. That should take care of most, if not all, problems with the project so that it builds and can deploy. Does that help?	1	1	https://issues.jboss.org/browse/TOOLSDOC-250
3653	TOOLSDOC-246	Isaac on irc indicated that this was probably an eclipse error versus a docs error. If it is an easy fix, could we consider adding to 4.1.1 release?	0	0	https://issues.jboss.org/browse/TOOLSDOC-246
3654	TOOLSDOC-246	Hey Nick, could you take a look at this please. The image is there when accessing the book through a web browser so it seems to be an issue with Eclipse over being a doc issue. Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-246

#	Requirement	Questions	Answered	Successfully Answered	URL
3655	TOOLSDOC-246	1. installed jbdevstudio-product-linux-gtk-x86_64-4.1.1.v201110060056R-H730-M2.jar 2. images load fine throughout the doc, including in section 2.2.5 of Struts Tools Tutorial where the flowchart appears. Multiple other images and icons also appear. Perhaps something was fixed between the 4.0 and 4.1 docs' release? Is the path to that image relative or absolute? I thought maybe it was an absolute path to a VPN-only URL, but after dropping off VPN other docs referring to that image all loaded fine. If the images are relative paths and they're contained inside a plugin, maybe the problem was the plugin's build.properties failed to include the images when building the binary? Anyway, looks good here. Cannot reproduce your screenshot.	1	1	https://issues.jboss.org/browse/TOOLSDOC-246
3656	TOOLSDOC-192	[~irooskov] Has this JIRA already been resolve? The links for information no longer work and I have no idea what this JIRA is about. Please can you comment.	0	0	https://issues.jboss.org/browse/TOOLSDOC-192
3657	TOOLSDOC-140	Hey Maxim, is there a write-up or can you give me a quick summary and maybe an image of what the final outcome was for the new wizard please? Reading through the JIRA there seems to be a lot of back and forth on what was actually done and I'm finding it hard to nail down what should be documented. Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-140
3658	TOOLSDOC-107	If you mean "how do you use Maven to do our JBT builds?" we have a few docs about this in the wiki already: http://community.jboss.org/wiki/JBossToolsTargetPlatformProvisioningorhowtoresolvemissingdependencieswhencompilingJBossTools http://www.jboss.org/tools/docs/building http://www.jboss.org/tools/docs/testing If you have other questions, please ask specific ones and I'll wiki/faq up what you need.	1	1	https://issues.jboss.org/browse/TOOLSDOC-107
3659	TOOLSDOC-100	So I'm guessing the confusion comes from the fact that I didn't manage to complete the thought... The "Simple Web Service" wizard provides a way to handle the Bottom-up case in Chapter 3 as well as the WTP Web Service wizard. It does so more easily than the WTP wizard. You talk about that in Chapter 2. But we don't mention in Chapter 3 that you can use the "Simple Web Service" for the Bottom-up case if you provide your own annotated classes for the Class and/or Application Class fields. Is there a way to refer back to Chapter 2 from Chapter 3 suggesting that using the "Simple Web Service" is another way to handle the bottom-up case?	1	1	https://issues.jboss.org/browse/TOOLSDOC-100
3660	TOOLSDOC-90	Hi Isaac, can you have a quick look at this doc to make sure I haven't told any whopping lies?	1	1	https://issues.jboss.org/browse/TOOLSDOC-90
3661	TOOLSDOC-77	New structure: 1. Introduction 1.1. What is ESB and JBoss ESB Tools? 1.2. Key Features of the JBoss ESB Tools 1.3. Requirements and Installation 2. Tasks 2.1. Creating a ESB Project 2.2. Using ESB Project Examples 2.3. Deploying a ESB Project 2.4. Creating a ESB File 2.5. Creating a ESB Action 2.6. Configuring ESB Runtime in Preferences 2.7. Using and Configuring SOA Platform 3. Reference 3.1. JBoss ESB Editor 3.1.1. Tree View 3.1.2. Source View 3.2. ESB Editor Features 3.2.1. ESB syntax validation 3.2.2. Support for XML Schema 3.2.3. Content Assist for ESB XML file 3.2.4. OpenOn for ESB XML file 3.2.5. Synchronized Source and Visual Editing 4. Summary 4.1. Other Relevant Resources on the Topic	0	0	https://issues.jboss.org/browse/TOOLSDOC-77
3662	TOOLSDOC-76	Could you let Isaac know about where the current Teeid eclipse tooling docs are located ?	1	0	https://issues.jboss.org/browse/TOOLSDOC-76
3663	TOOLSDOC-71	Can you recommend a book that this information might be best suited to? I can only think of the 'Getting Started Guide' but that is JBDS specific so the content wouldn't reach our JBoss Tools users.	1	0	https://issues.jboss.org/browse/TOOLSDOC-71
3664	TOOLSDOC-63	The following Guide structure is provided: 1.Introduction 1.1. Guvnor Tools Key Features 1.2. What is Guvnor Tools? 1.3. Requrements and Installation. 2. Guvnor Tools Tasks 2.1 Adding new Guvnor connection 2.2 Importing Guvnor Repository Resources 2.3 Getting Local Copies of Guvnor Repository Resources 2.4 Managing Guvnor Resources 3. Reference 3.1 Guvnor Perspective 3.2 Connection Wizard 3.3 Repositories View 3.4 Resource History View 3.5 Guvnor Preferences 3.5.1 Guvnor Repository Connection Preferences 3.5.2 Local Guvnor Repository Resource Decoration Properties. 4. Summary 4.1 Other relevant resources on the topic.	0	0	https://issues.jboss.org/browse/TOOLSDOC-63
3665	TOOLSDOC-43	JSF Reference Guide contains detailed description for all fields in the wizards, but they are missing from the Help tray for some reason. Could you please check why they are not picked up?	0	0	https://issues.jboss.org/browse/TOOLSDOC-43
3666	TOOLSDOC-37	Since the installer now only uses the ALL zip as part of the assembly process, but does not publish it (in favour of the Update zip), and since this is documented in the what's new (from M2)... is this issue done?	1	1	https://issues.jboss.org/browse/TOOLSDOC-37
3667	TOOLSDOC-27	reopen for clarification. What is xmldiffpatch.exe ? Where is it from ? What is the license ? How is it to be used ?	1	1	https://issues.jboss.org/browse/TOOLSDOC-27
3668	TOOLSDOC-25	This is a very broad JIRA - can you be more specific please	0	0	https://issues.jboss.org/browse/TOOLSDOC-25
3669	TOOLSDOC-18	what's ne on the topic - http://docs.jboss.org/tools/whatsnew/hibernate/hibernate-news-3.3.0.M3.html - "Hibernate specific annotations"	0	0	https://issues.jboss.org/browse/TOOLSDOC-18
3670	TOOLSDOC-17	Svetlana that page no longer exists, can you give any more details on what needs to be added? Thanks ;)	0	0	https://issues.jboss.org/browse/TOOLSDOC-17
3671	TOOLSDOC-15	Its standard XML Catalog behavior - do we really need to mention this for *every* XML that is possible to edit in the tools ? http://www.eclipse.org/webtools/community/tutorials/XMLCatalog/XMLCatalogTutorial.html	1	1	https://issues.jboss.org/browse/TOOLSDOC-15
3672	TOOLSDOC-15	Hey Shaun, this JIRA explains that the issue occurs when: "using a newer version of hibernate than what is in the core", and max points out that this is not only related to Hibernate. However, using Hibernate as an example, are you referring to someone using a newer version of the hibernate tooling or hibernate itself? Thanks, Isaac	1	1	https://issues.jboss.org/browse/TOOLSDOC-15
3673	TOOLSDOC-14	What's new details added. Updates applied.	0	0	https://issues.jboss.org/browse/TOOLSDOC-14
3674	TOOLSDOC-3	Seam Dev Tools Reference Guide: 1.3."Installation into Eclipse -this section specific for JBT." Chapter 3."Maven Integration for Seam Projects also specific for JBT." Figure 4.3. "Changing the Seam Facet Version".Some facets on the picture is specifig for JBT and are not included into JBDS.	0	0	https://issues.jboss.org/browse/TOOLSDOC-3
3675	TOOLSDOC-3	ESB Tools Reference Guide: 1.1. What is ESB? -in this section the info that the guide will provide installation process info is mentioned. But it's not requred for JBDS guide version because ESB is already installed there.So this sentence is specific for JBT. 2.1. ESB Tools Installation -this section is specific for JBT. 2.5. "Creating a ESB File" -there is a "Note" in this section that specifies which versions of JBoss ESB Tools supports what jboss-esb.xsd versions.This info needs some refactoring in the context of JBDS guide, because a user could have no idea what version of JBDS goes with what JBT.The same problem appears in 2.2. Creating a ESB Project section(the last step in configuring the project).	0	0	https://issues.jboss.org/browse/TOOLSDOC-3
3676	TOOLSDOC-3	Drools Tools Reference Guide: 1.1. What is Drools? -describes Drools tools only as a plugin to eclipse but not as a part of JBDS.	0	0	https://issues.jboss.org/browse/TOOLSDOC-3
3677	WFKDOC-110	Suz, Do you have any migration docs as an example? I'm asking you that because intuitively https://deltaspike.apache.org/migration-guide.html makes sense for me, but I know that it could be much better for people that are starting the first contacts with DeltaSpike. Plus, I never used CODI, so my idea is to forward those questions to dev@deltaspike.apache.org and ask for contributions. Having a migration page as an example can help people to understand how it should looks like.	1	1	https://issues.jboss.org/browse/WFKDOC-110

#	Requirement	Questions	Answered	Successfully Answered	URL
3678	WFKDOC-110	Suz, I'll propose the following: Once that we have some documentation on deltaspike git repository, I'll ask at deltaspike@dev.apache.org for help. Wdyt ?	1	1	https://issues.jboss.org/browse/WFKDOC-110
3679	WFKDOC-110	Sure [~rafabene]. I'll show you what we've got so far anyway, can you view this? https://github.com/sdorfiel/deltaspike/tree/WFKDOC-103/documentation	1	1	https://issues.jboss.org/browse/WFKDOC-110
3680	WFKDOC-110	At https://github.com/sdorfiel/deltaspike/blob/WFKDOC-103/documentation/overview.adoc#background-portable-cdi-extensions - I suggest you to replace whole section for this one only: http://deltaspike.apache.org/documentation.html#what-is-a-portable-cdi-extension - At https://github.com/sdorfiel/deltaspike/blob/WFKDOC-103/documentation/overview.adoc#about-deltaspike - These capabilities can be added to projects that already -do not- have CDI enabled, -providing a method of connecting existing projects with new customization-. - -This provides a method of integrating custom and third party code to existing projects- - Users can -inject- use modules - -These capabilities are added when a user inserts a module into the classpath of their project. Injected objects can be customized and have interceptors associated to them.- - -DeltaSpike uses CDI container integration provided by Arquillian for booting and controlling context lifecycles. Containers are activated by Maven, which utilizes the profiles JBoss Weld and OpenWebBeans-	1	1	https://issues.jboss.org/browse/WFKDOC-110
3681	WFKDOC-108	{quote} Missing Seam2 and Seam 3 info - are these both still relevant? {quote} We're currently experiencing serious issues while trying to migrate Seam2 and Seam3 applications from JBoss AS7 to WildFly 8. Given the fact that Seam2/3 is discontinued and recommends DeltaSpike as migration target it would be very much appreciated to have some documentation/examples in place for that kind of migration scenarios.	1	0	https://issues.jboss.org/browse/WFKDOC-108
3682	WFKDOC-106	[~rafabene], could you describe an example application that would have CDI only, and an example of how the user might want to extend that application and how DS can assist. The links you have are great for overviews of the individual plug-ins. But we're looking for a more concrete example about an application to provide the user. For example, http://aerogear.org/docs/unifiedpush/ups_userguide/overview/#_use_cases_and_scenarios For example, in one of those links you state "One of the frequently needed features is accessing ServletRequest in a CDI bean." Suz and I don't have sufficient knowledge to be able to put this into context. Could you describe an example application / scenario where this is useful. What type of app would it be/ what functionality would it have? What would be the difficulties for the app in just using CDI and what is the effect for the developer after adding this annotation in their code? Thanks.	0	0	https://issues.jboss.org/browse/WFKDOC-106
3683	WFKDOC-105	Topic: Using DeltaSpike without Maven - We do encourage them to use the Maven Instructions even the users being able to build DeltaSpike themselves. If they decide to not use Maven, they're required to do some additional steps like placing it manually inside WEB-INF/lib or <EAR_ROOT>/lib depending the type of the project: EAR, WAR. Topic: Adding a Server or Servlet Container CDI Implementation - I couldn't find this topic at the documentation to understand the context. Can you send me the link, please ? Thanks	1	1	https://issues.jboss.org/browse/WFKDOC-105
3684	WFKDOC-105	Thanks [~rafabene]. Using DeltaSpike without Maven Can you spell the extra steps out for me here? Is it something akin to *Copy the downloaded file into a project's WEB-INF directory *Run X command ? Adding a Server or Servlet Container CDI Implementation https://deltaspike.apache.org/documentation.html#deployment-mode Additionally, are these the correct system requirements? Maven 3 Java EE (and/or SE?) Preinstalled CDI (version number?) Red Hat JBoss Enterprise Application Platform	1	1	https://issues.jboss.org/browse/WFKDOC-105
3685	WFKDOC-105	Thanks for that [~rafabene]. So... Using DeltaSpike without Maven: After the user has put the downloaded files into the locations you noted, do they have to do anything else (run any commands, restart anything) or is that it in terms of configuration? Adding a Server or Servlet Container CDI Implementation: Yeah, I see what you mean about the varience there. Makes it a bit challenging. What if we're just working in a JavaSE environment? Where does the user put the code snippet in that link? (https://deltaspike.apache.org/documentation.html#deployment-mode Start a CDI container using Java SE) Does it go in the project's pom file?	1	1	https://issues.jboss.org/browse/WFKDOC-105
3686	WFKDOC-105	Thanks [~rafabene]. And do you only need Weld and OWB when using JavaSE?	1	1	https://issues.jboss.org/browse/WFKDOC-105
3687	WFKDOC-105	Thanks [~rafabene]. We also need a DeltaSpike profile to add to Maven's settings.xml. Something a user can copy and paste into their own settings file. (Akin to the example-settings.xml file in WFK's Maven repository.) On a similar note, what is the URL the user needs to use Maven Central if they're adding it to settings.xml? For example, https://maven.repository.redhat.com/techpreview/all/ or https://repository.jboss.org/nexus/content/groups/public-jboss/	1	1	https://issues.jboss.org/browse/WFKDOC-105
3688	WFKDOC-96	[~nids1409], [~Nidhis], [~aakanksha_writer], [~nshendye], [~themjboss], [~d_michael], [~bobjboss], will one of you please pick up the Deploy and View the Application topic for a peer review? Once you have decided on reviewing the topic you must re-assign the JIRA ticket to yourself (currently, I am the assignee) before you actually start the review. DocBuilder link to the topic: http://docbuilder.usersys.redhat.com/22729/#sect-Deploy_and_View_the_Application	0	0	https://issues.jboss.org/browse/WFKDOC-96
3689	WFKDOC-77	[~manaRH] - I had a look at the links provided but can't see any Seam migration info. Do you know where the info went, if there ever was any info?	1	1	https://issues.jboss.org/browse/WFKDOC-77
3690	WFKDOC-77	[~manaRH] - need this content asap if it is to make it into the Seam doc. what's your timeline for getting it done?	1	1	https://issues.jboss.org/browse/WFKDOC-77

#	Requirement	Questions	Answered	Successfully Answered	URL
3691	WFKDOC-69	Kit" chapter as the very first chapter. See Hibernate Search Guide for instance. Make sure to add the topics that have already been created to the Explore the Examples content spec - do not create new topics. * Chapter 1. Exploring the Examples Overview ** Change chp title to match doc title "About Explore the Examples". ** This section needs simplifying and the strength of the text improving. It feels rambling and repetitive. This should be a very clear concept topic with the subject in this case the Explore the Examples doc. Remember that a good concept topic states a high level definition of the subject, the purpose of the subject, what the subject consists of and any distinguishing characteristics. So "Explore the Examples is … Its purpose is … The Guide consists of …" For example, see the "About Start Developing" topic in the WFK 2.6 Start Developing doc. * Chapter 2. The contacts-mobile-basic Example ** Please review the content specs that we set out for each of the example chps prior to commencing writing. These did not have Overview sections and specified the summary topic of the example be used as Initial Text for the chp. Please implement this as we planned. The reason I am asking you to remove the Overview topic is that it duplicates the information provided in the summary topic of the example. Likewise for chp 3. * 2.2. The contacts-mobile-basic Example ** "The example can be deployed to a desktop browser to observe in a testing environment." >> what does this mean? "deploy to a desktop browser" does not make sense - an app is deployed to a server and viewed in a browser; how does a desktop browser make a testing environment? does the app look the same on desktop browsers and mobile device browsers? ** "The framework is used to add widgets to pages, such as headers, buttons and forms. These widgets create an interactive, mobile-friendly interface that is touchscreen enabled." > would be more accurate to express as "The framework is composed of widgets that enable interactive, mobile-friendly, touchscreen enabled application pages with items including headers, buttons and forms." ** mobile-friendly > it's more than mobile-friendly, the primary design is for mobile devices ** touchscreen enabled > does jQuery Mobile provision this? ** "This chapter examines how to build and deploy the example, explore how it uses the jQuery Mobile framework, and provide details on how widgets are used." > typos * 2.3.1. System Requirements ** The system requirements for the two examples are identical so let's simplify the doc instead of presenting duplicate information. Put the System Requirements topic in the chapter with the About Explore the Examples topic. See Chp2 of the Start Developing doc as an example. ** Simplify the information in this section; we want to limit the amount of info we have to update in this section between releases so make use of the Supported Configurations page more heavily; don't list the browsers, instead point users to the Supported Configurations page for this info. See the "Start Developing System Requirements" topic in the Start Developing doc for an example and consider copying the appropriate text and style from that topic. ** "Current version of Maven with an accessible repository." > what do you mean by an accessible repository? * 2.3.2. Obtaining the Example ** "unzipped into" > extracted into; don't use unzip (see IBM pg#372) ** Proc title needs to reflect the proc task, which in this case is more than downloading ** The downloading steps of the proc need simplifying since the CP has changed - please see "Obtain the Project Source Code" topic in Start Developing doc and copy the appropriate text and style from that topic. ** "To unzip the file on the command line, navigate to the directory you saved the ZIP file in and execute the following command: unzip" > This command is insufficient to extract the contents, it would need to be: unzip filename.zip; "navigate to the directory in which you save the .zip file"; don't use execute (see IBM pg#327), instead "enter the following command:" * 2.3.3. Building the Example ** "it must be built." > replace with "the source code must be built." ** "This is accomplished by using Maven to install the example locally." > Maven does not install anything, it builds things. ** "Execute the following command to build the example:" > don't use execute (see IBM pg#327), instead "Enter the following command:" * 2.3.4. Deploying the Example ** "To	1	1	https://issues.jboss.org/browse/WFKDOC-69
3692	WFKDOC-24	I have successfully committed the changes to svn last night. Can you check again and let me know if you are still having problems viewing the updates.	0	0	https://issues.jboss.org/browse/WFKDOC-24
3693	WFKDOC-23	[~maschmid], couple of questions: # Should [1] be referencing EAP 6.1.0, 6..1.1, 6.2.0 and 6.2.1? Or just EAP 6.2.0 and 6.2.1? Or just EAP 6.2.0? # Can/should the example in [2] of spring-3.0-bom be updated to spring 4.0? f so, is there anything in the <dependencies> that needs changing? [Seems since we are deprecating 3.0 we should use newer versions in the example] # Are there any specific migration notes? Can the existing info in [3] be removed or are these relevant for users migrating from 2.4.0 to 2.5.0 too? [1] https://access.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Web_Framework_Kit/2.4/html-single/Maven_Repository_User_Guide/index.html#chap-Installing_the_Repository [2] https://access.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Web_Framework_Kit/2.4/html-single/Maven_Repository_User_Guide/index.html#Using_JBoss_Stacks_in_Your_Project [3] https://access.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Web_Framework_Kit/2.4/html-single/Maven_Repository_User_Guide/index.html#Migration_Notes_for_JBoss_Stacks_from_WFK_2.3.0_and_Earlier	1	1	https://issues.jboss.org/browse/WFKDOC-23
3694	WFKDOC-23	1. Should [1] be referencing EAP 6.1.0, 6..1.1, 6.2.0 and 6.2.1? Or just EAP 6.2.0 and 6.2.1? Or just EAP 6.2.0? EAP 6.2.0 with a note that deprecated components (actually only Errai) requires both EAP 6.2.0 and 6.1.1 maven repositories. 2. Can/should the example in [2] of spring-3.0-bom be updated to spring 4.0? f so, is there anything in the <dependencies> that needs changing? [Seems since we are deprecating 3.0 we should use newer versions in the example] The best would be to update to spring-3.2-bom. All of those artifacts are in 3.2 also, so no other change is needed there. 3. Are there any specific migration notes? Can the existing info in [3] be removed or are these relevant for users migrating from 2.4.0 to 2.5.0 too? it is only relevant to users migrating from WFK 2.3 and older. I think it makes sense to keep it for such users, given the release frequency of WFK... AFAIK there are no other migration notes for WFK 2.5 specifically.	1	1	https://issues.jboss.org/browse/WFKDOC-23

#	Requirement	Questions	Answered	Successfully Answered	URL
3695	WFKDOC-22	1) Red Hat JBoss Web Framework Kit is a product that combines Suggestion: need not mention 'is a product that' 2) You can create, modify, and deploy Java applications using the framework. Suggestion: You can create, modify, and deploy Java applications using the frameworks (or probably 'Web Framework Kit') 3)It allows for the use of .... Suggestion: 'It' should be replaced with the actual thing (like frameworks or Web Framework Kit) to be explicit and clear Suggestion: ...allows the use of... ('for' can be dropped) 4) ...approach to making applications... Suggestion: ...approach to make applications.. 5) You can combine these popular frameworks... Suggestion: 'popular' can be dropped 2.1 6)....confirmed frameworks and... Suggestion: serial comma - confirmed frameworks, and 3.1 7) This task shows you how to use the former. Suggestion: 'this task' – which task 'former' – better to be specific in a technical documentation In fact we can just do without this sentence and have a meaning name for Procedure 3.1 8)Procedure 3.1. Task Suggestion: Name of the procedure should be something more meaningful than just saying 'task' maybe 'Installing Component Frameworks using Maven' or 'Installing Web Framework Kit using Maven'.... 9) Note : The source for the frameworks is available from the Red Hat Customer Service Portal for inspection. Suggestion: The information in the note is covered in the 1st sentence of section 3.1, except for the hyperlink which can be added there as well. 10)IMO sections 3.1.1 and 3.1.2 should not be subsections of 3.1 but at the same level as 3.1, because all these sections are on the same lines 11) Chapter 3 is comparatively huge and can be split into 2 chapters ---the 3 installation topics as one chapter (3.1, 3.1.1, 3.1.2) ---examples and quickstarts as another chapter (3.2, 3.3) 12) Procedure 3.1 -> step 2 Suggestion: Name of the icon should be bold --Clicking on the icon is not allowed, can click on either of the 2 buttons available so in this case – need to click the Download Software button on the icon --Complete the step by saying – log in to continue 13) Procedure 3.1 -> step 3 Suggestion: In the left pane, click Web Framework Kit. -- as you are saying 'current release of WFK' you need not mention the version number of WFK 14)When the files have been downloaded, unzip them. Suggestion: Unzip the downloaded files. 15) Procedure 3.1 -> step Suggestion: should have a link to the section about installing Maven repo (even if it is just the next section) 3.1.1. 3.1.1.1. Installing the Repository 16) There are three ways how the Red Hat JBoss Web Framework Kit Maven repository can be installed: Suggestion: There are three ways to install Red Hat JBoss Web Framework Kit Maven repository: 17) Use the option that best suits your environment. Suggestion: IMO this is obvious and can be dropped or maybe rephrased to provide more useful information to the user. 18) Instead of having 2 blocks for Important they can be clubbed into 1 block as Important 1. 2. 19) Local File System Suggestion: Installing Maven Repository on a Local File System Global Suggestion: IMO names of all section, subsections, titles, procedures....should have complete meaning 1. This ways the name will say it all, the reader need not refer the above of below text for details 2. It will also be useful while cross-referencing I am not suggesting long names or titles but names that are complete in themselves 20)This option is the most suitable for testing in a small team. Suggestion: Can we use 'method' instead of 'option'? Can this information be elaborated a bit? 21)This option is the most suitable for multi-user environments. Suggestion: Can we use 'method' instead of 'option'? Can this information be elaborated a bit? 22) Suggestion: 'the' should not be used before product names. 23) https://access.redhat.com/site/documentation/en-US/JBoss_Enterprise_Application_Platform/6.2/html-single/Installation_Guide/index.html#Installation_Methods_for_JBoss_Enterprise_Application_Platform_6 Procedure 4.1 at the above link - I personally think this is a good style of writing crisp and to the point procedures 24) Global suggestion: e.g. should be spelled out as 'for example'. 'e.g.' is actually Latin and should be avoided in technical English documentation. 3.1.1.2. 25) with JBDS, and manually Suggestion: remove comma 26) The procedure below demonstrates... Global Suggestion: avoid	1	1	https://issues.jboss.org/browse/WFKDOC-22
3696	WFKDOC-22	Excellent feedback [~snagarwa]! Make Sneha's changes, except for the following: * 1-5) Sneha will be modifying 'About WFK' in WFKDOC-49 * 9) Keep this note but make it more useful. Which files? And don't hide the link for 505 compliance - "available from <url> on the Red Hat Customer Portal" * 15) Do but add a link but as a 'Refer to' in the content spec * 17) Do - in which instance would the user choose each option? * 18) Do but use bullets instead of numbers as order is not essential * 22) can use 'the' before product name when that name is being used as an adjective, for example 'the JBoss EAP repositories' * 23) Yes and no, it is a good example in terms of succinctness but could be better, should put the 'where' first in each step so 'From the drop-down list, select blah' * 27) I disagree, using 'the' before common nouns is perfectly valid but usually not valid before proper nouns (exceptions exist, e.g. the Queen) * 35) Do but either drop 'need' or change to 'must' if it is crucial for user to do * 40) Do but either 'the .jar file' or 'the JAR file' (IBM pg#205) Make these additional changes: * 57) Don't write 'click the blah button', should be simply 'click blah' - check whole document (IBM pg#212) * 58) Procedures 3.12 and 3.13 need titles not just 'Task' - see Sneha's point 8) * 59) Section 3.2.1 Downloading the Snowdrop Example should be generalized to all examples	1	1	https://issues.jboss.org/browse/WFKDOC-22
3697	WFKDOC-22	Part of phase 2 review comments: 1)Procedure 3.1. Suggestion: It will be helpful to the users if we can add a tip/information box describing what the diferent types of downloads (demo, binary, source code) mean. 2)Procedure 3.1. -> Step 4 -> and click Download Suggestion: and click the corresponding Download link 3) Procedure 3.1. -> Step 4 -> When the files have been downloaded, unzip them. Suggestion: Unzip the downloaded files. 4) Procedure 3.1. -> Step 5 Suggestion: It is not clear which Maven repository and where it is. Either add more info here to add link to relevant section (like say – For details on Maven repository.....refer to section.....) ----Similarly for adding profile.... 5) Procedure 3.2. -> Step 2 -> ...will be ... Global Suggestion: replace with 'are' 6) Global suggestion: can we generalize the version numbers in all the procedures so that we have less version numbers to update every release? 7) Procedure 3.3. -> Step 3 Suggestion: Where is info on configuring the Apache Web Server? Are we assuming that the user should know it? 8) Procedure 3.4. -> Step 2 ->Select the Repositories section from the Views/Repositories menu to the left of your repository manager. Suggestion: Items should be in the sequence/order to be followed: -- In the left pane of the repository manager, go to Views/Repositories menu and select Repositories. 9) Procedure 3.4. -> Step 4 -> a. -> Add... Suggestion: should be 'Add' only; do not use '...' even if it is there in the UI 10)Procedure 3.4. -> Step 4 -> b. ->Give the new repository a name and ID. Suggestion: Assign a name and ID to the new repository 11) Procedure 3.4. -> Step 4 -> c. -> into Suggestion: Should be 'in' 12) Procedure 3.4. -> Step 5 ->The Nexus server is now configured to use the repositories. If you want the repositories to be available in a repository group, please follow these additional steps: Suggestion: drop 'now' -- For the repositories to be available in a repository group, execute the following steps: 13) Procedure 3.4. -> Step 5 -> a -> Suggestion: Select a repository group for the repositories and click the Configuration tab. 14) Procedure 3.4. -> Step 5 -> b -> Suggestion: -- 'drag' – the item can be moved either by dragging or by using the buttons so I think better to say 'Move the new ...' -- Move the JBoss Web Framework Kit repository from the Available Repositories list to the Ordered Group Repositories list. 15) Procedure 3.4. -> Step 5 -> Note Suggestion: -- 'Please note that' is not required inside a Note box -- The order of the repositories in the Available Repositories list determines the priority of repositories to be used when searching for Maven artifacts.	1	1	https://issues.jboss.org/browse/WFKDOC-22
3698	WFKDOC-18	[~ppitonak], is this purely about updating the enterprise docs or the community ones too?	1	1	https://issues.jboss.org/browse/WFKDOC-18
3699	WFKDOC-18	[~mmurray], I still see old link on [1]. How often is documentation generated? [1] http://documentation-devel.engineering.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Web_Framework_Kit/2.5/html/RichFaces_Component_Reference_Guide/chap-Component_Reference-Introduction.html	1	0	https://issues.jboss.org/browse/WFKDOC-18
3700	WFKDOC-15	[~mmurray], is there a suitable way to provide updated content for this guide ? Notwithstanding bug fixes, the feature could be treated complete.	1	1	https://issues.jboss.org/browse/WFKDOC-15

#	Requirement	Questions	Answered	Successfully Answered	URL
3701	WFKDOC-15	[~vineet.reynolds], not sure we have a great mechanism for updating content directly. Could you use DocBuilder to provide input? So http://docbuilder.usersys.redhat.com/16724/ is the guide that this jira is referring to. You can click the 'Report a Bug' link at the end of any of the sections and this will open semi-completed jiras in which you could provide feedback or the required text updates. I think primarily the sections you need to review are 1.4 Application in Action and 2.1 Client Source Code, and 2.2 Server Source Code if that too has changed. We'll test out the ER1 build ourselves and use it to take new screen captures as necessary.	1	1	https://issues.jboss.org/browse/WFKDOC-15
3702	WFKDOC-6	I misunderstood, I thought fixing the language was for Chapter 3 and not the entire guide. I will address this. I see the screencap you're referring to. When you say step 1 can be left within a screencap, did you mean the one originally used in step 1? (Which has now replaced the one in step 2 as requested.)	1	1	https://issues.jboss.org/browse/WFKDOC-6
3703	WFKDOC-5	Great work Suz - looks good. You were objective in picking out the important intro points in 2.3, something I was finding hard to do because I wrote one of them. Q: Was there a specific need to create the new topic for 2.1? I was thinking that you could use the existing JBDS topic [22879] but perhaps there is something I'm missing. Minor fixes in 2.3: * the condition for wfk needs to go in the <step> tag otherwise you end up with an empty but numbered step in the final copy for jbds * also spaces don't necessarily get stripped at the start and end of <phrase> tags so be sure to start the first word straight after the closing > of phrase and likewise at the end, e.g. <phrase condition="">The.</phrase> rather than <phrase condition=""> The. </phrase> * for the wfk version, we only want users to do step 2 or 3+4; you could collapse 3 and 4 into one step "To create a new SSH private-public key pair, click New and complete the fields and options as detailed:" but how can you direct the user to step 2 or 3 only? If you're not sure we can chat about it * Suggest rearranging the sentences in the intro for better reading: {quote} Secure Shell (SSH) keys are essential when working with OpenShift as *t*hey *are* used to establish the encryption between the IDE and the OpenShift server. *SSH keys* enable you to develop and access deployed applications *and they* are also used to control access of other contributors to your OpenShift applications. You can use *an* existing key *pair* or generate a new one in the IDE. SSH keys must be uploaded to the OpenShift server and, as detailed in the procedure below, OpenShift Tools can assist with both the generation and uploading of SSH keys to OpenShift. {quote}	1	1	https://issues.jboss.org/browse/WFKDOC-5
3704	WFKDOC-4	[~sdorfiel], good start with adding the conditions. I've noticed the following needs fixing: 1. The itemized list in Do you need help? is missing for wfk 2. The para following the itemized list in Do you need help? is missing for wfk 3. <application> tags in give us feedback need to surround the whole jira project title 4. jbt needs adding to the condition for 'JBoss Tools and Developer Studio (TOOLSDOC)' in give us feedback 4. space between the two conditional phrases in give us feedback needs removing else will have double space before wfk phrase when processed 5. (WFKDOC) needs adding to end of 'Documentation for WFK' phrase Also, by checking in the PressGang handbook I've learnt that when the same text is shared by multiple conditions, then the xml can be simplified. See multiple conditions on an element at the end of [1]. So for example, <phrase condition="jbds;wfk">customer portal</phrase> rather than <phrase condition="jbds">customer portal</phrase><phrase condition="wfk">customer portal</phrase>. Do you know how to preview text filled with conditional statements in PressGang? If not let me know and I can show you. I know that when there's lots of conditions it can be hard to figure out what will show up when. [1] http://docbuilder.usersys.redhat.com/13968/#Add_a_Conditional_Build_Statement	1	1	https://issues.jboss.org/browse/WFKDOC-4
3705	DROOLS-1663	A small clarification needed. Currently, DMN specification supports the following "types" of import: - DMN Import - import of decisions, BKMs, item definitions, etc. - XML Import - import of data tables defined in other XML files - XSD Import - import of XSD type definitions. This is stated in the "DMN Method and Style", B. Silver. Please, check the attached pages. !IMG_2197.jpg|thumbnail! !IMG_2198.jpg|thumbnail! !IMG_2199.jpg|thumbnail! The specification, however, does not mention any specific type (DMN, XML, XSD). Of course, you can find combination of words "XSD import" here and there in the Specs, but, again, as I said, there's nothing specific about what you can exactly import. Should we focus on DMN import only for the "beginning"?	1	0	https://issues.jboss.org/browse/DROOLS-1663
3706	DROOLS-1661	[~ge0ffrey] Is there a way to make these dependencies optional? so if the user is using the kie-server-client only for rules/process/solver he only needs the respective dependencies instead of everything?	0	0	https://issues.jboss.org/browse/DROOLS-1661
3707	DROOLS-1652	What should the URL look like in your opinion [~tirelli]? From a RESTful point of view, we're creating a new kie-container resource, so in theory it should be a PUT (create resource) on URL"/server/container/\{id\}" that accepts a DMN file. However, that URL already exist. We could provide support for different request bodies in that operation (e.g. a `KieContainerResource` and a `DmnKieContainerResource`), but that could become a bit nasty ... We can also use a POST on something like "/server/containers/\{id\}/resources" and have that create a container resource under "/server/containers/\{id\}" (the reason I don't like a PUT in this case is because we create the resource (container) under a different URL than our request URL). The reason for the name "resources" (plural) is that we could provide multiple resources in the body (e.g. multiple DMN files). Second, apart from the DMN XML file (or files), do we also want to be able to post more information in the body (e.g. like additional information we pass in `KieContainerResource`)? And finally, should this be part of the DMN extension or KIE-Server Commons? The argument for DMN would be that otherwise we would expect Commons to have DMN knowledge. On the other hand, it's just deploying a KIE Containerner based on a resource file. The fact that that's a DMN model does not really matter (in theory this could be any resource file). Personally KIE-Server Commons makes most sense.	1	1	https://issues.jboss.org/browse/DROOLS-1652
3708	DROOLS-1652	I don't think we should make this for jBPM as by nature it is stateful so once you create such resource and bounce the server it's gone while the data in db remains - kind of orphans. That in general applies to stateless/in memory stuff too, how would you deal with server restarts? Or is this intended for quick checks of resources? My vote is to include it in DMN extension or Drools if you would like to support other rule related resource types.	1	0	https://issues.jboss.org/browse/DROOLS-1652
3709	DROOLS-1652	I see two approaches to this: - if we consider it as demo only I would opt for having this to be as separate extension - so it can be deployed outside of the official distribution that way we avoid potential issues/bug reports that it won't survive server restarts - make it all the way down to survive the restarts by saving given kjar into maven repo locally and then simply register it as any other container. Not sure what is preferred as I don't know the origin of this request....	1	0	https://issues.jboss.org/browse/DROOLS-1652
3710	DROOLS-1561	Can you please check the container in the kie-server if you haven't enabled the kie-scanner by mistake? There is a configuration there to run the kie-scanner continuously or to deactivate it.	1	1	https://issues.jboss.org/browse/DROOLS-1561

#	Requirement	Questions	Answered	Successfully Answered	URL
3711	DROOLS-1561	[~tirelli] Thanks for your suggestion. The documentation is wrong. Its mentions in seconds. But its actually in millisecs. I have configured as 60000. It scans every minute now. Thats really a fantastic option. Sorry for the bug. Please close it. https://docs.jboss.org/drools/release/6.4.0.CR2/drools-docs/html/ch22.html Automatic Update: If you want a deployed Container to always have the latest version of your deployment without manually editing it, you will need to set the Version property to the value of LATEST and start a Scanner. This will ensure that the deployed provision always contains the latest version. The Scanner can be started just once on demand by clicking the Scan Now button or you can start it in the background with scans happening at a *specified interval (in seconds).*You can also set this value to LATEST when you are first creating this deployment. The Resolved Release Id in this case will show you the actual, latest version number.	1	1	https://issues.jboss.org/browse/DROOLS-1561
3712	DROOLS-1444	Implementing AutoCloseable is only useful if you're planning to use the AutoCloseable resource in a try-with-resources block, but to be honest I don't think it will be common at all to use a KieSession in that way. Do you have a use case where you may need (or want) to do that? If so can you please paste the code snippet where you tried to do this, so I could have a better idea regarding the possible implications?	1	0	https://issues.jboss.org/browse/DROOLS-1444
3713	DROOLS-1444	In that example you're showing how to use try-with-resource with the StringWriter. That makes perfectly sense to me but it is not what I was asking. I was asking for a practical example using the KieSession because in my opinion in that case it is not that useful and using a try-with-resource is more verbose and cumbersome than simply adding a dispose. I cannot add this to drools 6.x because it has to stay Java 6 compatible, but I could do this on master (drools 7) that is Java 8 based. However I really don't see the usefulness in doing so. Moreover your scenario seems to fit far better a StatelessKieSession which doesn't require any dispose. Did you give it a try?	0	0	https://issues.jboss.org/browse/DROOLS-1444
3714	DROOLS-1398	[~Rikkola] sorry, I am confused. Does this change need to be done in the workbench or in core?	1	0	https://issues.jboss.org/browse/DROOLS-1398
3715	DROOLS-1398	Ok. So for this ticket, I was just aiming at full builds and not incremental. Still if a plugin for example generates DRL or changes the code structure in a way that affects the build result. Then the incremental build would have to reflect on that. So turns out this is more complicated than I first thought. On top of this even if the KieBuilder manages the full Maven build. The workbench still has to pick up the KIE version specified in the pom.xml. Since we are now compiling against that. This can get tricky since for example our DMO currently relies on features the core did not have in the earliest 6.x versions.	0	0	https://issues.jboss.org/browse/DROOLS-1398
3716	DROOLS-1398	[~manstis] Nope, as you said we have talked about this before. I even think there is an older ticket about this. Just saw an email in drools-usage where a user edited the pom.xml and then wondered why the editing had no affect on the build. Probably where others also got the idea for the discussion. The more I think about the workbench features for example. The more it feels like having this would be the best thing the users can have. We could for example reuse Maven plugins for project management in the workbench.	0	0	https://issues.jboss.org/browse/DROOLS-1398
3717	DROOLS-1387	[~mfusco], Sorry about not having distilled the test case enough. I'll make sure to be careful about eval use. Btw, I have removed eval from the test case rules (Rule1.drl and Rule2.drl) and attached (brms-perf-comarison-02.zip). The result still tells that ReteOO is 30% faster than PHREAK in average. Am I still overlooking some eval use in the rules? Or could we say "eval" and "accumulate" and "exists" are all slow in PHREAK?	1	1	https://issues.jboss.org/browse/DROOLS-1387
3718	DROOLS-1387	Hi [~tkobayashi]. The resulting diff can be also influenced by thread contention. I see that you have some sync points in the reproducer. Could you please try to run benchmarks from my reproducer? If you use IDEA or Eclipse, there is JMH plugin for these IDEs. If you install it, you can run the benchmarks similarly as JUnit tests. If you have any questions about this, please let me know. You can find me also on IRC on #droolsdev channel. I will also try to rerun the benchmarks. We can compare the numbers then.	1	1	https://issues.jboss.org/browse/DROOLS-1387
3719	DROOLS-1385	Friends,Please comment your ideas	1	1	https://issues.jboss.org/browse/DROOLS-1385
3720	DROOLS-1378	[~tkobayashi] I don't understand why an user may want to put a text into a cell with Numeric type. Is there a valid use for this? In general we should stop trying to parse any sort of garbage thrown in by users. This sort of guesses and type coercions causes many more damages than benefit on the long term and once they are in you cannot get rid of them anymore, otherwise you'll break compatibility.	1	1	https://issues.jboss.org/browse/DROOLS-1378
3721	DROOLS-1378	[~mfusco] Sorry about not explaining the use case well. Please have a look at the attached NumericDisabled.xls. It contains a cell "6600" but it will result in "6600.000000000001" in a generated DRL. The cell is copied from customer's XLS which is generated by VBA macro. Users see "6600" as its String representation regardless of how Excel stores its value behind. Changing "Format Cells" has no effect. Does this make sense that users want to get "6600"? Another example is BasicWorkbook_with_low_values.xls which is tested by SpreadsheetCompilerUnitTest#testProcessSheetForExtremeLowNumbers() https://github.com/droolsjbpm/drools/blob/master/drools-decisiontables/src/test/java/org/drools/decisiontable/SpreadsheetCompilerUnitTest.java#L374-L389 This test intentionally asserts double values (e.g. "3.0E-4"). But this option allows users to get "0.0003" which is the as-is String representation of the cell.	1	1	https://issues.jboss.org/browse/DROOLS-1378
3722	DROOLS-1344	[~kdwivedi], I have talked to Mario about this issue (docs jira linked). We do not have any Aries BP docs in our current guides (upstream nor downstream). If this issue blocks our release, as suggested by [~rsynek], this should be addressed (and incorporated in the downstream docs) by 6.4 release. I am not sure whether Mario has the time to finish this jira on time, maybe prioritization is in order? WDYT?	0	0	https://issues.jboss.org/browse/DROOLS-1344
3723	DROOLS-1323	Here are 2 normal runs and 2 hacked (= without the notifyAll()) call runs on the benchmarker config (not the same as the one that's being profiled, running that one now): [^Selection_246.png] In this case, there is maybe a 1% improvement. Update: ON THE SMALL DATASET.	0	0	https://issues.jboss.org/browse/DROOLS-1323
3724	DROOLS-1319	Hi Edson, Could you please share me some links / sample configs so that i can try out the same in my project ? Thanks, Kavi sri K K	1	0	https://issues.jboss.org/browse/DROOLS-1319
3725	DROOLS-1270	Can you please clarify your request and provide an example?	1	1	https://issues.jboss.org/browse/DROOLS-1270
3726	DROOLS-1250	I fixed the issue by adding _*dialect "mvel" *_ at the two rules in Sample.drl. Is it a bug in the drools wizard that dialet is missing in the examples? or It supposes to work without dialet? Thanks.	1	1	https://issues.jboss.org/browse/DROOLS-1250
3727	DROOLS-1250	I am facing the same issue. But the code runs fine despite compile time error. Also the error message makes no sense. Why does it say com.sample.DroolsTest.Message resolves to a package when it clearly isn't.	0	0	https://issues.jboss.org/browse/DROOLS-1250
3728	DROOLS-1250	Hi Tianwei Chen, just followed your comment and still got the same error. Did you still use droolsjbpm tools 6.4 plugin? I have to add dialect mvel to make it work. Is it an issue of plugin or runtime?	0	0	https://issues.jboss.org/browse/DROOLS-1250
3729	DROOLS-1225	Hi [~kuniyasu], if I'm understanding correctly what you're asking here is adding those new SequentialMaxPriority and SequentialMinPriority optional fields in the "RuleSet" section of the XLS for each sheet and setting the salience of the generated rules accordingly. Can you confirm?	1	1	https://issues.jboss.org/browse/DROOLS-1225

#	Requirement	Questions	Answered	Successfully Answered	URL
3730	DROOLS-1221	Hi Youcef, Thanks for the slide set about the Italian HACEP project. It makes it much more clear. I intend to port the Italian HACEP solution to the Amazon cloud replacing Infinispan, ActiveMQ, Camel Route, etc. with Amazon equivalents. The grouping is done on Tenant (aka customer environment). The domain is Business Event Monitoring in the cloud. Are you going to use Apache Geode in your project as a replacement for Infinispan? If so, why?	1	0	https://issues.jboss.org/browse/DROOLS-1221
3731	DROOLS-1168	ok, but what can be the workaround for the problem? How can I force the execution of update (final FactHandle handle, BitMask mask, Class <?> ModifiedClass). Since you are in the development you could give me some guidance? Best Regards	1	1	https://issues.jboss.org/browse/DROOLS-1168
3732	DROOLS-1168	Did you try what I suggested in my last comment? Something like {code} DroolsUtil.completeEntity($entity, fields, $untMis); $entity.setUntMisVen($entity.getUntMisVen()); update($entity) {code}	1	1	https://issues.jboss.org/browse/DROOLS-1168
3733	DROOLS-1168	May I understand why you're insisting using property reactivity for this use case? Which advantages do you see?	1	1	https://issues.jboss.org/browse/DROOLS-1168
3734	DROOLS-1168	I'm sorry, but property reactivity doesn't fit your use case. Have you tried to disable it and see what happens? I don't see why it shouldn't work.	1	1	https://issues.jboss.org/browse/DROOLS-1168
3735	DROOLS-1137	Note that the RPM work is working around this, so RPM's shouldn't really impact the decision here. I do think it's an improvement nonetheless. In this proposal (let's call it proposal A), what are the changes? - Will it remove 1 POM (because kie-parent-with-dependencies and kie-platform-bom will become kie-third-party-bom)? - Child poms don't change, only top level multiproject parent poms.	1	1	https://issues.jboss.org/browse/DROOLS-1137
3736	DROOLS-1071	To have a line number (of the DRL source file) we'd need to do a similar trick as how JSP's stacktraces are translated back to JSP source file line numbers (thanks to Tomaz for info). {quote} The JSP compiler jastow (fork of tomcat's jasper) does this: https://github.com/undertow-io/jastow what you are looking for is "smap" generation and applying to compiler: https://github.com/undertow-io/jastow/blob/master/src/main/java/org/apache/jasper/compiler/JDTCompiler.java#L464 https://github.com/undertow-io/jastow/blob/master/src/main/java/org/apache/jasper/compiler/SmapUtil.java https://github.com/undertow-io/jastow/blob/master/src/main/java/org/apache/jasper/compiler/SmapGenerator.java {quote}	0	0	https://issues.jboss.org/browse/DROOLS-1071
3737	DROOLS-1071	Can you specify which DRL line exactly caused it? We want to improve the error message in a future version, so it mentions that exact DRL line.	0	0	https://issues.jboss.org/browse/DROOLS-1071
3738	DROOLS-1053	this has been recently added so you can set system property to enable text based date format for JSON marshaller: https://github.com/droolsjbpm/droolsjbpm-integration/blob/master/kie-server-parent/kie-server-api/src/main/java/org/kie/server/api/marshalling/json/JSONMarshaller.java#L66 though it does use certain data format pattern: https://github.com/droolsjbpm/droolsjbpm-integration/blob/master/kie-server-parent/kie-server-api/src/main/java/org/kie/server/api/marshalling/json/JSONMarshaller.java#L73 would this be sufficient or shall we have another system property to control actual date format pattern?	1	1	https://issues.jboss.org/browse/DROOLS-1053
3739	DROOLS-1052	Should I only look in the java files or in the drl files too? What should I do in this case: /** * Finds currently available server ports. * * @author Apache Directory Project * @see IANA.org */ There are a bunch of @authors in test java files, like i.e. public class StandaloneDocBuilder { public static void main(String[] args) throws FileNotFoundException, ParseException { String drl = ""; drl += "# important information\n"; drl += "# about this package\n"; drl += "# it contains some rules\n"; drl += "package org.kie.test\n"; drl += "global java.util.List list\n"; drl += "# Really important information about this rule \n"; drl += "# Another line because one was not enough \n"; drl += "# \n"; drl += "# @author: trikkola \n"; Should this be removed too?	1	1	https://issues.jboss.org/browse/DROOLS-1052
3740	DROOLS-1039	There's already a logged error message like: {code} 2016-04-14 10:36:00,259 [main] ERROR Unable to build kie service url=file:/home/mario/workspace/droolsjbpm/drools/drools-beliefs/target/classes/META-INF/kie.conf reason=[Error: unable to resolve class: org.drools.beliefs.bayes.assembler.BayesAssemblerServicex] [Near : {... assemblers' : [new org.drools.beliefs.bayes.assem ....}] ^ {code} and I'd prefer to keep it instead of throwing and Exception. What do you think?	1	1	https://issues.jboss.org/browse/DROOLS-1039
3741	DROOLS-1008	I am not sure of what you're asking here. I guess you're talking about the maven default behaviour of only fetching SNAPSHOTS once per day, but if this is the case I think you can configure maven to enforce of checking SNAPSHOTs at each request. Please let me know if this is case or try to better clarify this ticket.	1	1	https://issues.jboss.org/browse/DROOLS-1008
3742	DROOLS-1008	[~mfusco] [~andrew.collins] Any status on this. I am able to duplicate this error and have to clear out my local .m2 before the snapshot gets updated in business-central on a deployment. I have the update policy set to always in my settings.xml so I am pretty sure everything is configured correctly My scenario is slightly different but basically the same build and deploy kjar to remote nexus - using a seperate settings.xml for this build and a seprate local reposiotry as to not conflict with what BC is using deploy GAV to BC The deployment is still the old snapshot	1	1	https://issues.jboss.org/browse/DROOLS-1008
3743	DROOLS-972	Can you please send a PR adding a new example in drools-examples-cdi so I could start working from there?	1	1	https://issues.jboss.org/browse/DROOLS-972
3744	DROOLS-963	This should be fixed for 6.4.0.Final, does it sound possible?	1	1	https://issues.jboss.org/browse/DROOLS-963
3745	DROOLS-963	This will be fixed in 6.4.0.Final. It got resolved as part of Maven upgrade from 3.2.2 to 3.2.5.	0	0	https://issues.jboss.org/browse/DROOLS-963
3746	DROOLS-907	As kie-group website is not prepared for rake I have done a "short term solution" adding links to this page. Please look at the image:!kie-group.png|thumbnail! What do you think?	1	1	https://issues.jboss.org/browse/DROOLS-907
3747	DROOLS-905	Thanks for the pointers [~martenscs]. Is there any other prefix we need to worry about? Or does startsWith("bundle") cover all the needed cases?	1	0	https://issues.jboss.org/browse/DROOLS-905
3748	DROOLS-823	Hi Geoffrey, I don't know what shoud be the expected behaviour? i.e. I run in drools jdeps -jdkinternals ./drools-verifier/src/test/resources/org/drools/verifier/model.jar only for test purposes. It says: Picked up JAVA_TOOL_OPTIONS: -Dfile.encoding=UTF-8 ./drools-verifier/src/test/resources/org/drools/verifier/model.jar -> /usr/java/jdk1.8.0_20/jre/lib/rt.jar Hope thta this is OK. When an error is found, what would be the behaviour?	1	1	https://issues.jboss.org/browse/DROOLS-823
3749	DROOLS-823	I think 2) is ok and 1) is ignorable. 3) is definitely a true negative, something that we need to have an inventory of and see if we can reduce that inventory to 0. But those are sun.misc.Unsafe usages from our depended libraries (com.google.common in this case), not from ours ourself. Are there any from us directly (org.drools/org.kie/org.optaplanner/org.jbpm)? We should handle those first.	0	0	https://issues.jboss.org/browse/DROOLS-823
3750	DROOLS-729	Can you come onto irc and speak to Mario (mfusco). This will make it much easier to discuss all the related issues and for us to figure out how to integrate your changes. http://drools.org/community/chat.html	0	0	https://issues.jboss.org/browse/DROOLS-729

#	Requirement	Questions	Answered	Successfully Answered	URL
3751	DROOLS-687	After discussion with Marco we see several ways to implement this: Proposal A) In mvel add a SourceContext object during bootstrap, so when evaluations turn bad, it can include the context in the error message. Note: you might be able to reuse ParserContext directly instead of having to add a SourceContext interface and instance. {code} interface SourceContext { String buildErrorMessagePart(); } class ReflectiveAccessorOptimizer { ... SourceContext sourceContext; public ReflectiveAccessorOptimizer(...., SourceContext sourceContext) { this.sourceContext = sourceContext; } public Object evaluateExpression(...) { ... // Old code // if (curr == null) throw new NullPointerException(); // New code if (curr == null) throw new NullPointerException("NPE happened at " + sourceContext.buildErrorMessagePart()); } {code} Pro: no performance impact! (the if's, catches etc are already there. Only when it's broken it might take a few nanoseconds longer to build the exception, but nobody cares how slow a broken app is) Con: A lot of work (need to be applied throughout mvel, including jitted constraints) Proposal B) try-catch around every mvel calls in drools and somehow figure out the context from drools. This means that the encapsulating drools code needs to know the exact line of every mvel object it uses... Con: performance impact! Try-catches during runtime code, all the time. Con: easy to forgot try-catch when writing new drools code that uses mvel (maintenance prob) Proposal C) Introduce boolean assertMode that if true does B) (so a try catch around every mvel call) and otherwise not Con: User has to be educated about this assertMode Con: 2 paths of execution that need to be kept in sync. Personally, I prefer A) - anything worth doing is worth doing well ;)	0	0	https://issues.jboss.org/browse/DROOLS-687
3752	DROOLS-664	Do we have anything similar available in 6.2? I am looking to filter collections without using from.	0	0	https://issues.jboss.org/browse/DROOLS-664
3753	DROOLS-660	public Class<?> fastFindClass(final String name) { lock.lock(); Class<?> cls = findLoadedClass(name); if (cls == null) { final byte[] clazzBytes = this.store.read(convertClassToResourcePath(name)); if (clazzBytes != null) { String pkgName = name.substring(0, name.lastIndexOf('.')); if (getPackage(pkgName) == null) { definePackage(pkgName, "", "", "", "", "", null); } cls = defineClass(name, clazzBytes, 0, clazzBytes.length, PROTECTION_DOMAIN); } if (cls != null) { resolveClass(cls); } } lock.unlock(); return cls; } it that ok? i find the lock has been removed in 6.2.0.Final, is that means 5.6.0Final has dead lock problem？ can i use 5.6.0 in my project?	1	1	https://issues.jboss.org/browse/DROOLS-660
3754	DROOLS-660	Are you reporting a deadlock problem ??? If so, please change this ticket to bug - blocker AND provide a reproducer. Reading your comment, this is not a bug report, not a feature request, but a question. In this case, please close this ticket and ask your question here https://groups.google.com/forum/#!forum/drools-development (By the way, the classloader architecture was completely redesigned for 6.x, so you can't just compare 5.6 and 6.x. Is there any reason why you would prefer an old, no longer supported version for your project? This would be a discussion for here: https://groups.google.com/forum/#!forum/drools-usage) Thanks	1	0	https://issues.jboss.org/browse/DROOLS-660
3755	DROOLS-660	Could you try using 6.x? The 5.6 series is a community version and no longer maintained, so even with a reproducer it may take a long time to fix it and you would have to build a SNAPSHOT version anyway. If the issue was still present in 6.x, things would be very different Davide	1	0	https://issues.jboss.org/browse/DROOLS-660
3756	DROOLS-656	mbiarnes, would you mind creating a new repo under droolsjbpm, push some stuff to it and then force push some more stuff to it? That will answer the question above. Note: We can't on existing repo's without potentially breaking them which is not an option.	1	1	https://issues.jboss.org/browse/DROOLS-656
3757	DROOLS-652	===== FROM DAVIDE ===== Hi Edson, my quick 2 cents on this proposal Section 2: I like the idea of "defining" rather than just declaring datatypes and events. This problem actually overlaps much with "semantic" classes (actually semantic traits), where, in addition to declaring the type, ancestors and member fields, an "equivalentTo" LHS-like definition(s) may be provided. So I wonder if we could merge the two parts (syntax is tentative here) {code}declare event VoiceCallEvent extends SuperTypeEvent // promote event to first-class ? into(voiceCalls) // either keyword or annotation @into(voicecalls) ? as // equivalent-to definitions callStart : CallStartEvent() -> callEnd : CallEndEvent(callingNumber == callStart.callingNumber) $phone : MobileDevice(number == callStart.callingNumber) $cust : Customer(id == $phone.customerId) ; // fields. in DL would be considered anonymous superclasses callStart : CallStartEvent // ok: initialize when var name coincides callEnd : CallEndEvent customer : Customer = $cust // should we allow explicit initialization using other expressions, possibly involving local vars ? end{code} Section 3 : ok Section 4 : I'd either use infix A() -> B() or prefix notation, but I wouldn't mix the two seq(A() B()), as we do with and/or CEs. Unless the parser requires so :) Question: what about a "strict sequencing" operator A() => B(), meaning "B() after A() and nothing in between"? Question: what about the interval deadlines, like we have in "after[a,b]" ? Section 4 again :) : Another options, maybe uglier? {code}A+(...), B*(....), C[2:4](...){code} Yet another, maybe more explicit? {code}exists[+] A(), exists[*] B(), exists[2:3] C(){code} Section 4.1 Ok! this is again related to what I've been dreaming of writing even in more traditional DRL: {code}$x : exists X(){code} Section 5.5 I'd use "not" ... what would be the high-level difference between "no" and "not" from a conceptual perspective? (not talking about implementation here) Moreover, shouldn't we stick to the "not X(this after $y)" semantics we already have?	1	1	https://issues.jboss.org/browse/DROOLS-652

#	Requirement	Questions	Answered	Successfully Answered	URL
3758	DROOLS-652	Davide, thanks for the feedback. Regarding your comments: h5. Section 2 I considered that as well, but I think it is best to keep them as 2 separate pieces of syntax. The easiest way to see this is that a "declare type" statement is like declaring a class, while an "event" statement is like instantiating that class. While there can be only one "declare type" for each type, there can be multiple "event" statements to instantiate events of the same type based on different criteria. Just a dumb example to illustrate the idea: {code}declare type AccountDebit @role(event) accountId : String ammount : BidDecimal end event AccountDebit as ATMTransaction(type == TransactionType.Withdraw, accountId : accountId, ammount : totalAmmount) end event AccountDebit as InternetTransaction(type == TransactionType.TransferOut, accountId : accountId, ammount : totalAmmount) end {code} h5. Section 4 * I agree that we don't want both a delimiter and the infix -> CE, but I am still not sure the infix will be enough for our needs. Will have to proceed to the design and implementation and revisit this issue when we get there. * The "strict followed-by" (=>) operator remains a possible implementation, but the more I think about it, the less I think it will be used, because in most of the use cases, there aren't any assurances that the events won't be intermixed in the streams coming into the system. Unless we have a compelling use case, I think this is a nice to have, but lower priority feature. I will add it to the spec anyway. Let me know if you know of any specific use case we should consider. * One can still use temporal operators to further constrain sequences of events: {code}$a : A() -> $b : B(this after[0,10s] $a){code} I don't think it would be good to allow, at this point, further parameterization of the -> CE by itself. Basically, what the -> CE defines is a state machine/automaton, and that can be constrained as necessary by the patterns themselves with both temporal and non-temporal constraints. What do you think? * The problem of using the [] after the pattern type is that it become ambiguous with array types: A[*] ---> is this an array of A's with a typo or does it mean matching a repetition of A's? The "exists[+] A()" is feasible, just more verbose... do you think it would be worth writing a keyword there all the time? h5. Section 4.1 Yes, it is a step towards what you asked before... probably 80% of the way, but we would still need a few more things to get where you want... nevertheless, one feature that serves 2 purposes. h5. Section 5.5 The main difference is in semantics: event qualifiers only work with event sequencing. They are not Conditional Elements, like the "not" CE. That is the reason I can't decide between overloading the use of "not" or using a different keyword. In any case, all qualifiers are used in sequences of events, not in regular patterns. I actually need to check if there is any syntax ambiguity that I haven't identified yet.	1	1	https://issues.jboss.org/browse/DROOLS-652
3759	DROOLS-652	(WL2) Section 3 (a) There is no "end timestamp" in the definition of "event" in the "Fusion" manual. Is it the "duration" for the derived event that is computed from the difference between the latest timestamp and the oldest one? (b) What if the contributing events have duration? Then the "end timestamp" of the derived event must be computed by max/over all events(timestamp+duration).	1	1	https://issues.jboss.org/browse/DROOLS-652
3760	DROOLS-652	(WL5) Subsection 5.1, binding of tuples Given the "event" declaration (Section 2): How would this interact with bindings to event tuples selected by wildcards? I see no simple way of permitting such combinations; here explicit event creation (using new) and insertion must be done. Please clarify.	1	1	https://issues.jboss.org/browse/DROOLS-652
3761	DROOLS-652	(WL8) Subsection 6.2 (a) Who is going to want an operator where you don't know the result? (b) Given that "any A() -> any C()" has a "possible (!) result" [A1, C1], why has "every(any A() -> any C())" an exact result - note that "X1 = [A1, C1], X2 = [A4, C3]" is presented as a well-defined (!) result. (And why does this only produce 2 results - there are A1,A2,A3 and C1, C2, C3?)	1	1	https://issues.jboss.org/browse/DROOLS-652
3762	DROOLS-651	The documentation uses the term _binding variable_ 8 times, so bindings seems to be variables too (but of course special ones). I understand that these bindings are not yet cached because of constraint idempotence. However, in my view the implication should be, that especially a non changing result +can+ be cached! Rather than a _good practice_, the documentation says: {quote} "Property accessors +must+ not change the state of the object in a way that may effect the rules. Remember that the rule engine effectively +caches+ the results of its matching in between invocations to make it faster." {quote} So if kind of functional immutability is required and used inside the engine, I would indeed think about caching method calls (with variables) in constraints, as this saves potentially massive processor time and would be much more intuitive. I don't understand your sentence {quote} "If you can't or desire not to guarantee this property, Drools won't enforce it caching the first value returned." {quote} Is it possible, to enforce Drools to cache variable bindings? Drools shouldn't cache values, that change. Let me say that I'm still intrigued about the huge, powerful Drools framework and that we get this as Open Source including this helpful support. Feel free to change the issue-type to a feature request, as the behavior is not strictly incorrect but optimizable.	1	1	https://issues.jboss.org/browse/DROOLS-651
3763	DROOLS-599	Hi [~mbiarnes] , in reference to my previous comment above, and now discussing with [~mfusco] , we would like to highlight the availability of the Kie Archetype to end users. However, [since this archetype was first proposed|https://github.com/droolsjbpm/droolsjbpm-knowledge/pull/74/files#diff-ae073dbf7d2d212e8ef79caab795cb87], now [License file-headers are attached to all files|https://github.com/droolsjbpm/droolsjbpm-knowledge/commit/7172a50c4989800a8cda029b1e5efaec9c1d7dcb#diff-42a73ec3fc7beda342f8c8ad2bddec31]; with the actual consequence that once you use the archetype, as a user, your files are containing the Red Hat file-header license; and I believe this could be very confusing for the end-users, and also I believe is wrong. My proposal is to remove License file-headers, at least from all files which are serving as a template-generation. I have two question, please? # Do you agree with me reverting your change? # Did you notice the License file-headers are missing thanks to some automatic tool? Is there a way to manage exclusion, either in the tool or by adding a special header to the file itself? Thank you in advance for your feedback.	0	0	https://issues.jboss.org/browse/DROOLS-599
3764	DROOLS-566	Thanks, what about the other way around, eg DROOLS SNAPSHOT with mvel 2.1.9 - does that work?	1	1	https://issues.jboss.org/browse/DROOLS-566
3765	DROOLS-552	I double checked that all the tests in your pull request run successfully and merged it. I am not sure of what we could do to better support other encodings. Any suggestion?	1	1	https://issues.jboss.org/browse/DROOLS-552

#	Requirement	Questions	Answered	Successfully Answered	URL
3766	DROOLS-552	Hi Mario, I have found why those tests fail with -Dfile.encoding=MS932. DrlParser.parse() uses resource.getInputStream() + resource.getEncoding() instead of resource.getReader() which commit#c3a90d5 fixed. resource.getEncoding() returns null and then it uses platform encoding. https://github.com/droolsjbpm/drools/blob/master/drools-compiler/src/main/java/org/drools/compiler/compiler/DrlParser.java#L150 If our direction is "Use UTF-8 if encoding is not specified" then, we should explicitly set "UTF-8" to 'encoding' in Resource class constructors. However, then, there is another concern. If users have Drools applications which depend on platform encoding (= expecting resources are decoded with platform encoding other than UTF-8), they will face that their applications are broken. To keep backward compatibility, one idea is a system property switch. Here I introduce "drools.default.encoding.utf8.enabled" system property (by default, it's true). If it's true, default encoding would be "UTF-8". If it's false, default encoding would be platform encoding (= "file.encoding"). So if users want backward compatibility in Drools 6.2, they would need to set -Ddrools.default.encoding.utf8.enabled=false. This is the commit in my work branch. Now test cases run fine with both -Dfile.encoding=UTF-8 and MS932. Could you review and let me know your thought? This commit fixes only ClassPathResource and ByteArrayResource. To make a pull request, I would need to fix more many classes. https://github.com/tkobayas/drools/commit/162fabbe0a614dbcf9783c1e10c1aa13b45de81f	0	0	https://issues.jboss.org/browse/DROOLS-552
3767	DROOLS-529	Could you please submit the rules you used to benchmark the engine? Thanks!	0	0	https://issues.jboss.org/browse/DROOLS-529
3768	DROOLS-529	We'd like to try and fix this for 6.1, which is due within 2 weeks. But to do so we need some benchmarks, can you upload a zip for us?	1	1	https://issues.jboss.org/browse/DROOLS-529
3769	DROOLS-529	Give me some time (one day tops) and I'll clean some organizational IP from the project and send you a maven project with some instructions on how we did the profiling. Thanks for the quick consideration.	1	1	https://issues.jboss.org/browse/DROOLS-529
3770	DROOLS-507	Is there any update on this? i'm facing the same issue..	0	0	https://issues.jboss.org/browse/DROOLS-507
3771	DROOLS-507	The FireUntilHalt command is marked as unpersistable (for obvious reasons - it is usually executed in a dedicated thread), which causes the exception even before the command is attempted. Could you elaborate more on your use cases? In particular, what would you expect to happen if a running session was persisted and then restored?	0	0	https://issues.jboss.org/browse/DROOLS-507
3772	DROOLS-492	[~cmoulliard] Great! If the exception isn't being thrown anymore, it sounds like this issue (the exception, old class, etc) has been fixed. Does it make sense to assign the issue to you and close it?	0	0	https://issues.jboss.org/browse/DROOLS-492
3773	DROOLS-483	Thanks for reporting this, can you submit a PR? It will be easier to discuss the details Thanks! Davide	1	1	https://issues.jboss.org/browse/DROOLS-483
3774	DROOLS-483	Hi there. It's been more than 3 months and neither this issue nor the pull request has seen any kind of response, so I was wondering if something important is missing (other than time of the Drools team side). Should I continue rebasing the modification on master from time to time?	1	1	https://issues.jboss.org/browse/DROOLS-483
3775	DROOLS-424	The configuration part of this would be easy. e.g. in the xsd: <xsd:attribute name="type" use="required" type="knowledgeSessionTypeEnum"/> could be replaced with: <xsd:attribute name="type" use="optional" type="knowledgeSessionTypeEnum"/> <xsd:attribute name="scope" use="optional" type="knowledgeSessionScopeEnum"/> <xsd:simpleType name="knowledgeSessionScopeEnum"> <xsd:restriction base="xsd:string"> <xsd:enumeration value="request" /> <xsd:enumeration value="session" /> <xsd:enumeration value="application" /> </xsd:restriction> </xsd:simpleType> and then in KnowledgeSessionDefinitionParser, scope="request" would be equivalent to type="stateless", and scope="application" would be equivalent to type="stateful". The new option would be scope="session". However I'm not sure when the StatefulKnowledgeSession would actually be created as it requires access to the HTTP session to store it for future requests. Does the HTTP request get passed all the way through from Camel to the Drools classes? I tried tracing it through but lost it somewhere in the depths of Camel.	0	0	https://issues.jboss.org/browse/DROOLS-424
3776	DROOLS-365	I see what you are trying to do, but I am not sure that such an attribute would be feasible, or even solve your problem. What happens if a rule inserts (and modifies, and retracts...) more than one object? And what happens if the same object is insert by more than one rule? Have you tried using the EQUALITY assertion behavior and/or the @propertyReactive annotation? The former considers equivalent (i.e. equals()) object to be the "same" and thus do not activate rules multiple times, regardless of the rule that generated the objects. The latter allows you to control how rules fire in case of updates. Best Davide	1	1	https://issues.jboss.org/browse/DROOLS-365
3777	DROOLS-365	I don't see the problem with modifications to multiple facts: you could condition on the NOT of the logical AND of all facts. Or the logical OR; I think both approaches are defensible. As for your second example, what's the issue with multiple rules asserting the same fact? In any case, I think the biggest challenge to implementation is that the rule consequence is a snippet of Java, and thus can do absolutely anything. For example, it could conditionally assert a fact. How do you modify the rule then? Also conditionally? Based on runtime input? It seems really challenging. The example that actually drove me to file this bug looks something like this: {noformat} rule "1" when accumulate (Foo($type : type, $value : value), $total : sum($value)) then insert(new Bar($type, $total)); end rule "handle all Bars of type first" when $bar : Bar(type == 'first', $total : total) ... then retract($bar); end rule "handle all Bars of type second" when $bar : Bar(type == 'second', $total : total) ... then retract($bar); end rule "generic rule to handle all leftover Bars" salience -1 // let custom Bar rules fire first when $bar : Bar($total : total) ... then retract($bar); end {noformat} The flow here is to convert the sum of all Foos (of a given type) into a Bar. Subsequently, there's a rule to handle each type of Bar, with a generic rule to handle Bars without a custom rule. To get that fallthrough behavior I need: # Salience -1 on the generic rule, so it fires last. # Retraction of each Bar as it's processed. However, the retraction causes the summation rule to reactivate, unless I add an explicit {{not Bar(type == $type)}} condition to the rule. Granted, in this case, I suppose the feature request mentioned here wouldn't work, because it'd add {{not Bar(type == $type, total == $total)}}, which would allow the rule to reactivate. For this example, I don't think equality-based assertion will solve my problem; fundamentally the retraction of a Bar causes us to reevaluate the accumulate rule. Likewise, I don't think @propertyReactive will help either; the issue here isn't with modifying a single property, but with removing an entire fact.	1	0	https://issues.jboss.org/browse/DROOLS-365

#	Requirement	Questions	Answered	Successfully Answered	URL
3778	DROOLS-355	These days, you can check it like this: {code} <build> <plugins> <plugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-jdeps-plugin</artifactId> <version>3.0.0</version> <executions> <execution> <goals> <goal>jdkinternals</goal> <goal>test-jdkinternals</goal> </goals> </execution> </executions> <configuration> <dependenciesToAnalyzeIncludes> <include>*:*</include> </dependenciesToAnalyzeIncludes> </configuration> </plugin> {code} which gives: {code} INFO] --- maven-jdeps-plugin:3.0.0:jdkinternals (default) @ drools-compiler --- [INFO] *org.kie.internal.builder* (kie-internal-7.0.0-SNAPSHOT.jar) -> com.sun.tools.xjc not found *org.kie.internal.builder.help* (kie-internal-7.0.0-SNAPSHOT.jar) -> com.sun.tools.xjc not found *org.drools.core.builder.conf.impl* (drools-core-7.0.0-SNAPSHOT.jar) -> com.sun.tools.xjc not found *org.drools.compiler.builder.impl* (drools-compiler-7.0.0-SNAPSHOT.jar) -> com.sun.tools.xjc not found *org.drools.compiler.runtime.pipeline.impl* (drools-compiler-7.0.0-SNAPSHOT.jar) -> com.sun.codemodel not found -> com.sun.tools.xjc not found -> com.sun.tools.xjc.model not found -> com.sun.tools.xjc.outline not found {code}	0	0	https://issues.jboss.org/browse/DROOLS-355
3779	DROOLS-355	Jigsaw is now part of JDK 9. Raising priority of this issue.	0	0	https://issues.jboss.org/browse/DROOLS-355
3780	DROOLS-355	Hi, I'm Henry from Red Hat IoT engineering team. We are working on putting Drools on IoT gateway devices. I'd like to use kie-server-client and install it into Eclipse Kura server (which is really OSGi container based on Eclipse Equinox). I'm running into the same problem as drools-core is a dependency of kie-server-client. Do you have an estimate of when this issue could be resolved? Cheers!	1	1	https://issues.jboss.org/browse/DROOLS-355
3781	DROOLS-355	Hello [~hkonsek]. Would it possible to attach a reproducer which shows the error you are encountering? It would help us a lot. Thanks!	1	1	https://issues.jboss.org/browse/DROOLS-355
3782	DROOLS-355	I did test and merge your pull request. And that part works. However I am now stuck at the next bundle: {code} org.osgi.framework.BundleException: The bundle "org.kie.remote.client_6.4.0.201604141724 [26]" could not be resolved. Reason: Missing Constraint: Import-Package: org.kie.remote.services.ws.common; version="0.0.0" {code} However the JAR {{kie-remote-ws-common-6.4.0.Final.jar}} is not an OSGi bundle. I updated the repository, so you can test if you like. Should I open a new issue?	1	1	https://issues.jboss.org/browse/DROOLS-355
3783	DROOLS-355	And which is not available on Equinox ;-) So should I open a new issue for that? Or can you do that. I guess you can provide better information than I can.	1	1	https://issues.jboss.org/browse/DROOLS-355
3784	DROOLS-337	Ok, you are right for the backward compatibility. ==> ena"B"led is fine. What should i do for the next step ? I didn't understood. Should i select under the workflow button, the "Link Pull Request" choice ?	1	1	https://issues.jboss.org/browse/DROOLS-337
3785	DROOLS-315	Edson, I 've been running into this issue a lot (with scary overflow or rounding errors such as 0.09 + 0.01 != 0.10). Something needs to be done. I see 3 ways: A) Stuff drools-planner with a kit functions, such as sumInteger, sumLong, sumDouble, ... I'd even recommend people in the manual not to ever use sum. B) Stuff drools-compiler itself with those functions. C) Fail-fast overloading of accumulate functions: Only the following sum functions are available: sum(Integer), sum(Double), sum(Long) so when "the result type of the expression inside the function parenthesis" is not Integer, Double or Long a compiler error is thrown. This presumes that mvel is mostly hard typed. If we wouldn't have to worry about backwards compatibility, I believe C) would be the best solution: Using the wrong type to sum leads to bugs, so we can't afford to guess and just go with double, because double only works for double and int: it does NOT work reliably for long and BigDecimal. Wdyt?	0	0	https://issues.jboss.org/browse/DROOLS-315
3786	DROOLS-315	Hey Edson, What's the status of this issue? Do we have sumInt(), sumLong(), sumBigDecimal() available in Drools? Or better yet, support for overloaded methods support, so sum(int), sum(long), sum(BigDecimal) just do the right thing (instead of using sum(double)).	0	0	https://issues.jboss.org/browse/DROOLS-315
3787	DROOLS-300	Issue #2 In our discussions, we have remarked that drools is not real-time, so the scheduled time is a lower bound for the actual activation: i.e., expect the rule to activate any time AFTER the scheduled time. Should we have an activation-specific deadline? E.g., activate with a period of 1 day, but cancel any activation that can't be fired within 1 hour since the activation. If you miss the 1hr-window, just skip and go for the next day	0	0	https://issues.jboss.org/browse/DROOLS-300
3788	DROOLS-299	Do you mean I need to add the org.drools:drools-jboss-integration dependency to the project itself? If I do this it adds over 80mb of dependencies to the deployment.	0	0	https://issues.jboss.org/browse/DROOLS-299
3789	DROOLS-298	This issue is actually a bit more severe than it seems: 1) When "user.home" property is not set (for whatever reason), MavenRepository may fail with an obscure exception: java.lang.NoClassDefFoundError: Could not initialize class org.kie.scanner.MavenRepository at org.kie.scanner.ArtifactResolver.getMavenProjectForGAV(ArtifactResolver.java:78) 2) You will get the same exception when the POM you're trying to parse cannot be parsed. (Or when some of the dependencies cannot be located.) This comes from the way how .embedder.MavenProjectLoader parses the POM - when the parsing fails in the static inner class, you won't find out why (the POM error), you just get the above. 3) You should be able to provide a settings.xml to the parser; since that is (AFAIK) the only way to tell Maven that you have additional repositories with additional dependencies. (Think in-house applications.)	1	1	https://issues.jboss.org/browse/DROOLS-298
3790	DROOLS-297	[~mbiarnes] On the ip-bom list, Keith and Gary we're wondering how to do this. How did you fix this for us? We 'll want to extract that to the ip-bom, so the other projects benefit from it too.	1	0	https://issues.jboss.org/browse/DROOLS-297
3791	DROOLS-297	Does this need to be an extra profile? Can't the enforcer rule be smart enough to do the following "if the version of this module is not a SNAPSHOT, then it should have no SNAPSHOT dependencies"?	1	0	https://issues.jboss.org/browse/DROOLS-297
3792	DROOLS-292	Edson Tirelli <etirelli@redhat.com> made a comment on [bug 1016727|https://bugzilla.redhat.com/show_bug.cgi?id=1016727] Mark fixed this one: 6.0.x: https://github.com/droolsjbpm/drools/commit/a550d5ea54c8e816a5bade7d87d8e65e65ac37a2 master: https://github.com/droolsjbpm/drools/commit/be1444020f71277207743d708238c05253b494a9	0	0	https://issues.jboss.org/browse/DROOLS-292
3793	DROOLS-290	Constraints are handled differently in 5.5 than 5.0. In later versions, you can have "free-form" constraints as opposed to simpler <Field OP Value>. This is not a bug. Are you manipulating the Descr tree yourself? Those classes have never been part of the public API, as far as I know	0	0	https://issues.jboss.org/browse/DROOLS-290

#	Requirement	Questions	Answered	Successfully Answered	URL
3794	DROOLS-266	expected output: --- T E S T S --- Running org.overlord.dtgov.dtgovworkflows.KieTest Reading your .m2 repo artifact=org.overlord.dtgov.issue266:dtgov-workflows:jar:2.0.0-SNAPSHOT 0 [main] WARN org.kie.scanner.MavenClassLoaderResolver - Dependency artifact not found for: org.jbpm:jbpm-bpmn2:${jbpm.version} 413 [main] WARN org.kie.scanner.MavenClassLoaderResolver - Dependency artifact not found for: org.jbpm:jbpm-shared-services:${jbpm.version} 753 [main] WARN org.kie.scanner.MavenClassLoaderResolver - Dependency artifact not found for: org.jbpm:jbpm-persistence-jpa:${jbpm.version} 1097 [main] WARN org.kie.scanner.MavenClassLoaderResolver - Dependency artifact not found for: org.jbpm:jbpm-runtime-manager:${jbpm.version} 1416 [main] WARN org.kie.scanner.MavenClassLoaderResolver - Dependency artifact not found for: org.drools:drools-core:${drools.version} 1751 [main] WARN org.kie.scanner.MavenClassLoaderResolver - Dependency artifact not found for: org.drools:drools-compiler:${drools.version} 2108 [main] WARN org.kie.scanner.MavenClassLoaderResolver - Dependency artifact not found for: org.kie:kie-api:${drools.version} 2458 [main] WARN org.kie.scanner.MavenClassLoaderResolver - Dependency artifact not found for: org.jboss.netty:netty:${drools.version} java.lang.NullPointerException at sun.net.util.URLUtil.urlNoFragString(URLUtil.java:47) at sun.misc.URLClassPath.getLoader(URLClassPath.java:321) at sun.misc.URLClassPath.access$000(URLClassPath.java:62) at sun.misc.URLClassPath$1.next(URLClassPath.java:221) at sun.misc.URLClassPath$1.hasMoreElements(URLClassPath.java:232) at java.net.URLClassLoader$3$1.run(URLClassLoader.java:583) at java.net.URLClassLoader$3$1.run(URLClassLoader.java:581)	1	1	https://issues.jboss.org/browse/DROOLS-266
3795	DROOLS-166	With reference to: https://groups.google.com/d/msg/drools-usage/fbF8KhOPCKM/D9pCQvLdfIEJ Feedback by Edson Tirelli: {quote} As we moved to a minimal public API, the getConfiguration() method was moved to the internal API, but I think this can be reconsidered. May I ask you please to open a JIRA requesting this method be moved to the public API? Meanwhile, as a workaround, you will need to cast it to the internal API: InternalKnowledgeBase.getConfiguration(). {quote} As I understand what should be the proper solution based on Edson's feedback, I also was able to discover this bug on JIRA; so I'm going to re-use this record instead of creating a new one. I do confirm that the workaround in the meantime works for me, thank you very much Edson. @Petr, for KieSession you can do .getSessionConfiguration() of the public API interface, works for me at least.	1	1	https://issues.jboss.org/browse/DROOLS-166
3796	DROOLS-147	Hi,thanks for the reply. Couple of things i want to point out, 1) For the attached code which i had sent , i had already added the jBPM 5.4 dependency. 2) In the attached code i have commented the part wherein i am using knowledge builders ,knowledgepackages and stateless knowledge sessions. The code works fine for te same scenario when i use the above mentioned classes. The problem is with the usage of of PackageBuilders and that too in a peculiar way in which the rule flow is called. Initially we are calling a rule and from the rule we are invoking the rule flow using the drools.getWorkingMemory().startProcess(<rule flow id>) Kindly refer the drl and rf files attached. In the above scenario the first rule is executed successfully but the execution of 2nd rule fails(ie the call to rule flow).But the exact case works fine with the knowledge builders. Is it possible to provide a workaround / patch for the case to work with packagebuilders?	1	0	https://issues.jboss.org/browse/DROOLS-147
3797	DROOLS-146	This has worked well for months on the optaplanner repository. Michael, 1) All the other blessed repositories should now also get that .gitattributes file (copy it for the optaplanner repo). Possibly improve it first before copying. 2) Clean up all previous errors, as stated in the important above. Shells scripts like dos2unix etc can help here. All java files etc were dos2unixed 2 years ago, but we've seen many contributions since then ...	0	0	https://issues.jboss.org/browse/DROOLS-146
3798	DROOLS-113	Is some work being done on this feature?	1	0	https://issues.jboss.org/browse/DROOLS-113
3799	DROOLS-49	So what do you suggest to do ?	1	1	https://issues.jboss.org/browse/DROOLS-49
3800	DROOLS-49	Turns out the pull requests are versus 5.5.x, not master, which explains the <manifestFile>META-INF/MANIFEST.MF</manifestFile> stuff. It's unclear to which extent the changes on master already fix or don't fix this problem.	0	0	https://issues.jboss.org/browse/DROOLS-49
3801	DTGOV-215	The classifiers are directly related to the ontologies. For example "http://www.jboss.org/overlord/deployment-status.owl#InProd" Then what I am thinking for create an interceptor that reads when an ontology is inserted/updated/deleted from s-ramp and update the UIConfiguration List accordingly. I think this can be done better when the JMS S-ramp integration is working in DTGov. Then could be a subscriber for a topic that only check if the item is an ontology. What do you think Eric?	0	0	https://issues.jboss.org/browse/DTGOV-215
3802	DTGOV-174	Would you be willing to take a look or would you advise we report the issue to jbpm straight away?	1	1	https://issues.jboss.org/browse/DTGOV-174
3803	DTGOV-174	OK I just got around to trying this. I used EAP 6.3 to test and was unable to reproduce the behavior. [~virchete] - can you confirm or provide additional details?	1	1	https://issues.jboss.org/browse/DTGOV-174
3804	DTGOV-147	Can this be closed?	0	0	https://issues.jboss.org/browse/DTGOV-147
3805	DTGOV-126	David can you add more context to this jira? How can it be reproduced?	0	0	https://issues.jboss.org/browse/DTGOV-126
3806	DTGOV-112	Want me to take care of this during SRAMP-382?	1	1	https://issues.jboss.org/browse/DTGOV-112
3807	DTGOV-76	Should the installer put the modules in the soa layer if installing into a SY EAP?	1	1	https://issues.jboss.org/browse/DTGOV-76
3808	DTGOV-65	Ultimately what we need is a plugin approach to deployers. We need a new interface (e.g. Deployer) which encapsulates the logic of actually sending a binary deployment artifact to the target. We then need to modify the deployment REST service impl to use that Deployer interface, moving the specific logic out of the REST impl and into contributed implementations of Deployer. We likely also need a DeployerProvider interface which can be contributed at runtime. Maybe have a look at the following classes in the s-ramp repository for inspiration: DeriverProvider ArtifactDeriverFactory ArtifactDeriver	0	0	https://issues.jboss.org/browse/DTGOV-65
3809	CHE-286	[~jastrachan] we are moving all issues from jira to github. Do you happen to know if there any related issue opened in che / rh-che ? if no I will create one	0	0	https://issues.jboss.org/browse/CHE-286
3810	CHE-285	Hi [~mloriedo] I am still facing this issue : https://drive.google.com/a/redhat.com/file/d/0B1qUq_laG6COX1BGemg4V0F5ekE/view?usp=sharing can you please update !	1	0	https://issues.jboss.org/browse/CHE-285
3811	CHE-257	Could you elaborate what do you mean here? Does it mean to verify that a workspace created from a stack definition can start?	0	0	https://issues.jboss.org/browse/CHE-257
3812	CHE-257	Workspace started is part of it but not enough. I think we need to have a test framework where we can define the image to pull and tests to run on. The tests does not necessarily need to be running a che workspace but rather can check if the created container from the image has the correct items to run a workspace. Here is an example of things to check off top of my mind * Check JAVA_HOME set and point to a JAVA SDK, not JRE * Check JAVA version * Is MAVEN_HOME set * Is correct maven version installed * Check node.js version	0	0	https://issues.jboss.org/browse/CHE-257

#	Requirement	Questions	Answered	Successfully Answered	URL
3813	CHE-257	Even I would have a test to check stacks.json, it would not be enough. Env vars are specified in Dockerfiles of stacks. Checking env vars manually in bash script could not help because env var could still point to a symlink not to a folder with SDK, or it could be a pointing to a folder which could be called sdk but it still could not contain sdk just only jre. Our stack images are not build from scratch, but they are using upstream docker images. So these tests would have to run all the time, because of possible changes in upstream. And if something breaks, we would have to fix also upstream image. I think the first step would be to abstract our images (for rh che distribution) from upstream images to have full control over stack images. Again, as I wrote in previous comment, we can run a workspace on Che an drun EE tests on each of them. Because EE tests uses maven and maven for building requires properly set JAVA_HOME pointing to a JDK instead of JRE, it would fail and we will know that something is wrong. And it still have to be responsibility of stack author that a stack is working when it is submitted. We then wouldnt end up with issues such as CHE-213 which I have reported some time ago, when I guess the stack author did not checked/verified that the stack is actually running on Che. What I would suggest is - have a one stack for EE tests in verification job of Che. Have another tests, which could reuse code of EE tests after few changes (making them parameterized) and run those tests on changes performed to RH Che stacks/stack images repos to verify it is working. Before creating such parameterized suite for workspaces, we have to build our stack images from scratch to not be dependant on upstream images.	0	0	https://issues.jboss.org/browse/CHE-257
3814	CHE-243	Can't we just print instructions on how to get the logs of different commands? Something like: {code:bash} echo " Workspaces don't log to stdout. Instead agents log files can be fond inside the container: - wsagent logs open are in folder /home/user/che/ws-agent/logs - ls-bayesian logs can be found in /home/user/che/bayesian.log (c.f. https://github.com/fabric8-analytics/fabric8-analytics-lsp-server/pull/9) " {code} The PR sent by [~snjeza] is brilliant but has 2 drawbacks: 1. The code is error prone and it's run in container pid 1: if, for whatever reason, logging to stdout fails the workspace will crash 2. Besides wsagent there are still other agents logging (i.e. bayesian)	1	1	https://issues.jboss.org/browse/CHE-243
3815	CHE-231	Do we need to put some requirements to SPI work?	0	0	https://issues.jboss.org/browse/CHE-231
3816	CHE-230	[~amisevsk] It seems that we expect "Server-" to be there https://github.com/eclipse/che/blob/openshift-connector/plugins/plugin-docker/che-plugin-docker-machine/src/test/java/org/eclipse/che/plugin/docker/machine/ServerEvaluationStrategyTest.java#L135 ? Do we change this Server name evaluation strategy or the route creation strategy ? Somehow, these 2 strategies need to be the same IMO	1	1	https://issues.jboss.org/browse/CHE-230
3817	CHE-229	what has been done in https://github.com/eclipse/che/pull/5052 (compared to openshift-connector branch) * Skipped code actions (5 commits) already in master * Skipped Java Testing support Bug fixes and Enhancements - Step 2 already in master * Squashed 7f19381d813910897d48ce85e01f01ee69d39585 : Revert WorkspaceFilesCleaner binding (f3a9123be9f3e963635392a70cd4e01628913633) * Squashed Temporarly comment out usage of PV for workspaces and Rework PVC * rebase of course and some cleanup what does https://github.com/eclipse/che/tree/openshift-connector-wip include: * Add secret to workspaces for bayesian plugin * Set github token what still to do in https://github.com/eclipse/che/tree/openshift-connector-wip * -cherrypick Register ls by pattern 2 (#4846)- not anymore needed thanks to https://github.com/eclipse/che/pull/5156 what still need to do to merge https://github.com/eclipse/che/tree/openshift-connector-wip * *Temporary* fix for https://issues.jboss.org/browse/CHE-214 (will be fix by codenvy refactoring https://github.com/eclipse/che/issues/4679) * Add secret to workspaces for bayesian plugin -> should be done in rh-che using this. https://github.com/l0rd/che/commit/65030292589b48 (override injection) + create a ticket * set github token proper way that Sergii will accept -> CHE-262 should solve it * Rework Register ls by pattern 2 (#4846) so it is mergeable -> https://github.com/eclipse/che/pull/5156 * https://github.com/redhat-developer/rh-che/pull/54 stacks openshift.io * transplant https://github.com/eclipse/che/pull/4961 to rh-che stacks.json * https://github.com/eclipse/che/pull/5157 to rebase https://github.com/eclipse/che/pull/5191	1	1	https://issues.jboss.org/browse/CHE-229
3818	CHE-224	[~dfestal] > which is planned afaik What is the Jira issue ?	1	0	https://issues.jboss.org/browse/CHE-224
3819	CHE-224	What is the Jira issue ? It's possible that there isn't a Jira issue for now. But the idea is still that the CI build files will be moved to the `rh-che` repository. [~mloriedo] should we create an issue for this ?	0	0	https://issues.jboss.org/browse/CHE-224
3820	CHE-217	[~gercan] I was told by Ilya, that stacks are gonna be somehow resolved by pipelines. Could you elaborate on this topic? Reopening for clarification.	0	0	https://issues.jboss.org/browse/CHE-217

#	Requirement	Questions	Answered	Successfully Answered	URL
3821	CHE-180	[~snjeza] it would be easy to revert the https://github.com/eclipse/che/pull/4799, but I have a few comments inlined before we make an agreement of doing this: {quote} when a pod doesn't contain any logs, the OpenShift connector tries to read from a stream that doesn't exist which causes the "Pipe not connected" exception {quote} {quote} debug the che server, PipedInputStream.java, line 305 and start/stop a workspace. "Pipe not connected" isn't fixed, but only hidden. {quote} Indeed, pipe is not connected will be thrown, but it is done by design and not coupled with this issue (there are *no* logs of "Pipe not connected" in che-sever pod when the exception is thrown) - https://github.com/eclipse/che/blob/openshift-connector/plugins/plugin-docker/che-plugin-openshift-client/src/main/java/org/eclipse/che/plugin/openshift/client/OpenShiftConnector.java#L763 This issue is *not* coupled with the fact that that the pod logs are empty, because we do see logs in the che-server pod with "Pipe not connected" and after that che-server become unusable - not possible to create a workspace. {quote} getContainerLogs doesn't close a log watch the getEvents method doesn't close an event watch {quote} Yeah, agree - those issues need to be fixed, but are you sure that closing watches fix the main critical issue - after some time ~ 5 -10 mins it is simply not possible to create a worskpace {quote} https://github.com/eclipse/che/pull/4799 doesn't introduce a new issue, but doesn't fix anything and makes fixing the connection leak more difficult. {quote} In general closing client is what is done in all fabric8 examples e.g. - https://github.com/fabric8io/kubernetes-client/blob/master/kubernetes-examples/src/main/java/io/fabric8/openshift/examples/DeploymentConfigExamples.java#L96 OpenShift client should be treated as DB connection IMO - it implements Autocloseabe interface and must be closed once the operation on cluster is finished (the good think that in comparison with DB connection, creating OpenShift client is a "cheap" operation) One thing I notice, if OpenShift connector is created in method and not closed there will be warning - Resource leak 'openShiftClient' is never closed. Also, I do not understand how closing client in each method can actually "hide" any issue e.g what is difference between: {code} LogWatch watchLog = this.openShiftClient.pods().inNamespace(openShiftCheProjectName).withName(podName) .watchLog() ... LogWatch watchLog = new DefaultOpenShiftClient.pods().inNamespace(openShiftCheProjectName).withName(podName) .watchLog() {code} I believe that opening client and not closing it at all, as was implemented previously, might lead to some unpredictable consequences (like this issue). {quote} Some methods (getContainerLogs) require a Kubernetes connection as long as a workspace runs so we can't close the client. {quote} Why ? maybe I do not fully understand the code, but there is a while loop for pending logs (after it's execution I see no reason of having client opened) - https://github.com/eclipse/che/blob/openshift-connector/plugins/plugin-docker/che-plugin-openshift-client/src/main/java/org/eclipse/che/plugin/openshift/client/OpenShiftConnector.java#L756	1	1	https://issues.jboss.org/browse/CHE-180
3822	CHE-180	{quote} Indeed, pipe is not connected will be thrown, but it is done by design and not coupled with this issue (there are no logs of "Pipe not connected" in che-sever pod when the exception is thrown) - https://github.com/eclipse/che/blob/openshift-connector/plugins/plugin-docker/che-plugin-openshift-client/src/main/java/org/eclipse/che/plugin/openshift/client/OpenShiftConnector.java#L763 {quote} That is, probably, a bug in the Kubernetes client. {quote} OpenShift client should be treated as DB connection IMO - it implements Autocloseabe interface and must be closed once the operation on cluster is finished {quote} Agree. I have tested opening/closing the client. The first time it takes about 1,5 seconds and 30-50 ms every time after that. {quote} This issue is not coupled with the fact that that the pod logs are empty, because we do see logs in the che-server pod with "Pipe not connected" and after that che-server become unusable - not possible to create a workspace. Yeah, agree - those issues need to be fixed, but are you sure that closing watches fix the main critical issue - after some time ~ 5 -10 mins it is simply not possible to create a worskpace {quote} I can reproduce it with the current openshift branch, but can't with the snjeza/che-server:che180fixed image. {quote} Also, I do not understand how closing client in each method can actually "hide" any issue e.g what is difference between: {quote} Closing the client doesn't hide the issue. It is hidden in the getContainerLogs method (OpenShiftConnector.java#L763). {quote} Why ? maybe I do not fully understand the code, but there is a while loop for pending logs (after it's execution I see no reason of having client opened) - https://github.com/eclipse/che/blob/openshift-connector/plugins/plugin-docker/che-plugin-openshift-client/src/main/java/org/eclipse/che/plugin/openshift/client/OpenShiftConnector.java#L756 {quote} Che creates a separate thread for a container log. This thread runs as long as a workspace runs. If we close the client, we will get a connection leak. I have tried to read a container log synchronously, but haven't succeeded to find a solution that would work correctly. Maybe there is a bug in the kubernetes client.	1	1	https://issues.jboss.org/browse/CHE-180
3823	CHE-180	[~snjeza], yeah it looks like client has issues with watchers :/ {quote} I can reproduce it with the current openshift branch, but can't with the snjeza/che-server:che180fixed image. {quote} Could you please provide steps for reproducing this issue against the latest 1.0.91 [1] template, which uses docker image *rhche/che-server:e271ace* ? The only reason why PR https://github.com/eclipse/che/pull/4799 was merged is because neither I, nor reviewers were not able to reproduce this issue on minishift. I still can not reproduce it - creating / removing workspaces on minishift against different stacks does not break che server for me after a couple of hours of work {quote} Closing the client doesn't hide the issue. It is hidden in the getContainerLogs method (OpenShiftConnector.java#L763). {quote} PR https://github.com/eclipse/che/pull/4799 implements the *only* thing - creating / closing client in each method. The logic of OpenShiftConnector methods was not changed at all (at this point hiding getContainersLogs was done before this PR was merged, and it did not help - the issue was still reproducible). I do not insist on the fact that this PR fix anything and just fall back on two assumptions: - it is incorrect to open client and never close it (as it was done before) - after closing client correctly in each method the issue is not reproducible (for me) [1] http://central.maven.org/maven2/io/fabric8/online/apps/che/1.0.91/	1	1	https://issues.jboss.org/browse/CHE-180
3824	CHE-180	I can reproduce connection leak coupled with the watcher - there is a following trace in the console: {code} 2017-04-17 12:14:05,499[ConnectionPool] [WARN] [okhttp3.OkHttpClient 129] - A connection to https://kubernetes.default.svc/ was leaked. Did you forget to close a response body? To see where this was allocated, set the OkHttpClient logger level to FINE: Logger.getLogger(OkHttpClient.class.getName()).setLevel(Level.FINE); {code} But with this warn I still can create / remove workspaces without any issues, that is why I am not sure that watch issue is tightly coupled with pipe is not connected error, which results in making che-server unusable	1	1	https://issues.jboss.org/browse/CHE-180

#	Requirement	Questions	Answered	Successfully Answered	URL
3825	CHE-180	{quote} Could you please provide steps for reproducing this issue against the latest 1.0.91 [1] template, which uses docker image rhche/che-server:e271ace ? {quote} - create a workspace - start and stop it several times. In the che server log, you will get the following: {quote} 2017-04-17 13:20:07,208[ConnectionPool] [WARN] [okhttp3.OkHttpClient 129] - A connection to https://kubernetes.default.svc/ was leaked. Did you forget to close a response body? To see where this was allocated, set the OkHttpClient logger level to FINE: Logger.getLogger(OkHttpClient.class.getName()).setLevel(Level.FINE); 2017-04-17 13:20:09,975[ConnectionPool] [WARN] [okhttp3.OkHttpClient 129] - A connection to https://kubernetes.default.svc/ was leaked. Did you forget to close a response body? To see where this was allocated, set the OkHttpClient logger level to FINE: Logger.getLogger(OkHttpClient.class.getName()).setLevel(Level.FINE); 2017-04-17 13:20:09,983[ConnectionPool] [WARN] [okhttp3.OkHttpClient 129] - A connection to https://kubernetes.default.svc/ was leaked. Did you forget to close a response body? To see where this was allocated, set the OkHttpClient logger level to FINE: Logger.getLogger(OkHttpClient.class.getName()).setLevel(Level.FINE); 2017-04-17 13:20:14,691[ConnectionPool] [WARN] [okhttp3.OkHttpClient 129] - A connection to https://kubernetes.default.svc/ was leaked. Did you forget to close a response body? To see where this was allocated, set the OkHttpClient logger level to FINE: Logger.getLogger(OkHttpClient.class.getName()).setLevel(Level.FINE); 2017-04-17 13:20:17,420[ConnectionPool] [WARN] [okhttp3.OkHttpClient 129] - A connection to https://kubernetes.default.svc/ was leaked. Did you forget to close a response body? To see where this was allocated, set the OkHttpClient logger level to FINE: Logger.getLogger(OkHttpClient.class.getName()).setLevel(Level.FINE); {quote} - debug the che server and set a breakpoint to InputStreamPumper.java, line 305. You will see the "Pipe not connected" exception is thrown. After some time, you won't be able to create a workspace. The snjeza/che-server:che-180 image is the same as the rhche/che-server:e271ace image. It uses the current openshift branch and adds only a log so you can see what's happening. See https://github.com/snjeza/che/tree/che-180.	1	1	https://issues.jboss.org/browse/CHE-180
3826	CHE-180	[~snjeza] hmmm... can reproduce all except "After some time, you won't be able to create a workspace" :/ minishift has been running for ~ 5 hr non-stop and I have started / stopped / removed dozen of worrkspaces by now and still can create new one without problems. [~mloriedo], [~amisevsk] can you, guys, reproduce original issue against the latest openshift connector ?	1	1	https://issues.jboss.org/browse/CHE-180
3827	CHE-173	So a buffer on the ws-agent. I think we should also consider what happens when we move servers to run on their own container. What will be the role of agent on that architecture?	1	1	https://issues.jboss.org/browse/CHE-173
3828	CHE-142	Doesn't Che uses jgit ? then the missing piece is to have a CredentialsProvider that will read the .gitcredentials file or some other storage or API point that will provide that token given by platform.	0	0	https://issues.jboss.org/browse/CHE-142
3829	CHE-132	[~mloriedo], I was not able to find vertx stack there, I thought this one is top priority is it coupled with - https://issues.jboss.org/browse/CHE-128 ? (vertx stack is not going to be published on dockerhub ?)	1	1	https://issues.jboss.org/browse/CHE-132
3830	CHE-122	for longer time/with more request over time (e.g. soak scenario). I have following limits for user which runs Che starter and vertx server for performance: open files (-n) 1024 max user processes (-u) 62616 Exception: {code}2017-03-07 10:55:33.883 WARN 13595 --- [qtp21257599-76] o.e.j.u.t.QueuedThreadPool : Unexpected thread death: org.eclipse.jetty.util.thread.QueuedThreadPool$2@cec462f in qtp21257599{STARTED,8<=62<=200,i=39,q=0} Exception in thread "qtp21257599-43" java.lang.NoClassDefFoundError: org/apache/logging/log4j/core/impl/ThrowableProxy at org.apache.logging.log4j.core.impl.Log4jLogEvent.getThrownProxy(Log4jLogEvent.java:555) at org.apache.logging.log4j.core.pattern.ExtendedThrowablePatternConverter.format(ExtendedThrowablePatternConverter.java:61) at org.springframework.boot.logging.log4j2.ExtendedWhitespaceThrowablePatternConverter.format(ExtendedWhitespaceThrowablePatternConverter.java:50) at org.apache.logging.log4j.core.pattern.PatternFormatter.format(PatternFormatter.java:38) at org.apache.logging.log4j.core.layout.PatternLayout$PatternSerializer.toSerializable(PatternLayout.java:294) at org.apache.logging.log4j.core.layout.PatternLayout.toText(PatternLayout.java:195) at org.apache.logging.log4j.core.layout.PatternLayout.encode(PatternLayout.java:180) at org.apache.logging.log4j.core.layout.PatternLayout.encode(PatternLayout.java:57) at org.apache.logging.log4j.core.appender.AbstractOutputStreamAppender.directEncodeEvent(AbstractOutputStreamAppender.java:176) at org.apache.logging.log4j.core.appender.AbstractOutputStreamAppender.tryAppend(AbstractOutputStreamAppender.java:169) at org.apache.logging.log4j.core.appender.AbstractOutputStreamAppender.append(AbstractOutputStreamAppender.java:160) at org.apache.logging.log4j.core.config.AppenderControl.tryCallAppender(AppenderControl.java:156) at org.apache.logging.log4j.core.config.AppenderControl.callAppender0(AppenderControl.java:129) at org.apache.logging.log4j.core.config.AppenderControl.callAppenderPreventRecursion(AppenderControl.java:120) at org.apache.logging.log4j.core.config.AppenderControl.callAppender(AppenderControl.java:84) at org.apache.logging.log4j.core.config.LoggerConfig.callAppenders(LoggerConfig.java:447) at org.apache.logging.log4j.core.config.LoggerConfig.processLogEvent(LoggerConfig.java:432) at org.apache.logging.log4j.core.config.LoggerConfig.log(LoggerConfig.java:416) at org.apache.logging.log4j.core.config.LoggerConfig.log(LoggerConfig.java:402) at org.apache.logging.log4j.core.config.AwaitCompletionReliabilityStrategy.log(AwaitCompletionReliabilityStrategy.java:63) at org.apache.logging.log4j.core.Logger.logMessage(Logger.java:146) at org.apache.logging.slf4j.Log4jLogger.log(Log4jLogger.java:376) at org.eclipse.jetty.util.log.JettyAwareLogger.log(JettyAwareLogger.java:607) at org.eclipse.jetty.util.log.JettyAwareLogger.warn(JettyAwareLogger.java:431) at org.eclipse.jetty.util.log.Slf4jLog.warn(Slf4jLog.java:69) at org.eclipse.jetty.util.log.Slf4jLog.warn(Slf4jLog.java:64) at org.eclipse.jetty.util.thread.QueuedThreadPool$2.run(QueuedThreadPool.java:647) at java.lang.Thread.run(Thread.java:745) java.lang.NoClassDefFoundError:	1	1	https://issues.jboss.org/browse/CHE-122

#	Requirement	Questions	Answered	Successfully Answered	URL
3831	CHE-122	started to fail and produces following exception. Limits for this scenario were: open files (-n) 64000 max user processes (-u) 100000 stack size (kbytes, -s) 8192 {code}2017-03-07 11:13:34.921 WARN 26953 --- [qtp135184888-36] o.e.j.u.t.QueuedThreadPool : Unexpected thread death: org.eclipse.jetty.util.thread.QueuedThreadPool$2@2b919895 in qtp135184888{STARTED,8<=11<=200,i=0,q=1} Exception in thread "qtp135184888-36" java.lang.OutOfMemoryError: unable to create new native thread at java.lang.Thread.start0(Native Method) at java.lang.Thread.start(Thread.java:714) at org.eclipse.jetty.util.thread.QueuedThreadPool.startThreads(QueuedThreadPool.java:476) at org.eclipse.jetty.util.thread.QueuedThreadPool.access$200(QueuedThreadPool.java:49) at org.eclipse.jetty.util.thread.QueuedThreadPool$2.run(QueuedThreadPool.java:657) at java.lang.Thread.run(Thread.java:745) 2017-03-07 11:13:34.923 WARN 26953 --- [p135184888-5647] o.e.j.u.t.QueuedThreadPool : java.lang.OutOfMemoryError: unable to create new native thread at java.lang.Thread.start0(Native Method) ~[?:1.8.0_121] at java.lang.Thread.start(Thread.java:714) ~[?:1.8.0_121] at org.eclipse.jetty.util.thread.QueuedThreadPool.startThreads(QueuedThreadPool.java:476) ~[jetty-util-9.4.1.v20170120.jar!/:9.4.1.v20170120] at org.eclipse.jetty.util.thread.QueuedThreadPool.access$200(QueuedThreadPool.java:49) ~[jetty-util-9.4.1.v20170120.jar!/:9.4.1.v20170120] at org.eclipse.jetty.util.thread.QueuedThreadPool$2.run(QueuedThreadPool.java:635) [jetty-util-9.4.1.v20170120.jar!/:9.4.1.v20170120] at java.lang.Thread.run(Thread.java:745) [?:1.8.0_121] 2017-03-07 11:13:34.923 WARN 26953 --- [p135184888-5647] o.e.j.u.t.QueuedThreadPool : Unexpected thread death: org.eclipse.jetty.util.thread.QueuedThreadPool$2@2b919895 in qtp135184888{STARTED,8<=10<=200,i=0,q=1} Exception in thread "qtp135184888-5647" java.lang.OutOfMemoryError: unable to create new native thread at java.lang.Thread.start0(Native Method) at java.lang.Thread.start(Thread.java:714) at org.eclipse.jetty.util.thread.QueuedThreadPool.startThreads(QueuedThreadPool.java:476) at org.eclipse.jetty.util.thread.QueuedThreadPool.access$200(QueuedThreadPool.java:49) at org.eclipse.jetty.util.thread.QueuedThreadPool$2.run(QueuedThreadPool.java:657) at java.lang.Thread.run(Thread.java:745) 2017-03-07 11:13:34.925 WARN 26953 --- [qtp135184888-39] o.e.j.s.HttpChannel : /workspace org.springframework.web.util.NestedServletException: Handler dispatch failed; nested exception is java.lang.OutOfMemoryError: unable to create new native thread at org.springframework.web.servlet.DispatcherServlet.doDispatch(DispatcherServlet.java:978) ~[spring-webmvc-4.3.6.RELEASE.jar!/:4.3.6.RELEASE] at org.springframework.web.servlet.DispatcherServlet.doService(DispatcherServlet.java:897) ~[spring-webmvc-4.3.6.RELEASE.jar!/:4.3.6.RELEASE] at org.springframework.web.servlet.FrameworkServlet.processRequest(FrameworkServlet.java:970) ~[spring-webmvc-4.3.6.RELEASE.jar!/:4.3.6.RELEASE] at org.springframework.web.servlet.FrameworkServlet.doPost(FrameworkServlet.java:872) ~[spring-webmvc-4.3.6.RELEASE.jar!/:4.3.6.RELEASE] at javax.servlet.http.HttpServlet.service(HttpServlet.java:707) ~[javax.servlet-api-3.1.0.jar!/:3.1.0] at org.springframework.web.servlet.FrameworkServlet.service(FrameworkServlet.java:846) ~[spring-webmvc-4.3.6.RELEASE.jar!/:4.3.6.RELEASE] at javax.servlet.http.HttpServlet.service(HttpServlet.java:790) ~[javax.servlet-api-3.1.0.jar!/:3.1.0] at org.eclipse.jetty.servlet.ServletHolder.handle(ServletHolder.java:841) ~[jetty-servlet-9.4.1.v20170120.jar!/:9.4.1.v20170120] at org.eclipse.jetty.servlet.ServletHandler$CachedChain.doFilter(ServletHandler.java:1634) ~[jetty-servlet-	1	1	https://issues.jboss.org/browse/CHE-122
3832	CHE-120	Marian let me try to summarize what we should investigate: - If we can run a command as a super user in a non privileged container on OpenShift (e.g. "sudo service apache2 start") or what are the minimum needed privileges - Which command are run as super user when we bootstrap a vert.x workspace and how can we get rid of these (e.g. https://github.com/eclipse/che/blob/master/wsagent/agent/src/main/resources/org.eclipse.che.ws-agent.script.sh#L37 is one of these) And I think that these task can be done even if the Docker API hasn't been implemented yet in the OpenShift connector.	0	0	https://issues.jboss.org/browse/CHE-120
3833	CHE-60	For the remaining 2 we should be able to use a distinct PV then the one used by che-server volume. Can you test it and see if you get an error? I suspect something like folder does not exist...	1	1	https://issues.jboss.org/browse/CHE-60
3834	CHE-60	Che-server and a workspace can use the same PVC when they are in the same namespace. That means, we can use the claim-che-data PVC (from CHE-59) when creating a workspace. They will have the same /data directory (I have checked it) and we can just create the mentioned links. If we would use a distinct PVC (I suppose you meant PVC, not PV), we would have to ensure that they are shared between a server and a workspace. BTW Do you know how we can call a command in the workspace container?	1	1	https://issues.jboss.org/browse/CHE-60
3835	CHE-35	On Centos with Docker version 1.10.3 (package 1.10.3-46.el7.centos.14) and SELinux activated Docker provides an error when trying to mount a non existing folder: {code:bash} docker create -v /tmp/thisfolderdoesntexist/:/tmp/:Z busybox sh Error response from daemon: no such file or directory {code} This prevent to Che from creating new workspaces on CentOS with SELinux activated because Che uses flag 'Z' when mounting the new workspace folder. To avoid this problem Che should have a property and an environment variable that would let configurable the Z flag. The background of this particular Docker issue: {quote}Everything starts with this bugzilla https://bugzilla.redhat.com/show_bug.cgi?id=1274210. After that Dan (Walsh) opened an issue upstream in docker https://github.com/docker/docker/issues/17262. That issue has been closed at some point because docker deprecated the auto-creation of directories on the host when using bind mounts (https://github.com/docker/docker/pull/16349). Shortly after that, they realized that was too much of an issue for existing deployments and reverted that deprecation. So, Dan went ahead and implemented the directories auto-creation on host with https://github.com/docker/docker/pull/23026. The issue I see here is that we didn't backport https://github.com/docker/docker/pull/23026 to our 1.10.3 projectatomic/docker branches (yet?). I'm not sure we're going to ship this backport in 7.3.1 so I'd say this issue has been fixed in rhel/centos with docker-latest 1.12.1 (is centos shipping docker-latest?) {quote}	1	1	https://issues.jboss.org/browse/CHE-35
3836	CHE-35	I can reproduce the issue using centos/7 and centos/atomic-host with docker 1.10.3, but can't with docker >=1.12.1. centos/7 doesn't include docker-latest (1.12.1). It has to be separately installed. I am not sure what else I could help with. [~mloriedo], would you like me to create PRs for the 1.10.3 branches? There are two such branches in https://github.com/projectatomic/docker - fedora-1.10.3 and rhel7-1.10.3.	1	1	https://issues.jboss.org/browse/CHE-35
3837	CHE-34	We are working on the solution which has a similar feature but currently integrated as a part the build (ordering test execution). Is there any work around that started on your side?	1	0	https://issues.jboss.org/browse/CHE-34
3838	CHE-34	Which ITS it is? Can you share the link?	1	0	https://issues.jboss.org/browse/CHE-34
3839	CHE-32	[~amisevsk] that's great! We need to share this script with CentOS team. We can share that on eclipse-che@redhat.com but maybe it's worth create a github repository (for redhat che related source files) and put that script there. [~gercan] what do you think?	1	1	https://issues.jboss.org/browse/CHE-32
3840	CHE-31	PR was merged and this functionality should be available in Che 5.0.0-M7 release	0	0	https://issues.jboss.org/browse/CHE-31
3841	CHE-24	[~mloriedo] I was not able to finish this task in this sprint. Could you please drag it to the #123 during the next planning ?	0	0	https://issues.jboss.org/browse/CHE-24

#	Requirement	Questions	Answered	Successfully Answered	URL
3842	CHE-17	I had a chat with [~tmancini] and we agreed that we should concentrate testing effort around "unit testing for Java" for this release. CHE-11 is probably relevant to this work and may be a starting point. [~aslak], [~alrubinger] Can you think of additional requirements for Java unit testing on Che? Must support frameworks ?	1	0	https://issues.jboss.org/browse/CHE-17
3843	CHE-17	I would also like to start the discussion about how we should actually run those tests from the IDE. What should happen when we click "run test" or hit the shortcut? We can either do it like Eclipse and IDEA does, so having our own test runners or delegate it to the underlying build tool of choice and collect the results. If we go with the second approach we don't have to discuss the support for testing frameworks, but rather about build tools support in Che. As far as I know we are limited to Maven at this point. Do you see any disadvantages of delegating to build tool rather than writing IDE/framework specific test runner? From Arquillian experience, where [~aslak] can shed more light, IDE specific runners were always a struggle, but if we go with delegating to the build are we losing something? Speed? Incremental compilation? I've spent time reading aforementioned PR (what I assume is a GSoC project) and I've noticed it's a mix of both. For resolving test dependencies it [simply invokes Maven|https://github.com/Mirage20/che/blob/77ef09d02a2a68ded2b6ec872eabc4d201fa77c1/plugins/plugin-java-test-runner/che-java-testing-classpath/che-java-testing-classpath-maven-server/src/main/java/org/eclipse/che/ide/ext/java/testing/classpath/maven/server/MavenTestClasspathProvider.java#L38-L55], but for executing the tests [it is heavily reflection based|https://github.com/Mirage20/che/blob/77ef09d02a2a68ded2b6ec872eabc4d201fa77c1/plugins/plugin-java-test-runner/che-java-testing-junit/che-java-testing-junit-server/src/main/java/org/eclipse/che/ide/ext/java/testing/junit/server/JUnitTestRunner.java#L90-L148].	1	1	https://issues.jboss.org/browse/CHE-17
3844	CHE-17	[~gercan] Cool, let me know if you want some support from our side. Do the suggestions above make sense? I would also like to hear your opinion about few aspects mentioned in the thread. But no rush :)	1	1	https://issues.jboss.org/browse/CHE-17
3845	CHE-17	[~amisevsk] and I have been talking about this. Here is our thoughts. We plan to build on top of GSoC PR that was referred to. It seems to be doing most of what is discussed here but [~amisevsk] will take a pass and determine the gaps. We may also help with its merge if needed. [~bmajsak] I think most of the requirements were OK, a few questions and comments. * Can you elaborate what do you mean by _fork mode_? * Environment and VM paramters are a more general problem on Che that should be addressed for all run types not only tests. Angel will take a look and create suggestions on how to proceed with them. * We do not plan to depend on maven to run the tests because this will sever our ability to easily run a selected set of tests which is a requirement that was discussed above. * Created CHE-34 for the continuous testing plugin to Che, but I do not expect the work on it to start soon. [~alrubinger] I think build and deployment pipeline alignment is probably a larger scope issue then testing. I think in general the stacks used by Che and the pipeline should be same. See CHE-33 for more on that.	1	0	https://issues.jboss.org/browse/CHE-17
3846	CHE-17	{quote} Can you elaborate what do you mean by fork mode? {quote} Like in surefire for example (as well as in IDEA) you can define how you would like to run your tests in parallel. Classes - just classes are forked, methods - it also executes methods in parallel. {quote} We do not plan to depend on maven to run the tests because this will sever our ability to easily run a selected set of tests which is a requirement that was discussed above. {quote} How would this limit you? Just out of curiosity.	1	1	https://issues.jboss.org/browse/CHE-17
3847	CHE-17	What's the status? I see that epic child is resolved and closed, because the upstream has been merged. But [~amisevsk] has some pretty good points in the CHE-11 and the current user experience with java unit testing in Che is bad.	0	0	https://issues.jboss.org/browse/CHE-17
3848	ERT-391	[~amisevsk], pushed to master. One thing I noticed while checking the gerrit change is that Outline view is not refreshing on js file change. In order to update outline I have to close and open the file again. Do you experience the same behaviour ?	1	1	https://issues.jboss.org/browse/ERT-391
3849	ERT-222	Marking as RESOLVED since this has been fixed, and will be targeting 5.0 (RC3)	0	0	https://issues.jboss.org/browse/ERT-222
3850	ERT-74	No. JIRA and bugzilla have different users. It's impossible to match them. At least not always.	0	0	https://issues.jboss.org/browse/ERT-74
3851	ERT-68	[~dgolovin] Why do you closed the issue? Shouldn't we add ChDT features into WTP p2 repo?	0	0	https://issues.jboss.org/browse/ERT-68
3852	EJBTHREE-2122	Ondrej, I'm not responsible for modifying the community documentation, and as far as I know we don't have EJB3 docs shipping with EAP yet. Can you suggest anyone for this to be assigned to in the mean time? I'm not going to get the chance to look at anything extra until November or so.	1	0	https://issues.jboss.org/browse/EJBTHREE-2122
3853	EJBTHREE-2110	Patch applied to ejb3-core 1.3 branch http://anonsvn.jboss.org/repos/jbossas/projects/ejb3/branches/jboss-ejb3-core-1.3/ as well as EJB3 trunk.	0	0	https://issues.jboss.org/browse/EJBTHREE-2110
3854	EJBTHREE-2110	How can the patch be applied? Thank you.	0	0	https://issues.jboss.org/browse/EJBTHREE-2110
3855	EJBTHREE-2110	Is this patch embedded into JBoss 7 ?	1	0	https://issues.jboss.org/browse/EJBTHREE-2110
3856	EJBTHREE-2071	Did the AS workaround help? The missing JMXAdapter class is now include in the client jar.	0	0	https://issues.jboss.org/browse/EJBTHREE-2071
3857	EJBTHREE-2042	Remi, thank you for the profiler numbers. Can you provide more details about your application including the bean code and the invocation details?	0	0	https://issues.jboss.org/browse/EJBTHREE-2042
3858	EJBTHREE-2011	This is a BLOCKER for CDI integration in AS. Is there an estimate or a Fix Version that this is targeted for?	0	0	https://issues.jboss.org/browse/EJBTHREE-2011

#	Requirement	Questions	Answered	Successfully Answered	URL
3859	EJBTHREE-1851	names: 2009-06-23 21:00:30,589 DEBUG [org.jboss.ejb3.metrics.deployer.Ejb3MetricsDeployer] (RMI TCP Connection(10)-127.0.0.1) Attaching MC Bean with name "StatelessBean" .... This can lead to duplicate binding exceptions if there are multiple deployments containing the same EJB names: 2009-06-23 21:00:39,736 ERROR [org.jboss.kernel.plugins.dependency.AbstractKernelController] (RMI TCP Connection(10)-127.0.0.1) Error installing to Real: name=vfszip:/path/target/test-lib/stateful-test.jar/ state=PreReal mode=Manual requiredState=Real org.jboss.deployers.spi.DeploymentException: Error deploying: StatelessBean at org.jboss.deployers.spi.DeploymentException.rethrowAsDeploymentException(DeploymentException.java:49) at org.jboss.deployers.vfs.deployer.kernel.BeanMetaDataDeployer.deploy(BeanMetaDataDeployer.java:125) at org.jboss.deployers.vfs.deployer.kernel.BeanMetaDataDeployer.deploy(BeanMetaDataDeployer.java:51) at org.jboss.deployers.spi.deployer.helpers.AbstractSimpleRealDeployer.internalDeploy(AbstractSimpleRealDeployer.java:62) at org.jboss.deployers.spi.deployer.helpers.AbstractRealDeployer.deploy(AbstractRealDeployer.java:50) at org.jboss.deployers.plugins.deployers.DeployerWrapper.deploy(DeployerWrapper.java:171) at org.jboss.deployers.plugins.deployers.DeployersImpl.doDeploy(DeployersImpl.java:1440) at org.jboss.deployers.plugins.deployers.DeployersImpl.doInstallParentFirst(DeployersImpl.java:1158) at org.jboss.deployers.plugins.deployers.DeployersImpl.doInstallParentFirst(DeployersImpl.java:1179) at org.jboss.deployers.plugins.deployers.DeployersImpl.install(DeployersImpl.java:1099) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:348) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:1632) at org.jboss.dependency.plugins.AbstractController.incrementState(AbstractController.java:935) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:1083) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:985) at org.jboss.dependency.plugins.AbstractController.change(AbstractController.java:823) at org.jboss.dependency.plugins.AbstractController.change(AbstractController.java:553) at org.jboss.deployers.plugins.deployers.DeployersImpl.process(DeployersImpl.java:782) at org.jboss.deployers.plugins.main.MainDeployerImpl.process(MainDeployerImpl.java:702) at org.jboss.deployment.MainDeployer.deploy(MainDeployer.java:831) at org.jboss.deployment.MainDeployer.redeploy(MainDeployer.java:605) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at	0	0	https://issues.jboss.org/browse/EJBTHREE-1851
3860	EJBTHREE-1815	while processing container metadata for EJB: SimpleBean in unit: testInterceptor2.jar 16:01:01,086 ERROR [AbstractKernelController] Error installing to Real: name=vfsfile:/home/jaikiran/MyStuff/jboss-5.1.0.Beta1/server/default/deploy/test-1751.ear/ state=PreReal mode=Manual requiredState=Real org.jboss.deployers.spi.DeploymentException: Error deploying testInterceptor2.jar: Exception while processing container metadata for EJB: SimpleBean in unit: testInterceptor2.jar at org.jboss.ejb3.deployers.Ejb3Deployer.deploy(Ejb3Deployer.java:196) at org.jboss.ejb3.deployers.Ejb3Deployer.deploy(Ejb3Deployer.java:99) at org.jboss.deployers.vfs.spi.deployer.AbstractVFSRealDeployer.internalDeploy(AbstractVFSRealDeployer.java:45) at org.jboss.deployers.spi.deployer.helpers.AbstractRealDeployer.deploy(AbstractRealDeployer.java:50) at org.jboss.deployers.plugins.deployers.DeployerWrapper.deploy(DeployerWrapper.java:171) at org.jboss.deployers.plugins.deployers.DeployersImpl.doDeploy(DeployersImpl.java:1439) at org.jboss.deployers.plugins.deployers.DeployersImpl.doInstallParentFirst(DeployersImpl.java:1157) at org.jboss.deployers.plugins.deployers.DeployersImpl.doInstallParentFirst(DeployersImpl.java:1210) at org.jboss.deployers.plugins.deployers.DeployersImpl.install(DeployersImpl.java:1098) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:348) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:1598) at org.jboss.dependency.plugins.AbstractController.incrementState(AbstractController.java:934) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:1062) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:984) at org.jboss.dependency.plugins.AbstractController.change(AbstractController.java:822) at org.jboss.dependency.plugins.AbstractController.change(AbstractController.java:553) at org.jboss.deployers.plugins.deployers.DeployersImpl.process(DeployersImpl.java:781) at org.jboss.deployers.plugins.main.MainDeployerImpl.process(MainDeployerImpl.java:698) at org.jboss.system.server.profileservice.repository.MainDeployerAdapter.process(MainDeployerAdapter.java:117) at org.jboss.system.server.profileservice.repository.ProfileDeployAction.install(ProfileDeployAction.java:59) at org.jboss.system.server.profileservice.repository.AbstractProfileAction.install(AbstractProfileAction.java:53) at org.jboss.system.server.profileservice.repository.AbstractProfileService.install(AbstractProfileService.java:355) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:348) at	1	1	https://issues.jboss.org/browse/EJBTHREE-1815
3861	EJBTHREE-1721	Just wondering, isn't the {{@Asynchronous}} in JBoss AS complete because this issue is still open?	0	0	https://issues.jboss.org/browse/EJBTHREE-1721
3862	EJBTHREE-1669	Can you provide some context on the forums? The "service" Integration Tests and EJBTHREE-1530 are now working.	1	1	https://issues.jboss.org/browse/EJBTHREE-1669
3863	EJBTHREE-1659	Updated the tutorial to use the new API on the AsyncUtils, as explained in the forum thread. Works fine now.	0	0	https://issues.jboss.org/browse/EJBTHREE-1659
3864	EJBTHREE-1650	But we have multiple servers and we can add dynamically new servers (clustering) so we can't specify one explicit address. Is there a workaround ?	0	0	https://issues.jboss.org/browse/EJBTHREE-1650
3865	EJBTHREE-1650	Is there any workaround? If not when it is planned to be resolved?	0	0	https://issues.jboss.org/browse/EJBTHREE-1650

#	Requirement	Questions	Answered	Successfully Answered	URL
3866	EJBTHREE-1650	Andrew, Do I have to hardcode the clientBindUrl parameter into the RemoteBinding() annotation? This doesn't work: @RemoteBinding(clientBindUrl = "sslsocket://127.0.0.1:3843", jndiBinding="MessageServiceEJB/secure") and this doesn't work either: @RemoteBinding(clientBindUrl = "sslsocket://0.0.0.0:3843", jndiBinding="MessageServiceEJB/secure") and of course this doesn't work: @RemoteBinding(clientBindUrl = "sslsocket://${jboss.bind.address}:3843", jndiBinding="MessageServiceEJB/secure") This does works but isn't what I need: @RemoteBinding(clientBindUrl = "sslsocket://tommyboy:3843", jndiBinding="MessageServiceEJB/secure") Hardcoding the clientBindUrl will make it just about impossible to deploy my app. Any suggestions? Thanks, Bryant	1	0	https://issues.jboss.org/browse/EJBTHREE-1650
3867	EJBTHREE-1641	How come this is unscheduled? As it is now, while booting AS 5.1.0.Beta1 you get a bunch of: ... 17:56:28,765 WARN [Ejb3Deployment] [EJBTHREE-1641] NoOp setRemoteProxyFactoryRe gistry; developers may ignore this message; it will be removed when backwards-co mpatibility between EJB3 and AS is resolved 17:56:28,781 WARN [Ejb3Deployment] [EJBTHREE-1641] NoOp setRemoteProxyFactoryRe gistry; developers may ignore this message; it will be removed when backwards-co mpatibility between EJB3 and AS is resolved 17:56:30,578 WARN [Ejb3Deployment] [EJBTHREE-1641] NoOp setRemoteProxyFactoryRe gistry; developers may ignore this message; it will be removed when backwards-co mpatibility between EJB3 and AS is resolved 17:56:30,593 WARN [Ejb3Deployment] [EJBTHREE-1641] NoOp setRemoteProxyFactoryRe gistry; developers may ignore this message; it will be removed when backwards-co mpatibility between EJB3 and AS is resolved 17:56:30,781 WARN [Ejb3Deployment] [EJBTHREE-1641] NoOp setRemoteProxyFactoryRe gistry; developers may ignore this message; it will be removed when backwards-co mpatibility between EJB3 and AS is resolved 17:56:32,000 WARN [Ejb3Deployment] [EJBTHREE-1641] NoOp setRemoteProxyFactoryRe gistry; developers may ignore this message; it will be removed when backwards-co mpatibility between EJB3 and AS is resolved 17:56:32,000 WARN [Ejb3Deployment] [EJBTHREE-1641] NoOp setRemoteProxyFactoryRe gistry; developers may ignore this message; it will be removed when backwards-co mpatibility between EJB3 and AS is resolved 17:56:32,109 WARN [Ejb3Deployment] [EJBTHREE-1641] NoOp setRemoteProxyFactoryRe gistry; developers may ignore this message; it will be removed when backwards-co ...	1	1	https://issues.jboss.org/browse/EJBTHREE-1641
3868	EJBTHREE-1536	Alessio, would you please confirm r79904 properly handles these API changes? If so assign back and I'll remove my comments from the code.	0	0	https://issues.jboss.org/browse/EJBTHREE-1536
3869	EJBTHREE-1501	That's going to hold the map lock through a potentially long-running call to container.destroy(ctx) which calls any @PreDestroy. Before incurring that cost, what's the scenario where this is a problem, and can that scenario be solved another way, e.g. with a finer-grained lock, e.g. on the context rather than the cacheMap?	1	0	https://issues.jboss.org/browse/EJBTHREE-1501
3870	EJBTHREE-1492	http://repo1.maven.org/maven2/org/jboss/ejb3/jboss-ejb3-build/1.0.2/jboss-ejb3-build-1.0.2.pom [INFO] Unable to find resource 'org.jboss.ejb3:jboss-ejb3-build:pom:1.0.2' in repository central (http://repo1.maven.org/maven2) [INFO] ---------- [ERROR] FATAL ERROR [INFO] ---------- [INFO] Error building POM (may not be this project's POM). Project ID: null:jboss-ejb3-embedded:jar:1.0.0-SNAPSHOT Reason: Cannot find parent: org.jboss.ejb3:jboss-ejb3-build for project: null:jboss-ejb3-embedded:jar:1.0.0-SNAPSHOT for project null:jboss-ejb3-embedded:jar:1.0.0-SNAPSHOT [INFO] ---------- [INFO] Trace org.apache.maven.reactor.MavenExecutionException: Cannot find parent: org.jboss.ejb3:jboss-ejb3-build for project: null:jboss-ejb3-embedded:jar:1.0.0-SNAPSHOT for project null:jboss-ejb3-embedded:jar:1.0.0-SNAPSHOT at org.apache.maven.DefaultMaven.getProjects(DefaultMaven.java:432) at org.apache.maven.DefaultMaven.doExecute(DefaultMaven.java:300) at org.apache.maven.DefaultMaven.execute(DefaultMaven.java:137) at org.apache.maven.cli.MavenCli.main(MavenCli.java:356) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.codehaus.classworlds.Launcher.launchEnhanced(Launcher.java:315) at org.codehaus.classworlds.Launcher.launch(Launcher.java:255) at org.codehaus.classworlds.Launcher.mainWithExitCode(Launcher.java:430) at org.codehaus.classworlds.Launcher.main(Launcher.java:375) Caused by: org.apache.maven.project.ProjectBuildingException: Cannot find parent: org.jboss.ejb3:jboss-ejb3-build for project: null:jboss-ejb3-embedded:jar:1.0.0-SNAPSHOT for project null:jboss-ejb3-embedded:jar:1.0.0-SNAPSHOT at org.apache.maven.project.DefaultMavenProjectBuilder.assembleLineage(DefaultMavenProjectBuilder.java:1398) at org.apache.maven.project.DefaultMavenProjectBuilder.buildInternal(DefaultMavenProjectBuilder.java:823) at org.apache.maven.project.DefaultMavenProjectBuilder.buildFromSourceFileInternal(DefaultMavenProjectBuilder.java:508) at org.apache.maven.project.DefaultMavenProjectBuilder.build(DefaultMavenProjectBuilder.java:200) at org.apache.maven.DefaultMaven.getProject(DefaultMaven.java:632) at org.apache.maven.DefaultMaven.collectProjects(DefaultMaven.java:515) at org.apache.maven.DefaultMaven.getProjects(DefaultMaven.java:419) ... 11 more Caused by: org.apache.maven.project.ProjectBuildingException: POM 'org.jboss.ejb3:jboss-ejb3-build' not found in repository: Unable to download the artifact from any repository org.jboss.ejb3:jboss-ejb3-build:pom:1.0.2 from the specified remote repositories: central (http://repo1.maven.org/maven2) for project org.jboss.ejb3:jboss-ejb3-build at org.apache.maven.project.DefaultMavenProjectBuilder.findModelFromRepository(DefaultMavenProjectBuilder.java:605) at org.apache.maven.project.DefaultMavenProjectBuilder.assembleLineage(DefaultMavenProjectBuilder.java:1394) ... 17 more Caused by: org.apache.maven.artifact.resolver.ArtifactNotFoundException: Unable to download the artifact from any repository org.jboss.ejb3:jboss-ejb3-build:pom:1.0.2 from the specified remote repositories: central (http://repo1.maven.org/maven2) at org.apache.maven.artifact.resolver.DefaultArtifactResolver.resolve(DefaultArtifactResolver.java:228) at	0	0	https://issues.jboss.org/browse/EJBTHREE-1492
3871	EJBTHREE-1490	Does this represent improper use of EJB2.1 views defined in @Remote? A general case that's present throughout the testsuite or something more specific?	0	0	https://issues.jboss.org/browse/EJBTHREE-1490
3872	EJBTHREE-1480	This says resolved but for build: SVNTag=JBoss_5_0_0_GA date=200810310742, it still logs, "[JBoss] 08:19:36,599 WARN [LifecycleCallbacks] EJBTHREE-1480: LifecycleCallbackStack has not been defined for domain..." Is this forgotten, or when will it be removed?	0	0	https://issues.jboss.org/browse/EJBTHREE-1480
3873	EJBTHREE-1437	Is the fix for this in the EJB3 version currently used in AS trunk? If so, this issue can be closed	1	1	https://issues.jboss.org/browse/EJBTHREE-1437
3874	EJBTHREE-1337	I've added the Proxy implementation. Can someone take care of the injection?	1	1	https://issues.jboss.org/browse/EJBTHREE-1337
3875	EJBTHREE-1302	why are we not creating one ejb3 client jar rather than 10 different ejb3-xxx-client.jars?	0	0	https://issues.jboss.org/browse/EJBTHREE-1302
3876	EJBTHREE-1287	Does this mean the EJB3 project does not cover appclients in their test suite? I hope not. Wherever appclients are tested, service-ref injection should be tested as well.	0	0	https://issues.jboss.org/browse/EJBTHREE-1287
3877	EJBTHREE-1251	Anil, we're fixing up the TestSuite, can we get an update here? Current failures are due to Caused by: javax.ejb.EJBAccessException: Caller unauthorized	0	0	https://issues.jboss.org/browse/EJBTHREE-1251
3878	EJBTHREE-1248	Heiko: Who did you have in mind for this? Can you elaborate in this issue's description as to what needs to be done?	1	1	https://issues.jboss.org/browse/EJBTHREE-1248

#	Requirement	Questions	Answered	Successfully Answered	URL
3879	EJBTHREE-1248	Rajesh: Could you look at this for Heiko? I believe he's provided the sample configs, just need to get this into official QA.	1	1	https://issues.jboss.org/browse/EJBTHREE-1248
3880	EJBTHREE-1244	@Dragos: Can you provide more context on the EJB3 User Forum? Surprised to see this in 4.2.x. http://wiki.jboss.org/auth/wiki/BadPostMETOO	0	0	https://issues.jboss.org/browse/EJBTHREE-1244
3881	EJBTHREE-1096	What changes to JBoss AS or the EJB3 spec are required to achieve this? Please also comment on the difficulty and likely time frame for these changes.	1	0	https://issues.jboss.org/browse/EJBTHREE-1096
3882	EJBTHREE-1096	Any news on this? Will this require major modifications to JBAS or the EJB3 implementation?	1	0	https://issues.jboss.org/browse/EJBTHREE-1096
3883	EJBTHREE-1096	This does seem to be a popular issue. Yet, I see no comments regarding what would be required. Do any other app servers support this? If not, then wouldn't this be a major advantage to JBAS? If so, then shouldn't JBAS do the same?	1	0	https://issues.jboss.org/browse/EJBTHREE-1096
3884	EJBTHREE-1096	I would appreciate a response to Jason Long's questions directly above. This issue is not even assigned yet. This problem is a total disaster in terms of productivity. I've recently been restarting JBoss dozens of times a day during dev/test cycles for my Seam app. I use ant explode as default build for Seam apps. When you use ant deploy, it's faster for session bean dev/test cycles b/c the new version of the EAR will be re-deployed by JBoss (restart not required - which takes longer than re-deploy) but eventually you will have to restart Jboss/Eclipse due to strange problems and perm-gen memory issues even if you use -XX:PermSize=256M -XX:MaxPermSize=512M on a 2GB machine. Either way sucks. I've also seen seemingly invalid duplicate id errors when I save a change to a facelet and do an ant explode and refresh browser. These comments refer to Seam 2.0.0.GA, RF3.1.3.GA and JBoss 4.2.1.GA. Is this a EJB spec issue or JBoss EJB implementation issue or both?	1	0	https://issues.jboss.org/browse/EJBTHREE-1096
3885	EJBTHREE-1096	I am with you. I see several +1 from core JBoss developers, but is anyone actually taking this seriously? I realize that these developers are on other projects. Can we at least get this in for JBoss 5?	1	0	https://issues.jboss.org/browse/EJBTHREE-1096
3886	EJBTHREE-1096	And look at all the votes. Is there a way to search for issues by the number of votes? I would bet this it one of the top ones.	1	0	https://issues.jboss.org/browse/EJBTHREE-1096
3887	EJBTHREE-1096	Well I found Pete's answer from this thread I started on this issue: http://www.jboss.com/index.html?module=bb&op=viewtopic&p=4122726#4122726 me: Is it going to be targeted for release and coincide with EJB 3.1 release which apparently is going to remove local interface for session beans? Pete: Unlikely, this is implementation detail, not spec related. Also EJB3.1 is still a while away. So it sounds like it needs to be a part of a maintenance release for JBoss 5...	0	0	https://issues.jboss.org/browse/EJBTHREE-1096
3888	EJBTHREE-1096	I posted the following comment in Ken Saks' blog (Ken is the EJB 3.1 Spec Lead) when the EJB 3.1 Early Draft was announced: " Hi Ken. Thanks for the resume. I was working with EJB since 2.0 (limitations fixed by our own), then EJB 3. I really like the Java EE platform and I'm looking forward for the new features. What I would really like is the possibility to hot deploy EJB. This would be a very big advantage for EJB, specially for the developers. Currently the development process for EJB applications is a bit cumbersome, we need to redeploy application server every time an EJB is added or changed. Do you know about the hot deployment (for development process) support for EJB? Was it mentioned on JCP? Was it considered? Will it be? Thanks in advance!" His answer was: " Hi Fernando, Thanks for your feedback. Hot deployment has typically been left as a value-add for products as opposed to a spec requirement, so it's unlikely we'll address this specifically within EJB 3.1. There's a basic issue with the way classloaders are used in the EE platform that makes supporting hot deployment for any single component challenging. Typically most or all of the classes in an application are loaded by a single classloader, which means everything has to be reloaded when a single thing changes. From a glassfish perspective we have a couple things to ease redeployment. The first is directory-based deployment. That saves you having to re-issue an asadmin command and transfer the application to the server. We've also significantly reduced the overhead of deployment itself. In the case of EJB deployment we've eliminated all statically generated code from the deployment stage (except for legacy CMP 2.x apps). " So, as you can read, seems like the EJB (3 or 3.1) hot deployment is reserved to be a "value-add" for JBoss AS. I'm really sure this would be a really big advantage for JBoss AS. The post I mentioned is here: http://blogs.sun.com/kensaks/entry/ejb_3_1_early_draft Regards,	1	0	https://issues.jboss.org/browse/EJBTHREE-1096
3889	EJBTHREE-1096	Any news on this as it is killing productivity for all users of EJB3? Here http://www.jboss.org/index.html?module=bb&op=viewtopic&p=4122726#4122726 on Wed Jan 23, 2008 Pete Muir stated, "I spoke to Andy from EJB3 team about this a few days ago, he was going to try to get it scheduled." That was over a year ago and this is still dead. Would someone please comment on this issue so I can decide whether to drop EJB3 entirely or suffer longer until this can be fixed?	1	0	https://issues.jboss.org/browse/EJBTHREE-1096
3890	EJBTHREE-1096	Jevgeni Kabanov: "The big things to look for in the next release are integration with EJB containers supporting changes to EJB interfaces, full Struts 1.x and 2.x support, better support for enums and static values plus some stuff we'd like to keep secret for the moment :)" source: http://java.dzone.com/articles/saving-time-java-development Note "integration with EJB containers supporting changes to EJB interfaces". That is a very critical enhancement for the upcoming release of JRebel. What does it mean exactly? Will Weblogic, Glassfish or JBoss support this behavior in an upcoming release, for example? Pretty decent tool, we're using 2.1a currently...	0	0	https://issues.jboss.org/browse/EJBTHREE-1096
3891	EJBTHREE-1029	A one workaround I have found is to deploy a EJB2.1 facade to the EJB3.0 on the server, this I found I needed to compile with 1.4.2 even though it will end up running on 1.6.x. Then I use the client jar's from 4.0.x and download and use the 2.0.x jar from jbossremoting to replace the 1.x.x version that comes with 4.0.x, this seems to work ok on client jdk1.4.2 -> server jdk1.6.x. Is there a simpler way to do this?	0	0	https://issues.jboss.org/browse/EJBTHREE-1029
3892	EJBTHREE-1026	Is this work complete?	1	0	https://issues.jboss.org/browse/EJBTHREE-1026
3893	EJBTHREE-1017	Galder, did you run the ejb3 testsuite with this patch? If so, I'd appreciate your help as I'm running into some issues. Thanks.	0	0	https://issues.jboss.org/browse/EJBTHREE-1017
3894	EJBTHREE-1017	Can someone upload to it to network.jboss.com? Customer needs it. Cheers,	0	0	https://issues.jboss.org/browse/EJBTHREE-1017
3895	EJBTHREE-1017	What link are you refering to when you say that the link is not here? If you're refering to the support case link, it is already in the JIRA issue.	0	0	https://issues.jboss.org/browse/EJBTHREE-1017
3896	EJBTHREE-856	Can this issue be resolved? Since we have JbossCache 1.4.1.SP6 now	1	1	https://issues.jboss.org/browse/EJBTHREE-856
3897	EJBTHREE-689	How can i use the xmbean attribute in @Service?	0	0	https://issues.jboss.org/browse/EJBTHREE-689
3898	EJBTHREE-680	Hi, Can this issue be added into the EJB3 RC9. Since I did not see this issue listed on the RC9 roadmap. Thanks, Joe	1	1	https://issues.jboss.org/browse/EJBTHREE-680
3899	EJBTHREE-679	Hi, Can this issue be added into EJB3 RC9 please? I did not see this issue listed on RC9 roadmap. Cheers, Joe	1	1	https://issues.jboss.org/browse/EJBTHREE-679
3900	EJBTHREE-678	Hi, Can this issue be added into EJB3 RC9 please? It is not listed in RC9 roadmap. Cheers, Joe	1	1	https://issues.jboss.org/browse/EJBTHREE-678
3901	EJBTHREE-648	Is this reuqest will be realized in AS 7 ?	0	0	https://issues.jboss.org/browse/EJBTHREE-648

#	Requirement	Questions	Answered	Successfully Answered	URL
3902	EJBTHREE-646	Its false that a security-domain is unsupported at the ejb level. This shows up in the container-configuration: <!ELEMENT container-configuration (container-name , call-logging?, invoker-proxy-binding-name?, sync-on-commit-only?, insert-after-ejb-post-create?, call-ejb-store-on-clean?, store-not-flushed?, container-interceptors? , instance-pool? , instance-cache? , persistence-manager? , web-class-loader? , locking-policy? , container-cache-conf? , container-pool-conf? , commit-option? , optiond-refresh-rate? , security-domain?, cluster-config?, depends*)>	1	0	https://issues.jboss.org/browse/EJBTHREE-646
3903	EJBTHREE-585	Had a conf call last night with Brian, Bela, Steve and Emmanuel. Here are 3 points conclusion: Sounds good. One additional option we might want to look into is #4: invalidation. It may be interesting to see whether we can (a) replicate the full state including class and primary ID (if any) the first time and then (b) invalidate on write access (NOT read access !) subsequently. Invalidation would be nice because if we use JBossCache as 2nd level cache with Hibernate, then we're replicating each entity bean twice... Ben Wang wrote: Thanks for the time last time. Like I said, I'd like to check into item #3 first to see exacly what it takes to instrument Hibernate. I can see I will need Steve's help in this regards. If that doesn't work, I will then work with Steve to come up an API for #2. Cheers, -Ben ---------- From: Bela Ban Sent: Wednesday, June 28, 2006 11:28 PM To: Ben Wang; Brian Stansberry; Steve Ebersole; Emmanuel Bernard Subject: Minutes We essentially have 3 avenues to pursue: Full serialization for entity beans (replicate everything) Delta-based replication Hibernate provides a call which returns the deltas (changes) within the current scope (transaction, session?) We can tell Hibernate to apply the deltas (on the receiving side) Field-based replication Look into feasibility of instrumenting *everything* (persistence context etc) Steve: can we remove the 'transient' modifiers ? I don't think this should have an effect on performance, and documents that PojoCache needs to replicate those field	1	0	https://issues.jboss.org/browse/EJBTHREE-585
3904	EJBTHREE-582	Do not backport this to Branch_4_0. This will be done when we do a 4.0.5 release or another ejb3 release.	1	1	https://issues.jboss.org/browse/EJBTHREE-582
3905	EJBTHREE-436	Same problem: to stay in my example, 3 different xml's would be needed, each for one cluster and so it would still not be possible to deploy the same .ear on all clusters. How would you distribute a clustered ejb3 example? You force me to setup a cluster with the same partition name as yours? What if my 3 colleagues in the next room want to deploy your example, too on their computers. Would we have to setup individual subnetworks so they can all run jboss cluster wkith the same partition name?	1	0	https://issues.jboss.org/browse/EJBTHREE-436
3906	EJBTHREE-369	Should we raise this to the expert group on who should flush() container or persistence provider? BTW, flush_before_complete is set to default by default in the JBoss EJB 3.0 implementation.	1	1	https://issues.jboss.org/browse/EJBTHREE-369
3907	EJBTHREE-343	Hi There can be two ways of getting rid of INFO msgs: First one is to use other level of logging like DEBUG,ERROR in the source code instead of using INFO level for all the messages. Second one is to change the logging level in JBOSS .The log4j configuration is loaded from the jboss server conf/log4j.xml file. You can increase/descrease the level of logging for various jboss subsytems by adding category threshold statements to the conf/log4j.xml configuration file with the desired priority value: <category name="org.jboss.management"> <priority value="INFO"/> </category> The priority value can be one of: ERROR - error messages WARN - borderline error messages INFO - general status messages DEBUG - verbose debug level messages TRACE - very verbose Tell me if this is the one you expected? TIA. Jyoti.	1	1	https://issues.jboss.org/browse/EJBTHREE-343
3908	EJBTHREE-281	Would it be possible to add a default security domain on the jboss-app.xml as well? Since the ear usually represents an entire application, with several modules, this would be very useful, if possible.	1	1	https://issues.jboss.org/browse/EJBTHREE-281
3909	EJBTHREE-212	Apologies up front if this sounds stupid. I am using Struts-EJB3 on a new production project, but no injection using @EJB or @PersistenceContext seems to work with 4.0.4 so I can not create my SessionBean or use JTA (using oracle db). Is this related to this issue, or am I missing something - r is there some work around (service pojo mentioned elsewhere on forum). Any, repeat, any help, greatly appreciated. Paul	1	1	https://issues.jboss.org/browse/EJBTHREE-212
3910	EJBTHREE-204	This issue seems reopen. I downloaded jboss 4.0.4 GA and jboss-EJB-3.0_RC8-FD. I installed jboss and ejb3 support than I tryed to modify the standard port 3873 so I changed the file deploy/ejb3.deployer/META-INF/jboss-service.xml: <mbean code="org.jboss.remoting.transport.Connector" xmbean-dd="org/jboss/remoting/transport/Connector.xml" name="jboss.remoting:type=Connector,name=DefaultEjb3Connector,handler=ejb3"> <depends>jboss.aop:service=AspectDeployer</depends> <attribute name="InvokerLocator">socket://${jboss.bind.address}:10009</attribute> <attribute name="Configuration"> <handlers> <handler subsystem="AOP">org.jboss.aspects.remoting.AOPRemotingInvocationHandler</handler> </handlers> </attribute> </mbean> While the server listens correctly on new port the client has no way to contact the bean I deployed. Using strace I figured out that the client ignores the new port and tries to contact port 3873 as before RC3 of ejb3 support (while another installation with ejb3 rc3 works without problems). Is this a regression? Regards. Andrea	1	0	https://issues.jboss.org/browse/EJBTHREE-204
3911	EJBTHREE-142	can you elaborate on what you want for an API for this?	0	0	https://issues.jboss.org/browse/EJBTHREE-142
3912	EJBTHREE-105	I'd suggest to enable the deployer to provide the datasource name and the hibernate dialect in the ear archive (jboss-application.xml). How can one have one ear talk to a postgres and another talk to a oracle else?	0	0	https://issues.jboss.org/browse/EJBTHREE-105
3913	EJBBOOK-41	Build instructions are documented: http://community.jboss.org/docs/DOC-15565 What's nontrivial about "mvn clean install"? The build is already configured once Maven is installed.	1	0	https://issues.jboss.org/browse/EJBBOOK-41

#	Requirement	Questions	Answered	Successfully Answered	URL
3914	EJBCLIENT-161	----- T E S T S --- Running org.jboss.ejb.client.ClasspathConfigBasedSelectorTestCase Tests run: 1, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 0.759 sec <<< FAILURE! - in org.jboss.ejb.client.ClasspathConfigBasedSelectorTestCase testRemotingEJBReceiver(org.jboss.ejb.client.ClasspathConfigBasedSelectorTestCase) Time elapsed: 0.124 sec <<< ERROR! java.lang.ExceptionInInitializerError at org.jboss.ejb.client.ClasspathConfigBasedSelectorTestCase.testRemotingEJBReceiver(ClasspathConfigBasedSelectorTestCase.java:96) Caused by: java.lang.IllegalArgumentException: Invalid logger interface org.jboss.remoting3.remote.RemoteLogger (implementation not found in sun.misc.Launcher$AppClassLoader@75b84c92) at org.jboss.ejb.client.ClasspathConfigBasedSelectorTestCase.testRemotingEJBReceiver(ClasspathConfigBasedSelectorTestCase.java:96) Running org.jboss.ejb.client.MultipleConnectionsConfigBasedSelectorTestCase Tests run: 1, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 0.804 sec <<< FAILURE! - in org.jboss.ejb.client.MultipleConnectionsConfigBasedSelectorTestCase testRemotingEJBReceivers(org.jboss.ejb.client.MultipleConnectionsConfigBasedSelectorTestCase) Time elapsed: 0.12 sec <<< ERROR! java.lang.ExceptionInInitializerError at org.jboss.ejb.client.MultipleConnectionsConfigBasedSelectorTestCase.testRemotingEJBReceivers(MultipleConnectionsConfigBasedSelectorTestCase.java:95) Caused by: java.lang.IllegalArgumentException: Invalid logger interface org.jboss.remoting3.remote.RemoteLogger (implementation not found in sun.misc.Launcher$AppClassLoader@75b84c92) at org.jboss.ejb.client.MultipleConnectionsConfigBasedSelectorTestCase.testRemotingEJBReceivers(MultipleConnectionsConfigBasedSelectorTestCase.java:95) Running org.jboss.ejb.client.serialization.ProxySerializationTestCase Tests run: 1, Failures: 0, Errors: 0, Skipped: 0, Time elapsed: 0.569 sec - in org.jboss.ejb.client.serialization.ProxySerializationTestCase Running org.jboss.ejb.client.test.compression.CompressionTestCase Tests run: 7, Failures: 0, Errors: 7, Skipped: 0, Time elapsed: 0.754 sec <<< FAILURE! - in org.jboss.ejb.client.test.compression.CompressionTestCase testDataCompressionImplicitlyDisabled(org.jboss.ejb.client.test.compression.CompressionTestCase) Time elapsed: 0.357 sec <<< ERROR! java.lang.ExceptionInInitializerError at org.jboss.ejb.client.test.compression.CompressionTestCase.testDataCompressionImplicitlyDisabled(CompressionTestCase.java:120) Caused by: java.lang.IllegalArgumentException: Invalid logger interface org.jboss.remoting3.remote.RemoteLogger (implementation not found in sun.misc.Launcher$AppClassLoader@75b84c92) at org.jboss.ejb.client.test.compression.CompressionTestCase.testDataCompressionImplicitlyDisabled(CompressionTestCase.java:120) testDataCompressionExplicitlyDisabled(org.jboss.ejb.client.test.compression.CompressionTestCase) Time elapsed: 0.01 sec <<< ERROR! java.lang.NoClassDefFoundError: Could not initialize class org.jboss.remoting3.remote.RemoteLogger at org.jboss.ejb.client.test.compression.CompressionTestCase.testDataCompressionExplicitlyDisabled(CompressionTestCase.java:148) testDataCompressionEnabled(org.jboss.ejb.client.test.compression.CompressionTestCase) Time elapsed: 0.013 sec <<< ERROR! java.lang.NoClassDefFoundError: Could	1	1	https://issues.jboss.org/browse/EJBCLIENT-161
3915	EJBCLIENT-161	Which version are you trying to build?	1	1	https://issues.jboss.org/browse/EJBCLIENT-161
3916	EJBCLIENT-76	Hello [~jaikiran] can these setting be configured via xml like jboss-ejb3.xml to specify a method/methods to apply the settings?	1	1	https://issues.jboss.org/browse/EJBCLIENT-76
3917	EJBCLIENT-66	Dennis, how are you using that selector in your code? Or more specifically what does your code look like irrespective of whether you are directly using that selector or it's being picked up internally. Which version of EJB client API are you using?	1	0	https://issues.jboss.org/browse/EJBCLIENT-66
3918	EJBCLIENT-47	Will this be in EAP 7.0 or 7.1 release?	0	0	https://issues.jboss.org/browse/EJBCLIENT-47
3919	EJBCLIENT-34	This needs a bit more thought and discussion. Some things to consider are: 1) What happens when the application code has a jboss-ejb-client.properties file in the classpath (which triggers auto creation of a EJB client context) and the code also uses a map based InitialContext to setup a EJB client context? What's supposed to happen in that case? An application is expected to have just a single EJB client context. So what are we going to do with either of these client contexts that get created differently. 2) Consider the following case that can happen within the same application: {code} public void someMethod() { Properties props = new Properties(); props.put("endpoint.name", "foo-bar"); props.put("remote.connection.foo.host", "localhost"); props.put("remote.connection.foo.port", "4447"); Context = new InitialContext(props); {code} {code} public void someOtherMethod() { Properties props = new Properties(); props.put("endpoint.name", "foo-bar"); props.put("remote.connection.bar.host", "localhost"); props.put("remote.connection.bar.port", "5557"); Context = new InitialContext(props); } {code} What are we supposed to do here? We can't setup two different EJB client contexts. The client context lifespan will be the same as the application lifespan and there will be just one EJB client context. So when do we create it? The "first" initialcontext creation (which in itself isn't well defined)? There are many more things to consider but this is just a start. We need to make sure what we want to support and how that might conflict/impact with what we already have. We should avoid the confusion that's already there related to the various ways (remote-naming and ejb-client API) for remote EJB invocations.	1	1	https://issues.jboss.org/browse/EJBCLIENT-34
3920	EJBCLIENT-34	Could you consider to replace static selector with InheritableThreadLocal. This will solve a lot of problems and will make whole think Multithread ready!	1	1	https://issues.jboss.org/browse/EJBCLIENT-34
3921	EJBCLIENT-33	Is there any easy way to reproduce these clients getting stuck?	1	1	https://issues.jboss.org/browse/EJBCLIENT-33
3922	EJBCLIENT-33	In order to cherrypick this in the EAP ER2 build, we would need a resolution by Feb/27th. Is this doable?	0	0	https://issues.jboss.org/browse/EJBCLIENT-33
3923	EJBCLIENT-33	Jaikiran, I see you changed Coding In Progress, can you describe the strategy and design decisions here? Thanks.	1	1	https://issues.jboss.org/browse/EJBCLIENT-33
3924	EJBCLIENT-12	How do you provide the configuration without using a jboss-ejb-client.properties file then ? You pseudo-code does not appear to be suitable when using https://docs.jboss.org/author/display/AS71/EJB+invocations+from+a+remote+client+using+JNDI	1	1	https://issues.jboss.org/browse/EJBCLIENT-12
3925	EJBCLIENT-12	In ALL our applications the user is shown a login dialog with username, password, host address and port name. This is very essential for our customers because they have to be able to connect to different servers. How can this be achieved using just a static file? The file is just a property file so it would be fine if we can provide a Properties object somehow somewhere. Is there a way or is there at least a plan for doing so?	0	0	https://issues.jboss.org/browse/EJBCLIENT-12
3926	EJBCLIENT-12	{quote} In ALL our applications the user is shown a login dialog with username, password, host address and port name. This is very essential for our customers because they have to be able to connect to different servers. How can this be achieved using just a static file? {quote} Security aspects for a EJB remote client are currently being considered for 7.1.0.CR1. We'll take into account this usecase too.	1	1	https://issues.jboss.org/browse/EJBCLIENT-12

#	Requirement	Questions	Answered	Successfully Answered	URL
3927	EMBJOPR-362	The process for the admin-console release should be: 1) Tag RHQ 3.0.0.EmbJopr4 off http://git.fedorahosted.org/git/?p=rhq/rhq.git;a=shortlog;h=refs/heads/AdminConsole_EAP_5_1 2) Tag EmbJopr 1.3.4.SP5 off https://svn.jboss.org/repos/embjopr/branches/AdminConsole_EAP_5_1 after updating its root pom to depend on RHQ 3.0.0.EmbJopr4 Larry, when do we need this by?	1	1	https://issues.jboss.org/browse/EMBJOPR-362
3928	EMBJOPR-349	Will this be added?	0	0	https://issues.jboss.org/browse/EMBJOPR-349
3929	EMBJOPR-344	Will this get fixed please?	1	0	https://issues.jboss.org/browse/EMBJOPR-344
3930	EMBJOPR-344	Heiko W. Rupp <hrupp@redhat.com> made a comment on [bug 824773|https://bugzilla.redhat.com/show_bug.cgi?id=824773] Mazz, can you please investigate the effort of a fix here?	0	0	https://issues.jboss.org/browse/EMBJOPR-344
3931	EMBJOPR-326	Will this get fixed please?	0	0	https://issues.jboss.org/browse/EMBJOPR-326
3932	EMBJOPR-322	Will this get fixed please?	0	0	https://issues.jboss.org/browse/EMBJOPR-322
3933	EMBJOPR-321	Can you please tell us your reasons?	1	1	https://issues.jboss.org/browse/EMBJOPR-321
3934	EMBJOPR-245	Will this get fixed? It causes a minor issue - exceptions when user leaves this page in browser and the server goes down.	0	0	https://issues.jboss.org/browse/EMBJOPR-245
3935	EMBJOPR-205	Will this get fixed?	0	0	https://issues.jboss.org/browse/EMBJOPR-205
3936	EMBJOPR-201	Will this be implemented? These tiny textareas are really annoying...	0	0	https://issues.jboss.org/browse/EMBJOPR-201
3937	EMBJOPR-190	Stacktraces still dumped for most resources. Will this be fixed?	0	0	https://issues.jboss.org/browse/EMBJOPR-190
3938	EMBJOPR-186	Hi Ondrej, we need the server.log files rather than the boot.log files. Can you attach those instead? The boot.log files stop before any web apps have gotten deployed so they are not useful for showing the difference between deploying with and without the embedded console.	1	0	https://issues.jboss.org/browse/EMBJOPR-186
3939	EMBJOPR-182	patch for Revision 481. this patch eliminates hard-coded messages in .xhtml files.	1	1	https://issues.jboss.org/browse/EMBJOPR-182
3940	EMBJOPR-182	Thanks for this! Unfortunately, I was unable to apply the patch. I get "An unknown line type was found in line 94!" when I try to apply the patch via TortoiseSVN. I'm guessing it is because the patch was not created against the very latest from embjopr trunk. Would you please try recreating the patch against HEAD? Thanks, Ian	1	1	https://issues.jboss.org/browse/EMBJOPR-182
3941	EMBJOPR-160	Will this be implemented?	0	0	https://issues.jboss.org/browse/EMBJOPR-160
3942	EMBJOPR-108	I based the Resource types I exposed off of the ConnectionFactoryTypes enum from org.jboss.deployers.spi.management.KnownComponentTypes: /** * Enums for generic ConnectionFactory types */ public enum ConnectionFactoryTypes { XA("ConnectionFactory", "XA"), NoTx("ConnectionFactory", "NoTx"); private final String type; private final String subtype; private ConnectionFactoryTypes(String type, String subtype) { this.type = type; this.subtype = subtype; } public String type() { return type; } public String subtype() { return subtype; } public ComponentType getType() { return new ComponentType(type, subtype); } }; If there is also a LocalTX type, then this enum should probably be updated. Scott, can you comment?	1	1	https://issues.jboss.org/browse/EMBJOPR-108
3943	EMBJOPR-107	5.1.0.CR1/server/default/deploy/hellothere-DUD_RETRY.war/] 15:58:13,379 ERROR [WebModule] Could not create JSR-77 WebModule: hellothere-DUD_RETRY.war javax.management.InstanceAlreadyExistsException: jboss.management.local:j2eeType=WebModule,J2EEServer=Local,J2EEApplication=null,name=hellothere-DUD_RETRY.war already registered. at org.jboss.mx.server.registry.BasicMBeanRegistry.add(BasicMBeanRegistry.java:756) at org.jboss.mx.server.registry.BasicMBeanRegistry.registerMBean(BasicMBeanRegistry.java:233) at sun.reflect.GeneratedMethodAccessor96.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:138) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:140) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.mx.server.MBeanServerImpl$3.run(MBeanServerImpl.java:1431) at java.security.AccessController.doPrivileged(Native Method) at org.jboss.mx.server.MBeanServerImpl.registerMBean(MBeanServerImpl.java:1426) at org.jboss.mx.server.MBeanServerImpl.registerMBean(MBeanServerImpl.java:376) at org.jboss.management.j2ee.WebModule.create(WebModule.java:148) at org.jboss.management.j2ee.deployers.WebModuleJSR77Deployer.deployJsr77(WebModuleJSR77Deployer.java:50) at org.jboss.management.j2ee.deployers.WebModuleJSR77Deployer.deployJsr77(WebModuleJSR77Deployer.java:38) at org.jboss.management.j2ee.deployers.AbstractVFSJSR77Deployer.deployJsr77(AbstractVFSJSR77Deployer.java:46) at org.jboss.management.j2ee.deployers.AbstractJSR77Deployer.deploy(AbstractJSR77Deployer.java:173) at org.jboss.deployers.spi.deployer.helpers.AbstractSimpleRealDeployer.internalDeploy(AbstractSimpleRealDeployer.java:62) at org.jboss.deployers.spi.deployer.helpers.AbstractRealDeployer.deploy(AbstractRealDeployer.java:50) at org.jboss.deployers.plugins.deployers.DeployerWrapper.deploy(DeployerWrapper.java:171) at org.jboss.deployers.plugins.deployers.DeployersImpl.doDeploy(DeployersImpl.java:1439) at org.jboss.deployers.plugins.deployers.DeployersImpl.doInstallParentFirst(DeployersImpl.java:1157) at org.jboss.deployers.plugins.deployers.DeployersImpl.install(DeployersImpl.java:1098) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:348) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:1598) at org.jboss.dependency.plugins.AbstractController.incrementState(AbstractController.java:934) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:1062) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:984) at	1	1	https://issues.jboss.org/browse/EMBJOPR-107
3944	EMBJOPR-103	But shouldn't we use the Embedded Jopr logo, rather than the Jopr one?	0	0	https://issues.jboss.org/browse/EMBJOPR-103
3945	EMBJOPR-101	I've uploaded the latest snapshot of JSFUnit 1.2.0. This snapshot uses HtmlUnit 2.7, which was just released yesterday. http://snapshots.jboss.org/maven2/org/jboss/jsfunit/jboss-jsfunit-core/1.2.0.GA-SNAPSHOT/ Can we give this a try and see if the problem is still there? Also, it would be great if you guys can try out the new <rich:tree> methods in RichFacesClient. See http://community.jboss.org/wiki/JSFUnitWithRichFaces. I've uploaded a snapshot for that as well. http://snapshots.jboss.org/maven2/org/jboss/jsfunit/jboss-jsfunit-richfaces/1.2.0.GA-SNAPSHOT/	1	1	https://issues.jboss.org/browse/EMBJOPR-101

#	Requirement	Questions	Answered	Successfully Answered	URL
3946	EMBJOPR-85	Instead of looking for a solution in the individual frameworks you are using (Seam, Hibernate), shouldn't this be the job of the application server? So for example, you could have a switch in jboss-web.xml that delays bootstrapping the web.xml (listeners that start Seam and then Hibernate) until the first request for that context is received.	0	0	https://issues.jboss.org/browse/EMBJOPR-85
3947	EMBJOPR-85	See the attached admin-console-war-seam-init.log.. The scanning of the admin-console.war's WEB-INF/classes dir is taking less than a second. The remaining time is spent scanning the Seam jars, as well as a whole bunch of other jars from the war's WEB-INF/lib which definitely do not contain seam.properties. It seems like this may be a Seam bug. Is a seam.properties in a war's WEB-INF/classes dir supposed to indicate Seam should scan only the classes in WEB-INF/classes, or does it mean it should scan every single jar in WEB-INF/lib, even ones without seam.properties? I was assuming the former.	1	1	https://issues.jboss.org/browse/EMBJOPR-85
3948	EMBJOPR-85	The next issue is that somehow the entire .war deployment is being scanned from its location in the /tmp folder. This includes all .jar files in WEB-INF/lib and even jars not on the classpath e.g. in /plugins. This adds another about 7seconds to the startup time as you can see below. Is this a bug in Seam/VFS? Or is there a way to configure things so this doesn't happen? Assigning to Jason to see if he can get this answered. 18:06:23,265 TRACE [VFSScanner] Root url: file:/C:/opt/jboss-5.1.0.CR1-SNAPSHOT/server/default/tmp/aaaa-mk2ugc-fthpqrht-1-fthprbee-s/jbas5-admin-console.war/ 18:06:23,265 TRACE [VFSScanner] File: C:\opt\jboss-5.1.0.CR1-SNAPSHOT\server\default\tmp\aaaa-mk2ugc-fthpqrht-1-fthprbee-s\jbas5-admin-console.war 18:06:23,265 TRACE [VFSScanner] URL: file:/C:/opt/jboss-5.1.0.CR1-SNAPSHOT/server/default/tmp/aaaa-mk2ugc-fthpqrht-1-fthprbee-s/, relative: jbas5-admin-console. war 18:06:23,265 TRACE [VFSScanner] Top: FileHandler@16722742[path=jbas5-admin-console.war context=file:/C:/opt/jboss-5.1.0.CR1-SNAPSHOT/server/default/tmp/aaaa-mk2 ugc-fthpqrht-1-fthprbee-s/ real=file:/C:/opt/jboss-5.1.0.CR1-SNAPSHOT/server/default/tmp/aaaa-mk2ugc-fthpqrht-1-fthprbee-s/jbas5-admin-console.war/] 18:06:23,499 TRACE [Scanner] found css/console-style.css 18:06:23,515 TRACE [Scanner] found css/form-style.css 18:06:23,515 TRACE [Scanner] found error.xhtml 18:06:23,515 TRACE [Scanner] found errorLayout.xhtml ... 18:06:23,530 TRACE [Scanner] found plugins/jopr-jboss-as-5-plugin-2.2.0-SNAPSHOT.jar/lib/testng-5.7-jdk15.jar/META-INF/LICENSE.txt 18:06:23,530 TRACE [Scanner] found plugins/jopr-jboss-as-5-plugin-2.2.0-SNAPSHOT.jar/lib/testng-5.7-jdk15.jar/META-INF/MANIFEST.MF ... 18:06:24,937 TRACE [Scanner] found WEB-INF/lib/commons-digester-1.8.jar/org/apache/commons/digester/ExtendedBaseRules$1.class 18:06:24,937 TRACE [Scanner] found WEB-INF/lib/commons-digester-1.8.jar/org/apache/commons/digester/ExtendedBaseRules.class ... 18:06:26,796 TRACE [Scanner] found WEB-INF/lib/rhq-core-domain-1.2.0-SNAPSHOT.jar/org/rhq/core/domain/measurement/MeasurementBaseline.class 18:06:26,796 TRACE [Scanner] found WEB-INF/lib/rhq-core-domain-1.2.0-SNAPSHOT.jar/org/rhq/core/domain/measurement/MeasurementCategory.class ... 18:06:30,530 TRACE [Scanner] found WEB-INF/lib/richfaces-ui-3.3.0.GA.jar/org/richfaces/taglib/html/jsp/GraphValidator.class 18:06:30,530 TRACE [Scanner] found WEB-INF/lib/richfaces-ui-3.3.0.GA.jar/org/richfaces/ui.pack.js	1	1	https://issues.jboss.org/browse/EMBJOPR-85
3949	EMBJOPR-85	Can you try with JBoss-Seam-Int 5.2.0-SNAPSHOT?	1	1	https://issues.jboss.org/browse/EMBJOPR-85
3950	EMBJOPR-85	Jason's lazy init fix works nicely for me. The approximately 10s Seam init is deferred until the first request is made to the admin-console webapp. The only disadvantage is that the user must wait an extra 10s or so for the login page to load for the first request. Patch committed - r314. Ales, would you please attach a copy of the JBoss-Seam-Int 5.2.0-SNAPSHOT jar? When I test with it, what am I looking for - improved Seam init time due to fixes in the scanning code?	0	0	https://issues.jboss.org/browse/EMBJOPR-85
3951	EMBJOPR-72	(12:07:01 PM) ccrouch: ok first off, we had a big discussion a few months back with stan (12:07:22 PM) ccrouch: about how we should be using id's not xpath for identifying elements (12:08:00 PM) ccrouch: is that still QA's preference? or you good with xpath now? (12:11:04 PM) ozizka: ccrouch: There are cases when IDs are not an option (12:11:19 PM) ozizka: Like when working with tables etc. (12:11:29 PM) ccrouch: ok fair enough (12:11:48 PM) ccrouch: the change you used as an example there, looks fairly harmless (12:12:10 PM) ccrouch: as you say as long as it doenst impact form or function, i'm fine with it (12:12:53 PM) ozizka: ccrouch: Ok, thanks	1	1	https://issues.jboss.org/browse/EMBJOPR-72
3952	EMBJOPR-56	Hi Farah, could you take a look at what options we have to implement this. Richfaces presumably could help here.	1	1	https://issues.jboss.org/browse/EMBJOPR-56
3953	EMBJOPR-56	I added this label under the "Previous Operations" label on the Control tab. I've attached a screenshot. Does it look ok? Should I commit?	1	1	https://issues.jboss.org/browse/EMBJOPR-56
3954	EMBJOPR-56	I've been looking into why the a4j:poll component in resourceInstanceOperation.xhtml doesn't currently update the "Previous Operations" table as it should. Here's an update about what I've found so far: The toggle panel in the "Previous Operations" table has many conditionally rendered components. This may explain why the a4j:poll component isn't working: See section 1.59. "Why does reRendering fail? Hide/Show components using rendered." http://www.jboss.org/file-access/default/members/jbossrichfaces/freezone/docs/devguide/en/faq/faq.html I tried using a4j:outputPanel components to wrap conditionally rendered components, as suggested in the link. This allows the text to update (i.e. "In Progress" changes to "Successful"). However, the "Show/Hide Details" link stops working completely. Clicking on the link results in a redirect to the index page.	0	0	https://issues.jboss.org/browse/EMBJOPR-56
3955	EMBJOPR-37	Does the underlying DeploymentManager we are using have support for this?	1	1	https://issues.jboss.org/browse/EMBJOPR-37

#	Requirement	Questions	Answered	Successfully Answered	URL
3956	EMBJOPR-19	Discussion on this topic: (14:57:02) jess: ping spritz - which browser did you test the embedded config pages with? (14:57:32) spritz: I have Firefox 2.0.0.4 (14:57:47) jess: hmm, me too (14:58:04) jess: but when I go to http://10.11.14.163:8080/embedded-console/secure/resourceView.seam?path=-34 -- in Firefox - the page gets an invalid XML error (14:58:09) jess: but it works in opera (15:02:10) spritz: Funny, I know what is th ematter, it is something that was in our layout.xhtml which someone must have just recently overwritten ro took out. It was the same issue we had with s before (15:03:09) spritz: Yeah, it is no longer there. it was an attribute as part of the <f:view> tag (15:03:30) jess: oh - that xhtml thing (15:04:06) spritz: that contentType="text/html" (15:04:07) spritz: thing (15:05:24) ccrouch: the problem is that an extra div is getting rendered somehow (15:05:49) ccrouch: the way to remove the symptom is to get FF to not parse it as xhtml (15:06:07) ccrouch: but that doesn't remove the underlying problem (15:08:30) spritz: ccrouch: extra divs is under ips's jurisdiction here. :-D (15:10:36) ips: where are we seeing extra divs? (15:10:46) ccrouch: jess: if you right click on that error message you can get the full src, you might be able to chop of the xml declaration and get it to render (15:11:14) ips: the problem on the config pages were spans that were closed too early (15:11:22) jess: thx ccrouch I was able to add the line spritz mentioned and it fixed it (15:12:23) spritz: jess, can you put the error back to show ips? (15:12:35) ccrouch: it didn't fix it, it just doesnt tell you about the error anymore :-/ (15:12:45) ccrouch: spritz: +1 (15:12:54) ccrouch: we need to have a jira for that (15:13:17) jess: ok - I'll add a jira (15:14:19) ips: ccrouch: from what i read in a few places on the Web, it's recommended to have Firefox in trnasitional mode, rather than strict mode (15:14:44) ips: it renders certain things in wacky ways in strict mode (15:14:45) ccrouch: it is (15:14:45) ccrouch: <!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN" "http://www.w3.org/TR/xhtml1/DTD/xhtml1-transitional.dtd" > (15:15:03) ccrouch: this is the fact that the doc ain't well formed xml (15:15:40) ips: ccrouch: but FF bases what mode it uses on the content-type of the response (15:15:50) ips: not the doctype inside the response body (15:16:04) ips: so i think we want to keep the contentType="text/html" (15:16:34) ccrouch: how do you specify strict vs transitional if not through the DOCTYPE? (15:16:36) ips: instead, we should enable xhtml validation on the server side (15:16:51) ccrouch: ok, fair enough (15:16:57) ips: ie - via some Facelets setting, if possible (15:17:13) ips: yeah, you would think it would use doctype (15:17:23) ips: but it uses the content-type header (15:17:29) ccrouch: i just want it to be validated as well formed xml, dangling tags just means we've screwed up the rendering somewhere	1	1	https://issues.jboss.org/browse/EMBJOPR-19
3957	EMBJOPR-19	I'm pretty sure I fixed this last night (there was a missing closing td and tr tag). Jess or Charles, would you please verify?	1	1	https://issues.jboss.org/browse/EMBJOPR-19
3958	WOLF-57	How would you like to present wolf reports in jenkins like junit reports? In jenkins, there is summary view for tests results, which use table with tree columns (test name, duration and age) and for exception message is necessary to click on details. Which is fine for junit tests, but for wolf reports this would be uncomfortable. I can try encode all information into test name, but is it what you want?	1	1	https://issues.jboss.org/browse/WOLF-57
3959	WOLF-5	I walked through jar files, and found only one thing that I want ask. There are few jars that are not compiled, but contain .java sources in it. Is it intended to be, or they got from eap into repository by accident? Specifically few libraries from /seam/lib/src/ jars: jboss-seam-debug-2.2.2.EAP5-eap5_1_0.jar <-> jboss-eap-5.1/seam/lib/src/jboss-seam-debug-sources.jar jboss-seam-mail-2.2.2.EAP5-eap5_1_0.jar <-> jboss-eap-5.1/seam/lib/src/jboss-seam-mail-sources.jar jboss-seam-ui-2.2.2.EAP5-eap5_1_0.jar <-> jboss-eap-5.1/seam/lib/src/jboss-seam-ui-sources.jar jboss-seam-resteasy-2.2.2.EAP5-eap5_1_0.jar <-> jboss-eap-5.1/seam/lib/src/jboss-seam-resteasy-sources.jar jboss-seam-remoting-2.2.2.EAP5-eap5_1_0.jar <-> jboss-eap-5.1/seam/lib/src/jboss-seam-remoting-sources.jar For example jboss-seam-pdf-2.2.2.EAP5-eap5_1_0.jar is with class files from jboss-eap-5.1/seam/lib/jboss-seam-pdf.jar	0	0	https://issues.jboss.org/browse/WOLF-5
3960	WOLF-3	I see that license information was added to the POM, thanks! Can you also add a comment which includes the local path to the file in the EAP zip? For example, the pom for org.jboss.logging:jboss-logging-jdk should contain a comment with the file path. {code:xml} <!-- File: jboss-as/lib/jboss-logging-jdk.jar --> {code}	1	1	https://issues.jboss.org/browse/WOLF-3
3961	ENVERS-60	Would you be interested in contributing this feature?	0	0	https://issues.jboss.org/browse/ENVERS-60
3962	ENVERS-49	Hello, do you maybe have an example UserType that I can use for testing? Thanks, Adam	1	1	https://issues.jboss.org/browse/ENVERS-49
3963	ENVERS-26	Ah, I forgot about this Hibernate feature. So far I only tested mappings with @OneToMany, and without a @JoinColumn - in this case, Hibernate creates a third relationship table to maintain the association. Now, if you add a @JoinColumn, the database schema is exactly the same as if you used a @ManyToOne + @OneToMany(mappedBy="") combination, except that the association is owned by a different side. Is that what you wanted to achieve? I'll check the @OneToMany+@JoinColumn combination. -- Adam	1	1	https://issues.jboss.org/browse/ENVERS-26
3964	ENVERS-26	Hello, there is one problem with storing the relation in the table of the entity (as is the case when @OneToMany + @JoinColumn are used). Let's say that entity A is on the many end, and entity B on the one end. Then, B has a collection of A's. Now, if you add a new A to an instance of B, the property ID_B in table A is modified. To version it, a new revision has to be generated - for A. However, nothing in the A-bean changed, as the relation is unidirectional. All fields have the same values as before - so no revision should be generated. One solution to this problem is to store the relation in a third table. This wouldn't directly mimick the hibernate mapping, but it wouldn't generate unnecessary revisions. There is a problem here to: a new annotation would have to be introduced (@VersionsJoinTable), to be able to set the name of the third table. What do you think? Any other ideas? Adam	1	1	https://issues.jboss.org/browse/ENVERS-26
3965	ERRAI-1050	As far as I can tell, these dependencies are brought in through gwt-dev. Is it really a good idea to exclude these?	0	0	https://issues.jboss.org/browse/ERRAI-1050
3966	ERRAI-1005	It would be great to renaming classes (obfuscate or crypt them) because the requests between client are too verbose (security). Is there any option for that ?	0	0	https://issues.jboss.org/browse/ERRAI-1005
3967	ERRAI-983	I don't understand this request. Why get rid of dom4j? And what does getting rid of dom4j have to do with Java 8?	1	1	https://issues.jboss.org/browse/ERRAI-983
3968	ERRAI-979	"the client bus will negotiate for the token when first connecting." Is this the same mechanism as in https://github.com/google-web-toolkit/gwt.svn/blob/master/user/src/com/google/gwt/user/server/rpc/XsrfTokenServiceServlet.java which the official docs still declare as experimental and where you are told not to use it in production? My thought is this: is the first request of the xsrf token secured against XSRF attacks? It doesn't look like that to me.	1	0	https://issues.jboss.org/browse/ERRAI-979
3969	ERRAI-971	[~psiroky], I have fixed all the Errai JPA issues on [this branch|https://github.com/mbarkley/errai/commits/ip-bom-7.0.0.CR5]. Does the merger of this change need to be coordinated with the ip-bom upgrade in the workbench?	1	1	https://issues.jboss.org/browse/ERRAI-971

#	Requirement	Questions	Answered	Successfully Answered	URL
3970	ERRAI-960	Guys, Thanks for cleaning up the POMs! Now, i hope I don't have to filter out so many jars like in my Gradle builds any more: ========== subprojects { apply plugin: 'java' apply plugin: 'maven' ext.compileExcludes = { exclude group: 'com.google.gwt' exclude group: 'com.google.guava' exclude group: 'javax.enterprise' exclude group: 'javax.annotation' exclude group: 'javax.inject:javax' exclude group: 'org.keycloak' exclude group: 'org.jboss.resteasy' exclude group: 'dom4j' exclude group: 'org.lesscss' exclude group: 'org.jboss.spec' exclude group: 'javax.servlet' exclude group: 'javax.inject' exclude group: 'javax.el' exclude group: 'junit' exclude group: 'org.hibernate.javax.persistence' exclude group: 'org.jboss.spec.javax.interceptor' exclude group: 'org.jboss.spec.javax.el' exclude group: 'org.jboss.spec.javax.ws.rs' exclude group: 'org.jboss.spec' exclude group: 'org.hibernate' exclude group: 'asm' exclude group: 'org.apache.httpcomponents' } ========== There are still a few questions though: 1. Why can't we just declare dependency on the Java EE single jar instead of those proprietary JBoss integration BOMs? Here is what I have been doing for a while now in my Gradle builds with no problems: ========== .... //Java EE - has all APIs that Errai 4 needs providedCompile "javax:javaee-api:7.0" ========== 2. Guice. Can we get rid of it for good and declare Errai to be Java EE/CDI/Weld product. Guice is no longer maintained by google and brings in a nasty dependency of an old asm jar (v3), which can collide with the one used in GWT 2.8 (v5) 3. I see that Errai (code generators?) uses Apache Stombol - another obscure and dead-in-the-water product the was last updated in 2014. Not only it brings a tons of jars during the (usually painfully slow) gwt compilation, but also some ancient and nasty xml jars too (dom4j, xml-apis). My deep hunch is that Errai only need (the excellent) jsoup.jar here, but I have not dug into the code deep enough to be confident ... Just my 2c. Thanks!	0	0	https://issues.jboss.org/browse/ERRAI-960
3971	ERRAI-939	Hi [~mbarkley] - Im trying to run the tutorial app, and can see it is still referencing WF8 - but Im very interested in using WF10. Is there any docs on how to run Errai with WF10? Also, any plans to update the tutorial? Thanks	1	0	https://issues.jboss.org/browse/ERRAI-939
3972	ERRAI-932	Max, Here is the confusing diagnostic: *java.lang.IllegalArgumentException: Element to be wrapped must not be null - Did you forget to initialize or @Inject a UI field* It does not say; - which element - in which template/class	1	1	https://issues.jboss.org/browse/ERRAI-932
3973	ERRAI-932	[~hr.stoyanov], I've updated the first two error messages you described so that they reference the templated bean type and the data field name involved in the error. Regarding the third error, I'm not sure how or where this error occurs. Did you see this when you forgot to annotate a type as @Templated? I'm leaving this issue open for now in case you are able to provide more information about that particular error.	0	0	https://issues.jboss.org/browse/ERRAI-932
3974	ERRAI-925	Hi [~swiderski.maciej], This issue is fixed in Errai 4 (was closed in [this issue|https://issues.jboss.org/browse/ERRAI-887]), but wasn't back-ported to Errai 3.2 (though probably it should have been). It should be possible to use the EmbeddedWildFlyLauncher from Errai 4 with Errai 3 by changing the version of errai-cdi-jboss. Is this a good enough workaround?	1	0	https://issues.jboss.org/browse/ERRAI-925
3975	ERRAI-885	When can you push a release with the patch ... Keycloak is pretty unusable without it. Maybe just 3.2.2.Final ?	0	0	https://issues.jboss.org/browse/ERRAI-885
3976	ERRAI-875	Hi Slava, Could you please your use case? If you have a compelling reason for accessing these properties it might make more sense to have them available through some more public API (such as a new method on the DataBinder). Cheers.	0	0	https://issues.jboss.org/browse/ERRAI-875
3977	ERRAI-861	Sure is! Do you want to send us a pull request for this change? This will get the feature in faster and makes sure you get the credits for it ;).	1	1	https://issues.jboss.org/browse/ERRAI-861
3978	ERRAI-840	Hi Paul, Can you share the error message / stack trace with us? Using "agent" as a model property should actually work fine. We have a test for this: https://github.com/errai/errai/blob/master/errai-data-binding/src/test/java/org/jboss/errai/databinding/client/TestModel.java	1	1	https://issues.jboss.org/browse/ERRAI-840
3979	ERRAI-829	I tried the above but I can't successfully bind to those Widgets (I assume) because they don't implement {{HasValue<T>}} but only {{TakesValue<T>}}? Would it be possible to support binding {{TakesValue<T>}}? Obviously in this case the binding is one-way but that's not really a problem since these Widgets are display-only?	1	1	https://issues.jboss.org/browse/ERRAI-829
3980	ERRAI-822	This IS broken in 3.1.1 or the documentation is wrong: {code} <module> <inherits name="org.jboss.errai.common.ErraiCommon"/> <inherits name="org.jboss.errai.bus.ErraiBus"/> <inherits name="org.jboss.errai.ioc.Container"/> <inherits name="org.jboss.errai.enterprise.CDI"/> <set-property name="gwt.logging.logLevel" value="ALL"/> </module> {code} {code} @EntryPoint public class Client { private static final Logger JUL = Logger.getLogger(Client.class.getName()); @Inject org.slf4j.Logger SLF4J; @PostConstruct public void onReady() { LoggingHandlerConfigurator config = LoggingHandlerConfigurator.get(); config.getHandler(ErraiConsoleLogHandler.class).setLevel(Level.ALL); config.getHandler(ErraiSystemLogHandler.class).setLevel(Level.ALL); config.getHandler(ErraiDevelopmentModeLogHandler.class).setLevel(Level.ALL); // That doesn't even exist but is still documented... // config.getHandler(ErraiFirebugLogHandler.class).setLevel(Level.ALL); JUL.severe("### JUL"); JUL.warning("### JUL"); JUL.info("### JUL"); JUL.fine("### JUL"); SLF4J.error("### SLF4J"); SLF4J.warn("### SLF4J"); SLF4J.info("### SLF4J"); SLF4J.debug("### SLF4J"); ... {code} In SDM all I see is SEVERE/ERROR messages in the browser console and a breakpoint in GWT ConsoleLogHandler shows that only SEVERE is logged. Also, there is a typo in the doc: "ErraiSimpleFormmater" Could please make this logging hackery optional in a separate module so those of us who like JUL can continue with standard logging and lower blood pressure?	1	1	https://issues.jboss.org/browse/ERRAI-822
3981	ERRAI-785	Sounds good! How did you implement the dynamic redeployment of changes?	1	0	https://issues.jboss.org/browse/ERRAI-785
3982	ERRAI-784	Do you have the project on GitHub so we can take a look? Are the tests based on Arquillian?	1	0	https://issues.jboss.org/browse/ERRAI-784
3983	ERRAI-742	As I commented in ERRAI-741: I've checked generated proxy code, and looks like you are right, performance hit should not be big, because newVal.equals(oldVal) is executed on model level, without touching converters/widgets. So priority of this request is not high :-) Also once again the question form 741: Why {code:java} String getAaa1() {code} is considered a valid getter, even without corresponding {code:java} void setAaa1(String value) {code} setter declared? And why {code:java} String getAaa1(String someParams) {code} is not considered as valid getter? Because non-bean property format convention, right?	1	0	https://issues.jboss.org/browse/ERRAI-742

#	Requirement	Questions	Answered	Successfully Answered	URL
3984	ERRAI-736	Thanks for that input! To 1. Unfortunately, in both, the CDI and websocket specifications is missing how websocket sessions / message-processings could / should work with CDI builtin (HTTP related) scopes. Because of that it necessary to use Weld Implementation classes to solve this. As i see we have 3 options: 1. let it as it is. (Okay in my opinion, but non-JEE containers is a very good point.) 2. "CDI'fying" of websocket in a separate module. (Best choice in my opinion, because then we may can also support other CDI implementations like Apache Open Webbeans for example in the future). 3. Remove the CDI'fying. (From my point of view and when you google around, there are many people there they are missing this feature in other frameworks or plattforms.) I think about to fix that with an additional (optional) module to CDI'fying the JSR-356 websocket module (errai-bus-jee7-websocket). Is that okay? To 2. ... i will do it so.	1	1	https://issues.jboss.org/browse/ERRAI-736
3985	ERRAI-732	I am interested in using Errai Security for contextual role based access. The use case would be User X is a member of Role Z constrained to Project Y granting them permission to view a a particular field related to Project Y. For example using Errai data binding a particular Errai UI Template is bound to a model and I would like to restrict access to a particular DataField on the composite autobound to the model based on whether or not the user belongs to a role related to that particular model. Using the role provider approach implementations of RequiredRolesProvider would not have access to the CDI InjectionPoint like a traditional CDI producer would. How could the model of interest be passed from the Errai UI Template to the RequiredRolesProvider? Does the RequiredRolesProvider need to provide role implementations for all possible permutations? If so that would be impractical. The test RequiredRolesProvider in source control all provide static role lists which could be accomplished the same way by setting custom role implementations on the user in the AuthenticationService implementation. Perhaps what I am trying to say is that is there a way to combine role and context information to allow for more complex security constraints, similar to Apache Shiro Permissions? http://shiro.apache.org/authorization.html#Authorization-PermissionBasedAuthorization	1	0	https://issues.jboss.org/browse/ERRAI-732
3986	ERRAI-721	Can you share your test with us so we can reproduce this? If changing to a for loop improves the situation significantly and doesn't slow down other use cases we can definitely merge your changes! If I understood correctly you're saying the improvement for your 1000 entities is ~30% (from 12s to ~8s). That would be a big win. However, we can't use a string switch statement because we can't make Java 7 a minimum requirement before Errai 3.1 So, we'll have to go with if/else statements instead.	1	0	https://issues.jboss.org/browse/ERRAI-721
3987	ERRAI-721	for Errai 4.0, we can use Switch, correct? Can we look at revisiting this?	0	0	https://issues.jboss.org/browse/ERRAI-721
3988	ERRAI-703	Great news, Marius! Thanks for trying it right away! Can you create a separate JIRA for the bug involving Errai UI including all details you have (stacktrace etc.)? Thanks!	0	0	https://issues.jboss.org/browse/ERRAI-703
3989	ERRAI-693	Can you please elaborate on this? It is unclear what feature you are requesting based on just the title.	0	0	https://issues.jboss.org/browse/ERRAI-693
3990	ERRAI-630	Hi Christian, you're totally right. Took me a couple of days to figure it out. The BatchCaller has no control over the callbacks being called, the bus does that. I also seriously doubt that one can practically stop the queued requests. Since javascript is single-threaded, once flush() has been called and it returned, the requests have been placed on the bus. Any cancel() after that will have no effect. Maybe if we make flush() a RepeatingCommand, but then again, in the majority of cases flush() will complete by the time cancel() is called. We could prevent the callbacks from being called by 1. turn the RemoteCallback interface into an abstract class (or create a CancellableRemoteCallback that implements RemoteCallback) - add a field that stores the cancelled state (initialised as false) - isCancelled - add a cancel method that sets isCancelled to true 2. When calling the callback method with the received data, we check isCancelled first (various ways to achieve that elegantly) 3. In the batch caller, on cancel we can iterate through the pendingCallbacks and set them cancelled. The bus will still receive the responses, but we can block the callbacks' execution. As you said, the need for supporting cancelling is questionable. In my case I do have a very specific need for it, however, I see now that I can achieve it locally in my own code by wrapping the RemoteCallback and the BatchCaller.	0	0	https://issues.jboss.org/browse/ERRAI-630
3991	ERRAI-626	Thank you soch much for adding this feature. It works as excpected. However, event.getProperyName() does not report the full binding path. It only contains the name of the Property. Would it be useful to add a method event.getFullBindingPath()?	1	0	https://issues.jboss.org/browse/ERRAI-626
3992	ERRAI-612	There's a twist that hasn't been mentioned yet: the lookup of the User object is asynchronous. In general, it could take a long time--minutes, even. The @PageShowing method will have to return (and so the page will get displayed) before the User object has been resolved. What will we put there? How will we update the DOM when the User object comes back? How will we display errors in case the async request fails? I think a lot of apps would benefit from a nice clean way of implementing this asynchronous lookup pattern while displaying some sort of "loading" message or locally cached data in the mean time. Also, a customizable error handler that's on by default would make Errai-based apps feel less flaky, since they would generally do something sane when the server request fails. Ideas? Maybe a method annotated with @PageDataResponse that gets the results of an RPC/JAX-RS call? Here's a straw man to get us started: {code} @Page public class AccountPage { @Model private User user; // this method will be called with the response that a Caller<UserService> would get from findUser(userId) // so it could be an Errai RPC or a JAX-RS request @PageDataRequest(service=UserService.class, method="findUser", args={"userId"}) @ModelSetter private void setUser(User user) { this.user = user; } @PageShowing public void loadUser() { // Load the user here if necessary. } } {code} The error handling would probably best be done with a global handler (straw man again): {code} @Templated @PageDataRequestErrorHandler public class PageDataRequestErrorPanel extends Panel implements BusErrorCallback { @Inject @DataField Label errorLabel; @Inject @DataField Button retryButton; // <--- it would be nice if we could make a retry button that works. ideas? @Override public boolean error(Message message, Throwable throwable) { errorLabel.setText("Error fetching data from server: " + throwable.getMessage()); } } {code}	0	0	https://issues.jboss.org/browse/ERRAI-612
3993	ERRAI-560	To me, it would be more annoying to deal with a framework that makes me use different annotations (@Page, @DefaultPage, @LoginPage) to express the same idea: "this class is a page." Do you have any other ideas that could make us both less annoyed? :)	0	0	https://issues.jboss.org/browse/ERRAI-560
3994	ERRAI-547	Why this wont fix?	1	1	https://issues.jboss.org/browse/ERRAI-547
3995	ERRAI-538	Guys, can we possibly revisit this? It would be a big help to get this. Is it complicated, or is there a reason NOT to do it?	1	1	https://issues.jboss.org/browse/ERRAI-538
3996	ERRAI-538	We could use GWT's UncaughtExceptionHandler for this and make sure RPC errors that are not handled by an error callback are passed to it. {code} GWT.setUncaughtExceptionHandler(new UncaughtExceptionHandler() { @Override public void onUncaughtException(Throwable e) { } }); {code} We can also create a declarative API for it: {code} @UncaughtException private void onUncaughtException(Throwable t) { } {code} Josh, what do you think?	1	1	https://issues.jboss.org/browse/ERRAI-538

#	Requirement	Questions	Answered	Successfully Answered	URL
3997	ERRAI-516	I think we had a slight adjustment to that code since this issue was filed. It's possible that's why there isn't a ClassCastException anymore. But I agree, even without the ClassCastException, the lifecycle methods are not being called in a sane order by the framework in this case! I think for the lifecycle methods to be reliable, and to save boilerplate to the maximum extent possible, we should aim for the lifecycle methods to be called like this when an attempt to display PageWithRedirects is made (assuming PageWithRedirects has all 4 lifecycle methods): {code} PageWithRedirects.beforeShow() // navigateTo SomePage PageWithRedirects.afterShow() // called because transition to PageWithRedirects is canceled PageWithRedirects.beforeHide() // called because transition to PageWithRedirects is canceled PageWithRedirects.afterHide() // called because transition to PageWithRedirects is canceled SomePage.beforeShow() SomePage.afterShown() {code} This way, any side effects within PageWithRedirects.beforeShow() or PageWithRedirects.afterShow() will be cleaned up as usual by the PageWithRedirects.beforeHide() or PageWithRedirects.afterHide() methods. A major drawback to this idea is that the docs on @PageShown and @PageHiding say these methods will be called when the page is attached to the DOM. Do we actually want to do this? In some ways, it seems excessive to attach the page to the DOM when we know we don't want it to display. On the other hand, not doing so will make the lifecycle contract too loose, and inflate the amount of boilerplate code necessary to implement a page lifecycle method correctly.	1	1	https://issues.jboss.org/browse/ERRAI-516
3998	ERRAI-516	I think it's better to satisfy the contract to @PageShown and @PageHiding and live with the small overhead adding/removing it to the DOM. History has to be taken care of too. Should PageWithRedirects appear in the browser's history stack? So that going back from SomePage will end in PageWithRedirects. Imho adding it to the History stack is the cleanest and simplest	1	1	https://issues.jboss.org/browse/ERRAI-516
3999	ERRAI-423	What do you mean exactly? Because you could bind a UL or LI like this: @DataField private Element liElement = DOM.createElement("li");	0	0	https://issues.jboss.org/browse/ERRAI-423
4000	ERRAI-423	Where do changes need to be made for this to happen? I've got a use-case where it would be nice to be able to directly template list items but they get translated into s currently.	1	1	https://issues.jboss.org/browse/ERRAI-423
4001	ERRAI-406	don't know how I missed that. Will it work for subcomponents? ("foo.bar")?	1	1	https://issues.jboss.org/browse/ERRAI-406
4002	ERRAI-388	Dear Mike, I tried aliasing in 2.1.0.Beta1 and it didn't work, as per https://community.jboss.org/thread/204604, which I should've included when I submitted this. Has this been fixed and does it preserve the ordering as specified by the LinkedHashMap / LinkedHashSet etc.? Thanks.	0	0	https://issues.jboss.org/browse/ERRAI-388
4003	ERRAI-303	Mike, Any thoughts on this one? Is my suggested change not advisable? Maybe it could be put in properties, or somehow externalized. As it stands I have added the modified to ServerMessageBusImpl to my classpath, so will make my errai upgrades more difficult. Thanks. PS: using websockets has made a substantial improvement to the user experience of my app.	0	0	https://issues.jboss.org/browse/ERRAI-303
4004	ERRAI-293	RPC batching is implemented now. The current API looks as follows: {code} public class MyBean { @Inject BatchCaller batchCaller; public MyBean() { batchCaller.call(remoteCallback1, RemoteService1.class).method1(); batchCaller.call(remoteCallback2, RemoteService2.class).method2(); batchCaller.sendBatch(); //batchCaller.sendBatch(remoteCallback3); } } {code} Josh, does this meet all your needs? We could build a higher lever API based on your suggestion but we'd need to qualify the injected callers so they know that they're in batch mode. {code} public class ChainedCaller { @Inject BatchCaller batchCaller; @Inject @Batched Caller<Service1> service1; @Inject @Batched Caller<Service2> service2; public ChainedCaller() { batchCaller.batch(service1).call(responseHandler1).method1(); batchCaller.batch(service2).call(responseHandler2).method2(); batchCaller.sendBatch(); //batchCaller.sendBatch(remoteCallback3); } } {code} I am not sure if the latter API is required since it is not really shorter or more expressive.	1	1	https://issues.jboss.org/browse/ERRAI-293
4005	ERRAI-275	Is there a specific issue you can describe as the impediment to marshalling these classes? Classes which are not marshallable by default can, in fact, be mapped manually. Please refer to our docs on this: http://docs.jboss.org/errai/2.0.CR1/errai/reference/html_single/#sid-5931328	1	0	https://issues.jboss.org/browse/ERRAI-275
4006	ERRAI-275	Here is the code snippet for which i am having Marshalling problem for api classes. public class ClientDTO extends com.extjs.gxt.ui.client.data.BaseModel implements com.extjs.gxt.ui.client.data.TreeModel { } I want to send instance of this class to My ErraiService as part of Message. i provided the mappings in erraiApp.properties as below errai.bus.serializableTypes=com.extjs.gxt.ui.client.data.BaseModel \com.extjs.gxt.ui.client.data.TreeModel But when i submit the message getting marshalling exceptions.	1	1	https://issues.jboss.org/browse/ERRAI-275
4007	ERRAI-275	Can you please share the exception details? Also you will have to make sure that all types BaseModel and TreeModel depend on are also listed in ErraiApp.properties.	1	0	https://issues.jboss.org/browse/ERRAI-275
4008	ERRAI-275	Would I be able to see the MarshallerFactoryImpl.java from the .errai/ directory in the root of your project?	1	0	https://issues.jboss.org/browse/ERRAI-275
4009	ERRAI-260	Is this still a valid issue ? I was now some time hanging around Forge project and investigating how other jboss projects deal with Forge plugins. There is already existing Archetypes subproject https://github.com/errai/archetypes with availability to generate a template project examples and IMHO other jboss projects forge plugins (eg: richfaces, arquilian https://github.com/forge) are like creating template project examples, so like errai Archetypes. Does it bring any value to Errai project having a forge plugin in case it would make nearly the same stuff like errai Archetypes ? If yes, I am probably missing here some points, could you explain to me a bit detail, what exactly is the goal of errai forge plugin. One of the reasons, I ask this is I have checked already existing https://github.com/errai/forge-errai/ and it does not look like a "typical" Forge plugin (anotated Command methods ... like https://github.com/forge/plugin-richfaces/blob/master/src/main/java/org/richfaces/forge/RichFacesPlugin.java) - might be already outdated first shot ?	1	1	https://issues.jboss.org/browse/ERRAI-260
4010	ERRAI-260	Hi Pavel, That would be amazing if you could work on this. To answer your questions: "existing https://github.com/errai/forge-errai/ and it does not look like a "typical" Forge plugin" -- it is very old, and AFAIK it was never finished. If you have a better starting point, then disregard this old code. "Does it bring any value to Errai project having a forge plugin in case it would make nearly the same stuff like errai Archetypes ?" -- yes, for sure: 1. adding Errai to an existing project that does not already have it 2. generating UI from model objects (originally I said UiBinder, but Lincoln has now implemented a really nice Errai UI module that we'd probably want to target instead/also) 3. making a bunch of types @Portable in a single command (for example, all types used by a @Remote interface) 4. ability to remove Errai from a project :,(and more.. once the basic Errai plugin is working, we can probably add all sorts of amazing stuff!	0	0	https://issues.jboss.org/browse/ERRAI-260
4011	ERRAI-237	What would suggest for Spring?	0	0	https://issues.jboss.org/browse/ERRAI-237
4012	ERRAI-218	Why is there an empty ErraiApp.properties? What are the roles.properties and users.properties used for?	0	0	https://issues.jboss.org/browse/ERRAI-218

#	Requirement	Questions	Answered	Successfully Answered	URL
4013	ERRAI-166	org.jboss.errai.bus.server.servlet.AbstractErraiServlet.init(AbstractErraiServlet.java:98) at org.eclipse.jetty.servlet.ServletHolder.initServlet(ServletHolder.java:456) at org.eclipse.jetty.servlet.ServletHolder.doStart(ServletHolder.java:276) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:59) at org.eclipse.jetty.servlet.ServletHandler.initialize(ServletHandler.java:779) at org.eclipse.jetty.servlet.ServletContextHandler.startContext(ServletContextHandler.java:255) at org.eclipse.jetty.webapp.WebAppContext.startContext(WebAppContext.java:1212) at org.eclipse.jetty.server.handler.ContextHandler.doStart(ContextHandler.java:610) at org.eclipse.jetty.webapp.WebAppContext.doStart(WebAppContext.java:453) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:59) at org.eclipse.jetty.deploy.bindings.StandardStarter.processBinding(StandardStarter.java:36) at org.eclipse.jetty.deploy.AppLifeCycle.runBindings(AppLifeCycle.java:183) at org.eclipse.jetty.deploy.DeploymentManager.requestAppGoal(DeploymentManager.java:485) at org.eclipse.jetty.deploy.DeploymentManager.addApp(DeploymentManager.java:138) at org.eclipse.jetty.deploy.providers.ScanningAppProvider.fileAdded(ScanningAppProvider.java:142) at org.eclipse.jetty.deploy.providers.ScanningAppProvider$1.fileAdded(ScanningAppProvider.java:53) at org.eclipse.jetty.util.Scanner.reportAddition(Scanner.java:604) at org.eclipse.jetty.util.Scanner.reportDifferences(Scanner.java:535) at org.eclipse.jetty.util.Scanner.scan(Scanner.java:398) at org.eclipse.jetty.util.Scanner.doStart(Scanner.java:332) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:59) at org.eclipse.jetty.deploy.providers.ScanningAppProvider.doStart(ScanningAppProvider.java:118) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:59) at org.eclipse.jetty.deploy.DeploymentManager.startAppProvider(DeploymentManager.java:546) at org.eclipse.jetty.deploy.DeploymentManager.doStart(DeploymentManager.java:221) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:59) at org.eclipse.jetty.util.component.AggregateLifeCycle.doStart(AggregateLifeCycle.java:45) at org.eclipse.jetty.server.handler.AbstractHandler.doStart(AbstractHandler.java:53) at org.eclipse.jetty.server.handler.HandlerWrapper.doStart(HandlerWrapper.java:90) at org.eclipse.jetty.server.Server.doStart(Server.java:262) at org.eclipse.jetty.util.component.AbstractLifeCycle.start(AbstractLifeCycle.java:59) at org.eclipse.jetty.xml.XmlConfiguration$1.run(XmlConfiguration.java:1203) at java.security.AccessController.doPrivileged(Native Method) at org.eclipse.jetty.xml.XmlConfiguration.main(XmlConfiguration.java:1126) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at	0	0	https://issues.jboss.org/browse/ERRAI-166
4014	ERRAI-114	hbraun: IMO an application scoped dependency isn't a "passivation capable dependency" [1:32pm] hbraun: Am i wrong? [1:32pm] stalep: i compared with a snapshot i had from 12 nov, but that contained my optimizations that couldnt be used. [1:32pm] pmuir: an app scoped dependency is passivation capable if it is serializable [1:33pm] pmuir: otherwise not [1:33pm] hbraun: the error I get occurs on a application scoped dependency that's not serializable [1:33pm] pmuir: no no [1:33pm] hbraun: but let me double check [1:33pm] pmuir: i'm wrong [1:33pm] pmuir: all normal beans are passivation capable dependencies [1:33pm] pmuir: ala 6.6.2 [1:34pm] hbraun: "normal"? [1:34pm] pmuir: I mean all normal scoped of course [1:34pm] hbraun: in your terms: is application scoped = normal scoped? [1:34pm] pmuir: so an app scoped bean is a passivation capable dependency [1:34pm] stalep: pmuir, compared to a run we did at 20 nov the number of el invocation was greatly reduced [1:34pm] pmuir: so therefore it should be injectable into a passivation capable bean such as session [1:35pm] hbraun: but an application scoped instance is shared, no? how can it be passivated? [1:35pm] pmuir: it can't [1:35pm] hbraun: can't be shared or passivated? [1:36pm] pmuir: it is a passivation capable *dependency* not a passivation capable *bean* [1:36pm] hbraun: now we gte closer [1:36pm] pmuir: for a bean to be pasivation capable, all it's dependencies must be passivation capable [1:36pm] pmuir: it's possible weld is making a mistake here? [1:36pm] hbraun: well, to get back to my example [1:37pm] hbraun: a conversation scoped bean depends on an application scoped, non serialzable bean [1:37pm] hbraun: and the container chokes on it [1:37pm] hbraun: -> "declares passivating scope but has non-serializable dependency" [1:37pm] pmuir: right that seems like a bug in weld to me [1:38pm] pmuir: stalep: how do i get that invocation count screen in jprofiler? [1:38pm] hbraun: ok, so am not completely wrong [1:38pm] hbraun: let me double check [1:38pm] hbraun: i'll raise an issue you can omment on [1:38pm] stalep: pmuir, CPU Views, then choose Hot spot [1:39pm] stalep: pmuir, add View Filter: org.jboss.weld [1:39pm] plenyi left the chat room. (Quit: Leaving) [1:39pm] pmuir: ok [1:39pm] pmuir: thanks hbraun	1	0	https://issues.jboss.org/browse/ERRAI-114
4015	ERRAI-68	Even just a part of Errai? Why is it not feasible?	1	0	https://issues.jboss.org/browse/ERRAI-68
4016	ERRAI-48	Are you including new tool(s) for debugging in this Jira, or creating subtask associated with this Jira?	0	0	https://issues.jboss.org/browse/ERRAI-48
4017	ERRAI-10	Is this still important for BRMS? -- Posted from Bugbox for Android	1	0	https://issues.jboss.org/browse/ERRAI-10
4018	ERPL-7	Isn't this a duplicate of ERPL-4?	0	0	https://issues.jboss.org/browse/ERPL-7
4019	ERPL-3	This would be a fantastic idea. This would enable us to have a single servger impl for a in-house SWT app and our external customer facing GWT based app. Would this still allow push communications from the server? If so.... Nirvana :) Keep up the good work.	0	0	https://issues.jboss.org/browse/ERPL-3
4020	ESC-12	The only problem here is that you'd need Maven generate-sources to be executed everytime you want to test the code. Also, current IDE integrations allows you to simply right-click and execute. Would that be still possible?	1	1	https://issues.jboss.org/browse/ESC-12
4021	ESSC-1	Would it be possible to remove the port from the URL and do either some kind of port forwarding or just run it on standard port (80)? The URL is kind of ugly with the port.	1	1	https://issues.jboss.org/browse/ESSC-1
4022	ESSC-1	Why is there an incrementing number on the right end of an URL "?(X+1)"? Would it be possible to add a link to the home screen to the right menu (all products) / make the "Products" clickable? From purely subjective PoV, the green check mark before the titles is kind of useless. Personally I would remove it.	1	1	https://issues.jboss.org/browse/ESSC-1
4023	EXOJCR-1890	QA retested daily tests on [sha1:6e0bf47 TESTUKR-600|http://tests.exoplatform.org/issues/TESTUKR-600/single-thread-performance-testing/jcr.core/index.html] with reverted configuration (was used <property name="db-structure-type" value="single"/> instead of <property name="db-structure-type" value="isolated/>) Decreased TPS on [sha1:6e0bf47|http://tests.exoplatform.org/JCR/1.15.0-Beta01/05_18_6e0bf47/single-thread-performance-testing/jcr.core/index.html] was disappeared.	1	1	https://issues.jboss.org/browse/EXOJCR-1890

#	Requirement	Questions	Answered	Successfully Answered	URL
4024	EXOJCR-1829	Old database tables are not removed after migration via backup/restore. This, obviously, occurs only when we have different db types (e.g. single db -> isolated db), because in this case we create new tables instead of using the same tables (single db -> single db). The main question is: do we need to keep those tables for any reason, or they are useless and we should remove them.	1	1	https://issues.jboss.org/browse/EXOJCR-1829
4025	EXOJCR-1819	Which are ?	1	1	https://issues.jboss.org/browse/EXOJCR-1819
4026	EXOJCR-1819	> Which are ? < If you check the content of the 2 listed file above, you will see that the full XML content is duplicated. For example for the first file https://github.com/exodev/jcr/blob/master/exo.jcr.component.core.impl.infinispan.v5/src/test/resources/conf/standalone/test-configuration-mjdbc-ispn.xml, check line #665, you have: {code} <remove-configuration>org.exoplatform.services.scheduler.JobSchedulerService</remove-configuration> <!--<import>jar:/conf/database-configuration.hsql.xml</import> --> </configuration> <?xml version="1.0" encoding="ISO-8859-1"?> <!-- Copyright (C) 2009 eXo Platform SAS. This is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version. This software is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details. You should have received a copy of the GNU Lesser General Public License along with this software; if not, write to the Free Software Foundation, Inc., 51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA, or see the FSF site: http://www.fsf.org. --> <configuration {code}	0	0	https://issues.jboss.org/browse/EXOJCR-1819
4027	EXOJCR-1812	Hi Nicolas, Could you explain more details on this? We need to explain more clearly with reporter. Thanks, Minh	1	1	https://issues.jboss.org/browse/EXOJCR-1812
4028	EXOJCR-1812	To allow a user to add a mixin to a given node the JCR checks if it has the permissions SET_PROPERTY and ADD_NODE on that particular node. In other words addMixin is a special operation that requires special rights. In this particular case john doesn't have the permission ADD_NODE so it fails which is a normal behavior as it is mandatory. BTW why do you consider it as a bug?	1	1	https://issues.jboss.org/browse/EXOJCR-1812
4029	EXOJCR-1812	As I understand, to add a mixin to a given node, we just need to change the property *jcr:mixinTypes*. So only permision *SET_PROPERTY* is enough. Why you require user to have permission *ADD_NODE*?	1	1	https://issues.jboss.org/browse/EXOJCR-1812
4030	EXOJCR-1812	@Nicolas Filotto, Anatoliy Bazko: In our scenario, as you read in [https://jira.exoplatform.org/browse/ECMS-3545], user does not care about the *addMixin* funtion. He just notices that he has *SET_PROPERTY* and *REMOVE* permissions but can not edit the document. The question here is: - Do we need the *ADD_NODE* permission for all cases of *addMixin*? I think we only need it when the *mixin* node type contains at least 1 mandatory auto-created nodes. In case it does not have these things, *ADD_NODE* is not necessary.	1	1	https://issues.jboss.org/browse/EXOJCR-1812
4031	EXOJCR-1812	Hi Nicolas, IMO, that is a real use case in WCM and in the real life. He wants to edit a document so he just has edit permission on it. Why do we need *Add* permission? is it safe to let people has more permission than what he has to do? It is not logical for me. And if you consider this one could be done in JCR 1.15 then could you reopen it and add into 1.15 scope? Thanks, Minh	1	1	https://issues.jboss.org/browse/EXOJCR-1812
4032	EXOJCR-1812	>> Why do we need Add permission? is it safe to let people has more permission than what he has to do? It is not logical for me. << Maybe it is not logical but it is how it works from the beginning so changing it now would be a behavior change which is forbidden in a maintenance branch as you know >> And if you consider this one could be done in JCR 1.15 then could you reopen it and add into 1.15 scope? << Done BTW: The given unit test cannot work even with this improvement as making a node mix:versionable will add nodes to create the history of this node, so I guess it should be replaced with something like mix:lockable which seems to be the current issue in ECMS-3545	0	0	https://issues.jboss.org/browse/EXOJCR-1812
4033	EXOJCR-1812	What is the exception ?	0	0	https://issues.jboss.org/browse/EXOJCR-1812
4034	EXOJCR-1812	>> So there is no way for him to modify the child node which is created by him, right? right. >> In real life this is a real use case. >> I have a folder, It is mean you are owner, correct? >>I allow another one to add data into this folder, so he should be able to modify the files he created. If you allowed it is mean you add him some permissions, in case of modification files it should be SET_PROPERTY. Otherwise how it possible? Did I understand your correct?	1	1	https://issues.jboss.org/browse/EXOJCR-1812
4035	EXOJCR-1812	@Anatoliy Bazko: {quote} >>I allow another one to add data into this folder, so he should be able to modify the files he created. If you allowed it is mean you add him some permissions, in case of modification files it should be SET_PROPERTY. Otherwise how it possible? {quote} In case of modification files it should be SET_PROPERTY, yes. But I do not want the user to modify my folder. User can modify THEIR files inside my folder. Note that the files are created by him. I think the file owner can modify and should have full permission on his created file.	0	0	https://issues.jboss.org/browse/EXOJCR-1812
4036	EXOJCR-1793	I don't understand the purpose of this issue, what you describe should already be possible (I mean Could Workspace team can do that already), if not what is the issue met?	0	0	https://issues.jboss.org/browse/EXOJCR-1793

#	Requirement	Questions	Answered	Successfully Answered	URL
4037	EXOJCR-1791	Maybe there is some misunderstanding here. PortalContainer.getInstance() is not called directly from my JMXOrganizationInitializer. It's called from another class later in IDM layer (namely from class org.exoplatform.services.organization.idm.IDMUserListAccess), which is not instantiated by kernel and don't have reference to container. But what I am saying is, that during invocation of JMX component, should kernel layer pay attention to set correct portal container into ThreadLocal of ExoContainerContext. I think that this should be done in classes for JMX integration (namely in ManagementContextImpl.beforeInvoke or in ExoModelMBean.invoke). It's sure think that I have some possibilities to workaround it. What I am currently doing as workaround is like: {code} public class JMXOrganizationInitializer implements Startable { private ExoContainer exoContainer; public JMXOrganizationInitializer(ExoContainerContext exoContainerContext, OrganizationListenersInitializerService organizationListenersInitializerService, OrganizationService orgService) { this.exoContainer = exoContainerContext.getContainer(); this.initializerService = organizationListenersInitializerService; this.orgService_=orgService; } {code} And then in method launchAll, which is invoked by JMX, I can wrap the invocation with correct portal container this way: {code} public String launchAllListeners(@ManagedDescription("Check JCR Folders") @ManagedName("checkFolders") Boolean checkFolders) { // TODO: We need to encapsulate call with correct portal container by ourselves because JMX kernel layer is not doing it ExoContainer oldContainer = ExoContainerContext.getCurrentContainer(); try { ExoContainerContext.setCurrentContainer(this.exoContainer); boolean ok = initializerService.launchAll(checkFolders); String responseString = ok ? "All listeners executed successfuly." : "Error occured during execution of listeners."; return responseString; } finally { ExoContainerContext.setCurrentContainer(oldContainer); } } {code} But do you really think that it should be done this way? I really don't think that it's responsibility of calling component (in our case JMXOrganizationInitializer) to setup correct portal container to ThreadLocal. I think that this really should be already done by kernel. So in my case, variable "oldContainer" should be already set to correct container, which is not the case now. For normal invocations by Http request, we have SetCurrentIdentityFilter, which encapsulates whole Http request with correct Portal Container. I think that similar thing should be also done in JMX layer.	1	1	https://issues.jboss.org/browse/EXOJCR-1791
4038	EXOJCR-1791	Hi Nicolas, For that specific case with IDMUserListAccess: Previously it was done in the way, which you just suggest. IDMUserListAccess had references to Organization and IDM services. Problem is that IDMUserListAccess is Serializable and there were problems like "NotSerializableException" so Bolek fixed it by removing references. You can see diff here https://source.jboss.org/browse/GateIn/portal/trunk/component/identity/src/main/java/org/exoplatform/services/organization/idm/IDMUserListAccess.java?r2=7565&r1=7152 and JIRA here https://issues.jboss.org/browse/GTNPORTAL-2138 . Of course, this specific case can be solved in other way, like using some custom serialization mechanism. But I am pretty sure that there may be cases, where it's not easily possible to inject instance of service or ExoContainer to the calling component, and so you need to use ExoContainerContext.getCurrentContainer or PortalContainer.getInstance . For instance all WebUI components are using method UIApplication.getApplicationComponent for obtaining instances of some service, which in next turn invoke ExoContainerContext.getCurrentContainer. Another example can be ExtensibleFilters, which requires that GenericFilter must be last filter in the mapping (I remember that you personally created GTNPORTAL JIRA that GenericFilter mapping was before SetCurrentIdentityFilter and Filters could have bad instance of portal container because of it) I am sure there are many components and services, which relies on the fact, that ExoContainerContext.getCurrentContainer returns correct instance of PortalContainer. And during normal HTTP invocation, it always returns correct PortalContainer thanks to SetCurrentIdentityFilter. But during JMX calls, you can't be sure that ExoContainerContext.getCurrentContainer returns correct instance. And only 2 lines of code can fix the whole thing. Are you seeing any disadvantages of adding these 2 lines into kernel code to have more reliable JMX layer? Me not, i am seeing only advantages. I agree that it's not a bug but enhancement. Sorry for creating it as bug.	1	1	https://issues.jboss.org/browse/EXOJCR-1791
4039	EXOJCR-1782	Hi, reviewing the patch, I've noticed that you added new blocks to documents, for example, {code:xml} <itemizedlist> <listitem> <para><link linkend="JCR.Extensions.Concept">Concept</link></para> </listitem> <listitem> <para><link linkend="JCR.Extensions.Implementation">Implementation</link></para> </listitem> <listitem> <para><link linkend="JCR.Extensions.Configuration">Configuration</link></para> </listitem> <listitem> <para><link linkend="JCR.Extensions.Related_Pages">Related Pages</link></para> </listitem> </itemizedlist> {code} And that seemed to me quite a bit strange, because the contents of a chapter (a list of its sections along with links to them) are generated automatically during creation of documentation book. So I see no need to duplicate it by adding an {{itemizedlist}} in the very beginning of a chapter. In this case you will have two contents lists in a row. Could you please explain me your intends, if I have understood something wrong.	1	1	https://issues.jboss.org/browse/EXOJCR-1782
4040	EXOJCR-1778	Please provide concrete use cases what kind of error did you face ? what was the corresponding stack trace/issue?	0	0	https://issues.jboss.org/browse/EXOJCR-1778
4041	EXOJCR-1751	(from 2.9 to 3.5) Features for future investigation: * FieldCache? How can it be used? Expert API? * SearcherLifeTimeManager - should be used or not? Useful or not. * FST? (Lucene-3140) Can it improve similarity search or search by word part?. Seems it's internal feature and doesn't actually change the Lucene's behavior, only better RAM usage. * (Experimental) TwoPhaseCommitTool in Lucene. Can be useful when implementing new JCR SQL 2.0 Search Engine. * Facet search API in 3.4.0 - * IndexUpgrader - Compressed fields are not supported now. Old indexes with compressed fields will be "de-compressed" on optimize or segments merge. But also utility can do this. Possibly it can do smth more. * PathHierarchyTokenized - tokenizes path entries. Can be usefull if storing paths in index. * ParalellMultiSearcher - ? study. Changes and notices * IndexReader.openIfChanged() - consider using. Changed behavior comparing to reopen. * IndexSearcher.createWeight() -> createNormalizedWeight. * IndexWriter.addIndex() - doesn't merge segment any more. Need explicit call to .maybeMerge or optimize. * Searcher and Searcheble -> IndexSearcher. * Query.getSimilarity() - deprecated (lucene-2869). * TopScoreDocCollector throws exception when nDoc=0 (lucene-2785). * FSDirectory.open() can return MMapDirectory instance. * IndexWriter constructor accepts IndexWriterConfig instance * FSDirectory and LockFactory are not settable via SysProps (?).	0	0	https://issues.jboss.org/browse/EXOJCR-1751
4042	EXOJCR-1734	As solution for the "hot modification support" part I propose to expose a set of getter and setter through JMX (they will also be part of the interface of the service to allow UI to call them if needed) in order to give the ability to hotly change the parameters but the values won't be stored in the db, the customer will need to modify the configuration file manually. I prefer to keep it simple as I'm not sure it is a common use case. @tug: Is this proposal acceptable for you?	0	0	https://issues.jboss.org/browse/EXOJCR-1734

#	Requirement	Questions	Answered	Successfully Answered	URL
4043	EXOJCR-1729	{code} com.sun.org.apache.xml.internal.serializer.utils.WrappedRuntimeException: Could not load the propery file 'output_xml.properties' for output method 'xml' (check CLASSPATH) at com.sun.org.apache.xml.internal.serializer.OutputPropertiesFactory.getDefaultMethodProperties(OutputPropertiesFactory.java:324) at com.sun.org.apache.xalan.internal.xsltc.trax.TransformerImpl.setDefaults(TransformerImpl.java:1123) at com.sun.org.apache.xalan.internal.xsltc.trax.TransformerImpl.createOutputProperties(TransformerImpl.java:1084) at com.sun.org.apache.xalan.internal.xsltc.trax.TransformerImpl.<init>(TransformerImpl.java:250) at com.sun.org.apache.xalan.internal.xsltc.trax.TransformerImpl.<init>(TransformerImpl.java:241) at com.sun.org.apache.xalan.internal.xsltc.trax.TransformerFactoryImpl.newTransformer(TransformerFactoryImpl.java:607) at com.sun.org.apache.xalan.internal.xsltc.trax.TransformerFactoryImpl.newTransformerHandler(TransformerFactoryImpl.java:898) at org.exoplatform.container.configuration.ConfigurationUnmarshaller$2.run(ConfigurationUnmarshaller.java:287) at org.exoplatform.container.configuration.ConfigurationUnmarshaller$2.run(ConfigurationUnmarshaller.java:274) at java.security.AccessController.doPrivileged(Native Method) at org.exoplatform.commons.utils.SecurityHelper.doPrivilegedExceptionAction(SecurityHelper.java:307) at org.exoplatform.container.configuration.ConfigurationUnmarshaller.unmarshall(ConfigurationUnmarshaller.java:273) at org.exoplatform.container.configuration.ConfigurationManagerImpl.addConfiguration(ConfigurationManagerImpl.java:167) at org.exoplatform.container.configuration.ConfigurationManagerImpl.addConfiguration(ConfigurationManagerImpl.java:134) at org.exoplatform.container.configuration.ConfigurationManagerImpl.addConfiguration(ConfigurationManagerImpl.java:128) at org.exoplatform.container.RootContainer.buildRootContainer(RootContainer.java:458) at org.exoplatform.container.RootContainer.getInstance(RootContainer.java:521) at org.exoplatform.container.TestExoContainer.testRemoveComponent(TestExoContainer.java:64) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) {code}	0	0	https://issues.jboss.org/browse/EXOJCR-1729
4044	EXOJCR-1726	cr.framework.web:1[]src:1[]main:1[]java:1[]org:1[]exoplatform:1[]frameworks:1[]jcr:1[]web:1[]fckeditor:1[].svn:1[]text-base:1 1e27b3937f00000170d1f3c3354ec8ea. Cause >>>> Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)): Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)) (FtpCommandImpl.java, line 187) javax.jcr.InvalidItemStateException: [backup] DELETE NODE. Can not delete parent till childs exists. Item []:1[]exo.jcr.framework.web:1[]src:1[]main:1[]java:1[]org:1[]exoplatform:1[]frameworks:1[]jcr:1[]web:1[]fckeditor:1[].svn:1[]text-base:1 1e27b3937f00000170d1f3c3354ec8ea. Cause >>>> Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)): Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)): [backup] DELETE NODE. Can not delete parent till childs exists. Item []:1[]exo.jcr.framework.web:1[]src:1[]main:1[]java:1[]org:1[]exoplatform:1[]frameworks:1[]jcr:1[]web:1[]fckeditor:1[].svn:1[]text-base:1 1e27b3937f00000170d1f3c3354ec8ea. Cause >>>> Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)): Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)): Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)) at org.exoplatform.services.jcr.impl.core.SessionDataManager.commit(SessionDataManager.java:1528) at org.exoplatform.services.jcr.impl.core.ItemImpl.save(ItemImpl.java:713) at org.exoplatform.services.jcr.impl.core.SessionImpl.save(SessionImpl.java:1080) at org.exoplatform.services.ftp.command.FtpCommandImpl.removeResource(FtpCommandImpl.java:174) at org.exoplatform.services.ftp.command.CmdRmd.run(CmdRmd.java:47) at org.exoplatform.services.ftp.command.FtpCommandImpl.execute(FtpCommandImpl.java:78) at org.exoplatform.services.ftp.client.FtpClientCommandThread.run(FtpClientCommandThread.java:103) Caused by: org.exoplatform.services.jcr.impl.storage.JCRInvalidItemStateException: [backup] DELETE NODE. Can not delete parent till childs exists. Item []:1[]exo.jcr.framework.web:1[]src:1[]main:1[]java:1[]org:1[]exoplatform:1[]frameworks:1[]jcr:1[]web:1[]fckeditor:1[].svn:1[]text-base:1 1e27b3937f00000170d1f3c3354ec8ea. Cause >>>> Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)): Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)): Cannot delete or update a parent row: a foreign key constraint fails (`portal`.`JCR_SITEM`, CONSTRAINT `JCR_FK_SITEM_PARENT` FOREIGN KEY (`PARENT_ID`) REFERENCES `JCR_SITEM` (`ID`)) at org.exoplatform.services.jcr.impl.storage.jdbc.SQLExceptionHandler.handleDeleteException(SQLExceptionHandler.java:345) at org.exoplatform.services.jcr.impl.storage.jdbc.JDBCStorageConnection.delete(JDBCStorageConnection.java:748) at org.exoplatform.services.jcr.impl.dataflow.persistent.WorkspacePersistentDataManager.doDelete(WorkspacePersistentDataManager.java:825) at	0	0	https://issues.jboss.org/browse/EXOJCR-1726

#	Requirement	Questions	Answered	Successfully Answered	URL
4045	EXOJCR-1709	when volatile flushed, both non-system and system; * sync indexes on cluster node startup, both new and previously re-started node; * specifying RSync configuration for QueryHandler; * support of RSync Authentication; h1. How implemented Standard Shared Index strategy has much in common with new RSync-Based concept. Only difference is an RSync call when indexes changed on coordinator for synchronization with other cluster nodes. That's why RSync concept is implemented on the top existing JBossCache-based components. JBossCacheIndexChangesFilter updated to handle new type of IndexInfos, RSyncIndexInfos. Those are the classes that manage the list of indexes, calling MultiIndex when list changed. RSyncIndexInfos triggers call to system utility in order to retrieve the last snapshot of index folder. When process finished, then control flow returned to JBossCache-based index components. Also RSync utility invoked on startup to retrieve latest index state. h1. What tested Concept launched on 2 and 3-nodes cluster. Following features tested manually: * basic index synchronization, when coordinator flushes new index; * index retrieval on startup, when no local index present; * index re-synchronization on startup, when local index already present; * coordinator changes; * temporary outage of some cluster node. On each step new document was uploaded via Dav and tested launching queries on each node. Next task is a performance testing, comparing different approaches. h1. How to configure You need ordinary JBossCache index configuration as shown below: {code} <query-handler class="org.exoplatform.services.jcr.impl.core.query.lucene.SearchIndex"> <properties> <property name="index-dir" value="${exo.jcr.parent.dir:..}/temp/jcrlucenedb/production" /> <property name="changesfilter-class" value="org.exoplatform.services.jcr.impl.core.query.jbosscache.JBossCacheIndexChangesFilter" /> <property name="jbosscache-configuration" value="jar:/conf/portal/cluster/jbosscache-indexer.xml" /> <property name="jgroups-configuration" value="jar:/conf/portal/cluster/udp-mux.xml" /> <property name="jgroups-multiplexer-stack" value="false" /> <property name="jbosscache-cluster-name" value="JCR-cluster-indexer" /> <property name="jbosscache-shareable" value="true" /> <property name="max-volatile-time" value="60" /> </properties> </query-handler> {code} Make sure "index-dir" is pointing to local folder. Then need to add RSync configuration options: {code} <query-handler class="org.exoplatform.services.jcr.impl.core.query.lucene.SearchIndex"> <properties> <property name="index-dir" value="${exo.jcr.parent.dir:..}/temp/jcrlucenedb/production" /> <property name="changesfilter-class" value="org.exoplatform.services.jcr.impl.core.query.jbosscache.JBossCacheIndexChangesFilter" /> <property name="jbosscache-configuration" value="jar:/conf/portal/cluster/jbosscache-indexer.xml" /> <property name="jgroups-configuration" value="jar:/conf/portal/cluster/udp-mux.xml" /> <property name="jgroups-multiplexer-stack" value="false" /> <property name="jbosscache-cluster-name" value="JCR-cluster-indexer" /> <property name="jbosscache-shareable" value="true" /> <property name="max-volatile-time" value="60" /> <property name="rsync-entry-name" value="index" /> <property name="rsync-entry-path" value="/home/exo/tomcat" /> <property name="rsync-port" value="8085" /> <property name="rsync-user" value="rsyncexo" /> <property name="rsync-password" value="exo" /> </properties> </query-handler> {code} Let's start with authentication: "rsync-user" and "rsync-password". It is optional and can be missing if RSync Server allows anonymous access. Before reviewing other RSync index options need to have a look at RSync server configuration. Sample RSync Server (rsyncd) Configuration {code} uid = nobody gid = nobody use chroot = no port = 8085 log file = rsyncd.log pid file = rsyncd.pid [index] path = /home/exo/tomcat comment = indexes read only = true auth users = rsyncexo secrets file= rsyncd.secrets {code} This sample configuration shows	0	0	https://issues.jboss.org/browse/EXOJCR-1709
4046	EXOJCR-1691	Can you provide the related statistics?	1	1	https://issues.jboss.org/browse/EXOJCR-1691
4047	EXOJCR-1673	Hum, according to PLF-2493 you already choose solution #2 so why choosing solution #1 for exo:registry only? Can't you apply the same logic on this node at ECMS or PLF level?	0	0	https://issues.jboss.org/browse/EXOJCR-1673
4048	EXOJCR-1673	I checked this patch with PLF 3.1.0-GA rev 77495. It's ok for node exo:registry, exo:services. But exo:applications is not hidden (cf capture) What do you think ?	0	0	https://issues.jboss.org/browse/EXOJCR-1673
4049	EXOJCR-1661	What is the expected version? JCR 1.14 only or JCR 1.12 too?	0	0	https://issues.jboss.org/browse/EXOJCR-1661
4050	EXOJCR-1660	Warning what you propose is not backward compatible, we cannot remove a parameter just because it is not compatible with multitenancy, multitenancy is only one use case. What has already been implemented at NodeHierarchyCreator is this: {code} String repositoryName = hierarchyConfig.getRepository(); ManageableRepository repository = repositoryName == null || repositoryName.isEmpty() ? jcrService_.getCurrentRepository() : jcrService_.getRepository(repositoryName); {code} In other words if we define a repository name it uses the given repository otherwise it uses the current repository, it is compatible with multitenancy and it is backward compatible.	0	0	https://issues.jboss.org/browse/EXOJCR-1660
4051	EXOJCR-1652	Please review each query and rewrite those which are written to use JCR_IDX_SITEM_PARENT_NAME by using filters in the same order as the list of column defined in the index JCR_IDX_SITEM_PARENT_NAME. For example the query: {code} select I.NAME, V.DATA, V.ORDER_NUM from JCR_SITEM I, JCR_SVALUE V where I.I_CLASS=2 and I.CONTAINER_NAME=? and I.PARENT_ID=? and (I.NAME='[http://www.jcp.org/jcr/1.0]primaryType' or I.NAME='[http://www.jcp.org/jcr/1.0]mixinTypes' or I.NAME='[http://www.exoplatform.com/jcr/exo/1.0]owner' or I.NAME='[http://www.exoplatform.com/jcr/exo/1.0]permissions') and I.ID=V.PROPERTY_ID {code} first check I_CLASS, then CONTAINER_NAME, then the PARENT_ID, then the NAME exactly like the order of the column defined in the index definition which is "I_CLASS, CONTAINER_NAME, PARENT_ID, NAME, I_INDEX, VERSION DESC". The query should be rewritten to match with the order defined in the index JCR_IDX_SITEM_PARENT which is "CONTAINER_NAME, PARENT_ID, NAME, I_INDEX, I_CLASS, VERSION DESC"	0	0	https://issues.jboss.org/browse/EXOJCR-1652
4052	EXOJCR-1652	I removed the index JCR_IDX_SITEM_PARENT_NAME and launched {code} explain select I.*, P.NAME AS PROP_NAME, V.ORDER_NUM, V.DATA from JCR_SITEM I, JCR_SITEM P, JCR_SVALUE V where I.I_CLASS=1 and I.CONTAINER_NAME='collaboration' and I.PARENT_ID='collaborationac1b5f7fc0a8010218cb98a47d54b11b' and P.I_CLASS=2 and P.CONTAINER_NAME='' and P.PARENT_ID=I.ID and (P.NAME='[http://www.jcp.org/jcr/1.0]primaryType' or P.NAME='[http://www.jcp.org/jcr/1.0]mixinTypes' or P.NAME='[http://www.exoplatform.com/jcr/exo/1.0]owner' or P.NAME='[http://www.exoplatform.com/jcr/exo/1.0]permissions') and V.PROPERTY_ID=P.ID order by I.N_ORDER_NUM, I.ID {code} The result is not good, I get "Using where; Using temporary; Using filesort" for P so it seems that both indexes are needed: JCR_IDX_SITEM_PARENT is at least required for: {code} FIND_ITEM_BY_NAME = "select * from JCR_SITEM" + " where CONTAINER_NAME=? and PARENT_ID=? and NAME=? and I_INDEX=? order by I_CLASS, VERSION DESC"; {code} and JCR_IDX_SITEM_PARENT_NAME is at least required required for: {code} protected static final String FIND_NODES_BY_PARENTID_CQ_QUERY = "select I.*, P.NAME AS PROP_NAME, V.ORDER_NUM, V.DATA from JCR_SITEM I, JCR_SITEM P, JCR_SVALUE V" + " where I.I_CLASS=1 and I.CONTAINER_NAME=? and I.PARENT_ID=? and" + " P.I_CLASS=2 and P.CONTAINER_NAME=? and P.PARENT_ID=I.ID and" + " (P.NAME='[http://www.jcp.org/jcr/1.0]primaryType' or" + " P.NAME='[http://www.jcp.org/jcr/1.0]mixinTypes' or" + " P.NAME='[http://www.exoplatform.com/jcr/exo/1.0]owner' or" + " P.NAME='[http://www.exoplatform.com/jcr/exo/1.0]permissions')" + " and V.PROPERTY_ID=P.ID order by I.N_ORDER_NUM, I.ID"; {code} So we need to revert changes related to this issue and resolve it as won't fix	0	0	https://issues.jboss.org/browse/EXOJCR-1652

#	Requirement	Questions	Answered	Successfully Answered	URL
4053	EXOJCR-1646	What requested by PLF-1898 is already covered by EXOJCR-1374 even for the JCR locks, to benefit of the changes done for it you need to replace org.jboss.cache.loader.JDBCCacheLoader with org.exoplatform.services.jcr.impl.core.lock.jbosscache.JDBCCacheLoader in your JBC config file for the locks, here is an example https://svn.jboss.org/repos/exo-jcr/jcr/trunk/exo.jcr.component.core/src/test/resources/conf/standalone/test-jbosscache-lock.xml	0	0	https://issues.jboss.org/browse/EXOJCR-1646
4054	EXOJCR-1621	@tung do thanh What did you run?	1	1	https://issues.jboss.org/browse/EXOJCR-1621
4055	EXOJCR-1618	On which AS do you expect this use-case working?	1	1	https://issues.jboss.org/browse/EXOJCR-1618
4056	EXOJCR-1618	On which AS do you expect this use-case working? - Tomcat 6.0.32 or higher And on which context exactly this use-case is required? -cloud-ide.com application	1	1	https://issues.jboss.org/browse/EXOJCR-1618
4057	EXOJCR-1618	I know that it is for cloud-ide.com application what I meant is more why would the end-users need to modify the war files? or more generally in which context would you need to modify the war files?	1	1	https://issues.jboss.org/browse/EXOJCR-1618
4058	EXOJCR-1618	And finally does this use-case represent a common need or a very rare need?	1	1	https://issues.jboss.org/browse/EXOJCR-1618
4059	EXOJCR-1592	Do we need to restore locks?	0	0	https://issues.jboss.org/browse/EXOJCR-1592
4060	EXOJCR-1590	So for the second issue, you will be able to remove old locks created by JCR 1.12 once upgraded to JCR 1.14 (EPP 5.2) ? In this case, if we use shareable and same table for locks, customers will still have old N tables in their database. But if I understand it correctly, these tables will be empty then (because locks will be removed) and so it should not be a problem for migration ? Only thing is that customers will have empty tables for old locks but I guess this should not be a problem (only few useless additional tables). Is it correct? Especially I wanted to know if customers will be able to upgrade to shareable without doing any manual steps with their DB (aka deleting old lock tables). And can we assume that this can be fixed in JCR 1.14.3 (Sorry for bothering, but 1.14.3 is the version, which will be used in EPP 5.2.0.GA) ? Thanks a lot!	0	0	https://issues.jboss.org/browse/EXOJCR-1590
4061	EXOJCR-1561	What about JCR 1.12.10 and 1.14.0?	0	0	https://issues.jboss.org/browse/EXOJCR-1561
4062	EXOJCR-1559	What about JCR 1.12.10 and 1.14.0?	1	1	https://issues.jboss.org/browse/EXOJCR-1559
4063	EXOJCR-1522	Did you disable the JCR Cache before launching the tests? how many iterations did you configure for these tests?	1	1	https://issues.jboss.org/browse/EXOJCR-1522
4064	EXOJCR-1483	I'm wondering if the patch could cause inconsistency issue, indeed even with read operations as for some operations we launch several JDBC queries (for example get Item by id) it could be required to have autocommit to false to ensure consistency between all the queries. Could you please review all the read operations and set autocommit to false for all the read operations that requires several queries. If this new change doesn't allow to have the same perf as now (with this feature), simply revert everything and resolve it as won't fix.	0	0	https://issues.jboss.org/browse/EXOJCR-1483
4065	EXOJCR-1479	I see dependencies doesn't change. So what I have to do to safe use of slf4j 1.6.1 with ExoLoger ?	1	1	https://issues.jboss.org/browse/EXOJCR-1479
4066	EXOJCR-1471	There is strange usecase. Some non binary properties, for example primatyType, have binary values stored on file system. How to find such properties? {code:sql} select V.*, P.P_TYPE from JCR_SVALUE as V, JCR_SITEM as P where P.I_CLASS=2 and name='[http://www.jcp.org/jcr/1.0]primaryType' and V.PROPERTY_ID = P.ID and V.DATA is null {code} In our case it will return row {code} 241290, , collaboration608768af0afe1ae0198496cabb87c860, collaboration, 2 {code} We can find this value stored on file system, using shell script above. So the value is nt:file. Then fix the record in JCR_SVALUE table {code:sql} update JCR_SVALUE set DATA='[http://www.jcp.org/jcr/nt/1.0]file', STORAGE_DESC=null where (PROPERTY_ID='collaboration608768af0afe1ae0198496cabb87c860') and ORDER_NUM=0; {code}	0	0	https://issues.jboss.org/browse/EXOJCR-1471
4067	EXOJCR-1471	I have read link above, and description of CheckIndex tool {code} **WARNING**: -fix should only be used on an emergency basis as it will cause documents (perhaps many) to be permanently removed from the index. Always make a backup copy of your index before running this! Do not run this tool on an index that is actively being written to. You have been warned! Run without -fix, this tool will open the index, report version information and report any exceptions it hits and what action it would take if -fix were specified. With -fix, this tool will remove any segments that have issues and write a new segments_N file. This means all documents contained in the affected segments will be removed. This tool exits with exit code 1 if the index cannot be opened or has any corruption, else 0. {code} So, there is few points: - index must NOT be used during the check. So it definitely must be executed in standalone mode; - CheckIndex "fixes" index in very special way. If index-segment can not be read (some bytes are lost, etc) it removes this broken segment with all documents within. (I still investigating, how can we fetch broken documents list); - if we can not fetch docuemnts list from corrupted segment - there is no reason to make new tool with CheckIndex inside. Script will be enough. BTW I checked the index from crashed disk - its ok.	0	0	https://issues.jboss.org/browse/EXOJCR-1471
4068	EXOJCR-1406	According to the svn history, this lock exists since the class has been added to svn (the related JIRA is https://jira.exoplatform.org/browse/JCR-456) so we cannot know why it has been added. @peter As you seem to be the code owner, can you tell us more? If you can't since the QA did not detect any issue when the lock has been removed, I suggest to remove it completely from the code as long as we don't have good proof that it is required. If later we face issue because of this we will then implement a smarter mechanism based on a ConcurrentMap<String, Lock> where the String would be the path.	1	1	https://issues.jboss.org/browse/EXOJCR-1406

#	Requirement	Questions	Answered	Successfully Answered	URL
4069	EXOJCR-1361	Thanks to Anatoly, I understand why performance depends from getChildNodesData in ChildAxisQuery Lets check our usecase We want find node by path "/teststorage/root2/node2/node9/node5/1-9-5-2.txt/jcr:content" with query "select * from nt:base where dc:title = 'Exoplatform' and jcr:path like '/teststorage/root2/%'" 1) How it works with getChildNodesData * ChildAxisQuery$SimpleChildrenCalculator.getHits() takes root-node child nodes with getChildNodesData(). Result is "teststorage", and this node is cached. IMPORTANT: getChildNodesData assigned parent-nodes ACL to child nodes. Thats why "teststorage" cached WITH ACL. * ChildAxisQuery$SimpleChildrenCalculator.getHits() takes child nodes for "teststorage". Same algoritm and comments. "root2" with ACL is cached. * query does some job. Thats not impotant. * on - QueryResultImpl.getResuls(QueryResultImpl.java:329) - QueryResultImpl.isAccessGranted(QueryResultImpl.java:413) takes all result nodes - VersionableWorkspaceDataManager.getItemData is called * VersionableWorkspaceDataManager.getItemData takes a node and ACLInheritanceSupportedWorkspaceDataManager.getNearestACAncestorAcl(node) is called. * ACLInheritanceSupportedWorkspaceDataManager.getNearestACAncestorAcl(node) traverse all nodes from the end, until any node with acl is fetched. So, for node /teststorage/root2/node2/node9/node5/1-9-5-2.txt/jcr:content there is next calls: /teststorage/root2/node2/node9/node5/1-9-5-2.txt/jcr:content /teststorage/root2/node2/node9/node5/1-9-5-2.txt /teststorage/root2/node2/node9/node5 /teststorage/root2/node2/node9 /teststorage/root2/node2 /teststorage/root2 - and this is finish, because root2 is cached and contains ACL. 2) How it works without getChildNodesData Same steps. With single exception - "teststorage" and "root2" cached WITHOUT ACL. For node /teststorage/root2/node2/node9/node5/1-9-5-2.txt/jcr:content there is next calls: /teststorage/root2/node2/node9/node5/1-9-5-2.txt/jcr:content /teststorage/root2/node2/node9/node5/1-9-5-2.txt /teststorage/root2/node2/node9/node5 /teststorage/root2/node2/node9 /teststorage/root2/node2 /teststorage/root2 /teststorage / - yep, root node always have ACL There is two extended getItemData calls. And moreover, when "teststorage" and "root2" became fetched and cached, thay will be cached WITHOUT ACL. Thats why, redundant calls will be repeated all times. Why preformance in our usecase was so dramaticaly decreased? Because, search result contain 10 nodes on same level (/teststorage/root2/node2/node9/node5/1-9-5-2.txt/jcr:content - 7 nodes). For each node - 2 redundant getItemData calls on QueryResultImpl.getResuls() and fetching result with NodeIterator. So there is 40 redundatn calls per single usecase run.	0	0	https://issues.jboss.org/browse/EXOJCR-1361
4070	EXOJCR-1333	By the way, how do you do insert the text results in the comment?	1	1	https://issues.jboss.org/browse/EXOJCR-1333
4071	EXOJCR-1333	>> By the way, how do you do insert the text results in the comment? << Add it in a code block \{code\}my code\{code\} >> BTW why is it so important for you to optimize this method? Did you see on your profiler results that it was a bottleneck? It looks like a big patch for very small impact, don't you agree?	1	1	https://issues.jboss.org/browse/EXOJCR-1333
4072	EXOJCR-1333	Oh, yes, I totally agree that there's no need to optimize this method. I thought my first approach would work better than it did (thanks for the benchmark, btw). And, like I said, I was bored so I spent half an hour writing a slightly more optimized version. :) I do still wonder about memory usage of the current approach, though... What can I say, I like to optimize things when I can but I agree that in this case, it makes the code a lot more complex for no obvious gain.	1	1	https://issues.jboss.org/browse/EXOJCR-1333
4073	EXOJCR-1333	We already waste too much time on this so I will add it. WDYT about this compressed version of your optimized version: {code} public static boolean contains1(String profiles, String profileName) { if (profileName != null && profiles != null) { int i = profiles.indexOf(profileName); if (i != -1) { // look for next comma int nextComma = profiles.indexOf(',', i); int previousComma = profiles.lastIndexOf(',', i); return profiles .substring(previousComma == -1 ? 0 : previousComma + 1, nextComma == -1 ? profiles.length() : nextComma) .trim().equals(profileName); } } return false; } {code} And what about simply caching the Set of profiles? I don't think that the list of profiles will change much from one call to another?	0	0	https://issues.jboss.org/browse/EXOJCR-1333
4074	EXOJCR-1332	What files we need to store on svn?	0	0	https://issues.jboss.org/browse/EXOJCR-1332
4075	EXOJCR-1303	Hi! Can You, please, provide your JBoss AppServer bundle and quick steps to check the patch. I.e. some steps to reproduce and check that patched code works fine.	0	0	https://issues.jboss.org/browse/EXOJCR-1303
4076	EXOJCR-1295	I am not the one who created the question. I am just pointing the fact that the answer does not seem to answer. My bet on the meaning of the question is : Can 2 JCR instances be synchronised across a WAN ? That said. If JCR team doe not understand the question, I guess the question should be removed from their FAQ....	1	0	https://issues.jboss.org/browse/EXOJCR-1295
4077	EXOJCR-1293	What is also strange that Kernel's packaging contains reference on _commons-lang_ for {{commons}} module, but it doesn't present in Kernel Common POM: {noformat} module.commons = new Project("org.exoplatform.kernel", "exo.kernel.commons", "jar", module.version). addDependency(new Project("commons-lang", "commons-lang", "jar", "2.6")). ..... {noformat}	0	0	https://issues.jboss.org/browse/EXOJCR-1293
4078	EXOJCR-1293	What about _exo.kernel.commons_ and _commons-lang_ there? Do we need it in Kernel?	0	0	https://issues.jboss.org/browse/EXOJCR-1293

#	Requirement	Questions	Answered	Successfully Answered	URL
4079	EXOJCR-1275	(IndexMerger.java, line 399) 06.04.2011 17:36:54 *INFO * [Incoming-2,null] SearchIndex: Index closed: f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrluce nedb\production_system (SearchIndex.java, line 992) 06.04.2011 17:36:54 *INFO * [IndexMerger] IndexMerger: IndexMerger terminated (IndexMerger.java, line 399) 06.04.2011 17:36:54 *INFO * [Incoming-2,null] SearchIndex: Index closed: f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrluce nedb\production (SearchIndex.java, line 992) 06.04.2011 17:36:54 *INFO * [IndexMerger] IndexMerger: IndexMerger terminated (IndexMerger.java, line 399) 06.04.2011 17:36:54 *INFO * [Incoming-2,null] SearchIndex: Index closed: f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrluce nedb\backup (SearchIndex.java, line 992) 06.04.2011 17:36:54 *INFO * [IndexMerger] IndexMerger: IndexMerger terminated (IndexMerger.java, line 399) 06.04.2011 17:36:54 *INFO * [Incoming-2,null] SearchIndex: Index closed: f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrluce nedb\digital-assets (SearchIndex.java, line 992) 06.04.2011 17:38:48 *ERROR* [Thread-36] JobRepositoryRestore: Can't rollback changes (JobExistingRepositorySameConfigRestore. java, line 214) org.exoplatform.services.jcr.impl.backup.BackupException: java.io.IOException: Can't remove file : F:\data\claster_pgsql_8.2. 15\exo-tomcat\temp\jcrlucenedb\production_system_1_0.cfs at org.exoplatform.services.jcr.impl.backup.rdbms.DirectoryRestor.rollback(DirectoryRestor.java:169) at org.exoplatform.services.jcr.ext.backup.impl.JobExistingRepositorySameConfigRestore.restoreRepository(JobExistingR epositorySameConfigRestore.java:210) at org.exoplatform.services.jcr.ext.backup.impl.JobRepositoryRestore.run(JobRepositoryRestore.java:382) Caused by: java.io.IOException: Can't remove file : F:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrlucenedb\production_syste m_1_0.cfs at org.exoplatform.services.jcr.impl.util.io.DirectoryHelper.removeDirectory(DirectoryHelper.java:163) at org.exoplatform.services.jcr.impl.util.io.DirectoryHelper.removeDirectory(DirectoryHelper.java:151) at org.exoplatform.services.jcr.impl.util.io.DirectoryHelper.removeDirectory(DirectoryHelper.java:151) at org.exoplatform.services.jcr.impl.backup.rdbms.DirectoryRestor.rollback(DirectoryRestor.java:162) ... 2 more 06.04.2011 17:38:48 *INFO * [Incoming-2,null] SearchIndex: path=f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrlucenedb\digi tal-assets (SearchIndex.java, line 595) 06.04.2011 17:38:48 *INFO * [Incoming-2,null] SearchIndex: Index initialized: f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jc rlucenedb\digital-assets Version: 3 (SearchIndex.java, line 681) 06.04.2011 17:38:48 *INFO * [Incoming-2,null] SearchIndex: path=f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrlucenedb\back up (SearchIndex.java, line 595) 06.04.2011 17:38:49 *INFO * [Incoming-2,null] SearchIndex: Index initialized: f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jc rlucenedb\backup Version: 3 (SearchIndex.java, line 681) 06.04.2011 17:38:49 *INFO * [Incoming-2,null] SearchIndex: path=f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrlucenedb\prod uction (SearchIndex.java, line 595) 06.04.2011 17:38:49 *INFO * [Incoming-2,null] SearchIndex: Index initialized: f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jc rlucenedb\production Version: 3 (SearchIndex.java, line 681) 06.04.2011 17:38:49 *INFO * [Incoming-2,null] SearchIndex: path=f:\data\claster_pgsql_8.2.15\exo-tomcat\temp\jcrlucenedb\prod uction_system (SearchIndex.java, line 595) 06.04.2011 17:38:50 *ERROR* [Thread-36] JobRepositoryRestore: The restore was fail (JobRepositoryRestore.java, line 392) org.exoplatform.services.jcr.ext.backup.RepositoryRestoreExeption: Repository repository was not restored at org.exoplatform.services.jcr.ext.backup.impl.JobExistingRepositorySameConfigRestore.restoreRepository(JobExistingR epositorySameConfigRestore.java:218) at	1	1	https://issues.jboss.org/browse/EXOJCR-1275
4080	EXOJCR-1275	Changes to MultiIndex committed. Please, can You verify it?	0	0	https://issues.jboss.org/browse/EXOJCR-1275
4081	EXOJCR-1259	Thx Nadia, as far as I can see in your graph it is even worse in term of memory than the current code thus so far I don't see any reason to use the trove collections. @Dmitry Are you sure that the trove collections are used a lot during the daily tests ? * If not could you please check and measure how much it is used using a counter for example ? * If so do you have any idea why there is no benefit in term of perfs and memory?	0	0	https://issues.jboss.org/browse/EXOJCR-1259
4082	EXOJCR-1259	At least we can see some regression, is it not enough ?	0	0	https://issues.jboss.org/browse/EXOJCR-1259
4083	EXOJCR-1239	*Research and benchmark results* *JCR-1579* Benchmark test SearchNodesByPropertyUsingLikePathOperationTest runned query {code} select * from nt:base where dc:title = 'Exoplatform' and jcr:path like '/" + rootPath + "/%' {code} 1) clear JCR rev.4794 Index size 170 МБ (178 297 877 байт) SearchNodesByPropertyUsingLikePathOperationTest - 5457, 5580, 5755 tps 2) JCR-1579 patch Unchanged patch Index size 185 МБ (194 623 095 байт) SearchNodesByPropertyUsingLikePathOperationTest - 3133, 3053, 2961 tps 3) JCR-1579 + Path NO store Main idea is to reduce lucene-index size. So make PATH-field unstored. {code} doc.add(new Field(FieldNames.PATH, path.toString(), *Field.Store.NO* , Field.Index.NOT_ANALYZED_NO_NORMS)); {code} Index size 176 МБ (185 062 237 байт) SearchNodesByPropertyUsingLikePathOperationTest - 3834, 3839, 3842 tps 4) JCR-1579 - get all childs then take Path field Main idea is to avoid search by stored filed (Field.Store.YES). So we take all child documents by parent id then read PATH and compare {code} TermDocs docs = reader.termDocs(new Term(FieldNames.PARENT, uuid)); try { while (docs.next()) { int docId = docs.doc(); String path = reader.document(docId, FieldSelectors.PATH).get(FieldNames.PATH); if (path.equals(nameTest.getAsString())) { childrenHits.set(docId); } } } {code} also path filed size was reduced {code} doc.add(new Field(FieldNames.PATH, *node.getQPath().getName().getAsString()* , Field.Store.YES, Field.Index.NOT_ANALYZED_NO_NORMS)); {code} Index size 175 МБ (183 971 932 байт) SearchNodesByPropertyUsingLikePathOperationTest 4167 4221 4307 tps 4) New patch use NAMESPACE_URI and LOCAL_NAME instead PATH The main idea is to get rid PATH field and use NAMESPACE_URI and LOCAL_NAME fields like in example above. NAMESPACE_URI and LOCAL_NAME become stored. Index size 174 МБ (182 819 752 байт) SearchNodesByPropertyUsingLikePathOperationTest, 4592 4534 4565 tps || test query || *no patch* || *jcr-1579* || jcr-1579-No-store || jcr-1579-getPath || myPatch || *jcr-1571* || | select * from nt:base where dc:title like 'Exoplatf%' | 7550.442 | 7761.469 | 7062.517 | 7697.403 | 7680.896 | 7980.991 | | select * from nt:base where dc:title = 'Exoplatform' | 14199.297 | 18570.086 | 14375.355 | 14026.843 | 13619.727 | 14092.978 | | select * from nt:base where dc:title = 'Exoplatform' and jcr:path like '/" + rootPath + "/%' | 6982.468 | 3218.285 | 4410.023 | 4792.503 | 4960.121 | 6459.127 | | select * from nt:base where contains(*,'fournodes') and jcr:path like '/" + rootPath + "/%' | 7378.306 | 2709.797 | 3856.995 | 4197.39 | 4079.637 | 7514.903 | *Test with bloom filter* 1) original + bloom filter Index size 170 МБ (178 297 877 байт) SearchNodesByPropertyUsingLikePathOperationTest, 5697 5509 5569 2) JCR-1579 + bloom filter Index size 185 МБ (194 623 095 байт) SearchNodesByPropertyUsingLikePathOperationTest, 3356 3376 3295 tps || || original || JCR-1579 || | with BF | 5569 | 3295 | | no BF | 5509 | 3330 | *Summary* * bloom filter does not show any performace changes on this usecases. * JCR-1579 Can not find any algorithm that will use path form lucene-index and will not damage performance. * jcr-1571 Does not show performance regression on this usecases.	0	0	https://issues.jboss.org/browse/EXOJCR-1239
4084	EXOJCR-1216	What are the differences with the existing class?	0	0	https://issues.jboss.org/browse/EXOJCR-1216
4085	EXOJCR-1186	what about this result http://tests.exoplatform.org/exp/12_14/daily-performance-testing-results/jcr.core/ ? Can you clarify? Can the results that we get with JCR 1.12.x can be trusted? I thought that we mad changes on the server that caused those regressions, could you please relaunch the tests on JCR 1.12.x in order to be sure that we are on the exact same conditions?	0	0	https://issues.jboss.org/browse/EXOJCR-1186

#	Requirement	Questions	Answered	Successfully Answered	URL
4086	EXOJCR-1159	About upgrade ISPN to 4.2.1.FINAL (https://wiki-int.exoplatform.org/display/exoproducts/Review+all+the+third+party+libraries) The latest FINAL release is 4.2.0.FINAL, but last release 4.2.1.CR2 (https://repository.jboss.org/nexus/content/groups/public/org/infinispan/infinispan-core/4.2.1.CR2/) What version of ISPN will be used 4.2.1.CR2 or 4.2.0.FINAL?	0	0	https://issues.jboss.org/browse/EXOJCR-1159
4087	EXOJCR-1159	The dependency from article https://wiki-int.exoplatform.org/display/exoproducts/Review+all+the+third+party+libraries was updated, but with same new dependency we have follow problem: 1) com.ibm.db2:db2jcc - the version 9.7 was not present on maven repository (now this dependency was stored in local maven repository). 2) com.ibm.db2:db2jcc_license_cu - the version 9.7 was not present on maven repository (now this dependency was stored in local maven repository). 3) com.microsoft:sqljdbc - the version 3.0.1301.101 was not present on maven repository (now this dependency was stored in local maven repository). 4) commons-fileupload:commons-fileupload - now version in WS is 1.0, because one test failed with version 1.2.1 in project ws/trunk/exo.ws.rest.core. 5) javassist:javassist - the version 3.14.0.GA was not present on maven repository (now this dependency was stored in local maven repository). 6) jotm:jotm - the version 2.1.9 was not present on maven repository (now this dependency was stored in local maven repository). 7) jtidy:jtidy - now version is 4aug2000r7-dev , because same tests failed with version r938(and with version r918) in project core/trunk/exo.core.component.xml-processing. The version r938(and with version r918) was not present on maven repository (now this dependency was stored in local maven repository). 8) org.apache.poi:poi poi-oomxl poi-scratchpad - now version is 3.6, because same tests failed with version 3.7 in project core/trunk/exo.core.component.document. 9) org.apache.tika:tika-core tika-parsers - now version is 0.7, because project core/trunk/exo.core.component.document was not build with version 0.8. 10) org.apache.ws.commons:ws-commons-util - now version is 1.0.1, because version 1.0.2 was not exists (http://apache.cp.if.ua//ws/commons/util/binaries/) 11) org.aspectj:aspectjrt - now version is 1.6.10, because final version 1.6.11 not exists (http://www.eclipse.org/aspectj/downloads.php) 12) org.infinispan:infinispan-core infinispan-cachestore-jdbc - now version 4.1.0.FINAL, because one test (testDistributedCache(org.exoplatform.services.cache.impl.infinispan.TestAbstractExoCache)) failed with version 4.2.0.FINAL in project kernel/trunk/exo.kernel.component.ext.cache.impl.infinispan.v4. 13) org.slf4j:slf4j-log4j12 slf4j-api - now version is 1.5.8, because project kernel/trunk/exo.kernel.commons was not build with version 1.6.1. 14) org.w3c:sac - now version is 1.3, because version 1.3.1 was not exists (http://www.w3.org/Style/CSS/SAC/). 15) quartz:quartz - now version is 1.5.2, because project kernel/trunk/exo.kernel.component.common was not build with version 1.8.4. The version 1.8.4 was not present on maven repository (now this dependency was stored in local maven repository). 16) stax:stax-api - now version is 1.0.1, because version 1.2 was not exists. 17) xpp3:xpp3 - now version is 1.1.4c, because version 1.1.5 was not exists. Should the test and build issues be fixed for the latest library versions? Or it is better to leave versions unchanged for those libraries that caused problems? Should the dependency that not present on maven repositories be stored on our public maven repository?	0	0	https://issues.jboss.org/browse/EXOJCR-1159
4088	EXOJCR-1157	Successfully upgraded to lucene 2.9. What was changed/done: 1. Query class: changed definition (visibility modifier) of createWeight(..) method. 2. Interface Weight replaced with it's abstract type. 3. Weight.score(IndexReader) replaced with Weight.score(IndexReader, boolean, boolean). 4. IndexReader.getSequenceReaders() introduced and used in Lucene internals for per-segment search and caching. In our custom IndexReader should be overridden and return null (This is safe solution and confirmed within lucene's mailing list by Simon Willnauer). 5. FSdirectory.getDirectory(...) replaced with FSdirectory.open(...). But open(...) doesn't create a directory structure anymore. 6. Term can be null, when reader.termDocs(Term) invoked. Adding safe check to avoid NPE. 7. Got rid of HitCollector in favor of Collector class.	0	0	https://issues.jboss.org/browse/EXOJCR-1157
4089	EXOJCR-1156	How ConfigurationFormat class demands creation of SimpleExoLog? I'm interesting do we really need this staff in production?	1	1	https://issues.jboss.org/browse/EXOJCR-1156
4090	EXOJCR-1148	How can I reproduce usecase?	1	1	https://issues.jboss.org/browse/EXOJCR-1148
4091	EXOJCR-1144	Can you perform the same test as your first comment to see if it helps?	1	1	https://issues.jboss.org/browse/EXOJCR-1144
4092	EXOJCR-1105	Could you please give more details about your need because it is not clear for me?	0	0	https://issues.jboss.org/browse/EXOJCR-1105
4093	EXOJCR-1104	Do we need to load value data? or just empty property data?	1	1	https://issues.jboss.org/browse/EXOJCR-1104
4094	EXOJCR-1103	>> So we can not make simple query for pattern like - "e*o:name" << The idea would be to get all the properties/nodes that have a prefix and which are called name, then filter manually thanks to the ItemDataNamePatternFilter >> On other hand, What should we return for "exo*name" pattern? "exo:name"? << all the properties/nodes with no prefix which name match with this pattern So the expected behavior is: "na*me | e*o:name" 1) sql query '[]na%me' OR '[%]name' 2) check resulting set with ItemDataNamePatternFilter "exo:na*me" 1) sql query '[http://www.exoplatform.com/jcr/exo/1.0]na%me' 2) return result	0	0	https://issues.jboss.org/browse/EXOJCR-1103
4095	EXOJCR-1082	RepositoryServiceImpl is an internal class so you are not supposed to create any instance of it by your own, it is totally forbidden. Could you show me your code to allow me to understand what you did?	0	0	https://issues.jboss.org/browse/EXOJCR-1082
4096	EXOJCR-1080	Can't you send the query as a String and let the coordinator parses it for you? if parsing a query is not too resource consuming it could make sense	1	1	https://issues.jboss.org/browse/EXOJCR-1080
4097	EXOJCR-1080	About RPCManager from JBC. It look strictly implemented for internal need of JBC, it not the thing we need. First of all it's not clear how can we control marshaling/unmarshaling since it doesn't use JVM's serialization. It relies it's own internal tool. It's not clear how can we manage commands. I.e. in eXo RPCService, when command registered - we can inject whatever we need inside, using JBC's RPCManager it looks impossible in such a way. And it's completely not documented. Actual question is why should JBC's RPCManager be used, when we have our own?	1	1	https://issues.jboss.org/browse/EXOJCR-1080
4098	EXOJCR-1080	The problem with the RPCService is that it has not been well tested to be considered as production ready and it has not been designed to get the best performances. Moreover, the goal is to backport this feature in JCR 1.12.x and I did not plan to backport the RPCService. If you test your feature that is based on something on which we cannot really trust in term of performances how can you know if your results are relevant?	1	1	https://issues.jboss.org/browse/EXOJCR-1080

#	Requirement	Questions	Answered	Successfully Answered	URL
4099	EXOJCR-1080	To clarify the work the we did and we still need to do: What we did: 1. We try to execute the queries on the coordinator but it was a failure so we abort this study, so the related changes can even be reverted if. 2. We try to replicate Lucene Document instead of ids, the results are not as good as expected but at least they are better than what we currently have so we keep them The next steps are: 3. Instead of having one single volatile index in the coordinator, we will allow to have a volatile index in all the cluster nodes (even if it can cause consistency issues) but only the coordinator will be able to flush the volatile index into the disc. This way the other nodes will be aware of the latest changes sooner. When the coordinator flushes the volatile index, all the volatile indexes must be cleared 4. We will propose another search filter that will allow to have lucene indexes per cluster node, this approach is in fact 2 + 3 + the fact that all nodes will be able to flush their volatile indexes independently (like in JCR 1.10.x). The indexes will be local so we should be more scalable	1	1	https://issues.jboss.org/browse/EXOJCR-1080
4100	EXOJCR-1080	Why the RPC Service usage was discontinued? We can't use it no longer? How can we send commands (like "cleanup volatile now") through the cluster nodes?	1	1	https://issues.jboss.org/browse/EXOJCR-1080
4101	EXOJCR-1080	Why the RPC Service usage was discontinued? We can't use it no longer? > Because we have already something that works well without so I prefer to keep on using the same logic it is safer How can we send commands (like "cleanup volatile now") through the cluster nodes? > For me, the best way would be to change the method IndexUpdateMonitor.setUpdateInProgress(boolean updateInProgress, boolean persitentUpdate) into setUpdateInProgress(boolean updateInProgress, UpdateType type) in order to give more information about the current update, we will then propagate the id of the UpdateType that could be a byte over the cluster instead of the boolean (knowing that java uses a byte to serialize a boolean)	1	1	https://issues.jboss.org/browse/EXOJCR-1080
4102	EXOJCR-1022	Could you please describe more why providing a POJOValue would be helpful? maybe you could give a concrete example	1	1	https://issues.jboss.org/browse/EXOJCR-1022
4103	EXOJCR-1021	Can you point me to a test that do a conversion fron a POJO with jaxb annotations into json ?	0	0	https://issues.jboss.org/browse/EXOJCR-1021
4104	EXOJCR-1021	Do you expect JSON serialization to take in account @XmlTransient ? Which by idea "Prevents the mapping of a JavaBean property to XML representation. "	0	0	https://issues.jboss.org/browse/EXOJCR-1021
4105	EXOJCR-1021	Does your test case work on Jersey implementation? Personally, I don't find it really intuitive for a developer to use JAXB annotations to serialize JSON Objects knowing that most of them only makes sense for XML serialization which is actually what JAXB is used for.	1	1	https://issues.jboss.org/browse/EXOJCR-1021
4106	EXOJCR-1021	I did not tried with Jersey. We could also re-invent the wheel and create our own annotations. Actually supporting the JAXB annotation for the Json conversion does not prevent us from creating our own set of annotation if we think it has value. I just think that it would be nice to support JAXB annotation. The problem with inventing another annotation set is that for a bean that are converted in Json and XML, we need to have 2 annotations per property. And if we add another writer, let say something stupid: CSV, we will have to add more annotations again. {code} @JSONObjectValue(name = "bean") @XmlRootElement(name = "bean") public static class Bean { private String name; private String password; private String shouldNotSerialize = "SHOULD NOT APPEARS"; @JSONValue @XmlElement public String getName() { return name; } public void setName(String name) { this.name = name; } @JSONValue @XmlElement public String getPassword() { return password; } public void setPassword(String password) { this.password = password; } @JSONTransient @XmlTransient public String getShouldNotSerialize() { return shouldNotSerialize; } public void setShouldNotSerialize(String shouldNotSerialize) { this.shouldNotSerialize = shouldNotSerialize; } } {code} So to keep things simple, it would be better to reuse JAXB annotations. It worked well for my webservices.	0	0	https://issues.jboss.org/browse/EXOJCR-1021
4107	EXOJCR-1021	You were saying that JAXB annotation would not fit to JSON, so I deduced that you would prefer to develop non-standard annotation. Do you know a standard to do similar things with Json ?	0	0	https://issues.jboss.org/browse/EXOJCR-1021
4108	EXOJCR-1021	I am working on small improvement of our REST framework which will make possible to overwrite default providers (readers and writers). So, you will be able to replace out default JSON serializer (which is light and suitable for absolutely most the cases) w/ one from Jackson or whatever you need. This improvement will be available in several days in trunk (next CF release if you prefer stable), I'll let you know details then. Is it OK for everyone?	1	1	https://issues.jboss.org/browse/EXOJCR-1021
4109	EXOJCR-1021	>> You were saying that JAXB annotation would not fit to JSON << For me no since it is not intuitive at all, mixing JSON with JAXB/XML it is about the same as mixing tomatoes and potatoes >> so I deduced that you would prefer to develop non-standard annotation. << Yes because even if we decided to use JAXB annotations, it still won't be standard, so I rather prefer to use internal annotations and warn the developers in the doc that it is not standard so it should be used with caution >> Do you know a standard to do similar things with Json ? << Personally I don't If you still think that it is critical to have such feature, could you please give us a list of all the annotations that you would like to have to manage the serialization into JSON format? At least I guess that you want one annotation to indicate that the field is transient, maybe another one to be able to change the name of the field and/or object, what else?	1	1	https://issues.jboss.org/browse/EXOJCR-1021
4110	EXOJCR-999	Code research result: SessionCleaner - SessionCleaner is holded by SessionRegistry. SessionRegistry is per-repository. So, there is nothing to do with Session Cleaner. FileCleaner: - there is WorkspaceFileCleanerholder that uses single FileCleaner thread and may be used as per-repository (I did a commit already). - I have found a two FileCleaners that may be created and started, but does not make any usefull works . In SysViewWorkspaceInitializer and BackupWorkspaceInitializer. These initilizer now uses FileCleaner from WorkspaceFileCleanerholder; - FileValueStorage - each FileValueStorage have own FileCleaner. And there is a trouble with TreeFileValueStorage, it uses extended TreeFileCleaner. LockRemover: - each LockManager have own LockRemover, but this LockRemover executes removeExpired(). This operation may be heavy. Do we need single LockRemover thread for all LockManagers?	1	1	https://issues.jboss.org/browse/EXOJCR-999
4111	EXOJCR-999	Hi Sergey, Here are my remarks about your work: 1. WorkerService: Do we really need poolNumber? if so it should be static otherwise the value will always be 1. But I believe that ExoContainerContext.getName()/RepositoryName should be more helpful for debugging, don't you agree? 2. WorkerService: It could make sense to allow to create Deamons instead of normal thread to avoid blocking the AS when we stop it 3. WorkerService: executePeriodically should return object of type ScheduledFuture<?> instead of ScheduledFuture to avoid warnings 4. LockRemover: lockRemoverTask should be of type ScheduledFuture<?> 5. LockRemover: The method start should ignore cases where the lockRemoverTask is not null instead of calling stop method to prevent any issues 6. The FileCleaner should also be based on the WorkerService to be able to have more than one thread for the whole repository if needed 7. The new parameters should be described into the doc	1	1	https://issues.jboss.org/browse/EXOJCR-999

#	Requirement	Questions	Answered	Successfully Answered	URL
4112	EXOJCR-999	Multithread FileCleaner is a bit unclear for me. For now FileCleaner have a list of files that periodically cleaned. Each FileCleaner user may add files to this list How is it supposed to work as multithread? * There is many FileCleaner object , each have own file-list. These FileCleaners registered as tasks in WorkerService. What is bad - each call of getFileCleaner makes new class instance. Maybe you have own vision how it might work.	1	1	https://issues.jboss.org/browse/EXOJCR-999
4113	EXOJCR-998	Hi Sergey, Here are my first remarks: 1. Could you please remove old code instead of commenting it in order to have a code much more clear? 2. When you check if a given id is a NULL_ID, I don't think that it is necessary to use "equalsIgnoreCase", "equals" should be enough 3. In the method get(String parentId, QPathEntry name, ItemType itemType), if itemType == ItemType.PROPERTY and itemId equals NullNodeData.NULL_ID, it will return a NullNodeData instead of a Property. Which probably means that we need 2 other classes which will be NullItemData and NullPropertyData. NullNodeData should be used when itemType = UNKNOWN or NODE 4. The method "get(String id)" should return "getFromCacheById(id)" since those 2 methods do exactly the same things 5. When it is possible it would be cleaner to test the type of the item (instanceof NullItemData or NullPropertyData or NullNodeData) instead of testing if its path is null or its id equals NULL_ID 6. In putProperty(PropertyData prop, ModifyChildOption modifyListsOfChild), we should manage the case where the PropertyData is a NullPropertyData 7. Remove the method putNullNode, a manage null values directly in putNode and putProperty or ignore #6 and remove your changes related to null values in putNode, otherwise the code doesn't look consistent 8. In CacheableWorkspaceDataManager, when a node has no properties or sub nodes, store an empty list into the cache	1	1	https://issues.jboss.org/browse/EXOJCR-998
4114	EXOJCR-992	This is clearly not a bug at best it is an improvement but frankly speaking I don't believe that it is a good approach to manipulate Strings to create your result since it is really error prone, for example if your name contains a special character like '<' in the helloxml method or ':' in the hello method, your result will be incorrect. You should manipulate objects instead (which is for me much more convenient for developers) and let the framework de-serialize and serialize the objects for you. Andrey what do you think about that?	1	1	https://issues.jboss.org/browse/EXOJCR-992
4115	EXOJCR-991	Why need to set mixin via WebDAV? We can simple set directly to node via Node.addMixin(String mixinName) method.	0	0	https://issues.jboss.org/browse/EXOJCR-991
4116	EXOJCR-985	could you please describe more what you mean?	1	1	https://issues.jboss.org/browse/EXOJCR-985
4117	EXOJCR-934	Yes, I know that the way to add listeners is supposed to be via the component-plugin mechanism. This is actually what I used first. However this approach has a major drawback for my use case: I cannot retrieve the listener instance created by the kernel anymore and I need to get that instance so that I can actually pass it to other objects (the component is bridging GateIn session events so that the WSRP module can be made aware of them). More generally speaking, if a component is assigned a different name in the xml configuration, doesn't that mean that the component will be instantiated twice? I need the same component to handle both events, so that wouldn't fit my use case either. Globally, I think it's a rather dubious design to change the name of a component in a method that should be without side-effect to the passed object, so yes, I think that the addListener(String eventName, Listener listener) should be the proper way to add a listener and that it shouldn't change the passed object.	1	1	https://issues.jboss.org/browse/EXOJCR-934
4118	EXOJCR-920	Can someone assign this JIRA to me? thanks.	0	0	https://issues.jboss.org/browse/EXOJCR-920
4119	EXOJCR-887	2 nodes cluster usecases: 1) two nodes starts with interval (1-2 seconds) Second node always wait with message 17.11.2010 10:48:09 *INFO * [main] RepositoryCreationSynchronizer: Waiting to be released by the coordinator (RepositoryCreationSynchronizer.java, line 276) 2) two nodes starts at same time Some times, no one node become waiting but starts successfull (there was no messages like *INFO * [main] RepositoryCreationSynchronizer: Waiting to be released by the coordinator (RepositoryCreationSynchronizer.java, line 276)) 3 nodes cluster usecases: 1) all nodes starts with interval (1-2 sec) one by one Second and third node always wait with message 17.11.2010 10:48:09 *INFO * [main] RepositoryCreationSynchronizer: Waiting to be released by the coordinator (RepositoryCreationSynchronizer.java, line 276) 2) one node started , than 2 and 3 node starts at same time no errors 3) all nodes starts at same time: 1 node got exception 17.11.2010 10:34:36 *ERROR* [main] JDBCWorkspaceDataContainer: Error of init db org.exoplatform.services.jcr.impl.util.jdbc.DBInitializerException: Could not cr eate db schema of DataSource: 'system'. Reason: Duplicate entry ' ' for key 'PRI MARY'; . Last command: CREATE UNIQUE INDEX JCR_IDX_SREF_PROPERTY ON JCR_SREF(PRO PERTY_ID, ORDER_NUM) (CQJDBCWorkspaceDataContainer.java, line 258) org.exoplatform.services.jcr.impl.util.jdbc.DBInitializerException: Could not cr eate db schema of DataSource: 'system'. Reason: Duplicate entry ' ' for key 'PRI MARY'; . Last command: CREATE UNIQUE INDEX JCR_IDX_SREF_PROPERTY ON JCR_SREF(PRO PERTY_ID, ORDER_NUM) at org.exoplatform.services.jcr.impl.util.jdbc.DBInitializer.init(DBInitializer.java:457) but 2 and 3 node waits with message 17.11.2010 10:34:39 *INFO * [main] RepositoryCreationSynchronizer: Waiting to be released by the coordinator (RepositoryCreationSynchronizer.java, line 276) 4)all nodes starts at same time (another exceptions): 1-node got exception: 17.11.2010 10:38:16 *ERROR* [main] JDBCWorkspaceDataContainer: Error of init db org.exoplatform.services.jcr.impl.util.jdbc.DBInitializerException: Could not cr eate db schema of DataSource: 'system'. Reason: Table 'jcr_sref' already exists; . Last command: CREATE TABLE JCR_SREF(NODE_ID VARCHAR(96) NOT NULL, PRO PERTY_ID VARCHAR(96) NOT NULL, ORDER_NUM INTEGER NOT NULL, CONSTRAINT JCR_ PK_SREF PRIMARY KEY(NODE_ID, PROPERTY_ID, ORDER_NUM)) (CQJDBCWorkspaceDataCont ainer.java, line 258) org.exoplatform.services.jcr.impl.util.jdbc.DBInitializerException: Could not cr eate db schema of DataSource: 'system'. Reason: Table 'jcr_sref' already exists; . Last command: CREATE TABLE JCR_SREF(NODE_ID VARCHAR(96) NOT NULL, PRO PERTY_ID VARCHAR(96) NOT NULL, ORDER_NUM INTEGER NOT NULL, CONSTRAINT JCR_ PK_SREF PRIMARY KEY(NODE_ID, PROPERTY_ID, ORDER_NUM)) at org.exoplatform.services.jcr.impl.util.jdbc.DBInitializer.init(DBInitializer.java:457) 3-node got exception: 17.11.2010 10:38:16 *ERROR* [main] JDBCWorkspaceDataContainer: Error of init db org.exoplatform.services.jcr.impl.util.jdbc.DBInitializerException: Could not cr eate db schema of DataSource: 'system'. Reason: Duplicate key name 'JCR_IDX_SREF _PROPERTY'; . Last command: CREATE UNIQUE INDEX JCR_IDX_SREF_PROPERTY ON JCR_SRE F(PROPERTY_ID, ORDER_NUM) (CQJDBCWorkspaceDataContainer.java, line 258) org.exoplatform.services.jcr.impl.util.jdbc.DBInitializerException: Could not cr eate db schema of DataSource: 'system'. Reason: Duplicate key name 'JCR_IDX_SREF _PROPERTY'; . Last command: CREATE UNIQUE INDEX JCR_IDX_SREF_PROPERTY ON JCR_SRE F(PROPERTY_ID, ORDER_NUM) at org.exoplatform.services.jcr.impl.util.jdbc.DBInitializer.init(DBInitializer.java:457) 1-node keep strating , 2-node and 3-node become waiting with message 17.11.2010 10:38:19 *INFO * [main] RepositoryCreationSynchronizer: Waiting to be released by the coordinator (RepositoryCreationSynchronizer.java, line 276) Summary: there is no problems if cluster-nodes starts with small interval (1-2s), but if nodes started in same time, there is unexpected behaviour, no matter 2 or 3 nodes cluster.	1	1	https://issues.jboss.org/browse/EXOJCR-887
4120	EXOJCR-887	If so those errors are not a problem, maybe they should be warnings instead of errors, don't you agree? otherwise users could believe that it is an issue	1	1	https://issues.jboss.org/browse/EXOJCR-887
4121	EXOJCR-887	those exceptions happen because all the nodes try to create the tables and indexes at the same time. If they already exist, it means that the work has already been done so it is not a problem as soon as it doesn't prevent the node to start and work normally, don't you agree? If not what do you propose?	0	0	https://issues.jboss.org/browse/EXOJCR-887

#	Requirement	Questions	Answered	Successfully Answered	URL
4122	EXOJCR-886	Thnx. Also question about time-date format for spreadsheets. Currently "yyyy-MM-dd HH:mm:ss.SSSZ" is used, so that TIMEZONE of host machine added. May be timezone should be removed? WDYT? Because of it, jcr-index created on different machines can be a bit different (it include some data from host machine settings), tests for this are also a bit more complicated, since they can be launched in different time-zones.	1	1	https://issues.jboss.org/browse/EXOJCR-886
4123	EXOJCR-874	Could you get the same profiler results with JCR 1.12.x to identify the differences?	1	1	https://issues.jboss.org/browse/EXOJCR-874
4124	EXOJCR-783	I've just checked your patch and I did not see any relationship with this task, so I don't really understand why it is required only for this particular feature. Don't we have to apply it anyway? If so please create a dedicated JIRA. BTW please note that: 1. The fix versions are 1.12 and 1.14 and the privileged block is only for 1.14 2. The feature will be proposed only if it doesn't affect too much the performance on GateIn, if you need twice more time to launch GateIn, I really believe that it is useless, so we need to check it first	1	1	https://issues.jboss.org/browse/EXOJCR-783
4125	EXOJCR-783	What about the fix related to the "privileged block"? Is it always required or it is only required for this task?	1	1	https://issues.jboss.org/browse/EXOJCR-783
4126	EXOJCR-778	Protect means check permissions for each method mentioned? Do you mean one permission for all Repository methods or a dedicated for each method?	0	0	https://issues.jboss.org/browse/EXOJCR-778
4127	EXOJCR-778	<< Protect means check permissions for each method mentioned? >> Yes << Do you mean one permission for all Repository methods or a dedicated for each method? >> One generic RuntimePermission for all those methods could me enough such as for example "manageRepository"	1	1	https://issues.jboss.org/browse/EXOJCR-778
4128	EXOJCR-769	@Julien, Could you help please, where is the source code which need to be ported?	0	0	https://issues.jboss.org/browse/EXOJCR-769
4129	EXOJCR-769	What about the most important class aka RestManagementProvider.java? Why didn't you port it? Without the provider it cannot work	1	1	https://issues.jboss.org/browse/EXOJCR-769
4130	EXOJCR-752	I understand you point of view however don't forget that jcr 1.12 and 1.14 must remain compatible which means that the old method such as DocumentReader.getContentAsText must still be supported. I guess that we will need to create a new interface such as AdvancedDocumentReader that will extends DocumentReader that will propose new methods such as : * Reader getContentAsReader(InputStream is, String encoding) throws IOException, DocumentReadException * Reader getContentAsReader(InputStream is) throws IOException, DocumentReadException So in our code we will first check if the instance of DocumentReader found is a subclass of AdvancedDocumentReader, if so we will use getContentAsReader otherwise we will use getContentAsText. By default in JCR 1.14, we will propose only one implementation of AdvancedDocumentReader which will be the Tika implementation BTW in which format do we get the metadata with Tika? Is-it a Map or a Properties object? Because according to the type we could have to add another method in AdvancedDocumentReader for the metadata	1	1	https://issues.jboss.org/browse/EXOJCR-752
4131	EXOJCR-749	What I expect for this task: 1. An implementation of the new interface "AdvancedDocumentReader" for Apache Tika 2. This implementation will be registered into the configuration file (the old ones will be removed) with the same mimetypes in order to support at least the same file format 3. The configuration files of Tika must be retrieved thanks to the ConfigurationManager in order to rely on eXo kernel	0	0	https://issues.jboss.org/browse/EXOJCR-749
4132	EXOJCR-749	General thoughts, perhaps at JCR 1.15. Why do we need DocumentReaderService? What it can do this do not make Tika? In my opinion Tika API has a well sophisticated and well-supported structure. DocumentReaderService not added anything new to them. I believe that to maintain it in the future makes no sense and it would be logical to move to the exo-configuration only the configuration of Tika service.	0	0	https://issues.jboss.org/browse/EXOJCR-749
4133	EXOJCR-744	R&D "LIC-709" branch successfully upgraded to lucene 2.9. What was changed/done: 1. Query class: changed definition (visibility modifier) of createWeight(..) method. 2. Interface Weight replaced with it's abstract type. 3. Weight.score(IndexReader) replaced with Weight.score(IndexReader, boolean, boolean). 4. IndexReader.getSequenceReaders() introduced and used in Lucene internals for per-segment search and caching. In our custom IndexReader should be overridden and return null (This is safe solution and confirmed within lucene's mailing list by Simon Willnauer). 5. FSdirectory.getDirectory(...) replaced with FSdirectory.open(...). But open(...) doesn't create a directory structure anymore.	0	0	https://issues.jboss.org/browse/EXOJCR-744
4134	EXOJCR-739	I've added files: - product-exo-jcr-as-jboss-ear.xml - product-exo-jcr-as-jonas-ear.xml - product-exo-jcr-as-tomcat6.xml to <modules> in jcr/applications/pom.xml so maven will use them like simple modules and will update version automatically. Other changes related to this issue: 1. artifactId was changed in this files now it is: exo.jcr.applications.jboss exo.jcr.applications.jonas exo.jcr.applications.tomcat instead of exo.product.jcr 2. Packaging of poms chnged to pom (was jar) 3. A special profile for deploying named "deploy" added. This was done to reduce time of jcr build. We don't need to deploy application servers each time we build JCR 4. to build Application server run "mvn -f product-exo-jcr-as-tomcat6.xml clean install -P deploy" command.	0	0	https://issues.jboss.org/browse/EXOJCR-739
4135	EXOJCR-725	Which collation is proposed to be used in MySQL and MSSQL dialects?	0	0	https://issues.jboss.org/browse/EXOJCR-725
4136	EXOJCR-725	A collation usually set on a database not on a table or column, on MySQL our DBA executes: create database jcrtest default charset latin1 collate latin1_gneral_cs;. And from this point of view I manage this issue. JCR creates tables but not a database itself, it should be ready before first start of JCR service. If you haven't an opportunity to ask a DBA for a right db, we have a new service in 1.12 - org.exoplatform.services.database.creator.DBCreator. This service cares about a stuff JCR doesn't - a database creation (via ready script) - http://wiki.exoplatform.org/xwiki/bin/view/Core/Database+Creator (docs also in docbook on svn) An application can configure it: {code:xml} <component> <key>org.exoplatform.services.database.creator.DBCreator</key> <type>org.exoplatform.services.database.creator.DBCreator</type> <init-params> <properties-param> <name>db-connection</name> <description>database connection properties</description> <property name="driverClassName" value="com.mysql.jdbc.Driver" /> <property name="url" value="jdbc:mysql://localhost/" /> <property name="username" value="root" /> <property name="password" value="admin" /> </properties-param> <properties-param> <name>db-creation</name> <description>database creation properties</description> <property name="scriptPath" value="createdb.sql" /> <property name="username" value="testuser" /> <property name="password" value="testpwd" /> </properties-param> </init-params> </component> {code} Where createdb.sql will contains smth like: create database jcrtest default charset latin1 collate latin1_gneral_cs. Then create a dedicated database calling DBCreator.createDatabase(String name). This database should be bind to a JNDI with a DataSource name will be used in JCR. It's can be done via InitialContextInitializer.bind(String, String, String, String, Map<String, String>). Will it be useful?	1	0	https://issues.jboss.org/browse/EXOJCR-725

#	Requirement	Questions	Answered	Successfully Answered	URL
4137	EXOJCR-718	Can we add REST binding in addition to JMX ? REST has the advantage of only requiring a browser to invoke the operation. For support it is a great advantage.	1	1	https://issues.jboss.org/browse/EXOJCR-718
4138	EXOJCR-718	Is REST secured enough for such feature?	1	1	https://issues.jboss.org/browse/EXOJCR-718
4139	EXOJCR-716	Abstract of what do do : Leverage the ContainerLifeCyclePlugin inside the JobSchedulerService so that we can start the job only when all other components have finished starting. Step 1 : Hook the JobSchedulerService to ContainerLifeCyclePlugin public JobSchedulerService(ExoContainerContext ctx) { ExoContainer container = ctx.getContainer(); ctx.addContainerLifecylePlugin(new JobStartupController(this)); } Step 2 : JobStartupController This component will be notified when the container has started. It is responsible for accumulating all job initialization request and temporize them in a queue class JobStartupController extends BaseContainerLifecyclePlugin { @Override public void startContainer(ExoContainer container) throws Exception { processJobQueue(); // execute all accumulated jobs allowJobsToStart(); // the queue will no longer block } } Step 3 : the Job queue instead of forwarding operations to the underlying quartz instance, all operations are queued until a glag allows the queue to be processed.	0	0	https://issues.jboss.org/browse/EXOJCR-716
4140	EXOJCR-711	From commit comment: In J2EEServerInfo, a "System.out.println" has been replaced with a "log.info". At time of J2EEServerInfo use the logging system can be not yet initialized, e.g. on the system start. In this case we can lost log output. Did you checked this usecase with GateIn on Tomcat/JBoss?	1	0	https://issues.jboss.org/browse/EXOJCR-711
4141	EXOJCR-693	This article is already ported into docbook: https://svn.jboss.org/repos/exo-jcr/jcr/trunk/docs/reference/en/src/main/docbook/en-US/modules/jcr/multilanguage-support.xml Is that what you meant ?	1	0	https://issues.jboss.org/browse/EXOJCR-693
4142	EXOJCR-692	>> Method WorkspaceStorageCache.getChildNodesCount() does not help us. We in any case should get all childs from cache to be sure that all of them are not evicted. << Why? we don't care whether the nodes are evicted or not, we only want to know how many child nodes we have into the cache that can be known thanks to the total amount of keys found into the cache. Node.hasNodes is an important method for the JCR, so we need to try to improve the perfs, so please at least test this solution in order to know if it improves the perf or not.	1	1	https://issues.jboss.org/browse/EXOJCR-692
4143	EXOJCR-692	Good news, what did you do to get those results?	1	1	https://issues.jboss.org/browse/EXOJCR-692
4144	EXOJCR-684	Upgrade to 2.9 lucene on LIC-709 "Single index" was finished and NRT was tried for the first time. Result of first daily benchmark attached. In general, this POC shows good results, much similar to original index implementation, some tests even show 3x increase, but some of them show decrease. The reason of decrease in absolutely unknown up to now and should be profiled. * Single-Index currently has no RedoLog (comparing to trunk's version, have a RedoLog, that writes some bytes of data on FS, each time save() is called). Also some aspects of NRT usage are now under investigation. They are: * When commit is invoked and how can it be listened and/or controlled. * When optimize is called and how can it be controlled. * Performance decrease in some tests. * NRT interns and failover capabilities/external possibilities.	0	0	https://issues.jboss.org/browse/EXOJCR-684
4145	EXOJCR-684	What is known up to now: * autoCommit feature is deprecated in 2.9+ and should be controlled out of Lucene. So we can manage commit operation in the same way as before, providing the same "commit on idle", "commit when buffer filled" features. * There is an optimize operation, that re-builds index, making search faster. This procedure is / should be launched only manually and periodically; * failover facilities can be provided using the same RedoLog-based mechanism. Having control over the commit, we can easily manage and cleanup RedoLog in time; So SingleIndex with NRT (near real time) search can serve as an indexing subsystem in JCR replacing current MultiIndex implementation. Making it as production ready system needs some time to implement missing in POC features and have reasonable set of tests. Benefits: * Much simpler code, much easy to support; Disadvantages: * Time to implement/test; * Not tested in production, as existing impl. Under investigation: * Search should be a bit faster with larger indexes. * Possibly this implementation can give performance increase in concurrent environment, when multiple threads are writing and/or reading. Note: * is not directly affiliated with clustered implementation.	0	0	https://issues.jboss.org/browse/EXOJCR-684
4146	EXOJCR-680	What you propose won't work with FTP so for me it cannot be done this way. Moreover, you assume that the content is in HTML format which is not always the case. And finally this feature doesn't seem to be valid according to the JAX-RS specification. What do you want to do exactly? You would like to see eXo JCR index your HTML metadata when you add it though Webdav and/or FTP, is-it what you are trying to do?	1	1	https://issues.jboss.org/browse/EXOJCR-680
4147	EXOJCR-640	ok, but how can I test this JS to be sure everything is fine? Because PDFBox groupId in maven's repositories was changed from "pdfbox" to "org.apache.pdfbox" and this should be tested.	0	0	https://issues.jboss.org/browse/EXOJCR-640
4148	EXOJCR-634	The "replace forbidden chars" mechanism was ported. By default "replace forbidden chars" mechanism is enabled. By default the forbidden chars is ':', '[', ']', '*', '"', '\', '|'. By default the forbidden chars will be replaced to '_'. How to configuration in ftp config when need change defaults config : <value-param> <name>replace-forbidden-chars</name> <value>true</value> </value-param> <value-param> <name>forbidden-chars</name> <value>:[]*'"|</value> </value-param> <value-param> <name>replace-char</name> <value>_</value> </value-param>	0	0	https://issues.jboss.org/browse/EXOJCR-634
4149	EXOJCR-625	@Dmytro is-this task over? if not what else do we need to do?	1	0	https://issues.jboss.org/browse/EXOJCR-625
4150	EXOJCR-624	Since you won't fixed this issue. Can you let us know where/how you intend to implement it with Tika ?	1	0	https://issues.jboss.org/browse/EXOJCR-624
4151	EXOJCR-610	Which JCR features should be described first, i.e. which are most used in portal applications now?	1	0	https://issues.jboss.org/browse/EXOJCR-610
4152	EXOJCR-610	What kinds of guidance it should cover? - just API (huge doc, but means the spec retell) - any predefined usecases (dedicated to a task(s)) - integration questions (dedicated also, but to which other apps so)	0	0	https://issues.jboss.org/browse/EXOJCR-610

#	Requirement	Questions	Answered	Successfully Answered	URL
4153	EXOJCR-609	[Thread[pool-13-thread-6381,5,main]]! Lock held by [null] (NodeIndexer.java, line 403) org.jboss.cache.lock.TimeoutException: Unable to acquire lock on Fqn [/$ITEMS/$] after [20000] milliseconds for requestor [Thread[pool-13-thread-6381,5,main]]! Lock held by [null] at org.jboss.cache.mvcc.MVCCNodeHelper.acquireLock(MVCCNodeHelper.java:159) at org.jboss.cache.mvcc.MVCCNodeHelper.wrapNodeForWriting(MVCCNodeHelper.java:217) at org.jboss.cache.mvcc.MVCCNodeHelper.wrapNodeForWriting(MVCCNodeHelper.java:186) at org.jboss.cache.interceptors.MVCCLockingInterceptor.handlePutKeyValueCommand(MVCCLockingInterceptor.java:101) at org.jboss.cache.interceptors.base.PrePostProcessingCommandInterceptor.visitPutKeyValueCommand(PrePostProcessingCommandInterceptor.java:88) at org.jboss.cache.commands.write.PutKeyValueCommand.acceptVisitor(PutKeyValueCommand.java:100) at org.jboss.cache.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:116) at org.jboss.cache.interceptors.ReplicationInterceptor.handleCrudMethod(ReplicationInterceptor.java:148) at org.jboss.cache.interceptors.ReplicationInterceptor.visitPutKeyValueCommand(ReplicationInterceptor.java:107) at org.jboss.cache.commands.write.PutKeyValueCommand.acceptVisitor(PutKeyValueCommand.java:100) at org.jboss.cache.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:116) at org.jboss.cache.interceptors.base.CommandInterceptor.handleDefault(CommandInterceptor.java:131) at org.jboss.cache.commands.AbstractVisitor.visitPutKeyValueCommand(AbstractVisitor.java:65) at org.jboss.cache.commands.write.PutKeyValueCommand.acceptVisitor(PutKeyValueCommand.java:100) at org.jboss.cache.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:116) at org.jboss.cache.interceptors.TxInterceptor.attachGtxAndPassUpChain(TxInterceptor.java:301) at org.jboss.cache.interceptors.TxInterceptor.handleDefault(TxInterceptor.java:283) at org.jboss.cache.commands.AbstractVisitor.visitPutKeyValueCommand(AbstractVisitor.java:65) at org.jboss.cache.commands.write.PutKeyValueCommand.acceptVisitor(PutKeyValueCommand.java:100) at org.jboss.cache.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:116) at org.jboss.cache.interceptors.CacheMgmtInterceptor.visitPutKeyValueCommand(CacheMgmtInterceptor.java:119) at org.jboss.cache.commands.write.PutKeyValueCommand.acceptVisitor(PutKeyValueCommand.java:100) at org.jboss.cache.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:116) at org.jboss.cache.interceptors.InvocationContextInterceptor.handleAll(InvocationContextInterceptor.java:178) at org.jboss.cache.interceptors.InvocationContextInterceptor.visitPutKeyValueCommand(InvocationContextInterceptor.java:82) at	1	1	https://issues.jboss.org/browse/EXOJCR-609
4154	EXOJCR-592	What was done: query FIND_PROPERTIES_BY_PARENTID_CQ was updated in next order - fields in where constraint was sorted according to exist indexes. Oracle_GetChildNodesCQ graph shows that average execution time of old getChildNodes and getChildProperties queries are growing. Oracle_NewGetChildNodesCQ graph shows that average execution time of new getChildNodes query do not grows.	1	1	https://issues.jboss.org/browse/EXOJCR-592
4155	EXOJCR-583	What can works on JCR 1.12.1-GA-SNAPSHOT with expected conditions?	1	1	https://issues.jboss.org/browse/EXOJCR-583
4156	EXOJCR-551	To make the task JCR related please: 1. Describe the usecase in term of JCR only (i.e. use Node and/or Property without 'page') 2. Describe exact nodetype we should test and how to connect it to a Nodes/Properties given via parameters reads/writes. How those related methods parameters will be calculated or will be read from the configuration? 3. Should nodetype be configurable or hardcoded? 4. Is it a benchmark test or something else? 5. If case of a run on several physical JCRs, should it be synchronous start assuming that we need an initial node start or a manual start? 6. What the tests will measure? One operation time/tps (read or/and write?), whole tests time, other parameters. 7. How to manage the execution of the usecases: by time or by iterations? 8. How to treat results if run on several severs? Add all results, average result, other.	0	0	https://issues.jboss.org/browse/EXOJCR-551
4157	EXOJCR-551	About usage of JCR-sessions. Should session be one per usecase or one per scenario? I.e. 1) Session is one per usecase (per page in Your's terms) and it is opened-closed few times per scenario: read() { session = repository.login(); session.getNode(); session.logout(); } doRun() { read(); read(); write(); } 2) One session per-scenario read(Session session) { session.getNode(...); } doRun() { session = repository.login(); ... read(session); read(session); write(session); ... session.logout(); }	0	0	https://issues.jboss.org/browse/EXOJCR-551

#	Requirement	Questions	Answered	Successfully Answered	URL
4158	EXOJCR-551	what we have done. Test PageUsecasesTest from package org.exoplatform.jcr.benchmark.usecases.portal implements usecases mentioned in issue. On prepare it reads configuration and creates tree of nodes with configured depth. node1 node11 node111 node112 node12 node121 node122 node2 ... Configuration options are: * PREPARE parameters initialization depth, by default 2 is used: <!-- if parameter is not defined, default value "2" is used --> <param name="exo.prepare.depth" value="2" /> child nodes number <!-- if parameter is not defined, default value "10" is used --> <param name="exo.prepare.nodesPerLevel" value="10" /> String property length <!-- if parameter is not defined, default value "64" is used --> <param name="exo.data.stringLength" value="64" /> Number of values in multivalued property <!-- if parameter is not defined, default value "4" is used --> <param name="exo.data.multiValueSize" value="4" /> Size of randomly generated binary value or path to file. If nothing configured, random binary array of 1024 bytes is used <!-- if no binary parameter defined, default value "exo.data.binarySize=1024" is used --> <!--param name="exo.data.binaryPath" value="file.txt" /--> <param name="exo.data.binarySize" value="1024" /> Synchronization options. If nothing configured japex will run doRun as soon as latest doPrepare finishes. If "manual" is used, then after configured number ("exo.waitThread") of doPrepare finishes, it will ask to press a key to start doRun. "exo.waitThread" must be equals to the number of threads configured in japex driver. <!-- how to sunchronize cluster nodes. Default value is NONE --> <!--param name="exo.synchronisation" value="manual"/> <param name="exo.waitThread" value="5"/--> To define whether to remove or not the tree of nodes created for test <!-- Defines whether cleanup created nodes on finish or not. True by default. --> <param name="exo.finish.cleanup" value="false" /> Set scenario to perform. On each doRun only one "page" action will be performed. <!-- String containing scenario in form of usecases separated by ";". Each usecase can have a multiplier, defining that this usecase should be perform few times or may not have it, notifying single invocation. Each usecase has list of argument. Read and ReadAnon usecases has 2 arguments (Number of nodes and properties to read). Write usecase has 5 arguments: number of removed, added, modified properties and removed, added nodes. All strings are CASE INSENSITIVE and whitespaces between elements are skipped. --> <param name="exo.scenario.string" value="2*ReadAnon(2,3);8*Read(2,3);Write(2,3,1,1)" /> Existing actions: * ReadAnon(x,y), Read(x,y) Action that reads x nodes and for each node it reads y properties. ReadAnon uses anonymous session for operations. * Write(w,x,y,z) Action removes w properties, sets x properties, removes y nodes and adds z nodes. What is done in another way. 1. For stability of results, on each operation the node of latest level is used. (i.e. node121 on diagram above). For add/remove node the node of previous level is used (i.e node11). This is done in such a way, cause node.getNodes() (get all the child nodes) is rather heavy operation and if it would be called for different actions different number or times, the result of the test will have some chaotic manner. 2. JCR doesn't have "addProperty" and checking each time what properties are missing in node - will require read operation for all the 20 properties to find what is missing. so only "setProperty" operation is performed, instead of add and modify operations.	0	0	https://issues.jboss.org/browse/EXOJCR-551
4159	EXOJCR-536	How would you like to configure the directory name "exo-conf"? with a JVM parameter?	1	1	https://issues.jboss.org/browse/EXOJCR-536
4160	EXOJCR-535	It's not clear how current implementation - versions storage in system workspace is worse. Can you point the benefits visible from your side within a dedicated versionned-workspace? Only the sizing issue is visible now. And a migration question will be an another complicated task. At other hand JCR spec describes a version storage as a some place visible in each workspace at the path /jcr:system/jcr:versionStorage (8.2.2.1 jcr:versionStorage). I.e. we can store in hidden workspace or even directly in special persistence. But it's a big enough rework of JCR core. So, the benefits question interesting first.	1	0	https://issues.jboss.org/browse/EXOJCR-535
4161	EXOJCR-528	How do you define settings for the default portal ?	1	1	https://issues.jboss.org/browse/EXOJCR-528
4162	EXOJCR-528	<< How do you define settings for the default portal ? >> I don't understand your question, can you describe a concrete usecase?	0	0	https://issues.jboss.org/browse/EXOJCR-528
4163	EXOJCR-517	Jeremi, this issue not clear for our WS guys. I let Andrew to ask technical questions. But if the solution already exists (you link to japanese code) and you found it applicable, can you implement it or provide a patch with possible code? At least it will make the understanding of the task more clear.	0	0	https://issues.jboss.org/browse/EXOJCR-517
4164	EXOJCR-517	This japanese code is not application as it is. It needs to use the exo's user base. @Andrew, What is not clear?	1	1	https://issues.jboss.org/browse/EXOJCR-517
4165	EXOJCR-517	Is it what you mind or can we adapt existed code ?	1	1	https://issues.jboss.org/browse/EXOJCR-517
4166	EXOJCR-517	Have you looked at the the documentation I sent you, and at the code? The point is to validate a signature and if it's valid, to authenticate the user using opensocial_viewer_id. For exemple, it's something that could be really useful for the IDE since it would remove the need to do the basic auth.	1	1	https://issues.jboss.org/browse/EXOJCR-517
4167	EXOJCR-480	As a JDBC user (JCR here) it's not clear why the user cannot use JDBC API fully and be involved into XA transaction transparently. JCR manages JDBC transactions itself and right work with JTA should looks as a sequence of JDBC commits (or rollbacks) executed in global transaction, then the global transaction will be committed and all data will be persisted (including committed via JDBC connections; but rolled back data do not). Why only rollback causes a problem so? And why JDBC commit is ok so? Can be such behaviour (transparency) configured via JBoss Transaction Service or other XA related stuff?	0	0	https://issues.jboss.org/browse/EXOJCR-480
4168	EXOJCR-480	Did you configure the DataSourceProvider as described here http://docs.jboss.org/exojcr/1.14.0-CR4/developer/en-US/html/ch-data-source-provider.html ? Indeed by default all the datasources are considered as non managed.	1	1	https://issues.jboss.org/browse/EXOJCR-480
4169	EXOJCR-480	What is the full stack trace?	1	1	https://issues.jboss.org/browse/EXOJCR-480
4170	EXOJCR-480	Could you please provide your binary to allow me to test it?	1	1	https://issues.jboss.org/browse/EXOJCR-480
4171	EXOJCR-455	Have you seen this article? http://www-01.ibm.com/support/docview.wss?rs=2004&context=SSVHZU&dc=DB550&uid=swg1IZ36922&loc=en_US&cs=UTF-8&lang=en&rss=ct2004tivoli	0	0	https://issues.jboss.org/browse/EXOJCR-455
4172	EXOJCR-439	Why do we need that? since it will be fixed by EXOJCR-434. The variables like ${cacheloader.lock.datasource} can already be used in the repository configuration file. So if this configuration is in repository configuration file, we have no need to do extract task like this	0	0	https://issues.jboss.org/browse/EXOJCR-439
4173	EXOJCR-430	Actually, I would like to simplify that and avoid to have many files to maintain. So we need to have a discussion on how to make that easy at the configuration level. What I mean by that is the capability to have JCR "generate" the configuration for JBoss Cache programmatically because all the files are the same. The idea would be to have in repository-configuration.xml a setting where we would have <cluster>true</cluster> and it would configure JBoss Cache in the right fashion.	1	1	https://issues.jboss.org/browse/EXOJCR-430

#	Requirement	Questions	Answered	Successfully Answered	URL
4174	EXOJCR-422	Hi Alex, I have some remarks about your work: 1. in PDFDocumentReader.java, in your error message, you wrote "documet" instead of "document". 2. Your solution is not generic enough, it makes no sense to have a PDFDocumentReadException in a generic class such as NodeIndexer, so could you please for example rename PDFDocumentReadException in DocumentReadException to make it more generic	0	0	https://issues.jboss.org/browse/EXOJCR-422
4175	EXOJCR-420	What will be a criteria for a decision for "what is a best"? From my side it can be: - Better performance of API operations (current daily and load tests), but this tests don't cover usage described in the issue, so it's only for "ping" testing - Special tests written assuming keeping of JCR items on client side, memory consumption is a point there. We have to write such tests, so create a subtask for it.	1	1	https://issues.jboss.org/browse/EXOJCR-420
4176	EXOJCR-420	In which way to measure and in which tests (new?): 1. memory usage 2. cache misses 3. performances ? Currently we can use performance only from daily/weekly tests in TPSes.	1	1	https://issues.jboss.org/browse/EXOJCR-420
4177	EXOJCR-420	Hi Anatoliy, Could you please also test with the Weak Value HashMap from the google collections http://code.google.com/p/google-collections (maven repo: http://repo1.maven.org/maven2/com/google/collections/google-collections, blog entry: http://pietrowski.info/2010/01/guava-google-collections/). See the MapMaker http://google-collections.googlecode.com/svn/trunk/javadoc/com/google/common/collect/MapMaker.html.	1	1	https://issues.jboss.org/browse/EXOJCR-420
4178	EXOJCR-420	Nicolas, we have tested WeakValueHashMap and WeakHashMap with WeakReference value and the last shows better results of values GCing (ItemImpl instances released quickly here). What a reason of experiment with Google Collections here?	0	0	https://issues.jboss.org/browse/EXOJCR-420
4179	EXOJCR-392	Is-this still up to date?	0	0	https://issues.jboss.org/browse/EXOJCR-392
4180	EXOJCR-392	In addition to previous version of bug fix, there is another way. We can change ChangesLog of orderBefore. If this ChangesLog will contain ADDED and REMOVED states, for all affected nodes (reordered nodes and child nodes), cache will update all nodes correctly. What is bad - if reordered node will contain a big childs tree, changesLog will be huge. This method has slower performance than previous - because full node tree will be removed and added in cache, instead of list of child nodes will be fetched only if cache do not contains such one. What is good - this is very simple way to fix this bug.	0	0	https://issues.jboss.org/browse/EXOJCR-392
4181	EXOJCR-392	JDBCStorageConnection it's only an impl of WorkspaceStorageConnection interface, we have several implementations of it. Do you propose to change our internal API (WorkspaceStorageConnection) once again? Another point: the cache knows nothing about persistence, contract of cache is clear now. But with this change we'll mix the cache and persistence responsibilities. And we still will traverse a whole subtree in the cache. With such success we can just traverse a whole subtree and remove Items with descendant path of reordered node. It's the approach we have used in old impl cache. And it will be more efficient comparing to use of a database to fix cache content. So let select from two: - traverse cache and remove descendant items - fill changes log with all affected items WDYT?	0	0	https://issues.jboss.org/browse/EXOJCR-392
4182	EXOJCR-392	BTW how did we fix this issue with the old JCR Cache?	1	1	https://issues.jboss.org/browse/EXOJCR-392
4183	EXOJCR-392	Removing all nodes tree on CHILD_NODES_LIST eviction - do not guarantee cache consistence. Lets see the usecase. we have such nodes /node1[3]/child1/child2/child3 Cache do not have full nodes tree (those nodes was added many time ago, and info about those nodes already removed by eviction). Than: - node1[3] and child3 nodes are added to cache again; - node1[3] is reordered; There is no chains between node1[3] and child3, so updateTreePath can not find and update child3. So, there is two possible decision: - clean cache on sibling reordering; - check all items in cache - is it descendant of reordered node?. Second one is no so bad, as it seemed at first glance: - iteration is hard operation, but it will be executed only on samename siblings; - there is no additional operations on eviction; - we do not need to bother about links to child;	0	0	https://issues.jboss.org/browse/EXOJCR-392
4184	EXOJCR-391	Working under this task, I've faced such an impediment: Source: org.exoplatform.services.jcr.api.lock.TestLock Description: Node has child-node "bar" and property "bar". Later NodeIndexer (line 404) asks property from DM to get property's values. And it returns Node instead of property. Why it was working before? Before changing cache logic listChildNodes returned properties with values and everything was great. Status: Solved : use getItemData(String uuid) instead of getItemData(String parent, QName name).	0	0	https://issues.jboss.org/browse/EXOJCR-391
4185	EXOJCR-388	Do we have to synchronize all dependencies versions ?	0	0	https://issues.jboss.org/browse/EXOJCR-388

#	Requirement	Questions	Answered	Successfully Answered	URL
4186	EXOJCR-324	sun.nio.ch.FileChannelImpl.map(FileChannelImpl.java:731) at sun.nio.ch.FileChannelImpl.transferFromFileChannel(FileChannelImpl.java:540) at sun.nio.ch.FileChannelImpl.transferFrom(FileChannelImpl.java:603) at org.exoplatform.services.jcr.impl.storage.value.fs.operations.ValueFileIOHelper.copy(ValueFileIOHelper.java:263) at org.exoplatform.services.jcr.impl.storage.value.fs.operations.ValueFileIOHelper.copyClose(ValueFileIOHelper.java:321) at org.exoplatform.services.jcr.impl.storage.value.fs.operations.ValueFileIOHelper.writeValue(ValueFileIOHelper.java:168) at org.exoplatform.services.jcr.impl.storage.value.fs.operations.WriteValue.commit(WriteValue.java:104) at org.exoplatform.services.jcr.impl.storage.value.fs.FileIOChannel.commit(FileIOChannel.java:129) at org.exoplatform.services.jcr.impl.storage.jdbc.JDBCStorageConnection.commit(JDBCStorageConnection.java:305) at org.exoplatform.services.jcr.impl.dataflow.persistent.WorkspacePersistentDataManager$ChangesLogPersister.commit(WorkspacePersistentDataManager.java:278) at org.exoplatform.services.jcr.impl.dataflow.persistent.WorkspacePersistentDataManager.save(WorkspacePersistentDataManager.java:251) at org.exoplatform.services.jcr.impl.dataflow.persistent.ACLInheritanceSupportedWorkspaceDataManager.save(ACLInheritanceSupportedWorkspaceDataManager.java:214) at org.exoplatform.services.jcr.impl.dataflow.persistent.VersionableWorkspaceDataManager.save(VersionableWorkspaceDataManager.java:235) at org.exoplatform.services.jcr.impl.dataflow.session.TransactionableDataManager.save(TransactionableDataManager.java:298) at org.exoplatform.services.jcr.impl.core.SessionDataManager.commit(SessionDataManager.java:1133) at org.exoplatform.services.jcr.impl.core.ItemImpl.save(ItemImpl.java:649) at org.exoplatform.services.jcr.impl.storage.fs.TestRandomValueIO.testUpdate_BigToBigValue(TestRandomValueIO.java:434) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at junit.framework.TestCase.runTest(TestCase.java:164) at junit.framework.TestCase.runBare(TestCase.java:130) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:120) at junit.framework.TestSuite.runTest(TestSuite.java:230) at junit.framework.TestSuite.run(TestSuite.java:225) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.apache.maven.surefire.junit.JUnitTestSet.execute(JUnitTestSet.java:213) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:140) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:127) at org.apache.maven.surefire.Surefire.run(Surefire.java:177)	0	0	https://issues.jboss.org/browse/EXOJCR-324
4187	EXOJCR-302	- Properties may hav not value, so new getChildPropertiesData query must support this usecase. This query is incorrect because property may have not value GET_PROPERTIES_BY_PARENTID = "select I.*, V.ORDER_NUM, V.DATA, V.STORAGE_DESC from JCR_MITEM I, JCR_MVALUE V" + " where I.I_CLASS=2 and I.PARENT_ID=? and V.PROPERTY_ID=I.ID order by I.ID"; This works as expected: GET_PROPERTIES_BY_PARENTID = "select I.*, V.ORDER_NUM, V.DATA, V.STORAGE_DESC from JCR_MITEM I LEFT OUTER JOIN JCR_MVALUE V ON (V.PROPERTY_ID=I.ID)" + " where I.I_CLASS=2 and I.PARENT_ID=? order by ID, ORDER_NUM"; - Reference property may be multivalued and values may contain same uuid. So getReferencesData must return same property many times for this uuid. With new query it's impossible (it returns single property, but values count is multiplicated); FIND_REFERENCE_PROPERTIES = "select P.ID, P.PARENT_ID, P.VERSION, P.P_TYPE, P.P_MULTIVALUED, P.NAME, V.ORDER_NUM, V.DATA, V.STORAGE_DESC" + " from JCR_MREF R, JCR_MITEM P, JCR_MVALUE V" + " where R.NODE_ID=? and P.ID=R.PROPERTY_ID and P.I_CLASS=2 and V.PROPERTY_ID=P.ID order by P.ID, V.ORDER_NUM";	0	0	https://issues.jboss.org/browse/EXOJCR-302
4188	EXOJCR-286	Webdav usecases: AddNode - uses webdav PUT method GetNode - uses webdav GET method RemoveNode - uses webdav DELETE method GetNodeByUUID - can't be test via webdav, because this operation is never directly used in webdav CreateProperty - uses webdav PROPPATCH nethod SetProperty - uses webdav PROPPATCH nethod on previosly created property. RemoveProperty - uses webdav PROPPATCH nethod on previosly created property. GetPropertyValue - has the same semantiv as GetNode usecase (webdav GET method reads and returns the value of "jcr:data" property)	0	0	https://issues.jboss.org/browse/EXOJCR-286
4189	EXOJCR-267	Here are my remarks : * The query is incorrect because of setMaxRows that cannot work here mostly due to the fact that we don't have one row per node since we get the properties primaryType, mixinTypes, owner and permissions at the same time for each node. So if your limit is 10, and you have one node with 7 properties and another node with 5, you will get 2 nodes instead of 10 and the second node won't have its 2 last properties which causes an inconsistency issue. * The value of fromOrderNum must be determinist since we use it as key to store the page content in the cache. The value should be of type 0, 1, 2, 3 (aka the page index) and so on because if you make it depend on the result set if the result set change (due to a delete, move, add...) we could face unpredictable bugs. *So I would even say that we need to replace fromOrderNum with pageIndex.* * In WorkspaceStorageCache, we need to add a method of type isChildNodesByPageSupported to be able to call getChildNodesData(NodeData) and return false directly when we call CacheableWorkspaceDataManager.getChildNodesDataByPage otherwise for implementation like LinkedWorkspaceStorageCacheImpl that doesn't support this feature the result will never be stored into the cache so it will be very slow. * In SessionDataManager.getChildNodesDataByPage and TransactionableDataManager.getChildNodesDataByPage, I don't really understand the merge logic: Why the code is different in those two methods ? What about the state isUnchaged() ? Why state.isRenamed() is treated exactly like isAdded() ? In case of isMixinChanged and isUpdated why simply checking if the id exists in the map instead of using the OrderNumber? I mean instead of: {code} childs.remove(state.getData()); if (minOrderNum <= data.getOrderNumber() && data.getOrderNumber() <= maxOrderNum) { childs.add(data); } {code} why not simply {code} if (childs.remove(state.getData())) { childs.add(data); } {code}	1	1	https://issues.jboss.org/browse/EXOJCR-267
4190	EXOJCR-267	Hi Nikolay, here are my remarks about your last commit: * 75 as min page size is quite big, could you please decrease it to 20? * Looking at the code, it looks like if the page set is < min, we will get a warning at each session creation which is too annoying, we should have this warning only once at workspace startup.	1	1	https://issues.jboss.org/browse/EXOJCR-267
4191	EXOJCR-265	To make the issue more clearer please provide next answers: How JCR will be used via JSON in external services for integration (example of ESB, other)? What are disadvantages of integration via XML (WebDAV) making us to use JSON?	0	0	https://issues.jboss.org/browse/EXOJCR-265

#	Requirement	Questions	Answered	Successfully Answered	URL
4192	EXOJCR-265	How JCR will be used via JSON in external services for integration (example of ESB, other)? It cans be anything a google gadget for example or a Flex application What are disadvantages of integration via XML (WebDAV) making us to use JSON? JSON is much lighter than XML due to the XML tags, so most of the architects now, prefer to use JSON instead of XML	0	0	https://issues.jboss.org/browse/EXOJCR-265
4193	EXOJCR-258	Doing lazy loading via several transactions can made such feature implementation even not possible. E.g. 1. we have 10 child nodes, read 5 on first read, other transaction deletes 6 of them. Will it be a okay to tell - no one node exists anymore? I.e. we tell that this Node has 5 childs (it's not a truth). 2. or we have 10 child nodes, read 5 on first read, other transaction deletes the parent node, we should throw RepositoryException (or InvalidItemState) exception but we already returned 5 of the deleted (i.e. API user used them assuming consistency). So, I'm back to lazy fetching from the ResultSet obtained on initial transaction (on method executed). The ResultSet will be from one JDBC transaction, it will be wrapped into Iterator implementation and fetched on hasNext()/next() calls. In this case the database will be responsible for consistency (and it's works well for many RDBMS, except of MySQL/MyISAM). And user will fetch all he/she could see on the Node.getNodes(). WDYT?	1	1	https://issues.jboss.org/browse/EXOJCR-258
4194	EXOJCR-198	Unfortunately, this change causes many stack-traces during tests (e.g. in Core). Running org.exoplatform.services.security.TestRolesExtractor Error of default config init: java.lang.Exception: unsupport war uri in this configuration service at org.exoplatform.container.configuration.ConfigurationManagerImpl.getURL(ConfigurationManagerImpl.java:297) at org.exoplatform.container.configuration.ConfigurationManagerImpl.getURL(ConfigurationManagerImpl.java:269) at org.exoplatform.container.configuration.ConfigurationManagerImpl.addConfiguration(ConfigurationManagerImpl.java:112) at org.exoplatform.container.StandaloneContainer.initDefaultConf(StandaloneContainer.java:329) at org.exoplatform.container.StandaloneContainer.getInstance(StandaloneContainer.java:131) at org.exoplatform.container.StandaloneContainer.getInstance(StandaloneContainer.java:87) at org.exoplatform.services.security.TestRolesExtractor.setUp(TestRolesExtractor.java:56) at junit.framework.TestCase.runBare(TestCase.java:128) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:120) at junit.framework.TestSuite.runTest(TestSuite.java:230) at junit.framework.TestSuite.run(TestSuite.java:225) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.apache.maven.surefire.junit.JUnitTestSet.execute(JUnitTestSet.java:213) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:140) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:165) at org.apache.maven.surefire.Surefire.run(Surefire.java:107) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.apache.maven.surefire.booter.SurefireBooter.runSuitesInProcess(SurefireBooter.java:289) at org.apache.maven.surefire.booter.SurefireBooter.main(SurefireBooter.java:1005) StandaloneContainer initialized using: file:/D:/Projects/eXo/dev/projects/projects-jboss/exo-jcr/core/trunk/component/security/core/target/test-classes/conf/st andalone/test-configuration.xml	0	0	https://issues.jboss.org/browse/EXOJCR-198
4195	EXOJCR-172	In exo.core.packaging.module there is a file https://anonsvn.jboss.org/repos/exo-jcr/core/trunk/packaging/module/dependencies.txt which contains a full dependencies tree of exo.core.packaging.module. Is this file needed for some functionality or it can be safetly removed ?	0	0	https://issues.jboss.org/browse/EXOJCR-172
4196	EXOJCR-164	Did you try to use -Prun-tck,run-its in the build configuration ?	0	0	https://issues.jboss.org/browse/EXOJCR-164
4197	EXOJCR-65	It's not a RegistryService bug. It's QName issue and it's applicable for ALL JCR internal names compares. QName.equals() based on hashcodes of namespace and name of both QName instances. Yes, hascode doesn't guaranty full 100% equality for varies of Strings used in namespace and name. But we do that for performance staring from JCR 1.5 (JCR-15) and it works ok for 99.9% of human-readable Items names. yes, hashcodes for strings bBbB and aaaa are same, but is it a real usecase? Tung, is it a customer issue? Or where eXo uses similar names? I prefer to have valuable motivation of the issue. Anatoty, investigate how we can improve the QName equality procedure optimal to a performance.	0	0	https://issues.jboss.org/browse/EXOJCR-65
4198	FABRIC-1213	http://fabric8.io/v2/forge.html Under Forge Addons "figure out whats available" should use "what's".	1	1	https://issues.jboss.org/browse/FABRIC-1213
4199	FABRIC-1213	http://fabric8.io/v2/faqGeneral.html "What Java versions is supported?" should be "versions are supported".	1	1	https://issues.jboss.org/browse/FABRIC-1213
4200	FABRIC-1180	I think that would work. I was thinking more along the lines of separation of concerns within the fabric. So that a plain old user on a plain old container has the ability to do whatever basic operations needed for his container; but more sensitive operations (maybe external container interactions) or sensitive znode paths in the container would require an admin-type auth; and fabric-modifying operations would require a fabric-admin-type auth. Is something like that feasible?	1	0	https://issues.jboss.org/browse/FABRIC-1180
4201	FABRIC-1180	The point about not sharing passwords with non-trusted entities is part of the larger issue of the overall security of the system, which we are strengthening with this issue. This added security could be what keeps an attacker from getting control of an entire fabric instead of only one node, or just limited access to one node. And if we're looking at expanding the types of registry components we're using in the long term, then now is the best time to make sure we have a more secure design to model as we grow. Back to my original question, is the separation of zk roles/concerns a feasible way of moving forward?	0	0	https://issues.jboss.org/browse/FABRIC-1180
4202	FABRIC-1180	Hiram, you're right-- it's not required per se, but addressing this issue was a condition of Product Security's approval of the Planning Phase Exit. And while I understand that implementing a hardening feature for what seems like a longshot as a security risk is not anyone's first priority, we want to manage such issues sooner, before they become a very high priority for urgent reasons. If the role-separation approach is feasible, could we move forward with that for 6.2.0? Thanks!	0	0	https://issues.jboss.org/browse/FABRIC-1180
4203	FABRIC-1150	Could you please file issues with SCM sync project (and link them here) that describe what is not working. I'd like this to work out of the box.	1	1	https://issues.jboss.org/browse/FABRIC-1150
4204	FABRIC-1084	Does turning on this option, affect every optional import specified in a feature or multiple features? Is it global?	0	0	https://issues.jboss.org/browse/FABRIC-1084
4205	FABRIC-1084	Is fabric version 7.4.0.redhat-62 the version that went out with Red Hat JBoss Fuse version 6.1 GA release?	0	0	https://issues.jboss.org/browse/FABRIC-1084

#	Requirement	Questions	Answered	Successfully Answered	URL
4206	FABRIC-1077	[~chirino] Can this be marked as resolved or is there still work outstanding?	0	0	https://issues.jboss.org/browse/FABRIC-1077
4207	FABRIC-1063	Are the fix/affect versions for this Jira 7.4.0.redhat-62?	0	0	https://issues.jboss.org/browse/FABRIC-1063
4208	FABRIC-1057	All problems with building/compiling Fabric8 under JDK8 relates to ... Scala. Do we really need to depend in such (still) unstable and not backward-compatible language?	0	0	https://issues.jboss.org/browse/FABRIC-1057
4209	FABRIC-1056	This is fixed in Fuse 6.2 / Fabric8 1.2.	1	1	https://issues.jboss.org/browse/FABRIC-1056
4210	FABRIC-1054	And can you add a comment for release notes ? https://issues.jboss.org/browse/FUSEDOC-471	1	1	https://issues.jboss.org/browse/FABRIC-1054
4211	FABRIC-966	Is there anything that should be documented in the release notes for 6.1? From the description it looks like this issue is more of an FYI for when we document the APIs.	1	0	https://issues.jboss.org/browse/FABRIC-966
4212	FABRIC-896	I've added the following line to the cleanup job: find ${WORKSPACE}/.. -mindepth 1 -maxdepth 1 -type d -ctime +30 | xargs rm -rf Sorry for the delay, for some reason I did not get any email notifications on this bug.	0	0	https://issues.jboss.org/browse/FABRIC-896
4213	FABRIC-828	we've kinda got the basics of this here: https://vimeo.com/album/2635012/video/84674508 maybe we just edit the beginning bit out as a separate screen cast?	0	0	https://issues.jboss.org/browse/FABRIC-828
4214	FABRIC-812	__,_|___/___| JBoss Fuse (6.1.0.redhat-SNAPSHOT) http://www.redhat.com/products/jbossenterprisemiddleware/fuse/ Hit '<tab>' for a list of available commands and '[cmd] --help' for help on a specific command. Open a browser to http://localhost:8181 to access the management console Hit '<ctrl-d>' or 'osgi:shutdown' to shutdown JBoss Fuse. [main] DEBUG org.ops4j.pax.exam.container.remote.RBCRemoteTarget - Preparing and Installing bundle (from stream).. [main] DEBUG org.ops4j.pax.exam.rbc.client.RemoteBundleContextClient - Packing probe into memory for true RMI. Hopefully things will fill in.. JBossFuse:karaf@root> [main] DEBUG org.ops4j.pax.exam.container.remote.RBCRemoteTarget - Installed bundle (from stream) as ID: 150 [main] DEBUG org.ops4j.pax.exam.container.remote.RBCRemoteTarget - call [[TestAddress:PaxExam-5d6c1449-82c2-4d97-bd80-bfa4f88a1cd1 root:PaxExam-62db401f-2266-4a63-ad56-9f7d8ed24b48]] fabric:create -n Waiting for containers: [esb1] to successfully provision Container:esb1 Alive:false Status: SSH URL:null Container:esb1 Alive:false Status: SSH URL:null Container:esb1 Alive:false Status: SSH URL:null Container:esb1 Alive:false Status: SSH URL:null Container:esb1 Alive:true Status: SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:false Status:analyzing SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:resolving SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:downloading SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:resolving SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:false Status:finalizing SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:finalizing SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:finalizing SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:success SSH URL:ip-10-180-2-205:8103 Testing profile:jboss-fuse-minimal on container:[esb1] by adding feature:connector Waiting for containers: [esb1] to successfully provision fabric:profile-edit --features connector jboss-fuse-minimal Adding feature:connector to profile:jboss-fuse-minimal version:1.0 Waiting for containers: [esb1] to successfully provision Container:esb1 Alive:true Status:finalizing SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:finalizing SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:finalizing SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:finalizing SSH URL:ip-10-180-2-205:8103 Container:esb1 Alive:true Status:success SSH URL:ip-10-180-2-205:8103 fabric:profile-display jboss-fuse-minimal Profile id: jboss-fuse-minimal Version : 1.0 Attributes: parents: karaf feature-camel Containers: root esb1 Container settings ---------- Repositories : mvn:org.apache.servicemix.nmr/apache-servicemix-nmr/${version:servicemix-nmr}/xml/features mvn:org.jboss.fuse/jboss-fuse/6.1.0.redhat-SNAPSHOT/xml/features Features : connector camel-core camel-jms camel-spring camel-jetty camel-blueprint jasypt-encryption fabric-bundle camel-jdbc camel-jaxb Configuration details ---------- Other resources ---------- fabric:container-list [id] [version] [alive] [profiles] [provision status] root* 1.0 true fabric, fabric-ensemble-0000-1, jboss-fuse-minimal, unmanaged esb1 1.0 true jboss-fuse-minimal success container-connect -u admin -p admin esb1 osgi:list -s -t 0 Connected &#27;[31mCommand not found: &#27;[1mosgi:list&#27;[22;39m container-connect -u admin -p admin esb1 osgi:list -s -t 0 | grep geronimo-connector [main] DEBUG org.apache.karaf.tooling.exam.container.internal.KarafTestContainer - Shutting down the test container (Pax Runner) [main] ERROR org.ops4j.pax.exam.junit.JUnit4TestRunner - Exception org.ops4j.pax.exam.TestContainerException: [testFeatures(org.fusesource.esb.itests.basic.fabric.EsbProfileLongTest): Expected to find symbolic name:geronimo-connector] at	1	1	https://issues.jboss.org/browse/FABRIC-812
4215	FABRIC-812	Ideally, we would find and remove all places where the TCL is being set. Stuff that is supposed to run in an OSGi environment, should be able to do so without TCL. Any particular TCL value is likely going to be wrong. Does your fix work when you just unset the TCL?	0	0	https://issues.jboss.org/browse/FABRIC-812
4216	FABRIC-807	Do you mean freegeoip.net? It seems like the screenshot contains a typo (it says freegoip.net).	1	0	https://issues.jboss.org/browse/FABRIC-807
4217	FABRIC-796	The test fix is merged to master and 6.1.x branches	0	0	https://issues.jboss.org/browse/FABRIC-796
4218	FABRIC-739	An clarification / addition: A profile may also contain any kind of resources, e.g: i) a spring or blueprint descriptor (broker configuration, camel routes etc) ii) any kind of text resource an application may use (properties, xml, json yml etc). Those resources are displayed in profile-list in the section called "Resources". An additional option (--dispaly-resources) is available in 6.1 to display the content of those resources (the default behaviour is to just display their names). All those resources can be used by the app using the profile url handler.	1	1	https://issues.jboss.org/browse/FABRIC-739
4219	FABRIC-739	In the modifying profile configuration pid, there are the following cases we need to describe: i) Edit the pid in the editor (fabric:profile-edit --pid myPid myprofile will open myPid inside the cli text editor we have) ii) Edit the pid in line. When editing a pid in line there are the following cases we need to cover: i) value assignment to a specific key (is already covered). ii) appending a value to a delimited list (e.g. we want to add a url to a comma separated list of urls: fabric:profile-edit --append --pid myPid/myKey=newUrl myprofile) iii) removing a value to delemited list (e.g. we want to remove a url to a comma separated list of urls: fabric:profile-edit --remove --pid --pid myPid/myKey=newUrl myprofile) iv) removing the key completely (e.g. using the delete option) A perfect example of a case where we need to append/remove values to a delimited list is when adding maven repositories to a profile by appending/removing repository urls from the org.ops4j.pax.url.mvn.repositories key (usually directly in the default profile).	1	1	https://issues.jboss.org/browse/FABRIC-739
4220	FABRIC-730	Grzegorz, could you take a look? The biggest issue would be getting this to run in a reasonable amount of time... probably would have to boot up a container once and check all features against that. So a new test could go here: https://github.com/fabric8io/fabric8/tree/master/esb/esb-itests/extra/src/test/java/org/fusesource/esb/itests/extra and in that test you can use the features service to list all possible features to try: https://github.com/apache/karaf/blob/karaf-2.3.x/features/core/src/main/java/org/apache/karaf/features/FeaturesService.java#L72	0	0	https://issues.jboss.org/browse/FABRIC-730

#	Requirement	Questions	Answered	Successfully Answered	URL
4221	FABRIC-730	Also, when I've uninstalled *everything* (besides {{scr}} and {{eventadmin}}), I got: {noformat} JBossFuse:karaf@root> features:list -i State Version Name Repository Description [installed] [2.3.0.redhat-610337] eventadmin karaf-2.3.0.redhat-610337 OSGi Event Admin service specification for event-based communication [installed] [2.3.0.redhat-610337] scr karaf-2.3.0.redhat-610337 Declarative Service support JBossFuse:karaf@root> features:install process-fabric Error executing command: Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=io.fabric8.zookeeper.utils)(version>=1.0.0)(!(version>=2.0.0))) Fuse Process :: Fabric package:(&(package=io.fabric8.utils)) Fuse Process :: Fabric package:(&(package=io.fabric8.api.jcip)(version>=1.0.0)(!(version>=2.0.0))) Fuse Process :: Fabric package:(&(package=io.fabric8.internal)(version>=1.0.0)(!(version>=2.0.0))) Fuse Process :: Fabric package:(&(package=org.apache.curator)(version>=2.3.0)(!(version>=3.0.0))) Fuse Process :: Fabric package:(&(package=io.fabric8.zookeeper)(version>=1.0.0)(!(version>=2.0.0))) Fuse Process :: Fabric package:(&(package=org.apache.curator.framework)(version>=2.3.0)(!(version>=3.0.0))) Fuse Process :: Fabric package:(&(package=io.fabric8.service)(version>=1.0.0)(!(version>=2.0.0))) Fuse Process :: Fabric package:(&(package=org.mvel2.templates)(version>=2.1.0)(!(version>=3.0.0))) Fuse Process :: Fabric package:(&(package=org.mvel2)(version>=2.1.0)(!(version>=3.0.0))) Fuse Process :: Fabric package:(&(package=io.fabric8.api.scr)(version>=1.0.0)(!(version>=2.0.0))) Fuse Process :: Fabric package:(&(package=io.fabric8.api)(version>=1.0.0)(!(version>=2.0.0))) Fuse Process :: Fabric service:(service=io.fabric8.api.FabricService) Fuse Process :: Fabric package:(&(package=io.fabric8.api.jmx)(version>=1.0.0)(!(version>=2.0.0))) Fuse Process :: Fabric {noformat} Shouldn't features be declared with all dependencies (also the ones installed by default)?	1	1	https://issues.jboss.org/browse/FABRIC-730
4222	FABRIC-730	---------- package:(&(package=io.fabric8.utils)) Fuse Process :: Fabric package:(&(package=org.mvel2.templates)(version>=2.1.0)(!(version>=3.0.0))) Fuse Process :: Fabric package:(&(package=org.mvel2)(version>=2.1.0)(!(version>=3.0.0))) Fuse Process :: Fabric ======= feature: jclouds-rackspace-cloudblockstorage-us/1.6.2-incubating Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.jclouds.rackspace.cloudidentity.v2_0.config)(version>=1.6.2.incubating)) jclouds Rackspace Next Generation Cloud Block Storage US provider ======= feature: jclouds-rackspace-cloudblockstorage-uk/1.6.2-incubating Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.jclouds.rackspace.cloudidentity.v2_0.config)(version>=1.6.2.incubating)) jclouds Rackspace Next Generation Cloud Block Storage UK provider ======= feature: camel-disruptor/2.12.0.redhat-610337 Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=com.lmax.disruptor.dsl)(version>=3.2.0)(!(version>=4.0.0))) camel-disruptor package:(&(package=sun.misc)) Disruptor Framework package:(&(package=com.lmax.disruptor)(version>=3.2.0)(!(version>=4.0.0))) camel-disruptor ======= feature: mq-fabric/6.1.0.redhat-SNAPSHOT Could not start bundle mvn:org.jboss.amq/mq-fabric/6.1.0.redhat-SNAPSHOT in feature(s) mq-fabric-6.1.0.redhat-SNAPSHOT: Uses constraint violation. Unable to resolve bundle revision org.jboss.amq.mq-fabric [633.0] because it is exposed to package 'org.springframework.beans.factory' from bundle revisions org.springframework.beans [392.0] and org.springframework.beans [109.0] via two dependency chains. Chain 1: org.jboss.amq.mq-fabric [633.0] import: (&(osgi.wiring.package=org.springframework.beans.factory)(version>=3.1.0)(!(version>=4.0.0))) | export: osgi.wiring.package=org.springframework.beans.factory org.springframework.beans [392.0] Chain 2: org.jboss.amq.mq-fabric [633.0] import: (osgi.wiring.package=org.apache.activemq.command) | export: osgi.wiring.package=org.apache.activemq.command; uses:=org.apache.xbean.spring.context.impl org.apache.activemq.activemq-osgi [279.0] import: (&(osgi.wiring.package=org.apache.xbean.spring.context.impl)(version>=3.15.0)(!(version>=4.0.0))) | export: osgi.wiring.package=org.apache.xbean.spring.context.impl; uses:=org.springframework.beans.factory.support org.apache.xbean.spring [281.0] import: (&(osgi.wiring.package=org.springframework.beans.factory.support)(version>=2.5.0)(!(version>=4.0.0))) | export: osgi.wiring.package=org.springframework.beans.factory.support; uses:=org.springframework.beans.factory export: osgi.wiring.package=org.springframework.beans.factory org.springframework.beans [109.0] ======= feature: jclouds-rackspace-clouddns-us/1.6.2-incubating Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.jclouds.rackspace.clouddns.v1)(version>=1.6.2.incubating)) jclouds Rackspace Next Generation Cloud DNS US provider ======= feature: jclouds-rackspace-clouddns-uk/1.6.2-incubating Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.jclouds.rackspace.clouddns.v1)(version>=1.6.2.incubating)) jclouds Rackspace Next Generation Cloud DNS UK provider ======= feature: jclouds-rackspace-cloudservers-us/1.6.2-incubating Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.jclouds.rackspace.cloudidentity.v2_0.config)(version>=1.6.2.incubating)) jclouds Rackspace Next Generation Cloud Servers US provider ======= feature: jclouds-rackspace-cloudservers-uk/1.6.2-incubating Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.jclouds.rackspace.cloudidentity.v2_0.config)(version>=1.6.2.incubating)) jclouds Rackspace Next Generation Cloud Servers provider	1	1	https://issues.jboss.org/browse/FABRIC-730
4223	FABRIC-730	Hi [~grgrzybek], you found some new problems and some are already reported :) feature: jclouds-rackspace-cloudloadbalancers-us/1.8.0 _I saw problems with installing jclouds-rackspace-cloudbalancers-* in tests, but when i tried manual install with features:install ..., then it worked. Don't know where is problem._ camel-jcr tracked here: ENTESB-2763 camel-sip tracked here: ENTESB-2779 camel-rx: will most likely work only on Oracle Java, but not documented yet: ENTESB-2778 I report theese issues as subtasks to ENTESB-2663.	0	0	https://issues.jboss.org/browse/FABRIC-730

#	Requirement	Questions	Answered	Successfully Answered	URL
4224	FABRIC-712	]ERROR: OSGi API called while holding java locks: "FelixStartLevel" Id=13 in RUNNABLE at sun.management.ThreadImpl.dumpThreads0(Native Method) at sun.management.ThreadImpl.getThreadInfo(ThreadImpl.java:440) at org.apache.felix.framework.Felix.checkJavaLocks(Felix.java:2231) at org.apache.felix.framework.Felix.registerService(Felix.java:3619) at org.apache.felix.framework.BundleContextImpl.registerService(BundleContextImpl.java:346) at org.apache.felix.framework.BundleContextImpl.registerService(BundleContextImpl.java:320) at org.apache.aries.util.tracker.hook.BundleHookBundleTracker.open(BundleHookBundleTracker.java:141) - locked org.apache.aries.util.tracker.hook.BundleHookBundleTracker$Tracked@208b4e03 - locked org.apache.aries.util.tracker.hook.BundleHookBundleTracker@19aeb7a6 at org.apache.aries.util.tracker.RecursiveBundleTracker.open(RecursiveBundleTracker.java:92) at org.apache.aries.blueprint.container.BlueprintExtender$2.serviceFound(BlueprintExtender.java:114) at org.apache.aries.util.tracker.SingleServiceTracker.update(SingleServiceTracker.java:174) at org.apache.aries.util.tracker.SingleServiceTracker.findMatchingReference(SingleServiceTracker.java:132) at org.apache.aries.util.tracker.SingleServiceTracker.access$200(SingleServiceTracker.java:32) at org.apache.aries.util.tracker.SingleServiceTracker$1.serviceChanged(SingleServiceTracker.java:61) at org.apache.felix.framework.util.EventDispatcher.invokeServiceListenerCallback(EventDispatcher.java:934) at org.apache.felix.framework.util.EventDispatcher.fireEventImmediately(EventDispatcher.java:795) at org.apache.felix.framework.util.EventDispatcher.fireServiceEvent(EventDispatcher.java:544) at org.apache.felix.framework.Felix.fireServiceEvent(Felix.java:4675) at org.apache.felix.framework.Felix.registerService(Felix.java:3683) at org.apache.felix.framework.BundleContextImpl.registerService(BundleContextImpl.java:346) at org.apache.felix.framework.BundleContextImpl.registerService(BundleContextImpl.java:320) at org.apache.aries.proxy.impl.ProxyManagerActivator.start(ProxyManagerActivator.java:73) at org.apache.felix.framework.util.SecureAction.startActivator(SecureAction.java:645) at org.apache.felix.framework.Felix.doActivateBundle(Felix.java:2395) at org.apache.felix.framework.Felix$7.call(Felix.java:2333) at org.apache.felix.framework.Felix.runInContext(Felix.java:2188) at org.apache.felix.framework.Felix.activateBundle(Felix.java:2331) at org.apache.felix.framework.Felix.startBundle(Felix.java:2090) at org.apache.felix.framework.Felix.setActiveStartLevel(Felix.java:1317) at org.apache.felix.framework.FrameworkStartLevelImpl.run(FrameworkStartLevelImpl.java:304) at java.lang.Thread.run(Thread.java:722) ERROR: OSGi API called while holding java locks: "FelixStartLevel" Id=13 in RUNNABLE at sun.management.ThreadImpl.dumpThreads0(Native Method) at sun.management.ThreadImpl.getThreadInfo(ThreadImpl.java:440) at org.apache.felix.framework.Felix.checkJavaLocks(Felix.java:2231) at org.apache.felix.framework.Felix.registerService(Felix.java:3619) at org.apache.felix.framework.BundleContextImpl.registerService(BundleContextImpl.java:346) at org.apache.aries.blueprint.container.BlueprintContainerImpl.registerService(BlueprintContainerImpl.java:448) at	1	1	https://issues.jboss.org/browse/FABRIC-712
4225	FABRIC-712	I noticed quite a few related commits. Is this complete?	1	1	https://issues.jboss.org/browse/FABRIC-712
4226	FABRIC-704	Thomas can you give me admin access to Jenkins so I can work on this?	1	1	https://issues.jboss.org/browse/FABRIC-704
4227	FABRIC-690	Jon, do you know how to do this?	1	1	https://issues.jboss.org/browse/FABRIC-690
4228	FABRIC-690	We use retry failed builds plugin in hudson for the perfectus 6.1 job. See Schedule failed builds in http://ci.fusesource.com/hudson/job/perfectus-jboss-fuse-6.1.x/configure. Is there a particular job we want this set up for?	0	0	https://issues.jboss.org/browse/FABRIC-690
4229	FABRIC-683	Trying to spread the load. Can you please have a look?	0	0	https://issues.jboss.org/browse/FABRIC-683
4230	FABRIC-681	Trying to spread the load. Can you please have a look?	0	0	https://issues.jboss.org/browse/FABRIC-681
4231	FABRIC-674	Trying to spread the load. Can you please have a look?	0	0	https://issues.jboss.org/browse/FABRIC-674

#	Requirement	Questions	Answered	Successfully Answered	URL
4232	FABRIC-674	org/fusesource/mq/fabric/FabricDiscoveryAgent] at org.ops4j.pax.exam.invoker.junit.internal.JUnitProbeInvoker.invokeViaJUnit(JUnitProbeInvoker.java:128) at org.ops4j.pax.exam.invoker.junit.internal.JUnitProbeInvoker.findAndInvoke(JUnitProbeInvoker.java:94) at org.ops4j.pax.exam.invoker.junit.internal.JUnitProbeInvoker.call(JUnitProbeInvoker.java:77) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.ops4j.pax.exam.rbc.internal.RemoteBundleContextImpl.remoteCall(RemoteBundleContextImpl.java:86) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:322) at sun.rmi.transport.Transport$1.run(Transport.java:177) at sun.rmi.transport.Transport$1.run(Transport.java:174) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:173) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:556) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run0(TCPTransport.java:811) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:670) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at java.lang.Thread.run(Thread.java:744) Caused by: java.lang.NoClassDefFoundError: org/fusesource/mq/fabric/FabricDiscoveryAgent at io.fabric8.itests.basic.mq.MQProfileTest.waitForBroker(MQProfileTest.java:190) at io.fabric8.itests.basic.mq.MQProfileTest.testMQCreateBasic(MQProfileTest.java:101) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.ops4j.pax.exam.invoker.junit.internal.ContainerTestRunner.runChild(ContainerTestRunner.java:70) at org.ops4j.pax.exam.invoker.junit.internal.ContainerTestRunner.runChild(ContainerTestRunner.java:37) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at	0	0	https://issues.jboss.org/browse/FABRIC-674
4233	FABRIC-672	Trying to spread the load. Can you please have a look?	1	1	https://issues.jboss.org/browse/FABRIC-672
4234	FABRIC-670	Trying to spread the load. Can you please have a look?	0	0	https://issues.jboss.org/browse/FABRIC-670
4235	FABRIC-662	This issue was already fixed Unresolved constraint in bundle servicemix-eip Does it need to be merged to 6.1 branch?	1	1	https://issues.jboss.org/browse/FABRIC-662
4236	FABRIC-653	I found this {code} 405 ? Sl 20:59 /opt/jdk1.6.0_45/jre/bin/java -Xmx512M -Djava.awt.headless=true -jar /mnt/hudson/workspace/activemq-5.9.0.redhat-6-1-x-stable-dualjdk/activemq-unit-tests/target/surefire/surefirebooter4757849231766109607.jar /mnt/hudson/workspace/activemq-5.9.0.redhat-6-1-x-stable-dualjdk/activemq-unit-tests/target/surefire/surefire2173923929772988658tmp /mnt/hudson/workspace/activemq-5.9.0.redhat-6-1-x-stable-dualjdk/activemq-unit-tests/target/surefire/surefire817381951878609726tmp {code} It seems that there is a surefire process left over from a totally unrelated build job. Is it possible that the cleanup job interferes and perhaps deletes a tmp file that surefire needs in order to shut down correctly? If we cant find the root cause, we should include this condition in the kill-stale-processes.sh script	1	0	https://issues.jboss.org/browse/FABRIC-653
4237	FABRIC-646	The test should be split in 2: i) testJoin: Should be part of the smoke tests, as it test core fabric funcionality. ii) testJoinAndAddToEnsemble: Can be moved to an other test set, as it is actually a regression test.	0	0	https://issues.jboss.org/browse/FABRIC-646
4238	FABRIC-619	Thomas what is the status of this? What artifacts do you need published, and from what jobs?	0	0	https://issues.jboss.org/browse/FABRIC-619
4239	FABRIC-619	Still not available from any repo. Could you check it please? Remember that those artifacts are somehow no longer built in the nightly runs due to Thomas changes for the smoke test stuff. Those modules seem to be only built when using the release profile.	1	1	https://issues.jboss.org/browse/FABRIC-619
4240	FABRIC-619	Which job is publishing the artefacts?	1	1	https://issues.jboss.org/browse/FABRIC-619
4241	FABRIC-599	Would it be possible to force a reinstall of the bundle from Fabric whenever the bundle version changes? That would set the bundle location correctly but I am not sure what other drawbacks this approach would cause.	0	0	https://issues.jboss.org/browse/FABRIC-599
4242	FABRIC-517	Is this still used in IDE?	0	0	https://issues.jboss.org/browse/FABRIC-517
4243	FABRIC-508	Can we finally close this one? All the subtasks are resolved!	1	1	https://issues.jboss.org/browse/FABRIC-508
4244	FABRIC-502	After talking to Dave, we wanted to expand this feature a little bit to insure that clients who run in their own runtime environment can access resources in the fabric. For example, we have a rest service deployed in fabric across three machines. How can remote clients who can not user fabric-camel component connect? We believe the same mod_cluster/mod-proxy solution here will also cover this use case.	0	0	https://issues.jboss.org/browse/FABRIC-502
4245	FABRIC-409	Can you take a look at this one Willem?	1	1	https://issues.jboss.org/browse/FABRIC-409
4246	FABRIC-409	Hi Dejan, What's the error did you get? I tried to deploy the demo into jboss-fuse-6.0.0.redhat-015, I got the error {code} 2013-03-21 17:05:05,301 | ERROR | NAPSHOT-thread-1 | FabricLoadBalancerFeature | ? ? | 236 - org.fusesource.fabric.fabric-cxf - 7.2.0.redhat-015 | Cannot initialize the bus with FabricLoadBalancerFeature due to java.util.concurrent.TimeoutException: timeout reached while waiting on the lock: java.lang.Object@3582c145 {code} It looks like the FabricLoadBalancerFeature have some trouble to connect the zookeeper. I will keep digging it.	0	0	https://issues.jboss.org/browse/FABRIC-409
4247	FABRIC-391	Have u tried setting java.rmi.server.hostname system property to the same name as appears in the jmx url?	1	1	https://issues.jboss.org/browse/FABRIC-391
4248	FABRIC-391	That did it! This is kinda painful to do though when you have more than 5 separate machines I think. Is there any way we could maybe pass this on the command line automatically when starting containers? Ideally it should be set to whatever the resolver is I guess.	0	0	https://issues.jboss.org/browse/FABRIC-391

#	Requirement	Questions	Answered	Successfully Answered	URL
4249	FABRIC-382	Excellent work! Some things I would clarify: i) To open a pid in the editor you can also use --pid option and not just the --resource option (both are correct) ii) In the substitution section, its not just karaf.name, karaf.data, etc. It's all system properties, regardless of wether they are define in etc/system.properties or jvm options. iii) I like the part with the recursive substitution. I'd like to add an example, which pops up quite often from the cs and forums: "How do I refer to the path of resource, without being tied to the version: ${zk:/fabric/configs/versions/${zk:/fabric/configs/containers/${karaf.name}}/profiles/my-profile/my-resource}	0	0	https://issues.jboss.org/browse/FABRIC-382
4250	FABRIC-205	Can we close this old issue, as I think we have it covered in the current docs? It only makes sense to use an authentication system whose security data can be accessed throughout the fabric (e.g. also remotely). Currently, we have two JAAS modules that satisfy this requirement: the {{ZookeeperLoginModue}} (which is enabled by default) and the {{LDAPLoginModule}} (which is documented in the LDAP Authentication Tutorial in the Security Guide). So, I think the relevant cases are covered in the docs.	0	0	https://issues.jboss.org/browse/FABRIC-205
4251	FABRIC-183	There is the cluster-list command now. I guess we got some pieces for this already. So what is needed left? A simpler end-user API?	0	0	https://issues.jboss.org/browse/FABRIC-183
4252	FABRIC-150	What is the priority of this ticket? I can take a stab at it after finalising the camel-mongodb component, which will hopefully happen at the beginning of next week.	0	0	https://issues.jboss.org/browse/FABRIC-150
4253	FABRIC-107	The help should describe its optional and provide help on how to install it. I tried this and it failed in beta-051... {code} FuseESB:karaf@root> features:install fabric-jclouds Error executing command: Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.jclouds.sshj.config)(version>=1.3.0)) JClouds :: Karaf :: Services {code}	1	0	https://issues.jboss.org/browse/FABRIC-107
4254	FABRIC-48	mvn:org.fusesource.fabric/fabric-distro/1.1-SNAPSHOT/karaf/activemq but then could not connect to JMX from IDE. I suspect its missing the various authentication stuff? Caught: java.lang.SecurityException: Authentication failed java.lang.SecurityException: Authentication failed at org.apache.karaf.management.JaasAuthenticator.authenticate(JaasAuthenticator.java:101) at javax.management.remote.rmi.RMIServerImpl.doNewClient(RMIServerImpl.java:213) at javax.management.remote.rmi.RMIServerImpl.newClient(RMIServerImpl.java:180) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:305) at sun.rmi.transport.Transport$1.run(Transport.java:159) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:155) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:535) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run0(TCPTransport.java:790) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:649) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) at java.lang.Thread.run(Thread.java:680) at sun.rmi.transport.StreamRemoteCall.exceptionReceivedFromServer(StreamRemoteCall.java:255) at sun.rmi.transport.StreamRemoteCall.executeCall(StreamRemoteCall.java:233) at sun.rmi.server.UnicastRef.invoke(UnicastRef.java:142) at javax.management.remote.rmi.RMIServerImpl_Stub.newClient(Unknown Source) at javax.management.remote.rmi.RMIConnector.getConnection(RMIConnector.java:2329) at javax.management.remote.rmi.RMIConnector.connect(RMIConnector.java:279) at javax.management.remote.JMXConnectorFactory.connect(JMXConnectorFactory.java:248) at org.fusesource.fabric.service.AgentTemplate.createConnector(AgentTemplate.java:160) at org.fusesource.fabric.service.AgentCachingJmxTemplate.createConnector(AgentCachingJmxTemplate.java:39) at org.fusesource.fabric.service.JmxTemplate.getConnector(JmxTemplate.java:63) at org.fusesource.fabric.service.JmxTemplate.execute(JmxTemplate.java:33) at com.fusesource.ide.fabric.navigator.jmx.FabricConnectionWrapper.run(FabricConnectionWrapper.java:158) at com.fusesource.ide.jmx.core.tree.NodeUtils.createObjectNameTree(NodeUtils.java:56) at com.fusesource.ide.fabric.navigator.jmx.FabricConnectionWrapper.loadRoot(FabricConnectionWrapper.java:129) at com.fusesource.ide.fabric.navigator.jmx.FabricConnectionWrapper.loadChildren(FabricConnectionWrapper.java:145) at com.fusesource.ide.fabric.navigator.jmx.FabricConnectionWrapper.loadChildren(FabricConnectionWrapper.java:141) at	0	0	https://issues.jboss.org/browse/FABRIC-48
4255	FABRIC-45	Looks like zk:export and zk:import could allow this, right now you can do zk:export, copy paste the output to some file, make changes and then pass the URI of that file to zk:import, but they seem incomplete, for example if I pass a properties file to import I can't specify what path it should go in, so all the keys wind up at the root of the ZK config. Do you think though that if the import/export commands were more user friendly this would fit the bill? The problem with trying to do a live sync between the file system and ZK I can see is how to determine who has the definitive version of the config at a given time.	1	1	https://issues.jboss.org/browse/FABRIC-45
4256	FABRIC-45	So maybe we start with a recursive import and export; where we can import from a directory to a given path in the config - or export from a path in ZK to a directory? Once we've recursive import/exports, we should be able to build a 'if it changes in git/svn, update ZK' little heroku style example, where folks can hack files in source control then update ZK on a push. i.e. the git on commit handler would checkout git and do a recursive ZK import. Folks could even modify this so that only certain parts of the tree are imported from SCM etc. Extra bonus points - we could maybe have a .fabricignore file like a .gitignore where it lists the ZK paths to ignore when exporting maybe? On import we could maybe delete stuff not in ZK (unless the paths are in the .fabricignore) - so if you delete something in the file system, it'll be deleted in ZK. Then using the .fabricignore files; we can define 'holes' in the ZK tree where we don't import/export (or delete in ZK). The only worry with this I can see is if folks decide to setup a heroku style git commit hook to update ZK; if folks also setup a 'on ZK change export to file system' thingy; then automatically do a commit & push; we could end up with a ZK <-> git infinite loop. Maybe some code review may be required to move the 'runtime changes' branch into the 'heroku deploy branch' in git? I guess folks can figure out (outside of our fabric java code) how best they want to deal with this stuff. Or to say that another way - if we can provide flexible recursive import & export - with an easy way to ignore bits of the tree or certain named entries with wildcards etc - folks can build whatever file system / SCM / operational / workflow mechanism they wish to maintaining the fabric config.	1	1	https://issues.jboss.org/browse/FABRIC-45

#	Requirement	Questions	Answered	Successfully Answered	URL
4257	FH-3707	Hi [~davidmartin], [~jmadigan] , [~maleck13], [~wtrocki], This command has a TODO: To check if it still valid for FH3. Could you help me to know if we should keep this one or just delete? Following the line code. https://github.com/feedhenry/fh-fhc/blob/master/lib/cmd/common/git.js#L4	1	1	https://issues.jboss.org/browse/FH-3707
4258	FH-3687	With google not allowing OAuth from WebViews, shouldn't we look into Chrome Custom Tabs instead? https://developer.chrome.com/multidevice/android/customtabs	1	1	https://issues.jboss.org/browse/FH-3687
4259	FH-3687	Either way, what I mean is to not use WebViews, if you don't like Chrome Custom Tabs, then use intents to open a browser, but not WebViews. It could be also Chrome Custom Tabs with a fallback in case Chrome is not installed. In the Chrome Custom Tabs doc they say {quote} What happens if the user doesn't have a recent version of Chrome installed? Custom Tabs uses an ACTION_VIEW Intent with key Extras to customize the UI. This means that by default the page willopen in the system browser, or the user's default browser. If the user has Chrome installed and it is the default browser, it will automatically pick up the EXTRAS and present a customized UI. It is also possible for another browser to use the Intent extras to provide a similar customized interface. {quote} Where I understand that if Chrome Custom Tabs are not available it will use the default browser as a fallback, but then they mention in another part of the docs {quote} Provide a fallback for when Custom Tabs is not installed Although Custom Tabs is available for the great majority of users, there are some scenarios where a browser that supports Custom Tabs is not installed on the device or the device does not support a browser version that has Custom Tabs enabled. Make sure to provide a fallback that provides a good user experience by either opening the default browser or using your own WebView implementation. {quote} Which makes me wonder if they really provide the fallback or if you have to do it instead.	1	1	https://issues.jboss.org/browse/FH-3687
4260	FH-3687	Chrome Custom tabs are implemented by Chrome and exposed as a service that you bind to. {quote} How can I check whether Chrome supports Chrome Custom Tabs?# All versions of Chrome supporting Chrome Custom Tabs expose a service. To check whether Chrome supports custom tabs, try to bind to the service. If it succeeds, then custom tabs can safely be used. {quote} If chrome isn't installed then Chrome Custom tabs won't work. It is "probably" enough to have a factory that returns a Intent that can be used for authorization and have it be implemented by a ChromeCustomTabIntent or a regular Intent depending on the system's capabilities. We will however have to create our own convenience method for launchUrl and use it. (see https://github.com/GoogleChrome/custom-tabs-client/blob/master/customtabs/src/android/support/customtabs/CustomTabsIntent.java#L261)	0	0	https://issues.jboss.org/browse/FH-3687
4261	FH-3659	Would a planet site like http://fedoraplanet.org/ or http://planet.gnome.org/ do what we want? AIUI the source of the software used is here: http://www.planetplanet.org/	0	0	https://issues.jboss.org/browse/FH-3659
4262	FH-3658	Hi [~trepel], I did the PR, however, I don't know exactly how to run it locally. Could you help me test it?	0	0	https://issues.jboss.org/browse/FH-3658
4263	FH-3657	[~spahuta], could you please verify and approve the PR for we merge it?	0	0	https://issues.jboss.org/browse/FH-3657
4264	FH-3649	*Command :* fhc version --format=json This command[1] is not working as follows and I don't think it is useful. Could we remove it? Or what is the utility of this? {code:java} Camilas-MacBook-Pro:fh-fhc cmacedo$ fhc version --format=json (node) sys is deprecated. Use util instead. fhc ERR! Error: Error - not 2xx status code. (null) fhc ERR! undefined fhc ERR! fhc ERR! System Darwin 15.6.0 fhc ERR! command "/Users/cmacedo/.nvm/versions/node/v4.4.3/bin/node" "/Users/cmacedo/.nvm/versions/node/v4.4.3/bin/fhc" "version" "--format=json" {code} [1] - https://github.com/feedhenry/fh-fhc/blob/master/lib/cmd/common/version.js c/c [~jmadigan], [~davidfmartin], [~wtrocki]	1	1	https://issues.jboss.org/browse/FH-3649
4265	FH-3649	[~wtrocki] Tks for the answers. Have few changes after removing the support of V2 here[1]. I catcher the need, will be easier keep work on here after the merge of the following PR. I will add here in the description just to remove the duplicated command. [1] - https://github.com/feedhenry/fh-fhc/pull/342/files#diff-ce6080cd558d9dd7168590f83cf90f7f	1	1	https://issues.jboss.org/browse/FH-3649
4266	FH-3644	Following the commands {code:java} Camilas-MBP:fh-fhc cmacedo$ fhc session info (node) sys is deprecated. Use util instead. fhc ERR! Invalid arguments: fhc ERR! fhc session info <session-token> fhc ERR! fhc session verify <session-token> fhc Command executed with error Camilas-MBP:fh-fhc cmacedo$ {code} Following some points. * The `fhc session verify` call the URL "box/srv/1.1/admin/authpolicy/verifysession" which is not been used into NGUI, however, I could find into the `fh-mbaas-api` : https://github.com/feedhenry/fh-mbaas-api/blob/b31e688baaeed950ec7fe7ae5e470b491e1293e5/test/fixtures/auth.js#L22 * The `fhc session info` I couldn't find into N project. Should we keep these commands? Can we delete them? c/c [~jmadigan], [~davidfmartin], [~wtrocki]	1	1	https://issues.jboss.org/browse/FH-3644
4267	FH-3644	Hi [~bienbos-dmartin], [~jmadigan], I could not found valid data, session-token, to do the success test in this one. Where can we find the token? Is this a valid/useful command? Could you help me with it or who can with it? How can we check it?	0	0	https://issues.jboss.org/browse/FH-3644
4268	FH-3640	I couldn't find this functionality into Ngui. Should we really keep this command? I think it should be removed from FHC since the studio doesn't do the same. Could you help me with this decision? Command: https://github.com/feedhenry/fh-fhc/blob/master/lib/cmd/common/ota.js c/c [~davidmartin], [~weil], [~jmadigan] , [~wtrocki]	1	1	https://issues.jboss.org/browse/FH-3640
4269	FH-3640	[~cmacedo86] To be clear, the OTA function is in NGUI. Every time when you build an app, there is a QR code showing up on the screen to allow users to scan & download the app, and that is the OTA functionality. In NGUI, it is not using `ota.feedhenry.com` anymore, but instead using an endpoint in millicore: https://github.com/fheng/fh-ngui/blob/master/public/javascript/app/models/platform.model.js#L301-L315. Given that in NGUI, the OTA function is available as part of build function, we can probably do the same in fhc - When the build is completed, fhc should print out the OTA url of the binary (and QR code perhaps?). Then we can remove the OTA command from fhc. How does this sound?	1	0	https://issues.jboss.org/browse/FH-3640
4270	FH-3631	This command do a call (request) into the project. {code:java} Camilas-MBP:fh-fhc cmacedo$ fhc call (node) sys is deprecated. Use util instead. fhc ERR! fhc call <url> [<method> <params>] fhc Command executed with error {code} Is really good we have this command? Should we keep it or can we delete? c/c [~jmadigan], [~davidfmartin], [~wtrocki]	1	1	https://issues.jboss.org/browse/FH-3631
4271	FH-3626	Hi [~leahliu], do you know how it should work? I tried to test it and I could not find a scenario to test it? Could you help me understand better it? Command: https://github.com/feedhenry/fh-fhc/blob/master/lib/cmd/common/auth.js	0	0	https://issues.jboss.org/browse/FH-3626
4272	FH-3624	Hi [~davidmartin], [~jmadigan], [~peter.braun], [~maleck13], [~wtrocki] Should we keep this command? Is it really useful for the users? Could we remove it instanced to move for v3 standard? Line: https://github.com/feedhenry/fh-fhc/blob/master/lib/cmd/common/appinit.js	1	0	https://issues.jboss.org/browse/FH-3624

#	Requirement	Questions	Answered	Successfully Answered	URL
4273	FH-3624	Hi [~davidmartin], Thank you for your comment. However, I'd like to add few infos here. - The URL of the rest called here, <box/srv/1.1/app/init>, is not been used into the studio/ngui.[1]. - I understand that it should be used in the past to make a call for the cloud app with the parameters informed and it is not working currently as follows. {code:java} Camilas-MBP:fh-fhc cmacedo$ fhc appinit 2png2u5t7ewrtn4mdi5l5eyy 6c333a91bb63ea09cb43cd0505a683976f7da2bc (node) sys is deprecated. Use util instead. fhc ERR! Error: Error - not 2xx status code. (500) fhc ERR! {"date":"Mon Jul 10 12:05:36 UTC 2017 - 1499688336793","status":"error"} fhc ERR! Request ID: 4b5b829e-c0fb-4caa-a6a8-85e554db5e90 fhc ERR! fhc ERR! System Darwin 15.6.0 fhc ERR! command "/Users/cmacedo/.nvm/versions/node/v4.4.3/bin/node" "/Users/cmacedo/.nvm/versions/node/v4.4.3/bin/fhc" "appinit" "2png2u5t7ewrtn4mdi5l5eyy" "6c333a91bb63ea09cb43cd0505a683976f7da2bc" fhc Command executed with error Camilas-MBP:fh-fhc cmacedo$ {code} Is possible it be legacy either? Maybe [~leahliu] could help us to understand better it? [1] - https://github.com/fheng/fh-ngui/search?utf8=%E2%9C%93&q=box%2Fsrv%2F1.1%2Fapp%2Finit&type=	1	0	https://issues.jboss.org/browse/FH-3624
4274	FH-3619	thanks [~cmacedo86] - will take a look with [~davidmartin]	1	1	https://issues.jboss.org/browse/FH-3619
4275	FH-3619	Hi folks, During PR review, [~wtrocki] suggested that we should re-write all project and create a new one, fhc2, using `yargs` as we can find into digkins-node[1]. Tks for all your contributions and help so far Wojciech :-). IMHO re-writing is not the case since the current one still adherent with requirements and technology stack is not deprecated. What we have here are "technical debts"[2] matters, the main issue for me is that commands are not following the same patterns which are in some cases preventing horizontal improvements. Also, I think that rewriting the whole project will cause secondary effects like testing the whole thing and keep on both projects the features and fix bugs synchronized until we deprecate the previous version. *Why follow the V3 pattern and not adopt another?* This is the most present standard in the project. In this way, the cost to accomplish this task is lower than adopt another. Also, this standard works well and it is easy to understand, keep, maintained and so on. *What is the idea of this JIRA?* The goal is to apply the same standard for all commands, removing deprecated code and organizing things, after that we'll be able to do more improvements step by step. Hi [~davidmartin] and [~jmadigan], Please, let me know what do you think about it? Could you tell us what should we do, re-write the whole project or continue working on "technical debts"[2]? I am looking for your answer to keep working on this tasks. [1] - https://github.com/aerogear/digkins-node [2] - https://martinfowler.com/bliki/TechnicalDebt.html [3] - https://issues.jboss.org/browse/FH-3619	1	1	https://issues.jboss.org/browse/FH-3619
4276	FH-3588	Seeing this error during the soak testing, after running for about 5-6 hours. {noformat} [2017-05-10T15:20:31.118Z] uncaughtException: [TypeError: Cannot read property 'syncCompleted' of undefined] [2017-05-10T15:20:31.119Z] Internal error: [TypeError: Cannot read property 'syncCompleted' of undefined] TypeError: Cannot read property 'syncCompleted' of undefined at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/api-syncRecords.js:23:35 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/sync-metrics.js:54:25 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/storage/dataset-clients.js:256:16 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/storage/dataset-clients.js:227:12 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:359:16 at nextTask (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:5057:29) at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:5064:13 at apply (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:21:25) at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:56:12 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:843:16 at listRecords (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/storage/dataset-clients.js:216:16) at nextTask (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:5070:14) at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:5064:13 at apply (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:21:25) at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:56:12 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:843:16 Error when update dataset client with id myShoppingList4_e37707c42bbc82d7153b1b87acc84ef98305b6ce. error = {} [myShoppingList3] unknown dataset client datasetId = myShoppingList3 :: queryParams = { user: 'client_1754' } [2017-05-10T15:20:31.925Z] uncaughtException: [TypeError: Cannot read property 'syncCompleted' of undefined] [2017-05-10T15:20:31.925Z] Internal error: [TypeError: Cannot read property 'syncCompleted' of undefined] TypeError: Cannot read property 'syncCompleted' of undefined at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/api-syncRecords.js:23:35 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/sync-metrics.js:54:25 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/storage/dataset-clients.js:256:16 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/storage/dataset-clients.js:227:12 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:359:16 at nextTask (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:5057:29) at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:5064:13 at apply (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:21:25) at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:56:12 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:843:16 at listRecords (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/lib/sync/storage/dataset-clients.js:216:16) at nextTask (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:5070:14) at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:5064:13 at apply (/home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:21:25) at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:56:12 at /home/fh/.fh-npm-node4/fh-mbaas-api/7.0.5/node_modules/async/dist/async.js:843:16 Error when update dataset client with id myShoppingList4_08f7e79abbae70754ca0c8bad33eac6145aff5a8. error = {} {noformat}	0	0	https://issues.jboss.org/browse/FH-3588
4277	FH-3560	[~wtrocki] Maybe this is what you are looking for? https://gitlab.cee.redhat.com/red-hat-mobile-application-platform-documentation/RHMAPDocsNG/blob/master/guides/cloud-api/setRecordHashFn.adoc. This is available in the new sync APIs and it will allow developers to provide a different implementation to determine if the object has changed. You can just return some values from some of the fields of the object and concatenating them together, you don't have to hash them. Does it make sense? cc [~ndonnell]	1	1	https://issues.jboss.org/browse/FH-3560
4278	FH-3538	See it's working but not 100% sure what it is doing ??	0	0	https://issues.jboss.org/browse/FH-3538

#	Requirement	Questions	Answered	Successfully Answered	URL
4279	FH-3521	Added in the form to take the sprint name and added in the post route . {code:java} // post sprint name to change the data going to the graph router.post('/burndown', function(req, res){ console.log(req.body.sprint); let burndown = kujiraDataMiner.burndownReportData(message, req.body.sprint); fs.writeFile('./public/js/burndown.json', JSON.stringify(burndown, null, 4), function(err){ if(err){ console.log(err); }else { console.log('Success'); } }); if(!sess || !sess.username){ res.redirect('/'); } else { res.render('graphs', { title: 'Kujira graphs', fields: fields, message: message }); } }); {code} Works for good data ie IR271, fails for bad data IR289, dosen't fail gracefully {code:java} TypeError: Cannot read property 'startDate' of undefined at Object.exports.getSprintDates (C:\Users\austin\redhat-repo\kujira\node_modules\kujira-data-miner\utils\data.js:69:25) at Object.exports.actualBurndown (C:\Users\austin\redhat-repo\kujira\node_modules\kujira-data-miner\utils\burndown.js:9:27) at Object.exports.burndownReportData (C:\Users\austin\redhat-repo\kujira\node_modules\kujira-data-miner\utils\burndownChart.js:9:37) at C:\Users\austin\redhat-repo\kujira\controller\accounts.js:51:34 at Layer.handle [as handle_request] (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\layer.js:95:5) at next (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\route.js:137:13) at Route.dispatch (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\route.js:112:3) at Layer.handle [as handle_request] (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\layer.js:95:5) at C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\index.js:281:22 at Function.process_params (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\index.js:335:12) {code} Issue with kujira-data-miner	1	1	https://issues.jboss.org/browse/FH-3521
4280	FH-3483	[~mziccardi] [~peter.braun] Think about this again. I think it might be better to return a random generated value as the session id for security reasons. We probably don't want to perform authentication for each getChunk request. If this is the case, the sessionId can be used as kind of simple security mechanism. If we always use the file id, it could mean that somebody can bypass the authentication step in the "startDown" step and just use this endpoint to download data. WDYT?	0	0	https://issues.jboss.org/browse/FH-3483
4281	FH-3478	[~ngough] I think that we need to: * remove this from 3.x docs * remove from 4.4 docs What I'm not sure about is whether to remove it from: * Supported configs page https://access.redhat.com/node/2357761 (or should we move it to deprecated?) * 4.2, 4.3 docs, given that there was support for this in the past	1	1	https://issues.jboss.org/browse/FH-3478
4282	FH-3468	Chose Accordion https://semantic-ui.com/modules/accordion.html#/definition Issue where couldn't nest accordion fix found at https://github.com/Semantic-Org/Semantic-UI/issues/513 Currently having an issue with handlebars not rendering the inner each , tried switching the outer loop to a helper !helper.png! Didn't resolve the issue found something on relative paths at http://jsfiddle.net/ambiguous/PNTXw/ that seems to render data	0	0	https://issues.jboss.org/browse/FH-3468
4283	FH-3468	http://stackoverflow.com/questions/42658156/nested-each-with-inner-value-as-key-for-outer-object-handlebars got the data to render for each field. {code:html} Results {{#each fields}} {{name}} {{#each ../message}} {{lookup this ../name}} {{/each}} . {{/each}} {code}	0	0	https://issues.jboss.org/browse/FH-3468
4284	FH-3463	session seems to get destroyed but seems to allow allow access to restricted routes /home /query after render. Only seems to fully disallow access on clicking the / route directly. Also /home and /query routes don't fail gracefully after a restart of the app {code:java} TypeError: Cannot read property 'username' of undefined at C:\Users\austin\redhat-repo\kujira\controller\accounts.js:45:11 at Layer.handle [as handle_request] (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\layer.js:95:5) at next (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\route.js:137:13) at Route.dispatch (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\route.js:112:3) at Layer.handle [as handle_request] (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\layer.js:95:5) at C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\index.js:281:22 at Function.process_params (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\index.js:335:12) at next (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\index.js:275:10) at Function.handle (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\index.js:174:3) at router (C:\Users\austin\redhat-repo\kujira\node_modules\express\lib\router\index.js:47:12) {code}	1	1	https://issues.jboss.org/browse/FH-3463
4285	FH-3463	Looks like the /login route fails also if app is restarted while on login page {code:java} Listening at http://localhost:8000 C:\Users\austin\redhat-repo\kujira\controller\accounts.js:67 sess.username = req.body.username; ^ TypeError: Cannot set property 'username' of undefined at C:\Users\austin\redhat-repo\kujira\controller\accounts.js:67:21 at ChildProcess.exithandler (child_process.js:197:7) at emitTwo (events.js:106:13) at ChildProcess.emit (events.js:191:7) at maybeClose (internal/child_process.js:877:16) at Socket.<anonymous> (internal/child_process.js:334:11) at emitOne (events.js:96:13) at Socket.emit (events.js:188:7) at Pipe._handle.close [as _onclose] (net.js:493:12) Process finished with exit code 1 {code} Resolved by adding sess = req.session; before pushing username to cookie (/) {code:java} // /login if successful applies the current username to the cookie router.post('/login', function(req, res){ child = exec('jira-miner target https://' + req.body.url + ' --user ' + req.body.username + ' --password ' + req.body.password, function (error, stdout, stderr) { if (stdout.indexOf('Successfully targeted JIRA') >= 0){ sess = req.session; sess.username = req.body.username; res.render('home',{title:'Kujira Home'}); }else{ res.render('login',{title:'Login to Kujira'}); } }); }); {code}	1	1	https://issues.jboss.org/browse/FH-3463
4286	FH-3445	Isn't this rather 4.4.0 given OSCP MBaaS is stressed?	1	1	https://issues.jboss.org/browse/FH-3445
4287	FH-3420	[~davidmartin] I think this is done as part of FH-3419. Can you take a look at the PR of that ticket and see if the content there is enough?	1	0	https://issues.jboss.org/browse/FH-3420
4288	FH-3414	I'm not planning to merge with template, it's working now even without changes. Because of fix in FH-3490. As for cluster, I'm leaving it in my github branch for possible future use.	0	0	https://issues.jboss.org/browse/FH-3414

#	Requirement	Questions	Answered	Successfully Answered	URL
4289	FH-3409	h2. Break Kujira into Modules broke run-cli-jira-miner into modules (/) app.js calls the modules (/) target.js provides login credentials and base url for jira (/) populate.js pulls data from the project specified (/) query.js passes a jira key to search.js function and returns a result (/) e.g. populate.js {code:java} /** * Created by acunningham on 20/03/17. */ 'use strict'; var prompt = require('prompt'); var query = require('./query.js'); var colors = require('colors/safe'); var exec = require('child_process').exec; var child; //validation for prompt see https://www.npmjs.com/package/prompt var schema = { properties: { project: { description: colors.blue('Enter the project name e.g. "RHMAP or RAINCATCH"'), pattern: /^[a-zA-Z0-9._-]+$/, message: 'Must be a valid project', required: true, }, }, }; // Start the prompt prompt.start(); prompt.message = colors.green('-->'); prompt.delimiter = colors.green(':'); var populate = function (callback) { // Get three properties from the user: username , password and url and project //console.log('Enter project name e.g. "RAINCATCH or RHMAP"'); prompt.get(schema, function (error, result) { // Log the results. console.log(' project: ' + result.project); // pull data from the target with populate, to populate local loki db child = exec('jira-miner populate "project in (' + result.project + ')"', function (error, stdout, stderr) { console.log('stdout: ' + stdout); console.log('stderr: ' + stderr); if (error !== null) { console.log('exec error: ' + error); } //checks standard out is present and the callback is a function if (typeof callback === 'function' && stdout != false) { callback(); } }); }); }; module.exports = populate; {code} ref: https://www.impressivewebs.com/callback-functions-javascript/ http://stackoverflow.com/questions/5187968/how-should-i-call-3-functions-in-order-to-execute-them-one-after-the-other#5188020	0	0	https://issues.jboss.org/browse/FH-3409
4290	FH-3409	load {code:html} <!DOCTYPE html><html lang="en"><head><meta charset="utf-8"><meta htt p-equiv="X-UA-Compatible" content="IE=edge"><title>Oops - an error has occurred< /title><script type="text/javascript">contextPath = "";</script><link type=\'tex t/css\' rel=\'stylesheet\' href=\'/static-assets/metal-all.css\' media=\'all\'><script src=\'/static-assets/metal-all.js\'></script><!--[if lt IE 9]><link type= \'text/css\' rel=\'stylesheet\' href=\'/static-assets/metal-all-ie.css\' media=\ 'all\'><script src=\'/static-assets/metal-all-ie.js\'></script><![endif]--><!--[if IE 9]><link type=\'text/css\' rel=\'stylesheet\' href=\'/static-assets/metal- all-ie9.css\' media=\'all\'><![endif]--><meta name="decorator" content="none" /> </head><body class=" error-page error500"><script type="text/javascript">documen t.body.className += " js-enabled";</script> <header id="header" ro le="banner"></header><!-- #header --><section id="content" role="main"><div clas s="aui-page-panel"> <section class="aui-page-pa nel-content lowerContent"> <h1>Sorry, we had some technical problems during your last operation.</h1><h2 cl ass="technical-details-header">Request assistance</h2> Copy the content b elow and send it to your JIRA Administrator <div class="technical-details-con tent" contentEditable readonly><h2>Technical details</h2> Log 's referral number: 170fea3e-2a0b-44ca-9e7b-cf41e7f92fc8 <h2>Cause</h2> Referer URL: Unkn own <pre>java.lang.RuntimeException: Error during Lightblue user load.</p re><blockquote id="stacktrace" style="overflow-x: auto;"><pre>java.lang.RuntimeE xception: Error during Lightblue user load.\n\tat org.jboss.jira.login.Lightblue Authenticator.loadJbossOrgUser(LightblueAuthenticator.java:190) [classes/:?]\n\t at org.jboss.jira.login.AbstractDelegatingCrowdDirectory.authenticate(AbstractDe legatingCrowdDirectory.java:88) [classes/:?]\n\tat com.atlassian.crowd.manager.d irectory.DirectoryManagerGeneric.authenticateUser(DirectoryManagerGeneric.java:2 83) [crowd-core-2.8.8.jar:?]\n\tat com.atlassian.crowd.manager.application.Appli cationServiceGeneric.authenticateUser(ApplicationServiceGeneric.java:194) [crowd -core-2.8.8.jar:?]\n\tat com.atlassian.crowd.embedded.core.CrowdServiceImpl.auth enticate(CrowdServiceImpl.java:69) [embedded-crowd-core-2.8.8.jar:?]\n\tat com.a tlassian.crowd.embedded.core.DelegatingCrowdService.authenticate(DelegatingCrowd Service.java:37) [embedded-crowd-core-2.8.8.jar:?]\n\tat com.atlassian.crowd.emb edded.core.FilteredCrowdServiceImpl.authenticate(FilteredCrowdServiceImpl.java:5 1) [embedded-crowd-core-2.8.8.jar:?]\n\tat com.atlassian.jira.security.login.Jir aSeraphAuthenticator.crowdServiceAuthenticate(JiraSeraphAuthenticator.java:75) [classes/:?]\n\tat com.atlassian.jira.security.login.JiraSeraphAuthenticator.auth enticate(JiraSeraphAuthenticator.java:49) [classes/:?]\n\tat com.atlassian.serap h.auth.DefaultAuthenticator.login(DefaultAuthenticator.java:88) [atlassian-serap h-3.0.3.jar:?]\n\tat com.atlassian.seraph.filter.PasswordBasedLoginFilter.runAut hentication(PasswordBasedLoginFilter.java:133) [atlassian-seraph-3.0.3.jar:?]\n\ tat com.atlassian.seraph.filter.PasswordBasedLoginFilter.login(PasswordBasedLogi nFilter.java:78) [atlassian-seraph-3.0.3.jar:?]\n\tat com.atlassian.seraph.filte r.BaseLoginFilter.doFilter(BaseLoginFilter.java:110) [atlassian-seraph-3.0.3.jar :?]\n\tat com.atlassian.jira.web.filters.JiraLoginFilter.doFilter(JiraLoginFilte r.java:82) [classes/:?]\n\tat com.atlassian.core.filters.AbstractHttpFilter.doFi lter(AbstractHttpFilter.java:32) [atlassian-core-5.0.6.jar:?]\n\tat org.apache.c atalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java :240) [catalina.jar:8.0.33]\n\tat org.apache.catalina.core.ApplicationFilterChai n.doFilter(ApplicationFilterChain.java:207) [catalina.jar:8.0.33]\n\tat com.atla	0	0	https://issues.jboss.org/browse/FH-3409
4291	FH-3312	h2. *Create express server that creates endpoints to consume jira-miner* Express server Listening on Listening at http://localhost:8000 POST endpoint /api/target (/) Takes json input in this format url,username and password e.g. {code:json} { "url":"issues.jboss.org", "username":"austincunningham", "password":"password" } {code} {code:javascript} //jira-miner target takes json input url,username and password router.post('/target', function (req, res) { // execute jira-miner target to point at the source child = exec('jira-miner target https://' + req.body.url + ' --user ' + req.body.username + ' --password ' + req.body.password, function (error, stdout, stderr) { res.status(200).json(stdout); console.log('stderr: ' + stderr); if (error !== null) { console.log('exec error: ' + error); } }); }); {code} POST endpoint /api/populate (/) Takes json input in the following format project e.g. {code:json} { "project":"RAINCATCH" } {code} {code:javascript} //jira-miner populate takes json input project router.post('/populate', function (req, res) { // execute jira-miner target to point at the source console.log(req.body.project); console.log('jira-miner populate "project in (' + req.body.project + ')"'); child = exec('jira-miner populate "project in (' + req.body.project + ')"', function (error, stdout, stderr) { res.status(200).json(stdout); console.log('stderr: ' + stderr); if (error !== null) { console.log('exec error: ' + error); } }); }); {code} POST endpoint /api/query (/) Takes json input in the following format an array of objects that are field value pairs e.g. {code:json} [{ "field":"Assignee", "value":"'Austin Cunningham'" }, { "field":"key", "value":"RAINCATCH-630" }] {code} {code:javascript} // jira-miner query takes in an array of field value objects router.post('/query', function (req, res){ console.log(req.body); let searchString = ''; for (let i = 0; i < req.body.length; i++){ searchString += '--' + req.body[i].field +'='+ req.body[i].value +' '; console.log(searchString); } child = exec('jira-miner query search ' + searchString + ' --json', function (error, stdout, stderr) { res.status(200).json(stdout); console.log('stderr: ' + stderr); if (error !== null) { console.log('exec error: ' + error); } }); }); {code}	1	1	https://issues.jboss.org/browse/FH-3312

#	Requirement	Questions	Answered	Successfully Answered	URL
4292	FH-3260	100 Jira-Miner Functionality Utilisation of JIRA Miner JIRA Miner is a tool to gather data from JIRA and query them locally for better speed and more flexible query language than jql. First you need to target the data source {code:java} jira-miner target <JIRA uri> --user [username] --password [password] {code} To get a response like e.g. {code:java} Jira-miner target https://issues.jboss.org --user austincunningham --password ********** { url: 'https://issues.jboss.org', user: 'austincunningham' } 'Successfully targeted JIRA' {code} Which populates the Contents of .jira-miner {code:java} {"jira":{"url":"https://issues.jboss.org","user":"austincunningham","password":"passwordhere"}} {code} Once, you have targeted a jira, you can populate local database {code:java} jira-miner populate <jqlQuery> {code} To get a response like e.g. {code:java} jira-miner populate "project in (RAINCATCH)" Updated and stored collection default in \home\acunningham\.jira-minerdb {code} You can then use the data in the local database to run JIRA Miner queries Query the database JIRA miner provides an API to query the database. Query must be defined in one of following formats: {code:java} const query = ctx => { // ctx.acc contains result of previous call // ctx.collection contains Loki.js collection // ctx.args contains optional arguments, such as ones on command line } // or const query = [ctx => { return xyz }, ctx => { // ctx.acc now contains a result of previous function call }] // or const query = [{ 'age': {'$gt': 21 } /* Loki.js Mongo like query */ }] // or mix of function and Loki.js queries module.exports = query {code} Supposing that this file, called my-example-query.js is in current directory, you call it via following {code:java} jira-miner query my-example-query [--args] {code} {panel:title=CLI Help} usage: jira-miner query <file> [options] Queries locally populated database using query(ies) located in <file>. Query file is a Node.js module exporting "query(ctx)" object/array/function and optionally a "transform(ctx)" function. module.exports = { query: (ctx) => { // ctx.args contains command line arguments that can be used to parametrize query externally // ctx.acc contains result of the last subquery // ctx.collection contains collection to be queried }, transform: (ctx) => { // ctx.args contains command line arguments that can be used to parametrized transformation externally // ctx.data contains result of query function } } Note that query can be a array of functions or Loki.js query objects. In such subqueries are executed in sequence and users are encouraged to use "ctx.acc" object to chain execution. Since "transform(ctx)" function is optional by default jira-miner expect "query(ctx)" to return a JSON data and uses pretty print to output it to console. If this is not desired behavior, user can provide --json or --csv argument. In case that "transform(ctx)" is provide, it is responsible for handling output itself. Options: --help Show help [boolean] --debug Logs debug information to console output [boolean] [default: false] --version Show version number [boolean] --db, -d Database location [default: .jira-minerdb in HOME directory] --json, -j Print out query outcome in JSON format [default: false] --csv Print out query outcome in CSV format [default: false] {panel}	1	1	https://issues.jboss.org/browse/FH-3260
4293	FH-3220	[~davidmartin] the description links to [here|https://github.com/trentm/node-bunyan#level-suggestions] for guidelines. From there though, it says: bq. Trent's biased suggestions for node.js libraries: IMHO, libraries should only ever log at trace-level. Fine control over log output should be up to the app using a library. Having a library that spews log output at higher levels gets in the way of a clear story in the app logs. I'd have previously considered fh-mbaas-api to be a library, rather than a server app itself. Does that sound reasonable? Does that mean that we should only log at trace level?	1	1	https://issues.jboss.org/browse/FH-3220
4294	FH-3220	[~davidmartin] on these points: {quote} * if there is an error situation in the code, still log it below the default level as long as the error can be propegated back to the calling code * if there is an error situation in the code that can't be propagated back to the calling code, log it at error level e.g. if a background task like a sync worker hits an error {quote} This may not be as straightforward as I originally expected: when returning an error as the first parameter in a callback, should that be an error object, or a string with an error message? I'm not sure about the API that's exposed to consumers, but I think in the sync code as a whole, there's a mixture of these two approaches. This can make it more complex to wrap an error with the new information (that may have been logged at error level, but not actually returned to the consumer), because it's not always clear at a particular point, whether an Error object or a String (or something else) is being returned. At least initially, unless you have more information than I do (likely!), I'll just leave such log entries as error level.	0	0	https://issues.jboss.org/browse/FH-3220
4295	FH-3219	[~vsazel] Some notes from a conversation on bluejeans: * fh-mbaas-api uses its own 'consolelogger' implementation * it currently isn't possible to pass in a logger to fh-mbaas-api (the logger is a hardcoded instance of 'consolelogger') * a follow up issue should be created to allow passing in a logger to fh-mbaas-api * for this issue, this consolelogger implementation should be passed into the sync code, and sync should use it as its logger * liase with [~grdryn] on changes in https://issues.jboss.org/browse/FH-3220 to ensure the 'consolelogger' in fh-mbaas-api is sufficient for the type of logging we're doing in sync Does this sound correct?	0	0	https://issues.jboss.org/browse/FH-3219
4296	FH-3189	[~jmadigan] [~davidmartin] [~craig.brookes] [~weil] since you guys are long standing FeedHenries, can I ask you for a careful review of the JIRAs ? thanks!	1	1	https://issues.jboss.org/browse/FH-3189
4297	FH-3177	[~beve] given that you have worked on FH-3170, do you want to take this one as well?	1	1	https://issues.jboss.org/browse/FH-3177
4298	FH-3170	{quote} It should keep the data handlers overrides for each individual dataset in memory, or use the default global data handler it no overrides. {quote} After looking into this again I'm still not exactly sure about what we should do here (sorry). Should we still delegate to sync-DataSetModel as done [here|https://github.com/danbev/fh-mbaas-api/commit/c84e027451f26940fd40438cfb2e6c11d4dac69e].	1	1	https://issues.jboss.org/browse/FH-3170
4299	FH-3136	[~akeating1] As this is marked as Dev Complete, is there a PR for it we can link to the issue?	1	1	https://issues.jboss.org/browse/FH-3136
4300	FH-3123	Do the individual tickets in this epic need to have a fix version set?	1	0	https://issues.jboss.org/browse/FH-3123
4301	FH-3115	Should we write a deprecate notice on the blank-cordova-app README.md? and also close the issues related to that template?	1	1	https://issues.jboss.org/browse/FH-3115
4302	FH-3087	[~davidmartin] Can this issue be closed? I don't think it's relevant in the new sync framework? Now that we're using debug for the whole of fh-mbaas-api, should FH-3394 also be closed?	1	0	https://issues.jboss.org/browse/FH-3087
4303	FH-2996	Hi [~alanmoran] and [~jhellar] this is tagged for 3.16 and has an associated PR with it. Can we land it, verify it and merge it OR can we disassociate the 3.16 fix version. Alan I don't know if teams are aware of this ticket as you tagged it 3.16.0 but it missed the teams filters as we normally filter our Sprints on untagged release versions -- we tend to set the release version ourselves when the ticket is complete. So this one slipped between the cracks.	1	1	https://issues.jboss.org/browse/FH-2996
4304	FH-2993	I really like the fhc projects -mine idea. What do you think about something like {{fhc projects list amoran@redhat.com}}? It would be less restrictive this way and it gives the ability to search for other users projects too. The usage would be {{fhc projects list [<author-email>]}}	0	0	https://issues.jboss.org/browse/FH-2993
4305	FH-2992	Hi [~grdryn], I believe that is already fixed. Look the last PR as example: http://bob.feedhenry.net:8080/job/fh-fhc_pullrequests/663/console. Could you check ?	1	0	https://issues.jboss.org/browse/FH-2992

#	Requirement	Questions	Answered	Successfully Answered	URL
4306	FH-2981	Suggestion from [~fantasyxky] {quote} Allow counter resetting, parameters, tags, filtering (See how influxdb/grafana can do this) {quote} Should we extend fh.stats API? OR should we look at providing grafana & influxdb?	0	0	https://issues.jboss.org/browse/FH-2981
4307	FH-2978	[~bgallagh] [~fantasyxky] This request came in from the paltform feedback here https://docs.google.com/document/d/1Uw5-vz-d691o10iQfdrtdaflajGc-TRWnZu7Hov3H6k/edit Can you give some context around this feature request so we can better understand what its trying to solve, and for who?	1	1	https://issues.jboss.org/browse/FH-2978
4308	FH-2978	[~evanshortiss] just going through the backlog, is this still a valid request for TKE?	1	1	https://issues.jboss.org/browse/FH-2978
4309	FH-2956	[~corinnekrych] I'm wondering now why we aren't providing, at the most, a simple wrapper for Keycloak in the SDK? Something similar to `$fh.keycloak` which just exposes the JS Keycloak client and all of its functions, maybe add some small platform-specific helper functions if they'd be useful? Is there something I'm missing for why something like `$fh.auth({ policyId: "Keycloak "}, ...)` would be better?	1	1	https://issues.jboss.org/browse/FH-2956
4310	FH-2956	[~jfrizell] Sorry for getting back to this so late. I think that clears a lot of things up. Thank you. However, according to the PRD thing we are not looking at RHMAP Auth Policies at the moment and assuming the user has RHSSO / Keycloak hosted somewhere themselves. [Link to section in PRD|https://docs.google.com/document/d/1Delb4Nmi4GY1G1BkMArXQkxi1E0EvOiSpTtmS0V3GIg/edit#heading=h.kojksmf5p4db] If this Spike should take into account working with an RHSSO policy, maybe it would make sense to defer it until we are looking at an RHSSO Policy in a later phase? I, at least, was under the impression that this section of Phase 1 would involve exposing the Keycloak JS client in the client SDK WDYT? I don't know if I'm making much sense at this point? Need sugar!	0	0	https://issues.jboss.org/browse/FH-2956
4311	FH-2941	It looks like there is some kind of problem on the build phase for https://github.com/feedhenry-templates/welcome-windows/ which makes https://github.com/feedhenry-templates/welcome-windows/pull/8 flag an issue. Can someone check? [~wtrocki]	1	1	https://issues.jboss.org/browse/FH-2941
4312	FH-2941	[~corinnekrych][~supittma] https://github.com/feedhenry-templates/welcome-windows/pull/8 Corinne or Summers, this PR is being blocked because of an expired certificate. Updating the certificate should be fairly simple on a Windows machine with Visual Studio installed. Can you help me on this? I only have an apple machine and it keeps failing when I try to update it. thanks!	1	1	https://issues.jboss.org/browse/FH-2941
4313	FH-2941	[~wtrocki][~pmdarrow] Hey guys, can any of you review these PRs? Summers has already fixed the building issue :)	1	1	https://issues.jboss.org/browse/FH-2941
4314	FH-2940	Is 'platform' alone a good category name? Do we need to change anything else in any other service? [~wtrocki] any idea?	1	1	https://issues.jboss.org/browse/FH-2940
4315	FH-2939	How can I check this was working? (mysql access)	1	1	https://issues.jboss.org/browse/FH-2939
4316	FH-2928	How can I know which are the most used? [~wtrocki]	1	1	https://issues.jboss.org/browse/FH-2928
4317	FH-2927	Where can we do a review of naming convention for templates, samples etc as part of this epic?	1	1	https://issues.jboss.org/browse/FH-2927
4318	FH-2925	[~coneill] I am wondering what is missing here. - For CSV export, the user has the option to export by field name and field code. Do we need to add an option to export by field ID also? - For PDF export, we could add the field ID or field code to the exported PDF. The user could have the same option to view field ID, field code or both?	1	1	https://issues.jboss.org/browse/FH-2925
4319	FH-2895	Further question for the team - I can see these templates being updated from Android N branch to master branch. Are they landing in 3.15.0 tag? And what environment has been used to run verification? It should be QED1, that now contains 3.15.0-rc2.	1	1	https://issues.jboss.org/browse/FH-2895
4320	FH-2891	Could we at least do this for the ones that are already open-sourced?	1	1	https://issues.jboss.org/browse/FH-2891
4321	FH-2891	There are around ~200 repositories that aren't forks or mirrors under the 3 organizations: * FeedHenry - 94 repos * FeedHenry Templates - 98 repos * FeedHenry Engineering - 13 repos # Does every one of these need CONTRIBUTING.md? If not, specifically which ones do? # Should we create a story to clean up old, unused repositories? There's a lot of repositories that look out of date, e.g. https://github.com/feedhenry/FHActProxy.	1	1	https://issues.jboss.org/browse/FH-2891
4322	FH-2891	Added CONTRIBUTING.md for all mbaas components and dependencies outlined in RHMAP-10046, temporarily disabling branch protection where necessary to the the job done. Can finally put this to rest now!	0	0	https://issues.jboss.org/browse/FH-2891
4323	FH-2888	[~davidmartin] does this only come into effect for 3.16? does it affect 4.x? referring to https://gitlab.cee.redhat.com/red-hat-mobile-application-platform-documentation/doc-RHMAP_Docs/merge_requests/90/diffs	1	0	https://issues.jboss.org/browse/FH-2888
4324	FH-2886	Hello [~supittma], can we move that to 3.16.0 given 3.15.0 release date is slowly approaching?	1	1	https://issues.jboss.org/browse/FH-2886
4325	FH-2877	[~akeating1]Could you link to the JIRAs around adding nsp, security check? I couldn't find them in backlog? do they have clientsdks label?	0	0	https://issues.jboss.org/browse/FH-2877
4326	FH-2876	One part of the flow that may be worth adding, is how a developer can push changes to a git repo and have them reflected in his local MBaaS. - Can he push code to a repo and it will then trigger a build - Can he build a docker image and push it to the OSCP repo and have it redeploy	0	0	https://issues.jboss.org/browse/FH-2876
4327	FH-2875	[~wtrocki] Who is requesting this template? Should it be a template? A sample app sounds more appropriate. We have a old websocket sample client & cloud app already https://github.com/david-martin/websocket-demo It would probably need updating if we want to use that as the sample (the folder structure is very old with both client & cloud in the 1 repo)	1	0	https://issues.jboss.org/browse/FH-2875
4328	FH-2855	[~emilioicai] Initial review. As mentioned on slack {quote} The doc looks great. When you add the other section for 'adding mongo-express to an existing cloud app', might be best to put that as the first option? {quote}	1	1	https://issues.jboss.org/browse/FH-2855
4329	FH-2853	[~cbyrne] Initial change was done and merged. Can we include this in upcoming release. Please refer to the deploy notes.	1	1	https://issues.jboss.org/browse/FH-2853
4330	FH-2851	[~davidmartin] what about create a google form and send it out to the consulting team and gather feedback? This way we can ask some specific questions?	1	1	https://issues.jboss.org/browse/FH-2851
4331	FH-2851	h4. Status update: - Created sample forms app - Created request to get access to the us.feedhenry.com cluster for hosting h4. Proposed Questions Dear platform user. We planning to improve data access layer for RHMAP. To help us to improve it please find 2 minutes to provide answers for the form bellow. 1. Do you use DataBrowser mongo import and export functionality? 2. What's is the main reason for using/not using it? 3. What do you like or dislike in import and export functionality? 4. How useful for you would be get direct access to the mongo database? 5. Would you use data browser export functionality if it would export only the current query? 6. Do you store large amounts of data in mongodb (more than 1GB collections)? 7. How you manage your mongo datasets?	1	1	https://issues.jboss.org/browse/FH-2851
4332	FH-2851	Some thoughts: - I would add a "Why" for every Yes/No question - It would be nice to know how do they manage large amounts of data currently - Should we also get the general feeling of using the import/export functionality? I mean maybe a generic question like "How do you feel about the import/export functionality?" would drive more insight than a direct one as "Did you experience any problem?"	1	1	https://issues.jboss.org/browse/FH-2851

#	Requirement	Questions	Answered	Successfully Answered	URL
4333	FH-2851	My thoughts: * The link doesn't send me directly to the form but to the list of forms (pretty much what you said [~wtrocki]) * Strange permission requested: !Screen Shot 2016-11-23 at 1.22.34 PM.png|thumbnail! * Why does loading take so long that we need a progress bar? I feel like a simple form should load in < 200ms!	1	1	https://issues.jboss.org/browse/FH-2851
4334	FH-2851	[~pmdarrow] That's dogfooding experience for our forms. We wanted to use our internal forms to collect feedback. How about questions itself? Can you try to fill this as you were probably using data browser before.	1	1	https://issues.jboss.org/browse/FH-2851
4335	FH-2848	[~craig.brookes] I have been looking through the 3.14 documentation. I can see many references for "fh-mbaas-api" throughout the docs however I cannot find even one reference for "$fh.feed". I assume I should see something - am I doing something wrong, ie: should I broaden the scope of my search outside of 3.14 docs (if so, where should I look do you think)?	1	0	https://issues.jboss.org/browse/FH-2848
4336	FH-2842	There are many times when an Anonymous class seems slightly excessive and cumbersome and as of Java 8, a new structure has been introduced to this effect, the Lambda Expression. A Lambda Expression enhances the readability of code and also allows one to express instances of single-method classes more compactly. For more info, please refer to the Lambda Expression tutorial : https://docs.oracle.com/javase/tutorial/java/javaOO/lambdaexpressions.html. Within the context of FeedHenry, Lambda Expressions have an impact on Apps that use Jack and Java 8, specifically around the consumption of a Lambda Expression by the FeedHenry API. Be aware of the fact that as of Java 8, a Lambda Expression can be passed into a method that consumes an Anonymous Inner class. Given this, it is highly likely that Lambda Expressions will impact DefaultCallbacks as illustrated in the following code snippet: @FunctionalInterface public interface DefaultCallback extends FHActCallback { default void fail(FHResponse pResponse) { Log.e("DefaultCallback", pResponse.getErrorMessage()); }; @Override void success(FHResponse pResponse); } The following code snippet illustrates the use of a Lambda Expression within a FeedHenry context: @Override protected void onStart() { super.onStart(); FH.init(this, (DefaultCallback)(ignore) -> { findViewById(R.id.check_online).setVisibility(View.VISIBLE); findViewById(R.id.check_online).setOnClickListener(view -> Toast.makeText(getBaseContext(), FH.isOnline()?"Is online":"Is not online", Toast.LENGTH_LONG).show()); }); }	1	1	https://issues.jboss.org/browse/FH-2842
4337	FH-2826	[~pbrookes] thoughts on doing this in the dump too? Is it something we want to action?	1	0	https://issues.jboss.org/browse/FH-2826
4338	FH-2825	[~avibelli] we've decide during planning to go with the dependency on tar for 4.2.0, and investigate a pure-Go solution in 4.3.0. Therefore we need to update the rpmspec to depend on GNU tar and Gzip, if it doesn't already. Could you please create a JIRA for that work? Thanks.	0	0	https://issues.jboss.org/browse/FH-2825
4339	FH-2811	[~corinnekrych]Can we get this fixed for rc3, if not lets move it to 3.15.0	1	1	https://issues.jboss.org/browse/FH-2811
4340	FH-2801	Hey guys can you identify which platforms this issue is referring to, give an example project/gist, and the use case? Is this a cordova issue or iOS? Does it affect Android, etc.	0	0	https://issues.jboss.org/browse/FH-2801
4341	FH-2801	[~mwessendorf] [~tkonishi] Would it be possible to get some more information about this issue?	1	0	https://issues.jboss.org/browse/FH-2801
4342	FH-2801	[~akeating1],[~danielpassos], sorry for my late response. Cloud app sample: https://github.com/RHMAP-Support/apn-silent_notification-cloud Cordova app sample: https://github.com/RHMAP-Support/apn-silent_notification-cordova AeroGear's push object is referred here: https://github.com/RHMAP-Support/apn-silent_notification-cordova/blob/master/www/js/index.js#L49-L51 Sending a silent push: https://github.com/RHMAP-Support/apn-silent_notification-cloud/blob/master/lib/push.js#L17-L21 How to reproduce: # create a bare project on RHMAP # import these applications # modify the client app id in config.xml to fit to your certificate # enable iOS push on the cordova app # build it # install the app to an iOS device then run the app # connect the device to a Mac, start XCode then open device window to check a log of the device # deploy the cloud app # {{curl <the cloud app url>/push}} You can see messages like this: {code} Mar 30 12:43:41 myiphone Push Starter[34167] <Notice>: Background Fetch Mar 30 12:43:41 myiphone Push Starter[34167] <Notice>: didReceiveNotification Mar 30 12:43:41 myiphone Push Starter[34167] <Notice>: Notification received {code}	0	0	https://issues.jboss.org/browse/FH-2801
4343	FH-2791	What's this about ? Not sure we can reboot the fh-digger component very often, just because of some setting Or is this a tasks that needs to be done, when the AndroidN/Container feature is deployed to wild (aka production) ?	1	1	https://issues.jboss.org/browse/FH-2791
4344	FH-2791	[~supittma] [~danielpassos] so what's actually needed ? Work on the template/apps ?	1	1	https://issues.jboss.org/browse/FH-2791
4345	FH-2788	How about using the standard one that Visual Studio generates when creating a project: https://github.com/feedhenry-templates/helloworld-windows/blob/master/.gitignore	1	1	https://issues.jboss.org/browse/FH-2788
4346	FH-2788	This is difference between the standard one that Visual Studio generates and the https://github.com/github/gitignore: https://www.diffchecker.com/wwFk1CPh Would it make sense to merge it?	1	1	https://issues.jboss.org/browse/FH-2788
4347	FH-2787	There are other templates: * https://github.com/feedhenry/appforms-project-client * https://github.com/feedhenry-templates/helloworld-app * https://github.com/feedhenry-templates/fh-mobile-spec-app * https://github.com/feedhenry-raincatcher/raincatcher-demo-mobile Should not be those updated as well?	0	0	https://issues.jboss.org/browse/FH-2787
4348	FH-2782	Is this really needed ? :) I feel during one release there are way too many bumps of the version :)	1	1	https://issues.jboss.org/browse/FH-2782
4349	FH-2764	[~weili] [~mwessendorf] 'fh-mbaas-core' could be a bit misleading to Users and RHMAP developers. It sounds like something the 'fh-mbaas' component might use, but then has 'core' in the name which makes it sound like somthing for the 'fh-mbaas' talking back to the 'Core'. To avoid confusion maybe we could have a common prefix e.g. 'fh-mbaas-api-core', 'fh-mbaas-api-forms' wdyt?	1	0	https://issues.jboss.org/browse/FH-2764
4350	FH-2764	[~conoro] [~mwessendorf] I wonder are we missing a potential new cloud app template here on the back of this epic? What I'm thinking is we don't have any template that demonstrates an absolute minimal resource footprint, but still has fh api capabilities. If this new template was a goal of this epic, it might help drive how the server sdks could be made 'pluggable'. Also, the new template could be used as a reference for developers on best practices to keep a node.js app resource usage low. Perhaps we could get some input on it from the community (customers and consultants?) as well during development. What do you think?	1	0	https://issues.jboss.org/browse/FH-2764
4351	FH-2764	[~weil] This came up at triage today. Do you want to tag it and all issues as team-sustaining?	1	0	https://issues.jboss.org/browse/FH-2764
4352	FH-2764	[~davidmartin] this is really for the server sdk team. Which label they are using atm?	1	0	https://issues.jboss.org/browse/FH-2764
4353	FH-2741	Where is a link to the offending line : https://github.com/feedhenry/fh-android-sdk/blob/775dbc32ae47d7a1ed179823fc18628795cd3a97/fh-android-sdk/src/main/java/com/feedhenry/sdk/sync/FHSyncClient.java#L418	1	1	https://issues.jboss.org/browse/FH-2741

#	Requirement	Questions	Answered	Successfully Answered	URL
4354	FH-2739	[~akeating1] No worries at all - I'm a little unsure of what the status is too. [~supittma] What is the expectation re publishing this work? The title reads, "Document how lambdas should work and make epic for refactoring SDK" and to this point, I've documented "...how lambdas should work...". Do you/your team/someone on your side/etc publish this or do Docs publish this?	1	1	https://issues.jboss.org/browse/FH-2739
4355	FH-2739	[~supittma] : Where is this documentation supposed to land? Would it be customer facing documentation? ^ Updated by John Stafford ^	1	1	https://issues.jboss.org/browse/FH-2739
4356	FH-2739	[~kpiwko] Good to get feedback Karel. Re "this documentation should be kept close to Swift 2.3/3.0 one", not fully sure what you mean. I reviewed the table of contents and also searched the doc [https://docs.google.com/document/d/1UV2l9z-T3axQbMrqFlsXrVQsmvme0LgX4XVg5wstmpg/edit] for the keyword "Swift" and there is no mention of this. Is there another document for "Swift 2.3/3.0"? If so, can you provide me with a link to this "Swift 2.3/3.0" documentation?	1	0	https://issues.jboss.org/browse/FH-2739
4357	FH-2736	[~corinnekrych] Can this ticket be closed?	1	1	https://issues.jboss.org/browse/FH-2736
4358	FH-2734	[~ndonnell] this is something for the appforms work, and just having the serversdks label, because the code is hosted there? Is this work already planed to be done by your team ?	0	0	https://issues.jboss.org/browse/FH-2734
4359	FH-2732	[~ndonnell] this is something for the appforms work, and just having the serversdks label, because the code is hosted there? Is this work already planed to be done by your team ?	0	0	https://issues.jboss.org/browse/FH-2732
4360	FH-2728	[~ndonnell] this is something for the appforms work, and just having the serversdks label, because the code is hosted there? Is this work already planed to be done by your team ?	0	0	https://issues.jboss.org/browse/FH-2728
4361	FH-2727	[~ndonnell] this is something for the appforms work, and just having the serversdks label, because the code is hosted there? Is this work already planed to be done by your team ?	1	0	https://issues.jboss.org/browse/FH-2727
4362	FH-2724	Is not MBaaS-ART exactly that? https://github.com/fheng/mbaas-art	0	0	https://issues.jboss.org/browse/FH-2724
4363	FH-2721	IIUC, this is purely internal documentation, is that correct, [~ndonnell]? If so, could you assign the "tools/help" component instead of Documentation? Thanks.	0	0	https://issues.jboss.org/browse/FH-2721
4364	FH-2705	Where can I test this templates? qed1 an meh1 won't build them	0	0	https://issues.jboss.org/browse/FH-2705
4365	FH-2702	It will only use ChildBrowser if it's installed, which is not the case for new templates. It has an else to use InAppBrowser *if installed*, which isn't the case neither. Should we install InAppBrowser plugin in all Cordova templates?	1	0	https://issues.jboss.org/browse/FH-2702
4366	FH-2702	I have checked the code again and if ChildBrowser is not installed it will use window.open, but that window.open is for using inAppBrowser, not a regular window.open. The problem is that if the inAppBrowser is not installed, it will fail on next line where it adds an event listener for the inAppBrowser object. That works if the inAppBrowser is installed, but won't work on a regular window.open, and as that code is in a try-catch it will go to the catch and nothing will happen. So, should we document that the inAppBrowser is needed when using fh.auth with an OAuth2 auth? or should we just add inAppBrowser to all Cordova templates? Also, if the OAuth2 is for Google, they don't allow webviews anymore, so, should we update the fh-js-sdk to use a different approach?	0	0	https://issues.jboss.org/browse/FH-2702
4367	FH-2697	[~danielpassos] Are these tasks (Android, Windows, JS) still required/relevant? If so, would it be possible to expand a bit more on the task itself. I'm not sure if I'm understanding this correctly. Is this just adding a cordova build into the `travis.yml` file in the components to ensure it builds? Thanks :)	1	1	https://issues.jboss.org/browse/FH-2697
4368	FH-2653	Might be good to mention some limitations that GCM 3.0 brings, wdyt ? * AGPUSH-1625 * number of devices no longer represents number of active installations (we discussed this on IRC) * number of dispatched notifications is incorrect (dispatched only one for whole topic)	0	0	https://issues.jboss.org/browse/FH-2653
4369	FH-2649	What does it mean to support GCM-3 on Xamarian ? IMO it's unclear. Currently the Xamarian-SDK uses the .NET SDK for push, but has specific knowledge for other networks (e.g. GCM for Android): https://github.com/feedhenry/fh-dotnet-sdk/blob/master/FHSDK/Services/Network/PushBase.cs Topic support addition is needed	1	1	https://issues.jboss.org/browse/FH-2649
4370	FH-2631	[~supittma], [~mwessendorf] since it's already merged should we schedule that for 3.10 instead of 3.11?	1	1	https://issues.jboss.org/browse/FH-2631
4371	FH-2627	Templates are there. There are separate subtasks for verification of each template. [~jbalunas] for some reason blank-template is called "Xamarin Hello World Project" and there is another "Native Xamarin Hello World Project" which is actually what it is. Should we change it?	1	0	https://issues.jboss.org/browse/FH-2627
4372	FH-2619	[~kpiwko] Good point, I've labeled the issue to see it on agile board. Could it be something [~jhellar] could help us on?	1	1	https://issues.jboss.org/browse/FH-2619
4373	FH-2615	[~jbalunas] / [~corinnekrych] - Is there anyone else we could get to review this? [~rsmeral] has a lot on his plate between now and Friday, so if we could get someone else from Client SDK team to get involved here, that would be super.	1	1	https://issues.jboss.org/browse/FH-2615
4374	FH-2614	Hello [~corinnekrych], do I read correctly this is a blocker for 3.10.0 release?	1	1	https://issues.jboss.org/browse/FH-2614
4375	FH-2593	[~mmurphy_fh] - Can you follow up on this? Why was the app never added to the template list?	0	0	https://issues.jboss.org/browse/FH-2593
4376	FH-2593	[~supittma] [~davidmartin] This feels like invalid bug. Can you provide more details. How we can help here?	1	1	https://issues.jboss.org/browse/FH-2593
4377	FH-2587	[~supittma] [~danielpassos] given the code freeze at the end of sprint 272, and updates to millicore required this does not seem like something we can get done in 272? I'll shift it to 3.X-Future to eval later. Please feel free to comment/update if needed.	1	1	https://issues.jboss.org/browse/FH-2587
4378	FH-2561	[~corinnekrych] Some of the issues (5) in this Epic are scheduled for 3.9 release. Since we are now targeting iOS Swift SDK support in 3.10, do they need to be updated? If so, can you update them?	1	1	https://issues.jboss.org/browse/FH-2561
4379	FH-2559	[~edewit] [~corinnekrych] Can we shift this epic to either not have a version, or to 3.10.1, or 3.11?	1	1	https://issues.jboss.org/browse/FH-2559
4380	FH-2555	[~corinnekrych] - there seems to be a lot of unrelated tickets in this Epic. Can you please review and comment on each of the tickets and identify which ones will land as part of 3.7.0.	1	0	https://issues.jboss.org/browse/FH-2555
4381	FH-2535	[~corinnekrych] - is this going to land in 3.7.0? Seems like it is better suited to 3.8.0.	1	1	https://issues.jboss.org/browse/FH-2535
4382	FH-2482	[~supittma][~edewit] sorry for the delay on getting back to this. We may want to shelve this until we work through our own windows build farm support options. I know the Xamarin account is separate, but it feels like we can tackle all of this at once with the updates in a few sprints. Wdyt?	1	1	https://issues.jboss.org/browse/FH-2482
4383	FH-2479	[~corinnekrych] - is this going to land in 3.7.0? Seems like it is better suited to 3.8.0.	1	1	https://issues.jboss.org/browse/FH-2479
4384	FH-2475	[~corinnekrych] - is this going to land in 3.7.0? Seems like it is better suited to 3.8.0.	1	1	https://issues.jboss.org/browse/FH-2475
4385	FH-2471	Can you project the app link?	0	0	https://issues.jboss.org/browse/FH-2471
4386	FH-2467	[~corinnekrych] is there anything I'm doing wrong? Has MEH1 BF been updated?	1	1	https://issues.jboss.org/browse/FH-2467

#	Requirement	Questions	Answered	Successfully Answered	URL
4387	RAINCATCH-1057	wtrocki: Do you think we can find all console.statements in the code and replace them in this PR? This will be very quick considering your improvements in this PR and then we can totally close logging story as more loggers will be added on the fly. discussed with wtrocki and it is beyond the scope of this ticket , will leave the PR open and close this ticket and open a fresh ticket to replace the console with the new logger.	0	0	https://issues.jboss.org/browse/RAINCATCH-1057
4388	RAINCATCH-1055	Adding my reply to some of the points: - Ability to have generic UI process When going to new step user need to have method to render specific template. This will be for example method wrapping angular router > I've been thinking about this and it's definitely very hard to decouple rendering from an being angularjs specific, but we can try to restrain it to a single method on the task, and provide a couple of utilities to DRY custom task implementations - Correlation with status and actual step processing. While having status is good idea I think that we need to have clearly define when we moving to next step, by just having exact method. I do not think we should remove status, but just disconnect it from actual step definition > Yes from what I've seen in the existing UI I think statuses are only used for display purposes. In a WFM scenario I don't think there is much need for reporting partial progress at all. - Summary screen is really angular specific Summary screen should be really outside WFM. We should have this as some step that will take all results, but then I'm not sure how to mount that into html (js specific implementation) that's to specific Any ideas? Dropping summary screen from requirements it's not a runner. > I think we either have to define a strict set of fields for the Result so they can be rendered uniformly by a given UI implementation, and specific needs can be addressed by custom detail views, or give results the same ability (and complexity) of being as customizable as Tasks - Tasks are not strongly coupled to entire model Tasks needed to be added as separate module or code in demo application. I think that's the way to go because it depends on what we expect users to be more likely to modify for their customization. I'm not sure I understand this right, currently Tasks handle the logic for dealing with user interactions to specifying what next step in a Process should be run, but if we somehow split these responsibilities users would have to provide alternative implementations for other entities too, and ways to configure that. We do not need any for task as we can save them as array of ID's for Process. ProcessInstance in my opinion will just need to hold current step status and array of results assigned to steps. From what I've seen from jbpm and other engines the need for data instantiation is mainly to allow Processes to be modified at any time and even explicitly versioned for the benefit of analysts. What would happen if one was changed while there is an active ProcessInstance assigned to a user? Nevertheless we can keep it this way for 1.0 In relation to additional functionalities like json schema I think we may retain them. > In retrospect maybe reducing the feature set to the bare minimum of supporting the existing UI would allow us to deliver better. If the idea is starting integration with jbpm a lot of our custom implementations would be redundant.	0	0	https://issues.jboss.org/browse/RAINCATCH-1055
4389	RAINCATCH-1036	[~wtrocki][~jstaffor] some notes * Cutting and pasting content isn't going to help us * Repeating content isn't going to help us * There's a danger of content going out of date So, I propose criteria for moving content: * if it's still being developed (ie significant changes, tag it, and leave it in the readme) * if it's not related to the user story, leave it in the readme (tag it and/or see if it's covered under https://issues.jboss.org/browse/RAINCATCH-914) * check with PRD, make sure it's addressing a requirement before moving WDYT?	1	0	https://issues.jboss.org/browse/RAINCATCH-1036
4390	RAINCATCH-1027	[~wtrocki] Client interface is the exact same as the cloud interface? Description is misleading. Can we update to specify whats in the client interface?	1	1	https://issues.jboss.org/browse/RAINCATCH-1027
4391	RAINCATCH-1014	This is only for moving the generated javascript files. Other issues can be addressed as follow on tickets. [~wtrocki][~phajidec] Can you create these follow on tickets for a backlog refinement session.	1	1	https://issues.jboss.org/browse/RAINCATCH-1014
4392	RAINCATCH-923	[~ndonnell] In the spikes along with [~wtrocki] the pattern we were using was having something like: {code:java} import MongooseStore, { Store } from '@raincatcher/store'; {code} So, our supported implementation as the default export and the public, reimplementable api as a named export. Do you see them as separate modules for now? Or perhaps for the store specifically? It's a pattern that we'd need to agree upon for the other packages. Depending on the decision we can merge RAINCATCH-922 into this one and have a single package.	1	1	https://issues.jboss.org/browse/RAINCATCH-923
4393	RAINCATCH-923	[~phajidec] Would it mean though that anyone wanting to use our Store interface would require the raincatcher mongoose store as a dev dependency. Would it not make more sense to have the Store interfaces and classes as a separate module that has a much reduced dependency list? WDYT ?	1	1	https://issues.jboss.org/browse/RAINCATCH-923
4394	RAINCATCH-921	[~wtrocki] Perfect, will I redefine this ticket to be for Passport and move it to the passport epic? There is already a ticket for Keycloak in this epic.	0	0	https://issues.jboss.org/browse/RAINCATCH-921
4395	RAINCATCH-914	[~wtrocki] Just confirming, do you see the documentation repo having a dev dependency on typedoc and running it on top of the packages in the monorepo? Or running it through lerna and then 'exporting' the generated html into the doc repo?	1	1	https://issues.jboss.org/browse/RAINCATCH-914
4396	RAINCATCH-875	I was under impression that user profile is attached to request so we do not need to have this as part of the interface. Can you provide more info how you will see this implemented here? For passport: Both login and logout will need to be separate endpoints that will be handled by passport. We just need to build template for this. Logged user will be always attached to req.user. So effectively the only one method needed will be the actual security middleware.	0	0	https://issues.jboss.org/browse/RAINCATCH-875
4397	RAINCATCH-875	Oh.. I get ya. I think it's best to provide clean separation between server support (req.user available in handlers) and client. As we assumed that all profile data will be in fact stored on our side we can have user module function that will get user data. Something like that: {code} router.get("/profile", rsecMiddleware, function(req,res) { res.json(req.user); } {code} Is that also possible for keycloak?	1	0	https://issues.jboss.org/browse/RAINCATCH-875
4398	RAINCATCH-873	Thanks [~austincunningham], can you maybe sync with [~wtrocki] to get this documentation and examples onto https://feedhenry-raincatcher.github.io?	0	0	https://issues.jboss.org/browse/RAINCATCH-873
4399	RAINCATCH-871	[~wtrocki] Is the aim at this point to still include node_acl integration, or will it be a simpler role check based on a certain user id?	1	1	https://issues.jboss.org/browse/RAINCATCH-871
4400	RAINCATCH-844	I want to extend idea to talk also about supported API. To we can auto generate code using: https://github.com/swagger-api/swagger-codegen	0	0	https://issues.jboss.org/browse/RAINCATCH-844
4401	RAINCATCH-843	FYI [~davidmartin] - Do you see type definitions being part of the fh-mbaas-api or we should have separate repo for the moment?	1	1	https://issues.jboss.org/browse/RAINCATCH-843
4402	RAINCATCH-771	[~wtrocki] Thinking I might either re-open this ticket and update the tickets description or just create a new ticket instead. For an automated deploy script, i think this should include: # Running the Keycloak Server (Docker looks Good) # Creation/Configuration of Realms using the Keycloak Admin CLI - This would be a script that creates all the required Keycloak configurations to integrate with the Raincatcher demo solution. WDYT?	1	1	https://issues.jboss.org/browse/RAINCATCH-771
4403	RAINCATCH-769	[~tjackman] Hey. We have this ticket covered in our docs story. Do you have any updates on this? Maybe we can close and move some results from here.	1	1	https://issues.jboss.org/browse/RAINCATCH-769

#	Requirement	Questions	Answered	Successfully Answered	URL
4404	RAINCATCH-765	[~tjackman] Let's finish and close all non essential spikes outside the scope. h4. Draft of proposed set of tasks for Raincatcher 1.0 1. Create Cloud Interface for Security 2. Raincatcher rsec module implementation for Keycloak - Create the keycloak realm to seed the user/roles - docs + startup script + how to extend - ##Blocker No support for realm updates - ##Blocker Offline suport sync - spike - ##Blocker SaaS offering may not have keycloak exposed - Implement Linking user module users to keycloak users as part of express middleware required to expose all available data. - Add configuration option triggering offline authentication (separate epic) - Implementation for all supported Raincatcher protect them by roles in demo application. (This will be using interface only) - Investigate sync support and user filtering. - Investigate Web API module support - may be not used. - Complete user security story in demo - how the keycloak demo realm is configured - Demo cloud should support for the security interface. docs. - Custom login theming. Documentation 3. ClientSec Implementation for keycloak (May include elements that are not initial requirement) - JS module coming from server? no interface? - Docs surrounding hiding of UI elements by role - Investigation of online/offline as part Raincatcher - demo-mobile integration for client - supporting/testing verification of cookies - demo-portal integration for client - supporting/testing verification of cookies - custom login theming? 4. Raincatcher rsec cloud module implementation for passport.js (previously dummy authentication module) - Implementation for basic security interface - implement login endpoint - should have redirects to match keycloak - add more data to user module including passwords to mongo db - implement credentials handler that will match credentials with user ping [~bofarrel]	0	0	https://issues.jboss.org/browse/RAINCATCH-765
4405	RAINCATCH-760	h4. What was done I replaced our homegrown authentication service with keycloak for demo portal. h4. Quick demo: https://drive.google.com/a/redhat.com/file/d/0B4yds-KCN6cvYW15MWFZMWU5dzA/view?usp=sharing h4. Changes: https://github.com/feedhenry-raincatcher/raincatcher-demo-cloud/compare/master...wtrocki:keycloak h4. Documentation: https://keycloak.gitbooks.io/documentation/securing_apps/topics/oidc/javascript-adapter.html Suggested steps to fully migrate to keycloak: - Use keycloak client in mobile apps to retrieve tokens in mobile/cookie in browser. - Migrate sample users from auth service to keycloak - Remove auth service integration in clients and cloud - Remove auth service h4. Why remove security service - It provides zero additional value when keycloak solution is used. - Makes really hard to configure access origin for sessions and tokens (It's best to configure keycloak on the edge which is cloud app). - Having security service complicates raincatcher setup. Without security service we can setup raincatcher in 1 step instead of 6 steps. - It's not aligned with raincatcher modular architecture where modules can be placed everywhere.	1	1	https://issues.jboss.org/browse/RAINCATCH-760
4406	RAINCATCH-730	[~wtrocki] - Are you looking form some action from me here?	1	1	https://issues.jboss.org/browse/RAINCATCH-730
4407	RAINCATCH-711	Currently there are problems with the Drag and Drop functionality in the Selenium WebDriver. This prevents the use of drag and drop using protractor or accesing the WebDriver directly. A workaround for this is to use a drag and drop simulator such as [html-dnd|https://www.npmjs.com/package/html-dnd]. Example usage with protractor as follows: {code:javascript} var dragAndDrop = require('html-dnd').code; browser.executeScript(dragAndDrop, elementToMove, targetElement); {code}	1	1	https://issues.jboss.org/browse/RAINCATCH-711
4408	RAINCATCH-711	[~thnolan] test fails for me locally. Any clues? What cluster you use? {noformat} Running "ui:local" (ui) task Running "clean:0" (clean) task >> 0 paths cleaned. Running "clean:1" (clean) task >> 0 paths cleaned. Running "eslint:all" (eslint) task Running "config:local" (config) task Running "protractor:local" (protractor) task [10:18:08] I/local - Starting selenium standalone server... [10:18:08] I/launcher - Running 1 instances of WebDriver [10:18:09] I/local - Selenium standalone server started at http://10.202.15.139:56788/wd/hub testing scheduler functionality in demo app Test Datepicker Interactions Test Case - Read date from datepicker 1) check the current date from the date picker 0 passing (5s) 1 failing 1) testing scheduler functionality in demo app Test Datepicker Interactions Test Case - Read date from datepicker check the current date from the date picker: AssertionError: expected '12/04/2017' to equal '4/12/2017' + expected - actual -12/04/2017 +4/12/2017 From: Task: testing scheduler functionality in demo app Test Datepicker Interactions Test Case - Read date from datepicker check the current date from the date picker at Context.ret (node_modules/protractor/node_modules/selenium-webdriver/testing/index.js:128:10) [10:18:16] I/local - Shutting down selenium standalone server. [10:18:16] I/launcher - 0 instance(s) of WebDriver still running [10:18:16] I/launcher - chromeANY #01 failed 1 test(s) [10:18:16] I/launcher - overall: 1 failed spec(s) [10:18:16] E/launcher - Process exited with error code 1 >> Warning: Tests failed, protractor exited with code: 1 Use --force to continue. Aborted due to warnings. {noformat}	1	1	https://issues.jboss.org/browse/RAINCATCH-711
4409	RAINCATCH-708	[~phajidec] What needs to be updated to include the new sync server changes in fh-mbaas-api? Is it just the demo cloud app at https://github.com/feedhenry-raincatcher/raincatcher-demo-cloud? Or are there modules that need updating as well?	1	1	https://issues.jboss.org/browse/RAINCATCH-708
4410	RAINCATCH-692	[~wtrocki] In the description: "File endpoints are currently not secured. Anyone with link to the file can see them." Have you steps to reproduce? The portal app will return a 401 if not authenticated, and mobile users won't be able to see them as sync will only download files created by the logged in user. If you change the file identifier in the URL to one from another user, it won't show it, as that file was never synced down locally and can't render if its not there.	1	1	https://issues.jboss.org/browse/RAINCATCH-692
4411	RAINCATCH-636	[~wtrocki] I have updated the description of the ticket. Does this help give a clearer picture?	1	1	https://issues.jboss.org/browse/RAINCATCH-636
4412	RAINCATCH-634	POC Project repo - https://github.com/JameelB/raincatcher-workflow-angular-demo Findings: *Can the module be instantiated without any other modules?* Module was successfully instantiated without any other modules apart from raincatcher-mediator. It was able to work fully with all CRUDL operations and is able to begin and complete a workflow process. *What extra configuration would be needed if they can't?* [Subscribers|https://github.com/JameelB/raincatcher-workflow-angular-demo/blob/master/src/app/mediator-subscribers.js] to the mediator topics are required to be set up. Form templates to be used for the workflow process are required to be set up as well *Should the module stand alone or be merged with another module?* I believe that it should be a stand alone module to avoid limiting development options.	1	1	https://issues.jboss.org/browse/RAINCATCH-634
4413	RAINCATCH-633	POC project repo - https://github.com/witmicko/raincatcher-demo-workorder-angular Findings: *Can the module be instantiated without any other modules?* Module worked fully including all CRUDL operations and only needed RainCatcher-mediator module *What extra configuration would be needed if they can't?* I needed to add [mediator subscribers|https://github.com/witmicko/raincatcher-demo-workorder-angular/blob/master/src/app/mediator_subscribers.js] *Should the module stand alone or be merged with another module* I believe it should be a standalone module as merging it with another module might limit development options.	0	0	https://issues.jboss.org/browse/RAINCATCH-633

#	Requirement	Questions	Answered	Successfully Answered	URL
4414	RAINCATCH-620	[~phajidec], [~davidmartin] my understanding is that this change is already on master, but is not yet the version that's delivered with the latest RHMAP version. Should we start using this now in our load testing work, or hold off? Edit: I think we might already have it on our branch, as a result of some rebasing I did last week in an attemp to fix some issues seen in the cloud app. The new changes didn't fix the issues that I encountered. Should I go back to the old version, or keep pulling changes in from master periodically?	1	1	https://issues.jboss.org/browse/RAINCATCH-620
4415	RAINCATCH-619	currently. These tests were run with the raincatcher solution running locally on my laptop, and the test runner operating from there also. The cloud app appeared to be consuming 75-100% of the CPU time that was available to it for the duration. I think these results are useful: when running on a production-like RHMAP3, we don't reach the same concurrency -- 10-15 failures are seen at a concurrency of 30. Something to note is that in some cases of failure, this could be a problem in the load-test scripts: any time they go into a sync loop, they'll time out after 2 minutes if their success condition is not reached (i.e. the cloud app is responding as expected, but the test is still failing). However, the number of "Mediator request timeout..." messages in the logs appears to roughly match the number of test failures. I'm not sure if that validates the test script logic. I'm hesitant to say that it does. As I currently understand it, a real client would timeout a single sync request after 2 minutes, is that right? The test script times out the looping after two minutes -- so there could be timely responses from the cloud app, but because of the number of clients concurrently updating a dataset, a given script execution may just not be lucky enough to successfully match the hash that the server has for long enough. Does that make any sense? Another rare issue that we've seen, typically at high memory usage (when the memory leak was there, this occurred also, so it seems more tied to memory usage than concurrency), is that the cloud app would crash with the following uncaught exception output in the error logs: {code} uncaughtException: [Error: Can't set headers after they are sent.] Internal error: [Error: Can't set headers after they are sent.] Error: Can't set headers after they are sent. at ServerResponse.OutgoingMessage.setHeader (_http_outgoing.js:344:11) at ServerResponse.writeHead (_http_server.js:180:21) at endResponseCallback (/home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/node_modules/fh-mbaas-express/lib/mbaas.js:116:9) at /home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/node_modules/fh-mbaas-express/lib/mbaas.js:182:16 at timeWrappedCallback (/home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/lib/util/metrics.js:68:31) at /home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/lib/sync-srv.js:193:24 at timeWrappedCallback (/home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/lib/util/metrics.js:68:31) at /home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/lib/sync-DataSetModel.js:401:7 at invokeCb (/home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/lib/sync-DataSetModel.js:510:13) at /home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/node_modules/fh-mbaas-client/node_modules/fh-logger/node_modules/continuation-local-storage/node_modules/async-listener/glue.js:188:31 at nextTickCallbackWith0Args (node.js:420:9) at process._tickCallback (node.js:349:13) {code} [~davidmartin] and I have tried to understand this onen a bit, but we haven't gotten very far yet really. Another thing that's seen in the cloud app logs regularly at high concurrency is ECONNREFUSED to fh-mbaas: {code} {"name":"basic-logger","hostname":"testing-qed1-mbaas1-live","pid":20085,"level":50,"msg":"FH-MBAAS-CLIENT: Invalid Response Body { body: 'Service unavailable, check service exists and is running ok - { [Error: connect ECONNREFUSED 10.71.11.120:8819]\\n code: \\'ECONNREFUSED\\',\\n errno: \\'ECONNREFUSED\\',\\n syscall: \\'connect\\',\\n address: \\'10.71.11.120\\',\\n port: 8819 }' }","time":"2017-04-04T10:45:50.652Z","src":{"file":"/home/fh/apps/testing-t5smsumn2ninh3k4mrxnmugp-qed1-mbaas1-live/src/node_modules/fh-mbaas-api-metrics/node_modules/fh-mbaas-client/node_modules/fh-logger/lib/request-	1	1	https://issues.jboss.org/browse/RAINCATCH-619
4416	RAINCATCH-619	I'm curious as to the header-setting error, does it mean we have some asyncronous building of express responses?	1	1	https://issues.jboss.org/browse/RAINCATCH-619
4417	RAINCATCH-619	bq. Afterwards we could consider overriding that value to a larger one for network-related mediator topics in the raincatcher codebase. Yeah, that's one possible way. Another could be to have some init of each module, where things like this could be passed in from the app. That way, developers could configure different modules with different timeouts (or other properties), based on their app config file (or env vars that the app consumes and passes values to the relevant modules). bq. I'm curious as to the header-setting error, does it mean we have some asyncronous building of express responses? I'm not really sure what it means. It could be that, or maybe it could be some strange handler somewhere that's listening for some exceptional case and returning (or trying to return) a response even though execution continues as normal somehow also; or it could be that there's a missing {{return}} on a callback somewhere that lead to it being called more than once. In any case, if upgrading to the latest fh-mbaas-api might fix it, I'd rather try that instead of sinking lots of time into it to try to find. :)	0	0	https://issues.jboss.org/browse/RAINCATCH-619
4418	RAINCATCH-619	For the next test, everything was doubled: 800 executions at a peak concurrency of 80, ramping up over 160 seconds. The total time was 72 minutes, running between Thu Apr 27 15:29:25 and Thu Apr 27 16:41:20. There were 125 failures, all of which were timeouts in the load test scripts. To understand what this means, a bit of context from the test scripts is needed: There are a few places in the test scripts where we loop waiting for something for something from the server, checking every 5 minutes. If we haven't gotten what we want after 5 minutes, we fail that execution and move on. [Here's one example|https://github.com/feedhenry-raincatcher/raincatcher-demo-load-testing/blob/master/util/createRecord.js#L13..L30] - after sending a create or update of a record to the server, we wait until the server has told us that it has applied it before moving on. Looking at the times when these failures occurred is useful here: the first was at 17:28, in the final 15 minutes of the test. After this point, only 20 successful runs finished, all finishing at 17:30 (11 minutes before the end). This scenario was predicted in the last point in the previous comment. We should investigate whether this is due to too much load on the server, or bad assumptions on the test clients (are we not waiting long enough? How long is enough?). Some follow-on questons: * Will a soak test run experience these timeouts eventually at a lower concurrency? If we had only run 600 executions here, we wouldn't have seen any failures. * If we increased the timeout to 10 minutes or more, what's the next wall that we'd hit? * Are there any configurable options in the new sync that would improve this? * Would we continue to see the same behaviour if we scaled up the number of cluster workers in a multi CPU dyno, or the number of cloud app instances in an OCP cluster? Here are the grafana (with some lead-in and lead-out) and load-runner graphs, over the duration of the test: !raincatch-619_800_80_160-new-sync-grafana-fullpage.png|thumbnail! !raincatch-619_800_80_160-new-sync-lrgraph-fullpage.png|thumbnail!	0	0	https://issues.jboss.org/browse/RAINCATCH-619

#	Requirement	Questions	Answered	Successfully Answered	URL
4419	RAINCATCH-618	[~vsazel], [~phajidec], [~davidmartin], [~akeating1]: See the attached [^raincatch-618.csv] (which is created from running the collection script against the relevant runs in [^raincatch-618.tar.gz]), and the Grafana screenshot below. I ran this from the server that Aiden uses for the sync load testing server (the raincatcher-demo-load-testing repo is cloned there), then tarred the runs dir and downloaded it and the CSV, then stopped the instance again. Before the test I restarted both the cloud app and the auth service from the RHMAP studio. For the Grafana screenshot, I customized the view based on the times listed in the csv, just with a slight lead-in and lead-out to make them less sharp, and give an idea of what activity was like beforehand. !raincatch-618-fullpage.png|thumbnail! Some notes: * I ran the setup script from load runner with the {{--before}} option. I *think* that this leads to the setup script activity being part of the graph, but I might be wrong. I'm not sure which is better. * We don't have any visibility into the auth service activity during the run. I'll create a follow-on issue to instrument that and use it, and configure Grafana to show it. * Our load test scripts don't do the independent sync of each dataset at each interval: they just do it once, they just send some data to the sync endpoints. I'm not sure if that's important, or how much additional load that would add. From my understanding of how sync is used by raincatcher modules on devices, each dataset is synced every interval to *check if there's any new data on the server*, but when they've got data to *send* (like our load test scripts do), that's outside the regular interval sync. For the time that a script execution is running, that's probably currently fine -- executions should currently be quite quick, I'm not sure if that's good or bad or neither. How does this look? Is there anything clearly missing?	1	1	https://issues.jboss.org/browse/RAINCATCH-618
4420	RAINCATCH-602	I think I might've editted this issue's description back to an old version by mistake, removing an edit by [~davidmartin], if someone got the email with the diff could they revert it?	1	1	https://issues.jboss.org/browse/RAINCATCH-602
4421	RAINCATCH-582	[~phajidec] I'm proposing that we don't bother with this, and instead just restart the cloud app and auth service manually. What do you think? All of the stuff (users and workorders) that gets created by the setup script at least seems to get cleared up by doing that (and, depending on the number of users and workorders, this could be a much quicker way than getting the cloud app to process deletes for everything).	1	1	https://issues.jboss.org/browse/RAINCATCH-582
4422	RAINCATCH-582	bq. We'll be releasing the storage implementations that use mongo for persistency soon, so we won't be able to rely on this behavior How soon is soon? :) bq. We've also implemented a data reset http endpoint to make it easier to reset to the initial seeds, so in the future we can also make the script hit that one. Is that endpoint there now? If not, I propose that we leave the implementation of this script until it (and the persistent storage in MongoDB) is the default. Does that sound reasonable?	0	0	https://issues.jboss.org/browse/RAINCATCH-582
4423	RAINCATCH-575	[~tjackman] Is that still requirement. Can we close it and wait for more clear results of the architecture meetings?	1	0	https://issues.jboss.org/browse/RAINCATCH-575
4424	RAINCATCH-565	So from what I understand, the demo mobile client communicates to the cloud app as follows: - Login at the start - Each module (workflow, workorders, messages, result) performs a sync check at each interval - If there's new data, it will then do a syncRecords to get the records that it's missing - If there's data changed on the client, it always seems to be the *result* module that syncs that back to the server The portal client is similar, except that it can generate data for all of the different modules to sync back (e.g. adding a workflow, sending a message to a user, adding a workorder, etc). I think we could have the load test script cover both: some runs would act as a mobile user, and some would act as a portal user (we could configure the ratio of that using the run number I think). I've a feeling that mobile clients could have more reason to send data to the server, since they could send at each workflow step. I might be making the script too complex by thinking about it like this, I'm not sure. One thing that I'm not really clear on is how long each "client" (mobile or portal) should stay running. Just enough to carry out their own agenda, or should they continue to run after that to continue to receive updates from the server based on the updates from other clients? I guess once an instance of the script finished, it'll be replaced by another one, so there will always be N clients sending and receiving updates.	0	0	https://issues.jboss.org/browse/RAINCATCH-565
4425	RAINCATCH-564	My current understanding is that there would be regular sync requests from devices (e.g. every 5 seconds), but there won't always be new data to send to the cloud, and there may or may not be any new data on the cloud. Is this something that we'd need to have in there, or is that just implicit that we get by having many instances of the script running?	1	1	https://issues.jboss.org/browse/RAINCATCH-564
4426	RAINCATCH-564	[~grdryn] In reality there will be apps syncing with data & without data all the time (at sync frequency) I think having enough instances of the script running at a particular concurrency is enough to mimic the 'worst case' scenario. i.e. everyone is going through the 'happy path' and generating data for sync to be syncing. I don't think there's enough value in including sync traffic without data as it doesn't help answer the question of 'What is the maximum number of active users it can handle?' We can define an active user as someone working their way through a happy path as fast as programatically possible (worst case).	1	1	https://issues.jboss.org/browse/RAINCATCH-564
4427	RAINCATCH-564	How different from a load-testing standpoint it is to actually simulate a happy path at programatic speed as opposed to just hammering the datasets at a ratio and profile that we determine would reflect actual usage? i.e. results suffering a lot of additions and no editions, Workorders having a few orders of magnitude more changes than Workflows, etc. I feel the load-testing client setup might be simpler this way and allow for a higher speed of execution and easier extending as there we get more server-side load-testing points in the project. We could eventually replace educated guessing with actual usage reporting.	0	0	https://issues.jboss.org/browse/RAINCATCH-564
4428	RAINCATCH-540	[~supittma] I think the third option might be the right one. To me, mapping functionality seems like 2 things: - A map display module (raincatcher-map, raincatcher-map-angular?) to display "something" or a list of "something" where "something" is an object with locations, labels etc. It subscribes to "wfm:map:display, arrayOfSomething". - A workorder-map (workorder-map-angular?) module that displays workorders. It publishes a "wfm:workorder:list" topic to get the workorders, plucks out the data it wants to display and then publishes "wfm:map:display, arrayOfSomething". I think both of these modules have value. I can certainly see users wishing to integrate the raincatcher-map module into other stuff (e.g. where are my drivers at any time?).	1	1	https://issues.jboss.org/browse/RAINCATCH-540

#	Requirement	Questions	Answered	Successfully Answered	URL
4429	RAINCATCH-540	So pinching in with what I mentioned before about using typescript to annotate topics. [~supittma]'s third option fits very well with an optional addition of runtime checking of values (i.e. json-schema) to the mediator but also with a interface declaraion in typescript, which could provide additional information through jsdocs. an simple example interface for this topic would be: {code:java} interface IMapPoint { /** Point's latitude on Earth */ lat: number, /** Point's longitude on Earth */ lng: number, /** Optional label for display */ label?: string } {code} However the only way to leverage this through editor support is if we expose 'typed' utility functions: {code:java} class RaincatcherMap { add(point: IMapPoint) { this.mediator.publish('map:add', point); } } {code} These are akin to [action creators|http://redux.js.org/docs/basics/Actions.html#action-creators] in redux If we adopt a file structure/convention for adding these interfaces to the projects we could have some tooling to add them to the READMEs as well, so they can be seen alongside their related topics and used by non-typescript-able users. As a plus these bindings would still be picked up by VSCode's Intellisense even if the user was writing plain javascript.	0	0	https://issues.jboss.org/browse/RAINCATCH-540
4430	RAINCATCH-538	[~ndonnell] In https://github.com/feedhenry-raincatcher/raincatcher-demo-portal/blob/b9241d4a0d9cb5e681cdb6390926ad4169fe7a24/src/app/initialisation/run.js#L75 we call userClient.hasSession(). The userClient service/manager is provided by the fh-wfm-user module and is injected via angular. I think this is an example of where the appropriate module boundary is being respected, and emitting and waiting for a response to an event via the mediator may confuse things rather than make them clearer. This is also similar to places where userClient.list() is bound to a workers property. wdyt?	1	1	https://issues.jboss.org/browse/RAINCATCH-538
4431	RAINCATCH-536	Backing this out a bit I see in https://github.com/feedhenry-raincatcher/raincatcher-message/blob/89a7bea44e4f40082443a522a013c98921ad5248/lib/angular/directive.js the following : {code} .directive('messageForm', function($templateCache, mediator) { return { restrict: 'E' , template: $templateCache.get('wfm-template/message-form.tpl.html') , scope: { message : '=value' , workers: '=' } , controller: function($scope) { var self = this; self.model = angular.copy($scope.message); self.workers = $scope.workers; self.submitted = false; self.done = function(isValid) { self.submitted = true; self.model.receiver = JSON.parse(self.model.receiver); self.model.receiverId = self.model.receiver.id; self.model.status = "unread"; if (isValid) { mediator.publish('wfm:message:created', self.model); } }; } , controllerAs: 'ctrl' }; }) {code} Basically when the form is submitted the "created" event is emitted. However the message is not "created" yet but is submitted for creation. Should creation happen here or should there be another event handler which consumes a submit event and then emits the created event?	1	1	https://issues.jboss.org/browse/RAINCATCH-536
4432	RAINCATCH-518	So you are proposing the UI layer listen for read events published by the mediator instead of querying client services itself? Are the new UI directives implemented by extending a Raincatcher module instead of in the application module? Are there working demos (however trivial) that we can look at to compare/contrast the two approaches? I'm having trouble seeing how this proposed architecture makes the modules less reliant on Angular. Could you provide an example of how this Architecture would work with React or Angular 2?	1	0	https://issues.jboss.org/browse/RAINCATCH-518
4433	RAINCATCH-518	[~ndonnell] - All looks good and makes sense. 100% agree that we should split the core functionality and the presentation logic into separate modules - this would allow people to implement their own views (either in Angular or any other framework) without needing to touch the core functionality. Can you confirm with this approach that the core modules would not have any angular specific logic or code in them? If so, I think we are on to a winner here.	1	0	https://issues.jboss.org/browse/RAINCATCH-518
4434	RAINCATCH-518	[~ndonnell] [~supittma] Is splitting the UI parts into separate modules feasible? In my simplistic view of raincather, would it not make more sense for the Demo App to have the angular implementation in its repo (not in separate modules), requiring the core modules and using them? Having separate UI modules sounds like a very tricky thing to manage as they would all need to know there is a common view into where their things are rendered. There would need to be a UI management layer on top of all the modules that brings them all together as part of the single UI. Each UI module would need to understand where it sits in bigger UI picture so that it fits in appropriately. That bigger UI picture is going to be different depending on the solution which could immediately make UI modules unusable as is. If a module is unusable as is (e.g. needs to be forked, but never PR/merged back?) it probably shouldn't be a module.	1	0	https://issues.jboss.org/browse/RAINCATCH-518
4435	RAINCATCH-516	[~pmdarrow] These look all merged. Can this JIRA be closed?	1	0	https://issues.jboss.org/browse/RAINCATCH-516
4436	RAINCATCH-515	[~pmdarrow] These look all merged. Can this JIRA be closed?	0	0	https://issues.jboss.org/browse/RAINCATCH-515
4437	RAINCATCH-509	[~vchepeli] deployed to http://testing-rkbtsebhwbh6syl2gdsegmip-dev.phaji.skunkhenry.com/?url=http://testing-rkbtsegwgjr336vqus3lbsmu-dev.phaji.skunkhenry.com#/login with fixes from RAINCATCH-510, could you retest please? Project is at http://testing.phaji.skunkhenry.com/#projects/rkbtsecyj2d2cqbp2x5m5uw6/apps, auth service at http://testing.phaji.skunkhenry.com/#services/cpkwf63o24eed7jxl5jughdo/apps/cpkwf67qoepr4bylejy56h67/details.	1	1	https://issues.jboss.org/browse/RAINCATCH-509
4438	RAINCATCH-508	[~ndonnell] [~vchepeli] [~spriadka] Has any progress been made on this issue since the last comment? Anything we can do to help?	1	1	https://issues.jboss.org/browse/RAINCATCH-508
4439	RAINCATCH-489	[~ndonnell] I did some tooling setup and wrote a basic script for the video, but since the auth service didn't use the simple store before (it used pretty much equivalent code but not it as a dependency, see [here|https://github.com/feedhenry-raincatcher/raincatcher-demo-auth/blob/66349c79d085b139eaed55f14b199b4d169bfc61/lib/user.js#L7]) and we're still finishing adding the reset data and simple store to the cloud app+portal, I was thinking about putting this issue on hold and adding RAINCATCH-486 and RAINCATCH-451 as blockers so I can record it after that work is merged. WDYT?	1	1	https://issues.jboss.org/browse/RAINCATCH-489
4440	RAINCATCH-487	[~ndonnell] [~mogrodni] Would it be okay to just have the stores subscribe to `wfm:cloud:{dataset}:reset` like this https://github.com/feedhenry-raincatcher/raincatcher-simple-store/pull/5/files#diff-e208f5ddd2f6398fbbdf6ae5374688c0R24 instead of doing the listing/deleting/creating explicitly in demo-cloud?	1	1	https://issues.jboss.org/browse/RAINCATCH-487
4441	RAINCATCH-487	[~vchepeli] Works fine here !screenshot-1.png|thumbnail! Can you please clear storage in dev tools and/or try it in incognito mode?	1	1	https://issues.jboss.org/browse/RAINCATCH-487
4442	RAINCATCH-475	[~ndonnell] What did we agree to do on the tickets in this epic still targeted for 3.16 on the triage today?	0	0	https://issues.jboss.org/browse/RAINCATCH-475
4443	RAINCATCH-475	[~davidmartin] Is new version of the sync framework changing the way we should handle errors for sync. Do you have some early documentation we can verify before finishing this tickets?	1	1	https://issues.jboss.org/browse/RAINCATCH-475
4444	RAINCATCH-453	[~spriadka] We don't story point sub tasks. Can you remove the story points from this sub task and add to the parent task.	0	0	https://issues.jboss.org/browse/RAINCATCH-453
4445	RAINCATCH-452	[~spriadka] We don't story point sub tasks. Can you remove the story points from this sub task and add to the parent task.	0	0	https://issues.jboss.org/browse/RAINCATCH-452
4446	RAINCATCH-448	[~ndonnell] I removed the subtask for making a video as it's not critical path -- it's now a standalone task. Can you let me know if the ticket is still being worked on or if it's ready for QA? [~vchepeli] has added steps to verify which is leading me to think it is the latter. Leigh	1	1	https://issues.jboss.org/browse/RAINCATCH-448

#	Requirement	Questions	Answered	Successfully Answered	URL
4447	RAINCATCH-426	Awesome, thanks [~tjackman]! If that's PR based, I guess no extra jobs are needed. Do you want me to close this ticket ore do you want it open to keep tracking more vulnerabilities?	1	1	https://issues.jboss.org/browse/RAINCATCH-426
4448	RAINCATCH-416	Instead of implementing security ourselves.. would it make sense to evaluate Keycloak here?	1	1	https://issues.jboss.org/browse/RAINCATCH-416
4449	RAINCATCH-416	[~mwessendorf] this would be a nice opportunity for us to get familiar with the project! Relevant adapter for us would be [keycloak's nodejs adapter|https://keycloak.gitbooks.io/securing-client-applications-guide/content/v/2.4/topics/oidc/nodejs-adapter.html] Or perhaps the [oauth|https://github.com/jaredhanson/passport-oauth] or [oid connect|https://github.com/jaredhanson/passport-openidconnect] strategies for passportjs if we decide to go with it as a more generic client library and provide keycloak as a auth backend alternative. I think one of our main requirements with raincatcher is a low-friction deployment to RHMAP, is keycloak made available easily alongside the platform?	1	1	https://issues.jboss.org/browse/RAINCATCH-416
4450	RAINCATCH-416	Is the plan to use/provide a library for doing that or do we just modify demo-auth module to use best practices?	1	1	https://issues.jboss.org/browse/RAINCATCH-416
4451	RAINCATCH-412	[~spriadka] We don't story point sub tasks. Can you remove the story points from this sub task and add to the parent task.	0	0	https://issues.jboss.org/browse/RAINCATCH-412
4452	RAINCATCH-409	Will there be a crypto library as part of that? Or would be responsibility of each app developer to identify sensitive data?	0	0	https://issues.jboss.org/browse/RAINCATCH-409
4453	RAINCATCH-380	[~ndonnell] Could you check it? Could you set the fixVersion or tell us if it is already solved? Please, can you tell us if it is duplicate and you have another JIRA ?	0	0	https://issues.jboss.org/browse/RAINCATCH-380
4454	RAINCATCH-374	[~phajidec] New spike task blocking the collision handlers. If you are happy, can you mark as ready for story pointing.	0	0	https://issues.jboss.org/browse/RAINCATCH-374
4455	RAINCATCH-364	[~ndonnell] Could you check it? Could you set the fixVersion or tell us if it is already solved? Please, can you tell us if it is duplicate and you have another JIRA ?	0	0	https://issues.jboss.org/browse/RAINCATCH-364
4456	RAINCATCH-363	[~ndonnell] Could you check it? Could you set the fixVersion or tell us if it is already solved? Please, can you tell us if it is duplicate and you have another JIRA ?	0	0	https://issues.jboss.org/browse/RAINCATCH-363
4457	RAINCATCH-362	[~ndonnell] Could you check it? Could you set the fixVersion or tell us if it is already solved? Please, if it is duplicate and you have another JIRA in the public place can you tell us ?	0	0	https://issues.jboss.org/browse/RAINCATCH-362
4458	RAINCATCH-361	[~ndonnell] Could you check it? Could you set the fixVersion or tell us if it is already solved ?	0	0	https://issues.jboss.org/browse/RAINCATCH-361
4459	RAINCATCH-360	[~ndonnell] Could you check it? Could you set the fixVersion or tell us if it is already solved ?	0	0	https://issues.jboss.org/browse/RAINCATCH-360
4460	RAINCATCH-328	[~mogrodni] Could you clarify a bit the {{Things to consider}} section? I can't understand their implications precisely. i.e. - What should be the flow for using prerelease tags (when do we drop them) - How those tags would interact with the cli command - How to avoid using internal npm	0	0	https://issues.jboss.org/browse/RAINCATCH-328
4461	RAINCATCH-303	[~phajidec] Added detail for getting unit tests for the mediator. We should do this before load testing. If you are happy, can you mark as ready for story pointing.	1	1	https://issues.jboss.org/browse/RAINCATCH-303
4462	RAINCATCH-297	h1. Reopened. missing q in package json, its required in {{lib/mediator-service.js}} getting this err when running {code:java} grunt test {code} {code} ~/temp/raincatcher-map ⍰⍰ RAINCATCH-540 ⍰ grunt test Running "mochify:unit" (mochify) task Error: Cannot find module 'q' from '/Users/Mick/temp/raincatcher-map/lib' Error: Exit 1 Warning: `/Users/Mick/.nvm/versions/node/v4.4.7/bin/node /Users/Mick/temp/raincatcher-map/node_modules/grunt-mochify/node_modules/mochify/bin/cmd.js --reporter spec --colors test/mapController-spec.js test/mapDirective-spec.js test/mapMediatorService-spec.js test/mapService-spec.js` failed with code 1 Use --force to continue. Aborted due to warnings. {code}	1	1	https://issues.jboss.org/browse/RAINCATCH-297
4463	RAINCATCH-219	[~ndonnell] Could you add some details on How this should be achieved? i.e. - Passing in collision handlers from the raincatcher-sync API or also from raincatcher-workorder and the other consumers for passing in solution? - Brief sketch of how the API should be, or links to examples (searched for the example we saw during workshop but couldn't find them atm). - Suggestions for verification setup (unit testing / scenario for manually testing in a cluster) Can you also confirm if this only relates to allowing user-defined functions or also adding some default implementation?	1	1	https://issues.jboss.org/browse/RAINCATCH-219
4464	RAINCATCH-219	[~mogrodni] [~jameelb] just wondering can you give an update on this ticket? It looks to be holding steady for the past few days and I understand you had some blockers in it. Can you submit a PR so that we have visibility on it and can get more eyes on it if necessary. I'm going to suggest looking at putting in some sub tasks as well.	1	1	https://issues.jboss.org/browse/RAINCATCH-219
4465	RAINCATCH-205	[~vchepeli] If you have scope this sprint, can you take a look at this spike to see if there is a good way to load test the demo solution automatically.. I think it would be scenario-based but I think a spike of 3 - 4 days is required to see how this could be done. [~adamsaleh] is working on spiking how load / performance testing of the raincatcher modules in isolation could be done. It might be good to work together on this to get up to speed on Raincatcher. [~kpiwko] Thoughts? Any questions, let me know.	0	0	https://issues.jboss.org/browse/RAINCATCH-205
4466	RAINCATCH-187	[~phajidec] I just wanted to double check the goal of this task... Reading the header it says add "grunt tasks", but should this perhaps say create a grunt plugin? With a plugin the task whould become part of the projects own tasks, compared to just adding a grunt task to raincatcher-template-build which I think would still mean that we'd have to run them using either an npm script (like we currently do) or using something like: {quote}$ grunt --gruntfile node_modules/fh-wfm-template-build/Gruntfile.js build:test{quote} Let me know if I've misunderstood anything here. Thanks	1	1	https://issues.jboss.org/browse/RAINCATCH-187
4467	RAINCATCH-187	h2. Verification steps To test this you can pick an existing Raincatcher module that currently uses the npm scripts. For example raincatcher-user. h3. Install the plugin From the raincatcher-user directory run: {code} $ npm install --save-dev fh-wfm-template-build@0.1.1 --registry http://npm.skunkhenry.com $ npm install {code} Add the loading of the task to Gruntfile.js: {code} grunt.loadNpmTasks('fh-wfm-template-build'); {code} h3. Try the commands Try the _build_ task: {code} $ grunt wfmTemplate:build --debug --module=wfm.user.directives {code} Try the _watch_ task: {code} $ grunt wfmTemplate:watch --debug --module=wfm.user.directives {code} Touch one of the templates to trigger a rebuild: {code} $ touch lib/angular/template/worker-form.tpl.html {code} h3. Configure the plugin Try the above commands but this time with the configuration in Gruntfile.js (so no need to include the --module=<module name> option) {code} module.exports = function(grunt) { 'use strict'; grunt.initConfig({ eslint: { src: ["lib/**/*.js"] }, wfmTemplate: { module: 'wfm.user.directives', } }); grunt.loadNpmTasks('fh-wfm-template-build'); grunt.loadNpmTasks("grunt-eslint"); grunt.registerTask('default', ['eslint']); }; {code}	0	0	https://issues.jboss.org/browse/RAINCATCH-187
4468	RAINCATCH-152	[~ndonnell] This looks like it should be broken into at least two issues, do you agree? 13 points came from big scope and unknowns.	0	0	https://issues.jboss.org/browse/RAINCATCH-152

#	Requirement	Questions	Answered	Successfully Answered	URL
4469	RAINCATCH-143	[~ndonnell] [~wtrocki] [~vchepeli] Coveralls is not completely free. You have to pay monthly $5, but it can be used to track unlimited count of open-source repos. Do we have some sort of account associated with Coveralls? (Base feedhenry organization has, see: https://github.com/feedhenry/fh-ios-sdk and click on coverage badge)	1	1	https://issues.jboss.org/browse/RAINCATCH-143
4470	RAINCATCH-33	[~spriadka], I've made some comments on the review : https://github.com/feedhenry-raincatcher/raincatcher-appform/pull/29. What is the setup to test this? If you tell me which versions/branches of the other modules I need to use I can run it myself. Otherwise is this deployed to a cluster?	1	1	https://issues.jboss.org/browse/RAINCATCH-33
4471	RAINCATCH-33	[~spriadka], [~emilioicai] can you guys add details for ticket validation: cluster steps to validate? Thanks!	1	0	https://issues.jboss.org/browse/RAINCATCH-33
4472	RAINCATCH-33	[~spriadka] You would need to point your local setup to form that has checkbox. See your grunt configuration (FH_MBAAS_HOST etc.) Once form is updated make sure is deployed to environment you using in config. Last part would be to make sure that your workflow have your form reference id. To check that go to demo-cloud application and verify/change id's: https://github.com/feedhenry-raincatcher/raincatcher-demo-cloud/blob/master/lib/app/workflow.js#L18-L19 I do not see this process documented anywhere so maybe we can pull out this into: https://github.com/feedhenry-raincatcher/raincatcher-documentation/blob/master/running-locally.adoc WDYT?	0	0	https://issues.jboss.org/browse/RAINCATCH-33
4473	RAINCATCH-15	Advice from [~bgallagh] via e-mail: {quote} What I've done in the past is used a single dataset (workOrders) with query params to filter on the logged in user. The dataset structure looks like this: workOrder : { clientData: {}, // Data created on the client and never updated by customer cloudData: {}, // Data sent from customer - read only on client shared: {} // Data that can be updated on the client and from the customer } This makes handling collisions easier. I needed to write custom logic to 'merge' the shared properties on the client and cloud. Not sure if that helps at all. I think there is a balance, you need to be careful to not have too many datasets. Each one will have its own sync loop which could result in a lot of parallel http connections on device and in node. Are you duplicating data across the smaller datasets? If so, consider that all datasets are held in memory. There is a WFM style app coming shortly that requires supporting 1500 users. So, here is a basic calculation: 1500 Users * 60 second sync * 1 DataSet => 25 Http requests per second (minimum) 1500 Users * 60 second sync * 2 DataSet => 50 Http requests per second (minimum) 1500 Users * 60 second sync * 3 DataSet => 75 Http requests per second (minimum) etc etc {quote}	0	0	https://issues.jboss.org/browse/RAINCATCH-15
4474	RAINCATCH-1	Should this renaming also include the angular module names as well perhaps? For example, [lib/angular/sync-ng|https://github.com/feedhenry-raincatcher/raincatcher-sync/blob/a6fb3b6f61ece4132403375cf5ae2bdabc99e206/lib/angular/sync-ng.js#L7]	1	1	https://issues.jboss.org/browse/RAINCATCH-1
4475	FORGE-2752	[~claprun], what are the use cases for those examples? A feature that may help you is that Forge allows you to listen to command execution by implementing a CommandExecutionListener: https://github.com/forge/core/blob/master/ui/api/src/main/java/org/jboss/forge/addon/ui/command/CommandExecutionListener.java	1	0	https://issues.jboss.org/browse/FORGE-2752
4476	FORGE-2665	we could yeah - I"d be happy just doing the addon-remove / addon-Install TBH to start with ;) I guess JRebel would be much faster; but could end up with jboss modules issues possibly?	1	1	https://issues.jboss.org/browse/FORGE-2665
4477	FORGE-2665	Some questions: - What would be a good syntax for this command? (eg. {{addon-watch --coordinate a:b}}?) - How would you stop this command from executing? (another command? Specify a number of times after execution?)	1	1	https://issues.jboss.org/browse/FORGE-2665
4478	FORGE-2653	Totally agree. Had this in my todo list but haven't created a JIRA for it. I struggle a bit when trying to define it. Do you have any suggestions? Forge is a project management tool that allows you to automate repetititve, boring development tasks that you would probably google it, like setting up JPA in a Java eE 7 application using Maven for example (what files should be changed? What should go into those files?), . It sports a set of addons for the most used technologies nowadays and integrates with your favorite IDE to provide a seamless experience.	0	0	https://issues.jboss.org/browse/FORGE-2653
4479	FORGE-2645	Would it help to have two new options: {{-- idPropertyName}} (default: id) and {{-- versionPropertyName}} (default: version) ?	1	1	https://issues.jboss.org/browse/FORGE-2645
4480	FORGE-2637	Hi [~remmeier], Thank you for your feedback. They are really important. I am thinking of working on a JBoss Forge Users Guide and publishing it on https://www.gitbook.com (with a Github repository to allow easier contributions). I'd love to know what the community thinks about the beginner/intermediate/advanced labels. Would you mind posting a message in our user forums and ask for feedback(https://developer.jboss.org/en/forge) ? Once again, thank you very much!	1	1	https://issues.jboss.org/browse/FORGE-2637
4481	FORGE-2637	I noticed that one has been renamed from addon to forge-addon in 3.0.The websites mentions that forge 3.0 has been released. Is there an eclipse update site? because the jobs developer studio still has the old one. Some instructions in this regard would be helpful on the website. regarding command renaming: a deprecation mechanism (allow old command to be used for one major release, give warning on console) and stronger validation of input parameters could be helpful. In this example I used the command and ended up with just a war project.	1	1	https://issues.jboss.org/browse/FORGE-2637
4482	FORGE-2637	thx for the help. Got the basics of first plugin up and running (Apache Olingo). regarding documentation, a gradle example would be great. I installed it for both 2.x and 3.x. But did not manage to get it up and running. The 2.x version in the JBoss Develop Studio showed me only "Maven" as buildSystem option. And the following command gave me a pom without any errors/warnings: project-new --named demo --buildSystem Gradle What is the general status of the gradle support? I noticed that for 3.x it is still in Alpha stage. At my company, all new projects are soon supposed to use Gradle. And is there a possiblity to display the available parameters of a command from the command line? like "--help"? Similar to the wizard in the Dev Environments?	1	1	https://issues.jboss.org/browse/FORGE-2637

#	Requirement	Questions	Answered	Successfully Answered	URL
4483	FORGE-2615	to have the ui elements and the execute [15:11:10] <davsclaus> as today [15:11:17] <davsclaus> but being able to mark some of the uis [15:11:18] ggrzybek is now known as ggrzybek_bbl [15:11:20] <davsclaus> to be on page 2, 3 etc [15:11:29] <davsclaus> and then forge automatic turn that into a wizard with next/back [15:11:47] <davsclaus> so you dont have to move it into new UIStep classes and whatnot [15:11:51] <davsclaus> but can keep it all in 1 file [15:11:59] <@gastaldi> hmm [15:12:02] <davsclaus> for those simpler commands where [15:12:10] <davsclaus> you start with 3-4 ui\s [15:12:17] <davsclaus> and then add some more and then you have 8 or so [15:12:26] <davsclaus> and find out it makes the ui better if it was 2 pages [15:12:26] <@gastaldi> right, I think we can have a command that splits that into multiple UISteps internally [15:12:32] <davsclaus> yep [15:13:25] <@gastaldi> or a fluent language builder [15:13:49] <davsclaus> yeah [15:13:55] <@gastaldi> something like: Command.named("").having(UIInput, UIInput) [15:14:15] <davsclaus> builder.add(one).add(to).next().add(three).add(four).next().add(five); [15:14:27] <davsclaus> in that initializeUI method [15:14:31] <davsclaus> where you add the ui elements [15:14:34] <@gastaldi> ah you mean the UIBuilder [15:14:56] <davsclaus> yeah [15:14:58] <@gastaldi> hmm [15:15:13] <davsclaus> the pain today is to break up your command into N steps in the wizard [15:15:18] <@gastaldi> I see [15:15:23] <davsclaus> for simpler uis that is just straight line [15:15:37] <davsclaus> i can understand the complex ones you need more code and to code in the next() etc [15:15:44] <davsclaus> but for 90% ui's [15:15:54] <davsclaus> its N options you want to show the users [15:15:56] <davsclaus> but not all at once [15:16:02] <davsclaus> to not overload/confuse the user [15:16:07] <davsclaus> but spread out in N pages [15:16:21] <davsclaus> having all that in 1 java file [15:16:27] <@gastaldi> what about the conditions of showing an Input on a Nth page? [15:16:27] <davsclaus> makes it easier to code and "overview" [15:16:46] <davsclaus> from forge pov it is as if all the options are on the same page [15:16:56] <@gastaldi> right, I see [15:16:59] <davsclaus> but imageing people dont have big monitors [15:17:15] <davsclaus> so somehow it automatic overflow [15:17:27] <davsclaus> that can work for the simpler cases [15:17:38] <davsclaus> if you need conditional and branches, or whatnot [15:17:41] <davsclaus> then you can code it [15:18:30] <@gastaldi> by overflow you mean automatically adding a vertical scroll bar to the Wizard panel? [15:19:02] <davsclaus> if you add to many options to a command then today it will show a dialog [15:19:06] <davsclaus> that cannot fit the screen [15:19:10] <davsclaus> an you cannot see all the options [15:19:15] <davsclaus> and the next / finish button either [15:19:29] <davsclaus> that is okay if you keep to max 20-30 options etc [15:19:36] <davsclaus> but that happens seldom [15:19:40] <@gastaldi> in any IDE? or is that just in Eclipse? [15:19:45] <davsclaus> any of them [15:19:52] <@gastaldi> gotcha, that looks like a bug [15:20:06] <davsclaus> i dont think anyone builds ui like that [15:20:15] <davsclaus> we only see it in those camel wizards with camel endpoints [15:20:22] <davsclaus> where you have maybe 90+ options [15:20:37] <davsclaus> and we do not want 90 on one page either [15:20:45] <davsclaus> so that is why we have that dynamic wizard with the camel [15:20:46] <@gastaldi> right. What if you could split these options into fine grained commands? [15:20:53] <davsclaus> where we compute how many pages and all that [15:21:07] <davsclaus> i cannot its 100% dynamic [15:21:12] <@gastaldi> Right [15:21:23] <davsclaus> that is okay, its all the other commands i talk about really [15:21:31] <davsclaus> we often have commands with 4-10 options [15:21:36] <davsclaus> but would like 2 or 3 pages [15:21:42] <davsclaus> instead of 1 [15:21:50] <davsclaus> and today i have to code a fair bit to do that [15:21:56] <@gastaldi> I see [15:21:58] <davsclaus> hence why i would love to keep it in 1 java file [15:22:10] <@gastaldi> right [15:22:15] <davsclaus> and just tell forge that these X inputs = page 1, these Y inputs = page 2 [15:22:34] <@gastaldi> or have Forge automatically calculate that for you [15:22:36] <davsclaus> and optionally be able to set the title for each page as	1	0	https://issues.jboss.org/browse/FORGE-2615
4484	FORGE-2588	so one really simple workaround for this issue is to remove the pom.xml and src folder before the ArchtypeHelper processes the archetype jar; as for java / maven projects its going to replace whatever is there with the pom.xml thats inside the archetype anyway and will generate all the required java source folders with the correct package names or files from the archetype etc. Then if there's no pom.xml or src folders in the archetype jar, things just work beautifully without having to make major changes to ProjectType / NewProjectWizardImpl.java to defer decisions on what facets are inside a MavenArchetypeProjectType.java until after the archetype has been generated. Does this simple patch seem OK to you? It only affects projects created using MavenArchetypeProjectType. I did ponder about adding a custom "non-java" marker to the archetype to restrict the scope of this fix further; but thats harder right now as ArchetypeHelper is called with an InputStream of the archetype jar so its harder to pre-process the jar first to look for a magic file to delete pom.xml and src before processing it again. If you're happy with the above I can submit a PR?	1	1	https://issues.jboss.org/browse/FORGE-2588
4485	FORGE-2585	How about having a method like: {{T getSelectedRegionFor(Resource)}} and it could return a child element of the given resource ({{MethodSource}}, {{FieldSource}} when the resource is a {{JavaResource}} or {{Node}} if element is a {{XMLResource}} or even a {{String}} with the selected contents)? I am thinking on how useful would this API change be to make it easier for your code	1	1	https://issues.jboss.org/browse/FORGE-2585
4486	FORGE-2580	Hi [~davsclaus], Can you test against the latest master? I'll release Forge 3.0.0.Beta3 today if that works for you. Thanks!	0	0	https://issues.jboss.org/browse/FORGE-2580
4487	FORGE-2565	How does forge download from maven, do it uses a single thread? Maven itself used to be very slow as well, but they are now downloading concurrently. Maybe forge can do the same?	1	0	https://issues.jboss.org/browse/FORGE-2565
4488	FORGE-2553	In IDEA, even though the tooltip is being set in the component, it is not rendered (we use IDEA's {{TextFieldWithAutoCompletion}} class). I implemented the tooltip on labels and it's working fine now. Should be available on the plugin's 1.1.3.Final version.	0	0	https://issues.jboss.org/browse/FORGE-2553

#	Requirement	Questions	Answered	Successfully Answered	URL
4489	FORGE-2548	Index: 2, Size: 2 at java.util.ArrayList.rangeCheck(ArrayList.java:653) at java.util.ArrayList.get(ArrayList.java:429) at org.jboss.forge.addon.ui.impl.controller.WizardCommandControllerImpl.getCurrentEntry(WizardCommandControllerImpl.java:491) at org.jboss.forge.addon.ui.impl.controller.WizardCommandControllerImpl.getCurrentController(WizardCommandControllerImpl.java:502) at org.jboss.forge.addon.ui.impl.controller.WizardCommandControllerImpl.getInputs(WizardCommandControllerImpl.java:244) at org.jboss.forge.addon.ui.impl.controller.NoUIWizardControllerDecorator.getInputs(NoUIWizardControllerDecorator.java:189) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback$2.call(ClassLoaderAdapterCallback.java:123) at org.jboss.forge.furnace.util.ClassLoaders.executeIn(ClassLoaders.java:42) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback.invoke(ClassLoaderAdapterCallback.java:96) at org.jboss.forge.addon.ui.controller.CommandController_$$_javassist_c2571eee-14a3-454e-af9c-18361c1e0523.getInputs(CommandController_$$_javassist_c2571eee-14a3-454e-af9c-18361c1e0523.java) at org.jboss.forge.plugin.idea.ui.wizard.ForgeWizardStep.prepare(ForgeWizardStep.java:68) at com.intellij.ui.wizard.WizardModel.b(WizardModel.java:173) at com.intellij.ui.wizard.WizardModel.a(WizardModel.java:91) at com.intellij.ui.wizard.WizardModel.next(WizardModel.java:76) at com.intellij.ui.wizard.WizardAction$Next.actionPerformed(WizardAction.java:46) at javax.swing.AbstractButton.fireActionPerformed(AbstractButton.java:2022) at javax.swing.AbstractButton$Handler.actionPerformed(AbstractButton.java:2348) at javax.swing.DefaultButtonModel.fireActionPerformed(DefaultButtonModel.java:402) at javax.swing.DefaultButtonModel.setPressed(DefaultButtonModel.java:259) at javax.swing.plaf.basic.BasicButtonListener.mouseReleased(BasicButtonListener.java:252) at java.awt.Component.processMouseEvent(Component.java:6535) at javax.swing.JComponent.processMouseEvent(JComponent.java:3324) at java.awt.Component.processEvent(Component.java:6300) at java.awt.Container.processEvent(Container.java:2236) at java.awt.Component.dispatchEventImpl(Component.java:4891) at java.awt.Container.dispatchEventImpl(Container.java:2294) at java.awt.Component.dispatchEvent(Component.java:4713) at java.awt.LightweightDispatcher.retargetMouseEvent(Container.java:4888) at java.awt.LightweightDispatcher.processMouseEvent(Container.java:4525) at java.awt.LightweightDispatcher.dispatchEvent(Container.java:4466) at java.awt.Container.dispatchEventImpl(Container.java:2280) at java.awt.Window.dispatchEventImpl(Window.java:2750) at java.awt.Component.dispatchEvent(Component.java:4713) at java.awt.EventQueue.dispatchEventImpl(EventQueue.java:758) at java.awt.EventQueue.access$500(EventQueue.java:97) at java.awt.EventQueue$3.run(EventQueue.java:709) at java.awt.EventQueue$3.run(EventQueue.java:703) at java.security.AccessController.doPrivileged(Native Method) at java.security.ProtectionDomain$JavaSecurityAccessImpl.doIntersectionPrivilege(ProtectionDomain.java:76) at java.security.ProtectionDomain$JavaSecurityAccessImpl.doIntersectionPrivilege(ProtectionDomain.java:86) at java.awt.EventQueue$4.run(EventQueue.java:731) at	1	1	https://issues.jboss.org/browse/FORGE-2548
4490	FORGE-2548	Hi Claus, This sure looks annoying. How can I reproduce this error? Thanks	1	1	https://issues.jboss.org/browse/FORGE-2548
4491	FORGE-2526	Hi Claus, Would adding a {{UISelectOne.getSelectedIndex()}} method be enough to fulfill this requirement? I didn't understand the {quote}The same thing would be nice on the change listener, where its applicable{quote} part.	1	1	https://issues.jboss.org/browse/FORGE-2526
4492	FORGE-2508	Did it work after changing it to JDK 8? Check for the Java executable that executes your IDE. What UI are you using? (Eclipse, Shell)? In case you are using the shell (from the standalone ZIP) change the forge.bat to display the JAVA_HOME used to start Forge.	1	1	https://issues.jboss.org/browse/FORGE-2508
4493	FORGE-2492	Just added a {{org.apache.maven.log_transfer}} system variable in Furnace 2.22.4-SNAPSHOT. See http://github.com/forge/furnace/commit/4f61059ec Can you test with this version?	1	1	https://issues.jboss.org/browse/FORGE-2492
4494	FORGE-2490	How can we solve this issue? Can you provide a pull-request with the necessary fix? Thanks	0	0	https://issues.jboss.org/browse/FORGE-2490
4495	FORGE-2484	[~agoncal], could you review the PR?	1	1	https://issues.jboss.org/browse/FORGE-2484
4496	FORGE-2484	[~gastaldi] You should rename the parameter {{enumClass}} with {{targetClass}} : {code} [MyEnum.java]$ java-add-enum-const TAB --enumClass --named {code} I thought deprecated command would not appear when having code completion. It's important to state that the command is deprecated and will soon disappear (like in scripting), but should not be displayed in code completion. WDYT ? : {code} [proj]$ java-new-enum --named MyEnum [MyEnum.java]$ java-new-enum TAB java-new-enum java-new-enum-const [MyEnum.java]$ java-add-enum TAB java-add-enum-const [MyEnum.java]$ java-new-enum-const TEST ***WARNING*** The command 'Java: New Enum Const' is deprecated and will be removed in future versions. Use 'java-add-enum-const --named' instead ***SUCCESS*** Enum constants [TEST] created in org.proj.MyEnum {code} You mentioned the possibility to add several enums in a single command. I would be nice to still have it actually. This is not the case : {code} [MyEnum.java]$ java-add-enum-const --named VISA ***SUCCESS*** Enum constant 'VISA' added in org.myenum.MyEnum [MyEnum.java]$ java-add-enum-const --named VISA1 VISA2 VISA3 ***SUCCESS*** Enum constant 'VISA1' added in org.myenum.MyEnum // Only VISA1 is created {code} We have this behavior with methods in JPA entity listeners where we use plural : {{lifecycles}} instead of {{lifecycle}}. If we follow this naming convention, that would be like renaming all the commands {{name}} and {{names}} instead of {{named}}. So we could have : {code} $ java-add-enum-const --name VISA1 $ java-add-enum-const --names VISA1 VISA2 VISA3 $ jpa-new-entity-listener --lifecycles PRE_PERSIST PRE_REMOVE PRE_UPDATE {code} Too much refactoring :o(	1	1	https://issues.jboss.org/browse/FORGE-2484

#	Requirement	Questions	Answered	Successfully Answered	URL
4497	FORGE-2469	agoncal So the question of this JIRA is "how do I write a proper add command" gastaldi Right. We might need an AbstractAddxxx class agoncal (targetClass being the real pain) agoncal That's what I thought gastaldi Or we could have a class that creates and sets up these inputs properly gastaldi Something like: Inputs.createTargetClass() agoncal While we still use inheritance vs aggregation, I was thinking of having AbstractCDICommand at the top, and then AbstractCDINewCommand and AbstractCDIAddCommand agoncal yes, that goes back to the main idea of aggregating parameters gastaldi Yeah agoncal But is there a possible quick win ? We don't break everything (too expensive), keep on with inheritance, and get a nice pattern that can be reused in most of the add commands gastaldi Thats a good question gastaldi I'd like to minimize the number of inherited classes agoncal yes agoncal But I realize that most of the AbstractCDICommand, AbstractFacesCommand.... are good for new commands, not add gastaldi what if we moved these common inputs to another class and @Inject them in the add commands gastaldi I mean in an encapsulating class gastaldi @Inject CommonAddInputs gastaldi food for thought :) agoncal I would like it to be even more flexible : agoncal @Inject TargetPackage agoncal @Inject TargetClass gastaldi ah-ha agoncal @Inject Named gastaldi that's interesting agoncal This way you create a command that fits your needs gastaldi so what if we could come with an API that allows you to aggregate this input? gastaldi hmmm or maybe: agoncal This way we could even keep the inheritance (eg. AbstractCDICommand) and use these @Inject gastaldi @Produces public TargetPackage produceTargetPackage() { ...} gastaldi public interface TargetPackage extends UIInput<String> agoncal Also remember that targetClass can be customized agoncal example : agoncal java-add-annotation : target class is any class gastaldi sure, as it is an UIInput, you can call the methods you'd like agoncal cdi-add-producer : target class is most of the classes, except Entities gastaldi that would be in the initializeUI of the commands afaik gastaldi or we could have a factory class that could accept parameters agoncal Would you have time to give it a try in the coming days ? If yes and we find a nice pattern, then I could apply it to the Java EE addon and start developing new commands gastaldi I can try, sure gastaldi that would be awesome agoncal Yes, that would allow me to kiil a few "add" commands that are still waiting to be developped ;o)	1	1	https://issues.jboss.org/browse/FORGE-2469
4498	FORGE-2460	Have you tried calling {{org.jboss.forge.addon.ui.command.CommandFactory.getCommands()}} when the server starts? You need to lookup {{CommandFactory}} through the {{AddonRegistry}}.	1	0	https://issues.jboss.org/browse/FORGE-2460
4499	FORGE-2460	Interesting. Have you had a chance to profile to see where the bottleneck is? Last time I checked, it was the Maven integration's fault (It is a bit slow when accessed the first time).	1	0	https://issues.jboss.org/browse/FORGE-2460
4500	FORGE-2398	Thanks, George! :) When is it going to be available in Forge core?	1	0	https://issues.jboss.org/browse/FORGE-2398
4501	FORGE-2362	It already does for addon projects. Do you mean other projects as well?	1	1	https://issues.jboss.org/browse/FORGE-2362
4502	FORGE-2362	Yes. Each time you create a new project with {{project-new}} it should generate a basic readme (I think it's a good practice to have) with at least some technical information, such as the name of the project and the type (war, ear). Another idea is, each time a Facet is installed, an event is sent (we've talked about this for the Swarm addon ;o) and a new line is added in the readme. We could get some inspiration from [The Perfect Readme|http://fr.slideshare.net/netguru/perfect-project-read-me] and generate a structure as follow, with some information we can automatically add, but also leave blanks for the developper to add : {code} = ${projectName} == General information * ${projectName} * // blah blah description == Technologies * ${language} // this would be mainly Java * ${Database} // we can get this from JPA setup command * ${applicationServer} // We can get this from javaee setup command * ${List of installed facet} == Setup * ${buildSystem}} // Maven or gradle * ${howToBuild}} // If it's Maven then we can say "use mvn clean install".... * // blah blah about the config files (if any) passwords (if any) == Development * // blah blah Additional configuration to be able to develop * // blah blah How to use in your IDE == Tests * // blah blah How to test the code * ${Arquillian} // we know that from Arquillian setup command * ${Junit} or ${testNG} // we know that from arquillian-setup command == Common problems/bugs * // blah blah Solution to common problems == Access to staging/beta/production * // Git repository, branches, tags * // blah blah How can I access (user/passord) == Notes * ${codeStyle} // we can get this information from the code formatter * // blah blah Any note {code}	1	0	https://issues.jboss.org/browse/FORGE-2362
4503	FORGE-2331	agoncal Basically, when you write a command, you need attributes : agoncal named (for most of the commands), targetPackage (for the new commands), targetClass (for the add commands)... agoncal And then you also end up "duplicating" attributes : example the "qualifiers" are used in several CDI commands gastaldi right agoncal So the idea was giving up inheritance (AsbtractCDINewCommand, AbstractJPACommand....) in favor of aggragation agoncal And then gastaldi came with an idea agoncal gastaldi your turn ;o) gastaldi my idea was to introduce some classes with the common inputs containing @Inject UIInputxxx and then @Inject the aggregate class in the command gastaldi something like @Inject DefaultCDIInputs inputs; gastaldi and then in the initializeUI you could call inputs.initializeUI(builder) gastaldi and fetch the values agoncal gastaldi But also more specific ones ? Such as @Inject DefaultCDIQualifier, @Inject DefaultJPAField... ? gastaldi that could work too gastaldi hm, maybe we could introduce a "input-only" concept here gastaldi a class that has only injected inputs gastaldi that would save some memory, by lazy-injecting them gastaldi right now, every @Inject in the command is resolved every time a command is executed agoncal And one suggestion was to start with the "qualifiers" attribute (used in several new CDI commands) gastaldi agoncal, what do you mean by that, specifically? agoncal gastaldi Well, if we could start introducing the concept after this release, and if yes, the "qualifiers" attribute could be a good start gastaldi you mean having a separate class for the inputs? agoncal gastaldi Yes gastaldi yeah, we can think about it in the next version. Still needs some thought to become a pattern agoncal gastaldi Ok, let's keep on exchanging on the subject soro I didn't understand this yet. We have some issue of some document to explain about it or will be discussed in future? soro about this "qualifiers" agoncal soro "qualifiers" is just an example agoncal soro Basically, how do you aggregate common attributes into a command soro agoncal, ah... ok! soro nice!	1	1	https://issues.jboss.org/browse/FORGE-2331
4504	FORGE-2316	@Max why not rely on the {{pom.xml}} ? If you have {{javax.javaee-api}} version 7, then you know it's a Java EE 7 stack. Would you need more ?	1	1	https://issues.jboss.org/browse/FORGE-2316
4505	FORGE-2316	just a comment for my clarification - I thought Stacks would be something that managed one or more facets ? Is that what Stack does but is still seen itself as a Facet ? Could one have Stack's of Stack's then ?	1	1	https://issues.jboss.org/browse/FORGE-2316

#	Requirement	Questions	Answered	Successfully Answered	URL
4506	FORGE-2316	Not sure about the naming convention, but choices are mostly in uppercase, shouldn't it be JAVA_EE_7, JAVA_EE_6 : {code} [test]$ project-new --named MyProj --stack java-ee-6 java-ee-7 none {code} Talking about naming again, if the default choice is {{none}}, then maybe rename {{org.jboss.forge.addon.projects.stacks.NullStack}} to {{StackNone}} or {{NoneStack}} (or {{NoStack}}) When I start Forge, and press TAB (not even created a project) I get all these commands that shouldn't be appear : {{constraint-setup, ejb-setup, jms-setup}} What about filtering choices ? This can be tricky, but would help users. For example, if you setup a JAVA EE 7 stack, having GlassFish 3 and JBoss EAP 6 doesn't make sense {code} [test]$ jpa-setup --container CUSTOM_JDBC GLASSFISH_3 JBOSS_AS6 JBOSS_EAP6 NON_JTA TOMEE WILDFLY CUSTOM_JTA GLASSFISH_4 JBOSS_AS7 JBOSS_EAP7 SAP_HCP WEBLOGIC_12C {code} In the {{StackFacet}} I really don't get the priority. If it's just for ordering, name is fine. BTW, what is the use case of overriding the {{supports}} method ?	1	1	https://issues.jboss.org/browse/FORGE-2316
4507	FORGE-2315	So, what do you think better? Create a move or change the renameTo? For me: create move +1	1	1	https://issues.jboss.org/browse/FORGE-2315
4508	FORGE-2283	Ok, sounds good. Where can I find the latest release offline zip file?	0	0	https://issues.jboss.org/browse/FORGE-2283
4509	FORGE-2279	Done ? For all the parameters ? BTW, is there a way to list all the parameters of all the commands (something like {{command-list --show-parameters}}). It would be helpful to review all the naming conventions	1	1	https://issues.jboss.org/browse/FORGE-2279
4510	FORGE-2268	Is this something that'd be enabled in the CLI provider itself; or something we could add as an option inside the camel commands? If the latter, any pointers on examples I could use to try it out?	0	0	https://issues.jboss.org/browse/FORGE-2268
4511	FORGE-2246	Hum... So shall we "duplicate" the code of {{AbstractJavaSourceCommand}} (the target package, named and overwrite) into {{AbstractJavaEECommand}} ? Both {{AbstractJavaSourceCommand}} and {{AbstractJavaEECommand}} extend from {{AbstractProjectCommand}}, so {{AbstractJavaEECommand}} cannot extend from {{AbstractJavaSourceCommand}} (except we refactor the entire graph}. WDYT ?	1	1	https://issues.jboss.org/browse/FORGE-2246
4512	FORGE-2230	How much slow is it? Will it severely impact performance? I don't think this is really an issue because we'll be 1) Avoiding unnecessary NullPointerExceptions: (Eg: {{coordinate.getPackaging().matches("?ar")}}); 2) Promoting better code legibility (Eg: {{packaging.isEmpty()}} is better read than {{packaging != null}}) I am still not convinced we should return null, but an empty String instead. IMHO, statements like {{coordinate.setPackaging(null)}} should throw an {{IllegalArgumentException}} because {{null}} is not a valid packaging type (unless we want to assume JAR as the default packaging)	1	1	https://issues.jboss.org/browse/FORGE-2230
4513	FORGE-2224	What do you mean by "scripts" that already use this command ? Are we talking about Forge shell scripts ? A few months ago there was plenty of command renaming (see [FORGE-1387]). I'm sure it broke plenty of scripts (I am a big fan of Forge Shell scripts and I spend always a bit of time to update them)	1	1	https://issues.jboss.org/browse/FORGE-2224
4514	FORGE-2106	This facet is implemented in another addon: https://github.com/gastaldi/reflections-addon Should we bring it to core?	1	1	https://issues.jboss.org/browse/FORGE-2106
4515	FORGE-2102	Can this issue be closed? I guess the remaining tasks can be closed as well	1	1	https://issues.jboss.org/browse/FORGE-2102
4516	FORGE-2050	Strange thing is that addon is listed when running forge --list: ben@ben-K55VD:~/Development/playground$ forge --list /home/ben/Development/forge-distribution-2.12.0.Final/addons: org.jboss.forge.addon:addon-manager-spi,2.12.0.Final org.jboss.forge.addon:addon-manager,2.12.0.Final org.jboss.forge.addon:dev-tools-java,2.12.0.Final org.jboss.forge.addon:shell-spi,2.12.0.Final org.jboss.forge.addon:shell,2.12.0.Final org.jboss.forge.addon:scaffold-faces,2.12.0.Final org.jboss.forge.addon:parser-json,2.12.0.Final org.jboss.forge.addon:maven,2.12.0.Final org.jboss.forge.addon:addons,2.12.0.Final org.jboss.forge.addon:git,2.12.0.Final org.jboss.forge.addon:scaffold-spi,2.12.0.Final org.jboss.forge.addon:scaffold,2.12.0.Final org.jboss.forge.addon:templates-spi,2.12.0.Final org.jboss.forge.addon:templates,2.12.0.Final org.jboss.forge.addon:parser-java,2.12.0.Final org.jboss.forge.addon:dependencies,2.12.0.Final org.jboss.forge.addon:projects,2.12.0.Final org.jboss.forge.addon:parser-xml,2.12.0.Final org.jboss.forge.addon:environment,2.12.0.Final org.jboss.forge.addon:ui-spi,2.12.0.Final org.jboss.forge.addon:convert,2.12.0.Final org.jboss.forge.addon:ui,2.12.0.Final org.jboss.forge.addon:bean-validation,2.12.0.Final org.jboss.forge.addon:text,2.12.0.Final org.jboss.forge.addon:javaee-spi,2.12.0.Final org.jboss.forge.addon:rest-client,2.12.0.Final org.jboss.forge.addon:javaee,2.12.0.Final org.jboss.forge.furnace.container:cdi,2.12.0.Final org.jboss.forge.furnace.container:simple,2.12.0.Final org.jboss.forge.addon:facets,2.12.0.Final org.jboss.forge.addon:resources,2.12.0.Final org.jboss.forge.addon:configuration,2.12.0.Final org.jboss.forge.addon:database-tools,2.12.0.Final org.jboss.forge.addon:core,2.12.0.Final /home/ben/.forge/addons: reproduction path: http://forge.jboss.org/document/write-a-java-ee-web-application-basic OS: ubuntu 14.04 (64 bit) oracle java jdk 1.7.0_67-b01 (64 bit) maven 3.2.3 Stacktrace: [conference]$ scaffold-generate --targets org.conference.model.Speaker ***ERROR*** Error while executing 'Scaffold: Generate' org.jboss.forge.addon.facets.FacetNotFoundException: No Facet of type [interface org.jboss.forge.addon.javaee.faces.FacesFacet] is installed. at org.jboss.forge.addon.facets.AbstractFaceted.getFacet(AbstractFaceted.java:51) at org.jboss.forge.addon.scaffold.faces.FacesAccessStrategy.<init>(FacesAccessStrategy.java:27) at org.jboss.forge.addon.scaffold.faces.FacesScaffoldProvider.getAccessStrategy(FacesScaffoldProvider.java:440) at org.jboss.forge.addon.scaffold.faces.FacesScaffoldProvider.setupWebXML(FacesScaffoldProvider.java:555) at org.jboss.forge.addon.scaffold.faces.FacesScaffoldProvider.setup(FacesScaffoldProvider.java:211) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.jboss.forge.furnace.proxy.ClassLoaderInterceptor$1.call(ClassLoaderInterceptor.java:87) at org.jboss.forge.furnace.util.ClassLoaders.executeIn(ClassLoaders.java:42) at org.jboss.forge.furnace.proxy.ClassLoaderInterceptor.invoke(ClassLoaderInterceptor.java:103) at org.jboss.forge.addon.scaffold.spi.ScaffoldProvider_$$_javassist_cabe0c6b-761c-47e7-8b20-740c935b5b50.setup(ScaffoldProvider_$$_javassist_cabe0c6b-761c-47e7-8b20-740c935b5b50.java) at org.jboss.forge.addon.scaffold.impl.ui.ScaffoldExecuteGenerationStep.execute(ScaffoldExecuteGenerationStep.java:69) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at	1	1	https://issues.jboss.org/browse/FORGE-2050
4517	FORGE-2004	Did we discuss this issue on IRC?	1	0	https://issues.jboss.org/browse/FORGE-2004
4518	FORGE-1975	Now all deadlocks and thread leaks are fixed, but max build time has been reached (40 min). This is due to the tests, not the aesh upgrade. (This is the same with the old version) Can you increase the maximum time allowed for a build to validate this PR? Regards, Jeremie.	1	1	https://issues.jboss.org/browse/FORGE-1975
4519	FORGE-1970	I'm not convinced this is the source of the problems, and I think re-throwing the target exception would probably be a better solution anyway. Could you please provide steps to reproduce?	1	1	https://issues.jboss.org/browse/FORGE-1970
4520	FORGE-1953	Sure, create a PR and we'll be happy to merge it if it's ok. Does it allow you to set system properties too?	1	1	https://issues.jboss.org/browse/FORGE-1953

#	Requirement	Questions	Answered	Successfully Answered	URL
4521	FORGE-1953	Yes, actually we can type: - *sysprops* : to list system properties - *sysprops-get --propertyName foo* : to get the value of the system property 'foo' - *sysprops-set --propertyName foo --propertyValue bar* : to set the value 'bar' to property 'foo' What do you think about this definition? Kind regards, Mahmoud	1	1	https://issues.jboss.org/browse/FORGE-1953
4522	FORGE-1944	And what about an abstraction like the one we have with building (Maven or Gradle) ? Instead of {{junit-}} commands we could have something more abstract, such as {{test-}} and let us choose between Junit and TestNG. The Arquillian addon does that : {code} // Sets up a test framework test-setup --testFramework JUnit / TestNG --version // Creates a new test case test-new-test --named TheClassNameTest // Adds a new test to a test case test-add-test --named theMethodName --targetClass TheClassNameTes // Generates a test case for a given class (generates one test per public method) test-create-test --targets TheClassToGenerateTests ... {code}	0	0	https://issues.jboss.org/browse/FORGE-1944
4523	FORGE-1911	Hey Vineet, The Id isn't passed to the @PUT method, just the entity itself. How should this be handled then?	1	1	https://issues.jboss.org/browse/FORGE-1911
4524	FORGE-1897	[~gastaldi], http://img.ctrlv.in/img/15/02/28/54f1d0e53b6f6.png What I did? I added the getExtension in ProjectType interface. So, in AbstractProjectType I set the default value for null: @Override public String getExtension() { return null; } In each implementation, I need set this method. So.. I need to change my convert: type.setItemLabelConverter(new Converter<ProjectType, String>() { @Override public String convert(ProjectType source) { if (source == null) { return null; } if (source.getExtension() == null) { return source.getType(); } return source.getType().concat(" (").concat(source.getExtension()).concat(")"); } }); That seem good?	1	1	https://issues.jboss.org/browse/FORGE-1897
4525	FORGE-1897	This looks good, however using the String.concat method makes the code less readable :) Doesn't this also change the project-new command type option? What happens if you <TAB> the type option?	1	1	https://issues.jboss.org/browse/FORGE-1897
4526	FORGE-1881	Should this be a singleton instance?	1	0	https://issues.jboss.org/browse/FORGE-1881
4527	FORGE-1874	Is it possible to have this alias built in within Forge ?	1	1	https://issues.jboss.org/browse/FORGE-1874
4528	FORGE-1874	Hi [~agoncal], Have you tried {{build test --profile arquillian-wildfly-remote}} ? That should work afaik.	0	0	https://issues.jboss.org/browse/FORGE-1874
4529	FORGE-1874	What do you mean about treating Maven as an addon? Maven is already in it's own addon.	0	0	https://issues.jboss.org/browse/FORGE-1874
4530	FORGE-1867	Is there an existing use case for this? The calls to {{UIBuilder.add(...)}} is already a valid way of sorting UIInputs	0	0	https://issues.jboss.org/browse/FORGE-1867
4531	FORGE-1842	How do i get source code?? I am not getting from where to start??	1	1	https://issues.jboss.org/browse/FORGE-1842
4532	FORGE-1804	It's really weird. I've updated to the latest version of the code, executed the exact same commands, and ended up with badly formatted code. Could anyone else tried ? I'm running on a Mac with JDK 1.7_45	0	0	https://issues.jboss.org/browse/FORGE-1804
4533	FORGE-1802	What criteria to validate the script?	0	0	https://issues.jboss.org/browse/FORGE-1802
4534	FORGE-1802	I ran the above script and everything went fine: {code} [delme]$ run /tmp/crash.fsh project-new --named cdbookstore --topLevelPackage org.agoncal.training.javaee6adv --type war --finalName cdbookstore ; ***SUCCESS*** Project named 'cdbookstore' has been created. [cdbookstore]$ jpa-setup --persistenceUnitName cdbookstorePU ; ***SUCCESS*** Persistence (JPA) is installed. [cdbookstore]$ constraint-new-annotation --named ISBN ***SUCCESS*** Bean Validation Constraint Annotations org.agoncal.training.javaee6adv.constraints.ISBN was created [ISBN.java]$ jpa-new-entity --named Genre ; ***SUCCESS*** Entity org.agoncal.training.javaee6adv.model.Genre created [Genre.java]$ jpa-new-field --named name --length 100 ; ***SUCCESS*** Field name created [Genre.java]$ ***SUCCESS*** Field name created [delme]$ {code} Are there any issues here left to be fixed?	1	1	https://issues.jboss.org/browse/FORGE-1802
4535	FORGE-1802	OK, found the reason for the hanging shell Basically, when you run scripts, the Run command will execute the commands one after the other. It will wait until the command returns a result (success or failure) or until a timeout of 500 seconds passes. The issue doesn't come when the command *execution* throws an exception. Then the command execution listener's postCommandFailure method will be run and the run script command will not hang - it will report an error. The real issue occurs when the command *initialization* fails for some reason. In that case one of he ForgeCommandRegistry methods would throw an exception. The problem with that exception is that it gets lost somewhere in AeshShell and there is nobody to notify the command execution listeners about the failure. So the execute method in the RunCommand class will run for the full 500 seconds waiting in vain for a result. Now, I've hacked a quick solution to this problem, that works: https://github.com/ivannov/core/commit/5ed84f5f3ce5fedae2bfa100f7e270cc645a2994 I just need an advise from someone more knowledgeable: 1) Having so many nulls on line 132 is fine for RunCommand.ScriptCommandListener, but probably this does not apply for all the possible listeners that may come to this point. Obviously there is no Command object so far, but can I get somehow UIExecutionContext? I only got a UIContext 2) I just fire the failure notification when parsing of the command doesn't succeed. I can also spot some other candidates in this class. But are these all the locations that we must fire the event? Cheers, Ivan	1	1	https://issues.jboss.org/browse/FORGE-1802
4536	FORGE-1802	The thing is that after any exception is thrown here, we go to the Aesh shell. Would you propose a better place for firing those events?	0	0	https://issues.jboss.org/browse/FORGE-1802
4537	FORGE-1784	Strange, I've executed the same script again, and this time I have a stack trace {code} [agoncal-training-javaee6-advanced]$ scaffold-generate --webRoot admin --targets org.agoncal.training.javaee6adv.* ; Exception in thread "Aesh Process Loop 104588629" java.util.concurrent.RejectedExecutionException: Task AeshProcess{pid=86, manager=org.jboss.aesh.console.ProcessManager@150f7bf0, consoleCallback=org.jboss.aesh.console.AeshConsoleImpl$AeshConsoleCallbackImpl@2b91d45d, operation=ControlOperator: END, Buffer: scaffold-generate --webRoot admin --targets org.agoncal.training.javaee6adv.* } rejected from java.util.concurrent.ThreadPoolExecutor@730d3c03[Shutting down, pool size = 1, active threads = 1, queued tasks = 0, completed tasks = 85] at java.util.concurrent.ThreadPoolExecutor$AbortPolicy.rejectedExecution(ThreadPoolExecutor.java:2048) at java.util.concurrent.ThreadPoolExecutor.reject(ThreadPoolExecutor.java:821) at java.util.concurrent.ThreadPoolExecutor.execute(ThreadPoolExecutor.java:1372) at org.jboss.aesh.console.ProcessManager.startNewProcess(ProcessManager.java:44) at org.jboss.aesh.console.Console.processOperationResult(Console.java:628) at org.jboss.aesh.console.Console.processInternalOperation(Console.java:594) at org.jboss.aesh.console.Console.execute(Console.java:582) at org.jboss.aesh.console.Console.access$700(Console.java:60) at org.jboss.aesh.console.Console$6.run(Console.java:499) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at java.lang.Thread.run(Thread.java:744) {code}	1	1	https://issues.jboss.org/browse/FORGE-1784

#	Requirement	Questions	Answered	Successfully Answered	URL
4538	FORGE-1780	Hi, I did this issue, but I have a problem. I created a method to do this task, but I don't know where put my code, into Strings class of the roaster or furnace. What the better place?	1	1	https://issues.jboss.org/browse/FORGE-1780
4539	FORGE-1750	1. What about creating one command for both FORGE-1749, FORGE-1750 prompting the user while executing to choose one? 2. What builder pattern application to choose? setName vs withName, new Person() vs Person.newInstance()?	1	1	https://issues.jboss.org/browse/FORGE-1750
4540	FORGE-1729	Thanks, right. My question, though, is should this happen automatically when running in a test environment? Seems like this is something that you might expect to be compartmentalized for you. WDYT?	0	0	https://issues.jboss.org/browse/FORGE-1729
4541	FORGE-1729	Is there any system property we might check if it is being running from a unit test?	0	0	https://issues.jboss.org/browse/FORGE-1729
4542	FORGE-1729	Is there any system property we might check if it is being running from a unit test?	0	0	https://issues.jboss.org/browse/FORGE-1729
4543	FORGE-1712	Would this have a series of questions for people to answer, or is it simply delegating to the various commands?	1	1	https://issues.jboss.org/browse/FORGE-1712
4544	FORGE-1699	Hey Esteve, Sorry you are having this issue! Does this occur for all commands such as 'cd' and 'ls' as well, or is it limited only to specific commands? Thank you, Lincoln	1	1	https://issues.jboss.org/browse/FORGE-1699
4545	FORGE-1699	Hi Esteve, Could you replace Aesh with 0.49 manually in your {{addons/org-jboss-forge-addon-shell-spi-2-3-0-Final}} folder to confirm if this is an Aesh issue or something in the core? Don't forget to replace aesh-extensions in {{addons/org-jboss-forge-addon-shell-2-3-0-Final}} too. Thanks!	1	1	https://issues.jboss.org/browse/FORGE-1699
4546	FORGE-1699	Hey Esteve, Does slowness persist after the first TAB completes? It should be faster (but still not instant.) ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-1699
4547	FORGE-1699	What do you mean "dies?"	1	1	https://issues.jboss.org/browse/FORGE-1699
4548	FORGE-1699	Sure, actually. If you could do something, that would be great: I'm wondering if you could start a fresh forge instance, press <TAB>. Wait for the CPU usage to spike, then in another terminal, use 'kill -3 xxxxxx' where 'xxxxxx' is the process id of the running forge instance. (Can be obtained using 'ps -ef | grep forge' - it will be the PID of the Java process) Also, does the CLI ever return from a TAB? Or does it hang "forever"? I do notice a delay, but it always returns on my machine.	1	1	https://issues.jboss.org/browse/FORGE-1699
4549	FORGE-1699	Could you please check to see if this issue still occurs in Forge 2.4.1.Final ? Thank you! ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-1699
4550	FORGE-1699	From the graph, it looks like this took your computer about 30 seconds to complete. Is that the case? On my laptop it takes about 5 seconds to run all of the pasted commands. CPU usage is slightly high, but overall tolerable. How long does this take on your computer? What is your Java Version? What type of CPU do you have, and how much memory is allocated to your JVM? It's possible that garbage collection is running crazy.	1	1	https://issues.jboss.org/browse/FORGE-1699
4551	FORGE-1699	This doesn't seem related to Forge. Are you monitoring the Forge process only? I opened the same and my CPU never reaches above 60% on peak. It stays 20% most of the time.	1	1	https://issues.jboss.org/browse/FORGE-1699
4552	FORGE-1699	Curious, how fast and what type of CPU do you have?	0	0	https://issues.jboss.org/browse/FORGE-1699
4553	FORGE-1698	Isn't this already implemented? (https://github.com/forge/core/blob/master/maven/impl-projects/src/main/java/org/jboss/forge/addon/maven/projects/facets/MavenJavaCompilerFacet.java)	0	0	https://issues.jboss.org/browse/FORGE-1698
4554	FORGE-1680	This is caused because Furnace failed to start: {code} Mar 18, 2014 3:00:41 PM org.jboss.forge.furnace.impl.FurnaceImpl addRepository INFO: Monitoring repository [/Users/lb3/projects/jboss/forge/jbosstools-forge/plugins/org.jboss.tools.forge2.runtime/addon-repository] for file changes. Mar 18, 2014 3:00:41 PM org.jboss.forge.furnace.impl.FurnaceImpl addRepository INFO: Monitoring repository [/Users/lb3/.forge/addons] for file changes. Mar 18, 2014 3:00:41 PM org.jboss.forge.furnace.impl.FurnaceImpl start INFO: Furnace [2.1.1.Final] starting. Mar 18, 2014 3:00:41 PM org.jboss.forge.furnace.impl.FurnaceImpl start INFO: Detected changes in repository [/Users/lb3/projects/jboss/forge/jbosstools-forge/plugins/org.jboss.tools.forge2.runtime/addon-repository (Immutable)]. Mar 18, 2014 3:00:41 PM org.jboss.forge.furnace.impl.FurnaceImpl start INFO: Detected changes in repository [/Users/lb3/.forge/addons]. Mar 18, 2014 3:00:42 PM org.jboss.forge.furnace.impl.FurnaceImpl start SEVERE: Error occurred. java.lang.IllegalStateException: Cycle detected in Addon graph: [[org.jboss.forge.addon:scaffold-faces,2.1.1.Final], [org.jboss.forge.addon:javaee,2.1.2-SNAPSHOT], [org.jboss.forge.addon:scaffold,2.1.2-SNAPSHOT]] at org.jboss.forge.furnace.impl.graph.OptimizedAddonGraph.<init>(OptimizedAddonGraph.java:87) at org.jboss.forge.furnace.impl.addons.AddonLifecycleManager$4.call(AddonLifecycleManager.java:194) at org.jboss.forge.furnace.impl.addons.AddonLifecycleManager$4.call(AddonLifecycleManager.java:182) at org.jboss.forge.furnace.impl.LockManagerImpl.performLocked(LockManagerImpl.java:48) at org.jboss.forge.furnace.impl.addons.AddonLifecycleManager.forceUpdate(AddonLifecycleManager.java:181) at org.jboss.forge.furnace.impl.FurnaceImpl.start(FurnaceImpl.java:192) at org.jboss.forge.furnace.impl.FurnaceImpl$1.run(FurnaceImpl.java:123) {code} We probably need to figure out a way to fall back here and load only the internal repository if this occurs. Or to make Furnace more resilient.	0	0	https://issues.jboss.org/browse/FORGE-1680
4555	FORGE-1671	Another attempt... but looks like it's impossible to do with Roaster at the moment : {code} public class BeanValidationAnnotation { public static void main(String[] args) { // This is the inner annotation List final JavaAnnotationSource listAnnotation = Roaster.create(JavaAnnotationSource.class); listAnnotation.setName("List"); listAnnotation.addAnnotation(Retention.class).setEnumValue(RUNTIME); listAnnotation.addAnnotation(Target.class).setEnumValue(METHOD, FIELD, PARAMETER, TYPE); listAnnotation.addAnnotationElement("Email[] value()"); System.out.println(listAnnotation); // This is the annotation @Email final JavaAnnotationSource emailAnnotation = Roaster.create(JavaAnnotationSource.class); emailAnnotation.setPackage("org.agoncal.proj.constraints").setName("Email"); emailAnnotation.addImport(Payload.class); emailAnnotation.addAnnotation(Documented.class); emailAnnotation.addAnnotation(Retention.class).setEnumValue(RUNTIME); emailAnnotation.addAnnotation(Target.class).setEnumValue(METHOD, FIELD, PARAMETER, TYPE); emailAnnotation.addAnnotation(Constraint.class).setLiteralValue("validatedBy", "{}"); emailAnnotation.addAnnotationElement("String message() default \"wrong email address\""); emailAnnotation.addAnnotationElement("Class<?>[] groups() default {}"); emailAnnotation.addAnnotationElement("Class<? extends Payload>[] payload() default {}"); // I was expecting to have a method like that so I could add annotation inside another one // emailAnnotation.addAnnotationElement(emailAnnotation); System.out.println(emailAnnotation); } }{code}	1	1	https://issues.jboss.org/browse/FORGE-1671

#	Requirement	Questions	Answered	Successfully Answered	URL
4556	FORGE-1671	Cool, it works with the new {{addNestedType}} type of Roaster 2.1. I just need a bit of time to send a pull request ;o) {code} public class BeanValidationAnnotation { public static void main(String[] args) { // This is the inner annotation List final JavaAnnotationSource listAnnotation = Roaster.create(JavaAnnotationSource.class); listAnnotation.setName("List"); listAnnotation.addAnnotation(Retention.class).setEnumValue(RUNTIME); listAnnotation.addAnnotation(Target.class).setEnumValue(METHOD, FIELD, PARAMETER, TYPE); listAnnotation.addAnnotationElement("Email[] value()"); // This is the annotation @Email final JavaAnnotationSource emailAnnotation = Roaster.create(JavaAnnotationSource.class); emailAnnotation.setPackage("org.agoncal.proj.constraints").setName("Email"); emailAnnotation.addImport(Payload.class); emailAnnotation.addAnnotation(Documented.class); emailAnnotation.addAnnotation(Retention.class).setEnumValue(RUNTIME); emailAnnotation.addAnnotation(Target.class).setEnumValue(METHOD, FIELD, PARAMETER, TYPE); emailAnnotation.addAnnotation(Constraint.class).setLiteralValue("validatedBy", "{}"); emailAnnotation.addAnnotationElement("String message() default \"wrong email address\""); emailAnnotation.addAnnotationElement("Class<?>[] groups() default {}"); emailAnnotation.addAnnotationElement("Class<? extends Payload>[] payload() default {}"); emailAnnotation.addNestedType(listAnnotation); System.out.println(emailAnnotation); } } {code}	1	1	https://issues.jboss.org/browse/FORGE-1671
4557	FORGE-1644	The AddonRepository is associated with a Single Addon in a specific version isn't it? How would you deal with stored data across Addon versions? How do you get access to your own Addon info?	0	0	https://issues.jboss.org/browse/FORGE-1644
4558	FORGE-1625	Shoud an exception be thrown when an UIWizard is added?	1	1	https://issues.jboss.org/browse/FORGE-1625
4559	FORGE-1610	Ideally this feature would check the Forge website for the latest version, then prompt the user: "A new version of Forge is available, would you like to update? [Y,n]" and of course, once that is done: "Attempt to update automatically when Forge starts? [Y,n]"	0	0	https://issues.jboss.org/browse/FORGE-1610
4560	FORGE-1603	I'd like to bring this discussion to the channel or the forge-dev list. Does an overwrite flag make sense or should we just display the validation error?	0	0	https://issues.jboss.org/browse/FORGE-1603
4561	FORGE-1599	I think we could make our custom cache builder with additional options, and in case user wants to use it then he could create his own Cache producer which would use this extension. I don't know if that will work with many addons at the same time, is there some kind of addon scope in Forge CDI? Otherwise it will be necessary to use own custom annotations (or @Named) for caches other than default.	1	1	https://issues.jboss.org/browse/FORGE-1599
4562	FORGE-1593	What do you mean? Could you out an example? Thanks	1	1	https://issues.jboss.org/browse/FORGE-1593
4563	FORGE-1591	Thanks. Closing it	0	0	https://issues.jboss.org/browse/FORGE-1591
4564	FORGE-1590	[~abhijithtn], I just pushed the changes to the master branch. Since I don't have a proxy environment to test it, I need you to build from the sources and test it again. Can you do that for me? Thanks!	1	1	https://issues.jboss.org/browse/FORGE-1590
4565	FORGE-1590	Can you paste2.org the stacktrace please?	1	0	https://issues.jboss.org/browse/FORGE-1590
4566	FORGE-1590	SUCCESS [20.505s] [INFO] Forge - Addon Manager - Parent ..... SUCCESS [0.498s] [INFO] Forge - Addon Manager SPI ..... SUCCESS [37.105s] [INFO] Forge - Facets Addon Parent ..... SUCCESS [0.463s] [INFO] Forge - Facets Addon API ..... SUCCESS [36.820s] [INFO] Forge - Facets Addon Impl ..... SUCCESS [1.718s] [INFO] Forge - Facets Addon ..... FAILURE [28.864s] [INFO] Forge - Resource Addon Parent ..... SUCCESS [0.565s] [INFO] Forge - Resource Addon API ..... SKIPPED [INFO] Forge - Convert Addon Parent ..... SUCCESS [0.509s] [INFO] Forge - Convert API ..... SUCCESS [2.594s] [INFO] Forge - Dependencies Addon Parent ..... SUCCESS [0.466s] [INFO] Forge - Dependencies API ..... SKIPPED [INFO] Forge - Addon Manager API ..... SKIPPED [INFO] Forge - Environment Addon Parent ..... SUCCESS [0.555s] [INFO] Forge - Environment API ..... SUCCESS [1.330s] [INFO] Forge - Environment Impl ..... SUCCESS [1.729s] [INFO] Forge - Environment Addon ..... SKIPPED [INFO] Forge - UI Addon Parent ..... SUCCESS [0.554s] [INFO] Forge - UI Addon SPI ..... SKIPPED [INFO] Forge - UI API ..... SKIPPED [INFO] Forge - Convert Impl ..... SUCCESS [2.514s] [INFO] Forge - Convert Addon ..... SKIPPED [INFO] Forge - Resource Addon Impl ..... SKIPPED [INFO] Forge - Resource Addon ..... SKIPPED [INFO] Forge - Dependencies Impl ..... SKIPPED [INFO] Forge - Dependencies Addon ..... SKIPPED [INFO] Forge - UI Impl ..... SKIPPED [INFO] Forge - UI Addon ..... SKIPPED [INFO] Forge - Projects Addon Parent ..... SUCCESS [0.509s] [INFO] Forge - Projects API ..... SKIPPED [INFO] Forge - Addon Manager - Impl ..... SKIPPED [INFO] Forge - Shell Addon Parent ..... SUCCESS [0.593s] [INFO] Forge - Shell SPI Addon ..... SKIPPED [INFO] Forge - Projects Impl ..... SKIPPED [INFO] Forge - Projects Addon ..... SKIPPED [INFO] Forge - Addon Manager Addon ..... SKIPPED [INFO] Forge - XML Parser Addon Parent ..... SUCCESS [0.544s] [INFO] Forge - XML Parser Addon API ..... SKIPPED [INFO] Forge - XML Parser Addon Implementation ..... SKIPPED [INFO] Forge - XML Parser Addon ..... SKIPPED [INFO] Forge - Maven Addon Parent ..... SUCCESS [0.573s] [INFO] Forge - Maven API ..... SKIPPED [INFO] Forge - Maven Impl ..... SKIPPED [INFO] Forge - Java Parser Addon Parent ..... SUCCESS [0.599s] [INFO] Forge - Java Parser Addon API ..... SKIPPED [INFO] Forge - Java Parser Addon Implementation ..... SKIPPED [INFO] Forge - Java Parser Addon ..... SKIPPED [INFO] Forge - Text Addon ..... SKIPPED [INFO] Forge - JavaEE Parent ..... SUCCESS [11.772s] [INFO] Forge - JavaEE API ..... SKIPPED [INFO] Forge - Templates Addon Parent ..... SUCCESS [11.761s] [INFO] Forge - Templates Addon SPI ..... SKIPPED [INFO] Forge - Templates Addon API ..... SKIPPED [INFO] Forge - Templates Addon Impl ..... SKIPPED [INFO] Forge - Templates Addon ..... SKIPPED [INFO] Forge - JavaEE Impl ..... SKIPPED [INFO] Forge - JavaEE Addon ..... SKIPPED [INFO] Forge - Maven Projects Impl ..... SKIPPED [INFO] Forge - Maven Addon ..... SKIPPED [INFO] Forge - Shell Addon API ..... SKIPPED [INFO] Forge - Shell Addon Impl ..... SKIPPED [INFO] Forge - Shell Addon ..... SKIPPED [INFO] Forge - Shell Addon Test Harness ..... SKIPPED [INFO] Forge - Addon Manager - Tests ..... SKIPPED [INFO] Forge - Addon Development Addon Parent ..... SUCCESS [11.688s] [INFO] Forge - Addon Development Addon API ..... SKIPPED [INFO] Forge - Addon Development Addon Impl ..... SKIPPED [INFO] Forge - Addon Development Addon ..... SKIPPED [INFO] Forge - UI Test Harness ..... SKIPPED [INFO] Forge - Addon Development Addon Tests ..... SKIPPED [INFO] Forge - Convert Addon Tests ..... SKIPPED [INFO] Forge - Configuration Parent ..... SUCCESS	1	1	https://issues.jboss.org/browse/FORGE-1590
4567	FORGE-1590	Can you paste2.org the stacktrace?	1	0	https://issues.jboss.org/browse/FORGE-1590

#	Requirement	Questions	Answered	Successfully Answered	URL
4568	FORGE-1590	org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest Tests run: 8, Failures: 2, Errors: 3, Skipped: 0, Time elapsed: 657.64 sec <<< FAILURE! - in org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest testResolveDependencyHierarchy(org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest) Time elapsed: 233.761 sec <<< ERROR! org.jboss.forge.addon.dependencies.DependencyException: Could not resolve dependencies for addon [org.jboss.forge:resources:jar:forge-addon:2.0.0.Alpha3] at org.eclipse.aether.internal.impl.DefaultRepositorySystem.resolveDependencies(DefaultRepositorySystem.java:361) at org.jboss.forge.addon.maven.dependencies.MavenDependencyResolver.resolveDependencyHierarchy(MavenDependencyResolver.java:249) at org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest.testResolveDependencyHierarchy(MavenDependencyResolverTest.java:125) testResolveClassifiedArtifact(org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest) Time elapsed: 0.214 sec <<< ERROR! java.lang.RuntimeException: Could not resolve dependencies from Query [org.jboss.forge:resources::forge-addon:2.0.0.Alpha3 - null] due to underlying exception at org.eclipse.aether.internal.impl.DefaultRepositorySystem.resolveDependencies(DefaultRepositorySystem.java:361) at org.jboss.forge.addon.maven.dependencies.MavenDependencyResolver.resolveDependencies(MavenDependencyResolver.java:93) at org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest.testResolveClassifiedArtifact(MavenDependencyResolverTest.java:54) testResolveVersions(org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest) Time elapsed: 63.486 sec <<< FAILURE! java.lang.AssertionError: null at org.junit.Assert.fail(Assert.java:86) at org.junit.Assert.assertTrue(Assert.java:41) at org.junit.Assert.assertFalse(Assert.java:64) at org.junit.Assert.assertFalse(Assert.java:74) at org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest.testResolveVersions(MavenDependencyResolverTest.java:78) testResolveVersionsDependency(org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest) Time elapsed: 63.589 sec <<< FAILURE! java.lang.AssertionError: null at org.junit.Assert.fail(Assert.java:86) at org.junit.Assert.assertTrue(Assert.java:41) at org.junit.Assert.assertFalse(Assert.java:64) at org.junit.Assert.assertFalse(Assert.java:74) at org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest.testResolveVersionsDependency(MavenDependencyResolverTest.java:98) testResolveMetadata(org.jboss.forge.addon.maven.dependencies.MavenDependencyResolverTest) Time elapsed: 0.149 sec <<< ERROR! org.jboss.forge.addon.dependencies.DependencyException: Unable to resolve any artifacts for query [org.jboss.forge:resources:jar:forge-addon:2.0.0.Alpha3 - null] at org.eclipse.aether.internal.impl.DefaultUpdateCheckManager.newException(DefaultUpdateCheckManager.java:239) at org.eclipse.aether.internal.impl.DefaultUpdateCheckManager.checkArtifact(DefaultUpdateCheckManager.java:218) at org.eclipse.aether.internal.impl.DefaultArtifactResolver.gatherDownloads(DefaultArtifactResolver.java:599) at org.eclipse.aether.internal.impl.DefaultArtifactResolver.performDownloads(DefaultArtifactResolver.java:518) at org.eclipse.aether.internal.impl.DefaultArtifactResolver.resolve(DefaultArtifactResolver.java:436) at	1	1	https://issues.jboss.org/browse/FORGE-1590
4569	FORGE-1590	Do you have the proxy set in settings.xml?	1	0	https://issues.jboss.org/browse/FORGE-1590
4570	FORGE-1590	Hi Abhijith, Since I am not in a proxy environment, we'll need you to debug it as an attempt to find the root cause. Can you do that for us? You need to debug the MavenContainer class (where the proxies are set) to confirm that it is really a bug.	0	0	https://issues.jboss.org/browse/FORGE-1590
4571	FORGE-1582	I've been trying to create a Bean Validation constraint properly... in vain. What I want to achieve is the following (notice the @List annotation) : {code} @Target({METHOD, FIELD, ANNOTATION_TYPE, CONSTRUCTOR, PARAMETER}) @Retention(RUNTIME) @Documented @Constraint(validatedBy = {}) public @interface Email { String message() default "wrong email address"; Class<?>[] groups() default {}; Class<? extends Payload>[] payload() default {}; @Target({METHOD, FIELD, ANNOTATION_TYPE, CONSTRUCTOR, PARAMETER}) @Retention(RUNTIME) @interface List { Email[] value(); } } {code} This is the code I've written. Both annotation are generated separately... but I don't know how to add listAnnotation inside emailAnnotation. Any idea ? {code} public class Main { public static void main(String[] args) { // This is the inner annotation List final JavaAnnotationSource listAnnotation = Roaster.create(JavaAnnotationSource.class); listAnnotation.setName("List"); listAnnotation.addAnnotation(Retention.class).setEnumValue(RUNTIME); listAnnotation.addAnnotation(Target.class).setEnumValue(METHOD, FIELD, PARAMETER, TYPE); listAnnotation.addAnnotationElement("Email[] value()"); System.out.println(listAnnotation); // This is the annotation @Email final JavaAnnotationSource emailAnnotation = Roaster.create(JavaAnnotationSource.class); emailAnnotation.setPackage("org.agoncal.proj.constraints").setName("Email"); emailAnnotation.addImport(Payload.class); emailAnnotation.addAnnotation(Documented.class); emailAnnotation.addAnnotation(Retention.class).setEnumValue(RUNTIME); emailAnnotation.addAnnotation(Target.class).setEnumValue(METHOD, FIELD, PARAMETER, TYPE); emailAnnotation.addAnnotation(Constraint.class).setLiteralValue("validatedBy", "{}"); emailAnnotation.addAnnotationElement("String message() default \"wrong email address\""); emailAnnotation.addAnnotationElement("Class<?>[] groups() default {}"); emailAnnotation.addAnnotationElement("Class<? extends Payload>[] payload() default {}"); // I was expecting to have a method like that so I could add annotation inside another one // emailAnnotation.addAnnotationElement(emailAnnotation); System.out.println(emailAnnotation); } } {code}	1	1	https://issues.jboss.org/browse/FORGE-1582
4572	FORGE-1555	It looks like it's already working. Can you recheck this issue please?	0	0	https://issues.jboss.org/browse/FORGE-1555
4573	FORGE-1554	Is there any way we could override that behavior? I agree with Vineet that this makes Forge users productive.	1	1	https://issues.jboss.org/browse/FORGE-1554
4574	FORGE-1550	http://www.travis-ci.org Eg. https://travis-ci.org/aeshell/aesh WDYT?	1	1	https://issues.jboss.org/browse/FORGE-1550
4575	FORGE-1550	The project use travis: https://github.com/forge/core/blob/master/.travis.yml Do you want migrate to other service?	1	1	https://issues.jboss.org/browse/FORGE-1550
4576	FORGE-1550	I understand, the problem is builder on Windows. I don't know a free CI with windows. Doesn't viable execute builder on Windows in another server?	0	0	https://issues.jboss.org/browse/FORGE-1550
4577	FORGE-1509	No more freezing, however.. when I do: {code} jpa-setup (SUCCESS) ... remove the dependencies (no need to remove managed only) jpa-setup (SUCCESS) {code} the pom.xml is in the inconsistent state and the mvn install will not work. In other words, jpa-setup will not fix the pom.xml. Shouldn't be somehow fixed?	1	0	https://issues.jboss.org/browse/FORGE-1509
4578	FORGE-1498	Not sure I really follow this issue. What is the problem? Service lookups just return objects, so activation order could be handled by the consumer. -- Lincoln Baxter, III http://ocpsoft.org "Simpler is better."	1	1	https://issues.jboss.org/browse/FORGE-1498
4579	FORGE-1498	Ah, so you're saying that you want the listeners to be ordered? Do we ever need to execute them in order? If so, we should add a priority() method and/or make them Comparable/Provide a comparator. -- Lincoln Baxter, III http://ocpsoft.org "Simpler is better."	1	1	https://issues.jboss.org/browse/FORGE-1498

#	Requirement	Questions	Answered	Successfully Answered	URL
4580	FORGE-1498	Perhaps we could assume that Comparable instances could be automatically ordered in the Imported.iterator(). Then it could be up to the implementation to decide which strategy to use for ordering. WDYT? ----- Mensagem original ----- De: "Lincoln Baxter III (JIRA)" <issues@jboss.org> Para: forge-issues@lists.jboss.org Enviadas: Domingo, 26 de janeiro de 2014 16:59:29 Assunto: [forge-issues] [JBoss JIRA] (FORGE-1498) Allow ordering in the lookup of services in AddonRegistry [https://issues.jboss.org/browse/FORGE-1498?page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel&focusedCommentId=12938797#comment-12938797] Lincoln Baxter III commented on FORGE-1498: ---------- Ah, so you're saying that you want the listeners to be ordered? Do we ever need to execute them in order? If so, we should add a priority() method and/or make them Comparable/Provide a comparator. -- Lincoln Baxter, III http://ocpsoft.org "Simpler is better." -- This message is automatically generated by JIRA. If you think it was sent incorrectly, please contact your JIRA administrators For more information on JIRA, see: http://www.atlassian.com/software/jira _____ forge-issues mailing list forge-issues@lists.jboss.org https://lists.jboss.org/mailman/listinfo/forge-issues	1	0	https://issues.jboss.org/browse/FORGE-1498
4581	FORGE-1462	Hi Justin, We are still in the process of building our Roadmap for the future releases. Do you have any feature that would be essential for your adoption? Please use the Forums or the mailing list (forge-dev or forge-users) for questions. If you need some help, we are pleased to help you out on the #forge IRC channel. Best Regards, George Gastaldi	1	0	https://issues.jboss.org/browse/FORGE-1462
4582	FORGE-1452	Antonio, How can I reproduce that? This doesn't happen if you setup your project to Servlet 3.1	1	1	https://issues.jboss.org/browse/FORGE-1452
4583	FORGE-1432	What do you mean by that ? For me, the {{run}} command is there to run a script. I just give it a try but it doesn't work. Let's say I have the following {{test.fsh}} scrip : {code} @/* Country */; jpa-new-entity --named Country ; jpa-new-field --named isoCode --length 2 ; jpa-new-field --named name --length 80 ; jpa-new-field --named printableName --length 80 ; jpa-new-field --named iso3 --length 3 ; jpa-new-field --named numcode --length 3 ; {code} This is what I get when I run the following commands : {code} [Forge]$ project-new --named Tutu ***SUCCESS*** Project named 'Tutu' has been created. [Tutu]$ run ../test.fsh ***ERROR*** ../test.fsh: timed out. {code}	0	0	https://issues.jboss.org/browse/FORGE-1432
4584	FORGE-1432	Hey [~stalep] - Could you look at this class and tell me if you can think of any reason why it would freeze after executing the first line of a script? https://github.com/forge/core/blob/master/shell/impl/src/main/java/org/jboss/forge/addon/shell/command/RunCommand.java#L100 Thanks!	1	1	https://issues.jboss.org/browse/FORGE-1432
4585	FORGE-1432	How does this work ? I have a timeout. I've created a simple file with a few command lines : {code} echo before export export ACCEPT_DEFAULTS=true echo before project-new project-new --named testrun --topLevelPackage com.testrun --type war --finalName testrun echo script end {code} When I run the file, I have the following output : {code} [Forge]$ run proj.fsh [Forge]$ echo before export before export [Forge]$ ***ERROR*** proj.fsh: timed out. [Forge]$ {code}	1	0	https://issues.jboss.org/browse/FORGE-1432
4586	FORGE-1418	Could you please take a look at this? Thanks!	1	1	https://issues.jboss.org/browse/FORGE-1418
4587	FORGE-1407	Not quite working for the `scaffold-setup` wizard: {code} [jpa-demo]$ scaffold- scaffold-generate scaffold-setup [jpa-demo]$ scaffold-setup --facesVersion 2. 2.2 2.1 2.0 [jpa-demo]$ scaffold-setup --facesVersion 2. 2.2 2.1 2.0 [jpa-demo]$ scaffold-setup --facesVersion 2.2 {code} As you can see, the options that should be available are not, and the wizard moves immediately to the next setup step: {code} @Inject @WithAttributes(label = "Provider", required = true) private UISelectOne<ScaffoldProvider> provider; @Inject @WithAttributes(label = "Target Directory") private UIInput<String> target; @Inject @WithAttributes(label = "Overwrite existing files?", defaultValue = "false") private UIInput<Boolean> overwrite; {code}	1	1	https://issues.jboss.org/browse/FORGE-1407
4588	FORGE-1398	Already implemented the command + relationshipType. What about restricting user to select @Embedded relationshipType for not an @Embeddable object and vice versa? E.g not showing these relationTypes as suggestions?	0	0	https://issues.jboss.org/browse/FORGE-1398
4589	FORGE-1396	Good question. Can I get the opinion of [~crobson] again? What do you think Catherine? Should we go singular or plural, or both, and in what situations?	1	1	https://issues.jboss.org/browse/FORGE-1396
4590	FORGE-1392	Was this occurring during a running resource-transaction?	1	0	https://issues.jboss.org/browse/FORGE-1392
4591	FORGE-1381	What do you mean be uninstalling facets ? Does this also mean deleting the code that has been generated ? I'm using Forge to do some quick prototyping and sometimes just want to erase and start everything again, something like : {code} jpa-new-entity Customer jpa-new-field string name # got it wrong jpa-clear {code}	1	0	https://issues.jboss.org/browse/FORGE-1381
4592	FORGE-1346	I like this idea. I wonder [~stalep], is there any reason this could not be done that you can think of?	1	1	https://issues.jboss.org/browse/FORGE-1346
4593	FORGE-1346	well, atm this wont be possible in æsh since the completor never remove any text. it can only add. - ive planned to change this in an upcoming version, but i do not know when. imo i think a feature like this would create more confusion than it would be helpful. i just cant imagine an intuitive way of selecting which command to select. - if you're not thinking of similar to what zsh have?	1	0	https://issues.jboss.org/browse/FORGE-1346
4594	FORGE-1333	Example: {code} NAME jpa-setup -- Java EE/JPA SYNOPSIS jpa-setup [-m] [-c --container] PersistenceContainer [-j --jpaVersion] JPAFacet [-p --provider] PersistenceProvider DESCRIPTION Setup JPA in your project (multi-step wizard - some options may not be auto-documented in this man page.) -c container Container - [PersistenceContainer] (required) -j jpaVersion JPA Version - [JPAFacet] (required) -m configureMetadata Install a JPA 2 metamodel generator? - [Boolean] -p provider Provider - [PersistenceProvider] (required) RESOURCES forge: <http://forge.jboss.org/> COPYING Copyright 2013 Red Hat, Inc. and/or its affiliates. Licensed under the Eclipse Public License version 1.0, available at http://www.eclipse.org/legal/epl-v10.html Manual page jpa-setup(1) line 1 (press h for help or q to quit) {code}	1	1	https://issues.jboss.org/browse/FORGE-1333
4595	FORGE-1330	Do you really think this is more intuitive than new-project? I have doubts about this change in particular, singly because it is the first command that people need to type, and it would be an un-intuitive change from Forge 1. I know that this would be inconsistent with the way we are naming things in Forge 2, but I am not sure how to handle this. Let me loop in [~crobson] (our UI expert) and see what she thinks.	1	1	https://issues.jboss.org/browse/FORGE-1330
4596	FORGE-1330	[~crobson], The issue here is that our commands are generally named using the format: "technology-command" (See related issue: https://issues.jboss.org/browse/FORGE-944) But the "new-project" command is obviously different. What do you think we should do here?	1	1	https://issues.jboss.org/browse/FORGE-1330
4597	FORGE-1330	Yeah, this one is challenging. 'new-project' is certainly more intuitive from a first use standpoint - but project-new will become more intuitive over time once people are used to the Forge formatting of other object actions. Is there any reason why we cannot support both as valid entries to perform the action? This might be the best solution if it is possible.	0	0	https://issues.jboss.org/browse/FORGE-1330

#	Requirement	Questions	Answered	Successfully Answered	URL
4598	FORGE-1328	In Beta3, it should read: {{addon-install-build}}. Are you sure you're using Beta3 (tried emptying the ~/.forge before running forge --install core)?	0	0	https://issues.jboss.org/browse/FORGE-1328
4599	FORGE-1303	Does Forge or its addons use TCCL anyhow or do you only set it for Weld to work?	1	0	https://issues.jboss.org/browse/FORGE-1303
4600	FORGE-1250	Hey Matt! Just checking in. How is the JavaParser stuff going?	0	0	https://issues.jboss.org/browse/FORGE-1250
4601	FORGE-1229	This might not be a bad idea - Is the functionality in maven-jar-plugin documented/supported? If it's a bug or coincidence, I'd avoid making this change.	0	0	https://issues.jboss.org/browse/FORGE-1229
4602	FORGE-1228	I don't know what a pull request is. If it involves more than 1/2 hour of my time, I can't afford it right now. Possibly in a couple weeks however.....	1	1	https://issues.jboss.org/browse/FORGE-1228
4603	FORGE-1226	This is something we wanted for a long time. Wanna take a shot at it? Pull-Requests are more than welcomed ;)	0	0	https://issues.jboss.org/browse/FORGE-1226
4604	FORGE-1204	Hey Robb, Sorry you're having trouble, but this feature is definitely not broken (or nothing would work and the release would have failed hard) - Since I'm guessing you encountered this when upgrading to Alpha12, is it possible that the Resources Addon has not started yet? Remember to check all the versions in your Addon.waitUntilStartd(...) statements to give Furnace appropriate time to start up the Addons you need. ~Lincoln	1	0	https://issues.jboss.org/browse/FORGE-1204
4605	FORGE-1155	{quote} <gastaldi> right, but I need to set the values to the UIInput objects before that <gastaldi> should I do that in the Completer implementation for each option ? <stalep> aesh does that for you <stalep> atleast all values that are "finished" <stalep> so if you have: "foo --bar one --" then the bar field will be populated <stalep> eh, in that example aesh itself will try to populate the option name <stalep> but if it was trying to complete a value it would call the specific option's complete class <stalep> with the bar value populated <gastaldi> how would it set the values in the UIInput objects inside my UICommand? <gastaldi> org.jboss.forge.addon.shell.aesh.CommandLineUtil.populateUIInputs(CommandLine, Map<String, InputComponent<?, Object>>) does that <stalep> there is a Converter interface that will be called <vineetreynolds> bbl <stalep> public interface CLConverter<T> { <stalep> T convert(String input); <stalep> } <gastaldi> how does that work with multiple values ? <stalep> if its a list object its called for each value <stalep> the aesh adds it to the list <stalep> or set <stalep> etc {quote}	1	1	https://issues.jboss.org/browse/FORGE-1155
4606	FORGE-1138	How about: {code} 'forge update' only works when it is run outside from the IDE. The embedded Forge versions are automatically updated when there is a new version of the JBoss Tools Forge Plugin. If you want to use the latest version without waiting for a new plugin release, do the following: 1. Download the latest Forge version from http://forge.jboss.org and unzip in any place you may find convenient; 2. In Window->Preferences, look for Forge->Installed Forge Runtimes and add the path to your installation and make it the default runtime choice; 3. Start the forge console. You should see it is running the latest version. {code}	0	0	https://issues.jboss.org/browse/FORGE-1138
4607	FORGE-1129	hey max, you've mentioned jmerge before, but does that work with the eclipse (accept changes) merge editor?	1	0	https://issues.jboss.org/browse/FORGE-1129
4608	FORGE-1125	Gerry, Could you submit a Pull request with the changes? If done in time, we can still make it to 1.4.1.Final, to be released tonight or tomorrow morning. Best Regards, George Gastaldi	1	1	https://issues.jboss.org/browse/FORGE-1125
4609	FORGE-1113	Isn't this achieved by listening to PreShutdown and PostStartup events?	1	0	https://issues.jboss.org/browse/FORGE-1113
4610	FORGE-1067	Could this be handled a parameter? By default, we do not use it, however, there could be an additional parameter like scaffold from-entity model.* --overwrite --webjars	1	0	https://issues.jboss.org/browse/FORGE-1067
4611	FORGE-1067	I like the idea but I do have my concerns since: A) this would require webjar to have the exact version we support available. i.e. we would need to somehow QE that the official bootstrap.* files we support is actually the ones explicitly included into the files... B) initial blog says you need to add additional repositories to your project: <repositories> <repository> <id>webjars</id> <url>http://webjars.github.com/m2</url> </repository> </repositories> That is not good if true. B) looks like users would have to update all their specific version references of css/js in all their html pages when updates are made ? C) how does it work in m2e/eclipse and other IDE integrations ? If I read it right the actual javascript and css is not actually available in the filesystem making file:// viewing and things like livereload not working out of the box. D) Is this is only for serverside processed html (i.e. JSF) since the resources are on the classpath, not your project ? E) Since the resources are on the classpath my assumption is you won't see changes to these before restarting the module - this is probably not a big deal since it is just for a dependendent library but debugging/tweaking does becomes harder with this on first look at least. ...on the other hand I really like it since it does make dependency fetching much saner...but is it worth the loss of livereload, incremental updates, ide tooling ?	1	1	https://issues.jboss.org/browse/FORGE-1067
4612	FORGE-1043	Hey Hardy, Thanks for reporting this. Could you try with the most recent version of Forge? 1.3.3.Final? Thanks! You can get it here: http://forge.jboss.org/ ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-1043
4613	FORGE-1034	Although JBIDE-15281 is fixed, I don't see how the version of Maven that is used in order to build JBoss Tools affects the behavior of this plugin. Can you tell me more?	1	1	https://issues.jboss.org/browse/FORGE-1034
4614	FORGE-1034	Ok, but I still don't understand the link between forge-maven-plugin and JBT build. How does one depend on the other?	1	1	https://issues.jboss.org/browse/FORGE-1034
4615	FORGE-1034	Same question as [~mickael_istria] - why is jboss tools forge plugin build affected by this ? And is requiring Maven 3.1 okey knowing it actually causes breakage with existing 3.0 builds ?	1	1	https://issues.jboss.org/browse/FORGE-1034
4616	FORGE-1034	As replied in the e-mail: {quote} > > why is the forge-maven-plugin relevant for the eclipse plugins build ? The forge-maven-plugin retrieves the required addons in order to make Forge 2 usable in Eclipse and copy them to an addon-repository folder that is created during build time. This is to prevent binaries being pushed to the forge git repository. There is an older version of the plugin already doing this but for some strange reason behind the classworld land, Aether 1.13.1 (used to fetch the declared addons) does not work inside a maven plugin targeted for 3.0.5. > And isn't requiring Mvn 3.1 in Forge 2 going to break other plugins ? I tested and the build works with Maven 3.1 AFTER Mickael upgraded the tycho plugin version. So that is not an issue afaik. > afaics the mvn 3.1 release is getting many reports about incompatibilities... > > how does that affect our JBDS 7.0/7.1 where mvn 3.0 still needs to work too ? That should probably require some additional testing. I tested the JBT Forge build and the only plugin that was preventing this to happen was the tycho I as beforementioned, which is now fixed. {quote}	1	1	https://issues.jboss.org/browse/FORGE-1034

#	Requirement	Questions	Answered	Successfully Answered	URL
4617	FORGE-1034	> The forge-maven-plugin retrieves the required addons in order to make Forge 2 usable in Eclipse and copy them to an addon-repository folder that is created during build time. Can you please provide a link to the place in jbosstools-forge repository where this plugin is used? I had a quick look at som pom files and didn't find a reference to this plugin [~maxandersen]: With the move to Tycho 0.18.1, JBT build supports both Maven 3.0.x and Maven 3.1.0. The only benefit so far is that is makes developers' life easier (they won't have to deal with multiple versions of Maven and can use the one of their choice to build JBT). Moving to Tycho 0.18.1 is a way to let [~gastaldi] and friends keep only a single version of Maven to build both JBT and forge-maven-plugin.	1	1	https://issues.jboss.org/browse/FORGE-1034
4618	FORGE-1032	What about 1.x ? Would that be interesting to have ?	0	0	https://issues.jboss.org/browse/FORGE-1032
4619	FORGE-946	The current implemented addons check if the UISelection returned by UIContext is null. If not null, then a selection is available (that is, uiSelection.size() is always >=1). Unless we change this behavior, I think this method is not necessary, what do you think ?	1	1	https://issues.jboss.org/browse/FORGE-946
4620	FORGE-946	It feels strange that UISelection can be null, and that it is also a container/iterable, don't you think? I think we should re-evaluate this, and remove the ability for UISelection to ever be null.	1	1	https://issues.jboss.org/browse/FORGE-946
4621	FORGE-946	Do you think that we should throw an exception when UISelection.get() is called and no selection exists or should return null silently ?	0	0	https://issues.jboss.org/browse/FORGE-946
4622	FORGE-944	[~lincolnthree] mentioned on : https://community.jboss.org/message/843635#843635 {quote} So far, our command style is following this basic premise: jpa-new-entity (instead of entity --named) jpa-new-embedable jpa-new-mapped-superclass --named Person jpa-new-entity --named Vet --extends Person (spec alias or acronym)-new-(something) (spec alias or acronym)-(something)-from-(something else) We have a dilemma with faces and CDI, though. Should we call them 'faces' and 'rest' or "jsf" and "jaxrs", since we already have "jpa" "ejb" etc. I think JAX-RS is the real issue, because most people know it as REST. {quote} Today, most of the commands follow the proposed pattern, except for JAX-RS (rest), JAX-WS (soap), JPA (entity, persistence, field), JASF (faces) and Bean Validation (constraint, validation) . So to make it more explicit in some edges, here is what we have now and what could happen || Specification || Today's commands || Tomorow's commands ? || | CDI | beans | beans | | Bean Validation | constraint, validation | constraint | | EJB | ejb | ejb | | JSF | faces | jsf | | JPA | field, persistence, entity | jpa | | JMS | jms | jms | | JTA | jta | jta | | JAX-RS | rest | rest | | Servlet | servlet | servlet | | JAX-WS | soap | soap | And some other might come in the future such as Batch...	1	1	https://issues.jboss.org/browse/FORGE-944
4623	FORGE-943	Hi [~rajonjava], I believe this information should only appear when {code} set VERBOSE true {code} is run, otherwise, it will be kinda annoying. Thoughts ?	0	0	https://issues.jboss.org/browse/FORGE-943
4624	FORGE-911	Hi Hanine, Which Embedded ID bug are you referring to ? Could you please link to this issue ? Thanks, George Gastaldi	1	1	https://issues.jboss.org/browse/FORGE-911
4625	FORGE-911	(MSC service thread 1-4) HHH000059: Defining hibernate.transaction.flush_before_completion=true ignored in HEM 15:00:28,267 ERROR [org.jboss.msc.service.fail] (MSC service thread 1-4) MSC00001: Failed to start service jboss.persistenceunit."prototype.war#forge-default": org.jboss.msc.service.StartException in service jboss.persistenceunit."prototype.war#forge-default": Failed to start service at org.jboss.msc.service.ServiceControllerImpl$StartTask.run(ServiceControllerImpl.java:1767) [jboss-msc-1.0.2.GA.jar:1.0.2.GA] at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) [rt.jar:1.7.0_11] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:603) [rt.jar:1.7.0_11] at java.lang.Thread.run(Thread.java:722) [rt.jar:1.7.0_11] Caused by: javax.persistence.PersistenceException: [PersistenceUnit: forge-default] Unable to build EntityManagerFactory at org.hibernate.ejb.Ejb3Configuration.buildEntityManagerFactory(Ejb3Configuration.java:914) at org.hibernate.ejb.Ejb3Configuration.buildEntityManagerFactory(Ejb3Configuration.java:889) at org.hibernate.ejb.HibernatePersistence.createContainerEntityManagerFactory(HibernatePersistence.java:73) at org.jboss.as.jpa.service.PersistenceUnitServiceImpl.createContainerEntityManagerFactory(PersistenceUnitServiceImpl.java:162) at org.jboss.as.jpa.service.PersistenceUnitServiceImpl.start(PersistenceUnitServiceImpl.java:85) at org.jboss.msc.service.ServiceControllerImpl$StartTask.startService(ServiceControllerImpl.java:1811) [jboss-msc-1.0.2.GA.jar:1.0.2.GA] at org.jboss.msc.service.ServiceControllerImpl$StartTask.run(ServiceControllerImpl.java:1746) [jboss-msc-1.0.2.GA.jar:1.0.2.GA] ... 3 more Caused by: org.hibernate.AnnotationException: A Foreign key refering org.otika.prototype.model.Expertise from org.otika.prototype.model.Book has the wrong number of column. should be 2 at org.hibernate.cfg.annotations.TableBinder.bindFk(TableBinder.java:432) at org.hibernate.cfg.annotations.CollectionBinder.bindManytoManyInverseFk(CollectionBinder.java:1502) at org.hibernate.cfg.annotations.CollectionBinder.bindManyToManySecondPass(CollectionBinder.java:1403) at org.hibernate.cfg.annotations.CollectionBinder.bindStarToManySecondPass(CollectionBinder.java:691) at org.hibernate.cfg.annotations.CollectionBinder$1.secondPass(CollectionBinder.java:626) at org.hibernate.cfg.CollectionSecondPass.doSecondPass(CollectionSecondPass.java:66) at org.hibernate.cfg.Configuration.originalSecondPassCompile(Configuration.java:1586) at org.hibernate.cfg.Configuration.secondPassCompile(Configuration.java:1359) at org.hibernate.cfg.Configuration.buildSessionFactory(Configuration.java:1724) at org.hibernate.ejb.EntityManagerFactoryImpl.<init>(EntityManagerFactoryImpl.java:84) at org.hibernate.ejb.Ejb3Configuration.buildEntityManagerFactory(Ejb3Configuration.java:904) ... 9 more 15:00:28,477 INFO [org.jboss.as.server] (management-handler-thread - 4) JBAS015870: Deploy of deployment "prototype.war" was rolled back with failure message {"JBAS014671: Failed services" => {"jboss.persistenceunit.\"prototype.war#forge-default\"" => "org.jboss.msc.service.StartException in service jboss.persistenceunit.\"prototype.war#forge-default\": Failed to start service"}} 15:00:28,489 INFO [org.jboss.as.server.deployment] (MSC service thread 1-10) JBAS015877: Stopped deployment prototype.war in 12ms	1	1	https://issues.jboss.org/browse/FORGE-911
4626	FORGE-898	Would auto-complete of *org.jboss.forge.addon* in the groupId be sufficient in this case ?	1	1	https://issues.jboss.org/browse/FORGE-898
4627	FORGE-847	Good idea. This shouldn't be so hard to do. Feel like taking a shot at it, actually?	0	0	https://issues.jboss.org/browse/FORGE-847
4628	FORGE-835	Hey George, Thanks for mentioning that! I did a quick research and found that on some databases JPA creates the column with different type based on the 'length' attribute of the @Column annotation. Should we always set it to MAX_VALUE or we should let the user configure it with an option (defaulting to MAX_VALUE)? Something like: field lob --named <field-name> [--length <length>] Cheers, Ivan	1	1	https://issues.jboss.org/browse/FORGE-835
4629	FORGE-834	Thanks ! Unfortunately, I am not able to make it works with my Qooxdoo plugin ! I have re-compiled/installed the forge project (mvn clean install) with my changes (version 1.2.3-SNAPSHOT). Basically, it can not find the "org.jboss.forge.spec.javaee.rest.events.RestGeneratedResources". When I install the local forge distribution, run forge, go to plugin dir and run the forge source-plugin . command, I got an error: [ERROR] ... package org.jboss.forge.spec.javaee.rest.events does not exist Did I miss something ? Charles.	1	1	https://issues.jboss.org/browse/FORGE-834

#	Requirement	Questions	Answered	Successfully Answered	URL
4630	FORGE-815	Want to try this one on? Should be simple.	1	1	https://issues.jboss.org/browse/FORGE-815
4631	FORGE-811	Shouldn't this be added to the project plugin? Now there is a 'project add-dependency groupId:artifactId' already.	0	0	https://issues.jboss.org/browse/FORGE-811
4632	FORGE-806	We could however, introduce a "resetOutput" method in the AbstractShellTest class. would that be enough ?	0	0	https://issues.jboss.org/browse/FORGE-806
4633	FORGE-773	You are correct. The reflection part is not hard. I've done it mostly already. I implemented JavaClass using reflection. Method, Field, etc... also. URLClassloader can be used to load all the project's dependency jars. I don't think that is big either. Of course I wonder what the response time will be like. I think the big change is for ScaffoldPlugin to be changed so generateFromEntity takes a class name as string, not a JavaResource, alleviating ShellImpl from making the choice to get a class or a java resource; the plugin will find either the class on the dependency path or the java resource; and for Dependency plugin to give over a list of the compile time dependencies of the project. Are these kinds of changes to the forge system acceptable?	1	0	https://issues.jboss.org/browse/FORGE-773
4634	FORGE-773	Hey John! That's excellent. A working prototype gives me hope. As far as changing the scaffold plugin, it's fine to add to the API, but we can't remove or change anything that we have right now. So it's possible we could simply add another option / command in the ScaffoldPlugin that would kick things off using your reflection loading system. An issue that might be related here is the following: https://issues.jboss.org/browse/FORGE-563 Since you are already building an implementation of the JavaClass API on top of Reflection, would it be conceivable to use that information to create a type of cache for each object inspected in this way? The index could be parsed on-demand, and enable retrieving the information needed for the scaffold. It would also enable wildcard expansion in the JavaParser itself -> https://issues.jboss.org/browse/FORGE-424 Thoughts? Would enhancing the scaffold provider meet your goals?	1	1	https://issues.jboss.org/browse/FORGE-773
4635	FORGE-773	Yes, such an enhancement will meet the goal. I see the maven plugin/facet is using Aether, so I believe I can reliably obtain the dependency's jar file, though the applicable forge api is not known to me yet. I understand the generality to an interface that supports queries for a JavaClass(es) backed by a binary class. Such would be useful in the issues you mention. A more general implementation returns a JavaClass independent of whether it is found as a binary class or in the project's source directory. Does forge already have such an interface that returns JavaClass parsed from JavaSource?	1	0	https://issues.jboss.org/browse/FORGE-773
4636	FORGE-770	Here are the comments from David Bosschaert about OSGi versioning: {quote} Let's say you have a component that wants to use the org.acme.Foo service. Let's say you have 2 versions of the Foo service in your system: // version 1.0 interface Foo { public void bar(); } // version 2.0 interface Foo { public void bar(); public void bar(int blah); } Each of the above would be in a separate module. If I want to use the bar(int) API I clearly need version 2.0 of the API. So such a client would import org.acme at version range [2.0, 3.0), so it has visibility of the correct API. When looking up services they are normally looked up by API name, so org.acme.Foo. If there is both an 1.0 and 2.0 service in the system, how do you pick the right one? How does the service registry decide which one to return? Since it knows the client that is requesting the service it will check whether the interface requested as visible by the client is the same one as the interface implemented by the service. If it is not the same then it is a different version of the interface and the service should not be returned to the client. If it is assignable then we're good and the client gets the service instance. This is how OSGi frameworks do it, and it's actually quite a simple process, as the Class.isAssignableFrom() API gives you the tool you need to do it :) If you want automated support for installing additional addons when required you really need some sort of a repository that has all the addons and exposes metadata about what they provide. Then you need a resolver that knows how to use this information and installs the addons on request. This is potentially a (very) large piece of work. Especially resolvers are complex and can take a lot of time to write. The OSGi Enterprise R5 spec [1] and [2] defines a Repository spec that can be used to catalogue your addons in a machine-readable way. Then, the OSGi Resolver spec (also in [1] and [2]) implements a general purpose resolver that can work with the Repository to install just the additional modules that you want. If you want to play with this stuff the Reference Implementations are available as opensource, see here for more details: http://en.wikipedia.org/wiki/OSGi_Specification_Implementations I would think twice before implementing all this stuff, as it might seem simple when you start out but it could end up being more work than you think. OSGi implements all these things already. If you're inside AS8 then JBoss OSGi provides all this. If you are in a different environment you could simply use Apache Felix, which is about 500 kb and very embeddable. Cheers, David [1] http://www.osgi.org/Download/Release5 [2] http://www.osgi.org/javadoc/r5/enterprise/ {quote}	1	1	https://issues.jboss.org/browse/FORGE-770
4637	FORGE-766	Isn't this resolved? org.jboss.forge.maven.plugins.MavenPluginBuilder?	1	1	https://issues.jboss.org/browse/FORGE-766
4638	FORGE-734	Not sure I agree this makes sense. Capable is a more confusing term in this case, I think. Target is very explicit about "This thing can be the target of an annotation." Capable seems more ... loose. What does it do with annotations? It is capable of something with them, but not exactly sure what.	1	0	https://issues.jboss.org/browse/FORGE-734
4639	FORGE-734	There exists InterfaceCapable and GenericCapable right now. AFAIK, the operations on each of them are equivalent (add/remove/get). Is this concept different anyhow ?	0	0	https://issues.jboss.org/browse/FORGE-734
4640	FORGE-719	My biggest gripe/concern with the current Forge approach is that it is all very "(re)action based". i.e. plugins seem to follow the pattern of "ask user question 1,2,3,4 (based of the cmd annotations)", do something - realize something was not possible to ask about via cmd annotations "ask user question 2,3,4", do something, "ask user some more questions", do things. This is all great in a cmd-line world but really really hard for newbie users to grok (IMO) and when you actually have to do or combine multiple forge level operations things becomes really hard from an IDE user perspective - i.e. IDE users expect to be able to just press Cancel in a multi-step wizard and not see any change in his project from that, also when driving from the IDE, we'll need to be able to call out to Forge and query for the possible values etc. (something the current 'single Is this issue about solving/improving this challenge or "just" about providing an UI on top of the existing "cmd line/reaction pattern" ? Things to consider here is how forge can provide rich enough info to have "Browse.." buttons work for file/resource, database resource selection, etc. And how a multistep wizard could be created outlining pages needed to provide the questions ...and in a way that does not require fully running forge inside something like Eclipse.	1	1	https://issues.jboss.org/browse/FORGE-719
4641	FORGE-707	Is this already fixed?	1	1	https://issues.jboss.org/browse/FORGE-707

#	Requirement	Questions	Answered	Successfully Answered	URL
4642	FORGE-705	(03:31:20 PM) lincolnthree: kenfinnigan: https://github.com/forge/core/tree/2.0/plugin-container-api/src/main/java/org/jboss/forge/container (03:31:47 PM) lincolnthree: You'll want to probably figure out how people are going to interact with the Resource API (03:32:06 PM) lincolnthree: In Forge 1.x, you had to @Inject the ResourceFactory (03:32:13 PM) lincolnthree: That may still be the case in Forge 2.0 (03:32:18 PM) gastaldi: yeah (03:32:20 PM) lincolnthree: But just something to consider. (03:32:27 PM) lincolnthree: Something will need to set up the ResourceFactory (03:32:47 PM) lincolnthree: And if you need to do that, you'll probably want to @Observes Startup event. (03:33:02 PM) gastaldi: kenfinnigan: see the Aesh Addon example (03:33:11 PM) gastaldi: there are some stuff in there to help (03:33:41 PM) gastaldi: like the @Observes Startup event for instance (03:34:22 PM) kenfinnigan: so would resources api be in forge-resource-addon module under core? (03:34:46 PM) kenfinnigan: ie. create that module, or should the code live somewhere else? (03:34:56 PM) lincolnthree: yeah, create that module (03:35:01 PM) lincolnthree: for now we can put it all in the same repo (03:35:06 PM) lincolnthree: but it should eventually be split up	1	1	https://issues.jboss.org/browse/FORGE-705
4643	FORGE-694	Is this the same as https://issues.jboss.org/browse/FORGEPLUGINS-54 ?	1	1	https://issues.jboss.org/browse/FORGE-694
4644	FORGE-694	Hi Darren, You're right. I guess that FORGEPLUGINS-54 needed a better description after all. Could you please add your comments to that issue ? As a thumb of rule, we always create new issues related to plugins (new ideas, etc) in the FORGEPLUGINS project. The FORGE project in JIRA is related to core features only. :) Much appreciated ! George Gastaldi	1	1	https://issues.jboss.org/browse/FORGE-694
4645	FORGE-670	OK, I touched this as well :-) Maybe some hints would be helpful. So, FSHParser::captureLogicalStatement is the place, where we may detect that a redirect was issued. So there is maybe the place to somehow tell the LogicalStatement to redirect its output somewhere else. Just speculating... Next, I see also that there is a PipeOut interface that is passed around all the methods in the trace. But it is usually null and is never used inside org.jboss.forge.shell.command.Execution::perform. I see that this interface has just one implementation (org.jboss.forge.shell.util.PipeOutImpl), which uses the shell as output. However, the fact that it is never used inside Execution::perform makes me wonder whether creating a new implementation of PipeOut (RedirectPipeOutImpl for example) and passing it to that method would work at all. Can someone open my eyes for this Forge internals? :-) Thanks! Ivan	1	1	https://issues.jboss.org/browse/FORGE-670
4646	FORGE-664	In this case, I think using System.out.println() would work just fine. As long as your logger is also printing to STDOUT. However, I don't think this is really encouraged. The whole point of the Shell is that we can control this through a single API. I'm not sure this is something we should do - unless we create our own category, but before we do this, I suppose I should ask: Why do you need this? Maybe there is an alternative, or maybe we actually need it :) but we should discuss why you want to in the first place.	1	1	https://issues.jboss.org/browse/FORGE-664
4647	FORGE-652	Hi Jose, Any work on this ? If not, can we add it back to pool ? Thanks !	0	0	https://issues.jboss.org/browse/FORGE-652
4648	FORGE-616	Hey George! This is really really great. 1 comment, 1 question: Comment: Forge should ask the user if they would like to install before going ahead. Looks like it automatically does this right now? Question: Have you checked for other locations where Status.SUCCESS / Status.FAILURE to make sure that adding another enum won't cause issues? I ask because MISSING should probably behave like FAILURE if not otherwise specified. Thoughts?	1	1	https://issues.jboss.org/browse/FORGE-616
4649	FORGE-616	Hi Lincoln, I totally agree with your comment. Will implement that today. Also, I checked for other locations where the Status enum is used and is not used anywhere besides the Execution class and ForgePlugin itself. There are very few plugins which observes this CommandExecuted event and deals with the status. One of them is the Undo plugin, which Jevgeni already checks for it ("Create a commit only if Status == SUCESSS"). Anyway, I thought about that and I found out that it's a matter of interpretation: Should Status.FAILURE indicate a failure in the plugin execution or it can be used to indicate that a plugin is missing / wrongly typed ? The best way to detect that was with this new MISSING enum afaik. If you feel the need to change that, we can do it, no problem. Just let me know your thoughts about it. Best Regards, George Gastaldi	1	0	https://issues.jboss.org/browse/FORGE-616
4650	FORGE-604	Hey Chris! Thanks so much for reporting. Do you think you might be able to add a brief description of what Forge does incorrectly when operating on an EAR? Also perhaps attach a sample project that we can reproduce the functionality with? I know this is asking a lot, but it would really help us get things working! Thanks, Lincoln	1	1	https://issues.jboss.org/browse/FORGE-604
4651	FORGE-604	Hi Christopher, Forge 1.2.4.Final now supports EAR packaging as described in FORGE-864. Could you check if that is sufficient to close this issue ? Best Regards, George Gastaldi	1	1	https://issues.jboss.org/browse/FORGE-604
4652	FORGE-602	Isn't that already resolved ?	0	0	https://issues.jboss.org/browse/FORGE-602
4653	FORGE-601	Ryan, Can you provide us with a JUnit test case ? Best Regards, George Gastaldi	1	0	https://issues.jboss.org/browse/FORGE-601
4654	FORGE-598	Hey Rafael! I think you can do this in a subtle way: extend SimpleTokenCompleter @Inject ShellPrintWriter writer; writer.renderColor(ShellColor.WHITE, "CANDIDATE3"); Return candidates in the order in which you want them displayed. But we could certainly improve this experience, I think! Let me know if this works?	1	0	https://issues.jboss.org/browse/FORGE-598
4655	FORGE-588	Felipe, Which JDK version are you using ? Can you try using JDK 7 or add the root certificate from https://github.com to your cacerts keystore ? Thanks.	1	1	https://issues.jboss.org/browse/FORGE-588
4656	FORGE-583	We can also add PDF of wiki as well. something like JBoss-Forge-UserGuide-<version>.pdf or something similar. User can read and learn from the UserGuide. Make sure this must be updated and added in each release.	0	0	https://issues.jboss.org/browse/FORGE-583
4657	FORGE-580	Maybe it's a stupid question but what's wrong with Apache License? Is there a specific reason why Apache License doesn't work for Forge anymore? P.S. I always thought Forge was under GPL and not Apache. Source code headers I've seen only mention GPL. Btw, it might be a good idea to add forge license information here (empty now) : http://www.jboss.org/projects	1	1	https://issues.jboss.org/browse/FORGE-580
4658	FORGE-580	Which files are saying GPL ? That would be bad if it was ;)	1	1	https://issues.jboss.org/browse/FORGE-580
4659	FORGE-565	Isn't this issue already fixed ?	1	1	https://issues.jboss.org/browse/FORGE-565
4660	FORGE-561	Hi there :) Could you please describe the problem you are having in a little more detail? Also, if you could post the entity you used to generate scaffolding from, and the resultant xhtml files, that would be of great assistance in figuring out this problem. Thanks, Lincoln	1	0	https://issues.jboss.org/browse/FORGE-561

#	Requirement	Questions	Answered	Successfully Answered	URL
4661	FORGE-558	George, That's cool. I haven't used Twitter bootstrap so apologies if I'm speaking out of turn. Just so long as the developer doesn't look at the generated project and say 'what is this non-standard thing here? Oh it's Twitter bootstrap. I guess I'll have to go grok that project before I can edit this generated code'. IMHO this would apply even if Twitter bootstrap is just a bunch of CSS/JS/JQuery stuff. It's still stuff that the developer may not necessarily want in their end project. But it's your call. I only mention it as 'no external libraries' was an explicit requirement of the original UI scaffolding. And I'd think this would extend to the JS side too. Regards, Richard.	1	1	https://issues.jboss.org/browse/FORGE-558
4662	FORGE-558	That´s a valid point indeed. In fact, I think I didn´t take in count that requirement. Maybe we should discuss that on forge-dev list ? WDYT ? Regards, George	1	1	https://issues.jboss.org/browse/FORGE-558
4663	FORGE-508	Hey Pete, currently the as7 plugin just adds the jboss-as-maven-plugin to the build section of the project pom. Currently the version is hardcoded to be 7.1.1.Final. There is no further implication when setting up the plugin, i.e. the jboss-javaee-6.0 BOM is set up independently when setting up persistence. Am I right in assuming you want the persistence plugin and the as7 plugin to play together so that this is somehow automatically inferred? If yes, is there an easy way to find out which jboss-javaee-6.0 version is provided by each as7 version?	0	0	https://issues.jboss.org/browse/FORGE-508
4664	FORGE-508	Yes this makes sense, it is along the lines of my first comment. But the question of figuring out which BOM version is needed by each AS version still remains. Is this information available anywhere?	0	0	https://issues.jboss.org/browse/FORGE-508
4665	FORGE-508	Max and I had a plan to store this data in one place, and then consume it from all tools that need it (e.g. jdf site, JBDS, Forge etc etc.) My approach for this would be like Lincoln's for Forge plugins - https://github.com/forge/plugin-repository/blob/master/repository.yaml - a yaml file stored on git, that we update as needed, probably along with a utility class that people can copy into their project to parse and provide as an object structure, maybe along with simple queries. WDYT?	0	0	https://issues.jboss.org/browse/FORGE-508
4666	FORGE-508	That sounds like a good solution indeed. Where do you think this yaml file and utility should sit?	0	0	https://issues.jboss.org/browse/FORGE-508
4667	FORGE-505	No core plugins should be adding arquillian. Could you explain? Thanks.	1	0	https://issues.jboss.org/browse/FORGE-505
4668	FORGE-504	Could you provide a link to the BOM you want used? Thanks.	0	0	https://issues.jboss.org/browse/FORGE-504
4669	FORGE-504	Now I'm confused. Why do we have so many POMs for Hibernate? Forge currently uses the one from teh EE 6 BOM.	1	1	https://issues.jboss.org/browse/FORGE-504
4670	FORGE-504	Pete, is this still valid ? Or is the JDF plugin handling this ?	1	0	https://issues.jboss.org/browse/FORGE-504
4671	FORGE-500	Hey Burr, could you please explain this in a bit more detail? What does this flag do? Where should it be located in Forge? How should users access or control it? Thanks.	0	0	https://issues.jboss.org/browse/FORGE-500
4672	FORGE-488	Hi. Could you please explain this? Thanks.	1	0	https://issues.jboss.org/browse/FORGE-488
4673	FORGE-451	Should there be a prompt asking for the generation of these methods or only a warning when these methods are absent? Isn't this only facesScaffold specific?	1	1	https://issues.jboss.org/browse/FORGE-451
4674	FORGE-451	I think you could probably just add a checkbox to ScaffoldableEntitySelectionWizard to add missing equals and hashcode methods (if the object does not already have them, of course.) WDYT?	1	1	https://issues.jboss.org/browse/FORGE-451
4675	FORGE-451	Matej, is this issue complete?	1	1	https://issues.jboss.org/browse/FORGE-451
4676	FORGE-447	Hey Pete, Do you think you could provide a little bit more technical detail on what you want to have happen here? Perhaps some proposed JSF code that would do what you ask? That way we can have a better idea of exactly how you think this should work. I'm personally not quite following you yet. ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-447
4677	FORGE-422	Hey Pablo! Any update on this one? :) How's it going?	1	1	https://issues.jboss.org/browse/FORGE-422
4678	FORGE-422	Hey Linconln!!, I'm so sorry for no update it before!!. I was few busy, becaming a new job. So please, can you let me some days more?. Maybe, you would like to add this feature to forge CR1, wouldn't you?	1	1	https://issues.jboss.org/browse/FORGE-422
4679	FORGE-422	Hi, what's the status of this issue?	1	1	https://issues.jboss.org/browse/FORGE-422
4680	FORGE-422	Hi. Again, I'd like to ask whether it's possible to provide option "evaluate" in the next Forge release. I want to create automatized tests for RichFaces plugin, but it's impossible without this option. If there is nobody from Forge developers who is interested in this issue, would it be possible to navigate me to right places (java classes), so I can implement it on my own and send pull request? In this case It will take me a huge amount of time, because I don't know Forge structure.	1	1	https://issues.jboss.org/browse/FORGE-422
4681	FORGE-422	Maybe a stupid question, but where can I find the version providing the "evaluate" option? I've tried to download 1.0.3-SNAPSHOT (from JBoss repository), but it doesn't work. {code} [jpapouse@dhcp-27-156 tmp]$ forge -e "echo hi jan" Using Forge at /opt/forge-1.0.3-SNAPSHOT _____ | ___|__ _ __ __ _ ___ | |_/ _ | `__/ _` |/ _ \ \\ | _| (_) | | | (_| | __/ // |_| ___/|_| __, |___| |___/ [no project] tmp $ {code}	1	1	https://issues.jboss.org/browse/FORGE-422
4682	FORGE-421	Hi, We currently don't support Proxy configuration by any specific forge feature, but it's definitely something we need! Is this something you could look at? It's relatively simple, I think. Just need to set some system properties: http://docs.oracle.com/javase/6/docs/technotes/guides/net/proxies.html Otherwise, this could be done by setting FORGE_OPTS in your terminal environment before executing Forge itself (or adding them to your profile or windows settings): unix: {code}export FORGE_OPTS="-Dhttp.proxyHost=webcache.mydomain.com -Dhttp.proxyPort=8080"{code} windows: {code}set FORGE_OPTS="-Dhttp.proxyHost=webcache.mydomain.com -Dhttp.proxyPort=8080"{code} Let me know if this works? (The values will need to be chanced to match your proxy configuration.)	1	1	https://issues.jboss.org/browse/FORGE-421
4683	FORGE-421	Hi, Unfortunately I couldn't succeed as well while behind the proxy. My error was even less verbose: Connecting to remote repository https://raw.github.com/forge/plugin-repository/master/repository.yaml... **ERROR** [forge install-plugin] raw.github.com Does anyone know where can I read some logs or see some detailed exception traces? I saw a log file under ~/.forge, but it did not contain anything particularly interesting. Maybe I could somehow increase the logging severity? Cheers, Ivan	1	1	https://issues.jboss.org/browse/FORGE-421
4684	FORGE-421	Hey Ivan, I'm not sure if you can increase logging, but we don't have any explicit support for proxies at the moment. Are you looking for something to pick up? This could be a good issue for you since you can actually reproduce the problem. What do you think? It would really help us out a lot! ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-421
4685	FORGE-421	Hey Ivan! I've assigned this issue and given you permissions for future ones as well :) Sorry about that! In order to get more information, you might have to Debug Forge like so: https://docs.jboss.org/author/display/FORGE/Debugging+Forge Does that help?	1	1	https://issues.jboss.org/browse/FORGE-421

#	Requirement	Questions	Answered	Successfully Answered	URL
4686	FORGE-421	Hi guys, I have a question. Right now I found a way to setup programatically Apache HTTP client and JGit to work with a proxy. And this is done with just Forge instruments (no additional manual steps than configuring Forge). However, I just found that if you want to use maven behind a proxy (and Forge uses it for building the just downloaded plugin), then you have to change your settings.xml (http://maven.apache.org/guides/mini/guide-proxies.html). So, I had the following dilemma: maybe we should tell the forge users behind proxies that they should setup their settings.xml so that Forge can access external Maven repos? Argh, this sounds nasty: I hate manual configurations done by users!!! However, when given a second thought: you are behind a proxy... And you want to use open source software, that is most likely not located on your enterprise maven repository. Then you should definitely setup your settings.xml to access external sites! Otherwise you can live just with what you find on your www.proprietary.com intranet... Do you think the same or I miss something? Can we rely on the fact that Forge users that want to use our tool from within corporate environments have already setup their settings.xml to use their company proxy? Cheers, Ivan	1	1	https://issues.jboss.org/browse/FORGE-421
4687	FORGE-421	Hey Ivan, Sorry for the late reply, on vacation. We do some configuration of the maven settings internally in Forge as well. This should be a good place to try to avoid duplicate configuration for users. Thoughts? Can you do what is needed at runtime in Forge? ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-421
4688	FORGE-421	Hey Lincoln, How was your vacation? Did you drink plenty of Alt and Kolsch? :-) I also tried to exploit the Forge capabilities for configuring Maven. However, besides passing some parameters, like for example system properties, on the command line, I did find anything else. And to enable proxies in Maven, you need to touch settings.xml. Or at least this is what all the resources I found say (I actually tried some proposals for adapting MAVEN_OPTS, but they did not work). What I can do is download Maven source code and look there what exactly happens. Or maybe you can point me to the Forge code that manipulates Maven, as maybe I have missed something? Cheers, Ivan	1	1	https://issues.jboss.org/browse/FORGE-421
4689	FORGE-421	Have you tried integrating the settings at this point? This allows direct access to set system properties. https://github.com/forge/core/blob/master/project-model-maven/src/main/java/org/jboss/forge/maven/facets/MavenContainer.java I believe some proxy work is already done there? Also, any plans on adding username/password support to config.xml? ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-421
4690	FORGE-421	Hi Lincoln, Thanks for the idea of using the Maven Container! However, I'm not yet sure that the MavenContainer instance has any impact on the executeMaven method, that is eventually run upon building the plugin. As for the username/password support - thanks for reminding me! I thought of adding them, but then I forgot. It will not be a big deal... However, I'm afraid I will not be able to test it, because my company proxy does not use such stuff. Thanks also for merging my changes! They seem to be all there, even an experiment that I did some time ago and was not supposed to go into the upstream (it works, though). I am still afraid of using Git :-) BTW, something off topic. Do you have a good resource on how should we use Git when working with a Github project? I use mainly this: https://community.jboss.org/wiki/ArquillianDevelopmentAndContribution However, it does not answer an important question: how do I keep my origin repository (ivannov/core) up-to-date with the upstream one (forge/core). I tried to do that by pushing to ivannov/core after I pull from forge/core, but you saw what a mess I produced with that approach... Thanks, Ivan	1	1	https://issues.jboss.org/browse/FORGE-421
4691	FORGE-421	Hi, No luck. I've tried the solution but errors still persist. I am using the last version: JBoss Forge, version [1.0.2.Final] - JBoss, by Red Hat, Inc. [http://jboss.org/forge] The error is as follows: Connecting to remote repository [https://raw.github.com/forge/plugin-repository/master/repository.yaml]... ***ERROR*** [forge install-plugin] Connection to https://raw.github.com refused Can I get more stacktrace to upload? Regards	1	1	https://issues.jboss.org/browse/FORGE-421
4692	FORGE-398	Do you want this on all owning side relationships? Or all relationships as a whole?	1	1	https://issues.jboss.org/browse/FORGE-398
4693	FORGE-398	Update. Only needed on on owning side of @OneToOne relationships, but do we need this in other places as well? Why does @ManyToMany not need this on both sides of the relationship? How do we know which side to put it on? Both sides maybe? It's possible that the UI scaffolding may need to update the entity manually for its own needs, instead of doing this in the EntityPlugin.	1	1	https://issues.jboss.org/browse/FORGE-398
4694	FORGE-389	Can we please improve this a litte? At the moment it outputs: public String toString() { return "" + firstName + ", " + lastName + ", " + city + ", " + state + ", " + zip; } I'm not a big fan of the commas, but worse is that it outputs them even if the value is empty. So you end up with: "John, Smith, , , , " We need to: 1. definitely not append a field if it is empty 2. maybe not append commas Then we'd have: "John Smith"	1	1	https://issues.jboss.org/browse/FORGE-389
4695	FORGE-388	Hi Senthy, Right now Forge does not comprehend settings in multi-module projects. This is something we would like to support better. Do you have any ideas on how we can support non-trivial projects better? We do already support parent POM and dependency inheritance, but we don't do things like detect persistence.xml files in dependencies. Thoughts? ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-388
4696	FORGE-388	According to latest comments in FORGE-120 I see how to create "pom" type of project. Obviously the same way I can create "war" or "jar"projects. I can run something like {code} new-project --named foo --topLevelPackage com.bar --type pom new-project --named foo-data --topLevelPackage com.bar --type jar new-project --named foo-web --topLevelPackage com.bar --type war {code} but it is still needed to setup dependencies between sub-projects (modules) manually. In is possible to add some helper commands like {code} project add-subproject --named foo-data --type jar {code} But frankly speaking this is not really important. Instead we can have a simple one-page official documentation that explains people how to setup such projects. I would prefer having more documentation and less commands :-). Could you please add some documentation which covers this? Other issue is that is some cases in order to use some functionality you first needs to "setup" something. For example, in order to create new entity classes you need to run {code} persistence setup ... {code} It obviously makes seance. But if you really want to support multi-module projects this approach will not work. Default behavior should stay the same. But I propose to introduce alternative way to add "persistence awareness" to the project through plugins. Something like {code} forge install-plugin jpa {code} This "new" plugin will not add persistence.xml to the project but will rather add proper dependencies to pom.xml and will enable possibility to work with entities using seam commands. Same approach can be used to other plugins that require setup before usage. Another example I want to talk about is the automatic discovery functionality. For example {code} scaffold from-entity {code} and "tab" gives the list of all available entities. Forgive me for not trying this by myself but I had some issues last time when I tried to setup this. I am not sure if this list if only from "current" project or forge does scans also jars in dependencies to check if they have appropriate classes. But anyway the most important thing is some official documentation or guide how to work with multiple modules using seam. There is no need to support everything or all possible features. I am absolutely positive that even very limited (but documented and supported :-)) multi-module support would very helpful and absolutely ok in most cases. Thanks, Serhiy	1	1	https://issues.jboss.org/browse/FORGE-388

#	Requirement	Questions	Answered	Successfully Answered	URL
4697	FORGE-388	Maybe I am missing something but I was not able to find mentioned options in the prompt. Here is my console output (can't attache screenshot because issue is closed) {code} D:\tmp\forge-distribution-1.0.0.Beta3\forge-1.0.0.Beta3\bin>forge ____ _____ / ___| ___ __ _ _ __ __ ___	___|__ _ __ __ _ ___ ___ \ / _ \/ _`	'_ ` _ \	|_ / _ | '__/ _` |/ _ \ \\ ___)	__/ (_|								_| (_)			(_|	__/ // |____/ ___|__,_|_| |_| |_| |_| |_| ___/|_| __, |___| |___/ Windows? Really? Okay... [no project] bin $ new-project --named foo --topLevelPackage com.bar --type pom ? Use [D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo] as project directory? [Y/n] y ***SUCCESS*** Created project [foo] in new working directory [D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo] Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/pom.xml [foo] foo $ new-project --named foo-core --topLevelPackage com.bar ? Use [D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core] as project directory? [Y/n] y ***SUCCESS*** Created project [foo-core] in new working directory [D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core] Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core/pom.xml Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core/src/main/java Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core/src/test/java Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core/src/main/resources Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core/src/test/resources Wrote D:\/tmp/forge-distribution-1.0.0.Beta3/forge-1.0.0.Beta3/bin/foo/foo-core/src/main/resources/META-INF/forge.xml [foo-core] foo-core $ {code} I was asked only about project directory. And after I've added second project pom.xml of "foo" project is following {code} <?xml version="1.0" encoding="UTF-8"?> <project xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd" xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"> <modelVersion>4.0.0</modelVersion> <groupId>com.bar</groupId> <artifactId>foo</artifactId> <version>1.0.0-SNAPSHOT</version> <packaging>pom</packaging> <repositories> <repository> <id>JBOSS_NEXUS</id> <url>http://repository.jboss.org/nexus/content/groups/public</url> </repository> </repositories> <build> <finalName>foo</finalName> </build> </project> {code} So it is missing {code} <modules> <module>foo-core</module> </modules> {code} and respectively pom.xml of "foo-core" is missing "parent" definition. Do I need to add some special parameters to the commands? I really don't want to bother you but I can't find mentioned options or any documentation how to work with them. Thanks, Serhiy	1	0	https://issues.jboss.org/browse/FORGE-388
4698	FORGE-386	Thanks for doing this! But the version implemented is broken, I think. It basically says 'all unsaved entities will have the same hashCode' and 'all unsaved entities are equivalent'. It may even NullPointer? I would rather it delegate to super.hashCode and super.equals if (id == null). The same as my version above. But up to you.	1	1	https://issues.jboss.org/browse/FORGE-386
4699	FORGE-380	Hey Pete. This already works using the ';' character. Is that good enough? && already has a meaning in FSH (logical and)	0	0	https://issues.jboss.org/browse/FORGE-380
4700	FORGE-378	I agree with [~agoncal] that it should be possible to control JAVAEE separately from CONTAINER. I kinda see it as CONTAINER will specify the defaults for things like JAVAEE if user have not explicitly set it. With respect to referencing JavaEE api's instead of Java EE BOM then that was at least with Java EE 6 problematic since the referenced jars were incomplete thus making test runs using these jars impossible. This issue should be fixed in Java EE 7 (http://developinjava.com/2013/06/java-ee-7-fixes-java-ee-6s-absent-code/) but still, some of the spec jars require specific container content (i.e. JPA) thus something like using JBoss BOM still are necessary. Also reason for using BOM's is so you don't have to update all your projects individual dependency versions and just use properties provided by the BOM.	1	0	https://issues.jboss.org/browse/FORGE-378
4701	FORGE-378	What would happen if someone starts working with a specific version of EE, then sets the "environment" to a different version? While I agree with this concept in principal, I don't think that this should be the primary method of exposing it. As I explained in an email to Forge-Dev: Our current plan is this - https://github.com/forge/core/blob/2.0/projects/api/src/main/java/org/jboss/forge/addon/projects/ProjectType.java This is used in the new-project command like so: `new-project --named "blah" --type javaee-7 --version 1.0 --topLevelPackage org.blah` Now, the interesting thing is that depending on the project type, we can also specify additional parameters that will become activated: `new-project --named "blah" --type javaee-7 --version 1.0 --topLevelPackage org.blah --profile web/full` So we can do web profile selection as part of the javaee7 project type, or we could just have multiple project types for each profile. It could go either way.	1	0	https://issues.jboss.org/browse/FORGE-378
4702	FORGE-378	"What would happen if someone starts working with a specific version of EE, then sets the "environment" to a different version?" - you mean like you target JavaEE6 explicitly and then targets Wildfly 8 (which would default to JavaEE7?) I would just expect forge to ask something like: Current project is targeting JavaEE-6, the default type for Wildfly 8 is JavaEE-7. Do you wish to use JavaEE-7 or continue to use JavaEE-6 ?	0	0	https://issues.jboss.org/browse/FORGE-378
4703	FORGE-374	Okay cool. Could you please push Seam Render 1.0.0.Alpha5 into Maven? The Forge build is broken at the moment I think.	0	0	https://issues.jboss.org/browse/FORGE-374
4704	FORGE-372	Sorry for the late reply. I have been traveling. This is a valid concern. But is it enough just to delete one of the EntityManager Producers? It's not a simple yes or no because this deals with how plugins interact on the project together, and that complicates things. I'm not sure what is more correct, but I am leaning toward agreeing. Because this is part of the scaffold, which is an extension point for frameworks themselves, I think this would be OK. Is this something you're interested in prototyping? That would help get it done a lot sooner :) ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-372
4705	FORGE-372	Hi, In fact, I've test several solutions in the last few days and find some solutions. The main question is : do you want Forge plugins to use the EntityManager provided by Forge (in its package) or the "default" EntityManager of the application (define by the user) ? The first one is probably safer since you know exactly how the EntityManager is parametrized, but the second one seems more logic since it's the user's application, so if he wants an extended Seam EntityManager, good for him, the plugins should use it too (hoping all will be good). In the first case, easy to do : add a qualifier on the producer of the Forge EntityManager, add it also on the @Inject point, and you're done. In the second case, I have several propositions : 1) Delete the DatasourceProducer class in your package and generate it when the user initialize the project : in this configuration, the user can override it or delete it without any problem. A new command in the console like "restore datasourceProducer" would delete and re-create DatasourceProducer class with the default configuration provide by Forge. 2) Keep the DatasourceProducer class in your package but add an @Alternative annotation in the producer and put it in the beans.xml alternative list. So the day the user want to use it's own EntityManager, he just has to edit the beans.xml and remove the correct line. 3) There is also a solution by adding an @Veto annotation at runtime with Seam Solder but since, perhaps, the project does not require Solder, it's not good enough. In any case, you have to generate something in the user project in order to let him change it when he needs to. Hope my french-english will not kill your eyes.	0	0	https://issues.jboss.org/browse/FORGE-372

#	Requirement	Questions	Answered	Successfully Answered	URL
4706	FORGE-372	Hey Daniel, Sorry for the huge delay in getting you something else to chew on! Is this something you'd be interested in taking a shot at? ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-372
4707	FORGE-372	Hi, I think the second case with the second propositions sounds most reasonable to me. I would go with this solution. Is everybody in with me so I can go and realize it? Cheers, Daniel	1	1	https://issues.jboss.org/browse/FORGE-372
4708	FORGE-348	Actually I'd rather just go with metawidget-all. It's far cleaner for those learning Metawidget (just a single entry in the POM) and easier on Seam Forge too. I'd be much more interested in getting the AS 7 issue resolved, because this will impact all my tutorials and examples too. They all use metawidget-all, because it is an easier initial learning curve. Really only power users will want the fine-grained dependencies of metawidget-faces, metawidget-jpa etc. Does Jason have a link to the JIRA for the AS 7 issue?	1	1	https://issues.jboss.org/browse/FORGE-348
4709	FORGE-348	Using forge, it is expected that using Metawidget should be an option, not a mandatory feature. However this was found that when we use Forge to create a SEAM3 project from scratch, we must use metawidget. We have to understand the fact that Metawidget just creates some auto-generated xhtml codes from POJOs (good for simple CRUD based apps). But real life project is much more complicated. It is not always the direct table -> POJO - > xhtml mapping. What if we don't need metawidget at all, but we still want to use Forge to build up SEAM3 project using Richfcaes, JB7 etc... Do we have the option? Do we always need to use metawidget and then remove the config files (metawidget releted) manually to make our project metawidget free?	1	0	https://issues.jboss.org/browse/FORGE-348
4710	FORGE-345	Lincoln, if you get a chance, could you assign it to me?	1	1	https://issues.jboss.org/browse/FORGE-345
4711	FORGE-345	Hey Kevin, I've added you to the developer role. You should now be able to assign tickets to yourself :) Let me know if it doesn't work on this one and I'll see what's wrong. (I'll leave this unassigned so you can test it out.) Thanks Kevin! Lincoln	1	1	https://issues.jboss.org/browse/FORGE-345
4712	FORGE-343	I didn't see any j2ee references, is this for the javaee packaging? For example: ./src/main/java/org/jboss/forge/spec/javaee ./src/main/java/org/jboss/forge/spec/javaee/util ./src/main/java/org/jboss/forge/spec/javaee ./src/main/java/org/jboss/forge/spec/javaee/cdi ./src/main/java/org/jboss/forge/spec/javaee/jpa ./src/main/java/org/jboss/forge/spec/javaee/jpa/api ./src/main/java/org/jboss/forge/spec/javaee/jpa/container ./src/main/java/org/jboss/forge/spec/javaee/jpa/provider ./src/main/java/org/jboss/forge/spec/javaee/jsf ./src/main/java/org/jboss/forge/spec/javaee/servlet Would they all be moved to javaee6 directory/packaging? What about JavaEEDefaultContainer.java, the only java file I found with javaee in it's name? Should it change as well?	1	1	https://issues.jboss.org/browse/FORGE-343
4713	FORGE-340	the installFromGit() method already has a checkoutDir option that defines a directory in which to clone the repository. So, I have an implementation that uses this directory if defined, otherwise it prompts with "Would you like to save the source?". If this is true, then it saves to ~/.forge/plugins-src/ otherwise it saves to a temp directory that is removed when done. All in all, there are a total of 3 possibilities: 1) User defines the repository directory 2) User wants to save, but did not supply a directory 3) User does not want to save and a temp directory is used Does this make sense?	1	1	https://issues.jboss.org/browse/FORGE-340
4714	FORGE-340	Hey Rodney, think this one will make it in to Beta 3? (EoW or early next week hopefully)	1	1	https://issues.jboss.org/browse/FORGE-340
4715	FORGE-337	Minor point, but according to http://jcp.org/en/jsr/detail?id=344, it's JavaServer Faces. Maybe change "The default scaffold comes by default in any Forge distribution" to something like "The default scaffold comes as part of the Forge distribution" Should there be a discussion of the options for "scaffold setup" in the Getting Started section? The Customizing the Scaffold section tells you how to change what was generated. But, is there a way to change files in the plugin that would change all generated files? For example, change the css that is generated (used) for each application.	1	1	https://issues.jboss.org/browse/FORGE-337
4716	FORGE-330	It looks like this has already been pulled in. Did you work with Shane Bryzak to get this into the repository?	1	1	https://issues.jboss.org/browse/FORGE-330
4717	FORGE-327	imo for the plain scaffolding, there are not that many navigation cases to justify a navigation provider or something like that... or do I miss something? I'll make sure a nice index page is created that contains all the entities that are 'scaffolded'. This will be done by a 'menu' on the left of the screen, [kind of like|http://blog.mgm-tp.com/wp-content/uploads/2011/06/roo-web-application.png] :-) Agreed? Please assign this issue to me then	1	1	https://issues.jboss.org/browse/FORGE-327
4718	FORGE-325	Did you want this to target an EE appserver, or a servlet container?	1	1	https://issues.jboss.org/browse/FORGE-325
4719	FORGE-322	Done in new scaffolding implementation. Well, sort of. I don't really grok what these 'finder' methods are all about. What if an entity has 3 fields? Do I need: public Customer findByFirstname(...) { ... } public Customer findBySurname(...) { ... } public Customer findByAddress(...) { ... } public Customer findByFirstnameAndSurname(...) { ... } public Customer findByFirstnameAndSurnameAndAddress(...) { ... } public Customer findBySurnameAndAddress(...) { ... } public Customer findByFirstnameAndAddress(...) { ... } This seems unwieldy in Java? It works in, say, Ruby because they have 'method_missing'. But I don't like the idea of statically generating all these (when most will never be used). So instead I've implemented a simple Query By Example into the scaffolding using Criteria queries. Will that be sufficient? Please test!	1	1	https://issues.jboss.org/browse/FORGE-322
4720	FORGE-321	Could you please try a new snapshot and see if the problem persists? Thanks, I haven't been able to reproduce this: https://repository.jboss.org/nexus/content/groups/public/org/jboss/forge/forge-distribution/1.0.0-SNAPSHOT/forge-distribution-1.0.0-20110608.211641-40.zip	1	1	https://issues.jboss.org/browse/FORGE-321
4721	FORGE-321	Hi, The problem still persists with the version you provided: [forge-test] Customer.java $ scaffold from-entity ? No scaffold type was selected, use default (Metawidget & JSF)? [Y/n] Y ? Scaffold provider [metawidget] is not installed. Install it? [Y/n] Y ***ERROR*** [scaffold from-entity] promptChoice() Cannot ask user to select from a list of nothing. Ensure you have values in your options list. I am behind a proxy. Maybe it needs to be configured. Regards,	0	0	https://issues.jboss.org/browse/FORGE-321
4722	FORGE-321	Tried from my home laptop which is not behind a proxy and it works fine. So this could be due to some proxy settings. Is there any place to set the proxy settings for seam forge? Or does it take the Maven proxy settings in which case I guess we need to set it in maven settings.xml	1	1	https://issues.jboss.org/browse/FORGE-321
4723	FORGE-305	What error are you getting?	0	0	https://issues.jboss.org/browse/FORGE-305
4724	FORGE-302	Awesome! Glad you got this figured out. Since you already prototyped a solution, could you polish it up a bit and send a pull request? Thanks!	1	1	https://issues.jboss.org/browse/FORGE-302
4725	FORGE-289	To give some more info: I actually liked the use of this 'hidden' plugin because it provided me pretty much with all I needed. I think the only thing that needs to be provided is a way to really hide the command, maybe with a '@Hidden' annotation or something alike. Or is that too hacky?	1	1	https://issues.jboss.org/browse/FORGE-289

#	Requirement	Questions	Answered	Successfully Answered	URL
4726	FORGE-289	What do you think about something like this that gets evaluated before the command is parsed or executed? {code} /** * Allows the shell input stream to be intercepted, modified, or observed by extensions. * * @author Lincoln Baxter, III * */ public interface InputStreamInterceptor { /** * Intercepts a line of input and returns it along with any modifications. The returned value will be passed to the * Shell for invocation. */ public String intercept(String line); } {code} I hesitate to use hidden plugins, because I think we are already going to have plugin name conflicts soon (though we could qualify them) But this is probably a useful API anyway. Thoughts?	1	1	https://issues.jboss.org/browse/FORGE-289
4727	FORGE-286	quick top of my head comments: Creating yet another plugin/module system sucks ;) Using classifier of maven artifacts would be a "misuse" IMO - could we index it someway else ? Consider using something like P2 from eclipse.org to handle the installation/automatic updates?	1	1	https://issues.jboss.org/browse/FORGE-286
4728	FORGE-286	Max, Could you explain how you see that fitting in? I checked out the site but was not able to find much info about how to use it, what it actually does, etc... http://www.eclipse.org/projects/project_summary.php?projectid=rt.equinox.p2 Would it work with Maven artifacts/dependency management? Also, yes; we could: * Index assuming that we had control over the repository and it were dedicated to Forge plugins. * Use our own packaging type (fpg) or something, for which the artifact would be a JAR (is this possible?), but run a nightly cron-job to scan through all the artifacts. Personally I think we need a repository of which we have complete control. The repository would be responsible for containing necessary metadata, much like an eclipse update-site. I would prefer not to implement a repository system myself, but I am assuming that we will need to in order to provide this functionality ... Unless P2 works well enough that it would: A) not require significant modifications to our architecture, and B) not require more work/time than creating a simple system ourselves.	1	1	https://issues.jboss.org/browse/FORGE-286
4729	FORGE-285	On the forum they say to be working on a REST api. Let's put this on hold until the api is available, otherwise we're hacking around something that will proably brake later. Is it ok I assign this to myself? I'll monitor the forum and start working on it when available.	0	0	https://issues.jboss.org/browse/FORGE-285
4730	FORGE-285	Is this issue closeable?	1	1	https://issues.jboss.org/browse/FORGE-285
4731	FORGE-281	Lincoln, can you assign this to Paul Bakker to credit him as the contributor, then resolve it?	1	1	https://issues.jboss.org/browse/FORGE-281
4732	FORGE-273	Is this the one you're working on? Using FPak?	1	1	https://issues.jboss.org/browse/FORGE-273
4733	FORGE-273	Which archetype catalogs should we provide by default?	0	0	https://issues.jboss.org/browse/FORGE-273
4734	FORGE-268	I like the idea of Lombok, but the issue I see with it is that I think it requires plugins for the IDE in order to get things to work. (And not throw up lots of compile errors.) Can you tell me if that's still the case?	1	1	https://issues.jboss.org/browse/FORGE-268
4735	FORGE-254	Lincoln, does the Ctrl+D work also for e.g. dependency-add, new-field custom, repository-add, ... ? I mean their interactive mode. Now when I'm trying to escape from dependency-add with Ctrl+c/exit it's not possible. And another question... Shouldn't the fix version be higher? Because I'm testing 1.0.0.Alpha3 and it's not there. The only thing I get after pressing Ctrl+D is "***ERROR*** Exception encountered: null (type "set VERBOSE true" to enable stack traces)", the command really exits but it's not nice:)	0	0	https://issues.jboss.org/browse/FORGE-254
4736	FORGE-253	From mail on forge-dev May 11: well, is that info persisted somewhere so other plugins can read the info ? persistence.xml does not have it and as discussed at JBW I believe deplomentime DB info <> development DB info. What would be great is if the persistence plugin would have a single or even list of named JPADataSource's which have enough info for data query, reverse engineering etc. could be done withuot having to repeat it at every generate-entity call.	1	0	https://issues.jboss.org/browse/FORGE-253
4737	FORGE-239	What would be the best place for this? I was thinking in a new facet (ProjectExecutionFacet) in dev-plugins. The facet is mostly an API for plugin authors. We can also add a plugin that exposes the facet to end-users. Thoughts?	1	1	https://issues.jboss.org/browse/FORGE-239
4738	FORGE-239	JavaExecutionFacet perhaps? Not all projects would necessarily execute using Java. Wondering if this should be in the core API. Dev-plugins isn't a place for feature addition, so for now I'd say let's put it in the Shell-API itself (implemented by the Maven Project Impl)	0	0	https://issues.jboss.org/browse/FORGE-239
4739	FORGE-233	Would it make sense to integrate parts of the formatter plugin for this? https://github.com/thomashug/forge-formatter-plugin	1	1	https://issues.jboss.org/browse/FORGE-233
4740	FORGE-227	Hi. Thanks for working on this! Could you please send a pull request and link that URL? The link you have provided is just a link to your source tree. Thanks, Lincoln	1	1	https://issues.jboss.org/browse/FORGE-227
4741	FORGE-227	To send a pull request, navigate to https://github.com/includeeasy/core/tree/SEAMFORGE-162 and click the "send pull request" button. Then fill out the subject and description and click send. The encoding may pose a problem. If possible, could you change this to UTF-8 and re-submit? Also, if you could make the following refactoring to the findRange() method: From: {code}public <T> List<T> findRange(Class<T> type,int[] range) {code} To: {code}public <T> List<T> findAll(Class<T> type,int firstResult, int maxResults){code} I think that API will be a little easier to understand, and a little less user-error prone. Thanks! This is good stuff.	1	1	https://issues.jboss.org/browse/FORGE-227
4742	FORGE-227	Done! Thanks for this enhancement! If you have not already, could you please sign this contributors agreement and make sure you have an account on jboss.org? That way I will be able to give you credit for completed issues. https://cla.jboss.org/ Thanks, ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-227
4743	FORGE-227	Hantsy, if you'd like to see this feature. Could you please file a new issue describing your needs? This issue has been closed for some time. Thanks. This sounds like something that would better belong in the RichFaces plugin: https://github.com/forge/plugin-richfaces	1	1	https://issues.jboss.org/browse/FORGE-227
4744	FORGE-227	I have implemented this by add a field of "private T filter;",put filter parameter inside getItemsCount() and createPageDataModel(). But there is problems whith isolate richfaces/icefaces/primefaces and integrate richfaces & matewidget. So now you can write an plugin that implement scafford api. but how can those plugins work together?	1	0	https://issues.jboss.org/browse/FORGE-227

#	Requirement	Questions	Answered	Successfully Answered	URL
4745	FORGE-226	Discussion on irc: (03:38:13 PM) lincolnthree: imagine you have a PomXmlResource (there is such a thing in Forge 1) (03:38:16 PM) lincolnthree: you could say: (03:38:31 PM) lincolnthree: pomResource.getFacet(ProfilesFacet.class) (03:38:53 PM) lincolnthree: So in a sense, a project would just be a Resource! (03:39:32 PM) kenfinnigan: so if I understand it.. (03:39:34 PM) lincolnthree: webXmlResource.getFacet(ServletsFacet.class) (03:39:46 PM) kenfinnigan: JavaSourceResource.getFacet(EJBFacet.class).isBeanType(Type.MESSAGE_DRIVEN_BEAN) (03:39:55 PM) lincolnthree: webXmlResource.hasFacet(ContextParamsFacet.class) (03:39:56 PM) kenfinnigan: would return an EJB that was an MDB? (03:40:01 PM) lincolnthree: right (03:40:05 PM) kenfinnigan: ie the java source of the class that is (03:40:10 PM) lincolnthree: well (03:40:17 PM) lincolnthree: based on what you've written, i'd expect it to return a boolean (03:40:18 PM) lincolnthree: but (03:40:19 PM) lincolnthree: yes (03:40:42 PM) lincolnthree: javaResource.hasFacet(ArquillianFacet.class) (03:40:44 PM) lincolnthree: etc (03:41:22 PM) kenfinnigan: ok, I was reading JavaSourceResource as a static, but from what you're saying in that example it would actually be an instance (03:41:26 PM) kenfinnigan: correct? (03:41:44 PM) kenfinnigan: so asking that java source instance if it's an MDB (03:41:47 PM) lincolnthree: javaResource.getFacet(TransactionInterceptorFacet.class).getPropagationType(); (03:41:56 PM) lincolnthree: it would be an instance (03:41:57 PM) lincolnthree: yes (03:42:04 PM) kenfinnigan: ok (03:42:21 PM) lincolnthree: so the facet would be able to inspect/install/remove certain pieces of information from java code in this case (03:42:30 PM) lincolnthree: in a very granular/controlled manner (03:42:55 PM) kenfinnigan: hmm, so where would that control lie, in the Facet or Resource, or bit of both? (03:44:42 PM) lincolnthree: the resource would have no knowledge of the facets contained within (03:44:54 PM) lincolnthree: facets would be responsible for inspecting the resource and determining their own installation sate (03:44:56 PM) lincolnthree: state (03:44:59 PM) lincolnthree: just like Project facets (03:45:07 PM) aslak: kenfinnigan, what am i saying? :) (03:45:24 PM) kenfinnigan: ok, so Resource just needs to provide the means that Facet can do the kinds of things aslak wants? (03:45:29 PM) kenfinnigan: aslak: FORGE-226 (03:45:30 PM) jbossbot: jira [FORGE-226] Support for Facets construct on Resources [Open (Unresolved) Feature Request, Major, Unassigned] https://issues.jboss.org/browse/FORGE-226 (03:45:39 PM) lincolnthree: resource needs to provide a generic Facet API (03:45:47 PM) lincolnthree: like the Project class in forge 1.x does (03:45:47 PM) kenfinnigan: lincolnthree: ok (03:46:18 PM) lincolnthree: https://github.com/forge/core/blob/master/shell-api/src/main/java/org/jboss/forge/project/Project.java (03:46:31 PM) lincolnthree: something like that (03:46:39 PM) aslak: kenfinnigan, aa.. the basic idea is.. you have Facets on a Project (capabilities for a given project), the similar thing you would have on Resource level (03:47:15 PM) aslak: new XMLResource("arquillian.xml").getFacet(ArquillianFacet.class).addContainer() (03:47:51 PM) kenfinnigan: aslak: ok, so then the ArquillianFacet "knows" the correct XML to add back to the resource for addContainer()? (03:48:00 PM) aslak: yes (03:48:30 PM) kenfinnigan: ok, so Facets do all the work of modifying Resource, but Resource needs to know what kind of Facet it can deal with (03:48:33 PM) aslak: you need a 'supports(Facet)" or similar in ther ea swell (03:48:47 PM) aslak: it's basically a way to share code on resource level instead of just project level (03:48:59 PM) lincolnthree: yea (03:48:59 PM) kenfinnigan: ok, think that makes sense (03:49:18 PM) aslak: the EJB addon should know 'stuff' about a Class to determine if it's a EJB.. possible add EJB related annotations	1	1	https://issues.jboss.org/browse/FORGE-226
4746	FORGE-219	Hey Michael! Any more updates on this plugin? Can this issue be closed? Thanks. ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-219
4747	FORGE-219	Don't need to file core JIRA issues for creating new plugins. Have a blast :) Keep it up!	1	1	https://issues.jboss.org/browse/FORGE-219
4748	FORGE-213	Is it enough that someone be able to do this by typing "install beans" "install ejb" "install servlet"? Or do you want something more like an archetype?	0	0	https://issues.jboss.org/browse/FORGE-213
4749	FORGE-205	Kevin, if you think this is ready, I can begin merging it into the core JavaEE 6 plugins. This would mean that it will be part of the forge distribution, and you'll have to send pull requests to make updates. Is that OK? I'd really like for this to be a part of forge itself (You will of course still have credit for all of the work.) We already have a few JavaEE plugins: persistence faces beans (servlet facet) and this would add "validation" Thoughts? Lincoln	1	1	https://issues.jboss.org/browse/FORGE-205
4750	FORGE-202	[example] example $ project repository-add CUSTOM What is the name of the repository? OCPSOFT What is the URL of the repository? http://ocpsoft.com/repository/ Added repository [OCPSOFT->http://ocpsoft.com/repository/] [example] example $	0	0	https://issues.jboss.org/browse/FORGE-202
4751	FORGE-180	Hey Richard, could you take a look at this? Thanks!	1	1	https://issues.jboss.org/browse/FORGE-180
4752	FORGE-176	Hey Kevin, did you run the tests before you sent this pull request? From what I saw the tests didn't pass because there was a bug in the PromptType Enum. Anyway, this is great and the tests are very thorough. Keep up the good work!	1	1	https://issues.jboss.org/browse/FORGE-176
4753	FORGE-175	I agree that it would be best to add it as provided. Do you change that when pulling, or should I change and commit it?	1	1	https://issues.jboss.org/browse/FORGE-175
4754	FORGE-175	Now I know a little more about Forge I see you're right. I'll extend it a bit too when I make the changes: -Check if either cdi-api or jboss-specs are in the pom already -> if so, don't add anything -if not promt the user if cdi-api should be installed (defaults to Y) -if Y, promt the user what scope it should get, defaults to compile (probably it's a standalone project, otherwise the specs should have been on the classpath) Do you agree?	1	1	https://issues.jboss.org/browse/FORGE-175
4755	FORGE-169	Hey Ken, Just checking, were you working on this?	1	1	https://issues.jboss.org/browse/FORGE-169
4756	FORGE-169	Hey Ken, did you ever start on this one?	0	0	https://issues.jboss.org/browse/FORGE-169
4757	FORGE-167	Do you remember why we opened this issue, Paul? Something you want to knock out, maybe?	1	0	https://issues.jboss.org/browse/FORGE-167
4758	FORGE-163	Where would that name get injected from?	0	0	https://issues.jboss.org/browse/FORGE-163
4759	FORGE-161	The beauty of Forge is that we can support *everybody's* use case! :) I think this is a good addition, and will help streamline path to deployment on AS7, even if Forge actually handles unzipping itself (which I think may be the route I want to take, instead of using maven.) Forge can unzip servers to its local ~/.forge/ directory, and use those locations for deployments. But I think in the short term, using maven is a good option, too. Some people may want that, too.	0	0	https://issues.jboss.org/browse/FORGE-161
4760	FORGE-161	Thanks for picking this one up, paul! I think you have the right experience in order to do this right :) Do you think we can do startup/shutdown of the AS7 server as well? That would be harder but certainly doable (JBoss Tools does it.) That would mean forge could really stand on its own when it comes to developing webapps!	0	0	https://issues.jboss.org/browse/FORGE-161
4761	FORGE-157	Could you create a unit test that shows this failing? That would make it easy to fix :) Thanks!	1	0	https://issues.jboss.org/browse/FORGE-157
4762	FORGE-143	Could you please list out the actions that you think should be provided by base classes in specific? I'm willing to bet that most are already provided by Facets. E.g: JavaSourceFacet can be used to calculate the current package. But we could make that API a bit more powerful depending on what you need.	1	1	https://issues.jboss.org/browse/FORGE-143
4763	FORGE-142	To run forge from the IDE - you should ideally run the dist/ project, not the shell/ project which is what I think you are doing? Can you try that and let me know if it works?	1	1	https://issues.jboss.org/browse/FORGE-142

#	Requirement	Questions	Answered	Successfully Answered	URL
4764	FORGE-139	Do you want this issue to be closed?	0	0	https://issues.jboss.org/browse/FORGE-139
4765	FORGE-137	Hey Kevin, Do you have an example of such a plugin that we could incorporate? I'd be glad to add this; I think it's very useful. ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-137
4766	FORGE-137	Hey Lincoln, isn't this issue obsolete through the forge commands as7 deploy, as7 exploded, etc.? Regards//Kevin	1	1	https://issues.jboss.org/browse/FORGE-137
4767	FORGE-137	Hi Lincoln, I've a couple of ideas on how to solve this issue. The investigation below is focused on JBoss AS 7 only. I copied the scenario approaches from the JBoss AS 7 deployments README.txt. Within the pom.xml define variable for deployment mode, by default: - auto-deploy-zipped = "true" - auto-deploy-exploded = "false" h3. explode commands *maven compile -Pexplode* - Scenario 1: Add new unzipped content and deploy it {quote} 1. cp -r target/example.war/ $AS/standalone/deployments 2. (Manual mode only) touch $AS/standalone/deployments/example.war.dodeploy {quote} - Scenario 2: Live replace portions of currently deployed unzipped content without redeploying {quote} 1. (Auto mode only) touch $AS/standalone/deployments/example.war.skipdeploy 2. cp -r target/example.war/foo.html $AS/standalone/deployments/example.war {quote} One remark Scenario 2 only triggers if Scenario 1 has been executed first. *maven compile -Pre-explode* - Scenario: Replace currently deployed zipped content with a new version and deploy it: {quote} 1. cp target/example.war/ $AS/standalone/deployments 2. (Manual mode only) touch $AS/standalone/deployments/example.war.dodeploy {quote} *maven clean -Punexplode* - Scenario: undeploy and remove unzipped content {quote} 1. rm $AS/standalone/deployments/example.war.deployed 2. wait for $AS/standalone/deployments/example.war.undeployed file to appear 3. rm $AS/standalone/deployments/example.war {quote} h3. deploy commands *maven package -Pdeploy* - Scenario: Add new zipped content and deploy it {quote} 1. cp -r target/example.war/ $AS/standalone/deployments 2. (Manual mode only) touch $AS/standalone/deployments/example.war.dodeploy {quote} *maven package -Pre-deploy* - Scenario: Redeploy currently deployed content (i.e. bounce it with no content change): {quote} 1. touch $AS/standalone/deployments/example.war.dodeploy {quote} *maven clean -Pundeploy* - Scenario: undeploy and remove unzipped content {quote} 1. rm $AS/standalone/deployments/example.war.deployed 2. wait for $AS/standalone/deployments/example.war.undeployed file to appear 3. rm $AS/standalone/deployments/example.war {quote} I'm not sure whether the wait commands in unexplode and undeploy can be realized with maven. Therefore this might need a rework. What do you think? Any questions or suggestions? Can you think of any pitfalls with this approach? Regards//Kevin	1	1	https://issues.jboss.org/browse/FORGE-137
4768	FORGE-137	Hey Kevin, Sorry for the late reply. I've been on the road at conferences. These command breakdowns look good, but does this actually make this functionality available when using maven standalone? It looks like these are just Forge commands. Is that what you want? I'm not really clear how you plan to implement this approach (but that doesn't mean I'm against it, just that I want to know more.) Thoughts? ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-137
4769	FORGE-137	Hey Lincoln, no problem this gave me some time to think of a solution ;) Here is an example on how I thought to implement this issue for *mvn compile -Pexplode* {code:xml} <profile> <id>explode</id> <build> <plugins> <plugin> <artifactId>maven-war-plugin</artifactId> <configuration> <webappDirectory>${exploded.dir}</webappDirectory> <failOnMissingWebXml>false</failOnMissingWebXml> </configuration> <executions> <!-- execute goal war:explode with compile phase --> <execution> <id>explode</id> <phase>compile</phase> <goals> <goal>exploded</goal> </goals> </execution> </executions> </plugin> <plugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-antrun-plugin</artifactId> <version>1.7</version> <executions> <execution> <id>copy-as-deployment-descriptors</id> <phase>compile</phase> <goals> <goal>run</goal> </goals> <configuration> <target> <touch file="${exploded.dir}.dodeploy" /> </target> </configuration> </execution> </executions> </plugin> </plugins> </build> </profile> {code} Though this works it gets really ant like messy if you plan to implement this for all deploy and explode goals. Another approach I thought of is to use the the jboss as7 maven plugin. But as they only support deploying this isn't really a complete solution. This also leave the question open on how the implementation for Glassfish will look like? What do you think? //Kevin	1	1	https://issues.jboss.org/browse/FORGE-137
4770	FORGE-137	Hmmm, actually. If we do FORGE-194, do you think this is necessary? Would that be a solution that meets your use case? Also, the "as7 deploy" command now supports "as7 deployExploded". There's also FORGE-256 which is related to this I think. So many people asking for different things! Do all? One? Some?	1	1	https://issues.jboss.org/browse/FORGE-137
4771	FORGE-127	Do you think we should reconsider that? Feel free to reopen if that's the case. Thanks	0	0	https://issues.jboss.org/browse/FORGE-127
4772	FORGE-124	Thanks Paul! I'll get to these soon. In the future, could you create an individual pull-request for each issue? That will make it much easier to pull things in and test individually/ ~Lincoln	1	1	https://issues.jboss.org/browse/FORGE-124
4773	FORGE-124	Yes, I was looking for that on github. When I try to create another pull requests it says there already is one. How do you get around that, use feature branches for pull requests? I'm not that familiar with git yet, so still trying to figure out how to work best with it.	1	1	https://issues.jboss.org/browse/FORGE-124
4774	FORGE-121	[forge-distribution] dist $ [forge-distribution] dist $ [forge-distribution] dist $ prettyfaces ***ERROR*** No such command: prettyfaces [forge-distribution] dist $ [forge-distribution] dist $ [forge-distribution] dist $ forge git-plugin git://github.com/ocpsoft/prettyfaces-forge-plugin.git --ref 1.0.0.Alpha1 ***INFO*** Checking out plugin source files to [repo] via 'git' ***INFO*** Switching to branch/tag [1.0.0.Alpha1] ***INFO*** Invoking build with underlying build system. [INFO] [INFO] ---------- [INFO] Building prettyfaces-forge-plugin 1.0.0-SNAPSHOT [INFO] ---------- [INFO] [INFO] Copying 2 resources [INFO] [INFO] Compiling 3 source files to /tmp/forgetemp8366335975626723199/repo/target/classes [INFO] [INFO] No tests to run. [INFO] [INFO] Building jar: /tmp/forgetemp8366335975626723199/repo/target/prettyfaces-forge-plugin-1.0.0-SNAPSHOT.jar [INFO] ---------- [INFO] BUILD SUCCESS [INFO] ---------- [INFO] Total time: 2.334s [INFO] Finished at: Tue Mar 22 15:55:05 EDT 2011 [INFO] Final Memory: 32M/259M [INFO] ---------- ***SUCCESS*** Build successful. ***INFO*** Copying build artifact to [/home/lb3/.forge/plugins/com.ocpsoft.forge.prettyfaces_prettyfaces-forge-plugin.jar] [forge-distribution] dist $ [forge-distribution] dist $ [forge-distribution] dist $ prettyfaces setup mapping-add [forge-distribution] dist $ prettyfaces setup The [com.ocpsoft.prettyfaces] facet depends on the following missing facets: [WebResourceFacet, forge.spec.servlet]. Would you like to attempt installation of these facets as well? [Y/n]	0	0	https://issues.jboss.org/browse/FORGE-121
4775	FORGE-120	I've tried testing this with forge-distribution-1.0.0.Beta3.zip but there is no such functionality there. What has happened to this issue?	0	0	https://issues.jboss.org/browse/FORGE-120

#	Requirement	Questions	Answered	Successfully Answered	URL
4776	FORGE-120	This does not seems to be working anymore. Executing project-new --named demo --topLevelPackage org.demo. --type pom gives a war project. The commited source code has since than be refactored. NewProjectPlugin and NewProjectPluginTest are no longer available. Is there an example for this and general multi-module support? This would be fairly vital for any larger application.	1	1	https://issues.jboss.org/browse/FORGE-120
4777	FORGE-118	Actually, I think this is even easier than scanning the class. If the project packaging is WAR, then make a WebArchive, if it is JAR, then make a JavaArchive :) PS. Who are you?	0	0	https://issues.jboss.org/browse/FORGE-118
4778	FORGE-114	I abstract BackingBean as "CommonGenericBean",and then write an concrete bean "PersonBean" as following: public class PersonBean extends CommonGenericBean<Person> implements MenuItem{} When i test PersonBean on jboss-as7,two types of problems I encountered, 1.When write CommonGenericBean as floowing,enable interceptor class "org.jboss.seam.transaction.TransactionInterceptor" caused getItemType() returns Object,then disable the interceptor,getItemType() become return Peron class,so I can not use JTA. public abstract class CommonGenericBean<T> extends PersistenceUtil { private List<T> list = null; private T item = createInstence(); private long id = 0; private PaginationHelper<T> pagination; private transient Class<T> entityClass; public Class<T> getItemType() { if (entityClass == null) { entityClass = getSuperClassGenricType(getClass()); } return entityClass; } ... 2.When write CommonGenericBean as floowing,it says "No active JTA transaction on joinTransaction call",in this case,only "RESOURCE_LOCAL" type of persistence can run. public abstract class CommonGenericBean<T> { private List<T> list = null; private T item = createInstence(); private long id = 0; private PaginationHelper<T> pagination; private transient Class<T> entityClass; protected PersistenceUtil persit; protected abstract void setPersistenceUnit(PersistenceUtil persit); public Class<T> getItemType() { if (entityClass == null) { entityClass = getSuperClassGenricType(getClass()); } return entityClass; } ... Any one can fix those problem?	0	0	https://issues.jboss.org/browse/FORGE-114
4779	FORGE-111	im curious why does Java EE 6 war project require a web.xml ?	1	1	https://issues.jboss.org/browse/FORGE-111
4780	FORGEPLUGINS-211	agoncal The idea is when you setup a logger, you also generate the appropriate producer method... but that means the project relies on CDI 16:35:05 gastaldi Interesting 16:35:36 gastaldi Maybe this could be a separate addon 16:35:53 agoncal gastaldi Yes, I think so 16:36:09 gastaldi We could start by creating it outside of core 16:36:09 agoncal And what about the producer method and the CDI dependency ? 16:36:24 gastaldi Roaster can generate that 16:36:51 gastaldi It should depend on javaee facet 16:36:55 agoncal yes, I know, I'm just wondering if it's a good idea to always depend on CDI 16:36:57 gastaldi CDIFacet 16:37:04 agoncal (Spring projects for example) 16:37:21 gastaldi That depends on the facets installed in the project 16:37:23 agoncal I like CDI and would love to see it used everywhere... but it's not the case yet 16:37:36 agoncal or maybe : 16:37:40 agoncal logger-setup 16:37:49 agoncal logger-new-producer 16:37:51 agoncal or 16:37:54 agoncal logger-new-factory 16:38:12 gastaldi Yeah thats a good idea too 16:38:14 agoncal One uses CDI, the other doesn't.... but is that really useful ? 16:38:37 agoncal logger-new-producer would install the CDIFacet, the other no 16:38:39 gastaldi New factory could do the same as new producer 16:38:52 agoncal ok 16:39:01 gastaldi Hm. I think that it wouldnt need to install the CDIFacet 16:39:27 gastaldi Just enable the commands if the facet is already installed should be enough 16:40:28 agoncal Hum.... When we create a backing bean, the command faces-new-bean is enabled without or without CDI... and then, it installs it 16:40:47 agoncal I think the behevior should stay the same 16:40:57 gastaldi Ah right 16:41:05 gastaldi Yeah makes sense	1	0	https://issues.jboss.org/browse/FORGEPLUGINS-211
4781	FORGEPLUGINS-179	What setup? It's just the design of the endpoints.	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-179
4782	FORGEPLUGINS-171	This is excellent! Great idea for an addon! Do we know what the issue is here?	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-171
4783	FORGEPLUGINS-171	I don't think so because on the first step, the example_manual.ftl is copied to src/main/asciidoc and I used the ResourceFacet to do it and it works, so the folders src/main/resources and src/main/asciidoc are created. WDYT?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-171
4784	FORGEPLUGINS-171	The installation succeeded. Can you update the addon to use Forge 2.10.0.Final?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-171
4785	FORGEPLUGINS-171	It's weird, the plugin version is still not being added to the pom.xml. Do you see that happening?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-171
4786	FORGEPLUGINS-171	No. Do you specify a converter ?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-171
4787	FORGEPLUGINS-171	Indeed it's weird, this command works for me. Is it a new project or an existing project ?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-171
4788	FORGEPLUGINS-171	an existing project. Are you using Forge 2.10.1.Final?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-171
4789	FORGEPLUGINS-171	No. Can you try on a new project , I mean without any asciidoctor datas un the pom?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-171
4790	FORGEPLUGINS-171	Maybe it's something related to my .m2 local repository. Indeed the plugin version is not set in the code but it's resolved with the DependencyFacet, I may add a default version. I will try with a clean repo and debug too. WDYT?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-171
4791	FORGEPLUGINS-163	Isn't that done by windup?	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-163
4792	FORGEPLUGINS-97	[~pmuir] Did you have any idea in particular when you opened this issue ? George Gastaldi suggested split the container in categories: REMOTE, EMBEDDED, MANAGED so we can do something like: {code} arquillian setup --containerType REMOTE --containerName JBOSS_AS_4.2.X {code} What you think ?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-97
4793	FORGEPLUGINS-83	Hum, I wonder how would this client behave ? Main problem is that a REST endpoint is not fully described. You can pass any parameters you want. How would a client know which HTTP methods are supported (POST / GET/ DELETE/HEAD) ? Probably generating from a WADL file (if people use this anyway :)) would be much better. This should be something interesting to discuss about.	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-83
4794	FORGEPLUGINS-73	I wonder how Eclipse handle it's refactorings. Wouldn't that be on the AST level? In that case we could maybe use that directly.	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-73
4795	FORGEPLUGINS-69	Hi! Is it OK that a plugin extends another plugin? Or we want to definitely avoid inheritance? Cheers! Ivan	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-69
4796	FORGEPLUGINS-65	Although it's a common approach to use some marker syntax it's very breakable. If someone edits the file by hand there's a good change that the marker syntax is removed or the file structure is inconsistent. That's the problem with generating code in user-editable files. The Spring Roo guys solved this by putting all generated code in AspectJ Inter Type Declarations. The cool thing about it is that all generated code is outside of files editable by users. I don't like the heavy dependency on AspectJ however, but isn't something similar possible for Forge? For more advanced scaffolding etc. an elegant solution is necessary too.	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-65
4797	FORGEPLUGINS-64	Hey Mike, this looks like it is due to the FSH parser. Can you take a look and see if there is anything we can do about this? Thanks.	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-64

#	Requirement	Questions	Answered	Successfully Answered	URL
4798	FORGEPLUGINS-54	Feature *wish* list (starting with what would be desireable versus achievable): *Bi-directional Incremental change support. I do not recall where, but this was a black eye for Forge/Hibernate reverse engineering in one review ('destructive' was a word used to describe it, which is unfair when you talk reverse engineering tool versus change management tool) when compared to Grails and Roo. --change in database schema and have 'incremental' change support to domain code...without replacing existing code (example, if you modified the JPA design to include TABLE_PER_CLASS inheritance, and you add a new column to one of those tables or a separate class that as one-to-many to that inheritance, do not want to loose that work). --change the domain in code, incrementally handle db migration. Note that column renames need to be specifically addressed, and some type of ability to *track* other columns marked for deletion, but not actually purge unless explicitly defined. *SQL support to handle more database-specific work, like data migration work (if merging or breaking apart data in column(s) to other column(s)). *SQL specific schema support -- Specifically, there may be additional columns/relations in the table design that do not/are not part of the domain model. This should be supported with some type of 'ignore' feature when working on the code-domain but still exist/support/be upgraded as part of the upgrade process (new/improved indexes, constraints, spoc/functions, sql views for example). Not primary feature of this plugin, but should support it. *Control data population (setting default values, app-specific values, environment/client values, etc). *Application-deploy upgrade process, cluster aware (avoid multiple apps upgrading the schema at the same time) -- have support to (optionally) auto-upgrade database schema when deploying newer application. This should support multi-version upgrade (optional to control size/number of diff/changes -- if database schema was for application v3, and now deploying application v5, but able to upgrade through all changes). --caveat, one of the comments against this approach is the upgrade library/code and the upgrade scripts would 'always' be packaged with the application, versus only needed once per upgrade. There is no way around this without introducing additional external support (separate system) or separate artifact like an 'app-upgrade.war' sibling that can be undeployed without impacting the real application (potentially deploy both, have the real application have lightweight version to check db version matches need for application version, if not call/wait for upgrade artifact, then undeploy/remove the upgrade artifact). *Support both code-first and database-first kickstart (how/where to baseline the database). *(nice to have) - ability to know what changes are backward compatible from the application domain code and the database changes point-of-view, versus those that are not. For example, adding new columns that are nullable and no FK constraint are backward compatible allowing the prior applications to continue to function during the upgrade. This is useful for stepped upgrade processes for minimal downtime.	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-54
4799	FORGEPLUGINS-37	does those two lines mean you end up with the merged set of boms from both -with-tools and -with-hibernate or is this meant to be just one time per project ?	0	0	https://issues.jboss.org/browse/FORGEPLUGINS-37
4800	FORGEPLUGINS-37	To use a engine to retrieve the available stacks I suggest that jdf uses two different commands (one for install and another get the list of available stacks) like {code} jdf install-stack --stack jboss-javaee-web-6.0-with-tools --version 1.0.0.Final {code} and {code} jdf update-stacks || jdf update-available-stacks {code} Most users uses the autocomplete feature to discover what a plugin does. If we use --stack option in the default command I think that will be hard to users discover how they can update BOM list. What you think ?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-37
4801	FORGEPLUGINS-17	Pete, can you paste a minimal arquillian.xml to be created ? Regards, George Gastaldi	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-17
4802	FORGEPLUGINS-16	Just wondering. Did you have the richfaces plugin installed at the time?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-16
4803	FORGEPLUGINS-16	Sorry. I was (stupidly) expecting you to call 'richfaces setup' before 'scaffold setup'. Try again now?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-16
4804	FORGEPLUGINS-16	Why is it stupid to expect that one would call richfaces setup? Shouldn't that be required if you want richfaces functionality? PS. Nice job richard!	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-16
4805	FORGEPLUGINS-16	Further fixes for: CSS fixes. Adding multiple one-to-many at once. Required field handling. Nested collections. Note we appear to have a bit of an issue with Bean Validation's @NotEmpty. Pete could you please investigate? Basically: 1. Generate the app 2. Cheer how the date pickers no longer overlap 3. Cry because when you go 'Booking->Create New' you get validation errors on the search fields Now, I have already removed JSF's 'required=true' attribute from these fields. So it appears it is Bean Validation doing this. You cannot use 'immediate=true' on the button or else nothing will get POSTed at all. The 'correct' solution, I think, is to add f:validateBean disabled="true" to each of the required fields. I can do this no problem, but it doesn't seem to work (under AS 7.0.2.Final). Example: 1. Add 'f:validateBean disabled=true' to 'h:inputText id="bookingBeanSearchCancellationCode"' and 'h:inputText id="bookingBeanSearchContactEmail"' in 'booking/search.xhtml' 2. Go to the search screen and click 'Create New'. Works! Hurrah! 3. Go back to the search screen and click 'Search' and then 'Create New'. Now it fails. It appears f:validateBean does not survive POST-back? Is this a bug? Even if this *did* work, a bad side effect is that all validation is disabled for these fields. So for example it no longer validates the e-mail address format. We can probably live with that on the search screens, but it may be unacceptable on the one-to-many 'add' fields (where this is also a problem). Suggestions?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-16
4806	FORGEPLUGINS-16	I'm not 100% sure of what the design is here. Do you want Create New to take values from the search fields at all? I'm guessing so. I'm also assuming that you generally don't want @NotEmpty to apply to searches at all. In this case the answer is validation groups. I think you need a search group which ignores certain validations such as @NotEmpty. If you give this group a standard name, this will work for an existing domain model as well, as we can apply the group there.	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-16
4807	FORGEPLUGINS-16	Emmanuel, I agree most validations won't apply to search fields. I think an 'easy default' (i.e. without requiring people to modify their domain model) would be to just add <f:validateBean disabled="#{true}"/> to each search field. However this doesn't appear to survive POST-back (i.e. it works for one click, but then fails after that). Are you aware of this? Is it a bug? I attach a minimal test case. I'd be most grateful if you could look at it. If you agree it is a bug, could you raise it with the relevant team? I could do it myself, but you have the whole 'scary French genius' thing going on so can probably get people to react more quickly :) Regards, Richard.	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-16
4808	FORGEPLUGINS-16	Okay, opened as: https://issues.jboss.org/browse/AS7-3621 I guess FORGE-460 will remain in limbo until that's resolved? Pete can you please assign to Stan?	1	1	https://issues.jboss.org/browse/FORGEPLUGINS-16

#	Requirement	Questions	Answered	Successfully Answered	URL
4809	FURNACE-104	java.lang.reflect.Method] at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback.enhance(ClassLoaderAdapterCallback.java:1044) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback.enhance(ClassLoaderAdapterCallback.java:889) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback.enhanceResult(ClassLoaderAdapterCallback.java:294) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback.access$200(ClassLoaderAdapterCallback.java:37) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback$2.call(ClassLoaderAdapterCallback.java:124) at org.jboss.forge.furnace.util.ClassLoaders.executeIn(ClassLoaders.java:42) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback.invoke(ClassLoaderAdapterCallback.java:96) at org.jboss.forge.addon.ui.context.UISelection_$$_javassist_575aebd0-8e5c-4c06-906f-2ee8326b6cdd.getSelectedRegionFor(UISelection_$$_javassist_575aebd0-8e5c-4c06-906f-2ee8326b6cdd.java) at org.jboss.forge.addon.ui.example.commands.ShowSelectedRegionCommand.execute(ShowSelectedRegionCommand.java:64) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.jboss.forge.furnace.proxy.ClassLoaderInterceptor$1.call(ClassLoaderInterceptor.java:87) at org.jboss.forge.furnace.util.ClassLoaders.executeIn(ClassLoaders.java:42) at org.jboss.forge.furnace.proxy.ClassLoaderInterceptor.invoke(ClassLoaderInterceptor.java:103) at org.jboss.forge.addon.ui.example.commands.ShowSelectedRegionCommand_$$_javassist_36612cef-454a-4bb7-896c-5a0dde266598.execute(ShowSelectedRegionCommand_$$_javassist_36612cef-454a-4bb7-896c-5a0dde266598.java) at org.jboss.forge.addon.ui.impl.controller.SingleCommandControllerImpl.execute(SingleCommandControllerImpl.java:95) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback$2.call(ClassLoaderAdapterCallback.java:123) at org.jboss.forge.furnace.util.ClassLoaders.executeIn(ClassLoaders.java:42) at org.jboss.forge.furnace.proxy.ClassLoaderAdapterCallback.invoke(ClassLoaderAdapterCallback.java:96) at org.jboss.forge.addon.ui.controller.CommandController_$$_javassist_5f874824-26e5-47f7-be2e-2e2360216527.execute(CommandController_$$_javassist_5f874824-26e5-47f7-be2e-2e2360216527.java) at org.jboss.tools.forge.ui.internal.ext.wizards.ForgeWizard$1.run(ForgeWizard.java:113) at org.eclipse.jface.window.ApplicationWindow.run(ApplicationWindow.java:696) at org.jboss.tools.forge.ui.internal.ext.wizards.ForgeWizard.performFinish(ForgeWizard.java:97) at org.jboss.tools.forge.ui.internal.ext.dialog.WizardDialogHelper.openWizard(WizardDialogHelper.java:204) at org.jboss.tools.forge.ui.internal.ext.dialog.WizardDialogHelper.openWizard(WizardDialogHelper.java:102) at	0	0	https://issues.jboss.org/browse/FURNACE-104
4810	FURNACE-60	Would that be the prefix or the suffix ?	0	0	https://issues.jboss.org/browse/FURNACE-60
4811	FURNACE-26	Is this still an issue?	1	0	https://issues.jboss.org/browse/FURNACE-26
4812	FURNACE-21	Hi [~ozizka], Would you mind providing a failing test case illustrating this behavior? Thanks!	1	1	https://issues.jboss.org/browse/FURNACE-21
4813	FURNACE-12	Con you confirm that the maven central repo is overriden by a mirror in the settings.xml? If you could paste a sample of their settings.xml that would help also. Thanks	1	0	https://issues.jboss.org/browse/FURNACE-12
4814	FURNACE-12	Hi [~agoncal], Can you confirm this is still an issue with 2.15.0.Final? Thanks	1	0	https://issues.jboss.org/browse/FURNACE-12
4815	ESBSYSTEMTESTS-73	If we're extracting these classes, it might be a good idea to make a distinction between: - generic classes to access JMX, read files, ... that can be reused to test other projects (e.g. in kite/common) - ServiceMix-specific classes to ease building tests for FUSE ESB (e.g. in kite/servicemix) - the tests themselves and other artifacts we need for testing (e.g. in tests/*) Wdyt?	1	0	https://issues.jboss.org/browse/ESBSYSTEMTESTS-73
4816	FUSEFORGE-127	But point taken.. we should force https access to avoid credentials being passed around in the clear.	1	1	https://issues.jboss.org/browse/FUSEFORGE-127
4817	FUSEFORGE-118	should this be moved to a forge jira?	0	0	https://issues.jboss.org/browse/FUSEFORGE-118
4818	FUSEFORGE-115	How, other than as a "project directory mechanism" would people navigate the forge? I agree that all project pages should have some required common elements, but they should be customizable. I really don't understand the difference between sites/management and projects/MANAGEMENT though. Would a project have more than one site? Would a site have more than one project? (why?) If there's a 1-1 mapping between site and project then we should collapse 'em.	0	0	https://issues.jboss.org/browse/FUSEFORGE-115
4819	FUSEFORGE-115	My question "what's the difference" is perhaps misleading. I didn't mean to imply that they looked the same. In fact, the "sites" page actually looks good, and the "project" page in my opinion doesn't. At the same time, they offer pretty much the same options. I see the analogy with sourceforge, and I have to say I've never liked the way sourceforge does this - I've always found it confusing. At least initially, I would expect that people will come into the forge and navigate around rather than going to specific projects, so we should make that experience as clear as possible (and as a guy who's kind of walking in off the street, so to speak, my sense is that it's not all that clear right now - hence Jiras like this one). How's this (option 1): 1. For projects that don't choose to set up websites, the project page should look more like the "site" link rather than the "project" link in the above example. 2. If someone chooses to customize their project page to create a site, they should be able to do so and when I'm navigating the forge that's the site that I see. 3. We should have some control in place (like maybe a required navbar or something that has standard links like jira, wiki, source, downloads) or some such. Alternatively (option 2): 1. If we must have the "project" vs "site" distinction, can we make the "project" one really small in content, cleaner in design and clearly hold a pointer to the "site" if such a thing exists?	0	0	https://issues.jboss.org/browse/FUSEFORGE-115
4820	FUSEFORGE-104	So taking the management project as an example - are we getting rid of the generic page ? http://fusesource.com/forge/projects/MANAGEMENT And replacing the link off the projects page to be the custom pages - if there is one available ie . http://management.fusesource.org/	1	1	https://issues.jboss.org/browse/FUSEFORGE-104
4821	FUSEFORGE-104	The web site model that we are using on FUSE Forge is the same as SourceForge, where you will find the following: 1) A generic "project" site such as http://sourceforge.net/projects/docbook/, which is accessible from a list of projects (http://sourceforge.net/softwaremap/) 2) A "public" web site for the project, such as http://docbook.sourceforge.net I find this model confusing, since there is duplicate information on both sites. It's also hard to distinguish the purpose of each site. Contrast this with Apache, where if you go to the Apache home page (http://www.apache.org/) and click a project name, you are taken to the project's public home page (http://camel.apache.org/) So my question is, if we remove the generic "project" page entirely for public projects, as Amanda and Heather suggest, do we lose any functionality?	1	0	https://issues.jboss.org/browse/FUSEFORGE-104

#	Requirement	Questions	Answered	Successfully Answered	URL
4822	FUSEFORGE-79	Is this just for search engine optimization purposes?	0	0	https://issues.jboss.org/browse/FUSEFORGE-79
4823	FUSEFORGE-56	Can you suggest a better way to navigate people back to the project they came from then? There's no easy way to do that right now.	0	0	https://issues.jboss.org/browse/FUSEFORGE-56
4824	FUSEFORGE-50	Hi Jon, Could you start looking into this? This might be a good way to get you familiar with the forge rails app. Could you brainstorm what would be the best/easiest user interaction scenario to fulfill this requirement. Thx!	1	1	https://issues.jboss.org/browse/FUSEFORGE-50
4825	FUSEFORGE-50	Could we get the ICLA formatted more professionally, i..e. with progress letterhead etc. Also need it in PDF format. Also I don't understand the over 18 deal. Why can't that copy be in the ICLA?	1	1	https://issues.jboss.org/browse/FUSEFORGE-50
4826	FUSEFORGE-50	I don't understand how more explicit it can get? I doubt there anything more explicit than a legal contract with your signature on it saying your 18 right? How is an little web form click through going to get better than that? They do know I have no way of proving who actually clicked on the click through right? Regards, Hiram	1	1	https://issues.jboss.org/browse/FUSEFORGE-50
4827	FUSEFORGE-50	Aparently, some minor or their parents argued in court that they hadn't fully read the contract and and hadn't got to the bit about 18 so by forcing a click, we are covered against that. However, Jaime - can we split that bit out from this jira to a new one and add it for a future update of the forge and go live with what we have ? We have a lot to get through for Monday.	1	1	https://issues.jboss.org/browse/FUSEFORGE-50
4828	FUSEFORGE-42	how about you put a redirect into your static website - in a .htaccess file? Then you can update those redirects whenever you like in your git/svn repo?	1	1	https://issues.jboss.org/browse/FUSEFORGE-42
4829	FUSEFORGE-42	The site now supports .htaccess files for the project websites. The only problem is that WebDAV does not really let you edit them.	0	0	https://issues.jboss.org/browse/FUSEFORGE-42
4830	FUSEFORGE-27	We currently show a tag cloud on the "Search" page. Do you also want to show this tag cloud on the "Projects" page?	1	0	https://issues.jboss.org/browse/FUSEFORGE-27
4831	FUSEFORGE-27	ah yes so you do! :) Yes, lets show the tags on the projects page too (I'd not thought to look at the search page :) If its easy to do; showing the tag cloud at the top before showing the category boxes would be cool - maybe showing the tag cloud so its the width of both columns of 'category boxes' maybe? Also can we sort the tag cloud in alphabetical order?	0	0	https://issues.jboss.org/browse/FUSEFORGE-27
4832	FUSEFORGE-20	Hmm, weird. I did something called "Test user permissions" in phpbb. It lets me browse the forum using your permissions, and when I went to the board index, I did not see activeio, lightsabre, or jmstester. Can you try something for me? Log out of fusesource, delete your cookies and cache, and then login and see what the forum index screen shows?	1	1	https://issues.jboss.org/browse/FUSEFORGE-20
4833	FUSEFORGE-11	I am planning on changing the *external* website url for each project so that it would point to, for example, http://cloudmix.fusesource.org. But you bring up a good point, would this be the right thing to do for a non-opensource project? Do we need to capture whether the project is opensource, commercial, etc? I guess Muleforge does it by having the project admin select a license for their project, but we don't currently have that function implemented. Should we?	0	0	https://issues.jboss.org/browse/FUSEFORGE-11
4834	FUSEFORGE-11	BTW when stuff shows up at $PROJECT.fusesource.org - can we avoid the forge login/pwd stuff we've setup currently to keep the forge internal for now. i.e. as soon as the $PROJECT.fusesource.org capability is working - I'd like to start pointing folks on the web to some of the projects we've got in the forge - without necessarily having to wait for the full forge to go live. (i.e. I'm gagging to point folks at some of the projects we've got :)	1	1	https://issues.jboss.org/browse/FUSEFORGE-11
4835	FUSEFORGE-5	Debbie you going to work on this or what?	0	0	https://issues.jboss.org/browse/FUSEFORGE-5
4836	FUSEFORGE-5	Debbie, is there something you had in mind for this? Currently the text is "FUSE Forge is the developer center for hosting extensions to FUSE , the leading open source ESB and integration platform." which is a copy of MuleForge's welcome text. We go live next week so we should probably change this to avoid embarrassment :)	0	0	https://issues.jboss.org/browse/FUSEFORGE-5
4837	FUSEFORGE-5	Text supplied by Coco via email - Welcome to FUSE Forge! This is a community site for creating and sharing projects based on FUSE, the leading open source integration and messaging products. The Forge includes a complete infrastructure for developing software including a source code repository, integrated forums, a wiki, an issue tracker and licensing options to help you get the most out of Apache technology. More questions? View our FAQ [link]. Want to create a project or become a committer? View the contributor agreement [link] Search for a Project | Create a Project | Register	1	1	https://issues.jboss.org/browse/FUSEFORGE-5
4838	FUSEFORGE-5	Front page updated. Went with a slight variation of Coco's text: Focus on your application, not on software development infrastructure. The FUSE Forge is an open environment for the creation of integration projects and is designed to help you make the most out of FUSE technology. More questions? View our FAQ. Want to create a project or become a committer? View the contributor agreement	0	0	https://issues.jboss.org/browse/FUSEFORGE-5
4839	JANSI-37	Do you have a solution for my request please ? thanks	1	0	https://issues.jboss.org/browse/JANSI-37

#	Requirement	Questions	Answered	Successfully Answered	URL
4840	MOP-23	BTW the reason I introduced the "download" command was to differentiate the execution commands (jar/run/exec/war/spring/guice etc) from the install commands (install/upgrade/list/whatever else - the rpm/yum-ish stuff) - where download purely just fetches remote stuff into the local repo without doing any kind of execution or unpacking. e.g. you might want to download the jars/wars for stuff; so that in offline mode you can then execute stuff (without doing a binary install). If we were just building a kinda rpm/yum tool, things would be simple; we could just be like rpm/yum. Though I see mop as being a variety of things to different people * an alternative to mvn/ant/java command line tools for dev folks wanting to run some java, jar, war, spring, guice, ear, jbi, bundle, whatever from the command line * a neat way for developers to install/run open source software itself, example programs, test cases, soak tests, performance tests and so forth without having to delve into mvn/ant/shell stuff * a sysadmin-ish tool for installing/uninstalling/updating software packages like the binary distros of activemq, servicemix, jetty, tomcat et al The first two need to start in online mode unless the user turns it off. The sysadmin tool's install/update commands should definitely be in online mode by default. But a sysadmin person would mostly just be installing stuff right? e.g. installing packages & services - setting up daemons and not actually arbitrarily running jar/war/run/spring/guice stuff from the command line? So the only real difference the sysadmin role for mop is that the execution stuff tends to be in offline mode for sysadmins - if they use those commands at all?. I wonder if sysadmins would ever use the execution stuff? Wouldn't they prefer an alternative to run/exec/war/jar/spring/guice et al - where they create a daemon/service/script to run things (which are in offline mode - only using local stuff). Maybe we just need something like this... * mop [offline|online] to switch purely the execution stuff (everything but install/upgrade) to work in offline/online mode - this could then write to ~/.mop.properties the offline/online mode * install/upgrade default to online (would someone ever want an offline install - e.g. installing only from local repo?) - maybe someone who was super anal about what artifacts could be used could just limit the remote repos to a local-ish one on the intranet? * a sysadmin alternative to run/exec/jar/war/spring/guice et al - where instead of actually running things, we generate a kinda installation which has the same effect. For example; a sysadmin would install binary packages like the activemq broker; but if they wanted to run a load testing producer/consumer which is just a (say) executable jar or spring main - they may want to create an offline install of the executable. Not sure what to call it, lets maybe say something like "mop package someDir name <command>") e.g. {code} mop package fooDir123 runFoo jar myGroup:myArtifact:123 {code} Then in the directory *fooDir123* there would be a lib dir with all the jars required; then a script *runFoo* (with a .bat for windows) which sets the classpath and runs the jar?	0	0	https://issues.jboss.org/browse/MOP-23
4841	SCALATE-69	Hello, Could you tell us any information about this request?, do you need more information? Could you provide an estimate of time when you will give us an answers? Thanks in advance. Regards. Carlos Saavedra.	0	0	https://issues.jboss.org/browse/SCALATE-69
4842	SCALATE-69	Hi, Do you need more information?. Could you send us any answer? Regards. Carlos Saavedra.	0	0	https://issues.jboss.org/browse/SCALATE-69
4843	SCALATE-28	Would this work with singletons only? It would be nice if it also worked with default constructors of classes, then I could trigger dependency injection manually in the constructor (I think at least Guice allows that, but I haven't tried). Or maybe the uri scheme could be pluggable, for example "guice:org.acme.MyTags" and I could extend Scalate with a Guice provider. Or Spring etc.	1	0	https://issues.jboss.org/browse/SCALATE-28
4844	SCALATE-23	not sure about this one. The question is at what point does the value get put back into the context attribute map? Just doing it when the template finishes rendering would probably not be good enough since it would be more natural to expect it occur when the variable get set with the new value.	0	0	https://issues.jboss.org/browse/SCALATE-23
4845	SCALATE-15	ERB has a handy features where trailing newlines are stripped for code blocks terminated with: {code}-%>{code} How about we do the same? It also comes in handy for loop control statements. Example: {code} <% list.foreach { -%> ${_} <% } -%> {code}	1	1	https://issues.jboss.org/browse/SCALATE-15
4846	FRESH-10	Ales, it wouldn't be just one line :) The problem is that that ajaxgroovyshell deployes as servlet, but we don't have anything deployed as web app... Should we just provide sample web app that has this functionality as an add-on? How could we implement this so elegantly that it would just "magicly" apear in lets say jboss status page as additional "link"? The only question I have is packaging as functionality itself is not hard to implement. I would still like to keep simple deployment aka "drop one jar into deploy and everything works"	0	0	https://issues.jboss.org/browse/FRESH-10
4847	FRESH-5	I made it work with jboss threadPool that gets injected by mc Mare, can you please check if my fixes are ok? I am a bit worried about Process.startjoin() method where you do some kind of joining of processes...	0	0	https://issues.jboss.org/browse/FRESH-5
4848	JBFSWDOC-97	Anshu, I think you have done a really great job here. I really like the way you have provided a summary of each aspect to governance. Well done! Suggestions for improvement: - The rtgov-ui gadget server is being replaced in 6.1 and I don't think the following content will remain true - "Run-Time Governance uses a gadget server to display runtime governance information with a set of configurable gadgets. You can access these gadgets with the gadget server." Can you please confirm with Len (and cc me), or delete it. - S-RAMP - I think a lot of the information here is not interesting for customers. What do they really need to know about S-RAMP at this high-level? What's it for? Perhaps replace most of the content here with something like: "The S-RAMP repository is used as the storage mechanism for Design-Time Governance artifacts." Thanks, Ben	1	1	https://issues.jboss.org/browse/JBFSWDOC-97
4849	JBFSWDOC-97	Thanks, Anshu. I just have one more idea. We have at least two types of readers: one who wants to do minimal reading to get to what they are looking for, and one who wants to find out more about what governance has to offer. Do you think we can explore separating the governance information from the links to the consoles? What might that look like? Would it help our readers/users?	1	1	https://issues.jboss.org/browse/JBFSWDOC-97
4850	JBFSWDOC-96	Draft content - http://docbuilder.usersys.redhat.com/22489/#chap-JBoss_Fuse_Management_Console Waiting on content specific to what the console provides for SwitchYard. Will this be much the same as what is provided by the EAP Management Console?	1	0	https://issues.jboss.org/browse/JBFSWDOC-96
4851	JBFSWDOC-94	Thank you, Nichola! You have obviously put a lot of thought and effort into this which is much appreciated. Suggestions for improvement: Section - Administering JBoss Fuse Service Works Components - There are some instances of 'EAP'. These should be 'JBoss EAP'. - "Consoles are available for the following products:" -> "...components:" - There are also the CLIs and JBoss Operations Network. Do they belong here too? Section: Components of Red Hat JBoss Fuse Service Works - I have also created a table of components [1], but mine is also aiming to explain which component comes with what. To avoid confusion and inconsistencies, do you think it would be possible to consolidate/rearrange or something? - "Table 1.1. Add-user Command Arguments" - I think this table heading needs to be updated. [1] http://docbuilder.usersys.redhat.com/13666/#Application_Server_Compatibility	1	0	https://issues.jboss.org/browse/JBFSWDOC-94
4852	JBFSWDOC-31	Is this more relevant to administration?	0	0	https://issues.jboss.org/browse/JBFSWDOC-31
4853	JBFSWDOC-16	Do we have the option of choosing a JBoss Fuse server during JBDS installation?	1	0	https://issues.jboss.org/browse/JBFSWDOC-16

#	Requirement	Questions	Answered	Successfully Answered	URL
4854	FAB-13	we should allow a 'version range digit' header to be supplied to indicate how an exact version number gets translated to a range. e.g. for 2.5.6 of a dependency we could have possible values of... {code} 0 => "[2.5.6.qualifier,2.5.6.qualifier]" // exact only 1 => "[2.5.6.qualifier,2.5.7)" // exact only but allowing beta v ga releases etc 2 => "[2.5.6.qualifier,2.6)" // patch releases but not 2.6.x or later 3 => "[2.5.6.qualifier,3)" // patch and minor releases but not 3.x or later 4 => "2.5.6.qualifier" or "[2.5.6.qualifier,)" // the current OSGi default which just means anything at all >= 2.5.6 {code} A default value of *2* seems to make the most sense - in that we can then easily update bundles with emergency bug fixes but require a development / testing release cycle for more riskier updates	0	0	https://issues.jboss.org/browse/FAB-13
4855	FAB-13	the Fabric-Version-Range-Digits manifest entry is used now to define the number of variations of digits in version ranges used by default on all non-range OSGi versions on Import-Packages. The default value is 2 (i..e any micro increase in version)	0	0	https://issues.jboss.org/browse/FAB-13
4856	FAB-3	Why wouldn't the jar have the pom.xml inside it if it was built with mvn?	1	1	https://issues.jboss.org/browse/FAB-3
4857	ESB-1727	Further does aries.transaction.howl.bufferSizeKBytes rather need to be called aries.transaction.howl.bufferSize? [TransactionManagerService.java|https://svn.apache.org/repos/asf/aries/trunk/transaction/transaction-manager/src/main/java/org/apache/aries/transaction/internal/TransactionManagerService.java] suggests only the latter variable exists.	0	0	https://issues.jboss.org/browse/ESB-1727
4858	ESB-1522	xmlns:jaas="http://karaf.apache.org/xmlns/jaas/v1.0.0" xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-ext/v1.0.0"> <jaas:config name="karaf" rank="1"> <jaas:module className="org.apache.karaf.jaas.modules.ldap.LDAPLoginModule" flags="required"> initialContextFactory=com.sun.jndi.ldap.LdapCtxFactory connection.username=uid=admin,ou=system connection.password=secret connection.protocol= connection.url=ldap://localhost:10389 user.base.dn=ou=users,ou=system user.filter=(uid=%u) user.search.subtree=true role.base.dn=ou=groups,ou=system role.name.attribute=cn role.filter=(member=uid=%u) role.search.subtree=true authentication=simple </jaas:module> </jaas:config> </blueprint> {code} with a role {code} dn: cn=admin,ou=groups,ou=system objectClass: groupOfNames objectClass: top cn: admin member: uid=jdoe {code} and a user {code} dn: uid=jdoe,ou=users,ou=system objectClass: inetOrgPerson objectClass: organizationalPerson objectClass: person objectClass: top cn: John Doe sn: Doe uid: jdoe userPassword:: e1NIQX01ZW42RzZNZXpScm9UM1hLcWtkUE9tWS9CZlE9 {code} but ServiceMix will not authenticate jdoe in role admin: {code} 16:50:50,691 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Create the LDAP initial context. 16:50:50,691 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Bound access requested. 16:50:50,691 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Get the user DN. 16:50:50,691 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Initialize the JNDI LDAP Dir Context. 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Define the subtree scope search control. 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Looking for the user in LDAP with 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | base DN: ou=users,ou=system 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | filter: (uid=jdoe) 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Get the user DN. 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Bind user (authentication). 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Set the security principal for uid=jdoe,ou=users,ou=system 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Binding the user. 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | User jdoe successfully bound. 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Get user roles. 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | Looking for the user roles in LDAP with 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | base DN: ou=groups,ou=system 16:50:50,706 | DEBUG | NioProcessor-2 | LDAPLoginModule | 16 - org.apache.karaf.jaas.modules - 2.2.0.fuse-00-43 | filter: (member=uid=jdoe) 16:50:50,722 | DEBUG | NioProcessor-2 | KarafJaasPasswordAuthenticator | 17 - org.apache.karaf.shell.ssh - 2.2.0.fuse-00-43 | User authentication failed with User does not have the required role admin javax.security.auth.login.FailedLoginException: User does not have the required role admin {code} What am I missing?	1	1	https://issues.jboss.org/browse/ESB-1522
4859	ESB-1522	Hi Eric, Which LDAP server are you using? Nothing is jumping out at me but most likely it is something with the role.filter setting. Depending on which LDAP server you are using you may need an extra colon in there Also, I attached my ldif file which you can import and use the settings I was using. For ApacheDS, you may have to add a partition for dc=fusesource,dc=com. Here's the note I sent to Lowry on this: {quote} yes that's a "partition" that you have to add first. This link should help: http://hasini-gunasinghe.blogspot.com/2011/07/how-to-created-new-partition-in.html The server.xml file is a bit buried (./instances/default/conf/server.xml). I added a partition (me): <partitions> <!-- NOTE: when specifying new partitions you need not include those --> <!-- attributes below with OID's which are the system indices, if left -> <!-- out they will be automatically configured for you with defaults. --> <jdbmPartition id="example" cacheSize="100" suffix="dc=example,dc=com" optimizerEnabled="true" syncOnWrite="true"> <indexedAttributes> <jdbmIndex attributeId="1.3.6.1.4.1.18060.0.4.1.2.1" cacheSize="100"/> <jdbmIndex attributeId="1.3.6.1.4.1.18060.0.4.1.2.2" cacheSize="100"/> <jdbmIndex attributeId="1.3.6.1.4.1.18060.0.4.1.2.3" cacheSize="100"/> <jdbmIndex attributeId="1.3.6.1.4.1.18060.0.4.1.2.4" cacheSize="100"/> <jdbmIndex attributeId="1.3.6.1.4.1.18060.0.4.1.2.5" cacheSize="10"/> <jdbmIndex attributeId="1.3.6.1.4.1.18060.0.4.1.2.6" cacheSize="10"/> <jdbmIndex attributeId="1.3.6.1.4.1.18060.0.4.1.2.7" cacheSize="10"/> <jdbmIndex attributeId="dc" cacheSize="100"/> <jdbmIndex attributeId="ou" cacheSize="100"/> <jdbmIndex attributeId="krb5PrincipalName" cacheSize="100"/> <jdbmIndex attributeId="uid" cacheSize="100"/> <jdbmIndex attributeId="objectClass" cacheSize="100"/> </indexedAttributes> </jdbmPartition> <jdbmPartition id="faa" cacheSize="100" suffix="dc=dex" optimizerEnabled="true" syncOnWrite="true"> </jdbmPartition> <jdbmPartition id="me" cacheSize="100" suffix="dc=fusesource,dc=com" optimizerEnabled="true" syncOnWrite="true"> </jdbmPartition> </partitions> {quote} Hope that helps, Susan	1	1	https://issues.jboss.org/browse/ESB-1522
4860	ESB-1456	Hi Gert, The customer was wondering if we could make the message more specific, e.g: {code} XXXXXX> shutdown Do you really want to shutdown (yes/no): {code} Will thel user realize this message means all of servicemix will shut down? Thanks, Susan	1	0	https://issues.jboss.org/browse/ESB-1456
4861	ESB-1448	Tracked at Apache with https://issues.apache.org/jira/browse/KARAF-590 which will be included in Karaf 2.2.x branch	1	1	https://issues.jboss.org/browse/ESB-1448
4862	ESB-1441	Sections 1 and 3 looks ok. In the table 2.1, why can a Library "Maybe" import its own package ? I think if that's the case, it's mostly an API bundle then, not a library anymore.	0	0	https://issues.jboss.org/browse/ESB-1441

#	Requirement	Questions	Answered	Successfully Answered	URL
4863	ESB-1438	@Freeman: Could you also backport this fix into the karaf-2.1.x-fuse branch ? Thanks !	0	0	https://issues.jboss.org/browse/ESB-1438
4864	ESB-1432	We may want to investigate switching to the geronimo specs. They've rewritten the logic in order to add some improvements, but i can't recall which one exactly. But as a work around for the next patch releases, the BundleFactoryLoader should definitely cache the loaded class, but i'd rather keep the url as the main way to retrieve the class, as that does not need the bundle to be resolved.	1	1	https://issues.jboss.org/browse/ESB-1432
4865	ESB-1413	Can we have a bit more details about the use cases / scenario ? Do they have problem writing the manifest or at deployment time ?	1	0	https://issues.jboss.org/browse/ESB-1413
4866	ESB-1352	@Freeman: Could you take a look at this?	1	1	https://issues.jboss.org/browse/ESB-1352
4867	ESB-1315	Providing doc for the OSGi Event Admin service has been added to our list of doc tasks. It won't make it into the ESB 4.4 release, however.	0	0	https://issues.jboss.org/browse/ESB-1315
4868	ESB-1308	Is this only relevant for AMQ and SMX? What about CXF and Camel?	1	0	https://issues.jboss.org/browse/ESB-1308
4869	ESB-1304	Thanks to Lars for providing these patches -- both have been applied at Apache and fixes have been merged into the karaf-2.1.x-fuse branch.	0	0	https://issues.jboss.org/browse/ESB-1304
4870	ESB-1302	This issue is a bit vague. Do you know exactly what kind of problems they are having configuring Fuse ESB?	0	0	https://issues.jboss.org/browse/ESB-1302
4871	ESB-1301	I would like to confirm what kind of problem the user is encountering here. Is it something like this? Say we have bundle A that depends on bundle B and on bundle C. Bundle B depends on log4j version 1.1 and bundle C depends on log4j version 1.2. Now we have a conflict, because bundle A will try to pull in both version 1.1 and version 1.2 of log4j into the same Classpath. Is that the essence of the problem we are talking about here?	0	0	https://issues.jboss.org/browse/ESB-1301
4872	ESB-1271	Would the standard *{{Bundle-Description}}* suffice here? From the spec: "The Bundle-Description header defines a short description of this bundle." This header would already be visible in the default osgi:headers output	0	0	https://issues.jboss.org/browse/ESB-1271
4873	ESB-1220	Hey Gert, would it be possible to also move the external dependencies and properties to the parent/pom.xml file as well?	0	0	https://issues.jboss.org/browse/ESB-1220
4874	ESB-1211	Gertv this is implemented in the source code? Its just that it will be in FUSE ESB 4.3.0-fuse-01-00 release?	1	1	https://issues.jboss.org/browse/ESB-1211
4875	ESB-1185	Hey Guillaume, is there a snapshot of pax logging we can test out? Thanks!	1	1	https://issues.jboss.org/browse/ESB-1185
4876	ESB-1182	Gullaume, Couldn't we just make the Camel NMR endpoint implement the AsyncProcessor interface already? It would only be useful when doing a .toAsync("nmr:EndpointName") for now, but at least we could get that bit done for ESB 4.3... Regards, Gert	1	1	https://issues.jboss.org/browse/ESB-1182
4877	ESB-1176	Could you update the JBI component at the Camel wiki page?	1	0	https://issues.jboss.org/browse/ESB-1176
4878	ESB-1152	Chapter 7 comments page 34: Apache Camel 4.2 -> Apache ServiceMix 4.2 (and shouldn't it be FUSE ESB 4.2 ?) page 35: Customizing the Spring configuration. Sounds a bit unclear about the root replace? What you actually just have to do is to replace the namespace from the old 1.x style to the 2.x style. See more at: http://fusesource.com/wiki/display/ProdInfo/Fuse+Mediation+Router+2.2+Migration+Tips And its very annoying that you refer to another document for more details, as this is a mandatory migration thing and I think we should have all that detailed in the migration guide. Then for any optional or more in depth you can refer to other guides. page 35: Blueprint? Should we really mention blueprint at this time. Ask gnodet if that is ready for enterprise and something we support/recommend already. page 36: The commons-managment is already in central maven repo, so no need to mention it and that you need to add a maven repository to your pom. In fact do not mention it at all. page 36: Is this migration guide for running Camel in FUSE ESB 4.2? If so it comes with a camel-spring feature which means it install all the right bundles for your out of the box. Those .jar dependencies are only for users not using FUSE ESB 4.2, eg. running in Tomcat or standalone etc. page 40: The exchange copy methods is really only for internal usage or really hardcore advanced Camel stuff, mostly if you develop your own component and have a very special case. I would simply just remove this section as it has no relevant for regular Camel end users. page 45: I do not think we should mention Spring 3.0 support as its not officially supported in all our products. So I think this section should be removed. page 45: Interaction between filer and split: This was a new feature in 2.0 and was later removed. So there is no migration needed between Camel 1.x and 2.2 page 45: error handling. Camel 2.x uses a new DefaultErrorHandler instead of DeadLetterChannel as in 1.x, as the default error handler. This new DefaultErrorHandler does *not* catch and handle thrown exceptions, which means any exception thrown will be propagated back to the caller. page 45: Stream caching. The interaction between SMX and Camel is now handling streams in such a way that Camel does not have to use stream caching anymore. Its thus disabled by default.	1	1	https://issues.jboss.org/browse/ESB-1152
4879	ESB-1149	Complete draft of Fuse ESB 4.2 "Migration Guide" for review.	1	1	https://issues.jboss.org/browse/ESB-1149
4880	ESB-1149	Flintan I just want to note that users upgrading from FUSE ESB 4.1 to 4.2 may have problems with 2 bundles - camel-spring - camel-osgi Have been merged into 1 single bundle in FUSE ESB 4.2 - camel-spring-osgi Which is the bundle they *must* used, and the others should *not* be used. So when you upgrade there must be some procedure in place to ensure this. There may be some SMX / Karaf configuration file in ESB 4.1 which lists the old bundles. Which you must upgrade / ensure only lists the new single bundle. However this is just a hunch. Get in touch with the SMX / Karaf experts to be sure. (gnodet, gertv, willem jiang, ccustine, freeman etc.)	0	0	https://issues.jboss.org/browse/ESB-1149
4881	ESB-1138	Section "Defining Service Beans" notes : If you are using the Spring DM container version 2.0 or later, it is actually possible to mix both kinds of bean element under the beans root element, (as long as you define the bean elements using the appropriate namespace prefix). I was just temporarily confused by thinking it was referring to the commercial Spring DM 2.0 server...May be it is possible to change "If you are using the Spring DM container version 2.0 or later, it is actually possible to mix both kinds.." to something like "Note that Spring DM specification 2.0 allows for mixing both kinds..."	0	0	https://issues.jboss.org/browse/ESB-1138

#	Requirement	Questions	Answered	Successfully Answered	URL
4882	ESB-1134	the basic deployable unit by users already familiar with OSGi. It might be better to re-phrase as something like: "it is often convenient to aggregate inter-dependent or related bundles into a larger unit of deployment." --- p50. "Essentially, a feature is created by adding a new feature element to a special kind of XML file, known as a feature repository." Often referred to synonymously as as a "features URL". Neither terminology is really satisfactory, especially the former as it connotes a centralized location where all the bundles are stored (like a maven repository or the OBR) which of course it isn't. So it might make sense to emphasize the point that a features repository is really just an aggregation of _references_ to inter-dependent or related bundles, not a location where the bundles _themselves_ are stored. "Create a *local* feature repository" The reference to locality may be confusing as the standard features repos shipping with the ESB would also be available locally (under the <smx_install_dir>/system dir) as well as being remotely accessible via the fusesource maven repo. You might want to call it a _custom_ features repository instead. "<features> </features>" The top-level features element may optionally be decorated with a name attribute. See for example {code} <smx_install_dir>/system/org/apache/felix/karaf/apache-felix-karaf/1.4.0-fuse-01-00/apache-felix-karaf-1.4.0-fuse-01-00-features.xml: <features name="karaf-1.4.0-fuse-01-00"> {code} If the name isn't specified, Karaf will default to a generic name like "repo-0". "Add the local repository URL to the features service" Cart before the horse somewhat? Surely some feature elements should be defined in the new features repo _before_ running features:addUrl? "In order to make the new feature repository available to the Apache Felix *kernel* ..." Repeated used of 'kernel' in this sentence may confuse, given that we switched over from the old ServiceMix Kernel to Apache Felix Karaf for this release. May be better as: "In order to make the new feature repository available to the Karaf features service ..." --- p51. {code} karaf@root> features:listUrl mvn:org.apache.servicemix.nmr/apache-servicemix-nmr/1.1.0-psc-01-00RC1/xml/features mvn:org.apache.servicemix.camel/features/4.1.0-psc-01-00RC1/xml/features file:C:/Projects/features.xml mvn:org.apache.ode/ode-jbi-karaf/1.3.3-psc-01-00RC1/xml/features mvn:org.apache.felix.karaf/apache-felix-karaf/1.2.0-psc-01-00RC1/xml/features mvn:org.apache.servicemix/apache-servicemix/4.1.0-psc-01-00RC1/xml/features {code} Replace all "psc-01-00RC1" with "fuse-01-00". "You can optionally specify a version attribute on the feature element, if you want to keep track of the feature's version number." It would be more a case of assigning a non-zero version to the feature, as opposed to just keeping track of the version. When versions are used, a specific version of a particular feature can be specified, e.g. {code} features:install cxf-jaxrs 4.2.0-fuse-01-00 {code} "To check whether the *kernel* successfully parses the new feature entry" Similar comment to the above, probably best to refer to it as the features service. "but you should be able to find the entry for your new feature (in this case, example-camel-bundle) by scrolling back through the listing." or by simply shortening the list to only those features are actually installed via the \-i,\-\-installed option: {code} karaf@root> features:list -i {code} or by grep'ing for the particulat feature {code} karaf@root> features:list | grep example-camel-bundle {code} "The features:refreshUrl command forces the kernel to reread *all* the feature repositories" Unless you specify a URL argument, in which case only that repo is refreshed. --- p52. "Each child feature element contains the name of a feature on which the current feature depends." A version can be specified as well, e.g. {code} <feature name='camel-core' version='2.2.0-fuse-01-00'> <feature version="2.5.6.SEC01">spring</feature> ... <feature> {code} This encodes a dependency of the camel-core feature (v. 2.2.0-fuse-01-00) on the spring feature (v. 2.5.6.SEC01). "When you deploy a feature with dependent features, the dependency mechanism checks whether or not the dependent features are *available* in the container. If not, the dependency mechanism automatically *loads* the missing dependencies (and any recursive dependencies)." Maybe better to use "installed" and	1	1	https://issues.jboss.org/browse/ESB-1134
4883	ESB-1134	Giving a blueprint example of property-placeholder substitution sounds like a good idea, but I don't have any examples for it. The only think I could find on it is the following page from Karaf: http://felix.apache.org/site/45-security-framework.html This example seems to exploit a non-standard feature of Apache Geronimo. Is this the recommended approach?	1	1	https://issues.jboss.org/browse/ESB-1134
4884	ESB-1088	Ulhas From looking at the comments on DEV-1775, it looks like the format of the configuration filename should be: org.apache.servicemix.components.<component>.cfg For example, org.apache.servicemix.components.jms.cfg and org.apache.servicemix.components.file.cfg Am I correct?	1	1	https://issues.jboss.org/browse/ESB-1088
4885	ESB-1068	Has this been fixed in 4.2? How can i check? Pedro	1	0	https://issues.jboss.org/browse/ESB-1068
4886	ESB-963	Updated versions of commons-* libraries. Updated from ODE's custom build of saxon to 9.1.0.8. I have a patch to also update: - xerces to 2.9.1 - xmlbeans to 2.3 - wsdl4j to 1.6.2 but we haven't done extensive testing on this yet so it may be too late in the release cycle to apply this change.	1	1	https://issues.jboss.org/browse/ESB-963
4887	ESB-958	Hello Aileen, Do I need to start working on karaf? Also, Gert commented earlier about releasing FileInstall, what has to be done on that? Rahul	1	1	https://issues.jboss.org/browse/ESB-958
4888	ESB-958	Hello, Can Gert provided the instructions of releasing FileInstall? Rahul	1	1	https://issues.jboss.org/browse/ESB-958
4889	ESB-958	Hello, Karaf 1.2.0-psc-01-00M2 has been released. Repo URL: http://repo.open.iona.com/restricted/org/apache/felix/karaf/ SVN URL: http://fusesource.com/forge/svn/kernel/tags/karaf-1.2.0-psc-01-00M2/ Source: http://fusesource.com/forge/svn/kernel/trunk/ Can you please let me know when NMR is ready for release? Rahul	1	1	https://issues.jboss.org/browse/ESB-958
4890	ESB-958	Hello Aileen, I am using the mentioned component source tree. However, I am getting test failures while doing release:prepare. I am attaching the report and the prepare run log file for your reference. Can you have someone look at it? Rahul	1	1	https://issues.jboss.org/browse/ESB-958
4891	ESB-958	@Chris, you maybe right I ran it again and it failed at some other tests. So am trying it for the 3rd time and hopefully it completes successfully. @Aileen, I am working on components. Should I wait or at-least complete components? Rahul	1	1	https://issues.jboss.org/browse/ESB-958
4892	ESB-958	@Aileen, can you mention where can I find the source branch for BPEL? Rahul	1	1	https://issues.jboss.org/browse/ESB-958
4893	ESB-958	Hello, While releasing Features, SNAPSHOT was detected in one of the example poms. <ode-version>1.3.3-fuse-SNAPSHOT</ode-version> Can I change this to BPEL version released? <ode-version>1.3.3-psc-01-00M2</ode-version> Rahul	1	1	https://issues.jboss.org/browse/ESB-958
4894	ESB-958	Yes, can you change this to the bpel version released on Friday. Aileen	1	1	https://issues.jboss.org/browse/ESB-958
4895	ESB-958	Hello, While doing a release of features, there is a missing artifact. http://teamcity.boston.amer.iona.com:8111/viewLog.html?buildId=248856&tab=buildResultsDiv&buildTypeId=bt903 Artifact: org.apache.servicemix:apache-servicemix:pom:4.1.0-psc-01-00M2 This artifact has not been created as part of features deploy but its looking for it. So as part of the build we are missing, http://fusesource.com/forge/svn/fuseesb/smx4/features/trunk/cxf/cxf-transport-osgi/ Can some one let me know what has to be done here? Rahul	1	1	https://issues.jboss.org/browse/ESB-958
4896	ESB-944	If I understand this issue correctly, the customer has implemented its own JAAS LoginModule which basically makes all their users known to the container and they now want to limit access to the JMX Console to only a subset of those users (e.g. users that are a member of some group). Is this correct? Would a solution like the one proposed in attachment ESB-944.diff be sufficient or do we still need the per-call authorization check as is suggested in the second scenario in the word document?	1	1	https://issues.jboss.org/browse/ESB-944

#	Requirement	Questions	Answered	Successfully Answered	URL
4897	ESB-944	Ashwin, If the goal is to allow users access based on the group they belong to, that should be possible by just looking at the subject - after a successful login with the JAAS LoginModule, the Subject will have Principal added for the groups so we can just look at those to determine if a user is a member of the admin group or whatever group we configure. Your example uses a call to doPrivileged() for the subject, but I'm not entirely sure what the goal of that block of code is. Why does the code need to executed with additional permissions granted? Regards, Gert	1	1	https://issues.jboss.org/browse/ESB-944
4898	ESB-944	Dave, Do you have any more details concerning this problem in JDK 1.5 (what happens? are they seeing exceptions or is the authorization just not being handled correctly without any errors in the log? ...) Thanks, Gert	1	1	https://issues.jboss.org/browse/ESB-944
4899	ESB-882	> Camel 1.6.1.0M1-fuse (new release) Aileen is this a bit confusing? Are you saying it should use a Camel 1.x version? I thought it was the first release to use Camel 2.0	1	1	https://issues.jboss.org/browse/ESB-882
4900	ESB-882	Hello Aileen, We do not have any instructions to build Karaf. Can you let me know with whom can I work to get those instructions for releasing Karaf? Rahul	1	1	https://issues.jboss.org/browse/ESB-882
4901	ESB-882	Aileen, While building nmr, I an getting some snapshots detected error. ./pom.xml: <felix.karaf.version>1.1.0.0-fuse-SNAPSHOT</felix.karaf.version> ./pom.xml: <servicemix.eip.version>2009.02-fuse-SNAPSHOT</servicemix.eip.version> ./pom.xml: <servicemix.http.version>2009.02-fuse-SNAPSHOT</servicemix.http.version> ./pom.xml: <servicemix.jsr181.version>2009.02-fuse-SNAPSHOT</servicemix.jsr181.version> ./pom.xml: <servicemix.shared.version>2009.02-fuse-SNAPSHOT</servicemix.shared.version> ./nmr/fuse-management/pom.xml: <version>1.0-SNAPSHOT</version> Can you confirm to what these SNAPSHOTs are to be changed? Ref: http://teamcity.boston.amer.iona.com:8111/viewLog.html?buildId=235334&tab=buildResultsDiv&buildTypeId=bt902 Rahul	1	1	https://issues.jboss.org/browse/ESB-882
4902	ESB-882	Gert, while working on releasing Karaf, I am getting the error as attached in the JIRA. (kernel_mvn_prepare.log). The error is, [INFO] Failed to resolve artifact. GroupId: org.eclipse ArtifactId: osgi Version: 3.5.0.v20090520 I am running the following command at the root pom.xml level, mvn release:prepare -Prelease Can you let me know how can I resolve this artifact? Rahul	1	1	https://issues.jboss.org/browse/ESB-882
4903	ESB-882	Rahul, Guillaume has changed the karaf build so it no longer requires the org.eclipse.osgi sources. Could you give that another go? Thanks, Gert	1	1	https://issues.jboss.org/browse/ESB-882
4904	ESB-882	Hello Aileen, The earlier question is unanswered for NMR release. Karaf => 1.0.0.0M1-fuse and what about others? FYI..... Aileen, While building nmr, I an getting some snapshots detected error. ./pom.xml: <felix.karaf.version>1.1.0.0-fuse-SNAPSHOT</felix.karaf.version> ./pom.xml: <servicemix.eip.version>2009.02-fuse-SNAPSHOT</servicemix.eip.version> ./pom.xml: <servicemix.http.version>2009.02-fuse-SNAPSHOT</servicemix.http.version> ./pom.xml: <servicemix.jsr181.version>2009.02-fuse-SNAPSHOT</servicemix.jsr181.version> ./pom.xml: <servicemix.shared.version>2009.02-fuse-SNAPSHOT</servicemix.shared.version> ./nmr/fuse-management/pom.xml: <version>1.0-SNAPSHOT</version> Can you confirm to what these SNAPSHOTs are to be changed? Ref: http://teamcity.boston.amer.iona.com:8111/viewLog.html?buildId=235334&tab=buildResultsDiv&buildTypeId=bt902 Rahul	1	1	https://issues.jboss.org/browse/ESB-882
4905	ESB-854	@Guillaume How can Karaf know which version of Camel you are using if you have specified a versioned scheme? It all maps to the same {{CamelNamespaceHandler}} class. {code} http\://camel.apache.org/schema/spring=org.apache.camel.spring.handler.CamelNamespaceHandler {code} So I fail to see how, lets say you use {{camel-spring-2.3.0.xsd}} in bundle A, and then {{camel-spring-2.4.0.xsd}} in bundle B. How can Karaf know that bundle A should leverage Camel v2.3.0 only?	0	0	https://issues.jboss.org/browse/ESB-854
4906	ESB-813	will this fix make it into Apache as well ? Also, CXF 2.2.4 will be used in M2 so this fix will have to be done in CXF	1	1	https://issues.jboss.org/browse/ESB-813
4907	ESB-778	Hello, While updating the POMs file for components, I came across following SNAPSHOTs version for servicemix-shared-version. http://projects.open.iona.com/projects/svn/iona/servicemix/components/trunk/ bindings/servicemix-jms/pom.xml: <servicemix-shared-version>2009.01.0-fuse-SNAPSHOT</servicemix-shared-version> engines/servicemix-quartz/pom.xml: <servicemix-shared-version>2009.01.0-fuse-SNAPSHOT</servicemix-shared-version> Can you let me know what versions these SNAPSHOTs needs to be replaced with? Thanks, Rahul	1	1	https://issues.jboss.org/browse/ESB-778
4908	ESB-765	Can you start on 3.4 release build?	0	0	https://issues.jboss.org/browse/ESB-765

#	Requirement	Questions	Answered	Successfully Answered	URL
4909	ESB-754	following exception: {code} java.io.FileNotFoundException: \\.host\Shared Folders\.servicemix\servicemix.history (The system cannot find the path specified) {code} See the full stack trace: {code} org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'localShellWrapped' defined in URL [bundle://8.0:0/META-INF/spring/gshell-local.xml]: Cannot create inner bean 'org.apache.geronimo.gshell.wisdom.shell.HistoryImpl#150a0' of type [org.apache.geronimo.gshell.wisdom.shell.HistoryImpl] while setting bean property 'history'; nested exception is org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'org.apache.geronimo.gshell.wisdom.shell.HistoryImpl#150a0' defined in URL [bundle://8.0:0/META-INF/spring/gshell.xml]: Invocation of init method failed; nested exception is java.io.FileNotFoundException: \\.host\Shared Folders\.servicemix\servicemix.history (The system cannot find the path specified) at org.springframework.beans.factory.support.BeanDefinitionValueResolver.resolveInnerBean(BeanDefinitionValueResolver.java:230) at org.springframework.beans.factory.support.BeanDefinitionValueResolver.resolveValueIfNecessary(BeanDefinitionValueResolver.java:117) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.applyPropertyValues(AbstractAutowireCapableBeanFactory.java:1245) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.populateBean(AbstractAutowireCapableBeanFactory.java:1010) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.doCreateBean(AbstractAutowireCapableBeanFactory.java:472) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory$1.run(AbstractAutowireCapableBeanFactory.java:409) at java.security.AccessController.doPrivileged(Native Method) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.createBean(AbstractAutowireCapableBeanFactory.java:380) at org.springframework.beans.factory.support.AbstractBeanFactory$1.getObject(AbstractBeanFactory.java:264) at org.springframework.beans.factory.support.DefaultSingletonBeanRegistry.getSingleton(DefaultSingletonBeanRegistry.java:222) at org.springframework.beans.factory.support.AbstractBeanFactory.doGetBean(AbstractBeanFactory.java:261) at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:185) at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:164) at org.springframework.beans.factory.support.DefaultListableBeanFactory.preInstantiateSingletons(DefaultListableBeanFactory.java:429) at org.springframework.context.support.AbstractApplicationContext.finishBeanFactoryInitialization(AbstractApplicationContext.java:728) at org.springframework.osgi.context.support.AbstractDelegatedExecutionApplicationContext.access$1600(AbstractDelegatedExecutionApplicationContext.java:68) at org.springframework.osgi.context.support.AbstractDelegatedExecutionApplicationContext$4.run(AbstractDelegatedExecutionApplicationContext.java:343) at org.springframework.osgi.util.internal.PrivilegedUtils.executeWithCustomTCCL(PrivilegedUtils.java:85) at org.springframework.osgi.context.support.AbstractDelegatedExecutionApplicationContext.completeRefresh(AbstractDelegatedExecutionApplicationContext.java:308) at org.springframework.osgi.extender.internal.dependencies.startup.DependencyWaiterApplicationContextExecutor$CompleteRefreshTask.run(DependencyWaiterApplicati	1	0	https://issues.jboss.org/browse/ESB-754
4910	ESB-725	Thanks very much Freeman! I just wanted to verify the following: Will the next version of sf (with the codegen pom fix) definitely be released before the next version of esb 4.1 features is released? Will the next version of esb 4.1 features be upgraded to the sf with the fix (once sf is released)? Just wondering because I am doing a certification for aix 64bit, and when I see an error like this in the teamcity aix nightly build for instance, I don't want to assume that it will be fixed by the time the release happens if it is possible that esb will be released without sf being released, or without being upgraded to the new version of sf. Would be best for me to be able to see a green build in teamcity for the aix nightly for instance, well in advance of the release process. :) By the way, I asked in another jira, but didn't get much response, do you know if its documented anywhere exactly which FUSE branches are used for which releases? i.e. will the next version of activemq (5.3.0.3-fuse) be released from this branch: https://projects.open.iona.com/projects/svn/iona/activemq/branches/activemq-5.3.0-fuse/ or from trunk: https://projects.open.iona.com/projects/svn/iona/activemq/trunk/ Even if you only know the answer of which branches are used for the servicemix releases, or can only point me to who I can ask who might know the answers, that would be a great help!	1	1	https://issues.jboss.org/browse/ESB-725
4911	ESB-710	Can you add a link to the docs that have been updated. Thanks.	1	1	https://issues.jboss.org/browse/ESB-710
4912	ESB-710	HI Orlan, Ok and when will they be available in FUSE? And what would be the FUSE link. ALso, can the 'fix version' be set to 4.2. Currently there is no plan for a 4.2 and this may be misleading to FUSE customers. Thanks, Matt	1	1	https://issues.jboss.org/browse/ESB-710
4913	ESB-686	Hi Torsten, One problem of this patch is that if we don't use X509Certificate from ws-security to do the authentication, we just sign the message, then will get some exception like java.lang.SecurityException: Unable to locate a login configuration at com.sun.security.auth.login.ConfigFile.<init>(ConfigFile.java:97) at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:39) at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:27) at java.lang.reflect.Constructor.newInstance(Constructor.java:494) at java.lang.Class.newInstance0(Class.java:350) at java.lang.Class.newInstance(Class.java:303) at javax.security.auth.login.Configuration$3.run(Configuration.java:216) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.login.Configuration.getConfiguration(Configuration.java:210) at javax.security.auth.login.LoginContext$1.run(LoginContext.java:237) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.login.LoginContext.init(LoginContext.java:234) at javax.security.auth.login.LoginContext.<init>(LoginContext.java:367) at javax.security.auth.login.LoginContext.<init>(LoginContext.java:444) at org.apache.servicemix.jbi.security.auth.impl.JAASAuthenticationService.authenticate(JAASAuthenticationService.java:52) Seems we still need add a flag in CxfBcConumer to indicate if we need the X509Certificate to do the authentication, Regards Freeman	1	1	https://issues.jboss.org/browse/ESB-686

#	Requirement	Questions	Answered	Successfully Answered	URL
4914	ESB-612	Here's a first set of comments. I had been hoping to try the logisticx demo with 4.1 but I don't think a 4.1 version is available yet. 'Getting started with FUSE product family' - the legal notices talk about IONA Technologies Plc. I don't know if that should be Progress Software Inc if they want to get permission to re-use it. p1 - why Use Integration I found the first reason very difficult to read and understand. Is the second part 9'... and implementing') necessary. While It's in the dictionary all right, I'd never heard the word foundational before. But that could be because it's used more in American English p1. Using an ESB in Integration An ESB often uses XML but doesn't really have to I don't think. Later on, should the E of 'enterprise Java Beans' be capitalised. I don't know what the Progress style guide says about how to write Web service, I just remember IONA used to capitalised the W, so I just wanted to ask if we're now doing it in lower case. p15 FUSE ESB - very small issue. Do we want to put OSGi support before JBI since that's the way we want to push people? p16 - Usage in the Demo - 'to a client web application that provides a web browser for customer interface' confused me, should it be a web interface, not a web browser? p17 FUSE SF - REST seems out of places in the list of JAX-RS, JavaScript, REST, or POJOs, since it's more an approach than a technology like the others. p22 - FUSE Dev Models JMS messaging - XML isn't the only thing you can send on JMS, you can send streams of bytes, Java Objects amongst others. Later on in that paragraph, should it be message broker or messaging broker. p25 - should we encourage people to download ESB 3, or always suggest ESB 4 now? I don't know if JBI support in 4 is complete though. p26 - run ESB as a standalone - the standalone section sems to be FUSE ESB being the container, whereas in the intro for this section we talk about 'Standalone or in an FUSE ESB container' (and should that be a not an?) p26 - EIP (JBI Container) if it's deploying as JBI, then it's not really standalone is it? It's probably going ot be in ServiceMix/ESB 3 or another JBI container. p27 - Should we provide a brief paragraph or reference about Spring in case the user got this far without knowing what Spring is p27 - Can you deploy your JBI or OSGi apps directly in an app server? Are we in dnager of telling the user too many ways to deploy but not recommending some over others. p31 - why poll filing system	0	0	https://issues.jboss.org/browse/ESB-612
4915	ESB-612	Sorry this next bit took so long. I was hoping to get to try it in version 4.1 but I'm getting too many weird errors on startup. Anyway, here it is. I tried it in FUSE ESB 4.0.0.4 but have a couple of 4.1 related comments p31 - polling a file system seems to be an implementation issue rather than a business task, maybe some context p34 - Data directory doesn't always have \ could be / if on Unix. This appears on other pages too, e.g. p35 p34 - When we tell users to remove SAs deployed, is that not just SAs relating to this demo? They may have their own SA's, or are we assuming that if they are reading a getting started guide that they won't have anything important in there? p35 - deploying the app before servicemix has started up has caused me problems in the past (with the underlying system not being initiated in time) and there are problems like this in the 4.1 kit so maybe change the order if it's easy. [Update, they are hoping to fix this in 4.1 today, Apr 21st] p36 - The servicemix shell appears before startup is really complete in SMX4. it is still deploying most of the bundles in teh background. for example, the command 'jbi list' won't work when the console first appears for me, I need to wait a few seconds p36 - For ESB 4.1, command likes 'jbi list' are now 'jbi/list'. Not sure how you will handle it if it is meant to work in SMX 4.0 too p37 - I got the following error when I tried tomcat6.exe start (presumably because I didn't have catalina_home set or something similar [2009-04-21 12:22:41] [427 prunsrv.c] [error] The system cannot find the file specified. [2009-04-21 12:22:41] [1336 prunsrv.c] [error] Load configuration failed typing 'startup' did work though. This was with tomcat 6.0.0.18 (which admittedly is not completely up to date) p40 - Clicking the link to the webapp in the PDF results in the page being opened in the same web browser, thus you lose the ability to read the next instructions. You either have to go to the footnote and copy and paste or open it, copy and paste, then go back then open a new window. p42 - You refer to the web server (Tomcat or Jetty) as the client console, when I read it first I assumed you meant the web browser, not the web server. p42-p45 I didn't get a message that it was in stock in the UI even though it did appear in Servicemix. I just get told my order is confirmed, not any news about it being in stock so my output in the web console looks quite different p42 again - you talk about the FUSE ESB products and how each is used, but all I have installed is FUSE ESB. Maybe we need to make it clearer that FUSE ESB contains the other FUSE products p44 - I think the stockinventory directory location is in the directory from which you launched FUSE ESB - give you had been saying to launch it with bin\servicemix, the location for me was directly under the FUSE ESB 4 p48 - the web app container doesn't appear in the diagram. it shows the web browser talking directly to FUSE ESB. p50 - in the diagram CORBA should be all in capitals. I don't know if it's worth showing FUSE SF and FUSE MR in the diagram, maybe within the green FUSE ESB box	0	0	https://issues.jboss.org/browse/ESB-612
4916	ESB-587	Hi Paul, We just need to change the wording on the jars for servicemix 3.3.1 => 3.4 So under API Changes New jars: (Needs to change to the following) The Servicemix-jbi.jar was removed from the FUSE 3.4. The classes that were packaged in this jar were moved to a new jar called org.apache.servicemix.specs.jbi-api-1.0-1.2.0.1-fuse.jar The servicemix-services.jar was also removed. The classes in this jar were moved to a new jar called servicemix-utils-1.0.0.2-fuse.jar we need to maybe put x in for the versions here in the above like jbi-api-1.0-1.2.0.x and 1.0.0.x Next change - just a spelling mistake in ws-addressing under FUSE Services Framework Changes EnpointReferenceType => EndpointReferenceType Next changed under Product Version The camel version should be the same 1.5.0.x Are we doing a different guide for 4.0 as some more jars changed there ? We dropped servicemix-core jar of which the classes moved throughout or removed as it was a refactor/rewrite of the core and mixed into different projects such as the kernel, nmr etc Servicemix-common.jar was a new jar introduced.	0	0	https://issues.jboss.org/browse/ESB-587
4917	ESB-559	Is this fixed?	1	1	https://issues.jboss.org/browse/ESB-559
4918	ESB-526	From looking at all the plugins - yes they are all on different versions currently. We have the following which all have the same parent of maven-plugins-pom (1.0.0.0-fuse) checksum-maven-plugin => 1.0.0-fuse-SNAPSHOT depends-maven-plugin => 1.1-fuse-SNAPSHOT features-maven-plugin => 1.0.0-fuse-SNAPSHOT jbi-maven-plugin => 4.0.0.3-fuse-SNAPSHOT res-maven-plugin => 4.0.0-fuse-SNAPSHOT servicemix-build => 1.1.0-fuse-SNAPSHOT xfire-maven-plugin => 4.0.0-fuse-SNAPSHOT So I think it would be good to sync them all back up for 4.1. Do we need to update the maven-plugins-pom as well ? the jbi-maven-plugin is the only one that has dependencies on other servicemix code namely Servicemix-common (Here we only depend on 3 classes in common which could be moved to the jbi src too in order to elimate the cyclic dependency between components and the jbi plugin) Servicemix-core - it depends namely on the management classes and the JBIContainer class. The management classes now exist in the NMR also but I guess we share the plugins also with servicemix 3 so this dependency could remain for now. So to clean up some of the versioning what we could do is remove the cyclic dependency on the components and a proposal of the following versions for 4.1 bundles - individually versioned specs - individually versioned utils - 4.1.0.0-fuse Plugins - 4.1.0.0-fuse Components - 2009.01.0.0-fuse Kernel - 1.1.0.0-fuse NMR - 1.1.0.0-fuse Features - 4.1.0.0-fuse Comments ? Edell.	1	1	https://issues.jboss.org/browse/ESB-526
4919	ESB-514	Do we need the latest version of CXF\FUSE SF with the various patches to get this to work? Here's a test that is failing.	1	1	https://issues.jboss.org/browse/ESB-514

#	Requirement	Questions	Answered	Successfully Answered	URL
4920	ESB-512	I presume this isn't fixed in 3.3.0.12, since version is never likely to be released. Should this have a fix version of 3.3.1.13?	1	0	https://issues.jboss.org/browse/ESB-512
4921	ESB-473	what architecture does is AIX supported on? Are there any required OS patch levels?	1	1	https://issues.jboss.org/browse/ESB-473
4922	ESB-466	Hi Eric, Where can we find the attachment for reviewing? Regards Freeman	0	0	https://issues.jboss.org/browse/ESB-466
4923	ESB-425	Hi Freeman, Sabre made the following comment: "It seems that JLine has been separated, but Geronimo GShell is still bundled as an internal part of SMX 4.1 Kernel." Can you please clarify? Thanks, Pedro	1	0	https://issues.jboss.org/browse/ESB-425
4924	ESB-390	I've been researching this issue and believe that I can put forward the following analysis. Based upon SM-1554 there are two outstanding issues regarding this Improvement request; 1) SM-1607 : JBI container should be able to initialize all SAs first and then start them. And, 2) SM-1608 : Support inter-SA dependency declaration in order to tweak the SA startup order. SM-1607: should be possible to implement. SM-1608: The JBI spec does not specify a way to introduce dependencies between SAs, hence a solution to this issue may not be immediately forth coming. Would the implementation of SM-1607 be a suitable solution to meeting your immediate requirement for ESB-390?	1	1	https://issues.jboss.org/browse/ESB-390
4925	ESB-384	After reading through DEV-943 I see that there is no specific list of components to apply this to. I would hate to spend energy in the wrong area so it would be nice to document the specific components that are causing pain and address them directly. Would it be possible to get such a list?	1	0	https://issues.jboss.org/browse/ESB-384
4926	ESB-350	Is there any reason why this is 9MB smaller than the 3.3.1.2 repo Darren?	1	0	https://issues.jboss.org/browse/ESB-350
4927	ESB-331	On Page 15 - it should be "The following JBI Components....." and we no longer ship the jsr181 component either so we probably should add that to the list there. On page 17 - Is the java2wsdl still available in Cxf 2.1 - has this not changed to java2ws http://cxf.apache.org/21-migration-guide.html - If they do swich over to use the jars from servicemix 4 - I found that I could no longer use the cxf-codegen-plugin with java2wsdl - I had to change this to use cxf-java2ws-plugin with the goal as java2ws put an example at the end On Page 25 - it should be "In FUSE ESB 3 your...." "a level of persistence and failure tolerance out of ....." It might be useful to mention API Changes - One that I have hit with the Logisticx Demo is the MessageExchangeListener If you use the Servicemix-bean component and implement the MessageExchangeListener. This Class has now moved to a different package. It use to be in org.apache.servicemix and is now in org.apache.servicemix.jbi.listener I have got this change backported to Servicemix 3.3.1.x (3.3.1.10) but if people are on older version they will hit this. Another thing that might be worth mentioning is that you cannot put a dependency on servicemix-core and servicemix-jbi any longer as they no longer exist. The equivalent would be to put a dependency on servicemix-utils and servicemix-common The code is generally the same just moved to different jars. <plugin> <groupId>org.apache.cxf</groupId> <artifactId>cxf-java2ws-plugin</artifactId> <version>${cxf-version}</version> <dependencies> <dependency> <groupId>org.apache.cxf</groupId> <artifactId>cxf-rt-frontend-jaxws</artifactId> <version>${cxf-version}</version> </dependency> </dependencies> <executions> <execution> <id>generate-wsdl</id> <phase>process-classes</phase> <configuration> <className> com.iona.fuse.demo.logisticx.service.warehouse.Warehouse </className> <outputFile> ${project.build.directory}/generated-sources/jaxws/warehouse.wsdl </outputFile> <verbose>true</verbose> <argline>-classdir ${project.build.directory}/classes</argline> </configuration> <goals> <goal>java2ws</goal> </goals> </execution> </executions> </plugin> We mention that we don't support the wsn2005 component but we do ship that in Servicemix 4 GA - so we should probably remove that.	1	1	https://issues.jboss.org/browse/ESB-331
4928	ESB-331	On Page 15 - it should be "The following JBI Components....." and we no longer ship the jsr181 component either so we probably should add that to the list there. On page 17 - Is the java2wsdl still available in Cxf 2.1 - has this not changed to java2ws http://cxf.apache.org/21-migration-guide.html - If they do swich over to use the jars from servicemix 4 - I found that I could no longer use the cxf-codegen-plugin with java2wsdl - I had to change this to use cxf-java2ws-plugin with the goal as java2ws put an example at the end On Page 25 - it should be "In FUSE ESB 3 your...." "a level of persistence and failure tolerance out of ....." It might be useful to mention API Changes - One that I have hit with the Logisticx Demo is the MessageExchangeListener If you use the Servicemix-bean component and implement the MessageExchangeListener. This Class has now moved to a different package. It use to be in org.apache.servicemix and is now in org.apache.servicemix.jbi.listener I have got this change backported to Servicemix 3.3.1.x (3.3.1.10) but if people are on older version they will hit this. Another thing that might be worth mentioning is that you cannot put a dependency on servicemix-core and servicemix-jbi any longer as they no longer exist. The equivalent would be to put a dependency on servicemix-utils and servicemix-common The code is generally the same just moved to different jars.	1	1	https://issues.jboss.org/browse/ESB-331
4929	ESB-331	There appears to be differences in the trademark and copyright notices between the migration guide and jbi guide - should we have one to be consistent or is this because there are different dates for the these guides etc ?	1	1	https://issues.jboss.org/browse/ESB-331
4930	ESB-329	On Page 15 - it should be "The following JBI Components....." and we no longer ship the jsr181 component either so we probably should add that to the list there. On page 17 - Is the java2wsdl still available in Cxf 2.1 - has this not changed to java2ws http://cxf.apache.org/21-migration-guide.html - If they do swich over to use the jars from servicemix 4 - I found that I could no longer use the cxf-codegen-plugin with java2wsdl - I had to change this to use cxf-java2ws-plugin with the goal as java2ws put an example at the end On Page 25 - it should be "In FUSE ESB 3 your...." "a level of persistence and failure tolerance out of ....." It might be useful to mention API Changes - One that I have hit with the Logisticx Demo is the MessageExchangeListener If you use the Servicemix-bean component and implement the MessageExchangeListener. This Class has now moved to a different package. It use to be in org.apache.servicemix and is now in org.apache.servicemix.jbi.listener I have got this change backported to Servicemix 3.3.1.x (3.3.1.10) but if people are on older version they will hit this. Another thing that might be worth mentioning is that you cannot put a dependency on servicemix-core and servicemix-jbi any longer as they no longer exist. The equivalent would be to put a dependency on servicemix-utils and servicemix-common The code is generally the same just moved to different jars. <plugin> <groupId>org.apache.cxf</groupId> <artifactId>cxf-java2ws-plugin</artifactId> <version>${cxf-version}</version> <dependencies> <dependency> <groupId>org.apache.cxf</groupId> <artifactId>cxf-rt-frontend-jaxws</artifactId> <version>${cxf-version}</version> </dependency> </dependencies> <executions> <execution> <id>generate-wsdl</id> <phase>process-classes</phase> <configuration> <className> com.iona.fuse.demo.logisticx.service.warehouse.Warehouse </className> <outputFile> ${project.build.directory}/generated-sources/jaxws/warehouse.wsdl </outputFile> <verbose>true</verbose> <argline>-classdir ${project.build.directory}/classes</argline> </configuration> <goals> <goal>java2ws</goal> </goals> </execution> </executions> </plugin>	1	1	https://issues.jboss.org/browse/ESB-329

#	Requirement	Questions	Answered	Successfully Answered	URL
4931	ESB-329	**** COMMENTS For the JbI pdf **** Page 13 : The functional units are deployed into JBI components that "ARE" hosted within the JBI environement. Page 15: Service Engines - Service engines are components that provide some "OF" the logic needed to provide services, Page 16: service assembly - The root element of the JBI descriptor is "A" service-assembly service unit - The root element of the JBI descriptor is "A" service-unit Page 19: Second paragraph: spelling mistake form should be from Page 23: under Managing service units - shoutdown should be shutdown last line of the first paragraph = In the started state, the functionality defined by the service unit "is" accessible. Second para - "A" service ..... Page 27: The JBI specification defines a number of the Ant tasks that can be used "TO" manage ..... Are we supporting these Ant tasks in servicemix 4 ? Page 33 the description is in the wrong box in Required. There appears to be inconsitently with the tables in that in some we say yes/no for the required fields and others we say nothing. Page 42: Creating a JBI Maven Project - The archetype will generate the proper file structure for the projecct and a POM file containing the essential metadata "Remove for" required..... Page 50: The modules element contains one module element for each service unit - "and a module element" might just as easy to give an example <module>app1</module> Page 53: First paragraph: The contents of a service unit(S) ..... Page 55: Specifying the target components - "Remove You" Add a dependency child element..... Page 56: First paragraph: Maven "WILL" download..... You can only use one target ComponentName for a service unit so not sure that you can say it can target two components here.	1	1	https://issues.jboss.org/browse/ESB-329
4932	ESB-329	The sections Specifying a Fuse project is also listed in http://fusesource.com/issues/browse/ESB-245 Can you apply the changes for that jira to the same here as well. Changes needed for the "A service unit Project" thanks, Edell.	0	0	https://issues.jboss.org/browse/ESB-329
4933	ESB-324	Will this release include the most recent releases of camel, cxf, and activemq? camel -> 1.4.1.0 cxf -> 2.1.1.0 activemq -> 5.1.0.0	1	1	https://issues.jboss.org/browse/ESB-324
4934	ESB-295	Can this be closed?	1	1	https://issues.jboss.org/browse/ESB-295
4935	ESB-282	Can you do this upgrade - assuming that it doesn't introduce any incompatibilities?	0	0	https://issues.jboss.org/browse/ESB-282
4936	ESB-282	This has already been upgraded for the FUSE ESB 4.0 release. Can we also get this upgraded for the 3.4 release? I think it would be OK to break some backward compatibility there?	0	0	https://issues.jboss.org/browse/ESB-282
4937	ESB-279	We already have a more generic solution in place, which is using XA transactions instead of JMS transactions. It looks like the customer recently asked for some sample/demo code on using XA transactions for the JMS endpoints. Is there any reason why they can't use the solution proposed there?	1	0	https://issues.jboss.org/browse/ESB-279
4938	ESB-272	Hi Matt, There is now substantially more information available about how to use Camel in ServiceMix than there used to be--for example, see http://fusesource.com/products/enterprise-servicemix4/ . Do you think this is enough? Can we close this JIRA item now?	1	1	https://issues.jboss.org/browse/ESB-272
4939	ESB-257	Working on getting the smx4 samples running. Hit a few different snags depending on which version of smx4 is used. Get a build error with the apache smx4 trunk http://svn.apache.org/repos/asf/servicemix/smx4/features/trunk . It's unable to resolve a 5.2.0-SNAPSHOT version of activemq. A checkout of the fuse smx4 trunk builds fine, but I see errors like this: servicemix features> install nmr ERROR BundleException: Unresolved package in bundle 39: package; (&(package=javax.xml.stream.events)(version>=1.0.0)) which looks a lot like this jira: https://issues.apache.org/activemq/browse/SMX4-20 maybe the fuse version hasn't quite picked up that fix yet?	1	1	https://issues.jboss.org/browse/ESB-257
4940	ESB-257	What's the latest status - did you get past these build errors?	1	1	https://issues.jboss.org/browse/ESB-257
4941	ESB-257	Are these demos in 4.0 distribuition now?	1	1	https://issues.jboss.org/browse/ESB-257
4942	ESB-257	Can you also add these into CPI?	0	0	https://issues.jboss.org/browse/ESB-257
4943	ESB-229	Aileen, Has someone been assigned to look at this? Thanks, Eric	1	1	https://issues.jboss.org/browse/ESB-229
4944	ESB-227	Hi Torsten, Can you verify this fix ok with customer? We have a CxfBcSecurityJAASTest to illustrate it works Freeman	0	0	https://issues.jboss.org/browse/ESB-227
4945	ESB-227	Hi Torsten, What's your jdk and maven version? Did you run the command from servicemix-cxf-bc? Did you build the whole trunk before run the test? You should build trunk first mvn -Dmaven.test.skip=true -Pstep1 install mvn -Dmaven.test.skip=true -Pstep2 install then cd deployable/bindingcomponents/servicemix-cxf-bc to run the test mvn install -Dtest=CxfBcSecurityJAASTest	0	0	https://issues.jboss.org/browse/ESB-227
4946	ESB-213	Can this be closed now Darren?	1	1	https://issues.jboss.org/browse/ESB-213
4947	ESB-212	Umm... why would they be a camel-cxf endpoint? That seems strange. Why wouldn't it be a spring-dm module plugged into SMX4 as per gnodet's sample?	0	0	https://issues.jboss.org/browse/ESB-212
4948	ESB-203	Looks good Fintan. Can we add something about how you can return an exception. There is a brief section on this {quote} Fault message—a message that travels from a target endpoint back to a source endpoint for the purpose of indicating an exception or error condition (usually in response to an In message). {quote} It could do with an example to make it a bit more clear. I recently discovered you can add an exception in this manner on a exchange that is being returned: {code} org.apache.camel.wsdl_first.types.UnknownPersonFault personFault = new org.apache.camel.wsdl_first.types.UnknownPersonFault(); personFault.setPersonId(""); org.apache.camel.wsdl_first.UnknownPersonFault fault = new org.apache.camel.wsdl_first.UnknownPersonFault("Get the null value of person name", personFault); exchange.getFault().setBody(fault); {code}	1	1	https://issues.jboss.org/browse/ESB-203
4949	ESB-202	Eric, I notice you just closed this bug. Can you add a link to the documentation from here? It would be useful for anyone who comes across this through a search. Cheers, -M	1	0	https://issues.jboss.org/browse/ESB-202
4950	ESB-198	HI, I've sent you notices.txt by email. Can you add into kits?	1	1	https://issues.jboss.org/browse/ESB-198
4951	ESB-197	Hi Ulhas, have you had a chance to review this yet? When can you get comments back to Eric?	1	0	https://issues.jboss.org/browse/ESB-197
4952	ESB-197	Ulhas, Have you had a chance to look at this doc? I'd appreciate it if you could provide some feed back. Thanks, Eric	1	0	https://issues.jboss.org/browse/ESB-197
4953	ESB-194	These are the following examples I have run and they appear to be ok Bridge, Cluster Basic Wsdl-first Vfs-binding Rss-binding File-binding Cxf-wsdl-first Cxf-ws-security Problems: ========== Camel ? James is looking at this Wsn-http-binding - sent you a separate question on it Loan-broker - should I have to CTRL-C the client - ???? run: [echo] Running exsample client [java] Connecting to JMS server. [java] Sending request. [java] Response was: 6.110767707946193 from bank4 [java] Closing. [java] Java Result: 1 Ws-notification - When I attempt to run the second instance - it complains about stuff missing - need to re-look at this one - it also needs README updates. To Do: ========== Servicemix-web Loan-broker-bpel	0	0	https://issues.jboss.org/browse/ESB-194

#	Requirement	Questions	Answered	Successfully Answered	URL
4954	ESB-194	either of these. http://localhost:8192/Broker/ http://localhost:8192/CreatePullPoint/ I download the SoapUI project and launched the Test and it fails on step 4. 2) servicemix-web - Needs README updates 3) loan-broker-bpel - There is a version released now so you should no longer need to build the src but when I attempt to deploy the apache-ode-jbi-1.1.zip - it fails to deploy - complains that it can't find the deployment descriptor and the jbi.xml is within the zip. 4) loan-broker - Do I have to CTRL-C the client - if so then this demo is fine. 5) Camel - Hi, I was trying to run the Camel example that's in the FUSE ESB kit. (http://repo.open.iona.com/maven2-snapshot/com/iona/fuse/fuse-esb/3.3.1-SNAPSHOT/) I was following the instructions at http://servicemix.apache.org/camel-example.html Welcome to the ServiceMix Camel Example ========== This example shows how to use Apache Camel to deploy EIP routes You can run this example using an embedded ServiceMix install via cd camel-sa mvn jbi:servicemix => THIS THROWS THE ERROR Or if you prefer you can start a ServiceMix server (if not already started) by r unning bin/servicemix in the root dir of this ditribution. To start this sample, run: mvn install jbi:projectDeploy For more information on running this example please see: http://servicemix.apache.org/camel-example.html Edell. ---------- From: Strachan, James Sent: 21 April 2008 16:55 To: Nolan, Edell Cc: eng-fuse Subject: Re: Error running the camel example in todays 3.3.1 kit BTW which camel example are you trying? On 21 Apr 2008, at 14:49, Nolan, Edell wrote: Hi, I am trying to run the camel example in todays fuse-esb-3.3.1-SNAPSHOT kit and if I try to run the example in embedded Servicemix container mode I get the error below. It runs fine if I launch servicemix and then run mvn install jbi:projectDeploy Is this a known issue ? Thanks, Edell. <component-task-result xmlns="http://java.sun.com/xml/ns/jbi/management-message" > <component-name>servicemix-camel</component-name> <component-task-result-details> <task-result-details> <task-id>deploy</task-id> <task-result>FAILED</task-result> <message-type>ERROR</message-type> <task-status-msg> <msg-loc-info> <loc-token/> <loc-message>Could not deploy xbean service unit</loc-message> </msg-loc-info> </task-status-msg> <exception-info> <nesting-level>1</nesting-level> <msg-loc-info> <loc-token/> <loc-message>org.springframework.beans.f actory.BeanCreationException: Error creating bean with name 'camelContext:beanPo stProcessor': Cannot resolve reference to bean 'camelContext' while setting bean property 'camelContext'; nested exception is org.springframework.beans.factory. BeanCreationException: Error creating bean with name 'camelContext': Invocation of init method failed; nested exception is java.lang.NoSuchMethodError: org.apac he.camel.spring.SpringCamelContext.setName(Ljava/lang/String;)V</loc-message> <stack-trace><![CDATA[org.apache.xbean.k ernel.ServiceRegistrationException: org.springframework.beans.factory.BeanCreati onException: Error creating bean with name 'camelContext:beanPostProcessor': Can not resolve reference to bean 'camelContext' while setting bean property 'camelC ontext'; nested exception is org.springframework.beans.factory.BeanCreationExcep tion: Error creating bean with name 'camelContext': Invocation of init method fa iled; nested exception is java.lang.NoSuchMethodError: org.apache.camel.spring.S pringCamelContext.setName(Ljava/lang/String;)V at org.apache.xbean.kernel.standard.ServiceManagerRegistry.registerServi ce(ServiceManagerRegistry.java:424) at org.apache.xbean.kernel.standard.StandardKernel.registerService(Stand ardKernel.java:220) at org.apache.xbean.server.spring.loader.SpringLoader.load(SpringLoader. java:152) at org.apache.servicemix.common.xbean.AbstractXBeanDeployer.deploy(Abstr	0	0	https://issues.jboss.org/browse/ESB-194
4955	ESB-185	Is this already done?	1	1	https://issues.jboss.org/browse/ESB-185
4956	ESB-171	This was fixed in 3.3.1 not 3.3.0. Are you seeing the problem in 3.3.1?	1	1	https://issues.jboss.org/browse/ESB-171
4957	ESB-171	The JMS Spec does not allow "." in property names as outlined in sections 3.5.1 and 3.8.1.1 of the spec, so Camel and ServiceMix both check for valid Java identifier names when populating JMS properties. Current behavior ignores any property that has an invalid name and that is what is being exhibited here. Is there a particular problem with setting a non-default sequenceNumberKey on the default comparator? I would usually recommend doing this anyway, so I am wondering if there is a particular reason that this isn't acceptable?	0	0	https://issues.jboss.org/browse/ESB-171
4958	ESB-170	FWIW, we can make the process homogeneous for servicemix-bean, servicemix-jsr181 and servicemix-cxf-se, but other components do not use POJOs in the same way, so it would not be possible to have the same thing. However, https://issues.apache.org/activemq/browse/SM-1289 is related and applies to all servicemix components. It is currently solved in apache 3.3-SNAPSHOT only.	0	0	https://issues.jboss.org/browse/ESB-170
4959	ESB-155	Attached is our CustomTraceComponent which does what we want. Could you please review it and consider adding it to Fuse distribution? Technical comments: This class overrides TraceComponent, adding processing of properties. These changes should probably be moved to the TraceComponent class itself. Maybe a property specifying whether to print the properties could be also introduced. Additional questions: - is it enough to only convert values of class Source to String, or should Objects of other types be also transformed somehow (see attached source file)? - why is the TraceComponent class located in servicemix-core, not in servicemix-components? The package prefix is common for classes defined in servicemix-components and the TraceComponent - org.apache.servicemix.components.	1	0	https://issues.jboss.org/browse/ESB-155
4960	ESB-154	Adrian, can you review and apply patch?	0	0	https://issues.jboss.org/browse/ESB-154
4961	ESB-149	Adrian - can you review and apply patch?	1	1	https://issues.jboss.org/browse/ESB-149
4962	ESB-132	But those deploying the cxf-se will want to enable this on the xbean, not in any cxf.xml? How would you configure the cxf.xml picked up by the service engine?	1	1	https://issues.jboss.org/browse/ESB-132
4963	ESB-131	Can this one be closed?	1	0	https://issues.jboss.org/browse/ESB-131
4964	ESB-124	Hi, when is this likely to be complete? We are releasing SF 2.1 end May and need to have CPI in place to test demos. thanks, Aileen	1	1	https://issues.jboss.org/browse/ESB-124
4965	ESB-124	Can you prioritise CXF demos? we releasing 2.1 next week and it would be good not to have to test all the demos by hand!	1	1	https://issues.jboss.org/browse/ESB-124
4966	ESB-120	Could you provide an update on your status please Greg? Two weeks to implement may be okay, but you also had one week off :-) Do you have a target date for completion? What tests have been implemented to date?	0	0	https://issues.jboss.org/browse/ESB-120
4967	ESB-120	What's the latest status on these?	0	0	https://issues.jboss.org/browse/ESB-120
4968	ESB-120	Can we have someone look into the integration demos. I am getting failures running the demos by hand and as a result, cannot confirm that the cpi setup for them is completely correct. Thanks.	1	1	https://issues.jboss.org/browse/ESB-120
4969	ESB-120	What failures are you seeing?	1	1	https://issues.jboss.org/browse/ESB-120
4970	ESB-104	Eric, I notice you just closed this. Can you add a link to the updated docs from here? It would be useful for anyone who comes across this through a search. Thanks, - M	1	0	https://issues.jboss.org/browse/ESB-104
4971	ESB-96	Do you feel that this could be a little easier to do using the CXF BC in the 3.3.1 branch Bruce? I would imagine so.	1	0	https://issues.jboss.org/browse/ESB-96
4972	ESB-71	Eric, can you please re-save the PDF with editing features enabled?	1	1	https://issues.jboss.org/browse/ESB-71
4973	ESB-48	Oisin, can you please address this issue? We will be performing the next incremental release of the Fuse components next week and it would be nice to have the versioned Eclipse update sites available for the customer that requested access to older versions.	1	1	https://issues.jboss.org/browse/ESB-48
4974	ESB-27	I did quick scan of the code, and I don't see an HttpClient in the consumer endpoint. Do you mean the provider endpoint maybe?	1	1	https://issues.jboss.org/browse/ESB-27

#	Requirement	Questions	Answered	Successfully Answered	URL
4975	ESB-27	Hi Bruce, Just one last question, I'm not really sure, but is the soTimeout different from the exchangeTimeout in the HttpClient? Just making sure that the right property is being configured. :)	1	1	https://issues.jboss.org/browse/ESB-27
4976	HQ-386	What is a portss?	0	0	https://issues.jboss.org/browse/HQ-386
4977	HQ-360	Hi Jack, Can you ask her to clean data folder from HQ agent folder? If that doesn't solve the problem ask her to reinstall agent again. I am working on Linux so never tried it on Windows. If the problem doesn't get resolved by Monday I will give it a try on Windows. It is worth asking for HQ agent log to find out if anything visibly wrong could be found. Regards, Ulhas Bhole	1	1	https://issues.jboss.org/browse/HQ-360
4978	HQ-360	Hi Jack, Can you ask Debbi to delete data folder and call {code} hq-agent.bat install {code} and then on subsequent attempt to {code} hq-agent.bat start {code} will run initial setup and ask you for server ip and other details. Something link below: {quote} [Running agent setup] What is the HQ server IP address: mantra Should Agent communications to HQ always be secure [default=no]: What is the HQ server port [default=7080]: - Testing insecure connection ... Failure What is the HQ server IP address: 10.5.1.10 Should Agent communications to HQ always be secure [default=no]: What is the HQ server port [default=7080]: - Testing insecure connection ... Success What is your HQ login [default=hqadmin]: What is your HQ password: What IP should HQ use to contact the agent [default=10.5.1.26]: What port should HQ use to contact the agent [default=2144]: - Received temporary auth token from agent - Registering agent with HQ - HQ gave us the following agent token 1266835166690-1986179072184931521-7864810437137618626 - Informing agent of new HQ server - Validating - Successfully setup agent {quote}	1	1	https://issues.jboss.org/browse/HQ-360
4979	HQ-360	Hi Jack, Did you ask them to remove camel-plugin.xml from both server and agent? Are they seeing Camel 2.x being discovered? Logs from agent and server would help in this case to shed more light. I will see if I can run here again myself on windows. Regards, Ulhas Bhole	1	1	https://issues.jboss.org/browse/HQ-360
4980	HQ-320	Name of the Server currently observed is 'Sonic(1099/karaf-root)' by default for ESB4.1-M2. Is this all expected?	1	0	https://issues.jboss.org/browse/HQ-320
4981	HQ-319	Please see Attached Screenshot. Is this all that is expected?	0	0	https://issues.jboss.org/browse/HQ-319
4982	HQ-305	Based on feedback and ideas to date: * This screen will be attached as a sub-tab of the "Views" tab of HQ server resources of the type "KARAF" ("KARAF" is subject to change). * The sub-tab title will be "Features" * Tab will display a table with a row of action buttons above ** The Karaf admin MBean ObjectName will be retrieved from the config response using the key "karafFeaturesName" ** Values for table will be obtained by getting the "Features" JMX attribute * Table will have following columns (obtained using CompositData key from tabular data): ** Name ("Name") ** Version ("Version") ** Installed ("Installed") ** Bundles ("Bundles") .. this is an array of Strings so looking for ways to render .. maybe we show as a comma delimited list and if selected it appears as a popup table of all values (the delimited string could end in "..." if everything cannot be displayed within the cell) ** Dependencies ("Dependencies") .. this is a table (JMX TabularData) of feature "Name" and "Version" so again looking for ways to render .. maybe we show as a comma delimited list and if selected it appears as a popup table of all name/version values (the delimited string could end in "..." if everything cannot be displayed within the cell) ** Configurations ("Configurations").. this is a table (JMX TabularData) of configuration "PID" and "Elements" so again looking for ways to render .. this is further complicated by "Elements" being a table (JMX TabularData) of "Key" and "Value" pairs * Action buttons: ** "Install" *** applicable to a single selected instance [can't do multiples as the JMX API does not support it .. but should we support impression of multiples by making multiples calls to the backend?], calls the JMX operation "installFeature" with 2 params: **** param1 = the name of the selected feature **** param2 = the version of the selected feature ** "Uninstall" *** applicable to a single selected instance [can't do multiples as the JMX API does not support it .. but should we support impression of multiples by making multiples calls to the backend?], calls the JMX operation "uninstallFeature" with 2 params: **** param1 = the name of the selected feature **** param2 = the version of the selected feature	1	1	https://issues.jboss.org/browse/HQ-305
4983	HQ-304	Based on feedback and ideas to date: * This screen will be attached as a sub-tab of the "Views" tab of HQ server resources of the type "KARAF" ("KARAF" is subject to change). * The sub-tab title will be "Repositories" * Tab will display a table with a row of action buttons above ** The Karaf admin MBean ObjectName will be retrieved from the config response using the key "karafFeaturesName" ** Values for table will be obtained by getting the "Repositories" JMX attribute * Table will have following columns (obtained using CompositData key from tabular data): ** Name ("Name") ** URI ("Uri") ** Dependent repositories ("Repositories") .. this is an array of Strings so looking for ways to render .. maybe we show as a comma delimited list and if selected it appears as a popup table of all values (the delimited string could end in "..." if everything cannot be displayed within the cell) ** Included Features ("Features") .. this is a table (JMX TabularData) of feature "Name" and "Version" so again looking for ways to render .. maybe we show as a comma delimited list and if selected it appears as a popup table of all name/version values (the delimited string could end in "..." if everything cannot be displayed within the cell) * Action buttons: ** "Add" *** call the JMX operation "addRepository"; but first need to prompt for the following params: **** param1 = the URI of the repository to be added ** "Remove" *** applicable to a single selected instance [can't do multiples as the JMX API does not support it .. but should we support impression of multiples by making multiples calls to the backend?], calls the JMX operation "removeRepository" with 1 params: **** params1 = the URI of the selected repository to be removed	1	0	https://issues.jboss.org/browse/HQ-304
4984	HQ-281	Swathy...David says in his comment above the "normal" values in their install are ~24 hours, yet you say it's (15 mins currently) in your comment...can you clarify what the value is out-of-the-box? Thanks Amanda	1	1	https://issues.jboss.org/browse/HQ-281
4985	HQ-267	As well as the two buttons there needs to be a Monitoring Metrics state This should be either on or off - (if on would it be possible to have a link to the metrics tab for the resource ? - this is a nice to have) I have updated the graphic attached with the on/off state	0	0	https://issues.jboss.org/browse/HQ-267
4986	HQ-230	Implemented as in http://fusesource.com/wiki/display/FUSEINT/HQU_Plugins with some minor changes * Name is "Active Exchanges" instread of "Exchanges" * Action is "Cancel Exchanges" instead of "Cancel" * Status column is not there as there is no status returned from the API (only active exchanges are returned)	1	1	https://issues.jboss.org/browse/HQ-230
4987	HQ-222	Amanda, Is this still valid? Should we close it?	1	0	https://issues.jboss.org/browse/HQ-222
4988	HQ-196	We are able to see the skelton of the Routes tab but not able to get any data into this page. Refer - HQ-348. Amanda, can we close this leaving the issue to be tracked in HQ-348 ?	1	0	https://issues.jboss.org/browse/HQ-196
4989	HQ-163	Can you leave this one open until Murali investigates whether there is a plugin approach available to us that might allow this? Tx Amanda	1	0	https://issues.jboss.org/browse/HQ-163
4990	HQ-152	Murali We actually discussed this item on yesterdays call if you remember. Can you look at how you might include this in the Fuse-operator dashboard? Thanks Amanda	1	1	https://issues.jboss.org/browse/HQ-152
4991	HQ-147	Can this issue be marked Resolved?	1	0	https://issues.jboss.org/browse/HQ-147

#	Requirement	Questions	Answered	Successfully Answered	URL
4992	HQ-140	Agreed - there is no need for a Best Practice Guide for the Sonic 8 release. Should we sit down and go through the doc plan with the updated requirements? John From: David Porter Sent: 06 November 2009 11:16 To: Amanda Hall; John McHugh Cc: Orla Fitzgerald Subject: RE: Feature Freeze date looming - October 30th Amanda, John Since HQ is no longer "our" product, and since there will no longer be a separate HQ install guide. I don't think this is a requirement any more. David	1	0	https://issues.jboss.org/browse/HQ-140
4993	HQ-139	Can this be maked Resolved?	1	1	https://issues.jboss.org/browse/HQ-139
4994	HQ-101	As suggested I have tested if the help links work. 1.Selected Browse from the Resources masthead and clicked on the help link . It takes you to the hyperic help on Browse Resources. 2. Selected a resource (Fuse Message Broker) and set an Alert Definition and clicked on the help link - it takes you to the help link of hyperic on how to set alerts. 3. Created a new group from the resources screen (Tools Menu - New Group) and clicked on the help link. It gives you help on how to create group from Hyperic. Should this issue then be considered as closed?	1	1	https://issues.jboss.org/browse/HQ-101
4995	HQ-85	I just tried this fix on Windows. I copied mysql-connector-java-5.0.4-bin.jar to \fuse-hq-installer\installer-3.2.5.2-fuse\bin and I still get the same error. Am I doing something wrong?	0	0	https://issues.jboss.org/browse/HQ-85
4996	HQ-57	With the fix for HQ-332 the CXF services are now discovered using only their Mbeans. So all that is needed is the JMX url needs to be added to the etc/agent.properties fusesource.jmx.urls property .	0	0	https://issues.jboss.org/browse/HQ-57
4997	HQ-55	Is this issue still in progress or was it completed? Tx Amanda	1	0	https://issues.jboss.org/browse/HQ-55
4998	HQ-46	Gerard I'm not sure what you're asking for here. Does this write up from Heather help at all? http://fusesource.com/wiki/display/HDOX/Monitoring+multiple+CXF+servers+in+HQ Or are you looking for something different? Thanks	1	0	https://issues.jboss.org/browse/HQ-46
4999	HQ-28	The link used to monitor progress is a dead link..... Unclear on this process for migrating.... Are you going to convert the wiki mark-up to PDF, Word, and text? Or is this conversion from the XML? Why Word and text? IONA documentation is provided in HTML and PDF. All of the numbered statements save (2) talk about wiki marked up content. How does this relate to moving the content into XML? (3) -- the second one -- talks about articles still being in wiki format but rearranging the "include" statements. Do you mean the xincludes in docbook or simply links in the wiki?	1	0	https://issues.jboss.org/browse/HQ-28
5000	HQ-14	Hi, Eric- Where is the FUSE HQ doc source. I don't see it in http://svn.boston.amer.iona.com/repos/docs_inferno/trunk, which is where the other FUSE doc is. Are you working from home today? Howard	0	0	https://issues.jboss.org/browse/HQ-14
5001	OSFUSE-606	Susan, where is the cacert installed? Seems it's being picked up by oc since it's installed in the OS somewhere. Which OS was this on? Have you guys been able to re-create?	0	0	https://issues.jboss.org/browse/OSFUSE-606
5002	OSFUSE-598	Any considerations to use [agent-bond|https://github.com/fabric8io/agent-bond] to combine Jolokia and Prometheus agent (as we already do for the fabric8 upstream images) ? What's about security ? Don't think that Prometheus has SSL client cert support out of the box.	1	1	https://issues.jboss.org/browse/OSFUSE-598
5003	OSFUSE-598	Last I peeked at agent-bond, it seems to be just a packaging exercise so that user would only need to use 1 agent. Is that still correct? If that's the case I think it would be simpler to document and maintain if we just enable both the jolokia and prometheus agents separately on the java CLI. Part of this issue would should be to identify what are the options for securing prometheus.	1	0	https://issues.jboss.org/browse/OSFUSE-598
5004	OSFUSE-595	[~nibbio84] you mentioned on FIS call that there was a never release of f-m-p that should be used for Fuse 7. Can you follow up with [~tcunning] to get this into the Fuse 7 pipeline?	1	1	https://issues.jboss.org/browse/OSFUSE-595
5005	OSFUSE-595	What version of kubernetes client do we use for Fuse 7 0 is it 2.4 ?	1	1	https://issues.jboss.org/browse/OSFUSE-595
5006	OSFUSE-577	[~rhuss], can you comment on the impact on an update to jolokia 1.3.6 in the first FIS rollup patch after 2.0 release?	1	1	https://issues.jboss.org/browse/OSFUSE-577
5007	OSFUSE-577	[~djcoleman] Is Jolokia built from source? How much effort is it to upgrade for Rolllup?	1	1	https://issues.jboss.org/browse/OSFUSE-577
5008	OSFUSE-577	I've updated the Docker files to include the latest version of the jolokia-jvm agent jar already, but it looks like camel references jolokia as well, platforms/commands/commands-jolokia/pom.xml for example. I take it these references will need updating to the latest version as well, is that correct?	0	0	https://issues.jboss.org/browse/OSFUSE-577
5009	OSFUSE-557	Do you only refer to OSFUSE-492 or are there also other spots where we can improve error messaging ? For OSFUSE-492 I think I have good news, as there is now a 3.6.0 release so that we can f-m-p and kubernetes-client. The build error message problem should be fixed then.	0	0	https://issues.jboss.org/browse/OSFUSE-557

#	Requirement	Questions	Answered	Successfully Answered	URL
5010	OSFUSE-557	project camel-ose-springboot-xml: Execution default of goal io.fabric8:fabric8-maven-plugin:3.1.80.redhat-000013:build failed: Can't instantiate binary build, due to error reading/writing stream. Can be caused if output is stream closed by the server. Connection reset by peer: socket write error -> [Help 1] org.apache.maven.lifecycle.LifecycleExecutionException: Failed to execute goal io.fabric8:fabric8-maven-plugin:3.1.80.redhat-000013:build (default) on project camel-ose-springboot-xml: Execution default of goal io.fabric8:fabric8-maven-plugin:3.1.80.redhat-000013:build failed: Can't instantiate binary build, due to error reading/writing stream. Can be caused if output is stream closed by the server. at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:212) at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:153) at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:145) at org.apache.maven.lifecycle.internal.LifecycleModuleBuilder.buildProject(LifecycleModuleBuilder.java:116) at org.apache.maven.lifecycle.internal.LifecycleModuleBuilder.buildProject(LifecycleModuleBuilder.java:80) at org.apache.maven.lifecycle.internal.builder.singlethreaded.SingleThreadedBuilder.build(SingleThreadedBuilder.java:51) at org.apache.maven.lifecycle.internal.LifecycleStarter.execute(LifecycleStarter.java:128) at org.apache.maven.DefaultMaven.doExecute(DefaultMaven.java:307) at org.apache.maven.DefaultMaven.doExecute(DefaultMaven.java:193) at org.apache.maven.DefaultMaven.execute(DefaultMaven.java:106) at org.apache.maven.cli.MavenCli.execute(MavenCli.java:863) at org.apache.maven.cli.MavenCli.doMain(MavenCli.java:288) at org.apache.maven.cli.MavenCli.main(MavenCli.java:199) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.codehaus.plexus.classworlds.launcher.Launcher.launchEnhanced(Launcher.java:289) at org.codehaus.plexus.classworlds.launcher.Launcher.launch(Launcher.java:229) at org.codehaus.plexus.classworlds.launcher.Launcher.mainWithExitCode(Launcher.java:415) at org.codehaus.plexus.classworlds.launcher.Launcher.main(Launcher.java:356) Caused by: org.apache.maven.plugin.PluginExecutionException: Execution default of goal io.fabric8:fabric8-maven-plugin:3.1.80.redhat-000013:build failed: Can't instantiate binary build, due to error reading/writing stream. Can be caused if output is stream closed by the server. at org.apache.maven.plugin.DefaultBuildPluginManager.executeMojo(DefaultBuildPluginManager.java:145) at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:207) ... 20 more Caused by: io.fabric8.kubernetes.client.KubernetesClientException: Can't instantiate binary build, due to error reading/writing stream. Can be caused if output is stream closed by the server. at io.fabric8.kubernetes.client.KubernetesClientException.launderThrowable(KubernetesClientException.java:61) at io.fabric8.openshift.client.dsl.internal.BuildConfigOperationsImpl$1.writeTo(BuildConfigOperationsImpl.java:233) at okhttp3.internal.http.CallServerInterceptor.intercept(CallServerInterceptor.java:47) at	0	0	https://issues.jboss.org/browse/OSFUSE-557
5011	OSFUSE-557	Any update on this? Are there any improvements to error reporting that we could include in next patch?	0	0	https://issues.jboss.org/browse/OSFUSE-557
5012	OSFUSE-542	Jan, Karaf framework is not baked into the builder image {{jboss-fuse-6/fis-karaf-openshift}}. So it's not exactly the same as Java version environment variable. So technically, the same builder image could be used with a later Karaf framework version. Or, a user project could be misconfigured to use an older version of Karaf framework as well. In light of this fact, does it make sense to still add this variable? We could bake it in as part of the build pipeline, but I'm wondering what purpose would that serve exactly if the version can't be guaranteed?	1	0	https://issues.jboss.org/browse/OSFUSE-542
5013	OSFUSE-542	Perhaps it would make sense to include a "released with" version number for JBoss Fuse/Karaf framework. So users know which specific Karaf build it was released with, i.e. the preferred version to use with that image. Or, maybe we should instead include the FIS Karaf BOM version that was built with the image, which is actually meant to be used in user poms to import JBoss Fuse/Karaf dependencies. Does that make more sense?	1	0	https://issues.jboss.org/browse/OSFUSE-542
5014	OSFUSE-514	At this stage we are approaching code freeze and want to minimise change. Are there specific critical fixes you recommend we take in?	1	1	https://issues.jboss.org/browse/OSFUSE-514
5015	OSFUSE-514	[~iocanel] end of sprint is tomorrow, would be nice if these could be merged in before then. Could close this this out once that's done?	1	1	https://issues.jboss.org/browse/OSFUSE-514
5016	OSFUSE-510	[~roddiekieley] Does David's suggestion work for you? Can I close this?	1	1	https://issues.jboss.org/browse/OSFUSE-510
5017	OSFUSE-507	Certainly, but which mechanism will be used to update the Fuse version? * Will we release a new FIS Karaf BOM for every Fuse patch/rollup? * or should the customer instead prepend the patched Fuse BOM to the dependency management?	0	0	https://issues.jboss.org/browse/OSFUSE-507
5018	OSFUSE-504	This breaks Salesforce apps for FIS and iPaaS. Can we bump the priority? Please. :-)	1	1	https://issues.jboss.org/browse/OSFUSE-504
5019	OSFUSE-504	[~kcbabo], can you point me to your project source? I just checked the camel-2.18.x.redhat [branch|https://github.com/jboss-fuse/camel/tree/camel-2.18.x.redhat/components/camel-salesforce], and I don't see a JodaTime dependency in the component or the plugin.	0	0	https://issues.jboss.org/browse/OSFUSE-504
5020	OSFUSE-492	So, I added a more detailed message in the client: https://github.com/fabric8io/kubernetes-client/pull/607. In the particular case it will display something like: "Can't instantiate binary build, due to error reading stream." This is as descriptive as it can get from the client side, as in the particular case all the client see is that it can't read the tarball.	1	1	https://issues.jboss.org/browse/OSFUSE-492
5021	OSFUSE-492	[~chirino] as we now have a fix, how would be the procedure to get this into the v3.1.80-redhat branch ? Is the kubernetes-client also productized ? Or should the v3.1.80-redhat branch simply include the lates kubernetes-client as soon as the fix is included within a release ? To be more specific the kubernetes-client is a dependency of fabric8, which the productised f-m-p v3.1.80-redhat depends on as version 2.2.170 (current version is 2.2.192). So the question is more whether we should productis: * fabric8 * kubernetes-client * kubernetes-model (a dep from kubernetes-client) or depend on the latest version with this fix ?	1	1	https://issues.jboss.org/browse/OSFUSE-492
5022	OSFUSE-492	[~chirino] so what are the next steps now ? You would need that fix on branch v1.4.14.redhat now. [~iocanel] could you add an extra PR for cherrypicking your fix to this branch, too ? If this fix than is merged, does it then automatically propagate to fabric8v2 (v2.2.170.redhat) and then to fabric8-maven-plugin (v3.1.80-redhat) ? BTW, is the difference between ".redhat" and "-redhat" as suffix intentional ?	1	1	https://issues.jboss.org/browse/OSFUSE-492
5023	OSFUSE-492	Do we not know the Build object that was created? (its status and/or logs would be helpful to show to users...)	0	0	https://issues.jboss.org/browse/OSFUSE-492

#	Requirement	Questions	Answered	Successfully Answered	URL
5024	OSFUSE-492	Now that we found out why the FileInputStream was closed (it was because of a RetryInterceptor which then caused to read from an already closed stream a second time in the first retry), we can tackle the real problem ;-) With the latest update to the kubernetes-client the error message looks now like {code} [ERROR] Failed to execute goal io.fabric8:fabric8-maven-plugin:3.1.80.redhat-SNAPSHOT:build (default-cli) on project spring-boot-webmvc: Execution default-cli of goal io.fabric8:fabric8-maven-plugin:3.1.80.redhat-SNAPSHOT:build failed: Can't instantiate binary build, due to error reading/writing stream. Can be caused if input stream points to a missing resource or output stream closed by the server. Broken pipe -> [Help 1] {code} [~iocanel] I would maybe remove the part on the "input stream pointing to a missing resource" as this can be safely detected by a FileNotFoundException which already includes the warning. Since the server still sends a response even when he closes its own inputstream (and do not accept anything anymore on the HTTP request), he still sends a response on the other stream (his outputstream). Thats what the CLI shows. I wonder how the kubernetes-client could capture this ? Maybe simply ignoring a "broken pipe" exception and continue, then failing because of the error code in the HttpResponse ?	1	1	https://issues.jboss.org/browse/OSFUSE-492
5025	OSFUSE-492	[~chirino] [~aileenc] [~iocanel] How should we proceed here ? As said, we would need a new kubernetes-client with support for OkHttp 1.6.0-SNAPSHOT (current used version: 1.4.2). Unfortunately with 1.5.0 OkHttp totally changed the WebSocket handling so its a bit of work to update kubernetes-client for that (and that maybe for both branches ?). IMO there are two solutions: * We live with the somewhat unspecific error message in FIS 2.0 (which is still better than the original one). For FIS 2.1 (when OkHttp 1.6 hopefully is released) we then update the kubernetes-client for this more fine grained error message. * We "backport" the specific "Expect: 100-continue" handling to kubernetes-client directly. This is quite some effort and I'm not 100% sure whether its possible at all. Don't know whether we make it until Friday. My suggestion would be option 1, to delay this specific handling. wdyt ?	1	1	https://issues.jboss.org/browse/OSFUSE-492
5026	OSFUSE-492	"This late in the game it seems like we need to live with "the somewhat unspecific error message". Let put upgrading OkHttp as a new issue for Fuse 7." Can we reconsider it to include it perhaps not in FIS 2.0 but in FIS 2.1 or 2.0.1? Today [~tsedmik] is trying to use FIS, he has the same kind of error reported initially, this time it is not related to the docker insecure registry and we have no clue on what is happening.	0	0	https://issues.jboss.org/browse/OSFUSE-492
5027	OSFUSE-491	Is there anything more to be done for this issue?	0	0	https://issues.jboss.org/browse/OSFUSE-491
5028	OSFUSE-475	Is there any update on this? FIS 2.0 code freeze is in 1.5 weeks so we may need to defer.	1	1	https://issues.jboss.org/browse/OSFUSE-475
5029	OSFUSE-424	Or even better create a new Camel HA issue umbrella under which we can move ENTESB-6320 and this one ?	1	1	https://issues.jboss.org/browse/OSFUSE-424
5030	OSFUSE-418	Do we want to have these images preconfigured in the fabric8-maven-plugin, or will we just say in the docs to use -Dfabric8.generator.from= ?	1	0	https://issues.jboss.org/browse/OSFUSE-418
5031	OSFUSE-380	So I'll answer question with a question. :-) What's the estimated effort to do an offline Maven repo for FIS 2.0? AFAIK, the other xPaaS images provide one, but I think those are mostly (completely?) inherited from the corresponding standalone product.	1	1	https://issues.jboss.org/browse/OSFUSE-380
5032	OSFUSE-380	Can I close this now?	0	0	https://issues.jboss.org/browse/OSFUSE-380
5033	OSFUSE-371	Which Jenkins env do GSS use to build patches?	1	1	https://issues.jboss.org/browse/OSFUSE-371
5034	OSFUSE-371	[~janstey]Will you be able to help us set this up?	1	1	https://issues.jboss.org/browse/OSFUSE-371
5035	OSFUSE-359	[~djcoleman] Is this still an issue? Should I schedule for 7.0?	1	0	https://issues.jboss.org/browse/OSFUSE-359
5036	OSFUSE-358	it means extra effort to be portable across OS. Are we running all tests on Windows too?	0	0	https://issues.jboss.org/browse/OSFUSE-358
5037	OSFUSE-354	Is the question whether we support a plain CXF app in FIS 2.0 that doesn't use Camel? We definitely support use of Camel + CXF in a Spring Boot FIS app.	0	0	https://issues.jboss.org/browse/OSFUSE-354
5038	OSFUSE-342	I talked to Diogenes and he agreed this is something we should focus on from a top-level xPaaS standpoint. [~rettori] Where would you like to see this tracked?	1	1	https://issues.jboss.org/browse/OSFUSE-342
5039	OSFUSE-339	The question is, where would the karaf micro-container come from? If it is being built using the karaf maven plugin, it would be a lot easier to just add f-m-p to that project. Since Spring Boot is a third party framework, it's possible that customers could want to deploy stand alone jars from non-Fuse projects. Can you describe a reasoning for a Fuse customer using Karaf, just wanting to deploy a pre-built assembly? One concern is that users might misuse this ability to deploy monolithic bloated JBoss Fuse 6.x/Apache Karaf containers with a Broker, etc. in a docker image and shoot themselves in the foot.	1	1	https://issues.jboss.org/browse/OSFUSE-339
5040	OSFUSE-336	Should the f-m-p be able to get the current namespace from the oc configuration? (so that one does not have to explicitly configure users' namespace? (maybe we should add some new property to f-m-p to represent "current namespace"?	1	1	https://issues.jboss.org/browse/OSFUSE-336
5041	OSFUSE-334	Guillaume, would this be possible?	0	0	https://issues.jboss.org/browse/OSFUSE-334
5042	OSFUSE-334	Can we drop it as part of the pipeline build? Is it only brought in when the assemblies are created? Should we just skip building the Karaf assemblies?	0	0	https://issues.jboss.org/browse/OSFUSE-334
5043	OSFUSE-334	Is this stil an issue or has Gemini been removed on product branch?	1	1	https://issues.jboss.org/browse/OSFUSE-334
5044	OSFUSE-334	Gregor has just committed something to the Apache Karaf master branch: https://github.com/apache/karaf/commit/560147a78c1a3ac3a4f0cd0bd7c59f6429b511b8 Which karaf branch should that be backported to ?	1	1	https://issues.jboss.org/browse/OSFUSE-334
5045	OSFUSE-334	Can I mark this as Resolved now?	0	0	https://issues.jboss.org/browse/OSFUSE-334
5046	OSFUSE-332	How often do we update Jolokia? In the past we've done one off builds but if we update it frequently then it makes more sense to have it in pipeline.	1	1	https://issues.jboss.org/browse/OSFUSE-332
5047	OSFUSE-332	I was going by what Roland had done for the Java s2i image. Do you need me to change it to 1.3.5 released today?	0	0	https://issues.jboss.org/browse/OSFUSE-332
5048	OSFUSE-332	We are currently using jolokia-jvm-1.3.2.redhat-1-agent.jar in all images AFAIK , do we need to productise a newer 1.3.x one for anything?	1	1	https://issues.jboss.org/browse/OSFUSE-332
5049	OSFUSE-332	[~maschmid] correctly pointed out that the current version of jolokia used across MW products is 1.3.2.redhat-1. Is there a need for 1.3.6? I'd rather we standardised on one version please!	1	1	https://issues.jboss.org/browse/OSFUSE-332
5050	OSFUSE-332	Is this related to https://issues.jboss.org/browse/ENTESB-6052 ? [~sonicaaaa] had requested a version matching upstream 1.3.5. I don't see this as a problem but would this suite our needs for FIS 2.0? 6.3.0 was released with 1.3.2.redhat-2 (which is 1.3.2 built with an extra profile to produce sources jar, which was originally omitted)	1	1	https://issues.jboss.org/browse/OSFUSE-332
5051	OSFUSE-332	The S2I images now include the latest Jolokia Red Hat version built by the release pipeline. Is there any additional work required for this issue? If not, I'll mark it as resolved.	0	0	https://issues.jboss.org/browse/OSFUSE-332
5052	OSFUSE-320	Which version of Undertow are we planning to use?	0	0	https://issues.jboss.org/browse/OSFUSE-320

#	Requirement	Questions	Answered	Successfully Answered	URL
5053	OSFUSE-320	How are we going to override the version? Will we have a productized version of spring-boot-starter-undertow (and/or spring-boot-starter-web)?	0	0	https://issues.jboss.org/browse/OSFUSE-320
5054	OSFUSE-320	Can this be closed? I think we decided to stick with Tomcat.	1	1	https://issues.jboss.org/browse/OSFUSE-320
5055	OSFUSE-319	Note also https://issues.jboss.org/browse/OSFUSE-323 , we probably need to consider * image change triggers that already contain an explicit namespace in the image stream tag reference, to reference some common images (e.g. images installed in the "openshift" namespace, such as openshift/jboss-amq-6.2:1.3) * image streams in the user's namespace (different than the test one) (Should it always be a single namespace, or can there be more?) * image streams populated by the test itself (if we want to support such a thing) The fundamental question is: Should fabric8 tooling expect all the user's builds (image streams) are in the same namespace? * If we do, we can have a single property for f8-arquillian to import the imagestreamtags from. (note that users can still do "oc tag" on their own to import imagestreams from other namespaces into their namespace) * If we don't, we need users to provide a mapping between all artifacts and the namespaces in which they are being built, so that f8-arquillian can do proper "oc tag"	1	1	https://issues.jboss.org/browse/OSFUSE-319
5056	OSFUSE-311	[~lfuka], I and [~lburgazzoli] recently pushed some changes to the spring-cloud-kubernetes project and in the meantime we added much more documentation about the supported features in https://github.com/fabric8io/spring-cloud-kubernetes/blob/master/readme.md#kubernetes-propertysource. Do you think we should add more information (or doing something different)?	1	1	https://issues.jboss.org/browse/OSFUSE-311
5057	OSFUSE-310	Converted to enhancement since this is not really a bug. Users can manually create routes if they want to. Not all services have to be exposed to external users. But I do agree, it would be nice to exposing services to external users by default. This is actually slightly complex as we have to pick from creating a Route|LoadBalancer|Ingress|NodePort. The type you pick is based on the environment. The best way to do this is to annotate the service with expose and then run the exposecontroller pod. But this means we need to productize this go binary/Docker image: https://github.com/fabric8io/exposecontroller [~aileenc], [~kcbabo] What do you guys think? Should we add this to FIS?	1	1	https://issues.jboss.org/browse/OSFUSE-310
5058	OSFUSE-295	When camel-spring-boot-starter is used, should camel-core still be explicitly depended on? (amq and teiid quickstarts still contains explicit camel-core dependency even when using the starter, IIUC, camel-core dependency is redundant then...)	1	1	https://issues.jboss.org/browse/OSFUSE-295
5059	OSFUSE-290	Which part - adding user-supplied libs that are not Maven dependencies or specifically the SAP jars. If the latter, do we know where that filename check is that causes it to puke when adding the jar in the normal Maven way?	1	0	https://issues.jboss.org/browse/OSFUSE-290
5060	OSFUSE-289	Seems that if you change order of listed features, you can work around this. Can we make this work order, independent?	0	0	https://issues.jboss.org/browse/OSFUSE-289
5061	OSFUSE-277	We are in the process of migrating the old FIS 1.0 test suite to FIS 2.0 (spring boot conversion and updating the quickstarts). We can try to integrate them into the pipeline after we are done with the conversion and have them stabilised. Note however, that since our tests are mostly end-to-end tests, they require a properly configured OpenShift Enterprise environment. The best way of integrating our test suite with the build pipeline would probably be using the CI bus, so the tests would still run in our environment, but would be automatically launched by a CI bus event. (see https://mojo.redhat.com/docs/DOC-986839#jive_content_id_CI_Message_Hub , https://mojo.redhat.com/docs/DOC-948824 , https://mojo.redhat.com/docs/DOC-1023866)	1	1	https://issues.jboss.org/browse/OSFUSE-277
5062	OSFUSE-277	Marek, could you point us to git repo with a pipeline that we can use?	1	1	https://issues.jboss.org/browse/OSFUSE-277
5063	OSFUSE-276	What is in fabric8-annotations and fabric8-devops? Any concerns about supporting these?	0	0	https://issues.jboss.org/browse/OSFUSE-276
5064	OSFUSE-276	Can we just move the few annotations used by fabric8-arquillian to fabric8-arquillian?	1	1	https://issues.jboss.org/browse/OSFUSE-276
5065	OSFUSE-270	Hi [~chirino], looking at ER1 f-m-p. There's a few new additions to goal's mix. Can you please verify that everything should go to product? Already opened OSFUSE-414 for import goal. {code} Fabric8 Maven :: Plugin 3.1.80.fuse-000001 Fabric8 Maven :: Plugin This plugin has 17 goals: fabric8:apply Base class for goals which deploy the generated artifacts into the Kubernetes cluster fabric8:build Builds the docker images configured for this project via a Docker or S2I binary build. fabric8:debug Ensures that the current app has debug enabled, then opens the debug port so that you can debug the latest pod from your IDE fabric8:deploy This goal forks the install goal then applies the generated kubernetes resources to the current cluster. Note that the goals fabric8:resource and fabric8:build must be bound to the proper execution phases. fabric8:distro Generates a tarball of all the dependent kubernetes and openshift templates fabric8:help Display help information on fabric8-maven-plugin. Call mvn fabric8:help -Ddetail=true -Dgoal=<goal-name> to display parameter details. fabric8:import Imports the current project into fabric8 so that it can be automatically built via Jenkins and appears in the fabric8 developer console fabric8:log This goal tails the log of the most recent pod for the app that was deployed via fabric8:deploy To terminate the log hit Ctrl+C fabric8:push Uploads the built Docker images to a Docker registry fabric8:resource Generates or copies the Kubernetes JSON file and attaches it to the build so its installed and released to maven repositories like other build artifacts. fabric8:resource-apply This goal forks the compile goal then applies the generated kubernetes resources to the current cluster. Note that the fabric8:resource goal must be bound to the proper execution phases. fabric8:run This goal forks the install goal then applies the generated kubernetes resources to the current cluster, builds the images, waits for the new pod to start and tails it to the console. Pressing Ctrl+C will then terminate the application. You can think of this goal as like a combination of `fabric8:deploy`, `fabric8:logs` then `fabric8:undeploy` once you hit Ctrl+C Note that the goals fabric8:resource and fabric8:build must be bound to the proper execution phases. fabric8:start Starts the application that was previously created via fabric8:deploy and then stopped via fabric8:stop fabric8:stop Stops the application that was previously created via fabric8:deploy fabric8:undeploy Undeploys (deletes) the kubernetes resources generated by the current project. This goal is the opposite to the fabric8:run or fabric8:deploy goals. fabric8:watch Used to automatically rebuild Docker images and restart containers in case of updates. fabric8:watch-spring-boot Runs the remote spring boot application {code}	1	1	https://issues.jboss.org/browse/OSFUSE-270
5066	OSFUSE-266	Does fabric8 need Jetty and Zookeeper? Can these dependencies be removed?	1	1	https://issues.jboss.org/browse/OSFUSE-266
5067	OSFUSE-266	I've completed a first cut of this. The updated versions should show up on the next release. However, I spoke to Darren and it seems most of the redhat versions listed in the spreadsheet are not available in the usual Nexus repositories yet. So for now, I've got the build pipeline to plug in the equivalent community versions. Versions that I have not updated (yet): *Sundrio* - It's used to generate Maven BOMs and code for the Kubernets Client. Not something an end user will interact with. There's a bug which currently prevents the version used in the kubernetes-client project being updated. Darren was fine with leaving it as-is. *Jetty* - Do we really need to bump Jetty to 9.3? Can we not stick with 9.2? Camel compiled ok, but there were some strange test failures, which makes me nervous about other unforeseen changes that may go unnoticed.Why do we need to diverge from the version used in the community release?	1	1	https://issues.jboss.org/browse/OSFUSE-266

#	Requirement	Questions	Answered	Successfully Answered	URL
5068	OSFUSE-266	My main concern for jetty was the difference between the versions used by Camel and Fabric8 - 9.2.x vs 8.1.x. The redhat version was more of an advisory for the closest version that has been built internally. Are there any issues with Camel and Fabric8 using different jetty versions?	1	1	https://issues.jboss.org/browse/OSFUSE-266
5069	OSFUSE-265	Do we need to do this for the tech preview?	0	0	https://issues.jboss.org/browse/OSFUSE-265
5070	OSFUSE-265	Does this cover both the current 2.x CDK and the new 3.x version based on minishift?	0	0	https://issues.jboss.org/browse/OSFUSE-265
5071	OSFUSE-265	Did this get done? I think I saw that there were in CDK 3.	1	1	https://issues.jboss.org/browse/OSFUSE-265
5072	OSFUSE-260	Should this work already with the latest image (docker-registry.usersys.redhat.com/fabric8/fis-java-openshift:2.0 6ffda078502b), or do we need a rebuild?	0	0	https://issues.jboss.org/browse/OSFUSE-260
5073	OSFUSE-260	Doesn't seem to work with the current image (was working previously) Tried with both the current docker-registry.usersys.redhat.com/fabric8/fis-java-openshift:2.0 and the prod build ER2 image brew-pulp-docker01.web.prod.ext.phx2.redhat.com:8888/jboss-fuse-6/fis-java-openshift:2.0-0 getting {noformat} $ oc logs sti-sb-build-1-build -f Pulling image "brew-pulp-docker01.web.prod.ext.phx2.redhat.com:8888/jboss-fuse-6/fis-java-openshift:2.0-0" ... Pulling image "brew-pulp-docker01.web.prod.ext.phx2.redhat.com:8888/jboss-fuse-6/fis-java-openshift:2.0-0" ... Receiving source from STDIN as file sti-sb.jar ========== Starting S2I Java Build ..... Copying binaries from /tmp/src/deployments to /deployments ... cp: cannot stat '/tmp/src/deployments/*': No such file or directory Aborting due to error code 1 for copying /tmp/src/deployments to /deployments error: build error: non-zero (13) exit code from brew-pulp-docker01.web.prod.ext.phx2.redhat.com:8888/jboss-fuse-6/fis-java-openshift:2.0-0 {noformat}	0	0	https://issues.jboss.org/browse/OSFUSE-260
5074	OSFUSE-257	Is this for multi-module builds? If so, that would require more than just a directory name.	0	0	https://issues.jboss.org/browse/OSFUSE-257
5075	OSFUSE-257	Ok, which "auto generated" build configuration are you guys referring to exactly? What is it that generates them?	1	1	https://issues.jboss.org/browse/OSFUSE-257
5076	OSFUSE-245	Can this be closed now?	1	1	https://issues.jboss.org/browse/OSFUSE-245
5077	OSFUSE-243	[~maschmid], I just added these to the known issues in the Release Notes. I tweaked the text of the last issue as follows: {quote} The Java runtime does not respect the container's memory limits and this can be a problem, because OpenShift will kill a container that tries to allocate memory over the container limit. This problem can be mitigated by customizing the ratio of available memory for Java heap space using the `JAVA_MAX_MEM_RATIO` option. By default, FIS images are configured to allocate half the available memory for Java heap space, but the ideal ratio is application-specific. {quote} Does that make sense?	1	1	https://issues.jboss.org/browse/OSFUSE-243
5078	OSFUSE-243	Can this be closed now?	1	1	https://issues.jboss.org/browse/OSFUSE-243
5079	OSFUSE-233	What's the forge command to do this?	1	1	https://issues.jboss.org/browse/OSFUSE-233
5080	OSFUSE-233	Can you demo this on call on Friday? [~kcbabo] [~chirino] Do we want to make this default approach in quickstarts, doc etc? Seem much more user friendly.	1	0	https://issues.jboss.org/browse/OSFUSE-233
5081	OSFUSE-231	This includes version aligning to the Fuse 6.3 redhat ER versions.	1	1	https://issues.jboss.org/browse/OSFUSE-231
5082	OSFUSE-226	Hi, I've seen tried your camel-amq quickstart with Spring Boot. There's some problem during starting the context and the work directory is not created. Resulting in periodical exception from timer route. {code} 2016-06-28 13:56:11.070 INFO 1 --- [main] o.a.c.impl.DefaultStreamCachingStrategy : StreamCaching in use with spool directory: /tmp/camel/camel-tmp-ff956b99-a1a6-4756-b8d1-c245527ea793 and rules: [Spool > 128K body size] 2016-06-28 13:56:11.115 WARN 1 --- [main] o.a.camel.component.file.FileEndpoint : Cannot auto create starting directory: work/jms/input {code}	0	0	https://issues.jboss.org/browse/OSFUSE-226
5083	OSFUSE-226	It's very strange, I double checked it and cannot experience the problem. I tested the following scenarios: - running the "main" from the IDE - using the spring-boot maven plugin "spring-boot:run" from cmd - deploying to openshift using -Pf8-local-deploy - generating a sample project using the archetype and deploying it to openshift using -Pf8-local-deploy In every case, I did not have any issue and the directory structure is created normally. Can you provide me with more details about the environment you are using? The creation of files is not mandatory for this quickstart, it was inherited from the karaf example...	0	0	https://issues.jboss.org/browse/OSFUSE-226
5084	OSFUSE-221	[~rhuss], do you have any example projects using the fabric8 mvn plugin /w binary builds?	1	1	https://issues.jboss.org/browse/OSFUSE-221
5085	OSFUSE-221	[~rhuss] I tried building git@github.com:fabric8-quickstarts/spring-boot-camel.git with -Dfabric8.mode=openshift, but the os build failed with: {code} The value of ALLOWED_UIDS is [] The value of DROP_CAPS is [] Creating a new S2I builder with build config: "Builder Image:\t\t\tfabric8/java-alpine-openjdk8-jdk\nSource:\t\t\t\tfile:///tmp/s2i-build826140978/upload/src\nOutput Image Tag:\t\tdefault/spring-boot-camel-1:fcbf8f88\nEnvironment:\t\t\tOPENSHIFT_BUILD_NAME=spring-boot-camel-1,OPENSHIFT_BUILD_NAMESPACE=default\nIncremental Build:\t\tdisabled\nRemove Old Build:\t\tdisabled\nBuilder Pull Policy:\t\tif-not-present\nPrevious Image Pull Policy:\talways\nQuiet:\t\t\tdisabled\nLayered Build:\t\t\tdisabled\nWorkdir:\t\t\t/tmp/s2i-build826140978\nDocker NetworkMode:\t\tcontainer:3b51bbfdb796264e54e3859a862b09549e84b3b2e8f0a966754f68f029934c11\nDocker Endpoint:\t\tunix:///var/run/docker.sock\n" warning: Image does not contain a value for the io.openshift.s2i.scripts-url label Receiving source from STDIN as archive ... error: build error: failed to install [assemble run] {code} Are you planning to S2I enable the fabric8/java-alpine-openjdk8-jdk image? If not then we need an easy way to switch the base builder image that the fabric8 plugin defaults to.	1	1	https://issues.jboss.org/browse/OSFUSE-221
5086	OSFUSE-221	Have you removed the build config from the previous run (if running f-m-p [3.1.11, otherwise it will be automatically updated ? Otherwise it will reuse the old one (which points to a Docker image). Your best bet is {{-Dfabric8.build.recreate=all}} (see also docs). However, if you are using the latest f-m-p then the buildconfig should be automatically updated when you change the buildStrategy (e.g.from docker --] s2i). Which version are you using ? 3.1.11 is the latest.	1	1	https://issues.jboss.org/browse/OSFUSE-221
5087	OSFUSE-221	We could change the generator that it looks up a property from a resource property file for the default value, so that we can have different resource property files for the different branches of f-m-p (master and redhat branch have different prop fils with different base images). Could this be a solution ? In that case it would be backed into f-m-p and should refer not to a specific version (but 'latest' or a major version like '2' which constantly is moved along when new base images are created). So the switch would not be in the quickstart but in f-m-p directly.	1	1	https://issues.jboss.org/browse/OSFUSE-221

#	Requirement	Questions	Answered	Successfully Answered	URL
5088	OSFUSE-221	Verified that s2i karaf builds work with the last change to s2i scripts. The s2i build commands to test with {{jboss-fuse-camel-log}} were: {code} mvn clean install -Dfabric8.generator.karaf.from=fabric8/s2i-karaf -Dfabric8.mode=openshift -Dfabric8.build.recreate=all oc new-app --image-stream=jboss-fuse-camel-log {code} Note the following: * The _fabric8_ options had to be specified. I thought there were being defaulted. Will check in the f-m-p generator to see what's going on. * The new-app command was required to deploy the image stream generated by the f-m-p. [~pirx] any specific reason why the f-m-p didn't create the deployment config? Does it have to do with the f8 ci/cd workflow from fabric8 forge? The s2i script is working, and the resultant image contains an extracted Karaf container, which will be useful for debugging the image.	0	0	https://issues.jboss.org/browse/OSFUSE-221
5089	OSFUSE-218	[~fbolton] Should this be reassigned to some one in doc team? Do you have all the info you need?	1	1	https://issues.jboss.org/browse/OSFUSE-218
5090	OSFUSE-217	3.2.16.RELEASE --> 4.2.5.RELEASE (for all modules) - fixed log configuration to use spring-boot defaults (slf4j with logback) - fixed test dependencies on camel-spring3-test to camel-spring-test (needed to run unit tests with the correct spring version) - hacks to overcome https://issues.apache.org/jira/browse/CAMEL-10060 (they are needed in order to run unit tests in a fat-jar, not needed in production) Packaging: - module packaged in a spring-boot fat-jar - additional dependencies included in the fat jar (and transitive dependencies): -- camel-spring-boot-starter -- camel-core-xml -- shrinkwrap-api -- shrinkwrap-impl-base -- junit - test classes and resource left on file system and included in classpath Dependency Management: - camel dependency management only I've done some tests with spring-boot dependencies first, then camel ones. The result is a different "downgrading" maven hell (eg. camel components requiring Jetty 8, while spring-boot brings Jetty 9). As conclusion: when the dependency management is inverted, different problems arise, but a solution can, of course, be found (except for components included in point 4). 1) The following components are not supposed to work with Spring-Boot: camel-blueprint: it's OSGi-related. camel-cdi: afaik, CDI is not supposed to be used in Spring (while Spring beans can be used in a CDI container). camel-eventadmin: it's OSGi-related. camel-jibx: known compatibility issues with JDK 8, and we don't have S2I images with JDK 7. The only way to make it work with JDK 8 seems to be including a SNAPSHOT lib (BCEL-6.0-SNAPSHOT). camel-paxlogging: it's OSGi-related. 2) The following components work with some lib version hacks: camel-ahcws: Including optional libs grizzly-websockets and connection-pool (bom version). If not included, many unit tests *assertions* fail. camel-cassandraql: Set netty-all version at 4.0.27.Final, Removed dependecies com.codahale.metrics:*, Added io.dropwizard.metrics:metrics-core:3.1.0 camel-jclouds: Downgraded guice 4.0 --> 3.0 and guava 19.0 --> 17.0 camel-jcr: Downgraded lucene-core 5.4.1 --> 3.6.2.redhat-8 camel-mongodb: Added spring-data-mongodb v. 1.9.2.RELEASE camel-mongodb-gridfs: Added spring-data-mongodb v. 1.9.2.RELEASE camel-olingo2: Upgraded javax.ws.rs-api 2.0-m10 --> 2.0.1 camel-servlet: Added spring-boot-starter-web v. 1.3.3.RELEASE, upgraded spring-web to 4.2.5.RELEASE camel-spring-redis: Upgraded spring-data-redis 1.3.4.RELEASE --> 1.6.4.RELEASE camel-spring-ws: Added spring-boot-starter-web v. 1.3.3.RELEASE, upgraded spring-webmvc to 4.2.5.RELEASE camel-swagger (deprecated): Scala 2.11.7 --> Scala 2.10.4 3) the following components require a more sophisticated hack: camel-jpa --> The component has some classes marked as @Entity both in the test folder and in the main folder (eg. JpaTraceEventMessage or MessageProcessed, required to create an idempotent JPA consumer), but maven is configured to post-process only the test folder, so classes in the main folder are not "enhanced" at compile time. To enhance entity classes at runtime with OpenJPA, afaik, we need to pass a "javaagent:" configuration in the launch script, specifying the openjpa agent jar. So, the images at https://github.com/fabric8io-images/s2i/tree/master/java should be modified to download such agent and we need to provide at least a quickstart that uses the JAVA_AGENT or JAVA_OPTIONS property to include the openJPA agent. Why does it work when running tests in surefire or Idea? Don't know, but the integration test uses the "installed" artifact and classes included there are not enhanced. If the "main" classes are not enhanced, the error thrown at runtime is: org.apache.openjpa.persistence.ArgumentException: This configuration disallows runtime optimization, but the following listed types were not enhanced at build time or at class load time with a javaagent: " org.apache.camel.processor.interceptor.jpa.JpaTraceEventMessage" camel-ejb --> In theory, it can work in Spring, but the unit tests fail. *** Need to create a sample application *** camel-jgroups --> There are 5 Unit tests failing in surefire, IntelliJ-Idea and spring-boot integration tests. 4) The following components (most likely) will not work: camel-optaplanner --> Requires an old logback version 1.0.x, as it depends (internally, it doesn't have a logback.xml file) on a removed class:	1	1	https://issues.jboss.org/browse/OSFUSE-217

#	Requirement	Questions	Answered	Successfully Answered	URL
5091	OSFUSE-217	providing the wrong logback libraries. I've also run the full integration test again to check all other modules. *camel-spark* The logging problem can be solved by excluding log4j:apache-log4j-extras (imported by Hive), a library that ideally should provide additional appenders for log4j, but actually is a log4j replacement, as it embeds log4j classes, and they replace the ones in the slf4j bridge at runtime (goodbye OSGi, welcome Maven hell). So it's a standard logging problem (like the others that I have ignored because a simple workaround exists). *But Spark has some other issues, see point 2.1.* *1.1)* The following components are not supposed to work with Spring-Boot: camel-blueprint: it's OSGi-related. camel-cdi: afaik, CDI is not supposed to be used in Spring (while Spring beans can be used in a CDI container). camel-ejb: ejbs are not well suited to run in spring-boot. camel-eventadmin: it's OSGi-related. camel-paxlogging: it's OSGi-related. *1.2)* Modules that have known compatibility issues with JDK 8 (we don't have S2I images with JDK 7): camel-jibx: the only way to make it work with JDK 8 seems to be including a SNAPSHOT lib (BCEL-6.0-SNAPSHOT). *2.1)* These components are affected by a problem with dependency management. When a component requires a specific version of a library, different from the one specified in the *camel-parent* project, maven simply ignores it. The reason is that the *user* application pom should also import the dependency management from camel-parent, so Maven takes in consideration that version, ignoring the one required by the module (transitively). So, in order to make the following modules work, versions must be *overriden again* in the user application pom. camel-jclouds: guice 4.0 --> 3.0 and guava 19.0 --> 17.0 camel-jcr: lucene-core 5.4.1 --> 3.6.2.redhat-8 camel-spark: jackson-core, jackson-databind, jackson-annotations 2.6.3 --> 2.4.5 camel-swagger (deprecated): scala 2.11.7 --> scala 2.10.4 *2.2)* The following components have dependency issues like the ones in 2.1, but they are not related to transitive imports: camel-cassandraql: Set netty-all version at 4.0.27.Final, Removed dependecies com.codahale.metrics:*, Added io.dropwizard.metrics:metrics-core:3.1.0 camel-mongodb: Added spring-data-mongodb v. 1.9.2.RELEASE (because of a spring-boot issue fixed in 1.3.6.RELEASE) camel-mongodb-gridfs: Added spring-data-mongodb v. 1.9.2.RELEASE (because of a spring-boot issue fixed in 1.3.6.RELEASE) camel-servlet: Added spring-boot-starter-web v. 1.3.3.RELEASE, upgraded spring-web to 4.2.5.RELEASE camel-spring-redis: Upgraded spring-data-redis 1.3.4.RELEASE --> 1.6.4.RELEASE camel-spring-ws: Added spring-boot-starter-web v. 1.3.3.RELEASE, upgraded spring-webmvc to 4.2.5.RELEASE *3)* the following components require a more sophisticated hack: camel-jpa --> The component has some classes marked as @Entity both in the test folder and in the main folder (eg. JpaTraceEventMessage or MessageProcessed, required to create an idempotent JPA consumer), but maven is configured to post-process only the test folder, so classes in the main folder are not "enhanced" at compile time. To enhance entity classes at runtime with OpenJPA, afaik, we need to pass a "javaagent:" configuration in the launch script, specifying the openjpa agent jar. So, the images at https://github.com/fabric8io-images/s2i/tree/master/java should be modified to download such agent and we need to provide at least a quickstart that uses the JAVA_AGENT or JAVA_OPTIONS property to include the openJPA agent. Why does it work when running tests in surefire or Idea? Don't know, but the integration test uses the "installed" artifact and classes included there are not enhanced. If the "main" classes are not enhanced, the error thrown at runtime is: org.apache.openjpa.persistence.ArgumentException: This configuration disallows runtime optimization, but the following listed types were not enhanced at build time or at class load time with a javaagent: " org.apache.camel.processor.interceptor.jpa.JpaTraceEventMessage" camel-jgroups --> There are 5 Unit tests failing in surefire, IntelliJ-Idea and spring-boot integration tests. *4)* The following components are deprecated, so it's preferable not using them: camel-gae camel-ibatis camel-javaspace camel-jetty8 camel-kestrel camel-mina (tests failing, but we have a mina2 component, so I didn't check it) camel-	1	1	https://issues.jboss.org/browse/OSFUSE-217
5092	OSFUSE-213	Do you also need a datasource component? That is a component which is used when you call getConnection on a database or message broker that results in XAResource enlistment with the Transaction Manager.	1	1	https://issues.jboss.org/browse/OSFUSE-213
5093	OSFUSE-213	Are there any other specific use cases for this feature? (is it only about having a supportable transaction manager (narayana) with camel, or does it also enable any new use cases?) (should QE just verify existing functionality doesn't break with different TM, or are there any new use cases to cover with narayana?)	0	0	https://issues.jboss.org/browse/OSFUSE-213
5094	OSFUSE-213	[~nibbio84], In our last call you mention you think every thing should just works with exiting FIS 2.0 bits for the simple, non-scaled case with an attached PV. Could you validate and comment on this issue the list all the restrictions you would place around supporting that case?	1	0	https://issues.jboss.org/browse/OSFUSE-213
5095	OSFUSE-213	The standard Narayana service available on spring-boot 1.4.1 allows to configure automatically a transaction manager and recovery manager using a file system based transaction log. For a single pod, the location of the transaction log can be set using the 'spring.jta.log-dir' property to a mounted PV. The 'spring.jta.transaction-manager-id' should be set to a fixed value (e.g. '1', and not to the host name as in the "scalable" scenario). In this scenario, the transaction manager id should not change on pod restart. Openshift guarantees that only one pod is running for a given deployment config, but there are some situations where a pod may be considered "dead", allowing Openshift to create another pod. This results in a total of 2 pods running for the same deployment config, and it is a problem, since you'll have 2 recovery managers active for the same transaction manager id ('1'), and this is not allowed by Narayana. One common scenario where this can happen is when the node hosting the pod loses contact with the Openshift master (think to a network partition). After a given timeout, Openshift declares the node disconnected and the pod deleted, but the pod may be still running on that node. To avoid such situation, you need to do some "fencing" actions on the lost pod. Openshift does not currently support fencing (although it is in the roadmap). A working way to do fencing and avoid write concurrency between multiple recovery managers should be: # Using a RWO persistent volume # Use a storage provider that support native locking, to allow only one pod to mount the volume at time This way, even if two pods are active, only one is allowed to see the transaction log (and so, only one recovery manager should be able to recover transactions). Afaik RWO only is not sufficient, as Openshift does not implement any constraint if the storage provider does not have native locking. I've not investigated which providers/configurations provide these guarantees, but there should exist some good combination. Does it make sense?	0	0	https://issues.jboss.org/browse/OSFUSE-213
5096	OSFUSE-211	Is it evolving due to changes in the Kubernetes API itself or for other reasons? One way I've seen that handled in the past is to include the version number in the package name for your classes/interfaces. I wouldn't recommend it for something that is changing constantly though.	0	0	https://issues.jboss.org/browse/OSFUSE-211
5097	OSFUSE-211	Do we have a list of FIS artifacts / packages for which we do want to guarantee API compatibility?	0	0	https://issues.jboss.org/browse/OSFUSE-211
5098	OSFUSE-210	Does this layer on top of EP image from CE team?	1	0	https://issues.jboss.org/browse/OSFUSE-210
5099	OSFUSE-207	[~kcbabo] Hi Keith, can you clarify what you had in mind here?	1	1	https://issues.jboss.org/browse/OSFUSE-207
5100	OSFUSE-177	We're hoping there won't be a next build for GA. This does not look like a critical issue - can it come in the next update post-release?	0	0	https://issues.jboss.org/browse/OSFUSE-177
5101	OSFUSE-173	CR1 should be a release candidate, so all the image versions in pom.xml and templates should be whatever we expect to be the final GA name and versions. Will the released GA images be 1.0-7 ? Also, the templates should be updated as well. Currently there is still BUILDER_VERSION set to 1.2.5 in https://github.com/fabric8io/ipaas-quickstarts/blob/redhat/quickstart/cdi/cxf/quickstart-template.json#L41 , which is suspicious.	1	1	https://issues.jboss.org/browse/OSFUSE-173
5102	OSFUSE-173	Good point. I will set then version to 1.0-10 (which doesn't exist yet. is this a problem for testing ?) and the release this image version when going GA (when will this be ?) I'll also update the templates accordingly	1	1	https://issues.jboss.org/browse/OSFUSE-173

#	Requirement	Questions	Answered	Successfully Answered	URL
5103	OSFUSE-173	I wonder whether the final images that will be available via registry.access.redhat.com really need to have the same version tag as the ones we currently use? Can they just be retagged during the process of shipping the images to registry.access.redhat.com? (So they would end up with some nice version like 1.2 or something to align them versions with the other xPaaS images?)	1	1	https://issues.jboss.org/browse/OSFUSE-173
5104	OSFUSE-173	[~maschmid], I can change the image streams to use the 1.0-10 tag. Not sure whether/how the 1.0 stream corresponds to the docker tag version though. [~kconner], can you help us and explain how the OSBS stream and tag versions work for image streams?	1	1	https://issues.jboss.org/browse/OSFUSE-173
5105	OSFUSE-168	None. Would it fix OSFUSE-166 ?	1	1	https://issues.jboss.org/browse/OSFUSE-168
5106	OSFUSE-168	Would it really fix OSFUSE-166? (I would expect the transitive dependency to still be the same, it would just not be explicitly mentioned in our repository...)	0	0	https://issues.jboss.org/browse/OSFUSE-168
5107	OSFUSE-168	[~iocanel] this is now an old thread and in the meantime the _fabric8_project_bom_full_ has been used by _function_ and _portswizzler_ (an also by a CDI quickstart). Do you think that projects can be converted to use the platform bom?	1	1	https://issues.jboss.org/browse/OSFUSE-168
5108	OSFUSE-162	[~maschmid] Not sure I get the need for annotation processors here. Probes can be defined via the fabric8-maven-plugin configuration. Would this satisfy the requirement?	1	1	https://issues.jboss.org/browse/OSFUSE-162
5109	OSFUSE-149	[~janstey] wonder if you know how those -redhat builds of fabric8 works. Do you need to manually keep those camel/amq versions up to date yourself?	1	1	https://issues.jboss.org/browse/OSFUSE-149
5110	OSFUSE-144	Given then we will use different dependencies for spring boot (camel 2.18) and Karaf (Fuse 6.3), how many BOMs should we have in FIS? I guess we need * A single BOM file for spring boot + productized camel 2.18 + fabric8 stuff * A separate BOM with only the fabric8 stuff to be used with a Fuse 6.3 BOM for Karaf correct?	1	1	https://issues.jboss.org/browse/OSFUSE-144
5111	OSFUSE-143	I just checked the Fuse distribution and confirmed what Ioannis is seeing - Fuse uses 1.11. Is there a reason why we can't use that version? Agree with Nick's original comment that we need to sync on one version within Fuse Integration Services.	1	1	https://issues.jboss.org/browse/OSFUSE-143
5112	OSFUSE-143	So should we use 1.15 version for FIS ?	0	0	https://issues.jboss.org/browse/OSFUSE-143
5113	OSFUSE-142	What about https://github.com/fabric8io/fabric8/blob/redhat/pom.xml#L161 ?	1	1	https://issues.jboss.org/browse/OSFUSE-142
5114	OSFUSE-136	[~kcbabo], we can do this, but it would add a dependency on {{jboss-image-streams.json}} or {{xpaas-streams}} such that users have to install that template before they can use any of the quickstart templates. Need PM guidance on this. Also, [~chirino], what do you think?	1	1	https://issues.jboss.org/browse/OSFUSE-136
5115	OSFUSE-136	Is there any PR yet to put these into the ansible installer? https://github.com/openshift/openshift-ansible	1	1	https://issues.jboss.org/browse/OSFUSE-136
5116	OSFUSE-118	* Isn't ARTIFACT_DIR the final directory ? The MAVEN_BUILD_DIR is more specific as it points to the, well, maven build dir from where artefacts are copied to the target dir. Since we have all maven relevant configurations all prefixed with MAVEN_ it would make sense to keep that. But we could add ARTIFACT_DIR as alternative configuration to MAVEN_BUID_DIR (if it means the same, didn't check that yet) * The decision was not to allow the usage of a simplified mirror or proxy configuration, but to put that into the project to build so support for that is the responsibility of the project, not of the builder image. Please ask [~dbokde] and [~chirino] for a more detailed reasoning about this. * The difference between $MAVEN_ARGS and $MAVEN_EXTRA_ARGS are, that $MAVEN_ARGS overwrites the default (which doesn't make sense often, since e.g. certain plugins need to be called), $MAVEN_EXTRA_ARGS adds to the default instead of overwriting it, so it has some value. Since this is maybe only a minor difference, we could remove that, if decided. * Could you please elaborate on the point with the binary deployments ? Didn't get this, sorry. * We avoid /home/jboss because of certain permission problems when running under OpenShift. * Deleting the artifacts makes sense if you don't need incremental builds but want your images small (i.e. without a populate maven repository). remember, the image layer is checked in only after the assemble script has been finished. So that's also an extra feature. * BUILD_LOG_LEVEL support (i.e. setting 'set -x') could be added to karaf, right. will do that.	1	1	https://issues.jboss.org/browse/OSFUSE-118
5117	OSFUSE-118	Agreed that this shouldn't be a lonely decision (so what I stated above also solely expresses my personal opinion and mentions the decisions taken so fa ;-) * Ok I see the point. I wouln't mind to switch to ARTIFACT_DIR, too (maintaining two vars in parallel doesn't make sense probably) * Good point with MAVEN_ARGS_APPEND. If there are no objections I would change to this, too. * binary deployments: I see your point, that could be easily changed, too. * Are you sure that /tmp is not part of the commited image after the bulld ? My experience is that the complete source dir below /tmp is still there after the s2i build has finished and the pod is run from the then comitted image. * BUILD_LOG_LEVEL: Here I have the opposite opinion: BUILD_LOG_LEVEL is already used by s2i (and set when --log-level=5 is given for s2i or oc new-app), so why not hook into this s2i intrinsic mechanism to provide debug info on the finest grained log level ? btw, just remarked that is should be called BUILD_LOGLEVEL, so I will fix that (ref: https://docs.openshift.com/enterprise/3.0/dev_guide/builds.html#accessing-build-logs). This can also helps in (a) using the images and more importantly in (b) issue reports against the images.	1	1	https://issues.jboss.org/browse/OSFUSE-118
5118	OSFUSE-118	We don't currently use MAVEN_ARGS_APPEND nor MAVEN_CLEAR_REPO in the CE images, can you provide a link to the MAVEN_ARGS_APPEND modifications? I don't believe there will be a problem in adding them both. I'm not sure why we copy the source to the home directory but will find out, there doesn't appear to be any reason for it that I can determine.	0	0	https://issues.jboss.org/browse/OSFUSE-118
5119	OSFUSE-118	CE will use configuration/settings.xml for the location of the settings.xml (and not .m2/settings.xml as we do) --> https://gitlab.mw.lab.eng.bos.redhat.com/cloudenablement/jboss-dockerfiles/merge_requests/469 Should we do align our image (and then quickstarts), too ?	1	1	https://issues.jboss.org/browse/OSFUSE-118
5120	OSFUSE-109	Would it help, if I add an option e.g. '--probeRead <host>:<port>:<mBeanName>:<attribute>' to the JVM agent which: * Binds to the given IP and port (localhost:9779 in our case) * Allows only READ access to the given MBean and attribute That would be *safe* and *simple*, but maybe too simple as it couples the specific OS livenessCheck of how the image is started (i.e. a corresponding environment variable must be set for the image, too). Any opinion on this ? Would this be sufficient ? My goal is to have it as simple as possible on the Jolokia configuration side, since the number of available configuration options are already overwhelming.	1	1	https://issues.jboss.org/browse/OSFUSE-109
5121	OSFUSE-101	The fabric8-cdi is also able to inject a service via querying the k8s API if the environment variable doesn't exist (e.g. if the "injectee" service has been created after the "injector" has been deployed) Should this work also with the camel {noformat}{{service:name}}{noformat} ? Where does the {noformat}{{service:name}}{noformat} feature come from? Is it a camel core feature, or does it require a specific fabric8 bundle to be present?	0	0	https://issues.jboss.org/browse/OSFUSE-101
5122	OSFUSE-101	So, considering that we have the CDI+Camel and Karaf system properties support for accessing k8s service injection, can we mark this as a non-issue?	1	1	https://issues.jboss.org/browse/OSFUSE-101

#	Requirement	Questions	Answered	Successfully Answered	URL
5123	OSFUSE-90	[~dbokde] [~kconner] As mentioned previously, 3.0.5 is not good enough for the quickstarts (we need at least 3.2). How can we solve this ? Isn't there a newer, productised Maven version ? Any ideas ?	1	1	https://issues.jboss.org/browse/OSFUSE-90
5124	OSFUSE-81	From IRC {noformat} [11:08:23] <ro14nd> dsimansk: need some help to understand the issue :) so, the productised image fabric8/s2i-java-rhel will eventually in a registry registry.access.redhat.com ? [11:08:41] <dsimansk> ro14nd: yes [11:08:56] <ro14nd> but why are non-rhel user are supposed to access this image ? [11:09:26] <ro14nd> isn't fabric8/s2i-java targeted for them ? [11:10:03] <ro14nd> (which is available also on docker hub, the 'community' image) [11:10:44] <dsimansk> ro14nd: I would expect something like developers within the company not using rhel on their laptops etc. [11:11:28] <dsimansk> ro14nd: the thing is that fabric/s2i images can be used localy with docker plugin [11:11:42] <dsimansk> ro14nd: which is not the case for other products [11:11:56] <dsimansk> ro14nd: so I'm not sure how to tackle it actually :) [11:12:53] <ro14nd> ok, only on a rhel docker daemon, this registry is registered as a default registry. any other docker installation (fedora, docker) requires explicit qualified registries (as part of the same). Is this correct ? [11:13:04] <ro14nd> s/same/name/ [11:13:27] <dsimansk> ro14nd: yes [11:14:39] <ro14nd> what you can do is a 'mvn package docker:build -Pdocker.registry=registry.access.redhat.com' on such systems. [11:14:54] <dsimansk> ro14nd: sure [11:15:11] <dsimansk> ro14nd: so maybe mentioning it in documentation is enough [11:15:24] <ro14nd> this will transparently pull the base image from this registry and retags it locally to fabric8/s2i-java-rhel (hope so ;-) [11:15:34] <ro14nd> yeah, but I have no idea what documentation ;-) [11:16:10] <dsimansk> ro14nd: once there will be product docs rewritten from upstream, I hope :) [11:17:07] <ro14nd> We could put it into the README of the quickstart's redhat branch, but this will end up on github, too (which is probably not so nice because of internal information ?) [11:17:58] <dsimansk> ro14nd: hmm, I believe registry.access.redhat.com is public information [11:18:06] <ro14nd> Also, this '-Ddocker.registry=...' won't currently work if you do docker:build (which needs to pull the base image) and docker:push (which pushes the quickstart image to openshift) in the same step (e.g. with [11:18:14] <dsimansk> ro14nd: well at least you can't access it without subscription [11:18:16] <ro14nd> 'mvn -Pf8-deploy' [11:19:26] <ro14nd> haven't tested this use case (base image from other registry as the one the built image should pushed to) yet. I'll need probably a second configuration property for that so that both registries could be specified. [11:20:53] aileenc (~aileenc@vpn1-6-192.ams2.redhat.com) left IRC (Ping timeout: 612 seconds) [11:20:58] <ro14nd> dsimansk: ok, I'll update the issue to update the README documentation (but won't be able to do it on my own since i'll be on pto next week ...) {noformat}	0	0	https://issues.jboss.org/browse/OSFUSE-81
5125	OSFUSE-81	I can't seem to find the right value for -Ddocker.from for the productized images. Even falling back to "docker pull" with just the qualified image name fails for me: {code} docker pull fabric8-registry.usersys.redhat.com:5000/fabric8/s2i-java-rhel:1.2.0 {code} What am I screwing up here?	1	1	https://issues.jboss.org/browse/OSFUSE-81
5126	OSFUSE-81	Do you have fabric8-registry defined in insecured registries for docker deamon? It has to be with port also.	1	0	https://issues.jboss.org/browse/OSFUSE-81
5127	OSFUSE-77	Hi Marek, I thought about providing templates like you describe, but users would first have to use an archetype to generate a project then add it to a git repo before they could use it. As I have configured them now, users can directly point to the templates in GitHub like so: {code} oc create -f https://github.com/fabric8io/ipaas-quickstarts/blob/redhat/quickstart/cdi/camel/quickstart-template.json {code} I know we are not distributing the source, but it's not like users won't discover where the source lives in GitHub. Also, I agree that the templates are quite similar, but there are enough differences in parameter values between quickstarts to have a template per module. If we want to provide generic templates, we'll have to come up with a way to distribute them. One option could be to use a separate maven module for templates, and deploy the xml files in maven with a custom classifier, like openshift-template or something. Do you know how other RedHat projects are distributing theirs other than GitHub?	1	1	https://issues.jboss.org/browse/OSFUSE-77
5128	OSFUSE-77	Should the templates be present in the latest DR build in the archetypes already? (2.2.0.redhat-058) They don't seeem to be there in the archetypes, so I wonder if the templates commit antecedes the build tag, or if there might be some bug in the archetype assembly filter that skips the .json files.	1	1	https://issues.jboss.org/browse/OSFUSE-77
5129	OSFUSE-73	Can we remove this dependency?	0	0	https://issues.jboss.org/browse/OSFUSE-73
5130	OSFUSE-73	What are our options? Do we need to productise devops or can we move the functionality that ApiMan needs over into one of the projects that we are already productising?	0	0	https://issues.jboss.org/browse/OSFUSE-73
5131	OSFUSE-72	[~ceposta], can you check the file {{target/assembly/etc/org.ops4j.pax.logging.cfg}} to make sure it has the line: {code} # Root logger log4j.rootLogger=INFO, out, osgi:* {code} I actually do not see a difference in that file between plugin version 4.0.2 and 4.0.1, and version 4.0.2 seems to be logging correctly. Note that karaf s2i builder redirects the log output to {{stdout}} after extracting the assembly for the app docker image.	1	1	https://issues.jboss.org/browse/OSFUSE-72
5132	OSFUSE-72	[~dbokde] this seems fixed in 4.0.3 , should we close this jira ?	1	1	https://issues.jboss.org/browse/OSFUSE-72
5133	OSFUSE-69	Yep, the Karaf 4 plugin is released with the OS Fuse build. Branch is at https://github.com/jboss-fuse/karaf/tree/4.0.x.redhat-6-2-x-patch. [~davsclaus] where did you see that error?	1	1	https://issues.jboss.org/browse/OSFUSE-69
5134	OSFUSE-66	[~kcbabo] How important is this for 1st release?	0	0	https://issues.jboss.org/browse/OSFUSE-66
5135	OSFUSE-61	I guess the question is what is the point of having the MAVEN_MODULE in the first place? IMHO the point of MAVEN_MODULE should be to have an easy way of defining a specific submodule of a multi-module project as the thing that gets deployed. The mechanism doesn't need to be perfect and it can operate under a set of reasonable (and documented) constraints (such as directory names must match artifactIds, etc...) Perhaps instead of MAVEN_MODULE being the maven coordinates, it could instead mean the subdirectory of a multi-module project? Getting the maven coordinates from the given pom.xml would perhaps be easier than the other way around? (to consturct -pl argument) Having to specify both MAVEN_MODULE and OUTPUT_DIR will probably be OK too, however it does feel like a DRY violation (repeating same information twice)	0	0	https://issues.jboss.org/browse/OSFUSE-61
5136	OSFUSE-61	In the wild, this is rarely the case that artefactId and directory name correspond, especially not for multi module builds with many modules. These are typically grouped into other directory levels (e.g. the quickstarts project mentioned above). These extra directory levels are impossible to guess. The only authorative instance who can answer the question which Maven project belongs to which directory is the Maven reactor, which is queried with the call above. I suggest to use the mvn call above, and use the detected base directory if it exists. If not, the build will exit with an error message that the OUTPUT_DIR should be specified explicitely. wdyt ?	1	1	https://issues.jboss.org/browse/OSFUSE-61

#	Requirement	Questions	Answered	Successfully Answered	URL
5137	OSFUSE-53	There is a bunch of components at https://github.com/fabric8io/fabric8/tree/master/components Some of those should be disabled/not supported in OSFUSE. Could we come up with a list for that?	0	0	https://issues.jboss.org/browse/OSFUSE-53
5138	OSFUSE-53	[~iocanel] you have disabled some more. Are there anymore left or should we be good? https://github.com/fabric8io/fabric8/blob/redhat/components/pom.xml Ah you haven't disabled the devops stuff. Also [~rajdavies] should all of ActiveMQ/MQ be disabled as well? There is a mq module https://github.com/fabric8io/fabric8/blob/redhat/components/pom.xml#L53	0	0	https://issues.jboss.org/browse/OSFUSE-53
5139	OSFUSE-48	DR3 is ready now. * Which image should be used for "Java" applications using fabric8 workflow? * Which image should be used for "Karaf" applications using fabric8 workflow? * Which image should be used for "Java" applications using S2I workflow? * Which image should be used for "Karaf" applications using S2I workflow? * Is there any other image to be tested with DR3? (API Registry?, APIMan?, any other image that is planned but not done yet for this release 3.1 ?)	1	1	https://issues.jboss.org/browse/OSFUSE-48
5140	OSFUSE-48	Just found an issue in fabric8/sti-java-jboss-jdk8:1.0.12. Could you please test fabric8/sti-java-jboss-jdk8:1.0.13 ?	0	0	https://issues.jboss.org/browse/OSFUSE-48
5141	OSFUSE-48	Was java image pushed already too? I'm getting not found message. However, karaf image is present in the registry with 1.2.1 tag. {code} $ docker pull fabric8-registry.usersys.redhat.com:5000/fabric8/s2i-java-rhel:1.2.1 Trying to pull repository fabric8-registry.usersys.redhat.com:5000/fabric8/s2i-java-rhel ... not found Tag 1.2.1 not found in repository fabric8-registry.usersys.redhat.com:5000/fabric8/s2i-java-rhel {code}	0	0	https://issues.jboss.org/browse/OSFUSE-48
5142	OSFUSE-45	[~aileenc] [~chirino] - should the Fuse BOM manage dependency versions for Karaf?	1	1	https://issues.jboss.org/browse/OSFUSE-45
5143	OSFUSE-45	Is there anything more needs to be done for this?	0	0	https://issues.jboss.org/browse/OSFUSE-45
5144	OSFUSE-39	[~djcoleman] Are you familiar with this? Would it be useful for Fuse 7?	1	0	https://issues.jboss.org/browse/OSFUSE-39
5145	OSFUSE-33	I agree that for real world setup a fabric8.imagePullPolicySnapshot with 'Always' is absolutely useful. However, for a quick win when people are using the image with a local (meaning single node, possibly Vagrant based setup) this will require an otherwise unnecessary push to the OpenShift registry. Configuring the usage of the OpenShift registry is quite involved so it is likely that the 'quick' aspect gets lost. I will add some documentation to the (revamped) quickstart documentation on the fabric8.io site how the registry can be configured. But maybe with the new profiles added to all quickstarts today ('f8-build', 'f8-deploy' and 'f8-local-deploy' for a one-stop usage of all required plugin goals) one could add this property to the 'f8-deploy' profile and leave out for 'f8-build' and 'f8-local-deploy' ?	1	1	https://issues.jboss.org/browse/OSFUSE-33
5146	OSFUSE-12	[~janstey] Can you update handover doc?	1	1	https://issues.jboss.org/browse/OSFUSE-12
5147	OSFUSE-12	[~janstey] Can we run OS Fuse Perfectus with JDK8 to catch these issues earlier?	1	1	https://issues.jboss.org/browse/OSFUSE-12
5148	OSFUSE-12	Have a perfectus build going now with JDK 1.8.0_40. Let's see if everything else passes :-)	0	0	https://issues.jboss.org/browse/OSFUSE-12
5149	OSFUSE-8	The fabric8 project is using a maven plugin the generates 2 boms: i) artifacts we ship (pure bom). ii) artifacts we ship with all transitive dependencies (full). is there anything else that should be done in that front?	0	0	https://issues.jboss.org/browse/OSFUSE-8
5150	OSFUSE-1	& just to clarify what is the benefit of aligning with EAP dependencies? Seems unrelated to me.	0	0	https://issues.jboss.org/browse/OSFUSE-1
5151	OSFUSE-1	If we don't use the components from EAP, then will OpenShift Fuse commit to support the other versions? Or can EAP support e.g. Hibernate 5 for EAP 6.4 ?	1	1	https://issues.jboss.org/browse/OSFUSE-1
5152	OSFUSE-1	What CDI implementation does camel-cdi use when used as a plain java-main?	1	1	https://issues.jboss.org/browse/OSFUSE-1
5153	OSFUSE-1	Which implementation is supposed to be used in the product though? Do we want to support all of them? (I probably makes sense to limit the supported set to a single Weld version)	1	1	https://issues.jboss.org/browse/OSFUSE-1
5154	OSFUSE-1	Where do you expect to get this definitive list from? As far as I am aware there is no such list.	0	0	https://issues.jboss.org/browse/OSFUSE-1
5155	FUSEJON-126	[~tplevko], can you attach the log file? It would help to see the full stack trace. This looks like the fixed timeout of 60 seconds is not enough to wait for the server to start or stop. Perhaps it needs to be made configurable so that customers can adjust it to the amount of time their containers take to start/stop.	1	1	https://issues.jboss.org/browse/FUSEJON-126
5156	FUSEJON-126	[~tplevko], I ran some local tests, and it looks like the {{Failed with error code null:}} error will happen only when the start.bat/stop.bat script itself takes more than 15 seconds to finish. Unfortunately the 15 seconds is hardcoded. But more importantly, from your log output it looks like the actual Karaf JVM is being launched within your start script. The {{start.bat}} script I see in JBoss Fuse simply launches another shell in the background and returns immediately. So it is most likely a configuration issue and not a plugin issue. Can you verify whether your start script is configured to use start.bat, instead of karaf.bat and that no other parameters are being used that might launch the server within the {{start.bat}} script? I have opened an enhancement request FUSEJON-129 for the next release of the plugin to make the start/stop script and availability retries and timeouts configurable in the plugin.	1	1	https://issues.jboss.org/browse/FUSEJON-126
5157	FUSEJON-126	I attached the debug log file from rhq-agent. I am still hitting the described issue, the start script I used for starting the JBoss AMQ and also JBoss Fuse is the {{start.bat}} script. Can you please provide more information, how should the AMQ/Fuse be configured so the JON operation would be executed correctly?	1	1	https://issues.jboss.org/browse/FUSEJON-126
5158	FUSEJON-126	[~tplevko], the default start.bat script should just call the karat.bat script, launching another shell and quitting immediately. What happens when you try to run the start.bat script from a CMD console? Does it return immediately, or do you see Karaf being started in the same shell?	1	1	https://issues.jboss.org/browse/FUSEJON-126
5159	FUSEJON-122	Is this working now?	1	1	https://issues.jboss.org/browse/FUSEJON-122
5160	FUSEJON-94	Hi Dhiraj, I have just committed a84d74e , could you tell me if it's fine for you? Tested against AMQ 5.1 and 5.9	0	0	https://issues.jboss.org/browse/FUSEJON-94

#	Requirement	Questions	Answered	Successfully Answered	URL
5161	FMC-499	The problem with changing all the various profile ids (so close to 6.1 going out) is we have to test absolutely every profile again from scratch to check things don't break. I'm all for easier to browse/understand stuff; just right now its not a good time to completely change all the profile IDs (and ensure all the profile metadata, docs and demos are in sync). Though for 6.2 I'd like to replace the "-" is a folder hack we've got right now and use "/" for folders; plus I'd like to switch the import folder layout to the real git file layout (e.g. example/quickstart/cbr.profile versus the single directory containing "example-quickstart-cbr" which is the legacy layout). How about you just create a nice single wiki page that does what you need; and links to whatever you want? e.g. turn your issue into a wiki page; but use the current links/profile IDs)? Right now we have a simple top level page: https://github.com/jboss-fuse/fuse/blob/master/fabric/fuse-fabric/src/main/resources/distro/fabric/import/fabric/configs/versions/1.0/profiles/ReadMe.md Maybe we just need a full Overview.md page (maybe in the docs folder: https://github.com/jboss-fuse/fuse/tree/master/fabric/fuse-fabric/src/main/resources/distro/fabric/import/fabric/configs/versions/1.0/profiles/docs) which looks as you suggest with links to the actual profiles. See the various ReadMe.md files to see how to link to folders or link to profiles. e.g. here's the markdown to link to a particular quickstart: {code} [cbr](/fabric/profiles/example/quickstarts/cbr.profile) {code}	0	0	https://issues.jboss.org/browse/FMC-499
5162	FMC-478	Stan, you said you upgrade features by updated the feature repository version in the repositories tab, but I don't see how I can update the feature repository version. Are you saying that because there is only ever one version of a feature, Fabric will always use the latest one, based on what's available in the features repositories?	1	1	https://issues.jboss.org/browse/FMC-478
5163	FMC-339	There actually is a concept of abstract profiles in fabric and I just happened to add some filtering for abstract profiles so they can't be assigned to containers yesterday (go figure!). However at the moment, only the "mq-base", "mq-example-base" and the "fabric-ensemble-0000" profiles are abstract, and I don't have a toggle exposed in the UI to change it, so I'll keep this open. We should also open a ticket in FABRICPRV about setting some of the "base" profiles to abstract as well, I'll do that.	0	0	https://issues.jboss.org/browse/FMC-339
5164	FMC-320	Can you have a look at the install guide as well when you get a chance? Thanks!	1	1	https://issues.jboss.org/browse/FMC-320
5165	FMC-318	Luis, can you have a look at the admin guide as well when you get a chance?	1	1	https://issues.jboss.org/browse/FMC-318
5166	FMC-254	I think all the cases are covered... Will create additional tickets if anything else shows up.	0	0	https://issues.jboss.org/browse/FMC-254
5167	FMC-237	James: {quote} I wonder for the non-fabric version; should we have an easy way of running the Console + Fabric (and a single node ZK ensemble), then have a way to just register agents into the registry via a command line / web UI. Kinda like discovery, like HQ does. Then from Console's perspective everything is the same whether you're really using Fabric or not; its just if not using Fabric there's a little manual configuration step required. Even if folks are not using fabric to configure/provision stuff; it'd be much easier to just add the fabric-agent feature into each JVM, provide the ZK URL and then you don't have to manually configure agents to get Console to work. i.e. you can use Fabric purely as a discovery registry and not use any of the provisioning/configuration/elastic stuff. We should encourage folks to dip their toes in (or maybe walk barefoot on our soft Fabric is a better metaphor :) - particularly as I can see a really common use case of Fabric just being really easy discovery of JVMs/brokers, load balancing & Console registration - even if folks hardwire all their configs & deployment units, so its a fairly static system. {quote}	0	0	https://issues.jboss.org/browse/FMC-237
5168	MR-861	Can you show me the whole camel route? If your camel route never send response back to the client, the continuation on the CxfConsumer side will be blocked for ever. I added a new option in [camel 2.14.0.| https://issues.apache.org/jira/browse/CAMEL-7401] to specify the continuation timeout. But I'm not sure if it relates to your question.	0	0	https://issues.jboss.org/browse/MR-861
5169	MR-849	[~willem.jiang] Can this be marked resolved?	0	0	https://issues.jboss.org/browse/MR-849
5170	MR-825	Isn't there a context-suspend command? If so you can start it after a suspend	1	1	https://issues.jboss.org/browse/MR-825
5171	MR-822	Did you get a chance to work on this yet? I would like to get into next round of kits.	0	0	https://issues.jboss.org/browse/MR-822
5172	MR-814	Hooman - can you create PATCH request for the 2nd change related to quickfix component? To add the dependency on mina to 6.0 branches.	1	1	https://issues.jboss.org/browse/MR-814
5173	MR-796	Roman, can you verify that the fix works in the next fuse 6.1 build? Thanks, Dhiraj.	1	0	https://issues.jboss.org/browse/MR-796
5174	MR-791	Willem is there more work on this, or can this ticket be resolved?	1	1	https://issues.jboss.org/browse/MR-791
5175	MR-791	Willem is this fixed? And did you fix this at Apache also? https://issues.apache.org/jira/browse/CAMEL-6925 If so can you assign and resolve the Apache ticket?	0	0	https://issues.jboss.org/browse/MR-791
5176	MR-784	Wouldn't it be better to host Smooks component at Camel Extra, as we used to do? We release Camel Extra every time we release the Apache Camel. If Smooks will be hosted under Camel Extra, then Camel Extra team will make sure that there is a Smooks component working with every possible version of ASF Camel.	1	1	https://issues.jboss.org/browse/MR-784
5177	MR-596	I have attached a camel-jms JAR that has this improvement merged. If you for example use the latest Fuse MR 2.8 release, 2.8.0-fuse-03-06, then you can possible just renamed this -SNAPSHOT JAR as that version, and use it as a drop in replacement. Then you should have the replyToType option you can configure, accordingly to documentation at http://camel.apache.org/jms	0	0	https://issues.jboss.org/browse/MR-596
5178	MR-530	I checked the Jetty version of latest CXF (2.4.x, 2.3.x) , they are using 7.4.5.	0	0	https://issues.jboss.org/browse/MR-530
5179	MR-525	Yes Jetty 7.4.1 causes a NPE on startup. Using the 7.4.5 from the 2.8 onwards release works fine.	0	0	https://issues.jboss.org/browse/MR-525
5180	MR-414	Hi Lowry, Do you need anything else? Or can we close this JIRA?	1	1	https://issues.jboss.org/browse/MR-414
5181	MR-413	Willem can you take a look at this when you got a bit time.	1	1	https://issues.jboss.org/browse/MR-413
5182	MR-411	I've fixed the issue at Apache. Can we backport that in all our active camel branches?	1	1	https://issues.jboss.org/browse/MR-411
5183	MR-404	Is there any other dynamic placeholders thats useable, other than replyTo which would be the queue name defined in the jms endpoint, eg {code} <to uri="activemq:queue:foo"/> {code} That would mean the {{replyTo}} would be {{foo}}, and thus the example would use a persistent reply queue named {{foo.reply}} based on the pattern. So is there other placeholders we need to support in that name pattern?	0	0	https://issues.jboss.org/browse/MR-404
5184	MR-388	20549 merged to fuse 2.4 branch in rev: 22324.	0	0	https://issues.jboss.org/browse/MR-388

#	Requirement	Questions	Answered	Successfully Answered	URL
5185	MR-368	It was in fact Camel 2.4 which uses Spring 3.0.3 by default. Camel 2.5 now uses Spring 3.0.4 by default. Spring 3.0 is backwards compatible with Spring 2.5, and therefore you only need to change the namespace schemas if you want to take advantage of some of the new features in Spring 3.0. But its a good idea to keep the schema version in sync with the Spring version used. Order of dependency. TIP: The <camelContext> tag also has a depends-on attribute you can use. > The XML namespace for the Spring DSL schema has changed between Apache Camel 1.x and Apache Camel 2.2, as follows: A typo, it should be Camel 2.x instead of Camel 2.2 Or maybe just Camel 2.0 instead. > Apache Camel 2.2 uses the new DefaultErrorHandler, instead of the A typo, it should be Camel 2.x instead of Camel 2.2 Or maybe just Camel 2.0 instead.	1	1	https://issues.jboss.org/browse/MR-368
5186	MR-368	Thanks for the speedy review, Claus! Just one thing puzzles me: where is the reference to Spring 3.0.3 in the doc? In the draft I attached here, it only references 3.0.4.	0	0	https://issues.jboss.org/browse/MR-368
5187	MR-352	That is odd because CAMEL-3057 is only fixed recently and hence will be part of Camel 2.5.0. But I guess there has been other commits after the Camel 2.2.0 which could have resolved the customer issue. But you are 100% sure the attached zip contains the patch against MR 2.2.0 which fixes the customer issue? Do you have commit rights to the branch? If so you could most likely build the patched camel-cxf .jar and provided it to the customer.	1	1	https://issues.jboss.org/browse/MR-352
5188	MR-295	Willem have you implemented this? If so can you verify that the fix got synced to FUSE? I can see the wiki page has been updated with a new {{chuncked}} option you can set to {{false}} to force it to not use it http://camel.apache.org/jetty	1	1	https://issues.jboss.org/browse/MR-295
5189	MR-291	Jonathan what is status of this one?	1	1	https://issues.jboss.org/browse/MR-291
5190	MR-291	Isnt it possible to skip this then? Spring Javaconfig is going to be @deprecated and become standard in Spring 3.0 so we might just wait a bit till its mainstream and then base it on Spring 3.x .	1	1	https://issues.jboss.org/browse/MR-291
5191	MR-277	Hi Can you post a sample what you want please. And pin point out exactly what your problem is?	1	1	https://issues.jboss.org/browse/MR-277
5192	MR-273	Is the list okay Paul?	1	1	https://issues.jboss.org/browse/MR-273
5193	MR-262	Martin, and you want to reuse that session why? Is the TO endpoint on the same host? etc. I wonder why you want to reuse that session if you send something to a totally different sever? Or maybe MINA is totally capable of handling this using a single session? You can do request/reply like this anyway from(mina?sync=true) -> processor Since we have sync=true Camel will send the reply back from the OUT of the processor.	1	1	https://issues.jboss.org/browse/MR-262
5194	MR-260	Hi all, The latest Camel unit & ESB integration testing has passed. We still need the bundles thou... could someone add a note here when the bundles/specs are released? MR trunk needs to be updated to use these before it can be released. After that update we should be good to go though. Cheers, Jon	1	1	https://issues.jboss.org/browse/MR-260
5195	MR-260	Hello, Thanks Jonathan for changing the versions and checking them. I am waiting for the camel trunk checkin build to pass and then I will start working on releasing camel trunk. But can we ensure there are no further checkins while I will be working on the release? Thanks, Rahul	1	1	https://issues.jboss.org/browse/MR-260
5196	MR-252	Also check camel examples - do we need to ship them all - or just keep to core examples - like jms/cxf/osgi/spring etc	0	0	https://issues.jboss.org/browse/MR-252
5197	MR-252	Good report - are we doing the same for the other bits too - ESB, CXF and all that? Is there any way to consolidate them into a single report? And - how do I run it ? :) A lot of space gets taken up with Maven plug-in prerequisites too. I'd expect this for Enterprise to come in at > 500MB. This is the kind of stuff that Nexus doesn't supply, and it's fairly vital to controlling our dependencies, as well as producing legal trails for approval. I think that we still need to put the test dependencies in, right? People who are building the software will most likely need to test it :)	1	1	https://issues.jboss.org/browse/MR-252
5198	MR-250	What kind of persistent capabilities are customers looking for, to have out of the box? - Database using JPA? - JMS (I wonder how to do this nicely) - Local files For JPA we have some code for the Tracer and IdempotentRepository. So I think we should provide at least this one. JMS is interesting as most people using JMS messaging.	1	1	https://issues.jboss.org/browse/MR-250
5199	MR-246	Hmmm.. its downloaded fine for me. I just changed {code} --- parent/pom.xml (revision 16393) +++ parent/pom.xml (working copy) @@ -41,7 +41,7 @@ <activemq-version>5.3.0.5-fuse</activemq-version> <apacheds-version>1.5.4</apacheds-version> <axiom-version>1.2.8</axiom-version> - <cxf-version>2.2.2.2-fuse</cxf-version> + <cxf-version>2.2.4-psc-01-00</cxf-version> {code} But now it fails to compile and I'm not sure if its a simple fix or not. Looking at the trunk version of camel-cxf, a major refactor was done updating to a 2.2 SNAPSHOT of CXF. Wondering if a similar fix will be needed for the 1.x branch...	1	1	https://issues.jboss.org/browse/MR-246
5200	MR-246	Hello, There are SNAPSHOT dependencies which need to be resolved. Can someone have a look and let me know SNAPSHOTS are to be replaced with which version? components/camel-blueprint/pom.xml {code} <dependency> <groupId>org.apache.aries.blueprint</groupId> <artifactId>aries-blueprint</artifactId> <version>1.0.0-incubating-SNAPSHOT</version> </dependency> {code} Also, this needs to be changed from {code} <parent> <groupId>org.apache.camel</groupId> <artifactId>camel-parent</artifactId> <version>1.6.2-SNAPSHOT</version> </parent> {code} to, {code} <parent> <groupId>org.apache.camel</groupId> <artifactId>camel-parent</artifactId> <version>1.6.2.0</version> </parent> {code} platforms/karaf/features/pom.xml {code} <parent> <groupId>org.apache.camel</groupId> <artifactId>karaf</artifactId> <version>1.x-fuse-SNAPSHOT</version> </parent> {code}	1	1	https://issues.jboss.org/browse/MR-246
5201	MR-246	Hello, The deploy build for camel is failing, http://teamcity.boston.amer.iona.com:8111/viewLog.html?tab=buildLog&buildTypeId=bt845&buildId=254966 Can someone have a look at it? Rahul	1	1	https://issues.jboss.org/browse/MR-246
5202	MR-246	Jonathan can you take a look? Looks weird? Hadrian was able to cut a 1.6.2 release at Apache the other day	0	0	https://issues.jboss.org/browse/MR-246
5203	MR-220	What is the rational to route payloads that you cannot de serialize? As soon any content based router is hit then Camel needs access to the payload. So you can really only use it as a pass through scenario? And in many cases it can be an oversight that the class is needed. And thus you want this exception. If you want to send data without having the classes etc. then send it as a neutral payload such as byte[]. And assemble it manually on the client side that has the needed classes.	1	1	https://issues.jboss.org/browse/MR-220
5204	MR-215	Since most work will be done on the CXF side, I think it will not take much time to port it from Camel 2.0 to Camel 1.x branch.	0	0	https://issues.jboss.org/browse/MR-215
5205	MR-185	Is there a deadline when this should be done, the review? I suppose I am one of the reviewers :)	1	1	https://issues.jboss.org/browse/MR-185
5206	MR-185	Hi Claus, The sooner the better. Is a week enough time? Next Monday (7/13)? If you need more time, let me know. --Paul DiSabito	1	1	https://issues.jboss.org/browse/MR-185

#	Requirement	Questions	Answered	Successfully Answered	URL
5207	MR-185	*What are components* > Exchange objects also provide a means of correlating messages, because related messages are stored in the same Exchange. Hmm I dont see it like that. For instance using JMS messages are correlated using the JMSCorrelationID. Each message is stored in their own Exchange but they have a header with the CorrelationID so you can correlate them if needed. > The Consumer, Endpoint, Consumer, Producer, and Exchange classes comprise a component Consumer is duplicated	0	0	https://issues.jboss.org/browse/MR-185
5208	MR-185	*General Dependencies* Its the full list of .jars. Many of these .jars are component specific and thus not always needed. But yes they can be included anyway. camel-core.jar activation-1.1.jar commons-logging-1.1.1.jar Is the mandatory .jars at a minimum. If using Java 6 then the activation-1.1.jar is included with Java itself and thus can be omitted.	0	0	https://issues.jboss.org/browse/MR-185
5209	MR-185	General observations The consumer / producer is mentioned a lot. I think some more basic examples is needed for the components covered. Some very basic and then some more commonly used examples. How to poll a folder and its subfolder. Do some basic processing and send to a JMS queue How to consume from a JMS queue, use a processor to set a file name and write to a file How to consume from a JMS queue, use a processor to set a file name and upload to a FTP server How to poll a folder and upload to a FTP server How to use a Processor to do message transformation, maybe as part of the consume from JMS write to file How to poll a folder but only filter specific files, eg all .txt files. I think people best learn and understand Camel when they see small and a little bigger routing samples. So it makes sense to focus on good samples. The options for each component would be advised to divide into multiple tables. 1 table for commonly used options. And another for seldom/advanced usages. The .jars needed in CLASSPATH has too many. It includes stuff needed for unit testing only such as easymock etc. Maybe try reducing the number of .jars to what is actually needed. The command: {{mvn dependency:list}} can help as it can list what is needed for = compile = what you need in CLASSPATH Maybe we should add a sample how to use Camel JMS with Sonic MQ now that its also a Progress product.	0	0	https://issues.jboss.org/browse/MR-185
5210	MR-183	The intercept() works between the EIP. So its basically like this from -> interceptor -> to -> interceptor -> cbr (when A, when B, otherwise C) -> interceptor -> to A -> interceptor -> to B -> interceptor -> to C The intercept() is not designed to work *inside* the EIP patterns. What is your use case? Looks like you need somekind of debug logging or what? And besides the when .. when .. otherwise does not work like a java/C style case. There are *no* fall through. Camel will only pick *one* output from the CBR and go this direction.	0	0	https://issues.jboss.org/browse/MR-183
5211	MR-166	This was fine for 1.x, however now in 2.x sslSocketConnector has been replaced with a map containing a separate SSLSockectConnector instance per endpoint. So the work around doesn't work as well for 2.x., you've got to create a map and use the port your jetty endpoint listens on as the key: {code} <bean id="jetty" class="org.apache.camel.component.jetty.JettyHttpComponent"> <property name="sslSocketConnectors"> <map> <entry key="8043"> <bean class="org.mortbay.jetty.security.SslSocketConnector"> <property name="password" value="..." /> <property name="keyPassword" value="..." /> <property name="keystore" value="..." /> <property name="needClientAuth" value="..." /> <property name="truststore" value="..." /> </bean> </entry> </map> </property> </bean> {code}	1	1	https://issues.jboss.org/browse/MR-166
5212	MR-162	What are the the dependencies for this?	1	1	https://issues.jboss.org/browse/MR-162
5213	MR-161	Hi Ulhas, I didn't quit understand this issue, can you give me a test case or more specific example for it. Thanks, Willem	1	1	https://issues.jboss.org/browse/MR-161
5214	MR-161	Willem, thanks for sharing your findings. What exactly is the Format processor, its is some XPath/XML specify or just a regular Camel processor where we fix it by hand?	1	1	https://issues.jboss.org/browse/MR-161
5215	MR-136	I just added support for a basic set of operators in the simple language. http://issues.apache.org/activemq/browse/CAMEL-1338 However the equals operator will be like java, so you should use == (two =) Will be in included as standard in Camel 2.0 in the camel-core so no need for any special .jars on the classpath.	1	1	https://issues.jboss.org/browse/MR-136
5216	MR-122	Is this task focused on the Camel data handlers? Or is this for something more like implementing a message transform scenario?	1	1	https://issues.jboss.org/browse/MR-122
5217	MR-112	Perfect! I would add CXF to the list. Thanks, Jaime On 2/27/09 2:47 PM, "Eric Johnson" <EMJOHNSO@progress.com> wrote: Jaime, Does this look like a good list of components to focus on? Eric ---------- From: Linda Stinson Sent: Friday, February 27, 2009 2:33 PM To: Eric Johnson Cc: Adrian Trenaman; David Greco Subject: Top 10 components Hi Eric, Here is the top ten that we are including in the Camel course, give or take a few. Linda - Beans (camel-core) - E-mail (camel-mail), POP3 to receive and SMTP to send, these also work with JavaMail - Files (camel-core) - FTP (camel-ftp) - JMS (camel-jms) - use for ActiveMQ as well - JPA (camel-jpa) - using a database as a queue via OpenJPA, Hybernate or TopLink - Web services, HTTP to send and Jetty to receive - MINA (camel-mina) - used for working with Apache MINA - SEDA (camel-core) - used for creating processing pipelines in the same Camel context - Mock (camel-core) used for creating tests	1	1	https://issues.jboss.org/browse/MR-112
5218	MR-111	Perfect! I would add CXF to the list. Thanks, Jaime On 2/27/09 2:47 PM, "Eric Johnson" <EMJOHNSO@progress.com> wrote: Jaime, Does this look like a good list of components to focus on? Eric ---------- From: Linda Stinson Sent: Friday, February 27, 2009 2:33 PM To: Eric Johnson Cc: Adrian Trenaman; David Greco Subject: Top 10 components Hi Eric, Here is the top ten that we are including in the Camel course, give or take a few. Linda Beans (camel-core) E-mail (camel-mail), POP3 to receive and SMTP to send, these also work with JavaMail Files (camel-core) FTP (camel-ftp) JMS (camel-jms) - use for ActiveMQ as well JPA (camel-jpa) - using a database as a queue via OpenJPA, Hybernate or TopLink Web services, HTTP to send and Jetty to receive MINA (camel-mina) - used for working with Apache MINA SEDA (camel-core) - used for creating processing pipelines in the same Camel context Mock (camel-core) used for creating tests	1	1	https://issues.jboss.org/browse/MR-111
5219	MR-88	How about we take this one on over in eclipse tools land?	1	0	https://issues.jboss.org/browse/MR-88
5220	MR-88	Yeah definitely. Can we just move this issue I wonder? Really, I think this should just be added to our newly minted FUSE tooling. I believe I created this JIRA before it was out :) Would it easy to add templates to the FID? They are here: http://projects.open.iona.com/projects/svn/iona/camel/trunk/etc/eclipse/camel_java_templates.xml http://projects.open.iona.com/projects/svn/iona/camel/trunk/etc/eclipse/camel_xml_templates.xml	1	1	https://issues.jboss.org/browse/MR-88
5221	MR-88	I've linked ETOOLS-30 to this report. Can I confirm the licensing of these files? ASL 2.0 or EPL 1.0?	1	0	https://issues.jboss.org/browse/MR-88

#	Requirement	Questions	Answered	Successfully Answered	URL
5222	MR-85	Hi Fintan, I think you need to update this Service Activator Pattern , since there some enhancements[1] of Camel bean-integration in the apache camel 1.5.0, specially we added some annotations which could be much convenience for the end user. BTW, you didn't mention any Spring Remoting related things in the doc, could you add some descriptions about it ? [1]http://activemq.apache.org/camel/bean-integration.html Cheers, Willem	1	1	https://issues.jboss.org/browse/MR-85
5223	MR-83	Hi Fintan, Camel supports a pluggable DataFormat[1] to allow messages to be marshalled to and from binary or text formats. I think it is an implementation of the Messaging Mapper Pattern. Can you add some description about it in this charpter? After went through the EIP book's Messaging Mapper section, I found DataFormat[1] make a connection between domain objects and the messaging infrastructure and the user will not touch any dirty work of the mapping staff. [1] http://activemq.apache.org/camel/data-format.html Willem	1	1	https://issues.jboss.org/browse/MR-83
5224	MR-75	@Fintan, I am pretty sure this is taken care of. Can we close it?	1	1	https://issues.jboss.org/browse/MR-75
5225	MR-72	Current Camel will take care of the URI will even the URI has the unsafe codes, Please double check with your example :) It could be much easy for user with {code:language=java} from("jms:dispatcher?selector=CountryCode='US'").to("cxf:bean:replica01"); {code} instead of {code:language=java} from("jms:dispatcher?selector=CountryCode%3d%27US%27").to("cxf:bean:replica01"); {code}	1	1	https://issues.jboss.org/browse/MR-72
5226	MR-70	That's what I was referring to in the block labelled "Scheduled poll component". Should I just rename this block to "Scheduled poll consumer" to make it clearer what I am talking about?	1	1	https://issues.jboss.org/browse/MR-70
5227	MR-68	Thanks for this update, Willem! There is something else I meant to ask in connection with this pattern, though. I have no idea how to define a multicast route in Spring XML. Can you give me an example of the correct syntax? E.g. how would you translate your Java DSL route (above) into Spring XML?	1	1	https://issues.jboss.org/browse/MR-68
5228	MR-66	Thanks very much, Jonathan and Willem, for investigating the correctness of the XML examples and for fixing the bug. Just as a matter of interest, is it now possible to place the load balance strategy tag at _either_ location (beginning or end of the loadBalance element contents)?	1	1	https://issues.jboss.org/browse/MR-66
5229	MR-50	Fault message - how about adding a note to the effect that Fault messages are _deprecated_ and will be removed in a future release of Camel? That might be more useful for the current users of Camel who want to know what a Fault message is for. First mention of Java DSL - the Java DSL is covered in another book, "Defining Routes". Would it be Ok simply to insert a cross-reference to that book here? (The cross-reference should be mapped to a clickable link, if our book-building system behaves properly :-)	1	1	https://issues.jboss.org/browse/MR-50
5230	MR-50	In the block where I discuss "Message Structure", shouldn't I also mention attachments (I only found out about them recently)?	1	1	https://issues.jboss.org/browse/MR-50
5231	MR-49	Fixed in rev http://projects.open.iona.com/projects/trac/iona/changeset/45779 The fix employs a technique similar to that of the pipeline processor to preserve the exchange type. Basically, in order for a type of Exchange to be propagated, it must override the newInstance method so the routing slip has a way to create a clean instance of it (this is the case for the JbiExchange). The test case I came up with was purely routing slip specific. The stack shown above does not include the routing slip. Is there a more specific test case that the customer was hitting this stack against?	0	0	https://issues.jboss.org/browse/MR-49
5232	MR-30	General comments: - I didn't see any reference to the EIP book. Very important! Camel page (http://activemq.apache.org/camel/enterprise-integration-patterns.html) has a reference to the book. - Check unfinished parts with YOUR_ABSTRACT_HERE - Chapters 5-10 contain type hierarchies, but the role of different classes is not really clear. I would add at least links to javadocs. - Maybe more info on endpoint configuration would be helpful ? - a few typos * Preface | Open Source Project Resources - CXF no longer in incubator :). Links to be udpated * 1 | Exchanges - "On the request leg of the route (left to right), only the In message is accessible from the exchange and on the reply leg of the route (right to left), both the In message and the Out message (if any) are accessible." This is not actually true. At every moment the whole exchange is available. There isn't really a request and response path per se. The route is composed out of a series of processors. In general the in is the previous out before calling the processor and the out/fault gets populated by the processor. - believe setFault() is missing * 1 | Messages - Definition of marshalling/unmarshalling. Isn't it the other way around? In my interpretation, lower level is closer to the hardware. I would remove the analogy between type conversion and marshalling as it could get confusing. * 2 | Custom Processor - pre-post. Use an interceptor for that or you own extension of DelegateProcessor. * 2 | ExchangeHelper - I would not add big chunks of code, as it can change and can quickly get out of sync. I guess people don't read such big chunks of code anyway. I link to the javadocs on the apache site may be better. * 2 | RecipientList Processor Example - talks about the MulticastProcessor ?! - implementation should not be there - this is actually more useful: http://activemq.apache.org/camel/recipient-list.html * 3 | Implementing a Custom Type Converter - "Remember to import the org.apache.camel.Converter annotation type [...]". Not really necessary, as the compiler tells that. - You could note that the method name is not really relevant, the @Converter annotation and argument and return types are! * 4 | How to Implement a Component - see http://activemq.apache.org/camel/writing-components.html * 6 | Endpoint Interface - Endpoints are actually singletons per CamelContext - The code you're looking for is in DefaultCamelContext.getEndpoint(). The newly created endpoint is not added to the endpoints map if it's not a singleton, hence a subsequent getEndpoint won't find it and will create a new one. * 6 | Implementing the Endpoint Interfaces - again, i would not put implementations there * 9-10 - More detail on when/why one would implement her own Exchange/Message would help. Other than that it looks pretty good. There are some places where more details would help, and some inspiration could come from the questions users ask on the mailing lists http://www.nabble.com/Camel-f22882.html	1	1	https://issues.jboss.org/browse/MR-30

#	Requirement	Questions	Answered	Successfully Answered	URL
5233	MR-30	(http://activemq.apache.org/camel/enterprise-integration-patterns.html) has a reference to the book. - Check unfinished parts with YOUR_ABSTRACT_HERE - Chapters 5-10 contain type hierarchies, but the role of different classes is not really clear. I would add at least links to javadocs. - Maybe more info on endpoint configuration would be helpful ? - a few typos * Preface | Open Source Project Resources - CXF no longer in incubator :). Links to be udpated * 1 | Exchanges - "On the request leg of the route (left to right), only the In message is accessible from the exchange and on the reply leg of the route (right to left), both the In message and the Out message (if any) are accessible." This is not actually true. At every moment the whole exchange is available. There isn't really a request and response path per se. The route is composed out of a series of processors. In general the in is the previous out before calling the processor and the out/fault gets populated by the processor. [Fintan] I think there is a problem with my figure. I understand how an exchange works, but I think the figure gives the impression that there are two different channels for the request leg and the reply leg. I think I need to describe explicitly how processors are connected together in a chain. That should help to clear up any confustion. - believe setFault() is missing [Fintan] Actually it's missing in the source code(!) I should probably log a JIRA for this. You wouldn't necessarily notice it is missing, because it is defined in a sub-class (that is, in DefaultExchange). * 1 | Messages - Definition of marshalling/unmarshalling. Isn't it the other way around? In my interpretation, lower level is closer to the hardware. I would remove the analogy between type conversion and marshalling as it could get confusing. [Fintan] You're absolutely right. I got it backwards. Thanks for spotting this! * 2 | Custom Processor - pre-post. Use an interceptor for that or you own extension of DelegateProcessor. [Fintan] I'm still not entirely sure what I should do here. I'll send you an e-mail to ask you about it further. (By the way, this chapter is probably the weakest in the book. I really need to straighten it out.) * 2 | ExchangeHelper - I would not add big chunks of code, as it can change and can quickly get out of sync. I guess people don't read such big chunks of code anyway. I link to the javadocs on the apache site may be better. [Fintan] I will break this big chunk of code into smaller chunks and re-distribute it throughout the rest of the section. I will also include a link to the Javadocs. * 2 | RecipientList Processor Example - talks about the MulticastProcessor ?! [Fintan] Yes, I know this is confusing. Didn't I just say this is the weakest chapter? :-) - implementation should not be there [Fintan] Actually the main problem I have here is that I am not sure what to use as an example. But I will get back to you on that. - this is actually more useful: http://activemq.apache.org/camel/recipient-list.html * 3 | Implementing a Custom Type Converter - "Remember to import the org.apache.camel.Converter annotation type [...]". Not really necessary, as the compiler tells that. [Fintan] Hmm! I don't understand what you mean here (probably due to my ignorance of annotations). Do you mean the compiler will give you a warning? But surely you need to import the type anyway? - You could note that the method name is not really relevant, the @Converter annotation and argument and return types are! [Fintan] Good point! * 4 | How to Implement a Component - see http://activemq.apache.org/camel/writing-components.html [Fintan] This reminds me, I should remember to ask you how auto-discovery of components works in combination with Spring configuration of a component. * 6 | Endpoint Interface - Endpoints are actually singletons per CamelContext - The code you're looking for is in DefaultCamelContext.getEndpoint(). The newly created endpoint is not added to the endpoints map if it's not a singleton, hence a subsequent getEndpoint won't find it and will create a new one. [Fintan] Thanks, Hadrian! I had tried and failed to track down this snippet of code on my own. * 6 | Implementing the Endpoint Interfaces - again, i would not put implementations there [Fintan] Yes, but... there is no demo code that I can use in place of this example. I appreciate that implementation code is not ideal to use as a demo, but there don't seem to be any alternatives in this case. And I think it is helpful to see a complete example of an endpoint. Is this a reasonable	1	1	https://issues.jboss.org/browse/MR-30
5234	MR-25	How does it happen that so many jars are left out??	1	1	https://issues.jboss.org/browse/MR-25
5235	MR-25	Hi Aileen I know that the packaging of binary distributions versus the maven download needs to be resolved (esp for customers without Internet access). Should we be considering two separate modes of delivery? Eric	1	1	https://issues.jboss.org/browse/MR-25
5236	MR-25	Should you not be pointing at FUSE versions of AM & CXF rather than Apache?	1	1	https://issues.jboss.org/browse/MR-25
5237	MR-25	Hi Peter, what's the rationale behind which jars we include in the kit and which we pick up from other distribiutions?	0	0	https://issues.jboss.org/browse/MR-25
5238	MR-22	Are the route definition and the message containing the fault and/or a trace log available? Thanks.	0	0	https://issues.jboss.org/browse/MR-22
5239	MR-18	David, could you please add a few details about what the expectation is? I am mostly interested in the case where non camel producers/consumers are involved, say a native mina client of server. How would the correlation id be propagated? Should this address *only* the timeout situation?	1	1	https://issues.jboss.org/browse/MR-18
5240	MB-1066	HI Eric, So where did you all end up documenting "soWriteTimeout"? I checked here: http://fusesource.com/docs/broker/5.5/connectivity_guide/FMBConnectTransportRefWire.html I'm afraid I'm not seeing it ... Kindest regards, Susan	1	1	https://issues.jboss.org/browse/MB-1066
5241	MB-1066	Susan, I couldn't find it either. I've made the update for the 7.1 version of the content. Does this need to get pushed out to the Fuse Message Broker docs as well? Cheers, Eric	1	1	https://issues.jboss.org/browse/MB-1066
5242	MB-1009	I have applied the fix for this issue in the 5.5.x branch. I was also going to fix this in Trunk but noticed it was no longer there, is that expected?	1	1	https://issues.jboss.org/browse/MB-1009
5243	MB-988	@Claudio Here is a revised version of your patch. I added the ability for the servlet on the broker to inform the client transport that it accepts compression so that the client can also compress the outbound POST message payloads using gzip. Could you review and let me know what you think of this addition.	0	0	https://issues.jboss.org/browse/MB-988
5244	MB-967	Thinking on this a bit, the delivery flag needs to be treated like a message ack. Currently the broker becomes aware of redelivery when a consumer closes and provides the broker sequence id of the last message delivered to the consumer. Any inflight messages with id > that that sequence are deemed to be redelivered. In the even of a consumer failure, the close info from the consumer is lost, all unconsumed are deemed redelivered and there can be false positives. In the event of a broker failure, all redelivery information is lost. The change would have the consumer send an delivered_ack before a message is consumed, this would be a sync call that would require the broker to sync to disk. This will be as costly as a transaction commit! There is a chance of a false positive on the redelivered flag, if the consumer fails after the delivered ack is persisted by the broker but before delivery of a message. It may not be necessary to track updates in the index, in the event of failure we could always rebuild the index if this feature is enabled as it will be a rare event. Question: is fast recovery vital in this case?	0	0	https://issues.jboss.org/browse/MB-967
5245	MB-852	Awesome work Fintan. I think it's good as it is. Can you just let me know when you post it online, it'd be good to link it from Apache site and maybe blog/tweet about it.	1	1	https://issues.jboss.org/browse/MB-852
5246	MB-838	Is there a corresponding JIRA ticket to add documentation for a Fuse library JMX Management guide?	1	1	https://issues.jboss.org/browse/MB-838

#	Requirement	Questions	Answered	Successfully Answered	URL
5247	MB-822	Hi Gary, Thanks for all your help with this, we are still seeing a stack trace: {code} 2011-03-04 08:41:07,962 [38.171.43:58640] WARN ProtocolConverter - Exception occured processing: SUBSCRIBE ... {code} java.lang.SecurityException: User system is not authorized to read from: topic://c531dd48e123fd545516f496655781af at org.apache.activemq.security.AuthorizationBroker.addConsumer(AuthorizationBroker.java:119) at org.apache.activemq.broker.MutableBrokerFilter.addConsumer(MutableBrokerFilter.java:95) at org.apache.activemq.broker.TransportConnection.processAddConsumer(TransportConnection.java:549) at org.apache.activemq.command.ConsumerInfo.visit(ConsumerInfo.java:349) at org.apache.activemq.broker.TransportConnection.service(TransportConnection.java:310) at org.apache.activemq.broker.TransportConnection$1.onCommand(TransportConnection.java:184) at org.apache.activemq.transport.TransportFilter.onCommand(TransportFilter.java:69) at org.apache.activemq.transport.stomp.StompTransportFilter.sendToActiveMQ(StompTransportFilter.java:81) at org.apache.activemq.transport.stomp.ProtocolConverter.sendToActiveMQ(ProtocolConverter.java:140) at org.apache.activemq.transport.stomp.ProtocolConverter.onStompSubscribe(ProtocolConverter.java:425) at org.apache.activemq.transport.stomp.ProtocolConverter.onStompCommand(ProtocolConverter.java:188) at org.apache.activemq.transport.stomp.StompTransportFilter.onCommand(StompTransportFilter.java:70) at org.apache.activemq.transport.TransportSupport.doConsume(TransportSupport.java:83) at org.apache.activemq.transport.tcp.TcpTransport.doRun(TcpTransport.java:222) at org.apache.activemq.transport.tcp.TcpTransport.run(TcpTransport.java:204) at java.lang.Thread.run(Thread.java:619) {code} Was there any reason to leave this one? Ideally it would just be a one or two line message that the user is not authorized. Kindest regards, Susan	1	1	https://issues.jboss.org/browse/MB-822
5248	MB-728	Hi Tim, Do we have a snapshot build of this? Seán	1	0	https://issues.jboss.org/browse/MB-728
5249	MB-666	Thanks Dejan! Do you think this information will make it into the FUSE documentation set at some point?	1	1	https://issues.jboss.org/browse/MB-666
5250	MB-641	Hi Matt, yes, it is visible in the current fuse snapshot. Can you point David at my previous comment so we can validate that the proposed solution (implemented in AMQ-378) will meet their requirements. It would be great if they could validate a [5.4.0-fuse-SNAPSHOT|http://repo.fusesource.com/maven2-snapshot/org/apache/activemq/apache-activemq/5.4.0-fuse-SNAPSHOT/] with their use case and provide some feedback that it will or will not meet their strategic need. I am aware of the need for some JMX visibility (AMQ-2741) but I want to be sure we are on the same page and on the right track before completing that out. The relevant xml config for the new policy is:{code} <policyEntry ... ... <slowConsumerStrategy> <abortSlowConsumerStrategy /> </slowConsumerStrategy> ... </policyEntry>{code}	0	0	https://issues.jboss.org/browse/MB-641
5251	MB-594	Hello Aileen, Can you let me know how the notice file should be named? Rahul	1	1	https://issues.jboss.org/browse/MB-594
5252	MB-594	Hello Aileen, After unzipping the installer from previous version, http://repo.open.iona.com/maven2/org/apache/activemq/apache-activemq/5.3.0.4-fuse/apache-activemq-5.3.0.4-fuse-bin.zip I could only see a file called, NOTICE. (this also can be located at, http://fusesource.com/forge/svn/fusemq/branches/activemq-5.3.0-fuse/) Do you want me to replace this file or add the new file as "notices.txt"? Rahul	1	1	https://issues.jboss.org/browse/MB-594
5253	MB-590	Dejan, Can you update the 'fixed version' for this issue. It's only 5.3.x, correct? Thanks, Matt	1	0	https://issues.jboss.org/browse/MB-590
5254	MB-546	Hello Dejan, There is a SNAPSHOT dependency in one of the pom files. Can you let me know if this is right? File: activemq-5.3.1-fuse/assembly/pom.xml <!-- fuse web console --> <dependency> <groupId>org.fusesource</groupId> <artifactId>fuse-web-console</artifactId> <version>1.0-SNAPSHOT</version> <type>war</type> <classifier>war</classifier> </dependency> Rahul	1	1	https://issues.jboss.org/browse/MB-546
5255	MB-511	Could you provide specifics on what kind of corruption and when/how it it happens.	1	1	https://issues.jboss.org/browse/MB-511

#	Requirement	Questions	Answered	Successfully Answered	URL
5256	MB-511	file and it works for the first issue with "missing Journal files". However, it does not work for the second one and the broker still could not recover if a Journal file became corrupted. Here is some way to corrupt a Journal file: 1. Open it in an editor and randomly delete portions of the file. Since the file is binary, it doesnt really matter what I delete. What I saw was that the broker could not recover from this condition. Or 2. Randomly append some junk to the end of the file (with some output redirection like >>). This also results in a condition where the broker cannot recover and I have to delete all data files to make things work again. Below is exact steps I did: repeate steps from my previous comments; 5. wait for at least 3 or 4 db log files appear under "data\amq" directory; 6. stop the activemq broker and then open db-2.log file in a text editor. Then randomly delete some parts of the file and save the file again. 7. restart the activemq broker; 8. restart the consumer and the consumer continue receives messages till stops; Note, the consumer will only consume small part of messages before it stops and the broker starts to throw below exceptions: {quote} 2009-08-18 09:56:54,684 [ue://test.queue] ERROR Queue - Failed to page in more queue messages java.lang.OutOfMemoryError: Java heap space at org.apache.activemq.openwire.v5.BaseDataStreamMarshaller.looseUnmarshalByteSequence(BaseDataStreamMarshaller.java:639) at org.apache.activemq.openwire.v5.MessageMarshaller.looseUnmarshal(MessageMarshaller.java:229) at org.apache.activemq.openwire.v5.ActiveMQMessageMarshaller.looseUnmarshal(ActiveMQMessageMarshaller.java:101) at org.apache.activemq.openwire.v5.ActiveMQTextMessageMarshaller.looseUnmarshal(ActiveMQTextMessageMarshaller.java:101) at org.apache.activemq.openwire.OpenWireFormat.doUnmarshal(OpenWireFormat.java:368) at org.apache.activemq.openwire.OpenWireFormat.unmarshal(OpenWireFormat.java:279) at org.apache.activemq.store.kahadb.KahaDBStore.loadMessage(KahaDBStore.java:536) at org.apache.activemq.store.kahadb.KahaDBStore$KahaDBMessageStore$4.execute(KahaDBStore.java:230) at org.apache.kahadb.page.Transaction.execute(Transaction.java:728) at org.apache.activemq.store.kahadb.KahaDBStore$KahaDBMessageStore.recoverNextMessages(KahaDBStore.java:223) at org.apache.activemq.broker.region.cursors.QueueStorePrefetch.doFillBatch(QueueStorePrefetch.java:81) at org.apache.activemq.broker.region.cursors.AbstractStoreCursor.fillBatch(AbstractStoreCursor.java:227) at org.apache.activemq.broker.region.cursors.AbstractStoreCursor.reset(AbstractStoreCursor.java:99) at org.apache.activemq.broker.region.cursors.StoreQueueCursor.reset(StoreQueueCursor.java:157) at org.apache.activemq.broker.region.Queue.doPageIn(Queue.java:1279) at org.apache.activemq.broker.region.Queue.pageInMessages(Queue.java:1419) at org.apache.activemq.broker.region.Queue.iterate(Queue.java:1103) at org.apache.activemq.thread.DeterministicTaskRunner.runTask(DeterministicTaskRunner.java:84) at org.apache.activemq.thread.DeterministicTaskRunner$1.run(DeterministicTaskRunner.java:41) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:650) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:675) at java.lang.Thread.run(Thread.java:595) 2009-08-18 09:56:54,980 [ue://test.queue]	1	1	https://issues.jboss.org/browse/MB-511
5257	MB-488	Thanks very much Dejan! One question: Do you know if changes to apache trunk (like this one) will get merged back to the activemq-5.3.0-fuse branch? I believe that's the branch that 5.3.0.3-fuse will be released off of, and I just want to make sure. Thanks again for committing this!	1	1	https://issues.jboss.org/browse/MB-488
5258	MB-452	p71. "where the relevant user security data is stored in a pair of flag files" the use of 'flag file' is a little odd to me, property file? p84. On the Note: "The requirement for Broker A to be configured with credentials holds, even if Broker B is not configured to require authentication." - Is it possible to expand on this? I Don't understand what it adds. - Otherwise all great.	0	0	https://issues.jboss.org/browse/MB-452
5259	MB-451	typo: space needed, p12: Immunity to man-in-the-middle attacksbecause p13: Selfif (This space issue seems to be generic, may be in my viewer - a bullet point and an italic seems to not be followed by a space. Is this text correct?: I don't understand it: "There is no way for the attacker to obtain the target's private key, which makes it impossible for the attacker to authenticate itself to the client." On "a self-signed certificate", possibly mention that this is typically for a root cert and useless otherwise unless you have prior knowledge that the cert is good. On "Setting properties by programming" p24 see some additional info on setting up an SSLContext via code at: http://www.nabble.com/Alternatives-for-deprecated-setKeyAndTrustManagers-in-ActiveMQSSLConnectionFactory-td22724170.html#a22741202 and note the SSLContext thread local, accessed by set/getCurrentContext The activeMQ SSLContext is a thread local wrapper around a JSSE SSLContext, this is the same class that is configured through XML injection. Looks good otherwise,	0	0	https://issues.jboss.org/browse/MB-451
5260	MB-449	I used code-prettify project (suggested earlier) and made a few more modifications on the page. It should be more readable now. !Picture 4.png|width=800!	0	0	https://issues.jboss.org/browse/MB-449
5261	MB-365	Joe, can you give a few lines that describes "This way we would get exactly the behavior we expect the MB to have." given master/slave with jdbc persistence. start db, start master, start slave. master has lock, slave waits to aquire lock. DB dies. master see db failure and stops. Slave (I think) should see db failure and stop. Then using some intervention, DB restarts, master restarts, slave restarts. How different is the customer expectation. In what way does the proposed improvement fit into this scenario. to fully understand this I need to start with a test case so I will attempt to create one now, but I need to fully understand the expectations. Thanks for your help.	0	0	https://issues.jboss.org/browse/MB-365
5262	MB-331	If you look at the status page for the FUSE May releases, "activemq trunk" is building ActiveMQ 5.1.0-fuse-SNAPSHOT http://teamcity.boston.amer.iona.com:8111/statusPage.html?filterName=FUSE+May+Releases&Filter=filter So is this task done, or are you looking for something different? :)	0	0	https://issues.jboss.org/browse/MB-331
5263	MB-325	ok. good work. So the second send will block ? is it possible to have a failure on the second send, failIfLimitReached or something option? Can the large message get consumed? So long as the behaviour is deterministic up to a point we have something to test/verify/document. Without a systemUsage limit and with a very large message, what happens?	1	1	https://issues.jboss.org/browse/MB-325
5264	MB-319	Eric - we'd like to ship this tool in next FUSE MB patch. Can we get doc for it?	0	0	https://issues.jboss.org/browse/MB-319
5265	MB-319	There has been talk on and off about productising this tool for a long time, but it hasn't happened yet (has it really been four years?). Sure, we can close this issue. If the tool ever gets productised, we can open a doc issue for it then.	0	0	https://issues.jboss.org/browse/MB-319

#	Requirement	Questions	Answered	Successfully Answered	URL
5266	MB-315	hi Geoff, Eric, Some comments on the patch, first off the test looks good and second, we need Rob or Hriam to have a look at this one. My first impression is that the code that blocks in addMessageLast on reaching the limit should first check to see if it will block, then release locks and then block waiting for free space. So all the changes that make the cursor use a lock rather than sync methods make sense. However the: {code} while (!addMessageFirstImpl(node)) { Thread.sleep(1); }{code} idiom looks wrong to me. It is really a busy loop. Could it be more like:{code} if (systemUsage.getTempUsage().isFull()) { releaseLockThatStopsMessageComsumption(); systemUsage.getTempUsage().waitForSpace(); aquireLockThatStopsMessageComsumption(); } getDiskList().add(node);{code} I imagine this is not trivial but sort of what you are attempting to do. All other examples of Usage limits seem to be in the position where they can make some space available. In the TempStore case, the only thing that can make space available is to expireOldMessages() or to consume a Message. So these activities need to happen in parallel. Is there a case for adding a call to exporeOldMessages() before waiting? I guess that is another test! Hope this helps, Gary.	1	1	https://issues.jboss.org/browse/MB-315
5267	MB-315	Are the new issues logged anywhere in JIRA?	0	0	https://issues.jboss.org/browse/MB-315
5268	MB-309	Can you set these up to run against 5.0 branch as well?	1	1	https://issues.jboss.org/browse/MB-309
5269	MB-308	Have we made a definitive decision on how we are going to do these. Are they going to reflect the README template guidelines and provide more details in the README.txt files of the examples? Or are we going to just provide basic instructions of how to run the example with a link to FUSE or Apache documentation describing the example in detail, as it is now?	0	0	https://issues.jboss.org/browse/MB-308
5270	MB-304	We assumed that setAlwaysSyncSend was the property of the ActiveMQConnection, while trackMessages and maxCacheSize (not cacheSize) are the ones in FailoverTransport. Is there any documentation on the use of those properties ?	0	0	https://issues.jboss.org/browse/MB-304
5271	MB-304	do you think it is worth spending some time to try and make the test fail more consistently, all producers share the same proxy, if each had its own proxy that failed randomly over a period of a few seconds, would that help? I have run this a few times on windows but it always works :-(	0	0	https://issues.jboss.org/browse/MB-304
5272	MB-299	The information looks good though there is a lot of white space/blank pages in the document. Is it possible to reduce this ?	0	0	https://issues.jboss.org/browse/MB-299
5273	MB-289	Hi Aileen, this feature is now complete and in will make it into the 5.0.0.11 release but it needs doco. Could you create an issue for that and assign it to the right person? Thanks.	0	0	https://issues.jboss.org/browse/MB-289
5274	MB-280	Has this been fixed? 5.0.0.7-fuse has been released as of today.	1	1	https://issues.jboss.org/browse/MB-280
5275	MB-263	Hi Hiram, What is the create <broker-name> task suppose to do? Is it to allow running multiple broker instances in the same distribution? I guess an overview of the directory structure would help. :) Thanks!	1	1	https://issues.jboss.org/browse/MB-263
5276	SF-487	Is there any more work to be done on this or can it be closed?	1	1	https://issues.jboss.org/browse/SF-487
5277	SF-480	[16:57:42] <+tdiesler> ppires, the jaxrs runtime processes jaxrs annotations to setup the endpoint, right? [16:58:41] <+jstrachan> tdiesler once its been registered with the CXF service, yes - the issue is getting the DS component object into the CXF service so it can have its JAXRS stuff applied [17:00:17] <+tdiesler> jstrachan, can we not have a generic jaxrs component that, based on a set of jaxrs specific properties, sets up the jaxrs endpoint as part of the comp lifecycle? [17:01:08] <+jstrachan> tdiesler thats what that issue was for ;)	0	0	https://issues.jboss.org/browse/SF-480
5278	SF-463	Thanks super awesome thanks! :) Gonna start using it as soon as the next perfectus completes... Will this API call only work for WSDL endpoints; or should it work for both WADL and WSDL?	1	0	https://issues.jboss.org/browse/SF-463
5279	SF-396	Is this really the best we can do? If they have soap:faults in their WSDL contract, then that is an expected reply type. I would prefer that CXF have some easy to configure option to control the "noise". When people monitor their logs, and see huge stacktraces, they think its an error and want to have it fixed etc. For unexpected exceptions it may be desireable to have it log a stacktrace etc as that was not expected to happen.	1	1	https://issues.jboss.org/browse/SF-396
5280	SF-357	Can the fix version for this issue be updated to 2.3.x so this gets picked up in the 4.3.1 SP2 release. Thanks - Matt	1	0	https://issues.jboss.org/browse/SF-357
5281	SF-245	Attaching patch and test case. I ran the patch against the trunk and against the cxf-2.2.5-psc-01-xx-HF branch. The branch passes all RT tests. There were some test failures in the RT modules in the trunk, but I beleive them to be unrelated to the changes in the patch.	0	0	https://issues.jboss.org/browse/SF-245
5282	SF-239	The 2 tests that are failing are security related tests & Dan thinks we may need to add strong cypto to JDK for AIX We switched to JDK 6 recently Eamonn - can you investigate?	0	0	https://issues.jboss.org/browse/SF-239
5283	SF-239	I'm attaching the log of a successful run on pdrhas4-64 using ibm jdk 1.6 (SR7) with the unrestricted export jars installed in the JDK. As this test works fine with the ibm 1.5 jdk, can we just release note thet there is a bug in the ibm 1.6 jdk when using the restricted jars? If so will I put code into the test to skip it when the restricted jars are installed for the ibm 1.6 jdk?	1	1	https://issues.jboss.org/browse/SF-239
5284	SF-194	Is any progress being made on the review?	0	0	https://issues.jboss.org/browse/SF-194
5285	SF-194	Dan, do you have any free cycles to work on these technical reviews? Can you estimate a time of delivery? Cheers, Fintan	1	1	https://issues.jboss.org/browse/SF-194
5286	SF-172	Consumers making asynchronous requests will not receive the same exceptions returned *than* when they make synchronous requests. Should a ExecutionException not be an ExecutionException? In the section Catching an Exception, I think we should say that the get() method is used to get the response from the service and it is only at this point an exception is discovered. I guess this is a subjective thing but I think using the word exception after ExecutionException is a bit redundant. Like saying a basketball ball or a football ball..	1	1	https://issues.jboss.org/browse/SF-172
5287	SF-162	Is this for JDK 1.5?	1	0	https://issues.jboss.org/browse/SF-162
5288	SF-158	It should be possible to add this to FUSE SF 2.0.x, which will allow it to be included in FUSE ESB 3.3.x. If this is not possible, can we please state why?	0	0	https://issues.jboss.org/browse/SF-158
5289	SF-158	Are you saying that the support is already there Dan? In that case can we add some documentation around this?	0	0	https://issues.jboss.org/browse/SF-158
5290	SF-144	Hi Eric, I just went through this chapter, It looks good. Here are my comments for the chapter: 1. We need more examples to show how to configure the serviceName with the namespace and how to set the binding or servicefactory. It's difficult for the new user to write a right one without an example. 2. Can you add simple frontend configurations description such as <simple:server> <simple:client> ? Or you can tell some difference between <simple> and <jaxws>:) Cheers, Willem	1	1	https://issues.jboss.org/browse/SF-144
5291	SF-144	Willem, Can you point me to some good examples of configuring the binding and the serviceFactory? Thanks, Eric	1	1	https://issues.jboss.org/browse/SF-144
5292	SF-140	Was this resolved in the release 2.0.9.0? [CXF-1776|http://issues.apache.org/jira/browse/CXF-1776] is resolved.	1	1	https://issues.jboss.org/browse/SF-140
5293	SF-104	Jon, I've just sent you notices & license file - can you add into the kit? thanks, Aileen	1	1	https://issues.jboss.org/browse/SF-104
5294	SF-104	I didn't get the notices & license file... did you send them to janstey@iona.com?	1	1	https://issues.jboss.org/browse/SF-104

#	Requirement	Questions	Answered	Successfully Answered	URL
5295	SF-89	Is this still an issue for 2.1?	1	0	https://issues.jboss.org/browse/SF-89
5296	FUSEDOC-1704	Updated the PID prefix in the "Deploying into Apache Karaf" guide for the 6.3, 6.2.1, and 6.2 branches in the docs source. These changes should become visible in about a week, after we republish the docs.	0	0	https://issues.jboss.org/browse/FUSEDOC-1704
5297	FUSEDOC-1350	Fixed the doc source in the 6.2.1 and 6.3 branches. Replaced the description of the {{\-\-new-user-*}} options in {{fabric:container-create-ssh}} by the more detailed descriptions from {{fabric:container-create-cloud}}.	0	0	https://issues.jboss.org/browse/FUSEDOC-1350
5298	FUSEDOC-1192	Hi Jane, I marked this issue with Fix Version/s 6.3, because I understand that 6.3 is the version that will ship with the new UI interface (not 7.0). Is that correct?	1	1	https://issues.jboss.org/browse/FUSEDOC-1192
5299	FUSEDOC-954	Hopefully, this issue wil be fixed soon in the Publican build system. That is the ultimate solution to this problem. So, yes, I am in favour of waiting. By the way, do you have a reference to a Bugzilla bug to get the rendering fixed?	1	1	https://issues.jboss.org/browse/FUSEDOC-954
5300	FUSEDOC-924	These docs are affected: Fuse repo: janemurpheysmbp:Apache_Camel_Component_Reference jmurphey$ cd .. janemurpheysmbp:fuse jmurphey$ grep "technical preview" */*.xml */*/*.xml Installation_Guide/switchyard.xml: SwitchYard on JBoss Fuse is currently provided as a technical preview only. As such, it is not in scope of patching and is supported only for use in development environments. However, SwitchYard will be supported with the release of JBoss Fuse 6.2.1 with full documentation provided. Release_Notes/whats_new.xml: <para>The following features are provided on a technical preview basis and are not Tooling_User_Guide/new_project.xml: Camel, it is provided as a technical preview only. It is not ready Tooling_User_Guide/transformation_tooling.xml: <para> JBoss Fuse Transformation tooling is currently provided as a technical preview janemurpheysmbp:fuse jmurphey$ AMQrepo: janemurpheysmbp:amq jmurphey$ grep "technical preview" */*.xml */*/*.xml Configuring_Broker_Persistence/ReplicatedLeveldb.xml: <para>The Replicated LevelDB persistence adapter is provided for technical preview only. Migration_Guide/activemq.xml: <para>&prodnamefull; includes a technical preview of the new LevelDB store, which uses Migration_Guide/activemq.xml: <para>This is a technical preview only. The LevelDB store is currently not supported and is amq/Configuring_Broker_Persistence/ReplicatedLeveldb.xml: <para>The Replicated LevelDB persistence adapter is provided for technical preview only. amq/Configuring_Broker_Persistence/leveldb.xml: technical preview only.</para> Configuring_Broker_Persistence/leveldb.xml: technical preview only.</para> janemurpheysmbp:amq jmurphey$	1	1	https://issues.jboss.org/browse/FUSEDOC-924
5301	FUSEDOC-911	Hello [~jmurphey], I probably didn't describe stuff around encryption service correctly. Sorry for confusion. Can you please take look on following merge request [https://gitlab.cee.redhat.com/fuse-docs/amq/merge_requests/1] ?	1	1	https://issues.jboss.org/browse/FUSEDOC-911
5302	FUSEDOC-911	Can you merge the request then?	1	0	https://issues.jboss.org/browse/FUSEDOC-911
5303	FUSEDOC-900	{quote} Options Password encryption for JAAS login modules can optionally be enabled by setting the following login module properties: {quote} Where should users set those properties? {{$FUSE_HOME/etc/org.apache.karaf.jaas.cfg}} or somewhere else?	1	1	https://issues.jboss.org/browse/FUSEDOC-900
5304	FUSEDOC-900	[~jmurphey], about my [first comment|https://issues.jboss.org/browse/FUSEDOC-900?focusedCommentId=13100125&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13100125]: You might mention that users can either edit properties in {{$FUSE_HOME/etc/org.apache.karaf.jaas.cfg}} or deploy their own blueprint.xml as described in section "Example of a login module with Jasypt encryption". Can you let me know, when it's ready for review again? Thanks, Josef	1	1	https://issues.jboss.org/browse/FUSEDOC-900
5305	FUSEDOC-871	Freeman, in your description above, what exactly do you mean "without login" in "... to automatically encrypt the passwords in users.properties without login ..."?	0	0	https://issues.jboss.org/browse/FUSEDOC-871
5306	FUSEDOC-871	Freeman, where in the example blueprint file (Example of a login module with jasypt encryption) does this bean go? Assuming that is the place you were pointing to in the description.	0	0	https://issues.jboss.org/browse/FUSEDOC-871
5307	FUSEDOC-871	So what triggers the autoencryption event? When does it happen?	0	0	https://issues.jboss.org/browse/FUSEDOC-871
5308	FUSECDC-20	Hey Gert, What is meant by invalid profile? Are you referring to a missing profile or a profile with bad properties?	0	0	https://issues.jboss.org/browse/FUSECDC-20
5309	GTNMGMT-58	Scheduling for 2.1.0 as 2.0.0 is scheduled to be included into JPP 6.1 and it was more work then expected to move to the newer version of CRaSH. Curious what features/issues you are looking forward to in upgrading to newer version ?	1	0	https://issues.jboss.org/browse/GTNMGMT-58
5310	GTNMAVEN-8	I just kept the original assignment ;) Can you review the tasks you are assigned and don't plan to work on please ?	0	0	https://issues.jboss.org/browse/GTNMAVEN-8
5311	GTNPORTAL-3539	I find that this happens when we pass username with NULL value to UserHandler#findUserByName() method. It indirectly calls org.picketlink.idm.api.PersistenceManager.findUser(String) as below: {code} try { foundUser = session.getPersistenceManager().findUser(userName); } catch (Exception e) { handleException("Cannot obtain user: " + userName + "; ", e); } {code} The #findUser() method throws an exception, then the rollback of IDM transaction is performed inside the #handleException method (by indirectly calling #recoverFromIDMError method). That's why it reverts the new user was created before. This not only impacts to UserHandler#findUserByName() BUT also for every place that is using PersistenceManager#findUser() --> So the *workaround* is to remove the #handleException in the try-catch block. ========== However, what really concerns me is about the org.exoplatform.services.organization.idm.PicketLinkIDMOrganizationServiceImpl.flush() which is already called before PersistenceManager#findUser() performed. It looks like the pending changes should be saved at this time, then the transaction rollback should not be a problem. Is there any IDM expert who can help to confirm this ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-3539

#	Requirement	Questions	Answered	Successfully Answered	URL
5312	GTNPORTAL-3509	When importing automatically all applications, it only replaces slashes by underscores in Application Name and Display Name in [ApplicationRegistryServiceImpl.java|https://github.com/gatein/gatein-portal/blob/3.7.1.Final/component/application-registry/src/main/java/org/exoplatform/application/registry/impl/ApplicationRegistryServiceImpl.java#L382-L385] {code:title=ApplicationRegistryServiceImpl.java} // Need to sanitize portlet and application names in case they contain characters that would // cause an improper Application name portletApplicationName = portletApplicationName.replace('/', '_'); portletName = portletName.replace('/', '_'); {code} While adding an application to a category on UI, it must pass the rule of length limit and character validation. For example, in [UIApplicationForm.java|https://github.com/gatein/gatein-portal/blob/3.7.1.Final/portlet/exoadmin/src/main/java/org/exoplatform/applicationregistry/webui/component/UIApplicationForm.java#L58-L61] {code:title=UIApplicationForm.java} public UIApplicationForm() throws Exception { addUIFormInput(new UIFormStringInput("applicationName", "applicationName", null).addValidator(MandatoryValidator.class) .addValidator(StringLengthValidator.class, 3, 30).addValidator(NameValidator.class)); addUIFormInput(new UIFormStringInput("displayName", "displayName", null).addValidator(StringLengthValidator.class, 3, 30).addValidator(NotHTMLTagValidator.class)); addUIFormInput(new UIFormTextAreaInput("description", "description", null).addValidator(NullFieldValidator.class) .addValidator(NotHTMLTagValidator.class)); } {code}	1	1	https://issues.jboss.org/browse/GTNPORTAL-3509
5313	GTNPORTAL-3495	Yes, I know the default value is set to true. Sorry for the copy/paste in TestCase, but the first one attached in jbossexo ML didn't get the same page. Anyway, did you reproduce the problem using skipPaginationInMembershipQuery=false? (After test case change)	1	1	https://issues.jboss.org/browse/GTNPORTAL-3495
5314	GTNPORTAL-3495	Ok, I understand. So you intentionally didn't implement the membership pagination using LDAP store. (as seen in method LDAPIdentityStoreImpl.resolveRelationships). Can we make a change in GateIN so that in case we use LDAP/AD in PL IDM, the "skipPaginationInMembershipQuery" is switched automatically, or at least a "log.warn" to say: "Hey, you are using LDAP, switch this flag to 'true' ".	1	1	https://issues.jboss.org/browse/GTNPORTAL-3495
5315	GTNPORTAL-3397	What do you mean by "And SecureRandomService is configurable component" ? I see the service is being used in org.exoplatform.web.security.security.AbstractTokenService and OAuth providers implementation (such as facebook, google) The purpose of external configuration file gatein/conf/configuration.xml is to re-define/override default configuration. That means portal should still work w/o it.	1	0	https://issues.jboss.org/browse/GTNPORTAL-3397
5316	GTNPORTAL-3315	as COR-317 has been resolved, [~haint] could you please make a final check ?	1	1	https://issues.jboss.org/browse/GTNPORTAL-3315
5317	GTNPORTAL-3285	I am not sure what the last sentence of the previous comment means for the linked pull request. It means that the PR is unnecessary, right? Could you please comment and close the PR if so?	0	0	https://issues.jboss.org/browse/GTNPORTAL-3285
5318	GTNPORTAL-3283	Thanks for your comment, Marko. I vote for reporting this under the appropriate container project. Do you know which one is it?	1	0	https://issues.jboss.org/browse/GTNPORTAL-3283
5319	GTNPORTAL-3260	Please add (1) some string which uniquelly identifies the portlet, e.g. fq class name plus some kind of steps to reproduce (while doing what is this a problem?) (2) Why is this an issue? What is wrong now? What should be better after fixing? (3) One or two sentences how it could/should be done (if you know)	0	0	https://issues.jboss.org/browse/GTNPORTAL-3260
5320	GTNPORTAL-3145	Boleslaw Dawidowicz <bdawidow@redhat.com> made a comment on [bug 920553|https://bugzilla.redhat.com/show_bug.cgi?id=920553] Can you add documentation section about this too?	0	0	https://issues.jboss.org/browse/GTNPORTAL-3145
5321	GTNPORTAL-2852	We find that removing of *_en.properties and only keeping base ones will raise an issue: English translations would not be displayed anymore when specifically requested (by the browser or by the user setting) IF "en" is not the default locale of GateIn. The big issue by removing "_en" is that information is simply *lost*. The system doesn't know anymore that the base bundle file is English locale. Therefore we decided to keep *_en* bundle files only (do NOT generate base ones anymore) AND the *_en* locale will be taken as a last fallback locale by ResourceBundleService. So the overall sequence of candidate locales will be : - User locale - Portal locale - Base files - *_en* locale	0	0	https://issues.jboss.org/browse/GTNPORTAL-2852
5322	GTNPORTAL-2765	Hi Matt, I gave another try to the login page, I'll save it as an attachment in the next comment. Can you please review it? When it's fine and finished, I'll make a pull request to your mobile fork. The blue form header is now less detailed, if we want to make it the same as before, I can try to extract the pattern from original image. However, I'm not clear about some of the rules: 1, Any ideas what's the most sophisticated way how to do it? I've found this http://www.broken-links.com/2010/06/14/using-svg-in-backgrounds-with-png-fallback/, but I don't like that the bitmap is showed on the page anyway - covered with svg. 3, not really sure what this is for, if you have some jquery plugin/javascript lib/article discussing this, please share. For the record: * The registration portlet has ugly HTML markup, too. I should rewrite it after login screen is finished. * Nice presentation for everyone interested in responsive design: http://www.youtube.com/watch?v=H6z7vpgkCZw * Page for previewing the page in several sizes at once (works with localhost): http://mattkersley.com/responsive/	1	1	https://issues.jboss.org/browse/GTNPORTAL-2765
5323	GTNPORTAL-2765	The commit with my recent changes is located here: https://github.com/vrockai/gatein-portal-master/commit/a06f16a59a8d8fe99b94b65187da757f5e8209be I didn't test this on retina page. Do you use any emulator or tool for this purpose?	1	1	https://issues.jboss.org/browse/GTNPORTAL-2765
5324	GTNPORTAL-2765	What browsers are having troubles with svg for you? I'm testing with latest FF, latest Chrome and IE10.	1	1	https://issues.jboss.org/browse/GTNPORTAL-2765
5325	GTNPORTAL-2765	And does the key svg look smooth in all of them? or does it look pixelated in some?	1	1	https://issues.jboss.org/browse/GTNPORTAL-2765
5326	GTNPORTAL-2574	With basic JCR dependency upgrade testsuite is failing in component/test/jcr with: Caused by: java.lang.NoSuchFieldError: LUCENE_30 at org.exoplatform.services.jcr.impl.core.query.lucene.JcrStandartAnalyzer$1.run(JcrStandartAnalyzer.java:50) at org.exoplatform.services.jcr.impl.core.query.lucene.JcrStandartAnalyzer$1.run(JcrStandartAnalyzer.java:47) at org.exoplatform.commons.utils.SecurityHelper.doPrivilegedAction(SecurityHelper.java:290) at org.exoplatform.services.jcr.impl.core.query.lucene.JcrStandartAnalyzer.<init>(JcrStandartAnalyzer.java:46) at org.exoplatform.services.jcr.impl.core.query.lucene.SearchIndex.<init>(SearchIndex.java:559) ... 69 more	0	0	https://issues.jboss.org/browse/GTNPORTAL-2574
5327	GTNPORTAL-2574	With -DskipTests and packaged TC7 it complains about Caused by: java.lang.ClassNotFoundException: org.hibernate.service.ServiceRegistryBuilder at java.net.URLClassLoader$1.run(URLClassLoader.java:202) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at java.lang.ClassLoader.loadClass(ClassLoader.java:247) ... 44 more Is JCR 1.15 compatible with both Hibernate 3 and 4 or just the latter?	0	0	https://issues.jboss.org/browse/GTNPORTAL-2574

#	Requirement	Questions	Answered	Successfully Answered	URL
5328	GTNPORTAL-2574	{quote} Is JCR 1.15 compatible with both Hibernate 3 and 4 or just the latter? {quote} It is compatible with Hibernate 4 only	1	1	https://issues.jboss.org/browse/GTNPORTAL-2574
5329	GTNPORTAL-2533	Martin Weiler <mweiler@redhat.com> made a comment on [bug 846716|https://bugzilla.redhat.com/show_bug.cgi?id=846716] Description of problem: We have observed that when we close the "Edit Page" view (after we click on edit portlet (pencil icon) for a portlet but don't do any changes in the portlet's Edit mode dialog), we get the popup message - "Modifications have been made. Are you sure you want to close without saving?". Expected Result: The message "Modifications have been made. Are you sure you want to close without saving?" should only be shown if we made changes in any portlet (through its respective Edit Mode UI) or in the page containers.	0	0	https://issues.jboss.org/browse/GTNPORTAL-2533
5330	GTNPORTAL-2467	Peter Palaga <ppalaga@redhat.com> made a comment on [bug 830825|https://bugzilla.redhat.com/show_bug.cgi?id=830825] Esp. the .project and .classpath files make problems when importing the projects into Eclipse. The files need to be deleted and subsequently makked with svn:ignore.	0	0	https://issues.jboss.org/browse/GTNPORTAL-2467
5331	GTNPORTAL-2463	Hi, I cannot check the patch (access denied), could you attach it here also?	0	0	https://issues.jboss.org/browse/GTNPORTAL-2463
5332	GTNPORTAL-2406	Is this issue still valid?	1	0	https://issues.jboss.org/browse/GTNPORTAL-2406
5333	GTNPORTAL-2373	Kien is the patch commited in gatein? Could you confirm and resolve?	0	0	https://issues.jboss.org/browse/GTNPORTAL-2373
5334	GTNPORTAL-2373	printing "null" is an expected behavior from Groovy template engine world. I don't really like to change this default behavior (And I can see that JSP syntax does the same.) For me, the proper fix is to check non-NULL value before printing if we don't want to write "null". So I prefer to take "Won't Fix" for this. WDYT ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-2373
5335	GTNPORTAL-2358	I am adding new patch GTNPORTAL-2358-fullName.patch as previous patches can work only with "displayName" support in OrganizationService API on interface User, which is not option for now. New patch is using fullName as temporary persistent field, so fullName is placeholder for non-existing displayName. What exactly is done in patch is: - Adding new UI field with name "displayName" to all UI forms, where you can create or edit user (UIRegisterInputSet, UIAccountProfiles, UIAccountInputSet, UIAccountEditInputSet). So name of UI field is "displayName" and name of localization label is also "displayName". I added localizations only for "en" and "cs" properties bundles as I don't know other languages :) - UI components are communicating with User interface via "setFullName" and "getFullName". So fullName is now used as temporary placeholder until methods "getDisplayName" and "setDisplayName" will be available in Organization API. I added "TODO: GTNPORTAL-2358" in all pieces of code, which needs to be changed once displayName will be available in Organization API. This will allow us to use getDisplayName and setDisplayName and get rid of fullName as temporary placeholder. - IDM integration is changed to support persistent implementations of "getFullName" and "setFullName". For this purpose I forked class UserImpl from Organization API (I made a subclass of eXo UserImpl class with overriden methods getFullName and setFullName). Again, name of new attribute at DB level is "displayName", which will assure compatibility with future versions (after support for displayName in Organization API will be added).	1	1	https://issues.jboss.org/browse/GTNPORTAL-2358
5336	GTNPORTAL-2358	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 794235|https://bugzilla.redhat.com/show_bug.cgi?id=794235] Marek Posolda <mposolda@redhat.com> made a comment on jira GTNPORTAL-2358 I am adding new patch GTNPORTAL-2358-fullName.patch as previous patches can work only with "displayName" support in OrganizationService API on interface User, which is not option for now. New patch is using fullName as temporary persistent field, so fullName is placeholder for non-existing displayName. What exactly is done in patch is: - Adding new UI field with name "displayName" to all UI forms, where you can create or edit user (UIRegisterInputSet, UIAccountProfiles, UIAccountInputSet, UIAccountEditInputSet). So name of UI field is "displayName" and name of localization label is also "displayName". I added localizations only for "en" and "cs" properties bundles as I don't know other languages :) - UI components are communicating with User interface via "setFullName" and "getFullName". So fullName is now used as temporary placeholder until methods "getDisplayName" and "setDisplayName" will be available in Organization API. I added "TODO: GTNPORTAL-2358" in all pieces of code, which needs to be changed once displayName will be available in Organization API. This will allow us to use getDisplayName and setDisplayName and get rid of fullName as temporary placeholder. - IDM integration is changed to support persistent implementations of "getFullName" and "setFullName". For this purpose I forked class UserImpl from Organization API (I made a subclass of eXo UserImpl class with overriden methods getFullName and setFullName). Again, name of new attribute at DB level is "displayName", which will assure compatibility with future versions (after support for displayName in Organization API will be added).	0	0	https://issues.jboss.org/browse/GTNPORTAL-2358
5337	GTNPORTAL-2358	theute@redhat.com made a comment on [bug 794235|https://bugzilla.redhat.com/show_bug.cgi?id=794235] Nothing to add on the technical note, what should I do ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-2358
5338	GTNPORTAL-2073	Should we also do the same thing with other ones. For example UIFormStringInput or Table?	0	0	https://issues.jboss.org/browse/GTNPORTAL-2073
5339	GTNPORTAL-1939	Could you explain more why you tend to put some attributes to current PortalContainer?	1	1	https://issues.jboss.org/browse/GTNPORTAL-1939
5340	GTNPORTAL-1937	Have you tried using an empty producer configuration? {code} <producer-configuration/> {code}	1	1	https://issues.jboss.org/browse/GTNPORTAL-1937
5341	GTNPORTAL-1867	for Lucene 2.9.4 The 3.1 version is available. http://lucene.apache.org/ : "31 March 2011 - Lucene Core 3.1 and Solr 3.1 Available" would it be possible to have this version in the portal ? or it involves a major change, config and API changes, etc... and cannot be done for GateIn 3.2.0 ? Thanks, Antoine	1	0	https://issues.jboss.org/browse/GTNPORTAL-1867
5342	GTNPORTAL-1816	Dimitri, what do you want to do with this issue ?	1	0	https://issues.jboss.org/browse/GTNPORTAL-1816
5343	GTNPORTAL-1808	why i do not see any commit in GateIn trunk for this ? btw could you specify Fix Version/s for this issue ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-1808
5344	GTNPORTAL-1805	Can you verify if this patch is applied and resolve?	0	0	https://issues.jboss.org/browse/GTNPORTAL-1805
5345	GTNPORTAL-1805	Does WCI cover the REST API ({{/portal/rest/sso/authcallback/auth/_username_/_password_}})? I've looked at http://anonsvn.jboss.org/repos/gatein/components/wci/trunk/wci/src/main/doc/wci-authentication.pdf but it seems to me the answer is "no". So, in order to count calls to aforementioned REST URL where the return value is {{false}}, it seems to me {{UserDaoImpl.authenticate(username, password)}} would be the best place to register failed login attempts, much in the same way that {{User.setLastLoginTime(_now_)}} is called after succesful authentication. You could then expand {{org.exoplatform.services.organization.UserEventListener}} with the methods {{preAuthenticate(User user)}} and {{postAuthenticate(User user, boolean success}}, and trigger these at the start and end of {{authenticate}}. Then you could do any logic like sending mail or blocking the user in the listener, and developers would be able to easily extend this functionality.	0	0	https://issues.jboss.org/browse/GTNPORTAL-1805

#	Requirement	Questions	Answered	Successfully Answered	URL
5346	GTNPORTAL-1788	It seems like Chris' idea should work great. The portlets on all tabs need to be initialized, but the render method should only be invoked on the visible portlets. The non-visible portlets will then be able to process events and actions and serve resources, but avoid rendering which is obviously a waste of CPU time. This assumes the portlet developer is not changing any state in the render method, which is the best practice anyway. How difficult would it be to *not* call render on the non-visible portlets? That should be the behavior of any tabbed container.	1	0	https://issues.jboss.org/browse/GTNPORTAL-1788
5347	GTNPORTAL-1762	@Khoi Nguyen : Can you please fix this issue ASAP?	1	1	https://issues.jboss.org/browse/GTNPORTAL-1762
5348	GTNPORTAL-1746	Is this planned for release of GateIn 3.2-beta1 ?	1	1	https://issues.jboss.org/browse/GTNPORTAL-1746
5349	GTNPORTAL-1684	I am seeing change only in noncluster configuration but not in cluster profile configuration - http://anonsvn.jboss.org/repos/gatein/portal/trunk/web/portal/src/main/webapp/WEB-INF/conf/organization/picketlink-idm/jboss-cache-cluster.xml only in noncluster. Is it some issue why not add exp.algorithm here as well?	1	0	https://issues.jboss.org/browse/GTNPORTAL-1684
5350	GTNPORTAL-1662	could you give an usecase that needs such UI Component ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-1662
5351	GTNPORTAL-1570	I am testing with EPP 5.1 ER2 and I am seeing performance improvement after applying patch. More info in email. I am seeing that jira is now in state "Won't fix"? Is there any regression founded?	1	0	https://issues.jboss.org/browse/GTNPORTAL-1570
5352	GTNPORTAL-1539	In fact it breaks 1 test in the Shindig testsuite: <testcase time="0.006" classname="org.apache.shindig.gadgets.render.DefaultServiceFetcherTest" name="testReadConfigWithValidEndpoints"> <failure message="expected:<{http://localhost/api/rpc=[system.listMethods, service.get], gadgets.rpc=[messages.send, ui.resize], http://%host%/api/rpc=[do.something, delete.someting], http://%host%/socia l/api/rpc=[weather.get]}> but was:<{gadgets.rpc=[messages.send, ui.resize], http://localhost/api/rpc=[system.listMethods, service.get]}>" type="junit.framework.AssertionFailedError">junit.framework.Ass ertionFailedError: expected:<{http://localhost/api/rpc=[system.listMethods, service.get], gadgets.rpc=[messages.send, ui.resize], http://%host%/api/rpc=[do.something, delete.someting], http://%host%/social/a pi/rpc=[weather.get]}> but was:<{gadgets.rpc=[messages.send, ui.resize], http://localhost/api/rpc=[system.listMethods, service.get]}> at junit.framework.Assert.fail(Assert.java:47) at junit.framework.Assert.failNotEquals(Assert.java:277) at junit.framework.Assert.assertEquals(Assert.java:64) at junit.framework.Assert.assertEquals(Assert.java:71) at org.apache.shindig.gadgets.render.DefaultServiceFetcherTest.testReadConfigWithValidEndpoints(DefaultServiceFetcherTest.java:133) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at junit.framework.TestCase.runTest(TestCase.java:168) at junit.framework.TestCase.runBare(TestCase.java:134) at junit.framework.TestResult$1.protect(TestResult.java:110) at junit.framework.TestResult.runProtected(TestResult.java:128) at junit.framework.TestResult.run(TestResult.java:113) at junit.framework.TestCase.run(TestCase.java:124) at junit.framework.TestSuite.runTest(TestSuite.java:232) at junit.framework.TestSuite.run(TestSuite.java:227) at org.junit.internal.runners.JUnit38ClassRunner.run(JUnit38ClassRunner.java:79) at org.apache.maven.surefire.junit4.JUnit4TestSet.execute(JUnit4TestSet.java:62) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:140) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:127) at org.apache.maven.surefire.Surefire.run(Surefire.java:177) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.booter.SurefireBooter.runSuitesInProcess(SurefireBooter.java:345) at org.apache.maven.surefire.booter.SurefireBooter.main(SurefireBooter.java:1009) </failure> </testcase>	1	1	https://issues.jboss.org/browse/GTNPORTAL-1539
5353	GTNPORTAL-1539	I really am confusing on this. In my side, I can get service methods and then bug just appear because of parsing json data, so i have logs (see Description of bug), I guess it can be 'Network Proxy' that affect to retrieve osapi available services. In this bug, we only want to remove console logs because they are not the sign of a real problem. Could you please patch another 2010-10-18-GTNPORTAL-1539-shindig-update-comment-logs.patch. Thanks	1	1	https://issues.jboss.org/browse/GTNPORTAL-1539
5354	GTNPORTAL-1486	Could you make a patch or pack those files with directory structure? So assignee need not arrange the files in relevant folder	1	1	https://issues.jboss.org/browse/GTNPORTAL-1486
5355	GTNPORTAL-1468	What are missing ? can you clear ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-1468
5356	GTNPORTAL-1450	Is there any news please ?	1	1	https://issues.jboss.org/browse/GTNPORTAL-1450
5357	GTNPORTAL-1450	Is this patch validated? Can you please tell us when it will be integrated?	1	1	https://issues.jboss.org/browse/GTNPORTAL-1450
5358	GTNPORTAL-1397	Let see if the extension mechanism could help to do this ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-1397
5359	GTNPORTAL-1362	Marc, do you mean that this bug doesn't exist anymore in 5.2 or that you don't need to delete your DB and don't face it ? I'm not able to reproduce so far.	1	0	https://issues.jboss.org/browse/GTNPORTAL-1362
5360	GTNPORTAL-1297	@Prabhat Jha, could you provide a mockup to make this clear ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-1297
5361	GTNPORTAL-1290	spaces after comma is removed now	0	0	https://issues.jboss.org/browse/GTNPORTAL-1290
5362	GTNPORTAL-1134	Can this be made a priority as it will help with scalability?	0	0	https://issues.jboss.org/browse/GTNPORTAL-1134
5363	GTNPORTAL-1134	I will let you know at the beginning of next week. Are you fine with that?	1	1	https://issues.jboss.org/browse/GTNPORTAL-1134
5364	GTNPORTAL-1106	Actually this is a feature and it is still used. For instance, in WCM they configured two Portal templates are : Classic and ACME So we should revert this commit and configure a new basic portal template in GateIn default to show the feature. WDYT ?	1	1	https://issues.jboss.org/browse/GTNPORTAL-1106
5365	GTNPORTAL-1070	Check again images: 4.jpg, 3.jpg,11.jpg	1	1	https://issues.jboss.org/browse/GTNPORTAL-1070
5366	GTNPORTAL-1060	Why not using Application registry identifier ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-1060
5367	GTNPORTAL-1006	Why do we need that task ? I removed the notfound node before 3.0 GA release because it was not working correctly and there were many bugs.	0	0	https://issues.jboss.org/browse/GTNPORTAL-1006
5368	GTNPORTAL-992	Marek, did you eventually manage to get JOSSO work with 1.8.2-rev1573?	0	0	https://issues.jboss.org/browse/GTNPORTAL-992

#	Requirement	Questions	Answered	Successfully Answered	URL
5369	GTNPORTAL-918	For example, why is it possible to add a gadget to several categories at once directly from the gadget page but it's not possible to do so from a portlet page?	0	0	https://issues.jboss.org/browse/GTNPORTAL-918
5370	GTNPORTAL-916	Could you provide a mockup to show how the feature looks like in UI ? BTW, UI team could give some proposals for this too.	1	0	https://issues.jboss.org/browse/GTNPORTAL-916
5371	GTNPORTAL-888	And how/where exactly would these automatically imported portlets go? And why would we need to make an exception for the WSRP case and not for locally deployed portlets? There might be a way to get these portlets without them being imported into categories but it's not clear to me how to do it at this time. Maybe I just need more time to play with the Application Registry first. I can't help but feel that if it's not clear how to do this immediately, it won't be clear either for our users. :(	1	0	https://issues.jboss.org/browse/GTNPORTAL-888
5372	GTNPORTAL-870	on my opinion, the round-trip configuration have been asked and discussed a lot for JBoss Portal... and it is not really feasible. the trouble is : centralization. if you have several bags of rock stones, and you put all of them together in one bag... how do you separate them back in each of their origin bag ? especially if you add one or two stone directly in the big bag : in which original bag to you attribute it ? so... round-trip configuration could be done... but it would be a huge and complex tool... and it is not really useable then. on my opinion, the focus is : - configuration of a new portal and new portlets can be done in GateIn, even if a bit more complicated than in JBoss Portal : may be allow/add some more simple way. - globally maintain a focus and high priority on this way of making portal and portlet deployment. - hot deploy of these artifacts (cf jira GTNPORTAL-315) not sure there a lot to do in this jira itself.	1	1	https://issues.jboss.org/browse/GTNPORTAL-870
5373	GTNPORTAL-770	Note to ExoJBossCache configuration in common-configuration.xml. Now I am seeing "classpath:/conf/jbosscache/cluster/config.xml" as value of cache.config.template for both cases (cluster profile enabled and disabled). I guess it is typo and should be "classpath:/conf/jbosscache/local/config.xml" ? Maybe I am wrong since local configuration for jboss cache is empty and it has not specified cluster name. But then why is this specified with same value for both cases.	0	0	https://issues.jboss.org/browse/GTNPORTAL-770
5374	GTNPORTAL-741	- Now, with SNF_PRL_25,26,27 fixed and move to candidate directory (but have not right to commit) and I send file for Arthur - SNF_PRL_21: still fail with verifyTextNotPresent - SNF_PRL_29,30: I cann't run because have problem with gatein in Dashboard page and fixed by GTN_PORTAL- 768. I will return to check it tomorrow.	0	0	https://issues.jboss.org/browse/GTNPORTAL-741
5375	GTNPORTAL-721	Acknowledged. Major focus is screenshots, primarily re-XMLing to cater for different format renderings but also reshooting many. Also fixing some fairly catastrophic publican-build failures. Luc, where did you log the bugs you mentioned? (can't see anything in svn log)	0	0	https://issues.jboss.org/browse/GTNPORTAL-721
5376	GTNPORTAL-633	Fixed with attached patch file (2010-08-09-GTNPORTAL-633.patch) Solution: Don't display PageId field if create page, display PageId field (in Label type) if edit page	0	0	https://issues.jboss.org/browse/GTNPORTAL-633
5377	GTNPORTAL-509	Hi, can you please specify with which variables you are running the tests? I have following results on my side (only selenium folder tests): - localhost (JBoss): failures at tests 2, 7, 35 - hudson: - http://hudson.jboss.org/hudson/job/gatein_portal_ui_selenium_tests_jboss5/44/ (2,7,35) - http://hudson.jboss.org/hudson/job/gatein_portal_ui_selenium_tests_tomcat6/42/ (2,7,35) - http://hudson.jboss.org/hudson/job/gatein_portal_ui_selenium_tests_jboss5_db2/22/ (2,35) - http://hudson.jboss.org/hudson/job/gatein_portal_ui_selenium_tests_jboss5_mssql/19/ (33,35) - http://hudson.jboss.org/hudson/job/gatein_portal_ui_selenium_tests_jboss5_mysql/27/ (33,35) - http://hudson.jboss.org/hudson/job/gatein_portal_ui_selenium_tests_jboss5_postgresql/14/ (33,35) All tests were running on GateIn trunk version (rev. 2030). In hudson I run the tests with following command: mvn clean install -Pselenium -Dbrowser="*firefox" -Dselenium.port=6666 -Dselenium.timeout=100000 -Dselenium.speed=2000 Can you please verify my results? Thanks	1	0	https://issues.jboss.org/browse/GTNPORTAL-509
5378	GTNPORTAL-505	do you mean you want to overwrite the @Config template definition in UIComponent ?	1	1	https://issues.jboss.org/browse/GTNPORTAL-505
5379	GTNPORTAL-479	Should we apply that patch ?	1	1	https://issues.jboss.org/browse/GTNPORTAL-479
5380	GTNPORTAL-324	I did not see Miro in assignee list. Can you pls have him work on this? He should use the location you have for other UI tests.	0	0	https://issues.jboss.org/browse/GTNPORTAL-324
5381	GTNPORTAL-315	I don't understand why this is out of date ? Does the deployement (and re-deployement) of the artifacts of the "Extension Mechanism" can be hot ? You know, the feature of hot deployement of the portal definition (pages, navigation,...) that we had with JBoss Portal, before the reshape with Exo portal. This was a very important feature that is not anymore available. For development of portals in companies, getting this feature back would be a great dev productivity jump. Tug (of Exo) told me a few months ago that the feature was under thinking, even if not sure it could be done as full hot redeployment. What could be done is : - hot redeployement of the EAR, for update of the page and navigation definition (and locale resource bundle, etc...) - this is for an existing portal only, yet deployed What could not be done is : - not possible to hot redeploy to add or suppress a portal container or a portal (site) definition This would be great. The need is hot redeployement for Page/Nav development (with xml definition) and the tuning of all this. Creating a new portal container or portal (site) can be done with stop and restart of the server. This was Tug say,... if you want to check about what's on for this feature. ******** This feature is related with the update of the page/nav definition in the JCR, from the XML definition. As far as I know, this can be done only with suppression of the JCR content, and recreation,... for the time being. Did I missed something about this ? Thanks for all the dev of the product. Antoine www.sysemo.com	1	1	https://issues.jboss.org/browse/GTNPORTAL-315
5382	GTNPORTAL-307	How is that progressing ? I'm willing to help if needed (This is an important feature)	0	0	https://issues.jboss.org/browse/GTNPORTAL-307
5383	GTNPORTAL-94	Can we drop this one ? Sounds more like a WCM concern	1	1	https://issues.jboss.org/browse/GTNPORTAL-94
5384	GTNPORTAL-66	Thanks a lot for looking into this. I havent't followed too closely but have 2 questions: - Is there a workaround for the JDK6 bug ? Did you get Sun interest to fix it ? - Can this be used to precompile the groovy templates and reduce the first time access when unzipping and running GateIn ?	0	0	https://issues.jboss.org/browse/GTNPORTAL-66
5385	GTNPORTAL-66	- Is there a workaround for the JDK6 bug ? Did you get Sun interest to fix it ? I signalized the problem in the bugs resolution application of the sun and actually there is no response actually. Some day ago I've seen some interest to this in the sun forums (http://forums.sun.com/thread.jspa?threadID=5416155&tstart=0). I think in theese days I will know something more. Actually to resolve this problem you should delete the 'default' operator by the enums. An example of 'default' operator is inside the org.exoplatform.webui.event.Event.Phase class - Can this be used to precompile the groovy templates and reduce the first time access when unzipping and running GateIn ? It compiles the groovy scripts through the exo groovy compiler, so I suppose it builds new class files. I have verify if the achievement is the same as for the runtime compiling (verify where the class files are created)	0	0	https://issues.jboss.org/browse/GTNPORTAL-66
5386	GTNPORTAL-27	Is there someone on windows who can try to apply the attached patch on gatein portal trunk to see if fixes the issue ?	1	1	https://issues.jboss.org/browse/GTNPORTAL-27

#	Requirement	Questions	Answered	Successfully Answered	URL
5387	GTNSSO-34	[~juraci.costa] Can we close this "out of date"? Or should I assign to you? Ideally I just want GTNSSO JIRAs to be cleared from the list of my open jiras :)	1	1	https://issues.jboss.org/browse/GTNSSO-34
5388	GTNSSO-29	Could you resolve this issue as the PR is merged ?	0	0	https://issues.jboss.org/browse/GTNSSO-29
5389	GTNWCI-14	Normally you can't catch Errors. If an error is thrown it should just print the error message and stacktrace. Can you give an example of where an error gets swallowed with the Tomcat integration?	0	0	https://issues.jboss.org/browse/GTNWCI-14
5390	GTNWCI-11	Agreed with Matt, this should not be addressed by WCI layer. Is this javascript deployment performed by GateIn resource deployer ?	0	0	https://issues.jboss.org/browse/GTNWCI-11
5391	GTNWCI-11	I didn't see anything about WCI Generic mode and Native mode. Now i have to create a patch for wci-tomcat to make our products can run properly (at least with tomcat). Can you tell me more about these two modes ? How can i config it ?	0	0	https://issues.jboss.org/browse/GTNWCI-11
5392	GTNWCI-11	When i work with GateIn, i also got another problem, i raised a bug here GTNPORTAL-505, can you explain more for me ?	0	0	https://issues.jboss.org/browse/GTNWCI-11
5393	GTNWCI-11	We should try to depend the less we can on deployment ordering, can you describe the issue you have because of wrong deployment ordering ?	0	0	https://issues.jboss.org/browse/GTNWCI-11
5394	GTNWSRP-372	[~pmensik], would you be able to run a regression test with this commit, to check if other parts are affected by this change?	1	1	https://issues.jboss.org/browse/GTNWSRP-372
5395	GTNWSRP-340	theute@redhat.com made a comment on [bug 880729|https://bugzilla.redhat.com/show_bug.cgi?id=880729] Is it safe to put as default, or does it work only with a loadbalancer ?	0	0	https://issues.jboss.org/browse/GTNWSRP-340
5396	GTNWSRP-324	Could you please attach the response sent by WebSphere as a result of a getServiceDescription call from GateIn (in the WSRP admin UI, refreshing the consumer should trigger such a call)?	1	0	https://issues.jboss.org/browse/GTNWSRP-324
5397	GTNWSRP-324	Hi Ashwin, Thanks for the updates. Actually, could you please re-post server.log after activating TRACE level for FILE appender of your AS install (as explained at https://community.jboss.org/wiki/JBossAS7Logging)? Could you also please attach the message requests and responses for initCookie and getMarkup as well (with any potential cookies) if possible? Thanks in advance.	1	0	https://issues.jboss.org/browse/GTNWSRP-324
5398	GTNWSRP-324	Would it be possible for you to test a new version of WSRP to see if it helps with your issue?	1	0	https://issues.jboss.org/browse/GTNWSRP-324
5399	GTNWSRP-295	Honza Fnukal <hfnukal@redhat.com> made a comment on [bug 835375|https://bugzilla.redhat.com/show_bug.cgi?id=835375] Technical note updated. If any revisions are required, please edit the "Technical Notes" field accordingly. All revisions will be proofread by the Engineering Content Services team. Diffed Contents: @@ -1,9 +1,15 @@ -Is a RN required for this issue? If so, please write CCFR. - CAUSE: + * WSRP: getPortletProperties doesn't behave properly when en empty set of property names is passed + * WSRP type factories require a lang value to create a Property even though it's optional according to XSD +LONG DESCRIPTION: + * Using WSRP v1, the getPortletProperties method fails to return the list of portlet properties if the request contains an empty set of property names. If the names element contains the 'nil' attribute with value 'true', the request fails with a NullPointerException. + * Using WSRP v1, the setPortletProperties method fails if the 'properties' element does not contain the 'lang' attribute. The request gets rejected by the server with an IllegalArgumentException: 'Property requires a non-null, non-empty language'. According to the WSRP specification, the 'lang' attribute is optional. CONSEQUENCE: +The request gets rejected by the server with an IllegalArgumentException: 'Property requires a non-null, non-empty language' FIX: +Patch applied -RESULT:+RESULT: +Issue fixed, no exception is thrown	0	0	https://issues.jboss.org/browse/GTNWSRP-295
5400	GTNWSRP-285	Martin Weiler <mweiler@redhat.com> made a comment on [bug 824032|https://bugzilla.redhat.com/show_bug.cgi?id=824032] Closing as fixed in JPP 6.0.0, as GTNWSRP-285 has been fixed in WSRP 2.2.0.CR01. JPP 6.0.0 includes GateIn WSRP 2.2.2.Final-redhat-1.	0	0	https://issues.jboss.org/browse/GTNWSRP-285
5401	GTNWSRP-258	Hi Chris, I have tried with 2 cluster nodes (on different machines) + loadbalancer + consumer on other machine. I'm still getting same behavior - create consumer (with wsdl of loadbalancer), registration is done with node1, add portlet on page - ok, kill node1, refresh - portlet not displayed, refresh of consumer -> could not transmit message. I'm using EPP 5.2.2 ER03 with WSRP 2.1.7 included, is this fixed at following version of WSRP (2.2.0)? Thanks for information	1	0	https://issues.jboss.org/browse/GTNWSRP-258
5402	GTNWSRP-225	I'm getting another error when accessing page with remote portlet after failover. java.io.IOException: Could not transmit message (Connection refused) And when I try to refresh the consumer, I'm getting the same error (Problem getting service description for producer cluster, please see the logs for more information. Could not transmit message) My environment: - 2 EPP 5.1.1 instances in cluster (lab) - 1 loadbalancer (lab) - 1 EPP 5.1.1 local instance (using WSDL from loadbalancer to register consumer)	1	1	https://issues.jboss.org/browse/GTNWSRP-225
5403	GTNWSRP-219	How do you produce this error? The wsrp admin doesn't exactly work right now with AS6	1	1	https://issues.jboss.org/browse/GTNWSRP-219
5404	GTNWSRP-218	What do you mean by: "did not work"? Were you not able to run the app or did you it exhibit the issue? I'm not interested in the behavior in AS 7 at the moment, it needs to run in AS 6. Regarding the fact that the EL impl seems to be correct, I tend to agree with you since I doubt that an issue as big as this one would have gone undetected until now. This might have something to do with the PortletBridge... :/	1	1	https://issues.jboss.org/browse/GTNWSRP-218
5405	GTNWSRP-218	at org.gatein.test.ELBug.isB(ELBug.java:17) [classes:] seems to imply that isB is evaluated, which is definitely wrong. Am I missing something?	1	1	https://issues.jboss.org/browse/GTNWSRP-218
5406	GTNWSRP-192	Is this something GateIn specific, though?	0	0	https://issues.jboss.org/browse/GTNWSRP-192
5407	GTNWSRP-181	Are there steps to have GateIn work on AS 6 somewhere? It's failing way before WSRP gets a chance to fail here :)	1	1	https://issues.jboss.org/browse/GTNWSRP-181
5408	GTNWSRP-168	Do you mean display the URL at which the Producer publishes its WSDL? If yes, it might not be as easy as first thought since I am not sure if there's a standard way to retrieve that information from JAX-WS and I would rather not code something specific to a particular WS stack.	1	1	https://issues.jboss.org/browse/GTNWSRP-168
5409	GTNWSRP-69	Do we still need this? We do use an in-memory ExportPersistenceManager as part of the tests, and I don't see how an in-memory only implementation would do much good outside of the tests.	0	0	https://issues.jboss.org/browse/GTNWSRP-69
5410	GTNWSRP-67	Should now work properly. No need to deploy external jars anywhere. Added bonus NetUnity event test portlets should also work! :)	0	0	https://issues.jboss.org/browse/GTNWSRP-67
5411	GRAVEL-31	Would you ever want to read a document from someplace other than either the servlet root or the classpath? Which is to say, why a URL?	1	0	https://issues.jboss.org/browse/GRAVEL-31
5412	GRAVEL-28	I don't see this as a good idea - a given component can currently reasonably expect all of its children to be rendered within the same page as the component itself. By breaking this guarantee, I can only picture bad things happening - labels that refer to IDs that exist within the component tree, but not within the HTML document, for example. When would you use such a thing? Perhaps it would be better to provide a means to render a complete Facelet into a buffer instead. Can you give your use case(s)?	1	1	https://issues.jboss.org/browse/GRAVEL-28
5413	GRAVEL-28	The first use case can be done with existing facelets templating facilities pretty easily. I'm not interested in the second use case - a section of HTML cannot be assumed to be valid without its enclosing context. Sure it can be if you're careful, but don't want to go there. Any others?	1	1	https://issues.jboss.org/browse/GRAVEL-28

#	Requirement	Questions	Answered	Successfully Answered	URL
5414	GRAVEL-28	Facelets templating will cause evaluating same EL expressions many times in first case. I'm not sure why you are discarding second use case: Let's say you need to render a big table (which takes long time using JSF "as is") on a screen and send it via email using Seam email tag on same click. In that case you just could render it into buffer stored in backing bean and use it in 2 places. How could you do it better? ;) Thanks	0	0	https://issues.jboss.org/browse/GRAVEL-28
5415	GRAVEL-27	Can you give an example of a case where evaluating an expression at render time is preferable? I don't see any case where such a thing would be relevant for performance improvements.	0	0	https://issues.jboss.org/browse/GRAVEL-27
5416	GRAVEL-24	Perhaps this patch could help. It will add a targetFileName and targetContentType attributes to gd:inputFile. Also, it would be good to have a targetFileSize but I haven't find an easy way of doing it. I also have problems with tld documentation. I have added tld annotations. Do I need something else to add this attributes to tlds?	1	0	https://issues.jboss.org/browse/GRAVEL-24
5417	GRAVEL-18	Why is this not documented in TLD doc on gravel web site? Please, update the docs on web site to the docs from the latest release	1	1	https://issues.jboss.org/browse/GRAVEL-18
5418	GUVNOR-3305	[~manstis], I'm surprised that marshalling doesn't work with List<Object>. Do you have a simple example of it breaking I could look at?	1	1	https://issues.jboss.org/browse/GUVNOR-3305
5419	GUVNOR-3305	Hi, thank you for these investigations! The role of this extension is to allow end user to put additional data into their DMN model as needed. Mostly is for meta-data of the business. A pragmatical example: I might embed as byte[] in one of those extension, in order to attach BLOB of Words or PDF to any DMN node. However is not only that, because for instance other partner editor tooling of DMN diagrams, do use this extension mechanism to extend the DMN specification with additional execution semantics. There is a basic example in our test suite: https://github.com/kiegroup/drools/blob/master/kie-dmn/kie-dmn-core/src/test/java/org/kie/dmn/core/compiler/extensions/DMNExtensionRegisterTest.java#L121-L125 and related resource: https://github.com/kiegroup/drools/blob/master/kie-dmn/kie-dmn-core/src/test/resources/org/kie/dmn/core/compiler/extensions/0001-input-data-string-with-extensions.dmn#L36-L38 At this point I have a question, please? How can I make sure when I write an extension element, like for example the following: https://github.com/kiegroup/drools/blob/master/kie-dmn/kie-dmn-core/src/test/java/org/kie/dmn/core/compiler/extensions/FirstNameDescription.java that is portable?	1	1	https://issues.jboss.org/browse/GUVNOR-3305
5420	GUVNOR-3290	Hey [~manstis] What you mean by a "top-level" item? The palette is displayed as for a given palette model, this palette model actually provides 3 levels: categories, morph base types and finally the items. Actually you can drag a category or any of the items in the palette, so it means you can drag from the "big" palette icons and from each of the stencil icons as well. Saying that, what's a "top-level" item? Not sure if I'm following you sorry...	1	1	https://issues.jboss.org/browse/GUVNOR-3290
5421	GUVNOR-3290	Ah ok I think I undestand [~manstis]. Well first of all say that actually you can drag from the top-level/categories items into the canvas, so this part seems that already fits you needs. On the otrher hand, as commetned, the displayed palette is built from a palette model, so you can try to use a custom PaletteDefinitionBuilder class, that instead of generating items for each of your stencils, it skips that part and only generate a palette model that contain "categories", so the top-level menu items. Tbh I never tried this, not sure what happens if you provide a palette model only with categories, it should work... but not sure if the floating div will still appear even if empty. So that's what you can probably do... create a custom palette definition model builder, that builds a model which contains your categories (top-level) items, but do not have any children (no morph base types, no items at all). Doing that, you should be able to drag from the top-level icons, as it's actually supported, but probably you'll have to fix some bug on t he palette rendering, as if no children, no floating div should appear, it only has to display top-level items. Does that makes sense? Just brainstorming, feel free to change whatever you need.	1	1	https://issues.jboss.org/browse/GUVNOR-3290
5422	GUVNOR-3290	Oh ok I see, well this happens for me with, for example, my IDE, I have to change some of the default OS keyboard shortcuts to make the ones from the IDE work. So not sure if that's a bug in Stunner to be reported and fixed, or it's just about changing your OS keyboard shortcuts settings? WDYT? Did you try with Shift too, does the window manager intecepts it too? :(	1	1	https://issues.jboss.org/browse/GUVNOR-3290
5423	GUVNOR-3265	p.s. Google Maps has an on-demand slider as well, kinda nifty. Have folks discussed using a slider?	1	0	https://issues.jboss.org/browse/GUVNOR-3265
5424	GUVNOR-3265	[~roger600] * I reached out to another team who has worked on canvases. It looks like they are using a button control. They have posted some work on this topic to PF that isn't shown on patternfly.org, yet, but it's available as an Angular Patternfly example: http://www.patternfly.org/angular-patternfly/#/api/patternfly.canvas.directive:pfCanvas. * I also talked to [~mstevens1] who has some feedback that he'll be posting. * I couldn't find the UI controls in anything other than the GDT, so I'm at a little bit of a loss. If there's any pics you could post that would be super. * Would it be possible to put the zoom controls in or near the canvas? They seem a little hidden, or I'm not finding them. ;) In GDT, I'm not sure placing them within the menu item "View" is ideal. View> Zoom seems like a difference action, than View > Audit Log, so they might not even belong in the same menu. So pulling the control out of that menu and placing it on the canvas (or the canvas header) might help the user. [~mstevens1] Wdyt?	1	0	https://issues.jboss.org/browse/GUVNOR-3265
5425	GUVNOR-3264	[~roger600] Hi Roger, how does "pan" work in Stunner? I see {{PanControl}} however pressing {{alt}} and moving the mouse around appears to have no effect.	1	1	https://issues.jboss.org/browse/GUVNOR-3264
5426	GUVNOR-3264	[~roger600] I talked with [~mstevens1] about this and we're both aligned on these points: * Canvas should be infinite. It should grow when I reach the edges. * Users can use scrollbars to pan. * Users can use the little "preview" window to pan. * [Matthew] As a bonus, users could use keyboard. For this, I am accustomed to the spacebar. Though it looks like GenMyModel uses Shift. A side question that I have: Why does the panned selection get a little gray border after it's dropped? I get why it would be there during the panning, but not sure why it should persist?	1	0	https://issues.jboss.org/browse/GUVNOR-3264
5427	GUVNOR-3264	Hey [~uxdlc] [~mstevens1], Thanks for the feedback! I'm agree on the above points, just a few comments: * *Canvas should be infinite. It should grow when I reach the edges* -> Still not supported. [There already exists a task for this|https://issues.jboss.org/browse/JBPM-5536]. * *Users can use the little "preview" window to pan* -> That's a good improvement, I think it's neither identified as a task yet, so does not exist a ticket for this. Actually the preview and the main canvas are "decoupled", you can pan/zoom in both panels but have no effect between each other. Would be nice when doing zoom/pan in the preview panel, that it would reflect in the main one. For example being able to select a region in the preview panel, and that region gets automatically zoomed/pannel into screen's center for the main panel. Does it makes sense?	0	0	https://issues.jboss.org/browse/GUVNOR-3264

#	Requirement	Questions	Answered	Successfully Answered	URL
5428	GUVNOR-3222	Hi [~roger600] working on this now :-) I want most DMN Elements to show the "Name" text however one in particular ("Text Annotation") to show the "text" property of the definition and not "name". This seems quite easy to achieve by re-using most of {{CanvasNameEditionControlImpl}} and {{NameEditBoxWidget}} to read from and write to the correct property on the definition. I'd need to expose this new "Text" editor as a _new_ handler on {{CanvasControlRegistrationHandler}} and have it only register itself for the applicable types of Shape. Sound about right to you? Or.... I might make {{CanvasNameEditionControlImpl}} and {{NameEditBoxWidget}} support a configurable ({{String}}) property to edit and not tied to "Name". WDYT?	1	1	https://issues.jboss.org/browse/GUVNOR-3222
5429	GUVNOR-3222	[~roger600] Oh, as an aside... do you have a JIRA already for when the text to edit is an empty {{String}}? For example, if I edit the name of a node from "Subprocess" (or whatever) to simply "" then I can no longer edit the text "in place". Presumably it's because the event handlers for double-click etc are attached to a Lienzo {{Text}} Shape and when this represents an empty {{String}} there is nothing in the selection layer to respond to mouse events...	1	1	https://issues.jboss.org/browse/GUVNOR-3222
5430	GUVNOR-3222	Hey [~manstis], my comments :) * Yeah for me sounds good making the CanvasNameEditionControlImpl more generic, so as you say, for example, by adding the name of the property to edit (probably you can even change the class name then, it's not just for "name"...). So having the property name to editr, you can use the adapters (DefinitionAdapter etc) to obtain and to set the value for that property, * On the other hand, not sure if using that PropertyMetaTypes or whatever you think could work, but you need to know on the client side, for a given bean, which is that property to edit. So basically to figure out, in some way, when you have a bean, to obtain the name of the property you want to edit. * About the bug on empty text - good point! I don't remember any existing bug for this, so can you please open it if so? You can even put it on top of backlog, if you need this, this way we'll see it on next Monday's meeting and will take into account for fixing it sooner than latter :) And yeah, I think it's due to what you say, the handler is added to the text instance, so we should ensure the text, although empty, is being drawn in the selection layer with some size, sounds a bug to fix in lienzo side. Thanks for the feedback!	1	1	https://issues.jboss.org/browse/GUVNOR-3222
5431	GUVNOR-3222	Hi [~roger600] I've made the property used as the "display text" for nodes in the diagram and edited by the "in-place popup" configurable on a [branch|https://github.com/manstis/kie-wb-common/commit/710f9729694da97fbdfc599098574066e2cf69d6] in my fork. I have an adaptor factory that is used by {{BaseCanvasHandler}} and the "in-place popup" that returns read/write adaptors for different {{Elements}}. The default (and fallback) implementation of an adaptor uses the old code that was tied to the {{NAME}} meta-property. I have a custom implementation for DMN that reads/writes the {{String}} from a different field on the definition. What I am not sure about however is how to update a property. The {{propertyId}} I use is statically [defined|https://github.com/manstis/kie-wb-common/commit/710f9729694da97fbdfc599098574066e2cf69d6#diff-d03fe0c68d6601e5d56d002255e537f8R88]. Is this OK or should I be looking it up from one of Stunner's adaptors somewhere!?!	1	1	https://issues.jboss.org/browse/GUVNOR-3222
5432	GUVNOR-3222	See https://github.com/manstis/kie-wb-common/commit/1de061f1841a2317e8d9f0c335bb98dc2730efd3 This commit introduces a mechanism to lookup a "provider" of text to be used in the following by {{Element}} type:- * As the caption for nodes in the diagram * As the text edited in the "in-place" caption editor * As the text shown in Explorer The default (and fall-back, should no specific provider be found) falls back to the {{NAME}} MetaProperty.	1	1	https://issues.jboss.org/browse/GUVNOR-3222
5433	GUVNOR-3221	Hey [~manstis] Not sure if current features are enough to achieve your needs, I think so.. but lemme just share a few comments about how this SVG stuff works: * Therre is no resize support yet for rounded borders or for shapes based on canvas arcs. We're working on [this issue|https://github.com/ahome-it/lienzo-core/issues/140] in lienzo to natively support it. But due to this limitations stunner did some work-arounds.. * Due to the above limitation, and due to a WiresShape is composed by a MultiPath, currently Stunner a concrete [WiresScalableContainer|https://github.com/kiegroup/kie-wb-common/blob/master/kie-wb-common-stunner/kie-wb-common-stunner-client/kie-wb-common-stunner-lienzo/src/main/java/org/kie/workbench/common/stunner/client/lienzo/shape/view/wires/WiresScalableContainer.java] implementation, in case you create the SVG shape views by setting resizable = true, see [this flag here|https://github.com/kiegroup/kie-wb-common/blob/master/kie-wb-common-stunner/kie-wb-common-stunner-sets/kie-wb-common-stunner-bpmn/kie-wb-common-stunner-bpmn-client/src/main/java/org/kie/workbench/common/stunner/bpmn/client/resources/BPMNSVGViewFactory.java#L48]. * The WiresScalableContainer just check if the lienzo shape instance supports width and height attributes, if so, it sets the new ones (that makes the shape rendering size fine), otherwise, it just scales the instance. See [this logic here|https://github.com/kiegroup/kie-wb-common/blob/master/kie-wb-common-stunner/kie-wb-common-stunner-client/kie-wb-common-stunner-lienzo/src/main/java/org/kie/workbench/common/stunner/client/lienzo/shape/view/wires/WiresScalableContainer.java#L49]. * In case the SVG declaration contains more than one element, it's possible to specify a layout position and make the element scale as well on resize. This is done by adding a custom attribute on the SVG declaration, you can find here an example for [custom layout declaration|https://github.com/kiegroup/kie-wb-common/blob/master/kie-wb-common-stunner/kie-wb-common-stunner-sets/kie-wb-common-stunner-bpmn/kie-wb-common-stunner-bpmn-client/src/main/resources/org/kie/workbench/common/stunner/bpmn/client/resources/images/event/event-start.svg?short_path=3e4906c#L21] and for [an scalable element|https://github.com/kiegroup/kie-wb-common/blob/master/kie-wb-common-stunner/kie-wb-common-stunner-sets/kie-wb-common-stunner-bpmn/kie-wb-common-stunner-bpmn-client/src/main/resources/org/kie/workbench/common/stunner/bpmn/client/resources/images/event/event-intermediate.svg?short_path=aded512#L75]. So after all that comments.. hehe , let's focus on your issue - I understand you're using an svg path for the InputData svg file, which path displays rounded corners? Is this svg path element the first one is your svg file (if it is, it's considered as the shape's constructor multipath argyment)? In that case, I think that providing the right svg path, which renders the rounded corners using the right size as from the whole svg , and you create the SVGShapeView instance from by setting the resizable argument as true (in SVGViewFactory) -> the multipath will just scale as the user resizes the shape, so it will keep "ratio", isnt it? (This is because MultiPath does not contain width/height attributes, so WiresScalableContainer just scales it).	1	1	https://issues.jboss.org/browse/GUVNOR-3221

#	Requirement	Questions	Answered	Successfully Answered	URL
5434	GUVNOR-3221	Yeahhh [~manstis] good point!! That was why I added the WiresScalableContainer, to be able to do NOT scale, where possible, the shapes once resizing, because the borders were scaled too! So I moved to use of rectangles, circles etc for the main svg shape and added support for resizing those using attributes, rather than scaling. But did not realized about these kind of details once resizing, like updating corner sizes etc.... Well I'm just thinking on a quick solution for this, not sure if too overkilling, but what you think about adding support for the "CORNER_RADIUS" in addition to the actual WIDTH and HEIGHT, once scaling the svg shape? Look [here|https://github.com/kiegroup/kie-wb-common/blob/master/kie-wb-common-stunner/kie-wb-common-stunner-client/kie-wb-common-stunner-lienzo/src/main/java/org/kie/workbench/common/stunner/client/lienzo/shape/view/wires/WiresScalableContainer.java#L55]. So even not specifying a custom ratio update for this CORNER_RADIUS attribute, probably updating it as using same ratio as for WIDTH and HEIGHT works for you, because then you can move back to the use of the SVG rectangle, instead of the path, and the border will be not scaled. WDYT?	1	1	https://issues.jboss.org/browse/GUVNOR-3221
5435	GUVNOR-3212	Do we need such feature? Should Advanced users have a possibility to edit the pom in empty project? There might be a case when the users would like to edit the pom. [~paulovmr] [~manstis] What is your opinion to this?	1	0	https://issues.jboss.org/browse/GUVNOR-3212
5436	GUVNOR-3178	Does this mean we will have just two different WARs - worbench + controller, even for non-cloud deployments? Or will we have three WARs: the current one (combining both workbench and controller) and then the additional ones?	1	0	https://issues.jboss.org/browse/GUVNOR-3178
5437	GUVNOR-3168	Hi [~tirelli] and [~manstis]. This task seems to me as nice example of too big task for one sprint if we assume coding, testing, PR review. I would like to ask you to split this to smaller parts. My suggestion would be: - Write unmarshaller for Data Inputs - Write unmarshaller for Decisions - Write unmarshaller for Knowledge Sources - Write unmarshaller for Business Knowledge Models - Write unmarshaller for Connections - Write unmarshaller for Annotations? (Not sure if present) What do you think? Missed I some elements? Or should we split more for particular properties of given nodes?	1	0	https://issues.jboss.org/browse/GUVNOR-3168
5438	GUVNOR-3167	Hi [~tirelli] and [~manstis]. This task seems to me as nice example of too big task for one sprint if we assume coding, testing, PR review. I would like to ask you to split this to smaller parts. My suggestion would be: - Write marshaller for Data Inputs - Write marshaller for Decisions - Write marshaller for Knowledge Sources - Write marshaller for Business Knowledge Models - Write marshaller for Connections - Write marshaller for Annotations? (Not sure if present) What do you think? Missed I some elements? Or should we split more for particular properties of given nodes?	1	1	https://issues.jboss.org/browse/GUVNOR-3167
5439	GUVNOR-3162	Just to clarify: should keep the simple properties based authentication (like we have now)? Or should we completely move to Keycloak as default?	1	0	https://issues.jboss.org/browse/GUVNOR-3162
5440	GUVNOR-3105	Hi [~jomarko] what browser was this with?	1	1	https://issues.jboss.org/browse/GUVNOR-3105
5441	GUVNOR-3105	Thanks [~jomarko] Could you please advise whether the V&V panel state changes where you say shortcuts stop working? e.g. a rule becomes invalid/valid etc Keyboard operations require the canvas to have the "focus". I suspect when performing some operations the V&V panel updates causing focus to be lost.	1	1	https://issues.jboss.org/browse/GUVNOR-3105
5442	GUVNOR-3029	[~porcelli] Could you please provide more details? Will each project have such configuration? Will it be possible to edit this configuration in the UI, or will it be provided automatically and opaquely for every workbench project? What will be configuration possibilities available for use? Any Idea who will be working on this feature?	0	0	https://issues.jboss.org/browse/GUVNOR-3029
5443	GUVNOR-3028	[~porcelli] Could you please clarify the note regarding HA not being covered? What are the supported scenarios and what are the not supported? What's the impact on scalability?	1	1	https://issues.jboss.org/browse/GUVNOR-3028
5444	GUVNOR-3025	[~dward], [~porcelli] Is this somehow related to issues in this epic BPMSPL-656 ? Which of these new images do we plan to use in Guvnor ALA?	1	0	https://issues.jboss.org/browse/GUVNOR-3025
5445	GUVNOR-3025	[~porcelli] Thank you for your answer. It is good idea to have images also for community. Is there Jira(s) for these image(s)? Maybe it would be good to not duplicate effort and coordinate work with [~KrisVerlaenen] on OpenShift images. It seems that we will have at least 10 images produced by Kris. See subtasks for BPMSPL-656. In my opinion it will be hard to test and understand to this huge number of images. It will be even harder if we will have separate images produced specially for Guvnor ALA.	0	0	https://issues.jboss.org/browse/GUVNOR-3025
5446	GUVNOR-3023	Hello [~dward] I'm working on test automation for Guvnor ALA. Can you be more specific about kie-server provisioning? On which application server do you expect it will run (EAP, JWS, other platforms)? Where it will be possible to provision kie-server using Guvnor ALA (docker, OpenShift, Wildfly, EAP, others)? Can you describe basic scenario how Kie-Server will be used with Guvnor ALA? Will it be possible to somehow deploy kjar or kwar into this kie server and how?	1	1	https://issues.jboss.org/browse/GUVNOR-3023
5447	GUVNOR-3022	[~porcelli] Hello, do we have PL Jira or document about "kwar"? What is it and how to use it? What is the difference between kjar and kwar? Will be execution server embedded into kwar? Maybe this Jira is related to GUVNOR-3023. Do we expect to embed Kie server into kwar and then deploy it into app. server using Guvnor ALA. Or will we deploy kie server on app. server and then deploy kjars and kwars into it using Guvnor ALA?	1	0	https://issues.jboss.org/browse/GUVNOR-3022
5448	GUVNOR-3021	[~wmedvede] Hello Walter, I'm working on test automation for Guvnor ALA. Can you please be more specific about GUI. Will it be part of workbench or do you plan to create new standalone GUI for Guvnor ALA? What features will this GUI provide? Will it be possible to model and run application life cycle in this GUI? Is this Jira somehow related to GUVNOR-3028 or GUVNOR-2933?	1	0	https://issues.jboss.org/browse/GUVNOR-3021
5449	GUVNOR-2933	[~wmedvede] Hello Walter, which platforms do we plan to support? Currently Guvnor ALA is able to work with Docker and Wildfly. Which additional platforms do you plan to use?	0	0	https://issues.jboss.org/browse/GUVNOR-2933
5450	GUVNOR-2876	Team, Do you have any plans for this ?	1	0	https://issues.jboss.org/browse/GUVNOR-2876

#	Requirement	Questions	Answered	Successfully Answered	URL
5451	GUVNOR-2854	Hi [~Rikkola]. I played a little bit with guided decision tables using Resolved HP. I found out one use case I would like to discuss with yoy: # Lets say you have guided decision table like this: ||RC||Description||Has priority over|| |1|a| | |2|b|1| |3|c|1| # Now swap the first and second row and you will obtain ||RC||Description||Has priority over|| |1|b|1| |2|a| | |3|c|2| # However, if you save and reopen the table, you will obtain the expected: ||RC||Description||Has priority over|| |1|b| | |2|a| | |3|c|2| Do you think could be possible to achieve this without reopening the table? I had look also into the source code. For me are suspicious these two lines of code: [1] I think there shouldn't be *-1* because values stored in Number and Priority columns have no offset. Both assumes number 1 as the first row of the table. [2] I think we should compare also for equality like *oldValue >= rowNumber* because if the row X had priority over row Y (Y < X) but after moving X <= Y it is no needed to have in table that row X has priority over row Y because row Y is now under row X. I also admit, that my thoughts are wrong and the problem is on another place. Please tell me what do you think. [1] https://github.com/droolsjbpm/drools-wb/blob/master/drools-wb-screens/drools-wb-guided-dtable-editor/drools-wb-guided-dtable-editor-client/src/main/java/org/drools/workbench/screens/guided/dtable/client/widget/table/model/synchronizers/impl/PrioritySynchronizer.java#L67 [2] https://github.com/droolsjbpm/drools-wb/blob/master/drools-wb-screens/drools-wb-guided-dtable-editor/drools-wb-guided-dtable-editor-client/src/main/java/org/drools/workbench/screens/guided/dtable/client/widget/table/model/synchronizers/impl/PrioritySynchronizer.java#L71	1	0	https://issues.jboss.org/browse/GUVNOR-2854
5452	GUVNOR-2854	Hi [~Rikkola]. I was increasing the coverage for moving rows and I think that I managed to fix the issue form the comment above. Here is the PR: https://github.com/kiegroup/drools-wb/pull/425. Please could you have a look when you have a moment?	0	0	https://issues.jboss.org/browse/GUVNOR-2854
5453	GUVNOR-2832	Hi [~karreiro]. Few days ago, I discussed this a little bit with [~manstis] and he found out that there should be introduced two cases. The one that I am describing. It means that user is not interested in operator and the corresponding cells in the table should offer just check boxes. And the second case when user can define the operator directly in the table cell. [~manstis] proposed to name this like "per cell". He told me you are rewriting column definitions of dtable. Please could you incorporate these changes? [~manstis] please could you confirm I remember correctly our thoughts?	1	0	https://issues.jboss.org/browse/GUVNOR-2832
5454	GUVNOR-2832	[~jomarko] In this scenario we'd normally recommend defining the column as {{Player($s : score != null)}} and use the "is not null" operator. However as discussed on IRC there *may be* occasions when the User is also interested in {{null}} values. [~karreiro] Jozef and I discussed that you are refactoring the column definition *screens* ([~jomarko] not the column definitions themselves). My proposal was that we replace the "(no operator)" operator with "(per cell)" as this is a better indication of where the operator is expected. "No operator" would give what [~jomarko] is requesting BUT the "operator" in our model is a {{String}} where an empty of null value currently represents "(per cell)" so we do not have a way to specify "no operator" without a "magic string" (e.g. we store operator as "_none" and use this magic value to mean, erm, no operator. Personally, this type of JIRA just adds more support to removing all the "simple" columns all together and only supporting "BRL fragments" (as the internal representation; we still may have the Wizard showing "New field constraint" etc type stuff; but internally we build a "BRL Condition fragment" rather than a {{ConditionCol52}}.. anyway.. this hasn't been agreed or scheduled yet.. and if it does happen it'll probably be closed to 8.x).. I'm just rambling ;-)	0	0	https://issues.jboss.org/browse/GUVNOR-2832
5455	GUVNOR-2774	Hi [~Rikkola], I uploaded screenshot with PopUp dialogue where user creates new Guided Decision Table. Is it intentional that there is no _Label_ for _Hit Policy_ dropdown? I think it would be worth to add some _Label_ to describe purpose of the newly added dropdown. What do you think?	0	0	https://issues.jboss.org/browse/GUVNOR-2774
5456	GUVNOR-2673	[~mcimbora] Probably easier to discuss this on IRC. Question is did the handler not get invoked or the popup not displayed? The handler should definitely work in @AfterInitialization.	1	1	https://issues.jboss.org/browse/GUVNOR-2673
5457	GUVNOR-2637	[~karreiro] Hi, can this be closed?	1	1	https://issues.jboss.org/browse/GUVNOR-2637
5458	GUVNOR-2626	Can you please advise: 1) {{eval("$param" == "$param")}} etc is just to support conversion to Guided Decision Tables? 2) Do you have a XLS that works OK (when used as a XLS) in Drools? Can you provide it? There should be no need to "work around" things. 3) What purpose does, for example, B13:B26 serve? AFAIK this won't create a constraint when used as a XLS file; only columns C, E and F in "RuleTable Reason Codes LT Evaluation" seem to serve any function.	1	1	https://issues.jboss.org/browse/GUVNOR-2626
5459	GUVNOR-2626	Thanks. Is it possible to have you XLS without the workaround (use of $param==$param)? The issue is you are needing to try to workaround an issue with the conversion. You should not need to workaround anything; so I see the "fix" for this is to convert to a Guided Decision Table that works functionally equivalent to your unaltered XLS file.	1	1	https://issues.jboss.org/browse/GUVNOR-2626
5460	GUVNOR-2584	It is also possible to show everything in the white list. Then automatically add the imports into the file when a class is selected. This would require us to change all the data model selectors/lists. Most likely not a quick task.	0	0	https://issues.jboss.org/browse/GUVNOR-2584
5461	GUVNOR-2558	[~karreiro] Hi.. the crux is [this|https://github.com/droolsjbpm/kie-wb-common/blob/master/kie-wb-common-widgets/kie-wb-metadata-widget/src/main/java/org/kie/workbench/common/widgets/metadata/client/KieMultipleDocumentEditor.java#L547]. We iterate all "open documents" and save each one in turn.. calling to the {{SavePopup}}. Ideally we'd need to use {{SavePopup}} once and have the command passed to that then iterate all "open documents".. Some special care would be needed to handle - and report to Users - concurrent lock failures (when another User 2 had documents 'C' and 'D' open and edited, when User 1 has documents 'A', 'B', 'C' and 'D' open)...	1	0	https://issues.jboss.org/browse/GUVNOR-2558
5462	GUVNOR-2529	[~Rikkola] I believe this is done; can we close this (yes, sorry.. I'm looking for a JIRA and stumbled across this).	1	0	https://issues.jboss.org/browse/GUVNOR-2529
5463	GUVNOR-2507	For the first iteration; it's probably easiest to have a menu bar entry for (for example) "Open..." which launches a popup listing the other decision tables in the same package. When one is opened Conditions are scanned, checking for where they use a Field that's set in the Action of another (open) table; and vice-versa. These form the links. Things that need some tweaking:- * LockManager needs to operate per table * Version dropdown needs to operate per table * MetaData tab needs to operate per table * Data Objects needs to operate per table * BasicFileMenuBuilder is fixed to a single Path * .. any other tabs.....?	0	0	https://issues.jboss.org/browse/GUVNOR-2507
5464	GUVNOR-2494	A picture paints a thousand words.. even if the words are lost in translation ;)	1	1	https://issues.jboss.org/browse/GUVNOR-2494
5465	GUVNOR-2463	[~porcelli] Is this out of date? complete?	0	0	https://issues.jboss.org/browse/GUVNOR-2463

#	Requirement	Questions	Answered	Successfully Answered	URL
5466	GUVNOR-2460	Okay, I see your point. Is it feasible to make the fields at least ineditable, disabled? Similarly to the Save and Cancel buttons - they are disabled if the server template doesn't have process capabilities. Leaving them editable and at the same time not allowing the user to save the changes really is confusing.	0	0	https://issues.jboss.org/browse/GUVNOR-2460
5467	GUVNOR-2450	What webapp is this?	1	1	https://issues.jboss.org/browse/GUVNOR-2450
5468	GUVNOR-2450	[~jlocker] you may know, is this still valid issue?	1	0	https://issues.jboss.org/browse/GUVNOR-2450
5469	GUVNOR-2450	No idea. I never work with optaplanner-wb. Where can I get the WAR?	1	1	https://issues.jboss.org/browse/GUVNOR-2450
5470	GUVNOR-2385	Eh? Don't we look-up implementations with IOCBeanManager.getBeans(ResourceHandler.class) or something?!? How could NewResourceHandler try to instantiate a handler that cannot be instantiated?	1	1	https://issues.jboss.org/browse/GUVNOR-2385
5471	GUVNOR-2353	Toni, if you're looking at improvements to PNWL then there have also been a number of related JIRAs raised by QE recently. Did you want to consider them all?	1	0	https://issues.jboss.org/browse/GUVNOR-2353
5472	GUVNOR-2348	[~manstis] Hi Mike. In this case, I think it would be weird to have a dropdown to execute the actions "Add all" or "Add none", considering they are actions and not states. What do you think about adding a little cog button inside every cell to show those actions (similar to the one in the Project Explorer bar), and only when it is clicked this actions are shown?	1	1	https://issues.jboss.org/browse/GUVNOR-2348
5473	GUVNOR-2290	[~manstis] [~jsniezek], just realized, can' t really combine simple and advance search into the same screen. Problem is that the current simple search is context aware, that means it only filters out data that is currently displayed. Each view is responsible for reacting to when a search term is requested. How about its currently functionality but combining both options in the UI? Advanced search would be tied together with the simple search input box, via a dropdown, something like this: http://getbootstrap.com/components/#input-groups-buttons-segmented	1	1	https://issues.jboss.org/browse/GUVNOR-2290
5474	GUVNOR-2290	[~manstis] [~jsniezek] have a look on what I got so far, trying to keep it simple. Search box gets located at the primary nav bar, always visible, independently from the selected perspective. Find is removed from utility nav bar and moved into Authoring context nav bar, right aligned. We could also relabel Find to Advanced Search in there. What do you think?	1	1	https://issues.jboss.org/browse/GUVNOR-2290
5475	GUVNOR-2283	Hello [~Rikkola], this jira is marked as complete in BPMSPL-257 gdrive document. Could you please add pull requests which concerns this jira? Thanks.	1	0	https://issues.jboss.org/browse/GUVNOR-2283
5476	GUVNOR-2274	Is the goal to ignore WEB-INF/lib when building? This would make it possible to build legacy projects (or not, depends if the Maven KIE build plugin has support for using an older version). Editors might not work 100%, but the build would.	1	0	https://issues.jboss.org/browse/GUVNOR-2274
5477	GUVNOR-2224	Thanks again, is neither possible to send out the database userName?	1	1	https://issues.jboss.org/browse/GUVNOR-2224
5478	GUVNOR-2181	[~jomarko] Yes, I guess you are right... however we should probably stop adding new JIRAs to GUVNOR-2181 at some point (shouldn't we?!)	1	1	https://issues.jboss.org/browse/GUVNOR-2181
5479	GUVNOR-2107	Marco, I saw your comment about drools-wb-webapp being closer to the source; but in reality it is no closer to the source (guvnor-rest) than kie-drools-wb. drools-wb-distribution-wars is really not maintained (as nobody other than you!) uses it.. kie-drools-wb is what we rely on to "work". I'm still in favour of deleting drools-wb-distribution[-wars].	1	1	https://issues.jboss.org/browse/GUVNOR-2107
5480	GUVNOR-2094	Michael, Should the javax.security dependency in jbpm-human-task-core also be changed? See the pom.xml [here|https://github.com/droolsjbpm/jbpm/blob/master/jbpm-human-task/jbpm-human-task-core/pom.xml#L54].	1	1	https://issues.jboss.org/browse/GUVNOR-2094
5481	GUVNOR-2089	No, you cannot import a 5.5 repository.xml direclty into 6.0. There is a migration utility (https://github.com/droolsjbpm/drools-wb/tree/master/drools-wb-jcr2vfs-migration) that can convert a repository.xml into a git repository suitable for use with 6.0.x. The migration utility is part of the drools-workbench download available from https://www.jboss.org/drools/downloads	1	1	https://issues.jboss.org/browse/GUVNOR-2089
5482	GUVNOR-2067	Thanks but I tried, still no luck: java.lang.RuntimeException: Cannot find a default KieSession at org.drools.compiler.kie.builder.impl.KieContainerImpl.findKieSessionModel(KieContainerImpl.java:302) at org.drools.compiler.kie.builder.impl.KieContainerImpl.newKieSession(KieContainerImpl.java:295) at org.drools.compiler.kie.builder.impl.KieContainerImpl.newKieSession(KieContainerImpl.java:278) at com.austin.test.TestDemoStrata.main(TestDemoStrata.java:55) Cannot find a default KieSession I suspect something wrong in my kmodule.xml (shown in attachment) In fact I got RuleFlow running fine using 6.0.0.Final distribution (drools + guvnor). But with kie-drools-wb 6.1.0 created RuleFlow, it doesn't work. Or wonder if there is sample project based on 6.1.0?	1	1	https://issues.jboss.org/browse/GUVNOR-2067
5483	GUVNOR-2043	Jervis Liu <jliu@redhat.com> made a comment on [bug 996905|https://bugzilla.redhat.com/show_bug.cgi?id=996905] 6.0 has no problem to upload a jar file into Guvnor m2 repo. But it has problem to extend from a pojo class from a third-party jar. How to reproduce: 1. Upload includingSuper.jar (attached) into Guvnor m2 repo 2. Open demo -> uf-playground -> GuvnorM2RepoDependencyExample1, goto "Dependencies:dependencies list" editor, add includingSuper.jar as a dependency (use "add from repository"). 3. Use data modeller or modify Bean.java directly, to make it extends from com.sample.SuperClass. the result is like below: public class Bean extends com.sample.SuperClass implements java.io.Serializable { static final long serialVersionUID = 1L; @org.kie.api.definition.type.Position(value = 0) private java.lang.Integer value; public Bean() { } public Bean(java.lang.Integer value) { this.value = value; } public java.lang.Integer getValue() { return this.value; } public void setValue(java.lang.Integer value) { this.value = value; } } 4. project "build & deploy" fails.	0	0	https://issues.jboss.org/browse/GUVNOR-2043
5484	GUVNOR-1991	Is Apache Log4J on Guvnor's classpath? IIRC adding to your Application Server's /lib folder or uploading as an asset to Guvnor should both work.	1	0	https://issues.jboss.org/browse/GUVNOR-1991
5485	GUVNOR-1954	Prakash Aradhya <paradhya@redhat.com> made a comment on [bug 958396|https://bugzilla.redhat.com/show_bug.cgi?id=958396] Mike, can you take a look at this issue and see if ti can be fixed in 6.0 ? The work around is a little frustrating.	0	0	https://issues.jboss.org/browse/GUVNOR-1954
5486	GUVNOR-1948	Edson Tirelli <etirelli@redhat.com> made a comment on [bug 958398|https://bugzilla.redhat.com/show_bug.cgi?id=958398] Pedro, can you check if someone in your group can tackle this one? The core engine has a configuration option to define the severity of rule name duplication build result message: ERROR, WARN, INFO. Maybe all that is missing here is to properly catch that message and present to the user in the UI?	0	0	https://issues.jboss.org/browse/GUVNOR-1948
5487	GUVNOR-1945	Should this be on GUVNOR JIRA? That aside, I assume you need the ability to double-click the tab (or something like "maximize") and have the Panel expand to full screen. Double clicking reverts to the previous layout?	0	0	https://issues.jboss.org/browse/GUVNOR-1945
5488	GUVNOR-1940	Basically, the "workbench" needs to accept that developers' screens might be buggy and hence capture any java.lang.Exception and report it with a popup. Would you agree? This would apply to any method on a developers' screen that is invoked by the workbench; e.g. Constructors, @OnStart, @OnReveal etc etc.	0	0	https://issues.jboss.org/browse/GUVNOR-1940

#	Requirement	Questions	Answered	Successfully Answered	URL
5489	GUVNOR-1936	will the following use case be covered by this? On the RHS more attributes are set, than are used in the condition on the LHS (one could say the LHS shows the "unique key"). At the moment the element on the RHS is constructed manually and bound, in order that additional attributes can be set. //from row number: 1 //Hierarchie rule "Row 1 1_Anlage Gruppen & Hierarchiebaum" no-loop true dialect "mvel" when $h : Hierarchie(art == "RISK") not Hierarchie(name == "GR_FIL1" , parent == "GKMA_FIL1") then Hierarchie $h1 = new Hierarchie(); $h1.setArt($h.art); $h1.setName("GR_FIL1"); $h1.setParent("GKMA_FIL1"); $h1.setTyp("Gruppe EK"); $h1.setMinanzahl(1); $h1.setLevel(1); insert($h1); end	0	0	https://issues.jboss.org/browse/GUVNOR-1936
5490	GUVNOR-1908	How to make it work with ChangeSet.xml generated from Guvnor webapp (and not from classpath)?	1	1	https://issues.jboss.org/browse/GUVNOR-1908
5491	GUVNOR-1903	Sorry Toni, can you explain this a little more - so when we come to fix it we remember what the exact issue is?	0	0	https://issues.jboss.org/browse/GUVNOR-1903
5492	GUVNOR-1892	Can you please provide some more information? Guvnor (where this JIRA has been raised) is an authoring tool, and hence being able to insert a new event (presumably into working memory) doesn't make much sense - unless you're talking about from the right-hand-side of an authored rule or something, but you are not clear.	1	1	https://issues.jboss.org/browse/GUVNOR-1892
5493	GUVNOR-1875	@Edson, silly question possibly... changing from use of RuleBase to KnowledgeBase (and related KnowledgeSession implementations) is obvious. I am however a little confused over the use of Package[]. We extend PackageBuilder too and have references to Package[] floating around (PackageBuilder.getPackages). Is this part of the deprecated API too? Should we look to change to KnowledgeBuilder and Collection<KnowledgePackage>?	0	0	https://issues.jboss.org/browse/GUVNOR-1875
5494	GUVNOR-1837	Should we adjust the Ballroom theme to follow the new GWT changes or follow the current style? Ballroom might be unusable to us because of GWTP, but the styles were made for the custom components we have (popups, tabs, panels).	0	0	https://issues.jboss.org/browse/GUVNOR-1837
5495	GUVNOR-1834	One more thing... what is your expectation of how XMLGregorianCalendars might be used? It does not implement Comparable so normal operators on the object itself won't work - e.g. MyFact(myGregorianCalendarField > "01-01-2000") etc. You'd need to handle the date components separately - e.g. $m : MyFact(myGregorianCalendarField.eonAndYear > 2000)... $m.setEonAndYear(2012); etc. Did you envisage the coercion between object (used in the UI) and fields (in the DRL) to be handled implicitly - then things like this: "myGregorianCalendarField1 > myGregorianCalendarField2" become difficult in the LHS - mapping to "myGregorianCalendarField1.day > myGregorianCalendarField2.day && myGregorianCalendarField1.month > myGregorianCalendarField2.month && myGregorianCalendarField1.year > myGregorianCalendarField2.year && myGregorianCalendarField1.eon > myGregorianCalendarField2.eon" (plus time parts I've omitted here) etc. Also CEP functions become almost impossible: MyEvent( ) after $myBoundGregorianCalendarObject...	0	0	https://issues.jboss.org/browse/GUVNOR-1834
5496	GUVNOR-1834	Any ideas how you expected this to operate when XMLGregorianCalendar does not implement Comparable? I bumped into a couple of guys at the BRMS 2012 event in London, who said they were responsible for this requirement coming to light (they'd been liaising with you or something, I forget the detail). They were not entirely sure how this feature may actually work in practice either. Can you advise any further (otherwise I fear this JIRA runs the risk of being rejected - essentially a XMLGregorianCalendar becomes just another "business domain" class - for which we provide the ability for custom editors to be configured already - see Working Set definitions).	0	0	https://issues.jboss.org/browse/GUVNOR-1834
5497	GUVNOR-1809	@Bill Thanks. Gwendo (a community member) is giving this issue a go. The hope is that RESTEASY-704 will allow us to do what we need. Here's what I was struggling with (possibly fixed by 704) http://article.gmane.org/gmane.comp.java.resteasy.user/121/match=geoffrey {code} Hi guys, We're moving from cxf with Abdera for Atom to RestEasy with RestEasy's atom support. So far, most of the migration has been a breeze, but this part is unclear: In Cxf Abdera we got: Entry e = factory.newEntry(); ... //generate meta data ExtensibleElement extension = entry.addExtension(METADATA); ExtensibleElement childExtension = extension.addExtension(ARCHIVED); //childExtension.setAttributeValue("type", ArtifactsRepository.METADATA_TYPE_STRING); childExtension.addSimpleExtension(VALUE, p.isArchived() ? "true" : "false"); childExtension = extension.addExtension(UUID); childExtension.addSimpleExtension(VALUE, p.getUUID()); ... Content content = factory.newContent(); ... entry.setContentElement(content); How does the ExtensibleElement stuff translate to RestEasy's atom support? Entry e = new Entry(); ... // TODO ExtensibleElement stuff ... Content content = new Content(); ... e.setContent(content); The manual does seem to imply that's possibly, only just not how: "Feeds are composed of a number of items, known as "entries", each with an extensible set of attached metadata." http://docs.jboss.org/resteasy/docs/2.3.1.GA/userguide/html_single/index.html#Atom Thank in advance for any help. {code}	1	1	https://issues.jboss.org/browse/GUVNOR-1809
5498	GUVNOR-1802	6.0.x REST API does not support asset operations (i.e. get assets for a category). Asset operations should be conducting using a GIT library to manipulate the underlying VFS directly.	0	0	https://issues.jboss.org/browse/GUVNOR-1802
5499	GUVNOR-1785	What is the problem in this issue? Guvnor 5.3 uses XmlElement and builds fine with JDK 5 because there is this dependency in the parent pom: {code} <dependency> <groupId>javax.xml.bind</groupId> <artifactId>jaxb-api</artifactId> <version>2.2.1</version> </dependency> {code}	1	1	https://issues.jboss.org/browse/GUVNOR-1785
5500	GUVNOR-1785	I'd love to see you guys support Java 6. Is there a policy in place for JRE support for Drools in general? Or Maven versions?	1	0	https://issues.jboss.org/browse/GUVNOR-1785
5501	GUVNOR-1756	You can already build DRL akin to this:- 1. | rule "test" 2. | dialect "mvel" 3. | when 4. | $r : java.util.List( ) from collect (Applicant(age > 35)) 5. | then 6. | end You need to use the "from collect" option and ensure you have java.util.List added as a package import. You are unable to select properties of the List though (like "size") so I'll leave this JIRA open to fix that.	0	0	https://issues.jboss.org/browse/GUVNOR-1756
5502	GUVNOR-1746	Toni, can we reject this?	0	0	https://issues.jboss.org/browse/GUVNOR-1746
5503	GUVNOR-1741	Does the linked JIRA suffice? Can this be closed, or is there something more to it?	0	0	https://issues.jboss.org/browse/GUVNOR-1741
5504	GUVNOR-1735	Michael, is there anyway I can change the GWT defaults so that the >> button will work for rows less than 1000? If so where would this code change need to be made and then how would I submit this change to GUVNOR? -Wendell	1	1	https://issues.jboss.org/browse/GUVNOR-1735
5505	GUVNOR-1735	Hi Wendell, It isn't a complicated change - GuvnorSimplePager is the class and the zero parameter constructor called by GWT's UiBinder delegates to another taking the 'fast forward' step size. We can simply change this to a sensible size (1000 is quite large for Guvnor's general usage - I'd say 100, would you agree?) Don't forget I've already added the ability to go to the last page as this JIRA requested. I'll make the change - to 100 - this evening.	1	1	https://issues.jboss.org/browse/GUVNOR-1735
5506	GUVNOR-1735	Thank you Michael, How would I go about downloading this updated version with the change? Or, is it slated for the next release? I ask because I have a client who would like the update in their production environment as soon as possible. -Wendell	1	1	https://issues.jboss.org/browse/GUVNOR-1735

#	Requirement	Questions	Answered	Successfully Answered	URL
5507	GUVNOR-1733	This is not entirely clear: {code} create snapshot by name (should compile and create a snapshot with a specific name). getPackageBinary doesnt work in production because if a DTABLE takes some time to compile then it will halt the prod system while it compiles {code} Does this mean, that when calling this new method createPackageSnapshot(), you expect the method to return immediately (and do the compilation as async - which can cause other problems) or return after compilation?	1	1	https://issues.jboss.org/browse/GUVNOR-1733
5508	GUVNOR-1733	{code} @Get @Path("{packageName}/assets/{assetName}/binary") - is this the .pkg binary? or the .xls... the work "binary" is not explicit enough {code} This is certainly the pkg, which is considered the binary. Update: looks like it might not be returning the pkg {code} @Get @Path("{packageName}/assets/{assetName}/source") - is this a binary .xls source or a text drl? the word "source" is not explicit enough {code} This should be the source file (the .xls), but is currently always .drl (which is not always the real source file). Update: throws exception if there is no plain text source.	1	1	https://issues.jboss.org/browse/GUVNOR-1733
5509	GUVNOR-1723	You can always "view" the actual version to drill down into the real nitty-gritty meta-data. Is that possible today, or in the future? Last Contributor is the most important part.	0	0	https://issues.jboss.org/browse/GUVNOR-1723
5510	GUVNOR-1716	What would you propose? This JIRA is very subjective. "Chapter 5. Managing Assets" of the user guide covers the new screen and use. Perhaps this chapter should be moved nearer the front of the documentation?	0	0	https://issues.jboss.org/browse/GUVNOR-1716
5511	GUVNOR-1710	Looks good. How is the lifecycle assigned? For each asset in the module, each asset or can you change the lifecycle if you want to?	0	0	https://issues.jboss.org/browse/GUVNOR-1710
5512	GUVNOR-1700	Should we rename it to EAP instead? Problem is most community users don't know what EAP means...	0	0	https://issues.jboss.org/browse/GUVNOR-1700
5513	GUVNOR-1675	*Parse definition* {code} List<Map<String, Object>> workDefinitionsMaps = (List<Map<String, Object>>) MVEL.eval(content, new HashMap()); for (Map<String, Object> workDefinitionMap : workDefinitionsMaps) { if (workDefinitionMap != null) { WorkDefinitionImpl workDefinition = new WorkDefinitionImpl(); workDefinition.setName((String) workDefinitionMap.get("name")); workDefinition.setDisplayName((String) workDefinitionMap.get("displayName")); workDefinition.setIcon((String) workDefinitionMap.get("icon")); workDefinition.setCustomEditor((String) workDefinitionMap.get("customEditor")); Set<ParameterDefinition> parameters = new HashSet<ParameterDefinition>(); if(workDefinitionMap.get("parameters") != null) { Map<String, DataType> parameterMap = (Map<String, DataType>) workDefinitionMap.get("parameters"); if (parameterMap != null) { for (Map.Entry<String, DataType> entry : parameterMap.entrySet()) { parameters.add(new ParameterDefinitionImpl(entry.getKey(), entry.getValue())); } } workDefinition.setParameters(parameters); } if(workDefinitionMap.get("results") != null) { Set<ParameterDefinition> results = new HashSet<ParameterDefinition>(); Map<String, DataType> resultMap = (Map<String, DataType>) workDefinitionMap.get("results"); if (resultMap != null) { for (Map.Entry<String, DataType> entry : resultMap.entrySet()) { results.add(new ParameterDefinitionImpl(entry.getKey(), entry.getValue())); } } workDefinition.setResults(results); } if(workDefinitionMap.get("defaultHandler") != null) { workDefinition.setDefaultHandler((String) workDefinitionMap.get("defaultHandler")); } if(workDefinitionMap.get("dependencies") != null) { workDefinition.setDependencies(((List<String>) workDefinitionMap.get("dependencies")).toArray(new String[0])); } workDefinitions.put(workDefinition.getName(), workDefinition); } } {code} *Runtime registration* {code} ksession.getWorkItemManager().registerWorkItemHandler("Notification", new NotificationWorkItemHandler()); {code} *Runtime invocation* {code} WorkItemManager wim = ksession.getWorkItemManager(); KnowledgeSessionConfiguration sessionConfiguration = ksession.getSessionConfiguration(); Map<String, org.drools.runtime.process.WorkItemHandler> handlers = ((SessionConfiguration) sessionConfiguration).getWorkItemHandlers(); WorkItemHandler wih = handlers.get("Notification"); wih.executeWorkItem(workItem, wim); //TODO Where does workItem come from? {code}	0	0	https://issues.jboss.org/browse/GUVNOR-1675
5514	GUVNOR-1595	This can be accompished by using an external source for the enumeration's values. See http://docs.jboss.org/drools/release/5.5.0.Final/drools-guvnor-docs/html/ch04.html#d0e1927. The mentioned "SampleDataSource" could read from a database. Do you however request the means to achieve this with no code? e.g. you can create data sources in Guvnor directly and select tables\fields\stored procedures\views etc from which the data should be obtained.	0	0	https://issues.jboss.org/browse/GUVNOR-1595
5515	GUVNOR-1588	<ge0ffrey> manstis: ^^ someone should be appointed time to delete that old RestApiServlet really... <manstis> ge0ffrey, is there a JIRA for it? <ge0ffrey> manstis: yes <manstis> ge0ffrey, can you assign to 5.5.0.Beta2 and jliu?	1	0	https://issues.jboss.org/browse/GUVNOR-1588
5516	GUVNOR-1567	Only a thought, but having BRMS push to a HA store moves the single-point of failure from a single BRMS instance to a HA node in the system architecture, although admittedly you do eliminate the race condition you describe too. This http://community.jboss.org/wiki/PublishDroolsArtifactsFromAProductionEnvironment would achieve the same result; using an intermediate pull mechanism? Furthermore can the same resilience not be equally achieved by deploying BRMS in such a way that it is not the single-point of failure in the first place; the HA store becoming the JCR cluster fronted by a BRMS web-server cluster? http://blog.athico.com/2011/01/configuring-multiple-guvnor-instances.html?utm_source=feedburner&utm_medium=feed&utm_campaign=Feed%3A+DroolsAtom+%28Drools+-+Atom%29	1	0	https://issues.jboss.org/browse/GUVNOR-1567
5517	GUVNOR-1521	How did you migrate the rules from Derby to Oracle? I assume you exported the repository whilst running with Derby (to an XML file), re-configured and then imported into Oracle; both using the Admin Import\Export feature? If this is correct did you get any errors performing the import?	1	1	https://issues.jboss.org/browse/GUVNOR-1521

#	Requirement	Questions	Answered	Successfully Answered	URL
5518	GUVNOR-1521	Dear Michael, Firstly, before deploying the Drools guvnor in weblogic11g , I set the repository configuration in components.xml.below is the path i mentioned for repositoryConfiguration <component name="repositoryConfiguration"> <property name="homeDirectory">E:/GuvnorRepo/</property> Deployed the Drools guvnor in weblogic11g , and in E drive GuvnorRepo folder has created with 3 sub folders workspaces,repository,version and repository.xml. Now i updated the repository.xml for changing the database from derby to oracle. Below are the changes <PersistenceManager class="org.apache.jackrabbit.core.persistence.bundle.OraclePersistenceManager"> <param name="driver" value="oracle.jdbc.driver.OracleDriver"/> <param name="url" value="jdbc:oracle:thin:@rsa10gserver:1525:ifpub34"/> <param name="user" value="drools_repository"/> <param name="password" value="drools_repository"/> <param name="schema" value="oracle"/> <param name="schemaObjectPrefix" value="version_"/> <param name="tableSpace" value=""/> <param name="errorHandling" value=""/> </PersistenceManager> Now i restarted the weblogic11g and it successfully deployed and entered into the login page and but i could not see any rules in the left navigation page. Is there any other procedure rather than what I have done? Whether I should import/export new repository.xml ? Please help me regarding this.	1	0	https://issues.jboss.org/browse/GUVNOR-1521
5519	GUVNOR-1521	@Deepak, can you create new rules when configured to use Oracle? If so your configuration, I'd assume, to be correct. AFAIK the repository folder is only relevant if you're using JCR's file-system storage - if configured to use a database I doubt it has much relevance. So you'll need to export the repository XML first (your comment above implied you know how) from Derby, re-configure for Oracle and import. I don't know if the binary form stored by JCR on RDBMS's BLOB fields can be migrated used a database script - that would be a question for the JCR forums not Drools. Please also see Rikkola's comment about where this should be tracked - I am happy to continue the conversation there. @Rikkola, I agree. We shouldn't be tracking this on JIRA. Closing this one....	0	0	https://issues.jboss.org/browse/GUVNOR-1521
5520	GUVNOR-1509	How do we want to keep track of the valid flag, should we add it as a MetaData field to be stored and recalculated when we save the rule/enumeration? Where should the visual indicators be shown? In the asset list? When you open an asset? Next to status in the title bar? It seems like the unsaved changes indicator might be solved by https://issues.jboss.org/browse/GUVNOR-1263. So that leaves only the validation part for this issue.	1	1	https://issues.jboss.org/browse/GUVNOR-1509
5521	GUVNOR-1417	Can you elaborate or close?	0	0	https://issues.jboss.org/browse/GUVNOR-1417
5522	GUVNOR-1414	Pray tell, where might this issue be experienced? I tried the various DSL screens with which I am familiar and failed to identify where this issue might occur.	1	1	https://issues.jboss.org/browse/GUVNOR-1414
5523	GUVNOR-1404	In order to make search easy as possible, we would like to have automatic meta data collected during the process of importing rules from files or during the save process from UI, to repository back end. The whole idea here is to keep the concept of Category only for business purposes. Besides the category concept , we could have attributes set into the repository attached with its components. With that we mean, that we could automatically (The saver to repository is a suggesiton on doing this), to extract Columns (Could be OTN also?) from rules, and create lists of all column types a given rule uses and even other type of data could be extracted from rule files and definitions With that we could be able to link in searches, business desciptions and categories to real implementation, giving very usefull results (A Better search engine than Google).... This woudl apply to any rule change/save. The meta data would need to be refreshed. We could use a structure like RuleNode / ColumnNode -field list (String[]) -name	0	0	https://issues.jboss.org/browse/GUVNOR-1404
5524	GUVNOR-1382	Is this a duplication of GUVNOR-1381?	0	0	https://issues.jboss.org/browse/GUVNOR-1382
5525	GUVNOR-1381	Is this a duplication of GUVNOR-1382?	1	0	https://issues.jboss.org/browse/GUVNOR-1381
5526	GUVNOR-1356	Have you looked at "Working Sets"?	1	1	https://issues.jboss.org/browse/GUVNOR-1356
5527	GUVNOR-1355	Can you please explain what your request is in a little more detail? Do you mean to select many fields on the screen and change their value en-mass?	0	0	https://issues.jboss.org/browse/GUVNOR-1355
5528	GUVNOR-1355	The issue came up in the class/lab at JBW that a student had an enumeration drop down in the UI. In their use case they were looking at Cities, imagine a list of cities, and you want to select 10 of 100. In the current version of this I would need to create 100 when lines, selecting only one at a time, they were hoping that even if under the covers this could be managed, from a UI perspective, they would just have selected and unselected cities in that list of 100 without needing to create 10s of when lines. I hope this makes sense to you, if not let me know.	0	0	https://issues.jboss.org/browse/GUVNOR-1355
5529	GUVNOR-1322	Fixed row heights for all tables that can have the second "description" row. Obviously last page, if having fewer than the PAGE_SIZE of rows, still needs some mouse movement but otherwise original issue resolved. A happy compromise.	0	0	https://issues.jboss.org/browse/GUVNOR-1322
5530	GUVNOR-1247	Will the non checkin saves be public for all or just for the user who edited? This would help to make GUVNOR-1067 possible (Asset Editor: Auto save when editing a rule). It would also help to do big changes to the package/module without breaking everything for everybody else.	1	1	https://issues.jboss.org/browse/GUVNOR-1247
5531	GUVNOR-1220	Good idea, but there is a pitfall: If you add the modeshape impl jar to the war, the war might not work for jackrabbit on JDK 1.5 and guvnor must keep working on JDK 1.5 on jboss AS (at least for this year, let's hope we can loose JDK 1.5 soon after). => Test if this change still allows the war to run on a JDK 1.5 on jboss AS before pushing it.	1	1	https://issues.jboss.org/browse/GUVNOR-1220
5532	GUVNOR-1205	Isn't this almost complete for 5.3.0.Beta1?	1	1	https://issues.jboss.org/browse/GUVNOR-1205
5533	GUVNOR-1168	Geoffrey De Smet <gdesmet@redhat.com> made a comment on [bug 724493|https://bugzilla.redhat.com/show_bug.cgi?id=724493] For drools 5.4, several improvements by Mat Allen (the original author of the bulk importer) were merged into the bulk importer. Is this issue still valid?	0	0	https://issues.jboss.org/browse/GUVNOR-1168
5534	GUVNOR-1125	Hi Toni, I was thinking about how to use Guvnor to govern assets from other project, eg JBOSS ESB. Will this fall into the category of perspective as well? For example, Drools perspective and JBOSS ESB perspective. Or its sth bigger than perspective?	1	1	https://issues.jboss.org/browse/GUVNOR-1125
5535	GUVNOR-1103	Geoffrey.. this is getting escalated by many users.. please take a look	0	0	https://issues.jboss.org/browse/GUVNOR-1103
5536	GUVNOR-1076	Kris Verlaenen <kverlaen@redhat.com> made a comment on [bug 958408|https://bugzilla.redhat.com/show_bug.cgi?id=958408] This related to how metadata for files is stored in Guvnor for BRMS/BPMS6. Is metadata stored in an additional file on the file system (so it can be editor from Eclipse using for example EGit)?	0	0	https://issues.jboss.org/browse/GUVNOR-1076
5537	GUVNOR-1038	Marcos: can we have some more details on this?	1	1	https://issues.jboss.org/browse/GUVNOR-1038
5538	GUVNOR-1030	Interesting. I see this use case all the time. Are you proposing automatic WSDL generation? Another factor to consider is the possible need for ruleset parameters, a la ILOG JRules, since many of these services are backed by rule flows.	1	1	https://issues.jboss.org/browse/GUVNOR-1030

#	Requirement	Questions	Answered	Successfully Answered	URL
5539	GUVNOR-1029	Would it not be an idea to allow for some other JBOSS Rules DRL keywords to be available in the JBRMS? For example, this specific issue could be resolved by exposing "from" in the editor. This would allow for static lists, database queries and whatever else. exposing other keywords too (such as "accumulate" etc) would allow for sophisticated rules to be defined graphically - of course the GUI design would be the stumbling block!	1	1	https://issues.jboss.org/browse/GUVNOR-1029
5540	GUVNOR-1029	This comment follows on from my interpretation of Michael N's original post (25/4) which as I read it would result in a rule with a single value hard coded in the condition. Michael A's post, if I understand it correctly, is asking for some of the more powerful features like 'from' to be added which provide dynamic conditions derived at time of rule execution. This idea was prompted by the user requirement for drop down lists within JBRMS which we are proposing to provide by importing a java bean into the package. This enables a drop down list as shown in the attached document. Would it be possible to introduce shared code+descriptor table(s) that could be used both by JBRMS for the LHS rule conditions and the application that will be asserting objects for the rules to process. Rather than hard coding conditions in rules : Person(country=="IRQ") This would permit drop down lists using the descriptor Person(country==IRAQ) INDIA POLAND which when IRAQ was selected would generate the same rule Person(country=="IRQ") This would have 2 advantages : 1) the same code "IRQ" would be used in all rules for Iraq 2) the calling application would use the same table for populating country in the object to be asserted ie."IRQ" thus ensuring the rule would be fired.	1	1	https://issues.jboss.org/browse/GUVNOR-1029
5541	GUVNOR-1012	Email from Manik (from Michael). > > I had a question regarding jboss cache - we are looking at it as a mechanism to "publish" rule bases to anyone who is interested (ie runtime engines look in a distributed cache for a rulebase, and the rule server pushes out changes into the cache). I'm presuming a "rulebase" is represented by a simple Java object? > Firstly, I am making the assumption this is a sensible use of jboss cache (seems ok to me) - but do you have any thoughts? Yes, except JBC is more of a P2P system rather than a client-server type thing. Any node running JBC can put stuff in the cache it would be visible on every cache in the cluster. That said, one node in the cluster could always be designated as the "server" and be responsible for adding rulebases while the others just read rulebases from the cache. > Secondly, the "runtime engines" ideally will need to be cold started - in a sense that the rule server may not be up. So they need some persistent cache in a local database or something. Does JBoss Cache have an facilities to have local persistence? or is this best left up to a specific app? Cold starts can be dealt with in 2 different ways - persistence (which we do support - lots of plugins for different engines, see [1] and [2]) - or state transfer (see [3] and [4]). For the latter, this is when a cache instance starts up and asks its' neighbours for their in-memory state, which then gets streamed across the wire. [1] Crappy old documentation from the "current" 1.4.x series on cache loaders http://labs.jboss.com/file-access/default/members/jbosscache/freezone/docs/1.4.1.SP2/TreeCache/en/html/cache_loaders.html [2] Newer docs from the 2.0.0.BETA series on cache loaders http://labs.jboss.com/file-access/default/members/jbosscache/freezone/docs/2.0.0.BETA1/JBossCache-UserGuide/en/html/cache_loaders.html [3] Crappy old documentation from the "current" 1.4.x series on state transfer http://labs.jboss.com/file-access/default/members/jbosscache/freezone/docs/1.4.1.SP2/TreeCache/en/html/state_transfer.html [4] Newer docs from the 2.0.0.BETA series on state transfer http://labs.jboss.com/file-access/default/members/jbosscache/freezone/docs/2.0.0.BETA1/JBossCache-UserGuide/en/html/clustering.html#d0e1356 -- Manik Surtani Lead, JBoss Cache JBoss, a division of Red Hat	1	1	https://issues.jboss.org/browse/GUVNOR-1012
5542	GUVNOR-1002	Was talking to Fernando today about the ROM -> Descr conversion and we came to a conclusion. Besides all we talked about supporting directly ROM->Descr without serializing to DRL, if a rule is "contaminated" by a single DSL sentence, the whole rule MUST be expanded to a DRL and use the parser to convert it to a descriptor. That is because someone could write in the BRMS (mixing DSL and DRL sentences) something like: rule xyz when Cheese() // drl sentence - whose type is "brie" // dsl sentence then // RHS end In the above case, the DSL sentence is a constraint (type == "brie") that needs to be added to the Cheese() pattern. So, once contaminated by a DSL sentence, the direct conversion can't be done. Once we understand that we need to have a ROM -> DRL dumper for cases where the ROM is contaminated, the question that rises is "what good does us to have a ROM->Descr converter? does the performance gain of not using the parser justifies the disadvantage of code complexity and maintanance of another 'converter'?". I know I pushed for direct ROM->Descr conversion, but I didn't know about the DSL/DRL sentence mixing possibility. I would say now that we should simply: 1. drop the BRXMLtoDescrConverter 2. make BRLPersistence an interface and rename the current implementation to something like XMLBRLPersistence. 3. implement a DRLBRLPersistence to dump ROM to DRL format 4. create a new BRXMLtoDescrConverter that is simply a helper, higher level API that does ROM->DRL->Descr conversion using above class + drl parser. What do you think?	1	1	https://issues.jboss.org/browse/GUVNOR-1002
5543	GUVNOR-1002	I know I pushed for direct ROM->Descr conversion, but I didn't know about the DSL/DRL sentence mixing possibility. I would say now that we should simply: MN: yes its a bit unfortunate - would could drop the DSL for that, or just restrict it in a lot of ways (but either is painful) 1. drop the BRXMLtoDescrConverter 2. make BRLPersistence an interface and rename the current implementation to something like XMLBRLPersistence. MN: yes, thats a better name no matter what. 3. implement a DRLBRLPersistence to dump ROM to DRL format MN: OK, initially this would be one way of course - but eventually round trip (but not that concerned) 4. create a new BRXMLtoDescrConverter that is simply a helper, higher level API that does ROM->DRL->Descr conversion using above class + drl parser. MN: Yes that was my initial idea, but I was badly jetlagged at the time, so I still like your descr idea, but as you say, DSL messes it up. \ What do you think? Well, originally I invented this specifically to spit out DRL to show what it looks like, so I don't see a problem with this as long as its not too hard. I don't see a problem with that approach, I never thought the direct to Descr really helped that much for performance anyway, parsing is very fast, its the compiling that kills us.	1	1	https://issues.jboss.org/browse/GUVNOR-1002
5544	GUVNOR-1002	Michael, I added the persister capable of doing ROM -> DRL transformation. Although, the round trip is something that worries me a bit. Basically, we delliberatelly added some restrictions to the modeller and to the ROM (like no arbitrary nesting of CEs) that don't exist in the pure core engine (and in the DRL). How do you think we should handle such restrictions when "importing" DRLs into the BRMS? I'm asking that because the round trip code is where this needs to be implemented, right? []s Edson	1	1	https://issues.jboss.org/browse/GUVNOR-1002

#	Requirement	Questions	Answered	Successfully Answered	URL
5545	GUVNOR-1002	Michael I added the integration of the conversion to the BRXMLContentHandler and added the test case in ServiceImplementationTest. Look for the test named: testBinaryPackageCompileAndExecuteWithBRXML. The only thing I'm not sure is how you differentiate Strings from other constants, like numbers in the model. For instance, in the test, I wanted to set the integer "age" field of the person object. As a way to differentiate the value from a regular String value, I made it a "formula", adding the "=" as the first character. Although, I'm not sure that this is the best way of doing it. Can you plz advise what you imagined for that? Example: This: ActionSetField action = new ActionSetField("p"); ActionFieldValue value = new ActionFieldValue("age", "=42"); action.addFieldValue(value); Is converted to (in DRL): p.setAge(42); Thanks.	1	1	https://issues.jboss.org/browse/GUVNOR-1002
5546	GUVNOR-1002	Well we can augment the model to express if its literal etc... but then the issue is how to know what to put in the model - do we ask the user? or look at the type/variable names etc.	1	1	https://issues.jboss.org/browse/GUVNOR-1002
5547	GUVNOR-1002	Ok, commited and it is working. What else needs to be done under this ticket?	1	1	https://issues.jboss.org/browse/GUVNOR-1002
5548	GUVNOR-1002	OK, I woke up this morning in a panic (I tend to wake up like that - I hate the clarity and perspective that waking up brings): What about date literals? Dates are also kind of a special case I think. On the LHS in drl dates are pretty easy, we just slap quotes around them and the compiler takes care of it. But on the RHS we need to new java.util.Date() - although MVEL may take care of it. Some general questions I am not sure about (regarding the RHS): * Do we allow a field to be set to a bound variable? * We can deal with string/number literals, but what about dates? * Say we want to set a field to a calculation, for a non String, this should work as it is now right? so we have a "value" of "x *3.0 - someVar" then that will be: boundVariable.setSomeField(x * 3.0 - someVar) UNLESS it is a string, in which case it will try to wrap quotes around it (which is not really a problem, as you generall don't do calculations with a string). I guess my original question remains: shoudl we ask the user to classify if they are doing a literal, a formula or choosing from a list of variables? or is that overkill and unhelpful?	1	1	https://issues.jboss.org/browse/GUVNOR-1002
5549	GUVNOR-1002	My suggestion is the obvious one: lets start simple and build over it. > What about date literals? Dates are also kind of a special case I think. On the LHS in drl dates are pretty easy, we just slap quotes around them and the compiler takes care of it. But on the RHS we need to new java.util.Date() - although MVEL may take care of it. Yes, the old java Date issue. I think what we need to do is to define a general (maybe configurable) mask for dates and automatically parse strings into dates in the input fields. This can be based on the field type, as we already do manipulations for strings (adding quotes) and formulas (removing the = sign). > Do we allow a field to be set to a bound variable? Sure. From a code generation perspective, it is already supported. Just a matter of allowing users to chose or type variables in the Web GUI. > We can deal with string/number literals, but what about dates? See above. > Say we want to set a field to a calculation, for a non String, this should work as it is now right? so we have a "value" of "x *3.0 - someVar", then that will be: boundVariable.setSomeField(x * 3.0 - someVar), UNLESS it is a string, in which case it will try to wrap quotes around it (which is not really a problem, as you generall don't do calculations with a string). The way it is now, from code generation perspective, it is already working. The key thing is that if the field type is String and you want to do manipulations, you MUST use the '=' sign prefixing the expression so that the engine knows it is a formula. If it is not a String, that the '=' is optional. It will work both with and without it. > I guess my original question remains: shoudl we ask the user to classify if they are doing a literal, a formula or choosing from a list of variables? or is that overkill and unhelpful? Lets do something similar to what Excel does... if it is a literal, variable, etc, the user simply types/selects it. If it is a formula, user can start the expression with '=' (exactly like excel) or maybe use a wizard when we build one. Only thing is not there yet, but easy to implement as soon as we define how it will work is Date parsing. Regarding the "type" field in ActionSetField, I think it is a useful declarative information for the model and it is good to have it there, independent of the SuggestionCompletionEngine. I think we may not have the SCE always available everywhere to use and having the field type information makes it kind of "self-contained". So, what we do regarding Dates? Does JodaTime has anything useful we can use here?	1	1	https://issues.jboss.org/browse/GUVNOR-1002
5550	GUVNOR-1002	Yes, best thing would be to support a specific Date type in the SCE. Will look into doing that... do you see any problem? Yes, everything else is covered from an engine perspective, IMO.	0	0	https://issues.jboss.org/browse/GUVNOR-1002
5551	GUVNOR-996	Michael, Looking at the SuggestionCompletionEngine, it is a quite heavy object o populate. I imagine if a package has a model with just a few tens of classes, it will take a long time to inspect them, populate the map with all the attributes and transmit over the network, reparsing it in the browser side before presenting it to the user. Wouldn't be the case to do it on demand? Since each rule uses just a few facts, we could possibly inspect and offer completion for the classes effectivelly in use by the rule. What do you think?	1	1	https://issues.jboss.org/browse/GUVNOR-996
5552	GUVNOR-937	in the "then" part - if you put an "=" at the start (a bit like in excel) you can do a calculation (it won't be a text field). Is that what you want?	1	1	https://issues.jboss.org/browse/GUVNOR-937
5553	GUVNOR-937	The calculation box does not take in any alphabets,but it takes in special characters and numerals, as the "SI" is of float data type. If the variable is of string type it takes all the keys. can you please suggest me in how to remove the data type validation in jbrms. What my idea is, in the "when" part, we have an option of "bound with variable", similar to it can we have "computation on variables" in "then" part.The functionality will be same, just the difference in usage.	1	1	https://issues.jboss.org/browse/GUVNOR-937
5554	GUVNOR-937	I have not been able to reproduce it. If you type "=" as the first character in the text box (talking about "then" of course) then it WILL let you put in numbers (or should). If not, then it is a bug (and we should log it as such). Can you provide more details on this? Yes I like your idea of the "literal/formula" thing similar to the "when" part (but we want to keep it simple - as most of the time it is literal values or calculations).	0	0	https://issues.jboss.org/browse/GUVNOR-937
5555	GUVNOR-912	As GWT is still 1.4 RC1 (and a bit bugged) we should wait for a stable version before change anything. agree ?	1	1	https://issues.jboss.org/browse/GUVNOR-912
5556	GUVNOR-912	What happened when you tried to update it? Yes its still RC1, but we NEED to know what the issues are, so we can log them with the GWT team (I have already logged on on the Rich Text one) so they can do another RC. So please let me know what issues are so we can report them.	1	1	https://issues.jboss.org/browse/GUVNOR-912
5557	GUVNOR-901	Some concrete tasks: * i18n any class, using the eclipse "externalise strings" feateure * Take the generated Messages class, make it work with GWT (not sure exactly how, GWT has lots of ways) * test performance of rendering using this The above will prove that using the externalise feature can work well (GWT has another way, but it is not automated from Eclipse - perhaps there is another plug in to do that, I don't know).	1	1	https://issues.jboss.org/browse/GUVNOR-901
5558	GUVNOR-888	Barry - what do you mean by testing a rule? My take on this is that in different environments you have different rules included/excluded (not just different versions of rules) - my point of confusion is how this can work? How do you isolate testing to just a rule or rules - aren't you testing the system behaviour as a whole?	1	1	https://issues.jboss.org/browse/GUVNOR-888
5559	GUVNOR-888	How are the rules marked as "tested" and who sets the effective dates for testing? is it manual?	1	1	https://issues.jboss.org/browse/GUVNOR-888

#	Requirement	Questions	Answered	Successfully Answered	URL
5560	GUVNOR-888	> How are the rules marked as "tested" and who sets the effective dates for testing? is it manual? The rules will be manually updated with "Tested", Effective Date, Expiration Date, Test Effective Date and Test Expiration Date by the business user. In the long run I would like to automate the tested flag to true when the snapshot containing that rule is deployed to the test environment, but no time for that right now. :-) [Show »] Michael Neale [15/Aug/07 10:09 PM] How are the rules marked as "tested" and who sets the effective dates for testing? is it manual? > have a config file on the classpath that specifies the selectors available. They are a class that implements certain interfaces. Out of the box there > are none, but if you add one in, then you can choose it in the GUI when you go to build a package (it will show up a list of them, you pick the > appropriate one) - that allows users to define them in any way they see fit. > This does NOT however address fields that don't currently exist (like tested dates etc). This would probably work well for me, since the application will give my class control where I can look up in a custom database table or repository extension for my custom rule properties and values. > That or we can allow admins to configure a properties file with a fixed list of selectors to choose from - which option sounds better? The admin configuration of the properties file would be perfect. The selective deployment script are tightly tied to the rule management process and rule life cycle. These will be defined one time by IT and probably see little change after that.	1	1	https://issues.jboss.org/browse/GUVNOR-888
5561	GUVNOR-882	need more info - logs etc... what is it doing and why can't it load up.	0	0	https://issues.jboss.org/browse/GUVNOR-882
5562	GUVNOR-863	Hi Tihomir, Did the ant task you did help with this?	0	0	https://issues.jboss.org/browse/GUVNOR-863
5563	GUVNOR-859	Could we have some more info on what you do on production BRMS? (ie why you need the full BRMS in the production environment, not just .pkg or drl files).	0	0	https://issues.jboss.org/browse/GUVNOR-859
5564	GUVNOR-859	Is this capability actually implemented? In which release? How does it work? Through the Admin Import / Export/	0	0	https://issues.jboss.org/browse/GUVNOR-859
5565	GUVNOR-859	Isn't this related to the bulk-import-util module, which we want to make obsolete by implementing that feature in the GWT gui?	0	0	https://issues.jboss.org/browse/GUVNOR-859
5566	GUVNOR-859	Summary of what need to fix to be able to partial imports and re-imports (=> sync 2 guvnor repo's): - Modules (= packages) need to be versioned - Global package is just a module, so it needs to be versioned too - Versioned dependencies between modules (which includes a dependency of a package to the global package). Restrictions for the user: - Imports can only occur on module level granularity. This is actually good for a user, as assets depends on each other and it makes sense to import them in that granularity atomically. - Modules need to be released (= get a fixed, unique, unalterable version) before they can be imported into another repository. In later versions we can look into supporting timestamped SNAPSHOTS to avoid the need for explicit releases. Once we have this metadata, the partial import is a matter of transferring the needed version (and it's dependency versions) from the sourceRepo to the destinationRepo. Things that make it harder: - categories reused across all modules - how do they get synced? - permission nodes reused across all modules - how do they get synced?	0	0	https://issues.jboss.org/browse/GUVNOR-859
5567	GUVNOR-763	Hi Toni. Able to take a look at this? I am sure you used to be able to do this - there should be a create new option for functions.... otherwise its just a GUI stuff up.	0	0	https://issues.jboss.org/browse/GUVNOR-763
5568	GUVNOR-763	Did this change go into M3? I've uploaded a function, but still don't see any way to modify nor verify the function. Attaching a screen shot. Am I missing something?	1	1	https://issues.jboss.org/browse/GUVNOR-763
5569	GUVNOR-761	Permission check has been added. And corresponding test cases.	0	0	https://issues.jboss.org/browse/GUVNOR-761
5570	GUVNOR-725	1. and 3. are something that could use the improvements that you are talking about. Not sure how to do 1. without breaking backwards compatibility. 2. Can you describe the situation where you are missing the date picker? Just in case it is something that we have missed.	0	0	https://issues.jboss.org/browse/GUVNOR-725
5571	GUVNOR-711	Is this an optional thing - ie an option to disable ruleflows - or just ALWAYS disable them for scenarios?	0	0	https://issues.jboss.org/browse/GUVNOR-711
5572	GUVNOR-710	mic, Is this fixed?	0	0	https://issues.jboss.org/browse/GUVNOR-710
5573	GUVNOR-628	Please find a translated 5.1 file. I don't have the 5.1 code to check labels usage (because 5.1 fails to compile my rules, see GUVNOR-578), so I have some problems regarding the usage context of some labels : - l.1011 : Causes=Causes : Is it used as a verb ("something causes something else") or as a noun ("The causes of X are Y"). I suppose this is the second option - l.1045 : TemplateKey=Template key : don't know how to translate ... Is key like 'identifier' or something like that ? Tagged all previous remarks with comments in the translated file Otherwise, a lot of technical term have been kept 'as this', or translated with siplificiation, instead of finding some weird academician translation understood only by language ayatolas (who won't use drools anyway)... For instance, "package", "working sets", "rule flow", "rule flow group"... Some others, like "upload", are quite difficult to translate. In common french, most of us simply uses the english word (with a strong french accent). For this one, I chose "Envoyer", which means "Send" literraly ... "	1	1	https://issues.jboss.org/browse/GUVNOR-628
5574	GUVNOR-622	Hello, here is my pull request. From the package screen, the snapshot creation takes in account the build selector configuration as for the build. What it does : 1) It copies the whole workspace in the snapshot (as before) 2) it removes all asset which do not apply to the selector configuration So the snapshot only contains the good asset list. A second thing I corrected is in the status list, selecting more than one status was not possible : Q1,Q2 was looking for a state "Q1,Q2". Now all asset that are in status "Q1" or Q2" are selected. Thanks Nicolas	1	1	https://issues.jboss.org/browse/GUVNOR-622
5575	GUVNOR-612	Was that for two separate guides or for creating a part1 and part2 in the guvnor guide ? There was some email discussion about this. I haven't been able to look at these docs for a few months. Also when is this required by ?	0	0	https://issues.jboss.org/browse/GUVNOR-612
5576	GUVNOR-612	Ok, I'll provide a patch for just a reorganization of the guide into 2 parts, Administrator Guide & User Guide. I'm pretty occupied with releases for the next couple of weeks, but it shouldn't take too long. You guys are in git now ? Should I fork https://github.com/droolsjbpm/guvnor.git do the work in master and send a pull request when done ? Are there any particular git rules I need to abide by ? Also is it cool if I re-org the files a little to ensure they are compatible with Publican ? It will make future productization easier.	1	1	https://issues.jboss.org/browse/GUVNOR-612
5577	GUVNOR-605	Could you add an additional checkbox (or a radio button for the 3 states?) that does reverse ordering (or negates the number, same effect)? That way the rules would be evaluated top-down, instead of bottom-up. Also, thank you for implementing this!	1	1	https://issues.jboss.org/browse/GUVNOR-605
5578	GUVNOR-602	Unfortunately neither GWT nor Mosaic provides a TabPanel close to what Gwtext provides. Firstly, the tabPanel from GWT/Mosaic looks different (of course), not as fancy as Gwtext tabs. For what does the TabPanel look like in Gwt and Mosaic, please check http://gwt.google.com/samples/Showcase/Showcase.html#!CwTabPanel and http://mosaic.analytical-labs.com/#CwTabLayoutPanel. Secondly, tab from GWT and Masaic does not have a close icon/button. Mosaic can remove a tab programmatically, but its not obvious where to put this close button. The second problem is a blocker for us.	0	0	https://issues.jboss.org/browse/GUVNOR-602
5579	GUVNOR-584	Do you have plans for another 5.1.x release? If so, do you have a timeframe? I would like to get a formal release with this patch included if possible since this bug make Oryx useless for FF users.	0	0	https://issues.jboss.org/browse/GUVNOR-584

#	Requirement	Questions	Answered	Successfully Answered	URL
5580	GUVNOR-584	Does anybody know when there be a stable version of Oryx that works fine with drools 5.1 and all browsers?	0	0	https://issues.jboss.org/browse/GUVNOR-584
5581	GUVNOR-581	Do we really need sub-package? If the reason why you need sub-package is to use inheritance, then similar result can be achieved by using dependency.	0	0	https://issues.jboss.org/browse/GUVNOR-581
5582	GUVNOR-581	This was mainly because of the display "problem" (the parent package appearing twice when one of the package in path is not empty). I think that 5.2 does not have the problem anymore But your dependency function is very interresting (did not know it exists). How do you achieve that ?	1	1	https://issues.jboss.org/browse/GUVNOR-581
5583	GUVNOR-581	Vincent, can you please check you are happy this issue is resolved in 5.2 so we can close this JIRA? Discussions for "dependency" should be discussed either on one of the mailing lists or IRC, but not the history of a specific JIRA. Thanks. Mike	1	1	https://issues.jboss.org/browse/GUVNOR-581
5584	GUVNOR-580	Hi, 5.0.1 was created 21st of May 2009. Is it possible to get a patch for a later release or the trunk? SVN URL for the trunk http://anonsvn.jboss.org/repos/labs/labs/jbossrules/trunk/	1	1	https://issues.jboss.org/browse/GUVNOR-580
5585	GUVNOR-580	Is it for testing or for include it in future releases ? If it is only for testing, may be the faster is to provide a war that works this way, and then see if you are happy with that. (The war is big, and I don't think it can be placed in the JIRA. Are you ok to download it from a FTP server ?) If you are already convinced, I will try to get time to do that (otherwise it will be on my own time and then delays are longer :)) Should I include my other modification (GUVNOR-576 , which has an impact on the service classes) ?	1	1	https://issues.jboss.org/browse/GUVNOR-580
5586	GUVNOR-580	I have tried to download from trunk After fighting with some wrong pom files (referencing badly their parent pom) and let maven do its stuff, I was facing 1 problem : - drools-bpmn2 does not compile (some unknow methods on XML objects) So I simply go and download the generated jar from your maven repo and make drools-guvnor depends from it instead of the project After that, drool-guvnor is OK (no java problems identified by eclipse) But when I try to compile (using latest GWT 2.0.4 eclipse plugin, under Helios), I got the error : java.lang.NoSuchFieldError: warningThreshold at com.google.gwt.dev.javac.JdtCompiler.getCompilerOptions(JdtCompiler.java:312) ... I have found some forum posts that identify this problem as a JDT version package in GWT ... (or something like that). I aplied a solution based on the one described at end of http://code.google.com/p/google-web-toolkit/issues/detail?id=4479, using a JDT 3.4 file downloaded from google code, and then compilation runs. At that time, I still don't know if it works, but I find all these solution really ugly. Do you have some tips to make that cleaner ?	1	1	https://issues.jboss.org/browse/GUVNOR-580
5587	GUVNOR-580	Compilation is OK. Then I tried hosted mode, with firefox. Got the message indicating that firefox has to download the GWT plugin, installed it ... and now the same error message appears as error in hosted mode console... So I suppose that the GWT plugin for firefox comes with the same bad JDT as GWT SDK ... How do I change this ?	1	1	https://issues.jboss.org/browse/GUVNOR-580
5588	GUVNOR-580	GWT plugins and Drools use a different version of the JDT. You can solve this problem by loading the GWT SDK before Drools modules in the class path. How to do this: Go to Properties -> Java Build Path -> Order and Export, select GWT SDK and click Top. Click OK. It should work now. Some more help in case you need it: http://community.jboss.org/wiki/RulesRepositoryBuildInstructions	1	1	https://issues.jboss.org/browse/GUVNOR-580
5589	GUVNOR-580	Is it the normal guvnor skin (see guvnor appearance capture above)? I think I have still a problem (no border around block in package view is suspicious). Any ideas ?	1	1	https://issues.jboss.org/browse/GUVNOR-580
5590	GUVNOR-580	Done, based on a source last updated on Wednesday (20th october). What is the most convenient for delivery ? I will post "svn diff" patches. Hope it will work even if you have done modifications since this date...	1	1	https://issues.jboss.org/browse/GUVNOR-580
5591	GUVNOR-541	Hi Esteban, I believe this has been completed (by you?) a long time ago. Can we close this JIRA? Cheers, Mike	1	1	https://issues.jboss.org/browse/GUVNOR-541
5592	GUVNOR-485	Work around exists for Drools 5.0: https://community.jboss.org/wiki/DroolsGuvnor50Installation Drools 5.1 M2 works with JBOSS AS 5.1 out-of-box.	0	0	https://issues.jboss.org/browse/GUVNOR-485
5593	GUVNOR-451	Hi Len, I don't really understand the title. Could you please attach a screen-shot or something?	1	1	https://issues.jboss.org/browse/GUVNOR-451
5594	GUVNOR-419	<jliu> mic_hat, regarding the shareable artifact stuff. <jliu> mic_hat, from the prospective of gui, what u think a user need to do . <jliu> mic_hat, how he/she operates to share an artifact <mic_hat> jliu: good question <jliu> mic_hat, what i have in my mind is: 1. every artifact has a perm link. <mic_hat> yes <jliu> 2. say the mortgage module under mortgage package. if we want to share mortgate module for testpackge, <mic_hat> I think you choose to create an item of type "link" - and then you get to browser for what you are linking to... <mic_hat> ok <jliu> what we need to do is to copy the perm link of this module. <mic_hat> yes <mic_hat> jar you mean <mic_hat> model even ;) <mic_hat> I know what you mean <jliu> then open testpackage, click add a new model. <mic_hat> but yes <mic_hat> and we have an option to add in a link <jliu> then a new model can be created from scratch or imported from a perm link <mic_hat> yes <mic_hat> that sounds good <mic_hat> alternative is to allow drag and drop from one package to another <jliu> mic_hat, drag and drop? never tried that. sounds a fancy feature. ;-) does gwt supports this? <mic_hat> yeah - its doable <mic_hat> easier then it sounds <mic_hat> just an idea <mic_hat> might be dumb <mic_hat> dunno <mic_hat> but you get the idea <jliu> mic_hat, we may want to avoid browsering stuff: "I think you choose to create an item of type "link" - and then you get to browser for what you are linking to." <mic_hat> yeah <jliu> as potentially the browser can bring too many artifacts on one screen. can be very slow.	0	0	https://issues.jboss.org/browse/GUVNOR-419
5595	GUVNOR-392	As a workaround, would it make sense for you to use status instead of category? Selecting by status works fine.	0	0	https://issues.jboss.org/browse/GUVNOR-392
5596	GUVNOR-392	Does this solve your issue? Looks like it is possible to filter for categories.	0	0	https://issues.jboss.org/browse/GUVNOR-392
5597	GUVNOR-370	Hi Prem, I've fixed the multiple rules problem in archieved manager. so that ppl can now delete/restore multiple rules in one go. But I am not sure what this jira is proposing. Did you mean the AssetItemGrid that is used to display search result as well as the AssetItemGrid that is used to display assets under a package or category? If this is what u mean, are you suggesting that we add delete/copy operations in AssetItemGrid? (at the moment, the AssetItemGrid used by search results and package/category only has two buttons: refresh list/open selected).	1	0	https://issues.jboss.org/browse/GUVNOR-370
5598	GUVNOR-357	It doesn't work as expected. The timestamp displayed after successful build is the time of the package's last modification BEFORE this build. I think it'd be better just to update the "Last Modified" field after a successful build. When a user checks the tab to see if the build has finished and sees the green mark and the message "Package built successfully." then he would obtain the time of build from the updated "Last Modified" field. What do you think?	0	0	https://issues.jboss.org/browse/GUVNOR-357

#	Requirement	Questions	Answered	Successfully Answered	URL
5599	GUVNOR-355	Also, this is the comment that gets written out to the server log when an upload happens: 09:13:21,477 INFO [STDOUT] DEBUG 25-06 09:13:21,476 (ServiceImplementation.java:loadRuleAsset:491) Load time taken for asset: 2 Could this be improved please with more details such as file name, package name, etc. It could also be a nice log for tracking various events in a friendly log4j format, as it has nice tools such as ChainSaw etc to watch them.	0	0	https://issues.jboss.org/browse/GUVNOR-355
5600	GUVNOR-352	Should now be better in every other browser than IE.	0	0	https://issues.jboss.org/browse/GUVNOR-352
5601	GUVNOR-338	Works for lists. Nicolas, can you give an example of what you were trying to do and what is missing? It is possible that I hid something when I refactored the codes.	1	1	https://issues.jboss.org/browse/GUVNOR-338
5602	GUVNOR-338	Hi Toni, I will do a few screen captures and put it in attachement. Is it ok of I do this tomorrow ?	1	1	https://issues.jboss.org/browse/GUVNOR-338
5603	GUVNOR-318	does this work in regular DRL - or do you mean just in GUI?	1	1	https://issues.jboss.org/browse/GUVNOR-318
5604	GUVNOR-239	"Also - this link should be immutable - always the same, even if we import it into another repo (ideally). But most importantly: should this link always include a version number so it always returns the same content?" how to make the link immutable? For example, the URL http://servername:9000/repository/atom will be changed if we switch to another host or change the port name. In this case, is it still immutable? we can include the version number. If we dont have the version number in the URL, that means the latest version. Does that make sense?	1	1	https://issues.jboss.org/browse/GUVNOR-239
5605	GUVNOR-165	Mic, do you still remember what this patch is for? I mean what problem this patch is to fix? is this patch still valid? Thanks.	0	0	https://issues.jboss.org/browse/GUVNOR-165
5606	GUVNOR-164	So this feature is described in http://community.jboss.org/wiki/OptionalExternalcontentstores Was this feature implemented? Any details?	0	0	https://issues.jboss.org/browse/GUVNOR-164
5607	GUVNOR-158	GWT is on 1.5.3. , Jackrabbit is on 1.6. Is there a need to upgrade GWT-ext? (As we are moving away from GWT-ext anyway).	1	1	https://issues.jboss.org/browse/GUVNOR-158
5608	GUVNOR-113	OK I made it so it would work - but it has to show the actual value in the grid (not ideal). So part way there at least...	0	0	https://issues.jboss.org/browse/GUVNOR-113
5609	GUVNOR-109	+1 - would be easier to understand and to use. Would the users ones still be preloaded? I mean, would the user have to reload the webapp to have the custom selectors show up in the drop down list? Couldn't selectors be part of a package, hence created like a regular asset? Sure, you would have to create several duplicate selectors for every particular package then...	0	0	https://issues.jboss.org/browse/GUVNOR-109
5610	GUVNOR-109	This is kind of related to: https://jira.jboss.org/jira/browse/GUVNOR-90 Yes - I don't see why it can't be stored as a resource - I think with enough built in ones mostly people won't need to tweak their own ones. Can you think of a list of common things that would be needed for deployment?	1	1	https://issues.jboss.org/browse/GUVNOR-109
5611	GUVNOR-91	Here is the issue for sharing assets between packages. Would be nice to create a new asset by linking to an existing one. when linking, can choose a soft link (follow the latest) or hard link (lock it to a specific version - ie copy ?) - but when opening the "linked" item, be able to follow the link to the real one.	0	0	https://issues.jboss.org/browse/GUVNOR-91
5612	GUVNOR-91	This is the proposed way how it works: 1. create a global area. an asset can belong to a package intially. When we want to share it, we prompt it to the global area using node.move(). 2. The package will have a depenceny link attribute. This dependency attribute points to an asset under global area. When u do a package build, we iterate through package's depencency list, retrieve corresponding asset from whatever location the dependency link points to, to form a package bundle. 3.The asset (who lives in the global area) potentially can have a method called getDependant() which returns UUIDs or URLs of these packages that have a depencency link to this asset. 4. As we changed jcr structure, we may need to provide a migrate tool on start up to migrate existing data. 5. We need to change GUI as well. Add a root node called "common" on the same level of where "packages" node stays.	0	0	https://issues.jboss.org/browse/GUVNOR-91
5613	GUVNOR-90	Should have GUI for selectors - selectors can select on status, category, meta data. Those are stored as "smart views" - can then use those "smart views" to build from. Optionally, if rule in "wrong" status, can either warn, or ignore, or use "last known good version".	0	0	https://issues.jboss.org/browse/GUVNOR-90
5614	GUVNOR-86	Wow thanks david ! Patch applied. I am still in 2 minds on this - will the appearance of "this" confuse people? if all the other names have intrinsic meaning - "this" really only makes sense to a java programmer (or javascript, perhaps ?). Thoughts?	1	1	https://issues.jboss.org/browse/GUVNOR-86
5615	GUVNOR-84	I couldn't find it in the docs: https://hudson.jboss.org/hudson/job/drools/lastSuccessfulBuild/artifact/trunk/target/docs/drools-guvnor/html_single/index.html Is it possible to disable WebDAV access (GUI, xml)?	0	0	https://issues.jboss.org/browse/GUVNOR-84
5616	GUVNOR-2	Do you have any further explanation of this? I'm looking into the WebDav interface at the moment, and have some ideas, but don't want to re-state goals/tasks that are already ongoing.	1	1	https://issues.jboss.org/browse/GUVNOR-2
5617	GUVNORSOA-12	If the metadata type is in use, I think people would expect a warning to the effect of "This metadata type is being used on N artifacts. Are you sure you want to delete it?". If they still choose to delete the metadata, the system would clean it up (i.e., remove that property from any node that contains it).	0	0	https://issues.jboss.org/browse/GUVNORSOA-12
5618	HAL-1279	Are you sure it's already on HAL.next? I can't find it on any attribute.	1	1	https://issues.jboss.org/browse/HAL-1279
5619	HAL-1279	I see it now, thanks. I was thinking about the Management Model. In HAL.next the access type is shown. If there's read-write access for an attribute, a new text (requires reload/requires restart) could me added. The same would apply for operations. Should I move this discussion to the .next repository?	0	0	https://issues.jboss.org/browse/HAL-1279
5620	HAL-1278	I'm a bit confused about you not having the time when I'm the one trying to fix those issues (albeit not directly) by working on this. Is hal.next planned to be included in 7.2?	1	0	https://issues.jboss.org/browse/HAL-1278
5621	HAL-1253	I've suggested a modification of the transaction subsystem to add a warning to the operation response's header: {code:json} [standalone@localhost:9990 /] /subsystem=transactions:write-attribute(name=use-jdbc-store, value=true) { "outcome" => "failed", "response-headers" => { "warnings" => [{ "warning" => "Please note that only journal store or jdbc store can be used, thus 'use-journal-store' has been disabled.", "level" => "WARNING", "operation" => { "address" => [("subsystem" => "transactions")], "operation" => "write-attribute" } }], "process-state" => "restart-required" }, "failure-description" => "WFLYTX0032: jdbc-store-datasource must be defined if use-jdbc-store is 'true'.", "rolled-back" => true } {code} However, this does not lead to a message being displayed in HAL. Do you I need to use a specific message template for the message to be displayed in HAL ? Or is there some change required in HAL too ?	1	1	https://issues.jboss.org/browse/HAL-1253
5622	HAL-1226	Hi [~jstefl], can you provide a reproducer ? I did some deploys (managed and unmanaged), everything works.	1	0	https://issues.jboss.org/browse/HAL-1226
5623	HAL-1216	[~harald.pehl], TextBoxItem is a bit confusing, the message says "No whitespace, no special chars allowed" however the item doesn't check for any special characters apart from spaces. Does the validation need improving?	1	1	https://issues.jboss.org/browse/HAL-1216

#	Requirement	Questions	Answered	Successfully Answered	URL
5624	HAL-1216	What kind of improvement do you have in mind?	1	1	https://issues.jboss.org/browse/HAL-1216
5625	HAL-1198	I was thinking about it as well. It would be HAL specific added value so this would require HAL specific configuration in model. Is it doable? The configuration should be persistent and available for all users so we cannot use local browser storage.	0	0	https://issues.jboss.org/browse/HAL-1198
5626	HAL-1177	[~claudio4j], I see, I'll get on it. Should I file a new issue?	1	1	https://issues.jboss.org/browse/HAL-1177
5627	HAL-1177	Reverted the previous changes and moved the management to runtime under the Messaging subsystem. As the subsystem now shows more than metrics should the descriptions be changed? (if so to what?)	0	0	https://issues.jboss.org/browse/HAL-1177
5628	HAL-1147	> Claudio Miranda shouldn't this be marked as resolved? The wildfly upstream feature branch is not merged yet https://github.com/wildfly/wildfly-core/pull/1675 https://github.com/wildfly/wildfly/pull/9063 So the feature branch for HAL was not submitted https://github.com/claudio4j/hal-core/tree/HAL-1147	1	1	https://issues.jboss.org/browse/HAL-1147
5629	HAL-1141	Please set "Fix Version/s" for tracking purposes. btw, what's the ETA for 2.9.0 release + inclusion in WF master ?	1	0	https://issues.jboss.org/browse/HAL-1141
5630	HAL-1091	The server configuration in domain mode is able to configure some attributes (socket binding, port offset, server group), however those attributes doesn't exists in standalone mode. What do you think should be there in view mode for standalone ?	0	0	https://issues.jboss.org/browse/HAL-1091
5631	HAL-1064	So what about filtering by regex? May be with possibility to predefine them?	1	0	https://issues.jboss.org/browse/HAL-1064
5632	HAL-1056	Do you mean last field in same step 3/3 or in the last summary page before clicking finish button ?	1	1	https://issues.jboss.org/browse/HAL-1056
5633	HAL-976	Just one more suggestion, as I see work hasn't been finished. If I understand correctly, subdeployments are only relevant for EAR's. In that case, could it be disabled if the deployment is not an EAR?	1	1	https://issues.jboss.org/browse/HAL-976
5634	HAL-972	This has been merged into master but it's still missing on 2.8.x branch. Could you please apply the patch to it?	0	0	https://issues.jboss.org/browse/HAL-972
5635	HAL-972	The problem is, if i replace a deployment, the previously choosen 'runtime name' (in my case i choosed ROOT.war when i deploeyd the first time) isnt used again for the replacement. Instead, the runtime name is generated out of the war file i use to replace. do i am missing sth.? to replace a deployment i have to remove and re-add the deployment to be able to set the correct runtime name.	0	0	https://issues.jboss.org/browse/HAL-972
5636	HAL-738	[~heiko.braun] In my opinon, actions like Assign, View are the primary actions in the Finder view and it'd be better that we always show them. As a result, users can browse the preview content and still see the primary button in the column, so they know what they could do in the next step. One reason of bringing this up: I gave the prototype to one another designer in our team who run the same usability test at Summit. He said that when he played around with the prototype he accessed the Application view at the very beginning and then he forgot how to access it again and at that time he was locked in the Finder view. I think this might happen to a lot novice users who are not familiar with the Finder and Application model and we should make the primary buttons more prominent if that won't harm anything else on the screen. What do you think?	1	1	https://issues.jboss.org/browse/HAL-738
5637	HAL-621	Will this also be in the EAP 6.4 release?	0	0	https://issues.jboss.org/browse/HAL-621
5638	HAL-604	What's the content for the both Domain and standalone runtime?	1	1	https://issues.jboss.org/browse/HAL-604
5639	HAL-564	This seems to require a restart to change ANYTHING about a datasource. What happens to the Admin console ability to edit a datasource for testing/tuning runtime perfomance before making them permanent in the XML.?	1	0	https://issues.jboss.org/browse/HAL-564
5640	HAL-551	so, whats the issue here? do you guys want to keep re-assigning it? jim, please elaborate. i've said "there is a test connection button". what else do you need? what's the real issue?	1	0	https://issues.jboss.org/browse/HAL-551
5641	HAL-551	Where in the set of error messages is the problem in the JDBC driver in that I did not supply a password. Instead of a blank password mistype it. I get no idea in the thread dumps that I fat fingered the password. I would expect to see a MySQL driver exception around the lack of authenticating the database, but that never occurred in a reasonable fashion as it used to. About 1/3 of the way in the stack dump I get this: [Server:server-one] at java.lang.Thread.run(Thread.java:662) [rt.jar:1.6.0_32] [Server:server-one] Caused by: java.sql.SQLException: Access denied for user 'root'@'localhost' (using password: NO) But how is a user in a console, with a test connection button that didn't show this error supposed to find this. I would expect in domain mode, I might not have access to the local console, it is only through the contrived usage of this lab that the console is front and center. I am pretty sure I also got into a state where the test console button got this into a place I could not update this, as I first at to undeploy the datasource, make a change, and then redeploy it. Why can't I test it before it deploys? I just think there is/was some usability around datasources that needs to be thought through. This is the one area in the console someone will hit every time they use the console. Jim	1	0	https://issues.jboss.org/browse/HAL-551
5642	HAL-551	Jesper, would a failure during prefill result in a StartException from a service?	1	0	https://issues.jboss.org/browse/HAL-551
5643	HAL-551	So - everyone including Jim - is the answer to provide an optional prefill? (It sounds like we already have that.) Do we just need to document this more clearly for the users? I'm guessing that making this a mandatory prefill is way too restrictive, right?	0	0	https://issues.jboss.org/browse/HAL-551
5644	HAL-551	Comment from Jim: I do see your comment, but the problem to me is that I need to be "more locally" told that the password I put in is incorrect or not valid. Having that info buried in logs that when using the console I typically can not see is a big problem. Question - How can the error be displayed directly to the user via the console?	0	0	https://issues.jboss.org/browse/HAL-551
5645	HAL-551	[~jesper.pedersen] What's the conclusion on this? Should we close the issue?	1	0	https://issues.jboss.org/browse/HAL-551
5646	HAL-551	As the problem is related to the datasource usage feedback, there is nothing the web console could do to correctly show the cause. The datasource subsystem must provide the cause.To provide some context, I did three tests and its results, as logged in server.log 1) server not started {code}Caused by: org.postgresql.util.PSQLException: Connection refused. Check that the hostname and port are correct and that the postmaster is accepting TCP/IP connections.{code} 2) wrong username {code}Caused by: org.postgresql.util.PSQLException: FATAL: autentica????o do tipo password falhou para usu??rio "admin"{code} 3) wrong password {code}Caused by: org.postgresql.util.PSQLException: FATAL: autentica????o do tipo password falhou para usu??rio "timekeeper"{code} But the reason displayed by HAL and CLI are the same {code} "failure-description" => "WFLYJCA0040: failed to invoke operation: WFLYJCA0047: Connection is not valid", {code} The main issue is that the original cause is buried in log, do not show in the webconsole or CLI. The user must open the log and search for the cause. Therefore I suggest to open an issue for datasource subsystem (to show the original cause) and let this issue depends on that.	0	0	https://issues.jboss.org/browse/HAL-551
5647	HAL-550	Can we expect anything for 7.1? Otherwise we might re-schedule this issue.	0	0	https://issues.jboss.org/browse/HAL-550
5648	HAL-547	Should I report this issue to HAL project instead?	1	1	https://issues.jboss.org/browse/HAL-547
5649	HAL-547	Not reproducible in wildfly 10, Rémi, can you try it with wildfly 10 ?	1	1	https://issues.jboss.org/browse/HAL-547

#	Requirement	Questions	Answered	Successfully Answered	URL
5650	HAL-547	I see your point. I didn't know about the specific setting {{jbossws.undefined.host}}. I always set the reverse proxy hostname in {{wsdl-host}}, can you set this on your webservices configuration ? This is not a HAL (web console) bug, but a wildfly issue. I will talk to the team, about this.	1	1	https://issues.jboss.org/browse/HAL-547
5651	HAL-531	Hi [~hbraun]! Do you think such feature is not useful for now in Management Console?	1	1	https://issues.jboss.org/browse/HAL-531
5652	HAL-503	I can understand the use case and it make complete sense. But wouldn't the actual server name be sufficient? Why do you think do we need an extra title attribute?	1	1	https://issues.jboss.org/browse/HAL-503
5653	HAL-459	Can this be resolved?	1	1	https://issues.jboss.org/browse/HAL-459
5654	HAL-424	duplicate of #HAL-391 ?	1	0	https://issues.jboss.org/browse/HAL-424
5655	HAL-408	Can you attach a screenshot please? Also, you note in the JIRA that this is AS 7.1.1.Final. Have you tried against the latest released WildFly (JBoss AS has been renamed to WildFly) version http://wildfly.org/downloads/	0	0	https://issues.jboss.org/browse/HAL-408
5656	HAL-398	Can we close this as a duplicate of HAL-390?	1	0	https://issues.jboss.org/browse/HAL-398
5657	HAL-390	There are 39 screens to scroll through and there is no way to sort or filter and it isn't even in order. 5. How would the user like to achieve this? (List the functional requirements here) a) Either increase the list (or) b) Add a Filter/Sorter to sort the list	1	1	https://issues.jboss.org/browse/HAL-390
5658	HAL-380	Is this available in EAP 6.x?	0	0	https://issues.jboss.org/browse/HAL-380
5659	HAL-372	I don't know what your are talking about. Can you elaborate?	1	1	https://issues.jboss.org/browse/HAL-372
5660	HAL-372	As you can see the Server Runtime instances shows as expected, the server and profiles list do not work as shown. Is this a firefox issue or a gwt issue. If we know something does not work we should pop up a warning telling people that this combo does not work. Does this make sense?	1	0	https://issues.jboss.org/browse/HAL-372
5661	HAL-372	can you verify this happens with a clean browser cache?	0	0	https://issues.jboss.org/browse/HAL-372
5662	HAL-367	What "startup options"?	1	1	https://issues.jboss.org/browse/HAL-367
5663	HAL-367	Can you elaborate on this?	1	1	https://issues.jboss.org/browse/HAL-367
5664	HAL-367	The domain.xml has an ability to define a series of services that are started, for slimming purposes it would be nice to expose these in the UI. Data sources are set/stored in this file. Really wondering out loud if we are hiding too much of the mapping between domain.xml and what is in the UI. The UI does not call out domain.xml or that notion in any meaningful way. Does that make sense.	1	0	https://issues.jboss.org/browse/HAL-367
5665	HAL-363	Is there a system I can reference to review this?	1	1	https://issues.jboss.org/browse/HAL-363
5666	HAL-363	Thanks Heiko. I took a look and overall I really like the idea behind this page and the visual layout of the content. The issues I noted are below from most severe to least severe, and I also attached a PDF with comments embedded so the comments can be seen in context of the element they are referring to. I can turn these into new separate JIRAs if that would help. 1. The pagination jumps up and down the page vertically as you use it based on the length of the content of that page. This makes the user "chase" the pagination rather than being able to leave their mouse in one place and click it without really looking at it to jump through multiple pages. Possible solutions: - have pagination at the top of the table as well - Make the table fixed size and it will scroll vertically if there is more content than can fit, and whitespace at the bottom if there is less. (not a visually pleasing option) - Use a horizontal scroll bar to scroll through all of the hosts rather than pagination. This means the table will automatically be loaded at the maximum length it requires. (SUGGESTED) 2. The page does not meet 508 compliance. The table is not tagged appropriately for headers and content, the actions are not actionable in non-styled mode which screen readers use (disable CSS), and the layout of content in non-CSS view would likely not translate well. 3. The 'Stop Group, Start Group, Restart Group' functionality in the groups row headers are hidden and some users might miss this feature. Options: - Always show options, but this adds a lot of clutter when they are not always needed - Add an "Actions" icon or dropdown that allows users to hover over/click and choose these actions. (SUGGESTED) 4. More likely that users will end up hovering over the server instance blocks and see the actions show up, but still a bit hidden. Suggestions: - Maybe use the same "Actions" icon/dropdown as in the row header. 5. One column header has a star. What does this star mean? It doesn't show up on every controller. Is this the "master" or something? A hover-over label of this star should be added for better understanding. 6. When a "Restart is required" (text from hover-over of an exclamation point icon) - why don't we have the Restart action rather than the stop action?	1	1	https://issues.jboss.org/browse/HAL-363
5667	HAL-354	I would like to help on this. There is the following attribute, should be displayed alongside each server name, see attachment. It will show in units: seconds, minutes, hours, days /host=master/server=srv-1/core-service=platform-mbean/type=runtime:read-attribute(name=uptime) Is that good ?	1	1	https://issues.jboss.org/browse/HAL-354
5668	HAL-354	Should it be displayed alongside each server in topology view or in jvm runtime view ? I think the server topology view can display the uptime for all servers. If displayed in the jvm runtime view, the user must choose each server individually.	1	1	https://issues.jboss.org/browse/HAL-354
5669	HAL-354	Should the uptime be displayed under "Operating System" and "Processors" ?	1	1	https://issues.jboss.org/browse/HAL-354
5670	HAL-322	You mean on the <svg .../> element itself, or the nested elements (<g><rect/></g>) as well? What exactly do you need to test? Just the existence of the svg elements?	0	0	https://issues.jboss.org/browse/HAL-322
5671	HAL-311	Why was this closed? I don't see the feature in Wildfly 8.1 either.	0	0	https://issues.jboss.org/browse/HAL-311
5672	HAL-293	Exposing the provider attribute in the console, leads to the question whether and how to restart the domain and refresh the console in order to make use of the new provider. I can see several options: # Let the user change the provider, but do not automatically restart. Instead notify the user that she must manually restart everything and after that refresh the console. Most of the work will be left over to the user. In addition the steps to activate the new provider need to be executed in a given order. # Let the user change the provider and ask whether to restart all hosts. However restarting *all* hosts (particularly in large domains) is an operation which might be problematic to do using the console. It might be that the restart must happen in a given order (first DC than all other host - or vice versa), some hosts might not come up again, there might be a network error, etc. It's hard to control the restart using the console. # Show the current provider and give the user some kind of checklist what steps are necessary to change the provider. Changing the provider is quite a complex workflow which needs some preparation and one of the last steps in this workflow is to change the provider attribute. So all in all I think we should go with option 3. WDYT?	0	0	https://issues.jboss.org/browse/HAL-293
5673	HAL-265	Should it better to submit a single PR with all enhancements or a individual PR for each enhancement ? I prefer a single PR, as most of them enhancements are very small (2 to 5 lines).	0	0	https://issues.jboss.org/browse/HAL-265
5674	HAL-196	Is this done?	1	1	https://issues.jboss.org/browse/HAL-196

#	Requirement	Questions	Answered	Successfully Answered	URL
5675	HAL-174	Brian, can you comment on the last section "system properties (addressable=false)" and then re-assign to me?	1	1	https://issues.jboss.org/browse/HAL-174
5676	HAL-174	Can you explain what you mean by "cannot access the system properties (addressable=false)"? There are different system property resources; children of the server group and otherwise, so I'm not clear on the context. And from looking at the SensitivityClassification they should be addressable. I can see how the wildcard handler would lead to issues. I'll think a bit on that one, but probably won't be fixed today.	0	0	https://issues.jboss.org/browse/HAL-174
5677	HAL-98	addition of runtime metrics for runtime services in JGroups (channels) and Infinispan (cache-managers, caches) - addition of cluster-related information and operations (cluster membership, shared transport, other cluster-specific operations) 1.1 For JGroups: /subsystem=jgroups/stack=* /subsystem=jgroups/channel=* /subsystem=jgroups/cluster=* The stack resource exists already and represents a persistent stack configuration. It exposes transport and protocol configuration attributes. The channel resource will be new and each child resource instance represents a run-time Channel instance created from a given stack. It will expose transport and protocol read-only,runtime metrics as well as read-only versions of transport and protocol configuration attributes from the associated stack (a la JMX in previous versions of AS). The cluster resource will represent a cluster (a set of JGroups channels, possibly on different nodes). Attributes of a cluster may expose the group view or other cluster-wide attributes; operations may perform cluster-wide operations. Cluster children represent cluster members and expose the channel=* resources on that node. /subsystem=jgroups/channel=X /subsystem=jgroups/cluster=X/member-address=192.168.0.100 /subsystem=jgroups/cluster=X/member-address=192.168.0.101 The way in which cluster members are identified may require more than IP address (e.g. two JGroups channels on the same node). Cluster differs from channel in that channel exposes only the JGroups resources on a single node, whereas cluster exposes resources across the cluster. The user may thus view cluster-wide information from a single CLI window. A JGroups RPC mechanism and appropriate handler can be used to obtain management info from other nodes, in domain or standalone modes. Because JGroups starts services on-demand, channels and clusters need to be implemented as custom resources which are presented iff the runtime services which back them up have been started. 1.2 For Infinispan /subsystem=infinispan/cache-manager=* /subsystem=infinispan/cache-manager=*/local-cache=* /subsystem=infinispan/cache-manager=*/invalidation-cache=* /subsystem=infinispan/cache-manager=*/replicated-cache=* /subsystem=infinispan/cache-manager=*/distributed-cache=* /subsystem=infinispan/cluster=* The cache-manager and x-cache resources already exist and represent persistent cache manager and cache configurations. Unlike channels, which are represent runtime instances and whose configuration is tied to a stack definition, cache-managers and caches can represent both a configuration as well as a run-time instance. So existing cache-manager and cache resources may be extended with read-only, runtime metrics without defining any new resources. As in the case of JGroups, an Infinispan cluster can represent cluster-wide attributes associated with the cluster (e.g. transport) and cluster wide operations. Cluster children represent cluster members and expose the cache-manager=* resources (and cache children) on that node. /subsystem=infinispan/cluster=X /subsystem=infinispan/cluster=X/member-address=192.168.0.100 /subsystem=infinispan/cluster=X/member-address=192.168.0.101 Why have a cluster=* resource defined in two subsystems? An Infinispan cluster presents a view of a cluster from the Infinispan level(transactions, evictions, etc); a JGroups cluster from the JGroups level(RPCs, bytes sent, etc). 1.3 Comments Introduction of the cluster resource is experimental only and presented now for discussion. I think its natural for a user to want to see what is happening in a cluster, as opposed to only on a single node. Feedback appreciated. Introduction of the channel resources and extension of the cache-manager and cache resources is planned and needs to be completed asap. 2. Implementation Issues (Brian, Kabir, Tomas) 2.1 Paul and I discussed the placement of channel=* and felt that locating them on the same level as stack=* was better than as a child, as when looking at a channel, we don't really care which stack is being used. This is instead of locating channel as children of stack. Any objections? 2.2 In the case of channel, I would like to expose all of the stack attributes (transport/protocols) in read-only fashion, and append the runtime metrics to that. Why? So that (i) for example, all protocol=MPING information, both	1	1	https://issues.jboss.org/browse/HAL-98
5678	HAL-98	Dennis, thanks for the feedback! > That should already be available in domain mode, or using any other tool (JON) that looks at multiple nodes. At least for now, IIUC, a server group and a cluster are really orthogonal concepts. A server group is a set of servers, which may or may not be clustered, and a cluster is also a set of servers, which may or may not equate to a single server group. Whether or not the notion of a cluster as described above is rejected until such time as it gets integrated into domain mode is another issue altogether. Brian? Also, there are already features in the management api (starting/stopping server groups, deploying to server groups) which look at multiple nodes and also appear in JON. This would simply be yet another. Not everyone who wants to manage servers as a group need use JON, right?. > And it's not specific to the JGroups/Infinispan subsystems. The same argument of wanting to see data across nodes would apply to most subsystems. Hmm. I think the need is greater when the nodes interact - as in a cluster - where the behaviour of one node can directly impact the other. So can we say this applies to all clustering subsystems, as opposed to subsystems in general? > I don't think that functionality should be duplicated in a different way. Agreed. > Don't duplicate the information already available under stack=! The only reason for duplicating the configuration info was to avoid the user having to swap back and forth between the configuration for a channel (in the stack=* resource) and the runtime information for the channel (in the channel=* resource). This was the state of affairs in the old JMX view of channels where the two were integrated. You could see all configuration elements as well as runtime elements. If you feel its too much, it could be left out. > What would be useful under the runtime part would be exposing the actual JGroups protocol info (the current data, which could be different from the configuration if things were changed at runtime). This is what is planned - the channel=* resource's purpose is to expose the runtime info of each channel instance. > It would also be nice to see the runtime part not be limited to channels created by the AS JGroups subsystem, but also include any custom JGroups channels (maybe limited to those channels using the shipped JGroups version to avoid version mismatch complications). IIUC, only EJB, ISPN and other clustering subsystems should be creating channels and this is done via the ChannelFactoryService. The channel=* resources will expose all channels so created.	1	1	https://issues.jboss.org/browse/HAL-98

#	Requirement	Questions	Answered	Successfully Answered	URL
5679	HAL-98	> At least for now, IIUC, a server group and a cluster are really orthogonal concepts. Yes, my argument against the cluster=* section is really that management and clustering should stay orthogonal. Leave the management tasks (including aggregating data across nodes) to the management layer, unrelated to the cluster. > Hmm. I think the need is greater when the nodes interact - as in a cluster - where the behaviour of one node can directly impact the other. > So can we say this applies to all clustering subsystems, as opposed to subsystems in general? I agree that you may care more about data across nodes that interact. But not that it's limited to clustering subsystems. As an example, maybe you'd want to look at the stats of JDBC pools across the cluster (an undersized pool could end up bringing down the cluster). I agree that the functionality is wanted, just at a higher level in the management layer, not specific to clusters or clustering subsystems. > The only reason for duplicating the configuration info was to avoid the user having to swap back and forth between the configuration > for a channel (in the stack=* resource) and the runtime information for the channel (in the channel=* resource). I don't have anything against putting the configuration and runtime data near each other. I'm just against the duplication of the same data in different places (more work to maintain, QA, document, support, etc). I'd like to see it in just one place (wherever that is). > This is what is planned - the channel=* resource's purpose is to expose the runtime info of each channel instance. Sorry, my explanation was kind of vague in this section. What I was specifically referring to was to make sure we show the runtime version of the protocol list and configurations, as opposed to just a copy of the stack= configurations. Quite a few protocol settings could be changed at runtime for a specific channel, and protocols themselves could even be added or removed (don't think we currently do that in AS subsystems, but JGroups supports it). > IIUC, only EJB, ISPN and other clustering subsystems should be creating channels and this is done via the ChannelFactoryService. > The channel=* resources will expose all channels so created. I was mainly referring to any channels used by customer code. We do have a handful of customers using JGroups for other things outside the AS uses. (and that one would really be a "nice to have", certainly not critical).	0	0	https://issues.jboss.org/browse/HAL-98
5680	HAL-98	Re Richard's comment at 05/Dec/12 9:18 PM: * As you mentioned in 1.1/1.2, we're exposing a JGroups and Infinispan cluster. But, perhaps, we should expose the cluster as a top level subsystem, and JGroups channels / Infinispan caches as children ? E.g. /subsystem=cluster/name=session-cluster/channel* /subsystem=cluster/name=session-cluster/caches=Session*/attr=numOwners * With probe.sh, we can currently modify attributes and invoke operations cluster-wide. Is this something you want to support, too ?	1	1	https://issues.jboss.org/browse/HAL-98
5681	HAL-96	Is this still being considered ?	0	0	https://issues.jboss.org/browse/HAL-96
5682	HAL-96	Current state: - Client side: Deprecated branch https://github.com/hpehl/console/tree/cloud needs to be copied into master / new HAL-96 branch. This old branch contains the server selection described here: http://hpehl.info/independent-jboss-admin-console.html - Server side: Branch https://github.com/hpehl/wildfly/tree/keymaker is pretty much wip. The basic implementation for keymaker is in place. What's missing is the configuration and the precedence of digest over keymaker PRs have yet to be created.	1	1	https://issues.jboss.org/browse/HAL-96
5683	HAL-64	from #wildfly-management hpehl: darranl: If a user is assigned to multiple roles, which role will be used? Or will all roles be taken into account? darranl: hpehl, at the moment I believe the default is actually multiple roles will be used, the role mapping maps all it can and then there is a config option for the authorizer to decide if this is acceptable (although the config attribute may not have been exposed) darranl: hpehl, will need to check with Brian I think as we were going to support both modes - personally I don't see a reason to prevent a user having multiple roles ... hpehl: darranl: Is there an operation to read all roles a user belongs to? darranl: hpehl, I am going to be updating the whoami operation shortly to report the current users roles	1	1	https://issues.jboss.org/browse/HAL-64
5684	HAWKULAR-1207	@hrupp - you were involved in this on the agent side I believe. Do you think we need to create another agent-level property called "container-id" for this? Or am I missing something?	0	0	https://issues.jboss.org/browse/HAWKULAR-1207
5685	HAWKULAR-1202	I understand these flags should be set in an wildfly/eap container with the agent installed. Should these also be set in the hawkular-services container?	1	1	https://issues.jboss.org/browse/HAWKULAR-1202
5686	HAWKULAR-1200	How MiQ can manage certificates? Even if the certificate is pre-created (or created using MiQ) and stored somewhere in the MiQ, how to propagate it to the hawkular server and import to the keystore/truststore file or vault? With ansible? Or this is the jira talking about? I.e. via mounting the volume into docker then it wouldn't work for the normal case when we don't use any docker, is this desired? So is there a silent assumption that it will work only for containerized h-services at first? What about OpenShift? Can we assume it's running in openshift so that the Secret could be used?	1	1	https://issues.jboss.org/browse/HAWKULAR-1200
5687	HAWKULAR-1189	[~pilhuhn], What is the difference between this and HAWKULAR-1188? I sent a PR to 1188 (waiting for your review) but maybe it should have been this one?	1	1	https://issues.jboss.org/browse/HAWKULAR-1189
5688	HAWKULAR-1183	On this I was wondering if: * ManageIQ "detect" if a MW server is running in a container and then "hide" the operations or * We configure the agent with a new property "immutable:yes/no" and then let ManageIQ decide to show/hide operations IMO the second solution may be better since: * It's a bit more generic (an agent on baremetal may need to be immutable) * More "secure", the agent can forbid operations that happen through the command gateway * More "flexible", if someone really want a container to be mutable (knowing he would lose his changes on new image deployment), he could do so... What do you think ? [~mazz] [~pilhuhn] [~jkremser]	1	1	https://issues.jboss.org/browse/HAWKULAR-1183
5689	HAWKULAR-1183	There might be also a hybrid (container and bare metal) case in case of managed domain. Where one host can live in container and other doesn't. In this case it's safer to set the immutable flag I guess. I like the idea of making this configurable rather than implicit auto-magic. Having this immutable flag hard coded during the image build time is one option, other option would be the system env property that is read by the agent. Oh and there is also the jolokia scenario. Correct me if I am wrong, but so far it's only read-only pull-based mechanism so this should be ok, right? If we support also invoking some jmx ops via the jolokia, perhaps it make sense also to introduce some attribute in that remote-jmx tag (https://git.io/vDGyb) !20170203_124020.jpg|thumbnail!	1	1	https://issues.jboss.org/browse/HAWKULAR-1183
5690	HAWKULAR-1183	This makes sense to me. It'd be easy to run the WF with immutable agent (-Dhawkular.agent.immutable=true) and with a little bit more effort this JVM param can be added if the env property is defined in the system. So that one can run both mutable and immutable containers based on the env property (I can add that.) [~mazz] and how to read this flag in MiQ? Can you store it as a configuration of the feed itself in the inventory? Or some additional metadata in the /status endpoint? Hmm, the latter doesn't make much sense, because 1 h-services can have multiple agents connected that are either mutable or immutable.	1	1	https://issues.jboss.org/browse/HAWKULAR-1183

#	Requirement	Questions	Answered	Successfully Answered	URL
5691	HAWKULAR-1183	[~jkremser] > and how to read this flag in MiQ? I could add this as a resource configuration property on the agent resource itself - https://github.com/hawkular/hawkular-agent/blob/master/hawkular-wildfly-agent-feature-pack/src/main/resources/subsystem-templates/hawkular-wildfly-agent.xml#L824 Or I could add a CLI operation to the agent "getImmutable" which you could execute via the command gateway ExecuteOperation command. Let me know which you prefer. [~pilhuhn] > In MiQ we need both immutable and in-container. that is a separate JIRA: https://issues.jboss.org/browse/HAWKULAR-1185	1	1	https://issues.jboss.org/browse/HAWKULAR-1183
5692	HAWKULAR-1183	Why doing a shutdown of EAP inside a container ?	0	0	https://issues.jboss.org/browse/HAWKULAR-1183
5693	HAWKULAR-1183	I'm going to make all power ops the agent supports to be disabled if immutable. I think it would be confusing to people to have SOME power ops disabled but not others. "So, if the agent is in a container, I can shut down EAP but not restart it. Or is it I can restart it but not shut it down? Can I reload it or not?" For consistency, if we want some power ops to be disabled, all of them should, IMO. The user should go through the container to perform power ops (if you want to restart EAP, restart the container. If you want to shut it down, shutdown the container).	0	0	https://issues.jboss.org/browse/HAWKULAR-1183
5694	HAWKULAR-1182	Where the yaml template should go? 1) new repo jboss-dockerfiles/hawkular-services 2) hawkular/hawkular-services	1	1	https://issues.jboss.org/browse/HAWKULAR-1182
5695	HAWKULAR-1182	Put that somewhere in hawkular/hawkular-services and make it work with the pilhuhn/hawkular-services for now. Once we have automated hawkular/hawkular-services builds we switch over. Create a parameter for that image (with the pilhuhn/.. one as default)	1	1	https://issues.jboss.org/browse/HAWKULAR-1182
5696	HAWKULAR-1176	[~pilhuhn] that is certainly possible, but that is what I meant when I said, "won't cause us to do alot of work and to alot of refactoring and splitting out code. " - because what you propose is not how the current backport_eap6 branch is structured. We'd have to do a lot more work to get things refactored out into separate parts. We can do it, but it would take time. So, the question is: How much time do we have? When are we required to have an EAP6 agent build and EAP7 agent build (with all the installer enhancements that need to go with them?)	1	1	https://issues.jboss.org/browse/HAWKULAR-1176
5697	HAWKULAR-1161	Can someone add details to this JIRA? The agent already has two sets of metrics that appear relevant - "WildFly Aggregated Web Metrics" and "Undertow Metrics" which include the following: WildFly Aggregated Web Metrics: - Aggregated Active Web Sessions - Aggregated Max Active Web Sessions - Aggregated Expired Web Sessions - Aggregated Rejected Web Sessions - Aggregated Servlet Request Time - Aggregated Servlet Request Count Undertow Metrics: - Active Sessions - Sessions Created - Expired Sessions - Rejected Sessions - Max Active Sessions See: https://github.com/hawkular/hawkular-agent/blob/master/hawkular-javaagent-wildfly-feature-pack/src/main/resources/featurepack/content/standalone/configuration/hawkular-javaagent-config.yaml#L93-L154	1	1	https://issues.jboss.org/browse/HAWKULAR-1161
5698	HAWKULAR-1161	@[~mazz] iiic based on your comments above, the agent already collects the required metrics in both eap 6 and 7. is this correct? Should there anything additional be done on agent/hawkular side to report/enable those metrics (so miq can request/collect them), or this task is completed on agent/hawkular side?	1	0	https://issues.jboss.org/browse/HAWKULAR-1161
5699	HAWKULAR-1155	[~karel.hala] [~pilhuhn] [~theute] [~jdoyle] this happened during my leave - I see that this is the only sub task that was not completed out of the 4 subtasks for https://issues.jboss.org/browse/JMAN4-117 was the development stopped intentionally? if yes, should it be resumed or should this JMAN be closed or deferred?	0	0	https://issues.jboss.org/browse/HAWKULAR-1155
5700	HAWKULAR-1154	I just added the command *update_collection_intervals*, as used them on my EAP6 testing. Is there any other command that needs to be supported on the gem?	1	1	https://issues.jboss.org/browse/HAWKULAR-1154
5701	HAWKULAR-1153	The requeriment does not explicitly define but I read that it implies that also an user may want to show/use in reports/graphs the selected metrics. Today, Manage UI miq_report engine uses static yaml files to define the supported layout and metrics. If this requirement also implies to dynamically build reports and graphs, the current use of miq_report needs to be revisited and probably study/define an alternative implementation. Can you confim ?	1	1	https://issues.jboss.org/browse/HAWKULAR-1153
5702	HAWKULAR-1153	From the past discussion on IRC I guess it is clear that this requirement does not involve that an user can create a custom chart with the metrics selected to collect. Please [~jdoyle], [~theute], [~pilhuhn], can you confirm ?	1	1	https://issues.jboss.org/browse/HAWKULAR-1153
5703	HAWKULAR-1141	BTW, I got a question: Is this Jira limited to work only for WildFly servers hosted in managed VMs, hosts, etc.? or should it also work if the underlying host is not in ManageIQ inventory? cc [~alissabonas]	1	0	https://issues.jboss.org/browse/HAWKULAR-1141
5704	HAWKULAR-1093	To get started it will be helpful to have an idea of what sort of auto-detectable stuff can be used to tie together clustered resources. Is there anything we can grab from a WF Server's config? jgroups settings, ispan cache info? I'll start looking. And also to get a better idea of what we want to generate relationship-wise. If we already have some fairly firm ideas of what we want [~pilhuhn], [~jdoyle2] can you add your thoughts or requirements.	1	1	https://issues.jboss.org/browse/HAWKULAR-1093
5705	HAWKULAR-1085	Is that still a thing?	0	0	https://issues.jboss.org/browse/HAWKULAR-1085
5706	HAWKULAR-1081	[~jmazzitelli], I removed the maven-war-plugin part from the Agent Download WAR and got this diff. Are those entries needed by something (like the installer?), or may this maven-war-plugin part be removed entirely? {code} $ diff ../tmp1/META-INF/MANIFEST.MF META-INF/MANIFEST.MF 7d6 < Specification-Vendor: Red Hat, Inc. 9d7 < Specification-Title: Hawkular Agent: WildFly Agent Download WAR 11c9 < Built-On: 2016-06-02T17:25:10+02 --- > Built-On: 2016-06-02T17:27:46+02 15d12 < Specification-Version: 0.18.1.Final-SNAPSHOT {code}	0	0	https://issues.jboss.org/browse/HAWKULAR-1081
5707	HAWKULAR-1081	Pull 197 was merged long time ago. Is this Jira still a thing - especially after Mazz' last comment?	0	0	https://issues.jboss.org/browse/HAWKULAR-1081
5708	HAWKULAR-1034	The new sparkline charts have ALL of the capability of the normal charts just that they are scaled down. For instance, you can see the pink alert areas and if you click and drag the time range actually changes to what you have selected. One code base, multiple purposes.	0	0	https://issues.jboss.org/browse/HAWKULAR-1034
5709	HAWKULAR-1022	[~uxdlc] sure. Your proposal is much better. This part was made without design, that's why it ended up like that :P Based on your design, I propose one change, which is to remove the Hide/show metrics and put it as a collapse / expand arrow before the panel header. This interaction pattern is well-known, saves us vertical space and looks better. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-1022
5710	HAWKULAR-1021	[~mithomps], can u sign it to someone from your team?	1	0	https://issues.jboss.org/browse/HAWKULAR-1021
5711	HAWKULAR-983	An iteration on this flow might be to do the same, but all within a dialog window So it would be: * Click Export JDR... * Dialog window pop-ups: It has a title, explanatory text and a spinner. * Once the file is ready for download, replace the spinner with File Dialog fields. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-983
5712	HAWKULAR-983	I am not sure how it is done on the server, but can the server somehow either send progression while generating the JDR, or guess the time? So we can show toastr message and then loading bar at the top similar to youtube's loading bar. I guess you saw it on couple of sites, we can actually use this http://victorbjelkholm.github.io/ngProgress/ Once it is finished, we can either prompt the user to immediately download the file (just as we do now). Or we can place some download button like: !Selection_013.png|thumbnail!	1	1	https://issues.jboss.org/browse/HAWKULAR-983

#	Requirement	Questions	Answered	Successfully Answered	URL
5713	HAWKULAR-983	And should I somehow dismiss the "Generating..." message once it's completed?	0	0	https://issues.jboss.org/browse/HAWKULAR-983
5714	HAWKULAR-983	[~karel.hala] Are these 3 stacked messages all relating to the same process? Can we dismiss/replace one toast with another as the process goes through its cycles? Please refer to: https://issues.jboss.org/browse/HAWKULAR-870	1	1	https://issues.jboss.org/browse/HAWKULAR-983
5715	HAWKULAR-960	[~uxdlc], [~mithomps], [~pilhuhn] what do you think? Did I missed anything? Is there anything else you would like to see?	1	1	https://issues.jboss.org/browse/HAWKULAR-960
5716	HAWKULAR-960	+1 * I like that these messages are only for the current session, otherwise there might be confusion/overlap with Alert Center or other notifications. * We might want to consider another icon for the "drawer." I found these at font-awesome: icon-notification & icon-flag-alt. Or we could use a text label instead ("Notifications", "Message" or etc.)? Wdyt? * Fyi: if you don't already have it, here's a link to the PatternFly pattern for toasts (though it doesn't cover this use case): https://www.patternfly.org/patterns/toast-notifications/	1	1	https://issues.jboss.org/browse/HAWKULAR-960
5717	HAWKULAR-960	[~karel.hala] RE: "merging" to some degree we have already confronted this issue with: https://issues.jboss.org/browse/HAWKULAR-870. For a single user action, that might have multiple events tied to it (like Export JDR), it could make it easier for the user if one corresponding toast updated itself - rather than stacking a series of event updates for the same user action. Totally unrelated tasks/events would still stack up in the queue though. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-960
5718	HAWKULAR-954	[~mithomps] Question: * This applies to what you initially see on lote, with the reflective server view? * And/or after an agent install process? In the case of the agent install, perhaps the redesign of this view should also be included/considered as part of #918 work flow. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-954
5719	HAWKULAR-954	Do we have any sense of what an average load time might be? Or are we able to calculate progress as the data is loading. I'm trying to figure out whether we should include either an indeterminate or determinate progress indicator (aka spinner or progress bar). Thanks!	1	1	https://issues.jboss.org/browse/HAWKULAR-954
5720	HAWKULAR-930	This is done. It's now possible to select a date for the expiration of the token. We might iterate this again, to fine tune it. For instance, do we want to make it easier to select a specific time as well, or date is fine? Also, slider or date picker? And so on. For now, I'm resolving this, as the base code is in Hawkular master already.	0	0	https://issues.jboss.org/browse/HAWKULAR-930
5721	HAWKULAR-929	Would this be more of a title, or comment/description (a few words, or a paragraph description)? I ask, because I'm wondering if I could use the token "title" in the drop-down for the agent-install flow (rather than key/secret pair)?	1	1	https://issues.jboss.org/browse/HAWKULAR-929
5722	HAWKULAR-929	By the way, we could have two ways of displaying the token info on a dropdown: * My token name (3b3f4f08-be83-11e5-9ee9-3c970e9648d7) * Token created at Tue 19. Jan 09:04:27 CET 2016 (4b8a027c-be83-11e5-90e0-3c970e9648d7) The first one has a custom title (My token name), the second doesn't. I realize that this is somewhat big (specially because of the size of the token key itself), but I think it provides enough context to the user. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-929
5723	HAWKULAR-929	Make it required/unique only on the UI, right? Tokens can also be provisioned on demand, so, it doesn't make much sense to require a name field there, as it would be an extra effort for an admin to come up with a "meaningful" random name for the token (changes are that they will just fill the name field with garbage). I'd do it one step at a time, and see if we really have this naming problem in the real world.	1	1	https://issues.jboss.org/browse/HAWKULAR-929
5724	HAWKULAR-929	[~jpkrohling] The token-key is basically not human-readable. What do you mean with "on demand"? In the context of creating a customized agent-installer? We came to the conclusion today that we don't want ad-hoc creation in that workflow, but the user should do that before on the tokens screen. We can't prevent the user from entering garbage in the title field, but then it's users labeling system. Forcing the user to remember which token key is what, is not user friendly at all.	0	0	https://issues.jboss.org/browse/HAWKULAR-929
5725	HAWKULAR-929	I still think my previous suggestion would be better than just displaying the key/code. {quote} By the way, we could have two ways of displaying the token info on a dropdown: * My token name (3b3f4f08-be83-11e5-9ee9-3c970e9648d7) * Token created at Tue 19. Jan 09:04:27 CET 2016 (4b8a027c-be83-11e5-90e0-3c970e9648d7) {quote} Attached is a screenshot from the "Tokens" screen. Note that two of the tokens have no name nor description, and an "auto generated" name for those are used (Token generated ...). Is that a viable solution?	1	0	https://issues.jboss.org/browse/HAWKULAR-929
5726	HAWKULAR-929	[~juraci.costa] - I'm sorry, could you recap your "previous suggestion"? :-} - The two ways of currently showing the token in the dropdown, aren't really very readable (user friendly) due to the long string length. - Thanks for sharing the attached screenshot: The general design seems fine, although some questions: * What does the Edit button do: Edit only the name or other fields? * Will the design break if the Name or Description field text wraps? * If you can "auto generate" a name for this, and we can refer to an auto-generated name in a dropdown, then why can't the token name be required? I would much rather refer to the name in the dropdown than these really long text strings?	1	0	https://issues.jboss.org/browse/HAWKULAR-929
5727	HAWKULAR-929	{quote} I'm sorry, could you recap your "previous suggestion"? :-} The two ways of currently showing the token in the dropdown, aren't really very readable (user friendly) due to the long string length. {quote} Sure, the suggestion was: {quote} By the way, we could have two ways of displaying the token info on a dropdown: * My token name (3b3f4f08-be83-11e5-9ee9-3c970e9648d7) * Token created at Tue 19. Jan 09:04:27 CET 2016 (4b8a027c-be83-11e5-90e0-3c970e9648d7) The first one has a custom title (My token name), the second doesn't.{quote} Those are big, but they can shortened. The code itself, for instance, could show only the first part (everything before the first dash). So, it would be something like this: {code} - My token name (3b3f4f08-...) - Token generated at Jan 27, 2016 2:33:45 PM (a632102c-...) {code} If they are still too big, we can try to use the ISO representation for the date, such as "2016-01-27T02:33:45Z" {quote} * Thanks for sharing the attached screenshot: The general design seems fine, although some questions: * What does the Edit button do: Edit only the name or other fields? * Will the design break if the Name or Description field text wraps? {quote} The edit button opens the name, description and expiration date fields. It currently breaks the screen and it gets really ugly. The JIRA HAWKULAR-979 is for fixing those quirks on the screen. {quote} * If you can "auto generate" a name for this, and we can refer to an auto-generated name in a dropdown, then why can't the token name be required? I would much rather refer to the name in the dropdown than these really long text strings? {quote} I would also prefer that we'd have an usable name and description for every token, but realistically, users will add a meaningful name and description to only a few tokens. So, it doesn't make sense to require the user to enter a name. Specially for environments where the token is created in a big batch, or via an automated system (ie: not via the UI), this is an unnecessary burden to the system administrators. I also absolutely agree that referring to tokens by their code is not a good way to present the information. Hence, the "auto generated name". If the user hasn't bothered setting a name to a token, we can at least show something like "Token generated at Jan 27, 2016 2:33:45 PM", as a way for the user to have an idea which token to select (was it the one I created yesterday? the one I just created?).	0	0	https://issues.jboss.org/browse/HAWKULAR-929

#	Requirement	Questions	Answered	Successfully Answered	URL
5728	HAWKULAR-921	* Yes, the admin receives an email notification about an user requesting to join the organization * You mean, on the screen shown in screenshot-4? I was thinking that we could change the listing of organizations to use a Card layout. What do you think ? * Users can apply again to join an organization, in case they first got rejected. Agree about marking the rejected as label-danger. Would you change it and add to the PR? The code would be like this: on organizations.html , http://paste.fedoraproject.org/317408/54405175/ * Yes, sounds reasonable. Are you able to change that as well ?	1	1	https://issues.jboss.org/browse/HAWKULAR-921
5729	HAWKULAR-921	* Yes, I meant the screenshot-4. The table has only the "Name" column, we could add "Privacy" and a new column (without heading) that shows in a "label label-info" the info "1 request pending". I would keep using the table. We are currently using cards when there are charts. * Can you apply the label-danger to rejected status? In the code you are using a tag, you'll need to split in two conditions, which I don't know how to do :) * Updated the select in a new PR + other small fixes, just needed to add pf-select :)	1	1	https://issues.jboss.org/browse/HAWKULAR-921
5730	HAWKULAR-918	this is a hawkular-accounts token key (it is NOT the user's username; it is the user or organization token's key value) 2. "security-secret" - this is a hawkular-accounts token secret (is it NOT the user's password; it is the user or organization token's secret value). 3. "encryption-key" - this is a passcode that the user will provide in order to encrypt the security-secret inside the installer. This is any string the user wants to give. You will need to know this when you then run the installer. These can be provided as FORM POST name/value pairs (if accessing the installer download URL via HTTP POST), or they can be query string parameters (if accessing the installer download URL via HTTP GET). The installer download URL that the UI must access is: {noformat} http://<hawkular-server>:8080/hawkular/wildfly-agent/installer {noformat} When downloading the installer via the URL above using a browser or some other utility like wget or curl, you store that file somewhere on your file system and call it "hawkular-wildfly-agent-installer.jar" (that is the content-disposition name the installer download URL will provide - so if you download using a browser, that will be the filename you get automatically). If you store it in the same directory where your WildFly server instance is, it makes the command line to run the installer simpler (because you won't need to tell it a --target-location - the default is the current working directory). But suppose you saved the installer jar in /tmp and your WildFly Server where you want to install the agent is located in /opt/wildfly and the passcode (encryption-key) you passed into the installer download URL is "foobar" - then the installer command you run to install the agent inside that WildFly server would be: {noformat} java -jar /tmp/hawkular-wildfly-agent-installer.jar --target-location=/opt/wildfly --encryption-key=foobar {noformat} If you copy the jar file in that same location where your WF server is installed (/opt/wildfly), then you can just do this: {noformat} cd /opt/wildfly java -jar hawkular-wildfly-agent-installer.jar --encryption-key=foobar {noformat} NOTE: if you only say "--encryption-key" rather than also give it its value, the installer will ask you on the console for the encryption key passcode. This way, it won't be stored in your bash history. There are optional values you can pass into the installer (in addition to the three required ones mentioned above). Whether we want the user to be able to specify these or not, I do not know. Up to the UXD team to figure out what's best. Perhaps the less amount of data to ask for, then the easier it is for the user to understand what is going on. But these would help pre-configure the installer further which could make it easier for the person running the installer. * managed-server-name - the name of the local WildFly endpoint the agent will monitor. This is the name of the <local-dmr> as defined in the agent configuration file. The default is "Local". * target-location - the location where your WildFly Server is located and where you want the agent to be installed. If all wildfly servers in the user's organization are always installed in the same place (say, "/opt/wildfly" is a standard within their organization) you can pre-configure the installer to use this. This would then allow you to omit the "--target-location" command line argument you pass to the installer command line (so the command would look like "java -jar /tmp/hawkular-wildfly-agent-installer.jar --encryption-key=foobar") * enabled - if specified, you are telling the installer to configure the agent to either be enabled or disabled. The default is true - the agent will be enabled and will deploy and run when you next restart the WildFly server. If you pass in a value of false the agent will not run when the WildFly Server is restarted. You can later manually enable the agent later. Note that all of the properties (even the two required security-key and security-secret ones) can be overridden by the user running the installer by simply passing in the property on the command line. For example, suppose in the UI you specified "target-location" of "/opt/wildfly" so the installer will install the agent in WF servers it finds in /opt/wildfly. But suppose you are on a machine where your WildFly server is located someplace different - like "/usr/share/wf". You can tell the installer to install in that location by simply passing in a different value for --target-location on the command line: "java -jar /tmp/hawkular-wildfly-agent-	1	1	https://issues.jboss.org/browse/HAWKULAR-918
5731	HAWKULAR-918	[~mazz] Is the "encryption-key" a one-time thing? If the user runs this again, will they need to generate a new encryption-key each time?	1	0	https://issues.jboss.org/browse/HAWKULAR-918
5732	HAWKULAR-918	[~mazz] [~pilhuhn] I attached an interim design proposal, as something to use while we're working on the final wizard flow. I ran it by Mike T. already and have incorporated his feedback. There's one variant proposal in which I've changed the title to "WildFly Server Agent," to try to make things more clear in the short-term. Wdyt?	1	0	https://issues.jboss.org/browse/HAWKULAR-918
5733	HAWKULAR-918	Good. How did u think the interaction of "Create token" in step 2? It is just added to the list automatically?	1	1	https://issues.jboss.org/browse/HAWKULAR-918
5734	HAWKULAR-918	[~gcardoso] That would match the interaction in the "Tokens" section, so that seems reasonable. One question is whether to allow Token edits in this view or require the user to go to the Tokens section. Wdyt?	1	0	https://issues.jboss.org/browse/HAWKULAR-918
5735	HAWKULAR-918	Your solution regarding the name sounds good. The need for login at this stage is not optimal however. At least some users will not understand why they have to login again if they are already logged in. Isn't there another solution that runs this login in background, like you do by just going to the accounts area, where you don't need to put login information again?	1	1	https://issues.jboss.org/browse/HAWKULAR-918
5736	HAWKULAR-918	How does this process of redirecting the user to Keycloak and redirect back appears to the final user? Just a loading in the modal or would he see something different?	1	1	https://issues.jboss.org/browse/HAWKULAR-918
5737	HAWKULAR-909	[~lkrejci] any idea, how should we name the new Commons Maven module that will contain CanonicalPath?	1	1	https://issues.jboss.org/browse/HAWKULAR-909
5738	HAWKULAR-901	Hi [~mazz] [~uxdlc] FYI, the patternfly team has added this issue to the backlog and will be designing this pattern in a future sprint. Please see it here: https://patternfly.atlassian.net/browse/PTNFLY-639 Could you please let me know if there is a determined timeframe of implementing this feature? If this is urgent, I can communicate to the PatternFly team and ask the product owner to reprioritize it. Thanks!	1	0	https://issues.jboss.org/browse/HAWKULAR-901
5739	HAWKULAR-901	[~pilhuhn] Are you envisioning a "message center" that is distinct from the "alert center" (by location and etc.), or something that is part of it? And which types of notifications are toasts, versus other types of messaging?	1	0	https://issues.jboss.org/browse/HAWKULAR-901
5740	HAWKULAR-885	Hi, I probably was not very clear, but I didn't mean to use different widgets together in the same panel. I meant to pick one of them and try to apply to our context. Your proposal presents the information in a clear way but I didn't see anything similar in patternfly (with the lines dividing, one heading above, chart in the middle and legend below). I attached a prototype that follows more the Utilization Trend Card. It feels more like patternfly to me, and also presents the information clearly. WDYT? Thanks	1	1	https://issues.jboss.org/browse/HAWKULAR-885

#	Requirement	Questions	Answered	Successfully Answered	URL
5741	HAWKULAR-884	Let's say I want the Heap Usage metric, how do I query inventory for such thing without hard-coding a string which identifies it ? All I know it's the ID, which is "MI~R~[bucephalus/Local~~]~MT~WildFly Memory Metrics~Heap Used".. so that pretty much will need to be concatenated/hardcoded somewhere. I can query all metrics for the resource "bucephalus/Local" and then look for the ID I care about (eg: "~MT~WildFly Memory Metrics~Heap Used") but IMHO, it increases the queries while not providing a major improvement since if the metric id changes, it still breaks. I tried to minimize the damages of changes by using this MetricsService.getMetricId(type: string, feedId: FeedId, resId: ResourceId, metricId: MetricId) method (https://github.com/hawkular/hawkular/blob/master/console/src/main/scripts/plugins/metrics/ts/services/metricsService.ts#L94-L104) so that it is less error prone and we can easily change if the metric id changes.	0	0	https://issues.jboss.org/browse/HAWKULAR-884
5742	HAWKULAR-882	Do we want to perform it via UI ? I can add the support for the component and un-assign this for an UI tasks once component support it.	0	0	https://issues.jboss.org/browse/HAWKULAR-882
5743	HAWKULAR-881	The original problem with "adding an user to an organization" is that it allows for users to know about the existence of other users. If this is not seen as a problem, then this could be implemented. About the new user requesting access to an organization: should then the list of organizations be visible to everyone? Or should there be a setting, with the organization's visibility ? Should there also be a "public" possibility, so that people would join without the need for an approval?	1	1	https://issues.jboss.org/browse/HAWKULAR-881
5744	HAWKULAR-870	[~mazz] We can talk to the PatternFly people to see if they're working on a "message ctr" pattern. In the meanwhile, do we have a clear definition around what "messages" are (versus alerts and alert notifications)? Just want to make sure we don't get into the place in the future where the user is confused about where to look to get messaging and alert information. [~jdoyle] John, do you think we should start working up designs for a message center, I imagine one that's similar to the EAP message queue in the upper-right?	1	1	https://issues.jboss.org/browse/HAWKULAR-870
5745	HAWKULAR-870	[~alexandrem] Thanks for minimizing the stacking. WRT to stacking etc., we're trying to align with PF wherever possible and reasonable, when they do have stated recommendations. This is a patterns where they do provide clear guidance: * "Toast notifications should be transient and stay on the screen for 8 seconds, so that they do not block the information behind them for too long, but allows the user to read the message. Ideally, the user can decide what kinds of notifications appear and how long they remain on the screen. Toast notifications should remain on the screen when the user is hovering on them.Generally, there shouldn't be more than one toast notification shown at the same time. If two notifications are triggered at the same time, they should be shown one by one." https://www.patternfly.org/patterns/toast-notifications/ What would be great if rather than having Hawkular deviate from these recommendations, if you could raise them with Patternfly. I'm including the pattern lead on this thread. [~joqian] Josephine, is there some way that Alex can contribute his concerns/suggestions to the pattern work? That said, perhaps the message queue (Jira #901) might also help address the issues/concerns that you raise?	1	1	https://issues.jboss.org/browse/HAWKULAR-870
5746	HAWKULAR-870	* [~alexandrem] Alex, is one of your points that a toast shouldn't persist for 8 seconds, if a new toast is generated from a user action? If so, I agree with that, the new toast should replace the old one even if its 8 secs aren't up. I don't think there's anything it the pattern that precludes that? * [~mazz] John, the message ctr discussion should probably go on #901 :) I filed this Jira for the "quick fix" with the hope of moving the larger message ctr discussion to another Jira. I think we do need to discuss messaging in the UI holistically (toasts, email, inline messages, alerts, flags, badges and etc.).	0	0	https://issues.jboss.org/browse/HAWKULAR-870
5747	HAWKULAR-870	Change request: Align more closely with PF recommendation not to stack. If a compromise position is needed (until there's a message ctr or some such): Perhaps we could show more than one, by allowing the old message to phase out while a new message fades in, like a vertical carousel. I think if there's some transition effects that might help communicate that these things are coming and going, wdyt? I think that could help resolve the case where many are coming in at the same time, but not stacking them. Here are some examples: * https://amazingcarousel.com/examples/jquery-vertical-scroller-id15/ * https://amazingcarousel.com/examples/jquery-carousel-id5/ * "Example 1 or 2" http://wmh.github.io/jquery-scrollbox/ * http://jsfiddle.net/spliter/g2PE7/ [~alexandrem] How about this for a compromise/ interim solution? Perhaps this carousel queue could store up to 3-5 toasts max, but only show one at a time. And the whole toast carousel could clear after a minute or two. If only one toast is spawned in any given 1 - 2 minute interval, then just show the regular toast. wdyt?	0	0	https://issues.jboss.org/browse/HAWKULAR-870
5748	HAWKULAR-858	[~gcardoso] Thanks for picking this up! Could we also resolve the multi toast presentation in this Jira? According to PF: "Generally, there shouldn't be more than one toast notification shown at the same time. If two notifications are triggered at the same time, they should be shown one by one." Currently we present a whole bunch at the same time, I attached a screen grab.	1	1	https://issues.jboss.org/browse/HAWKULAR-858
5749	HAWKULAR-855	Is that still a thing?	0	0	https://issues.jboss.org/browse/HAWKULAR-855
5750	HAWKULAR-850	* Are you interested in an alternative presentation of the text inside the tooltip? * One question on the hovers themselves - would it be possible to present the hover over any part of the vertical axis line in the chart (regardless of whether the bar is there or not)? See: https://www.patternfly.org/widgets/#basic-charts	1	1	https://issues.jboss.org/browse/HAWKULAR-850
5751	HAWKULAR-833	[~gcardoso] In working on moving the trigger definitions back into a dialog, as we discussed, I realized that we need to address alert details as well. Because in the current UI clicking the (>) chevron icon takes the user to the Alert Details, not the triggers. So I was thinking that along the lines of what we discussed, perhaps clicking the (>) could take the user to the dialog - which could include the details as well as the triggers. We could separate these by using tabs, and extend that further by using a tab for "notifications" as well. Wdyt?	1	0	https://issues.jboss.org/browse/HAWKULAR-833
5752	HAWKULAR-833	[~gcardoso] Cool, thanks for your help! [~mithomps] Mike, which release version do you think this change could get in?	1	0	https://issues.jboss.org/browse/HAWKULAR-833
5753	HAWKULAR-816	Isn't this task done? Or should there be something more implemented? Currently there is Response Time and Availability as separate tabs and sidebar with alerts in it. [~uxdlc] or do you think that there should be something more to be done?	1	1	https://issues.jboss.org/browse/HAWKULAR-816
5754	HAWKULAR-815	Hi [~uxdlc], I implemented the delete on: http://cardosogabriel.github.io/hawkular-client/urls-list.html To do so I moved the alerts icon to the left, since the action makes more sense in the right. Consistency wise, we should probably use the same approach for the datasources tile, but with the 3 dots instead of the trash can. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-815
5755	HAWKULAR-815	[~gcardoso] +1, to alerts badge on the left, actions on the right. And I agree that we should make this change to Datasources as well - using some kind of actions icon (either the three dots or a dropdown arrow). Do you know if they're using either action menu icon in PF tiles or cards?	1	0	https://issues.jboss.org/browse/HAWKULAR-815
5756	HAWKULAR-815	Is there a JIRA for the delete button ? I see that comments were added after it was actually implemented and resolved. And the delete icon is still missing	1	1	https://issues.jboss.org/browse/HAWKULAR-815
5757	HAWKULAR-803	[~uxdlc] [~gcardoso] How should this new capability of context chart/zoom be incorporated into the existing screens. I'm assuming a fairly simple insert of it somewhere.	1	1	https://issues.jboss.org/browse/HAWKULAR-803

#	Requirement	Questions	Answered	Successfully Answered	URL
5758	HAWKULAR-803	I believe we had some early designs for this, I'll search around for them. Could you move forward with the implementation and then we could refine the presentation/interaction after?	1	0	https://issues.jboss.org/browse/HAWKULAR-803
5759	HAWKULAR-803	Very nice Mike! In the video, you have a Context Chart that you use to zoom in and the other charts reflect the zoom. How would that work for Hawkular? For example, in the Response Time screen, would you need to have two charts, one to manipulate and another with the changes? In terms of UI, I'd initially hide the chart controls (threshold, data points...) under some button that open these options.	1	1	https://issues.jboss.org/browse/HAWKULAR-803
5760	HAWKULAR-803	I'm looking to UxD for design for this, so whatever you think. We can have individual controls on charts, if that makes more sense (and can even make the context chart smaller if need be). Should these be top/bottom. Static/expanding?	1	1	https://issues.jboss.org/browse/HAWKULAR-803
5761	HAWKULAR-803	[~mithomps], I discussed with [~uxdlc] this simple approach for chart zooming. The approach would be the same as Google Maps, so basically you zoom in/out (also with the slider), and can drag the chart left/right to see backwards/forward. WDYT about this approach, is it feasible? Tried to make it very simple, so we don't need this second contextual chart. Also add a ballon that explains how to use in the first time.	1	1	https://issues.jboss.org/browse/HAWKULAR-803
5762	HAWKULAR-803	Question - does "zooming" change the date range selection for just the one chart or everything on the page?	1	0	https://issues.jboss.org/browse/HAWKULAR-803
5763	HAWKULAR-796	[~gcardoso] First pass. Similar to the URL view, the user could click on a tile to view a details screen. The details screen would show what's in the expand panel now - the metrics and the Alerts side bar. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-796
5764	HAWKULAR-796	Prototype on: http://cardosogabriel.github.io/hawkular-client/application-datasources-new.html [~mithomps], will we use c3 for this chart? If so, I'll investigate how to fix the hover issue that only displays the tooltip inside the tile (thus hiding part of the tooltip)	1	1	https://issues.jboss.org/browse/HAWKULAR-796
5765	HAWKULAR-796	[~gcardoso] I'm still seeing the scrollbars, FF 43.0.1 on OSX 10.8.5. P.S. Is the title set to truncate or wrap to a second line when its too long?	1	1	https://issues.jboss.org/browse/HAWKULAR-796
5766	HAWKULAR-796	[~gcardoso] I have noticed that there is problem with number of available connections. They are not shown because there is fixed height of each tile to 300px, changing this to height: 100% fixes this problem, but I am not sure if it won't break anything else. Could you double check, please?	1	1	https://issues.jboss.org/browse/HAWKULAR-796
5767	HAWKULAR-796	Another thing I noticed with tiles: When you hover over part of chart, the hover is not visible and is hidden. This is done because of overflow: hidden and overflow: auto for classes .tiles-view-pf .tile-content. I disabled these two options and it is good, but we should add word-wrap: break-word; for .hk-tile-info .hk-data, so long info at bottom of tile will not go outside tile. What do you think [~gcardoso] are these things viable or could it break anything? I will PR this without these changes and let them to be discussed later on, since it is not that much related to this task.	1	1	https://issues.jboss.org/browse/HAWKULAR-796
5768	HAWKULAR-792	Are the data being collected? I am looking trough the responses from server and I can't see any bucket with data inside them. Or if you could point me out what the REST URL should look like.	1	1	https://issues.jboss.org/browse/HAWKULAR-792
5769	HAWKULAR-792	[~pilhuhn] I have shown the data for platform and I have been changing the UI. For now I have two possible ways to show filesystems graphs. Option A: !option_a_collapsed.png|thumbnail! !option_a_expanded.png|thumbnail! Option B: !option_b_collapsed.png|thumbnail! !option_b_expanded.png|thumbnail! Option A looks better expanded, but option B looks better collapsed. What do you think would be better?	1	1	https://issues.jboss.org/browse/HAWKULAR-792
5770	HAWKULAR-792	I guess I can wrap up the platform tab. Just couple of things which might need some tweaking, but that is more or less back-end stuff. * The CPU usage for whole processor is wrong, it shows all data combined together, rather than average. * File systems can't show data for less than hour, still don't know why. * There is no way how to create trigger alert for combined resource, so alert triggers for processor are created for each core. This might be a bit clunky, but you have more power over individual core. I will keep the branch up to date with master and once you say [~pilhuhn] that it is ready to go I will PR this.	0	0	https://issues.jboss.org/browse/HAWKULAR-792
5771	HAWKULAR-790	[~pilhuhn] Could it be a central "settings" link off of the Server Overview screen? If so, it could link to /open up a "Metrics Settings" dialog (or etc.), which would have toggle buttons for each subsystem: Metrics Settings * WEB [ON/OFF] * TRANSACTIONS [ON/OFF] * DATASOURCES (for each): [ON/OFF] * CACHE (for each): [ON/OFF] * though I don't know what this is in the UI? * EJB3 (for the subsystem as a whole):[ON/OFF] * is this the whole thing? * JMS MESSAGING (for each server):[ON/OFF] Am I understanding this correctly?	1	1	https://issues.jboss.org/browse/HAWKULAR-790
5772	HAWKULAR-790	[~mazz] Hi, I attached a first pass, mockup. I placed the toggles up-front on the overview page, but not sure it works like that... I was thinking that the user could hit the "settings" link, to bring up a dialog with more granular settings. Wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-790
5773	HAWKULAR-790	[~mazz] Thanks for checking it out, and for the backend jira. Does it need the "settings" link & dialog to fine-tune the timing intervals for each metric? If so, I'll mock that up along with the tabs view.	1	1	https://issues.jboss.org/browse/HAWKULAR-790
5774	HAWKULAR-790	[~gcardoso] * Good point, in the individual tab view currently you can only turn ON metrics. We could add a toggle when metrics are displaying to turn them off at the tab level. wdyt? * +1 to rewording the title to "Web Collection is turned off."	1	1	https://issues.jboss.org/browse/HAWKULAR-790
5775	HAWKULAR-790	For the sake of simplicity, I would not add one more control in the tab, but rather leave the possibility of turning of in the overview (as it is). We can observe if there is really a need to have it in the screen. One good practice would be to display a message when we first turn on a collection, like: "You can turn off metrics collection in the overview tab." wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-790
5776	HAWKULAR-774	Prototype on: http://cardosogabriel.github.io/hawkular-client/application-url.html [~uxdlc], don't we need a link for "Alert Settings" on this screen? - The kind of cards that patternfly provide for our needs (number + label below) are these smaller ones. But if you think they should be taller I can customize a bit. - Updated the date-picker box, placing the icon inside at the right (as it was suggested in that thread about this component) - I'll update the header area with your white background as another task. Can you create a JIRA for that? Can you check how it looks for all screens (including organizations, user settings)?	1	1	https://issues.jboss.org/browse/HAWKULAR-774
5777	HAWKULAR-773	[~uxdlc], how would be the interaction to delete the URL? For datasources, we are using a dropdown in the title. We could use the same approach, my only concern is that urls can be long. What are your thoughts?	1	0	https://issues.jboss.org/browse/HAWKULAR-773

#	Requirement	Questions	Answered	Successfully Answered	URL
5778	HAWKULAR-761	what I did today: * Made a couple of changes [1] in hawkular-parent, i.a. to upgrade to WF 10 * moved nest to new repo [2] to solve the circular dependency introduced by moving it to HK main * Started rebuilding the components one by one with the new parent ** encountered no problems with commons [3], bus [4] and the new nest component ** Found out that there is no WF10 compatible Keycloak release ** spoke with mposolda and mstruk figuring out that they have a wf10 branch https://github.com/keycloak/keycloak/commits/wildfly-10 ** they did not name any specific date when the work in the branch will be released, so Thomas (probably) should speak with them to make sure that they release ASAP ** For now, I am trying to build with a Keycloak srcdep from that branch ** Anyway, I am quite skeptic that I will be able to build a working Accounts without Juca's help ** This is what I have so far in accounts: [4] [1] https://github.com/hawkular/hawkular-parent-pom/pull/55 [2] https://github.com/hawkular/hawkular-nest [3] https://github.com/hawkular/hawkular-commons/pull/21 [4] https://github.com/hawkular/hawkular-accounts/pull/64	0	0	https://issues.jboss.org/browse/HAWKULAR-761
5779	HAWKULAR-755	This is already supported to some degree in the Alerts component. It's possible to define a CompareCondition using 2 metrics such that if (metric1 <oper> (somePercentage * metric2) then the condition matches. The metrics can be from any resource. Some examples: If activeSessionsDS1 > 150% of activeSessionsDS2 then match If heapUsedJVM1 > 80% heapMaxJVM1 then match Is this the sort of thing you're looking for? Of course there is no real provision for anything like this in the UI right now, which is very single-resource, single-condition oriented.	1	1	https://issues.jboss.org/browse/HAWKULAR-755
5780	HAWKULAR-737	Do you believe they want to manage organization and review user settings being an organization? If so, we leave it. But the text can be reduced a bit, we don't need to explain where is the switch, they already used it to switch to an organization :)	1	0	https://issues.jboss.org/browse/HAWKULAR-737
5781	HAWKULAR-732	[~john.sanda] the migration to Artemis has been reached and the present issue is resolved, is it not?	1	1	https://issues.jboss.org/browse/HAWKULAR-732
5782	HAWKULAR-732	I see that the one PR marked as open (in the jira) has been closed on GitHUb. Where do we stand here? If the PR is not needed, then please remove it from the jira and forward the jira into the resolved state.	1	1	https://issues.jboss.org/browse/HAWKULAR-732
5783	HAWKULAR-729	[~gcardoso] Gabriel, what do you think of iconic buttons? Example: "Apply2AlertC.jpg"	1	1	https://issues.jboss.org/browse/HAWKULAR-729
5784	HAWKULAR-729	[~gcardoso] Gabriel, yes I think going with the drop-down might be a safer bet from a consistency standpoint. I also mocked up Heiko's suggestion of adding all the actions in an Actions drop-down (Apply2AlertDb.jpg), I think this could work as well. Wdyt? [~mithomps] I totally hear your point about a consistent widget to initiate the action, but in the context of a table view having a dropdown is a pretty standard model so I think we'll be OK. Once PF has some/more guidelines around actions toolbars, perhaps this could be something to explore in the future.	1	1	https://issues.jboss.org/browse/HAWKULAR-729
5785	HAWKULAR-727	[~gcardoso] Great, thanks for checking it out! For the apply interaction - could we have some that behaves something like this: https://twitter.github.io/typeahead.js/examples/	1	1	https://issues.jboss.org/browse/HAWKULAR-727
5786	HAWKULAR-727	[~jdoyle] I'd like to mark this as done, and pass it along to Gabriel. Does it look OK to you too?	1	1	https://issues.jboss.org/browse/HAWKULAR-727
5787	HAWKULAR-727	[~theute] Thomas, should I mark this as done so Gabriel can start coding or wait for additional feedback? Thx!	1	1	https://issues.jboss.org/browse/HAWKULAR-727
5788	HAWKULAR-721	Do we want to use Alerts tagging feature to view/manage custom groupings of Alerts?	1	0	https://issues.jboss.org/browse/HAWKULAR-721
5789	HAWKULAR-721	AFAIK, Tagging is still largely undefined, but in the longer term I do that that we would want to be able to sort and a filter by tags. Is there something we can do in the coming milestone?	1	1	https://issues.jboss.org/browse/HAWKULAR-721
5790	HAWKULAR-715	Thanks, though I'm still confused. I understand that the only user actions for System generated labels would be Apply/Remove. But aren't we supporting users being able to Add (create) a label? If so, then they can Create a label but they cannot Delete it?	1	1	https://issues.jboss.org/browse/HAWKULAR-715
5791	HAWKULAR-715	[~jayshaughnessy]Is it possible to add tags to an alert, or are they truly static?	1	1	https://issues.jboss.org/browse/HAWKULAR-715
5792	HAWKULAR-715	[~jdoyle] I could imagine that one might want to group alerts as part of generating a trouble report, is that true?	1	1	https://issues.jboss.org/browse/HAWKULAR-715
5793	HAWKULAR-712	Can we please rename this epic to "trouble report" , [~jdoyle] ?	0	0	https://issues.jboss.org/browse/HAWKULAR-712
5794	HAWKULAR-711	There is a list of drivers and those can be deleted individually, but it doesn't specify the version. Should we add the version to the list or is it expected to have more details ?	1	0	https://issues.jboss.org/browse/HAWKULAR-711
5795	HAWKULAR-710	[~gcardoso] Gabriel, I wanted to run them by you before marking them as done. I made them a while back, and I'm concerned they might be a little diff than the current screens. Wdyt? Thanks.	1	0	https://issues.jboss.org/browse/HAWKULAR-710
5796	HAWKULAR-710	I comment about both screens in different posts. exampleurl.com - Breadcrumbs are nice :) - date picker and last update seem to be in inverted position. For server overview you placed them differently. - IMO this design have too similar styles for the url (exampleurl.com), widget heading (availability metrics) and the status boxes. It is hard to differentiate the information and know straight away what is important. I'd try some variations on font-sizes / weight / opacity. For the status boxes, I'd try to use the blue top borders that they have on patternfly. I believe the font-size of the status numbers should be bigger than the box heading. Not sure if we need the horizontal / bottom borders, probably spacing should resolve this. Somehow the layout seem to have too many boxes. Maybe because all information are in boxes, we don't need the white background at the top? Can you try a variation? - chart title is too distant from the chart (availability), information doesn't seem related. - in server overview screens you presented the time range information inside the box title, at the right. We should keep it consistent. I'd say remove from the chart heading. - [~mithomps], when clicking on "update" button, does it update only one chart, both charts, or also the status bar numbers?	1	1	https://issues.jboss.org/browse/HAWKULAR-710
5797	HAWKULAR-710	myapplication server - somehow the url widgets feel too small, it is expected to monitor many urls and show as much as possible in the screen? I'd make them bigger (probably 3 per line on this width). the way we present urls are: http://www.goawake.org , so we might need more space (it is not ideal to break the line of the url) - the circle around the up/down icon feel like a button to me (download), my first impression - what was the reason to hide the last downtime? - would place the alert badge at the top right (we are usually placing at the right side, also app icons are usually top right, more commom location) - not sure how useful is the graph without a reference (minimum and maximum value). I know that it looks nice though. [~mithomps], in terms of performance, how is to load all this graphs at the same time? - is the iP/apache information more important than the response time? I don't think so. hierachy should be inverted. - I'd use another design for the hover effect (maybe background color, blue border). this layer above with one arrow above another is confusing. we can also make the box grow a bit and make the shadow darker (think material design uses something like this :) - since URLs is highlighted in the menu, do we need the heading below? Can you try a variatiion?	1	1	https://issues.jboss.org/browse/HAWKULAR-710

#	Requirement	Questions	Answered	Successfully Answered	URL
5798	HAWKULAR-710	Response to your feedback: * I agree, it's not ideal to break the URL onto a second line but I don't think a larger tile is the answer. For one, larger tiles means less URLs per view, and 2 it's just a work around. If displaying the URL name on one line is a priority then I think we should move back to a list view. [~jdoyle] John, are you OK with using a list view instead of a tile view for the URL's list? * Icon + circle was used in the design that this was derived from. * Last downtime is now on the details view. * Moving badge to right sounds good. * I think Mike is already working on sparkline charts. * Apache/IP is at the bottom of the tile to avoid having too much text at the top. I'll play around with it a bit though, depending on which layout we decide to go with. * Talked to one of the visual designers about the hover, he just said to make the click icon opaque which should resolve any potential confusion. * I used the URL heading, because it's on other secondary pages. It would be good to get consensus on list versus tile view before I spend too much time revising the tile presentation details.	1	0	https://issues.jboss.org/browse/HAWKULAR-710
5799	HAWKULAR-710	Nice, URL details feels better. The only thing it is not cool here is the visual style of the agregates (boxes with shadows inside a box). If we could improve this presentation would be great! For the URL tile, in terms of hierarchy, the status (up) and response time are presented as having less importance than the IP. I thought it was the opposite, is this correct? I'm still not really confortable with the hover. Is that patternfly style? Otherwise I'd like to try a variation. Thanks!	1	1	https://issues.jboss.org/browse/HAWKULAR-710
5800	HAWKULAR-710	In the spirit of progress, could we please move forward with implementation and fine tune the visual details in an iterative fashion? These screens have been through a few reviews/revisions and I think we're in a good place to move forward.	1	0	https://issues.jboss.org/browse/HAWKULAR-710
5801	HAWKULAR-710	I attached a pair of alternate/additional mockups: * Tile view without URL details. * Details view with URL details added. We should probably provide the URL Details on the details view, to resolve the linked Jiras. wdyt?	1	0	https://issues.jboss.org/browse/HAWKULAR-710
5802	HAWKULAR-710	[~uxdlc], yes, I saw that the two tabs were consolidated in one screen, but I think the usage of panels for each one is not optimal, since as a result we have a design which is not consistent with other screens. Also, placing the aggregates below the charts again breaks our consistency. I attached a prototype that consolidate the two tabs in one screen but keeps the consistency of hawkular. I presented it to Andres and he also believes it is a better solution due to consistency matters. How about heading this direction?	1	1	https://issues.jboss.org/browse/HAWKULAR-710
5803	HAWKULAR-710	[~gcardoso] Hi, I saw the CSS version of the URL tile view. It looks nice, I like the hover effort. I noticed the circles around the boxes turned to squares. Something doesn't feel right about that, I can't put my finger on it. How about just showing the icon, no box or circle? I attached an example. WDYT?	1	0	https://issues.jboss.org/browse/HAWKULAR-710
5804	HAWKULAR-710	[~uxdlc], thanks! The hover effect comes from patternfly itself :) Weird, it works here. Can you tell me your browser version? Attached a screenshot of what I see.	1	1	https://issues.jboss.org/browse/HAWKULAR-710
5805	HAWKULAR-710	Sure we can talk. Actually I don't see 3 level tab-sets. We have the primary nav (black top bar), secondary nav (tb below the page heading), and some contextual nav (like go to details etc.). What I'm proposing is to keep the secondary nav for the detail screens. Practical example: I'm inside an application URL and want to see datasources. Currently I need to click "All URLs", and then the tab that contains the link to datasources is available and I can get there (2 clicks). However, if the tab is available in the details level, it is only one click. Also good because inform the user where he is. Not sure why we need a back link. Let's talk a bit about this after our normal meeting on Thursday?	1	1	https://issues.jboss.org/browse/HAWKULAR-710
5806	HAWKULAR-710	Yes, the sparkline graphs were intended to be for "Response Time." It is a variant design, based on the following PatternFly pattern: https://www.patternfly.org/patterns/utilization-trend-card/. [~mithomps] Mike, are the sparklines something we are able to use now?	1	0	https://issues.jboss.org/browse/HAWKULAR-710
5807	HAWKULAR-710	I guess this task can be changed to unassigned, since we don't have URLs specific for one server. Do we?	0	0	https://issues.jboss.org/browse/HAWKULAR-710
5808	HAWKULAR-703	[~ppalaga] you assigned that to me - what do you need?	1	1	https://issues.jboss.org/browse/HAWKULAR-703
5809	HAWKULAR-692	And the alert definition would be retained, but in a disabled state? That seems reasonable to me. Where you would incorporate this? Any UI that has a list of alerts? Alert Center, WildFly subsystem pages, etc..	0	0	https://issues.jboss.org/browse/HAWKULAR-692
5810	HAWKULAR-679	[~uxdlc], regarding this detail screen, how did you designed the interaction to change state, the user change in the dropdown AND needs to Save, or doesn't need to Save? Since this is not really a form, I'd suggest to not require saving to change the status (as it happens in the alert list). I'd not present the buttons Save/Cancel upfront. For the description, I'd do like jIRA, first you present only the text, if the user wants to edit, he clicks on a "Edit" icon next to the text, it becomes an textarea component, and he confirms the edition. Actually we are using this interaction in some modal (maybe deployments).	1	1	https://issues.jboss.org/browse/HAWKULAR-679
5811	HAWKULAR-670	[~mithomps], can you put them back in the deployments wizard? I can work on the design then :)	1	1	https://issues.jboss.org/browse/HAWKULAR-670
5812	HAWKULAR-670	[~mithomps], what is the name of the wizard component that you used? I can move the steps to the top via css, but would like to check if the component allow specifying the position. Thanks	1	1	https://issues.jboss.org/browse/HAWKULAR-670
5813	HAWKULAR-665	What happens when the user clicks on: - the > icon next to Open/Acknowledged ? - see all alerts link? - display resolved alerts?	1	1	https://issues.jboss.org/browse/HAWKULAR-665
5814	HAWKULAR-665	We're using this icon on the Alert Center (to view details): https://issues.jboss.org/secure/attachment/12393242/AlertCenter3.jpg John preferred the icon, for the visual "call to action." I think we should be consistent across the screens, and in particular the Alert Center screen. Is there another icon that you think could work better?	1	1	https://issues.jboss.org/browse/HAWKULAR-665
5815	HAWKULAR-665	[~gcardoso] Gabriel, I reassigned this to you in the hope that you could begin the basic screen layout set up. Is this possible, or do you need to wait for high fidelity mockups for each tab?	1	1	https://issues.jboss.org/browse/HAWKULAR-665
5816	HAWKULAR-665	Marked as "done". [~gcardoso] * When making the updates, would it be possible to reduce the padding between the left and right panels so there is less of a margin between them? It seems like if: "div class "col-md-9" had padding-left, and padding-right, of just 10px that would work. wdyt? * And could we make the overall screen content (container?) area wider, so that there's less empty space on the left and right of the screen? With the new screen layouts, I think we'll need more screen area... Not sure what the right way is to accomplish that, but hoping you can help. Thanks!	1	1	https://issues.jboss.org/browse/HAWKULAR-665
5817	HAWKULAR-665	Updated prototype with these changes: http://cardosogabriel.github.io/hawkular-client/application-jvm.html Some differences: - Put last update in the same line as the heading (as in applications overview) - Kept the Alert settings link on the UI - Didn't update the style of the status panels, due to your new proposal for uRL screen inside application. Are we moving to that style? [~uxdlc], please test also how the layout behave for smaller screen (up to tablet), since there was no specification I did something. It's good also to compare with the application overview and alert details screens, which also have this sidebar.	1	1	https://issues.jboss.org/browse/HAWKULAR-665

#	Requirement	Questions	Answered	Successfully Answered	URL
5818	HAWKULAR-665	[~gcardoso] * RE: Alert Settings. There's a Jira request view Trigger Definitions (Alert Settings) through the Alert Details instead: https://issues.jboss.org/browse/HAWKULAR-617. From there, I believe that there was an interest in redirecting the user to specific Trigger definition details screen. This would be good to check with JD about, and there are some navigation issues to discuss as well. * Not sure I understand the question about the status panels. The URL view proposal is unique to URL's if that helps. * As far as tablet view goes, I guess the potential concern is that the alerts panel drops to the bottom of the view. Perhaps that could be mitigated if the Alerts total in the particular status panel jump-linked to its related Alert panel. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-665
5819	HAWKULAR-665	[~uxdlc], did not understand what you are talking about alert settings inside this JIRA... Can u clarify? - I looked at our current views and your new design for URL details and would like to ask you why we need different views for them, since the elements are the same (status tiles + chart/charts). IMO we should propose a solution that works for both cases, this consistency is better in terms of implementation and also for the user. Otherwise we will have two styles to present the same kind of information, + the datasources screen, I think it was not specified how it would look if we add the margins to the tiles. - Sorry, could not understand the proposal for your concern (last topic). Can you clarify? Thanks!	1	1	https://issues.jboss.org/browse/HAWKULAR-665
5820	HAWKULAR-662	[~uxdlc] Can you attach the designs for this?	1	1	https://issues.jboss.org/browse/HAWKULAR-662
5821	HAWKULAR-662	Can I get more of a description? I want to make sure that I post the right design.	1	0	https://issues.jboss.org/browse/HAWKULAR-662
5822	HAWKULAR-657	[~gcardoso] That sounds good, thanks! P.S. I noticed that the subhead font sizes differ, in the modal. I attached a screenshot. Is there any technical barrier to making those consistent.	1	1	https://issues.jboss.org/browse/HAWKULAR-657
5823	HAWKULAR-655	What is the behavior of Cancel button?	0	0	https://issues.jboss.org/browse/HAWKULAR-655
5824	HAWKULAR-653	[~theute] [~jayshaughnessy] Thomas, Jay - it's just a first pass, but would like to get your feedback early. Does it seem to be headed in the right direction? Thanks.	1	1	https://issues.jboss.org/browse/HAWKULAR-653
5825	HAWKULAR-653	I'm assuming that the list view, will just provide the most important basic information. Then we will show the details on a dedicated definitions details screens, which includes actions. The details screen would also provide an overview of the conditions configurations, which would be editable. Similar to the Master/Details model used for Alerts: https://issues.jboss.org/browse/HAWKULAR-617 The Jira for Definition details is: https://issues.jboss.org/browse/HAWKULAR-655 Does that sound about right?	1	1	https://issues.jboss.org/browse/HAWKULAR-653
5826	HAWKULAR-653	[~pilhuhn], [~mithomps] I'm trying to follow Lucas's and Jay's comments. I thought we had decided to use "hard-coded" definition settings for the next/current milestone, rather than the generic/dynamic ones. Is there a question/concern about that?	0	0	https://issues.jboss.org/browse/HAWKULAR-653
5827	HAWKULAR-653	Lucas is expanding on what I was talking about as he is also starting to discuss triggers that the user will be able to create "from scratch", via a wizard. Ignoring that use case for a moment we have the issue I tried to describe. The trigger definition screens in the UI today are customized to the resource and metric they are dealing with. For JVM Heap the code knows which metric to query in order to get the Max value that it then uses to let the user define upper and lower range condition values. And implicit in that statement is that the code knows it wants to use a Range condition type (as opposed to Threshold, Availability, etc). Finally, the views are both modal and in some cases, like app server, tabbed. There is no url that will bring you directly to a desired trigger definition page. Hopefully this is more clear than my first comment. Now consider the Alert Center's Trigger List view. It is a list of all Triggers and for each we want to have a link to get to the trigger detail. The way to do this is not immediately obvious. We could do some really hacky thing and, using the trigger id or some other information we store in the trigger context, perhaps pop up the proper, existing modal. Maybe it's worth doing this as a quick interim solution. But this approach will break down when we add wizard-generated triggers that have no existing definition page. And it will require maintenance as we add new built-in triggers. The solution is likely to create a generic trigger generation view based on the trigger characteristics. This would likely drive off of Trigger context information we store at trigger creation time. Let's take the JVM Heap trigger as an example. We'd likely need to store information something like this, in the trigger context: hawkular.condition-count: 1 hawkular.condition1-type: RangeByPercentOfMax hawkular.condition1-resource: <theResourcePath> hawkular.condition1-metric: <metricIdForJvmHeap> Given this info we might be able to generate the proper detail view. You can ignore the count=1, that's just maybe thinking ahead to having multiple conditions. The really important thing is that we introduce these hawkular-level condition types, like 'RangeByPercentOfMax', 'ResponseTime', etc. Something that indicates how we want to test the given metric. A RangeByPercentOfMax is basically what we do today. We then know to present a widget that will take the Max of the relevant metric and generate a Range condition based on the user's desired percentage setting for high/low, with the ability to disable either. In short, we want a trigger detail page that renders properly for any trigger that we send to it. It may need to query inventory and/or metrics to present itself. If we have that we should then be able to use it everywhere and get rid of the existing, custom, trigger definition code that we currently have buried in the specific javascript code for each metric we support. Feasible? Desirable? Thoughts?	1	1	https://issues.jboss.org/browse/HAWKULAR-653
5828	HAWKULAR-651	[~pilhuhn] For this specific story, could you please define "resources"? Is that EAP Application Servers, and their related components/metrics only?	1	1	https://issues.jboss.org/browse/HAWKULAR-651
5829	HAWKULAR-651	@Thomas John is out this week, would you please update this story to include those requirements? It wasn't clear to me whether it should be a group of resources, and which resources specifically.	0	0	https://issues.jboss.org/browse/HAWKULAR-651
5830	HAWKULAR-651	Revised draft, with the following changes: * Removed "description" field (is system generated.) * Added Search input box. * Added Category dropdown, although there is one resource type to choose from: Application Servers.	0	0	https://issues.jboss.org/browse/HAWKULAR-651
5831	HAWKULAR-651	[~jayshaughnessy], [~rutlucas], do you guys want me to generate a static html or you guys want to try to implement it and I can review it later?	1	1	https://issues.jboss.org/browse/HAWKULAR-651
5832	HAWKULAR-651	Before embarking on this I need to confirm that in face this will offer the user something they don't already have. We have a lot of "built-in" triggers for the types, metrics, and avail we already discover and report on. Is there in fact a decent population of types/metrics that don't have built-in triggers and for which none are planned? It only makes sense, I think, to add an ad-hoc trigger creation mechanism if in fact there is stuff we monitor that is not covered by a built-in.	1	0	https://issues.jboss.org/browse/HAWKULAR-651

#	Requirement	Questions	Answered	Successfully Answered	URL
5833	HAWKULAR-651	Even adding one ad-hoc trigger is a new feature. But the essence of my question is whether we are offering anything new in that we already have built-in triggers for the resources we are discovering and monitoring. For example, for the App Server JVM we already add triggers for heap used, non-heap, GC duration etc... Do we have a significantly high number of (or any) resources and metrics that will "appear" in this list that we don't already have built-in triggers for? Perhaps I'm just thrown off by the example in the PDF showing JVM Heap as the selected metric, because we already have that built-in and adding it as ad-hoc doesn't make sense to me. Also, I have to mention that in reality the impl can't match the PDF mock. In this case it's showing Heap used as a percentage range against Heap Max. That's not actually trivial. To do what the mock is showing you'd need to: # Select that you want to perform a percentage range sort of alerting. # Specify the metric against which you are comparing (Heap Max). This complexity is hidden in our built-in Heap Max handling for App Server JVM. But in an ad-hoc scenario this sort of complexity, I think, can not be hidden. Anyway, I suppose we can move forward on the assumption that there will be resource metrics that we monitor, for which we do not plan to have built-in triggers, and on which users may want to perform alerting.	0	0	https://issues.jboss.org/browse/HAWKULAR-651
5834	HAWKULAR-650	Can we already get that token somehow?	1	1	https://issues.jboss.org/browse/HAWKULAR-650
5835	HAWKULAR-646	[~gcardoso], which icon should be used for that? The name of the FA icon would be helpful.	1	0	https://issues.jboss.org/browse/HAWKULAR-646
5836	HAWKULAR-646	Is this still relevant? As part of HAWKULAR-661 / HAWKULAR-67 , there will already be an icon next to the username.	1	0	https://issues.jboss.org/browse/HAWKULAR-646
5837	HAWKULAR-638	[~mithomps], with whom could I talk about the implementation? I found an component that provides what we need (autocomplete, multiple items, tag the item selected): https://brianreavis.github.io/selectize.js/ Go to the email contacts example. It seems that there are options in angular also :)	1	1	https://issues.jboss.org/browse/HAWKULAR-638
5838	HAWKULAR-638	[~gcardoso]why don't you work with [~alexandrem] on this one.	1	0	https://issues.jboss.org/browse/HAWKULAR-638
5839	HAWKULAR-638	Sounds good! [~alexandrem], do u want me to do something or perhaps you implement and if needed I refine?	1	1	https://issues.jboss.org/browse/HAWKULAR-638
5840	HAWKULAR-638	That can work, [~gcardoso]. Not sure if a user can access the list of users.. or is this for users belonging to the same organization ? [~juraci.costa], can you tell what are the possibilities here ?	1	0	https://issues.jboss.org/browse/HAWKULAR-638
5841	HAWKULAR-638	OK, that's what I think that makes sense as well, users from the same organization. Do we need to do anything on the accounts side to achieve that currently ? I know from organizations we can list the users belonging to it, but maybe that's just for the owner ?	0	0	https://issues.jboss.org/browse/HAWKULAR-638
5842	HAWKULAR-638	As of now, the only solution would be to get a list of organizations, then a list of memberships for the organization. But I can add an endpoint at GET /hawkular/accounts/users that would retrieve the visible users based on the current user's permissions. One extra information: we do not have the user's email on the database, unless each user has added the email as the "preferred email" in the user settings. So, don't count on the user object containing an email address :) I can investigate how feasible would it be to get the user's emails from Keycloak.	0	0	https://issues.jboss.org/browse/HAWKULAR-638
5843	HAWKULAR-623	Don't we have that now?	1	0	https://issues.jboss.org/browse/HAWKULAR-623
5844	HAWKULAR-622	Hi, While working on a related task (HAWKULAR #665), I made a layout change request. Just wondering if we should make that here as well? Here's what I proposed: "would it be possible to reduce the padding between the left and right panels so there is less of a margin between them? It seems like if: "div class "col-md-9" had padding-left, and padding-right, of just 10px that would work. wdyt?"	1	1	https://issues.jboss.org/browse/HAWKULAR-622
5845	HAWKULAR-622	[~rutlucas] I think your suggestion of "once the user select Ack or Resolve, the Open state is not show in the select option list." works well. The Open option could either be disabled in the list, or simply not shown. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-622
5846	HAWKULAR-621	[~uxdlc], regarding this detail screen, how did you designed the interaction to change state, the user change in the dropdown AND needs to Save, or doesn't need to Save? Since this is not really a form, I'd suggest to not require saving to change the status (as it happens in the alert list). I'd not present the buttons Save/Cancel upfront. For the description, I'd do like jIRA, first you present only the text, if the user wants to edit, he clicks on a "Edit" icon next to the text, it becomes an textarea component, and he confirms the edition. Actually we are using this interaction in some modal (maybe deployments).	1	1	https://issues.jboss.org/browse/HAWKULAR-621
5847	HAWKULAR-621	I have just seen the addition of the Edit button. I have the opinion that in this case it can be redundant. The Edit means that the Alert could be editable, and it's not the case, we can change the state and add comments. In the first implementation I've uploaded into hawkular, the comments textarea and the save button is inactive until the state change from open to acknowledged or resolved. That can save clicks to an user. What do you think ?	1	1	https://issues.jboss.org/browse/HAWKULAR-621
5848	HAWKULAR-619	[~uxdlc], some doubts about your prototype: - The comment icon used is usually used as "attachment" by other apps. How about using the "chat bubble" icon to indicate comment? - How is the flow to acknowledge/close an item an add a comment? - What was the reason to put Acknowledge/Resolve in a dropdown? Before was easier imo.	1	1	https://issues.jboss.org/browse/HAWKULAR-619
5849	HAWKULAR-619	[~gcardoso], I've updated the JIRA as in the 616 umbrella there are two JIRAs per task, one per UxD work, and other for UI development. I was assigned to both of them, I update the UxD to Liz or you, just to tidy up and validate we don't have any missing thing. I guess with the proptotype is mostly done for UxD, now it's on my side. Please, could you update the status as the UX codes used (UXD provided or what do you think is appropiate) ? Thanks,	1	1	https://issues.jboss.org/browse/HAWKULAR-619
5850	HAWKULAR-617	Can we have a couple of small changes for this design ? - We have added in the alerts backend the possibility to add generic comments on an alert, so instead to have one comment, we might have a list of then with the data (comment, user, timestamp), if we have some design I can add it to the UI nicely to have it on MS6 cut. - Normally a notification occurs on a state (open, ack, resolved) so it can we worth to add the state also in the notification list. I can start implementing this on the UI and we can refine the design later, or wait until you have some design proposal. I think these two modifications were commented on the UX calls so it can be worth to implement something.	1	1	https://issues.jboss.org/browse/HAWKULAR-617
5851	HAWKULAR-617	[~rutlucas] Just a couple of questions about the enhancements: RE: Comments with user, timestamp and etc. * I thought we decided to add this in the next rev., do you remember why? * If it's a list, do you plan to show it in a scrolling input box, or etc., so the list presentation has a max height? Or are you thinking of showing all of the list content with no scrolling, so it would increase as much as it need and then increase the overall page height? RE: Notification & state: * So you are suggesting that we present the alert state in the notification? It's not that notifications have their own state? Thanks	1	1	https://issues.jboss.org/browse/HAWKULAR-617
5852	HAWKULAR-605	Based on Mazz's comment on the mailing list, can we just refer to Resource Name?	0	0	https://issues.jboss.org/browse/HAWKULAR-605
5853	HAWKULAR-601	Also, in a preliminar scope, I can see that alertManagers.ts could be simplified. There are some methods that are duplicated like addEmailAction() / createAction() or with similar function like createTrigger() / createAlertDefinition(). I think this can be revisit and unified to make it more easy to work with and maintain. What do you think [~mithomps] and [~alexandrem] ?	1	1	https://issues.jboss.org/browse/HAWKULAR-601

#	Requirement	Questions	Answered	Successfully Answered	URL
5854	HAWKULAR-590	In my opinion, in your prototype it is less noticeable being so close to the headings / tab. As we discussed in the call, I believe that the results you got from the test where a reflection of the question (to change the date/range in the chart context). Actually we found that people expect to navigate directly on the chart (for this use-case), but I'd not say that the issue is the position of the date picker (from what I heard from sj). As I suggested, I would investigate more if that's a real issue. I would test with the following scenario: "You accessed hawkular and is in the availability screen. How do you do to see the information related to the last 12h?". My bet is that the date-picker position won't be an issue, but let's check.	1	1	https://issues.jboss.org/browse/HAWKULAR-590
5855	HAWKULAR-590	Two alternatives for presenting date/time presets: # "Preset" buttons. (ref. JON) # Dropdown menu - which was inspired by: Consider Whether You Really Need Customers to Specify Dates - See more at: http://www.uxmatters.com/mt/archives/2011/08/date-filters-successful-calendar-design-patterns.php#sthash.3goz9kAo.dpuf * Note: I also moved the location of the date widget to try to increase its visibility. Wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-590
5856	HAWKULAR-590	In response to: 1) I'm not sure I understand the comment. If a preset button is selected, wouldn't that tell users which date/time range is selected? 2) You can check with SJ but some numbers of users did not understand that the text field was interactive. RE: Usability results. Please confirm with Sj. I observed all of the sessions, and noticed users having the following problems with the date picker: * Afore-mentioned problem with it not being understood as interactive. * Due to its location in the upper-right, users often didn't notice it. * The text was difficult for users to parse, particularly wrt it expressing a date RANGE.	1	1	https://issues.jboss.org/browse/HAWKULAR-590
5857	HAWKULAR-590	[~mithomps] Clicking on the calendar icon does expose the date picker buttons. Are you suggesting making this icon a button instead?	1	0	https://issues.jboss.org/browse/HAWKULAR-590
5858	HAWKULAR-587	Prototype on http://cardosogabriel.github.io/hawkular-client/application-datasources.html Some aspects of the interaction/visual design to be defined/clarified: - what happens when the user clicks on Edit? - how is the feedback for deleting it? I would guess a transition... - Inform visually which ones are disabled, so the user doesn't need to go into actions to know?	1	1	https://issues.jboss.org/browse/HAWKULAR-587
5859	HAWKULAR-587	RE: 1) Edit - I need to add that. I should be able to post it today. 2) Some feedback would be good. A transition effect when being removed from the list would be nice. wdyt? 3) Yes, inform visually. Serena suggested using "disabled" text color for the datasource name. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-587
5860	HAWKULAR-586	[~theute] How do you feel about using a dashboard style layout for this? Also, this "Overview" tab is replacing the server "Alerts" tab but it includes the alerts info.	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5861	HAWKULAR-586	Something along those lines sounds great to me. Few comments/suggestions after a quick look: * I would not allow for personalization anytime soon (So I would remove the part about "Customize tasks"), but this doesn't prevent that dashboard style layout. * We'll need to clear what charts we put there (what is useful)l, we can start small, maybe with what we already show in subsequent tabs. * We need to be able to edit the tags. * Export JDR report may be part of the common tasks, even though it would directly propose to save the file. * Should we reuse the date picker ? (Or change it everywhere)	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5862	HAWKULAR-586	Great, I'll keep iterating on this design then. Thanks for the comments/suggestions, some followup questions: * I'll remove the "Customize" link. * RE Charts: Yes starting small sounds good. * RE Label editing - Do you think it should be inline like the OS example that you posted? Or would it be useful to have a dialog that could be re-purposed for assigning/editing labels for other resources as well? * RE JDR button: I added it to the common tasks section as well. * RE Date picker: TBD. There is a JIRA for revisiting some Date Pickers issues: https://issues.jboss.org/browse/HAWKULAR-590 As part of that, we're thinking about alternate date picker presentations that showcase date presets. One idea is to use a drop-down, and another using preset buttons. Additionally, maybe it would make sense to use different widgets for different contexts/use cases? For example: - Using the preset dropdowns, for in-context views where space is limited. Or when changing dates for a subset of data instead of the whole view. - Using preset buttons, for where a user might want to quickly toggle between date ranges. - Using a fully featured custom-date, picker for more detailed troubleshooting or reporting.	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5863	HAWKULAR-586	re: 90% being good are bad, you have a good point. Can we skew the numeric values on the doughnut such at the 90 - 100 persent section is large and the smaller values take less and less space? I'm not wedded to a doughnut, but do want an eye catching graphic. I'm OK with the alerts on the right. We're building an alert center that will present existing alerts in a good amount of real estate. We should be able to link directly to that view from the server summary such that the view limits the alerts to alerts on the individual server. Let's reuse that view. In the summary I think we should prioritize critical issues.	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5864	HAWKULAR-586	Attached a version using gauges instead of doughnuts for URL metrics. And placing the alerts in a sidebar. I'm thinking the other tabs should also place the alerts in the right sidebar. What do others think?	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5865	HAWKULAR-586	[~uxdlc], didn't look in detail, but visually looks nice :) I see that the bg color of the top area changed to white. Also the breadcrumb changed, plus some spacing below the nav items. How will these changes affect other screens?	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5866	HAWKULAR-586	[~gcardoso] I noticed a little nit - the panel titled "Notifications" should be called "Common Tasks." And do you think we should use a badge in the header of the Alerts panel that presents the total number of alerts [3] ?	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5867	HAWKULAR-586	Ah ok, I updated it in the prototype. Regarding the badge, I think it makes sense to count the opened alerts. But how many will we list there? Another question, are they related to only the selected range, or all opened until now? I didn't put it cause we had alert badges in other places but it was not implemented, so if we want it we need to check with the dev team if they have this capability (counter).	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5868	HAWKULAR-586	What's missing for the labels ?	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5869	HAWKULAR-586	Btw any idea what kind of alerts should be visible in right panel? Are we going to have specific alerts defined for overview tab or should I display all opened alerts for this machine? I have merged all alerts that happened (for current machine and it's datasources) to the status overview panel (center panel which will show bar chart of alerts that were triggered), but I am not sure about the right panel with alerts. I guess we have 3 options * All alerts that are not resolved (including alerts for datasources) * All alerts for current machine that are not resolved * We will have definition which will handle alerts for overview panel And also, I guess the chart with data sources should show used connections for individual datsrouce (if there will be more than one datasource, there will be more than one chart with used connections) What to do you think [~theute], [~uxdlc], [~mithomps]	1	1	https://issues.jboss.org/browse/HAWKULAR-586
5870	HAWKULAR-581	bus and agent has been released with the new auth stuff (using the new accounts auth stuff). All that is left is the UI to now start passing in credentials via the "authentication" node in the JSON of all messages sent over WebSockets Assigning to Mike so he can figure out how the UI team will implement this. What this means is that all JSON messages from UI to the server via WebSockets needs to have something like this now: "authentication":{"username":"jdoe","password":"password"} or "authentication":{"token":"blah-blah-blah"}	1	1	https://issues.jboss.org/browse/HAWKULAR-581

#	Requirement	Questions	Answered	Successfully Answered	URL
5871	HAWKULAR-555	Taking the registered users into account, can we provide an "autocomplete" solution when inserting the email? Is that possible [~jpkrohling], [~mithomps]?	1	1	https://issues.jboss.org/browse/HAWKULAR-555
5872	HAWKULAR-555	[~gcardoso], would it be possible to annotate the mockups with the name of the components and/or with a sample markup? For instance: "this icon is called 'users' in FA, this is a 'well'/'jumbotron'/... wrapped in a .row > col-sm-6 " , or even "please use the layout 'xyz' from PatternFly".	1	1	https://issues.jboss.org/browse/HAWKULAR-555
5873	HAWKULAR-555	[~juraci.costa], the way we usually work is, I create a static html prototype + css, then mike/alex take that and integrate. Would you like to try that way? Then you don't need to guess html elements ;)	0	0	https://issues.jboss.org/browse/HAWKULAR-555
5874	HAWKULAR-552	Hi Liz. It's nice to see the effort to consistency. Suggestions inline: - As you said yesterday, the red background worked well for people. I believe we should keep it, specially for the open alerts, they call attention comparing with the other white boxes. - Another suggestion, to save space, you could remove the "created" text. If you only put the date, it is still understandable (like we have currently). - If I understood correctly, it was said that seeing the resolved alerts is not the focus of the user. In this sense, I would do the opposite, not displaying the Resolved alerts initially, and presenting an action "Show resolved alerts". I'm attaching Wunderlist's design as reference, which I consider a good solution. We have a style for this in the Datasources tab, the toggle "View/Hide Metrics". I think it could be positive to use the same interaction style. That way we also don't need to move the Alert Settings link. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-552
5875	HAWKULAR-552	If we were to take your idea a step further, I think there might be an argument for using alert "notifications" in this context (area) instead of a filtered alerts list. And by that I mean something more akin to the toaster notifications, rather than a static list that's a subset of the list alert center. If we went with something more like a notification, then we could deviate from the alert center list format. From what I've observed from the usability studies, the toaster notifications aren't noticed a lot of the time. And in the patternfly work, there has been some discussion about moving them into the main area of the screen. So maybe we can bring all of this together for hawkular. And by doing so it would help resolve JIRA 439. Wdyt?	0	0	https://issues.jboss.org/browse/HAWKULAR-552
5876	HAWKULAR-552	What is the little down arrow after the status? The URL area will need to be much larger than what is shown.	1	1	https://issues.jboss.org/browse/HAWKULAR-552
5877	HAWKULAR-552	My question is more why to use such component (checkbox) for filtering? How about a select dropdown, that we already use in other places? And if we want to allow a more powerful filter, it could be even the multi-select, where the user could choose open, acknowledged or resolved, or a combination of some. https://www.patternfly.org/widgets/#bootstrap-select	1	1	https://issues.jboss.org/browse/HAWKULAR-552
5878	HAWKULAR-551	[~uxdlc], I looked at the pattern and actually they don't have something like us (object + label). What they have is object (top) + part of the object / action (content). I personally think that our current presentation is easier to read than what you proposed, and considering that they don't have something like us, I'd check with Serena, but it should be fine to leave like this. What is beneficial and make as closer to PF is to have the spaces between the tiles. I personally like the blue stripes above for aggregators in page view (which is kind of a dashboard), it calls attention in a positive way for aggregators IMO. As for the url list for example, they don't make sense. How about trying to see? Regarding the larger font, I'm considering to reduce the narrative values. "a few seconds ago" could be "a few ago", and "2 minutes ago" could be "2 min ago". wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-551
5879	HAWKULAR-551	Example attached: * Reduced Value text to 20px. * Reduced Alert text to 12px, with reduced margins to 10px. * Reduced alert icon size to 18px. * Added margin between tiles, but not as big as patternfly. Made it 5px, to preserve space for tiles. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-551
5880	HAWKULAR-551	[~uxdlc], you can check on the livingontheedge instance. I didn't apply the new style in two designs that are different: - URLs list - Application Datasources For the urls list, it makes sense to migrate to the new tiles design. The application datasources might need some adjustments. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-551
5881	HAWKULAR-544	[~gcardoso] Gabriel, can you verify that this doesnt break any of the layouts visually? Also verify that controls like date picker still work.	1	1	https://issues.jboss.org/browse/HAWKULAR-544
5882	HAWKULAR-544	[~mithomps], updated to the last patternfly (2.2.0) in my prototype, made a few adjustments and it looks fine, anything broken. What u have to do on dev: - update patternfly - update links in the <head>. Some paths changed, and also the less/css files, due to a patternfly update - in hawkular.less, import patternfly-additions.less instead of patternfly.less: - @import "../bower_components/patternfly/less/patternfly.less"; + @import "../bower_components/patternfly/less/patternfly-additions.less"; - update metrics.less With that it should work. Please let me know if you need something else.	1	1	https://issues.jboss.org/browse/HAWKULAR-544
5883	HAWKULAR-544	[~gcardoso] Why don't you take a stab at this in a branch it can't hurt anything. And even if Alex or I need to tweak, it will be close.	0	0	https://issues.jboss.org/browse/HAWKULAR-544
5884	HAWKULAR-541	What is the status of this, [~ppalaga]	1	1	https://issues.jboss.org/browse/HAWKULAR-541
5885	HAWKULAR-540	Is it a datasource related to JDBC? I'm trying to figure out how related this flow would be to the "Add JDBC" driver flow: https://issues.jboss.org/browse/HAWKULAR-516	1	1	https://issues.jboss.org/browse/HAWKULAR-540
5886	HAWKULAR-540	By studying your design, I believe we should provide inline actions for datasources, as we provide for the URLs list. The edit button doesn't work well with multi-select, given that you only edit one datasource at a time. Moving to inline would keep consistency with similar screen (url). Space is not a problem, we can adjust to fit. Also, we would not need to introduce a new visual style for Add, we could keep using the button with text. Could you try a variation [~uxdlc]?	1	1	https://issues.jboss.org/browse/HAWKULAR-540
5887	HAWKULAR-540	Its an atypical use case, and I'm not saying that we should optimize for it. But we should make considerations for *many* datasources to ensure we can scale well. In this design I'm not sure how even 30 datasources would scale. Is that in the design currently?	1	1	https://issues.jboss.org/browse/HAWKULAR-540
5888	HAWKULAR-540	Implemented the modal for Add Datasource: http://cardosogabriel.github.io/hawkular-client/application-datasources.html [~uxdlc], we usually have the labels at the side of the inputs. Should we keep the first screen with labels above? I'll start implementing the datasources screen. Will you provide a hi-fi prototype?	1	1	https://issues.jboss.org/browse/HAWKULAR-540
5889	HAWKULAR-540	[~ppalaga] Peter, what is the correct Datasource Type list for the drop down in step 1?	1	1	https://issues.jboss.org/browse/HAWKULAR-540
5890	HAWKULAR-540	[~ppalaga] Peter, which step/screen of the wizard did you end up adding "xa-datasource-class" to? Thanks.	1	1	https://issues.jboss.org/browse/HAWKULAR-540
5891	HAWKULAR-540	Is this UxD done, [~uxdlc], [~gcardoso] ?	1	1	https://issues.jboss.org/browse/HAWKULAR-540
5892	HAWKULAR-538	[~uxdlc]: how is the interaction to go from the list to the details screen? Right now when the user clicks anywhere in the row (but the checkbox) it becomes selected. Maybe we should do like gmail, click on the row, it goes to details, and select only on the checkbox?	1	1	https://issues.jboss.org/browse/HAWKULAR-538
5893	HAWKULAR-538	Could you run this idea by Patternfly/Serena? I didn't see anything in her wireframe doc that outlined the selection behavior so I can't say what her recommendation for the selection behavior is. PF states this about table selection: https://www.patternfly.org/patterns/table/	1	1	https://issues.jboss.org/browse/HAWKULAR-538

#	Requirement	Questions	Answered	Successfully Answered	URL
5894	HAWKULAR-538	RE: Alert Details screen, alert details section - the HTML prototype does not match the wireframe. Uses a 2-col. layout center-aligned text, instead of single column left-aligned... Were there technical issues?	1	1	https://issues.jboss.org/browse/HAWKULAR-538
5895	HAWKULAR-538	I would like to title the alert details with the Alert ID, so I asked Jay if alerts have unique ID's. Here's what Jay said (edited snippets): “they do. Every alert has an alertId. It's currently generated as a concatenation of the triggerId+createTime. What makes an alert unique is the tenant, triggerId, alert create time. The tenant is likely an implied value that already limits your view. The triggerId is not really for display purpose, So, I'd say the trigger name and the alert create time. In the table you linked to me I don't see a trigger name anywhere. I'm not 100% sure what the "Alert Message" column is supposed to be but I presume it is supposed to be any note that has been put on the alert for the current State. So, for Open I'm not sure, maybe it's unset or maybe it's the trigger description. For Acknowedged/Resolved it would be the optional not added by the ack-er or the resolver. if you needed to save real estate I'd probably drop the Alert Message column and add the trigger name. Or perhaps combine the created and resolved columns into just a single State Time, which would show the timestamp associated with the current state. First, it makes a big assumption that there is a message, which ties it very much to the actions performed for the alert. Second, the alert does no carry that information around, so you'd somehow have to make separate calls to get the info. third, it's likely big and would never fit nicely in a column, it could even be html. ...but mainly, I just don't think it's available..." --jshaughn So a proposal for the Alerts Center table then is to: - Remove/reword the "Alert Message" column and add a "Trigger Name" column. - Combine the created and resolved columns into just a single "State Time", which would show the timestamp associated with the current state. Feedback?	1	1	https://issues.jboss.org/browse/HAWKULAR-538
5896	HAWKULAR-538	[~mithomps][~jdoyle] Based on Jay's input, we need to make some modifications. Mike, does this match your understanding? 1) Alert Center > Alerts List screen: Data columns for alert center should be changed to: * State * Severity * Ancestry (replaces host name & resource type) * Trigger Name (replaces Alert message) * Created Date (remove Resolved date) 2) Alert Center > Alert Details screen: * Should be titled with Trigger Name. * Include ancestry path details. * Link to specific trigger settings screen. ("nice to have").	1	1	https://issues.jboss.org/browse/HAWKULAR-538
5897	HAWKULAR-538	[~uxdlc], as soon as you clarify, can you provide a protoype?	1	1	https://issues.jboss.org/browse/HAWKULAR-538
5898	HAWKULAR-538	[~gcardoso] I created a mockup using the current presentation that you've implemented, here's the change-list: * Revised the columns based on new requirements. * Added Acknowledge" button. * Removed color coding for alert states after a discussion with John. We don't need to call attention to the states so dramatically any more. In addition, the new PF hover/select colors might conflict. So instead, I simply bolded the open state. * Added an example ancestry path - we'll need to allow for truncating, but the string should be truncated in the middle if possible. * I added an icon to guide users to click to view more details. Wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-538
5899	HAWKULAR-538	[~mithomps], saw that you closed it. How about the last updates that I didn't implement?	1	0	https://issues.jboss.org/browse/HAWKULAR-538
5900	HAWKULAR-537	Is that still a thing?	1	0	https://issues.jboss.org/browse/HAWKULAR-537
5901	HAWKULAR-527	While I like the idea of linking, I wonder what the semantics and value of the link is in the described case of {quote} access the URL itself, so if Hawkular reports the URL as being down, I could just click on it and access the target server. {quote} When e.g. http://hawkular.org is reported as down, how would accessing it via browser (in new tab!) improve the situation? How would we deal with urls with a local part like https://issues.jboss.org/browse/HAWKULAR-527 -- link to the full url, the server or .. ? Especially as in theory the server could be up, but that specific local part could have vanished (e.g. corresponding web-app undeployed)	0	0	https://issues.jboss.org/browse/HAWKULAR-527
5902	HAWKULAR-519	From de backend perspective, does the server need to wait for one action to finish to start another one, or can the user start various at the time? That's important to define the interaction.	1	1	https://issues.jboss.org/browse/HAWKULAR-519
5903	HAWKULAR-519	I'm looking at my prototype, and the uxd team had decided that in that table we would not have inline actions. So basically we will only have the top "actions" button, that work when one item is selected at least. Considering this approach, I suggest us to block the action(s) until the process/processes are finished. The visual representation would be the spinner inside the actions button and make it disable. Important, that depends on how the WildFly server behaves with multiple actions. [~mithomps], could you check how does it work? If it doesn't, we should get rid of the multiple actions and use only inline actions...	1	1	https://issues.jboss.org/browse/HAWKULAR-519
5904	HAWKULAR-516	RE: the wireframes that I posted, I was principally concerned with making sure the flow and fields are correct. The chevrons are in there because that's one of the canned indicators that I found in Balsamiq. Any effective step indicator should be fine. If the general flow is ok, we could port the flow to whatever wizard framework makes the most sense. Do you have more screenshots, or a pointer, to an interactive example?	1	1	https://issues.jboss.org/browse/HAWKULAR-516
5905	HAWKULAR-516	Yes, but probably next week only. I'm finishing something else now. When would you need to have it?	1	1	https://issues.jboss.org/browse/HAWKULAR-516
5906	HAWKULAR-516	Thanks for the screens, [~uxdlc] and [~gcardoso]. Two things: (1) To create a minimal driver, we need at least one more field called "Driver Class". WF wizard has it too. The screen "Step 3 of 5: Datasource Driver Add Driver: Define Parameters" seems to be the best place for it. (2) There is a couple of other things that are configurable on the backend, when adding a new JDBC driver. The question is how good do we want to cover all options supported by the backend? Let me explain it in detail: Creating a JDBC Driver on WF is a two step process: (2.1) Create a JBoss Module containing the driver.jar and (2.2) Register the driver in the Datasources subsystem. (2.1) Create a JBoss Module containing the driver.jar is basically a filesystem operation: one copies one or more jar files to {{$JBOSS_HOME/modules/my/module/name/main}} directory and creates a {{module.xml}} file there that describes the module. ATM the screens do not cover the following options (see https://docs.jboss.org/author/display/MODULES/Module+descriptors for what they mean): * {{slot}} - used so that one can have two variants for one module name. The dafault slot is "main" * {{resources}} - typically a list of jar files, we have field for a single jar file, which is IMO typically enough for JDBC drivers. * {{dependencies}} - a list of other modules this module depends on. The user should IMO be able to be able to edit this. It can be a simple * {{properties}} - I have no idea if we need them. (2.2) Register the driver in the Datasources subsystem - see http://wildscribe.github.io/Wildfly/9.0.0.Beta1/subsystem/datasources/jdbc-driver/index.html > Operations add. The options, we do not have ATM: * {{driver-class-name}} the "Driver Class" I named in (1) - we need it. * {{driver-datasource-class-name}} The fully qualified class name of the javax.sql.DataSource implementation - no idea if we need this. * {{driver-xa-datasource-class-name}} The fully qualified class name of the javax.sql.XADataSource implementation - no idea if we need this. * {{profile}} Domain Profile in which driver is defined. Null in case of standalone server - no idea if we need this. [~theute], [~pilhuhn] and [~jdoyle] could you please comment?	1	1	https://issues.jboss.org/browse/HAWKULAR-516

#	Requirement	Questions	Answered	Successfully Answered	URL
5907	HAWKULAR-516	[~gcardoso] Per Peter's input - we need to add a "Driver Class" form fields to "Step 3 of 5: Datasource Driver Add Driver: Define Parameters." Do I need to update the wireframes or are you good with the change? P.S. I still need to look at your prototype, but the changes sound good. Thanks.	1	1	https://issues.jboss.org/browse/HAWKULAR-516
5908	HAWKULAR-516	[~uxdlc], which list are u talking about?	1	1	https://issues.jboss.org/browse/HAWKULAR-516
5909	HAWKULAR-501	this might be difficult right now because there isn't a DMR resource to obtain this from that I know of. So there is no way to to provide <resource-type-dmr> or <resource-config-dmr> metadata to which you can attach this. It might require some custom "ad-hoc" coding inside MonitorService or... somewhere. I don't know where.	1	1	https://issues.jboss.org/browse/HAWKULAR-501
5910	HAWKULAR-488	[~theute] and/or [~pilhuhn], could you please confirm that "IP Address" and "Powered By" should really be added to Application Servers table? In case they should, it would help me a lot to get some kind of a definition of the data to show. The info we present in the App servers list ATM comes from the Agent (cc [~mazz]) and I guess, it should be the Agent providing this info rather than Pinger. The Agent probably can get the list of IP addresses the AS is bound to using internal APIs of WildFly. However, there is no guarantee that any of these addresses would make sense to present in the list of Application Servers in the UI. I doubt that showing 127.0.0.1 tells the user something interesting? Private addresses 10.0.0.x, 192.168.0.x, etc. can be interesting to some extent and there is no doubt that showing (one or more) public addresses accessible globally would make sense. So, please provide some directions which IP addresses to show. As for "Powered by", it makes no sense to me at all. If you insist on displaying it, please define where to get that data.	1	1	https://issues.jboss.org/browse/HAWKULAR-488
5911	HAWKULAR-483	right now, I think the prototype is actually able to work. Requires bus 0.3.2.Final and agent 0.3.2.Final UI needs to pass the following over the websocket: {code}ExecuteOperationRequest={"resourceId":"the-id", "operationName":"the-op"}{code} I'm working on getting this to work better. Needs to probably utilize inventory API so I can get the feed ID (right now I rely on the resource ID to be prefixed with the feed ID followed by "~" but that means it only works with inventory from wildfly agents. Also eventually will need to hook into alerts. How/where to store the operation results? Don't know the answer to that either.	0	0	https://issues.jboss.org/browse/HAWKULAR-483
5912	HAWKULAR-478	I provided my input above: Date/time stamp, for the reasons stated, and referring to #336 for formatting. Mike - if we need Catherine's input, could you please add her to the watchers? Thanks.	1	1	https://issues.jboss.org/browse/HAWKULAR-478
5913	HAWKULAR-478	[~uxdlc], in my understanding this is something that most services are using, while they could all use timestamps. Are timestamps really more readable ? can a user quickly tell that "2015-07-26 16:00:00" is different from "2015-07-28 16:00:00", or is it easier to distinguish "3 days ago" and "3 hours ago" when you read/glance over it ? Obviously at some point, "X time ago" is not helping, thus we move to a more concrete date as you've shown in the screenshot with "April 1".. and that's what other major services do (eg, see HAWKULAR-463 Dates on the column to the right.. created is a full timestamp, updated is fuzzy). The way I see it, the page should be talking to the user, I read that as the user asking "\[When was\] Last downtime?" and the page answering "\[It was\] 2 days ago", as a person would naturally answer, not with a timestamp.. unless it was asked, "But when exactly?". As for not providing precise information, I think these boxes are about providing an overview, a first glance at data.. the user should know if he sees the last downtime as "2 days ago" if he expects it or not, and if not go investigate it where the actual full, precise data is at. Because we are actually just showing the last downtime, which even may be something he expected (like an upgrade or reboot), but before that last one, there were several others that are unexpected and are not shown.. so these values should not be taken as the place for precise data. Regarding the tooltips, 1.5s is the time a "regular" tooltip takes to display in Chrome, I just did as Chrome. I'm fine with lower delay, as long as it's not immediate tooltip, as it's not good to be moving the mouse around and tooltips popping out.	1	0	https://issues.jboss.org/browse/HAWKULAR-478
5914	HAWKULAR-474	Prototype done on http://cardosogabriel.github.io/hawkular-client/url.html, tabs Availability and Response Time. [~uxdlc], I adjusted the design to be consistent with the other modal boxes, so the on/off switch is placed at the side of the time input. We should define what happens when the switch is off. You said to turn the input disabled (which I agree), but in the case when all switches are off, it doesn't make sense to expose Alert / Notification. WDYT about hiding these areas in this case?	1	1	https://issues.jboss.org/browse/HAWKULAR-474
5915	HAWKULAR-468	About password confirmation: on the usual scenario, it should be possible. But what about the scenarios where the user is logged in via SAML or Kerberos for instance? On this situation, neither we nor Keycloak have access to the user's password. What should be done in such cases?	1	1	https://issues.jboss.org/browse/HAWKULAR-468
5916	HAWKULAR-468	[~uxdlc], it was suggested to have this when we have the alert center. Do you have some prototype for it already, or something that can give some basis to think about the settings screen?	1	1	https://issues.jboss.org/browse/HAWKULAR-468
5917	HAWKULAR-468	I agree that alerts is diff in that it's in context. Could you post a wireframe/prototype of the sidebar idea, to illustrate the full interaction?	1	1	https://issues.jboss.org/browse/HAWKULAR-468
5918	HAWKULAR-468	Created https://issues.jboss.org/browse/HAWKULAR-638. [~theute], are the screens in accordance to what we spoke?	1	1	https://issues.jboss.org/browse/HAWKULAR-468
5919	HAWKULAR-468	what "our side" means here ? Where will it be stored ?	0	0	https://issues.jboss.org/browse/HAWKULAR-468
5920	HAWKULAR-468	I am trying to follow the thread and I have some doubts. In the triggers, we are going to define email as actions/users/both ? I understand that we might have something like email:jdoe@hawkular.org or user:jdoe. For the second one there is defined a JIRA https://issues.jboss.org/browse/HWKALERTS-80 to work on it. Basically the preferred idea would be a service that a plugin can invoke to resolve if user:jdoe should be an email, and sms, or other preferred transport. I guess that perhaps today we want something simple but in the future preferences could be more complex (i.e. email by day, sms by night, or other kind of rules). Also, I don't know if this is the proper JIRA to comment an important feature of alerts that might be interesting for the UX. Today, the UI creates an action per email in a very simple way, and in the forms look like you have to specify the email always. Alerts offer to create an action definition (i.e. group-of-admins) and specify in the definition the details (recipients or other properties supported by the plugin). In that way, when you are creating a trigger, you might want to select an existing group. This has the advantage that tomorrow you can add/remove recipients of an action in a more easy way.	0	0	https://issues.jboss.org/browse/HAWKULAR-468
5921	HAWKULAR-468	Are we OK then with the approach of storing this data on Accounts? If so, I'll assign this JIRA to me and work on both the backend and frontend (based on the screen provided by [~gcardoso]). Other components would then be able to retrieve a "UserSettingsService" via CDI, that would return a Map<String, String> with all settings, or a more specialized method returning a simple "String" for a given user and given setting name (ie: "hawkular.settings.alerts.notificationemail").	1	1	https://issues.jboss.org/browse/HAWKULAR-468
5922	HAWKULAR-468	Well, I guess that with this change, then the alerts definition form will change. Does anybody know if with the email defined in the user settings, then the alert definition form will let the user to choose between an specific email or implicit user ? If this is discussed in other JIRA, please let me know the # as perhaps this is off-topic for this one. But from alerts perspective is interesting to know to decide which implementation is more appropiate and what we should ask to accounts for it .	0	0	https://issues.jboss.org/browse/HAWKULAR-468

#	Requirement	Questions	Answered	Successfully Answered	URL
5923	HAWKULAR-468	Guys, this JIRA is related to https://issues.jboss.org/browse/HAWKULAR-638. Can you take a look at this other one?	1	1	https://issues.jboss.org/browse/HAWKULAR-468
5924	HAWKULAR-468	PR submitted for the Hawkular console, based on the HTML provided by [~gcardoso]. I'm not sure what the "Cancel" button should do, though. Perhaps should I remove it, as it was possibly done with the idea that this settings screen would be outside of the console? For Accounts, it can be used via the UserSettingsService . You can also inject a setting via CDI. There's some documentation in the format of javadoc, as well as a sample here: https://github.com/hawkular/hawkular-accounts/blob/master/sample/src/main/java/org/hawkular/accounts/sample/boundary/SampleEndpoint.java#L96-L101	0	0	https://issues.jboss.org/browse/HAWKULAR-468
5925	HAWKULAR-465	[~uxdlc], we don't have these two levels of headings in the other modals. Also, they basically replicate the content of the tab. Can we do like you did for JVM, an introductory text? Or improve the labels in tabs?	1	1	https://issues.jboss.org/browse/HAWKULAR-465
5926	HAWKULAR-463	Gabriel - Mike suggested that we use spin boxes instead of sliders, so I'm attaching a revised mockup for non-heap. Wdyt? If you agree, I'll update Heap as well.	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5927	HAWKULAR-463	[~uxdlc], as an improvement, I would make more clear that an alert is a message that appears in the UI, and a notification is an email sent. How about changing the label "Create Alert" to "Display Alert in the UI", and "Notify" to "Notify via Email"? How did u define the behaviour for Save? Does it apply to the current tab, or all tabs?	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5928	HAWKULAR-463	Implemented, prototype on: http://cardosogabriel.github.io/hawkular-client/application.html > JVM > JVM Alert Settings [~uxdlc], I implemented with small changes, trying to improve when I found something confusing/inconsistent. Please let me know if there is something that you disagree. We are proposing some default values (usage percentages, alert duration, garbage duration). Where these numbers put intentionally or just to illustrate?	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5929	HAWKULAR-463	* Thanks for adding the "Save All," I didn't update the wireframes with that change. I noticed that you changed the button order (Cancel is left now.) Either way works for me as long as it is consistent throughout. * Could we change the form label weight from 600, to 500? Please see attached image. At 600, I think they visually compete with the other labels. * New spin boxes look good to me. * When Toggle is in OFF state, could we disable the corresponding form input? * Garbage Collection - I would prefer if we kept with the wireframe design because I would like to avoid wrapping form labels where possible. And I think it's helpful to emphasis Major vs. Minor. I made a mockup, attached.	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5930	HAWKULAR-463	* PF button documentation is conflicting then, please see: https://www.patternfly.org/wikis/patterns/forms/form-layout/ Again, I am agnostic on this point so whatever PF aligns on. * Can I see your latest version? I would like to verify its readable. PF doesn't cover every use case, so I need to verify it works for ours. * Please see the attached mockup for Major.	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5931	HAWKULAR-463	- Yes, I'm using the patternfly widgets (code), that should be more recent than the documentation itself. - You can see my latest version in the same link for my prototype that I posted above, just refresh it. - Will update according to major. We need a text for Garbage Collection, it is a placeholder now. Can you send it to me Liz?	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5932	HAWKULAR-463	Regarding "Considering that we will not have minor GC, we don't need to have the major/minor headings anymore. Updated the prototype without the minor input": We don't know if the WF 'Accumulated GC Duration' is combined major + minor GC or... Can someone confirm this? [~mazz] [~pilhuhn] [~alexandrem]	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5933	HAWKULAR-463	Could someone verify if the current HTML, or other coded version, includes the checkbox for the email notification. Or any other type of ON/OFF setting?	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5934	HAWKULAR-463	[~uxdlc], can you please file a separate JIRA for the on/off notification thingy?	1	1	https://issues.jboss.org/browse/HAWKULAR-463
5935	HAWKULAR-439	For error I think it is nice, cause the user can go back and try again. But for success not so cool, cause the user already finished what he had to do in that modal (and still has to close the modal). I would rather close it when clicking on Save and show the toastr of success, it is more fluid interaction. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-439
5936	HAWKULAR-431	I personally like this dark tooltip, which is the same style we already have in the console and similar to what patternfly is proposing for graphs. However, [~uxdlc] was proposing a white tooltip for response-time graph. IMO it would be good to use the same visual style in both graphs. Liz, in which direction do you want to go? The white one is more suitable if we have the graph zoomed-in. Mike, will this feature be available for response time charts?	1	1	https://issues.jboss.org/browse/HAWKULAR-431
5937	HAWKULAR-431	Could you post a link to the prototype? Thanks!	1	1	https://issues.jboss.org/browse/HAWKULAR-431
5938	HAWKULAR-431	[~mithomps], for the white widget, do you want me to provide the html / css of the widget? Please let me know which way is nice to work.	1	1	https://issues.jboss.org/browse/HAWKULAR-431
5939	HAWKULAR-431	[~uxdlc], [~mithomps], what do you think about using a tooltip with white background for this graph, in order to have the same visual style as the response-time tooltip? chart-availability-hover.png [~mithomps], will we have the tooltip for the response time graph also in this milestone?	1	1	https://issues.jboss.org/browse/HAWKULAR-431
5940	HAWKULAR-431	Yes. So what do you propose?	1	1	https://issues.jboss.org/browse/HAWKULAR-431
5941	HAWKULAR-429	Looks like {{com.github.marschall:jdeps-maven-plugin}} does not have a goal that makes the build fail in case of internal deps. The last of the three plugins {{org.codefx.maven.plugin:jdeps-maven-plugin}} has just this: https://github.com/CodeFX-org/JDeps-Maven-Plugin	1	1	https://issues.jboss.org/browse/HAWKULAR-429
5942	HAWKULAR-422	While I think it's a good idea, and I prefer the visual communication most of the times, I don't think it will be trivial to communicate some things.. eg, 99.999% is not a "bad" thing. 100% is perfect, but, 99.999% is still good. How do we differentiate ? The same applies for the "Last Downtime", it may have been 6 months ago, which again, is not perfect (Always Up) but is not bad, and should not be marked in red or as something less positive.	1	0	https://issues.jboss.org/browse/HAWKULAR-422
5943	HAWKULAR-399	Is that still relevant?	0	0	https://issues.jboss.org/browse/HAWKULAR-399
5944	HAWKULAR-398	The wildcard is handled "special" inside the agent. Are you using an older agent build by any chance? This should be working in the latest agent release, I looked at it yesterday and I was seeing non-zero values getting sent to metrics.	0	0	https://issues.jboss.org/browse/HAWKULAR-398
5945	HAWKULAR-393	This happened when we adopted the gray color for the down status of an application server (see attachment). [~uxdlc] had a rationale for that. Does this apply to the urls list?	1	1	https://issues.jboss.org/browse/HAWKULAR-393
5946	HAWKULAR-393	[~pilhuhn], why was the PR declined?	0	0	https://issues.jboss.org/browse/HAWKULAR-393
5947	HAWKULAR-389	What do you mean by alert dialog boxes? Can you elaborate the usecase?	1	1	https://issues.jboss.org/browse/HAWKULAR-389
5948	HAWKULAR-389	When you click on 'Alert Settings' link it pulls up a dialog box to allow the user to set various alerts on metrics. We never decided what those screens should look like? Probably too late now for this sprint.	0	0	https://issues.jboss.org/browse/HAWKULAR-389

#	Requirement	Questions	Answered	Successfully Answered	URL
5949	HAWKULAR-389	Revised "add deployments" alert settings to remove dampening. Doesn't seem to make sense in this case. wdyt?	1	1	https://issues.jboss.org/browse/HAWKULAR-389
5950	HAWKULAR-378	[~uxdlc] I'm not familiar with Jira - how do I see the linked Jira with new designs? Thanks.	1	0	https://issues.jboss.org/browse/HAWKULAR-378
5951	HAWKULAR-370	Was this flow already discussed? Let me know if or when I can start.	1	1	https://issues.jboss.org/browse/HAWKULAR-370
5952	HAWKULAR-370	[~mithomps], sent the PR:https://github.com/hawkular/hawkular/pull/437. The html structure changed a bit. I just did small changes to fit better in our modals for now, since we don't have a proper design for it. In step 2 you need to hide the "discard/confirm buttons" after clicking on editing. Please check the behavior in my implementation: http://localhost:1337/application-deployments.html . Also, the tooltips are not working for the buttons. [~uxdlc], a visual style for this new steps component need to be defined (see attachment). Does patternfly have something? Probably makes sense to place it at the top and remove our current heading (step 1 of 3), which is redundant right now.	1	1	https://issues.jboss.org/browse/HAWKULAR-370
5953	HAWKULAR-367	[~jayshaughnessy] or [~lponce], could you please give me some exact URLs and payloads to create (trigger -> dampenings -> condition for) such scenario?	1	1	https://issues.jboss.org/browse/HAWKULAR-367
5954	HAWKULAR-367	Thanks [~jayshaughnessy] for you comments and again for the [docs|http://www.hawkular.org/docs/dev/alerts.html] which are really helpful. [~mithomps], I can see now two ways to go: # Provide a default {{autoResolve}} dampening/condition for threshold and availability. The threshold value would be based on the existing {{firing}} dampening/condition value. The time would be a constant right now (let's say, N minutes). There are several use cases for that: #* *Threshold _Every time the maximum time is exceeded_*: No {{autoResolve}} condition would be provided and we will fire alerts the same way as we do now. #* *Threshold _Only when exceeded for greater than_*: The {{autoResolve}} condition would be: _when below threshold for N minutes_. #* *Downtime*: The {{autoResolve}} condition would be: _when UP for N minutes_. # Make the {{autoResolve}} part configurable by users. Since we don't have UI mockups I would vote for #1 with N = 5 minutes. Once #1 is implemented, #2 would be very easy to implement. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-367
5955	HAWKULAR-367	Do you need some UXD for this milestone? Can you clarify what needs to be done?	1	1	https://issues.jboss.org/browse/HAWKULAR-367
5956	HAWKULAR-362	Should we revert that to "new" as 1.3.0 is busted and then re-submit an updated-PR?	1	1	https://issues.jboss.org/browse/HAWKULAR-362
5957	HAWKULAR-361	Would it be possible to have a regular scale (starting at -0-, with the y axis including the threshold), but offer a "macro lens" tooltip to zoom in on detailed info? Please see attached mockup. My concern with the logarithmic (as I understand it), is that it might still cause some confusion: https://en.wikipedia.org/wiki/Misleading_graph#Improper_use_of_logarithmic_scaling Zooming aside - I think we could also make the chart easier to "read" if we used background shading to illustrate "exceeded threshold" rather than line fills.	1	0	https://issues.jboss.org/browse/HAWKULAR-361
5958	HAWKULAR-356	Is that still an issue [~gcardoso] ?	0	0	https://issues.jboss.org/browse/HAWKULAR-356
5959	HAWKULAR-355	Definitely more info on hover is necessary. Thinking about the user's need, I think it is valuable to inform the exact time and exact response time on hover. WDYT?	1	0	https://issues.jboss.org/browse/HAWKULAR-355
5960	HAWKULAR-355	It looks better. I uploaded a suggestion with the following changes: - added a circle that explicits the intersection between time and response time in the chart (that should be the color of the line also...) - replicated this same circle in the tooltip - updated the color of the vertical line (hover) to match the color of the line inside the tooltip - removed the black circle from the time axis. WDYT?	1	0	https://issues.jboss.org/browse/HAWKULAR-355
5961	HAWKULAR-355	How would this look in a multi-line chart?	1	1	https://issues.jboss.org/browse/HAWKULAR-355
5962	HAWKULAR-355	Last mockup didn't come. I like your last comment because it is exactly what patternfly did. Have you seen this? https://rawgit.com/patternfly/patternfly/master/tests/charts.html	1	0	https://issues.jboss.org/browse/HAWKULAR-355
5963	HAWKULAR-355	We agreed to use the white background in both to be consistent. [~mithomps], do you need something from me?	1	0	https://issues.jboss.org/browse/HAWKULAR-355
5964	HAWKULAR-355	[~mithomps], how do we proceed here, it was assigned to me.	1	0	https://issues.jboss.org/browse/HAWKULAR-355
5965	HAWKULAR-351	One issue issue is during startup, where C* is not yet up and the pinger is pinging already 17:09:01,300 INFO [com.datastax.driver.core.RequestHandler] (New I/O worker #11) Query SELECT triggerId, triggerMode, type, conditionSetSize, conditionSetIndex, conditionId, dataId, operator, data2Id, data2Multiplier, pattern, ignoreCase, threshold, operatorLow, operatorHigh, thresholdLow, thresholdHigh, inRange, tenantId FROM hawkular_alerts.conditions WHERE tenantId = ? AND triggerId = ? ORDER BY triggerId, triggerMode, type is not prepared on /127.0.0.1:9042, preparing before retrying executing. Seeing this message a few times is fine, but seeing it a lot may be source of performance problems 17:09:01,330 INFO [com.datastax.driver.core.RequestHandler] (New I/O worker #11) Query SELECT tenantId, triggerId, triggerMode, type, evalTrueSetting, evalTotalSetting, evalTimeSetting, dampeningId FROM hawkular_alerts.dampenings WHERE tenantId = ? AND triggerId = ? is not prepared on /127.0.0.1:9042, preparing before retrying executing. Seeing this message a few times is fine, but seeing it a lot may be source of performance problems 17:09:02,193 WARN [org.hawkular.component.pinger] (EJB default - 2) Start: Tue Jun 16 17:09:01 CEST 2015, t2: 664, t3: 664 17:09:02,194 WARN [org.hawkular.component.pinger] (EJB default - 2) Got status code 200 from http://bsd.de in [671]ms 17:09:05,169 INFO [com.datastax.driver.core.RequestHandler] (New I/O worker #2) Query UPDATE data USING TTL ?SET m_tags = m_tags + ?, n_value = ? WHERE tenant_id = ? AND type = ? AND metric = ? AND interval = ? AND dpart = ? AND time = ? is not prepared on /127.0.0.1:9042, preparing before retrying executing. Seeing this message a few times is fine, but seeing it a lot may be source of performance problems Here it took 671ms in total for a socket read from the remote in ~28ms This may be less of an issue with an external C* instance	0	0	https://issues.jboss.org/browse/HAWKULAR-351
5966	HAWKULAR-343	Attached is the design for a response time alert notification. Besides an email for availability alert, will we send any other emails?	1	1	https://issues.jboss.org/browse/HAWKULAR-343
5967	HAWKULAR-343	[~rutlucas], I had understood that we would NOT send an email when an alert is resolved, in alpha2. [~theute], can you clarify? I sent you an email that is a suggestion for the plain text. For the email experience, is it possible: - To have the sender defined as "Hawkular"? - To display the url and problem already on the subject? E.g: "Alert message: http://www.google.com is down" / "Alert message: http://www.google.com has response time above the threshold" Thanks	1	1	https://issues.jboss.org/browse/HAWKULAR-343
5968	HAWKULAR-343	Attached a prototype for success. [~theute], I opted to display the total time alert duration and when it was recovered. Think we can leave the initial time in this email if we have these 2 pieces of information. Also, the initial time is in the email that the user received previously. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-343
5969	HAWKULAR-343	[~gcardoso], +1 on the designs [~rutlucas] I am still unclear on how alerts/events are managed and why we would keep "intervals", it seems related to HAWKULAR-130 ? Is there a part of this JIRA that can be delivered in Alpha 2 ?	1	1	https://issues.jboss.org/browse/HAWKULAR-343
5970	HAWKULAR-328	Is this still relevant [~ppalaga] and [~lkrejci] ? The mentioned inventory jira seems in and the pinger seems still to work?	0	0	https://issues.jboss.org/browse/HAWKULAR-328
5971	HAWKULAR-321	[~alexandrem] Does that suit your needs? Is that reflected in the UI somewhere?	1	1	https://issues.jboss.org/browse/HAWKULAR-321

#	Requirement	Questions	Answered	Successfully Answered	URL
5972	HAWKULAR-288	When I discussed this topic the other day with Heiko, it got thinking about the driver's Session object. Should each Hawkular service, i.e., metrics, alerts, etc., be creating its own Session, or should it be shared? I think it should be one Session object per physical cluster. The fact that each of our services might be using their own keyspaces is not as important as the fact that we are sharing the same physical cluster. There was a related discussion on an internal mailing list about having something like "NoSQL" data sources. We need to investigate this.	1	0	https://issues.jboss.org/browse/HAWKULAR-288
5973	HAWKULAR-264	Which tests are targeted ? All ? There are tests where verbosity is useful, I agree that we can remove calls to System.out with a proper log4j or similar. I can review this for alerts component (as it was scheduled). But, is this creating some conflict (excess of verbosity on tests?) or it is to control where to store them (configure the appender) ? Thx	0	0	https://issues.jboss.org/browse/HAWKULAR-264
5974	HAWKULAR-259	Status: {code} <jpkroehling> stianst, morning! I'm building a PoC for the tokens filter/proxy that we discussed the other day, and I have a question... our frontend is a single-page app built on top of hawt.io, and uses keycloak.js for auth, and the backend is configured as a bearer-token client <jpkroehling> it seems that the refreshToken is not available for the backend (for bearer token clients) <jpkroehling> I can kinda of imagine a reason, but I think it's worth asking if it's indeed intended * msavy (~msavy@redhat/jboss/msavy) has joined #keycloak <stianst1> it's indeed intended and it wouldn't make sense to sense refresh token to the service <jpkroehling> stianst1, so, I should change the backend to not be a bearer token <stianst1> jpkroehling: the proxy should use direct grant api to obtain the token on behalf of the agent <jpkroehling> I think the refreshToken is not available for basic auth as well <stianst1> jpkroehling: it would give the agent a reference id, not the full token <jpkroehling> stianst1, for the direct grant, I'd need the user's password, wouldn't I ? <jpkroehling> let me check the docs again <stianst1> jpkroehling: basically the proxy manages tokens on behalf of the agents <stianst1> jpkroehling: ah, yes that's one piece of the puzzle I didn't consider :/ <jpkroehling> stianst1, I can work it around, as we either receive a request from the frontend (and we can send the refreshToken along with the create-token request), or we get a backend call with user:pass <jpkroehling> so, I have a way to get the refresh token, *but*, that feels bad :) <stianst1> jpkroehling: that's not correct either though <jpkroehling> no? <jpkroehling> ah, rihgt <stianst1> jpkroehling: you shouldn't share the refresh token - it should be created specifically for the agent <stianst1> jpkroehling: when you link an agent you should use direct grant api to create a refresh token for that "link" <jpkroehling> right, that's what you mean by "not correct" :) agreed <stianst1> so, only way to do it atm is to ask the user for username/pass again :(<stianst1> lemme think about this a bit ;) <jpkroehling> okay <stianst1> jpkroehling: I got it :) <jpkroehling> that was quick <stianst1> jpkroehling: you don't even need to use direct grant and it'll support sso <stianst1> jpkroehling: now, just need to explain it * lfryc is now known as lfryc|lunch <jpkroehling> stianst1, we do can a hangout, if it's easier <stianst1> jpkroehling: think so {code}	0	0	https://issues.jboss.org/browse/HAWKULAR-259
5975	HAWKULAR-259	Status: a local branch with some draft commits was made, but I'm currently "blocked" on a conceptual issue. Will ping stian again later for clarification. What is already done on the branch: - token-store-api with models - token-store with endpoint, that generates a secret/key and stores it along with the refresh token - undertow-filter , which gets the refresh token and tries to convert it into an access token To consume it, a module (like, "sample") can just add undertow-filter as a dependency. This filter detects that the incoming request has secret/key credentials and transparently injects an appropriate set of KC credentials into the incoming request, if the secret/key is valid. Right now, the token-store is a module on Accounts and we are also storing the persona information, but other than that, it could be split into a separate project and reused by other projects. Once the first iteration is done, we can consider splitting it. {quote} <jpkroehling> stianst, context: that proxy/token store for offline tokens (for hawkular agents) <jpkroehling> you might remember the discussion we had before I went for PTO, regarding storing the refresh token and getting a new access token based on it <jpkroehling> the issue I'm having right now is: when the user hits the token-store endpoint, the endpoint requests a refresh token from the KC adapter, and it comes with the token-store client on it <jpkroehling> on the undertow filter, I read the deployment's keycloak.json , and use it for replacing the refresh token by an access token <jpkroehling> this seems to not be appropriate, as I get an exception here: https://github.com/keycloak/keycloak/blob/88f14b39a71e679fe7ed1699d7b2acb8c1c4742a/services/src/main/java/org/keycloak/protocol/oidc/TokenManager.java#L136 <jpkroehling> so, which would be appropriate: to have the token-store to share the client credentials with the hawkular backend, or to make the undertow filter to use the token-store client? <jpkroehling> I think the second option would hit the same problem later on the request, but wanted to understand which would be the correct scenario in your view <jpkroehling> (btw, hawkular backend is a bearer only, so, perhaps it's indeed appropriate to have the token-store client on the undertow-filter){quote}	0	0	https://issues.jboss.org/browse/HAWKULAR-259
5976	HAWKULAR-245	Seems done and released https://github.com/hawkular/hawkular-agent/tree/0.1.0 ?	0	0	https://issues.jboss.org/browse/HAWKULAR-245
5977	HAWKULAR-232	the tests don't pass because WF9 changed the way attribute names can be named (from the way it was in 8.2). I didn't get a response to this in #wildfly, but I confirmed by changing code in your branch thomas and it passed (I changed the attribute names so they had a "-" instead of a ".") 03:23:11 PM) mazz: I need a sanity check please. I'm trying to run unit tests on my module subsystem (via my subclass of org.jboss.as.subsystem.test.AbstractSubsystemBaseTest). My test DID work on WF8.2 but FAIL when I move to WF9 CR2. (03:23:11 PM) mazz: The error is : (03:23:11 PM) mazz: failure description: "WFLYCTL0201: Unknown attribute 'org'" (03:23:11 PM) mazz: in my custom subsystem, I have several attributes of the form "org.abc.xyz" such as: (03:23:11 PM) mazz: <element org.hawkular.package.name.foo="test1" org.hawkular.another.package.blah-blah="test2" /> (03:23:44 PM) mazz: could it be possible that attribute names can no longer have "." in their names? (03:24:11 PM) mazz: I have no attributes called "org" but plenty that look like "org.some.name.here" (03:47:24 PM) mazz: uh-oh. Yeah, I thikn that's a problem (03:47:48 PM) mazz: in WildFly 9 CR2, it is no longer possible to have a dot (".") in attribute names. (03:47:50 PM) mazz: This makes me sad (03:48:06 PM) mazz: this was possible in WildFly 8.2.Final (03:48:36 PM) mazz: I just changed all my "." to "-" in my custom module subsystem extension (in the java code and the .xml) and now my tests pass (03:50:06 PM) mazz: I don't know if that is also true for element names. (03:50:25 PM) mazz: I have "org.abc.xyz" element names as well. but the error message I got just mentioned attribute name	1	1	https://issues.jboss.org/browse/HAWKULAR-232
5978	HAWKULAR-226	Shouldn't they go onto the details screen? IP could perhaps go in fine print to the right of the url , but all the other things (powered by can have more than 1 item) would be too much for here.	0	0	https://issues.jboss.org/browse/HAWKULAR-226
5979	HAWKULAR-226	Details screen yes that is a possibility as well. Will the users want to sort the urls list by these attributes? For instance, grouping by Powered by so all Apache servers are grouped together? [~uxdlc] [~gcardoso] WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-226

#	Requirement	Questions	Answered	Successfully Answered	URL
5980	HAWKULAR-226	Thanks for the specs. Just to make sure that I understand correctly: I should implement what is on URL-list_plus.jpg [12391301], disregarding all other mocks, right? [~uxdlc]: > For table column order, proposing: State, App name, Server IP, Type, Powered By, Host. I do not understand this statement - was it missplaced maybe? This is about list of URLs where App name and Type make no sense. In URL-list_plus.jpg, I can see only Powered by and Server IP. Could you please clarify that?	1	1	https://issues.jboss.org/browse/HAWKULAR-226
5981	HAWKULAR-222	[~jdoyle] I made it one chart just to group all of the Topics chart info in one view, but they can certainly be separated. Are you envisioning them stacked or side by side? Other than the Topic chart, or charts, I had figured that all of the other queue charts would stack on top of each other, running vertically down the page.	1	1	https://issues.jboss.org/browse/HAWKULAR-222
5982	HAWKULAR-222	[~mithomps] Mike, I think this is in pretty good shape. I marked it as "uxd done", do I need to reassign it to Karel or Gabriel?	0	0	https://issues.jboss.org/browse/HAWKULAR-222
5983	HAWKULAR-222	Wow, that was fast! P.S. I had posted some alternative visuals for the tiles: https://issues.jboss.org/secure/attachment/12395240/JMS3b.jpg Did you decide to go with the no icon version instead?	0	0	https://issues.jboss.org/browse/HAWKULAR-222
5984	HAWKULAR-222	Updated the prototype. Actually I thought with icons was a first design, and without was a newer design. I updated our clickable icons in the console to patternfly blue (as patternfly does). I kept the informative icons (non-clickable) in light gray. Considering the icons from this tile are not clickable, shouldn't they be light gray?	0	0	https://issues.jboss.org/browse/HAWKULAR-222
5985	HAWKULAR-221	Another mockup using a % based chart, stacked bar. [~mithomps] wdyt, does this work?	1	1	https://issues.jboss.org/browse/HAWKULAR-221
5986	HAWKULAR-220	[~uxdlc], is this the last version?	1	1	https://issues.jboss.org/browse/HAWKULAR-220
5987	HAWKULAR-220	Do enable and timestamp need to be in the same cell? Otherwise it is easier to have it in two cells.	1	1	https://issues.jboss.org/browse/HAWKULAR-220
5988	HAWKULAR-219	I attached a proposal with a more minimal half-donut graph. Think we don't need the 0% and 100%. Looks cleaner and still easily understandable to me. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-219
5989	HAWKULAR-219	Yes, % should be rounded off. Do you think it should even offer the decimal, or simply rounded to the whole: "5%"?	1	1	https://issues.jboss.org/browse/HAWKULAR-219
5990	HAWKULAR-217	RE: Badge in the tab. Would the tally in the badge change when you change the date range? And does changing a date range change it for all the tabs (and their badges)?	0	0	https://issues.jboss.org/browse/HAWKULAR-217
5991	HAWKULAR-217	Looks good! I saw that you added the date in the "JVM Status" and "JVM Memory Metrics" headings. I understand, our current date picker is not prominent at all. I worked on a more prominent date-picker (attached), I think we solve this problem with the change. WDYT?	0	0	https://issues.jboss.org/browse/HAWKULAR-217
5992	HAWKULAR-215	Does that also "unify" the schema creation and keyspaces? I can imagine that the schema definition data should stay in the respective projects, while the code that applies those definitions should be in the common code(?)	1	1	https://issues.jboss.org/browse/HAWKULAR-215
5993	HAWKULAR-203	Which certificates do you consider trustworthy for a pinger? Do you anticipate a configuration property for that? We are talking about Pinger. Its main goal is to find out whether some address can be pinged or not. IMHO it is not Pinger's responsibility to tell which certificate to trust. How would you imagine its output in a case when it was possible to ping the target but it had an untrusted certificate? If we were to add this to the Pinger, I would suggest changing its API to not only verify the ping ability but to also throw an exception when the certificate is untrusted (passing a "special" value in the HTTP response status code would be a dirty solution ;-)). Having said that, I believe that Pinger does its work properly and it is not up to its responsibility to judge the certificates.	0	0	https://issues.jboss.org/browse/HAWKULAR-203
5994	HAWKULAR-194	Draft 1 was published here: https://github.com/jpkrohling/hawkular-accounts/tree/HAWKULAR-194 [~jkremser], would you say that it's better to have one event object for each event, or would you say that having one event object, with different "action" properties is enough? This here would be the CDI event: https://github.com/jpkrohling/hawkular-accounts/blob/2ad015f0a0dab1a82b4e28fc532cfefc894d1938/events-listener/src/main/java/org/hawkular/accounts/events/listener/AccountsEvent.java And this is an example of a topic listener: https://github.com/jpkrohling/hawkular-accounts/blob/2ad015f0a0dab1a82b4e28fc532cfefc894d1938/events-listener/src/main/java/org/hawkular/accounts/events/listener/internal/AccountsMessageListener.java	0	0	https://issues.jboss.org/browse/HAWKULAR-194
5995	HAWKULAR-189	Question: The header on the initial (welcome) URL list screen should remain vertically centered within the page, is that still the case? It's the screen that you see after registering. I can't bring up the instance to verify.	0	0	https://issues.jboss.org/browse/HAWKULAR-189
5996	HAWKULAR-185	This is blocking UI automation in Hawkular MS1. Can we get this addressed in MS2??	0	0	https://issues.jboss.org/browse/HAWKULAR-185
5997	HAWKULAR-178	[~alexandrem], which approach do you think would be better?	1	1	https://issues.jboss.org/browse/HAWKULAR-178
5998	HAWKULAR-158	Is this still valid?	1	1	https://issues.jboss.org/browse/HAWKULAR-158
5999	HAWKULAR-156	Wasn't that already done some time ago?	0	0	https://issues.jboss.org/browse/HAWKULAR-156
6000	HAWKULAR-156	Is that still a thing?	0	0	https://issues.jboss.org/browse/HAWKULAR-156
6001	HAWKULAR-155	Where do we stand here?	1	0	https://issues.jboss.org/browse/HAWKULAR-155
6002	HAWKULAR-155	The issue that I believe is still outstanding is that the current Artemis bus is not secured. Anyone can put messages on the bus (AFAIK). This JIRA should address the question "should we put KeyCloak/Accounts in front of bus as an authentication mechanism forcing a bus client to log in so it can put messages on the bus and read messages off the bus" I haven't looked into how Artemis supports authentication plugins, or how the WildFly integration of Artemis can help here.	0	0	https://issues.jboss.org/browse/HAWKULAR-155
6003	HAWKULAR-152	Is this intended for the current MVP run?	1	1	https://issues.jboss.org/browse/HAWKULAR-152
6004	HAWKULAR-152	[~nstefan], I've been looking into the failing tests issue. The problem may be in the CORS support on the server side. i.e. when trying to create a gauge, the browser makes a ["preflight" OPTIONS request|http://stackoverflow.com/questions/12111936/angularjs-performs-an-options-http-request-for-a-cross-origin-resource/12112294#12112294] which doesn't seem to handled right and thus resulting into a CORS issue. My guess is, the problem may in the response of the OPTIONS request, which is not a json object: bq. HEAD, POST, GET, OPTIONS What do you think?	1	1	https://issues.jboss.org/browse/HAWKULAR-152
6005	HAWKULAR-152	I was able to fix the issue on the server side. After adding {{"hawkular-tenant"}} to CORS supported headers list, we got only two failing tests in the {{hawkular-ui-services}} rest metrics test-suite but no more strange CORS issues. Now, I need to revisit the test-suite itself, fix failing stuff and find a better way to set the {{"hawkular-tenant"}} header, so it doesn't effect all {{$http}} requests made from AngularJS. The PR with the CORS fix is here: https://github.com/hawkular/hawkular-metrics/pull/248 [~nstefan], is it possible to pass the hawkular-tenant as an argument through URL parameters?	1	1	https://issues.jboss.org/browse/HAWKULAR-152
6006	HAWKULAR-152	Is this already in Hawkular-master?	1	1	https://issues.jboss.org/browse/HAWKULAR-152

#	Requirement	Questions	Answered	Successfully Answered	URL
6007	HAWKULAR-132	Question is: should OrganizationService automatically remove the sub-organizations, or should the caller be responsible for that? I think the caller should make sure that there are no sub organizations before calling the "delete" method, so, the "delete" method should just remove the sub organizations automatically, in the assumption that the caller really wants to delete this organization, no matter what.	0	0	https://issues.jboss.org/browse/HAWKULAR-132
6008	HAWKULAR-130	You are talking implementation, I am talking user experience :) I don't understand your a) b) since it's based on the implementation. Would be clearer if we stick to user experience, I am not sure how to properly explain again but let me try (and ignore some "corner case" issues for now) Hypothesis: - response time is collected every minute - I have a threshold set for 1s with alert set if the threshold is above that value for 5 min that sends me an email - from t=0 to t=2h the response time is always measured as 2s, from t=2h to t=4h it is measured at 0.5s From a *user* point of view: - I should receive an email at t=5min - If I go to the admin console at t=2min there should be no notification there (only the graph could tell me that I am above the threshold) - If I go to the admin console at t=1h I should have a notification in the UI that a problem is currently happening at started at t=0 is still open (and may require immediate action) - if I go to the admin console at t=3h I should have a notification in the UI that a problem happened between t=0 and t=2h but resolved by itself. Do we agree on that ? If so, the "corner cases" such as response time being below the threshold for a short period of time should be handled without affecting this basic scenario. As of now, if the previous scenario happens, a user gets an alert entry every 40s in my environment. This is likely again related to the notion of alerts and events, I think those are separate concerns that we need to address. We will also need to add escalation capabilities (if the issue still exist at t=1h and noone "acked" we could send an SMS to the manager)	0	0	https://issues.jboss.org/browse/HAWKULAR-130
6009	HAWKULAR-129	[~uxdlc], any ideas how to progress? Do we have a mock-up or any other information which can help me to resolve this JIRA? Thanks.	1	1	https://issues.jboss.org/browse/HAWKULAR-129
6010	HAWKULAR-129	Hi [~uxdlc], is this JIRA still relevant?	1	1	https://issues.jboss.org/browse/HAWKULAR-129
6011	HAWKULAR-128	Those are there by default, from Keycloak. Should I keep only the fields from the previous version of Keycloak (username, email, first name, last name, password, password confirmation), or do you have a set of fields that you'd want to keep ?	0	0	https://issues.jboss.org/browse/HAWKULAR-128
6012	HAWKULAR-127	[~theute] I don't see any resources for this in Alpha6. Push to Alpha7? CC: [~pilhuhn]	1	1	https://issues.jboss.org/browse/HAWKULAR-127
6013	HAWKULAR-127	[~uxdlc], are you aware if patternfly updated the search/filter behavior? Serena was working on something...	1	1	https://issues.jboss.org/browse/HAWKULAR-127
6014	HAWKULAR-127	[~gcardoso] it looks good, but it bothers me that whenever you switch from text search to select list (in the example from name to birth month) it flickers. Won't this be a problem?	1	1	https://issues.jboss.org/browse/HAWKULAR-127
6015	HAWKULAR-127	I have created new directive for searching over tables, this directive can be used with any table. Filter which matches 0 fields: !Selection_001.png|thumbnail! Filter which matches some fields: !Selection_003.png|thumbnail! What do you think, should I create also the basic/advanced mode? Where you would type something like 'state = up and type=wildfly and byText=local' which would create 3 filters. Switching between basic and advanced would not change current filters. We should agree on some kind of search form. I guess only some simple search values would be possible (e.g. no !=, OR, LIKE, IS, IN ... operands). If I am overthinking this just let me know, I will implement just some simple search in advanced mode. Also I would need [~gcardoso] new .less rule, as you can see there is top border with the active filters, which is a bit stretched and the border is bigger than the table itself. I was unable to break the filters apart, so you would have to select what filter you want to use (State, Type) and then either type or select value. I have used the http://angular-patternfly.rhcloud.com/#/api/patternfly.filters.directive:pfSimpleFilter as Gabriel suggested but there is slight bug when switching between text and select filter. I have looked into the source codes and I think I know why is this happening, but it looks like it is desired behavior.	1	1	https://issues.jboss.org/browse/HAWKULAR-127
6016	HAWKULAR-127	Very nice Karel! How does this "By text" works? Does it search in all columns and filter? If not, this is something desirable, for the user to be able to type something and it filters in all columns (like it is the current search input). Regarding the advanced search, I'll extend this question to [~mithomps], but I would start with a simple and add it later if we feel the need. Once you merge this on master I can fix this line issue :) Completely fine that you didn't break the filters apart, since this is the patternfly approach (my design was prior to them).	1	1	https://issues.jboss.org/browse/HAWKULAR-127
6017	HAWKULAR-127	That's great! My question is specifically regarding the label "By text", did u created this text or it came from patternfly? Cause to me it is a bit confusing, it is not clear that it filters in all fields, u know what I mean? Yes, inform them. Thanks	1	1	https://issues.jboss.org/browse/HAWKULAR-127
6018	HAWKULAR-125	As a separate thing, can you look into sizing the table pagination to the screen height. So if there is more height we can see more rows. Is this correct @Lclayton@redhat.com?	1	1	https://issues.jboss.org/browse/HAWKULAR-125
6019	HAWKULAR-111	Liz, Jay has some good ideas here that the durations are confusing for the simple view and should probably be an advanced option. WDYT?	0	0	https://issues.jboss.org/browse/HAWKULAR-111
6020	HAWKULAR-108	[~ppalaga] Our artifacts seem to show up on Central right now - so is this still needed? I understand that signing has it own benefits, but is it required?	1	0	https://issues.jboss.org/browse/HAWKULAR-108
6021	HAWKULAR-97	What exact steps are needed to see the header before the redirection? Everything looks fine when I try to test it (with latest Chrome, FF).	0	0	https://issues.jboss.org/browse/HAWKULAR-97
6022	HAWKULAR-91	The messaging subsystem has been deprecated... not sure if we should spend time on this: http://wildscribe.github.io/Wildfly/9.0.0.Beta1/subsystem/messaging/index.html "Deprecated Since 1.4.0 The messaging subsystem is deprecated and may be removed or limited to managed domain legacy server use in future versions."	0	0	https://issues.jboss.org/browse/HAWKULAR-91
6023	HAWKULAR-85	My one concern is performance. Firing all those queries but it is async so maybe not an issue. WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-85
6024	HAWKULAR-77	How does that work with say 100 urls already in place? Browsers usually allow scrolling though the list. Wouldn't the "Add .." be sort of hidden in that case?	0	0	https://issues.jboss.org/browse/HAWKULAR-77
6025	HAWKULAR-71	Would be interesting to investigate if it's possible to have two .gitignore files . A default one, marking the files in dist/ as ignored, and a "release" one, that accepts the files in dist/. This way, a developer won't commit changes to the dist/ files during the usual development cycle, while the release process still generate those files, which are required in order to be consumed by other modules as bower dependency.	0	0	https://issues.jboss.org/browse/HAWKULAR-71
6026	HAWKULAR-71	I agree this is a problem (and should be solved). How about instead of .gitignore the dist folder just *not* have the connect-* task update the dist folder. Instead, you need to publish, to update the dist folder. So do a 'gulp publish' which I guess would just be a 'gulp build' (and maybe its a default task).	1	0	https://issues.jboss.org/browse/HAWKULAR-71
6027	HAWKULAR-70	[~vrockai] Can you update with git pull # and resolve if this is done. Thanks.	1	1	https://issues.jboss.org/browse/HAWKULAR-70
6028	HAWKULAR-68	[~gcardoso], I just noticed that there's no update to the "Account Management" screen. Is it intended? You can access this by booting the hawkular-accounts distribution and opening the following URL: http://localhost:8080/auth/realms/hawkular/account	1	1	https://issues.jboss.org/browse/HAWKULAR-68
6029	HAWKULAR-68	[~juraci.costa], thanks for the reminder. Just to know, is this for June 3rd or future release?	1	1	https://issues.jboss.org/browse/HAWKULAR-68
6030	HAWKULAR-68	[~theute], do you have any idea about when we'll expose accounts?	1	1	https://issues.jboss.org/browse/HAWKULAR-68

#	Requirement	Questions	Answered	Successfully Answered	URL
6031	HAWKULAR-67	Hi Juca! What is an organization? Is this somehow related to Keycloak?	1	1	https://issues.jboss.org/browse/HAWKULAR-67
6032	HAWKULAR-65	{quote} <jpkroehling> stianst, morning! on the keycloak-dev list, you said that "If you're requiring programmatic setup directly to Keycloak you're doing it wrong IMO." <stianst> jpkroehling: yup <jpkroehling> stianst, I wonder if the context applies to an older thread on keycloak-users, where Marek mentioned this approach as a possibility to change the master admin password <jpkroehling> http://lists.jboss.org/pipermail/keycloak-user/2015-March/001770.html <stianst> jpkroehling: nah, for that you should only do it if KC is embedded in your project <jpkroehling> stianst, it's not "embedded in the project", but it's embedded in the distribution we are currently building, ie, we provide the possibility of using an external KC <stianst> jpkroehling: what I mentioned to Stan was that he's working on a simple way to get started with the WF adapter, in which case it should work for local and remote servers, so should use rest api's <jpkroehling> as far as Hawkular is concerned, KC can be anywhere, but we'd want to either disable the master's admin user or change its password on a first boot <stianst> jpkroehling: omayto, tomahto ;) <stianst> jpkroehling: tomayto even <jpkroehling> O_o <stianst> jpkroehling: when you embed it you want to disable master admin and should do so by "changing" the embedded KC war <stianst> jpkroehling: for embedding the idea is that you build your own war and "modify" it to fit your needs <jpkroehling> but that's the point: we are not embedding it as part of our project, we are embedding it as an easy way to quick start using hawkular <jpkroehling> for a production usage, KC will probably be a remote instance <stianst> jpkroehling: irc confusion setting in here ;) <stianst> jpkroehling: hawkular will have a dist/download that contains both hawkular and kc right? for devs and quick starts and such <jpkroehling> what I'm currently toying with is to ship with master-realm.json + master-users-0.json a,d import it on first boot, just to disable the admin's user <jpkroehling> right <stianst> jpkroehling: right, that should be built like: https://github.com/aerogear/aerogear-unifiedpush-server/tree/master/servers/auth-server <stianst> jpkroehling: and you can add whatever initial config you want like this: https://github.com/aerogear/aerogear-unifiedpush-server/blob/master/servers/auth-server/src/main/java/org/jboss/aerogear/unifiedpush/keycloak/UpsKeycloakApplication.java <jpkroehling> does that work also for remote KC instances ? <stianst> jpkroehling: no, that's for building the embedded KC <stianst> jpkroehling: connecting to a remote KC instance is a separate issue <stianst> jpkroehling: for that you obviously don't disable master admin user, but instead you want to import your app and such into an existing realm <stianst> jpkroehling: I'm looking at creating a rest endpoint that lets you "link" hawkular (or any app) with a keycloak realm <stianst> jpkroehling: thing is you want to add an app to an existing realm, not create a new realm - as otherwise you don't get sso between consoles <jpkroehling> stianst, good point /me jpkroehling thinking about how that would affect multi tenancy <jpkroehling> but yes, you are right <stianst> jpkroehling: you'll need to explain me more about how multi tenancy works in hawkular <stianst> jpkroehling: but I imagine you'd have to "link" each tenant to a kc realm <jpkroehling> hawkular is not using multi tenancy features from KC, I'm thinking more about multi tenancy in KC itself :) <jpkroehling> but that's a problem for later, I guess <stianst> jpkroehling: now I'm really confused <jpkroehling> stianst, forget about the multi tenancy part :) it just just me thinking out loud about multi tenancy in KC, not related to Hawkular at all <stianst> jpkroehling: I reckon by default the mechanism to link hawkular to KC would take the hostname of KC, it would then link with master realm <stianst> jpkroehling: but, you could also link to a specific realm <stianst> jpkroehling: for example you may have one realm with KC admin, 2 x wildfly servers and a hawkular server <stianst> jpkroehling: then another realm with 2 x wildfly servers and a hawkular server <stianst> jpkroehling: in most cases though I'd say all servers would be managed through the same realm <jpkroehling> make sense <jpkroehling> stianst, is that automatic linking planned for 1.x? <stianst> jpkroehling: I hope to have some mechanism to make it easy to link a jboss project	0	0	https://issues.jboss.org/browse/HAWKULAR-65
6033	HAWKULAR-53	I verified this works. Shows proper version and git sha for hawkular-ui-components. How difficult would this be to add for hawkular repo git sha as well? So we can make sure the kettle and ui-components are matching up properly.	0	0	https://issues.jboss.org/browse/HAWKULAR-53
6034	HAWKULAR-47	[~alexandrem], can you please test this branch, if it fulfills your requirements? https://github.com/vrockai/hawkular-ui-components/tree/HAWKULAR-47 If you're happy with it, I'll make a PR. Thanks.	1	1	https://issues.jboss.org/browse/HAWKULAR-47
6035	HAWKULAR-47	[~mithomps], can you please test this? If it works for you, you can cherry-pick this commit and I'll set the JIRA status to resolved. https://github.com/vrockai/hawkular-ui-components/commit/ce732fdd43b58d719556e7de82f4041f743b9863	1	1	https://issues.jboss.org/browse/HAWKULAR-47
6036	HAWKULAR-27	Sure. It is already implemented in the HTML prototype: http://cardosogabriel.github.io/hawkular-client/application.html I sent an email with updates on Friday, did you receive it?	1	1	https://issues.jboss.org/browse/HAWKULAR-27
6037	HAWKULAR-24	This is my concept: https://github.com/vrockai/hawkular-ui-services/commit/1bb925881c32aec83d49eb4696e72915f19fc289 I'll add it other gulp projects. Should I add it to the slush generator, too?	1	1	https://issues.jboss.org/browse/HAWKULAR-24
6038	HAWKULAR-24	Can the file begin with the reference or does it need to have copyright first? For example: /// <reference path="metricsPlugin.ts"/> Can that be at the top or need to be after header?	1	1	https://issues.jboss.org/browse/HAWKULAR-24
6039	HAWKULAR-24	I've never seen files with something before the license itself. Is there a problem with that in TS?	1	1	https://issues.jboss.org/browse/HAWKULAR-24
6040	HAWKULAR-16	Is that still a thing?	1	1	https://issues.jboss.org/browse/HAWKULAR-16
6041	HAWKULAR-13	I was able to run the rhq-metrics server, so I'm able to look at the REST API by myself. I've got several issues. Some of the examples from the doc are not working for me. Like using the {{"numeric.5min": 100m}} when creating tenant, or accessing the metrics itself like "http://localhost:8080/rhq-metrics/{tenantId}/metrics/availability/{id}/meta" (what's the ID of metrics? in the examples, it has just name, and it won't allow me to create metrics with an "id" field). Is there something from Hawkular REST available yet? Can I try that interface by myself? By testing, do you mean integration tests? For unit tests, I think it's usual to use some mock back-end (because the situation above would be actually testing the server-side, not the client side in my opinion).	1	1	https://issues.jboss.org/browse/HAWKULAR-13
6042	HAWKULAR-13	Cool, so to conclude, I will: 1. Create ngResource services for all known endpoints (described in the readme file above). 2. Create a test-suite, which will actually check if those services work against a live server. I can create a separate gulp task for this. We can run it on our own, once we start the REST server, but it won't be part of the build. For that I need to know: Do we have a preferred project structure (where to put services)? Do we have preferred testing framework?	1	1	https://issues.jboss.org/browse/HAWKULAR-13
6043	HAWKULAR-13	I'm thinking that there will be many more services to come so we should probably have a 'services' directory at the root that services like this reside in. services/hawkular-rest-service WDYT?	1	1	https://issues.jboss.org/browse/HAWKULAR-13
6044	HAWKULAR-4	This does not yet work as expected / described snert:metrics hrupp$ npm install -g bower gulp typescript /usr/local/bin/tsc -> /usr/local/lib/node_modules/typescript/bin/tsc > v8flags@1.0.8 install /usr/local/lib/node_modules/gulp/node_modules/v8flags [...] snert:metrics hrupp$ gulp [13:05:03] Local gulp not found in /h/hawkular/ui/console/metrics [13:05:03] Try running: npm install gulp snert:metrics hrupp$ npm install gulp [..] snert:metrics hrupp$ gulp Error: Cannot find module 'wiredep' at Function.Module._resolveFilename (module.js:338:15) at Function.Module._load (module.js:280:25) at Module.require (module.js:364:17) at require (module.js:380:17) at Object.<anonymous> (/h/hawkular/ui/console/metrics/gulpfile.js:2:15) at Module._compile (module.js:456:26) at Object.Module._extensions..js (module.js:474:10) at Module.load (module.js:356:32) at Function.Module._load (module.js:312:12) at Module.require (module.js:364:17)	0	0	https://issues.jboss.org/browse/HAWKULAR-4

#	Requirement	Questions	Answered	Successfully Answered	URL
6045	HWKAGENT-180	The underlying error is: org.hawkular.dmr.api.DmrApiException: java.io.IOException: java.net.ConnectException: WFLYPRT0053: Could not connect to http-remoting://127.0.0.1:9990. The connection failed whose cause is: org.xnio.http.RedirectException: XNIO000816: Redirect encountered establishing connection We build the client via: new org.jboss.as.controller.client.ModelControllerClientConfiguration.Builder().setProtocol("http-remoting")....build() which gets the config which is passed to: org.jboss.as.controller.client.ModelControllerClient.Factory.create(theconfig)	0	0	https://issues.jboss.org/browse/HWKAGENT-180
6046	HWKAGENT-165	I was JUST thinking about address this today! And then you created this JIRA. Great minds think alike! :) I was thinking more along the lines of waiting for the remote server to be available (so for every local-dmr, remote-dmr, local-jmx, remote-jmx defined, try to connect to the server and if it can't just wait). But then your idea of waiting for an MBean made me realize that for things like local-jmx the MBeanServer is always going to be there so that won't help. [~tcunning] how do you propose we do this for JMX? Do we define a specific MBean(s) for the agent to wait for in remote-jmx or local-jmx entry? Something like: {code} managed-servers: local-jmx: name: Local JMX enabled: true wait-for: - my.domain:type=MyMBean - my.other.domain:type=MyOtherMBean {code} Where the "wait-for" defines a list of MBeans that must be registered before that managed server begins discovery and collecting data?	1	1	https://issues.jboss.org/browse/HWKAGENT-165
6047	HWKAGENT-160	I will submit this PR shortly - but I wanted to document what you will now see if the credentials are incorrect: {noformat} 12:40:33,182 WARN [org.hawkular.agent.monitor.service.MonitorService] (Hawkular WildFly Agent Startup Thread) HAWKMONITOR010081: Cannot get Hawkular Server status - does the agent have proper credentials? (401/Unauthorized) {noformat}	0	0	https://issues.jboss.org/browse/HWKAGENT-160
6048	HWKAGENT-151	What JMX support specifically are you looking for? If this was discussed on a mailing list can you point me to the thread?	0	0	https://issues.jboss.org/browse/HWKAGENT-151
6049	HWKAGENT-151	Do we need Jolokia support or JMX support ? And do you 'ship' Jolokia with Fuse on EAP ?	1	1	https://issues.jboss.org/browse/HWKAGENT-151
6050	HWKAGENT-140	"rate" metrics are not persisted, they are derived. Nothing for the agent to do. Closing this issue. From: http://www.hawkular.org/hawkular-metrics/docs/user-guide/#_metric_types "Rate data points are not persisted. They are computed at query time."	0	0	https://issues.jboss.org/browse/HWKAGENT-140
6051	HWKAGENT-85	[~jayshaughnessy] and [~bobmcwhirter] are currently very close to getting a swarm build for the agent that only requires the agent subsystem snippet .xml rather than a full standalone.xml. It also is a much smaller bundle than a full blown WildFly (should be just a simple uber-jar that runs with the "java -jar" command with some configuration file passed in perhaps on the command line). I don't know all the details yet, once we get something working, we'll post here with some details.	1	1	https://issues.jboss.org/browse/HWKAGENT-85
6052	HWKAGENT-80	[~mwringe] can you give more details on what these bearer tokens look like and where we get them? e.g. are they just simple strings that someone defines in the standalone.xml in the agent's <remote-jmx> or <remote-dmr> elements (like we do today with username/password)? Or do we need to involve SSL/security realms in Wildfly? Right now the DMR and JMX endpoints fully support SSL as long as someone configures the agent wildfly instance with the proper security realm definitions (basically just a reference to the truststore). They also support basic auth (at least they should, you can provide a username/password in the configs). I don't know what other functionality we need. The Prometheus endpoints do NOT support anything but http right now - not even basic auth. We'd just have to add some code to the connection stuff.	1	1	https://issues.jboss.org/browse/HWKAGENT-80
6053	HWKAGENT-79	[~mwringe] - can you give more details about what you are looking for? Before I figure out how to label the metrics, I have to figure out how the user is going to tell me what labels to attach to what metric. My first question would be - are these labels dynamic or can they be statically defined in the standalone.xml settings? I have a feeling these labels are dynamic - if so, what are they dependent upon and what do they look like and how do I determine what they are? Are these labels defined on a per-metric basis or per-resource basis? In other words, will ALL metrics from a specific resource have the SAME labels? Or will each individual metric have different labels, even if those metrics come from the same resource? Are there going to be more than one label per metric? Or is there only a single label per metric? Any other details let me know. I don't know what I don't know - so I may not be asking all the right questions.	1	1	https://issues.jboss.org/browse/HWKAGENT-79
6054	HWKAGENT-79	When I talk about 'labels' here what I really mean is 'tags' in the Hawkular Metrics sense. Not using the right term may have been confusing things here. Eg see http://www.hawkular.org/docs/rest/rest-metrics.html#PUT__gauges__id__tags In essence you can add a string key and value pair to a metric. This allows for specifying metadata about the metric and allows for querying based on this data. bq. My first question would be - are these labels dynamic or can they be statically defined in the standalone.xml settings? I have a feeling these labels are dynamic - if so, what are they dependent upon and what do they look like and how do I determine what they are? The idea here is that for a normal situation, they would specify these in the standalone.xml file. This would allow an admin to specify some information about the metrics they are gathering. On the OpenShift integration side, nothing really comes from the standalone.xml (except pointing to where the OpenShift server is). In essence this piece would interface with the java objects to programtically specify what would normally be done in the standalone.xml file. I am currently hacking around in the code right now to pull out the tag information from a running OpenShift cluster. bq. Are these labels defined on a per-metric basis or per-resource basis? In other words, will ALL metrics from a specific resource have the SAME labels? Or will each individual metric have different labels, even if those metrics come from the same resource? Labels would be applied on a per-metric basis. Each individual metric would have different labels. bq. Are there going to be more than one label per metric? Or is there only a single label per metric? More than one per metric. See the tags support in Hawkular Metrics	1	1	https://issues.jboss.org/browse/HWKAGENT-79
6055	HWKAGENT-79	[~mwringe] Can the tags be the same across all metric instances of a specific metric type? Or must the tag names be unique across even metric instances? For example, consider a EAP/WildFly case. There is a resource type called "WildFly Server" that represents WildFly app servers. This "WildFly Server" resource type has an associated metric type called "Heap Memory Used". If I am monitoring three WildFly Server instances (thus they are all of the type "WildFly Server"), each of those instances will have one metric instance of metric type "Heap Memory Used" (call the metric instances "Heap Memory Used #1" for WildFly Server #1, "Heap Memory Used #2" for WildFly Server #2, and "Heap Memory Used #3" for WildFly Server #3). Thus there is one metric type but because I have three resource instances being managed, I have three metric instances. Can the tags of those three metric instances be the same? Or must they somehow be unique among all three? If I want to tag my Heap Memory Usage metric with "foo=bar", can "Heap Memory Used #1", "Heap Memory Used #2" and "Heap Memory Used #3" all have the tag "foo=bar" or must they be something like "foo1=bar1" for "Heap Memory Used #1", and "foo2=bar2" for "Heap Memory Used #2", etc.??	1	1	https://issues.jboss.org/browse/HWKAGENT-79
6056	HWKAGENT-78	Question: why don't we just send this 'fancy id' to inventory when registering the metric? Clients should not assume the M~foo%bar[baz]~ format, but query inventory	1	1	https://issues.jboss.org/browse/HWKAGENT-78

#	Requirement	Questions	Answered	Successfully Answered	URL
6057	HWKAGENT-71	According to Brian, we can now run the agent subsystem in a host controller or domain controller. We need to add a call to setHostCapable() in our extension initialization. From HipChat: John Mazz 2:08 PM @BrianStansberry hey - I know I've asked you this several months ago (and several months before that :)) but I'm being asked about it again to see if its changed, so doing my due diligence, I'm asking for confirmation. The Domain Controller and Host Controller - they are NOT extensible by injecting my own custom subsystem in it (like I could do in a standard WildFly installation). Correct? I have a subsystem that I can configure in WildFly by adding my own <subsystem> element to standalone.xml and putting the module code under modules/ - I assume Domain and Host Controllers still cannot support such an extension. Brian Stansberry 2:08 PM @JohnMazzGuest you can add extensions to an HC; our standard host.xml has one for jmx John Mazz 2:09 PM really???? this is new. Brian Stansberry 2:09 PM since WF 9, maybe 10 John Mazz 2:09 PM ok. glad I asked. this is NOT true for domain controler, correct? or is it? Brian Stansberry 2:09 PM there's no difference a DC is just a HC with some config settings John Mazz 2:10 PM ok, so both domain and host controllers can take a standard WildFly subsystem extensions. I think you just made my life easier :) Brian Stansberry 2:11 PM there are some tweaks necessary; @KabirKhan are there docs? tweaks meaning we need to be sure the extension was written for this use case but I don't think they are a big deal Brian Stansberry pokes to refresh his memory John Mazz 2:12 PM ok, any docs you guys can provide would be appreciated. Brian Stansberry 2:13 PM @JohnMazzGuest I think this is the key bit: https://github.com/wildfly/wildfly-core/blob/master/jmx/src/main/java/org/jboss/as/jmx/JMXExtension.java#L101 javadoc of that method: /** * Mark a subsystem as host capable. This will only take effect if running on a HC, and be ignored on a standalone or managed server. * The standard behaviour on the HC is to register the subsystem resource definitions in the domain model under the {@code /profile=*} address. * If this method is called, the standard behaviour happens, but in addition the resource definitions are registered in the local host model * so that the subsystem can be added there as well. */ void setHostCapable(); John Mazz 2:14 PM nice. actually... awesome :D Brian Stansberry 2:15 PM if you look at our standard host.xml files, they include the jmx subsystem so that's kind of an example Kabir Khan 2:21 PM That is a good start. Sorry am afk There is also a test case https://github.com/wildfly/wildfly-core/blob/master/testsuite/domain/src/test/java/org/jboss/as/test/integration/domain/suites/HcExtensionAndSubsystemManagementTestCase.java	1	1	https://issues.jboss.org/browse/HWKAGENT-71
6058	HWKAGENT-62	HipChat between me/Peter and Brian: === Peter Palaga 9:37 AM @BrianStansberry we'd need some advice where in the WF installation tree our Hawkular Monitoring Agent should be placed John Mazz 9:37 AM ^^ refering to https://issues.jboss.org/browse/HWKAGENT-62 Peter Palaga 9:37 AM @BrianStansberry are you the best addressee for this kind of questions? Brian Stansberry 9:38 AM @PeterPalaga yes it should not be in the base layer the base layer is owned by the WildFly or EAP distribution Peter Palaga 9:39 AM @BrianStansberry OK @BrianStansberry the Agent should go to base layer even though it will not be maintained/patched by EAP/WF team? John Mazz 9:39 AM the question we then had is - is it a layer or add-on. The current thinking is "add-on" Brian Stansberry 9:40 AM @PeterPalaga it should *not* go in the base layer because it's not part of the base dist Peter Palaga 9:41 AM @BrianStansberry ah, overseen the "not". Sorry Brian Stansberry 9:41 AM np :) John Mazz 9:41 AM @BrianStansberry the agent can be installed in any WF instance (normal out-of-box WF distro, EAP product, BRMS, blah blah) - so I assume the agent shouldn't be in its own layer but rather an add-on. correct? Brian Stansberry 9:41 AM @JohnMazz the key difference between a layer and an addon is the modules in a layer take precedence over those in the base, and can override John Mazz 9:42 AM @BrianStansberry and yo have to identify it in layers configuration file right? Brian Stansberry 9:42 AM yes John Mazz 9:42 AM add-ons do not need to be identified anywhere? correct? or did I miss the add-ons config file? :) Brian Stansberry 9:42 AM right John Mazz 9:42 AM ok. makes sense. then that JIRA is correct then, that it should go in add-ons Brian Stansberry 9:43 AM layers.conf isn't so much about identifying, it's about declaring the precedence order since layers can take precedence add-ons can't so there's no add-ons.conf	1	1	https://issues.jboss.org/browse/HWKAGENT-62
6059	HWKAGENT-18	Is that still a thing?	0	0	https://issues.jboss.org/browse/HWKAGENT-18
6060	HWKAGENT-18	I can remember that i-testing reload was tricky because during the reload, the container closes the management connection which AFAIK the wf-maven-plugin was not able to handle properly (it somehow force-failed the whole build on reload). Are we still using the wf-maven-plugin to start the container? I can remember we were migrating to pure Arquillian in some components, but I am not sure if that was the case here in Agent too. I any case, the reload functionality should be properly i-tested and one should look again how wf-plugin or Arquillian can help.	0	0	https://issues.jboss.org/browse/HWKAGENT-18
6061	HWKAGENT-6	Can you explain this further? I don't know what the description means.	1	1	https://issues.jboss.org/browse/HWKAGENT-6
6062	HWKAGENT-6	Is it understandable?	1	1	https://issues.jboss.org/browse/HWKAGENT-6
6063	HWKAGENT-4	So what needs to be done on the invnetory side to enable HTTPS? Does it work in metrics, so that I can try to copy some config over? I believe this should be a stadalone.xml/web.xml config thing, right?	1	1	https://issues.jboss.org/browse/HWKAGENT-4
6064	HWKALERTS-265	[~rutlucas], what do you think? Assuming my thinking here is correct, I think this could be pretty important for a truly federated trigger consisting of event-style conditions (like logging). There are a couple of ways we could go about it, and I don't think it would be too bad to do. We could have a single setting on the trigger, or I suppose we could have condition-level granularity where condition could override a global trigger setting. We'd basically throw out condition evals that were too old, likely in a lazy way when checking for a firing.	0	0	https://issues.jboss.org/browse/HWKALERTS-265
6065	HWKALERTS-240	This seems like a good approach. Just to make sure I understood correctly, I think there was a typo in your comment. Did you mean: node/node1/disc.network.in.bytes[eth0] {name: disc.network.in.bytes, node: node1, adapter: eth0} node/node1/disc.network.in.bytes[eth1] {name: disc.network.in.bytes, node: node1, adapter: *eth1*} node/node2/disc.network.in.bytes[eth0] {name: disc.network.in.bytes, node: *node2*, adapter: eth0} node/node2/disc.network.in.bytes[eth2] {name: disc.network.in.bytes, node: *node2*, adapter: *eth1*} We Are not tagging the dynamic metrics at the moment but we can, and should (Implementation is still under discussion) So if we tag all the metrics and write the tag query as described, we would get a different alert for each of the adapters? and have a way to identify which node/adapter triggered the alert? Another question, can we set auto-resolve query for such alerts? Will it fire on every metric that satisfies the condition, or only once until resolved, etc?	1	1	https://issues.jboss.org/browse/HWKALERTS-240

#	Requirement	Questions	Answered	Successfully Answered	URL
6066	HWKALERTS-240	[~zgalor], yes, you are totally correct, I'm sorry for the confusion in the example tagging. Each metric would be tagged with name-value pairs that reflect the metric name. It is probably clear but to be explicit, the tagging would be applied on the metric definitions, not each data point. Tagging on each data point is possible, but it adds a lot of overhead, it is not meant for repeated tagging of the same information. By verbosely tagging the metric definitions it adds a lot if query flexibility, not just for this feature but in general. > So if we tag all the metrics and write the tag query as described, we would get a different alert for each of the adapters? I should mention that everything is still constrained by tenant. So, the query would be for all of the adapters under a tenant. We could further narrow the set of adapters by adding more criteria, as allowed by hMetrics' tag query language. By using the 'evalType: EACH' clause we are telling the alerter to test each adapter individually, so each offender generates an alert. If we used 'evalType: ALL' the alerter would combine all of the metric data and perform the test one time against the aggregate. > ...and have a way to identify which node/adapter triggered the alert? In this approach we have only the single trigger defined in alerting, and it can generate alerts for any adapter. The offending adapter can be found in the context data of the event generating the alert, The context data is basically a map with information lending context to an alert, or the data that fired the alert. In this case you'd have a context map entry like, {{context=\{"MetricName":"node/node1/disc.network.in.bytes\[eth0\]"\}.}} > can we set auto-resolve query for such alerts? Will it fire on every metric that satisfies the condition, or only once until resolved, etc? Right, this is a critical point in the design difference between the alerter approach and the group trigger approach. The alerter approach uses a single trigger that defers evaluation work to the external alerter. The group trigger approach creates a member trigger for each metric. The alerter approach is nice because it has only one trigger to manage regardless of whether metrics come or go. And that trigger does not require any forwarding of data from metrics to the alerting engine. Instead, we query the persisted data on a schedule. The group trigger approach introduces the overhead of synchronizing our member trigger population with the potentially large number of metrics we want to alert on. But it does have the advantage of more granular control. Using our example, you can use auto-resolve with the alerter-based trigger; but with the realization that if it fires for *any* adapter it would go into autoresolve mode and will not fire alerts until the autoresolve conditions are met (in this case probably a MissingCondition ensuring no more offending events are being received). If you don't use auto-resolve or auto-disable, you could get a repeating alert for, say, the same bad adapter. If you do use auto-resolve or auto-disable, you could potentially miss alerts for different adapters. So, the alerter approach does have this potential drawback. Using the group trigger approach, each member trigger can be handled separately, and so auto-resolve would probably work more as expected. A useful enhancement may be to add a sort of auto-disable feature into the metrics alerter. For example, we could add something like, "quietCount: N" and "quietDuration: Duration" options into the expression. The "quietCount" feature would mean: after a send, stay quiet (no resend) until after we see N false evaluations. The "quietDuration" feature would mean: after a send, stay quiet (no resend) until after we see false evaluations for the specified duration. For the ALL option is would stop any send. For the EACH option (our use case) it would stop sends for the relevant metrics. Perhaps this would be enough for us, it would stop repetition, although it would not automatically set alerts to RESOLVED, that would be done manually. Thoughts?	1	1	https://issues.jboss.org/browse/HWKALERTS-240
6067	HWKALERTS-240	> I think this could collides with the purpose of the Dampening... I agree that it is not an important extension when using "EvalType=ALL" on the Metrics Alerter condition. In this case you have a single trigger and the event generated by the alerter applies to the data for the entire set of target metrics. In this case I would recommend using dampening, autoDisable and autoResolve as usual, and there probably is no need to use quietCount. The critical use case is for "EvalType=EACH". This is the use case for this Jira, In this situation the external condition evaluates the condition expression for each target metric individually, and could generate events for multiple offenders. In essence this is like having a condition for multiple dataIds. We can't really use dampening, autoDisable or autoResolve effectively in this situation because we've got a single trigger dealing with multiple metrics. Adding quietCount to the alerter falls somewhere in between autoDisable and autoResolve. It's not really dampening because it doesn't delay a trigger firing. Instead, it prevents multiple firings *for the same metric* until we return to a steady state for that metric, as defined by a specified number of false evaluations. During this time the trigger remains active in order to service events for other metrics in the target set. Originally I had proposed quietCount and quietDuration but quietDuration isn't really useful. We already have evaluation frequency, so duration is just count * frequency. So, at this point I only propose quietCount. [~rutlucas], does this help explain the motivation?	1	1	https://issues.jboss.org/browse/HWKALERTS-240
6068	HWKALERTS-219	Hi Daniel, The update of single conditions was discussed in preliminar versions of Alerting and it has some implications in the internal process. What is the use case ? Take into account that a put operation is not heavy, and in terms of efforts it is even costly to update an attribute rather to update the full conditions set.	1	1	https://issues.jboss.org/browse/HWKALERTS-219
6069	HWKALERTS-201	Do you have some tests where we can check before/after ? Changes sound good, but it would be good to have something to get measures. Something like Perf*ITest could be good to follow a convention for these scenarios.	1	1	https://issues.jboss.org/browse/HWKALERTS-201
6070	HWKALERTS-137	Docs say Endpoint /actions/plugin/{actionPlugin} Find all action ids of an specific action plugin How can I distinguish between * no id available for this plugin -> 200 + empty colleciton * plugin not deployed -> 404 not found	1	1	https://issues.jboss.org/browse/HWKALERTS-137
6071	HWKALERTS-123	I haven't thought about this very much. Is this tied to a feature request for Hawkular? Anyway, I agree that conditions and dampening are possible paths but I'm thinking we may want to pull this up higher, and be able to "time-decorate" triggers. Although adding time aspects to individual conditions is intriguing, it could also get complicated to manage. And for dampening, even more complicated. If it gets too difficult to understand why a trigger didn't fire then value may be diminished. Instead, I can imagine letting a trigger fire, and then handling it in time-dependent ways like suppressing it, changing severity, or performing a different action set depending on the time. This would keep the existing model the same and fairly simple. The trigger severity and actions would act as defaults which could optionally be overridden by a time-decorator, or something like that. We need to look at some use cases. (it would be nice if Jira showed who was logged in, this comment is actually from [~jayshaughnessy])	1	1	https://issues.jboss.org/browse/HWKALERTS-123
6072	HWKALERTS-107	I think we don't have a clear way to not use the VFS library. In our scenario, a .jar deployed in a .war, should have visibility of the WEB-INF/classes to find specific annotations. So, classes are vfs resources and I don't know another way to inspect the vfs resources without using the propietary library. [~richard_opalka] or [~ropalka] what do you think ? This code is tested against EAP62, Wildfly9 and Wildfly10 versions. Would you suggest a different way to inspect classes folder package as VFS resource ?	0	0	https://issues.jboss.org/browse/HWKALERTS-107
6073	HWKALERTS-80	[~rutlucas], Is this still relevant or can it be closed?	1	0	https://issues.jboss.org/browse/HWKALERTS-80
6074	HWKALERTS-68	Can't accounts be used independently of the bus, and of hawkular? Would keeping accounts give standalone users an option for security? If we kept accounts how easy is it to ignore if a consumer doesn't care?	1	1	https://issues.jboss.org/browse/HWKALERTS-68

#	Requirement	Questions	Answered	Successfully Answered	URL
6075	HWKALERTS-65	Can you explain this a little bit further? Why would an action plugin have any dependency on alert status? I thought action plugins supported the definition of an action and that actions were assigned to triggers (currently, implicitly for open status). The actions were then requested by the engine and performed by the plugin. Why does the plugin care about the alert status in that model?	1	1	https://issues.jboss.org/browse/HWKALERTS-65
6076	HWKALERTS-65	Are you saying that a single action would perform differently based on alert status? I was assuming that different actions would be defined. It's not obvious to me that a single action should perform different tasks, this may be something to discuss in a broader group if that is the design.	1	1	https://issues.jboss.org/browse/HWKALERTS-65
6077	HWKALERTS-57	I don't think it is necessary to provide anything but a generic alert message out of box, with the ability for the user to create custom messages for specific triggers. The alerts component won't know about different kinds of triggers. Of course Hawkular does have different kinds of built-in triggers, for thresholds and avail, so Hawkular could define more messages. What is the "alertLink" in the generic message? What may be good is the ability to really pretty-print the Alert, or portions of the Alert (header, evals), in the message.	0	0	https://issues.jboss.org/browse/HWKALERTS-57
6078	HWKALERTS-52	Interoperability, great. Have you considered a particular system? Nagios checks maybe?	1	1	https://issues.jboss.org/browse/HWKALERTS-52
6079	HWKALERTS-49	> Type - this is used to categorise the situations, e.g. Exception, SLA Violation, etc. > Subject - the subject of the Situation, generally relates to a service These do sound like Tags to me. The values seem open-ended. If the Trigger were to be tagged the subsequent alerts could be accessed via the trigger's tags. > Timestamp > Assigned to We timestamp alerts at each lifecycle point: create/ack/resolved. And I think the ackBy field could be the assignedTo value. > Severity This should likely be an enhancement. low,medium,high,crtitical seem like good values. We can discuss and likely generate an enhancement Jira. > Link to associated activity event id - used to retrieve any associated message I'm not sure about this. It would differ for each Alert, I think, and so a Trigger tag is probably not appropriate. Perhaps it could be derived from the information already on the alert, like the condition evaluation information? It's not clear to me where this information would come from but perhaps it would require the need for alert-specific tags, something we don't have at the moment. > Are tags the equivalent concept? Should there be a way to define a tag as private? Tags can be flagged as visible=false, although the actual semantic for that is unclear. The thinking was that it should not be presented in a gui. * Properties (e.g. order item code, order value, customer id) Again, this seems like it would be alert-specific. And again, I'm not sure at which point this would be assigned. We'll need to discuss further. * Resolution State (unresolved, in progress, resolved) Yay, an existing mapping! * Context/Correlation Ids (e.g. order id) Perhaps already in the captured Alert information? * Parent alert id This is also something we should discuss. This sounds like a resubmission of a business transaction? If so, it would seem the alert id would be stored on some new BTResubmission object?	1	1	https://issues.jboss.org/browse/HWKALERTS-49
6080	HWKALERTS-49	Does the Action have access to the source event(s) that triggered the Alert? Otherwise not sure the Action would be able to establish the link. To give a concrete example, if the source event was an object representing a business transaction trace, and it had an id, then it would be useful to place that id within the Alert in a properties field - then (for example) the BTM UI could retrieve Alerts containing the 'btm id' property to present to its users, and then when the user selects a particular alert, the "btm id" property value can be used to retrieve the business transaction trace. Currently it appears the Alert does not have a properties type field - would it be ok to add this, to enable adhoc information to be added to Alerts?	1	1	https://issues.jboss.org/browse/HWKALERTS-49
6081	HWKALERTS-47	One behavioral point that needs to be formalized. Currently Tag.category is optional and can be null. The question is whether a null category when fetching should act as a value or as a wildcard. Do we match tags that are explicitly null:name or tags that are *:name ?	0	0	https://issues.jboss.org/browse/HWKALERTS-47
6082	HWKALERTS-43	Hi Mike, On which items would be interested ? Just new alerts ? or also some change in the definitions ? We have like two domains in alerts: the definitions part, that is mostly static, and the "live" part fhat is the alerts/actions as a result of a trigger. If you can ellaborate the idea, I can work on that as well. Thanks, Lucas	1	1	https://issues.jboss.org/browse/HWKALERTS-43
6083	HWKALERTS-43	Ok, ¿ in the technical side, which preferred mechanism could be for the implementation ? From what I understand from webhooks the point is that as a result of a new alert we should invoke a defined endpoint. right ? Do you have a previous webhook what I can take a look ? Normally, the notification/callback mechanism is made via actions-plugins, my first try will be to create a new one but if you point me how in other parts is doing today I can see what is the preferred way. Thanks !	1	1	https://issues.jboss.org/browse/HWKALERTS-43
6084	HWKALERTS-29	[~rutlucas] first draft jmeter script is attached. It just sends 1 request per endpoint. Could you please take a look and provide any feedback? Currently 2 issues noticed: 1. Endpoint /actions/send -> POST -> Parameter body Required false -> request without body returns Internal Server Error 2. If we have >1 trigger conditions for the same trigger: GET http://127.0.0.1:8080/hawkular/alerts/triggers/1/conditions [: { : : "triggerId":"1", : : "triggerMode":"SAFETY", : : "type":"STRING", : : "conditionId":"1-SAFETY-8-1", : : "dataId":"10", : : "ignoreCase":false : }, : { : : "triggerId":"1", : : "triggerMode":"SAFETY", : : "type":"STRING", : : "conditionId":"1-SAFETY-8-2", : : "dataId":"10", : : "ignoreCase":false : }, : { : : "triggerId":"1", : : "triggerMode":"SAFETY", : : "type":"STRING", : : "conditionId":"1-SAFETY-8-3", : : "dataId":"10", : : "ignoreCase":false : }, : { : : "triggerId":"1", : : "triggerMode":"SAFETY", : : "type":"STRING", : : "conditionId":"1-SAFETY-8-4", : : "dataId":"10", : : "ignoreCase":false : }, : { : : "triggerId":"1", : : "triggerMode":"SAFETY", : : "type":"STRING", : : "conditionId":"1-SAFETY-8-5", : : "dataId":"10", : : "ignoreCase":false : }, : { : : "triggerId":"1", : : "triggerMode":"SAFETY", : : "type":"STRING", : : "conditionId":"1-SAFETY-8-6", : : "dataId":"10", : : "ignoreCase":false : }, : { : : "triggerId":"1", : : "triggerMode":"SAFETY", : : "type":"STRING", : : "conditionId":"1-SAFETY-8-7", : : "dataId":"10", : : "ignoreCase":false : }, : { : : "triggerId":"1", : : "triggerMode":"SAFETY", : : "type":"STRING", : : "conditionId":"1-SAFETY-8-8", : : "dataId":"10", : : "ignoreCase":false : }] I try to update one condition (by PUT): PUT http://127.0.0.1:8080/hawkular/alerts/triggers/1/conditions/1-SAFETY-8-10 PUT data: { "conditionId" : "8", "type" : "STRING", "dataId" : 10, "triggerId" : "1", "triggerMode" : "SAFETY" } The following response is received: { : "errorMsg":"Condition 1-SAFETY-1-1 doesn't exist" } Response code: 400	0	0	https://issues.jboss.org/browse/HWKALERTS-29
6085	HWKALERTS-18	Lucas, is this now done, did it go in with the cassandra impl?	1	1	https://issues.jboss.org/browse/HWKALERTS-18
6086	HWKALERTS-17	Is this still valid? If so this needs some more information as to what the motivation is, and what the priority may be.	1	1	https://issues.jboss.org/browse/HWKALERTS-17
6087	HWKALERTS-17	Perhaps, with the status of the project this JIRA has not much sense. I think that this JIRA can be re-arrange to study if we can unify the packaging from a .ear to a .war, as we can add the EJB3 used as .jar. This has no priority at the moment as there is not issue related, it's just a suggestion to follow same package (if possible as other components).	1	1	https://issues.jboss.org/browse/HWKALERTS-17
6088	HWKALERTS-9	Would a simple Sauce Labs automated test validate this as well?	0	0	https://issues.jboss.org/browse/HWKALERTS-9

#	Requirement	Questions	Answered	Successfully Answered	URL
6089	HWKALERTS-4	So I've spent some time this morning looking at how PagerDuty we could integrate with PagerDuty. h5. Email The simplest solution is to use their email integration. In your PagerDuty account, you create a new email integration, and they give you a special email address. Then you simply configure the Hawkular email sender to use this address. PagerDuty has a rule system to help the user determine when an email is the signal for a new or existing incident. h5. API They also have a REST API. In this case, the caller must provide an incident key. I'm not sure yet how we could provide this information. If we use the triggerId then all actions will be sent to PagerDuty as the same incident. Which is wrong. On the other hand, it would also be a bad idea to send every action as a new PagerDuty incident. Is it possible with the current API to determine when a trigger (in the Hawkular sense) goes into the "alerted" state?	1	1	https://issues.jboss.org/browse/HWKALERTS-4
6090	HWKALERTS-4	[~jayshaughnessy] [~lponce] : I don't know if you get notifications for all JIRAs. Can we talk about that next week? Here, or maybe better on the phone? Thanks.	1	1	https://issues.jboss.org/browse/HWKALERTS-4
6091	HWKALERTS-2	Is it natural to think in some kind of polling ? I am thinking in something like this: 1.- I have a trigger checking if disk is full. 2.- Once disk is full, an alert is thrown and trigger is disabled to avoid multiple and unnecessary alerting. 3.- Then we can set up a rule like "if trigger was disabled by an alert, check each X interval if it should be re-enable". With [3] I see that an user doesn't need to define a explicit rule to enable a trigger, and this rule can be also generic.	1	1	https://issues.jboss.org/browse/HWKALERTS-2
6092	HWKAPM-892	Doesn't openshift (and possibly kubernetes) have a way to understand dependencies between the services - and therefore only indicate service is 'ready' only if its dependencies are also ready? Reason for question is to understand whether the server readiness probe only needs to indicate its own state, or also internally interrogate the readiness of its storage (i.e. cassandra)? If it does take into account the readiness of its dependencies, then can't we just keep the readiness of each component simple and (if a matter of milliseconds) just report readiness and liveness as both being when ready to handle requests?	1	1	https://issues.jboss.org/browse/HWKAPM-892
6093	HWKAPM-891	So is it not possible to get the client Provider automatically registered? If it is possible then would be good to support that aswell, to make simple instrumentation of Jax-rs app as non-intrusive as possible.	1	0	https://issues.jboss.org/browse/HWKAPM-891
6094	HWKAPM-873	The Hawkular project has a grafana datasource (https://github.com/hawkular/hawkular-grafana-datasource) for presenting metrics information, so potentially it would be possible for us to provide a datasource above Jaeger (regardless of underlying datastore). Would it be possible for you to outline the type of information you were looking to display in Grafana, so that we could assess what might be possible?	1	1	https://issues.jboss.org/browse/HWKAPM-873
6095	HWKAPM-873	What we want to achieve is, to store all logging information via span.log, so it would be amazingly easy to investigate issues when having the traceId. Following the semantic conventions of opentracing, the statuscode of http requests (i.e. 500) would be the starting point and we could list all logs of all services involved in this operation. Another important aspect is a high level view, that just shows failed requests with statuscode 4xx or 5xx per hour. Allowing a drill down with a view per service name / operation name. Grafana is our current choice, cause we also could build all kinds of metric dashboard with the same tool. If it would be possible to query jaeger with grafana via hawkular datasource, this would be awesome. One thing that comes to my mind is to build a high level dashboard which aggregates the before mentioned data and also having a drilled down dashboard per trace, which could be visualized with https://github.com/jdbranham/grafana-diagram (hehe just thinking loud, maybe it makes no sense). Grafana has some nice templating features, so other diagrams on the detailed dashboard could automatically show metrics of involved services. What do you think about my ideas? Maybe you guys have some better recommondations? Martin	1	0	https://issues.jboss.org/browse/HWKAPM-873
6096	HWKAPM-873	Hi Martin, thanks for the additional information. Your first requirement seems a better fit for the tracing UI itself, which would allow you to filter out the trace instances you were interested in and then view the individual spans and their logs/tags. Is there any particular reason why you would want this all in the Grafana UI, rather than finding a way to use both UI's together? The second requirement is more metrics based - Jaeger provides the ability for its clients to provide a way to capture certain metrics, and I believe in the near future it will be possible to provide these via a Prometheus endpoint, which would enable metric systems (such as Hawkular) to capture this info - and in the case of Hawkular present those details via Grafana. So in this particular case I don't think we even need a Grafana datasouce sitting above Jaeger itself. The grafana-diagram looks interesting, and if having close integration between a tracing UI (e.g. Jaeger) and Grafana was not possible, then it might provide a future solution - but would need to investigate other options first.	1	1	https://issues.jboss.org/browse/HWKAPM-873
6097	HWKAPM-873	Hey Gary, now I can come back to you with a more sophisticated answer :) Yeah, I agree on your points. The support team will probably not dig that deep into analysis when monitoring system performance. To switch UIs for in-depth analysis is okay, I guess. While tihinking about best approaches some other thoughts came to mind and you already answered som in your response: In general there are 3 types of metrics one would want to gather: - Timebased Metrics - Tracing information - Logging information And there are different use cases, where a different metric type applies best: 1) general system performance - distribution of response times (<200ms, >200, >1000, etc.) - distribution of http status codes (2xx, 3xx, 4xx, 5xx) - alerting rules 2) in-depth analysis, when alerts are triggered - find requests that have unwanted characteristics (i.e. 5xx server error, response is slow) - find the causing service in the chain of all invoked services for this request - find the root cause by analysing all logs that are attached to the identified trace So the idea is to have alle information for use case 1) in grafana, so we could build dashboards that show the state of our services, maybe group them by labels and so on. So for use case 1 a metrics system would be sufficient. My first idea was to aggregate tracing information for use case 1), but since tracing can be sampled it is probably not the best idea.Tracing and in-depth analysis is not needed at this point. For use case 2) we need the tracing information and all attached logs linked together. This is what Jaeger is already doing, so this should work fine. If i get this right, a connection from grafana to a Jaeger datasource/Cassandra is not needed at all for this constellation. We could build a scripted grafana dashboard which provides a link that prefills for example time range and desired filters like response status codes for the jaeger ui, if the ui supports it (didn't check it yet), so we could achieve a quite good integration. Can you provide more details about the jeager metrics and a possible prometheus endpoint? Sounds really good. Thanks for your time. This discussion helps a lot to refine our tracing concept.	1	0	https://issues.jboss.org/browse/HWKAPM-873

#	Requirement	Questions	Answered	Successfully Answered	URL
6098	HWKAPM-873	bq. We could build a scripted grafana dashboard which provides a link that prefills for example time range and desired filters like response status codes for the jaeger ui, if the ui supports it (didn't check it yet), so we could achieve a quite good integration. If it doesn't then just create an issue in the https://github.com/uber/jaeger-ui repo, as this sounds like a useful feature. bq. Can you provide more details about the jeager metrics and a possible prometheus endpoint? Sounds really good. I've just created an issue for this on the Jaeger Java client: https://github.com/uber/jaeger-client-java/issues/172 Its something we are interested in supporting in the near future - so feel free to add your ideas to that issue. Once Java is implemented, we could look at other languages such as Node.js.	0	0	https://issues.jboss.org/browse/HWKAPM-873
6099	HWKAPM-861	It seems those are the places that need a change: {code} $ mvn clean >/dev/null && grep jdoe . -r ./README.md:-> go to http://localhost:[8080|9411]/ and log in with jdoe:password ./README.md:* `-Pdev` - create user jdoe:password ./dist/src/main/security/application-users.properties:jdoe=35c8096d943bf6555d0a39f35fd5c7e0 ./dist/src/main/security/application-roles.properties:jdoe=read-write,read-only ./dist/src/main/resources/apm/setenv.sh:export HAWKULAR_APM_USERNAME=jdoe ./server/processors/src/test/java/org/hawkular/apm/server/processor/nodedetails/NodeDetailsDeriverTest.java: Property p2 = new Property(Constants.PROP_PRINCIPAL, "jdoe"); ./examples/vertx-opentracing/README.md: export HAWKULAR_APM_USERNAME="jdoe" ./examples/polyglot-zipkin/README.adoc:$ curl -ivX POST -H 'Content-Type: application/json' 'http://localhost:3001/nodejs/createUser' -d '{"name": "jdoe"}' ./examples/polyglot-zipkin/README.adoc:$ curl -ivX POST -H 'Content-Type: application/json' 'http://localhost:3000/dropwizard/users' -d '{"name": "jdoe"}' ./examples/polyglot-zipkin/README.adoc:$ curl -ivX POST -H 'Content-Type: application/json' 'http://localhost:3003/wildfly-swarm/users' -d '{"name": "jdoe"}' ./tests/app/polyglot-opentracing/js/index.js: recorder: new hawkularAPM.HttpRecorder('http://hawkular-apm:9080', 'jdoe', 'password'), ./tests/app/polyglot-opentracing/js/index.js: res.end(JSON.stringify({name: "jdoe"})); ./tests/app/polyglot-opentracing/java/README.md:$ export HAWKULAR_APM_USERNAME=jdoe ./tests/app/polyglot-opentracing/java/src/main/java/org/hawkular/apm/tests/app/polyglot/swarm/rest/UsersHandler.java: return Response.ok("{name: \"jdoe\"}").build(); ./tests/instrumentation-framework/src/test/resources/apmagent-env-variables.properties:HAWKULAR_APM_USERNAME=jdoe ./tests/instrumentation-framework/src/test/resources/apm-env-variables.properties:HAWKULAR_APM_USERNAME=jdoe ./tests/instrumentation-framework/src/test/resources/apmotagent-env-variables.properties:HAWKULAR_APM_USERNAME=jdoe ./tests/container/standalone/src/test/java/org/hawkular/apm/tests/client/jav/ClientJavaMainITest.java: private static final String TEST_USERNAME = "jdoe"; ./tests/container/wildfly/pom.xml: <javaOpt>-DHAWKULAR_APM_USERNAME=jdoe</javaOpt> ./tests/container/jetty/pom.xml: -DHAWKULAR_APM_USERNAME=jdoe ./tests/container/jetty/src/test/java/org/hawkular/apm/tests/client/jetty/ClientJettyStreamITest.java: private static final String TEST_USER = "jdoe"; ./tests/container/jetty/src/test/resources/realm.properties:jdoe: password,user ./tests/dist/pom.xml: <argument>jdoe</argument> ./tests/dist/src/test/java/org/hawkular/apm/tests/dist/AbstractITest.java: public static final String HAWKULAR_APM_USERNAME = "jdoe"; ./tests/instrumentation/pom.xml: -DHAWKULAR_APM_USERNAME=jdoe ./tests/agent-opentracing/src/test/java/org/hawkular/apm/tests/agent/opentracing/server/http/JavaxServletSyncServerITest.java: private static final String TEST_USER = "jdoe"; ./tests/agent-opentracing/src/test/resources/realm.properties:jdoe: password,user ./client/agent/dependency-reduced-pom.xml: -DHAWKULAR_APM_USERNAME=jdoe ./client/agent/pom.xml: -DHAWKULAR_APM_USERNAME=jdoe {code} It got me thinking: what to do with the old tutorials and instructions? We'll break those all if we just change from "jdoe" to "admin". Is it worth it?	1	1	https://issues.jboss.org/browse/HWKAPM-861
6100	HWKAPM-861	How about adding admin:password as a second user? so jdoe is still supported.	0	0	https://issues.jboss.org/browse/HWKAPM-861
6101	HWKAPM-847	[~objectiser] I see the list being updated, the only situation where that doesn't happen is having a service selected (in the dropdown) and a new buildstamp exists, it will need to select a different service and then back to the previous to see the new buildstamps. Is this what you are observing as well ?	0	0	https://issues.jboss.org/browse/HWKAPM-847
6102	HWKAPM-837	I will start as always with my why questions... Are users(developers) of OT interested in a text representation of returned HTTP 500 codes? Why is it good to capture this text message returned to API users? In my opinion this text often say something like "Some error happened and we are working on it...", but internally you might want log different data/text/error code to find a real problem..... Is returned text information helpful for operators to find the root of the error? Probably not, but if they look at trace data they would find more information how the request flowed through system, even better if all exceptions are logged. This is an instrumentation issue right? In jaxrs a business error would be probably represented by an exception. Which is handled by exception mapper. It would make sense to add a "error" log "event" with payload like exception message and stack trace? I did this for spring instrumentation: https://github.com/opentracing-contrib/java-spring-web/pull/1/files#diff-caa4c5580cd6b8b30a1362ec751c474cR91 Would it make sense to use log instead of `fault` tag? Would it be possible to write rules to capture exceptions? [1]: https://docs.oracle.com/javaee/7/api/javax/ws/rs/ext/ExceptionMapper.html	0	0	https://issues.jboss.org/browse/HWKAPM-837
6103	HWKAPM-835	I think that it's important to keep performance impact of instrumentation as low as possible. In other projects they even care about the performance of spans serialization... Data on UI can be processed easily to the most appropriate form (e.g. timeline view or nodes) while keeping data model clear and consistent and self-describing. I don't see that this approach will put more work on the server. Other thing about the consideration should be processing of reported data on the server. Concretely we use JMS retry mechanism to wait for "all" fragments being reported to be able to calculate trace completion time. Processing of the spans could be simplified to keep as minimal state as possible e.g. first span arrives we can derive so called "TraceCompletionTime" and publish it. Next span arrives, use data from previous event and publish next event. When claiming putting more work on the server I would like to see why. The current solution is not performance tuned either. I'm bit scared how JMS with our retry approach would perform when monitoring big systems. Yes, I'm talking about our own data model supporting all OT features including the possibility of multiple references (parents). This model should replace the whole Trace model (step by step). We could start simple: 1. well define model, Think about all aspects, especially about our "analytics" events, what impact it will have. 2. implement REST/DB handling 3. analytics handling on top of span model. + unit/integration test for it. 4*. use model in our providers js+java 5*. teach UI to use new model 6. merge branches from 4+5 and remove Trace model.	0	0	https://issues.jboss.org/browse/HWKAPM-835

#	Requirement	Questions	Answered	Successfully Answered	URL
6104	HWKAPM-835	> Support for baggage... whether the baggage and tags should be treated as separate concepts Correct me if I'm wrong: * Tags are designed to add data to the span which are reported to the server and are not passed to children spans either available for reading to the API consumer. * Baggage is designed to add data to the span which is available for reading by API consumer and it's not reported to the server. Does it answer the question? > Trace id We discussed this, currently OT spec assumes all spans are assumed from one traceId. About the transactions. For me it makes sense that one trace traverses multiple transactions. E.g. in microservices each team is responsible for a particular service and can define transactions for their service. Then usually to fulfill some business task more "transactions" can be involved (get users->get content of shopping cart->). Does it make sense?	0	0	https://issues.jboss.org/browse/HWKAPM-835
6105	HWKAPM-835	Could you point me at the time in the video when there is baggage shown in UI? What I saw was {{log}} payload, concretely mutex queue on fryer which contains ordered dunuts favours with client ids.	0	0	https://issues.jboss.org/browse/HWKAPM-835
6106	HWKAPM-830	I'm not UX expert by it seems reasonable to me that filter would stay if I selected something. I will play more with it and see how it goes... Proposed tab names look good. Could we find a better name for "Distributed Tracing? The whole tool is for distributed tracing. I'm proposing: Latest Traces +2, Traces +1. Tab would by default sort traces from latest to oldest.	0	0	https://issues.jboss.org/browse/HWKAPM-830
6107	HWKAPM-799	Is this the right way to do it? Sort ascending on server side and on client side re-sort do a descending order?	1	1	https://issues.jboss.org/browse/HWKAPM-799
6108	HWKAPM-779	Looks good to me - have you asked Heiko for comment? I haven't had any experience with fractions - does it provide some way to provide configuration, e.g. sampling, recorder?	1	1	https://issues.jboss.org/browse/HWKAPM-779
6109	HWKAPM-779	So it would not be possible to have a generic opentracing fraction where users would provide their favorite tracer implementation. The best what we can do is to implement opentracing-hawkular fraction. Should I proceed with it?	1	1	https://issues.jboss.org/browse/HWKAPM-779
6110	HWKAPM-764	Just to clarify, will the sampling decision be applied prior to the trace fragment being submitted to the recorder (which may potentially be batching fragments)?	1	1	https://issues.jboss.org/browse/HWKAPM-764
6111	HWKAPM-762	Just to clarify, are you thinking the HAWKULAR_APM_SERVICE_* can be used in place of the HAWKULAR_APM_URI? So we update the opentracing providers to be aware of these vars aswell? Also - the OPENSHIFT_BUILD_NAME could be used as a default 'service' tag/property - but should it be qualified by the project name (OPENSHIFT_BUILD_NAMESPACE) - I guess this will depend upon how the projects are used.	1	1	https://issues.jboss.org/browse/HWKAPM-762
6112	HWKAPM-759	Just to clarify, are you saying that we can rename the book, and it will update the URI to the book? but is not dependent upon the gitrepo name?	1	1	https://issues.jboss.org/browse/HWKAPM-759
6113	HWKAPM-758	What are the cons/pros compared to Apache Kafka?	1	1	https://issues.jboss.org/browse/HWKAPM-758
6114	HWKAPM-754	[~pavol_loffay] Would be good to get a consistent Sampler api and implementations across providers. Is the proposed API ok? {noformat} Sampler { SampleLevel getSampleLevel(Trace); } {noformat} (probably will rename ReportingLevel to SampleLevel) Also does the rate limiting implementation make sense as the initial simple implementation that we provide out of the box?	1	1	https://issues.jboss.org/browse/HWKAPM-754
6115	HWKAPM-754	Could you please define `SampleLevel` semantics? Sampler should basically tell if the trace should be reported e.g. passed to the underlying reporter or not. What you designed returns a level and client code would have to decide whether report the data or not. Something like this would make more sense to me: {code:java} Sampler { boolean isSampled(Trace); } {code}	1	0	https://issues.jboss.org/browse/HWKAPM-754
6116	HWKAPM-754	Sure no problem. Could you please describe that your rate limiting impl in some pseudo code?	1	0	https://issues.jboss.org/browse/HWKAPM-754
6117	HWKAPM-753	Discussion from opentracing gitter channel: {noformat} Pavol Loffay @pavolloffay 16:36 Hey, I have a question: Is it valid to pass spans from different traces to tracer.startSpan("foo").addReference(span1).addReference(span2)? Gary Brown @objectiser 18:26 The spec says "A Span may reference zero or more other Spans that are causally related. OpenTracing presently defines two types of references: ChildOf and FollowsFrom. Both reference types specifically model direct causal relationships between a child Span and a parent Span". It does not directly say whether those spans belong to the same or different trace instances. But I believe in a 'join' scenario, the spans could belong to different trace instances. It would be good to get clarification from the spec authors. bhs @bensigelman 23:43 @objectiser I agree that the spec should be clearer. By definition, Spans are in the same trace if they have transitive references to each other. The fact that Dapper/Zipkin choose trace_ids at the root and expect them to be unique throughout a trace is a little problematic but not a dealbreaker here… I personally really regret this aspect of Dapper's design, FWIW, as it limited the design space and for no (very) good reason I've filed an issue about improving the spec: opentracing/specification#19 {noformat}	1	0	https://issues.jboss.org/browse/HWKAPM-753
6118	HWKAPM-743	Is OpenTracing making use of the ServiceLoader API? If so, that's enough for an implementation to be chosen at runtime by just plugging it in. As a provider, we might chose to provide an additional CDI integration, but I would keep those separate: the actual implementation in a library (with META-INF/services/io.opentracing.Tracer) and a CDI integration. As a consumer, I might have several reasons why I'd want to have my own CDI producer for the Tracer.	1	1	https://issues.jboss.org/browse/HWKAPM-743
6119	HWKAPM-743	[~juraci.costa] For clarity, could you provide an example of why a developer would want their own CDI producer? The current opentracing impl can be obtained via service loader, but the opentracing spec (or abstract java impl) does not make any recommendations on how the tracer should be obtained by the app.	1	1	https://issues.jboss.org/browse/HWKAPM-743
6120	HWKAPM-743	One case could be to provide it in a different scope as the one provided by the implementation: I might want to make it request-scoped or session-scoped in a multi tenant application, but as application scoped for a single-tenant solution. Or I might want to report different data to different servers. Or I might just want to wrap it on my own implementation, to workaround a possible bug. {quote}The current opentracing impl can be obtained via service loader, but the opentracing spec (or abstract java impl) does not make any recommendations on how the tracer should be obtained by the app.{quote} How does it look like? I tried to find an example on OpenTracing's repo, but couldn't find one.	1	1	https://issues.jboss.org/browse/HWKAPM-743
6121	HWKAPM-736	"Do we want to make it easier for an APM user to create alert triggers? If so, what types of trigger should be supported, and how should they be managed in the UI (i.e. to enable edit/deletion)." Bringing some experience from ManageIQ integration. In ManageIQ we use the concept of alert template which maps with a groupt trigger (in Hawkular Alerting terminology), this means that trigger is created once and propagated in several instances with the benefit to modify the group trigger and see the change propagated at members level. This can also helps to reduce the steps to deal with triggers definitions.	1	1	https://issues.jboss.org/browse/HWKAPM-736
6122	HWKAPM-736	Could you provide some examples of what an alert template may look like, based on the scenario I have mentioned? e.g. creating an alert trigger to check if the completion time is above a particular threshold, possibly also with some property filter? For example, if the trigger was for: transaction=Place Order url=/x/y/z operation=POST property(name=item,value=car), duration over 50ms	1	0	https://issues.jboss.org/browse/HWKAPM-736

#	Requirement	Questions	Answered	Successfully Answered	URL
6123	HWKAPM-736	Can we somehow reuse those things from MiQ? IIRC, it was done as web components, so, there's a chance we might be able to reuse a lot of the stuff from there.	0	0	https://issues.jboss.org/browse/HWKAPM-736
6124	HWKAPM-736	"Can we somehow reuse those things from MiQ? IIRC, it was done as web components, so, there's a chance we might be able to reuse a lot of the stuff from there." I don't think so. ManageIQ integrates Hawkular Alerting from backend-to-backend. So ManageIQ UI is only talking with ManageIQ model of Alerts/Alerts profiles and their own Events semantic. What we do, is to translate those models into Hawkular Alerting. Some of them are pure 1to1 (i.e. MiQ Alert into a Group Trigger) and others need some adjustments.	1	1	https://issues.jboss.org/browse/HWKAPM-736
6125	HWKAPM-736	"For example, if the trigger was for: transaction=Place Order url=/x/y/z operation=POST property(name=item,value=car), duration over 50ms" On this case, how are modelling the events ? I mean, this transaction is coming from several services / applications ? Do we want to be notified per service/app ? Or we want a single trigger per all the possible events of or possible services ? I think that is the main point to discuss how we want to model. In a nutshell, Hawkular Alerting operates at dataId level. That means that a condition is linked with a dataId. If it is too generic (i.e. dataId="apm") a condition will process everything, and it we personalize it (i.e. dataId="service-a", dataId="service-b"). We can operate to such degree of granularity to define what we want to evaluate for alerting.	1	1	https://issues.jboss.org/browse/HWKAPM-736
6126	HWKAPM-736	[~jayshaughnessy] [~rutlucas] We are going to be looking at a change to the way we generate trace completion events - currently a single event is generated when the end to end trace is built and the time derived - however there is a delay in reporting the event, as we need to be sure we have received all possible data related to the trace. We are considering taking a different approach where we may produce a trace completion event earlier, and in the case where further trace data is received we then send one or more update events. So potentially for a single trace instance, there may be a series of events - but ideally from an alerts perspective we would like it to be processed as a single event? For example, the event has an id - so if the alerts engine receives multiple events for the same id - does it re-process the event and any subsequent effect - i.e. if the first event didn't meet the criteria, but the second one does, then it triggers after the second event - but then if a third update for the same event arrives that means the criteria is no longer valid, is the alert (that was triggered) cancelled? Alternatively, if during the updates, the alert is triggered, is it possible to then suppress further processing of the event, so that the alert actions are not performed for each subsequent update?	1	0	https://issues.jboss.org/browse/HWKAPM-736
6127	HWKAPM-731	Looks good - I think just using the single unit on the service dependency diagram is probably best. How do the durations show up in the tooltips? In the same way as in the nodes?	1	1	https://issues.jboss.org/browse/HWKAPM-731
6128	HWKAPM-731	[~objectiser], I've attached two screenshots for the DT and Instance diagrams, with single and two units, for comparison. WDYT ? For me, I'm not sure about this (and it's kind of a general feeling I'm having with this change) as sometimes it's hard to compare due to different units.. maybe with a different range, like not switching immediately to seconds after 1000ms, but a higher value, like 5000.	1	1	https://issues.jboss.org/browse/HWKAPM-731
6129	HWKAPM-731	It would not need to apply the same pattern for all units ([~gcardoso], can you confirm?), so we could have different rules for different units. I think above some threshold (let's say when we enter the minutes), second unit make sense.. 1.25 minutes is not clear that is actually 1m 15s. Before that, I think we can go with single unit (as 1.234 seconds is actually 1s 234ms). My concern is when we see a 700 value, we need to be carefully looking if it's 700µs or 700ms.. same for ms vs s.	1	1	https://issues.jboss.org/browse/HWKAPM-731
6130	HWKAPM-722	[~gcardoso], looks good to me. Is there any patternfly component we might (re)use for this or should be done from scratch ?	1	0	https://issues.jboss.org/browse/HWKAPM-722
6131	HWKAPM-722	I'm going to add some my thoughts: * it should be possible to expand/collapse spans from one service -- EDIT actually it should be possible to expand/collapse any span which has children. * logs(event in time) should be highlighted in span. How to deal with events which start/end time is out of span time span? We can show them at beginning/end of span or show them at right place within the trace - in case they are in trace <start, end> interval, or show them at beginning/start of trace (in line of given span) if they are outside of span timespan. * Highlight critical path * consider showing a vertical line across all expanded spans. It should help to see which spans started at the same time.. Color scheme: * each span with different color or the same color for spans from the same service?	0	0	https://issues.jboss.org/browse/HWKAPM-722
6132	HWKAPM-721	[~gcardoso] Hi Gabriel - we are thinking about adding an extra level - do you have any thoughts on another colour that could be introduced in the scale?	1	1	https://issues.jboss.org/browse/HWKAPM-721
6133	HWKAPM-721	Currently we have 4 colors. Do u want to go for a 5 color scheme? The extremes would still be green and red, correct?	1	1	https://issues.jboss.org/browse/HWKAPM-721
6134	HWKAPM-721	Yes - we have in interim solution which adds pf-orange-500 just before the red. Is this an acceptable addition, or should a different set be used?	1	1	https://issues.jboss.org/browse/HWKAPM-721
6135	HWKAPM-721	[~gcardoso], I guess you've missed attaching the image. As I've discussed with Gary before, I think it's harder to distinguish between what is a good/better green than between darker yellow/orange. I'm attaching a sample with the currently used colors, [Get Account] is the lighter orange (gold), Account Manager is the middle orange and AccountsDB[RetrieveAccount] is the newly introduced darker orange. !8TscfQPPP.png|thumbnail! WDYT ?	1	0	https://issues.jboss.org/browse/HWKAPM-721
6136	HWKAPM-720	I uploaded two images using the same color code based on patternfly colors. Also we should use the same font sizes / colors / icon sizes for both charts. We can even define same spaces inside the boxes for both (they are slightly different now) if possible. What do u think about the colors?	1	1	https://issues.jboss.org/browse/HWKAPM-720
6137	HWKAPM-720	Difficult to say - I think the colours are fine, but personally prefer the single left strip rather than the whole node being in the colour. But think we need to get a wide range of views to see what is preferred. Would it be possible to provide an image using these colours but just down the left hand side of the symbols, to provide a comparison?	1	1	https://issues.jboss.org/browse/HWKAPM-720
6138	HWKAPM-720	Good, thank you. What do you mean by this minor change? Perhaps I took and old screenshot to make my proposal. I don't intent to change anything but the color strip and the background color itself (all gray now).	1	1	https://issues.jboss.org/browse/HWKAPM-720
6139	HWKAPM-711	What do you exactly mean by logging provider? Span Logs contextually enrich trace instance. Why would one want to send/store logs in a separate system? Could you please create a separate jira to discuss this.	1	1	https://issues.jboss.org/browse/HWKAPM-711
6140	HWKAPM-703	A tooltip has the downside of not being easy to copy it's content (due to it's nature of show/hide on mouser over or even on click). Maybe it'd help if we truncate the properties to a max number of chars (eg, 20) and then on click expand them, would that be good ? but still showing the multiple instances of it.	0	0	https://issues.jboss.org/browse/HWKAPM-703
6141	HWKAPM-693	Should this be put on hold, because of the possible integration with Metrics? Or will we also be sending data to Alerts, even with integration with Metrics?	1	0	https://issues.jboss.org/browse/HWKAPM-693

#	Requirement	Questions	Answered	Successfully Answered	URL
6142	HWKAPM-683	The http query string will now be in a property called "http.query". If the query string contains multiple property/value pairs, and you would (also) like these individual values defined as separate properties, could you raise a separate jira so that the issue can be discussed.	1	1	https://issues.jboss.org/browse/HWKAPM-683
6143	HWKAPM-664	Hi Gary, Since I'm not so familiar with APM, I could not find what you are talking in the UI. Could you send some screenshots, including of Zipkin? If u want we can also schedule a meeting and we can discuss this.	1	1	https://issues.jboss.org/browse/HWKAPM-664
6144	HWKAPM-630	Issue may be the length of the list if a large number of services are involved, and may not be clear to the end user which list items represent the entry points. What about if the list contains the endpoints that are initially used to trigger a trace?	0	0	https://issues.jboss.org/browse/HWKAPM-630
6145	HWKAPM-602	I think milliseconds should be the lowest unit. If the value is over 1 second, then use seconds. If the value is over 60 seconds use minutes If the value is over 60 mins then use hours. If value is over 24 hours then days Although probably never used, might aswell support years awell? All values to 3 decimal places? NOTE: This issue is dependent upon HWKAPM-638 which is standardising all time units on microseconds. [~pavol_loffay] ok with you?	1	1	https://issues.jboss.org/browse/HWKAPM-602
6146	HWKAPM-602	In terms of distributed tracing and analyzing timing data, when you want to find the service which took the longest time maybe it's better to start with micro seconds. If you are running some highly optimized services written in Go/C++ I think It can happen that you expect durations in microseconds, then you can easily distinguish which took longer(milli) than expected unit (micro). What do you think? The upper bound probably won't happen a lot. My concern is about using years/days/hours/minutes with decimal numbers. Isn't it strange representation? e.g. 54,235 hours? Would'n be better to say 54h Xmin Y.YYYs? I was also thinking about using only seconds. Here are some numbers: 60mins = 3600s 24h= 86400s 1year=31536000s	1	1	https://issues.jboss.org/browse/HWKAPM-602
6147	HWKAPM-602	Is there any reason why not to start with micro seconds? Imagine your services run in micro seconds, then when there is ms somewhere in the UI you can find it quicker.	1	1	https://issues.jboss.org/browse/HWKAPM-602
6148	HWKAPM-602	We should think more about usability, and also what would we prefer to see in UI, what can make UI more useful..	1	1	https://issues.jboss.org/browse/HWKAPM-602
6149	HWKAPM-598	Currently the node does not have an explicit id, but could create and use an id based on the relative position of the node in the hierarchy. Do we need the fragment id to be stored separately - is there a benefit of being able to identify these separately? If a composite id, then would need to extract the fragment before being able to locate the fragment and navigate to the node, so more efficient if separate. Only issue is to ensure that these relative ids are not unstable (i.e. if a change occurs to the structure after an id has been derived and used).	0	0	https://issues.jboss.org/browse/HWKAPM-598
6150	HWKAPM-577	I'm wondering which would be the use cases for it. Is it to make it easier for developers to replace their dev instances of zipkin with APM? Or for production? For production, it's usual to have a router/reverse proxy in front of backends, to serve as SSL termination points and to "map" HTTP ports to backend ports (80 -> 9411, 80 -> 8080). I couldn't find on OpenShift a place to configure the port mapping, but I suspect there's such an option. For the developer case: would it suffice to provide a "zipkin-compat.cli" , where the port mapping would be applied to Wildfly?	0	0	https://issues.jboss.org/browse/HWKAPM-577
6151	HWKAPM-562	Could we have a quick chat to discuss this? I think a far simpler, easier to implement and consume would be to just add an HTTP header to the response, like: "Hawkular-APM-Missing-Records: 1". This is generic enough to be easily implemented by any endpoint, and generic enough to have the UI show a toast to the user if any responses contains this: "Note that the results you are seeing are not complete. We estimate that 1 record could not be shown."	0	0	https://issues.jboss.org/browse/HWKAPM-562
6152	HWKAPM-562	So no exception, just the endpoint reporting the number of missing records? How would the AnalyticsService API change? Have a special list type that includes the extra field?	1	1	https://issues.jboss.org/browse/HWKAPM-562
6153	HWKAPM-556	Potentially later on but we could have a static version for client side initially if we wanted to replace the PropertyUtil mechanism? For now maybe keep the PropertyService as a backend server component that uses PropertyUtil to retrieve values it does not have in the cache. And client instrumentation code use PropertyUtil. We can then review this later. Yes infinispan cache.	0	0	https://issues.jboss.org/browse/HWKAPM-556
6154	HWKAPM-555	This is not clear to me: "Add method that can return all comms details with an id field of a particular value" -> but there is already {{public CommunicationDetails get(String tenantId, String id)}} or should it return all comms details which id contains some string?	1	1	https://issues.jboss.org/browse/HWKAPM-555
6155	HWKAPM-495	Given the focus on cloud, it may be interesting to explore options around the test applications being provided as registered docker images - so the test framework/harness would take a (list of) docker image(s) and create an instrumented version which is then deployed and tested. Things to consider: * how would the version of technology be choosen? possibly via an environment variable supplied when building the image - but it may be that the version is simply associated with the image version - so to use a new version will require a new docker image version * instrumenting a docker image is discussed here (http://www.hawkular.org/blog/2015/12/01/hawkular-btm-instrument-docker-image.html) and (https://hawkular.gitbooks.io/hawkular-apm-user-guide/content/installation/CLIENT.html#_instrumenting_a_docker_image) - although this could be done by creating a derived image for each technology app image being tested, it would be good to see if there is a general purpose way of doing this (as a separate jira) that could be provided as a tool to users to help them automate the instrumentation of their own images	0	0	https://issues.jboss.org/browse/HWKAPM-495
6156	HWKAPM-462	This may not be necessary - should be able to run the 'add-user.sh' script externally on the container - but needs to be tested. But how will this work when the image is deployed in openshift/kubernetes?	1	1	https://issues.jboss.org/browse/HWKAPM-462

#	Requirement	Questions	Answered	Successfully Answered	URL
6157	HWKAPM-449	server. Spans contain identifying information such as traceId, spandId (id), parentId, and RPC name. I have tried only dropwizard-zipkin [2] with http transport. It seems that instrumented app reports only complete traces (It report them as complete list of spans). Annotation: Is used to record an occurance in time cs - Client Start. The client has made the request. This sets the beginning of the span. sr - Server Receive: The server has received the request and will start processing it. The difference between this and cs will be a combination of network latency and clock jitter. ss - Server Send: The server has completed processing and has sent the request back to the client. The difference between this and sr will be the amount of time it took the server to process the request. cr - Client Receiver: The client has received the response from the server. This sets the end of the span. The RPC is considered complete when this annotation is recorded. lc(binary annotation)- is the component or namespace of a local span. While structurally the same, local spans are fundamentally different than RPC spans in how they should be interpreted. For example, zipkin v1 tools center on RPC latency and service graphs. Root local-spans are neither indicative of critical path RPC latency, nor have impact on the shape of a service graph. By flagging with "lc", tools can special-case local spans. Binary annotation: Does not have timestamp, provide extra information about RPC Scenario 1: Client -> Service1 -> Service2 [https://github.com/smoketurner/dropwizard-zipkin/blob/4adc5435e28b3b7f8284f31585bd2d14e617c422/zipkin-example/src/main/java/com/example/helloworld/resources/HelloWorldResource.java#L24] produces: {code:java} [{ "traceId": "3ad0b2acb841c26c", "name": "get", "id": "064c292e02db027c", "parentId": "3ad0b2acb841c26c", "timestamp": 1467297761131000, "duration": 785000, "annotations": [{ "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 }, "timestamp": 1467297761131000, "value": "sr" }, { "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 }, "timestamp": 1467297761916000, "value": "ss" }], "binaryAnnotations": [{ "key": "http.status_code", "value": "204", "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 } }, { "key": "http.url", "value": "http://127.0.0.1:8888/result", "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 } }] }, { "traceId": "3ad0b2acb841c26c", "name": "get", "id": "064c292e02db027c", "parentId": "3ad0b2acb841c26c", "timestamp": 1467297761128000, "duration": 789000, "annotations": [{ "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 }, "timestamp": 1467297761128000, "value": "cs" }, { "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 }, "timestamp": 1467297761917000, "value": "cr" }], "binaryAnnotations": [{ "key": "http.url", "value": "http://127.0.0.1:8888/result", "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 } }] }, { "traceId": "3ad0b2acb841c26c", "name": "get", "id": "3ad0b2acb841c26c", "timestamp": 1467297760278000, "duration": 1639000, "annotations": [{ "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 }, "timestamp": 1467297760278000, "value": "sr" }, { "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 }, "timestamp": 1467297761917000, "value": "ss" }], "binaryAnnotations": [{ "key": "http.status_code", "value": "204", "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 } }, { "key": "http.url", "value": "http://localhost:8888/fetch", "endpoint": { "serviceName": "dropwizard-zipkin", "ipv4": "127.0.0.1", "port": 8888 } }] spanId (span name) -> parent id timestamp	1	1	https://issues.jboss.org/browse/HWKAPM-449
6158	HWKAPM-449	For the second scenario, I've just produced an abbreviated version of the hierarchy: {noformat} Trace Consumer (no correlation id as parent id is null) (not sure what the URI should be, or whether it related to http, as there is not enough info) Component (lc should be mapped to a component - ideally of type Database, but not sure if enough info here) Component (type = Database, query = "select distinct .....") - normally I would say cs/cr would map onto a Producer, but in special cases (e.g. jdbc.query) it may be best to just map it onto a Component with the Database type? Component (type = Database, query = "select distinct .....") Component (type = Database, query = "select distinct .....") {noformat} When building this type of Trace structure, where it is a multi level hierarchy, it will be necessary to keep track of the span id associated with a node, so that new nodes can be attached as child nodes based on the parent id. The only case where parent id should not be used in this way is where the annotation type is sr/ss, where a new Trace should be created with it being mapped to the top level Consumer node.	0	0	https://issues.jboss.org/browse/HWKAPM-449
6159	HWKAPM-449	The challenges I see with building these Trace structures from the spans, is: 1) If multiple proxy servers are running in a cluster, how to construct the traces when potentially the different spans required for a specific trace may be received by different servers - so possibly requires some caching - but that gets complicated when needing to update the same nodes, requiring synchronisation/transactions - which would impact performance. 2) How to know when a particular Trace fragment has been completed? May need to be some time delay after last update - not ideal. The other option is to deal with each span as a separate asynchronous action, which would result in one Trace per span. I'll see if I can provide an alternate example of each scenario encoded in this way. It would overcome the issues above, but the downside is it would create more Trace objects than necessary.	0	0	https://issues.jboss.org/browse/HWKAPM-449
6160	HWKAPM-449	An alternative plan - as mentioned above, one approach would be to create Trace fragments using a simple non-clusterable solution, but this may run into issues if needing to scale. The other approach is to not create the Trace fragments, but instead use the span information to create the derived information which is currently created from the Trace fragments. The Trace fragments themselves are not extensively used at the moment, but will be used more when focusing in on displaying the individual end to end Trace for a single invocation of the application. The other purpose of the Trace fragment is to record business properties and message content (if extracted) both of which are not relevant for zipkin clients. If we could this approach, then the spans could be used to create (and publish) the following pieces of information: NodeDetails - an instance is created per node in the trace fragment - so in the case of a zipkin span, it would be one per span. The type of the Node would be determined based on the examples shown above. FragmentCompletionTime - this records the time taken within a particular fragment, which is equivalent to a service. As zipkin primarily seems to deal with synchronous applications currently, this could simply be recorded for each Consumer node (i.e. sr/ss) with the duration between those two annotations. TraceCompletionTime - this is similar to the fragment completion time above, except only record for the span that has a null parent id - again assumes a synchronous invocation, but that seems to be the case at the moment - will need to investigate more complex solutions (as implemented in APM) when dealing with asynchronous cases. CompletionDetails - this information represents the latency between client and server - so this will require some information to be cached, as it will be necessary to calculate the different between the client's span duration and the server's span duration for a particular shared span id. Using this approach it will short cut creating the Trace at this stage - but would mean it may need to be constructed at a later stage when displaying information about an individual application invocation. Performance is less of an issue in this case, so this is probably the better solution. Have a think about both options and we can discuss if you want.	0	0	https://issues.jboss.org/browse/HWKAPM-449

#	Requirement	Questions	Answered	Successfully Answered	URL
6161	HWKAPM-445	The other issue is how we determine a property has numerical values - should this be inferred or recorded in the config somewhere - but then if the config changes, or the business txn config is removed, this info would be lost. This may be a good motivation to change the way properties are recorded, so rather than map of name/value, could be list of objects with name/value/type enumerator. This does assume that all properties values recorded for the property name have a consistent type though.	0	0	https://issues.jboss.org/browse/HWKAPM-445
6162	HWKAPM-444	Following chat with Andrew Dinn, there are a couple of issues we need to take into account: 1) The keys used to link the application hooks to the rules (at runtime) are uniquely created. Normally as the application instrumentation, and initialisation of the rules, are performed at the same time (using a javaagent when the application is executed), this is not an issue - i.e. they will share the same value. In this case, the hooks are being introduced at build time, and rules initialised at runtime, so the keys could potentially be different. From initial experiments, the keys used appear to be consistent based on rule script name and position of rule in the script, but Andrew warned that if a particular jar is duplicated in multiple wars (for example), then each jar (i.e. its classes) will have a different key. So we will need to ensure that each jar is only presented for processing once. 2) Classloading - as the rule compilation uses the classloader of the instrumented class to associate the helper, this would be a problem initialising the rules are runtime. Therefore at the moment it may only work if the system classloader is being used to run the app - so does swarm just use the system classloader?	0	0	https://issues.jboss.org/browse/HWKAPM-444
6163	HWKAPM-423	Currently an 'application' equates to a business transaction fragment triggered by consuming a message - so the application is effectively known by its URI. A business transaction configuration can be created to 'name' an end to end business transaction comprised of multiple fragments (i.e. application invocations). So we need to decide how best to 'name' an application (i.e. the fragment associated with a particular URI). Business transaction names are configured (along with other info used for pre-processing a business transaction) in the server and pushed out to clients environments being instrumented, so in this case the 'name' is applied as the information is captured. We need to decide whether the same should apply for applications - as we capture a fragment associated with a particular URI, do we name the application (i.e. record the name in the fragment), or do we take a late binding approach and just have the association between URI and application remain in server side configuration? I think it will depend upon what types of queries we may want to perform on the data - whether knowing the URI for an application is enough, or whether actually querying the name (if used for multiple URIs for example) may be of benefit. Once we are able to associate names against URIs, then the new communication graph could leverage those names when displaying the nodes/links. We would also probably want/need a separate application UI, containing the list of applications, some metrics and this new Apdex index.	1	1	https://issues.jboss.org/browse/HWKAPM-423
6164	HWKAPM-423	Just to help understand the aims, when we refer to an "application", are we primarily interested in a micro-service world, and therefore just focusing on things that consumer messages (and potential interact with other 'applications')? If so, generally (although not necessarily) a micro-service would have more than one endpoint. Should the Apdex be applied to the endpoint or to the collection of endpoints associated with the 'application'? The reason for asking is that performance of individual (e.g. REST) endpoints may be considerably different as they are performing different tasks involving varying levels of work, so my assumption is that we want to apply the Apdex at the application/endpoint level?	1	1	https://issues.jboss.org/browse/HWKAPM-423
6165	HWKAPM-386	If internal producer/consumer nodes are created to represent concurrent activities, then need to ensure that the producer/consumer nodes are marked as internal, and have a URI that correctly scopes the node to the outer (non internal) fragment URI, to ensure aggregation tasks correctly distinguish the internal nodes as being part of the outer (fragment). However when visually displaying the information about the internal fragments, it may be useful to strip away the scoping to just represent each fragment by its unique part (e.g. if the internal fragment is scoped to /trades:scheduler, then the initial fragment would be associated with the URI /trades and then it could be linked to a second fragment with name 'scheduler'.	0	0	https://issues.jboss.org/browse/HWKAPM-386
6166	HWKAPM-386	Current approach uses a thread count associated with the fragment - however its possible that the 'spawned' activities may be completed before the next activity is started - so when the next activity starts it may seem like there is no concurrent threads, so will directly record the activity in the parent fragment. Is this an issue? Would it be better to be consistent, and once the activities have started being recorded concurrently, then it should continue? Need to try out using a fork/join example - but possibly as a separate issue.	0	0	https://issues.jboss.org/browse/HWKAPM-386
6167	HWKAPM-269	There will be two types of events - the 'fragment' completion times which represent the durations for the individual services being invoked in a chain, and the "trace" completion times which are for the end to end instance. Both may be required, as the user may be interested in monitoring the individual service times or the overall "business transaction" (i.e. trace instance) time.	0	0	https://issues.jboss.org/browse/HWKAPM-269
6168	HWKAPM-161	This issue may be occurring because, even if inclusion filters are defined, if non-named txns are allowed, then the initial components are still being recorded when it reaches the Producer, which could potentially have named the txn. If non-named txns are not permitted, then the initial components would not have been recorded, and when the Producer is reached, it matches the inclusion filter, so begins recording. So the issue is, how do we overcome the first described scenario? - while at the same time avoiding doing constant checks for the inclusion filters.	0	0	https://issues.jboss.org/browse/HWKAPM-161
6169	HWKAPM-66	Has not been necessary when building end to end view, so closing.	0	0	https://issues.jboss.org/browse/HWKAPM-66
6170	HWKAPM-1	[~partenon] Would you be able to review the PR? Main changes are adding accounts to the REST service, getting the persona and propagating the persona.id as tenant id through all of the backend components - which are all placeholders currently so don't actually do anything with it, and then updating the integration tests to supply the basic auth credentials.	1	1	https://issues.jboss.org/browse/HWKAPM-1

#	Requirement	Questions	Answered	Successfully Answered	URL
6171	HWKINVENT-198	One other very important feature I forgot to mention. We need to be able to ignore roots, even of the same type of entity. For example, a feed could do the following (this is what the hawkular wildfly agent does) - it could independently manage different roots and their children and thus needs to sync different roots at different times. For example, suppose a feed's entire inventory tree looks like this: {code} Feed "foo" |___OperatingSysetm | |___ File System [/opt] | |___ CPU 1 | |___ Memory |___ WildFly Server ABC |___ Deployment hello-world.war |___ Stateless EJB MyEjb |___ Stateless EJB MyOtherEjb |___ Datasource MyDatabase |___ Transaction Manager |___ WildFly Server XYZ |___ Deployment my-app.war |___ Deployment another.war |___ Stateless EJB SomeEjb |___ Transaction Manager {code} We need to support the ability of a feed to ONLY sync SOME of those roots but NOT ALL. In other words, I would like to say, "Here are SOME roots of the tree - sync the roots I give you (including all the children) but DO NOT delete any roots that are missing. You should ignore any missing roots." This allows the feed to sync some roots but not all. In other words, now my feed can send up just information about WildFly Server XYZ but no others: {code} Feed "foo" |___ WildFly Server XYZ |___ Deployment my-app.war |___ Deployment another.war |___ Stateless EJB SomeEjb |___ Transaction Manager {code} Sending that tree as part of the sync should NOT cause WildFly Server ABC nor Operation System roots to be deleted - they and their children should remain as-is. You only sync those resources at or below the WildFly Server XYZ level. What if a root really is missing? The feed would have to sometimes to a full sync (with the entire tree) or the feed or some outside agency (like a user in some UI screen) would have to manually delete roots. It would be up to the feed or a user to know to delete missing roots. This can be done outside of sync if need be.	1	1	https://issues.jboss.org/browse/HWKINVENT-198
6172	HWKINVENT-198	[~lkrejci] BTW: just so I'm clear - today, resource types do NOT have parent-child relationships correct? In other words, all resource types must be peers at the root level. For example, this is not supported: {code} + WildFly Server Type |___ Deployment Type |___ Stateless EJB Type |___ Stateful EJB Type {code} Instead, all types are root types, with no children, even though logically some types are always children of another: {code} + WildFly Server Type + Deployment Type + Stateless EJB Type + Stateful EJB Type {code} Correct?	1	1	https://issues.jboss.org/browse/HWKINVENT-198
6173	HWKINVENT-166	Should this be a strict match on type name/id or do we require substring matching? The latter will require some research into Titan indices and cooperation with Elastic Search, because titan cannot do substring matching on just Cassandra. Currently we start ES locally within the Hawkular JVM because we currently don't really use any Titan features that require ES but should we start doing that, we should externalize ES so that all Hawkular cluster members can reach that. From POV of inventory, this is just a config change (+ some new fancy indices and queries) but from the POV of managing the hawkular cluster, this is a potentially big new component - apart from our C* cluster, we'd also need an ES cluster.	0	0	https://issues.jboss.org/browse/HWKINVENT-166
6174	HWKINVENT-165	I believe this has been sort of done. Am I right [~pavol_loffay]?	1	1	https://issues.jboss.org/browse/HWKINVENT-165
6175	HWKINVENT-165	SwitchElementType is still in base.spi package. Should I move it to api.filters?	0	0	https://issues.jboss.org/browse/HWKINVENT-165
6176	HWKINVENT-123	What is the status of this, [~pavol_loffay]? I think this has been already fixed on the inventory side, but not sure. Could you please try again with the latest inventory if this still is an issue?	1	1	https://issues.jboss.org/browse/HWKINVENT-123
6177	HWKINVENT-63	Can we make this a bit more explicit? This jira is for *operation definitions*. Meaning neither actual parameters nor returned values are stored in this data structure. Actual operation outcome will be discussed in a different jira.	0	0	https://issues.jboss.org/browse/HWKINVENT-63
6178	HWKINVENT-49	What is the preferred solution? (a) Single entity add only (b) Multiple entity add only (c) Both Single entity add and multiple entity add	1	1	https://issues.jboss.org/browse/HWKINVENT-49
6179	HWKINVENT-34	[~mazz], could you comment on the above? I think the app servername, type, hostname, and running state can be or even are already all published by your agent. Could you provide Mike with some instructions on the names of the properties, etc.? Total number of servers I think can be achieved by reading the info from the paged results when searching for all resources of the app server type. The number of servers with given running state is not currently possible as that is search query which has not yet been implemented in inventory.	1	0	https://issues.jboss.org/browse/HWKINVENT-34
6180	HWKINVENT-34	Is this still valid, [~mithomps]?	0	0	https://issues.jboss.org/browse/HWKINVENT-34
6181	HWKINVENT-33	Is there some agreed upon way of doing pagination? Headers? Query params? Both? Response formatted as {code:js} {"data": {...}, "paging": {...}} {code} or something else? I am not too well versed in REST so I'd rather follow some agreed upon way of doing things than invent my own here..	1	1	https://issues.jboss.org/browse/HWKINVENT-33
6182	HWKINVENT-21	no qe activity ATM - will be covered with integration rest tests (end-to-end scenarios)	0	0	https://issues.jboss.org/browse/HWKINVENT-21
6183	HWKINVENT-7	Why have you closed this as having no QE activity?	1	0	https://issues.jboss.org/browse/HWKINVENT-7
6184	HWKMETRICS-693	I have a question about the last example. Doing the {{a1 = 'defg'}} query first makes sense. I believe it hits metrics_tags_idx. The results we get back from metrics_tags_idx do not have the full set of tags for each metric. We still need a query against metrics_idx for the {{nost hostname}} query. We can avoid avoid a full table scan though. While using the CQL clause is generally discouraged, it is probably ok when querying against a single partition as would be the case here. The query would be something like: {noformat} select ... from metrics_idx where tenant_id = $tenantId and type = $type = and metric in (/* metric names from the a1 = 'defg' result set */) {noformat} Is my analysis here correct?	1	1	https://issues.jboss.org/browse/HWKMETRICS-693
6185	HWKMETRICS-682	[~rvargasp] what is the status on this? We have another meeting with the OpenShift ops team on Wednesday, and I am hoping that we will have something to share and/or discuss. Don't hesitate to ping me if you have questions. Let me you have other tasks you are working on, and I can take over the ticket if need be. Thanks.	0	0	https://issues.jboss.org/browse/HWKMETRICS-682
6186	HWKMETRICS-680	Can you let me know what the final prometheus.yaml should look like? We need to trim this down to only the metrics we are interested in. Do you also know if we can secure this endpoint using the prometheus java agent?	1	1	https://issues.jboss.org/browse/HWKMETRICS-680
6187	HWKMETRICS-668	Ops would like to consume the metrics with Zabbix, are the metrics exposed in a format that can be pulled by them ? We will need a doc to explain what metrics are available and what they mean (what to look for to identify an issue)	1	1	https://issues.jboss.org/browse/HWKMETRICS-668
6188	HWKMETRICS-668	Since we have decided to store the metrics in the same Cassandra cluster in which we store user metrics, they will be accessible using Hawkular Metrics REST API. They will be stored under the admin tenant though which means client requests will also need to supply the admin token in addition to the admin tenant id. Is the REST API suficient? Should we consider for example building something on top of the [Go client|https://github.com/hawkular/hawkular-client-go], or something that could integrate more directly with Zabbix? Whatever metrics we expose will certainly be documented. Right now I am still working through some of the implementation details as well as trying to figure out what all to expose.	1	1	https://issues.jboss.org/browse/HWKMETRICS-668

#	Requirement	Questions	Answered	Successfully Answered	URL
6189	HWKMETRICS-668	There is not a 1-to-1 mapping between metric types in Hawkular Metrics and in DropWizard. Hawkular Metrics only supports a subset of those in DropWizard. I want to point out how things from DropWizard get mapped in Hawkular Metrics *Counter* Counters map directly into Hawkular Metrics. *Gauge* DropWizard gauges can have anything as a value. Right now we are only using doubles which map directly to Hawkular Metrics. We could also support strings. *Meter* Suppose we have a DropWizard meter named MyMeter. This will map to the following in Hawkular Metrics: || Generated Metric || Description || | MyMeter-1min | 1 minute moving average, stored as gauge | | MyMeter-5min | 5 minute moving average, stored as gauge | | MyMeter-15min | 15 minute moving average, stored as gauge | | MyMeter-mean | average since start, stored as gauge | | MyMeter | count, stored as counter | *Timer* Suppose we have a timer named MyTimer. This will map to the following in Hawkular Metrics: || Generated Metric || Description || | MyTimer-1min | 1 minute moving average, stored as gauge | | MyTimer-5min | 5 minute moving average, stored as gauge | | MyTimer-15min | 15 minute moving average, stored as gauge | | MyTimer-mean | average of timed values, stored as gauge | | MyTimer | count, stored as counter | | MyTimer-median | median of timed values, stored as gauge | | MyTimer-min | min of timed values, stored as gauge | | MyTimer-max | max of timed values, stored as gauge | | MyTimer-stdDev | standard deviation of timed values, stored as gauge | | MyTimer-75p | 75th percentile of timed values, stored as gauge | | MyTimer-95p | 95th percentile of timed values, stored as gauge | | MyTimer-98p | 98th percentile of timed values, stored as gauge | | MyTimer-99p | 99th percentile of timed values, stored as gauge | | MyTimer-999p | 99.9th percentile of timed values, stored as gauge | *Histogram* Suppose we have a histogram named MyHistogram. This will map to the following in Hawkular Metrics: || Generated Metric || Description || | MyHistogram-mean | average of timed values, stored as gauge | | MyHistogram | count, stored as counter | | MyHistogram-median | median of timed values, stored as gauge | | MyHistogram-min | min of timed values, stored as gauge | | MyHistogram-max | max of timed values, stored as gauge | | MyHistogram-stdDev | standard deviation of timed values, stored as gauge | | MyHistogram-75p | 75th percentile of timed values, stored as gauge | | MyHistogram-95p | 95th percentile of timed values, stored as gauge | | MyHistogram-98p | 98th percentile of timed values, stored as gauge | | MyHistogram-99p | 99th percentile of timed values, stored as gauge | | MyHistogram-999p | 99.9th percentile of timed values, stored as gauge | A timer maps to 14 different time series. A meter maps to 5 different time series. This will add up quickly. Cassandra also uses DropWizard Metrics; so, we face a similar situation for collecting/store Cassandra metrics. Do we want to store all of these metrics. If not, which ones do we store?	1	1	https://issues.jboss.org/browse/HWKMETRICS-668
6190	HWKMETRICS-666	It sounds reasonable. What kinds of endpoints would you be adding?	1	0	https://issues.jboss.org/browse/HWKMETRICS-666
6191	HWKMETRICS-650	Is there a productization requirement to update to older versions and newest bugfixed version wasn't going to be used in 0.26.x?	1	1	https://issues.jboss.org/browse/HWKMETRICS-650
6192	HWKMETRICS-650	Is this still relevant? master is at 1.3.0	1	1	https://issues.jboss.org/browse/HWKMETRICS-650
6193	HWKMETRICS-628	[~nstefan] can you provide some more details on what you have in mind and the use case(s) that you are trying to address? What are you proposing that cannot already be accomplished with tags, be it metric or data point tags? Granted data point tags only support string values currently; however, we easily extend that to support arbitrary values.	1	0	https://issues.jboss.org/browse/HWKMETRICS-628
6194	HWKMETRICS-622	Can you please create a ticket and/or PR for origin-metrics to upgrade to 3.0.12 and link it to this ticket?	1	1	https://issues.jboss.org/browse/HWKMETRICS-622
6195	HWKMETRICS-614	What about out of order writes? Someone writes something to the past during the next day for example.	1	1	https://issues.jboss.org/browse/HWKMETRICS-614
6196	HWKMETRICS-614	{quote}What about out of order writes?{quote} [~burmanm] Not sure yet. How do we currently handle this scenario with the compression job? And once the compression block has passed, any data received for that block would be considered out of order right?	1	1	https://issues.jboss.org/browse/HWKMETRICS-614
6197	HWKMETRICS-597	I want to provide some additional background since my initial description is rather short on details. I am not sure that reactive backpressure is what we need here. As I understand backpressure is a mechanism for dealing with situations where an observable is emitting items faster than they can be consumed downstream. A good example for backpressure would be we execute a query that returns a big result set, too big for available memory. We are then performing long running tasks on rows we get back. Rather than the driver return all rows at once, we would request rows in chunks/pages via reactive back pressure. The situation this ticket aims to address is a bit different though. We want to deal with the situation when we try to execute a query and the C* driver does not have sufficient resources to execute it. It is not the downstream operator/subscribe that needs to slow down but rather the upstream observable. The first question we need to answer is, why do we need to do this. Consider the following example. We execute our queries using asynchronously via {{Session.executeAsync}}. Suppose we have a burst of writes as is the case with ingesting metrics. We can very easily fire off enough queries asynchronously such that we wind up with queries failing with the infamous NoHostAvailableException. Let's say that if we try to execute 1,000 writes concurrently we wind up with a NoHostAvailableException on the 900th write. If we can determine that we need to slow down a bit when we get to the 900th write, we should be able to execute all queries without error. Assuming we can detect when the C* driver is at capacity what should we do? We could simply reject the request with a better exception that is more descriptive than NoHostAvailable (which can mean different things). Alternatively we can queue up requests until driver resources become available. Even with the queue approach we might still have to reject requests if the queue gets full (assuming we use a bounded queue which I think we want to do). The queue approach however should allow us to avoid failures in some of the temporary spike scenarios. We ought to use a priority queue. In general, user facing, i.e., real time, requests should take higher priority than requests being executed by a background job that is not as sensitive too per request latency spikes. Possibly the most important thing we need to do is log when we hit or exceed capacity. By log, I mean record the event. We should write something meaningful to the log file and also store metrics. We are already doing some of the latter with HWMETRICS-668. {{CassandraDriverMetrics.java}} captures some load metrics. The metric names look like {{localhost|org.hawkular.metrics|Load_127.0.0.1:9042}} and {{localhost|org.hawkular.metrics|MaxLoad_127.0.0.1:9042}}. We should be able to use the Load and MaxLoad metrics (and possibly) others when the driver is at or over capacity. And if those metrics are getting persisted, we can look at past data to determine whether or not changes are needed. I covered a lot of detail because there is a lot of background. I am in no way suggesting we tackle everything in this ticket.	0	0	https://issues.jboss.org/browse/HWKMETRICS-597
6198	HWKMETRICS-597	Can we break this up into more manageable smaller pieces? I think the first thing to do is to figure out how to monitor Cassandra to know if its under too much pressure or not. And then make sure we are logging and collecting metrics for this. I think after that we can then determine how we want to handle things based on this knowledge. We may also want a priority queue for writes over reads	1	1	https://issues.jboss.org/browse/HWKMETRICS-597
6199	HWKMETRICS-563	Is this functionality already covered by HWKMETRICS-373?	1	0	https://issues.jboss.org/browse/HWKMETRICS-563
6200	HWKMETRICS-563	Hi John, I've tried with both counters and gauges, still the same, 204. I've looked, but can't see any endpoints which return datapoints with their tags, only metric definitions with their tags. Does this endpoint exist? i want to try this to prove that the tags should match.	1	1	https://issues.jboss.org/browse/HWKMETRICS-563

#	Requirement	Questions	Answered	Successfully Answered	URL
6201	HWKMETRICS-563	Notice in my previous example that I used {{method:*}} which essentially means that a separate bucket will be created for each different value of the {{method}} tag that is found. If you only wanted buckets for {{DELETE}} and {{PUT}}, you would instead use: {{method:DELETE|PUT}} {noformat} curl -X GET http://localhost:8080/hawkular/metrics/counters/my_counter/stats/tags/method:DELETE%7CPUT -H "Content-Type: application/json" -H "Hawkular-Tenant: com.acme" {noformat} and get back: {code:javascript} { "method:PUT": { "tags": { "method": "PUT" }, "min": 18, "avg": 18, "median": 18, "max": 18, "sum": 18, "samples": 1 }, "method:DELETE": { "tags": { "method": "DELETE" }, "min": 10, "avg": 10, "median": 10, "max": 10, "sum": 10, "samples": 1 } } {code} You can also group by multiple tags. I hope this helps.	1	0	https://issues.jboss.org/browse/HWKMETRICS-563
6202	HWKMETRICS-561	[~mwringe] do you think we might see additional gains from batching across multiple tenants? I realize that doing so goes against the concept of multi-tenancy to a degree, but if we can achieve better throughput then I think it is worth considering. Maybe the endpoint could require a special header. [~nstefan] what do you think?	1	0	https://issues.jboss.org/browse/HWKMETRICS-561
6203	HWKMETRICS-561	Wouldn't HTTP/2 help in this ? A single HTTP connection with multiple inserts.	0	0	https://issues.jboss.org/browse/HWKMETRICS-561
6204	HWKMETRICS-537	Is is already done and can be closed, right?	1	1	https://issues.jboss.org/browse/HWKMETRICS-537
6205	HWKMETRICS-525	This kinda defeats the whole purpose of tenancy if you can avoid it. Then it's no longer a safeguard, without authentication/authroization at least. Why does ManageIQ need to fetch multiple projects at once?	1	1	https://issues.jboss.org/browse/HWKMETRICS-525
6206	HWKMETRICS-525	Oh, sorry, I created this JIRA on metrics instead of Alerting. "Why does ManageIQ need to fetch multiple projects at once?" Because ManageIQ wants to fetch alerting from all the projects (potentially high) at single step. We are proposing alternatives to this, as for alerting, multi-tenancy is a strong partition too. P.S.: I am going to move this JIRA to Alerting. It was a mistake from my side to create it here.	0	0	https://issues.jboss.org/browse/HWKMETRICS-525
6207	HWKMETRICS-523	I just read through the [kubernetes labels doc|http://kubernetes.io/docs/user-guide/labels/] which provides a lot of context. Labels in kubernetes according to the aforementioned doc are very similar if not the same as our key/value tags. All labels are stored under a single tag named {{label}}. Tags are a formal concept in hawkular-metrics. Labels are not, and there is work being done in this ticket to implement informal semantics for labels. This concerns me a bit as it could be a source of problems down the road. Suppose other clients decide to store a {{label}} tag that is a simple key/value pair. Will that break anything? Should we introduce labels as formal concept to avoid ambiguity? And/or can we simply flatten a label storing each of the key/value pairs separately and possibly enabling us to use our existing tag filtering capabilities more easily? I have not fully thought this through, but consider this idea. Suppose I want to query for all metrics tagged with 'type=pod_container' and 'descriptor_name=memory/usage' and where the labels match 'deployment=frontend' and 'deploymentconfig=frontend' and 'name=frontend'. As things our implemented today, we will have a partition in metrics_tags_idx with label (as part) of the partition key. Why not flatten the labels storing each one in separate partition where the tname column will have a prefix of something like {{label#}} followed by the label name. I need to go through some of the query use cases to see if/how this might work, but I think it is worth considering.	1	1	https://issues.jboss.org/browse/HWKMETRICS-523
6208	HWKMETRICS-500	How does job scheduler allow providing back pressure to prevent next run starting before previous one has finished (although I don't think this is what you meant, but maybe that should be tackled at the same time)	1	1	https://issues.jboss.org/browse/HWKMETRICS-500
6209	HWKMETRICS-482	{quote}Doing so would finally allow querying by tag in a restful manner, instead of the current POST endpoint on ".../raw/query".{quote} Could this replace the POST endpoint?	1	1	https://issues.jboss.org/browse/HWKMETRICS-482
6210	HWKMETRICS-480	What use case could that possibly be? The only category of the query is that it does not have a certain tag? That could return tens of thousands or hundreds of thousands of metric ids.	1	1	https://issues.jboss.org/browse/HWKMETRICS-480
6211	HWKMETRICS-447	Hawkular Metrics. Since I was running against a 0.8.2 schema, this scenario simulates trying to apply schema changes while {{upgradesstables}} runs. Interestingly the tests ran without error, but {{upgradesstables}} aborted with the following error: {noformat} error: Unknown column n_value in table hawkulartest.data -- StackTrace -- java.lang.AssertionError: Unknown column n_value in table hawkulartest.data at org.apache.cassandra.db.LegacyLayout.readLegacyAtom(LegacyLayout.java:1013) at org.apache.cassandra.db.UnfilteredDeserializer$OldFormatDeserializer$AtomIterator.readAtom(UnfilteredDeserializer.java:518) at org.apache.cassandra.db.UnfilteredDeserializer$OldFormatDeserializer$AtomIterator.hasNext(UnfilteredDeserializer.java:504) at org.apache.cassandra.db.UnfilteredDeserializer$OldFormatDeserializer$UnfilteredIterator.readRow(UnfilteredDeserializer.java:447) at org.apache.cassandra.db.UnfilteredDeserializer$OldFormatDeserializer$UnfilteredIterator.hasNext(UnfilteredDeserializer.java:423) at org.apache.cassandra.db.UnfilteredDeserializer$OldFormatDeserializer.hasNext(UnfilteredDeserializer.java:289) at org.apache.cassandra.io.sstable.SSTableSimpleIterator$OldFormatIterator.computeNext(SSTableSimpleIterator.java:145) at org.apache.cassandra.io.sstable.SSTableSimpleIterator$OldFormatIterator.computeNext(SSTableSimpleIterator.java:97) at org.apache.cassandra.utils.AbstractIterator.hasNext(AbstractIterator.java:47) at org.apache.cassandra.io.sstable.SSTableIdentityIterator.computeNext(SSTableIdentityIterator.java:100) at org.apache.cassandra.io.sstable.SSTableIdentityIterator.computeNext(SSTableIdentityIterator.java:30) at org.apache.cassandra.utils.AbstractIterator.hasNext(AbstractIterator.java:47) at org.apache.cassandra.db.rows.LazilyInitializedUnfilteredRowIterator.computeNext(LazilyInitializedUnfilteredRowIterator.java:100) at org.apache.cassandra.db.rows.LazilyInitializedUnfilteredRowIterator.computeNext(LazilyInitializedUnfilteredRowIterator.java:32) at org.apache.cassandra.utils.AbstractIterator.hasNext(AbstractIterator.java:47) at org.apache.cassandra.utils.MergeIterator$TrivialOneToOne.computeNext(MergeIterator.java:484) at org.apache.cassandra.utils.AbstractIterator.hasNext(AbstractIterator.java:47) at org.apache.cassandra.db.rows.UnfilteredRowIterators$UnfilteredRowMergeIterator.computeNext(UnfilteredRowIterators.java:419) at org.apache.cassandra.db.rows.UnfilteredRowIterators$UnfilteredRowMergeIterator.computeNext(UnfilteredRowIterators.java:279) at org.apache.cassandra.utils.AbstractIterator.hasNext(AbstractIterator.java:47) at org.apache.cassandra.db.transform.BaseRows.hasNext(BaseRows.java:112) at org.apache.cassandra.db.ColumnIndex.buildRowIndex(ColumnIndex.java:93) at org.apache.cassandra.io.sstable.format.big.BigTableWriter.append(BigTableWriter.java:161) at org.apache.cassandra.io.sstable.SSTableRewriter.append(SSTableRewriter.java:135) at org.apache.cassandra.db.compaction.writers.DefaultCompactionWriter.realAppend(DefaultCompactionWriter.java:65) at org.apache.cassandra.db.compaction.writers.CompactionAwareWriter.append(CompactionAwareWriter.java:141) at	0	0	https://issues.jboss.org/browse/HWKMETRICS-447

#	Requirement	Questions	Answered	Successfully Answered	URL
6212	HWKMETRICS-418	Considering most of our deployments are likely going to be on Linux (I of course am the odd ball here working on a Mac) I think this makes sense. Does this mean we will have to have a separate distro that includes this, or we can include it without harm on all platforms?	1	1	https://issues.jboss.org/browse/HWKMETRICS-418
6213	HWKMETRICS-393	[~tsegismont] what do you have in mind for the metric id parameters?	1	1	https://issues.jboss.org/browse/HWKMETRICS-393
6214	HWKMETRICS-375	Can we hold off on this? I am finally starting on HWKMETRICS-360. There is some overlap and it looks like it'll probably be easier for me to remove this stuff during that work. I am going to be reusing some of the classes.	1	1	https://issues.jboss.org/browse/HWKMETRICS-375
6215	HWKMETRICS-369	It's not clear to me yet, would we create a new metric type, like "Increment"? Or would we use the existing "Counter" type?	1	0	https://issues.jboss.org/browse/HWKMETRICS-369
6216	HWKMETRICS-369	How would the server know if the counter to be created/updated is an increment or an absolute number?	1	0	https://issues.jboss.org/browse/HWKMETRICS-369
6217	HWKMETRICS-369	If we use the existing counter type, then we need to store some meta data indicating that the metric is an incrementing counter. We would need to do this on each write since we do not require metrics to be explicitly created before storing data for them. We can probably store this in metrics_idx and possibly as a system-defined tag. Furthermore, it would be the client's responsibility to ensure that it does not mix update types. I do not think this is a big issue, but it is something about which users will need to be aware. When fetching data points, do we return the raw data points which are the increments? I am inclined to think that users are typically interested in the total count and rate. If we want to return counts, we could compute and return per-minute counts by default. Rate queries would then also use those per-minute counts. We will need to do this additional work regardless of whether or not we store incrementing counters as a new, distinct metric type. If we treat it as an implicit type using the existing counter type, then we will have to do a little extra work at query time to determine that the metric is an incrementing counter. The endpoint that supports adding data point for mixed metric types will need to be modified. If we go with an implicit type, then something like, {code:javascript} { "gauages": { ... }, "availabilities": { ... }, "counters": { "set": { ... }, "increment": { ... } } } {code} versus {code:javascript} { "gauages": { ... }, "availabilities": { ... }, "counters": { ... } "increment-counters": { ... } // or whatever we would choose to call them } {code}	0	0	https://issues.jboss.org/browse/HWKMETRICS-369
6218	HWKMETRICS-369	Incrementing counters by definition are not idempotent. Using PUT is probably a bad idea since it is supposed to be idempotent. Since we already use POST for our current counters, I think the best course of action is to introduce a new metric type. In lieu of a better name, I am going to go with inc-counter. I am still undecided as to whether we should store the raw increment values. The upside is that we can provide the count at any given point in time. The downside is that this could result in a lot of additional storage. A middle ground solution might be something like, {noformat} CREATE TABLE counters (tenantId text, id text, bucket timestamp, count counter, PRIMARY KEY ((tenantId, id), bucket)); {noformat} The bucket column is a timestamp that by default could rounded down to the minute. This would keep the number of data points fixed and still allow a good level of granularity for finding the count at any point in time. There are a couple of other considerations. HWKMETRICS-368 added support for per-data point tagging. Outside of the primary key, counter tables cannot have non-counter columns. This is an unfortunate limitation that make Cassandra's native counters less than ideal. Of course per-data tagging might be a bit different with per-minute counts. For example, let's say we receive three data points for the same minute, each having some tags. How do we handle the tags? We could merge them together. If the data points include tags with the same tag key though, what should we do? Presumably we do not want any of the tag values to be overwritten. One solution could be to include the timestamp in the tag key. We probably want to use a time UUID to avoid timestamp collisions. Now back to how/where counters are stored. If we want to support data point tagging (and I think we do), then Cassandra's native counters are not an option. And storing the raw increments concerns me due to the potentially large increase in storage. We could use an in-memory distributed cache for maintaining the per-minute counts. The cache should effectively be write through so that any time we update a counter we have the counter value in long term storage.	0	0	https://issues.jboss.org/browse/HWKMETRICS-369
6219	HWKMETRICS-368	One thing I have not considered is what we want to do with the tags at query time other than HWKMETRICS-373. If the client is querying for raw data points, then it certainly makes sense to include any tags. What do we do though with bucketed data points and with counter rate data points? A bucketed data point is made up of one or more data points. Maybe we don't need to do anything. If we need to do something, I am inclined to think we can do it as a follow up ticket.	0	0	https://issues.jboss.org/browse/HWKMETRICS-368

#	Requirement	Questions	Answered	Successfully Answered	URL
6220	HWKMETRICS-361	be summarized as follows, * Enable us to perform schema changes incrementally. Currently we have to reset the schema to apply new changes. ** Incremental changes is key for application upgrades * Version the schema changes so that we can easily determine the state of a schema * Verify the schema version by checking that changes are applied in the expected order * Support multiple upgrade paths Everything except the last item is more or less already handled by Cassalog. Supporting multiple upgrade paths and also verifying that changes are applied in the correct order gets tricky, and that is what I want to discuss. I need to present some examples to illustrate. Suppose the schema in master 1, 2, 3 (these are the version numbers assigned to the changes). We then release v1.0 of our application, which means we now have a v1.x branch. We then make schema change 4 in master. The next release from master will be v2.0. Now we need a schema change in v1.0 that will necessitate a v1.1, which will be released out of the v1.x branch. The schema change is also needed for v2.0, so we make schema change 5 in master and back port it to the v1.x branch. *Scenario* For a user installing v2.0, the schema will be 1, 2, 3, 4, 5. For a user upgrading from v1.0 to v1.1, the schema will be 1, 2, 3, 5. Everything is fine so far. Cassalog applies changes in the order that they are declared, and since schema change 5 is back ported to the v1.x branch, it is expected to come after change 3. Now let's say we want upgrade from v1.1 to v2.0. The expected order of changes for v2.0 is 1, 2, 3, 4, 5, but coming from v1.1 the order is 1, 2, 3, 5. Cassalog will complain about this. We can adopt a convention with tags along with some changes to Cassalog to handle this scenario. Each schema change will be tagged with the application version branch in which it is being applied. Changes 1, 2, and 3 will be tagged with v1.x. Change 4 will be tagged with v2.x. Change 5 will be tagged with both v1.x and v2.x. New installations will run Cassalog without specifying any tags which should result in all changes being applied in the order declared. For upgrades Hawkular Metrics specifies some version tags when running Cassalog. Cassalog will only look at and execute the changes with the specified tags. When upgrading from v1.0 to v1.1, we specify the v1.x tag to indicate that we want any new changes on the v1.x branch to be applied. When we upgrade to v2.0 from v1.x, we first run Cassalog with the v1.x tag to make sure we have the latest v1.x schema. Cassalog will find that changes 1, 2, 3, 5 are tagged with v1.x. This matches are schema so we are good to go. Then we run it with the v2.x tag. Only change 4 has the v2.x tag, so only it is applied, which is what we expect. There is a question of how the tags that get passed to Cassalog are specified. We do not want to impose that on the user as it could be error prone. It should be done programmatically within Hawkular Metrics. See the section at the end on determining tags to pass to Cassalog. *Scenario* Let's extend the previous scenario by adding a v1.2 which requires a schema changes that also should go into v2.0. This will be change 6. I am currently on v1.1. I want to upgrade directly to v2.0, skipping v1.2. Recall that our v1.1 schema is 1, 2, 3, 5. In v2.0 we now have 1, 2, 3, 4, 5, 6. Changes 5 and 6 are tagged with v1.x and v2.x. The schema verification is still good to go. When applying updates, changes 4 and 6 are applied since they are both tagged with v2.0. We end up with changes 1, 2, 3, 5, 4, 6. If we had upgraded from v1.1 to v1.2 to v2.0 our schema would be 1, 2, 3, 5, 6, 4. I think that the difference in order here is ok as long as both 4 and 6 come after 5. *Scenario* We have v1.0 already released with schema changes 1, 2, 3. The we have v2.0 already released with 1, 2, 3, 4, 5. The next release in master will be v3.0. We are providing patch release support out of both the v1.x and v2.x branches. Schema changes 6 and 7 will go into v3.0. We have a bug in v1.0 that requires a v1.1 release and requires schema change 8. We need the change in master and in the v1.x branch, but we do not need it in the v2.x branch. Since we do not need the change in v2.x, we want to avoid forcing a v2.x release since putting out patch releases is costly. If we do not put the change in the v2.x branch, then what does this mean? In master/v3.0 we will have 1, 2, 3, 4, 5, 6, 7, 8. In v1.1 we will have 1, 2, 3, 8. In v2.0 we have 1, 2, 3, 4, 5. Changes 6 and 7 will be tagged with v3.x.	0	0	https://issues.jboss.org/browse/HWKMETRICS-361
6221	HWKMETRICS-361	I think I follow although the tag terminology is a little confusing. Combining scenarios 1 and 2, and expressing the tags as a map, I can envision mappings like this: { 2.0: { schemaBase: 6, schemaPatches: nil }, 1.2: { schemaBase: 3, schemaPatches: [5,6] }, 1.1: { schemaBase: 3, schemaPatches: [5] }, 1.0: { schemaBase: 3, schemaPatches: nil } } To get from one version to another you just take the increments between schemaBase versions, add patches for the new version, and remove patches for the old version. So, from V1.0 to V2.0 you must increment from baseVersion 3 to 6, so [4,5,6]. Add V2.0 patches []. Remove V1.0 patches []. So [4,5,6] + [] - [] = [4,5,6]. From V1.1 to V1.2 you must increment from baseVersion 3 to 3, so []. Add V1.2 patches [5,6]. Remove V1.1 patches [5]. so [] + [5,6] - [5] = [6]. So, from V1.1 to V2.0 you must increment from baseVersion 3 to 6, so [4,5,6]. Add V2.0 patches []. Remove V1.0 patches [5]. So [4,5,6] + [] - [5] = [4,6]. Overall, I think this could work and in map form like this the upgrade matrix is fairly clear. Of course there is danger in that the matrix must be updated as you go, but without a lot of patches it would be fairly straightforward. You've mentioned only schema changes I think. Are you considering also that an upgrade step could apply some sort of data change? Also, just to be clear, I assume that a new install would simply get a new schema and not have to have any upgrade steps applied.	1	1	https://issues.jboss.org/browse/HWKMETRICS-361
6222	HWKMETRICS-361	{quote} think I follow although the tag terminology is a little confusing. Combining scenarios 1 and 2, and expressing the tags as a map... {quote} What is confusing about the tag terminology. The tags are simply labels. Don't confuse them with the key/value pair tags that exist in Hawkular Metrics or in Hawkular Alerts. Some of the concepts here, like tagging, are adopted from how things are done in Liquibase. And by "here", I am referring to Cassalog, not Hawkular Metrics. The use of tags as I have previously described allows us to support non-linear upgrade paths and still verify that changes have been applied in the expected order. The first scenario is a pretty typical upgrade scenario. We are on the latest v1.x release and want to upgrade v2.0. In the second scenario we are *not* on the latest v1.x release and want to upgrade to v2.0. {quote} You've mentioned only schema changes I think. Are you considering also that an upgrade step could apply some sort of data change? {quote} Yes, we have considered data changes. And more importantly changes that we might prefer to do at runtime vs installation time such as data migrations that could take an arbitrary amount of time. I will try to provide an example of a scenario that involves making changes post-installation a little later today when I have some more time. {quote} Also, just to be clear, I assume that a new install would simply get a new schema and not have to have any upgrade steps applied. {quote} There aren't really separate steps for new installations vs upgrades. The version from which we upgrade is important so that we can support non-linear upgrades. Since there is no previous version for a new install, we apply changes in linear order.	1	1	https://issues.jboss.org/browse/HWKMETRICS-361
6223	HWKMETRICS-361	> What is confusing about the tag terminology Probably just me reading it. Represented it in the map above was just my way of representing the tag idea in a format that was really clear to me. The use cases fell out pretty easily after that. > There aren't really separate steps for new installations vs upgrades. Are you saying that a new install will actually start with schema V1 and apply every upgrade step? And not just start with/apply the most current full schema definition?	1	1	https://issues.jboss.org/browse/HWKMETRICS-361
6224	HWKMETRICS-361	{quote} Are you saying that a new install will actually start with schema V1 and apply every upgrade step? And not just start with/apply the most current full schema definition? {quote} That is correct. New installations will get the latest schema, but the schema is always applied incrementally so that we can version the the changes and support various upgrade paths.	0	0	https://issues.jboss.org/browse/HWKMETRICS-361

#	Requirement	Questions	Answered	Successfully Answered	URL
6225	HWKMETRICS-358	For the most part the time on the Hawkular Metrics server doesn't really matter too much as its the client's sending the timestamps. But, we do use the Hawkular Metrics' time when handling things like when the start or end time are not specified and we resort to using defaults. So there is some precedent here in using the servers clock. Would it be possible to expose the current timestamp on the status page? I think we should also add the option to specify the 'start' or 'end' time relative the server time by having this value start with a '+' or '-'. That way we don't even need to introduce new parameters, we can just use the current ones. Eg: start=-30000 -> the start time is 30 seconds behind the current server's clock time start=1456843111234 -> the UTC time like normal Any thoughts?	1	1	https://issues.jboss.org/browse/HWKMETRICS-358
6226	HWKMETRICS-342	Do we need to be concerned with accounts integration any more?	1	1	https://issues.jboss.org/browse/HWKMETRICS-342
6227	HWKMETRICS-335	Can't this be done with two requests using the limit and order params introduced in HWKMETRICS-334? Let's say my start time is x and my end time is y. I could do one request with start=x&limit=1&order=desc to give me the closest data point at or before x. My second request could be with start=y&limit=1&order=asc to give me that cloest data point at or after y. [~mwringe] is this what you want to do?	0	0	https://issues.jboss.org/browse/HWKMETRICS-335
6228	HWKMETRICS-334	[~nstefan] I noticed you started working on this. Could you update the REST documentation as part of this PR ({{api/metrics-api-jaxrs/src/main/rest-doc/base.adoc}}) ? Thanks!	1	1	https://issues.jboss.org/browse/HWKMETRICS-334
6229	HWKMETRICS-330	verified asynchronously. It seems it won't be straightforward as RestEasy doesn't like you to start the {{AsyncContext}}. Here's what I get when running {{curl --basic -u thomas:thomas -H "Accept: application/json" -H "Hawkular-Tenant: default" http://127.0.0.1:8080/hawkular/metrics/gauges}}: {noformat} 15:53:00,309 TRACE [org.hawkular.metrics.api.jaxrs.exception.mappers.ExceptionMapperUtils] (default task-45) Turning org.jboss.resteasy.spi.BadRequestException into a 400 response: org.jboss.resteasy.spi.BadRequestException: RESTEASY003320: Failed processing arguments of org.jboss.resteasy.spi.metadata.ResourceMethod@1d932ca4 at org.jboss.resteasy.core.MethodInjectorImpl.injectArguments(MethodInjectorImpl.java:106) at org.jboss.resteasy.core.MethodInjectorImpl.invoke(MethodInjectorImpl.java:114) at org.jboss.resteasy.core.ResourceMethodInvoker.invokeOnTarget(ResourceMethodInvoker.java:295) at org.jboss.resteasy.core.ResourceMethodInvoker.invoke(ResourceMethodInvoker.java:249) at org.jboss.resteasy.core.ResourceMethodInvoker.invoke(ResourceMethodInvoker.java:236) at org.jboss.resteasy.core.SynchronousDispatcher.invoke(SynchronousDispatcher.java:395) at org.jboss.resteasy.core.SynchronousDispatcher.invoke(SynchronousDispatcher.java:202) at org.jboss.resteasy.plugins.server.servlet.ServletContainerDispatcher.service(ServletContainerDispatcher.java:221) at org.jboss.resteasy.plugins.server.servlet.HttpServletDispatcher.service(HttpServletDispatcher.java:56) at org.jboss.resteasy.plugins.server.servlet.HttpServletDispatcher.service(HttpServletDispatcher.java:51) at javax.servlet.http.HttpServlet.service(HttpServlet.java:790) at io.undertow.servlet.handlers.ServletHandler.handleRequest(ServletHandler.java:85) at io.undertow.servlet.handlers.FilterHandler$FilterChainImpl.doFilter(FilterHandler.java:129) at org.hawkular.openshift.auth.BasicAuthentication.doFilter(BasicAuthentication.java:92) at org.hawkular.openshift.auth.OpenShiftAuthenticationFilter.doFilter(OpenShiftAuthenticationFilter.java:119) at io.undertow.servlet.core.ManagedFilter.doFilter(ManagedFilter.java:60) at io.undertow.servlet.handlers.FilterHandler$FilterChainImpl.doFilter(FilterHandler.java:131) at io.undertow.servlet.handlers.FilterHandler.handleRequest(FilterHandler.java:84) at io.undertow.servlet.handlers.security.ServletSecurityRoleHandler.handleRequest(ServletSecurityRoleHandler.java:62) at io.undertow.servlet.handlers.ServletDispatchingHandler.handleRequest(ServletDispatchingHandler.java:36) at org.wildfly.extension.undertow.security.SecurityContextAssociationHandler.handleRequest(SecurityContextAssociationHandler.java:78) at io.undertow.server.handlers.PredicateHandler.handleRequest(PredicateHandler.java:43) at io.undertow.servlet.handlers.security.SSLInformationAssociationHandler.handleRequest(SSLInformationAssociationHandler.java:131) at io.undertow.servlet.handlers.security.ServletAuthenticationCallHandler.handleRequest(ServletAuthenticationCallHandler.java:57) at io.undertow.server.handlers.PredicateHandler.handleRequest(PredicateHandler.java:43) at	1	1	https://issues.jboss.org/browse/HWKMETRICS-330
6230	HWKMETRICS-324	Can you explain where the issue is? We expect users to send data points with timestamps as milliseconds since January 1, 1970, 00:00:00 UTC. See http://www.hawkular.org/docs/rest/rest-metrics.html#_timestamps Where is the value messed up?	0	0	https://issues.jboss.org/browse/HWKMETRICS-324
6231	HWKMETRICS-319	Could you point me to the report where the performance drop was noticed? I am aware of two pull requests related to bus integration: - https://github.com/hawkular/hawkular-metrics/pull/393 - https://github.com/hawkular/hawkular-metrics/pull/398 The corresponding perf job results: - http://209.132.178.156/report394.html - http://209.132.178.156/report416.html The reports linked above do not show any regression or big drop.	0	0	https://issues.jboss.org/browse/HWKMETRICS-319
6232	HWKMETRICS-300	[~tsegismont] and [~john.sanda] : Sorry for the late response, but yes I am positive to look into this issue. Thanks for the hint :-)	0	0	https://issues.jboss.org/browse/HWKMETRICS-300
6233	HWKMETRICS-297	Unless this is a requirement for the OpenShift work, can we move the fix version to 0.9.0?	0	0	https://issues.jboss.org/browse/HWKMETRICS-297
6234	HWKMETRICS-291	Unless this is a requirement for the OpenShift work, can we move the fix version to 0.9.0?	1	0	https://issues.jboss.org/browse/HWKMETRICS-291
6235	HWKMETRICS-290	Is there anything we are planning to get done for this for 0.8.0? Or can this be moved to 0.9.0 instead?	1	1	https://issues.jboss.org/browse/HWKMETRICS-290
6236	HWKMETRICS-263	What would be a good address for this functionality? {{GET /hawkular/metrics/gauges/tags/<tags>/data}}?	1	1	https://issues.jboss.org/browse/HWKMETRICS-263
6237	HWKMETRICS-263	{quote} What would be a good address for this functionality? {{GET /hawkular/metrics/gauges/tags/<tags>/data?}} {quote} I would say no in large part because the filtering might be expanding beyond tags. For instance, Matt and I briefly querying with the ids specified. If the client already knows the metric ids, then it would be much more efficient to specify the ids rather than a tags query.	1	1	https://issues.jboss.org/browse/HWKMETRICS-263
6238	HWKMETRICS-263	[~mwringe] can you elaborate on the expectations? Do we want the REST interface to return down sampled data for multiple metrics in a single request? Or do we want to, somehow, "merge" the data coming from the metrics?	1	1	https://issues.jboss.org/browse/HWKMETRICS-263

#	Requirement	Questions	Answered	Successfully Answered	URL
6239	HWKMETRICS-263	{quote}This ticket only covers gauges. Matt Wringe do we need similar functionality for counter rates?{quote} Yes, of course, a 'gauge' or a 'counter' is still essentially both a data point. I do think we need to handle some functionality differently between a gauge and a counter, but something like this should be handled equally in both situations. With the delta function we can applied against a counter, its essentially a gauge. If for whatever reasons counters need to be tracked separately, the counter jira is: https://issues.jboss.org/browse/HWKMETRICS-264 {quote}What would be a good address for this functionality? GET /hawkular/metrics/gauges/tags/<tags>/data{quote} I would have thought something like /hawkular/metrics/gauges/data?{queryString} Where the query string could be the same as a tag query, or it could be something like metrics=foo,bar,baz Doing with this tag queries is probably the better option, but we may want to leave this open to listing the metrics directly for future support. What ever it is, it should be consistent with our existing we REST interface. {quote}Matt Wringe can you elaborate on the expectations? Do we want the REST interface to return down sampled data for multiple metrics in a single request? Or do we want to, somehow, "merge" the data coming from the metrics?{quote} An example could be something like I have network errors. I have tx errors being stored in metricA and rx error being stored in metricB, what I want is to be able to graph all network errors (eg tx and rx) as a single metric. Right now if I wanted to do this I would have to do a query on the server based on a tag, then perform two separate get requests to get the metrics, and then merge the data on the client side. We are trying to make this a bit more streamlined and perform it in one request. I am not entirely sure if that answers your question or not. On the server side I would expect this to be just as simple as performing the separate requests to Cassandra to get the data, and then just merging the buckets together before doing things like calculating the average/min/max etc, but I don't know if it can be handled more directly in with a Cassandra query or not.	1	1	https://issues.jboss.org/browse/HWKMETRICS-263
6240	HWKMETRICS-263	[~mwringe] thanks for the lengthy clarification post. I have one more question about the merging process. Let's say the downsampling is applied over three metrics. How does downsampling work when one of the metrics does not have data in a bucket or one metric has more data points in a bucket? 1) Just work with the data in a bucket (do not care about individual data point weight); bucket data first and then apply downsampling 2) Take into account the number of metrics first; downsample at metric level first, bucket, and then downsample again over downsampled data (one data point per metric) The last URL proposed makes perfect sense. My URL was based on the assumption that is to be applied to tags only. And we can probably implement this for tags and ids at the same time. After we get this done for gauges, I expect to be a fairly simple exercise for counters. I would like to focus on this ticket first and then replicate most of the stuff to counters afterwards.	1	1	https://issues.jboss.org/browse/HWKMETRICS-263
6241	HWKMETRICS-263	{quote}How does downsampling work when one of the metrics does not have data in a bucket or one metric has more data points in a bucket?{quote} Hmm, I don't know how to handle an empty bucket. If a bucket is empty for one metric, It might make sense to leave the aggregated bucket as empty (since we don't have enough information for the aggregated bucket). This would at least indicate to the client that some data is missing rather than trying to ignore the empty bucket and returning potential confusing results. I don't know if one bucket has more data points or not should make a difference. I view this essentially as just getting the buckets for each individual metric, and then 'adding' them together (well, something along those lines, it might not be as simple as that) Does that make sense? Eg in bucket1, metricA has an average of 50, metricB has an average of 100, so the average for A+B is 150?	1	1	https://issues.jboss.org/browse/HWKMETRICS-263
6242	HWKMETRICS-263	Why would you aggregate/down sample each metric first? That sounds wrong to me. If we want the 95th percentile, are you going to take the 95th percentile of all the 95th percentiles? As the user I would not expect that. I think that we should execute the queries to determine the metrics that match the tag filters. As those results arrive, we start fetching the data points. Then merge those results and then do the down sampling. Going back to Matt's example, I would expect each bucketed data point to contain an average of all network errors, not an average of averages. It might useful to provide some meta data for the data points in each bucket. With max and min for example, it might be useful to know from which metrics own the respective data points. I am suggesting this be done right now, but something to think about.	1	1	https://issues.jboss.org/browse/HWKMETRICS-263
6243	HWKMETRICS-260	What are the steps involved with setting up a remote connect? Will I first need to obtain the certs from the Cassandra container so that I can generate a key? Whether it is a local or remote connection, I am hoping we can minimize the number of manual steps involved. How much effort is involved with this? I do not consider this a high priority, but it would be nice to have before 0.7.0 if it is not too much work in order to facilitate development/testing.	1	1	https://issues.jboss.org/browse/HWKMETRICS-260
6244	HWKMETRICS-260	{quote} What are the steps involved with setting up a remote connect? Will I first need to obtain the certs from the Cassandra container so that I can generate a key? Whether it is a local or remote connection, I am hoping we can minimize the number of manual steps involved. {quote} You would need to generate your certificate, add it to the Cassandra trust store (available in the Cassandra secret), update the Cassandra secrets to include the updated trust store, and then restart al the Cassandra instances. (or cheat and use either the already existing Cassandra or Hawkular certificate which is already authorized)	1	1	https://issues.jboss.org/browse/HWKMETRICS-260
6245	HWKMETRICS-251	Why do you think system wide defaults need to be associated to a system tenant ? Can't it be stored with no relation to a tenant?	1	1	https://issues.jboss.org/browse/HWKMETRICS-251
6246	HWKMETRICS-251	{quote} Why do you think system wide defaults need to be associated to a system tenant ? Can't it be stored with no relation to a tenant? {quote} We have a tiered system for configuration settings - system default, tenant, metric. If no data retention is specified for a tenant, the system wide default is used. If I create a metric and specify the data retention, then that will override the tenant setting; otherwise, the tenant setting is used. Tenant settings are stored in the tenants table. Per metric settings are stored in metrics_idx. I proposed the system tenant because this is where we are storing other, similar configuration settings. I think of a system tenant as something for internal, system wide configuration settings that should not be visible to users. I am not committed to a particular solution at this point. What do you have in mind?	1	1	https://issues.jboss.org/browse/HWKMETRICS-251
6247	HWKMETRICS-247	[~mithomps] I don't remember discussing tagging particular datapoints for the UI, it doesn't mean it didn't happen but I don't remember a UI usecase But IIRC, the idea was that instead of having metrics such as hostname1_cpu1,hostname1_cpu2,... we would just have a cpu metric and each data points would be tagged as hostname=X, cpu=Y and we would query a single metric to get the values about all CPUs of a single host for instance (by filtering down on the hostname only). This also doesn't require to know how many CPUs in advance. So that was more about the architecture and Metrics API IMO (Should get the input from [~pilhuhn]). Apparently BTM may have needs	0	0	https://issues.jboss.org/browse/HWKMETRICS-247
6248	HWKMETRICS-247	Just for clarification, when you say "transaction" are you talking about the transaction type (i.e. an order management process) or the instance (i.e. joe bloggs ordering a laptop)? If you are referring to 'transaction' as the type, then when retrieving the whole metric id, it would imply I could then filter the information associated with that metric id to determine which data points belong to a particular transaction instance. Is that correct? I thought that was the purpose of the tag on the data point? or is there another mechanism to filter the data points?	1	1	https://issues.jboss.org/browse/HWKMETRICS-247

#	Requirement	Questions	Answered	Successfully Answered	URL
6249	HWKMETRICS-247	Thats fine with me - I was primarily looking for recommendations on how BTM should be reporting the metrics, but that can be done as a separate exercise, and once understand create the appropriate requirements. Should I create a separate jira to discuss the BTM use case?	1	1	https://issues.jboss.org/browse/HWKMETRICS-247
6250	HWKMETRICS-237	what is the update on this?	1	1	https://issues.jboss.org/browse/HWKMETRICS-237
6251	HWKMETRICS-237	[~mwringe] can you provide some more details on what exactly is needed here. I looked over the references you cited, but it still is not clear to me. Do we need to verify that we can query for container metrics using a tag filter query?	1	1	https://issues.jboss.org/browse/HWKMETRICS-237
6252	HWKMETRICS-218	Just to be sure, there should be more "N/A" cells in the "Rate" column, right? Because the current table does not reflect the explanation coming after it or the code in the pull request (invalid tuples are ignored). Or did I misunderstand the table?	1	1	https://issues.jboss.org/browse/HWKMETRICS-218
6253	HWKMETRICS-218	Sorry, my explanation is a bit unclear. The table only illustrates my originally proposed approach. I switched to using tuples to illustrate the approaches taken in the PR because I thought that viewing the data as tuples is easier to understand. For the purposes of this ticket and the changes introduced, forget about the tables. The window operator gives us the following stream of values, (1, 2), (2, 3), (3, 1), (1, 2), (2, 3), (3, 1), (1, 2), (2, 3) We want to filter out any tuple in which the second value is less than the first. This yields, (1, 2), (2, 3), (1, 2), (2, 3), (1, 2), (2, 3) for the example.	1	1	https://issues.jboss.org/browse/HWKMETRICS-218
6254	HWKMETRICS-211	I started to make the changes that allow the Golang client to send different tenantId per fetched metric, but after thinking about the Heapster integration, there's a one serious flaw I don't have an answer to. The process of Heapster starts so that the Hawkular sink is initialized. At this init we fetch from the metrics all the metric definitions that have been set and then update the tags if necessary (and cache the results since they're needed everytime we write something to keep the definitions up-to-date). Now, this works fine with a single tenant that we now in the start. But in the case of label[project.id] -> tenantId, how do we pre-fetch all the metrics? I'm not sure if there's a way from the Heapster instance to request all the available projectIds from the Openshift. Or are there? If there isn't, there's a quite a bit more job to do (and it'll cause delays on the first time after Heapster restart a new projectId is met, which I would like to avoid) as we need to monitor the seen projectIds and when we met a new one we need to fetch all the metrics definitions, cache them, update them etc before we can continue with the metrics insertion. All in all, while label.Value -> tenantId is a fine idea and should work fine in theory, this process will increase the amount of REST-calls a lot. So much that I'm not sure if the synchronous JAX-RS 1.1 implementation will keep up without increased CPU requirements.	0	0	https://issues.jboss.org/browse/HWKMETRICS-211
6255	HWKMETRICS-208	[~mwringe] I saw that you bumped this to 0.9.0, but I almost wonder if we should get this into 0.8.0. If we introduce this into a later release (and I do think we need this), do we have to working about introducing changes that could break 0.8.0 installations when they upgrade? I am inclined to say no. The changes or restrictions would be around writing data. We would not allow ids, metric names, tag names, etc. that include reserved characters; however, I do not a problem with querying for data from a tenant/metric whose name includes reserved characters.	1	1	https://issues.jboss.org/browse/HWKMETRICS-208
6256	HWKMETRICS-206	I know that there was a good bit of discussion on IRC about some of the changes. I was unable to keep up with those discussions. If possible, can you provide a brief summary? It will help provide some context while reviewing the PR.	1	1	https://issues.jboss.org/browse/HWKMETRICS-206
6257	HWKMETRICS-205	We are going to support aggregated metrics. The plan is to allow for the aggregation functions used to be configurable. Can the cpu usage metrics be identified by tag? One of the things I want to support is specifying the metrics by tags in addition to the function(s) to apply.	0	0	https://issues.jboss.org/browse/HWKMETRICS-205
6258	HWKMETRICS-200	We're not implementing this in 0.7, are we?	0	0	https://issues.jboss.org/browse/HWKMETRICS-200
6259	HWKMETRICS-195	Using Infinispan with a replicated cache seems like an obvious choice here; however, I consider/discuss at least one other alternative approach for comparison. Is there another caching mechanism we could implement/utilize just using HTTP/REST for example? Or what about implementing some sort of gossip protocol between H-Metrics nodes?	0	0	https://issues.jboss.org/browse/HWKMETRICS-195
6260	HWKMETRICS-193	Is this really going in 0.7 ?	1	0	https://issues.jboss.org/browse/HWKMETRICS-193
6261	HWKMETRICS-191	What's up here? If this is no dev-done, then click on "reopen" and/or remove the wrongly linked PR. Also keep in mind that in RHQ, the delete per schedule was very expensive in ways of wall clock time (which was one reason to go with TTL in the first place).	1	0	https://issues.jboss.org/browse/HWKMETRICS-191
6262	HWKMETRICS-184	I had tested the running service case like this: {noformat} mvn clean install -Drunning.service {noformat} When you do that, both the {{rest-tests-integration-tests}} and {{wildfly.deployment}} profiles are activated. So Maven disables profiles active by default only if you explicitly activate another one (with {{-P}}). Well done, Maven. And sorry Matt, I wasn't aware of that I'm looking into another solution right now. Do you happen to know how we can use the enforcer plugin? I've googled a bit and couldn't find an existing rule.	0	0	https://issues.jboss.org/browse/HWKMETRICS-184
6263	HWKMETRICS-179	Wasn't this solved with [~john.sanda]'s documentation update?	1	1	https://issues.jboss.org/browse/HWKMETRICS-179
6264	HWKMETRICS-166	I'm not seeing any efficiency gains here. We still do the same 3 different writes on our side, we've forced the client to do the difficult and time consuming sorting and using named array on top of the common entity classes. What have we won here compared to doing three different REST-calls? Also, the errors happening have to be from the same as from three different calls as we can't leave the client without that information. Do we stop at the first encounter of the error? Do we resume to the next type? What parts were written and which should be repeated by the client?	1	1	https://issues.jboss.org/browse/HWKMETRICS-166

#	Requirement	Questions	Answered	Successfully Answered	URL
6265	HWKMETRICS-160	>FullStackITest.shouldFindMetricsOnServer:107 Expected: Point[name='jvm.memory.HeapMemoryUsage.committed', timestamp=1435324901000, value=1.22683392E8] Point[name='jvm.memory.HeapMemoryUsage.committed', timestamp=1435324903000, value=1.22683392E8] Point[name='jvm.memory.HeapMemoryUsage.committed', timestamp=1435324904000, value=1.22683392E8] Point[name='jvm.memory.HeapMemoryUsage.committed', timestamp=1435324904000, value=1.22683392E8] Point[name='jvm.memory.HeapMemoryUsage.init', timestamp=1435324901000, value=1.27926272E8] Point[name='jvm.memory.HeapMemoryUsage.init', timestamp=1435324903000, value=1.27926272E8] Point[name='jvm.memory.HeapMemoryUsage.init', timestamp=1435324904000, value=1.27926272E8] Point[name='jvm.memory.HeapMemoryUsage.init', timestamp=1435324904000, value=1.27926272E8] Point[name='jvm.memory.HeapMemoryUsage.max', timestamp=1435324901000, value=1.797259264E9] Point[name='jvm.memory.HeapMemoryUsage.max', timestamp=1435324903000, value=1.797259264E9] Point[name='jvm.memory.HeapMemoryUsage.max', timestamp=1435324904000, value=1.797259264E9] Point[name='jvm.memory.HeapMemoryUsage.max', timestamp=1435324904000, value=1.797259264E9] Point[name='jvm.memory.HeapMemoryUsage.used', timestamp=1435324901000, value=8467568.0] Point[name='jvm.memory.HeapMemoryUsage.used', timestamp=1435324903000, value=8467568.0] Point[name='jvm.memory.HeapMemoryUsage.used', timestamp=1435324904000, value=8467568.0] Point[name='jvm.memory.HeapMemoryUsage.used', timestamp=1435324904000, value=8467568.0] Point[name='jvm.memory.NonHeapMemoryUsage.committed', timestamp=1435324901000, value=2.7590656E7] Point[name='jvm.memory.NonHeapMemoryUsage.committed', timestamp=1435324903000, value=2.7590656E7] Point[name='jvm.memory.NonHeapMemoryUsage.committed', timestamp=1435324904000, value=2.7656192E7] Point[name='jvm.memory.NonHeapMemoryUsage.committed', timestamp=1435324904000, value=2.7656192E7] Point[name='jvm.memory.NonHeapMemoryUsage.init', timestamp=1435324901000, value=2555904.0] Point[name='jvm.memory.NonHeapMemoryUsage.init', timestamp=1435324903000, value=2555904.0] Point[name='jvm.memory.NonHeapMemoryUsage.init', timestamp=1435324904000, value=2555904.0] Point[name='jvm.memory.NonHeapMemoryUsage.init', timestamp=1435324904000, value=2555904.0] Point[name='jvm.memory.NonHeapMemoryUsage.max', timestamp=1435324901000, value=-1.0] Point[name='jvm.memory.NonHeapMemoryUsage.max', timestamp=1435324903000, value=-1.0] Point[name='jvm.memory.NonHeapMemoryUsage.max', timestamp=1435324904000, value=-1.0] Point[name='jvm.memory.NonHeapMemoryUsage.max', timestamp=1435324904000, value=-1.0] Point[name='jvm.memory.NonHeapMemoryUsage.used', timestamp=1435324901000, value=2.6468544E7] Point[name='jvm.memory.NonHeapMemoryUsage.used', timestamp=1435324903000, value=2.6461736E7] Point[name='jvm.memory.NonHeapMemoryUsage.used', timestamp=1435324904000, value=2.6495592E7] Point[name='jvm.memory.NonHeapMemoryUsage.used', timestamp=1435324904000, value=2.6499464E7] Point[name='jvm.operatingsystem.FreePhysicalMemorySize', timestamp=1435324902000, value=2.13663744E8] Point[name='jvm.operatingsystem.FreePhysicalMemorySize', timestamp=1435324903000, value=2.12885504E8] Point[name='jvm.operatingsystem.FreePhysicalMemorySize', timestamp=1435324904000, value=2.1174272E8] Point[name='jvm.operatingsystem.FreePhysicalMemorySize', timestamp=1435324904000, value=1.7186816E8] Point[name='jvm.operatingsystem.FreeSwapSpaceSize',	0	0	https://issues.jboss.org/browse/HWKMETRICS-160
6266	HWKMETRICS-160	Did not happen since the issue was reopened. Closing (hopefully for good).	0	0	https://issues.jboss.org/browse/HWKMETRICS-160
6267	HWKMETRICS-131	[~mazz] I had a quick look, it seems that in master, most of metrics artifact ids start with {{hawkular-metrics}} I wouldn't change the Openshift related module coordinates, nor those of {{ptrans}}. But I'm fine with renaming {{clients}} and {{clients-common}} to {{hawkular-metrics-clients}} and {{hawkular-metrics-clients-common}}. Is that ok for you?	1	1	https://issues.jboss.org/browse/HWKMETRICS-131
6268	HWKMETRICS-120	Sounds like we still need to investigate the shutdown issue. [~lkrejci] what version of metrics are you using? In earlier versions we may not have been closing the {{Session}}. We do now though. And alerts was not shutting down their {{Session}} either. I recently submitted a PR for alerts.	1	0	https://issues.jboss.org/browse/HWKMETRICS-120
6269	HWKMETRICS-113	{quote}There is still some tight coupling between the REST API and the core API which I think should be broken. I think we introduced some of that coupling in the past to avoid what we perceived as an unnecessary mapping back and forth between the REST API model and the core domain model which in turn leads to greater code reuse as both the REST API and core API share a set of common model classes. This presents problems though. When we change the core model implementation, we should not have to worry about it breaking the REST API, provided we are not changing the core model API. {quote} +1000 In the end the core modules should have no dependency on Jackson and Swagger. {quote} The Metric.dpart field needs to be altogether removed. It should be a calculated value based on the date partition size for the metric and the timestamp of its data points. The MetricData.timeUUID field should not be exposed in the API classes. We support unix timestamps in the API. The usage of time UUIDs is an implementation detail that should not be exposed to clients. The same goes for the writeTime and ttl fields. Those are specific to the storage implementation and should not be visible in the API. These are implementation details that should be changeable without impacting clients. That is not possible right now though. {quote} Wouldn't a REST dedicated model solve this isssue? Or maybe you think that even core library users shouldn't see those details?	1	1	https://issues.jboss.org/browse/HWKMETRICS-113
6270	HWKMETRICS-113	{quote} In the end the core modules should have no dependency on Jackson and Swagger. {quote} Right. Jackson and Swagger are specific to the REST API; therefore, those dependencies do not belong in the core. I suppose a counter argument could be made that we are only support JSON and a REST API and so having those dependencies is not a big deal. The problem with that thinking though is sooner or later other aspects of the REST API leak in the core that could prove more problematic. \\ \\ {quote} Wouldn't a REST dedicated model solve this isssue? Or maybe you think that even core library users shouldn't see those details? {quote} Right, even clients of metrics-core should not be concerned with these details. They are very much specific to the core implementation. For example, we might find that it is more efficient to manually purge data rather than relying on Cassandra's TTL feature. This is a very plausible scenario. TTL requires an additional 4 bytes overhead per cell and at least in earlier versions of Cassandra could lead to problems with compaction. If we were to change how we purge/expire data, that should not change clients' ability to to configure a metric's data retention. The REST API should have its own dedicated model. That does not preclude reusing core API classes, but the REST model should be driven by its API and not by the core API and vice versa.	1	1	https://issues.jboss.org/browse/HWKMETRICS-113

#	Requirement	Questions	Answered	Successfully Answered	URL
6271	HWKMETRICS-113	{quote} I have pushed a commit to the PR branch the removes the Jackson dependencies from metrics-core. The model classes in both the core and in the jaxrs modules are immutable for the most part. I was not going out of my way to do that, but it just kind of fell into place. I think it is a good idea though as the majority of metrics code is async. {quote} I just had a brief look so far, and it seems there are not "completely" immutable. Is that intended? {quote} We still have the swagger annotations in core-api. I am not sure of a good way to do it. Thomas Segismont do you have any suggestions? {quote} There are two ways around this, AFAIK: 1. Create a specific model for the REST API In this case, we can annotate the specific model with Swagger annotations straight away. Then we need a data binding tool (preferably MapStruct for performance). Given the little difference expected between the core model and rest model, databinding should require very little code/config. 2. Use mixins and write our own Swagger models Have a look at: org.hawkular.metrics.api.jaxrs.swagger.filter.JaxRsFilter#JaxRsFilter You can force a Swagger model by hand writing the JSON file and declaring it in the conversion filter. My preference goes to #1 as: - in both cases we have to write a class (a JavaBean for #1, a Mixin for #2) - maintaining external JSON files is error prone: any change in the model and you'll forget to update it	1	1	https://issues.jboss.org/browse/HWKMETRICS-113
6272	HWKMETRICS-113	{quote} I just had a brief look so far, and it seems there are not "completely" immutable. Is that intended? {quote} I was aware of one or two places like {{Metric.addData}}. It is not intended. I will go through and make the necessary changes. {quote} My preference goes to #1 as: * in both cases we have to write a class (a JavaBean for #1, a Mixin for #2) * maintaining external JSON files is error prone: any change in the model and you'll forget to update it {quote} I have only looked at MapStruct very briefly (since I read your comment about it). Maybe I do not understand correctly, but wouldn't we have to introduce the mapper interface in core-api? If so I do not see that as being any better off than having Swagger in core-api. And for option 2, we would writing extra JSON files, not Java classes, right? Just want to make sure I understand what you are proposing.	1	1	https://issues.jboss.org/browse/HWKMETRICS-113
6273	HWKMETRICS-103	Hey Matt, in order to help closing issues I'd like to test this locally. Is the Openshift v3 setup and metrics deployment documented somewhere? If it's not, could you help me to get started? Thanks!	0	0	https://issues.jboss.org/browse/HWKMETRICS-103
6274	HWKMETRICS-90	Same here, can you help me to setup a test environment on my box? Thanks	1	1	https://issues.jboss.org/browse/HWKMETRICS-90
6275	HWKMETRICS-74	"we may want to change endpoints to do streaming" - what endpoints? How would that change look like?	0	0	https://issues.jboss.org/browse/HWKMETRICS-74
6276	HWKMETRICS-50	There have already been discussion around endpoint changes. And as the UI is the primary client for reading data, we need to take into consider what makes sense for the UI in terms of those changes. I have a couple questions. 1) What availability states do we need to represent? 2) There was no bucketing for HWKMETRICS-36, just an aggregation. Is that what we want here as well?	0	0	https://issues.jboss.org/browse/HWKMETRICS-50
6277	HWKMETRICS-35	Right now we (still) have the issue that we don't really have the concept of availability. For the pinger we have http status codes (and should perhaps even have some pseudo codes). And classic availability of 0=down, 1=up. So how to map the http status codes (and perhaps ping times) into 0 and 1? There is an email "Computed resource state" on hawkular-devel that deals with this. For the moment we an probably assume 1xx,2xx,3xx = up , otherwise down. But this won't fit for the future. And as we discussed today at the Watercooler, we may in the future even have more sources to compute availability from. If it helps, we can improve the pinger to also deliver availability with the above rule into metrics and put on the bus for alerts.	1	1	https://issues.jboss.org/browse/HWKMETRICS-35
6278	HWKMETRICS-35	Could we continue this conversation on hawkular-dev? My opinion is that availability should not be part of the Metrics. I believe we should open the series capabilities to store more than just numeric points. And then H* components could store availability series as it makes most sense. I'm really not sure Metrics should try to model what availability means.	1	1	https://issues.jboss.org/browse/HWKMETRICS-35
6279	HWKMETRICS-34	_avg_ is already supported at _/\{tenantId\}/metrics/numeric/\{id\}/data_ if you provide a _buckets_ param. By default, _buckets_ is zero, and data points are returned as they are stored. [~mithomps] : if we add _median_ and _95percentile_ to the statistics object, would the endpoint API fit your needs?	1	1	https://issues.jboss.org/browse/HWKMETRICS-34
6280	HWKMETRICS-34	I am not sure that I agree that the current endpoint is a problem with respect to the data format. It is common enough for an endpoint to support different formats - xml, json, yaml, etc. If we introduce a second endpoint, e.g., /data/statistics, that endpoint is going to have different data formats, so does it really simply the API from a client perspective? We also need to consider what things will look like once we start adding support for aggregate metrics. Should /metric/Id/data only return raw data? Suppose for some metric, we have 5 min, 15 min, 1 hr, etc. aggregate metrics. Should all of those be handled by the /data/statistics endpoint? I don't think so, and here is a compelling reason. Let's say I send a request /metricId/data with a date range of a year. That could wind up being too many data points. We should by default perform downsampling to reduce the granularity of the data. How exactly we do the downsampling has yet to be determined, but it is something to consider.	0	0	https://issues.jboss.org/browse/HWKMETRICS-34
6281	HWKMETRICS-32	[~jkandasa], is this still an issue for you?	0	0	https://issues.jboss.org/browse/HWKMETRICS-32
6282	HWKMETRICS-31	Stefan is already in the process of refactoring and cleaning up the REST API. Do we really need another ticket and two people working on it? The API is not that large.	1	1	https://issues.jboss.org/browse/HWKMETRICS-31
6283	HWKMETRICS-31	Can this ticket be closed? [~nstefan] what do you think?	1	1	https://issues.jboss.org/browse/HWKMETRICS-31
6284	HWKMETRICS-24	In addition to retrieving data for a metric, I think we also need to be able to retrieve the metric's meta data independently of fetching any data. Meta data current includes tags (for the metric itself, not for individual data points) and data retention. It might also be useful to include some information about any pre-computed aggregates that get generated from the metric. I suppose we could have something like {{/hawkular/tenants/\{tenantId\}/metrics/numeric/\{id\}}} for retrieving data, and something like {{/hawkular/tenants/\{tenantId\}/metrics/numeric/\{id\}/meta}} for retrieving the meta data. What do you think?	0	0	https://issues.jboss.org/browse/HWKMETRICS-24
6285	HWKMETRICS-24	I have a few thoughts to share. First, since this is a breaking API change, we probably ought to give a chance for some discussion on the dev list before merging. Secondly, the change results in an increase in endpoints. At what do we (or should we) become concerned about the API becoming too complex due to the number of endpoints? I am necessarily suggesting that it is a problem, but I think it is something we should consider whenever we are making these sorts of API changes. Lastly, the last part does not make sense to me. I thought that the idea of a {{/raw}} endpoint was to facilitate querying raw data points. Why would we also use that same endpoint for querying aggregated data points? Right now we only compute aggregates at query time. In the future we will add the ability to pre-compute them. I really think that is an implementation detail. If the user wants to see a fixed number of aggregated data points, we have an endpoint for that. Whether those data points are computed at query time or pre-computed is an implementation detail with which the user should not be concerned IMO.	0	0	https://issues.jboss.org/browse/HWKMETRICS-24

#	Requirement	Questions	Answered	Successfully Answered	URL
6286	HWKMETRICS-7	I had a look at both insertion methods. To make sure I've understood the requirement correctly, what we want is to avoid calling {{DataAccess#insertData}} and {{DataAccess#updateMetricsIndex}} multiple times for the same metric, correct? And this should be done by grouping data points in case the same metric was supplied multiple times by the client? In this case, while grouping data of the same metric, what should be done with the tags and interval fields? I'm asking because these fields are part or the row specification in the update query. Perhaps, if there's no reason to specify this every time we want to insert points and only points, we should change DataAccess and REST API?	1	1	https://issues.jboss.org/browse/HWKMETRICS-7
6287	HWKMETRICS-7	Let me describe the current behavior of {{MetricsService.addGaugeData}}. The behavior is basically the same for {{addAvailabilityData}}, despite the the inconsistencies introduced from the RxJava refactoring. Each metric can have one more more data points. We execute a separate batch insert for each metric. The we call {{DataAccess.updateMetricsIndex}} which does another batch insert for all of the metrics. We are already batch inserts that are getting written to the same partition, so to some extent this ticket is a little out of date. On the other hand, there are some questions that this ticket brings up that we need to answer. When if at all should we be used batch inserts? The conventional thinking I hear coming from the Cassandra community is that batch statements should be avoid and doing async inserts in parallel should be preferred. I do not think that batch inserts should altogether avoided, particularly when all of the mutations are for the same partition (which means that will be applied atomically). If a metric has only a single data point, should we still use a batch insert? I am not sure if there is any overhead for an unlogged batch statement. If there is, then that overhead could easily outweigh the benefits of batching when we are only inserting a single data point. How does cluster topology affect the usage of batch inserts? I mentioned the conventional thinking on batch inserts vs async, parallel inserts. I think there is an implicit assumption there of a multi-node cluster. With respect to Cassandra we need to support a somewhat atypical use case of a single node cluster. It might be for example, that batch inserts offer better performance for a small cluster, like up to 3 nodes, but for larger clusters async, parallel inserts are better. Do we need to call {{DataAccess.UpdateMetricsIndex}} every time we insert data? It is true that writes are cheap in Cassandra, and it is common to perform multiple writes of the same data in order to optimize for reads; however, that does not mean those writes come for free. We use (or will use) the metrics_index table to answer different queries, like what metrics exist. There is no requirement to explicitly create a metric before storing data for it. By updating metrics_index on each write, we can answer queries about metrics that have not been explicitly created. I am not sure this is necessary now because in reality at least with Hawkular and with the Heapster client, metrics are explicitly created. I think we need to get a better feel for what the overhead is of performing that index update on every write. Maybe we do want to update the index for the purpose of knowing about which metrics exist, but may we can relax how often we do the update. Maybe we update it periodically in a background job. I do not expect these questions to be answered with this ticket. In fact I should probably create a separate ticket for each one. With respect to this ticket, I think we need to with trying to answer some of these questions.	1	0	https://issues.jboss.org/browse/HWKMETRICS-7
6288	HWKMETRICS-3	[~burmanm] What is the status of this work? I want to get this done in plenty of time for 0.7.0. I know that you are juggling a things right now. If you do not think you will be able to get to this soon, could you reassign it? Thanks.	0	0	https://issues.jboss.org/browse/HWKMETRICS-3
6289	HWKMETRICS-1	Can we please first discuss use cases for the pre-computed aggregates? We should not implement them just because RHQ has them. I am not saying that they are not useful, but we should try to better understand what we get from them, when they are used (e.g. you have 1h, 6h, 1d avg and 7 days raw retention. What data from what aggregate (raw, 1d, 6h,..) is used for e.g. a display range of 15days ago to 6days ago; this clearly spans the whole range. Also do we know that min/max/avg is really what we want? What about 99%ile or such?	0	0	https://issues.jboss.org/browse/HWKMETRICS-1
6290	HAWKULARQE-8	FYI: Update chat on the Jira Alissa Bonas added a comment - 6 days ago Where does this requirement stand? I don't see any activity on it, but it's still targeted to 4.5 Permalink pilhuhn Heiko Rupp added a comment - 9 hours ago This is in my opinion mostly a documentation issue on how to point the external MM instead of the one inside of OCP. As resources inside OCP can talk to IPs outside, I think this is solved engineering-wise Permalink jdoyle John Doyle added a comment - 1 hour ago Yes, I agree Heiko.	1	0	https://issues.jboss.org/browse/HAWKULARQE-8
6291	HIBERNATE-132	Hi, can I get a comment about this issue? It is resolved within ten minutes or so :-) Regards A	0	0	https://issues.jboss.org/browse/HIBERNATE-132
6292	HIBERNATE-88	Should this also go in the release notes?	0	0	https://issues.jboss.org/browse/HIBERNATE-88
6293	HIBERNATE-85	I don't know what led you to believe Cache 1.3 was used. Here is the information from the log, courtesy of Mike (from Release Engineering): 0:jboss-cache-1.4.1-4.SP8_CP01.1jpp.ep1.1.el4.noarch You can see what was used in the build in: http://porkchop.devel.redhat.com/brewroot/packages/hibernate3/3.2.4/1.SP1_CP02.0jpp.ep1.1.el4/data/logs/noarch/root.log On the other hand, we did use JBoss 4.0.5 so we could build it with JDK 1.4: Also, from the same log, we haven't used the JBoss AS at all, see same log above. We have this in the spec file: # FIXME Some circular dependencies here... #BuildRequires: jboss4-jmx >= 0:4.0.2 #BuildRequires: jboss4-system >= 0:4.0.2 as emmanuel gave us the following incantation on irc: ant init jar javadoc refdoc dist and it works fine. Perhaps we need those for testing? In any case, we only built Hibernate with JDK 1.4 due to a request from Dimitris that is probably not any longer applicable. So we can rebuild with JDK 1.5 if you want. This will cause a hige number of alarms to set off on QE tests (the change of JDK 1.4 to JDK 1.5), so please let Rajesh know it is comming. Do you want us to add the jboss-mx and jboss-system as build dependencies? They don't seem to be used. Perhaps we should now use a different ant target to build hibernate (different from the above)?	0	0	https://issues.jboss.org/browse/HIBERNATE-85
6294	HIBERNATE-55	I just upgraded from hibernate 3.0.5 to 3.2.4.sp1 (included with JBoss 4.2.1) and now my application which has been running for two years is getting the "Closing a connection for you. Please close them yourself" log message and a STACKTRACE. I am also running hibernate as an MBean. 1) Is there a workaound for this issue? 2) Is it OK to just call hsession.close() at the end of the EJB method getting this warning? Or is this indicative of some other problem. Thank you.	0	0	https://issues.jboss.org/browse/HIBERNATE-55
6295	HIBERNATE-17	Shouldn't we simply add a couple of "attributes" in the hibernate-service.xml? like: <attribute name="PostUpdateListerner"> net.objectlab.safemargin.persist.util.AuditInterceptor </attribute> and the Hibernate.java would do the following: private void buildSessionFactory() throws Exception { log.debug("Building SessionFactory; " + this); Configuration cfg = new Configuration(); .... cfg.setlistener("post-update", getPostUpdateListener()); ... etc for all listeners? I'm quite surprised this is not supported... I am really looking for this functionality... Can I help? Thanks Benoit	1	1	https://issues.jboss.org/browse/HIBERNATE-17
6296	HIBERNATE-17	Thanks Steve, Would be good to have an example of how you see it working for say, an EJB application that requires a post-update listener... where would you put that code? How could it be set without the application knowing? The app is using a jboss-service.xml from HAR; all .hbm.xml files are automatically detected so my application code does not know anything about listeners... Moving away from configuration file to programmatic setup raises a few questions in my mind... help... Many thanks Benoit.	1	0	https://issues.jboss.org/browse/HIBERNATE-17

#	Requirement	Questions	Answered	Successfully Answered	URL
6297	HIBERNATE-17	Huh? <mbean code="org.jboss.hibernate.jmx.Hibernate ...> ... <attribute name="listenerInjector">MyListenerInjector</attribute> </mbean> MyListenerInjector implements ListenerInjector { public void injectListeners(ObjectName objectName, Configuration configuration) throws DeploymentException { configuration.setListener("post-update", getPostUpdateListener()); } } Not sure why your code needs to know about listeners. This class is orthogonal to your application. But regardless, you are missing the point, I think. Say that not only do you have a post-update listener, but also a flush listener that needs to know about your post-update listener so that it can get information from it. How can that possibly gracefully be acheived by code within the MBean instantiating them? Or what if your post-update listener needs access to a seperate datasource. Again, how could generic instantiation code inside the MBean ever handle that gracefully?	1	0	https://issues.jboss.org/browse/HIBERNATE-17
6298	HIBERNATE-17	Thanks Steve, I now see how it is working, I did not get the fact that the interceptor would have been an attribute to the MBean; that was why I was wondering how you'd do it programmatically. May be it was clear on the forum but not from the comments here... I understand your comments on the need of having more flexibilty in defining the listeners, we also need 'easy' ways to deploy simple listeners. Anyway, this looks good! Thanks! Will it be part of JBoss 4.0.3? Benoit	1	0	https://issues.jboss.org/browse/HIBERNATE-17
6299	HIBERNATE-14	I think this feature is very easy to implement and would help many developers to simplify configuration files maintenance. Why it was removed from all versions?	0	0	https://issues.jboss.org/browse/HIBERNATE-14
6300	HORNETQ-1535	There is not much we can do with just 2 phisical nodes. There's is a voting schema between the broker to ellect if the node is really down.. but with just 2 nodes there's not much to be done. We could add a feature maybe to live to stop live whenever the connection to the backup goes down.. but even though it would have a tradeof to have live going down if backup is down. The only possible way to fix this would be to either add more brokers, or to add a sentinel to the topology (a new feature).. but such thing would require a 3rd and 4th nodes to look over the network and decide when things are down. It would be a new feature and I don't think it could be done on 2.3.x.	1	1	https://issues.jboss.org/browse/HORNETQ-1535
6301	HORNETQ-1535	Hi guys, I have an interest in this issue (unfortunately I dont have much knowledge on how the replication cluster works). The direction here seems to be focused on preventing the backup becoming live. In general when the connectivity between the backup an live is broken the backup activates. Here is possibly naive suggestion. Is there any merit is making the live-backup continually ping for the originally configured live and to revert to backup mode if it rediscovers the originally configured live....? I guess some form of message reconciliation would also be required for this revert-to-backup mode event was triggered. Also your suggestion [~clebert] seems similar to the way VCS NetworkHosts attribute works to monitor NICs "List of hosts on the network that are pinged to determine if the network connection is alive" ..? your comment "if the network is gone.. all nodes will be gone. (which is better than a split brain)." ....possibly..? , but only if its not possible to dynamically reform the cluster.	1	1	https://issues.jboss.org/browse/HORNETQ-1535
6302	HORNETQ-1535	(Well... ok, we could keep pinging the addresses and start whenever is active again.. problem solved :)) This is just an idea I had over coffee yesterday. I didn't see any other similar features out there.. so I'm still learning the possibilities.	1	1	https://issues.jboss.org/browse/HORNETQ-1535
6303	HORNETQ-1461	What does this mean: "activemq6 which is now upstream for hornetq"?	1	0	https://issues.jboss.org/browse/HORNETQ-1461
6304	HORNETQ-1452	It's optional.. and we have a requirement for compilation and whatever is available at runtime. Perhaps you just need the dependency to be runtime? Do you have a fix? we could merge it. With activemq6 this is moot as we are removing the dependency from the codebase and placing it elsewhere. Check the activemq6 repo: https://github.com/apache/activemq-6/	1	0	https://issues.jboss.org/browse/HORNETQ-1452
6305	HORNETQ-1452	Thanks for getting back at me. If It's an optional dependency, adding {{<optional>true</optional>}} for that dependency in the {{hornetq-jms-server}} pom would be enough. Are you going to release a maintenance release for it still?	0	0	https://issues.jboss.org/browse/HORNETQ-1452
6306	HORNETQ-1428	this would fit better on the user's forum instead of a bug report... The receive(0, true) will make a roundtrip to the server what will then immediately go to the server and return if there is any message available. I would need a couple more information about your case.. such as version of hornetq, your testcase.. etc. It could be an old issue that was already fixed on blocking producers... or any other number of factors... I will close this now... we certainly don't ignore timetous on receive.. if there is any issue this is not the issue. so, please open an user forum (or support ticket if you have EAP support) with some more information about your case.	0	0	https://issues.jboss.org/browse/HORNETQ-1428
6307	HORNETQ-1420	Andy is it not sensible to show that the node is continuously trying to reconnect? I could see this being useful when trying to figure out what is happening in the cluster and to ensure that the node is still retrying?	0	0	https://issues.jboss.org/browse/HORNETQ-1420
6308	HORNETQ-1387	command line errors? hornetQ doesn't have a command line tool... what are you referring to? Wildfly CLI?	1	1	https://issues.jboss.org/browse/HORNETQ-1387
6309	HORNETQ-1383	Is there any doc/webpage that explains how is this going to work? What qualifies as a slow consumer? Is it a consumer that is hung and has no associated connection (we do have occasional abandoned/orphaned consumers), or is it a consumer that is simply not able to keep up with the message publish rate. If its the former, to close abandoned consumers, then this would be a great feature. But if we would end up reaping slow but active consumers, then this is not desirable. What does it mean to remove a slow consumer? Will the server close the connection or just that specific consumer? How will this impact other consumers/producers that are associated with the same session or with the same JMS connection? In the case of a reaping an active (but slow) consumer, what exception would be thrown back to the JMS client? When HQ is used in a normal enterprise deployment with multiple independent producers and multiple independent consumers on multiple destinations, paging is a fact of life. You could have multiple destinations that are paging simultaneously - these could be JMS queues or durable subscribers on a JMS topic. And the server should be able to handle this. All indications are that it already does. My suggestion is that if there are certain functions that are slow when paging is in effect, then we should address the efficiency of the paging design. A JMS server should never unilaterally close any active JMS resource such as a consumer or a connection. Thanks Aspi Engineer	1	1	https://issues.jboss.org/browse/HORNETQ-1383

#	Requirement	Questions	Answered	Successfully Answered	URL
6310	HORNETQ-1383	bq. Is there any doc/webpage that explains how is this going to work? Clebert provided one link. Here's another: https://github.com/hornetq/hornetq/blob/master/docs/user-manual/en/queue-attributes.xml#L160 bq. What qualifies as a slow consumer? Is it a consumer that is hung and has no associated connection (we do have occasional abandoned/orphaned consumers), or is it a consumer that is simply not able to keep up with the message publish rate. A slow consumer is a consumer whose rate of message acknowledgement falls below the specified {{slow-consumer-threshold}}. This assumes that the message production rate meets or exceeds the {{slow-consumer-threshold}} as well. I'm not clear about what you mean when you say, "a consumer that is hung and has no associated connection (we do have occasional abandoned/orphaned consumers)." Can you clarify that a bit more? My understanding is that a consumer must be associated with a connection on the broker. bq. If its the former, to close abandoned consumers, then this would be a great feature. But if we would end up reaping slow but active consumers, then this is not desirable. Actually, reaping slow-but-active consumers is exactly what some users want. bq. What does it mean to remove a slow consumer? Will the server close the connection or just that specific consumer? The connection will potentially be closed. However, you can configure instead for a notification to be sent to the management address with details about the slow consumer so that different action can be taken if desired. bq. How will this impact other consumers/producers that are associated with the same session or with the same JMS connection? If the broker is configured to close the consumer's connection then any other session on that connection would be closed as well. bq. In the case of a reaping an active (but slow) consumer, what exception would be thrown back to the JMS client? The consumer would receive an {{org.hornetq.api.core.HornetQObjectClosedException}}. bq. When HQ is used in a normal enterprise deployment with multiple independent producers and multiple independent consumers on multiple destinations, paging is a fact of life. What you consider "normal" is not the same for all users. Some users want to avoid paging for slow consumers. bq. You could have multiple destinations that are paging simultaneously - these could be JMS queues or durable subscribers on a JMS topic. And the server should be able to handle this. The server can and does handle this, but there are performance implications which some users wish to avoid. bq. My suggestion is that if there are certain functions that are slow when paging is in effect, then we should address the efficiency of the paging design. Paging is slower than not paging by it's very nature. I'm sure there are efficiency gains to be made in our design and I'm open to investigating that angle. However, I believe that some users still would want to avoid the performance penalty associated with paging for slow consumers. bq. A JMS server should never unilaterally close any active JMS resource such as a consumer or a connection. This feature is configurable. If you don't want it to close the connection of slow consumers it doesn't have to.	1	1	https://issues.jboss.org/browse/HORNETQ-1383
6311	HORNETQ-1383	Looking in Git, it appears this fix was cherry-picked into the 2.3 release branch and featured in 2.3.24. Should the fix versions in this issue be updated? Will it be cherry-picked into the next 2.4.x release too?	0	0	https://issues.jboss.org/browse/HORNETQ-1383
6312	HORNETQ-1383	This feature has been implemented in the 2.3.x branch but not in the 2.4.x branch. Is it normal? It does not seem correct to have new features in older branches that are not in the upstream branch.	0	0	https://issues.jboss.org/browse/HORNETQ-1383
6313	HORNETQ-1374	Can the specified address be a wildcard?	1	1	https://issues.jboss.org/browse/HORNETQ-1374
6314	HORNETQ-1327	Hm how can I assign this feature request to Justin Bertram as requested by him in the forum thread ? Any ideas ?	0	0	https://issues.jboss.org/browse/HORNETQ-1327
6315	HORNETQ-1307	I'm also having this issue running embedded. I could alternatively use xml file for configuration, but cannot find any clear documentation on the xml config with embedded, and have tried a few different, and cannot connect (creating a client session). Could anyone point me to some help on that? {color:red}Update:{color} [this | http://stackoverflow.com/questions/21397534/embedded-hornetq-with-persistence-enabled-fails-to-start-with-programmatic-confi/21398118#21398118] is the workaround for programmatic config running embedded.	1	1	https://issues.jboss.org/browse/HORNETQ-1307
6316	HORNETQ-1265	Hi, Do you have any error messages? If you download EAP, you shouldn't have any problem installing HornetQ, the script should always work as far as I know. Thanks Howard	1	1	https://issues.jboss.org/browse/HORNETQ-1265
6317	HORNETQ-1265	You need to get approval from the PM to backport a feature... The JIRA should be open on the EAP proper channel (either Bugzilla or JIRA for EAP5), not on hornetq, since it's a backporting on EAP. This is a feature.. so likely to be rejected.. EAP5 is a stable branch.. no features.. bug fixes only!	0	0	https://issues.jboss.org/browse/HORNETQ-1265
6318	HORNETQ-1243	@Jim, From a HornetQ pov we don't delete any messages or the journal so its not an issue with HornetQ. However, Splitting hairs aside and considering my above comments how would you see either of the above as a possible solution, taing into consideration the constraints we have?	1	1	https://issues.jboss.org/browse/HORNETQ-1243
6319	HORNETQ-1243	I'm operating a number of layers above HornetQ itself (using immutant), but it honestly never occurred to me that a "clustered" MQ wouldn't be able to meet its obligations given changing cluster membership. My deployment target is EC2, using autoscaling of prebaked AMIs for deployment. Whether you consider that a PaaS or not (I personally don't, but that's neither here nor there), nodes do simply disappear. And, even given perfect "upstream" operations (i.e. AWS not yanking bare metal out from undeath you, which of course happens), the objective of autoscaling automation (whatever the provider) is as much to minimize overallocation as it is to cope with usage spikes. e.g. when half my users go away on Friday afternoon, I surely don't want to be paying for those CPU hours all weekend. I feel like a bit of an ass quoting from [the HornetQ docs|http://docs.jboss.org/hornetq/2.2.5.Final/user-manual/en/html/messaging-concepts.html#d0e398], but: {quote} With reliable messaging the server gives a guarantee that the message will be delivered once and only once to each consumer of a queue or each durable subscription of a topic, *even in the event of system failure*. {quote} How is message loss acceptable given that objective, even in a relatively "hostile" (but, common, and increasingly so) deployment scenario?	1	0	https://issues.jboss.org/browse/HORNETQ-1243
6320	HORNETQ-1243	If a node failed, you have to bring it back... either by a backup node... or by returning the node back. when you bring a node down.. you have to send us some option to tell us the node is going down and we would then do the proper cleanup. A fully synced cluster is useless on messaging.. just use a database if that's what you want. Message are supposed to be asynchronous.. and they stay on the journal when persistent. I don't mean to be pedantic as well.. but If you don't provide us persistent on the cloud.. then the problem is on the cloud.. we can do IO and we need the IO mechanism to guarantee delivery.	1	1	https://issues.jboss.org/browse/HORNETQ-1243

#	Requirement	Questions	Answered	Successfully Answered	URL
6321	HORNETQ-1243	[~jcrossley] as I said, we have a storage IO for the messages. You have to provide HornetQ a storage folder. if you provide me an ephemeral storage, that's what you get.. ephemeral messages! HornetQ is not a database where messages can be replicated to every other node. That would require you a sync cluster as I said and that woudn't work. Not only I don't want to redesign HornetQ as a sync cluster (as it doesn't work).. nor I have resources to do it (which I wouldn't anyways). We can do our part on the design for scale down nodes.. but you also have to do your part same way you would do with any other database (such as NoSQL, mySQL.. or anything else.. you have to keep the IO somehow for these guys) where you would provide storage IO.	1	1	https://issues.jboss.org/browse/HORNETQ-1243
6322	HORNETQ-1243	There are plenty of databases that are robust to loss of primary storage by using N-way replication to ensure durability. Full cluster sync is not the only alternative to reliable, persistent storage. For example, is there a reason you can't de-couple slaves and physical nodes? For example, if I'm willing to pay per-node bandwidth, could a given physical node act as both a master and a slave to N other nodes (where N is small and much less than the total cluster size)? In this way, any node going down causes one active node to become a live slave that could then choose to do a clean shutdown via the existing shutdown hook. If durability is more important than absolute network throughput (often the case for the EC2 / JBoss setting I think folks here are talking about), then the extra message replications at each node are not necessarily a critical performance issue. There are also replication strategies used by many replicated data stores which use rings (e.g. Riak) where durability is guaranteed by registering copies received in memory in N >= 2 servers; you don't need fsync guarantees except perhaps in a master node. I'm not sure this would work with the current HornetQ internal architecture of course.	1	1	https://issues.jboss.org/browse/HORNETQ-1243
6323	HORNETQ-1243	>> should not assume they can somehow be reactivated. That's impossible.. we can't do magic! We can't sync every message to every other node... you would maximize the throughput of a single node. JBoss MQ acted like that and it sucked on performance! a message belongs to the node where it came from. If you provided the storage up until the folder is graceful normally, then we would transfer messages. Otherwise you would need backup nodes.. or make sure the node will reapear after failed (say a VM crashed). As I said a few times... if you want persistent messages, you need to provide a persistent folder or a backup server with replication. otherwise you get what you provide!	1	1	https://issues.jboss.org/browse/HORNETQ-1243
6324	HORNETQ-1243	>>For example, is there a reason you can't de-couple slaves and physical nodes? For example, if I'm willing to pay per-node bandwidth, could a given physical node act as both a master >>and a slave to N other nodes (where N is small and much less than the total cluster size)? In this way, any node going down causes one active node to become a live slave that could >>then choose to do a clean shutdown via the existing shutdown hook. you can already do this, a jboss/Wildfly application server can have n hornetq servers, we call this the colocated topology and we have users who leverage this, altho typically its in a stable cluster, i.e. 2 jboss servers with a live/backup hornetq pair. How you would decide what nodes back up which live servers would be problematic. >>If durability is more important than absolute network throughput (often the case for the EC2 / JBoss setting I think folks here are talking about), then the extra message >>replications at each node are not necessarily a critical performance issue. Most of our serious messaging customers/users want both, durability and performance which is why HornetQ was architected the way it is.	1	1	https://issues.jboss.org/browse/HORNETQ-1243
6325	HORNETQ-1243	Is there a fundamental reason you can't have a master node also serving as a backup node for another master so we don't have to maintain a separate tier of backup servers? I think it's the doubling of nodes and figuring out how to recover from fail over that is tough for people. Can we simplify the administration of live backup nodes via some other method?	0	0	https://issues.jboss.org/browse/HORNETQ-1243
6326	HORNETQ-1243	> Is there a fundamental reason you can't have a master node also serving as a backup node for another master so we don't have to maintain a separate tier of backup servers? I think it's the doubling of nodes and figuring out how to recover from fail over that is tough for people. Can we simplify the administration of live backup nodes via some other method? Right now we offer the collocated topology, which I reckon needs a lot of manual config, which we already have a task to improve on the next release, where a node can act as a backup as well with manual config. a node would have two sets of journals and would assume the other node by transferring the messages in case of a failure. What I was planning to tackle this before next release (2.5) was: - collocated topology, but with improved configs - Maybe we still need scale down where you can gracelly shutdown a node with messages transfer. My schedule for this was still 2.5 as I talked to you guys earlier. Maybe we could squeeze something earlier (like a snapshot for you guys to test it and collaborate with us). Again: I'm all up to collaborate.. The discussion so far here wasn't being very productive with pointing database expectations which wouldn't work on messaging. Maybe we should reboot the discussion and start over? ;)	1	1	https://issues.jboss.org/browse/HORNETQ-1243
6327	HORNETQ-1243	You will need backup nodes on each node... you will need at least two nodes to preserve messages. The way we suggested you is - configure a backup on each node.. - use collocated topology. If you had an option to provide a persistent folder (that will survive a shutdown of the node, you won't need collocated topology.. maybe just a restart on a node). With these two options things will already work well with the current setup. What we plan on going forward: - Facilitate configuration on the collocated backup - Implement graceful shutdown with message transfers.	0	0	https://issues.jboss.org/browse/HORNETQ-1243
6328	HORNETQ-1243	So, let's drill into those three points more. 1) Reasonable perf. What's reasonable? Is there a number of messages per second, minute, etc. per node that we should hit? 2) No message loss I believe this is assuming no persistent messages are lost. Applications could be using volatile messages, and they would expect to have message loss. 3) Ephemeral network What's provided here? There must be some storage capability here. What is it, and is it only available locally i.e. visible only on that node? If we get more detail on the requirements, we may be able to come up with a temporary solution, through configuration, while we come up with the long-term solution that makes this much easier and smoother, out-of-the-box.	0	0	https://issues.jboss.org/browse/HORNETQ-1243
6329	HORNETQ-1225	If we kick the connection how should we handle this with HA? still disconnect, disconnect without recreating the consumer on failover or should we maybe just close the consumer rather than disconnecting the whole connection?	1	1	https://issues.jboss.org/browse/HORNETQ-1225
6330	HORNETQ-1210	Possibly a blocker.. we will talk about it	0	0	https://issues.jboss.org/browse/HORNETQ-1210
6331	HORNETQ-1208	Yong Hao Gao <hgao@redhat.com> made a comment on [bug 959216|https://bugzilla.redhat.com/show_bug.cgi?id=959216] Hi Mirek, Actually your test is not a valid one. You have to initialize your char array before constructing a string with it, like: char[] chars = new char[1024]; for (int i = 0; i < 1024; i++) { chars[i] = 'a'; } String text = new String(chars); Otherwise the char array may have contain invalid values that cause the Stomp connection to be reset. Can you adjust your test? Howard	1	1	https://issues.jboss.org/browse/HORNETQ-1208

#	Requirement	Questions	Answered	Successfully Answered	URL
6332	HORNETQ-1208	Yong Hao Gao <hgao@redhat.com> made a comment on [bug 959216|https://bugzilla.redhat.com/show_bug.cgi?id=959216] Hmm, strange. Can you attach the new test code? The line-number seems to point to the old test at org.hornetq.tests.integration.stomp.StompTest.testSendManyMessages(StompTest.java:63)	1	1	https://issues.jboss.org/browse/HORNETQ-1208
6333	HORNETQ-1208	Hi Mirek, What you need to do is replace this line: String text = new String(new char[32 * 1024]); with the following char[] chars = new char[32*1024]; for (int i = 0; i < 32*1024; i++) { chars[i] = 'a'; } String text = new String(chars); And the test will pass. If you intend to send binary data in the frame body, you have to add the content-length header so that the server can correctly extract the body of the message. Howard	1	1	https://issues.jboss.org/browse/HORNETQ-1208
6334	HORNETQ-1193	Martin Svehla <msvehla@redhat.com> made a comment on [bug 962264|https://bugzilla.redhat.com/show_bug.cgi?id=962264] Netty 3.6.6.Final is part of EAP 6.1.1.ER3 build 7a6066a78659581facb35874915bf1f9 modules/system/layers/base/org/jboss/netty/main/netty-3.6.6.Final-redhat-1.jar	1	1	https://issues.jboss.org/browse/HORNETQ-1193
6335	HORNETQ-1179	Is this really correct? The {{[javax.jms.Session|http://docs.oracle.com/javaee/6/api/javax/jms/Session.html]}} JavaDoc states, "If a client desires to have one thread produce messages while others consume them, the client should use a separate session for its producing thread." Isn't {{org.jboss.as.test.smoke.messaging.client.jms.JmsClientTestCase}} doing exactly this albeit implicitly with the MessageListener?	1	0	https://issues.jboss.org/browse/HORNETQ-1179
6336	HORNETQ-1179	am I missing something, or this is exactly what Justin said, meaning... breaking session constraint on threading? you have the listener for receiving and still using the same session on producing? That could lead to other issues.. that's why we added the warning. I know there are valid scenarios where the warning is appearing, but I don't think this is a valid case. Again: I may be missing something... let me know on that case.	1	1	https://issues.jboss.org/browse/HORNETQ-1179
6337	HORNETQ-1152	baranowb <bbaranow@redhat.com> made a comment on [bug 917606|https://bugzilla.redhat.com/show_bug.cgi?id=917606] There is such beauty: @LogMessage(level = Logger.Level.WARN) @Message(id = 222133, value = "There is no connector deployed with name '{0}'. The broadcast group with name '{1}' will not be deployed." , format = Message.Format.MESSAGE_FORMAT) void broadcastGroupNoConnector(String connectorName, String bgName); It seems like its intended behavior to skip offending group. Seems like patch Ive created for https://issues.jboss.org/browse/HORNETQ-1152 is a way to go. Clebert?	1	1	https://issues.jboss.org/browse/HORNETQ-1152
6338	HORNETQ-1145	Not quite sure about the setup(since quickstart guide content did not change). All(tested trunk, CR1 and Alpha) are affected. Runnign the run.sh: {noformat} ls: cannot access ../lib/*.jar: No such file or directory ********** java -XX:+UseParallelGC -XX:+AggressiveOpts -XX:+UseFastAccessorMethods -Xms512M -Xmx1024M -Dhornetq.config.dir=../config/stand-alone/non-clustered -Djava.util.logging.manager=org.jboss.logmanager.LogManager -Dlogging.configuration=../config/stand-alone/non-clustered/logging.properties -Djava.library.path=. -classpath ../config/stand-alone/non-clustered:../schemas/ org.hornetq.integration.bootstrap.HornetQBootstrapServer hornetq-beans.xml ********** Could not load Logmanager "org.jboss.logmanager.LogManager" java.lang.ClassNotFoundException: org.jboss.logmanager.LogManager at java.net.URLClassLoader$1.run(URLClassLoader.java:202) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:247) at java.util.logging.LogManager$1.run(LogManager.java:167) at java.security.AccessController.doPrivileged(Native Method) at java.util.logging.LogManager.<clinit>(LogManager.java:157) at java.util.logging.Logger.getLogger(Logger.java:287) at sun.management.snmp.util.MibLogger.<init>(MibLogger.java:57) at sun.management.snmp.util.MibLogger.<init>(MibLogger.java:42) at sun.management.jmxremote.ConnectorBootstrap.<clinit>(ConnectorBootstrap.java:760) at sun.management.Agent.startAgent(Agent.java:135) at sun.management.Agent.startAgent(Agent.java:239) Exception in thread "main" java.lang.NoClassDefFoundError: org/hornetq/integration/bootstrap/HornetQBootstrapServer Caused by: java.lang.ClassNotFoundException: org.hornetq.integration.bootstrap.HornetQBootstrapServer at java.net.URLClassLoader$1.run(URLClassLoader.java:202) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:247) Could not find the main class: org.hornetq.integration.bootstrap.HornetQBootstrapServer. Program will exit. {noformat}	1	0	https://issues.jboss.org/browse/HORNETQ-1145
6339	HORNETQ-1031	This is a customer request... Basically a server will replicate to two servers.. it's backup failover, and to another server where there's is no automatic failover... That server would have to be manually activated in case of a disaster (that means a real disaster.. say.. Hurricanes, Earthquakes... etc). So the data will be backed up to an offline site. in case of a disaster that site will be manually activated. It can be part of the topology for clients to reconenct but the start will be manual (by administrators). This is a tentative for 2.3.0.GA. That means.. if we can do it.. otherwise it can wait 2.3.1.GA... I'm leaving it on 2.3.0.Beta for now, but it will be certainly moved.	1	1	https://issues.jboss.org/browse/HORNETQ-1031
6340	HORNETQ-960	It exists a JMX mbean to pause an resume a destination. it works very well with Queue but not with topic. Is it supposed to work with topic ?	0	0	https://issues.jboss.org/browse/HORNETQ-960
6341	HORNETQ-942	iirc DeliveryActive was a MDB property in JBoss Messaging to prevent MDB to receive messages on deployment. It was possible to call start/stopDelivery() on the MDB MBean to toggle message delivery. There is no similar properties in HornetQ RA. Is there a workaround for it or is it something to implement?	1	1	https://issues.jboss.org/browse/HORNETQ-942
6342	HORNETQ-921	The issue you linked to as a duplicate is marked as fixed. As far as I can tell this issue is not fixed, or else I wouldn't have opened it in the first place. Can you please confirm with a link to the maven artifact in question?	1	1	https://issues.jboss.org/browse/HORNETQ-921
6343	HORNETQ-909	@Howard: I believe this issue is done, right? I will ask you two things> Talk to Jeff Mesnil and look if there is anything to be done on the configs at AS7. Also: make sure this jira is marked on the upgrades at the EAP, so QA will be able to track it down. Close it when you finish these two things please?	0	0	https://issues.jboss.org/browse/HORNETQ-909

#	Requirement	Questions	Answered	Successfully Answered	URL
6344	HORNETQ-908	Usecase: A web application running in Jboss7 needs to send to or receive from a queue on a remote HornetQ server. The javaee/jca-config example seems to address this. What we can do is define a pooled connection factory in jboss7: {code} <connectors> <connector name="remote-jmsxa"> <factory-class>org.hornetq.core.remoting.impl.netty.NettyConnectorFactory</factory-class> <param key="host" value="<remote_host_ip>"/> <param key="port" value="5445"/> </connector> <netty-connector name="netty" socket-binding="messaging"/> <netty-connector name="netty-throughput" socket-binding="messaging-throughput"> <param key="batch-delay" value="50"/> </netty-connector> <in-vm-connector name="in-vm" server-id="0"/> </connectors> <pooled-connection-factory name="remote-hornetq-ra"> <transaction mode="xa"/> <connectors> <connector-ref connector-name="remote-jmsxa"/> </connectors> <entries> <entry name="java:/RemoteJmsXA"/> </entries> </pooled-connection-factory> {code} This allows use to access the remote connection factory with a local jndi lookup: {code} Context ic = new InitialContext(); ConnectionFactory = ic.lookup("RemoteJmsXA"); {code} However, there's no way to lookup the queues we want to send messages to or receive messages from. Currently the javaee/jca-config example seems to work around this by manually constructing the Queue objects from java code. {code} HornetQQueue destQueueA = new HornetQQueue("A"); HornetQQueue destQueueB = new HornetQQueue("B"); {code} This makes the code less portable because we're not programming against the API but against the implementation. In Jboss.4/5 we could lookup queues by importing the remote jndi as an external context: {code} <mbean code="org.jboss.naming.ExternalContext" name="jboss.jndi:service=ExternalContext,jndiName=remote/pangaea"> <attribute name="JndiName">remote/pangaea</attribute> <attribute name="Properties"> java.naming.factory.initial=org.jnp.interfaces.NamingContextFactory java.naming.provider.url=jnp://remote_host_ip:1099 </attribute> <attribute name="RemoteAccess">true</attribute> </mbean> <mbean code="org.jboss.jms.jndi.JMSProviderLoader" name="jboss.jms:service=JMSProviderLoader,name=PangaeaJMSProvider"> <attribute name="ProviderName">PangaeaJMSProvider</attribute> <attribute name="ProviderAdapterClass">org.jboss.jms.jndi.JNDIProviderAdapter</attribute> <attribute name="FactoryRef">remote/pangaea/XAConnectionFactory</attribute> <attribute name="QueueFactoryRef">remote/pangaea/XAConnectionFactory</attribute> <attribute name="TopicFactoryRef">remote/pangaea/XAConnectionFactory</attribute> </mbean> <tx-connection-factory> <jndi-name>PangaeaJmsXA</jndi-name> <xa-transaction/> <rar-name>jms-ra.rar</rar-name> <connection-definition>org.jboss.resource.adapter.jms.JmsConnectionFactory</connection-definition> <config-property name="SessionDefaultType" type="java.lang.String">javax.jms.Topic</config-property> <config-property name="JmsProviderAdapterJNDI" type="java.lang.String">java:/PangaeaJMSProvider</config-property> <security-domain-and-application>PangaeaJmsXARealm</security-domain-and-application> <max-pool-size>20</max-pool-size> </tx-connection-factory> {code} But this is not possible with JBoss.7. Redhat support pointed us to this jira issue to raise our vote and explain our usecase. One thing that comes to mind us to construct an InitialContext via the url of the remote hornetq server: {code} java.naming.factory.initial=org.jnp.interfaces.NamingContextFactory java.naming.provider.url=jnp://remote_host_ip:1099 {code} but that depends on what the remote server is (jboss5, jboss7, standalone, etc) and possibly conflicting jar versions when using remoting for example.	0	0	https://issues.jboss.org/browse/HORNETQ-908
6345	HORNETQ-894	Why is this reopened if it's not an issue?	0	0	https://issues.jboss.org/browse/HORNETQ-894
6346	HORNETQ-887	We had that discussion in the past.. there isn't really a real need for that. We have had a few JIRAs that were closed with a similar discussion in the past. We have recently added a XMLExporter that will give us better managment of the data on the JOurnal. Wouldn't that cover any concern users have about it?	1	1	https://issues.jboss.org/browse/HORNETQ-887
6347	HORNETQ-887	In addition to Clebert's question about the XMLExporter, I'm curious to know about the use-case here. What specific functionality are you supporting which requires the use of DB persistence? There are multiple strong architectural reasons why we don't support persistence to a database so I'm keen to here any counter-argument.	0	0	https://issues.jboss.org/browse/HORNETQ-887
6348	HORNETQ-887	I scoured every jboss documentation about the jdbc persistance in hornetq. but there are two discussion related with that. even if I read those things I found, I couldn't understand the reason why this requirements for jdbc data store is rejected. it had discussed for two years. but why did the request is rejected eventually? please let me know if you have any document about rejecting jdbc persistance in hornetq.. I'm urgent.. help me.. thanks in adavance...	0	0	https://issues.jboss.org/browse/HORNETQ-887
6349	HORNETQ-854	We have several possibilities to speed up HornetQ UUID generation 1. we could ditch org.hornetq.utils.UUIDGenerator and uses java.util.UUID instead (+) we remove the network resolution which is the speed bottleneck (-) java.util.UUID is Java 1.5 only (aiui, it's not an issue for AS7 integration) (-) we switch from a name-based UUID (based on the MAC address) to a pseudo-random UUID. Collision can happen... 2. we could parallelize the network resolution. For each network interface, we try to check it in a thread and wait for the 1st matching one. (+) we speed up the network resolution (+) we keep the 1.4 compatibility (+) we continue to use name-based UUID (-) some minor thread management wdyt?	1	1	https://issues.jboss.org/browse/HORNETQ-854
6350	HORNETQ-854	interesting read one important thing that I did not mention is that we must generate UUID *fast* (each message generates a message ID). For that reason alone, I don't think we can use random UUID (using either Random or SecureRandom). The best way for that would be to migrate to the latest version of JUG. Previously, we forked it into HornetQ code base. Do we want to do the same (and streamline the code) or do we add another dependency to hornetq-code and use JUG from the Maven repository? *But* even using the latest code, this would not solve the boot time issue. JUG is looping around the network interfaces the same way than HornetQ does. In any case, we need to change that bit of code to parallelize the code finding a suitable network interface.	1	1	https://issues.jboss.org/browse/HORNETQ-854
6351	HORNETQ-854	Tomaz, I have provided a patch to HornetQ to speed up the loop finding the hardware address. The performance are not conclusive on my laptop (linux with 3 network interfaces). I have not modified anything else (still using the same UUID generator) I've attached the patched hornetq-core.jar, could you give it a go on you Windows machine please and tells me if it improves AS7 boot time?	1	1	https://issues.jboss.org/browse/HORNETQ-854
6352	HORNETQ-854	Great! What's the Fix Version?	1	0	https://issues.jboss.org/browse/HORNETQ-854
6353	HORNETQ-831	AndyTaylor: ok, what would the default be? [4:39pm] clebert: AndyTaylor: JCA/MDBs/inVM = unblocked... remote = Blocked [4:39pm] clebert: remote Factories outside of JCA I meant	1	1	https://issues.jboss.org/browse/HORNETQ-831
6354	HORNETQ-804	Can dev team take this JIRA on high priority? Thanks.	0	0	https://issues.jboss.org/browse/HORNETQ-804
6355	HORNETQ-763	Will this be is in the next EAP 5 roll-up patch due in Feb?	0	0	https://issues.jboss.org/browse/HORNETQ-763
6356	HORNETQ-759	BTW: We are already doing some work to stop delivering messages on the cluster once the node is dead. But for messages already scheduled to be delivered to that node are as good as on that node, hence you have to restart the node (or a copy of the node) to receive the messages. This issue is likely a duplicate unless you find a better description that tells me this is something different. Can you summarizethe issue?	1	1	https://issues.jboss.org/browse/HORNETQ-759

#	Requirement	Questions	Answered	Successfully Answered	URL
6357	HORNETQ-759	Hi Clebert, Iin simpler words, I need some mechanism in HornetQ so that the messages in the queue of one node of the cluster should be redelivered to the other node of the cluster if in case one node of the cluster goes. We cannot afford to wait to till the node restarts because that is something which is uncertain. It can/cannot happen. BTW i donot understand why you are avoiding live to live delivery like Jboss messaging. Please let me know what you are proposing in future releases of HornetQ which can address this problem. Thanks alot, Vineet	1	1	https://issues.jboss.org/browse/HORNETQ-759
6358	HORNETQ-721	> All you would need is message redistribution when slow consumers are not being fast > enough. That's why I renamed the JIRA here Agreed. > If you configured the MDB remotely, you would have an instance on each box, but > I don't think that would fix your issue as you would still have the MDBs created > evently across the cluster. I don't quite understand this. It doesn't fix the issue from a Single Point of Failure perspective (because now there's only one queue), but it *does* fix the issue from a 'MDBs consume as fast as possible' point of view. Because some MDBs run faster than others. What do you mean by 'you would still have the MDBs created evenly across the cluster'?	1	1	https://issues.jboss.org/browse/HORNETQ-721
6359	HORNETQ-720	Hi, Can you take a look at ReplicatedFailoverTest.testNonTransacted()? Right now, I've already fixed all steps that I understand from the HA code set-up. However, now the test will timeout while waiting for a {{SessionFactory}} from {{ServerLocator}}. (actually it times-out while waiting for the topology during {{SessionFactory}} creation). My problem is that I don't know where to go from here, and neither am I sure whether this is caused by something missing from the 'test set-up' or at the actual code.	1	1	https://issues.jboss.org/browse/HORNETQ-720
6360	HORNETQ-698	I just went to look into this. All classes with names *FileFactory had already been moved to hq-journal. Is there something left to move?	0	0	https://issues.jboss.org/browse/HORNETQ-698
6361	HORNETQ-695	@Andy: Is this really an issue?	0	0	https://issues.jboss.org/browse/HORNETQ-695
6362	HORNETQ-675	Hi Clebert, Thanks for the quick response. Not sure if the issue is related to Glassfish or our application. My guess is that this problem affects any container using XA transactions with the resource adapter. Do you know if anyone is actually using the hornetq-ra for XA ? Thanks Thomas	1	0	https://issues.jboss.org/browse/HORNETQ-675
6363	HORNETQ-663	What's the different between this and HORNETQ-662 ?	0	0	https://issues.jboss.org/browse/HORNETQ-663
6364	HORNETQ-662	It doesn't hang any more. There are cursors on the page-files now. However dead subscriptions will definitely cause performance issues, as the messages will build up on the folder. However, I'm not sure we can reap dead subscriptions automatically like that. Is this something you mean for rest only?	0	0	https://issues.jboss.org/browse/HORNETQ-662
6365	HORNETQ-653	This is because we changed the functionality to not call listeners on failover, how ever ive been thinking off adding an interface, say FailoverExceptionListener, that extends ExceptionListener so users can still be notified of a failover event. wdyt?	1	1	https://issues.jboss.org/browse/HORNETQ-653
6366	HORNETQ-651	Is HQ 2.3.0 (or later) the version that will go into EAP 6 ?	0	0	https://issues.jboss.org/browse/HORNETQ-651
6367	HORNETQ-651	This issue is marked as fixed in 2.2.8 but also unresolved. Which is it?	0	0	https://issues.jboss.org/browse/HORNETQ-651
6368	HORNETQ-650	What's going on with this issue guys?	0	0	https://issues.jboss.org/browse/HORNETQ-650
6369	HORNETQ-650	Hi guys, I'm not sure I have any idea about how to do this. Can anyone shed a bit more light on this? Any example in AS7? Thanks Howard	0	0	https://issues.jboss.org/browse/HORNETQ-650
6370	HORNETQ-650	What workaround? Seriously? How hard is it to add a callback interface?	0	0	https://issues.jboss.org/browse/HORNETQ-650
6371	HORNETQ-650	why do we have to, its already taken care of, see JMSQueueService	1	1	https://issues.jboss.org/browse/HORNETQ-650
6372	HORNETQ-650	Hi guys, I didn't try this 2.2.5 version yet, but is this change backwards compatible with JBoss6? Do we need to change our hornetq-jboss-beans.xml?	0	0	https://issues.jboss.org/browse/HORNETQ-650
6373	HORNETQ-640	Hi Andy, could you provide more info what was done here please? I've just got task to look at this and develop tests for it if necessary. Thank you, Mirek	1	0	https://issues.jboss.org/browse/HORNETQ-640
6374	HORNETQ-640	Thanks Andy for quick answer. Is there also a situation when xa recovery resource is unregistered from transaction manager?	1	0	https://issues.jboss.org/browse/HORNETQ-640
6375	HORNETQ-626	I've just noticed that IDEA has shuffled around the import statements a fair bit. I used the existing IDEA project file that is checked in to svn, but this appears to have code style settings that don't match the source formatting. Is this an issue? Does someone have an (or is able to update the existing) ipr file with the correct settings? If so I'll update the patch with the corrected imports.	0	0	https://issues.jboss.org/browse/HORNETQ-626
6376	HORNETQ-580	Clebert, what's the JIRA id for the duplicate? Thanks.	1	0	https://issues.jboss.org/browse/HORNETQ-580
6377	HORNETQ-569	Clebert, can you please check if this could go into 2.3? The effort seems to be rather small but the benefit is much higher. We need this for statistical purposes. Otherwise we must patch this by ourselves what we want to avoid. Thanks!	0	0	https://issues.jboss.org/browse/HORNETQ-569
6378	HORNETQ-569	Justin, many thanks for jumping in an providing a quality implementation including proper documentation and unit tests! Please let me ask a few questions regarding the implementation. Is my general understanding of the interceptors (incoming and outgoing) correct that the return value is used to combine two separate semantics in one: a) determine whether the next interceptor in the chain shall be called (if present) b) determine whether the packet processing shall be continued or aborted So for example in order to skip processing of an incoming message (packet), an interceptor has to return false. What are use cases to do the same on the outgoing message? I noticed the different implementations of blocking and non-blocking processing of outgoing packets #send versus #sendBlocking regarding interceptor evaluation. While in the non-blocking send the execution is just skipped using a return statement, in the non-blocking implementation an IllegalStateException is thrown if an interceptor returns false. Who is supposed to handle/receive this exception and what is the intended purpose? Is it considered to be a programming error if an interceptor returns false to abort any blocking send call (and this runtime exception will rush through the different layers up to the API consumer) or will any above layer handle the runtime exception and convert it to a checked HornetQException? In any case I think the error message "Interceptor failed" could be a bit misleading. If there is an error in the interceptor implementation causing an exception, this will only be logged and have no impact on both interceptor chain and packet processing - it will just be continued. If the interceptor intentionally returns false to abort, I think in both cases the message could give a better indication about the actual purpose.	1	1	https://issues.jboss.org/browse/HORNETQ-569
6379	HORNETQ-563	Please, always discuss issues on the user's forum before open them... Also: can ou provide a testcase?	1	1	https://issues.jboss.org/browse/HORNETQ-563
6380	HORNETQ-562	This description makes it sound like additional tests for STOMP are some sort of feel-good measure. I think my tickets HORNETQ-726 and HORNETQ-742 show that clearly there are important cases not covered by the tests. Is hornetQ STOMP support not intended for production use?	0	0	https://issues.jboss.org/browse/HORNETQ-562
6381	HORNETQ-556	Isn't this fixed?	0	0	https://issues.jboss.org/browse/HORNETQ-556
6382	HORNETQ-547	When will the Version2.3 come to release? Hope the Version2.3 will include the function.	0	0	https://issues.jboss.org/browse/HORNETQ-547

#	Requirement	Questions	Answered	Successfully Answered	URL
6383	HORNETQ-542	What is your expected semantics for this? Having the sender's username or the current user's. If it's the current user on the session of the receiver (like you're proposing), I see no point on this feature. As for the sender's user... you could achieve this with an interceptor.	1	0	https://issues.jboss.org/browse/HORNETQ-542
6384	HORNETQ-515	You could just have three values of signature, i.e. signature="queue|topic|generic" where the default is generic. I'm guessing xa should be an element rather than an attribute. Wdyt? Also, I don't think delegation is necessary. Simply keep the current HornetQConnectionFactory class, but create a set of sub classes, HornetQQueueConnectionFactory, HornetQTopicConnectionFactory, HornetQXAConnectionFactory, HornetQXAQueueConnectionFactory, HornetQXATopicConnectionFactory which do not have any methods but simply implement the appropriate interfaces.	1	1	https://issues.jboss.org/browse/HORNETQ-515
6385	HORNETQ-510	This may not be a bug, but it definitely is restriction that serves no purpose. The spec does state that setClientId must be the first thing that must be called. But its vague in the sense that it does not explicitly define what cannot be called before a call to setClientId is made. For example, would a call to "<connection>.toString()" also prevent subsequent calls to setClientId? I am guessing that the intent of the spec is to ensure that setClientId() cannot be called after any sessions have been opened from the connection. Restricting the sequence of getClientId followed by setClientId serves no purpose. This seems to be same conclusion that folks at Weblogic, SunMQ, and Tibco EMS have come to.	0	0	https://issues.jboss.org/browse/HORNETQ-510
6386	HORNETQ-462	Andy, do you have an outline for this chapter? Can I help in any way?	0	0	https://issues.jboss.org/browse/HORNETQ-462
6387	HORNETQ-447	What is this task? It seems very hand wavy and general. Please can we have a proper title and description, I have no idea what it involves.	1	1	https://issues.jboss.org/browse/HORNETQ-447
6388	HORNETQ-446	Clebert, the only thing I see these changes impacting is the release notes. Do you see any of these subtasks adding/removing from the current HornetQ User Guide?	1	1	https://issues.jboss.org/browse/HORNETQ-446
6389	HORNETQ-426	Hi guys, as far as I know tim is not working on this anymore and the 2.1.1 fix version seems weird to me since hornetQ 2.1.1 is released. Can someone plus reassign and update this ticket? Also, deploying that on the main maven repository would be needed.	1	0	https://issues.jboss.org/browse/HORNETQ-426
6390	HORNETQ-426	Cool, good to hear! Can you change the fix version then?	1	0	https://issues.jboss.org/browse/HORNETQ-426
6391	HORNETQ-426	come on guys. Is it that hard to get the proper poms metadata? Can I help?	1	0	https://issues.jboss.org/browse/HORNETQ-426
6392	HORNETQ-416	To add more details around this feature request. What is desired is that ability to retrieve back: - a list of connections and connection data that includes connection id (a unique identifier for this connection), client IP, jms user id with which client connected, JMS clientid - a list of sessions and session data that includes the connection id, session id (a unique identifier for this session), session mode (transacted or non-transacted) - a list of producers and producer data that includes session id and/or connection id, destination name, producer id - a list of consumers and consumer data that includes session id and/or connection id, destination name, consumer id, boolean indicating if its durable or not, durable subscription name.	0	0	https://issues.jboss.org/browse/HORNETQ-416
6393	HORNETQ-402	Andy, it looks like Tim has outlined the features you plan to implement in the thread - do you have an outline of the changes to your documentation with these updates?	0	0	https://issues.jboss.org/browse/HORNETQ-402
6394	HORNETQ-394	Why is this appropriate?	0	0	https://issues.jboss.org/browse/HORNETQ-394
6395	HORNETQ-393	Ive been playing around with this today by pulling out network cables etc and cant seem to recreate the issue. Each lookup creates a new initial context and does a lookup. If there are network issues at this point then this will throw a network exception with a no route to host and retry as per the normal code. For this to hang the network would have to go down in a very small window, after the lookup had hit the server but before the response had been returned. I'm not sure if this would actually hang either as its difficult to recreate. Terry, how were you actually seeing this happen?	1	1	https://issues.jboss.org/browse/HORNETQ-393
6396	HORNETQ-392	I know Andy added more tests, so I'm closing this as done. @Andy: You can never have enough tests ;-) but if you want more, can you add a new task for 2.3.0 please?	0	0	https://issues.jboss.org/browse/HORNETQ-392
6397	HORNETQ-374	there is already documentation for STOMP (in interoperability.xml). Do you want to move it to its own section?	0	0	https://issues.jboss.org/browse/HORNETQ-374
6398	HORNETQ-340	Andy, Finally got around to testing this under TRUNK, and the patch seems to fix the problem. Many thanks. Is this going to be added to the next release?	1	1	https://issues.jboss.org/browse/HORNETQ-340
6399	HORNETQ-329	@Bill: I believe you created some sort of Pool at your spring integration (as far as i remember from the SVN commits) Can you assign it me back please if you didn't? Or close it if you did it please?	0	0	https://issues.jboss.org/browse/HORNETQ-329
6400	HORNETQ-323	It's not a patent issue. Patents and copyrights are quite different things. What does Red Hat use in other documentation? Are you 100% sure there are no free fonts available?	1	1	https://issues.jboss.org/browse/HORNETQ-323
6401	HORNETQ-323	Did you ask the Red Hat documentation team what they use?	1	1	https://issues.jboss.org/browse/HORNETQ-323
6402	HORNETQ-323	Did you ask the Red Hat documentation team what they use? No I didn't. I'll ask them if they can help.	0	0	https://issues.jboss.org/browse/HORNETQ-323
6403	HORNETQ-316	Hi Howard! Do you have an outline of the config documentation this will require? Can I give you a hand with anything?	1	1	https://issues.jboss.org/browse/HORNETQ-316
6404	HORNETQ-316	JGroups (and so Infinispan) have implemented various strategies which work fine on AWS (and others), and more are coming, wouldn't it make sense to reuse the same code? In all cases I'd love to use HornetQ with applications which also use Infinispan and JGroups, so it would be great if the cluster discovery was consistent across all involved technologies.	0	0	https://issues.jboss.org/browse/HORNETQ-316
6405	HORNETQ-316	Hi Paul, I haven't got time working on it yet. If you have a working solution, that's great. Would you mind put your here for discussion purpose? Thanks, Howard	0	0	https://issues.jboss.org/browse/HORNETQ-316
6406	HORNETQ-316	Hi Clebert, that sounds great. I've been trying to work the JGroups Discovery through trunk, but it's very tightly coupled. A Discovery abstraction will be very useful. What's the timeline for releasing with the Discovery abstracted? Also, where is the abstracted Discovery code at so I can code up a JGroups implementation for you? Thanks, Paul.	1	1	https://issues.jboss.org/browse/HORNETQ-316
6407	HORNETQ-316	Cool... where's the dev thread at? Thanks.	0	0	https://issues.jboss.org/browse/HORNETQ-316
6408	HORNETQ-316	Will this be released with 2.3? Is there any work that needs to be done that I can help with?	0	0	https://issues.jboss.org/browse/HORNETQ-316
6409	HORNETQ-302	How is this different to creating queues using JMX, Jconsole connected to HornetQ broker? I don't remember having to specify explicitly the configuration attributes.. but I could be forgetting	0	0	https://issues.jboss.org/browse/HORNETQ-302
6410	HORNETQ-302	Judging from the issue on the JMS Spec, it seems that they agreed to the changed in the docs for JMS 2.0. How does this leave this issue?	0	0	https://issues.jboss.org/browse/HORNETQ-302
6411	HORNETQ-279	Isn't this already done?	1	1	https://issues.jboss.org/browse/HORNETQ-279
6412	HORNETQ-273	How is this different from HORNETQ-255 ? Why has a new JIRA been opened?	1	1	https://issues.jboss.org/browse/HORNETQ-273

#	Requirement	Questions	Answered	Successfully Answered	URL
6413	HORNETQ-255	there is indeed a problem with the server resources. here is the scenario: the client deletes the temp queue -> the server will handle the destuction and remove the resources associated to the queue (credit managers, paging store) the client closes the producer -> the client producer will release its outstanding credits -> the server session will handle the request and calls getCreditManagerHolder() -> this method will recreate a paging store and the CreditManagerHolder for the temp queue which was deleted These resources will then remain in the server forever and accumulate until a OOME occurs. I'm not sure what is the proper fix but we should be more careful about creating server resources when calling getCreditManagerHolder(). E.g. shouldn't we check that there are queues bound to the address before calling this method?	1	0	https://issues.jboss.org/browse/HORNETQ-255
6414	HORNETQ-243	Thanks for responding Tim. Maybe, I am having difficulty finding the information I need then in the documentation. Other than specifying hornetq-queues.xml within the classpath when starting up JBoss, can you explain how HornetQ can load 2 separate (different named) XXX-jms.xml and possibly 2 separate XXX-configuration.xml files that complement eachother? Again, we are trying to keep all of our application specific configurations separate from the default that installed in the deploy\hornetq.sar deployment in JBoss. Thanks, Dave	0	0	https://issues.jboss.org/browse/HORNETQ-243
6415	HORNETQ-229	Hi Andy, Excellent work ... pulled it down and am playing with it now. Why have I reopened the ticket? The second example in the messages.xml file examples.message-group2 isn't in the examples.xml (I'm guessing on this) file so when the pdf doc generates you get a ??? instead of a reference to the example. Please let me know if you need any more information. Cheers, Brian.	1	1	https://issues.jboss.org/browse/HORNETQ-229
6416	HORNETQ-216	I am confused about what you are suggesting here. Can you outline what test you would like to see, then I can tell you if we already have it. Not sure what you mean by "There is no sync on sendMessages on the Bridge".	1	1	https://issues.jboss.org/browse/HORNETQ-216
6417	HORNETQ-216	By "There is no sync on sendMessages on the Bridge" I mean, the Bridge will use setBlockOnTransaction(false). I was referring this will be probably be ok since it uses the SendAckHandler. Also, All I'm saying is I didn't find a test validate message consistency after failover on the Bridge. So that would mean.. we don't have a test to support it works. This task is about making sure we have a test. All the failover tests I found so far were about reconnecting the bridge, and not over consistency on the message transfers. Then is simple.. or we have a test, on this case I just couldn't see it and we don't have an issue and we just reject this issue. (no postponing needed) Or we don't have it on which case we need to write tests where a failure happens *while* messages are being transferred through the bridge and we validate consistency here. On which case I wouldn't suggest postponing.	1	0	https://issues.jboss.org/browse/HORNETQ-216
6418	HORNETQ-195	isn't there a chicken-and-egg thing here? We need the Ant libs to run the Ant script which invokes Maven which will download the library files. We cannot remove the Ant libs and let Maven download them only when (if?) our build setup use Maven as its top build tool.	1	0	https://issues.jboss.org/browse/HORNETQ-195
6419	HORNETQ-195	The build script would call mvn initialize first which downloads the packages, you then set ant home to .m2/and whatever and then run ant, i think? tbh, i think we should totally mavenize the build when we have time.	1	1	https://issues.jboss.org/browse/HORNETQ-195
6420	HORNETQ-195	should this really be done for 2.1.0.CR1 or can it be postponed?	1	0	https://issues.jboss.org/browse/HORNETQ-195
6421	HORNETQ-193	wow.. you're fast! I will review this weekend.. I have taken a quick look and it looks good.	1	1	https://issues.jboss.org/browse/HORNETQ-193
6422	HORNETQ-191	Is this part of the Failover improvements? It should probably be a sub task then?	1	1	https://issues.jboss.org/browse/HORNETQ-191
6423	HORNETQ-189	All in all the patch looks good. I had to tweak a few things to get your example to work, classpath etc and actually it never received the incoming tweet. u may need to check this. also the post office api had changed so i had to change your code slightly. Here are a few minor issues 1. Make sure all your parameters and fields are final where appropriate, i noticed a few places where they weren't 2. I notice in OutgoingTweetsHandler constructor you set up quite a but stuff, maybe this would be better in the start method. Also in OutgoingTweetsHandler whats the reason for creating a temp queue if it doesn't exist, should this just throw an error? 3. i noticed a few format differences with '}catch(Exception e)' should be on a new line. 4. in OutgoingTweetsHandler handle you are ignoring a few exceptions, is there a better way to do this. 5. IncomingTweetsHandler same as point 2 6. a few magic numbers, i.e. this.paging.setCount(1); and ServerMessage msg = new ServerMessageImpl(this.storageManager.generateUniqueID(), 50);, make these and any others constants or reuse existing constants. 7. The example and test use the titter hornetq account i guess u created. The tests may not work on hudson because of firewall restrictions and i am not sure whether having a public account is a good idea. Is there any way to use a fake twitter URL that we can fake. Might be best to raise this issue on the dev forums and w can all discuss it.	0	0	https://issues.jboss.org/browse/HORNETQ-189
6424	HORNETQ-176	Do you think you could provide a testcase for it? Also.. there is some code for that on the HornetQMessage. Basically, we just need to move that to the CoreMessage. IMO the code should also deal with conversions. and the test should also do it. This should be a quick change. Let me know if you can do it...	0	0	https://issues.jboss.org/browse/HORNETQ-176
6425	HORNETQ-174	Does the paging manager already contain this information? If not, how will you compute it?	0	0	https://issues.jboss.org/browse/HORNETQ-174
6426	HORNETQ-174	Bijith, could you attach a patch instead of the files? svn diff > HORNETQ-174.diff It is easier for us to review and apply your changes with a patch.	1	1	https://issues.jboss.org/browse/HORNETQ-174

#	Requirement	Questions	Answered	Successfully Answered	URL
6427	HORNETQ-173	> > Not sure what you mean "the default behaviour is significantly flawed". > Then run the attached program with the line "cf.setBlockOnPersistentSend(true);" commented out and see yourself what happens after the HornetQ server is shut down. > > Bear in mind that: > > a) HornetQ is not just a JMS system. Many users are not interested in strict JMS compliance. However, do not forget that many users are interested in strict JMS compliance! Actually I think that if the users choose to use the JMS API, they usually appreciate the proper implementation (according the the JMS specification). > > 2) With the default setting of block-on-persistent-send set to false, as I have mentioned before, coupled with asynchronous send acknowledgements (read the user manual) you can get much higher throughput without losing any of the once and only once guarantees you get with blocking. This is certainly not "flawed" Of course that in this scenario there is nothing flawed. I was writing about a different scenario - behaviour of MessageProducers when the connection to the server fails. > Different users require different things, some want JMS compliance, some don't. The point of HornetQ is that it is configurable to all these types of users, but to say one way is "correct" and another is not is ridiculous. > I completely agree with you that it is nice to have a configurable behaviour, so that users can configure what suits them best. If somebody is willing to change the architecture of the application to be able to benefit from the improved performance of asynchronous send acknowledgements, let them do it. Feel free to make the default parameters whatever you think is the best for your proprietary core API. However, when you claim that HornetQ supports an industry standard API (JMS), I do not understand why you don't provide the proper default behaviour which is required by the API specification. Should I understand it so that you are knowingly trading improved performance for decreased reliability and violation of the JMS API specification? That would seem to be a terribly wrong decision to me... Let me sum it: I think that if you claim that HornetQ is JMS compliant, you should do your best to have the default settings for the JMS clients as conforming to the JMS specification as possible. Sure, make it configurable to be able to do many weird things, many advanced users appreciate it, but do not let the basic users who use the default settings to fall into some stupid trap like loosing messages because of some ridiculous performance improvements... The reliability should be the basic requirement of any serious messaging system IMO.	1	1	https://issues.jboss.org/browse/HORNETQ-173
6428	HORNETQ-161	fixed in the trunk. (see changelog) The jms/embedded example demonstrates the POJOification of the JMS server config. Kenny, could you have a look at the jms/embedded example to check that this new way to configure would work fine with Spring? Please close this issue if that's fine for you.	1	1	https://issues.jboss.org/browse/HORNETQ-161
6429	HORNETQ-161	Would this mean that the use of hornet-jms.xml would no longer be necessary? The jms/embedded example still uses hornet-jms.xml, but then instantiates a Destination POJO which refers to the destination that had been previously configured in that file.	1	1	https://issues.jboss.org/browse/HORNETQ-161
6430	HORNETQ-157	Did you envision to use a different mbean server for each embedded hornetq server? I understand your use case but I think this would make the API more brittle as the JMX clients would have to be aware of the specific domain name as opposed to only connect to different mbean servers. wdyt?	0	0	https://issues.jboss.org/browse/HORNETQ-157
6431	HORNETQ-128	Is there any news on that implementation? IMO web-socket is still not production ready and comet (for example with netty nio) fits the current mainstream need well.	1	0	https://issues.jboss.org/browse/HORNETQ-128
6432	HORNETQ-126	I believe that the spec has been released is there any more talk about implementing amqp? We are looking at ways we can implement messaging between a .net system and a java system. We originally chose hornetq with this in mind, because at that point it was on the HornetQ's future radar and we didn't need it at the time. Now about a year later we are starting to get to the point that it is on our radar to do. An all of the possibilities are getting put on the table (Using Stomp, Replacing HornetQ, Use rest). I personally would not like to have to throw away our investment around hornetq. We are sending serialized protocol buffer objects as the body of the messages, and would really like a protocol the fully supported binary messaging. Are there any thoughts or timeframes?	1	0	https://issues.jboss.org/browse/HORNETQ-126
6433	HORNETQ-89	Jeff does it mean that the persistence part will be implemented for 2.3.0.Final?	0	0	https://issues.jboss.org/browse/HORNETQ-89
6434	HORNETQ-82	In release for JBoss AS 6.1.0, this release seems to be broken again, as although the JMX console and the Administrative console both report that the queue is in a Paused state, the subscribers are still able to consume from the queue. Is this behavior as expected? Looking back at JbossMessaging, when "Start/StopDelivery" modes were activated, the queue could still receive objects, but those objects could not be consumed.	1	0	https://issues.jboss.org/browse/HORNETQ-82
6435	HORNETQ-77	Is this issue fixed in 2.3.0 RC1 ? Rajika	1	0	https://issues.jboss.org/browse/HORNETQ-77
6436	HORNETQ-72	Thanks for that Tim. Let me run it back at you to be sure I understand the caveat. When we use this "local VM object ref" feature, it would be wise to destroy the message immediately after send to avoid any misadvertant change to the message, correct? P.S. When do you anticipate 2.0 Beta to be available? And will 2.0 be shipping with EAP 5.0? Cheers	1	1	https://issues.jboss.org/browse/HORNETQ-72
6437	HORNETQ-72	Would be great to have this evolution in a next version of HornetQ (another JMS implementation offers this kind of feature). Performances are really important in my current project. Is it possible to reopen this bug ?	0	0	https://issues.jboss.org/browse/HORNETQ-72
6438	HORNETQ-67	Hmmm, Jeff Mesnil sent me an email saying: "Hi Galder, We're preparing the release of JBoss Messaging and there is this old issue you reported. Is it still a concern for you? If it is, I'll fix it for JBM 2.0 otherwise, can you close it? thanks, jeff -- Jeff Mesnil <jmesnil@redhat.com> Core Developer (JBoss Messaging) JBoss by Red Hat" It's not a concern for me any more, hence I closed it :) Maybe Jeff should talk to you instead? :)	1	1	https://issues.jboss.org/browse/HORNETQ-67
6439	HORNETQ-67	What is the target ant version? JBMESSAGING-725 refers to version 1.65. Is that still the plan?	1	1	https://issues.jboss.org/browse/HORNETQ-67
6440	HORNETQ-67	Thanks Ben- Could you attach patch files for the changes? as it's a harder to see what has changed otherwise. Thx	1	1	https://issues.jboss.org/browse/HORNETQ-67
6441	HORNETQ-67	Ben - question on your patch- Why are you requiring that JAVA_HOME is set?	1	1	https://issues.jboss.org/browse/HORNETQ-67
6442	HORNETQ-67	Ben- I have made the changes so all the examples now also use the distributed ant. I also updated all the examples readmes that referred to the old way of building and also the documentation (user manual and quick start). I don't actually have Windows here, so I was only able to test on Linux. Could you test everything seems ok on Windows - including running the examples both from inside TRUNK and the distro? If all is ok, we can then close the task.	1	1	https://issues.jboss.org/browse/HORNETQ-67
6443	HORNETQ-62	Can we close this ticket? JBPAPP-7083 shows the trace is added.	0	0	https://issues.jboss.org/browse/HORNETQ-62
6444	HORNETQ-36	Hi Tim, We have a similar requirement and I see you mentioning, this is already possible MDBs - Can you please throw some light on this. Awaiting your response. My issue is similar to the above and as mentioned in here - https://community.jboss.org/message/748220 From WS: http://pic.dhe.ibm.com/infocenter/wasinfo/v6r0/index.jsp?topic=%2Fcom.ibm.websphere.nd.doc%2Finfo%2Fae%2Fae%2Ftjt0007_.html -Jain	0	0	https://issues.jboss.org/browse/HORNETQ-36

#	Requirement	Questions	Answered	Successfully Answered	URL
6445	HORNETQ-34	One thought here. Would it be possible to have this functionality working across a cluster too? I.e. all messages sent using a particular value of JMSXGroupid go to the same consumer on the cluster?	1	1	https://issues.jboss.org/browse/HORNETQ-34
6446	HORNETQ-32	Sorry to Correct, The current functionality is Queue - exception is thrown with a descriptive message when the message is put on the queue. Topic - Message is dropped due to the implementation of the Topic. That is the only case where we actually just drop the message. The current functionality works, but I see the following problems with it... 1. Hard to execute custom logic when you hit MaxSize a. When an exception is thrown for a Queue add and the max hits, a very generic jms exception is thrown. There's little way to know what the problem is without having to parse the exception message.. b. when an exception is thrown for a Topic, since each subscription is a queue, the way that it happens gives no way to find out that you hit max size in any of the subscriber queues 2. I'm not sure hitting max size represents an exception after thinking about it. IMO: Hitting max size is a rule, not an exception state. 3. Using an exception for max size will possibly make users embed some part of their business flow into an exception logic. For instance if they want to send a mail or put the message in another queue, they will have to embed that business flow whereever they have logic to put something on a queue. This would lead to spreading logic in a lot of different places. For these reasons, it seems that some cross cutting solution needs to be done. I can only think of the two already listed(make an overflow queue, add a place where we could put a point cut for aop). The question that I have is for a topic. Since each subscriber queue will get the Max size, if I have 100 subscriptions, will that mean if we 1. use the overflow, we will put a message in the overflow for each subscription(there will be 100 of them) 2. use the overflow, would we need to keep track of the subscriptions that didn't get it and the ones that did? 2. use the aop, we will call the pointcut 100 times. I'm sure these are some of the difficulties that this design has also raised in the past. Jay:)	1	0	https://issues.jboss.org/browse/HORNETQ-32
6447	HORNETQ-32	When the "drop-messages-when-full" option is enabled, will it be possible to detect when messages are dropped? If not, then would it be possible to please reopen this Jira? The motivation behind the request is as follows. For non-durable subscribers, we would typically prefer to drop messages, rather than page, for performance reasons. Therefore, I expect the "drop-messages-when-full" setting would be useful here. However, we need to detect this situation and take the appropriate corrective action (e.g. raise a system alarm). Also, more generally, for any subscriber (durable or non-durable) we ultimately need to limit the use of physical resources, whether that be a database or file system. Therefore, at some point, messages will be dropped, with a similar requirement to detect that this has occurred.	0	0	https://issues.jboss.org/browse/HORNETQ-32
6448	HORNETQ-31	From a user perspective I have a problem on my hand if I cannot accurately determine the number of messages on a single JMS destination. And it should be a single operation. That is one should not have to get the MessageCount from the address and the MessageCount from the JMS queue or topic to come up with the "real" count. Assuming that there are no filtering rules in place, the message count attribute for the JMS queue or topic should be the total number of messages that were published to the JMS destination. Your original plan was to add a MessageCount attribute to the Address mbean with the intent that a user could make the calculation: Total Messages on Single JMS destination = MessageCount from Queue/Topic + MessageCount from Address Why can we not change this so that the MessageCount attribute of the Queue/Topic includes the MessageCount from the Address?	0	0	https://issues.jboss.org/browse/HORNETQ-31
6449	HORNETQ-31	How would this MessageCount attribute work for a address corresponding to a JMS topic that has multiple durable subscriptions? At this time there is only one address for each JMS topic, irrespective of the number of durable subscriptions. So would there be just a single MessageCount for the entire address or would we get different counts for each durable subscription?	0	0	https://issues.jboss.org/browse/HORNETQ-31
6450	HORNETQ-31	If counting the messages isn't possible that easily, what about some kind of processing time (i.e. time a message was queued) measuring? JMS messages already have the JMSTimestamp field (HornetQMessage appear to have something similar, too), all we need would be some code that calculates the processing time as soon as a message is being received resp. acknowledged/committed (TBD) by the client. The values could be exposed per address/queue/topic-subscriber via JMX. It would be nice, if the average processing time in the past 1, 5, 15 minutes or so could be reported to get some kind of trend information. It would be nice as well, if there would be dedicated values for messages, that haven't been scheduled and for messages that have been scheduled to get realistic values.	1	0	https://issues.jboss.org/browse/HORNETQ-31
6451	DEVELOPER-4315	[~paul.robinson] [~adelaarreola] [~jessk] - This is important...can we get it in the sprint ASAP?	1	1	https://issues.jboss.org/browse/DEVELOPER-4315
6452	DEVELOPER-4315	[~james.cobb] when does it need to go live? [~adelaarreola] can you make this a high priority once its ready to implement?	1	1	https://issues.jboss.org/browse/DEVELOPER-4315
6453	DEVELOPER-4315	[~paul.robinson] - ASAP... Can we have it for Monday (July 24th)? [~jessk] - This ok with you? [~adelaarreola] - Here is the mockup: http://design.jboss.org/redhatdeveloper/marketing/eap7_cheatsheet/website/images/RHDP_hero_introEAP7cheatsheet_r1v1.png	1	1	https://issues.jboss.org/browse/DEVELOPER-4315
6454	DEVELOPER-4315	[~james.cobb] I added the horizontal repeat and forced the height to be same as image, can you check that it looks as it should on smaller windows? Here is a link to review: https://developers-pr.stage.redhat.com/pr/1958/export/ [~jessk] I used /cheat-sheets/advanced-red-hat-jboss-enterprise-application-platform-7/ as the CTA link, is this the correct one? thanks :)	1	1	https://issues.jboss.org/browse/DEVELOPER-4315
6455	DEVELOPER-4315	[~adelaarreola] - The desktop looks good but on the smaller sizes, they aren't designed to repeat. Can we use the existing scaling method for tablet and mobile?	1	1	https://issues.jboss.org/browse/DEVELOPER-4315
6456	DEVELOPER-4281	Because of the current status of the Adobe Implementation we are not able to do this. A few words from Marcia on this subject "We use Clicktale to track scrolling on a page. We use Adobe Analytics to capture previous page scroll. If we did use Adobe Analytics, it would mess up the Bounce rate on the page. If someone came to an Article and scrolled all the way down and then left, that would be considered a Bounce because the didn't interact with the page to move to another page. If there was a link tracking server call that fires at the bottom of the page, then Adobe doesn't consider that a bounce because two server calls fired. One on page load and one on link track to capture scroll. Also, we would have to put something in place to only fire that link tracking once as it could fire more than once if the person scrolled up and back down again. "	1	1	https://issues.jboss.org/browse/DEVELOPER-4281
6457	DEVELOPER-4281	How about we have a 'page read' event that the page fires (once) when the Javascript thinks the person has read the page? Would this work around the Workspace limitation of no "greater than" comparisons?	1	0	https://issues.jboss.org/browse/DEVELOPER-4281
6458	DEVELOPER-4280	Is it possible to take the feedback on Mondays and implement them the same day? If not I can hold on sending the deck out until you are able to implement up until, say, Wednesday. I'd like to think that most of these things are a top priority and I believe many of the requests are do-able within a few hours. That requires you have the time put aside to do them Monday (and possibly Tuesday), let me know if you don't have that time.	1	1	https://issues.jboss.org/browse/DEVELOPER-4280

#	Requirement	Questions	Answered	Successfully Answered	URL
6459	DEVELOPER-4280	I don't have any free time to do this on Mondays. Could we discuss to move our Friday meeting to Thursday to have a bit more time to have able to prepare any remaining items before Monday? Thus we probably won't need to make changes on the same day as the Monday Metrics meeting.	0	0	https://issues.jboss.org/browse/DEVELOPER-4280
6460	DEVELOPER-4254	[~ldary]: expecting a user to have a YouTube channel, in order to join the chat, seems to much to expect. Am I miss-understanding?	1	1	https://issues.jboss.org/browse/DEVELOPER-4254
6461	DEVELOPER-4243	[~hwindzio][~jessk] - I moved this to DEVELOPER because it's website related. On the URL: Currently, the URL "/events/javaone/" redirects to "/events/javaone/2015/" [~jessk] - The graphic header was intentionally year neutral. Do we want to reuse this or replace it with one for 2017?	1	1	https://issues.jboss.org/browse/DEVELOPER-4243
6462	DEVELOPER-4238	Yep. [~visuale] - Can you take a look at fixing this? Should be straight forward. Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-4238
6463	DEVELOPER-4238	Sorry, I forgot to note that I had changed this script several days ago. The 'cancel reply' links now start with /blog. Have the 404 errors ceased?	1	1	https://issues.jboss.org/browse/DEVELOPER-4238
6464	DEVELOPER-4231	[~tnolan] Hello Tiffany, does the blog hyperlink color scheme match the rules now? I think most of this was already up to spec but the blue link color now applies to the sidebar as well. I'm not exactly sure if I'm getting the accessibility-focus correct, though.	1	1	https://issues.jboss.org/browse/DEVELOPER-4231
6465	DEVELOPER-4228	Thanks [~visuale]. Can you let us see a preview of what this will look like, when you get chance? Thanks.	1	0	https://issues.jboss.org/browse/DEVELOPER-4228
6466	DEVELOPER-4219	[~lkoranda] Lukas, To make it easy for me to develop the Ansible playbooks that provide the configuration to our main Jenkins slave, I was going to create a simple development environment that can be used locally for easy development and testing. I was planning on using a Vagrant environment with a CentOS 7 VM and the Ansible provisioner to be able to do this. Vagrantfile would be committed into GitHub so that it can be easily shared with other team members. Are you happy with that approach? cheers, Rob	1	1	https://issues.jboss.org/browse/DEVELOPER-4219
6467	DEVELOPER-4218	[~visuale], [~paul.robinson] - Sorry I totally missed this JIRA. Agree with Paul I don't think we need it here. Would it be helpful for bloggers and audiences to be able to list out expertise areas?	1	0	https://issues.jboss.org/browse/DEVELOPER-4218
6468	DEVELOPER-4218	[~visuale] Thinking they should be able to pick from our standard categories but also have option to add some in. Will that work?	1	0	https://issues.jboss.org/browse/DEVELOPER-4218
6469	DEVELOPER-4204	[~ldary] Notice that this Epic now has a high-ish PDD priority. Please can you take this into account when doing sprint planning?	0	0	https://issues.jboss.org/browse/DEVELOPER-4204
6470	DEVELOPER-4202	[~hwindzio] I have the first draft of some content for when a user comes to the Search page without entering a search yet. For example, this could happen if a user "hacks" the URL and just tries to go to /search. Currently, if this happens, a search is performed that pulls in ALL the results across the site. I do not feel like this is a good use of server resources and it is not helpful to the user to just get everything if they haven't even tried to search yet. Would you like to either (1) wait until the Marketing/Content team has a chance to review this or (2) go live with my suggestion after you do a quick review and provide feedback and we can then revisit this as we look at the voice and tone on the site? I will add some notes to the InVision Project and @ mention you for further review and comments. The InVision project is linked to the JIRA issue, but here is also the direct link: https://redhat.invisionapp.com/share/SPB9XRKGD#/238490555_12-Search_-Empty_State-	1	1	https://issues.jboss.org/browse/DEVELOPER-4202
6471	DEVELOPER-4200	[~james.cobb] I think this will need some minor V+D. Can you add your self or [~brock] to the Team info in the description?	1	0	https://issues.jboss.org/browse/DEVELOPER-4200
6472	DEVELOPER-4200	[~james.cobb] This is something different than the back to results link. We"re looking for a way that user can quickly get to the top of the page to use filters, change search criteria etc. With the lazy loading, it's possible the user may be way down the page and this would ease the pain of having to manually scroll all the way back to the top. The Customer Portal has an example of this type of interaction on their KB pages. Example: https://access.redhat.com/articles/63269 !2017-06-19_12-09-36.png|thumbnail! Would you be ok with re-using this style?	1	0	https://issues.jboss.org/browse/DEVELOPER-4200
6473	DEVELOPER-4200	[~ihamilto] Is it enough to say that this should look and function exactly how it does on the Customer Portal? What additional details would you need beyond that?	1	0	https://issues.jboss.org/browse/DEVELOPER-4200
6474	DEVELOPER-4200	[~ldary] - [~tnolan] and I walked through this and cleared up my misunderstanding of how this functions. Reviewing the wireframes, I was under the impression that this worked like pinterest where the page always adds new results to the bottom and we never see the footer. I'm ok with this pagination free version where the user selects to get additional results. As far as the access use of the up arrow, I'm comfortable with that too. Tiffany and I agreed that it should match the styling of access for a more seamless transition between sites. They use the font awesome "icon-circle-arrow-up" for the icon. Their code for styling queues: height: 44px; width: 44px; padding: 10px; background-color: #333; background-color: rgba(0,0,0,0.8); -webkit-transition: opacity 0.218s; -o-transition: opacity 0.218s; transition: opacity 0.218s; font-size: 24px; line-height: 24px; vertical-align: middle; color: #fff; text-decoration: none; cursor: pointer;	0	0	https://issues.jboss.org/browse/DEVELOPER-4200
6475	DEVELOPER-4188	[~akandel] "This version in Klipfolio doesn't take into compensate with the % of trackable. I do believe we should show the raw numbers in klipfolio without alteration." I don't know what you're asking. Are you saying that the current view in Klipfolio is taking into account the % trackable?	1	0	https://issues.jboss.org/browse/DEVELOPER-4188
6476	DEVELOPER-4186	During review Heinz noticed that first iteration was using the registrations Row 19 "target" instead of active users in row 10 "target # active". In https://docs.google.com/spreadsheets/d/1ylY1Zu2xdraMKfy84upNXfliqeNzqerWl2A1SCrWjWE This has now been corrected. Can this be now closed [~hwindzio]?	1	1	https://issues.jboss.org/browse/DEVELOPER-4186
6477	DEVELOPER-4185	[~lkrzyzanek] Is there a way to split JBD Registrations by channel in order to split them into: * Net New * Existing Red Hat Account * Existing RHD Member Registrations Similar to what is done with RHT user profile? https://github.com/searchisko/configuration/blob/master/data/query/rht_user_profile_metrics.md	1	1	https://issues.jboss.org/browse/DEVELOPER-4185
6478	DEVELOPER-4154	[~ihamilto] / [~paul.robinson] Not sure I see the value in doing a spike to get the requirements for this story. For me it's as simple as we don't have all of the requirements for this story and as such it can't be played. Putting in a spike to solve that would (in my opinion) simply result in the person that owns the spike going around various members of the team and asking them how to integrate a proposed plan of action. The plan of action however I believe should be agreed by anyone impacted on this and then we can decide if we need to spike any of it (chance are we won't need to, we can simply decompose the work down into stories as I don't expect many surprises here). So in conclusion, is it just easier to call a time-boxed meeting and agree on an implementation and then create the implementation stories from there?	1	1	https://issues.jboss.org/browse/DEVELOPER-4154
6479	DEVELOPER-4154	Thinking about it, wouldn't this process be done during issue refinement, before the issue gets added to the sprint? I think this is what Rob is saying also.	1	1	https://issues.jboss.org/browse/DEVELOPER-4154

#	Requirement	Questions	Answered	Successfully Answered	URL
6480	DEVELOPER-4154	[~ldary] [~ihamilto] I think we need to agree on the version of Node to be used here. I think the project is currently using 10.0.45 (I need to check this). I know it's out-of-date, but does that work for us [~ldary] until we can schedule an upgrade? [~ihamilto] A couple of extra tasks that I took from that meeting: * Create a PR configuration for the test framework to default to headless, single run test execution * Select and configure a test reporter to provide clear output of the test results (and ensure this is archived by the build) I also think it's worth us having a think about any more code that gets written to run these tests. We've already kinda duplicated the code for running the acceptance tests for running Blinkr. If we're going to duplicate it again, then that is saying to me that we should re-factor it into a common approach that can be used to execute our code, or look at a simpler way of running these things. cheers, Rob	1	1	https://issues.jboss.org/browse/DEVELOPER-4154
6481	DEVELOPER-4149	Is there any value in doing anything with this besides simply leaving it as HTML?	1	0	https://issues.jboss.org/browse/DEVELOPER-4149
6482	DEVELOPER-4148	Is there any value in doing anything with this besides simply leaving it as HTML?	1	0	https://issues.jboss.org/browse/DEVELOPER-4148
6483	DEVELOPER-4147	Is there any value in doing anything with this besides simply leaving it as HTML?	1	1	https://issues.jboss.org/browse/DEVELOPER-4147
6484	DEVELOPER-4145	[~adelaarreola] [~ldary] [~james.cobb] Can this one be marked as done and any additional items added to DEVELOPER-4205?	1	1	https://issues.jboss.org/browse/DEVELOPER-4145
6485	DEVELOPER-4143	Hi [~ecurley] – sure, easy to do. Is the '?intcmp=' URL parameter included to record the # of clicks from the blog to the signup page?	1	1	https://issues.jboss.org/browse/DEVELOPER-4143
6486	DEVELOPER-4143	[~visuale] - Thanks. Do you know what date that went live? I'm seeing 53 leads in salesforce. Will be interesting to see how that grows. :)	1	1	https://issues.jboss.org/browse/DEVELOPER-4143
6487	DEVELOPER-4125	Thanks [~visuale], I think this is a lot easier to read. Can you ask [~ecurley] to give the final review?	1	1	https://issues.jboss.org/browse/DEVELOPER-4125
6488	DEVELOPER-4122	[~visuale] I've put this in review, as I think [~ecurley] and I need to review it. Can you follow this workflow for other issues, and notify the appropriate reviewer to take a look?	1	1	https://issues.jboss.org/browse/DEVELOPER-4122
6489	DEVELOPER-4122	This looks good to me. Can you email [~ecurley] to review too? She probably won't see this notification as JIRA sends a lot of emails!	1	1	https://issues.jboss.org/browse/DEVELOPER-4122
6490	DEVELOPER-4122	Thanks [~visuale] and [~paul.robinson]. I'm super excited to see these updates! !Screen Shot 2017-06-02 at 8.52.21 AM.png|thumbnail! !Screen Shot 2017-06-02 at 8.52.31 AM.png|thumbnail! Two quick questions and wondering if it doing it only on my end. The featured image on some of them seem to have greatly enlarged when moving to the new layout. Are you seeing this as well? (attached screenshot) If yes, is there anything we can do to have them not be so large? And when we scroll down past where the right column has material the blog has a good amount of white space (see attached). Is there a way to make what is in the right hand column adjust down as the infinite scrolling happens? Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-4122
6491	DEVELOPER-4122	[~ecurley] I'm not sure the exact behaviour of making the right-bar adjust down. Given that the right-bar is quite long, when will that scroll? It's a bit hard for me to explain here. Maybe you have an example website that does this behaviour?	1	1	https://issues.jboss.org/browse/DEVELOPER-4122
6492	DEVELOPER-4118	Axel, can you do a mock up of the chart for the weekly dashboard deck? I believe we may need a couple of adjustments to get it best visualized before thursday's deck that you send out	1	1	https://issues.jboss.org/browse/DEVELOPER-4118
6493	DEVELOPER-4114	Thanks [~visuale]. I think this is a good start, but I worry that the widget is a little hidden. Should we also add a version of it towards the top of the post? Notice how DZone has it in the header of the post: https://dzone.com/articles/three-rules-of-my-api-communication-strategy. Do you have flexibility in the layout of the widget? I.e. can you change what the vote up and down icons look like? Can you change what text is displayed?	1	1	https://issues.jboss.org/browse/DEVELOPER-4114
6494	DEVELOPER-4114	[~visuale][~paul.robinson] - Are these something we can edit?	1	0	https://issues.jboss.org/browse/DEVELOPER-4114
6495	DEVELOPER-4114	[~james.cobb] do you mean like the which kind of icons can we use? Yes, the current ones have some CSS tweaks applied as it is but they can be swapped for something else entirely.	1	0	https://issues.jboss.org/browse/DEVELOPER-4114
6496	DEVELOPER-4111	[~hwindzio] I think you said you had no more requirements. Correct? [~ldary] Can you try prioritise DEVELOPER-4197 towards the top of the stories on this Epic?	1	0	https://issues.jboss.org/browse/DEVELOPER-4111
6497	DEVELOPER-4109	Thanks [~ihamilto]. Does this story need anything else then so it can be groomed and added to a sprint?	1	0	https://issues.jboss.org/browse/DEVELOPER-4109
6498	DEVELOPER-4109	[~tnolan] Looks fine to be added to a sprint. Will we be working on this or will it be [~lkrzyzanek]?	1	1	https://issues.jboss.org/browse/DEVELOPER-4109
6499	DEVELOPER-4105	[~ecurley] [~jessk] - Can you review and add some additional details if I'm missing something? Is this the same as the salesforce id? It looks like redhat.com uses something called an internal campaign id (intcmp). Is that the direction you'd like to take as well?	1	1	https://issues.jboss.org/browse/DEVELOPER-4105
6500	DEVELOPER-4105	[~jessk] Ok, thanks for the details. I remember [~ecurley] talking about this recently in a meeting so that makes sense. My schedule is pretty packed this week, so I'm more than ok with not being included in the meeting. I just said that I would create a JIRA story for this, so I did. Once you figure out which one or if you need both, can you just update this JIRA with that info? I think Drupal can and should be able to support having a field for this info that is then used to create the links for the buttons.	1	1	https://issues.jboss.org/browse/DEVELOPER-4105
6501	DEVELOPER-4082	[~ecurley] [~andreacarmona] - Does this email come from Eloqua? If so, who would work on these changes? !RHDP-welcome-email.png|thumbnail!	1	1	https://issues.jboss.org/browse/DEVELOPER-4082
6502	DEVELOPER-4079	[~paul.robinson] This should be a quick edit. Can someone on your team complete this during the current sprint?	1	1	https://issues.jboss.org/browse/DEVELOPER-4079
6503	DEVELOPER-4078	[~paul.robinson] Thanks for including that list. Very helpful. Yes, cheat-sheet and video should be fixed. I checked .com too and they also do use the plural where appropriate. If we could get the cheat-sheets one prioritized first, I think that would assist with our transition to the Drupal format. Is there something that should be done on the backend for that or is it just a matter of updating the URLs in Drupal?	1	0	https://issues.jboss.org/browse/DEVELOPER-4078
6504	DEVELOPER-4078	Current work is done at https://github.com/lightguard/developers.redhat.com/tree/feature/cheatsheet-alias. Do we need to change things that area already published? What about redirects?	1	1	https://issues.jboss.org/browse/DEVELOPER-4078
6505	DEVELOPER-4078	[~lightguard] Yes, we will need one redirect for the Advanced Linux Commands one, since that one is published already. Is that something you would do? Current URL - /cheat-sheet/advanced-linux-commands-cheatsheet New URL (not live yet) - /cheat-sheets/advanced-linux-commands	1	0	https://issues.jboss.org/browse/DEVELOPER-4078
6506	DEVELOPER-4078	[~paul.robinson] RE: bq. The person working on this issue needs to create an RHDENG ticket, requesting the redirect. They should then work with Lukas Koranda to coordinate the changes so as to minimise 404s. Does this mean that Jason or I should do this?	1	0	https://issues.jboss.org/browse/DEVELOPER-4078
6507	DEVELOPER-4078	[~lightguard] Is the cheat sheet alias live in Drupal? Is there something I should approve to get that rolled out? I just tried creating a new one and it did not use the new format.	1	0	https://issues.jboss.org/browse/DEVELOPER-4078
6508	DEVELOPER-4072	[~paul.robinson] would it be us that needs to write these, or do we need to find someone else to do this?	1	0	https://issues.jboss.org/browse/DEVELOPER-4072

#	Requirement	Questions	Answered	Successfully Answered	URL
6509	DEVELOPER-4059	[~paul.robinson] - This has been approved by [~ecurley]. [~ecurley] - Can you please supply the date which you want to go to this hero? Mockup: http://design.jboss.org/redhatdeveloper/marketing/reactivemicroservices_book/website/images/RHDP_hero_reactivemicroservices_r1v1.png header: Get Reactive! subhead: Building Reactive Microservices in Java. CTA: Download Book Link: https://developers.redhat.com/promotions/building-reactive-microservices-in-java Desktop: http://static.jboss.org/images/rhd/hero/RHDev_homehero_desktop_reactivemicroservices_25apr2017.jpg Tablet: http://static.jboss.org/images/rhd/hero/RHDev_homehero_tablet_reactivemicroservices_25apr2017.jpg Mobile: http://static.jboss.org/images/rhd/hero/RHDev_homehero_mobile_reactivemicroservices_25apr2017.jpg	1	1	https://issues.jboss.org/browse/DEVELOPER-4059
6510	DEVELOPER-4058	[~ecurley] - Are you still wanting this to go live June 6?	1	1	https://issues.jboss.org/browse/DEVELOPER-4058
6511	DEVELOPER-4058	[~ecurley] James has approved the VD for this hero but wasn't sure if the content is still current. Can you let me know if this is ready to go live or if I should change anything? thanks. https://developers-pr.stage.redhat.com/pr/1885/export	1	1	https://issues.jboss.org/browse/DEVELOPER-4058
6512	DEVELOPER-4055	[~mikgue] - Can we get the final approved text for [~adelaarreola]?	1	1	https://issues.jboss.org/browse/DEVELOPER-4055
6513	DEVELOPER-4052	# Search for "JBoss Sandbox BOMs" # The second item should be labeled as BOM, instead of QUICKSTART The URL of the result has /boms/ in it: https://developers.redhat.com/quickstarts/sandbox/boms/ -- Is this labeled incorrectly in the DCP for the type? That may be an issue outside this.	1	1	https://issues.jboss.org/browse/DEVELOPER-4052
6514	DEVELOPER-4052	[~paul.robinson] Thanks for the info. Who's making the decision to remove those? Is that [~hwindzio]? If we're removing content from the site, I fell like that should probably be part of a larger discussion around how long content should stay on the site and when/if it should be archived. I'm not sure we've ever documented a strategy for this. I think these have some good content, especially the EAP BOMs, so I'm not sure the best strategy is for us to get rid of content that could be helpful just because they have low pageviews. These also seem to tie into some of the tasks related to Job 1. Perhaps there's a better way to present these...	1	1	https://issues.jboss.org/browse/DEVELOPER-4052
6515	DEVELOPER-4050	[~tnolan] we currently have this scheduled for the next sprint (starting 2017-06-26). Will you have the wireframe ready for end of this week, or should I push back a sprint?	1	1	https://issues.jboss.org/browse/DEVELOPER-4050
6516	DEVELOPER-4050	[~paul.robinson] I'm assigning this one to [~gdoyle] for the UX. [~james.cobb] Initial thought is that we review what is being done on the Customer Portal and come up with something similar. I think it's likely that we would look into using an icon for this. If we went that direction, is this something you think you or [~brock] would assist with as the visual designer?	1	0	https://issues.jboss.org/browse/DEVELOPER-4050
6517	DEVELOPER-4050	[~james.cobb] I talked to the customer portal team about their choice of icon (patternfly key : pficon-key), asking if there is a reason they use that instead of something like a lock. They said it is a more friendly representation of premium content. To be consistent with customer portal, UX thinks it would be best to stick with a key for the icon (doesn't have to be the same one). What are your thoughts on this?	1	1	https://issues.jboss.org/browse/DEVELOPER-4050
6518	DEVELOPER-4041	[~unibrew] Can you update this file on the file management server? I've just got a new laptop and I forgot to copy across my SSH keys! They are in my backup, but that's at home, so I won't be able to this until the weekend.	1	1	https://issues.jboss.org/browse/DEVELOPER-4041
6519	DEVELOPER-4033	[~visuale]: can you do a quick investigation to scope this out? We are curious if it's an easy fix (install a plugin), or something more complex. Also, can you provide screenshots of what it would look like? I'd like [~ecurley] and [~tnolan] to see what it would look like before it goes live.	1	1	https://issues.jboss.org/browse/DEVELOPER-4033
6520	DEVELOPER-4033	[~paul.robinson] My thinking with using a 0 for thumbs-down is that the score would then be a simple percentage score of total likes out of total votes cast. Would the idea for a thumbs-down equalling -1 be that popularity is determined by deviation instead? Also with regard to a "trending" column, whether all-time or month-by-month, how would an article's overall popularity vs. its rating factor into its place in the list? For example, _Article ABC_ could have a 90% positive rating but only have 200 total votes whereas _Article XYZ_ could have a 72% positive rating but more than 4000 total votes. To me, it would seem that the latter post has the bigger impact.	1	1	https://issues.jboss.org/browse/DEVELOPER-4033
6521	DEVELOPER-4032	[~dcoughlin1], [~adelaarreola]: Can you update the description with the steps required to achieve this?	1	1	https://issues.jboss.org/browse/DEVELOPER-4032
6522	DEVELOPER-4032	How long to be available in: https://developers.redhat.com/stack-overflow/#!q=? *Camila Macedo | Red Hat Mobile* *IRC:* cmacedo Phone: +55 11 3529-6052 *Mobile: *+55 11 96151-3631 [image: Red Hat] On Wed, Apr 26, 2017 at 7:29 PM, Daniel Coughlin (JIRA) <issues@jboss.org>	1	1	https://issues.jboss.org/browse/DEVELOPER-4032
6523	DEVELOPER-4031	I added links to the wireframes in the Description. I have alternate (MVP) versions of this page as well. Can you let me know what's feasible? If we do the preferred version, [~james.cobb] or [~brock] should be able to provide the mock-up and graphics by Friday, April 21 at the latest. Let me know if you would need it sooner.	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6524	DEVELOPER-4031	Tiffany - great stuff! Question: During an event, there's commonly a short line of people waiting to enter their emails. What's the best way to reset the form? Set it as Home? Or maybe add a "Try again"?	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6525	DEVELOPER-4031	[~james.cobb] That works for me as long as we can easily see it! It would be great if it was large and green/red for yes/no so it's easy to see. My old eyes can't read text on mobile devices. haha! Any chance you or Jason could work on a mock-up early this week? If not, would you be ok with me taking a stab and sending it to you for review?	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6526	DEVELOPER-4031	[~ldary] can you take a look at this? Can you see if someone is available to work on it this week?	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6527	DEVELOPER-4031	Success and failure look the same to me. What am I missing?	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6528	DEVELOPER-4031	Thanks James. I think those work. Can you do the mobile design too for font sizing, since that's how most of our users will be accessing this? I didn't see them but I'm on my phone so maybe I missed them. ~ TIFFANY	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6529	DEVELOPER-4031	They still look the same - what am I missing?	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6530	DEVELOPER-4031	helpful - thanks. Should we create a short link for this so people can easily get to it on their phone? Maybe "bit.ly/RHDP-verify" or whatever.	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6531	DEVELOPER-4031	Have we settled on the final URL: https://developers.redhat.com/verify-member ?	1	1	https://issues.jboss.org/browse/DEVELOPER-4031
6532	DEVELOPER-4030	[~mikgue] Does this promo count for this ticket or would you like something in addition? http://design.jboss.org/redhatdeveloper/marketing/summit_2017/images/RHDev_homepromo_summit2017.png	1	1	https://issues.jboss.org/browse/DEVELOPER-4030
6533	DEVELOPER-4030	I see this is due to go live tomorrow. We need everything confirming today if we want to hit this deadline. CC: [~ldary], [~adelaarreola].	1	1	https://issues.jboss.org/browse/DEVELOPER-4030
6534	DEVELOPER-4030	[~mikgue][~ecurley][~jessk] - Shouldn't we be promoting the live streamed keynote? Headline: Red Hat Summit is for developers Sub-head: Live Keynote feed, major news, and deep dives < I'm spit balling here. :P CTA: Learn more Target URL: https://developers.redhat.com/summit-is-for-developers/	1	1	https://issues.jboss.org/browse/DEVELOPER-4030

#	Requirement	Questions	Answered	Successfully Answered	URL
6535	DEVELOPER-4030	[~mikgue] - Updated the mockup... please refresh. http://design.jboss.org/redhatdeveloper/marketing/summit_2017/website/images/RHDP_hero_summit_r1v1.png Do we have the landing page staged?	1	1	https://issues.jboss.org/browse/DEVELOPER-4030
6536	DEVELOPER-4030	Yes - see here: https://developers.redhat.com/summit-is-for-developers/ I used a promo page format - is that the correct one or do I need to change it?	1	1	https://issues.jboss.org/browse/DEVELOPER-4030
6537	DEVELOPER-4020	Isn't this a duplicate of DEVELOPER-4019?	1	1	https://issues.jboss.org/browse/DEVELOPER-4020
6538	DEVELOPER-4019	[~dhladky] Can you provide example data for what is stored in some of the variables? I'm trying to decide which ones to use. Specifically: * tcUser * tcDownloadFileName	1	0	https://issues.jboss.org/browse/DEVELOPER-4019
6539	DEVELOPER-4017	[~paul.robinson] [~hwindzio] I don't think I see this one on the PDD Reprioritization list and it does not have a priority. However, it looks like this Epic is slated for Q2. Should this be added to the spreadsheet?	1	1	https://issues.jboss.org/browse/DEVELOPER-4017
6540	DEVELOPER-4017	[~paul.robinson][~hwindzio] From our discussion the other day, I believe we decided to still give this one a priority, but not include it in the PDD. Heinz - Can you assign a Q2 priority number to this one here and in the prioritization doc? I added a row for it as the last item in the Q2 list.	1	1	https://issues.jboss.org/browse/DEVELOPER-4017
6541	DEVELOPER-4017	[~tnolan] I assign this a priority of 4, since last week it seemed like you had been asking about next epics to work on and I think the timing is right to start on this one. We are seeing people leave the site after viewing the cheatsheet page: https://docs.google.com/presentation/d/1ciZ7NPhfP9WXKtdaYBUQ1V6OeKXh5kaIVjFH-RJcFUc/edit#slide=id.g1d607e393c_0_12 From your comments in the previous weekly dashboard : "We have an Epic that is focused on improving this experience https://issues.jboss.org/browse/DEVELOPER-4017 and we planned to add a section to the cheat sheet page for "more cheat sheets"." Can we find other ways to keep them on the site as well? I would think that developers want to stay on topic, so maybe provide related links?	1	0	https://issues.jboss.org/browse/DEVELOPER-4017
6542	DEVELOPER-4017	For example, this one for containers - should we maybe provide links to the Containers topic and also a download for the CDK? https://developers.redhat.com/promotions/docker-cheatsheet/	1	0	https://issues.jboss.org/browse/DEVELOPER-4017
6543	DEVELOPER-4012	If it's essentially the same, does it make sense for it to be a new content type? Sounds like it might be easier to include a dropdown in the current content type to specify what it is or where it should go. We can then determine which layout to use based on that info.	1	1	https://issues.jboss.org/browse/DEVELOPER-4012
6544	DEVELOPER-4012	[~lightguard] That was what i was thinking originally too. However, are you able to change the default URL based on what the user would select from the drop-down? Also, they may be the same now, but I'm not sure that will always be the case. For example, I'd still like to look into some options to embed the actual cheat sheet on the page. This is not something we would want to do with an eBook. Can you come up with a quick list of benefits/drawbacks for doing a new content type vs. adding some logic to the current one? I'm just not sure what the overall difference would be to recommend/request one over the other for usability and future-proofing.	1	1	https://issues.jboss.org/browse/DEVELOPER-4012
6545	DEVELOPER-4012	[~tnolan]are there fields that are only required for ebooks? Are there fields that should not be included for print books or ebooks? How do we want those displayed? Do we want the form to change once you tick the ebook checkbox?	1	1	https://issues.jboss.org/browse/DEVELOPER-4012
6546	DEVELOPER-4012	[~tnolan] [~james.cobb] thanks! should the page also follow the same layout as cheat-sheets for mobile and tablet? and can you let me know which are the must-haves, and lessed priority items like you did with the cheat sheets? Also, just to be sure, should the published date be the date the book's page was published, link in the cheat sheets, or the date the book was published?	0	0	https://issues.jboss.org/browse/DEVELOPER-4012
6547	DEVELOPER-4012	Don't we already have this? http://developer-drupal.web.prod.ext.phx2.redhat.com/node/add/books	1	1	https://issues.jboss.org/browse/DEVELOPER-4012
6548	DEVELOPER-4010	[~adelaarreola] Sorry, but I'm not able to see the downloads section via the provided URL. Is there another way for [~hwindzio] and I to review? !eap-no-downloads.png|thumbnail!	1	1	https://issues.jboss.org/browse/DEVELOPER-4010
6549	DEVELOPER-4010	[~tnolan] Here is a screenshot https://www.dropbox.com/s/crjv5s17r4kxbkh/Screenshot%202017-04-12%2008.39.34.png?dl=0 The 'Try It' only shows up on pages with download links, so for example, it wouldn't show in /openshift/download page. Is that correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-4010
6550	DEVELOPER-4010	[~adelaarreola] The font size for TRY IT in the screen capture looks slightly larger than the ALL DOWNLOADS. These should be the same and use the existing ALL DOWNLOADS size. Approved once this change is made. No need for additional review/sign-off. bq. The 'Try It' only shows up on pages with download links, so for example, it wouldn't show in /openshift/download page. Is that correct? Correct. Middleware products are the main focus of this for now.	1	1	https://issues.jboss.org/browse/DEVELOPER-4010
6551	DEVELOPER-4008	[~lkoranda] This issue captures the work to get the configuration for our Jenkins slave in the Boston labs managed by Ansible. Is there anything else that you would like to add to this story? cheers, Rob	1	1	https://issues.jboss.org/browse/DEVELOPER-4008
6552	DEVELOPER-4002	[~paul.robinson], how much do we care about having links to the contributor's/author's JBoss Developer profile? We actually don't have that information currently in Drupal and going out and fetching it like we do in awestruct is probably going to be complicated.	1	1	https://issues.jboss.org/browse/DEVELOPER-4002
6553	DEVELOPER-4002	[~paul.robinson] also, do we want to display the commit history as well?	1	1	https://issues.jboss.org/browse/DEVELOPER-4002
6554	DEVELOPER-3997	Didn't someone in IT or DevOps build something already that monitors our page score? I thought they were saying they wanted to build out email notifications too.	1	0	https://issues.jboss.org/browse/DEVELOPER-3997
6555	DEVELOPER-3995	Not a lot we can do with this until we have some context. Is this going into the next sprint or shall we bump it?	0	0	https://issues.jboss.org/browse/DEVELOPER-3995
6556	DEVELOPER-3994	[~lightguard], I've updated the description. Does that make more sense now? I'm not fully sure what Caffelli are suggesting we do. Hoping you or [~ldary] understand this.	1	1	https://issues.jboss.org/browse/DEVELOPER-3994
6557	DEVELOPER-3994	[~lightguard], [~paul.robinson]: Is this a Drupal or an export thing?	1	1	https://issues.jboss.org/browse/DEVELOPER-3994
6558	DEVELOPER-3994	As component lead, I think [~ldary] needs to refine this issue. [~ldary]: Can you refine before the end of this week?	1	1	https://issues.jboss.org/browse/DEVELOPER-3994
6559	DEVELOPER-3991	Are you sure you want to take this one, Paul?	1	1	https://issues.jboss.org/browse/DEVELOPER-3991
6560	DEVELOPER-3991	[~lightguard] I think this is a Drupal thing, right? Can you schedule it for completion before the end of the quarter? That's when we want all the SEO fixes done by.	1	1	https://issues.jboss.org/browse/DEVELOPER-3991
6561	DEVELOPER-3978	You can see the videos I created at https://drive.google.com/drive/u/1/folders/0Bwx86SOh3ez-TEotVjlXYTJheIU [~mikgue] has looked at the first one I did. [~paul.robinson], [~mikgue] what else would you two like to see?	1	1	https://issues.jboss.org/browse/DEVELOPER-3978
6562	DEVELOPER-3974	[~adelaarreola] I'm not sure if you've pinged [~james.cobb] about this yet, but do you need more specific details for this beyond what is documented in the Unified Style Guide?	1	1	https://issues.jboss.org/browse/DEVELOPER-3974

#	Requirement	Questions	Answered	Successfully Answered	URL
6563	DEVELOPER-3974	[~tnolan] I think I got it :) [~james.cobb] Here is a link to review the CTAs: http://stumpjumper.lab4.eng.bos.redhat.com:36819/ctas The two things that were not clear from the guide were if the black CTA's have a hover color and the same thing for the white *light* CTA. Please let me know if I need to change anything. That /ctas page is only for the review of this Jira, it will be removed but the styles will stay once approved.	1	1	https://issues.jboss.org/browse/DEVELOPER-3974
6564	DEVELOPER-3974	[~adelaarreola] [~james.cobb] Thanks Adela. Heavy and Medium CTAs - The button sizes look off to me as the Medium CTA ends up being larger in height (46px vs 37px). I know you based this on the specs in the Style Guide so I'd like a few days to take this back to the UX Working group for review. I think we need to update the guide to include more information. The line height seems to be making the biggest difference for these. Light CTAs - Please make these bold (font-weight: 800;). I just added this info to the guide. I think the white should use the lighter hyperlink blue (#73BCF7) and black should use the default hyperlink blue (#0066CC) for the hover. Refer to the Tried.Tested. Trusted section on .com for examples: https://www.redhat.com/en	1	1	https://issues.jboss.org/browse/DEVELOPER-3974
6565	DEVELOPER-3974	[~adelaarreola] I heard back from the .com team and they are using line-height: 1.44 for both heavy and medium CTA, which is their default body style. We're updating the guide to reflect this. The difference between the medium CTA and the heavy CTA then is the the medium one is normally just the border like the Medium CTA outline ones you currently have on the ctas example page. I'm not sure we have plans to use a medium CTA solid at this time on this site. If we did, these would be a different color palette than the heavy CTAs. [~james.cobb] - Do you agree that we don't need to define the medium CTA solid styles/colors as part of this story?	1	1	https://issues.jboss.org/browse/DEVELOPER-3974
6566	DEVELOPER-3974	[~adelaarreola] - desktop looks fine. On mobile though... 1) Shouldn't wrap text. 2) [~tnolan] - I'm a fan of full width buttons on the mobile experience although RH.com doesn't do this. Should we advocate that or default to their style?	1	1	https://issues.jboss.org/browse/DEVELOPER-3974
6567	DEVELOPER-3974	bq. I'm a fan of full width buttons on the mobile experience although RH.com doesn't do this. Should we advocate that or default to their style? [~james.cobb] I'd like to go with the recommendation from the Unified Style Guide (maintain the button size across devices) for now so we're consistent with the other sites. If this is something you feel strongly about, it would be great to have you come to the next UX Working Group meeting to bring it up. I can see both sides. Other option is to use the recommendations for now and A/B test this. That way we'd have data to support a change if the results came back in favor of the full width button. This is something we could execute on pretty quickly if we found a good page to test this with. I can also check with Patty Behre to see if they've done testing already about the button. The .com team has a Mojo page for their tests which has some really interesting info.	1	1	https://issues.jboss.org/browse/DEVELOPER-3974
6568	DEVELOPER-3973	[~james.cobb] [~ecurley] [~adelaarreola] Nice! I really like the v8, but the "Interested in..." is still getting lost on me. What would it look like if you removed that from the blue box and put it underneath?	1	1	https://issues.jboss.org/browse/DEVELOPER-3973
6569	DEVELOPER-3973	[~adelaarreola][~tnolan][~ecurley]- We could remove the text out of the excerpt box but we'll have to be careful of the wording so we aren't redundant with "download" OPTION 4: Want to see more? Get the full cheat sheet. [LOG IN TO DOWNLOAD] Personally, I don't think we need the text as we say "Cheat Sheet Excerpt" with a quote mark/styled box and a CTA button. If folks can't figure that out, they may not be smart enough to use the cheat sheet. :P	1	1	https://issues.jboss.org/browse/DEVELOPER-3973
6570	DEVELOPER-3970	[~akandel], [~visuale] can I get an update on this please?	1	1	https://issues.jboss.org/browse/DEVELOPER-3970
6571	DEVELOPER-3970	[~paul.robinson] I'm still stuck trying to figure out how to get this banner to show. As I pointed out in an email yesterday or day before, I do actually cause it to show up, but only when I hit the 404 page (using local, non-EU proxies with or without the VPN on). Is it possible this is related to the fact that I see two cache headers at odd with one another (private and public)?	1	1	https://issues.jboss.org/browse/DEVELOPER-3970
6572	DEVELOPER-3970	[~paul.robinson] if I can't make the RH/Adobe-based cookie banner work for this, would it be ok to go with a PHP/JS-based workaround? For instance, I saw this plugin while poking around the other day: [https://wordpress.org/plugins/creare-eu-cookie-law-banner/|https://wordpress.org/plugins/creare-eu-cookie-law-banner/]	1	1	https://issues.jboss.org/browse/DEVELOPER-3970
6573	DEVELOPER-3970	[~ldary] I don't know it's something I'm totally missing but that same error as before continues and as a consequence the banner is still not coming up on the test link that was provided. I can't find anything in Stack Overflow about this other than it might be a timing issue. I'm not sure how I might be able to adjust that as it's coming from an external source. Is there anything you could think of that I might be missing? If all else fails, I've tested out "plan B" a little (on a test site), wherein I use a WP plugin with formatting to look consistent with the RH/Adobe EU Cookie banner.	1	1	https://issues.jboss.org/browse/DEVELOPER-3970
6574	DEVELOPER-3970	Yes, I think that errors gone but Adobe might be doing something on their end and once fixed we should know one way or another whether it works or if I go to plan B. Am I right about this [~ldary] ?	1	1	https://issues.jboss.org/browse/DEVELOPER-3970
6575	DEVELOPER-3966	[~ldary] I've fleshed this out a bit now. Some of the issues still need some figuring out before they can be implemented though. Can you take over the coordination of this now please?	0	0	https://issues.jboss.org/browse/DEVELOPER-3966
6576	DEVELOPER-3946	I think it;'s something for my team. Basically we need to make a page like https://github.com/redhat-developer/developers.redhat.com/blob/master/register.html.slim	1	1	https://issues.jboss.org/browse/DEVELOPER-3946
6577	DEVELOPER-3946	[~mikgue] when do you need this for?	0	0	https://issues.jboss.org/browse/DEVELOPER-3946
6578	DEVELOPER-3933	[~tnolan], [~hwindzio]: This needs a link to the PDD and a priority. Do we have this information?	0	0	https://issues.jboss.org/browse/DEVELOPER-3933
6579	DEVELOPER-3922	[~ldary] I agree with your concerns. I'm a little confused as to why we needed this? Was this to help fix for the failing downloads in PR environment? i.e. not being redirected to PR environment when attempting to download? (and download not triggering). If so, it still fails for me: {code} Not Found The requested URL /pr/1796/export/downloads/products/cdk/hello-world/ was not found on this server. Additionally, a 404 Not Found error was encountered while trying to use an ErrorDocument to handle the request. {code} {code} https://developers-pr.stage.redhat.com/pr/1796/export/downloads/products/cdk/hello-world/?tcUser=Wendell+Braun&tcWhenSigned=January+17,+2017&tcWhenEnds=January+17,+2018&tcEndsIn=308&tcDuration=365&tcDownloadFileName=cdk-2.4.0.zip&tcRedirect=5000&tcSrcLink=https://developers.stage.redhat.com/download-manager/content/origin/files/sha256/ec/ec9169717d9785d3df6254d55a340f2a9af25b9dc446a680b376a5303028ac0c/cdk-2.4.0.zip&p=Red+Hat+Container+Development+Kit+(CDK)&pv=2.4&tcDownloadURL=https%3A%2F%2Faccess.cdn.redhat.com%2Fcontent%2Forigin%2Ffiles%2Fsha256%2Fec%2Fec9169717d9785d3df6254d55a340f2a9af25b9dc446a680b376a5303028ac0c%2Fcdk-2.4.0.zip%3F_auth_%3D1489616157_8a83ba5b8910cf38bb5ebdc36a7884e8 {code}	1	1	https://issues.jboss.org/browse/DEVELOPER-3922
6580	DEVELOPER-3906	[~hwindzio] [~paul.robinson] I wonder if this one can be marked as done since the .io team is handling what I believe this epic was supporting for Summit? Do we have any specific work for the RHDP site?	1	0	https://issues.jboss.org/browse/DEVELOPER-3906

#	Requirement	Questions	Answered	Successfully Answered	URL
6581	DEVELOPER-3906	i have to admit i have no visibility into the openshift.io work, and how the site should tie into this. I'm assuming openshift.io will have a landing page post-summit, but that doesn't deserve an epic. If we close this, can we reopen it later if needed? If yes, i suggest we close it.	1	0	https://issues.jboss.org/browse/DEVELOPER-3906
6582	DEVELOPER-3904	hrm...i think it should be either [~tnolan] or I. And Tiffany seems exceptionally busy and could use the help. Is it possible to have both of us responsible?	1	1	https://issues.jboss.org/browse/DEVELOPER-3904
6583	DEVELOPER-3900	Where should the CTA link to?	1	1	https://issues.jboss.org/browse/DEVELOPER-3900
6584	DEVELOPER-3894	BOMs and Architypes complicated this a bit. [~paul.robinson] [~tnolan] [~james.cobb] where can I see the UI and UX for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-3894
6585	DEVELOPER-3881	Should this issue be moved to https://issues.jboss.org/projects/RHDEVDOCS or DEVELOPER? cc: [~paul.robinson] [~rkratky]	1	0	https://issues.jboss.org/browse/DEVELOPER-3881
6586	DEVELOPER-3881	[~unibrew] I think you are tracking this work in RHDENG. Is that right?	1	1	https://issues.jboss.org/browse/DEVELOPER-3881
6587	DEVELOPER-3881	[~paul.robinson] I'm wondering what the next steps are for this epic? Will your team be working with James on getting the header and footer in the docs? Ideally we'll want this completed before the docs show up in the search results page that [~tnolan] is working on.	1	0	https://issues.jboss.org/browse/DEVELOPER-3881
6588	DEVELOPER-3881	Yes, I think that's what needs doing by my team. It looks like we didn't plan this into the work for this quarter. Do you want to add it in and re-prioritise the current work?	1	1	https://issues.jboss.org/browse/DEVELOPER-3881
6589	DEVELOPER-3880	Thanks Axel. I will be sure to let Caffelli know this. Can we correct this today? I will then have Caffelli send a report. Do you think a meeting with FB help would be advantageous?	0	0	https://issues.jboss.org/browse/DEVELOPER-3880
6590	DEVELOPER-3878	[~rblake] My plan for the broken link check was to do something similar to what you did with the acceptance tests. 1. create _blcheck directory - create an environments dir containing docker-compose file. - Create Dockerfile containing the nmp download of blc and stuff it requires. 2. Create ruby scripts with a `_blcheck/run_blc.rb` and `blc_test_options.rb` that is used to run the tests (either run within docker or outside?). {code} opts.on('--update-github-status SHA1', String, 'The SHA to report test results for on GitHub') do |sha1| test_configuration[:github_sha1] = sha1 end {code} 3. Unit tests. 4. Catch up with you regarding the job creation. Initially pointing the results of the tests to the console output until we have better reporting. Was this what you had in mind? I'll catch up with you tomorrow.	1	1	https://issues.jboss.org/browse/DEVELOPER-3878
6591	DEVELOPER-3878	[~paul.robinson] Just an update on this ticket since our last 1-1: I decided to resurrect Blikr I think for what we need it does exactly that. It just needs a little work. I am currently refactoring the code, I added RuboCop (static code analyser) and there were > 700 offences. Mainly due to the syntax being more for Ruby 1.8.7. Given we're now on Ruby 2.4.1 it's time to update. There are also no unit-tests so It's been difficult to refactor it without knowing if I've broken it. What am I doing and what do I plan to do? In priority order: 1. Fix Rubocop errors (almost done). 2. Add unit-tests (should have been #1 however thought best update code first) 3. Add options for --ignore-external, --ignore-internal links. 4. Dockerize 5. Update Gem README.md and release an update version.	0	0	https://issues.jboss.org/browse/DEVELOPER-3878
6592	DEVELOPER-3878	[~paul.robinson] the work for this is complete. However, after a discussion with [~rblake] we recommends that I finish the work on [fixes/improvements to Blinkr|https://issues.jboss.org/browse/DEVELOPER-4138]. Shall I put this in blocked?	0	0	https://issues.jboss.org/browse/DEVELOPER-3878
6593	DEVELOPER-3877	[~tnolan] can you provide the UX for this? Also, what page types should this be applied to? All?	1	1	https://issues.jboss.org/browse/DEVELOPER-3877
6594	DEVELOPER-3873	[~tnolan] Dan is working on getting the cheat sheets on the dcp so I should be able to start working on this ticket soon. Is the wireframe for this approved?	1	0	https://issues.jboss.org/browse/DEVELOPER-3873
6595	DEVELOPER-3873	[~paul.robinson] I added a link to the InVision project. Do you know who on your team will be working on this? Both [~james.cobb] and I can add some developer notes to the InVision mockup and tag him/her.	1	1	https://issues.jboss.org/browse/DEVELOPER-3873
6596	DEVELOPER-3870	[~lightguard], [~rblake], [~ldary]: Who do you think should own the requirements gathering for this work? I'm thinking it's one of you three?	1	1	https://issues.jboss.org/browse/DEVELOPER-3870
6597	DEVELOPER-3868	[~hwindzio] can you review these requirements. Are they approved?	1	0	https://issues.jboss.org/browse/DEVELOPER-3868
6598	DEVELOPER-3868	[~hwindzio] I'd like us to push on with this when [~akandel] returns from PTO (8th May). Can you review the requirements by then? It's hard for us to press on without firm requirements.	1	0	https://issues.jboss.org/browse/DEVELOPER-3868
6599	DEVELOPER-3868	A couple of comments: * The contact sales button brings you here: https://www.redhat.com/en/about/contact/sales. It is part of the sub-header on the product pages. Not sure why that link shows the link to the 'Help' tab. * "Register for contributor program" - Agreed, ideally we see when they register. But, if we can't attach to submit button then just starting the process is probably good enough. Later, we can see the drop off rate for the form submissions. * "CTA for consulting, training, certification" - In addition, I think we will want to see the engagement broken out for each of these. A couple of additional things we wanted to track, based on the original spreadsheet: *Viewing a StackOverflow question on devs.rh.com. I would think that is similar to viewing a KCS article. Can we track that as a level of engagement? https://developers.redhat.com/stack-overflow/#!q= *Posting a StackOverflow question. We should have their stackID if they are logging in using social log in - StackOverflow. We can then compare that to the post that is being slurped in to see if they have posted the question for a red hat product tag. *Viewing a video at least 25%: https://developers.redhat.com/resources/#!type=video	1	0	https://issues.jboss.org/browse/DEVELOPER-3868
6600	DEVELOPER-3868	{quote} * The contact sales button brings you here: https://www.redhat.com/en/about/contact/sales. It is part of the sub-header on the product pages. Not sure why that link shows the link to the 'Help' tab. {quote} The example page, provides an example location of where this activity can be started from. I've clarified the description, which should make this clearer. {quote} * "CTA for consulting, training, certification" - In addition, I think we will want to see the engagement broken out for each of these. {quote} Done. {quote} *Viewing a StackOverflow question on devs.rh.com. I would think that is similar to viewing a KCS article. Can we track that as a level of engagement? https://developers.redhat.com/stack-overflow/#!q= {quote} Added. {quote} *Posting a StackOverflow question. We should have their stackID if they are logging in using social log in - StackOverflow. We can then compare that to the post that is being slurped in to see if they have posted the question for a red hat product tag. {quote} This will be hard as we don't have a way of getting this data into Adobe Analytics. Also, I would argue that the user is engaging with StackOverflow, not the RHDP site. {quote} *Viewing a video at least 25%: https://developers.redhat.com/resources/#!type=video {quote}	0	0	https://issues.jboss.org/browse/DEVELOPER-3868
6601	DEVELOPER-3862	Just happened to see this JIRA. One caution from UX is that many users have their full email address as their username. What happens with longer "usernames" on Mobile? What use cases/scenarios have you explored for this? FYI [~hwindzio]	1	1	https://issues.jboss.org/browse/DEVELOPER-3862
6602	DEVELOPER-3862	[~velias] Would you be able to attach a screenshot? I'm just curious how it looks.	1	1	https://issues.jboss.org/browse/DEVELOPER-3862

#	Requirement	Questions	Answered	Successfully Answered	URL
6603	DEVELOPER-3854	[~adelaarreola] [~ldary] [~dcoughlin1] [~rblake] would you all try running the following SQL and letting me know if everything is correct as per the comments? {code:sql} -- Removing videos BEGIN; CREATE TEMPORARY TABLE video_id (nid INT(10), INDEX(nid)) AS SELECT SUBSTR(source, 7) AS nid FROM `url_alias` WHERE alias LIKE "/video%"; -- These should all return >= 1 (1 due to node_revision__body) SELECT COUNT(1) FROM `node` WHERE nid IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node_field_revision` WHERE nid IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node__body` WHERE entity_id IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node_field_data` WHERE nid IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node_revision__body` WHERE entity_id IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node_revision` WHERE nid IN (SELECT nid FROM video_id); SELECT COUNT(1) AS nid FROM `url_alias` WHERE alias LIKE "/video%"; DELETE FROM `node` WHERE nid IN (SELECT nid FROM video_id); DELETE FROM `node_field_revision` WHERE nid IN (SELECT nid FROM video_id); DELETE FROM `node__body` WHERE entity_id IN (SELECT nid FROM video_id); DELETE FROM `node_field_data` WHERE nid IN (SELECT nid FROM video_id); DELETE FROM `node_revision__body` WHERE entity_id IN (SELECT nid FROM video_id); DELETE FROM `node_revision` WHERE nid IN (SELECT nid FROM video_id); DELETE FROM `url_alias` WHERE alias LIKE "/video%"; -- All of these should return 0 SELECT COUNT(1) FROM `node` WHERE nid IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node_field_revision` WHERE nid IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node__body` WHERE entity_id IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node_field_data` WHERE nid IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node_revision__body` WHERE entity_id IN (SELECT nid FROM video_id); SELECT COUNT(1) FROM `node_revision` WHERE nid IN (SELECT nid FROM video_id); SELECT COUNT(1) AS nid FROM `url_alias` WHERE alias LIKE "/video%"; DROP TABLE video_id; COMMIT; {code}	1	1	https://issues.jboss.org/browse/DEVELOPER-3854
6604	DEVELOPER-3854	[~lightguard] I assume you want this tested against the full database dump and not the "lite" version? My other comment here is a more general "how does the project make database schema changes now and moving forwards". As with most things, what I'd like to see is automation of schema changes and testing of changes before they get applied to production. I'd also like visibility and an audit trail of changes to production schema. So: * Changes to the production database schema are stored in version control using a defined naming pattern * Changes to database schema are raised as pull requests that are reviewed before being applied * Changes to database schema are applied in a testing environment first before being applied to a production database * Changes to database schema are applied using an automated tool and never by hand Coming from a Java background, I have always used [Liquibase|https://www.liquibase.org] to perform database migrations, but I see that PHP offers something similar e.g. [Phinx|https://phinx.org/]. My suggestion would be to use a PHP tool so that we can keep the Drupal container footprint to a minimum. From initial reading it also looks like Phinx is installed using Composer as well, so again that fits in nicely with the way we manage Drupal. Using a tool like Phinx would mean that we can store the DB migrations in Git and then have them automatically applied as part of the Drupal boot sequence. So with some tweaks to the drupal_install_checker.rb, update flow could be: # Start the container # run drush cim # run drush updb # run phinx schema migrations Doing that means that we can automate schema migrations and have them tested in our normal PR build pipeline before applying them to the production database. It feels beyond the scope of this story to do the above. However I think we need to consider it sooner rather than later so that we have an audit trail of the changes that have been made to the Drupal database over time and is it worth playing another user story before this one to setup a database schema migration pattern for Drupal?. [~paul.robinson] what are your thoughts here?	1	1	https://issues.jboss.org/browse/DEVELOPER-3854
6605	DEVELOPER-3847	[~ecurley] Looking at the styling for the excerpt block that [~james.cobb] provided compared (blue background) with your selected excerpts from the cheat sheets (also has a mix of blue background), we may need to modify the excerpt block background color on the page. Also, with the tables from the actual cheat sheets, this does add some complexity as to who can enter that info in the Drupal field. I'm guessing HTML knowledge may be needed to do this. I think it would be good to circle back with [~adelaarreola] on this next week. Adela - Maybe work on the most difficult looking of these first so we can see the worst case scenario and adjust things if needed? Who knows...maybe this isn't an issue?	1	1	https://issues.jboss.org/browse/DEVELOPER-3847
6606	DEVELOPER-3847	[~adelaarreola] [~james.cobb] [~tnolan] So trying to understand correctly, no tables, no headings, nothing that has a two column layout. Is that correct? If yes, that is going to be a bit challenging because many of these don't have paragraphs and the paragraphs they do are often just a basic intro info - not a great excerpt. Thoughts?	0	0	https://issues.jboss.org/browse/DEVELOPER-3847
6607	DEVELOPER-3846	[~paul.robinson] [~adelaarreola] - I updated the staging query to show what the cheatsheets dcp config looks like. Can you review at http://dcp.stage.jboss.org/v2/rest/search?sys_type=cheatsheet&field=_source? Staging doesn't seem to have cheatsheets by default, so you may need to add one to test if a merge overwrites the cheatsheet I added.	1	1	https://issues.jboss.org/browse/DEVELOPER-3846
6608	DEVELOPER-3846	Can this be closed? I see it up in both staging and production.	0	0	https://issues.jboss.org/browse/DEVELOPER-3846
6609	DEVELOPER-3844	[~paul.robinson] I fixed a spelling error in the URL for "cheet". Since Cheat Sheet is 2 words, shouldn't the URL be "/cheat-sheet/" rather than "cheatsheet". If we make it one word in the URL, I think it might be confusing and add more confusion to the "is it one word or two words" question. I know I get confused.	1	1	https://issues.jboss.org/browse/DEVELOPER-3844
6610	DEVELOPER-3844	[~adelaarreola] - Can you tell us the differences between the widget for www and the plugin on wordpress? Can we port that one over... or style the widget on www to match the blog?	1	1	https://issues.jboss.org/browse/DEVELOPER-3844
6611	DEVELOPER-3843	[~james.cobb] Which image usage? Can this be resolved today? I'd like the engineering work to go into the sprint starting Monday.	1	1	https://issues.jboss.org/browse/DEVELOPER-3843
6612	DEVELOPER-3843	OK. Sounds like it shouldn't hold up Engineering work? Do you expect the UX/VD to change beyond what image appears there?	1	1	https://issues.jboss.org/browse/DEVELOPER-3843
6613	DEVELOPER-3843	[~james.cobb] [~tnolan] where can I find the VD or layout (if styles don't change) for the tablet and mobile sizes? Also, can I get the specifications of text sizes, images heights/widths, colors, if buttons are heavy CTAs or light, etc so I can implement it as best as possible?	1	1	https://issues.jboss.org/browse/DEVELOPER-3843
6614	DEVELOPER-3840	In Drupal's core developers' infinite wisdom (and after how many rewrites??) they have decided not to give you a means of getting all revisions from an entity. What you need to do is create an query to pull in all the revisions (because....?). Here's the basic idea: {code:php} $node_storage = \Drupal::entityTypeManager()->getStorage('node'); $nids = \Drupal::entityQuery('node')->condition('type','page')->execute(); $nodes = $node_storage->loadMultiple($nids); foreach ($nodes as $node) { $result = $node_storage->getQuery() ->allRevisions() ->condition($node->getEntityType()->getKey('id'), $node->id()) ->sort($node->getEntityType()->getKey('revision'), 'DESC') ->pager(50) ->execute(); $revisions = array_keys($result); foreach ($revisions as $vid) { // version id $node_storage->deleteRevision($vid); } } {code} We could do this in a controller and create a route only admins can get to. We may also need to do that in a batching manner.	0	0	https://issues.jboss.org/browse/DEVELOPER-3840
6615	DEVELOPER-3838	[~ldary] will be working on this. Can you try discuss it with him today? We need to get the sprint planned for early Monday.	1	1	https://issues.jboss.org/browse/DEVELOPER-3838

#	Requirement	Questions	Answered	Successfully Answered	URL
6616	DEVELOPER-3818	Paul Robinson, thanks for the ordering fix! For pushing artifacts to DM I am creating issues in RCM Jira (like this one RCM-12018). AFAIK there is an "unpush" process for such artifacts. I was cautious to ask RCM for unpush to not end with a broken link for this artifact on "Downloads" page. Assuming that such changes are propagated to DM automatically, will it be convenient if I will create a RCM issue to unpush this specific artifact (8u111-3 ZIP bundle)?	1	1	https://issues.jboss.org/browse/DEVELOPER-3818
6617	DEVELOPER-3817	More work than you probably think is my guess. The proposed solution looks good and seems fine, but trying to come up with the correct implementation could take a bit. Balancing the ease of use for content creators and development work to implement it is not going to be easy. Do we add this as a widget you can use in the editor? Do we add another paragraph (to all of the steps) you can fill out? I think those two things need to be figured out first.	1	1	https://issues.jboss.org/browse/DEVELOPER-3817
6618	DEVELOPER-3817	[~lightguard] [~paul.robinson] Just a reminder that the MVP for this is some working HTML/code that can be inserted in the existing form field. Sure, it would be great to have something easier down the road, but I don't want that to hold up getting something out soon. The only page we currently need this on is RHEL. CDK has something similar, but with only 3 options vs 22, we don't need the drop-down there. Is there existing show/hide functionality like what is written for the blue bars on that page that can be leveraged? Will the text editor strip out class tags for the divs if we use that as part of the solution? Just did a quick test and the editor did not strip class tags from div. However it did remove the <option> tag. [~lightguard] Can that one be added to the RHD Approved list?	1	1	https://issues.jboss.org/browse/DEVELOPER-3817
6619	DEVELOPER-3817	[~ldary] This is looking really good. I have a working version on the RHEL page in this environment now: http://stumpjumper.lab4.eng.bos.redhat.com:36748/products/rhel/hello-world#fndtn-bare-metal_build-something-today The only thing that is missing the the divider line between the drop-down and the content. Can this be added to the styling for the drop-down? !2017-02-09_10-40-12.png|thumbnail! Once that's added, I think this is good to go after a quick review/final approval from [~james.cobb] for V+D.	1	1	https://issues.jboss.org/browse/DEVELOPER-3817
6620	DEVELOPER-3810	Have you contacted Wordpress VIP? It's a real pain we can't redirect https://developerblog.redhat.com.	1	1	https://issues.jboss.org/browse/DEVELOPER-3810
6621	DEVELOPER-3810	Can the admin console be accessed by using a different domain/IP in the URL?	1	1	https://issues.jboss.org/browse/DEVELOPER-3810
6622	DEVELOPER-3804	[~adelaarreola] I've marked you as the lead on this Epic. Can you work with [~lkrzyzanek] when he is ready to add this to the forums. Please help him to ensure the right collector is used in the footer vs 404 page.	0	0	https://issues.jboss.org/browse/DEVELOPER-3804
6623	DEVELOPER-3803	Hi [~visuale]. This is implemented on the main site pages. Can you make the same change on the blog?	1	1	https://issues.jboss.org/browse/DEVELOPER-3803
6624	DEVELOPER-3803	Actually the link is there, but it's not working properly. Can you coordinate with [~adelaarreola] and [~paul.robinson] to get this fixed?	1	1	https://issues.jboss.org/browse/DEVELOPER-3803
6625	DEVELOPER-3803	[~visuale] this is still broken, can you take a look?	1	1	https://issues.jboss.org/browse/DEVELOPER-3803
6626	DEVELOPER-3803	[~visuale] this doesn't seem to work anymore. Can you investigate?	1	1	https://issues.jboss.org/browse/DEVELOPER-3803
6627	DEVELOPER-3795	[~james.cobb] Per notes, I've revised this script. Would the attached screenshot from a local test be the correct presentation for list items? !Screen Shot 2017-05-21 at 5.35.04 PM.png|thumbnail!	1	1	https://issues.jboss.org/browse/DEVELOPER-3795
6628	DEVELOPER-3792	[~james.cobb] I believe this behavior is already in place across the site. Do you know anywhere in specific where this behavior is needed?	1	1	https://issues.jboss.org/browse/DEVELOPER-3792
6629	DEVELOPER-3792	[~adelaarreola] - The current buttons on our site are set to 100% for width on desktop. They should be actually have a fixed padding on the left or right as per the document. the "Heavy CTA" = Red buttons. Button dimensions: • Button width flows with length of text • Padding: Top is 9 pixels; bottom is 9 pixels; sides are 40 pixels Examples where it's not working correctly... 1) Homepage "join" button at the bottom (has overrides) 2) Product page "Hello World" section (has overrides) 3) Product detail download page (has overrides) 3) Downloads page "Popular Downloads" (has overrides) 4) Downloads page "Download" buttons (has overrides) -These should actually be medium CTAs. 5) Events page "learn more" buttons (has overrides) etc... etc.	1	0	https://issues.jboss.org/browse/DEVELOPER-3792
6630	DEVELOPER-3791	[~james.cobb] I believe this behavior is already in place across the site. Do you know anywhere in specific where this behavior is needed?	0	0	https://issues.jboss.org/browse/DEVELOPER-3791
6631	DEVELOPER-3781	[~ldary] I've marked you as lead on this Epic. Can you keep an eye on the issues and ensure a few of them get done each sprint? I'd like to see this done by the end of Q1FY17 if we can. You'll also need to work with the Blog, Forums and Keycloak engineers to get their changes in. Also note, that this has a very low PDD priority, so don't let it delay other Epic-level work.	0	0	https://issues.jboss.org/browse/DEVELOPER-3781
6632	DEVELOPER-3780	[~adelaarreola] [~paul.robinson] Thanks. I moved the doc (a copy) to: https://drive.google.com/drive/folders/0B_QaDtRa6sAkTTB2V2VFQmgwWUU Can one of you provide edit access so I can add a new tab and/or some notes?	1	1	https://issues.jboss.org/browse/DEVELOPER-3780
6633	DEVELOPER-3780	Thanks [~tnolan], can you work with [~adelaarreola] on this?	0	0	https://issues.jboss.org/browse/DEVELOPER-3780
6634	DEVELOPER-3768	When you say "mysql" are you talking about it starting up or actually running? If it's startup time, which it very well may be, that SQL file we're loading is nearly 7G.	1	1	https://issues.jboss.org/browse/DEVELOPER-3768
6635	DEVELOPER-3753	[~ldary] can you check if we need this or not? If not, I'd like to understand why not.	0	0	https://issues.jboss.org/browse/DEVELOPER-3753
6636	DEVELOPER-3749	[~paul.robinson] I believe this is in the actual code for the Drupal pages and not editable via the form for the product content type. If I'm setup with GitHub, is this a change that I could make myself via a PR?	1	1	https://issues.jboss.org/browse/DEVELOPER-3749
6637	DEVELOPER-3749	Thanks [~adelaarreola]. Can you add a comma before "which" to make this more grammatically correct? That was my bad for not seeing it sooner. Otherwise, this looks great. Approved. "and conditions of the Red Hat Developer Program{color:red}*,*{color} which provides no-cost subscriptions for development use only."	1	1	https://issues.jboss.org/browse/DEVELOPER-3749
6638	DEVELOPER-3738	[~paul.robinson] [~lightguard] Re: should every get-started page have it's first section marked as complete, when visiting it via a download? - No, since not all the products have a Get Started where the first step is to Download. In the future, I would like to get all of products using that format though. The checkbox Jason provided should give us the flexibility needed currently, as well as allow for future changes.	1	1	https://issues.jboss.org/browse/DEVELOPER-3738
6639	DEVELOPER-3738	[~tnolan] I see the 'Pending Download Section' option now. Have you confirmed that this does what you need it to do?	1	1	https://issues.jboss.org/browse/DEVELOPER-3738
6640	DEVELOPER-3714	how do we close JIRAs as duplicate and/or other?	1	1	https://issues.jboss.org/browse/DEVELOPER-3714
6641	DEVELOPER-3714	[~hwindzio] I've removed what I see as the duplicate tasks from this issue. See the description. There are still three places that the OpenShift product name appears outside of the drupalised product pages. Can this change go live now, or does it need to be done at the same time as the name change on the OpenShift product page?	1	1	https://issues.jboss.org/browse/DEVELOPER-3714
6642	DEVELOPER-3714	Yes we can do it at the same time. Can we try get this out today?	1	1	https://issues.jboss.org/browse/DEVELOPER-3714
6643	DEVELOPER-3708	Great, Can you create the Story(s) on this Epic and schedule them?	1	1	https://issues.jboss.org/browse/DEVELOPER-3708
6644	DEVELOPER-3704	[~hwindzio] when do you need this by? It will need a Mockup and maybe UX work. CC [~tnolan], [~james.cobb].	1	1	https://issues.jboss.org/browse/DEVELOPER-3704

#	Requirement	Questions	Answered	Successfully Answered	URL
6645	DEVELOPER-3702	[~adelaarreola] can you review this issue and makes sure you have all you need to implement it on Monday? It needs to go live on the 14th of Dec, do you have the time to achieve this?	1	1	https://issues.jboss.org/browse/DEVELOPER-3702
6646	DEVELOPER-3702	[~hwindzio] please review the deprecated text added to each forum: https://developers-pr.stage.redhat.com/pr/1663/export/forums/ Do you want the latest forum posts removing from the help pages of these four products also? I.e just show SO content on the help pages of these 4 products.	1	1	https://issues.jboss.org/browse/DEVELOPER-3702
6647	DEVELOPER-3702	[~hwindzio] This is ready to be merged. Do you want it to go live today or tomorrow?	0	0	https://issues.jboss.org/browse/DEVELOPER-3702
6648	DEVELOPER-3696	What technology tab? What page?	1	1	https://issues.jboss.org/browse/DEVELOPER-3696
6649	DEVELOPER-3696	[~lightguard] we should probably update the Product category name option here: https://github.com/redhat-developer/developers.redhat.com/blob/c37c7b1c9e5f6d3a8b410c6a2004769eae2930c5/_docker/drupal/drupal-filesystem/web/config/sync/field.storage.node.field_product_technology_group.yml#L30. Otherwise it's going to get confusing when assigning a Product to a category. Does this cause problems with any existing Product content using the old "Private Cloud" field name?	1	1	https://issues.jboss.org/browse/DEVELOPER-3696
6650	DEVELOPER-3688	[~rblake] had noticed a few things didn't match up in the PR with the "Definition of Done", good catch Rob, thanks! [~paul.robinson] where are we with the redirects?	1	1	https://issues.jboss.org/browse/DEVELOPER-3688
6651	DEVELOPER-3678	Can you add a description for why this is blocked?	1	1	https://issues.jboss.org/browse/DEVELOPER-3678
6652	DEVELOPER-3678	Is this still blocked?	0	0	https://issues.jboss.org/browse/DEVELOPER-3678
6653	DEVELOPER-3678	[~ldary] It may not be related, however I am still having issues getting results on /resources. On my local machine both inside and outside of docker testing against a pr-site preview it is fine! However when the tests run on Jenkins (PR build)I do not get results. I'm baffled! One thing I noticed with the stubbed responses was the search and stack overflow pages have an access-control-origin of `https://developers-pr.stage.redhat.com` However the /resources page has an access-control-origin of http://developers-pr.stage.redhat.com. Should that be https?	0	0	https://issues.jboss.org/browse/DEVELOPER-3678
6654	DEVELOPER-3677	Is this still blocked?	0	0	https://issues.jboss.org/browse/DEVELOPER-3677
6655	DEVELOPER-3673	[~paul.robinson] some products have multiple videos, is there a list or something where I can see which video is a product intro?	1	1	https://issues.jboss.org/browse/DEVELOPER-3673
6656	DEVELOPER-3673	I added an extra step above: {quote} Change the thumbnail to a frame that has the person's name displayed and one in which the interviewee doesn't look silly. You will likely need to pause the video at a suitible point and upload a screenshot to use as the thumbnail. I couldn't find a better way to do this with the YouTube API. {quote} Can you make this change too?	0	0	https://issues.jboss.org/browse/DEVELOPER-3673
6657	DEVELOPER-3670	When can you get this done? The Getting Started instructions now point to a broken link as 9.0.2 was removed.	1	1	https://issues.jboss.org/browse/DEVELOPER-3670
6658	DEVELOPER-3655	[~lightguard] is this feasible? Can you schedule it.	1	0	https://issues.jboss.org/browse/DEVELOPER-3655
6659	DEVELOPER-3655	[~mikgue] we can now archive items. Everything that was "deleted" before has now been archived. Content Admins can bring content back from being archived. Are you happy with this?	1	1	https://issues.jboss.org/browse/DEVELOPER-3655
6660	DEVELOPER-3655	[~mikgue] how does https://drive.google.com/open?id=0B9M2ZSLcm2xXcFpLY1MxTU5QQU0 help?	0	0	https://issues.jboss.org/browse/DEVELOPER-3655
6661	DEVELOPER-3650	Sooner than later. Is it possible for this or next sprint?	0	0	https://issues.jboss.org/browse/DEVELOPER-3650
6662	DEVELOPER-3643	[~ldary] I've just merged the PR. Is this ready to mark as done?	1	1	https://issues.jboss.org/browse/DEVELOPER-3643
6663	DEVELOPER-3643	[~ldary] I think this is ready to resolve. Is that right?	0	0	https://issues.jboss.org/browse/DEVELOPER-3643
6664	DEVELOPER-3642	[~ldary] is this now done? If so, can you update the issue status please.	1	1	https://issues.jboss.org/browse/DEVELOPER-3642
6665	DEVELOPER-3642	[~ldary] with https://github.com/redhat-developer/developers.redhat.com/pull/1600 merged, is this now done?	0	0	https://issues.jboss.org/browse/DEVELOPER-3642
6666	DEVELOPER-3628	I think it's good enough for you or Rob to execute. [~rblake], do you agree?	0	0	https://issues.jboss.org/browse/DEVELOPER-3628
6667	DEVELOPER-3626	[~adelaarreola], [~lightguard]. If we stop pushing the homepage via Awestruct (as we have planned), how will the hero background image be changed? I believe it is currently part of the style. Can this particular style be in-lined in the HTML, such that it can be edited via Drupal when editing the homepage HTML?	1	1	https://issues.jboss.org/browse/DEVELOPER-3626
6668	DEVELOPER-3626	[~james.cobb] we can up the JBD banner also as the build is now working. Can your team do that via a PR?	1	1	https://issues.jboss.org/browse/DEVELOPER-3626
6669	DEVELOPER-3622	[~adelaarreola] do you know why this behaviour has been reverted? I'm sure you fixed this a while ago. Do you remember?	1	1	https://issues.jboss.org/browse/DEVELOPER-3622
6670	DEVELOPER-3622	When we switched to Drupal, the default email reverted to the old email. Could this be the same cause?	1	1	https://issues.jboss.org/browse/DEVELOPER-3622
6671	DEVELOPER-3617	I've just re-ran the report. See: https://docs.google.com/spreadsheets/d/16bqTsaUr-PAJMvM2RdWlEsXnPED0YWanIk6G-8uW82I/edit#gid=1350294889. It looks like the articles are still showing http links. Can you take a look? Maybe something funny is happening with the export?	1	1	https://issues.jboss.org/browse/DEVELOPER-3617
6672	DEVELOPER-3614	What is email needed for? Is it part of the content change workflow?	0	0	https://issues.jboss.org/browse/DEVELOPER-3614
6673	DEVELOPER-3587	[~ihamilto] I understand the slowness of this page is affecting the tests. How urgent is it that we fix this?	1	1	https://issues.jboss.org/browse/DEVELOPER-3587
6674	DEVELOPER-3580	[~ldary] are you working on this for this sprint? If so, Can you add to the current sprint. We don;t normally add stuff after the start, but I prefer the sprint to reflect reality.	1	1	https://issues.jboss.org/browse/DEVELOPER-3580
6675	DEVELOPER-3578	[~tnolan] and/or [~james.cobb] could we get the information for these colors and icons for this (background, borders, text, fonts, etc.)? I didn't see them listed in the linked issue. Also, closing "X" vertical alignment? Seems like middle is the one on the Patternfly example, but wasn't sure because that wasn't specified in the RH Unified Style Guide (which I've made a comment on).	1	1	https://issues.jboss.org/browse/DEVELOPER-3578
6676	DEVELOPER-3578	[~james.cobb] Can you double check the way you labeled the colors? I think you have the Error and Warning mixed up above with the terms. I think it should be Warning (Yellow) and Error (Red). Color-wise, these all look good to me, except the yellow one. Is there more of a shade of orange rather than brownish that can be used for the icon/border? Since it's a warning, I feel like it should have a slightly brighter icon color to draw attention. The brownish gets kinda lost to me.	1	0	https://issues.jboss.org/browse/DEVELOPER-3578
6677	DEVELOPER-3578	[~ldary] Any update on this one? Is there anything else you need from the design team to get this implemented?	1	0	https://issues.jboss.org/browse/DEVELOPER-3578
6678	DEVELOPER-3578	[~paul.robinson] [~ldary] [~hwindzio] - Can we get this story added to the next sprint? I think the styles are already defined and approved by [~james.cobb] and it's just a matter of implementing it. I am going to need these styles as part of the progressive profiling work I'm doing, and the page designs for the registration flows.	1	0	https://issues.jboss.org/browse/DEVELOPER-3578
6679	DEVELOPER-3578	[~paul.robinson] [~ldary] Can you provide an update on this? Is it falling through the cracks since it has an old Sprint number already assigned to it? It would be very helpful to have these so we can utilize for both the progressive profiling work and the new site outage alerts.	1	0	https://issues.jboss.org/browse/DEVELOPER-3578

#	Requirement	Questions	Answered	Successfully Answered	URL
6680	DEVELOPER-3578	Thanks [~ldary]. I see that [~adelaarreola] is assigned to work on this. :) It would be great to have a document or mojo location where we documented these so that we can easily refer to them for implementation from other systems such as the Blog and KeyCloak. Do you have any suggestions for this?	1	0	https://issues.jboss.org/browse/DEVELOPER-3578
6681	DEVELOPER-3567	[~jessk] - I assigned this to [~brock] . Do we also want to add this to JBoss Developer's hero since it applies to JBoss?	1	1	https://issues.jboss.org/browse/DEVELOPER-3567
6682	DEVELOPER-3567	Thank you [~james.cobb] and [~brock] Do you think it will be hard to read? I'll check traffic to it with Tiffany in a week or so and see how it's doing. I do like that it matches the event page header though. [~paul.robinson] I will open a jira to implement this.	1	1	https://issues.jboss.org/browse/DEVELOPER-3567
6683	DEVELOPER-3560	[~adelaarreola] It looks really nice. Thanks so much for doing it quickly (so quickly!) One question - The pages that host the videos (like http://developers-pr.stage.redhat.com/pr/1481/build/2365/video/youtube/tEvzJBFIzV4/) have title and speaker names running together. See attached screenshot. Can we adjust the font for the speaker name or some other formatting to distinguish between the two? Thanks so much! !screenshot-1.png!	1	1	https://issues.jboss.org/browse/DEVELOPER-3560
6684	DEVELOPER-3560	Also, would it look better if "Red Hat Developers at JavaOne 2016" was white, larger and in the header graphic? I'm not sure how amenable the header is for this, due to the variation in colours used.	1	1	https://issues.jboss.org/browse/DEVELOPER-3560
6685	DEVELOPER-3560	[~adelaarreola]It looks great. Thank you. A few questions/comments... * I removed speakers names from the YouTube video titles. I moved them to the description area so we still had them. * The landing page that hosts the videos - hasn't updated yet....guessing it will just take some time? (http://developers-pr.stage.redhat.com/pr/1481/build/2372/video/youtube/35_yWad7IiQ/) * On the landing page that hosts the video I noticed there's a speaker field but we don't have speaker names there. How is that populated? Assuming it's being fed from the youtube data also? If the default is the YouTube account name (Red Hat Developers), Lets just remove it. Thanks again!	1	1	https://issues.jboss.org/browse/DEVELOPER-3560
6686	DEVELOPER-3555	[~paul.robinson] This looks like it should be a fairly quick task. Can it get added to a sprint and completed before March 1? We're still going to need this for when we combine the search/resources pages.	1	1	https://issues.jboss.org/browse/DEVELOPER-3555
6687	DEVELOPER-3554	[~paul.robinson] The SOLP is definitely the priority. However, we do have a lot of ads floating around that are driving people to the EAP Downloads page, specifically. If this could get implemented by the end of next week, I think that would be fine. Does that work from your end?	1	1	https://issues.jboss.org/browse/DEVELOPER-3554
6688	DEVELOPER-3553	[~james.cobb] what pages are you referring to, and which titles? Do you mean titles like: https://www.dropbox.com/s/0ct24nv1ndputik/Screenshot%202016-11-23%2012.48.28.png?dl=0 on http://developers.redhat.com/projects/?	1	1	https://issues.jboss.org/browse/DEVELOPER-3553
6689	DEVELOPER-3550	[~paul.robinson] It would be great if we could get this work prioritized for this week/next week. The urgency is because we'd like for this to be in place before all the promos for Don's book start going out. There's a big opportunity here to improve conversions. One other thing to note is the "Email Link" button. [~ecurley] and I talked about this along with Mark's team. For the initial release of that button, we'd like for it to be a mailto link that has some body text which includes a link back to the promo page. [~ecurley] - Would be able to provide the copy? Paul or someone on his team should be able to let you know what all the parameter options are such as subject line, body text, etc.	1	1	https://issues.jboss.org/browse/DEVELOPER-3550
6690	DEVELOPER-3550	[~paul.robinson] I'm not sure how to get links for [~tnolan] and [~james.cobb] to review, [this PR |https://github.com/redhat-developer/developers.redhat.com/pull/1572/files] has the CSS changes, and I have the html changes on my laptop. In order to get Tiffany a staging link to preview, I would need my PR to be merged into staging, but not production (it would possibly break the current microsite styles) and then I could go into the staging drupal gui and update the HTML pages. In other words, how should I handle the review of pages that were migrated as pure HTML?	1	1	https://issues.jboss.org/browse/DEVELOPER-3550
6691	DEVELOPER-3550	[~adelaarreola][~paul.robinson] Agreed on only having to implement the change in Drupal, not the other HTML pages. Do you have a URL [~james.cobb] and I can review?	1	1	https://issues.jboss.org/browse/DEVELOPER-3550
6692	DEVELOPER-3550	[~james.cobb] which version of fireforx are you running? It looks fine on mine https://www.dropbox.com/s/jklt9qug79666l6/Screenshot%202016-11-04%2007.46.39.png?dl=0	1	1	https://issues.jboss.org/browse/DEVELOPER-3550
6693	DEVELOPER-3550	[~adelaarreola][~tnolan] Desktop: the image usage is very large. We're going to set the rule that we size the image at 270px wide as the source. We'll need to adjust the columns to 6/18 so the spacing is ~30px between the image and the content. Mobile: 1) I've attached a screenshot to illustrate but the vertical spacing is quite inconsistent. We need to adjust the space above and below the content and the titles so it is consistent with the product mobile views. 2) Can we center the "email link" button below the download button. Global: We shouldn't force all caps on the title. Is this a drupal template change?	1	1	https://issues.jboss.org/browse/DEVELOPER-3550
6694	DEVELOPER-3550	[~adelaarreola] Just getting to this as I was in training today. What promotions are live with this change and is there HTML changes that need to occur in order for things like the email link to show up? Is this info documented anywhere so we don't have to bug you about it? :)	1	1	https://issues.jboss.org/browse/DEVELOPER-3550
6695	DEVELOPER-3550	[~tnolan] We have that doc here: https://docs.google.com/spreadsheets/d/1v-DNZvjpyw69Gx_2t2rgeALWjbt2_dnme6k2M5PXLS8/edit#gid=2016133270. Is this what you were thinking about?	1	1	https://issues.jboss.org/browse/DEVELOPER-3550
6696	DEVELOPER-3550	Thanks [~james.cobb]. Did you delete the old versions of the page too?	1	1	https://issues.jboss.org/browse/DEVELOPER-3550
6697	DEVELOPER-3540	[~james.cobb] - Is this still an issue? I've looked through a couple of the product pages and compared them on a local build vs. production. The pages look like they are pulling in the correct graphics, maybe we're having caching issues?	1	1	https://issues.jboss.org/browse/DEVELOPER-3540
6698	DEVELOPER-3536	http://static.jboss.org/_root/images/rhd/topic/RHDev_topics_primopromo_microservicesicon_20sept2016.png How is this? [~james.cobb] [~lincolnthree]	1	1	https://issues.jboss.org/browse/DEVELOPER-3536
6699	DEVELOPER-3503	For 10.0, I was required to hard code [0] a bunch of links in this file [1] ... will that be required for 10.1? If so, will the people creating the new docs also be responsible for adding/replacing links in that file so they'll appear here [2] ? [0] https://github.com/redhat-developer/developers.redhat.com/commit/3bfc172ec4129d70d57e2f30b61eac2c0b03d6ec [1] https://github.com/redhat-developer/developers.redhat.com/blob/master/products/devstudio/_common/product.yml#L17-L68 [2] http://developers.redhat.com/products/devstudio/docs-and-apis/	1	1	https://issues.jboss.org/browse/DEVELOPER-3503
6700	DEVELOPER-3503	Internal doc looks good (at least as far as there's a 10.1 section which duplicates the 10.0 content). Redirect still not in place. Will these be live before EOD today? Announcements are supposed to be tomorrow.	1	1	https://issues.jboss.org/browse/DEVELOPER-3503
6701	DEVELOPER-3503	[~nickboldt] I'd prefer to go live today, as long as it's not too soon. There's been some outages the last few days that have impacted our ability to push changes. I wouldn't want these to come back tomorrow and prevent this change going through. Who needs to make the call?	1	1	https://issues.jboss.org/browse/DEVELOPER-3503
6702	DEVELOPER-3498	[~paul.robinson] Happens automatically? Then what is this variable used for? https://github.com/redhat-developer/developers.redhat.com/pull/1442/files#diff-58cf26cee5226e736168698e99de542eL4	1	0	https://issues.jboss.org/browse/DEVELOPER-3498

#	Requirement	Questions	Answered	Successfully Answered	URL
6703	DEVELOPER-3496	I don't see the change live on http://developers.redhat.com/products/devstudio/overview/ ... does some other magic have to happen to make it go to production?	1	1	https://issues.jboss.org/browse/DEVELOPER-3496
6704	DEVELOPER-3496	[~dcoughlin1] I made a mistake and supplied the wrong URL link. Can you update it to this please? http://developers.redhat.com/products/cdk/overview/ (LMK if I should create a new ticket)	1	1	https://issues.jboss.org/browse/DEVELOPER-3496
6705	DEVELOPER-3496	The source I put into the page (via my PR) had this: {code}link:../../devsuite/overview/{code} But when the page is rendered, the link is bad -- it includes a 3rd ../ segment and sends the user to [1] instead of [2]. {code}<a href="../../../devsuite/overview/">{code} Is it expected that the adoc, when rendered as html, injects another ../ segment? [1] http://developers.redhat.com/devsuite/overview/ [2] http://developers.redhat.com/products/devsuite/overview/ ... I see this issue is fixed. Must have been looking at an old cached copy.	1	1	https://issues.jboss.org/browse/DEVELOPER-3496
6706	DEVELOPER-3496	What exactly was deployed?	1	1	https://issues.jboss.org/browse/DEVELOPER-3496
6707	DEVELOPER-3496	Should we link to the devsuite page so we can tell people there's an online installer for devstudio (which is inside the devsuite installer) ? cc: [~dgolovin] If so, PR: https://github.com/redhat-developer/developers.redhat.com/pull/1462 for review/approval.	0	0	https://issues.jboss.org/browse/DEVELOPER-3496
6708	DEVELOPER-3495	[~james.cobb] Would you or someone on your team be able to make 2 mock-ups of the alert message using one of the product pages, like: http://developers.redhat.com/products/fuse/overview/ I think we could either put the alert directly above or directly below the "Hello World" banner, but I'm not sure which would look better without seeing it. I was originally thinking that we could place the alert at the top of the first "white" background section, so it's not above the page title banner, but then I saw that a lot of the pages being redirected to are currently the product pages.	1	1	https://issues.jboss.org/browse/DEVELOPER-3495
6709	DEVELOPER-3495	[~paul.robinson] - Why is this marked done? I haven't seen this in action. [~ldary] - Was this marked as done by accident?	1	1	https://issues.jboss.org/browse/DEVELOPER-3495
6710	DEVELOPER-3495	Sorry [~lightguard] but we all don't monitor the PRs on Github for changes. It's been customary for links to page/elements for review to be posted in the JIRA comments for the non-devs to view. Can we please get a link to a page with this element in place?	1	1	https://issues.jboss.org/browse/DEVELOPER-3495
6711	DEVELOPER-3495	[~ldary] Can you update the JS to cover every page on the site? They will most likely all be redirected to as part of the JBD to RHD migration.	1	1	https://issues.jboss.org/browse/DEVELOPER-3495
6712	DEVELOPER-3495	Looks like [~tnolan]'s request has not yet been addressed: http://stumpjumper.lab4.eng.bos.redhat.com:36563/events/devnation/2016?referrer=jbd Otherwise, is this ready to go once the tests pass?	1	1	https://issues.jboss.org/browse/DEVELOPER-3495
6713	DEVELOPER-3491	current page shows version "10.0" - should be 1.0 It also refers to running DevSuite on OpenShift - but it should likely be running DevStudio on OpenShift. What is "JBoss AS application server on OpenShift"? Is this redundant?	1	1	https://issues.jboss.org/browse/DEVELOPER-3491
6714	DEVELOPER-3491	Paul, Thanks. Three things: 1. Can you also remove the "get a production subscription" part? It's inconsistent with the other products. 2. Updated text: To download DevSuite you must have an account and accept the terms and conditions of the Red Hat Developer Program which provides $0 subscriptions for development use only. 3. why does it say "Restricted"?	0	0	https://issues.jboss.org/browse/DEVELOPER-3491
6715	DEVELOPER-3477	[~paul.robinson] Will this include all possible page filters (Type, keyword, product, date(s), sort and # results) or a subset of those?	1	1	https://issues.jboss.org/browse/DEVELOPER-3477
6716	DEVELOPER-3477	[~adelaarreola] It's great to see this progressing. I'm curious about the minimum of 2 parameters though as that makes the URL "messy". I'd like to see sys_type displayed as just "type" in the URL. What is the minimum 2 parameters and is there a way to just have a default value if one of them is not included? Ideally, I'd like to see something like "type=blogpost" or "keyword=eap". I'm still in all-day meetings tomorrow (Thurs) but could make some time on Friday to chat more about this. I totally understand system constraints, but maybe we can come up with some alternative options here to make this a little more user friendly?	1	1	https://issues.jboss.org/browse/DEVELOPER-3477
6717	DEVELOPER-3477	[~adelaarreola] Sorry, I meant to send this sooner, but my laptop was constantly crashing on Friday. Thanks for scheduling the meeting. As discussed, please make the following name changes to the parameters in the URL: (this is required for UX sign-off) * change change sys_type to "type" * change project to "product" * change sortBy to "sort" * For the sorting, change most-recent to "most_recent" (for consistent use of the underscore vs a dash) * change query to "q" so it's consistent with search Nice to have as it's very "ugly": (not required for UX sign-off) * For code artifact, can we combine all the jboss developer parameters and just have that as “type=code” or type=”code_artifact”?	1	1	https://issues.jboss.org/browse/DEVELOPER-3477
6718	DEVELOPER-3471	[~paul.robinson] - Can you reset the cache when you get a chance? The project.properties have been updated.	1	1	https://issues.jboss.org/browse/DEVELOPER-3471
6719	DEVELOPER-3471	Which cache do you mean?	1	1	https://issues.jboss.org/browse/DEVELOPER-3471
6720	DEVELOPER-3471	[~paul.robinson] - The project.properties file was updated by the project manager. Don't we need to remove Undertow from DCP and reindex?	1	1	https://issues.jboss.org/browse/DEVELOPER-3471
6721	DEVELOPER-3455	[~james.cobb] Sounds great! Do you happen to know the best way to see what pages currently have alerts other than poking around? As one of the next steps, and once we have the designs, I'd like to document which versions (info, warning, etc.) should be used where. Right now, I think almost all of our alerts use the same style.	1	1	https://issues.jboss.org/browse/DEVELOPER-3455
6722	DEVELOPER-3455	[~james.cobb] [~morgan_thomas] I think this is a great first pass. Can we use the same background colors and text colors (if they pass contrast tests) as the Portal? Those should be brand compliant. I think we could also go a little smaller with the icon sizes to allow for enough padding (whitespace) between them and the border. For the error one, I think that one should be an "x" instead of reusing the "!".	1	1	https://issues.jboss.org/browse/DEVELOPER-3455
6723	DEVELOPER-3455	Hello Hello, Here is v2 with edits, can you review [~tnolan] [~james.cobb] ? http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/pages/images/RHD_alertstyle_r1v2.png	1	1	https://issues.jboss.org/browse/DEVELOPER-3455
6724	DEVELOPER-3455	[~tnolan] Maybe we can keep those dismissible alerts as an "info" and then they can have the "x". Is the image you attached already apart of RHD?	1	1	https://issues.jboss.org/browse/DEVELOPER-3455
6725	DEVELOPER-3455	[~james.cobb] I think this is ready to go to the implementation side. I'm guessing that we will need JIRAs in both RHDENG and DEVELOPER to make sure we get these updated, as well as the location for the main site, and Keycloak pages. Are you aware of any existing alerts on the blog or in the forums?	1	1	https://issues.jboss.org/browse/DEVELOPER-3455
6726	DEVELOPER-3451	[~paul.robinson] I took a look at the online checker and I think that anything smaller that our base font size (16px) should be considered as "smaller text". [~james.cobb] - Do you agree?	1	1	https://issues.jboss.org/browse/DEVELOPER-3451
6727	DEVELOPER-3441	So how would the images be added to the system when a new Topic is being created? Sounds like a PR would need to be raised on the CSS before the content creator can make the topic. Also, I guess the image would break when a content creator changed the title (such that it no longer matches an image name)?	0	0	https://issues.jboss.org/browse/DEVELOPER-3441
6728	DEVELOPER-3416	[~lightguard] [~paul.robinson] - Are we putting redirects into place for this or is there another solution? Just curious as this has been blocking for a quite a bit now.	0	0	https://issues.jboss.org/browse/DEVELOPER-3416
6729	DEVELOPER-3411	I like it. Is there a place to try it?	1	1	https://issues.jboss.org/browse/DEVELOPER-3411

#	Requirement	Questions	Answered	Successfully Answered	URL
6730	DEVELOPER-3408	I generally understand your comment, [~rblake], but missing some details. I've found this, but nothing is working for me and possibly a user error/issue: https://www.aspedia.net/faq/how-do-i-create-anchor-link-drupal Can you point to a file that has this in place? Thanks, Mike	1	1	https://issues.jboss.org/browse/DEVELOPER-3408
6731	DEVELOPER-3408	Can I assume we can do this in Drupal HTML mode? If so, cool. Can you supply a sample?	1	1	https://issues.jboss.org/browse/DEVELOPER-3408
6732	DEVELOPER-3408	Yes, you can do this with Drupal. When you're editing something you can switch to source mode, it's a button on the right of the WYSIWYG editor. What you need to do is create a link with an id attribute: {code:html} Google {code} Next, when you want to link to that within the same document you'd created text similar to: {code:html} Link to Google {code} You'll notice the `href` in the second link starts with a `#` it is then followed by the id of the first link. This actually works for any element within a document. You can link to any HTML element that has an id.	1	1	https://issues.jboss.org/browse/DEVELOPER-3408
6733	DEVELOPER-3408	[~mikgue], does my previous comment fulfill what you need?	1	1	https://issues.jboss.org/browse/DEVELOPER-3408
6734	DEVELOPER-3408	I'm still missing some context. What would really help is a sample RHDev drupal page where I can see the code and results. Would you mind assembling something simple? Thanks, Mike	1	1	https://issues.jboss.org/browse/DEVELOPER-3408
6735	DEVELOPER-3408	Ids can only be added to h2, h3, h4, h5, and h6 currently. Should we decide we need to allow other elements to have an id we can make that change within Drupal.	1	1	https://issues.jboss.org/browse/DEVELOPER-3408
6736	DEVELOPER-3394	[~mbenitez] do you need anything from my team now, or are you all set?	1	1	https://issues.jboss.org/browse/DEVELOPER-3394
6737	DEVELOPER-3394	[~paul.robinson] [~mbenitez] I see I've been assigned this issue. What is the selected solution? It's a bit rambling and I'm not sure what's going on with this or if it applies to me.	1	1	https://issues.jboss.org/browse/DEVELOPER-3394
6738	DEVELOPER-3382	Hi [~adelaarreola] Can we take the word 'free' out of this line: Stop by booth 5301 to snag a free t-shirt, stickers or other Red Hat freebies we'll be handing out. So it's: Stop by booth 5301 to snag a T-shirt, stickers or other Red Hat freebies we'll be handing out. Also, Can we do something to denote sessions that are on different days? Like a bold "Tuesday" and a horizontal line or something similar to separate the days of the week in our schedule? Does that make sense?	1	1	https://issues.jboss.org/browse/DEVELOPER-3382
6739	DEVELOPER-3377	Is this a duplicate of https://issues.jboss.org/browse/DEVELOPER-3274?	0	0	https://issues.jboss.org/browse/DEVELOPER-3377
6740	DEVELOPER-3368	[~tnolan] there's some naming discrepancies going on here too. The first link says "Enterprise Application Server" and the link to customer portal under the table says "JBoss EAP". Which correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-3368
6741	DEVELOPER-3368	[~tnolan] From all download pages?	1	1	https://issues.jboss.org/browse/DEVELOPER-3368
6742	DEVELOPER-3337	[~jessk] when does this need to go live?	1	1	https://issues.jboss.org/browse/DEVELOPER-3337
6743	DEVELOPER-3337	[~jessk] - Do we have a marketing direction for this? Some content direction?	1	1	https://issues.jboss.org/browse/DEVELOPER-3337
6744	DEVELOPER-3337	[~jessk]: Isn't this dependent on a landing page for cheat sheets? If not, where would the primary CTA go?	1	1	https://issues.jboss.org/browse/DEVELOPER-3337
6745	DEVELOPER-3337	[~jessk]: Do you have an update on when you need this by?	1	1	https://issues.jboss.org/browse/DEVELOPER-3337
6746	DEVELOPER-3337	[~adelaarreola] this is now merged. Can you add the SP for this please?	1	0	https://issues.jboss.org/browse/DEVELOPER-3337
6747	DEVELOPER-3321	Did you try the linkchecker from https://github.com/linkcheck/linkchecker as well? This looks like a fork basically because the maintainer of the original abandoned it and others in the community have picked it up. Try this one as well and let us know how it stacks up with the other one. I think either linkchecker or brokenlinkchecker are my top picks. Not sure why VPN is an issue for monkeytest. Couldn't we run it on a server inside our vpn or is this a hosted service?	1	1	https://issues.jboss.org/browse/DEVELOPER-3321
6748	DEVELOPER-3321	[~ihamilto] I'm slightly skeptical that all the broken links on the site have a single cause. I had expected us to have way more. Are you sure BLC is finding them all? For example, does it find broken links that are created by JS? E.g in the results of /resources?	1	1	https://issues.jboss.org/browse/DEVELOPER-3321
6749	DEVELOPER-3321	There's a bunch of 404s and 403s on here: https://developers.redhat.com/products/dotnet/hello-world/. Does BLC find them?	1	1	https://issues.jboss.org/browse/DEVELOPER-3321
6750	DEVELOPER-3321	[~paul.robinson] how are you testing to get the 403 & 404s? I am ignoring external url's, otherwise there would be 50+ more broken links. blc does find the 403 on this page as reported in DEVELOPER-3819. All occurrences where the "DOWNLOAD" button is visible on mobile layout fail with a 403 on staging. {code} Getting links from: https://developers-pr.stage.redhat.com/pr/1736/export/products/dotnet/hello-world/ ├─BROKEN─ https://developers-pr.stage.redhat.com/pr/1736/export/products/dotnet/downloads/ (HTTP_403) ├─BROKEN─ https://developers-pr.stage.redhat.com/pr/1736/export/products/dotnet/download/ (HTTP_403) Finished! 120 links found. 118 excluded. 2 broken. {code} Does this match what you see? Currently there is no way of automated (that I can find) link checking JS created links. This can be done with the ScreamingFrog premium standalone app.	1	1	https://issues.jboss.org/browse/DEVELOPER-3321
6751	DEVELOPER-3321	I use https://chrome.google.com/webstore/detail/check-my-links/ojkcdipcgfaekbeaelaapakgnjflfglf. You can then browse to https://developers.redhat.com/products/dotnet/hello-world/ and click the 'check my links' button, to see: https://www.dropbox.com/s/7wppoygyuzry78x/Screenshot%202017-02-03%2011.25.45.png?dl=0 and https://www.dropbox.com/s/f51nszzgxbwwq9f/Screenshot%202017-02-03%2011.26.18.png?dl=0 and https://www.dropbox.com/s/conkgmx2nw5hhxq/Screenshot%202017-02-03%2011.26.30.png?dl=0. We had issues with Blinkr when checking JS generated links also. I think this was because we needed to run PhantomJS or similar to render every page. I'm happy to mark this Spike as done. Can you create a story for incorporating BLC into our build?	0	0	https://issues.jboss.org/browse/DEVELOPER-3321
6752	DEVELOPER-3312	I know the details, but I'm honestly not sure how to describe that it's done :(Do we say we have groups in place and people can follow a prescribed workflow? Do we say content must flow in a prescribed workflow to be published?	1	1	https://issues.jboss.org/browse/DEVELOPER-3312
6753	DEVELOPER-3312	Maybe describe the manual tests we need to follow in order to check it's working? We need that anyway in order to do a review.	0	0	https://issues.jboss.org/browse/DEVELOPER-3312
6754	DEVELOPER-3287	[~tnolan] what do you suggest here?	1	1	https://issues.jboss.org/browse/DEVELOPER-3287
6755	DEVELOPER-3287	[~paul.robinson] I'd like to understand more about requirements for products like CDK. They have to start with some sort of download, even though it's not the actual "product", right? Where is this request coming from and do you happen to know who would be the best resource to figure requirements?	1	1	https://issues.jboss.org/browse/DEVELOPER-3287
6756	DEVELOPER-3287	[~paul.robinson] We do have a default download for CDK now, as of May. Is this an issue still for any other products?	1	1	https://issues.jboss.org/browse/DEVELOPER-3287
6757	DEVELOPER-3279	[~drieden] when does this need to go live? We need a variety of content pieces to do this. I would suggest working with [~lincolnthree] to get this together, then my team can implement the page on the site. You can take a look at http://developers.redhat.com/products/eap/ for an example of what we need. There's a number of pages, all linked from the left-hand navigation.	1	1	https://issues.jboss.org/browse/DEVELOPER-3279
6758	DEVELOPER-3273	I think this is going to be quite a lot of work to change. What's the impact of leaving this unchanged?	1	1	https://issues.jboss.org/browse/DEVELOPER-3273

#	Requirement	Questions	Answered	Successfully Answered	URL
6759	DEVELOPER-3273	It's confusing for users, and I'm pretty sure that corporate marketing wouldn't be too happy about products being presented using different version specifications than those that have been chosen by product teams. That said, the mandatory z-numbers can be overlooked, I guess. But the ".GA" at the end is really weird. It looks like something internal, not intended for public eyes. When I showed it to people, the almost universal reaction was "What's that GA doing there?". If we could get rid of the ".GA", that'd be great. (I understand this already works, only the number must be in quotes because the system expects a string.)	1	1	https://issues.jboss.org/browse/DEVELOPER-3273
6760	DEVELOPER-3273	Ok, we can probably drop the GA relatively easily. Do any of the products need the GA suffix?	1	1	https://issues.jboss.org/browse/DEVELOPER-3273
6761	DEVELOPER-3273	[~rkratky] - Can you please review http://developers-pr.stage.redhat.com/pr/1273/build/1762/? If you notice any issues, please let me know.	1	1	https://issues.jboss.org/browse/DEVELOPER-3273
6762	DEVELOPER-3256	[~jessk] - Do you approve of this text change?	1	1	https://issues.jboss.org/browse/DEVELOPER-3256
6763	DEVELOPER-3236	Ack to Andy's target for this link (https://access.redhat.com/documentation/en/red-hat-jboss-a-mq/?version=7.0-alpha). Can we add this link to http://www.jboss.org/products/amq/overview/ please.	0	0	https://issues.jboss.org/browse/DEVELOPER-3236
6764	DEVELOPER-3213	[~unibrew] Great thanks for reviewing it so quickly! Also, Paul was aiming to have this in production by tomorrow, do you think you will be finished with the boosting updates today? or when should we expect those changes?	1	1	https://issues.jboss.org/browse/DEVELOPER-3213
6765	DEVELOPER-3207	Isn't this work already covered by: DEVELOPER-3197, DEVELOPER-3157 & DEVELOPER-3158?	1	1	https://issues.jboss.org/browse/DEVELOPER-3207
6766	DEVELOPER-3207	PRs to set current_version = 10.0.0.GA: https://github.com/redhat-developer/developers.redhat.com/pull/1204 https://github.com/jboss-developer/www.jboss.org/pull/1039 [~paul.robinson] does this change still need to happen? Is this how the Resources page [1] decides which doc to show? Or does it just load the latest, making this variable (and these PRs) irrelevant? [1] https://www.jboss.org/products/devstudio/resources/	1	1	https://issues.jboss.org/browse/DEVELOPER-3207
6767	DEVELOPER-3207	I've added some doc links (as collected from [1]) into the https://github.com/redhat-developer/developers.redhat.com/pull/1204 PR. Is that what you were expecting? [~mhusnain] please review. I can apply the same list of docs for the second PR, too. [1] https://access.redhat.com/documentation/en/red-hat-jboss-developer-studio/?version=10.0	1	1	https://issues.jboss.org/browse/DEVELOPER-3207
6768	DEVELOPER-3206	Alternatively, we could provide a service that returns the HTML for the header and footer. That way any changes to them will automatically propagate to the other services. The customer portal team do this. They call it "The Chrome Service". Would this work better for your needs?	1	0	https://issues.jboss.org/browse/DEVELOPER-3206
6769	DEVELOPER-3179	Can this item be closed?	1	1	https://issues.jboss.org/browse/DEVELOPER-3179
6770	DEVELOPER-3172	[~james.cobb], I see you assigned this to me. I believe this is just a SOLP, can your team handle it?	1	1	https://issues.jboss.org/browse/DEVELOPER-3172
6771	DEVELOPER-3161	[~dcoughlin1] Did you push these updates to RHD as well? The .NEt content page on RHD has the old content. Can you check for me? Thank you so much!	1	1	https://issues.jboss.org/browse/DEVELOPER-3161
6772	DEVELOPER-3161	[~jessk] - Can you review http://developers-pr.stage.redhat.com/pr/1362/build/1966/dotnet/ when you get the chance?	1	0	https://issues.jboss.org/browse/DEVELOPER-3161
6773	DEVELOPER-3159	Can I close this? This is live in production.	0	0	https://issues.jboss.org/browse/DEVELOPER-3159
6774	DEVELOPER-3143	[~jessk] [~burrsutter] where can I find the content for the third white box?	1	1	https://issues.jboss.org/browse/DEVELOPER-3143
6775	DEVELOPER-3143	[~james.cobb] where can I find the image for the topic's grey top banner/hero?	1	1	https://issues.jboss.org/browse/DEVELOPER-3143
6776	DEVELOPER-3143	[~adelaarreola] is this just blocked on that image, or are other things needed?	1	1	https://issues.jboss.org/browse/DEVELOPER-3143
6777	DEVELOPER-3141	[~velias] Thanks for the review I have now updated scenarios based on your comments: Login Page: https://github.com/RubyGemTSL/developers.redhat.com/blob/9469d034442841d7bb3fc8a8acf02feb045870e7/_cucumber/features/login_register/basic_login.feature Registration Page: https://github.com/RubyGemTSL/developers.redhat.com/blob/9469d034442841d7bb3fc8a8acf02feb045870e7/_cucumber/features/login_register/basic_registration.feature Social Registration Page: https://github.com/RubyGemTSL/developers.redhat.com/blob/9469d034442841d7bb3fc8a8acf02feb045870e7/_cucumber/features/login_register/social_registration.feature Can you please have a quick review if they are OK give me a thumbsup on the PR? https://github.com/RubyGemTSL/developers.redhat.com/blob/9469d034442841d7bb3fc8a8acf02feb045870e7/_cucumber/features/login_register/social_registration.feature Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-3141
6778	DEVELOPER-3130	[~james.cobb] Can you add the latest mockups here so [~hwindzio] can take a look?	1	1	https://issues.jboss.org/browse/DEVELOPER-3130
6779	DEVELOPER-3130	Thanks [~james.cobb]. I think I remember seeing a version for when the user isn't logged in. Do you have that mockup handy too?	1	1	https://issues.jboss.org/browse/DEVELOPER-3130
6780	DEVELOPER-3130	Great, there is the single sheet of paper as the Knowledgebase icon I was looking for! I guess Tiffany and I were looking at the lock icon when guests are locked out (which is really not the icon for Knowledgebase but for being locked out). I'm just noticing this, but in the link James sent we have another 'Article' link/icon at the top. Does that bring you to the Knowledgebase articles (is it a dup?), or is another type of content with the same name?	1	1	https://issues.jboss.org/browse/DEVELOPER-3130
6781	DEVELOPER-3130	[~hwindzio][~tnolan] - To answer your second question but yes, we have another type called "Article" and now we're adding "Knowledgebase Article/Solution". It is a bit confusing... Should it just say "Knowledgebase"?	1	1	https://issues.jboss.org/browse/DEVELOPER-3130
6782	DEVELOPER-3128	Just stumbled across this ticket. I assume this should apply to all RHD articles, Get Started, etc. content that links to any RHCP protected articles? Should there be a macro similar to #{cdn(site.base_url)} to make it easy to put the right links in content that will do SSO to the right prod/staging environment?	1	1	https://issues.jboss.org/browse/DEVELOPER-3128
6783	DEVELOPER-3128	[~rterzi] yes It does apply to all links. I agree, a function would be useful, can you raise a separate JIRA ticket for that? You make a good point about the other places that link to these articles. We should search for such occurrences.	1	1	https://issues.jboss.org/browse/DEVELOPER-3128
6784	DEVELOPER-3112	Attached the files, but whats next? [~james.cobb]	1	1	https://issues.jboss.org/browse/DEVELOPER-3112
6785	DEVELOPER-3108	[~tnolan] how important is this for summit?	1	1	https://issues.jboss.org/browse/DEVELOPER-3108
6786	DEVELOPER-3108	[~james.cobb] Looking good! A few things: * Font size for the outside location: Sorry, it's hard to tell from the mockup, but is the font size also slightly smaller from the main size once it moves into the upper position? I think it would help to see an example with input plus the label on the outside too for comparison. * With KISS in mind, I'm not sure we need the "x" or the 'Show" button either. I know the Warby Parker example had these, but I feel like they add unnecessary visual clutter for features that are probably not used very often. I can try to see if there is some UX research on their usage.	1	1	https://issues.jboss.org/browse/DEVELOPER-3108

#	Requirement	Questions	Answered	Successfully Answered	URL
6787	DEVELOPER-3108	[~james.cobb] Sounds good. Can we also change the border color of the active field to a blue? I'm not sure if this is a standard form style, but right now, the active field gets a blue border and a soft "glow". The soft glow may not work here, but I'd still like to see the blue border to indicate an active field. Thoughts?	1	1	https://issues.jboss.org/browse/DEVELOPER-3108
6788	DEVELOPER-3108	[~james.cobb] +1 to the grey, especially since we'll probably have a red button for the primary action. Is the "Case Sensitive" styled on that one as well? If so, I'd like to see that with less contrast (lighter grey) so it does not conflict with the field or draw the eye too much. Other than that, I think this is good to go!	1	1	https://issues.jboss.org/browse/DEVELOPER-3108
6789	DEVELOPER-3108	[~adelaarreola] It's done? Should I review it? Thanks.	1	1	https://issues.jboss.org/browse/DEVELOPER-3108
6790	DEVELOPER-3108	[~adelaarreola] I'm finishing implementation on keycloak and here are few things: 1. min-width - remove it. It's breaking layout on mobile. https://github.com/adelasofia/developers.redhat.com/blob/2011b71ee84aee03b5f036629aedd399811ca6e4/stylesheets/_adaptive-placeholder.scss#L10 2. label is not correctly positioned for <select> form field can you add that "bottom" magic for "select" as well? Thanks. 3. text-align for last element: https://github.com/adelasofia/developers.redhat.com/blob/2011b71ee84aee03b5f036629aedd399811ca6e4/stylesheets/_adaptive-placeholder.scss#L69 what for this is? Preview here: https://rhdssodev-lkrzyzan.itos.redhat.com/auth/realms/rhd/account/	1	0	https://issues.jboss.org/browse/DEVELOPER-3108
6791	DEVELOPER-3105	[~sgilda] when do you want this live? ASAP?	1	1	https://issues.jboss.org/browse/DEVELOPER-3105
6792	DEVELOPER-3099	[~ganandan-Redhat] I see switchyard here: https://www.dropbox.com/s/u1bkmdumhp2xyke/Screenshot%202016-05-29%2013.00.39.png?dl=0 on http://developers.redhat.com/products/fuse/community/. Did you expect it to appear somewhere else?	1	1	https://issues.jboss.org/browse/DEVELOPER-3099
6793	DEVELOPER-3099	[~tcunning] - Can you please provide a one-line description for SwitchYard?	1	1	https://issues.jboss.org/browse/DEVELOPER-3099
6794	DEVELOPER-3098	[~aileenc], [~ganandan-Redhat], what in particular would you like the focus of this video to be? Otherwise we will just come up with a script of questions to have Hiram answer.	1	1	https://issues.jboss.org/browse/DEVELOPER-3098
6795	DEVELOPER-3098	I think the topics covered in the current video with James are good, but the answers will have changed a bit in the last 3 years since that video was shot. What is Fuse? What can you do with the current release of Fuse? What is the focus for the next releases? Adding [~chirino] for his input.	1	1	https://issues.jboss.org/browse/DEVELOPER-3098
6796	DEVELOPER-3096	[~ganandan-Redhat] I was unable to find the videos you are referring to. Could you verify that they are still there? Also, if I understand you correctly, you want to remove all resources from the learn page, beginning with "JBoss Fuse on OpenShift - How to Connect to Twitter" including any "older" resource?	0	0	https://issues.jboss.org/browse/DEVELOPER-3096
6797	DEVELOPER-3096	Please remove anything from 2 years and older. Also, can we add Christina Lin's videos from the Getting Started pages? Please also ensure https://vimeo.com/150873903 is present - its new and relevant as we moved forward in the product line.	1	0	https://issues.jboss.org/browse/DEVELOPER-3096
6798	DEVELOPER-3095	The currently linked page does have a link to https://access.redhat.com/products/red-hat-jboss-fuse/, but I agree it would be better to go direct. Fixing this requires a change to the layout and product configuration. I see this as a nice-to-have prior to Summit. Agree?	1	1	https://issues.jboss.org/browse/DEVELOPER-3095
6799	DEVELOPER-3095	[~tnolan] Please can you do this as part of your migration of this content to Drupal?	0	0	https://issues.jboss.org/browse/DEVELOPER-3095
6800	DEVELOPER-3094	[~tnolan] Please can you do this as part of your migration of this content to Drupal?	1	1	https://issues.jboss.org/browse/DEVELOPER-3094
6801	DEVELOPER-3093	Who on Docs team is building the REST API, is there a tracker for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-3093
6802	DEVELOPER-3093	Is there any update on this?	1	1	https://issues.jboss.org/browse/DEVELOPER-3093
6803	DEVELOPER-3093	If you can move to 6.3 easily then that works for me. I only suggested the redirect as I thought we were stalled waiting for the REST API. When you have draft of the Developer site updates can you share it for review?	1	1	https://issues.jboss.org/browse/DEVELOPER-3093
6804	DEVELOPER-3087	[~rblake] - The module seems to be working as intended; however, I didn't make any changes here. Jason suggested that I update the README. Is there anything else that needs to be completed here, or should I just plan on adding a quick walkthrough of how to view the diffs? Also, I don't know that I have access to that Phase2 documentation.	1	1	https://issues.jboss.org/browse/DEVELOPER-3087
6805	DEVELOPER-3085	What metrics and in which reporting tools are we talking about in this context?	1	1	https://issues.jboss.org/browse/DEVELOPER-3085
6806	DEVELOPER-3080	[~paul.robinson] I have a bit of a problem with product pages. DCP reported on missing descriptions and titles for quite a list of paths under /products/* but almost all of them are about index files that automatically redirect to overview.html e.g. http://developers.redhat.com/products/amq/ . I believe we shouldn't have those index files in sitemap as giving them descriptions/titles and making searchable doesn't make sense IMO. Moreover, the overview.html files were not reported as those with missing descriptions because they are automatically generated by this code: {code} - page.description ||= truncate_para(page._overview, 150) {code} Still, I'd argue we should provide better descriptions for those. WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-3080
6807	DEVELOPER-3080	I think the right approach is: 1) Put the redirects into Apache for redirecting http://developers.redhat.com/products/X/ to http://developers.redhat.com/products/X/overview. Please use a wildcard/variable for 'X' so that it works for all current and future products. 2) Delete all the html redirects that are now redundant. Will this issue be complete after this work?	1	1	https://issues.jboss.org/browse/DEVELOPER-3080
6808	DEVELOPER-3080	Well staging is part of the standard releasing process so I was counting this in already. Nevertheless, I have no idea how your PR envs are done. Is Lukas K. the person to handle this as well or somebody from your team? As for now I have this task created for Lukas https://issues.jboss.org/browse/ORG-3312	1	1	https://issues.jboss.org/browse/DEVELOPER-3080
6809	DEVELOPER-3080	Sounds good to me. I started thinking though. What's the point of having overview pages if by the content they are really the index pages? Why lengthen the url with additional word?	1	1	https://issues.jboss.org/browse/DEVELOPER-3080
6810	DEVELOPER-3077	I assume that the only way to tell that a page is a topic page is by checking that its title contains "topic" word. What if somebody will use this word in any other random title? Or is there a convention that forbids that and we should be good ? Otherwise, it would require quite a bit of changes to add some new field in document that tells whether this is a topic page and then propagate it with data.	1	1	https://issues.jboss.org/browse/DEVELOPER-3077
6811	DEVELOPER-3077	[~unibrew]: I agree, testing for 'Topic' in the title is going to be brittle. However, I'm not sure of another way to programmatically detect a topic page from the DCP documents. Is there a way to hard-code the boosts? There won't be that many topic pages added to the site, so we might be able to handle that manually.	1	1	https://issues.jboss.org/browse/DEVELOPER-3077
6812	DEVELOPER-3066	[~tnolan] - I appears that the changes to the product overview pages don't involve either the new StackOverflow or Forums integration. Isn't this out of scope for the Summit release?	1	1	https://issues.jboss.org/browse/DEVELOPER-3066

#	Requirement	Questions	Answered	Successfully Answered	URL
6813	DEVELOPER-3066	A quick question about the "Latest Resources" section. Currently it's represented with just blogs in the wireframe. Is it correct to assume this will cover all resources? I assume so but just want to clarify so functionality is correctly built.	1	1	https://issues.jboss.org/browse/DEVELOPER-3066
6814	DEVELOPER-3060	[~tnolan] Where can I find the content and links for the "Professional services and support" section of the Product Help pages?	1	1	https://issues.jboss.org/browse/DEVELOPER-3060
6815	DEVELOPER-3060	[~tnolan] can you point me to a wireframe or whatever that needs the content? Also, when is it needed?	1	1	https://issues.jboss.org/browse/DEVELOPER-3060
6816	DEVELOPER-3059	[~james.cobb] where can I find the icons/graphics for this page?	1	0	https://issues.jboss.org/browse/DEVELOPER-3059
6817	DEVELOPER-3046	[~burrsutter] you're the reporter, would you mind checking http://developers-pr.stage.redhat.com/pr/1115/build/1220/dotnet/ to see if it meets your criteria please?	0	0	https://issues.jboss.org/browse/DEVELOPER-3046
6818	DEVELOPER-3045	[~rblake], [~lightguard]: After talking to Phase2, it seems like this is not needed anymore? I'm thinking this due to: 1) The user's session is stored in the DB. 2) Changes to the Drupal config are done in-place without restarting the web server. Am I missing something?	1	1	https://issues.jboss.org/browse/DEVELOPER-3045
6819	DEVELOPER-3044	[~rterzi] This should now be fixed in master as a by-product of DEVELOPER-2933. Could you please re-test when you have time. cheers, Rob	1	1	https://issues.jboss.org/browse/DEVELOPER-3044
6820	DEVELOPER-3039	Can you login to the site manually, when it's running on Drupal on your computer? I suspect you just need 'localhost' adding to the staging KeyCloak instance. [~lkrzyzanek] should be able to help.	1	1	https://issues.jboss.org/browse/DEVELOPER-3039
6821	DEVELOPER-3039	[~lkrzyzanek] - Can you add http://docker:4242? I think this should solve the keycloak issues on local docker instances.	1	1	https://issues.jboss.org/browse/DEVELOPER-3039
6822	DEVELOPER-3039	[~lkrzyzanek] - Can we add http://stumpjumper.lab4.eng.bos.redhat.com:32778/ as well? I noticed that the login is broken here as well. Also, is it possible to add all of the docker ports in the 32000 range (or at least 32700 - 32900)? Sorry to ping you with all these, but the ports are assigned dynamically. And can the same be done for http://stumpjumper.lab4.eng.bos.redhat.com?	1	1	https://issues.jboss.org/browse/DEVELOPER-3039
6823	DEVELOPER-3039	This is my error log from the acceptance test target...I can't seem to get it to work properly locally. [~ihamilto] - have you been able to get tests running locally?	1	1	https://issues.jboss.org/browse/DEVELOPER-3039
6824	DEVELOPER-3035	[~james.cobb] Thank you James, can I get the static urls for the images you attached?	1	1	https://issues.jboss.org/browse/DEVELOPER-3035
6825	DEVELOPER-3035	[~adelaarreola] - Can we adjust the link color for the "Need Help" strip? .stack-help a { color: #08C0FC; Hover: color: #46a1da;	0	0	https://issues.jboss.org/browse/DEVELOPER-3035
6826	DEVELOPER-3022	Feels like we need the docker-compose per environment for this. Maybe we simply create new files and edit the control.rb script for v1? [~rblake] what do you think?	1	1	https://issues.jboss.org/browse/DEVELOPER-3022
6827	DEVELOPER-3019	[~james.cobb]: The page is being requested to go live on the 15th of May. Can you get this image ready CoB Thursday 12th May?	1	1	https://issues.jboss.org/browse/DEVELOPER-3019
6828	DEVELOPER-3019	[~ecurley] I think James is going to need more details. What do you want on the image?	1	1	https://issues.jboss.org/browse/DEVELOPER-3019
6829	DEVELOPER-3012	Maybe just #1 until GA? WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-3012
6830	DEVELOPER-3007	[~dcoughlin1] I was working on this last week and noticed some fields are not on the DCP, * Number of votes for each question * Number of views * Name of the answer's author is it possible to add these? [~paul.robinson] [~tnolan] I noticed that most questions don't have an 'accepted answer'. So, does the answer need to be 'accepted' in order for it to show up on this page? If not, should I skip that question?	1	1	https://issues.jboss.org/browse/DEVELOPER-3007
6831	DEVELOPER-2988	I like it. Do we use a period at the end of Survey? [~rayploski-launch] - thoughts?	1	1	https://issues.jboss.org/browse/DEVELOPER-2988
6832	DEVELOPER-2988	[~james.cobb]Do we have social banners for the Developer Survey also? Can we add text to the promo banner to say, "How do you stack up?" ? Thank you! Jess	0	0	https://issues.jboss.org/browse/DEVELOPER-2988
6833	DEVELOPER-2987	[~ecurley] [~paul.robinson] what should the URL be for this page?	1	1	https://issues.jboss.org/browse/DEVELOPER-2987
6834	DEVELOPER-2987	[~adelaarreola] - Link we will link to for document is TBD. [~newtonm] is working on getting it up in download manager. [~paul.robinson] - what URL would you like? promotions/developerssalarysurvey2016?	1	1	https://issues.jboss.org/browse/DEVELOPER-2987
6835	DEVELOPER-2977	Is this for the RHSCL 2.2 release? Do you want all the existing content to be updated for the new release? Or do you want to provide instructions on how to access tha beta until the release happens?	1	1	https://issues.jboss.org/browse/DEVELOPER-2977
6836	DEVELOPER-2977	Sounds good. What's the rough/approximate schedule for beta and GA? Do you want the RHSCL/DTS TTHWs updated for GA?	1	1	https://issues.jboss.org/browse/DEVELOPER-2977
6837	DEVELOPER-2977	Red Hat Software Collections 2.2 Beta is now available and includes these additions and updates: open source databases (MariaDB 10.1, MongoDB 3.2, and PostgreSQL 9.5); open source languages (Node.js v4.2.2, Python 3.5, Ruby 2.3, and Rails 4.2); and Maven 3.3; plus updates of MongoDB 2.6, Python 2.7, Apache 2.4, and Thermostat 1.4. Instructions for installing them can be found here[https://access.redhat.com/site/documentation/en-US/Red_Hat_Software_Collections/2-Beta/html-single/2.2_Release_Notes/index.html#sect-Installation-Subscribe].	1	1	https://issues.jboss.org/browse/DEVELOPER-2977
6838	DEVELOPER-2977	Rob - looks great. Can you rephrase "Release highlights" to something like "New additions". We have a number that are brand new and more that are updated, so I want to distinguish these as brand new. Use some other wording if you prefer. Otherwise, good to go.	1	1	https://issues.jboss.org/browse/DEVELOPER-2977
6839	DEVELOPER-2976	[~mikgue] Can you add the specifics for this and let [~paul.robinson]'s team know exactly what needs to be added. I'm not sure what's involved here.	1	1	https://issues.jboss.org/browse/DEVELOPER-2976
6840	DEVELOPER-2976	[~rterzi] you've shared a number of emails on this. Do you already know the answers here? ^^ If not, I'll look into it. Thanks.	1	1	https://issues.jboss.org/browse/DEVELOPER-2976
6841	DEVELOPER-2976	Sorry, I'm not sure if I understand what is being asked. I'll be glad to update the page once I'm sure what you want. The CDK downloads page, http://developers.redhat.com/products/cdk/download/, Step 3 links to RHCP for the vagrant boxes. Do you want me to replace those with download manager links like the Get Started Page has? What's the plan to keep these pages updated? The Downloads page is mostly an out of date clone of the Get Started page. If there has to be a Download link/page, maybe it should just redirect to the Get Started page?	1	1	https://issues.jboss.org/browse/DEVELOPER-2976
6842	DEVELOPER-2970	Do you have a recommendation on which way we should redirect? Should we favour https over http?	1	1	https://issues.jboss.org/browse/DEVELOPER-2970
6843	DEVELOPER-2970	[~martymartin]: I'm not sure redirecting all traffic to use HTTP is the right approach. We have some activities that must be done over HTTPS (e.g. login, registration, etc). If we wanted to use HTTPS for everything, why would we need to switch to using relative URLs? Couldn't we just ensure that the fully-qualified URLs use HTTPS?	1	1	https://issues.jboss.org/browse/DEVELOPER-2970
6844	DEVELOPER-2970	What is going to happen with http redirects created on developers.redhat,com? For example we use in DevSuite online installer: * http://developers.redhat.com/redirect/to/virtualbox-5.0.26.download * http://developers.redhat.com/redirect/to/virtualbox-macos-5.0.26.download Will they continue to work? The problem is virtualbox download is not available over https. Can we configure https redirect for http URL? Is that possible? Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-2970
6845	DEVELOPER-2970	[~dgolovin] the HTTPS version of the links work for me, downloading via a web browser: * https://developers.redhat.com/redirect/to/virtualbox-5.0.26.download * https://developers.redhat.com/redirect/to/virtualbox-macos-5.0.26.download Is the problem that these files are being downloaded by something else (i.e not a web browser) that can't handle HTTPS?	1	1	https://issues.jboss.org/browse/DEVELOPER-2970

#	Requirement	Questions	Answered	Successfully Answered	URL
6846	DEVELOPER-2970	[~paul.robinson] Thank for creating the redirects! They work for me. I have one question though. How does it work after this issue is fixed and developers site force to use https? I assume when devsuite installer hits http://developers.redhat.com/redirect/to/whatever redirect it is going to be redirected to https://developers.redhat.com/redirect/to/whatever and if latter is configured as redirect to external URL it would trigger download. Does that sound right? If it does then current release has already taken care of with https redirects in place, but we still have to take care of four previous releases and create https redirects to match listed below: 1. http://developers.redhat.com/redirect/to/virtualbox-5.0.8.download 2. http://developers.redhat.com/redirect/to/vagrant-1.7.4.download When those two are created DevSuite Instsallers should have no issues with using redirects on developers.redhat.com.	0	0	https://issues.jboss.org/browse/DEVELOPER-2970
6847	DEVELOPER-2969	[~martymartin] thanks for the report. Can you explain in more detail what the problem is here, as I'm not seeing it.	1	1	https://issues.jboss.org/browse/DEVELOPER-2969
6848	DEVELOPER-2960	Hi [~paul.robinson], Are these two scenarios that David outlined still valid? According to David a user should NOT need be logged in, in order to download JBoss/Red Hat Developer Studio. However this is not the case currently. Could you confirm whether or not these two scenarios are still valid? Also do we really need to be logged in to download cheat sheets? Both cases currently require a site visitor to log in or register.	1	1	https://issues.jboss.org/browse/DEVELOPER-2960
6849	DEVELOPER-2951	[~lightguard] Apparently it will be MariaDB 10.0.14 that we'l be provided with. Can you update from MySQL to a version of Maria DB close to 10.0.14? Sorry for messing around :-(	1	1	https://issues.jboss.org/browse/DEVELOPER-2951
6850	DEVELOPER-2942	[~paul.robinson]/[~lightguard] Things that still are unclear about this issue that need to be resolved before we can estimate: * What is the actual deployment mechanism e.g. is the do we simply update the current Drupal instance in-place or is there a blue/green deployment scenario * What is acceptable to content editors for v1 when a deployment happens e.g: ** The site may go down briefly why it happens and they lose content ** They absolutely cannot lose content they are working on and therefore deployments take place at scheduled times	1	1	https://issues.jboss.org/browse/DEVELOPER-2942
6851	DEVELOPER-2915	[~tnolan] I think this is covered by your re-working of the topics page UX. Correct? If so, I suggest closing this as a duplicate.	1	1	https://issues.jboss.org/browse/DEVELOPER-2915
6852	DEVELOPER-2913	Oh perfect. That is much easier! Do you want as Visual JIRA's or Website JIRA's?	1	1	https://issues.jboss.org/browse/DEVELOPER-2913
6853	DEVELOPER-2911	There's a related issue to this: how do we handle a/b testing and metrics javascript? I don't know of a way to do for a dev/staging/prod check as that doesn't exist in drupal (that I'm aware of anyway). I'm also worried about this possibly messing up metrics if it's included in a non prod environment.	1	1	https://issues.jboss.org/browse/DEVELOPER-2911
6854	DEVELOPER-2905	Can this be done today [~adelaarreola] [~paul.robinson]? The site has been updated for CDK 2 GA, but the download links point to beta (which are no longer available). cc [~rterzi]	1	1	https://issues.jboss.org/browse/DEVELOPER-2905
6855	DEVELOPER-2905	What's needed: 1. Download manager needs links for CDK 2.0.0 GA files. (cdk-2.0.0.zip and the 3 vagrant boxes) 2. The Downloads page needs to updated to point to cdk-2.0.0.zip instead of (-beta5.zip). http://developers.redhat.com/downloads/	1	1	https://issues.jboss.org/browse/DEVELOPER-2905
6856	DEVELOPER-2904	[~rblake], [~lightguard]. Is this still an issue, now that we don't run Blinkr?	1	1	https://issues.jboss.org/browse/DEVELOPER-2904
6857	DEVELOPER-2901	[~dhladky] - There will be a new GA version of CDK 2 - targeted for Tuesday. Will that have a new URL? Also, per the question above, I submitted a ticket here: https://issues.jboss.org/browse/ORG-3028. It needs a better ticket name as that was just a guess when I created it earlier this week. Regardless, please see David's questions there and correct/clarify anything I had put in. BTW - in addition to the 2 downloads mentioned in the description (JBDS and CDK), did we discuss this morning adding a 3rd item for a "JBDS for container development" download thingy? This is what ORG-3028 would relate to. Mike	1	1	https://issues.jboss.org/browse/DEVELOPER-2901
6858	DEVELOPER-2901	[~paul.robinson] The version for JBoss Dev Studio updated as expected today but the date is showing April 2 instead of April 19. See: https://db.tt/fuSzt0BY The date looks correct though on: http://developers.redhat.com/products/devstudio/download/ Are we using a different date field on the main downloads page than is on the individual product download page?	1	1	https://issues.jboss.org/browse/DEVELOPER-2901
6859	DEVELOPER-2900	[~tnolan] I don't know which download URLs you're referring to but the latest configuration in the Download Manager for JBDS is 9.1.0.Beta 2. Who made the request to create the Download URLs as they should know what the status is.	1	1	https://issues.jboss.org/browse/DEVELOPER-2900
6860	DEVELOPER-2900	Can someone clarify what this means? [~tnolan] [~adelaarreola] {quote}For the Tech preview link, the url should take the user through the Full User flow. Other URLs stay the same and use the same flow as they do now.{quote} Does that mean "require SSO login & check entitlements before allowing the download", like we do for the JBDS EAP installer (but not for the other zips & jars, which are 100% free)?	1	1	https://issues.jboss.org/browse/DEVELOPER-2900
6861	DEVELOPER-2900	Is JBDS-3665 == "the same workflow as CDK and RHEL which makes sure they are a full user" ?	1	1	https://issues.jboss.org/browse/DEVELOPER-2900
6862	DEVELOPER-2900	[~paul.robinson] How can we make the docs listed on http://developers.redhat.com/products/devstudio/docs-and-apis/ (shows latest 9.1.0.GA docs) also appear on https://www.jboss.org/products/devstudio/resources/ (still showing old 9.0.0.GA docs) ? Seems like the code's the same in both .adoc pages and in the _common/product.yml files for both repos.	1	1	https://issues.jboss.org/browse/DEVELOPER-2900
6863	DEVELOPER-2896	[~ecurley], [~james.cobb]: Do we have the graphics for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-2896
6864	DEVELOPER-2892	[~ecurley] - I also have graphics for jboss.org. Do we want a promo there too? [~paul.robinson][~adelaarreola]	0	0	https://issues.jboss.org/browse/DEVELOPER-2892
6865	DEVELOPER-2886	I have completed a first draft of the IoT topics page on the noted Google doc. @Lincoln, Please review and feel free to provide any text edits or recommendations (you can ignore the current comments, these are reference from the internal review process with IoT team stakeholders). Can you please provide a timeline as to when we think we can get this first iteration of the page published? Also, Note that the tags that should be used to populate the dynamic content areas of the page should be any/all IoT related tags: * iot * IoT * internet of things * Internet of Thins * Internet of Things (Industry)	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6866	DEVELOPER-2886	I suggest the following process: 1. Michelle and Lincoln work together to finalise the content 2. Adela builds the IoT Topic page from the finalised content 3. Any further content changes are made directly by Lincoln or Michelle. 4. The Page is pushed to production once Lincoln is happy that it is good enough. 5. Future edits will be made directly via the content team. Essentially, I'd like a process where Adela touches this once. Does this work?	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6867	DEVELOPER-2886	[~lincolnthree], [~mdiserio]: can you confirm that you are happy with my above proposed approach?	1	1	https://issues.jboss.org/browse/DEVELOPER-2886

#	Requirement	Questions	Answered	Successfully Answered	URL
6868	DEVELOPER-2886	Hi Paul, I am happy with the above approach, but to clarify #5. Future edits will be made directly via the content team. Is the process that I submit a Jira request? Any additional details I should be aware of? Additionally, it would be super awesome if we could go live with the page by the end of next week because we will then promote develepers.redhat.com on the collateral, etc. at IoT World Congress May 10 (10k attendees). Ambitious or do-able? [~lincolnthree] I made some updates to the draft content per your recommendation (forcing James Falkner, TMM, to assist with reducing the marketing tone). As part of that, I changed the content in one of the white boxes in order to make the content that was there the first link in the Blue Box (because I wanted the CTA in the text to be immediately relevant). Please let me know what you think at your earliest convenience. -Michele	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6869	DEVELOPER-2886	[~lincolnthree] Is the content for this page still in progress, or is it ready?	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6870	DEVELOPER-2886	[~mdiserio] Do you know which Red Hat Products you'd consider to be related to IoT? Could you list them here? Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6871	DEVELOPER-2886	[~adelaarreola] is this still blocked? If so, can you add an explanation so we can chase it up?	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6872	DEVELOPER-2886	In the bottom 2 rows of the dynamically populated content boxes there are duplicate entires ([screenshot attached|^IoT_Topic_duplicate-entries.png]). The 2 blue circles show the duplicates for "Re-Watch: Messaging for IoT - Virtual JBoss User..." and the red circles show the duplicates for "An Enterprise Developer's Journey to the IoT". Are these duplicates because they have more than one of the IoT tags? Can we fix this to avoid duplicate content appearing? Thanks again, the page looks great!	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6873	DEVELOPER-2886	Hi Adela, When I click on the Orange box it just re-loads the IoT Topics page. The URL should be http://developers.redhat.com/video/youtube/mli5c-oTN1U/ Is that happening for you? Thanks so much for the quick turn-around, it is greatly appreciated. -Michele	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6874	DEVELOPER-2886	Looks great, thank you. What are the next steps? When will this page go live?	1	1	https://issues.jboss.org/browse/DEVELOPER-2886
6875	DEVELOPER-2880	[~mikgue] it looks like you have the Ascidoc, can you raise a PR for it? [~adelaarreola] should be able to help you out with it if you have problems.	1	1	https://issues.jboss.org/browse/DEVELOPER-2880
6876	DEVELOPER-2874	[~mikgue] we don't do camel case for URLs, does http://developers.redhat.com/promotions/devnation-discount-code work for you?	1	1	https://issues.jboss.org/browse/DEVELOPER-2874
6877	DEVELOPER-2870	[~lincolnthree] [~dcoughlin1] Do we have a clear idea of what items originate from *.redhat.com domains? If so could they be added to the definition of done?	1	1	https://issues.jboss.org/browse/DEVELOPER-2870
6878	DEVELOPER-2869	[~adelaarreola] I've added the content in PR #988 [https://github.com/redhat-developer/developers.redhat.com/pull/988] Do any of the templates need to change to add the 5th tab?	1	1	https://issues.jboss.org/browse/DEVELOPER-2869
6879	DEVELOPER-2869	[~adelaarreola] Mike found some typos. I tried to edit on GitHub, but it keeps marking entire sections as changed that I haven't been touched. (The section starting with the alternate install instructions (all 5 copies of it gets marked as changed and I think it's loosing the headings/links in the process. Can you please pick up the file from [https://github.com/rctrh/rh-tthw/blob/master/install-rhel7-combined.adoc] (Raw: https://raw.githubusercontent.com/rctrh/rh-tthw/master/install-rhel7-combined.adoc). I verified that there were no changes against the `get-started.adoc` in your PR. You can view the commit to see the diff here: [https://github.com/rctrh/rh-tthw/commit/90ad209850c7854b8973ed3c7dd116537d7013cf] Thanks & sorry for the trouble.	1	1	https://issues.jboss.org/browse/DEVELOPER-2869
6880	DEVELOPER-2868	I think this issue has been resolved. I don't see the book error on Blinkr anymore. EDIT: [~paul.robinson] - Can this be closed? http://developers-pr.stage.redhat.com//pr/1004/blinkr/887/report.html	0	0	https://issues.jboss.org/browse/DEVELOPER-2868
6881	DEVELOPER-2866	[~james.cobb] If I understand #1 correctly, you want the word 'download' on the 'Try Red Hat Enterprise Linux' box to by a link to download RHEL... However, since the box itself is a link, should I remove the box's link? [~ecurley] Please let me know when the content is ready so I can add it :)	1	1	https://issues.jboss.org/browse/DEVELOPER-2866
6882	DEVELOPER-2866	[~adelaarreola] - Can the supplied text be formatted as a smaller <p> tag?	1	1	https://issues.jboss.org/browse/DEVELOPER-2866
6883	DEVELOPER-2862	what is the content for the new hero?	1	1	https://issues.jboss.org/browse/DEVELOPER-2862
6884	DEVELOPER-2847	So you've now got a URL which will cause the right sections to be visible? Do you change the change the "current URL" when a tab is selected? How does the download manager get the URL to return to?	1	1	https://issues.jboss.org/browse/DEVELOPER-2847
6885	DEVELOPER-2847	[~adelaarreola] Links to the tabs looks and works great! One edit/question: * On the second big box at the top (the one that changes) now, can the title on the corresponding blue bar (further down the page) be updated as well? Right now it says "2. Install, Configure, and Verify" and does not change with the tab selection like the top big blue box does.	1	1	https://issues.jboss.org/browse/DEVELOPER-2847
6886	DEVELOPER-2837	Wow - big section. Is this really a "troubleshooting". Should there be a point above that says "If you see something different, blah, blah...go to this section..."?	1	1	https://issues.jboss.org/browse/DEVELOPER-2837
6887	DEVELOPER-2830	[~rterzi] I was expect 4 steps. How do we get there?	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6888	DEVELOPER-2830	[~paul.robinson] I'm happy to do PRs. However last time I submitted one you decided since I couldn't see the output I should just be handing the changes to Adela. I don't think automatic CI builds were happening on PRs then. How much effort is it to setup a development environment? I'm comfortable with docker/vagrant boxes/RHEL, etc. I have a RHEL VM box that I do all my Red Hat VPN work in.	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6889	DEVELOPER-2830	[~rterzi] are you all set? Need anything?	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6890	DEVELOPER-2830	I thought [~adelaarreola] was working on this to deal with the template/design questions at least for the getting-started pages. Did it ever get integrated in stage? I haven't seen anything. (The 4 CDK TTHWs changes are straightforward, I can do a PR for those if necessary)	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6891	DEVELOPER-2830	Is there a branch for the CDK updates?	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6892	DEVELOPER-2830	Ok, what's the new target date for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6893	DEVELOPER-2830	[~adelaarreola] - Should the tabs still have the blocks [.large-17.columns.recommended] and [.large-17.columns.recommended]	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6894	DEVELOPER-2830	[~rterzi] Adela needs to know: do we want all links including internal ones to open in a new window? She doesn't understand your asciidoc.	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6895	DEVELOPER-2830	[~adelaarreola] - Yes, the links should open in new tabs similar to RHEL. I've updated [https://github.com/rctrh/rh-tthw/blob/master/cdk-get-started-combined.adoc] Just like the RHEL getting started guide the first download should NOT open in a new window, it should trigger the download manager / registration flow. Sorry, I forgot about the new download flow when updating this. Where there any other problems with the AsciiDoc?	1	1	https://issues.jboss.org/browse/DEVELOPER-2830

#	Requirement	Questions	Answered	Successfully Answered	URL
6896	DEVELOPER-2830	[~adelaarreola] I picked up your changes. What are the many syntax errors? I see you removed the sections I left in, but commented out, is // not valid for leaving comments in the AsciiDoc? I don't get an error when I run it through asciidoctor or in github's renderer. This is the diff between your version and mine. [https://github.com/rctrh/rh-tthw/commit/007f52f328803ac0c56c1e092eaf1af424278eb8] Note: Most of the changes are related to the _window='_blank' which I had added to my version. Where do you want me to submit the changes to add back the window='_blank'? [https://github.com/rctrh/rh-tthw/blob/master/cdk-get-started-combined.adoc] is now up to date. Let me know if you want me to submit a PR and which branch/repo I should submit it against at this point. Also, What tool are you using for editing AsciiDoc? Do you have something that is checking syntax? Thanks	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6897	DEVELOPER-2830	Yes, the tabs in step 4 look fine. Should "NODEJS" be changed to "NODE.JS"? Otherwise, good by me.	1	1	https://issues.jboss.org/browse/DEVELOPER-2830
6898	DEVELOPER-2829	[~paul.robinson] [~adelaarreola] - I've tweaked a screenshot to reflect the new color bands. Can we put these changes in the queue for FE?	1	1	https://issues.jboss.org/browse/DEVELOPER-2829
6899	DEVELOPER-2829	[~james.cobb] this looks like a small tweak, is it something you or your team can do?	1	0	https://issues.jboss.org/browse/DEVELOPER-2829
6900	DEVELOPER-2828	[~adelaarreola], [~wesbos]. I don't think it's feasible to return the user to the previous section. However, I think we can make the tab-selection more sticky. How about: 1. When a user selects a tab, or the page is loaded with a tab specified in the URL: Save the selected tab in local browser storage. 2. When the user visits the page (and no tab is specified in the URL): Select the tab specified in local browser storage (if one set for that product).	1	1	https://issues.jboss.org/browse/DEVELOPER-2828
6901	DEVELOPER-2828	This is certainly better than what we have. Would it be possible to return people to the section within a tab (1 through 4)?	1	1	https://issues.jboss.org/browse/DEVELOPER-2828
6902	DEVELOPER-2819	[~paul.robinson], can we put these up on Jboss.org? http://static.jboss.org/images/hero/hero_desktop_1175x400_devnation2016.png http://static.jboss.org/images/hero/hero_mobile_480x500_devnation2016.png	1	1	https://issues.jboss.org/browse/DEVELOPER-2819
6903	DEVELOPER-2809	[~james.cobb] Thanks a lot! [~ecurley] what should the url be?	1	1	https://issues.jboss.org/browse/DEVELOPER-2809
6904	DEVELOPER-2808	[~james.cobb] Where can I find the graphic?	1	1	https://issues.jboss.org/browse/DEVELOPER-2808
6905	DEVELOPER-2805	Just to confirm, so we want to link to Customer Portal for the Red Hat Enterprise Linux 7.2 CDK Vagrant Box? Is there a way to include that on our site too at some point? If we are going to the Portal, I'd like for the link to have the little external link indicator and then open in a new window/tab.	1	1	https://issues.jboss.org/browse/DEVELOPER-2805
6906	DEVELOPER-2803	[~mikgue] Where should I place this content? Also, GoogleDocs does not allow me to save images from it's documents.	1	1	https://issues.jboss.org/browse/DEVELOPER-2803
6907	DEVELOPER-2803	Sorry - I thought I had written it down - this will be at the very beginning of Step 2. For all tabs. How shall I get you the image?	1	1	https://issues.jboss.org/browse/DEVELOPER-2803
6908	DEVELOPER-2803	Is there a way to add a local file to a Jira ticket?	1	1	https://issues.jboss.org/browse/DEVELOPER-2803
6909	DEVELOPER-2789	Does this still need to be done?	1	1	https://issues.jboss.org/browse/DEVELOPER-2789
6910	DEVELOPER-2786	This made sense a month ago, but no longer (longer story). Can we change the text to read: Got a technical CDK question? Please join the CDK beta mailing list, container-tools@redhat.com, and get answers from the Red Hat CDK Team. Thanks, Mike	1	1	https://issues.jboss.org/browse/DEVELOPER-2786
6911	DEVELOPER-2784	Are you happy with the left-side nav and content jumping up, when on the get-started page? E.g: view https://www.dropbox.com/s/lxp7wqu26m8b9aa/Screenshot%202016-03-18%2010.42.54.png?dl=0 then view: https://www.dropbox.com/s/2u0bkuz39eu7dhm/Screenshot%202016-03-18%2010.43.14.png?dl=0	1	1	https://issues.jboss.org/browse/DEVELOPER-2784
6912	DEVELOPER-2780	[~ecurley][~ray.ploski] - Can you please review? ^^^^	1	1	https://issues.jboss.org/browse/DEVELOPER-2780
6913	DEVELOPER-2779	Is this complete?	1	1	https://issues.jboss.org/browse/DEVELOPER-2779
6914	DEVELOPER-2776	What 'associated content' are you expecting at this stage? I didn't think there would be anything beyond the individual topic pages?	1	1	https://issues.jboss.org/browse/DEVELOPER-2776
6915	DEVELOPER-2769	Design approved: https://issues.jboss.org/browse/DEVELOPER-2768 Approved design: http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/pages/images/RHD_downloads_r2v5.png psd: http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/pages/working/RHD_downloads_r2v5.psd Need any graphics processed?	1	1	https://issues.jboss.org/browse/DEVELOPER-2769
6916	DEVELOPER-2769	[~wesbos] Feel free to ping me on HipChat if you have any questions. For the product names, order and description text, you can use the wireframe: https://redhatdevelopers.mybalsamiq.com/projects/octopus-build-2016/Downloads [~mikgue] Can you or someone on your team provide the version and date information for each download? Is one easier than the other to get? I believe we have a different story to make this automated, so we would just need something for the interim.	1	1	https://issues.jboss.org/browse/DEVELOPER-2769
6917	DEVELOPER-2769	[~tnolan]how should I capture them? BTW - I know nothing about these version numbers and dates so will need to find someone who knows.	1	1	https://issues.jboss.org/browse/DEVELOPER-2769
6918	DEVELOPER-2769	[~mikgue] Ok, let's not worry about them for this version of the page then. [~wesbos] For the sections where we have a version or date being displayed, can you add placeholders for them in the code, but not display any information for now?	1	1	https://issues.jboss.org/browse/DEVELOPER-2769
6919	DEVELOPER-2769	That's great news about the versioning and dates. It looks like the PR finished and the page can be previewed via: http://developers-pr.stage.redhat.com/pr/965/build/681/downloads/ [~wesbos] Is this ready for feedback or are you still working on a few things?	1	1	https://issues.jboss.org/browse/DEVELOPER-2769
6920	DEVELOPER-2769	[~wesbos] Are the descriptions being pulled in dynamically from somewhere? [~mikgue] provided ones to be specifically used on this page. I also noticed that the banner background does not extend past 1800px now. Can this be fixed? Additional updates to the text: * JBOSS DEVELOPMENT AND MANAGEMENT should be: ACCELERATED DEVELOPMENT AND MANAGEMENT * CLOUD PRODUCTS should be "PRIVATE CLOUD" * The products within each section should be ordered alphabetically now as shown in the wireframe. [~james.cobb] "I'm for having the link there but it goes to the actual download. Folks see the name, click on the name and get the download." I'd like to remove the link for now as I requested and will keep an eye on heatmaps especially the first few weeks this is live to see if users try to click it.	1	1	https://issues.jboss.org/browse/DEVELOPER-2769
6921	DEVELOPER-2769	[~tnolan] - 1) All of the header graphics on the main pages are designed to tile horizontally. [~wesbos] - Can we change the css to have it repeat for larger sizes? 2) The title link - sure... monitoring the heat maps sounds good to me.	1	1	https://issues.jboss.org/browse/DEVELOPER-2769
6922	DEVELOPER-2769	im not sure then - the code is there and working locally. I can't see staging and no one seems to know why. I'd just merge this and worry about the bg	1	1	https://issues.jboss.org/browse/DEVELOPER-2769
6923	DEVELOPER-2768	I like the first one r2v1 better for the All section since that puts more focus on the download button. The header styles well to separate the categories so I'm ok with the icons being removed. ;) Can you come up with an alternate option for the Most Popular? I'm losing that section in the header as it looks like it's part of the header as opposed to page content.	1	1	https://issues.jboss.org/browse/DEVELOPER-2768
6924	DEVELOPER-2767	Can we aim to have the review completed today on this? This would allow us to ideally start the implementation next week.	0	0	https://issues.jboss.org/browse/DEVELOPER-2767

#	Requirement	Questions	Answered	Successfully Answered	URL
6925	DEVELOPER-2767	I like it and good by me. Well done [~tnolan] Will this page become editable via Drupal down the road?	1	1	https://issues.jboss.org/browse/DEVELOPER-2767
6926	DEVELOPER-2767	[~mikgue] Thanks for adding these so quickly. Can you make a few minor wording tweaks to have them all start with nouns so they are consistent? I think a few of them start with an action verb instead. If it's easier to keep the short with using a verb, then we can do that across the board too. Either way, I think they should use a consistent format.	1	1	https://issues.jboss.org/browse/DEVELOPER-2767
6927	DEVELOPER-2762	[~ray.ploski] - Can you please confirm this content is complete and accurate?	1	1	https://issues.jboss.org/browse/DEVELOPER-2762
6928	DEVELOPER-2761	[~ray.ploski] and [~akandel], it looks like one of the promos either has been changed out already or hasn't been changed out yet. It's not prescribed in the Google doc, but would you like me to remove Red Hat from the CDK image as well?	1	1	https://issues.jboss.org/browse/DEVELOPER-2761
6929	DEVELOPER-2761	Thanks Ray. [~paul.robinson][~akandel] - How do we proceed from here?	1	1	https://issues.jboss.org/browse/DEVELOPER-2761
6930	DEVELOPER-2757	[~morgan_thomas] - Please note this goes out before the .NET related work.	1	1	https://issues.jboss.org/browse/DEVELOPER-2757
6931	DEVELOPER-2757	[~ecurley] I can do the top banner for JBoss, is there a Jira for that?	1	1	https://issues.jboss.org/browse/DEVELOPER-2757
6932	DEVELOPER-2753	[~ecurley] Have these links been created? If so, where can I find them?	1	1	https://issues.jboss.org/browse/DEVELOPER-2753
6933	DEVELOPER-2753	Put them on the .NET mock up: https://docs.google.com/a/redhat.com/presentation/d/1hNJ4v8dTxhns1yZmvm5arQ-bDgSY_s2mrstJIeu3F9c/edit?usp=sharing	1	1	https://issues.jboss.org/browse/DEVELOPER-2753
6934	DEVELOPER-2750	Thanks [~adelaarreola]. Is it possible to change the font color for the side links to highlight which section you're in as the user scrolls down the page, like on the Quickstarts? http://developers.redhat.com//quickstarts/unifiedpush/push-helloworld-ios/index.html Other than that, this looks great!	1	0	https://issues.jboss.org/browse/DEVELOPER-2750
6935	DEVELOPER-2734	Thanks [~adelaarreola] A few minor changes: Get an invite as Red Hat co-hosts cocktail hour. - remove "." Can we pull the YouTube link? And the register for DevNation is hitting a redirect page - is that because it's in staging?	1	1	https://issues.jboss.org/browse/DEVELOPER-2734
6936	DEVELOPER-2734	[~ecurley] here's the latest, can you review? http://developers-pr.stage.redhat.com/pr/903/build/581/events/msbuild/2016/ Once you approve, I'll push it into production.	1	1	https://issues.jboss.org/browse/DEVELOPER-2734
6937	DEVELOPER-2734	[~adelaarreola] Link wasn't showing updates. Am I missing them?	1	1	https://issues.jboss.org/browse/DEVELOPER-2734
6938	DEVELOPER-2734	[~ecurley] That's very odd, are you sure you were looking at the right preview link? I know I sent you too many yesterday :D Here is a screenshot of what has changed https://www.dropbox.com/s/j2oswov1u47dbh1/Screenshot%202016-03-24%2009.01.16.png?dl=0	1	1	https://issues.jboss.org/browse/DEVELOPER-2734
6939	DEVELOPER-2732	Content available here: https://docs.google.com/a/redhat.com/document/d/1orWHd6ELOGYqExh2LWjrvtuyoCL5lkSc1DozWS4PVSg/edit?usp=sharing [~ray.ploski] - Mike did a once through. Do you want to review or shall [~adelaarreola] move forward?	0	0	https://issues.jboss.org/browse/DEVELOPER-2732
6940	DEVELOPER-2716	Thanks [~dcoughlin1]. Does your import process handle sub directories? I don't see many new quickstarts there. The beginner, camel, cxf and switchyard directories (at https://github.com/jboss-fuse/quickstarts/tree/6.2.1.GA) contain numerous quickstarts each.	1	1	https://issues.jboss.org/browse/DEVELOPER-2716
6941	DEVELOPER-2716	[~dcoughlin1] this latest version looks great - thanks! Could you pass on some details on the QSTools checks we need to pass for the switchyard quickstarts?	1	1	https://issues.jboss.org/browse/DEVELOPER-2716
6942	DEVELOPER-2716	[~rafabene] - The fuse quickstarts have a nested file structure. Is it possible to run qstools on that? We are currently having problems rendering the switchyard directory because the READMEs don't follow the qstools conventions.	1	1	https://issues.jboss.org/browse/DEVELOPER-2716
6943	DEVELOPER-2716	Hey [~aileenc], would you like me to point the "Demo Source" link to https://github.com/jboss-fuse/quickstarts/tree/master/beginner?	1	1	https://issues.jboss.org/browse/DEVELOPER-2716
6944	DEVELOPER-2711	This is really looking great, [~adelaarreola]! Some updates for the next iteration: # Change the Thank you message to "Thank you for downloading Red Hat Enterprise Linux." (first line) and "Your download should start automatically. If you have any problems with the download, please use the [direct link]." (second line). All other text should be removed. # Virtual Box is one word "VirtualBox" # Step 1 boxes should be "Download Red Hat Enterprise Linux". # Remove the times (10 minutes, etc) from the boxes, both up top and within the page text. Only on this Get Started page not throughout the site. # Collapse step 1 text if the Thank you message is there. # Change the step 4 layout to use vertical tabs. See the latest wireframe. https://redhatdevelopers.mybalsamiq.com/projects/octopus-build-2016/RHEL%20-%20Get%20Started%20%28post-download%29. [~james.cobb] Can you provide the styling for these? # For the step 4 content, can we keep them all the same height? I think Python is the tallest one, so maybe use that height as a gauge? # We talked about moving the Troubleshooting and FAQ at the bottom of the page elsewhere. [~mikgue] Is that something we still plan to do?	1	1	https://issues.jboss.org/browse/DEVELOPER-2711
6945	DEVELOPER-2711	We're really getting there! Thanks [~adelaarreola] Some additional edits/fixes: # The Thank you message at the top looks like it repeats "Red Hat Enterprise Linux" in the text. I see "Thank you for downloading Red Hat Enterprise Linux {color:red}Red Hat Enterprise Linux{color}". # We need a space between _the_ and _direct_ in the additional sentence in the Thank you message. "please use thedirect link." # Step 4 content: [~mikgue] It looks like we're using "CPP" now instead of C/C++. Is that intentional? # Step 4 layout/style: A few things and [~james.cobb] please chime in. A) I'd to see a different font color for "Recommended" to highlight it and make it stand out more. B) I'd like to see a larger font size used for the first item in the list, the recommended one, to differentiate it more from the others. See the Perl tab. C) Can we add lines or some other separation between the items like on the Perl tab?	1	1	https://issues.jboss.org/browse/DEVELOPER-2711
6946	DEVELOPER-2710	[~paul.robinson] Paul, what are you looking for me to do? It looks like you just re-assigned this to me.	1	1	https://issues.jboss.org/browse/DEVELOPER-2710
6947	DEVELOPER-2710	I set you as the assignee, as you are the one working on it. Are you asking Tiffany to review it? If so, can you mark it as 'in review'. See: https://www.dropbox.com/s/se0c3laxpauk8qr/Screenshot%202016-03-14%2013.57.17.png?dl=0 for a pointer on how to do it.	1	1	https://issues.jboss.org/browse/DEVELOPER-2710
6948	DEVELOPER-2710	Can we aim to have the review completed today on this? This would allow us to do the implementation next week.	1	1	https://issues.jboss.org/browse/DEVELOPER-2710
6949	DEVELOPER-2710	[~adelaarreola] what time should we get updates to you to be part of Prod on Monday?	1	1	https://issues.jboss.org/browse/DEVELOPER-2710
6950	DEVELOPER-2709	Due to some changes to the content and learning that both steps 2 and 3 could change depending on the user's desired configuration, I made a new version of the wireframe and am placing this sub-task back into review status. Here is the proposed new wireframe: https://redhatdevelopers.mybalsamiq.com/projects/octopus-build-2016/RHEL%20-%20Get%20Started%20%28post-download%29 Because I'd like to keep this to a single page for the RHEL "Get Started", I also made a few modifications: # I added a primary tab bar to the top of the page. - It is very important that we are able to swap the content for the tabs WITHOUT reloading the entire page. This is due to people hitting this page directly after download. We want to retain the "Thanks" as well as showing step one completed no matter which primary tab they select after the download. [~paul.robinson] - Is this possible? # We would need a new secondary tab style for the Languages section. [~james.cobb] is looking into some options for this.	1	1	https://issues.jboss.org/browse/DEVELOPER-2709

#	Requirement	Questions	Answered	Successfully Answered	URL
6951	DEVELOPER-2709	Thanks [~james.cobb]. Do you consider that style a primary or secondary style then? We need to go two levels deep on the RHEL Get Started pages. Tab set one: Bare Metal, Hyper-V, VirtualBox, VMware Each of those has multiple steps (1-4) Within Step 4, we would need another set of tabs for the language selection, which I think would be considered a secondary set of tabs. So it would go Primary Tabs > Steps > Secondary Tabs. Feel free to ping me if this doesn't make sense (which is quite likely, lol)	1	1	https://issues.jboss.org/browse/DEVELOPER-2709
6952	DEVELOPER-2709	Why are we trying to put all of this content on one page? We already have a multi-language selection method for Getting Started as seen here: http://developers.redhat.com/products/rhel/get-started-rhel7-perl/ Is there a business need to have all of this content on "one page"? Seems to me that we're adding visual complexity and load time as we'll need to load all of the tabbed content.	1	1	https://issues.jboss.org/browse/DEVELOPER-2709
6953	DEVELOPER-2709	Can we aim to have the review completed today on this? This would allow us to do the implementation next week.	1	1	https://issues.jboss.org/browse/DEVELOPER-2709
6954	DEVELOPER-2703	[~mikgue] Where can I find this content?	1	1	https://issues.jboss.org/browse/DEVELOPER-2703
6955	DEVELOPER-2703	[~mikgue] is this the issue you are working on here: https://github.com/redhat-developer/developers.redhat.com/pull/964 I think the plan was for [~adelaarreola] to take Mike's code from that PR and fix the problems. [~adelaarreola] do you have everything you need to do that? For example, do you know what it is intended to look like?	1	1	https://issues.jboss.org/browse/DEVELOPER-2703
6956	DEVELOPER-2703	A few things: # delete all of the old stuff - everything above "New to RHEL" (except the title) # can you standardize on the table widths? # can you align the columns of the various tables (similar to item above) # can you remove a space/line above each table so that it's closer to the text above it? # For the Perl table, can you replace "docker pull rhscl/perl-520-rhel7" with this "docker pull registry.access.redhat.com/hscl/perl-520-rhel7" Looks great! Thanks. Mike	1	1	https://issues.jboss.org/browse/DEVELOPER-2703
6957	DEVELOPER-2699	[~paul.robinson] Did it change from what we had in design? $0 in multiple languages/currency Text below: Getting RHEL is easier than ever. New $0 RHEL Developer Suite. CTA: Learn more. 3- 5 assets: Get started with a Read Me. Download now Using the languages you know. GOOGLE DOC: https://docs.google.com/a/redhat.com/document/d/1AswVoPh1gAxrcc2m3fYtoxxfjYn-YHzCE_WvPllAdXo/edit?usp=sharing Scroll down past blue bar.	1	1	https://issues.jboss.org/browse/DEVELOPER-2699
6958	DEVELOPER-2699	[~adelaarreola] Paul has marked this as done. Might you know when it will show on Prod? cc [~ecurley]	0	0	https://issues.jboss.org/browse/DEVELOPER-2699
6959	DEVELOPER-2698	[~ecurley] - Do we have content for the .NET topic page?	1	1	https://issues.jboss.org/browse/DEVELOPER-2698
6960	DEVELOPER-2698	[~james.cobb] [~brock] where can I find the graphics?	1	0	https://issues.jboss.org/browse/DEVELOPER-2698
6961	DEVELOPER-2696	[~brock] Survey report isn't done yet. Sorry. Let's swap with a blog that has been getting tons of hits: http://developers.redhat.com/blog/2016/02/24/10-things-to-avoid-in-docker-containers/ Work for you [~lincolnthree]?	1	1	https://issues.jboss.org/browse/DEVELOPER-2696
6962	DEVELOPER-2695	Here are the homepage backgrounds for .NET [~james.cobb] do you approve? Also the desktop version does pattern.	1	1	https://issues.jboss.org/browse/DEVELOPER-2695
6963	DEVELOPER-2695	[~ecurley] - Do we have new content for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-2695
6964	DEVELOPER-2691	[~paul.robinson] I can take a look during my next sprint which starts mid-week next week. Does that work ok for timing?	1	1	https://issues.jboss.org/browse/DEVELOPER-2691
6965	DEVELOPER-2691	[~tnolan] No problem, waiting makes sense as we should consider uniformity across RH properties. What throw-away work are you thinking about?	1	1	https://issues.jboss.org/browse/DEVELOPER-2691
6966	DEVELOPER-2686	Hey Jason, this looks interesting so far! Do you want to sit us down for a demo at some point? It's a little hard to guess with confidence at what this means. Wdyt?	0	0	https://issues.jboss.org/browse/DEVELOPER-2686
6967	DEVELOPER-2685	[~lightguard] what's the status of this? Do you have a URL at which I can view the site?	1	1	https://issues.jboss.org/browse/DEVELOPER-2685
6968	DEVELOPER-2685	Is it this job that will build the pilot environment of Drupal? If so, I think it will be triggered on a change to master, in the same way our current production job is rebuilt. So, it's not really a 'nightly' job. Perhaps the scope of this issue has changed since it was first envisaged?	1	1	https://issues.jboss.org/browse/DEVELOPER-2685
6969	DEVELOPER-2685	Random permission issue * http://atholl.eng.hst.ams2.redhat.com:33777/projects -- Fixed with new data dump * JS not firing on http://atholl.eng.hst.ams2.redhat.com:33777/community/contributor/ -- Still an issue * Only one (badly styled) connector on: http://atholl.eng.hst.ams2.redhat.com:33777/products/fuse/connectors -- JS issue is breaking this, not sure what though * No latest content on http://atholl.eng.hst.ams2.redhat.com:33777/ -- Probably just needs to be updated * No Events on http://atholl.eng.hst.ams2.redhat.com:33777/events * Links on left wrong on http://atholl.eng.hst.ams2.redhat.com:33777/products/cdk/get-started * Javascript on topics pages for adding news items doesn't work * Acceptance tests failing: {code} @site_user @logout @javascript Scenario: Customer verifies their account # _cucumber/features/login_register/basic_registration.feature:10 Given I am on the Registration page # _cucumber/features/step_definitions/common/navigation_steps.rb:1 Expected page title to be Register | Red Hat Developers, but was Red Hat Developers | Red Hat Developers (RuntimeError) ./_cucumber/lib/pages/base_page.rb:37:in `loaded?' ./_cucumber/lib/pages/register_page.rb:26:in `rescue in open' ./_cucumber/lib/pages/register_page.rb:22:in `open' ./_cucumber/features/step_definitions/common/navigation_steps.rb:2:in `/^(I am|they are) on the ([^"]*) page$/' _cucumber/features/login_register/basic_registration.feature:11:in `Given I am on the Registration page' And I complete the registration form # _cucumber/features/step_definitions/register_steps.rb:9 Then I should be logged in # _cucumber/features/step_definitions/common/navigation_steps.rb:10 HTML screenshot: _cucumber/screenshots/screenshot_2016-03-17-20-35-36.334.html Image screenshot: _cucumber/screenshots/screenshot_2016-03-17-20-35-36.334.png HTML screenshot: _cucumber/screenshots/screenshot_2016-03-17-20-35-52.330.html Image screenshot: _cucumber/screenshots/screenshot_2016-03-17-20-35-52.330.png The test failed on page: Red Hat Developers \| Red Hat Developers @site_user @logout @javascript Scenario: Newly registered site visitor navigates to product Download page and clicks on download, accepts Redhat T&C's should initiate download. # _cucumber/features/downloads/featured_downloads.feature:12 Given I register a new account # _cucumber/features/step_definitions/register_steps.rb:1 And I am on the Product Download page for eap # _cucumber/features/step_definitions/downloads.steps.rb:1 When I click to download the featured download of "Enterprise Application Server" # _cucumber/features/step_definitions/downloads.steps.rb:16 @site_user @javascript Scenario: Password validation # _cucumber/features/login_register/basic_registration.feature:17 Given I am on the Registration page # _cucumber/features/step_definitions/common/navigation_steps.rb:1 Expected page title to be Register \| Red Hat Developers, but was Red Hat Developers \| Red Hat Developers (RuntimeError) ./_cucumber/lib/pages/base_page.rb:37:in `loaded?' ./_cucumber/lib/pages/register_page.rb:26:in `rescue in open' ./_cucumber/lib/pages/register_page.rb:22:in `open' ./_cucumber/features/step_definitions/common/navigation_steps.rb:2:in `/^(I am\|they are) on the ([^"]*) page$/' _cucumber/features/login_register/basic_registration.feature:18:in `Given I am on the Registration page' When I try to enter passwords that do not match # _cucumber/features/step_definitions/register_steps.rb:18 Then I should see a warning that "Passwords don't match" # _cucumber/features/step_definitions/register_steps.rb:24 HTML screenshot: _cucumber/screenshots/screenshot_2016-03-17-20-36-16.042.html Image screenshot:	1	1	https://issues.jboss.org/browse/DEVELOPER-2685

#	Requirement	Questions	Answered	Successfully Answered	URL
6970	DEVELOPER-2685	Can you assign those issues to the 'Drupal Live' release? If it's a blocker for the pilot, then put it on the 'Drupal Pilot' release. Thanks!	0	0	https://issues.jboss.org/browse/DEVELOPER-2685
6971	DEVELOPER-2681	How is this different than a link?	1	1	https://issues.jboss.org/browse/DEVELOPER-2681
6972	DEVELOPER-2681	[~lightguard] did you mean to mark as 'implemented'? If so, what's the rational?	1	1	https://issues.jboss.org/browse/DEVELOPER-2681
6973	DEVELOPER-2680	[~paul.robinson] [~dhladky] When will RHEL have a DM link from here? http://developers.stage.redhat.com/downloads/ Or is it somewhere else that I don't see.	1	0	https://issues.jboss.org/browse/DEVELOPER-2680
6974	DEVELOPER-2679	[~wesbos] is this feasible? [~lkrzyzanek] when do you need this by?	1	1	https://issues.jboss.org/browse/DEVELOPER-2679
6975	DEVELOPER-2679	Is it pushed to production already ?	1	1	https://issues.jboss.org/browse/DEVELOPER-2679
6976	DEVELOPER-2677	[~mikgue] - Does this work for you? ^^	1	1	https://issues.jboss.org/browse/DEVELOPER-2677
6977	DEVELOPER-2676	Can you let me know what the URL will be for this page?	1	1	https://issues.jboss.org/browse/DEVELOPER-2676
6978	DEVELOPER-2669	[~rblake] it seems like #2 relates to DEVELOPER-2541? Can you tidy these issues up, as you are suggesting? Thanks.	1	1	https://issues.jboss.org/browse/DEVELOPER-2669
6979	DEVELOPER-2666	"Developers have multiple options when seeking to obtain a Red Hat Enterprise Linux developer subscription - see the details here[link]" - Where should I direct the link on this?	1	1	https://issues.jboss.org/browse/DEVELOPER-2666
6980	DEVELOPER-2666	Do I need to make any changes to this?	1	1	https://issues.jboss.org/browse/DEVELOPER-2666
6981	DEVELOPER-2666	[~mikgue] I submitted a new PR for this, but I think I'm a little confused on the requirement. Is it supposed to mirror the RHEL downloads page and include that extra line?	1	1	https://issues.jboss.org/browse/DEVELOPER-2666
6982	DEVELOPER-2666	[~dcoughlin1] This looks good. Can you make the first lines bold to make sure people read it? other than that - good to go.	1	1	https://issues.jboss.org/browse/DEVELOPER-2666
6983	DEVELOPER-2665	Is http://developers.redhat.com/downloads/ really the right place? Isn't http://developers.redhat.com/products/rhel/download/ better?	1	1	https://issues.jboss.org/browse/DEVELOPER-2665
6984	DEVELOPER-2662	[~paul.robinson] We'd like to link to 2 different PDFs from this page. What's the process to get these added and indexed on RHD? PDF 1: http://static.jboss.org/rhd/docs/rhel_developer_getting_started_guide.pdf PDF 2: http://static.jboss.org/rhd/docs/selinux_for_red_hat_developers.pdf	1	1	https://issues.jboss.org/browse/DEVELOPER-2662
6985	DEVELOPER-2662	Can we aim to have the review completed today on this? This would allow us to do the implementation next week.	1	1	https://issues.jboss.org/browse/DEVELOPER-2662
6986	DEVELOPER-2662	[~tnolan] You've assigned this to me - what is it you need me to do? :)	1	1	https://issues.jboss.org/browse/DEVELOPER-2662
6987	DEVELOPER-2656	[~mikgue] how will this link from the overview page? Is it a link on the left-hand side? If so, what text should be displayed on the link?	1	1	https://issues.jboss.org/browse/DEVELOPER-2656
6988	DEVELOPER-2653	A couple of things. 1. Shouldn't DTS and SCL stuff be on their own respective Update pages? 2. The other items are don't feel like they'd be part of a RHEL update page, as they really fall under "latest news" (like the page is called). Do we have site plans for "news" items"?	1	1	https://issues.jboss.org/browse/DEVELOPER-2653
6989	DEVELOPER-2653	[~mikgue] We can remove the 2 that are related to DTS and SCL, but that only leaves 2 on the page then, unless you have others to add? How old are these updates anyway and is it really worth including them other than not losing legacy content? For a DTS and SCL Updates tab, I don't think it's worth it to add that tab right now for just one item on those tabs. Same question as above...how old are these? I was also thinking about a "news" item but we do not have that on the site yet. It's possible the blog could fill that need.	1	1	https://issues.jboss.org/browse/DEVELOPER-2653
6990	DEVELOPER-2653	I'm getting confused. Isn't the point for an Updates tab to show history? What's there for DTS and SCL are from the Fall timeframe. If one is insufficient, then we can add a prior release, too. The two other items are less important by themselves, but RHEL does have an obvious history as well.	1	1	https://issues.jboss.org/browse/DEVELOPER-2653
6991	DEVELOPER-2653	[~mikgue] Ok, thanks for the additional info about the dates on those. It was hard to tell how recent they were from the CP site. We can add updates tabs for those other 2 then, but would need a separate JIRA for each and we would need to make sure we maintain those too as updates roll out. For RHEL Updates, we would only display 2 items for now, unless you have new ones that we can add. Please let me know if we need more items on this page than the 2. With the DTS and SCL landing pages being removed from the menu, how are users supposed to get to those pages now anyway? This is probably a side convo. :)	1	1	https://issues.jboss.org/browse/DEVELOPER-2653
6992	DEVELOPER-2653	So I guess I should create 2 Jira's for SCL and DTS Updates pages, yes? I'll assemble the content. As for the RHEL Update, I'll see what else I can find. I don't follow you on this one: "With the DTS and SCL landing pages being removed from the menu, how are users supposed to get to those pages now anyway? Being removed from which menu. If you mean, the CP site, it's up to us. (that's why I started this convo. :)) Much of the content is also available in a PDF/product brief but I recall that you don't like PDFs - would like understand why someday. Maybe during our side convo. :)	0	0	https://issues.jboss.org/browse/DEVELOPER-2653
6993	DEVELOPER-2644	[~paul.robinson] I can't define the definition of done for this ticket as there is no description. Could you please provide some context?	1	1	https://issues.jboss.org/browse/DEVELOPER-2644
6994	DEVELOPER-2644	[~adelaarreola] the content is here: https://issues.jboss.org/browse/DEVELOPER-2699. Can you confirm that you have all you need to implement this today?	0	0	https://issues.jboss.org/browse/DEVELOPER-2644
6995	DEVELOPER-2642	[~adelaarreola] This looks very good. Can you insert a line under "Download Red Hat Enterprise Linux 7"? I want the text underneath to not be associated with the link. After that, good to go.	1	1	https://issues.jboss.org/browse/DEVELOPER-2642
6996	DEVELOPER-2632	[~lightguard] this seems unnecessary now. WDYT? please close if you agree.	0	0	https://issues.jboss.org/browse/DEVELOPER-2632
6997	DEVELOPER-2617	[~mikgue] do you want any changes to the Spotlights or 'FEATURED CONTENT' area of the homepage? If so, can you create new issues under this Epic.	1	1	https://issues.jboss.org/browse/DEVELOPER-2617
6998	DEVELOPER-2617	[~ecurley] For Octopus, we need move the JEE book to a spotlight. What should come off?	1	1	https://issues.jboss.org/browse/DEVELOPER-2617
6999	DEVELOPER-2615	Where will this link from?	1	1	https://issues.jboss.org/browse/DEVELOPER-2615
7000	DEVELOPER-2615	I'm not sure what you mean by 'resource'. Specifically what type of resource? Blog post, article, or something else? Can you provide a link to the page you would like it linking from and indicate where on the page you want the link? This might help me to understand what your requirements are.	1	1	https://issues.jboss.org/browse/DEVELOPER-2615
7001	DEVELOPER-2615	[~paul.robinson] - What branch to use as the base? Where should images (screenshots) go?	1	1	https://issues.jboss.org/browse/DEVELOPER-2615

#	Requirement	Questions	Answered	Successfully Answered	URL
7002	DEVELOPER-2615	FYI - for anyone tracking this. There are three pieces that are part of this: # Install RHEL 7 (physical machine) - Is this a ticket for this one? # Install RHEL 7 VirtualBox VM - DEVELOPER-2613 # Install RHEL 7 VMware, this one, DEVELOPER-2615 This one, VMware, will be the *last one published*. The first one will be the Install RHEL 7. (located here: [https://github.com/rctrh/rh-tthw/blob/master/get-started-rhel7-install.adoc]. However, These were originally formatted to be TTHWs with the 3 sections/3 boxes. I will reformat and update for an article layout since that is the new direction. On the TTHWs there was an AsciiDoctor macro I think for locating the image. Is it the same for articles? It doesn't work in the local repo, it just gives broken images. Is there a trick for that? I've cloned the repo, but I don't have a development site setup to test against, so I won't be able to see what it looks like until it gets staged. {{ image:#{cdn(site.base_url + '/images/products/rhel/rhel7-install/rhel_7_1_a_0.png')}[Screenshot Boot Splash]}}	1	1	https://issues.jboss.org/browse/DEVELOPER-2615
7003	DEVELOPER-2615	Is there a naming and/or directory hierarchy for articles? These three articles will be for RHEL Thanks.	1	1	https://issues.jboss.org/browse/DEVELOPER-2615
7004	DEVELOPER-2614	What link are you waiting for? I suggest getting the PR in if you don't get the link before the end of the week.	0	0	https://issues.jboss.org/browse/DEVELOPER-2614
7005	DEVELOPER-2614	([~mikgue] I should probably take this to a separate email thread. That sentence needs a little minor improvement. You will need a Red Hat Enterprise Linux 7 system with a current Red Hat subscription that allows you to download software and updates from Red Hat. If you don't have a Red Hat Enterprise Linux subscription, register on developers.redhat.com [https://developers.redhat.com] and download Red Hat Enterprise Linux Developer Suite[https://developers.redhat.com/products/rhel/download/] Or maybe make it simpler: You will need a Red Hat Enterprise Linux 7 system with a current Red Hat subscription that allows you to download software and updates from Red Hat. If you don't have a Red Hat Enterprise Linux subscription, get Red Hat Enterprise Linux Developer Suite[https://developers.redhat.com/products/rhel/download/] after registering at developers.redhat.com [https://developers.redhat.com] (both gloss over that you will get/still need a subscription, but it's $0/no-cost/free for developers.) Note: This also should go in the FAQ. It's in all the others, but I left it out of the Java one until we figured out what the text should be. Old: I don't have a current Red Hat subscription, can I get an evaluation? + If you don't have a Red Hat Enterprise Linux subscription, you can try it for free. Get started with an evaluation at link:https://access.redhat.com/products/red-hat-enterprise-linux/evaluation[]. Developers should select the Red Hat Enterprise Linux Developer Workstation option to ensure your evaluation includes additional tools from the Red Hat Developer Toolset and Red Hat Software Collections. New: As a developer, how can I get a Red Hat Enterprise Linux subscription? + If you don't have a Red Hat Enterprise Linux subscription, register on developers.redhat.com [https://developers.redhat.com] and download Red Hat Enterprise Linux Developer Suite[https://developers.redhat.com/products/rhel/download/]	1	1	https://issues.jboss.org/browse/DEVELOPER-2614
7006	DEVELOPER-2614	As discussed, use these: You will need a Red Hat Enterprise Linux 7 system with a current Red Hat subscription that allows you to download software and updates from Red Hat. If you don't have a Red Hat Enterprise Linux subscription, get Red Hat Enterprise Linux Developer Suitehttps://developers.redhat.com/products/rhel/download/ after registering at developers.redhat.com https://developers.redhat.com. As a developer, how can I get a Red Hat Enterprise Linux subscription? + If you don't have a Red Hat Enterprise Linux subscription, register on developers.redhat.com https://developers.redhat.com and download Red Hat Enterprise Linux Developer Suitehttps://developers.redhat.com/products/rhel/download/	1	1	https://issues.jboss.org/browse/DEVELOPER-2614
7007	DEVELOPER-2611	Can you double check the redirect? I'm getting "Not Found - The requested URL /solutions/ was not found on this server." instead of being redirected to the homepage.	1	1	https://issues.jboss.org/browse/DEVELOPER-2611
7008	DEVELOPER-2611	Can we change this back to In-Progress until the redirect is added? I don't want us to lose track of this one, since it's marked as Done currently.	1	1	https://issues.jboss.org/browse/DEVELOPER-2611
7009	DEVELOPER-2607	Good deal! That works. Is there an .svg or .eot file for jboss.org? We need to replace the font icon on http://www.jboss.org/products/webserver/overview/ as well.	1	1	https://issues.jboss.org/browse/DEVELOPER-2607
7010	DEVELOPER-2607	[~brock] - Can you take care of this?	1	1	https://issues.jboss.org/browse/DEVELOPER-2607
7011	DEVELOPER-2607	[~james.cobb] and [~dcoughlin1] Here's a full color .svg: http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/illustrations/svg/products_apache_v2.svg Do we want to make it grayscale to match the other icons?	1	1	https://issues.jboss.org/browse/DEVELOPER-2607
7012	DEVELOPER-2607	[~wesbos] [~adelaarreola] - do you know integrate this into images/icons/fontcustom on jboss.org? I am having some difficulty getting this to work properly.	0	0	https://issues.jboss.org/browse/DEVELOPER-2607
7013	DEVELOPER-2602	[~james.cobb] Where do I put the carousel graphics in static?	1	1	https://issues.jboss.org/browse/DEVELOPER-2602
7014	DEVELOPER-2598	[~dcoughlin1] - Can you share a link were it can be reviewed?	1	1	https://issues.jboss.org/browse/DEVELOPER-2598
7015	DEVELOPER-2598	[~james.cobb] - Should be reviewable here: http://developers-pr.stage.redhat.com/pr/837/build/367/	1	1	https://issues.jboss.org/browse/DEVELOPER-2598
7016	DEVELOPER-2586	[~dcoughlin1] - Hey Dan... this is marked as in review, where can I review it?	1	1	https://issues.jboss.org/browse/DEVELOPER-2586
7017	DEVELOPER-2586	[~paul.robinson][~wesbos] [~dcoughlin1] - The spacing looks fine between the promos and the associated text but another related question came up. At desktop sizes, the width of the promos are much narrower that the rest of the page. I looked into the css and I question why we have a calculated width for these elements. Current CSS... .spotlight-item { width: calc(20% - 16px); I changed it in my inspector to this and it looked great... .spotlight-item { width: calc(20%); Why are we adding the " - 16p" at desktop sizes? Granted, I'm using chrome on mac and there may be a browser concern for this.	1	1	https://issues.jboss.org/browse/DEVELOPER-2586
7018	DEVELOPER-2586	[~james.cobb] if the above turns out to be an issue, can you create a new JIRA?	1	1	https://issues.jboss.org/browse/DEVELOPER-2586
7019	DEVELOPER-2584	[~mikgue] where should this article be linked from? Is due to go live on Project Octopus launch or ASAP?	1	1	https://issues.jboss.org/browse/DEVELOPER-2584
7020	DEVELOPER-2568	Both the items in the screenshot are in the spreadsheet with 2016 dates. If these are wrong, can you update the dates in the spreadsheet?	1	1	https://issues.jboss.org/browse/DEVELOPER-2568
7021	DEVELOPER-2568	Was digging into this - and noticed a few things that aren't pulling up. Not letting me add attachment so will send to you. But dates for DevNation are off. Saying June 25-27th on the site. http://developers.redhat.com/events/ JavaLand should be March 8 - 10th but on site saying the 7th - 17th. And the month listed above the cards isn't matching up with the card. Would it be helpful to screenshare?	1	1	https://issues.jboss.org/browse/DEVELOPER-2568
7022	DEVELOPER-2561	Should I move this to blocking in the meantime?	0	0	https://issues.jboss.org/browse/DEVELOPER-2561

#	Requirement	Questions	Answered	Successfully Answered	URL
7023	DEVELOPER-2561	Yes, I have some notes from a meeting yesterday regarding the breadcrumbs and will be creating a few JIRAs as a result later today. For this story specifically, we do not want to rename the breadcrumbs but to remove them instead. [~paul.robinson] - Should this Story be removed and a new one created or should this one just be revised? Our plan is to remove all breadcrumbs from the site because the content does not go deep enough now to have users get “lost” as to where they are in the sitemap. The new menu and IA really flattened the structure so breadcrumbs no longer add significant value.	1	1	https://issues.jboss.org/browse/DEVELOPER-2561
7024	DEVELOPER-2558	[~wesbos] it looks like the switch to using the REST API on the DL-Man didn't consider the green button. Can you add this pack in please? It should download the featured artefact for that product.	0	0	https://issues.jboss.org/browse/DEVELOPER-2558
7025	DEVELOPER-2557	[~lincolnthree] what new content? Is this now added to the staging build? Can you review it?	0	0	https://issues.jboss.org/browse/DEVELOPER-2557
7026	DEVELOPER-2550	Are you still creating one?	1	1	https://issues.jboss.org/browse/DEVELOPER-2550
7027	DEVELOPER-2545	[~lightguard] have you got enough info to know what the issue is? I know we discussed it on the call. If so, can you confirm with Olly that there is already a feature in Awestruct to do this. Otherwise, it might be worth adding the feature at the same time as the logging improvements you are doing this week.	1	1	https://issues.jboss.org/browse/DEVELOPER-2545
7028	DEVELOPER-2545	Is that working alright for you [~oshaw]?	1	1	https://issues.jboss.org/browse/DEVELOPER-2545
7029	DEVELOPER-2544	[~tnolan] is this in progress? If so, can you update the status please?	1	1	https://issues.jboss.org/browse/DEVELOPER-2544
7030	DEVELOPER-2532	Article or 1 page Microsite would work here i think. When does this need to go live?	1	1	https://issues.jboss.org/browse/DEVELOPER-2532
7031	DEVELOPER-2532	[~mikgue] - Are you moving forward with a Blog post for this content instead? I think we discussed that as an option instead of using an Article.	1	1	https://issues.jboss.org/browse/DEVELOPER-2532
7032	DEVELOPER-2532	[~paul.robinson] this would be needed in a couple of weeks. Don't have a date handy. [~tnolan] I'm afraid I don't recall that. If it's a blog, everyone would see the discount code and we don't want that. With an article, we'd want it behind the Reg wall to force people to register. We can use a blog to announce the discount code, however. Which reminds me how are articles assigned to the front/behind the reg wall?	1	1	https://issues.jboss.org/browse/DEVELOPER-2532
7033	DEVELOPER-2532	[~mikgue] [~lincolnthree] - We haven't talked about where articles would fit with the reg wall. I figured that they would be like most of the content, which does not require you to be logged in, right? Do you have a doc in progress for what should/shouldn't require registration? I don't think I've seen that plan yet or that we have a way to support that from the back-end. I understand the reluctance to use a blog for this, as what you say about not revealing the code to the world makes sense. What's to prevent someone from just sharing the code elsewhere, though? I'm not sure what we can do short-term, but long-term I think we should have a "promotion" content type that would require users to be logged in to see it. I still don't think an article is the right place for something like this that is short-lived content.	1	1	https://issues.jboss.org/browse/DEVELOPER-2532
7034	DEVELOPER-2532	Have we done anything like this in the past? One other option would be for the user to fill out a form, verify they are a member and we email them the code.	1	1	https://issues.jboss.org/browse/DEVELOPER-2532
7035	DEVELOPER-2532	I'm not aware of us doing this before. The form/email would be good. Should we get approval from Harry if he wants to collect more info for the discount?	1	1	https://issues.jboss.org/browse/DEVELOPER-2532
7036	DEVELOPER-2532	[~mikgue] [~paul.robinson] I think we have a promo page live for this now. Is that correct? Can this be closed?	0	0	https://issues.jboss.org/browse/DEVELOPER-2532
7037	DEVELOPER-2529	[~crobson], [~tnolan]: Can you review that this is something we want to do?	1	1	https://issues.jboss.org/browse/DEVELOPER-2529
7038	DEVELOPER-2520	[~mikgue] - Should I plan on just removing those product folders? Or, do I need to remove references to these products as well?	0	0	https://issues.jboss.org/browse/DEVELOPER-2520
7039	DEVELOPER-2511	[~paul.robinson][~wesbos] The new home page promos are ready for implementation. What is RHD? http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_whatisrhd.png DevNation : Registration open http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_devnation_registrationopen.png Dive into Java http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_diveintojava.png EAP7 beta is here http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_eap7betaishere.png New Containers Course http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_containercourse.png Content and links can be found here: https://issues.jboss.org/browse/DEVELOPER-2510	1	1	https://issues.jboss.org/browse/DEVELOPER-2511
7040	DEVELOPER-2511	* Previously What is Red Hat Developer about? Change to: About Red Hat Developers. * DevNation with dates - change to DevNation 2016. Text below - Why wait when you can register now? * Change Java one to Hot Topics - text below Dive into Enterprise Java * Change EAP image to "Download JBoss EAP 7 Beta" * Change text below containers course to "New no-cost containers course from Red Hat Training" [~james.cobb]	1	1	https://issues.jboss.org/browse/DEVELOPER-2511
7041	DEVELOPER-2510	[~james.cobb] What is Red Hat Developers about? <<developers.redhat.com/about>> DevNation registration <<https://rh2016.smarteventscloud.com/portal/createAccount.ww?>> Dive into Java <<http://developers.redhat.com/enterprise-java/>> JBoss EAP 7 Beta is here <<http://developerblog.redhat.com/2016/01/18/jboss-eap-7-beta-is-now-available/>> New containers course <<http://wp.me/p2WBxv-1KWJ>>	1	1	https://issues.jboss.org/browse/DEVELOPER-2510
7042	DEVELOPER-2500	As this ticket currently stands, is it still valid?	1	1	https://issues.jboss.org/browse/DEVELOPER-2500
7043	DEVELOPER-2493	This is done when all problems are in JIRA, that are discovered from: 1. Acceptance tests 2. Visiting all major pages on the site and checking they look right, and the key functionality works I appreciate the above is still a bit vague.	0	0	https://issues.jboss.org/browse/DEVELOPER-2493
7044	DEVELOPER-2491	[~lightguard], [~oshaw]: can you review this description please. Does it give us what we need to review where the bottlenecks are in the system?	1	1	https://issues.jboss.org/browse/DEVELOPER-2491
7045	DEVELOPER-2491	Could we log these to a different file instead? Logging to the console slows things down, it's been observed before, Doing a perf logging in .awestruct probably isn't a bad idea in general, but I'm not sure how much we want to log to the console. Awestruct used to log every page as it was created, but others in the community didn't want that level of granularity in their builds. There's a verbose flag (`-w`) currently, but it doesn't do much if you're not in debug. We could add some additional logging to the console that way without too much trouble.	0	0	https://issues.jboss.org/browse/DEVELOPER-2491
7046	DEVELOPER-2491	I agree logging every page to the console is too much. We have about 1200 pages. What do you think about logging the extension just before it executes?	0	0	https://issues.jboss.org/browse/DEVELOPER-2491

#	Requirement	Questions	Answered	Successfully Answered	URL
7047	DEVELOPER-2491	Sure, verbose is fine. There would be ~50 lines of output if you included all the extensions, helpers and transformers registered in pipeline.rb. Would verbose log any other stuff that we don't want?	0	0	https://issues.jboss.org/browse/DEVELOPER-2491
7048	DEVELOPER-2491	What would we use the perf.log for? I'm wondering if we need that now? The verbose logging of when an extension, helper or transformer runs fulfills the requirements I think. [~oshaw] agreed?	1	1	https://issues.jboss.org/browse/DEVELOPER-2491
7049	DEVELOPER-2491	Do we need it now? I don't know, maybe not but if we're trying to diagnose performance problems (I've had at least one person in the community ask for help) it would be good to have that option. We wouldn't log it every time. While I've worked on the RHD site I've had a few times where I needed to measure performance. I hacked the branch I was working on, but that's as far as it ever went.	0	0	https://issues.jboss.org/browse/DEVELOPER-2491
7050	DEVELOPER-2491	Can you expand on what this perf.log is? Is it a dump from ruby to be analysed by some tooling, or is it a tonne of log statements, kind of like a TRACE level log in Java?	0	0	https://issues.jboss.org/browse/DEVELOPER-2491
7051	DEVELOPER-2488	[~lkrzyzanek], I think we want a named query like: http://dcp2.jboss.org/v2/rest/search/developer_materials. That also includes sys_type of 'blogpost' and sys_type of 'book'. We'd also need to be able to filter by tags (optional) and a query string (optional). Can you create this for us please?	1	1	https://issues.jboss.org/browse/DEVELOPER-2488
7052	DEVELOPER-2488	Shouldn't this be in DONE state? I finished implementing this last week, it was reviewed and pushed to all possible DCP instances.	0	0	https://issues.jboss.org/browse/DEVELOPER-2488
7053	DEVELOPER-2487	BTW, I see the big green "download now" buttons don't exist anymore. Is that a consequence of this change to REST API, too?	1	1	https://issues.jboss.org/browse/DEVELOPER-2487
7054	DEVELOPER-2487	[~paul.robinson] Seems my PR failed to build... can you tell me what's wrong with those URLs I added? https://github.com/jboss-developer/www.jboss.org/pull/1023 Or is the problem this? {code} fatal: could not read Username for 'https://github.com': No such device or address {code} -- https://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/www.jboss.org-pull-player-executor/2331/console	1	1	https://issues.jboss.org/browse/DEVELOPER-2487
7055	DEVELOPER-2487	I've submitted a new version of PR 1023 with the links to Beta1 (not Beta2) docs updated so we can at least see what that looks like on staging. Of course that might not be strictly necessary if we just use http://www.jboss.org/products/devstudio/resources/ with its "See More Documentation" link to https://access.redhat.com/documentation/en/red-hat-jboss-developer-studio/ ... which shows by default the Beta1 docs (later this week, the Beta2 docs). [~mhusnain] What do you think? Do we want to call out the Beta docs on http://www.jboss.org/products/devstudio/resources/ or just "hide" them under the "See More" link? Here's what the resources page would look like if my PR above is applied to change all the links to 9.1.Beta1 versions instead of 9.0.0.GA versions: http://www-stg.jboss.org/pr/1023/build/2333/products/devstudio/resources/	1	1	https://issues.jboss.org/browse/DEVELOPER-2487
7056	DEVELOPER-2487	Looks like with the latest import of 9.1.0 Beta(space)2 we lose the filesize of the EAP bundled installer and the "2" of Beta2 wraps to the next line. Can we fix those? http://developers.redhat.com/products/devstudio/download/ !alignment-and-filesize-problems.png! https://www.jboss.org/products/devstudio/download/ !alignment-and-filesize-problems2.png!	1	1	https://issues.jboss.org/browse/DEVELOPER-2487
7057	DEVELOPER-2480	[~james.cobb] can you review: http://it-developers-pr.stage.redhat.com/pr/715/build/678/products/. It looks done to me. I believe the mockup is based on an old wireframe where the primary CTA was for 'Downloads'. Correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-2480
7058	DEVELOPER-2473	[~mikgue] I'm also confused by this. Can you explain what you want doing with this adoption page? We could convert it to an article maybe?	1	1	https://issues.jboss.org/browse/DEVELOPER-2473
7059	DEVELOPER-2472	[~tnolan], [~crobson], can you advise on where this content should go?	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7060	DEVELOPER-2472	[~paul.robinson] This git hub source is in the TTHW format, so is it not the RHEL product "get started" page?	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7061	DEVELOPER-2472	[~tnolan] I agree, it should go as another square (in the first position, before 'CPP') on http://developers.redhat.com/products/rhel/get-started. What text and image should be on the button?	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7062	DEVELOPER-2472	Can we get this one on staging ASAP? The other 2 (in separate JIRAs) are a lower priority but we'd like to have this one to test the user flow. Also, the title of it should be: *Install Red Hat Enterprise Linux for Development*. We're not going to use quick start or quickstart for this since we have other content that use those terms.	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7063	DEVELOPER-2472	Hi [~rterzi], I noticed that sections for the updated https://github.com/rctrh/rh-tthw/blob/master/install-rhel7.adoc still have FIX ME lines. Is that content still in progress or what should I change those to? cc: [~mikgue]	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7064	DEVELOPER-2472	Yes, it's still a WIP, but we need to see it in template. * Note: Before registration can proceed, you will be required to agree to the terms and conditions at www.redhat.com/wapps/ugc/ if you haven't already done so. [*FIXME* hopefully this can be deleted.] -- *Keep this for now * * Step 4 Content - [*FIXME* - Where do the tabs come from?] This should be replaced by the tabs you are integrating. * FAQ question 4: {noformat} link:http://developers.redhat.com/products/developertoolset/get-started-rhel7-cpp/[Get started developing with C++ and Eclipse from the Red Hat Developer Toolset] {noformat} * FAQ question 5 - [FIXME Note: I didn't give this much testing], that can be deleted.	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7065	DEVELOPER-2472	Note: I will supply some additional VirtualBox FAQs that need to be added to the troubleshooting and FAQ section. [~adelaarreola] [~paul.robinson] - Do you have any suggestions on better formatting/markup for the FAQs? In particularly to differentiate the questions from the answers?	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7066	DEVELOPER-2472	[~rterzi] Thanks for the clarification!... Also as I was adding the content for https://github.com/rctrh/rh-tthw/blob/master/install-rhel7-vbox.adoc I noticed that you have multiple broken links for images. I looked in your images folders and could not find some of the missing images (ubuntu-vbox-06-settings1.png and ubuntu-vbox-04-settings4.png). Can you email them to me so I can add them to the site? Thanks.	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7067	DEVELOPER-2472	Was there an updated build last night where I can see how it all looks?	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7068	DEVELOPER-2472	[~adelaarreola] Ok, that looks good. I have a bunch of edits that I need to make across all the sections of the RHEL getting started guide. What's the easiest way to do this? Is PR #922 the most recent version?: [https://github.com/adelasofia/developers.redhat.com/blob/DEVELOPER-2472-b/products/rhel/get-started.adoc] I could make a copy of that in my repo, and then commit edits against it which would like you pick up diffs or the whole file.	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7069	DEVELOPER-2472	[~adelaarreola] - when do you need this to make it into Stage tonight? cc [~rterzi]	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7070	DEVELOPER-2472	[~adelaarreola] - I need to make some edits to the rhel-install guide. Did you make any changes? Which repo/branch has the latest version?	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7071	DEVELOPER-2472	Looks good. Can it be promoted to staging for testers?	1	1	https://issues.jboss.org/browse/DEVELOPER-2472
7072	DEVELOPER-2471	[~mikgue] what do you want the label to be under each item on http://developers.redhat.com/products/rhel/get-started/? Also, what image should we use?	1	1	https://issues.jboss.org/browse/DEVELOPER-2471
7073	DEVELOPER-2471	[~tnolan] [~james.cobb] I'm not sure where or how to place these, is there a wireframe or mockup I can look at?	1	1	https://issues.jboss.org/browse/DEVELOPER-2471

#	Requirement	Questions	Answered	Successfully Answered	URL
7074	DEVELOPER-2471	They need to link from: http://developers.redhat.com/products/rhel/get-started/ I suggest these labels: * Containerized Python * Containerized PHP * Containerized Node.js * Containerized Ruby [~mikgue] do you approve these labels? If not, can you approve a different set? [~mikgue] can you also confirm what you want the images to look like? You suggested the technology's logo, somehow incorporated with a shipping container. [~james.cobb] is that achievable this week?	1	1	https://issues.jboss.org/browse/DEVELOPER-2471
7075	DEVELOPER-2471	Are these using the regular format for getting started or the new Article format? regular getting started format: http://developers.redhat.com/products/rhel/get-started-rhel7-perl/ If we are combining two different types of page layouts here, should we make some sort of visual indicator? [~tnolan] ^^ Thoughts? [~paul.robinson] We can get the four images this week.	1	1	https://issues.jboss.org/browse/DEVELOPER-2471
7076	DEVELOPER-2471	I'm not sure. Can we have a quick meeting to discuss when [~mikgue] is available?	1	1	https://issues.jboss.org/browse/DEVELOPER-2471
7077	DEVELOPER-2471	Sorry, I'm confused. Are these 4 new Getting Started pages being linked from the CDK page or on the RHEL page?	1	1	https://issues.jboss.org/browse/DEVELOPER-2471
7078	DEVELOPER-2471	Where do you mean by "from here"?	1	1	https://issues.jboss.org/browse/DEVELOPER-2471
7079	DEVELOPER-2471	[~rterzi] Rob, can you review these? They look good to me.	1	1	https://issues.jboss.org/browse/DEVELOPER-2471
7080	DEVELOPER-2471	These look good to me. What links to these? Will the new Step 4 in RHEL be the only place?	0	0	https://issues.jboss.org/browse/DEVELOPER-2471
7081	DEVELOPER-2470	[~wesbos] how can we review this?	0	0	https://issues.jboss.org/browse/DEVELOPER-2470
7082	DEVELOPER-2468	OK. Can you ping the reviewer please. is it [~james.cobb]?	0	0	https://issues.jboss.org/browse/DEVELOPER-2468
7083	DEVELOPER-2457	[~crobson] [~ray.ploski] [~ecurley] [~mikgue] [~james.cobb] [~paul.robinson] [~burrsutter] Here's a first draft of the updated center content section for the homepage: https://redhatdevelopers.mybalsamiq.com/projects/hard-launch/Home Some method to my madness... *New CTAs* When I was in Raleigh a few weeks ago, Ellie and Laura from RH.com talked with us one day and made an comment that really struck home with me. I've heard this from a few other people as well. When a user comes to our homepage, how do they know who we are and what we offer? Yes, we have the new about page which will answer this is more detail, but there isn't any static, consistent content so that users know they are in the "right" place. I think it's great that we have a lot of news, events, etc highlighted, but there is nothing there to anchor the user and give them that initial context. I made an attempt to help solve this by creating a place for 3 primary Call-to-Action cards. I am open to other ideas here too and just wanted to get the conversation going. I am ok with not including this section for Hard Launch. *New header and intro text for the content section* I added in the content Emily provided a while ago. However, I feel like this could also use a header to that entire section to help separate it and again to give the user some context. *Suggested* I gave this section more real estate, as well as some room for a description and thumbnail. I think these are more important than the Latest items, especially since this is a curated list, and thus gave them more space and details. *Latest* I felt like this could appear more as a news feed, with just an icon, title and date. It also seemed to be less important as the suggested content, so we can show more of them but with less detail. One final thought when looking at the homepage is that the Join RHD is way down the page. I can tell you from the heatmaps so far, that people are not making it down there. If we had this as a much smaller banner (less height) we could potentially move this right under the 5 featured items. GitHub developer does a great job with this: https://developer.github.com/	1	1	https://issues.jboss.org/browse/DEVELOPER-2457
7084	DEVELOPER-2457	I think we're straying away from the original intent of the site. We've gone away from trying to get folks to the information they need quickly (remember our commitment to TTHW?) and have decided we're going to market to the developers... or more accurately spoon feed stuff WE think is important. I look at our successful competitors and they aren't overloading their home page of content, they are directing traffic to areas of interest by their users. https://developer.github.com/ https://developers.google.com/ https://developer.salesforce.com/ https://developer.apple.com/ https://developers.facebook.com/ I keep thinking back to my discussions with developers at JavaOne. They HATE being marketed to and I'm afraid we may be slipping down that slope further.	1	1	https://issues.jboss.org/browse/DEVELOPER-2457
7085	DEVELOPER-2445	Do you need any more details? I just added a description.	1	1	https://issues.jboss.org/browse/DEVELOPER-2445
7086	DEVELOPER-2442	[~paul.robinson] What is the best URL to use for the most recent version of search?	1	1	https://issues.jboss.org/browse/DEVELOPER-2442
7087	DEVELOPER-2426	[~adelaarreola] is this still blocked? If so, do you know why?	1	1	https://issues.jboss.org/browse/DEVELOPER-2426
7088	DEVELOPER-2420	[~adelaarreola] is this still blocked? If so, do you know why?	1	1	https://issues.jboss.org/browse/DEVELOPER-2420
7089	DEVELOPER-2414	[~adelaarreola] is this still blocked? If so, do you know why?	1	1	https://issues.jboss.org/browse/DEVELOPER-2414
7090	DEVELOPER-2405	Are these approved or in review? Can you update the issue status accordingly please.	1	1	https://issues.jboss.org/browse/DEVELOPER-2405
7091	DEVELOPER-2403	[~adelaarreola] looks like this is unblocked now. Correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-2403
7092	DEVELOPER-2393	Thanks for getting this sorted! Where abouts can we review it? Can you add this to Wes' branch for the HardLaunch work?	1	1	https://issues.jboss.org/browse/DEVELOPER-2393
7093	DEVELOPER-2392	I just noticed that the title word/s "About" or "About us" are not in the hero. Are we still planning to include those or are they being added in implementation instead?	0	0	https://issues.jboss.org/browse/DEVELOPER-2392
7094	DEVELOPER-2391	Is this close enough for [~james.cobb] and team to also complete VD by the 29th?	1	1	https://issues.jboss.org/browse/DEVELOPER-2391
7095	DEVELOPER-2391	[~tnolan] was this approved? If so, can you mark as done. FYI, there' another issue for the copy, so this just covers the wireframe.	1	1	https://issues.jboss.org/browse/DEVELOPER-2391
7096	DEVELOPER-2385	[~adelaarreola] is this still blocked? If so, do you know why?	1	1	https://issues.jboss.org/browse/DEVELOPER-2385
7097	DEVELOPER-2379	[~adelaarreola] is this still blocked? If so, do you know why?	1	1	https://issues.jboss.org/browse/DEVELOPER-2379
7098	DEVELOPER-2371	[~adelaarreola] - This is marked as blocked... what do we need to do to move it forward?	1	1	https://issues.jboss.org/browse/DEVELOPER-2371
7099	DEVELOPER-2371	[~adelaarreola] Sorry... I thought Wes was working on this and tagged him instead. The are ready... https://issues.jboss.org/browse/DEVELOPER-2511 Specifically... The new home page promos are ready for implementation. What is RHD? http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_whatisrhd.png DevNation : Registration open http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_devnation_registrationopen.png Dive into Java http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_diveintojava.png EAP7 beta is here http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_eap7betaishere.png New Containers Course http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/makerfile/images/RHDev_homepromo_containercourse.png Content and links can be found here: https://issues.jboss.org/browse/DEVELOPER-2510	1	1	https://issues.jboss.org/browse/DEVELOPER-2371

#	Requirement	Questions	Answered	Successfully Answered	URL
7100	DEVELOPER-2371	[~ecurley] where can I find the 'titles' (or text that goes under the image) for each of the spotlights?	1	1	https://issues.jboss.org/browse/DEVELOPER-2371
7101	DEVELOPER-2370	*Main CTA - Book: Modern Java EE Design Patterns. Download now. *Side CTA: What is RHD all about? 20-min live demo on containers - register now Developers Salary Survey (delete for launch if the quality isn't what we expect) Become content contributor Complementary containers course <<this new from RH Training>> *Video: Getting started with JBoss EAP 7 Beta http://developerblog.redhat.com/2016/01/20/screen-cast-getting-started-with-jboss-eap-7-beta/	1	1	https://issues.jboss.org/browse/DEVELOPER-2370
7102	DEVELOPER-2369	[~james.cobb] *Main CTA - Book: Modern Java EE Design Patterns. Download now. <<http://developers.redhat.com/promotions/distributed-javaee-architecture/>> *Side CTA: What is Red Hat Developers all about? <<developers.redhat.com/about>> 20-min live demo on containers - register now <<coming soon>> Become content contributor <<http://developers.redhat.com/community/contributor/>> Complementary containers course <<http://wp.me/p2WBxv-1KWJ>> *Video: Getting started with JBoss EAP 7 Beta <<http://developerblog.redhat.com/2016/01/20/screen-cast-getting-started-with-jboss-eap-7-beta/>>	0	0	https://issues.jboss.org/browse/DEVELOPER-2369
7103	DEVELOPER-2366	[~paul.robinson], I think that since the Resources page updates are a lower priority than other items for Hard Launch, we can probably wait until Build. Agree?	1	1	https://issues.jboss.org/browse/DEVELOPER-2366
7104	DEVELOPER-2359	[~lightguard], [~oshaw]. Can you review please.	0	0	https://issues.jboss.org/browse/DEVELOPER-2359
7105	DEVELOPER-2355	[~ecurley] [~ray.ploski] Can you let us know by EOD today (Mon.) or first thing tomorrow if the wireframe is approved and ready for Front-end and Visual Design? Any changes you'd like to be made?	1	1	https://issues.jboss.org/browse/DEVELOPER-2355
7106	DEVELOPER-2355	Great! Thanks [~ecurley]! I'm going to mark this as "Completed", then. [~paul.robinson] - Will this trigger emails to people that own sub-tasks then? Or do I still need to notify people that the wireframe is done and work can begin on Front-end and Visuals?	1	1	https://issues.jboss.org/browse/DEVELOPER-2355
7107	DEVELOPER-2353	[~rayploski-launch], where should this image go now that the get-started page is no longer present?	0	0	https://issues.jboss.org/browse/DEVELOPER-2353
7108	DEVELOPER-2352	Latest is up: http://developers-pr.stage.redhat.com/pr/133/build/306/mobile/overview/	0	0	https://issues.jboss.org/browse/DEVELOPER-2352
7109	DEVELOPER-2352	Thanks, I think it looks good. Just a few minor update requests: 1. Can you please get rid of the : after the titles of the five narratives? We don't need them since the titles are bold. 2. Can you please make the black call out box the same height as the left hand overview text, so that the two are vertically aligned, rather than the same size horizontally? I've sent this link to Steve O'Keefe and Conor O'Niel for one last review.	1	1	https://issues.jboss.org/browse/DEVELOPER-2352
7110	DEVELOPER-2352	[~rebecca1021] What did you have in mind for those links?	1	1	https://issues.jboss.org/browse/DEVELOPER-2352
7111	DEVELOPER-2350	[~dhladky] is this also covered by our review earlier today? Can it be marked as resolved?	1	1	https://issues.jboss.org/browse/DEVELOPER-2350
7112	DEVELOPER-2349	[~rebecca1021] There was a lot of changes made to the solution pages in the run-up to the summit release. Can you let me know if there is anything outstanding on this issue or wether it can be resolved?	0	0	https://issues.jboss.org/browse/DEVELOPER-2349
7113	DEVELOPER-2345	[~paul.robinson] do you really need a new registered query? You can use a minor modification of your URL example to get what you need: {{http://dcp.stage.jboss.org/v2/rest/search?field=sys_title&field=sys_url_view&type=rht_website&query=cdi&size=10&query_highlight=true}} If the data were available on existing dcp the URL would be: {{http://dcp.jboss.org/v1/rest/search?field=sys_title&field=sys_url_view&type=rht_website&query=cdi&size=10&query_highlight=true}} It is possible to create a new registered query but if the above URL gives you all you need then I think you can start using it right away without need for new query. WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-2345
7114	DEVELOPER-2340	Duh. Question, [~paul.robinson]: While these are "tags" at the RHD end, are you referring to the "categories" or "tags" in WordPress? Or both? My recollections is the former.	1	1	https://issues.jboss.org/browse/DEVELOPER-2340
7115	DEVELOPER-2339	[~james.cobb] Can you review this screenshot: https://www.dropbox.com/s/ufz8ft5z9scw7px/Screenshot%202015-05-26%2009.23.34.png?dl=0 Is that the right thing to show when a user is logged in? Clicking on the user's name will take them to their profile page.	1	1	https://issues.jboss.org/browse/DEVELOPER-2339
7116	DEVELOPER-2337	[~ecurley] I believe there is a requirement for this Microsite. Can you provide the content please?	0	0	https://issues.jboss.org/browse/DEVELOPER-2337
7117	DEVELOPER-2335	[~james.cobb] I think this is a duplicate? Do you have link to the issue(s) it duplicated?	1	1	https://issues.jboss.org/browse/DEVELOPER-2335
7118	DEVELOPER-2333	[~rebecca1021] Can you review the above link please?	1	1	https://issues.jboss.org/browse/DEVELOPER-2333
7119	DEVELOPER-2333	[~rebecca1021] What did you have in mind for those links?	1	1	https://issues.jboss.org/browse/DEVELOPER-2333
7120	DEVELOPER-2312	[~adelaarreola] Can you take a look at the container solution (now merged) and: 1) Add a link to it to the bottom of the homepage. See other solution links for an example. 2) Improve the L&F on the overview and adoption pages 3) Add an illustration to the overview if you have time and one is available.	0	0	https://issues.jboss.org/browse/DEVELOPER-2312
7121	DEVELOPER-2297	Can you look into adding logic to render _partials/solution-left-menu.html.slim selectively like you did with the learn and buzz?	1	1	https://issues.jboss.org/browse/DEVELOPER-2297
7122	DEVELOPER-2289	Okay, just one quick change. Can you please change the slash (/) in the first sentence of the intro to a comma (,)? Thanks	1	1	https://issues.jboss.org/browse/DEVELOPER-2289
7123	DEVELOPER-2287	[~adelaarreola] how was it incorrect? Intentionally it's almost the same as JBDS.	1	1	https://issues.jboss.org/browse/DEVELOPER-2287
7124	DEVELOPER-2282	I created the necessary graphics for Fedora and OpenShift Commons (both regular sized and featured size). I put them on the design site. http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_1_0/upstream/images/fedora_500x400.png http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_1_0/upstream/images/fedora_200x150.png http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_1_0/upstream/images/openshiftcommons_500x400.png http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_1_0/upstream/images/openshiftcommons_200x150.png [~paul.robinson] Paul, should we host them on static.jboss.org/foo? I tried to post them there but I don't have the access privileges to create new directories at the root level. Should they be posted there or do we need to have an equivalent for RHD (static.developers.redhat.com)?	1	1	https://issues.jboss.org/browse/DEVELOPER-2282
7125	DEVELOPER-2280	This has been merged into master - can close.	0	0	https://issues.jboss.org/browse/DEVELOPER-2280
7126	DEVELOPER-2278	What are we to do about the mobile view of this banner? I'll create a mobile size image based on the sizes we use for jboss developer.	0	0	https://issues.jboss.org/browse/DEVELOPER-2278
7127	DEVELOPER-2275	[~dhladky] the chrome on http://developers.redhat.com/ is current. Can you update the download-manager please?	1	1	https://issues.jboss.org/browse/DEVELOPER-2275
7128	DEVELOPER-2275	Hi [~james.cobb], I was not able to find the place you mentioned in the html. Also I should not be the one modifying the CSS. I would be a little worried, that I screw up other development. Will you, please, ping me when you are on-line so we can have a look on it?	0	0	https://issues.jboss.org/browse/DEVELOPER-2275
7129	DEVELOPER-2273	[~crobson], [~kpassmore]: Do you know what categories we decided to use here?	1	1	https://issues.jboss.org/browse/DEVELOPER-2273
7130	DEVELOPER-2273	Thanks, Are we copying these categories from somewhere? For example, how do we know which product to put in which category?	0	0	https://issues.jboss.org/browse/DEVELOPER-2273
7131	DEVELOPER-2267	By mobile do we also mean tablets (Ipad, Chromebooks)? Or do we test them as desktop? Should we have a separate issue for tablets?	0	0	https://issues.jboss.org/browse/DEVELOPER-2267

#	Requirement	Questions	Answered	Successfully Answered	URL
7132	DEVELOPER-2267	Looks good! Can we have it trigger at tablet size (~975px) as the current header starts collapsing ungracefully at around that size and the mobile nav doesn't switch over until around 475px. Are there any changes to the html that we'll have to make for the login screens or are all the changes just in our CSS files?	1	1	https://issues.jboss.org/browse/DEVELOPER-2267
7133	DEVELOPER-2267	What is the widht for mobile device? Is it going to be same as "small" column widht? I'm asking because zurb foundation comes with small, medium and large column sizes and PR build has different min-widht compared to min-widht for large columns. I would say that new compacted header should be only on width where small column size is active. Header on medium size should be similar to large but slightly compacted by less padding/margins.	1	1	https://issues.jboss.org/browse/DEVELOPER-2267
7134	DEVELOPER-2267	Or are you pushing the change into production right now ?	1	1	https://issues.jboss.org/browse/DEVELOPER-2267
7135	DEVELOPER-2267	[~lkrzyzanek] Okay- will send another PR fixing this	1	1	https://issues.jboss.org/browse/DEVELOPER-2267
7136	DEVELOPER-2267	Thanks Libor... can you share this with David so it gets pushed to the T&C's page as well?	1	0	https://issues.jboss.org/browse/DEVELOPER-2267
7137	DEVELOPER-2264	OK, can you investigate the feasibility? Then we can decide on wether it's worth implementing now or later.	1	1	https://issues.jboss.org/browse/DEVELOPER-2264
7138	DEVELOPER-2252	Was it ever discussed how we would gather the email addresses of the people that submit this form? Given that we show this when everything is screwed, what infrastructure can we use?	1	1	https://issues.jboss.org/browse/DEVELOPER-2252
7139	DEVELOPER-2252	I got a note from Ray the other day asking if I could support this with an Eloqua form, the answer is yes. He said this was not Lead gen but he did want the email addresses to be stored in Eloqua. Because is not Lead gen we don't need to worry about picking up Campaign tracking codes out of the cookies and passing them in the form submit nor an Offer ID. This means that submitting this form will not result in a Lead be created in SalesForce nor will it result in a Campaign Membership record in SalesForce. But this does beg the question, what do you want to do with the people who submit this form? Usually, in situations like, this we just have 3 processing steps in the form in Eloqua: 1- Save the email address to the database 2- Add the email address to a Shared List so you can email all of the people, if needed 3- Send an email to someone @redhat.com telling them about the form submission Its seems like we are going to need #1 and #2, do you want #3? If so, what email address? Thanks, Greg	1	1	https://issues.jboss.org/browse/DEVELOPER-2252
7140	DEVELOPER-2252	[~gregmanetti], I suppose this choice doesn't affect the way this form will be integrated with Eloqua. We're very short on time here so I'd like to integrate it ASAP, while we wait for an answer from Ray, Paul or Mark if they want #3. My guess however would be that only #1 and #2 will be needed. Either way, could you tell me how to integrate this ?	0	0	https://issues.jboss.org/browse/DEVELOPER-2252
7141	DEVELOPER-2252	[~james.cobb] is this still needed? Is it out of date, or maybe already resolved?	1	1	https://issues.jboss.org/browse/DEVELOPER-2252
7142	DEVELOPER-2237	Should this be in the Review board? Note that this is for indexing non-middleware developer materials, which i don't yet think we have a plan for.	1	1	https://issues.jboss.org/browse/DEVELOPER-2237
7143	DEVELOPER-2235	Hi [~crobson], This page has been marked as approved on both the deliverables spreadsheet and in the wireframe itself. Is it ready to be handed over to front-end? Thanks, Adela	1	1	https://issues.jboss.org/browse/DEVELOPER-2235
7144	DEVELOPER-2230	A ticket has been filed to web@redhat.com (SN Incident INC0270155) for the initial redirect of the top pages. We also need to redirect from developer.redhat.com. There are MANY subpages that will require redirects - that will take time to sort through and I suggest post-launch. Are there other redirects you're thinking about?	1	1	https://issues.jboss.org/browse/DEVELOPER-2230
7145	DEVELOPER-2217	How should this work? Is this really back-end part of the website?	0	0	https://issues.jboss.org/browse/DEVELOPER-2217
7146	DEVELOPER-2217	This is usually a checkbox during registration. I thought this could be added to User Profile in KC (Not implemented yet - KEYCLOAK-311). [~pmuir] What is requirement from UX perspective? How (un)subscribe should work?	1	1	https://issues.jboss.org/browse/DEVELOPER-2217
7147	DEVELOPER-2217	In that case we should remove both checkboxes on account page. Subscribe happens automatically and unsubscribe happens on link in mail. Is it correct ?	1	1	https://issues.jboss.org/browse/DEVELOPER-2217
7148	DEVELOPER-2217	If you'd like to keep those checkboxes what is the API that we should call when user will "doing" subscribe and unsubscribe by clicking checkboxes? If there is no such API I would prefer to replace those checkboxes by text explaining how it works something like "You're automatically subscribed to newsletter. You can unsubscribe by clicking on link in newsletter"	1	1	https://issues.jboss.org/browse/DEVELOPER-2217
7149	DEVELOPER-2217	Can we find all email addresses for people that have the "Subscribe to Newsletter" ticked in the profile?	1	1	https://issues.jboss.org/browse/DEVELOPER-2217
7150	DEVELOPER-2217	Also, can we ensure the two tick-boxes default to ticked? In my account they are both un-ticked. I don't think i changed them.	1	1	https://issues.jboss.org/browse/DEVELOPER-2217
7151	DEVELOPER-2217	[~paul.robinson] Answering to your question "Can we find all email addresses for people that have the "Subscribe to Newsletter" ticked in the profile?" Yes it's possible. I can query Keycloak DB to get such list. But you mentioned earlier including information who signed Terms & Conditions. If you'd like to get list of e-mails who have "subscribe to newsletter" and have singed T&C then it's *not* technically possible to do that.	1	1	https://issues.jboss.org/browse/DEVELOPER-2217
7152	DEVELOPER-2207	[~dhladky] from talking to you yesterday, it seems like setting this up is in-progress. Would it be more accurate to move this to the 'back-end' board and mark as 'in progress'? If so, can you make the change please.	1	1	https://issues.jboss.org/browse/DEVELOPER-2207
7153	DEVELOPER-2194	[~rebecca1021], Ray was saying you've done this. Can you provide us a link to the copy so that we can put it on the site?	1	1	https://issues.jboss.org/browse/DEVELOPER-2194
7154	DEVELOPER-2193	Hi Paul, I just wanted to ask, what is the status of this ticket and I found out it was automatically assigned to me. Which is not the proper person to solve it. I am not sure, if it should be handled by you or by the Visual team. Whoever should be responsible for this page, it would be useful for me, if such page existed. I do not need its final version, but I would like to add it to DM staging configuration, so we can see how it works.	1	1	https://issues.jboss.org/browse/DEVELOPER-2193
7155	DEVELOPER-2193	The thank you pages are created. Here's the one for EAP: http://rhd.jboss.org/products/eap/get-started/#!project=eap We just need to have the download manager redirect to here for Red Hat Developer downloads. Do you have docs that tell me how to do this?	1	1	https://issues.jboss.org/browse/DEVELOPER-2193

#	Requirement	Questions	Answered	Successfully Answered	URL
7156	DEVELOPER-2193	Hi Paul, I have not updated the documentation yet, but Download manager behaves the same way as in the case of JBoss Developer thank you pages. You as the administrator only need to fill the additional field in the product page of the administration, and the user needs to request the file using RHD domain. One example from my local environment. I entered [http://rhd.jboss.org/products/eap/get-started?] to the product dialog of Enterprise Application Server product. After some debugging I found out you add the part after hash and if I do that myself, your page does not work properly. So this is it for admin. As the user I requested [http://developer-stg.redhat.com:8080/download-manager/file/jboss-eap-6.4.0.Beta.zip] (developer-stg.redhat.com is my testing domain for RHD requests). Download manager investigated the requested file and found out, that this file is to be served with "No terms and conditions" model, which results in following redirect: [http://rhd.jboss.org/products/eap/get-started?tcUser=David+Hladky&tcDownloadFileName=jboss-eap-6.4.0.Beta.zip&tcRedirect=5000&tcSrcLink=http%3A%2F%2Flocalhost%3A8080%2Fdownload-manager%2Fcontent%2Forigin%2Ffiles%2Fsha256%2Ff0%2Ff019aeac85191bb6bb4aefceb5856b37744ae57350b25738dd20444e39a6faea%2Fjboss-eap-6.4.0.Beta.zip&p=Enterprise+Application+Server&pv=6.4.0+Beta&tcDownloadURL=https%3A%2F%2Faccess.cdn.redhat.com%2Fcontent%2Forigin%2Ffiles%2Fsha256%2Ff0%2Ff019aeac85191bb6bb4aefceb5856b37744ae57350b25738dd20444e39a6faea%2Fjboss-eap-6.4.0.Beta.zip%3F_auth_%3D1431554550_50d2215d6c67e06065e0c3c34c2676d2] Your web page is working as well, because the download started for me. But let us investigate the link a bit further: - I realized, that I incorrectly encoded parameters using URI encoding in the past, that is not correct for query parameters. So I fixed it - I noticed you are adding #!project=eap to the end of the requests. This is not correct, because you do this even for links, that are not EAP (tested on Data Virtualization). That is why I added two new parameters p=product name and pv=product version name. I suppose those two can actually help you solving this issue and eventually help wit the customization of the page as well. Instead of product name, I could give you product code, if it would be better for you (the same code Eloqua uses) - for RHD-64 we created a feature called Featured artifact - see RHD-151. I suppose this link might be useful here as well? Now when it is already implemented it would be trivial to add it here. WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-2193
7157	DEVELOPER-2192	[~paul.robinson] for the pages with the old layout, should the developer materials at the bottom also still show?	1	1	https://issues.jboss.org/browse/DEVELOPER-2192
7158	DEVELOPER-2171	[~ray.ploski] I think you said this was done. Is that right? You need to move it through the workflow to mark sit as done. "Handover to review" -> "start work" -> "Done"	0	0	https://issues.jboss.org/browse/DEVELOPER-2171
7159	DEVELOPER-2165	Wireframed approved. Should use same template as learn page for Products.	0	0	https://issues.jboss.org/browse/DEVELOPER-2165
7160	DEVELOPER-2159	Is this ready to go then? Are we using the svg or the png version? Anywhere specific you want this on the overview page?	0	0	https://issues.jboss.org/browse/DEVELOPER-2159
7161	DEVELOPER-2036	[~ecurley], [~newtonm]: is this now done?	1	1	https://issues.jboss.org/browse/DEVELOPER-2036
7162	DEVELOPER-2036	[~newtonm] did you add this to GTM? If so, who's the best person to debug this? Axel?	0	0	https://issues.jboss.org/browse/DEVELOPER-2036
7163	DEVELOPER-2031	Ok [~lvlcek]. Are you waiting on anything to start the indexing? I appreciate that we need to tidy up the sitemap, but can that come slightly later?	1	1	https://issues.jboss.org/browse/DEVELOPER-2031
7164	DEVELOPER-2031	[~paul.robinson] I created a new indexer for RHD website and installed it on staging (http://dcp.stage.jboss.org/). It indexed 1192 pages for the first (full-reindex) run. Do you think you can have a look at the data and tell me if it looks ok to you? Ping me on IRC if you need mot help with this.	1	1	https://issues.jboss.org/browse/DEVELOPER-2031
7165	DEVELOPER-2031	[~lvlcek] could yo please clarify what river you have installed in dcp.stage ? Could you please push it to dcp2 as [~paul.robinson] needs it in production (since whole website in production already uses dcp2). Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-2031
7166	DEVELOPER-2030	[~brock] are these graphics approved? [~c.wong]: where should this link to? I'm assuming http://developers.redhat.com/products/eap/download/.	1	1	https://issues.jboss.org/browse/DEVELOPER-2030
7167	DEVELOPER-2030	[~paul.robinson]Not approved as of yet. [~c.wong] and [~mikgue], could you weigh in? Let me know if you'd like something different.	1	1	https://issues.jboss.org/browse/DEVELOPER-2030
7168	DEVELOPER-2030	This one looks like it's missing a space between "JBoss" and "Enterprise". Or is it just me? https://issues.jboss.org/secure/attachment/12396642/hero_desktop_1175x400_eap_01112016.png BTW, there's a "Wildfly 9 Now Available" banner on JBoss.org. Will this replace the Wildfly banner? Probably makes most sense... Could we also place a banner on the Wildfly site? Or is that something I need to bring up directly with Wildfly team members? We might want to tweak the wording for the Wildfly site. Something like, "Like Wildfly? Try Red Hat JBoss Enterprise Application Platform for your enterprise Java applications." Maybe I should open up a separate JIRA for that? Thanks, Christina	1	1	https://issues.jboss.org/browse/DEVELOPER-2030
7169	DEVELOPER-2030	Looks good Adela. Will this be the same one for RHDev?	1	1	https://issues.jboss.org/browse/DEVELOPER-2030
7170	DEVELOPER-2029	Hi Jason - Can we see something with a guy in it? I know odd request but 90%+ of our contributors will likely be male and I'd like something they will more identify with. Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-2029
7171	DEVELOPER-2029	Thanks Jason. Can we ditch the tie look? These contributors will be more developers contributing content. Audience this will target: Enterprise developers - some with presentation skills but more computer focused skills. They love a good t-shirt at a conference and likely will grab stickers. They will be primarily writing blogs, maybe creating a few videos, and submitting screen captured demos.	1	1	https://issues.jboss.org/browse/DEVELOPER-2029
7172	DEVELOPER-2029	[~ecurley][~paul.robinson], Should we assign this over to the front-end dev team to put it up?	1	1	https://issues.jboss.org/browse/DEVELOPER-2029
7173	DEVELOPER-2029	What text do you want under the image, and what spotlight should it replace?	1	1	https://issues.jboss.org/browse/DEVELOPER-2029
7174	DEVELOPER-2029	[~ecurley] : Can you supply the content requested by Paul?	0	0	https://issues.jboss.org/browse/DEVELOPER-2029
7175	DEVELOPER-2027	Perhaps I missed the meeting on Mega Menus but when was it decided we were to have a mega menu? It was my understanding that we were using simple flyouts. Current flyout design: http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/pages/images/RHD_home_final_flyout.png I have real problems with them for our site. 1) The "featured content" on the right comes off as marketing fluff which I believe our developer audience would likely reject. I personally feel navigation needs to be about way finding only. 2) The current mockups in Balsalmiq don't display any overview page for the sections. This would mean the user goes directly to subpages or we have to include a link to the overview pages. RedHat.com does this by adding a button for the main sectional pages. I find this VERY frustrating as it adds a click between me and the content. I think we'll just be creating barriers. 3) This sort of mega menu is a nightmare on Mobile. This Redhat.com site is AWFUL on mobile as it injects an additional hamburger icon for these navigation pages. It's a nightmare.	1	1	https://issues.jboss.org/browse/DEVELOPER-2027

#	Requirement	Questions	Answered	Successfully Answered	URL
7176	DEVELOPER-2023	What do you mean by 'landing page'? Is this an article, microsite, product page, topic, or something else?	1	1	https://issues.jboss.org/browse/DEVELOPER-2023
7177	DEVELOPER-2023	I'm adjusting the scope of this issue as we are not sure yet what the output of this issue will be. Some things that would help to know: * What is the business objective for this section of the site? * What are the driving factors for the timing? * Do we have an existing page that is similar to what you're looking to do? * How do we expect and want users to get to this content? * What do we want users to do next? Once we know/discuss this information, I can work on some user flows to support this.	1	1	https://issues.jboss.org/browse/DEVELOPER-2023
7178	DEVELOPER-2023	[~paul.robinson] I can't seem to be able to move 2 existing issues to this Epic. Can you take a look? Both https://issues.jboss.org/browse/DEVELOPER-2023 and https://issues.jboss.org/browse/DEVELOPER-2584 should be listed here.	1	1	https://issues.jboss.org/browse/DEVELOPER-2023
7179	DEVELOPER-2022	[~morgan_thomas] what's the status of this? Is it done? CC [~james.cobb]	1	1	https://issues.jboss.org/browse/DEVELOPER-2022
7180	DEVELOPER-2021	[~crobson], [~tnolan]. Do you have an ETA on this? I'm in the process of planning [~wesbos]' work.	1	1	https://issues.jboss.org/browse/DEVELOPER-2021
7181	DEVELOPER-2021	Ok, that sounds good [~paul.robinson]. [~ecurley] - Can you update the wireframe text in Balsamiq to fill out the: * sub-title text * the brief descriptions for each of the resource types Instructions, but feel free to ping me if you'd like me to walk you through it. # Go to: https://redhatdevelopers.mybalsamiq.com/projects/hard-launch/Resources%20%28more%20information%29 # Click the Edit link located near the top right # Click on the text you would like to change. It will then get a box around it. Only the text that you can edit is clickable. All other elements are locked. # When finished, click the save icon that has a cloud at the top left of the screen. # You can add a note then in the pop-up as to your changes and check the Send email notifications option if you're all done and want to let the group know.	1	1	https://issues.jboss.org/browse/DEVELOPER-2021
7182	DEVELOPER-2021	[~mikgue] - Can you assist with the content? - see above.	1	1	https://issues.jboss.org/browse/DEVELOPER-2021
7183	DEVELOPER-2021	I have to ask this question. How is this page different than the search page? If we present them with a full list of the content, then ask the user to filter down until they find it useful, it sounds a lot like "search" to me.	1	1	https://issues.jboss.org/browse/DEVELOPER-2021
7184	DEVELOPER-2021	[~tnolan] - I get where you are coming from on the browsing nature of this page being how folks will intend to use it but don't agree with the current presentation. When a person gets to this page, they will HAVE to filter it to get ANY useful information if we don't luckily strike on some content of interest with the default. If they are clicking and filtering anyway, why don't we give them categories to select from (can't by topic as that's what the "Topics" section should cover). In the current design we show this sort of prefiltering at the top but it looks like it's in a hidden pane. To me, the content in the hidden pane should be the primary way for folks to browse and therefore given primary status on the page. If a user clicks on "Videos" then they land on a page of default videos that has the keyword and other filter options.	1	1	https://issues.jboss.org/browse/DEVELOPER-2021
7185	DEVELOPER-2021	[~tnolan] I reviewed what RH.com did with their resources and I find it much more logical and intuitive. I say we lift that pattern. I strongly disagree with pushing this if we aren't 100%. This would mean front end would have to implement this page TWICE. We have seen a pattern of "release this now and fix it later" that never gets re-addressed because the teams are too busy to circle back. Can we discuss this at the Wednesday meeting?	1	1	https://issues.jboss.org/browse/DEVELOPER-2021
7186	DEVELOPER-2020	Do we have videos for RHEL that we can add to the site?	1	1	https://issues.jboss.org/browse/DEVELOPER-2020
7187	DEVELOPER-2020	[~paul.robinson] Paul, I think I'll be needing privs to the youtube account (at least it appears, per the instructions, that I should not index videos via the group account. Also, can we use videos from ANY YouTube playlist? Add to Octopus launch please.	1	1	https://issues.jboss.org/browse/DEVELOPER-2020
7188	DEVELOPER-2019	[~ecurley] [~ray.ploski] [~mikgue] - Do we have content suggestions for these three new promos?	1	1	https://issues.jboss.org/browse/DEVELOPER-2019
7189	DEVELOPER-2019	#1 - This will be on the top 12 articles: http://developerblog.redhat.com/2015/12/14/12-top-blog-articles-of-2015/ #2 How about learn Node.js? Not sure we have enough articles to assemble for this. #3 Should people be directed to the DN site? Or could we get away with having them register first? (I'm thinking no, but asking anyway)	1	1	https://issues.jboss.org/browse/DEVELOPER-2019
7190	DEVELOPER-2019	1) Devnation, we need the dates and location for the event. We'll address this again once we work out the revised branding. 2) Node.JS... Do we want to the copy to be more directed to the vacationing developers? "Have some downtime? Learn Node.js" 3) on the 12 top posts graphic... use he darker red verses the grey for the bars (ribbon simulation).	0	0	https://issues.jboss.org/browse/DEVELOPER-2019
7191	DEVELOPER-2019	I'd like to use the gray ribbon for the blog article and will schedule it for next week. If someone needs to manually post the banner, then let's wait until the first week of January. Can I download the pic as is, or does anyone need to do anything to it first?	0	0	https://issues.jboss.org/browse/DEVELOPER-2019
7192	DEVELOPER-2019	Do we have links. [~newtonm] - Do you need a separate JIRA for this to push them live?	1	1	https://issues.jboss.org/browse/DEVELOPER-2019
7193	DEVELOPER-2018	[~brock] are these approved or in review? Can you update the issue status accordingly please.	1	1	https://issues.jboss.org/browse/DEVELOPER-2018
7194	DEVELOPER-2018	[~adelaarreola] is the CDK page just going to be another product detail page? Is there going to be more than one CDK page?	1	0	https://issues.jboss.org/browse/DEVELOPER-2018
7195	DEVELOPER-2018	[~adelaarreola] No problem! Would you like me to append the header images to this Jira or wait for one to come down the line for the CDK page?	0	0	https://issues.jboss.org/browse/DEVELOPER-2018
7196	DEVELOPER-2018	Will the CDK be in the hard launch?	0	0	https://issues.jboss.org/browse/DEVELOPER-2018
7197	DEVELOPER-2008	[~crobson], [~james.cobb]. What do you think about this? This idea came from feedback by Mark Little and Julien Ponge.	1	1	https://issues.jboss.org/browse/DEVELOPER-2008
7198	DEVELOPER-2003	[~james.cobb] where can I find the assets for this page? i.e. the dotted graphic and the 'Join the red hat developer...' graphic? Thanks!	1	0	https://issues.jboss.org/browse/DEVELOPER-2003
7199	DEVELOPER-1997	So the first 2, I understand and like. For the 3rd, it feels like the "X" is telling me that videos will be hidden. Is the "X" displayed quickly so it would not be interpreted this way?	1	1	https://issues.jboss.org/browse/DEVELOPER-1997
7200	DEVELOPER-1997	This is looking great! I just have a few things though when comparing the design to the wireframe. *Initial view* * Code icon - I'm not familiar with the one being used here. Will all our developers understand this one (including .netters)? * Get Started icon - What other options do we have here? A lightbulb conveys more of an "idea" to me. * The sorting elements should be under the horizontal line. This allows for the # of results to be displayed above the line and so it matches the positioning on the search results page. * The wireframe had a 1 px border box around each of the individual results to help scanability. I'm open to other ideas than a border, but would like to see some clearer separation here. * Description text for each result should only show 3 lines of text, the VD show 4. 3 lines allows for plenty of room to get an idea of what the resource is about, and reduces the amount of space used for each result. It would also match up better, space-wise with the ones that have a thumbnail. *Filtered view* * Same comments as above * I'm concerned about the spacing for the Products drop-down. Are we going to have enough space if we plug in the longest titled one? * Actually, I'm not sure if we will have any data to fill in the Programming Language drop-down, so we may have more space. [~paul.robinson] - Are we able to use the Language currently? * The spacing around the Custom publish date fields seems odd to me. Do we have enough room to have these next to the date field instead of under it...like the wireframe? [~mikgue] - I believe the type filters would work just like on RH.com. You can see how they work: http://www.redhat.com/en/resources	1	1	https://issues.jboss.org/browse/DEVELOPER-1997

#	Requirement	Questions	Answered	Successfully Answered	URL
7201	DEVELOPER-1993	[~adelaarreola] was this ever done? Is it still needed?	1	1	https://issues.jboss.org/browse/DEVELOPER-1993
7202	DEVELOPER-1977	Who needs to review this? Also, is the next step to produce the actual header graphic?	1	1	https://issues.jboss.org/browse/DEVELOPER-1977
7203	DEVELOPER-1973	We added "Developer Blog" under "Community" once we noticed it was missing last week. It's on the wireframe but doesn't look like it's made it into the build yet. Can you add this in? It should be listed first. https://redhatdevelopers.mybalsamiq.com/projects/hard-launch/Navigation	1	1	https://issues.jboss.org/browse/DEVELOPER-1973
7204	DEVELOPER-1973	[~wesbos]: When the menus appear, and I click on a link (such as the for the EAP product page), nothing happens. Can you fix this as a priority, as it's making it harder for other people to review the rebrand and topics work.	1	1	https://issues.jboss.org/browse/DEVELOPER-1973
7205	DEVELOPER-1967	I think the issue here was setting page specific stuff in the global header (i.e 'chrome'). For example, the Facebook and Twitter meta-data is set by the page, and is not global across the site. This allows us to set, for example, a video page's thumbnail as the twitter card image. Have you got this working in Drupal?	1	1	https://issues.jboss.org/browse/DEVELOPER-1967
7206	DEVELOPER-1967	Is this done, or a dup?	1	1	https://issues.jboss.org/browse/DEVELOPER-1967
7207	DEVELOPER-1946	We need to decide how to handle the categories. In particular, is the current list correct, and how do we specify which project goes into which category?	1	1	https://issues.jboss.org/browse/DEVELOPER-1946
7208	DEVELOPER-1945	[~newtonm] Who's working on this? Can you update the Assignee?	0	0	https://issues.jboss.org/browse/DEVELOPER-1945
7209	DEVELOPER-1940	[~wesbos] do you need any visual design work doing for the global re-brand, or will it just be needed for the specific changes to specific pages?	1	1	https://issues.jboss.org/browse/DEVELOPER-1940
7210	DEVELOPER-1934	[~paul.robinson] and [~tnolan] I have added first draft of acceptance criteria for your review. I also have the following questions: 1. What are the Publish date drop down items? 2. What are the sort results drop down items? 3. Tags, are they clickable? What happens when you click on the tag? Should each have zero or more tags? 4. Contributor, is it clickable? where does it go? 5. What is the maximum number of characters that can be entered into the keyword search box?	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7211	DEVELOPER-1934	{quote} Scenario: All results must have a label that relates to that product. Given I am on the Resources page When I enter "Red Hat Container Development Kit" into the Keyword's box Then the results should contain the label "Red Hat Container Development Kit" {quote} What do you mean by 'label' here? Is it 'tag' that you mean? If so, I don't think it's a requirement that all results must be tagged with a related product. Not all our content has this tag.	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7212	DEVELOPER-1934	{quote} # IH: I'm not sure about this one, needs more detail from Paul or Tiffany Scenario: Videos should show thumbnail only - if they are unique. Given I am on the Resources page When I click to filter results by "Video" Then the results should contain a thumbnail {quote} The requirement here is that items of type video must have a thumbnail, others must not.	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7213	DEVELOPER-1934	[~paul.robinson] Scenario: All results must have a label that relates to that product. Given I am on the Resources page When I enter "Red Hat Container Development Kit" into the Keyword's box Then the results should contain the label "Red Hat Container Development Kit" PR Q: What do you mean by 'label' here? Is it 'tag' that you mean? If so, I don't think it's a requirement that all results must be tagged with a related product. Not all our content has this tag. IH A: This was a scenario that we discussed in our meeting?, By the Author link there is a product label in the wireframes of "Red Hat Container Development Kit"?	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7214	DEVELOPER-1934	[~tnolan], [~paul.robinson] and I have completed the scenarios for this ticket. Can you have a quick review of the AC so that [~adelaarreola] can start work on it this week.	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7215	DEVELOPER-1934	Okay just pushed the latest. A few outstanding issues: 1) Contributors endpoint is always empty - is this still working I've hidden the contributors for now. 2) Not sure how to turn off case sensitivity. Pinging [~lkrzyzanek] about this. 3) We need to add 'blog' and 'book' to the developer materials query. [~lkrzyzanek] can you do this for us? 4) We need to add `activity_date_interval` to the developer materials query. Again [~lkrzyzanek] can we add this? I commented on GitHub	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7216	DEVELOPER-1934	[~unibrew], OK, maybe I misunderstood. What additional content will come in? Will there be more than just Blogposts?	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7217	DEVELOPER-1934	Because one of the sys_types is 'article' then you'll also get all documents from https://developer.jboss.org. Other sys_types names that are listed for the query are not used in jbossorg provider. Would this be ok for you?	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7218	DEVELOPER-1934	So you mean, things like this will appear: https://developer.jboss.org/people/tremes/blog/2016/06/03/tracking-your-cdi-extension-actions? What about: https://developer.jboss.org/wiki/StudentResearchProjects. Would that appear?	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7219	DEVELOPER-1934	What is PR build URL for review please ?	1	1	https://issues.jboss.org/browse/DEVELOPER-1934
7220	DEVELOPER-1934	I now see blog post, but no books. Outstanding issues: 1. No books returned 2. if I do a search, select a result, then hit 'back', I see: https://www.dropbox.com/s/18rh57tnkb8h1rx/Screenshot%202016-06-20%2016.18.28.png?dl=0 3. Dates appear to render incorrectly. See 'October 9, 113' on the first item when visiting http://developers-pr.stage.redhat.com/pr/1100/build/1449/resources/ I think all three of the above are blockers for release. Also, I spotted these regressions: 1. Video thumbnails are no longer displayed in the result (where present, e.g videos) 2. The result type icon is missing from the results 3. I'm seeing results of type 'jbossorg_contributor_profile'. Notice how Libor and Rysiek's profiles appear at the top of the results: https://www.dropbox.com/s/n2c7eb62f6w7vzj/Screenshot%202016-06-21%2008.16.54.png?dl=0 And finally, I notice that you are now using 'http://dcp.stage.jboss.org/v2/rest/search?' instead of the registered query 'http://dcp2.jboss.org/v2/rest/search/developer_materials?'. Has that always been the case or was it a casualty of the merge?	0	0	https://issues.jboss.org/browse/DEVELOPER-1934
7221	DEVELOPER-1931	Below is the product description which can be used for OpenShift online. Is there other content needed for this page? If so, please give me the specific requirements and I will create content Thanks, Rebecca Faster Application Delivery with Red Hat OpenShift Online Break down IT productivity barriers. Trying to deliver more apps faster is often met by chronic productivity roadblocks. We believe the path to better IT efficiency and innovation lies in the cloud—via public Platform-as-a-Service (PaaS) with OpenShift Online by Red Hat®, or private PaaS with Red Hat's award-winning OpenShift Enterprise.	1	1	https://issues.jboss.org/browse/DEVELOPER-1931
7222	DEVELOPER-1929	Who needs to review this? Also, is the next step to produce the actual header graphic?	1	1	https://issues.jboss.org/browse/DEVELOPER-1929
7223	DEVELOPER-1925	[~paul.robinson] - is this issue a duplicate of 384?	1	1	https://issues.jboss.org/browse/DEVELOPER-1925
7224	DEVELOPER-1922	Is this ready for review? If so, can you update the status please.	1	1	https://issues.jboss.org/browse/DEVELOPER-1922
7225	DEVELOPER-1922	[~brock], I really like the darker banner in r2v1. This is looking great! One of my concerns though is how we can visually differentiate the cards from the Follow Red Hat block. Have you guys talked about maybe adding a shadow, small image or watermark to the cards that would help identify these as cards and be used throughout the site?	1	1	https://issues.jboss.org/browse/DEVELOPER-1922

#	Requirement	Questions	Answered	Successfully Answered	URL
7226	DEVELOPER-1922	[~james.cobb] or [~brock] Where can I find the assets these pages, like the page header images and the background images for the primary promos? or are those in a separate upcoming jira?	1	0	https://issues.jboss.org/browse/DEVELOPER-1922
7227	DEVELOPER-1921	We'd never planned to pull it from the DCP. It just comes from categories.yml. Can you take a look at how we implement it now and advise on wether the a-z ordering is possible?	0	0	https://issues.jboss.org/browse/DEVELOPER-1921
7228	DEVELOPER-1916	That job failed. Did you manage to review this with another job?	1	1	https://issues.jboss.org/browse/DEVELOPER-1916
7229	DEVELOPER-1916	Does that mean that this issue is blocked on DEVELOPER-2367?	1	1	https://issues.jboss.org/browse/DEVELOPER-1916
7230	DEVELOPER-1916	Is this resolved now? If not, what needs to be reviewed?	1	1	https://issues.jboss.org/browse/DEVELOPER-1916
7231	DEVELOPER-1914	[~james.cobb]: Is the issue here to just remove this text: "Let's walk through everything you need to build your first application.". Or is there some UX required to "rethink the arrangement of the Buy it and Try it elements if we reduce the content"?	1	1	https://issues.jboss.org/browse/DEVELOPER-1914
7232	DEVELOPER-1911	Shall we stand up our own blog instance or use the one that Mike guerette has?	0	0	https://issues.jboss.org/browse/DEVELOPER-1911
7233	DEVELOPER-1905	Who needs to review this? Also, is the next step to produce the actual header graphic?	1	1	https://issues.jboss.org/browse/DEVELOPER-1905
7234	DEVELOPER-1876	[~paul.robinson]/[~lightguard] - Can either of you please confirm that this is required for v1. I'm also not 100% on what this epic/story is trying to deliver. cheers, Rob	1	1	https://issues.jboss.org/browse/DEVELOPER-1876
7235	DEVELOPER-1830	@This is served up in Magnolia. When does it make sense to port this old content?	1	1	https://issues.jboss.org/browse/DEVELOPER-1830
7236	DEVELOPER-1828	[~paul.robinson] This is done, I just assumed I had to send all of the SlideShare work under the same PR, should I create one just for this part?	1	1	https://issues.jboss.org/browse/DEVELOPER-1828
7237	DEVELOPER-1828	[~adelaarreola], can you give an update on your findings. We can mark this as Done once that is reviewed.	0	0	https://issues.jboss.org/browse/DEVELOPER-1828
7238	DEVELOPER-1828	After speaking to [~adelaarreola]: 1. Embed a Slide deck in a page (Ensures we can embed the content on the RHD site) [WORKS] 2. Add slide decks to an album (is this even possible?) using the SlideShare site interface (Ensures that we can create product albums, for associating slides with a product) [NO ALBUM FUNCTIONALITY, NEED TO USE A SPREADSHEET TO LIST DECKS] 3. Pull slide meta-data via a REST API for a given slide deck (specified by a URL) [WORKS] 4. Pull slide meta-data via a REST API for all slide decks in an album (specified by a URL) [NO ALBUM FUNCTIONALITY, NEED TO USE A SPREADSHEET TO LIST DECKS]	0	0	https://issues.jboss.org/browse/DEVELOPER-1828
7239	DEVELOPER-1824	You say this is implemented. Where did you put the docs?	0	0	https://issues.jboss.org/browse/DEVELOPER-1824
7240	DEVELOPER-1821	[~mikgue] Thanks! [~tnolan] Please review for general consistency with the CDK page and let me know if you see anything we need to fix up before handing to front end. [~james.cobb] I'd like you to check out the callout in the header for visual styling. We will be using the callout area both on the containers page, and on the CDK product page in place of the try/buy bar (https://jboss.hotgloo.com/projects/73926/editor#/ux228a33xxft2bq). Do you want to use the same styling as the Try/Buy Bar, or would you want a unique styling for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1821
7241	DEVELOPER-1821	I'm working with [~mikgue] to finalize this content today. Google doc: [https://docs.google.com/a/redhat.com/document/d/15EJu9DWzK4f8PWRHgx_1YTbyvczLQKt_fGBDBmuZw5k/edit?usp=sharing]	1	1	https://issues.jboss.org/browse/DEVELOPER-1821
7242	DEVELOPER-1821	[~paul.robinson], can you let [~mikgue] and I know what's missing before handing over to the Front-end? The google doc (linked above) is updated with the content changes.	1	1	https://issues.jboss.org/browse/DEVELOPER-1821
7243	DEVELOPER-1821	[~crobson] [~tnolan] [~paul.robinson] Wireframe links are for hotgloo which I no longer have access to. Where can I find the current wireframe links for these pages?	1	1	https://issues.jboss.org/browse/DEVELOPER-1821
7244	DEVELOPER-1821	[~mikgue] I believe the only changes are to the Overview content (I took the new stuff from the google doc mentioned in the comments). Did I miss something?	1	1	https://issues.jboss.org/browse/DEVELOPER-1821
7245	DEVELOPER-1821	Can you add this list before the callout? Learn more Architecting Containers Part 1: Why Understanding User Space vs. Kernel Space Matters [http://rhelblog.redhat.com/2015/07/29/architecting-containers-part-1-user-space-vs-kernel-space/] Architecting Containers Part 2: Why the User Space Matters [http://rhelblog.redhat.com/2015/09/17/architecting-containers-part-2-why-the-user-space-matters-2/] Architecting Containers Part 3: How the User Space Affects Your Application [http://rhelblog.redhat.com/2015/11/10/architecting-containers-part-3-how-the-user-space-affects-your-application/] Can't We Just Run Boot2Docker in Production? [http://developerblog.redhat.com/2015/09/24/cant-we-just-run-boot2docker-in-production/]	1	1	https://issues.jboss.org/browse/DEVELOPER-1821
7246	DEVELOPER-1821	[~mikgue] what do you mean by 'callout'? Is it this: https://www.dropbox.com/s/xhgpwoxauoh0adi/Screenshot%202016-02-17%2007.51.12.png?dl=0	1	1	https://issues.jboss.org/browse/DEVELOPER-1821
7247	DEVELOPER-1821	yes. Is there a more common term for it??	1	1	https://issues.jboss.org/browse/DEVELOPER-1821
7248	DEVELOPER-1820	Handing over to Visual Design. [~james.cobb] - Can you create an ad or card treatment style for the "Need help?" It will be located under the left nav. [Wireframe | https://jboss.hotgloo.com/projects/73926/editor#/ux228a33xxft2bq]	1	1	https://issues.jboss.org/browse/DEVELOPER-1820
7249	DEVELOPER-1820	How about this [~paul.robinson], [~tnolan] ? "Got a technical CDK question? Head on over to Stack Overflow (linux-containers) to ask and Red Hatters will be on the lookout." Also, what about if we added something like "Got Stack privileges? Thank you for creating a new 'cdk' tag!" Mike	1	1	https://issues.jboss.org/browse/DEVELOPER-1820
7250	DEVELOPER-1820	We discussed this on HipChat. We felt this was better: {quote} Got a technical CDK question? Head on over to Stack Overflow (tag: linux-containers) to ask and Red Hatters will be on the lookout. {quote} Also, is 'Red Hatters' the right language?	1	1	https://issues.jboss.org/browse/DEVELOPER-1820
7251	DEVELOPER-1820	I see where you're headed now with this. How about this? Got a technical CDK question? You can submit questions on Stack Overflow (tag: linux-containers) and get answers from the Red Hat CDK Team.	1	1	https://issues.jboss.org/browse/DEVELOPER-1820
7252	DEVELOPER-1820	[~tnolan], [~mikgue]. Are we still happy with this approach?	1	0	https://issues.jboss.org/browse/DEVELOPER-1820
7253	DEVELOPER-1820	[~burrsutter] Burr - what do you think on this? Stack overflow, or the mailing list, or both? Or should the Stack portion be deferred until CDK GA (March)? To promote the CDK 2 beta mailing list, here's some text we can use: "We want your CDK beta feedback so join the upstream mailing list, container-tools@redhat.com. Go and join at https://www.redhat.com/mailman/listinfo/container-tools." Thoughts?	1	0	https://issues.jboss.org/browse/DEVELOPER-1820
7254	DEVELOPER-1809	Beyond the speed and testing (very important!) are there any business drivers that we're unable to achieve without this?	0	0	https://issues.jboss.org/browse/DEVELOPER-1809
7255	DEVELOPER-1809	[~ldary] so it seems like this is done. In addition there should be a JIRA issue for having Drupal select what JS to deliver on a per-page basis? Can you create an issue for that please?	0	0	https://issues.jboss.org/browse/DEVELOPER-1809
7256	DEVELOPER-1806	Is this now resolved?	1	1	https://issues.jboss.org/browse/DEVELOPER-1806
7257	DEVELOPER-1805	Can this be closed? Should be fixed by https://github.com/awestruct/aweplug/commit/a05d111bc1cdced301610bbfe75ea37cdc29d3c2	0	0	https://issues.jboss.org/browse/DEVELOPER-1805

#	Requirement	Questions	Answered	Successfully Answered	URL
7258	DEVELOPER-1805	Check the tags on the video links. Why aren't they showing Red Hat tags?	1	1	https://issues.jboss.org/browse/DEVELOPER-1805
7259	DEVELOPER-1799	[~paul.robinson] - So, it doesn't look like the blogs are being pushed to dcp1 anymore (my guess is because most of the Rivers are configured for DCP2). Should I incorporate dcp2 for this? Both RHD and jboss.org use dcp2 for their buzz sections.	1	1	https://issues.jboss.org/browse/DEVELOPER-1799
7260	DEVELOPER-1799	So, it looks like https://github.com/searchisko/configuration/blob/resources_query/data/query/resources.md can now be used on dcp2. Should I implement this on the page, or should I wait until it has been merged into the master searchisko configuration?	1	1	https://issues.jboss.org/browse/DEVELOPER-1799
7261	DEVELOPER-1799	[~wesbos] is this still blocked? If so, what's blocking it?	0	0	https://issues.jboss.org/browse/DEVELOPER-1799
7262	DEVELOPER-1797	I added a small block to the Video Presentation wireframe. Basic idea is to reuse the style from the Hello World on the product pages and add it to the top right of the page. We may need to play around with the height depending on the amount of content. I'm hoping that we can keep these succinct though. [~crobson], any additional feedback or comments before we move this issue along in the workflow? [https://jboss.hotgloo.com/projects/73926/editor#/u9busf7g0w885f8]	1	1	https://issues.jboss.org/browse/DEVELOPER-1797
7263	DEVELOPER-1797	[~paul.robinson], I updated the wireframe with some more realistic content. The idea is that we add the title of the related item in the first line, then reuse the same middle text and button action, that are on the full width bar. I can add in the sales-y stuff too if you think we need it, but would rather keep this simple. Does that help? https://jboss.hotgloo.com/projects/73926/editor#/u9busf7g0w885f8	1	1	https://issues.jboss.org/browse/DEVELOPER-1797
7264	DEVELOPER-1797	Can we also include a hover effect of going to solid red so it will look like the current buttons?	1	1	https://issues.jboss.org/browse/DEVELOPER-1797
7265	DEVELOPER-1795	This is coming along. Thanks for your work on this. Additional data: Can we get the publish date added and an author (if available)? Descriptions: # What is the number of characters being returned right now for the descriptions? We should be able to go up to at least 3 lines. # The descriptions also seem like they are being truncated oddly. For example, in a "jboss fuse" search, I see a result with the title of "Watch JBoss Fuse : Loan Broker Demo : on laptop". The description displayed in the search results starts with "1 Early Access running on a laptop." while the actual description starts with "This demo shows JBoss Fuse 6.1 Early Access". # Extra spaces are being removed after sentences for some reason. Example: "CDI has proven itself to be a great asset for Java.The many features it provides " is missing a space between Java and The. There are others that look like this too.	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7266	DEVELOPER-1795	Further questions that need to be added to AC: 1. If there is no search text entered, will the 'Search' button be disabled? 2. Do we have wireframes for mobile? Will there be a search field? If so how should it look?	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7267	DEVELOPER-1795	[~tnolan] We would need content approval on the below scenario (I just made the message up myself) Scenario: I search for something should return *no* entries, such as "bfehwfbhbn" Given I am on the Home page When I search for "bfehwfbhbn" Then I should see a message "Sorry there were no results found matching bfehwfbhbn" And there will be no results displayed	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7268	DEVELOPER-1795	[~paul.robinson] [~tnolan] We should also have a scenario to check the maximum characters for the search box. Can you provide me with that number?	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7269	DEVELOPER-1795	Is there a URL I can use to review this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7270	DEVELOPER-1795	[~james.cobb] For the search term highlight color in the description text, I see that the Customer Portal uses #fcf8e3, are you ok with using this same color or do you have an alternative you'd like to be used here?	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7271	DEVELOPER-1795	[~tnolan] - should the search box on the Search page also contain a Search and X button? The wifeframe includes them? https://redhatdevelopers.mybalsamiq.com/projects/hard-launch/Search	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7272	DEVELOPER-1795	[~tnolan] - Do we have an example of the usage?	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7273	DEVELOPER-1795	[~james.cobb] do we have a visual design for the search box in the chrome? See: https://www.dropbox.com/s/o8fc0nypzywpdd3/Screenshot%202016-05-26%2014.49.52.png?dl=0 Also, [~tnolan]: What should the text in the site-chrome search-box show? It's currently 'Filter by Keyword', which I don't think is right.	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7274	DEVELOPER-1795	[~tnolan] I think this is looking good. Do you have any other feedback (other than the quality of the results)?	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7275	DEVELOPER-1795	[~paul.robinson] [~wesbos] This is definitely looking good! I reviewed this as the latest build: http://developers-pr.stage.redhat.com/pr/1086/build/1284/ * The dates with each result: They all say "about 19 hours ago". I think we should only have dates appear on resource-type content, such as videos, blog posts etc. and not appear for general pages like product overview etc. I'm not sure what date options you have access to, but this should be a publish date or something similar. The dates should also be formatted as mon day year (like we use on the homepage) instead of using "time ago". * [~james.cobb] Is there a slightly larger font size we can use for the date on the results? Do you have a recommendation? * Can we fill in 2 full rows of the description? The second line seems to stop only after a few words. Maybe increase the max length for the data being returned here? Other than those minor things. Looks good!	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7276	DEVELOPER-1795	[~wesbos], [~james.cobb]: Do you have an update on [~tnolan]'s feedback?	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7277	DEVELOPER-1795	[~wesbos] [~paul.robinson] Here's some additonal feedback based on the most recent implementation. I queried containers: http://developers-pr.stage.redhat.com/pr/1086/build/1327/search/?q=containers # The date being used for both pages (like CDK - Updates) and resources (like cmt: Container Managed Transactions (CMT)) are all showing up as the same date. Is there another date field option that can be used here? The cmt: Container Managed Transactions (CMT), when you click through, has a publish date of Apr 22, 2016. The web pages (Like CDK -Updates) don't need to show a date as I'm guessing these would all be the same. # I'm not seeing a change in the font size for the date. Per [~james.cobb], this should be 14px. #-Many of the pages being returned for containers are all using the same description. "Red Hat Container Development Kit Hello World! Build Something Today Let's walk through everything you need to build your first application. Get Started" : Can we content from the page to create the description instead?- Nevermind on this one. I see why this is happening. [~paul.robinson] : Can we have the Get Started band removed in this build so we can see how these pages show up in search?	1	1	https://issues.jboss.org/browse/DEVELOPER-1795
7278	DEVELOPER-1795	[~tnolan] we'd need to remove the 'Get Started' bar from production to have that updated in the search results. Can we do that now, or do we need to wait for the new product overview pages (currently scheduled for post-summit).	0	0	https://issues.jboss.org/browse/DEVELOPER-1795
7279	DEVELOPER-1793	[~lvlcek] - Just looking at this now - the pagination seems to be working - did you implement this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1793
7280	DEVELOPER-1793	[~lvlcek] - I'm a bit confused as to the difference between `size` and `from`. It seems the API only returns 9 documents at a time? If so, what is the point of having `size25=true` if it only gives a max of 9? Are we supposed to make multiple requests if we want to show more than 9 at once? Would we have to do the math to figure out how to show page 25 items on page 5? Seems a overly complicated if that is true.	1	1	https://issues.jboss.org/browse/DEVELOPER-1793
7281	DEVELOPER-1793	[~wesbos] which server are you testing your queries against? Try http://dcp2.jboss.org/v2/rest/search/developer_materials?size15=true	1	1	https://issues.jboss.org/browse/DEVELOPER-1793

#	Requirement	Questions	Answered	Successfully Answered	URL
7282	DEVELOPER-1793	Ah - I'm using this url: http://dcp2-searchisko.rhcloud.com/v2/rest/search/developer_materials?size25=true Should I not be?	1	1	https://issues.jboss.org/browse/DEVELOPER-1793
7283	DEVELOPER-1791	Looks like the plugin is in place. What's left to be done on this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1791
7284	DEVELOPER-1788	[~dhladky] Does this sound feasible, or will we hit problems doing this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1788
7285	DEVELOPER-1788	Ideally staging PR builds should use an instance of the dl-man as close as possible to production. This will ensure that when we tests PR build, it behaves the same as it will in production. For this reason, we've used production instances of DL-man and SSO up until now. However, a problem came when we started doing regular testing of login and downloads on every PR build. These extra requests started to interfere with statistics in production data. Therefore we decided it best to move PRs to use staging instances of the back-end services. To make this work we need: 1) Data to be synced regularly from production dl-man to staging dl-man. Ideally nightly. 2) The staging instance of DL-man to be reliable. I.e it needs to stay up and functional. 3) New versions of DL-man to be backward compatible with client's configured to use previous version. This ensures that PR builds continue to work when a new version of the DL-man is put into staging prior to a release. Does this sound sensible?	1	1	https://issues.jboss.org/browse/DEVELOPER-1788
7286	DEVELOPER-1788	{quote} It does sound sensible, however it would require implementing the automatic sync between production and staging, which is not trivial, especially due to firewall problems. {quote} This will be less of a problem once we start using the REST API. At least then we will not have broken links. How often do you currently push production data back to the staging instance? {quote} Another approach would be adding a specific query parameter, that would opt out request from being processed by statistics. {quote} This is possible. However, we have similar issues for SSO as well. If we move to using the staging instance of SSO, we also need to use staging instance of DL-man.	1	1	https://issues.jboss.org/browse/DEVELOPER-1788
7287	DEVELOPER-1788	Is this okay to switch over now? What are the staging urls for the DL manager?	0	0	https://issues.jboss.org/browse/DEVELOPER-1788
7288	DEVELOPER-1788	[~dhladky] we are nearly done with tidying up the data in the production download manager. When this is done, can you push the changed data to the staging download manager?	1	1	https://issues.jboss.org/browse/DEVELOPER-1788
7289	DEVELOPER-1788	[~vladimir_v] are you able to copy the production data for dl-man to the staging instance? This will help our testing as we'll be using up-to-date data. Ideally this task would be done nightly, is that possible?	1	1	https://issues.jboss.org/browse/DEVELOPER-1788
7290	DEVELOPER-1788	I created this ticket for syncing the data: https://engineering.redhat.com/rt/Ticket/Display.html?id=386065 [~dhladky]: can you please confirm that you are happy with my suggested approach?	1	1	https://issues.jboss.org/browse/DEVELOPER-1788
7291	DEVELOPER-1788	Thanks [~dhladky], As i also mentioned in the ticket, a one-off data copy would get us a long way here. Do you now have the permissions to do that?	0	0	https://issues.jboss.org/browse/DEVELOPER-1788
7292	DEVELOPER-1786	[~james.cobb], Can you add the event card header image to the right location on static.jboss.org? I believe we need 2 sizes, one for RHD and the other for JBD?	1	1	https://issues.jboss.org/browse/DEVELOPER-1786
7293	DEVELOPER-1786	[~james.cobb]: can you point me to the location of these on the filemgmt server's file system? Also, don't we need a JBD version too? I'll try write up some docs around this, as there seems to be some confusion among the team.	1	1	https://issues.jboss.org/browse/DEVELOPER-1786
7294	DEVELOPER-1784	Thanks [~lvlcek]. Do you know how this compares to JBD? Also can you provide a list of the offending pages?	1	1	https://issues.jboss.org/browse/DEVELOPER-1784
7295	DEVELOPER-1784	[~lvlcek] using the jboss-developer account, I don't have access to the queries above. Can you save off the results from the queries above (just the ones for the data_rht_website content)? Hopefully we'll just need to run it once, so probably not worht giving me admin access.	1	1	https://issues.jboss.org/browse/DEVELOPER-1784
7296	DEVELOPER-1784	Thanks [~lvlcek], Might be easier to make it public. Are there any security concerns here?	1	1	https://issues.jboss.org/browse/DEVELOPER-1784
7297	DEVELOPER-1783	[~velias] can you review Adela's PR? Also, does this need to go live at the same time as the CDK 2 Beta download, or can it go live before?	1	1	https://issues.jboss.org/browse/DEVELOPER-1783
7298	DEVELOPER-1783	[~velias] it was done in error. I thought I fixed it shortly after though. Can you check this issue and RHD-1072 to make sure they make sense?	1	1	https://issues.jboss.org/browse/DEVELOPER-1783
7299	DEVELOPER-1777	[~crobson], is RHEL the only one that has/will have "children"? Maybe, since that's such a big product, we style that one a little different and then keep the sub-product ones (now being indented in the design) like the others on the page? I would love to hear more ideas from [~james.cobb] and team though.	1	1	https://issues.jboss.org/browse/DEVELOPER-1777
7300	DEVELOPER-1777	Thanks [~tnolan]. Is it ready for front-end dev? If so, please handover to front-end (using buttons above): https://www.dropbox.com/s/yyqh4w87ss8qnqu/Screenshot%202015-11-24%2009.01.39.png?dl=0	1	1	https://issues.jboss.org/browse/DEVELOPER-1777
7301	DEVELOPER-1777	Idea for the Note text style for the Get Started page that we talked about in the CDK meeting this morning: [https://db.tt/MZvOz6mt] [~james.cobb], can you provide the border and/or background colors for this based upon our site theme?	1	1	https://issues.jboss.org/browse/DEVELOPER-1777
7302	DEVELOPER-1777	[~tnolan] is this for all of the 'Notes' in the multiple CDK Get Started pages? or is this a new style for the 'Notes' in ALL products' get started pages?	1	1	https://issues.jboss.org/browse/DEVELOPER-1777
7303	DEVELOPER-1777	Will this callout only appear on the Get Started pages then?	0	0	https://issues.jboss.org/browse/DEVELOPER-1777
7304	DEVELOPER-1776	Add video game to 404 page? https://github.com/leereilly/games https://mattsurabian.com/duckhunt/ http://www.dougmcinnes.com/html-5-asteroids/	0	0	https://issues.jboss.org/browse/DEVELOPER-1776
7305	DEVELOPER-1776	Ping [~lightguard] - how do we add a 404?	1	1	https://issues.jboss.org/browse/DEVELOPER-1776
7306	DEVELOPER-1776	Hey [~newtonm] - looks like we need to add a line to the Apache .htaccess - something like `ErrorDocument 404 /404.html` Could we do this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1776
7307	DEVELOPER-1772	Can you provide a link to the PR in here please. We need a way to track the review process.	0	0	https://issues.jboss.org/browse/DEVELOPER-1772
7308	DEVELOPER-1771	What's the status of this?	0	0	https://issues.jboss.org/browse/DEVELOPER-1771
7309	DEVELOPER-1771	I think ideally we would pull these from `app.namespace`. We could have these set on page load in the body, right before we load the JS files. Is that doable?	1	1	https://issues.jboss.org/browse/DEVELOPER-1771
7310	DEVELOPER-1771	Are you hard coding localhost here? Or is that just generated?	1	1	https://issues.jboss.org/browse/DEVELOPER-1771
7311	DEVELOPER-1771	I don't want to generate any JavaScript in .js files - just the code that is embedded into a script tag on the page- does that make sense?	0	0	https://issues.jboss.org/browse/DEVELOPER-1771
7312	DEVELOPER-1770	Pete, When is this needed by? Can you give it a fix version?	1	1	https://issues.jboss.org/browse/DEVELOPER-1770
7313	DEVELOPER-1770	[~pmuir], [~maxandersen]: Is this still a requirement? I'm not saying it's something we plan to work on soon, just deciding wether to close as 'out of date' or not.	1	1	https://issues.jboss.org/browse/DEVELOPER-1770
7314	DEVELOPER-1736	[~james.cobb], [~myfear]: When does this need to go live?	1	1	https://issues.jboss.org/browse/DEVELOPER-1736

#	Requirement	Questions	Answered	Successfully Answered	URL
7315	DEVELOPER-1736	What is the URL? It seems, nothing is working: http://jboss.org/events/events/judcon/2015/brno/ http://jboss.org/events/events/judcon/2016/brno/ <= the expected one :) http://developers.redhat.com/events/events/judcon/2015/brno/ http://developers.redhat.com/events/events/judcon/2016/brno/ Thanks, m	1	1	https://issues.jboss.org/browse/DEVELOPER-1736
7316	DEVELOPER-1733	What do you mean: they are no longer issues with it as all the current product have featured artifact (this can change when a new product will be added) or you check this state and the page does not break up when the featured artifact is not present. If it is the second case, then this ticket can be closed. If it is the first case, then it needs to be fixed still.	1	1	https://issues.jboss.org/browse/DEVELOPER-1733
7317	DEVELOPER-1729	[~crobson] is this approved yet? If not, is it close enough for [~adelaarreola] to start work on it?	1	1	https://issues.jboss.org/browse/DEVELOPER-1729
7318	DEVELOPER-1729	[~paul.robinson] Wireframe is complete and approved. Should I pass off to Visual Design or Front End? There is a banner needed by visual team at the top, and the pictures for the speaker may need to be created/edited.	1	1	https://issues.jboss.org/browse/DEVELOPER-1729
7319	DEVELOPER-1729	Thanks Adela. Comments: * Are these white bits expected? https://www.dropbox.com/s/86l7zzsw00mcty4/Screenshot%202015-10-12%2017.09.21.png?dl=0. Looks like the top of another word is being chopped off. * Should the heading be 'full bleed'? It looks quite 'intense' at the moment. However, I'l defer to someone with a 'design eye'. * Remember to add content to share button Otherwise, looks great!	0	0	https://issues.jboss.org/browse/DEVELOPER-1729
7320	DEVELOPER-1729	Is the resolution on the banner high enough? I'm viewing it on a retina display, so that might be the issue. Here's a screenshot, not sure if you will see the change in sharpness between the banner text and the rest of the page: https://www.dropbox.com/s/6vh6ovndoazz63u/Screenshot%202015-10-14%2011.14.51.png?dl=0	1	1	https://issues.jboss.org/browse/DEVELOPER-1729
7321	DEVELOPER-1718	If we're going to use these on the site too, should they be the illustration styling, or will we make separate images for that use? (example: http://developers.redhat.com/products/eap/overview/)	0	0	https://issues.jboss.org/browse/DEVELOPER-1718
7322	DEVELOPER-1718	Are these ready to go on the CDK product page?	1	1	https://issues.jboss.org/browse/DEVELOPER-1718
7323	DEVELOPER-1717	What about this: https://origin-it-developers.redhat.com/auth/js/keycloak.js ?	0	0	https://issues.jboss.org/browse/DEVELOPER-1717
7324	DEVELOPER-1716	Markus, did they supply you with a mockup of the cover we can use in the banner?	0	0	https://issues.jboss.org/browse/DEVELOPER-1716
7325	DEVELOPER-1716	[~brock] Those images are sized for JBoss Developer. Can you create a set for RHD?	1	1	https://issues.jboss.org/browse/DEVELOPER-1716
7326	DEVELOPER-1716	Are we to run this promo on both RHD and jboss.org?	0	0	https://issues.jboss.org/browse/DEVELOPER-1716
7327	DEVELOPER-1716	[~brock] The links for the RHD mockups don't work, is it different than the JBoss.org mockup?	1	1	https://issues.jboss.org/browse/DEVELOPER-1716
7328	DEVELOPER-1713	David, We have a series of yaml files containing this data. Here's an example for EAP: https://github.com/jboss-developer/www.jboss.org/blob/master/products/eap/_common/product.yml#L26-L171 Would that be enough to recreate most of the data?	1	1	https://issues.jboss.org/browse/DEVELOPER-1713
7329	DEVELOPER-1713	I am afraid it is not readable for me. I do not see a way how to map the things in it to the actual files. Would you be able to create let us say a spreadsheet, that contains lines with columns: product, product version, file name (including the access.cdn path), description (optional), TC model (anonymous/no tc/Zero Dollar) (optional)?	1	1	https://issues.jboss.org/browse/DEVELOPER-1713
7330	DEVELOPER-1713	[~dcoughlin1] - Have you inputted the data for this DL manager?	1	1	https://issues.jboss.org/browse/DEVELOPER-1713
7331	DEVELOPER-1713	Also, there are quite a few that are being displayed on http://it-developers-pr.stage.redhat.com/pr/332/build/826/products/devstudio/download/ (e.g. https://www.dropbox.com/s/bdblkdtoxv22q5q/JBoss_Developer_Studio_-_Downloads.jpg?dl=0) that are not displayed on http://developers.redhat.com/products/devstudio/download/. Should I just be ignoring those for now?	1	1	https://issues.jboss.org/browse/DEVELOPER-1713
7332	DEVELOPER-1713	[~paul.robinson] [~dhladky] I can link the products together. Should I add those products and set them as master for the zero dollar/No T&C?	1	1	https://issues.jboss.org/browse/DEVELOPER-1713
7333	DEVELOPER-1713	Can this be closed?	1	1	https://issues.jboss.org/browse/DEVELOPER-1713
7334	DEVELOPER-1712	[~ecurley] Is the content for this microsite ready? If so, where can I find it? Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-1712
7335	DEVELOPER-1712	Hi Paul - Here was the feedback I had sent Adela. My apologies as we were communicating in hipchat back and forth. Can we add a line space in the red hat booth band between the 17:00? and the "Meet the ..." Devoxx for Kids - should say "Devoxx4Kids" And that band will need to be hidden until Wednesday of the conference In the "Meet Markus" headline. Can we swap "BUY" for "GET"?	1	1	https://issues.jboss.org/browse/DEVELOPER-1712
7336	DEVELOPER-1701	[~mikgue] Can you respond to Libor's query?	1	1	https://issues.jboss.org/browse/DEVELOPER-1701
7337	DEVELOPER-1701	Will this work? https://en.support.wordpress.com/export/	1	1	https://issues.jboss.org/browse/DEVELOPER-1701
7338	DEVELOPER-1701	[~mguerri] just read the answer from wordpress and I'd like to give a try the export file: https://en.support.wordpress.com/export/ Could you please export whole blog and send me to my RH e-mail? Thanks.	1	1	https://issues.jboss.org/browse/DEVELOPER-1701
7339	DEVELOPER-1701	[~mikgue] Thanks a lot. I had to fix few character issues but it works for me. [~paul.robinson] here is output of all entries from feed. Please check: https://planet-stg.jboss.org/#tags=feed_group_name_nonmiddleware The archive is perfect but I just realised that current feed (https://developerblog.redhat.com/feed/) doesn't provide whole content and has only "description" which is short version of it. Planet doesn't index whole content then (e.g. https://planet-stg.jboss.org/post/behind_the_internet_of_things_middleware_demo_at_red_hat_summit). [~mikgue] is it possible to change the setting of this actual feed to serve whole content as well?	1	1	https://issues.jboss.org/browse/DEVELOPER-1701
7340	DEVELOPER-1701	Also, will this process ensure the posts are re-indexed as new tags are added? Someone will need to go through all the existing posts and assign product tags to the ones we need to appear on a product page.	1	1	https://issues.jboss.org/browse/DEVELOPER-1701
7341	DEVELOPER-1701	[~lkrzyzanek] I've created a PR for 4. https://github.com/redhat-developer/developers.redhat.com/pull/353 Note, that it will pull data from dcpbeta-searchisko.rhcloud.com. Is that a problem?	1	1	https://issues.jboss.org/browse/DEVELOPER-1701
7342	DEVELOPER-1701	Actually the last link should be working (you need to be on the VPN): http://it-developers-pr.stage.redhat.com/pr/353/build/1294/products/softwarecollections/buzz/ Does that look right?	1	1	https://issues.jboss.org/browse/DEVELOPER-1701
7343	DEVELOPER-1701	Could you please create jira issue for it and link it. Thanks. I'll create separate issue about using "snippet" in planet.	1	1	https://issues.jboss.org/browse/DEVELOPER-1701
7344	DEVELOPER-1701	Can we try one out? Does UX need to do anything??	0	0	https://issues.jboss.org/browse/DEVELOPER-1701
7345	DEVELOPER-1697	Why is this a 'nice to have', isn't it critical to the Q1FY16 release? Or did you have an alternative way to view entitlements?	1	1	https://issues.jboss.org/browse/DEVELOPER-1697
7346	DEVELOPER-1697	I changed the class of work based on Deliverables table. This similar one RHD-63 is Must have.	0	0	https://issues.jboss.org/browse/DEVELOPER-1697
7347	DEVELOPER-1697	Thanks David. I notice you also released the mock API too. Do we need the Mock one, now that this is ready?	1	1	https://issues.jboss.org/browse/DEVELOPER-1697

#	Requirement	Questions	Answered	Successfully Answered	URL
7348	DEVELOPER-1697	Hi [~wesbos], I just added the cors header to the available downloads response. I allowed it for all origins and GET requests (other methods are not supported anyway). The RPM is being built As for 2 it is definitely a bug, that should be addressed. I will create a ticket for it and fix it as soon as possible. Ideally today. [~paul.robinson], will you mind, if I simplify and fix the response? I am not sure if the service is already used somewhere as it is.	1	1	https://issues.jboss.org/browse/DEVELOPER-1697
7349	DEVELOPER-1697	[~dhladky] what do you mean by "simplify and fix the response"? Neither JBoss Developer or Red Hat Developers are using this API yet in production. So I don't have a problem with you changing the API.	1	1	https://issues.jboss.org/browse/DEVELOPER-1697
7350	DEVELOPER-1697	How does it differ to the previous version of the response?	1	1	https://issues.jboss.org/browse/DEVELOPER-1697
7351	DEVELOPER-1697	[~paul.robinson] I've just pushed another version - right now I'm linking to the generic release notes page because there is not way to dynamically generate this url - they are all too different. We need the data to go into the DCP. [~dhladky] - Could we get the ability to add a "download" to the download manager that isn't a file, but a link? We need each release to have a link to the release notes.	1	1	https://issues.jboss.org/browse/DEVELOPER-1697
7352	DEVELOPER-1697	Why do you need a download manager for it? Do you want to enforce some subscription on the release notes as well? If so, you can put the release notes to access.cdn just like any other file. If you do not need the release notes to be protected, then you do not need DM for the release notes.	1	1	https://issues.jboss.org/browse/DEVELOPER-1697
7353	DEVELOPER-1697	Hi, so how about adding the link to the product version then (I suppose this is the right place for the info) and adding it to the available downloads rest service (another property of product version). Are you aware of any other links you would like to be added? Please, create a new ticket for this feature as it is definitely not part of this ticket.	1	1	https://issues.jboss.org/browse/DEVELOPER-1697
7354	DEVELOPER-1696	OK, Some additional comments. Carousel... 1) The switcher should be turned off once a person clicks on a link in the carousel bar. As it works now, I click on an item in the bar that looks interesting and it goes to that slide (yeah!) then after 5 seconds it starts rotating again (boo!). 2) The block quote is actually a super light blue and the bold copy is white. The blue hex color: #eafbff 3) The job title /company spacing is still far to much "Vice President, Software Engineering / Red Hat Inc." The company should be much tighter to the job title. 4) Hmmm... the download icon is titled "Icon_RH_Software_Transfer-Download_RGB_Shiny.svg" and I'll separately attach it to this jira. 5) "The SVG icon can't have a gradient in the BG unless we use images." Not following you here. What icon's have gradients? Do you mean the photo in the carousel?	1	1	https://issues.jboss.org/browse/DEVELOPER-1696
7355	DEVELOPER-1683	I just checked a bunch of these and the dates and ordering seems correct - can we close this?	0	0	https://issues.jboss.org/browse/DEVELOPER-1683
7356	DEVELOPER-1667	[~paul.robinson] For the descriptions, should I just plan on pasting in information from the about sections of the websites? I noticed that most of the descriptions start with "Join Redhat..." and have been customized. For now, I am just leaving the descriptions blank.	0	0	https://issues.jboss.org/browse/DEVELOPER-1667
7357	DEVELOPER-1663	[~james.cobb], [~crobson]: should this live at http://developers.redhat.com/resources?	1	1	https://issues.jboss.org/browse/DEVELOPER-1663
7358	DEVELOPER-1663	[~james.cobb], [~crobson]: Can you review: http://it-developers-pr.stage.redhat.com/pr/419/build/837/. Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-1663
7359	DEVELOPER-1662	[~brock] Do you have an update on how this is going? Are any of the icons available?	1	1	https://issues.jboss.org/browse/DEVELOPER-1662
7360	DEVELOPER-1662	Sure thing, Paul! I've completed a handful but I haven't made them available yet. DevNation has been taking up a lot of my focus this week. Please let me know if there is are any that you need more than others so that I can prioritize them accordingly.	0	0	https://issues.jboss.org/browse/DEVELOPER-1662
7361	DEVELOPER-1656	Do I close this? I am moving on to pushing this into the staging environment. There are some things that need to be done beforehand (mainly just getting Rysiek approve the template), but I think I'm ready to move onto the next step here.	1	1	https://issues.jboss.org/browse/DEVELOPER-1656
7362	DEVELOPER-1652	Please can you add a description?	0	0	https://issues.jboss.org/browse/DEVELOPER-1652
7363	DEVELOPER-1647	[~paul.robinson] Hey Paul, can you provide me with access to this Youtube account?	1	1	https://issues.jboss.org/browse/DEVELOPER-1647
7364	DEVELOPER-1647	You were invited to be a manger a while ago. Can you search your email, maybe in the spam.	1	1	https://issues.jboss.org/browse/DEVELOPER-1647
7365	DEVELOPER-1647	[~paul.robinson] So, many of the generic "Virtual: JBUG" are presentations that either don't have slides or don't actually have a good view of the presentation. I have changed a few, but is it acceptable to pull a snapshot of just the presenter or should I plan to leave the generic thumbnail?	1	1	https://issues.jboss.org/browse/DEVELOPER-1647
7366	DEVELOPER-1644	[~ecurley] can you take a look at the copy on the Wireframe. The JBoss EAP one talks about RHEL. Can you also check you are happy with the rest of the text?	1	1	https://issues.jboss.org/browse/DEVELOPER-1644
7367	DEVELOPER-1644	[~lkrzyzanek] once we create the page, is it you that will wire it into the registration flow?	1	1	https://issues.jboss.org/browse/DEVELOPER-1644
7368	DEVELOPER-1644	Conceptually I'm good with it but the text will need to change on them. OSE doesn't have a getting started experience. Shall I change it here or in hotgloo?	1	1	https://issues.jboss.org/browse/DEVELOPER-1644
7369	DEVELOPER-1644	[~ecurley], [~ray.ploski]: are you happy with this page? See: http://it-developers-pr.stage.redhat.com/pr/545/build/1130/confirmation/ I've checked the process, and the user is directed to this page after verifying their email.	0	0	https://issues.jboss.org/browse/DEVELOPER-1644
7370	DEVELOPER-1644	It looks like the text gets cut off behind the button for the EAP get started. Can we set it up so that the text has an upper size limit and becomes ellipsis (...) if it overflows that limit?	1	1	https://issues.jboss.org/browse/DEVELOPER-1644
7371	DEVELOPER-1641	Can you tell me where this wrong link is?	0	0	https://issues.jboss.org/browse/DEVELOPER-1641
7372	DEVELOPER-1638	Keycloak login server doesn't provide a single URL right now. However I think we can achieve this UX requirement by a single static page http://developers.redhat.com/register which would either have some text plus button "Register Now" or display nothing and put javascript that will immediatelly redirect user to register page. WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-1638
7373	DEVELOPER-1636	[~velias] The issue reporter field was mandatory. I set it to [~dcoughlin1] as he'll be working on setting this up. Can you make it optional, or create a new user for this. It's not ideal that the crazy comments, that we'll inevitably get, will be accredited to dan ;-)	1	1	https://issues.jboss.org/browse/DEVELOPER-1636
7374	DEVELOPER-1636	can you use my email address initially. We can re-think that later.	1	1	https://issues.jboss.org/browse/DEVELOPER-1636
7375	DEVELOPER-1636	[~crobson], [~james.cobb]: can you review: http://it-developers-pr.stage.redhat.com/. Notice the feedback link in bottom right. I went for the least obtrusive option. We can select something with more visibility if you prefer. Maybe we walk through the options in a quick call?	1	1	https://issues.jboss.org/browse/DEVELOPER-1636
7376	DEVELOPER-1636	[~paul.robinson][~crobson] - So I click the link and only see the homepage that doesn't have a "got feedback option. Is there a direct link to a page where I can see this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1636
7377	DEVELOPER-1634	Are we looking to do this now, or after initial launch?	1	1	https://issues.jboss.org/browse/DEVELOPER-1634
7378	DEVELOPER-1633	I think we should change the icons to illustrations. Shall we create a separate JIRA?	1	1	https://issues.jboss.org/browse/DEVELOPER-1633
7379	DEVELOPER-1631	[~crobson] is this approved? Can it be handed over to front-end?	1	1	https://issues.jboss.org/browse/DEVELOPER-1631

#	Requirement	Questions	Answered	Successfully Answered	URL
7380	DEVELOPER-1631	These look great. And agree on the card version. [~crobson] What area of the headers are you looking at? My two cents - There is a lot going on up top on this one: http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_1_0/images/RHDev_products_getstarted_dualpath_r1v1.jpg A couple of options (would love both of your thoughts) -Delete the "let's walk through..." in the Hello World section. With 4 things there it is hard to know what to do. -I'm normally a fan of allowing folks to easily contact sales but I wonder if it is premature up top. This is a first getting started. Maybe something down below? On this page: http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_1_0/images/RHDev_products_getstarted_dualpath_r1v1.jpg I would consider shortening "How to get started with Red Hat..." to just "Get started with Red Hat..." small change but more actionable and will reduce some of the wrap. Are you in need for shorter header copy from what exists now? I can trim if needed.	1	1	https://issues.jboss.org/browse/DEVELOPER-1631
7381	DEVELOPER-1628	Can this go into production now?	1	1	https://issues.jboss.org/browse/DEVELOPER-1628
7382	DEVELOPER-1628	[~lkrzyzanek] Can we push the RHD change live now, and the JBD change later today or tomorrow?	1	1	https://issues.jboss.org/browse/DEVELOPER-1628
7383	DEVELOPER-1623	Do we still plan to run these tests against production? Not sure it makes sense anymore?	1	1	https://issues.jboss.org/browse/DEVELOPER-1623
7384	DEVELOPER-1603	Hi Paul. I'm working to get the Hello world approved this week and if yes we will need a RH Mobile section within products. We are planning on promoting the Hello World for Mobile at NodeConf on Sept. 6th. Do you have a URL that we will use when that section goes up? James and I are working on a postcard to hand out directing people there. Thanks! Emily	1	1	https://issues.jboss.org/browse/DEVELOPER-1603
7385	DEVELOPER-1603	{quote} Is there a way to make the 4 icons in a row across the bottom or would 2 on each row work? {quote} For consistency with other products, I think it would be better to go with 3. Otherwise we'll end up with having to resize the content. Can one of those items be dropped?	1	1	https://issues.jboss.org/browse/DEVELOPER-1603
7386	DEVELOPER-1603	Also, would it be better to move the 3 key feature further up? We added those to the other product pages as a way to quickly grab the reader and use it as a way of getting them interested enough to read the more detailed information below. I would suggest adding it after the opening paragraph.	1	1	https://issues.jboss.org/browse/DEVELOPER-1603
7387	DEVELOPER-1603	[~james.cobb], [~brock]: Can you advise on what images to use on here: http://it-developers-pr.stage.redhat.com/pr/381/build/790/products/mobileplatform/overview/. I'm looking at these ones that are scaling badly: https://www.dropbox.com/s/b4qsh8i9o21gsvb/Screenshot%202015-09-03%2020.30.40.png?dl=0 Do we have illustrations?	1	1	https://issues.jboss.org/browse/DEVELOPER-1603
7388	DEVELOPER-1603	[~ecurley]: should the link to Mobile learn, really filter on 'iOS'? See this link on the overview page: {quote} See Mobile in action http://developers.redhat.com/mobile/learn/#!query=iOS {quote}	1	1	https://issues.jboss.org/browse/DEVELOPER-1603
7389	DEVELOPER-1603	Thanks Jason, I don't know wether we want scalable .svg or transparent .png. Can you check with [~james.cobb]? We want to push this live tomorrow, not sure if it's feasible to get the illustrations in that time-frame?	1	1	https://issues.jboss.org/browse/DEVELOPER-1603
7390	DEVELOPER-1599	Given that we plan to use Drupal's CDN integration, do we still need this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1599
7391	DEVELOPER-1595	Can we list the languages in alphabetical order: node, Python, Ruby? Otherwise, good to go. Mike	1	1	https://issues.jboss.org/browse/DEVELOPER-1595
7392	DEVELOPER-1592	[~brock] Can you review Adela's build: http://it-developers-pr.stage.redhat.com/pr/613/build/197	1	1	https://issues.jboss.org/browse/DEVELOPER-1592
7393	DEVELOPER-1464	One minor addition, can we add the month and year under the title? This was on the CDK wireframe and I just noticed it was missing here. It should be the smaller font size and a slightly lighter grey color. [https://db.tt/VHXDP7aN]	1	0	https://issues.jboss.org/browse/DEVELOPER-1464
7394	DEVELOPER-1463	Content from Markus [about judcon] JUDCon is coming to town! After some time of silence, it is about time to have another JUDCon. 2016 will start off with JUDCon Brno and we're very proud to be part of the DevConf.cz conference with a full day track of Microservices, Mobile, Cloud and Java EE. The JBoss Users and Developers Conferences are developer gatherings held around the globe to give JBoss users the chance to talk to and collaborate with leaders in the JBoss Community, and to allow the project leads and engineers to learn from and collaborate with end users. The JBoss Core Developers, along with the open source community, create and support the projects that drive innovation and help lead development in standards bodies like the Java Community Process, the Apache Software Foundation, OASIS, W3C and other open standards organizations. [JUDCon Schedule] Friday, February 5th 2016 9:50 Knowledge driven micro services - Maciej Swiderski [Microservices] 10:40 SilverWare: Microservices - Martin Swiech [Microservices] 11:30 Docker, Vagrant and Kubernetes walks into an Eclipse'd bar - Max Rydahl Andersen [Mobile & Cloud] 12:20 Java Enterprise'ish development without hassle - Jakub Marchwicki [JavaEE] 13:10 Getting S#it Done With JBoss Forge: A Gentle Path into Java EE 7 Applications - Koen Ars + Sebastien Blanc [Java EE] 14:00CI/CD with Openshift and Jenkins - Ari LiVigni, Aaron Weitekamp, Michal Fojtik [Mobile & Cloud] 14:50 Red Hat Mobile Application Platform + OpenShift Online - John Frizelle [Mobile & Cloud] 15:40 You've Got Microservices Now Secure Them - Steven Pousty [Microservices] 16:30 Taming Microservices Testing with Docker and Arquillian Cube - Aslak Knutsen [Mobile & Cloud] 17:20 WildFly Swarm: does my fatjar look big in this? - Mark Little [JavaEE] 18:10 Are you ready for a new WildFly security? Elytron! - Josef Cacek [JavaEE] [talk category description] Microservices As an architectural pattern, Microservices has graduated from being an esoteric buzzword and has great potential in many organizations - ranging from startups and large consumer internet organizations to enterprises. The microservices talks help bring in real world experience and first hand information from the engineers working on the future. Mobile & Cloud While mobile left the real hype and turned into commodity for many users, the combination with cloud and changing requirements in today's enterprises still influence the way we built mobile applications. The mobile and cloud talks by our engineers help you to make the best out of both. JavaEE As the backbone of Red Hat JBoss Middleware and used by literally every enterprise in the world we added the most interesting talks about latest developments to this category and invite you to brush up your knowledge and experience first hand information from our engineers. [more] Information about the city and venue: http://devconf.cz/city-and-venue Link to the code of conduct: http://devconf.cz/code-conduct	1	1	https://issues.jboss.org/browse/DEVELOPER-1463

#	Requirement	Questions	Answered	Successfully Answered	URL
7395	DEVELOPER-1463	constantly growing business demands more and more focus is put on gathering business knowledge and making use of it in IT. Knowledge Is Everything - that's a slogan that lots of business domains is catching up right now. On the other hand, recent trends in software development are shifting into micro services environments where loosely coupled, standalone and independent services provide access to various business assets. With that in mind, knowledge driven micro service idea came into live to combine couple of great technologies together to satisfy such needs: - vert.x - toolkit for building reactive applications on JVM - jBPM and Drools flexible process and rule engines that provide excellent execution environment for business assets - the knowledge - hazelcast - in memory data grid - various programming languages (Java, JavaScript, Scala, Groovy) to attract various development teams This session is about illustrating how easily and fast you can build knowledge driven micro services and scale them to almost no limits. All in reactive and polyglot fashion with powerful clustering. Scale individual knowledge micro services without a hassle and allow various languages to easily interact with business knowledge without a need to know where they are located. It all goes via clustered event bus provided by Vert.x, backed by deployable knowledge archives provided by jBPM and Drools, clustering delivered by hazelcast data grid and allow polyglot client side access. It's all about business knowledge that can be delivered as software solution with: - flexibility - speed of development - isolated scalability - manageability Come and see how to transform your IT solutions into knowledge driven micro services with Vert.x, jBPM, Drools and Hazelcast. This session will be full of life demos to actually show the running micro services driven by business knowledge. Speaker info: Maciej Swiderski Red Hat, Senior Software Engineer Speaker image: https://static.sched.org/a7/950423/avatar.jpg.320x320px.jpg?e9e _____ Session #2: SilverWare: Microservices 10:40 - 11:20 Description: SilverWare.io is a minimalistic, yet powerful, modularized, highly configurable and container-free implementation of Microservices. Builds on technologies you are familiar with. Start creating the Internet of Things today! Speaker Info: Martin Swiech I am working for RedHat since 2014 as middleware QE. Speaker image: None _____ Session #3: Docker, Vagrant and Kubernetes walks into an Eclipse'd bar 11:30 - 12:10 Description: Todays software industry is booming with tools like Docker, Kubernetes, Vagrant and more to enable faster turnaround times via virtualization. How does Eclipse the IDE deal with that ? Can it deal with it ? How should it and its users deal with it ? Come to this talk and hear how and what Eclipse Linux Tools and JBoss Tools are doing in the area. It will give a brief overview on the Docker and Vagrant tooling in Linux Tools and then how that is used and combined in Kubernetes and OpenShift Tools. The talk will show how you can use Docker and Vagrant practically today with plain Eclipse to improve your own local development experience and then shortly cover how it all comes together in the cloud and container space. Some basic understanding of the container and virtualization technologies are recommended, but not an absolute requirement. Speaker Info: Max Rydahl Andersen Born and raised in Denmark, worked on health care software systems for some time. In that work I bumped into this small project called Hibernate and had to fix a couple of things in it to make it useful Since then I've been working on Hibernate Core, Hibernate Tools, Seam and now lately JBoss Tools and Developer Studio Speaker image: https://static.sched.org/a7/2534492/avatar.jpg.320x320px.jpg?512 _____ Session #4: Java Enterprise'ish development without hassle 12:20 - 13:00 Description: This happens in every community: You don't need Rails, to build a web application in Ruby. You don't need templating engine for PHP - PHP is a templating language! We over engineer and over engineer, again and again. Same happens in Java. Hibernate + EJB + JSF is not the only blessed stack for development of web application… one can argue it's the	1	1	https://issues.jboss.org/browse/DEVELOPER-1463
7396	DEVELOPER-1463	[~myfear]: do you want the above content adding to the page?	1	1	https://issues.jboss.org/browse/DEVELOPER-1463
7397	DEVELOPER-1461	I can add code to pull this image into the site; however, all non-default card banners are currently coming from http://static.jboss.org/images/rhd/events/event_{*_picturename_*}.png. If possible, can you place the file under http://static.jboss.org/images/rhd/events? Also, is it possibles to rename the file to *_event_*_makerfile_judcon2016brno.png instead of events_makerfile_judcon2016brno.png? Sorry for all the changes. If it's not possible, please let me know and I'll add in the javascript logic to do this.	1	1	https://issues.jboss.org/browse/DEVELOPER-1461
7398	DEVELOPER-1442	[~dcoughlin1] Can you work on this today please?	1	1	https://issues.jboss.org/browse/DEVELOPER-1442
7399	DEVELOPER-1437	Hi [~james.cobb], can the text on these graphics be moved a little to the left? Right now it gets cut off once the window becomes smaller because the text is positioned outside the page's content area	1	1	https://issues.jboss.org/browse/DEVELOPER-1437
7400	DEVELOPER-1436	Can we make this as a developers.redhat.com microsite?	1	1	https://issues.jboss.org/browse/DEVELOPER-1436
7401	DEVELOPER-1435	Looks good overall. It would have been nice if the new section included the standalone installer from 9.0.0.GA too so users wouldn't have to click show downloads. Is that possible ?	1	1	https://issues.jboss.org/browse/DEVELOPER-1435
7402	DEVELOPER-1435	How about we add a new "release notes" entry, which links to the blog? PR: https://github.com/jboss-developer/www.jboss.org/pull/1006 This mirrors what EAP does, eg., on this page http://www.jboss.org/products/eap/download/	1	1	https://issues.jboss.org/browse/DEVELOPER-1435
7403	DEVELOPER-1434	[~sparulka] when is this due by? We have some other high priority stuff on at the moment. Would late this week or early next week work?	1	1	https://issues.jboss.org/browse/DEVELOPER-1434
7404	DEVELOPER-1434	This is due today. Product GA was yesterday Dec 7 and I had communicated this earlier and had already submitted the content changes except the actual links to the bits. Very minor content changes. I received the CDN links for the GA bits yesterday night (Dec 7). I sent the GA availability announcement today at noon (Dec 8). So appreciate if this is posted by atleast Dec 9 EOD. Doable?	1	1	https://issues.jboss.org/browse/DEVELOPER-1434
7405	DEVELOPER-1434	[~sparulka] Can you review the build ^^	1	1	https://issues.jboss.org/browse/DEVELOPER-1434
7406	DEVELOPER-1434	[~sparulka] Adela is working on the fixes know. Are you going to be available to do a final review over the next 2hrs? I'll keep an eye out and merge the PR once you approve.	1	1	https://issues.jboss.org/browse/DEVELOPER-1434
7407	DEVELOPER-1428	Should I plan on adding this to the download table?	1	1	https://issues.jboss.org/browse/DEVELOPER-1428
7408	DEVELOPER-1418	Can this be closed?	0	0	https://issues.jboss.org/browse/DEVELOPER-1418
7409	DEVELOPER-1408	Where does the /content/... link suppose to reside? Is that a location where the CSP puts their files?	1	1	https://issues.jboss.org/browse/DEVELOPER-1408
7410	DEVELOPER-1398	Why is it out of date?	0	0	https://issues.jboss.org/browse/DEVELOPER-1398
7411	DEVELOPER-1379	Hey [~nickboldt] - I've got an updated ascii-doc and image for this - I need credentials to test (running into the vault): Request the following values from the JBoss Developer team: vimeo_access_token_secret vimeo_client_secret vimeo_access_token dcp_user dcp_password google_api_key Who could provide these tokens?? Thkx	1	1	https://issues.jboss.org/browse/DEVELOPER-1379
7412	DEVELOPER-1371	[~nickboldt] Hey Nick, can you provide the download links (i.e. /content/origin/files/sha256/3d/3d115879e3be646a563d05452c3a22bac394652e40260a3b43d9dd5c3ba7121e/jboss-devstudio-9.0.0.Beta1-installer-eap.jar) when you get the chance?	1	1	https://issues.jboss.org/browse/DEVELOPER-1371
7413	DEVELOPER-1371	[~cpob] The download manager should be working now. Can you try and add the last remaining download?	1	1	https://issues.jboss.org/browse/DEVELOPER-1371
7414	DEVELOPER-1371	Has anyone tried it again? I responded on the same issue in the ticket tracker. Based on the hour it happened I suspect the ongoing reconfiguration of the network might have caused that.	1	1	https://issues.jboss.org/browse/DEVELOPER-1371
7415	DEVELOPER-1349	Can this be closed?	0	0	https://issues.jboss.org/browse/DEVELOPER-1349

#	Requirement	Questions	Answered	Successfully Answered	URL
7416	DEVELOPER-1348	[~paul.robinson] or [~pmuir], Should this be added to downloads.jboss.org or can I link directly to dropbox? Thanks! Dan	1	1	https://issues.jboss.org/browse/DEVELOPER-1348
7417	DEVELOPER-1348	Who can add the download to downloads.jboss.org since it isn't stored with download-manager (https://www.jboss.org/download-manager/ui/administration/products.jsf?renew=1)? Is Paul the only one? The current one should be added, which is on dropbox, PREVIEW-DataVirtWebUI-11Jun2015.war. Prior JIRAs for updating webui: https://issues.jboss.org/browse/DEVELOPER-1305 https://issues.jboss.org/browse/DEVELOPER-1224	1	1	https://issues.jboss.org/browse/DEVELOPER-1348
7418	DEVELOPER-1348	[~lightguard] Hey Jason, do you know who would be able to upload this to downloads.jboss.org?	1	1	https://issues.jboss.org/browse/DEVELOPER-1348
7419	DEVELOPER-1347	Thanks [~nickboldt]. Please can you email me when they are ready? I've got quite a lot of other things to coordinate over the coming days, so don't want to miss this. [~dcoughlin1] can you prepare the same change for http://developers.redhat.com please.	1	1	https://issues.jboss.org/browse/DEVELOPER-1347
7420	DEVELOPER-1340	Can this be closed?	1	1	https://issues.jboss.org/browse/DEVELOPER-1340
7421	DEVELOPER-1328	Pete - Since this was deployed, shall I go in and create a new macro for this new Data Layer Variable "product_code" within GTM? After that, I will update the macro {{Eloqua High Value URL}} with the updated URL values (with new macro) in the Lookup Table. After that, I will create a new version and have you review and Publish. Does this sound ok?	1	1	https://issues.jboss.org/browse/DEVELOPER-1328
7422	DEVELOPER-1324	[~paul.robinson] I know where to update it, but is there a good place to pull it. I just want to make sure that I have the correct profile information.	1	1	https://issues.jboss.org/browse/DEVELOPER-1324
7423	DEVELOPER-1324	Can this be merged and closed?	0	0	https://issues.jboss.org/browse/DEVELOPER-1324
7424	DEVELOPER-1316	[~paul.robinson] I have the PR ready, should I just hold onto it and rebase when requested?	1	1	https://issues.jboss.org/browse/DEVELOPER-1316
7425	DEVELOPER-1309	Do you have a web page that represents the Devoxx4Kids stuff done by your JUG? We can put your details in this area: https://www.dropbox.com/s/wgqi0jwzfvla244/Screenshot%202015-04-17%2009.14.33.png?dl=0 of http://www.jboss.org/devoxx4kids/. A link isn't essential (the Vancouver one doesn't have one). But if you have it then great!	1	1	https://issues.jboss.org/browse/DEVELOPER-1309
7426	DEVELOPER-1309	Sorry Paul, my bad. Our event was described here: http://rivierajug.org/xwiki/bin/view/Main/201411-devoxx4kids, we don't have a dedicated page for it yet but the whole team wants to do it on a regular basis so it won't be the last. I did not quite understand what you told me: how do I actually edit that page? I mean, I do see where it should go, but not how I can get the page's source.	0	0	https://issues.jboss.org/browse/DEVELOPER-1309
7427	DEVELOPER-1304	[~paul.robinson] That sounds good. Which banner should I replace on the main site?	1	1	https://issues.jboss.org/browse/DEVELOPER-1304
7428	DEVELOPER-1300	[~paul.robinson] We're already late on our release schedule (mea culpa), so is it something you can process ASAP? Do you have any ETA? Thanks for your help!	1	1	https://issues.jboss.org/browse/DEVELOPER-1300
7429	DEVELOPER-1295	Product: JBoss BRMS 6.1.0 Description: Minor release in the JBoss BRMS 6.x series Tag_Line: TBD Date of Release: 16-Apr-2015 CDN: (The set of files match what we currently upload for JBoss BRMS) /content/origin/files/sha256/2d/2db626b179f62a66911416ba173d1fd8c9c71c2f7285cbc45336c6a35bc673de/jboss-brms-6.1.0.GA-quickstarts.zip /content/origin/files/sha256/36/36d1f9ad97cec59168ff7812e71bed723cbf679f3e4086063fe7f19e73603564/jboss-brms-6.1.0.GA-sources.zip /content/origin/files/sha256/5c/5cd3c155110c5560ea26d7bfb6276e2d5f9504b54fa0494b3b2e5e4f6b80229c/jboss-brms-6.1.0.GA-deployable-eap6.x.zip /content/origin/files/sha256/85/85cee66d0dc3a50970b4b59c2230783fd31191d353a66c24d3c6fd82df201857/jboss-brms-6.1.0.GA-planner-engine.zip /content/origin/files/sha256/c1/c1c2fbbfdb95c2f74b71fe9b95b200d777cd956e7f244d206ee0406aaa4bbaf5/jboss-brms-6.1.0.GA-supplementary-tools.zip /content/origin/files/sha256/de/de5d17fb5ac679d175c1967b8375f48fbc7c6b427c20dd034937347e21a49501/jboss-brms-6.1.0.GA-installer.jar With this release, we no longer bundle EAP within the BRMS installer. To get started BRMS users will also need access to EAP 6.4 zip. Can this also be included in the downloads section? /content/origin/files/sha256/27/27a6fd62a8bc4f660970ab282a4bc013934275e47a850a974db6c7d2c62cc50e/jboss-eap-6.4.0.zip Release Notes Link: https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_BRMS/6.1/html-single/6.1.0_Release_Notes/index.html Date to go live: Now	1	1	https://issues.jboss.org/browse/DEVELOPER-1295
7430	DEVELOPER-1295	Product: JBoss BPM Suite 6.1.0 Description: Minor release in the JBoss BPM Suite 6.x series Tag_Line: TBD Date of Release: 16-Apr-2015 CDN: (The set of files match what we currently upload for JBoss BPM Suite) /content/origin/files/sha256/2d/2db626b179f62a66911416ba173d1fd8c9c71c2f7285cbc45336c6a35bc673de/jboss-bpmsuite-6.1.0.GA-quickstarts.zip /content/origin/files/sha256/36/36d1f9ad97cec59168ff7812e71bed723cbf679f3e4086063fe7f19e73603564/jboss-bpmsuite-6.1.0.GA-sources.zip /content/origin/files/sha256/85/85cee66d0dc3a50970b4b59c2230783fd31191d353a66c24d3c6fd82df201857/jboss-bpmsuite-6.1.0.GA-planner-engine.zip /content/origin/files/sha256/a5/a5aec379a6a9205e14660ffa42b3abab9d54e6a02d0b7f1f5b3118db4bf5f813/jboss-bpmsuite-6.1.0.GA-installer.jar /content/origin/files/sha256/c1/c1c2fbbfdb95c2f74b71fe9b95b200d777cd956e7f244d206ee0406aaa4bbaf5/jboss-bpmsuite-6.1.0.GA-supplementary-tools.zip /content/origin/files/sha256/e0/e03ac975ea18fb8c34e2c56de0509c78a2cf5294eb14fcdcfe6c0fc0264152f4/jboss-bpmsuite-6.1.0.GA-deployable-eap6.x.zip With this release, we no longer bundle EAP within the BRMS installer. To get started BRMS users will also need access to EAP 6.4 zip. Can this also be included in the downloads section? /content/origin/files/sha256/27/27a6fd62a8bc4f660970ab282a4bc013934275e47a850a974db6c7d2c62cc50e/jboss-eap-6.4.0.zip Release Notes Link: https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_BPM_Suite/6.1/html-single/6.1.0_Release_Notes/index.html Date to go live: Now	1	1	https://issues.jboss.org/browse/DEVELOPER-1295
7431	DEVELOPER-1295	Rajesh, Can you give a link to what tag line was used for 6.0 ?	0	0	https://issues.jboss.org/browse/DEVELOPER-1295
7432	DEVELOPER-1295	{quote} With this release, we no longer bundle EAP within the BRMS installer. To get started BRMS users will also need access to EAP 6.4 zip. Can this also be included in the downloads section? {quote} I think we need to update the get-started instructions to include this step. This is where we normally put installation instructions like this.	1	1	https://issues.jboss.org/browse/DEVELOPER-1295
7433	DEVELOPER-1295	Instructions to download EAP under get-started works as well. It would fit well under the "Installation Instructions" section. This would probably be step no.2. Could we link to the EAP 6.4 downloads page on jboss.org? Does the wording below look good? "JBoss BPM Suite from 6.1 release onwards does not bundle EAP inside its installer. Visit the EAP downloads page and obtain EAP version 6.4 before using the JBoss BPM Suite installer"	1	1	https://issues.jboss.org/browse/DEVELOPER-1295
7434	DEVELOPER-1295	Did we manage to get in that note about EAP ?	1	1	https://issues.jboss.org/browse/DEVELOPER-1295

#	Requirement	Questions	Answered	Successfully Answered	URL
7435	DEVELOPER-1295	Should "JBoss BPM Suite from 6.1 release onwards does not bundle EAP inside its installer. Visit the EAP downloads page and obtain EAP version 6.4 before using the JBoss BPM Suite installer" be posted under both BRMS and BPM Suite? I have it added to BPMSuite right now, but I'll wait to submit the PR in case you'd like it under both.	1	1	https://issues.jboss.org/browse/DEVELOPER-1295
7436	DEVELOPER-1295	[~rrajesh], I think having it as step 2 is confusing. The first step tells the user to download the installer. Many will then try and run the installer, which won't work as step 2 is not done yet. Also, I think it needs to be made clearer that EAP needs installing as well as downloading. How about: 1. JBoss BRMS from 6.1 release onwards does not bundle EAP inside its installer. Visit the EAP downloads page to download and then install EAP version 6.4 before using the JBoss BRMS installer. 2. Visit the JBoss BRMS download page and download the installer. 3. Make sure you have a JDK of your choice installed. We recommend OpenJDK or Oracle JDK. 4. You can now (optionally) install the Integration Stack for JBoss Developer Studio. This provides IDE support for developing JBoss BRMS applications. For more details, see the installation instructions.	1	1	https://issues.jboss.org/browse/DEVELOPER-1295
7437	DEVELOPER-1285	Do you know when you expect the CDN links to be available?	1	1	https://issues.jboss.org/browse/DEVELOPER-1285
7438	DEVELOPER-1273	[~kpeeples] Can you provide the full information that we need? There's a handy checklist here: https://github.com/jboss-developer/www.jboss.org/blob/master/CONTRIBUTING.md#how-to-request-a-new-product-download	1	1	https://issues.jboss.org/browse/DEVELOPER-1273
7439	DEVELOPER-1273	Are there source or maven repo links that need to be added as well, or should I just add the installer?	0	0	https://issues.jboss.org/browse/DEVELOPER-1273
7440	DEVELOPER-1273	I'd assume it's just a single download, as Kenny would have said otherwise. Can you proceed on this understanding. Kenny, please correct us ASAP if this isn't the case.	0	0	https://issues.jboss.org/browse/DEVELOPER-1273
7441	DEVELOPER-1271	Yes we can do this. Can you work with [~james.cobb], who's team will produce the graphics.	1	1	https://issues.jboss.org/browse/DEVELOPER-1271
7442	DEVELOPER-1271	[~philsimpson] Are we looking at two different releases here? BPM Suite 6.1 BRMS 6.1 How do we want to word this exactly?	1	1	https://issues.jboss.org/browse/DEVELOPER-1271
7443	DEVELOPER-1271	These look great! Couple of things - - Can we add JBoss to the product name? JBoss BPM Suite 6.1 - For the CTA button can we have "Learn More" instead of "Try It Now"? Since its not going directly to a download page. Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-1271
7444	DEVELOPER-1271	[~paul.robinson] Should I plan on adding this in place of another banner or in addition to the others? Also, what should I put for the "alt" on the image links?	1	1	https://issues.jboss.org/browse/DEVELOPER-1271
7445	DEVELOPER-1271	Looks like the release is delayed a couple days due to dependency on EAP 6.4. Can we push this banner to April 16? Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-1271
7446	DEVELOPER-1270	Which pages in particular do you want to change? http://www.jboss.org/products/bpmsuite/overview/ is based on: https://github.com/jboss-developer/www.jboss.org/blob/master/products/bpmsuite/overview.adoc. If you look carefully you can see where the various sections of the asciidoc file are inserted into the page layout. You could edit the adoc file maybe?	1	1	https://issues.jboss.org/browse/DEVELOPER-1270
7447	DEVELOPER-1269	[~paul.robinson] why is this still trying to do deployments? We released Beta1 on Mar 5, and CR1 last week. Surely this shouldn't be still staging old PRs?	1	1	https://issues.jboss.org/browse/DEVELOPER-1269
7448	DEVELOPER-1266	Looks like you've had a copy-and-paste catastrophe ;-) Can you provide the right link?	0	0	https://issues.jboss.org/browse/DEVELOPER-1266
7449	DEVELOPER-1266	Hi [~mfusco], It seems like the link you provided isn't working. Both safari and chrome say the server for the page isn't responding. Is there another link I can use? (I'll try it again later to see if it works) Adela	1	1	https://issues.jboss.org/browse/DEVELOPER-1266
7450	DEVELOPER-1266	Do you mean this url? http://www.ated.ch/corsi/8/devoxx4kids_3142.html Weird, I just tried it both Firefox and Chrome and it works for me. Please let me know. Thanks, Mario	1	1	https://issues.jboss.org/browse/DEVELOPER-1266
7451	DEVELOPER-1266	Hi [~mfusco] Sorry to bother you again. I am adding the event to the http://jboss.org/events/ page, but I need a description for the thumbnail. Can you send me one when you get a chance so I can add the event to this page as well? Thanks!	0	0	https://issues.jboss.org/browse/DEVELOPER-1266
7452	DEVELOPER-1266	How about you create a generic Devoxx4Kids image that we can use for all such events?	0	0	https://issues.jboss.org/browse/DEVELOPER-1266
7453	DEVELOPER-1258	In general, this seems sensible to me. I think we can work-around the expensive nature of doing this indexing, by doing it infrequently and outside of the main indexing. Do you need access to quickstart versions prior to those available today?	1	1	https://issues.jboss.org/browse/DEVELOPER-1258
7454	DEVELOPER-1258	Doesn't it make sense to address this as part of the issue we have to add versionning to all developer materials, thus allowing JBDS to select the materials by target product version?	1	1	https://issues.jboss.org/browse/DEVELOPER-1258
7455	DEVELOPER-1258	[~pmuir], [~tmancini], [~maxandersen], [~akazakov]: dear fearless leaders, what priority should we put on to making quickstarts/examples available *per* runtime version, in Devstudio (and the website)? As currently only the latest version is available. I'm asking because [~paul.robinson] and I need to do some sensible planning for implementing this feature.	1	1	https://issues.jboss.org/browse/DEVELOPER-1258
7456	DEVELOPER-1249	Can this be deleted?	0	0	https://issues.jboss.org/browse/DEVELOPER-1249
7457	DEVELOPER-1248	this is about right. As I mentioned in JBDS-3273, couldn't we get a quickstart_id collection instead? AFAIK, some quickstarts are meant to be used together (typically client & server). Also, how can we add more metadata for each client? I'm thinking JBoss Tools specific things, like eclipse plugin/server requirements	1	1	https://issues.jboss.org/browse/DEVELOPER-1248
7458	DEVELOPER-1248	Can you give me some specific examples of where you need the "quickstart_id collection"?	1	1	https://issues.jboss.org/browse/DEVELOPER-1248
7459	DEVELOPER-1248	{quote} Also, how can we add more metadata for each client? I'm thinking JBoss Tools specific things, like eclipse plugin/server requirements {quote} This would be harder. We'd need a deterministic way of picking that content out of the quickstart code. For example, in this file: https://github.com/jboss-developer/jboss-eap-quickstarts/blob/6.4.x-develop/bean-validation/README.md. For the other meta-data we look for required text patterns (e.g: 'Author: Karel Piwko').	1	1	https://issues.jboss.org/browse/DEVELOPER-1248
7460	DEVELOPER-1248	Daniel, how do you get these fields show up in the response to http://dcp.jboss.org/v1/rest/search?content_provider=jboss-developer&content_provider=rht&&field=github_repo_url&field=sys_title&field=sys_tags&query=sys_type:(jbossdeveloper_quickstart)&size=1 ?	1	1	https://issues.jboss.org/browse/DEVELOPER-1248
7461	DEVELOPER-1248	[~fbricon] what should the value of the 'download' attribute contain?	1	1	https://issues.jboss.org/browse/DEVELOPER-1248
7462	DEVELOPER-1248	This link downloads the entire repo containing the QS, not the individual QS. Is this what you want? E.g: https://github.com/jboss-developer/jboss-eap-quickstarts/archive/6.3.0.GA.zip BTW, there's no way, that I know of, to download just the quickstart dir from Github as a zip.	1	1	https://issues.jboss.org/browse/DEVELOPER-1248
7463	DEVELOPER-1248	[~paul.robinson] the git_commit is being pushed through for all of them that I can tell. I ran a straight print out to a local file and double-checked these items. It's peculiar that the first half of them don't seem to get the updates. Could it have to do with the rest endpoint not actually updating the database? Would Lukas know how to go about this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1248

#	Requirement	Questions	Answered	Successfully Answered	URL
7464	DEVELOPER-1248	Can you provide a reproduction steps for us to investigate?	1	1	https://issues.jboss.org/browse/DEVELOPER-1248
7465	DEVELOPER-1248	Can you provide an update on where you are with this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1248
7466	DEVELOPER-1246	[~paul.robinson] Currently, the /events page has calendar ranges that can be used and by default is showing only current/future events. Is this issue still relevant, or can I close it?	0	0	https://issues.jboss.org/browse/DEVELOPER-1246
7467	DEVELOPER-1242	When will this go into production?	1	1	https://issues.jboss.org/browse/DEVELOPER-1242
7468	DEVELOPER-1218	[~manaRH] We still need Ticket Monster updating. Can you create a separate issue for that?	1	1	https://issues.jboss.org/browse/DEVELOPER-1218
7469	DEVELOPER-1204	Can you assign to Adela, once the wireframes are approved.	0	0	https://issues.jboss.org/browse/DEVELOPER-1204
7470	DEVELOPER-1187	[~paul.robinson] So the Red Hat Event does show up, but the shadowman logo is not appearing next to it (https://www.dropbox.com/s/e9m8n4dlpkvj0p1/Events_featuring_JBoss.png?dl=0). We eventually want it to display like https://www.dropbox.com/s/vezzw1s4gahko1s/jbossdeveloper_getinvolved_events-1200px_png__1728%C3%972160_.png?dl=0, right? Should this be done by styling? The ruby block already differentiates between Red Hat and non-Red Hat events, the css doesn't show any references to the logos.	1	1	https://issues.jboss.org/browse/DEVELOPER-1187
7471	DEVELOPER-1181	Which dcp instance is this? I'm wondering if we have lingering quickstarts where we've changed an id or something. Chances are they aren't returned anymore when people search for them.	0	0	https://issues.jboss.org/browse/DEVELOPER-1181
7472	DEVELOPER-1180	https://github.com/jboss-developer/jboss-stacks/commit/02149f933af5d3b3df2ca40823875e1af6f99efc fixed it for me. Thanks a lot Rafael!	1	1	https://issues.jboss.org/browse/DEVELOPER-1180
7473	DEVELOPER-1170	Yes, it is possible. I would suggest to update relevant mapping configurations and store it into a new branch of Searchisko. I will be happy to review it and then we elaborate a plan of releasing it into production (and staging of course) and merge it. WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-1170
7474	DEVELOPER-1166	I think I see what you mean: 1. Visit: http://www.jboss.org/products/datavirt/get-started/ 2. Click on the link to 'Download JBoss Data Virtualization 6.0.0.GA Jar (Installer).' 3. This downloads jboss-dv-installer-6.0.0.GA-redhat-4.jar, yet the instructions tell you to unzip. Is this the issue you are referring to? If so, it's already happening in production, so not relevant to this particular issue. I've opened a new issue for it: https://issues.jboss.org/browse/DEVELOPER-1182 DEVELOPER-1182 aside, is this change good to merge?	0	0	https://issues.jboss.org/browse/DEVELOPER-1166
7475	DEVELOPER-1165	Where abouts would you like this text putting?	1	1	https://issues.jboss.org/browse/DEVELOPER-1165
7476	DEVELOPER-1165	Actually, we already have a feature in place to show this, although there are no examples live at the moment. This data normally goes in the blue early release box on the right-hand side of the product page. However, you don't get as much space as you would need for the text above. How about we put a summary in the blue box and then add the text above to the body of the overview page? I can't easily provide you with a mockup at the moment, so this is probably best holding until January. However, we are not pushing any more changes to the site until the 5th of January anyway so it makes little difference. [~dcoughlin1], if you understand what I'm getting at, and you have time, you could raise a PR with these changes: 1) Add the full text above to the body of the product overview page. Maybe straight after the 3 icon block: https://www.dropbox.com/s/8kmn8hdj9yaqes0/Screenshot%202014-12-19%2017.17.53.png?dl=0 2) Add 250 characters of "lorem ipsum" to the tag_line property on the 6.1 Beta download (once you've added that download, as requested in the another issue) in the DV product.yml. See here for an example: https://github.com/jboss-developer/www.jboss.org/blob/master/products/eap/_common/product.yml#L49. 250 characters might be too much, but it will give us an idea of what it looks like and how much is a sensible amount of text.	1	1	https://issues.jboss.org/browse/DEVELOPER-1165
7477	DEVELOPER-1165	Also, let us know if you want the full text (in the description) adding to the overview page. If you do want it, I have some queries about the text: 1) The text "Big Data technology including Cloudera Impala ..." seems a bit confusing to me. It sounds like DV includes these technologies. I think you mean something like "New Integration with Big Data technologies including ...". Can you clarify this? 2) This sentence is confusing me: "Improving Productivity and Deployment with a new Security audit log dashboard, simplified custom translator, EAP 6.3, RHEL 7, MariaDB and Azul JVM". I think these are separate new things: * Improving Productivity and Deployment with a new Security audit log dashboard * simplified custom translator * EAP 6.3 * RHEL 7 * MariaDB * Azul JVM In which case I understand how the first two items are new things in DV. However, it's not clear to me what you mean with the following four items. Is it support for these technologies that is new? If you do decide to have this text on the main body of the overview page, then you would have a bit more space available than you probably thought you had when producing the original text. Therefore you could make use of lists to clarify the text if you want.	1	1	https://issues.jboss.org/browse/DEVELOPER-1165
7478	DEVELOPER-1165	[~dcoughlin1] yes please. Can you add it add the end of the 'detail' section: https://github.com/jboss-developer/www.jboss.org/blob/master/products/datavirt/overview.adoc#detail so that it appears before resources.	1	1	https://issues.jboss.org/browse/DEVELOPER-1165
7479	DEVELOPER-1165	How about this: https://www.dropbox.com/s/vetwhfs5eeiz0sl/Screenshot%202015-01-15%2009.36.11.png?dl=0 Changes: * Added the title: "What's new in JBoss Data Virtualization 6.1 Beta?" * Moved the current title down to the beginning of the section * Collapsed the bullets into the opening sentence: "For JBoss Data Virtualization 6.1 Beta, we focused on Big Data, Cloud and other Productivity and deployment improvements.". [~kpalko] and [~kpeeples]: are you happy with these changes?	1	1	https://issues.jboss.org/browse/DEVELOPER-1165
7480	DEVELOPER-1165	This screen shot looks good. Will this text be displayed close to the "download" button so that users will know what they are getting when they download? thanks, Kim	1	1	https://issues.jboss.org/browse/DEVELOPER-1165
7481	DEVELOPER-1165	Hi Syed, Can you please help craft a short (100 - 250 char) text around What's New in JDV 6.1 for the beta download button? thanks, Kim	1	1	https://issues.jboss.org/browse/DEVELOPER-1165
7482	DEVELOPER-1165	[~srasheed], We've had a think about this and we think the best we can offer at this moment is to add the the ~250 chars description, that you provided above, to the right-hand-bar, with a "read more" link that points to the full list of new features. Can you create a customer portal article for the full feature list? This would be better than linking to the PDF. I appreciate that this might not be exactly what you were hoping for. In which case, can you open a new issue explaining your requirements in some detail. We'll do our best to get to it, but I can say that with our current workload it's unlikely to be in the next few weeks.	0	0	https://issues.jboss.org/browse/DEVELOPER-1165
7483	DEVELOPER-1159	This is in reference to the ../index reference that has been giving jboss.org issues with all of the quickstarts.	1	1	https://issues.jboss.org/browse/DEVELOPER-1159
7484	DEVELOPER-1154	Those links seem to have an erroneous trailing '0'. However, looking at http://www.jboss.org//products/fsw/overview/ and http://www.jboss.org//products/fsw/ I don't see the broken links. Can you double check? It might be fixed now.	1	1	https://issues.jboss.org/browse/DEVELOPER-1154
7485	DEVELOPER-1148	[~paul.robinson] Can I comment this out in Blinkr.yaml?	1	1	https://issues.jboss.org/browse/DEVELOPER-1148
7486	DEVELOPER-1148	I've just given permission to edit via your gmail address. Can you re-check?	1	1	https://issues.jboss.org/browse/DEVELOPER-1148
7487	DEVELOPER-1148	[~paul.robinson] None of these are on the spreadsheet anymore. Do I need to hit Lukas up for a reindex?	1	1	https://issues.jboss.org/browse/DEVELOPER-1148

#	Requirement	Questions	Answered	Successfully Answered	URL
7488	DEVELOPER-1144	{quote} modify the instructions for zip vs install {quote} In step 3, we tell the user how to install from a zip. Are you saying we should add instructions for how to install from an installer too? Ideally we would display different instructions based on what artefact is downloaded, but again we'd need DEVELOPER-712 to do this. In the short term we could add extra instructions for the installer. Is that what you would like?	1	1	https://issues.jboss.org/browse/DEVELOPER-1144
7489	DEVELOPER-1144	Does JBDS 8 and 7, with the IS plugin, work with both datavirt 6.0.0.GA and 6.1.0.ALPHA? If so, given that we can't yet customise the get-started page by download, we could just link to the JBDS 8 and IS instructions: http://www.jboss.org/products/devstudio/get-started/#ExtraSection. Could that work?	1	1	https://issues.jboss.org/browse/DEVELOPER-1144
7490	DEVELOPER-1139	[~paul.robinson] My plan is to comment this out until a new link is provided. Is that okay? If not, who should I contact to get the correct link?	1	1	https://issues.jboss.org/browse/DEVELOPER-1139
7491	DEVELOPER-1131	Are we wanting to use the README.md files for this or something else? The README file isn't quite up to par with the rest of our developer materials.	1	1	https://issues.jboss.org/browse/DEVELOPER-1131
7492	DEVELOPER-1103	{quote} JBoss Developer Studio provides an easy-to-use integration development environment built on Eclipse. The JBoss Developer Studio Integration Stack provides tooling for developing, configuring and deploying Fuse applications. {quote} Is it this text that you want moving? Also, I'm not sure what you mean by the "runtime download button". Can you explain or provide a screenshot?	1	1	https://issues.jboss.org/browse/DEVELOPER-1103
7493	DEVELOPER-1103	Aileen, the jbds page outlines that. Not much room if want some simple button added. Got a suggestion on how to fix it ? /max http://about.me/maxandersen	1	1	https://issues.jboss.org/browse/DEVELOPER-1103
7494	DEVELOPER-1103	Add this tex tunder the JBDS download button? The JBoss Developer Studio Integration Stack provides tooling for developing, configuring and deploying Fuse applications.	1	1	https://issues.jboss.org/browse/DEVELOPER-1103
7495	DEVELOPER-1103	Is it intentional that the 2nd download button is a different size and colour? I'm not a UX person, but would have kept them consistemt.	1	1	https://issues.jboss.org/browse/DEVELOPER-1103
7496	DEVELOPER-1075	[~sparulka] How important is it to change the text under "JBoss xPaaS services for OpenShift"? This is currently generic across all products, and the text you proposed wouldn't make sense for EAP. We can make this more customisable, but it would make maintenance harder and more error-prone.	1	1	https://issues.jboss.org/browse/DEVELOPER-1075
7497	DEVELOPER-1075	[~dcoughlin1] Can you work on this today please? Don't worry about the change to the download page, I can do that.	1	1	https://issues.jboss.org/browse/DEVELOPER-1075
7498	DEVELOPER-1075	[~paul.robinson] For the resources portion, this information is rendered from /_layouts/product-resources.html.slim. I can remove all of the entities in the product.yml and leave the 'See More Documentation' (https://access.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Fuse)? Is there a way to do this without changing the template .slim file? ~Dan	1	1	https://issues.jboss.org/browse/DEVELOPER-1075
7499	DEVELOPER-1063	[~paul.robinson] Would it be better if I address only JBoss Developer in this? For now, I am going to list each authors errors in separate Issues. Ex: "Vimeos with Duplicate/Truncated Descriptions - Eric Schabell" With each of the issues.	1	1	https://issues.jboss.org/browse/DEVELOPER-1063
7500	DEVELOPER-1057	Sameer, The page you are referring to provides results from a search. In this case you are searching for all developer materials that are related to the Fuse product. We don't currently provide customisable ordering, but we could in the future. However, it would more likely be the ability to order by recency, difficulty, duration etc. What ordering algorithm did you use to establish the above order?	1	1	https://issues.jboss.org/browse/DEVELOPER-1057
7501	DEVELOPER-1057	Ok, 1) is now live: http://www.jboss.org/products/fuse/get-started/ 2) We are the process of improving the entire developer-materials UX. In the process we will design how to do DEVELOPER-885. So it's in progress, but it will be a while. Do you have a urgent requirement for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-1057
7502	DEVELOPER-1057	(2) as soon as possible please..lot of pressure to improve the overall getting started experience for Fuse..DEVELOPER-1053 was very important and I am glad we have done that..however we still need to make the videos easily consumable.	1	1	https://issues.jboss.org/browse/DEVELOPER-1057
7503	DEVELOPER-1057	Given that we have alternatives, I think this current issue is a "won't fix". Could you create a new issue for the guides? If you can identify roughly what you want to appear in each guide, we can use it as input for the general design work around the guides infrastructure.	0	0	https://issues.jboss.org/browse/DEVELOPER-1057
7504	DEVELOPER-1056	Jason, Can you provide examples of what you think the RDF should look like for these items? I wasn't totally happy with containsProjects and containedWithinProjects, but couldn't think of anything better. I guess the problem is that it sounds like the codebase is contained within another project, rather than the other project simply uses the other project as a dependency and maybe ships with it in its binary. I think "consumes" and "consumedBy" is better, but feel like there's something clearer that is eluding us. [~pmuir], you are usually good at coming up with names, what do you think?	1	1	https://issues.jboss.org/browse/DEVELOPER-1056
7505	DEVELOPER-1053	After discussion with Sameer and Christina we would like to do the following. I attached the markdown file as well: Replace the 'Try a quickstart or example. Watch a video' with 'Build your first application' and then the content below which is the markdown I tested in a markdown editor: Follow the below series to build your first application. You will build some simple examples as well as build a Home Loan Application using JBoss Fuse to integrate different enterprise systems. The Home Loan Application accepts a loan request from the customer and provides a loan decision including a home mortgage rate. The series walks you through each step of the application development lifecycle from start to finish. |Demonstration Description | Material| |----------|---------| | What is JBoss Fuse? | Demo Source: https://github.com/kpeeples/jboss-fuse-websockets-demo Video 1: http://vimeo.com/user16928011/fusetting-started-part1 Video 2: http://vimeo.com/user16928011/fuse-getting-started-part2 Video 3: http://vimeo.com/user16928011/fuse-getting-started-part3 | | What is Apache Camel? What are Enterprise Integration Patterns (EIP)? | Demo Source: https://github.com/kpeeples/fuse-eip-quickstart | | Build the Home loan application using EIP's | Demo Source: https://github.com/weimeilin79/homeloan-part1 Video 1: https://vimeo.com/99901083 | | Business requirements change. Extend the home loan application. New SaaS applications? No problem.. | Demo Souirce: https://github.com/weimeilin79/homeloan-part2] Video 1: https://vimeo.com/100685377] Video 2: https://vimeo.com/100872412] | | Business expanding. Extend the home loan application to partners, suppliers. | Demo Source: https://github.com/weimeilin79/homeloan-part3 Video 1: https://vimeo.com/101265926 | | Innovate further. Try, experiment, test, deploy. | Video 1: https://vimeo.com/101266094 |	1	1	https://issues.jboss.org/browse/DEVELOPER-1053
7506	DEVELOPER-1053	When will this be live?	0	0	https://issues.jboss.org/browse/DEVELOPER-1053

#	Requirement	Questions	Answered	Successfully Answered	URL
7507	DEVELOPER-1047	Some more issues: * Tabbing though the items still (regularly) selects the 'required' text. Reproduce by trying to select the first text-box on the first page and then pressing tab a few times. Notice that it keeps on selecting the 'required' text field. * Hitting on the 'down arrow' on the 'team member role selector' creates a blue rectangle, rather than the drop-down. Clicking further to the left works as expected. This also happens if you click on a person's name in the right hand box. See here: https://www.dropbox.com/s/uxd6412dv1rave3/Screenshot%202014-10-20%2009.19.51.png?dl=0. This isn't a major issue, so just fix if it's simple. * Add this text to the 'Versions' page to make it clearer what we expect here: "Use this page to provide details of recent releases. There's no need to provide every single release for the project. For example, you could provide the latest GA release, the latest early-access release from an upcoming GA release and maybe the latest GA release that is available in a current product. Looking at the RDF produced for Narayana: https://github.com/jbosstm/project-metadata/blob/master/narayana.rdf * Should the 'Project Lead' be in a 'maintainer' block? * Despite adding two releases only one (the second one I added to the form) is present. * The description for the JTA specification is missing I temporarily added a Twitter account (Narayana has no social media). This is the result: {code} <foaf:account> <foaf:OnlineAccount> <rdf:type rdf:resource="http://xmlns.com/foaf/0.1/OnlineChatAccount" /> <foaf:accountServiceHomepage rdf:resource="http://www.twitter.com" /> <foaf:accountName>removeThisSection</foaf:accountName> </foaf:OnlineAccount> </foaf:account> {code} Should the "rdf:resource" be set to "http://xmlns.com/foaf/0.1/OnlineChatAccount"? I'm not sure how important this is.	1	1	https://issues.jboss.org/browse/DEVELOPER-1047
7508	DEVELOPER-1043	[~james.cobb] - do you know when you will be ready to pass this over to me?	1	1	https://issues.jboss.org/browse/DEVELOPER-1043
7509	DEVELOPER-1043	Yes, only some is available at the moment. Does it look OK with the limited info?	0	0	https://issues.jboss.org/browse/DEVELOPER-1043
7510	DEVELOPER-1040	Dan, can you create these as I still don't have access myself.	1	1	https://issues.jboss.org/browse/DEVELOPER-1040
7511	DEVELOPER-1012	I've noticed some of the requests made for contributors are being aborted because there are others going on. Don't we want all of those to go through? cc [~wesbos]	1	1	https://issues.jboss.org/browse/DEVELOPER-1012
7512	DEVELOPER-1012	Don't we basically have this done?	0	0	https://issues.jboss.org/browse/DEVELOPER-1012
7513	DEVELOPER-1012	I think so, [~pmuir] as the reporter, can you confirm?	1	1	https://issues.jboss.org/browse/DEVELOPER-1012
7514	DEVELOPER-1011	I don't think that makes any sense. What do we fallback to?	0	0	https://issues.jboss.org/browse/DEVELOPER-1011
7515	DEVELOPER-1009	[~paul.robinson] Apparently, I can't assign this to you. What are the next steps?	1	1	https://issues.jboss.org/browse/DEVELOPER-1009
7516	DEVELOPER-996	[~james.cobb] will the homepage banners always be either png or jpg, or could they be either in the future?	1	1	https://issues.jboss.org/browse/DEVELOPER-996
7517	DEVELOPER-994	[~pmuir] Where did we decide to host our own extension schema, jboss.org? [~paul.robinson] I'm a bit confused by your #4. What do you mean we don't allow projects to provide a logo?	1	1	https://issues.jboss.org/browse/DEVELOPER-994
7518	DEVELOPER-994	[~lightguard], [~pmuir]: We currently have James make an image at the right size, quality and dimensions and put it on static.jboss.org manually in a predictable location. This is not ideal, but it ensures we get good quality logos that fit correctly on the page. My comment was more a statement of what we do now, to open up discussion. Rather cryptic, I agree ;-) Anyway, if we keep the logo field, what is it used for? I see these options: 1) Grab the logo they suggest and do what we can with it automatically. Maybe we could validate the minimum dimensions in the DOAP form. 2) Have James use the logo as a suggestion. He can chose to use as is, modify it, or recreate entirely, then manually add it to static.jboss.org. I appreciate that 1) is more in-line with our goals around self-service, but I fear we will get a real mess of logos.	1	1	https://issues.jboss.org/browse/DEVELOPER-994
7519	DEVELOPER-994	Also, # Can the description field be made larger? # We should probably limit the number of characters for "Project Name", "Short Description" and "Description". [~james.cobb] do you have any limits imposed by your mockup? # What is "Browse URL"? # Make it clear which are optional and which mandatory. # What is "Community Account"? # Allow people to be removed, currently they can just be added. If you make a mistake, I think you have to start again with the entire form. # The logo URL should not be mandatory # User Forum and Developer forum # IRC/Chat room address # Why does the form need both a "Generate DOAP" and "Download DOAP" link? Isn't the latter enough?	1	1	https://issues.jboss.org/browse/DEVELOPER-994
7520	DEVELOPER-994	I have no character limits in place at this point. Do you have a suggestion for the max number of characters?	1	1	https://issues.jboss.org/browse/DEVELOPER-994
7521	DEVELOPER-994	Before I forget... Should the "release name" be mandatory?	1	1	https://issues.jboss.org/browse/DEVELOPER-994
7522	DEVELOPER-994	Sorry, we have had no requests for design on this page but we can. Do we want this to also have the Developer Look and Feel?	1	1	https://issues.jboss.org/browse/DEVELOPER-994
7523	DEVELOPER-994	Is this a project somewhere? Perhaps we can tweak the css a bit.	1	1	https://issues.jboss.org/browse/DEVELOPER-994
7524	DEVELOPER-994	[~lightguard], [~james.cobb]: Sorry for being unclear. I meant to ask: {quote} James Cobb: do you have any limits imposed by your *upstream projects page* mockup? {quote} The mockup of the upstream projects page has several areas of text. If we allow projects to provide too much text, it will not fit in those areas as currently designed. We can truncate, but better to have the project leads provide something that fits in the first place. I agree, we don't want any visual-design work done on this page as it is internal only. We might need UX input though as a lot of people are going to be filling this in without having one of us sat by their side explaining how to do it.	0	0	https://issues.jboss.org/browse/DEVELOPER-994
7525	DEVELOPER-979	[~paul.robinson] I'm not sure that I fully understand this. So I found the different README.mds that are generating the errors. Should I create separate issues per quickstart and list the respective broken links?	1	1	https://issues.jboss.org/browse/DEVELOPER-979
7526	DEVELOPER-979	How are you getting on with this?	1	1	https://issues.jboss.org/browse/DEVELOPER-979
7527	DEVELOPER-958	I notice you re-opened this 2 days ago. Is there something outstanding on this issue?	1	1	https://issues.jboss.org/browse/DEVELOPER-958
7528	DEVELOPER-958	Should this jira be assigned to [~lvlcek] also ?	1	1	https://issues.jboss.org/browse/DEVELOPER-958
7529	DEVELOPER-958	[~rafabene] Can you specify exactly which items to be removed? A DCP query is probably the best if you can construct one. [~lvlcek] does this make sense, or would you like the list in some other format?	1	1	https://issues.jboss.org/browse/DEVELOPER-958
7530	DEVELOPER-958	[~rafabene] the above link currently selects 29 documents in production DCP. Do you want me to remove all these documents from DCP (production and if available on staging then also on staging)?	1	1	https://issues.jboss.org/browse/DEVELOPER-958
7531	DEVELOPER-958	[~rafabene] I deleted 29 documents from production and 59 from staging. Shall we move the ticket to RESOLVED state now?	1	1	https://issues.jboss.org/browse/DEVELOPER-958
7532	DEVELOPER-956	[~lkrzyzanek] what about urls like these: https://community.jboss.org/people/tomjenkinson https://community.jboss.org/en/datagrid http://community.jboss.org/wiki/MavenGettingStarted-Users Will they be moved to developer.jboss.org? Will redirects be put in place?	1	1	https://issues.jboss.org/browse/DEVELOPER-956
7533	DEVELOPER-949	How about we set it manually and then document that this link needs updating whenever a new TiMo release is added to the site? The documentation would go as an extra step in https://github.com/jboss-developer/www.jboss.org/blob/master/CONTRIBUTING.md#updating-developer-materials-versions. I agree this is still error-prone, but it's simpler and doesn't preclude a more rigorous approach should it prove problematic in the future.	0	0	https://issues.jboss.org/browse/DEVELOPER-949

#	Requirement	Questions	Answered	Successfully Answered	URL
7534	DEVELOPER-935	Does it need to be Font-Awesome? What about this icon: http://theblogicon.com/	1	1	https://issues.jboss.org/browse/DEVELOPER-935
7535	DEVELOPER-933	[~pmuir] - can you create an experiment in Google analytics? The instructions are here: https://developers.google.com/analytics/solutions/experiments-client-side#configure We are doing a javascript swap-out so we don't need urls, just insert fake ones as it says.	1	1	https://issues.jboss.org/browse/DEVELOPER-933
7536	DEVELOPER-933	[~pmuir] did they give you an experiement code? Is it moRTWrZ2RgCt2XT5fJPqrQ?	1	1	https://issues.jboss.org/browse/DEVELOPER-933
7537	DEVELOPER-908	[~wesbos] AM i right in thinking that we now expire this data after a few mins/hrs? If so, do we do that on RHD, JBD or both?	1	1	https://issues.jboss.org/browse/DEVELOPER-908
7538	DEVELOPER-907	What do we want each of those to say, do we have approved text somewhere?	1	1	https://issues.jboss.org/browse/DEVELOPER-907
7539	DEVELOPER-893	Hi [~pmuir], Do you want the link in the technologies secondary navigation to take you to the current /research page or to the current research/centre page? I was a bit confused. * http://www.jboss.org/research/ * http://www.jboss.org/research/centre/	0	0	https://issues.jboss.org/browse/DEVELOPER-893
7540	DEVELOPER-864	Nick, We need the downloads configuring in the download manager, before we can put this on the site. Can you provide us with the CDN urls for each file download you added? We'll get the download-manager updated to handle them.	1	1	https://issues.jboss.org/browse/DEVELOPER-864
7541	DEVELOPER-864	[~paul.robinson], I don't see beta 3 in customer portal yet. Is there another place that I should look to get the SHA-256 numbers?	1	1	https://issues.jboss.org/browse/DEVELOPER-864
7542	DEVELOPER-864	[~nickboldt], [~dcoughlin1]: Thanks, I've merged it now. Did either of you check that the download file sizes are correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-864
7543	DEVELOPER-863	Here's my list of pages: {code} _site/products/datagrid/overview/index.html (wording change?) _site/products/operationsnetwork/community/index.html _site/products/operationsnetwork/developer-materials/index.html _site/products/operationsnetwork/download/index.html _site/products/operationsnetwork/get-started/index.html _site/products/operationsnetwork/overview/index.html _site/products/operationsnetwork/resources/index.html _site/quickstarts/operationsnetwork/index.html _site/quickstarts/operationsnetwork/rhq-client-cli/index.html _site/quickstarts/operationsnetwork/rhq-client-ejb/index.html _site/video/vimeo/16171314/index.html _site/video/vimeo/16182001/index.html _site/video/vimeo/19992280/index.html _site/video/vimeo/19992280/index.html (remove?) _site/video/vimeo/25855285/index.html (is this ON?) _site/video/vimeo/76471307/index.html {code} Let me know what to do with the ones I had questions about	1	1	https://issues.jboss.org/browse/DEVELOPER-863
7544	DEVELOPER-863	What about the quickstarts and videos, those are staying?	0	0	https://issues.jboss.org/browse/DEVELOPER-863
7545	DEVELOPER-848	Can we use this one? http://www.redhat.com/about/events-webinars/webinars/20140619-data-driven-business-decisions-made-easy	0	0	https://issues.jboss.org/browse/DEVELOPER-848
7546	DEVELOPER-819	So, I have been having issues finding the Narayana broken link. Can I use GateIn or another project as the test?	1	1	https://issues.jboss.org/browse/DEVELOPER-819
7547	DEVELOPER-819	Just so that I don't add the links incorrectly into blinkr.yaml, should it look like this? https://www.dropbox.com/s/isz8z2x1pxf1hjc/blinkr_regex.jpg	1	1	https://issues.jboss.org/browse/DEVELOPER-819
7548	DEVELOPER-819	[~dcoughlin1] I'm not sure that ignore is correct. I would have thought you could list the URL verbatim, without having to use a regex or escape sequence. See: https://github.com/pmuir/blinkr/blob/master/README.md#usage. I would do: {code} - url: https://github,com/jbosstm #BLACKTIE-449 message: Couldn't resolve host name {code} @pmuir is this correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-819
7549	DEVELOPER-819	Can I continue creating JIRA's in the respective projects? I am just checking to make sure that no changes need to be made.	1	1	https://issues.jboss.org/browse/DEVELOPER-819
7550	DEVELOPER-807	Sounds good. I have been searching for the actual training demonstrations on Google. It looks like they have been changed over to Eric Schabell's .org page.	1	1	https://issues.jboss.org/browse/DEVELOPER-807
7551	DEVELOPER-807	[~paul.robinson] I don't see any links to howtojboss within the project anymore, it looks like new links were added. Should this be closed?	1	1	https://issues.jboss.org/browse/DEVELOPER-807
7552	DEVELOPER-807	[~paul.robinson] I'm not sure if I should've created a new issue for this DEVELOPER-1117; however, it looks like the domain for howtojboss.com has expired. Should Eric Schabell be looped in on this?	1	1	https://issues.jboss.org/browse/DEVELOPER-807
7553	DEVELOPER-806	[~dcoughlin1]: You need to fix the filename to point to the bpmsuite download. This is already done in brms/product.yml; see this line "filename: jboss-bpmsuite-6.0.1.GA-source.zip". Can you raise a PR , to change the links for each dup. It's a known issue that clicking on the link redirects to the wrong getting-started page.	1	1	https://issues.jboss.org/browse/DEVELOPER-806
7554	DEVELOPER-805	[~c.wong] what would you like us to do with that link? remove it, fix it or replace it?	1	1	https://issues.jboss.org/browse/DEVELOPER-805
7555	DEVELOPER-804	[~dcoughlin1] can you remove these links. [~sparulka] if you decide to add some new ones, after this is resolved, can you create a new issue.	1	1	https://issues.jboss.org/browse/DEVELOPER-804
7556	DEVELOPER-791	I don't see the link at the top of the page. Can you provide a screenshot of the link?	0	0	https://issues.jboss.org/browse/DEVELOPER-791
7557	DEVELOPER-790	[~lightguard]: Sure, I appreciate that we can't get a release of Awestruct just for our stuff. Do you know roughly when the next release is due? I'm thinking it might be safer to wait for that. The current build time of 13m is still pretty good.	1	1	https://issues.jboss.org/browse/DEVELOPER-790
7558	DEVELOPER-782	[~lightguard] We need to be able to identify which projects appear in the featured area. I don't think it makes sense to put this data in Searchisko as the project data is shared, so we shouldn't be adding stuff specific to JBoss Developer in there. How about we make a .yml file that lists the three featured projects? The featured projects are of the page is then generated at site build time, with the "All Projects" area continuing to be populated with JavaScript?	1	1	https://issues.jboss.org/browse/DEVELOPER-782
7559	DEVELOPER-782	[~paul.robinson] it was [~james.cobb] - I think they were 5 max, we can do three no problem. Which three should we feature right now?	1	1	https://issues.jboss.org/browse/DEVELOPER-782
7560	DEVELOPER-782	[~james.cobb] - we are going to use three featured projects - Narayana, Weld and DeltaSpike. Can you gives me the images for those? Or ideally host them? Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-782
7561	DEVELOPER-781	Do I need access to dcpbeta-searchisko to add the detail for this webinar? I see that the JSON is coming from here, but I'm not sure how to add the content information.	1	1	https://issues.jboss.org/browse/DEVELOPER-781
7562	DEVELOPER-738	[~maxandersen] My problems with a single overview page were: 1) Where do we weave it into the current navigation? 2) Is it really what a user wants to see? Seems like information overload to me. My preference would be to slice it by product and display the appropriate slice on each product page.	1	1	https://issues.jboss.org/browse/DEVELOPER-738
7563	DEVELOPER-738	[~maxandersen] do you mean https://developer.jboss.org/maven or https://www.jboss.org/maven? Did you have a particular system in-mind to host this page? Either way, I think we need to revisit how we are handling Boms and Archetypes on the developer portals (www.jboss.org and developers.redhat.com). I also take your point that a single page showing everything would be useful as well as the "product view" on the product pages.	1	1	https://issues.jboss.org/browse/DEVELOPER-738
7564	DEVELOPER-733	[~petemuir] [~pmuir] Can you please push this to production ASAP? Thanks!	1	1	https://issues.jboss.org/browse/DEVELOPER-733

#	Requirement	Questions	Answered	Successfully Answered	URL
7565	DEVELOPER-717	Yes, I suppose we could further define what "Stand-alone" means, but when you juxtapose it next to "With EAP" I would like to think that most users would see "ah, it's the same thing but one has EAP and one does not". Explaining the update sites would take more than a couple words... if you've never used JBoss Central within JBDS, a picture is much more helpful. http://tools.jboss.org/documentation/whatsnew/jbosstools/4.2.0.Beta2.html#ui-changes Maybe we could change that column into optional links to a walkthru install guide [1] / more info [2] / getting started guide [3]? URL in the products.yml could vary but the link label could just be "More Info" or "Documentation" ... maybe [~mmurray] or [~maxandersen] have an opinion on how to best use that space? [1] http://tools.jboss.org/documentation/videos/jbds5_installation.html (needs updating to JBDS 8) [2] http://tools.jboss.org/documentation/index.html or https://access.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Developer_Studio/8.0/ [3] https://access.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Developer_Studio/7.1/html/Getting_Started_Guide/index.html (needs a new version for JBDS 8)	0	0	https://issues.jboss.org/browse/DEVELOPER-717
7566	DEVELOPER-701	We've gone back and forth on this topic many times it feels like. I never remember the final result. At one point we discussed no longer showing these - as they should exist on MeetUp and we would just pull in the events. I think Mark N. advocated for still having some information about the locations of user groups on JBoss Developer, which I could see. If we want something on JBoss Developer, can we talk as a group about what the requirements are - what information do we think JBoss Developer should own in this space? Then it can get assigned to me.	1	1	https://issues.jboss.org/browse/DEVELOPER-701
7567	DEVELOPER-683	[~paul.robinson] new the data has been re-indexed as part of upgrade to 1.1.0. Can you check please? Resolving...	1	1	https://issues.jboss.org/browse/DEVELOPER-683
7568	DEVELOPER-677	I'm not saying stay with the status quo - I'm saying we should add blog title, but not get rid of the blurb. Part of writing good blog posts should be writing the meta-description that search engines use (and we're pretty much acting as a search engine here). http://blog.hubspot.com/marketing/how-to-write-blog-post-simple-formula-ht (meta description section). http://www.socialmediaexaminer.com/26-tips-for-writing-great-blog-posts/ (#4 - description). Now - I can't disagree with the fact that a lot of people don't do this. We can send out of "how to write good blog posts" snippet to the core. But even if they still don't write good descriptions - are we losing anything? If we drop down to just title & name - it means that you don't get any extra useful information even for those bloggers who take the time to do this. If we leave the blurb, it means that some of them are not that useful if people don't use them right (and maybe having useless blurbs shown will help them start to do this right!), but you still have the title/author to look at as we would change all blog posts to show the title that we are not today.	0	0	https://issues.jboss.org/browse/DEVELOPER-677
7569	DEVELOPER-673	Would DEVELOPER-1082 solve this issue?	0	0	https://issues.jboss.org/browse/DEVELOPER-673
7570	DEVELOPER-672	The Vimeo player is HTML5 in browsers that support it (around 65%), and automatically degrades to flash in browsers that don't support it. I suspect that [~ssiddique]'s customers are using old versions of Internet Explorer that doesn't support HTML5 anyway. Vimeo offers the option to download the file using MP4 (270p, 380p, 720p) and MOV (720p), if you visit the vimeo page. We need to understand more from [~ssiddique] about what in particular is locked down in this environment - is it old browsers? Is it the ability to install Flash? Is it that particular websites (inc vimeo) are banned?	0	0	https://issues.jboss.org/browse/DEVELOPER-672
7571	DEVELOPER-668	Which JS are we talking about? I'm a bit confused there. Would this be displayed when they download something?	1	1	https://issues.jboss.org/browse/DEVELOPER-668
7572	DEVELOPER-668	I'm seeing js that references thank you, maybe it's the same as the T&C, can anyone verify?	1	1	https://issues.jboss.org/browse/DEVELOPER-668
7573	DEVELOPER-668	I've read these three or four times now and I'm still not 100% sure what needs to be done here, sorry (maybe because the title and description no longer have much to do with the current ideas in the comments). We have a javascript call that does something for thankyou after we download, and we have a REST call to determine the status of the T&C for the logged in user. Are we saying we need to check the status, display the T&C if they haven't agreed, log that somewhere then proceed to the download?	1	1	https://issues.jboss.org/browse/DEVELOPER-668
7574	DEVELOPER-668	My understanding is that we need to add a banner of some-sort to here: http://www.jboss.org/terms-and-conditions/. The banner would identify if the user has already signed the $0 subscription. This information is discovered by a REST call, to Libor's service. I think Libor's service requires you to specify a download file to see if the user has agreed to the licence on that file. I guess we could pick a file with the $0 licence, but that seems a bit brittle as that file's license could change or it could be removed. Am I missing something here?	1	1	https://issues.jboss.org/browse/DEVELOPER-668
7575	DEVELOPER-668	What we need is: 1) User logs in 2) User visits /terms-and-conditions in their web browser by typing in the address 3) If user has signed the t&c show, using a banner, that they have signed the t&c, and on what date. Specifically, Libor's service is one you can query based on the logged in user. David's service requires you to do a download to get thie info.	1	1	https://issues.jboss.org/browse/DEVELOPER-668
7576	DEVELOPER-668	[~lightguard] Make sure you login to the right SSO server. I got this to work: 1) Visit https://community-stg.jboss.org 2) Login 3) Visit https://community-stg.jboss.org/api/custom/v1/account/info However, I have not signed the agreement in the community-stg environment. [~lkrzyzanek] how would I do that?	1	1	https://issues.jboss.org/browse/DEVELOPER-668
7577	DEVELOPER-668	[~lkrzyzanek], does this endpoint support CORS? Looks like I'm blocked by not actually being able to make the call to the service.	1	1	https://issues.jboss.org/browse/DEVELOPER-668
7578	DEVELOPER-668	My understanding is that this page is on Magnola: http://www.jboss.org/developer-program/termsandconditions. [~lkrzyzanek] is this correct? [~lightguard] is there some other T&C page that you were thinking about and if so, where is it linked from?	1	1	https://issues.jboss.org/browse/DEVELOPER-668
7579	DEVELOPER-664	[~petemuir] ok, it is logo :-) If I am not mistaken then all is needed is something like: {code} var template = "<li class=\"upstream\">" + "<div class=\"defaultprojectimage\">" + "<img onerror=\"this.style.display='none'\" src=" + "\"http://static.jboss.org/" + (props.specialIcon || props.sys_project) + "/images/" + (props.specialIcon || props.sys_project) + "_200x150.png\"> " + "<h3 class=\"solution-name\">" + props.projectName + "</h3> "; {code} [here|https://github.com/jboss-developer/www.jboss.org/blob/master/javascripts/projects.js#L76]. Of course this is hack-ish but illustrates the idea. WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-664
7580	DEVELOPER-664	so specialIcon essentially overrides the project name in this case? Wouldn't it just be better to put the correct icon in the correct place on static?	1	1	https://issues.jboss.org/browse/DEVELOPER-664
7581	DEVELOPER-664	{quote} Wouldn't it just be better to put the correct icon in the correct place on static? {quote} From MPOW yes.	1	1	https://issues.jboss.org/browse/DEVELOPER-664

#	Requirement	Questions	Answered	Successfully Answered	URL
7582	DEVELOPER-664	[~dhladky] [~lvlcek] we don't seem to have specialIcon field available in http://dcpbeta-searchisko.rhcloud.com/v1/rest/search?sys_type=project_info&field=_source&query=NOT+archived%3Atrue+AND+hibernate&size=500 at all. Do you know why?	1	1	https://issues.jboss.org/browse/DEVELOPER-664
7583	DEVELOPER-664	BTW, the thing with {{project_info}} data is that we do not have fully automated way of updating it right now. The way how the data gets from the source (like property files on github and Magnolia) to DCP is quite complicated right now and some scenarios need careful manual interaction: - like renaming project or project keys - what should happen if one project is removed from the sources? We probably should not remove it from DCP automatically - etc ... I do not like it but can not think of quick fix for now. It can be quite a pain to keep all DCP instances in sync due to this as well. This all being said we would appreciate if you can report us all changes to {{project_info}} data that you are aware of. This can help us in keeping the data in shape. Hopefully we will have better and more transparent way of getting this data to DCPs.	0	0	https://issues.jboss.org/browse/DEVELOPER-664
7584	DEVELOPER-655	How can I filter by tags via the DCP? Seems this url (a regular buzz url + sys_tags on the end), doesn't filter for that tag http://dcpbeta-searchisko.rhcloud.com/v1/rest/search?field=sys_url_view&field=sys_title&field=sys_contributors&field=sys_description&field=sys_updated&field=author&field=sys_tags&query=&size=8&sys_type=blogpost&sortBy=new-create&sys_tags=EAP	0	0	https://issues.jboss.org/browse/DEVELOPER-655
7585	DEVELOPER-635	Hi Pete: Attached is image file. I am still waiting for URL for whitepaper from web team to upload to resource center. Is there a place where I can review the new page? Thanks you so very much. Regards, Syed Rasheed Ph: 978.589.1032(o) | 206-696-6266(m) Business Innovation Powered By Red Hat JBoss Middleware	1	1	https://issues.jboss.org/browse/DEVELOPER-635
7586	DEVELOPER-632	{quote} (6) Under the downloads page https://www.jboss.org/products/fsw/download/ Unde the 'Prefer to use Upstream product version? Add Apache Camel (6) Under community https://www.jboss.org/products/fsw/community/ Include links to following projects (in addition to the existing projects): Apache Camel, Apache ActiveMQ, Apache CXF {quote} Requested [~rdavies] to submit the data so we can link to it.	0	0	https://issues.jboss.org/browse/DEVELOPER-632
7587	DEVELOPER-628	Is this a duplicate of DEVELOPER-628? Or maybe it should be rolled into that issue.	1	1	https://issues.jboss.org/browse/DEVELOPER-628
7588	DEVELOPER-626	Do the table cells have to be so narrow as to cause line wrapping? Why not make *Version* and *Release Date* narrower and use the recovered horizontal space to not have to wrap "Stand- alone)" ? In fact if you drop the *Description* column entirely, since it's redundant to the *Download* column, you will likely have enough horizontal space to not have to line-wrap anything in the *Download* column. Is that doable?	1	1	https://issues.jboss.org/browse/DEVELOPER-626
7589	DEVELOPER-620	[~ray.ploski] what do you mean by "Out of sync"? Can you provide some specific examples?	0	0	https://issues.jboss.org/browse/DEVELOPER-620
7590	DEVELOPER-605	I am confused on this. The fusebyexample link redirects to http://www.jboss.org/products/amq/resources/. What information should I be adding?	1	1	https://issues.jboss.org/browse/DEVELOPER-605
7591	DEVELOPER-587	We have it in the address bar, but that's probably not intuitive for people. Should we add this to the metadata section?	0	0	https://issues.jboss.org/browse/DEVELOPER-587
7592	DEVELOPER-565	Is this relevant any longer?	0	0	https://issues.jboss.org/browse/DEVELOPER-565
7593	DEVELOPER-565	[~dcoughlin1] What did you have in mind for this issue?	1	1	https://issues.jboss.org/browse/DEVELOPER-565
7594	DEVELOPER-565	[~ray.ploski] I know you wanted me to add this a while back. Did you want documentation on how to add new releases to the products pages?	0	0	https://issues.jboss.org/browse/DEVELOPER-565
7595	DEVELOPER-552	[~kpeeples] Are you providing Fuse docker stuff? If not, do you know who is?	1	1	https://issues.jboss.org/browse/DEVELOPER-552
7596	DEVELOPER-549	[~kpeeples] are you providing A-MQ docker stuff? If not, do you know who is?	1	1	https://issues.jboss.org/browse/DEVELOPER-549
7597	DEVELOPER-517	Can you fix this [~wesbos] please?	1	1	https://issues.jboss.org/browse/DEVELOPER-517
7598	DEVELOPER-505	[~crobson] This is on the site now, but we still need it in the download manager (http://www-beta.jboss.org/download-manager) - as [~paul.robinson] is away, can you add it there? Assign back to me when done.	1	1	https://issues.jboss.org/browse/DEVELOPER-505
7599	DEVELOPER-505	[~james.cobb] - Can you take a look at the beta info in the right rail? It appears we really could use some sort of grouping mechanism in this case so the description & the downloads look like a single message in the right rail, not two separate entities.	1	1	https://issues.jboss.org/browse/DEVELOPER-505
7600	DEVELOPER-491	I think most of the code is there to do this now (if we have the info on their profiles of course). However, we don't have a standard way to display this. [~crobson] what do you think?	1	1	https://issues.jboss.org/browse/DEVELOPER-491
7601	DEVELOPER-491	I put together a quick mock-up of how I'm thinking this should work from an interactive standpoint. I think this works for the quickstarts, videos, etc where you have authors, and the get involved page quite well. We might need to consider how this would apply to to technology overview page, videos on the product overview pages, etc. where we have more dedicated space today to show the user information. [~james.cobb] - can you take this and consider the visual styling?	1	1	https://issues.jboss.org/browse/DEVELOPER-491
7602	DEVELOPER-491	Will do some design work on this. [~wesbos] - Would this be a fixed overlay box size? If so, we'll need to make it large enough to handle long names. [~lkrzyzanek] - Libor, is there any way to find out our longest username?	1	1	https://issues.jboss.org/browse/DEVELOPER-491
7603	DEVELOPER-468	[~crobson] [~wesbos] - I reviewed the changes in hotgloo and I think you're right and we should be able to incorporate those elements we already have on other pages. Does that work for you Wes or do you want a specific PSD of this?	1	1	https://issues.jboss.org/browse/DEVELOPER-468
7604	DEVELOPER-465	[~crobson] [~james.cobb]: Can you review?	1	1	https://issues.jboss.org/browse/DEVELOPER-465
7605	DEVELOPER-462	Hrm, i see we are specifying _source as the field, does that mean it searches all the subfields? I was getting partials back when doing the searchbar	0	0	https://issues.jboss.org/browse/DEVELOPER-462
7606	DEVELOPER-462	Cool, [~lvlcek] or [~velias], is there something we need to do in our searching to make that work without always doing a wildcard search behind the scenes, or is that just the way need to make this work?	1	1	https://issues.jboss.org/browse/DEVELOPER-462
7607	DEVELOPER-462	I think we want to wildcard either side of the query string. So a search for "cloak" should return "KeyCloak". Correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-462
7608	DEVELOPER-462	Don't think this feature is now used on the preview site built but I would like to point out regarding the second API (project name suggestions api) it can also return something like "did you mean" if the user query is not an exact match of project name substring. For example: Search: fie Results: Infinispan, Blacktie, JBoss Profiler However, we should check that the project name suggestion API does not return projects that we do not want to list on upstream projects page (for example archived projects?). I do not think this API now takes any restrictions into account, IIRC it works with all available project info data equally. If any restrictions are needed we would need to extend this API.	1	1	https://issues.jboss.org/browse/DEVELOPER-462
7609	DEVELOPER-462	[~paul.robinson] in that case we need to open a new ticket for project name suggestion API. Currently it does not allow to filter baed on {{archived}} field (or any other field). How critical is this?	1	1	https://issues.jboss.org/browse/DEVELOPER-462

#	Requirement	Questions	Answered	Successfully Answered	URL
7610	DEVELOPER-462	[~lvlcek] We seem good for now. Although I am confused. This approach: http://www-stg.jboss.org/pr/392/build/397/projects/ using project name suggestions api, doesn't return archived projects. Try searching for Narayana and notice that only one result comes back This approach: http://www-stg.jboss.org/pr/392/build/396/projects/ using wildcard on both sides of the keyword search, does return archived projects. Try searching for Narayana and notice that two result comes back (one is the archived jbossts project). Here it seems the suggestion API is obeying the archived filter. Correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-462
7611	DEVELOPER-443	Did DEVELOPER-412 fix this?	1	1	https://issues.jboss.org/browse/DEVELOPER-443
7612	DEVELOPER-443	I can't reproduce this issue. 1) If I control-click a link, it just opens a context sensitive menu (see http://guides.macrumors.com/control-click). I'm assuming you are actually saying it's broken when you try to open a link in a new tab 2) If I select open in a new tab from the context senstive menu, it opens a new tab with the link in it in Chrome, Safari and Firefox. Can you let me know how to reproduce properly?	1	1	https://issues.jboss.org/browse/DEVELOPER-443
7613	DEVELOPER-437	Not a fan of having the callout box floating up above the line for the page content title (we don't do it anywhere else). How about we put them side by side like the modified screen grab?	0	0	https://issues.jboss.org/browse/DEVELOPER-437
7614	DEVELOPER-414	I have added the following downloads to the beta download manager: EAP 6.3.0 Alpha A-MQ 6.1.0 GA Fuse 6.1.0 GA JBDS 8.0.0 Beta (no T&C's) JBDS 8.0.0 Beta (with EAP - Zero dollar sub/T&C's needed)	1	1	https://issues.jboss.org/browse/DEVELOPER-414
7615	DEVELOPER-414	Dan, can you add these new downloads to the site and create appropriate short urls in here: http://www-beta.jboss.org/download-manager. I don't think it's changed much, if at all, since you last did it. But if you are unsure, you could do one first and have me review it.	1	1	https://issues.jboss.org/browse/DEVELOPER-414
7616	DEVELOPER-414	[~crobson] what should the download page look like, when a Alpha or Beta is available? Does the Alpha or Beta appear on that page? Is it before or after the latest GA?	1	1	https://issues.jboss.org/browse/DEVELOPER-414
7617	DEVELOPER-414	[~petemuir] - Will the above approach work since the right rail will go away when the alpha/beta period is over - or will this not work because we need to keep those downloads available to people after the period is over? Not sure based on discussions we've had about alpha/beta/promoted build availability what the right direction is here!	0	0	https://issues.jboss.org/browse/DEVELOPER-414
7618	DEVELOPER-414	Are alphas and betas removed or kept once the GA is out?	1	1	https://issues.jboss.org/browse/DEVELOPER-414
7619	DEVELOPER-414	Hey Paul, Is this what you are looking for on the website portion? https://github.com/Dantheman720/www.jboss.org/commit/7d9be44e17fdbbd06d0930c306201b52e56deb43	1	1	https://issues.jboss.org/browse/DEVELOPER-414
7620	DEVELOPER-414	[~dcoughlin1], this looks right to me. Have you checked the file sizes are right?	1	1	https://issues.jboss.org/browse/DEVELOPER-414
7621	DEVELOPER-414	[~dcoughlin1], i've merged the PR and updated www-stg.jboss.org. Can you go make the links work? BTW, it's a known issue that some of the downloads appear out of order.	1	1	https://issues.jboss.org/browse/DEVELOPER-414
7622	DEVELOPER-414	[~dcoughlin1] are you adding the alphas and betas? Can you give me an ETA?	1	1	https://issues.jboss.org/browse/DEVELOPER-414
7623	DEVELOPER-408	Pete, what do we need to achieve in order to consider this issue resolved?	1	1	https://issues.jboss.org/browse/DEVELOPER-408
7624	DEVELOPER-372	[~crobson], with respect to pt. 3, what would you recommend as the optimum number, both for mobiles and desktops ? I'll try and adapt the items in the page to suit both.	0	0	https://issues.jboss.org/browse/DEVELOPER-372
7625	DEVELOPER-360	[~lightguard], [~pmuir]: Do we now have target_product set for all types now?	1	1	https://issues.jboss.org/browse/DEVELOPER-360
7626	DEVELOPER-357	[~pmuir] I nearly finished the mapper. But to use it we have to upgrade content type configurations in jboss-developer.json But to do this I need to know structure of data you push there to be able to correctly update configurations. Do you plan to use it for {{jbossdeveloper_quickstart}} type only as we talked over IRC, or for any other type also? Send me structure of input data please.	1	1	https://issues.jboss.org/browse/DEVELOPER-357
7627	DEVELOPER-297	[~wesbos] - Right now we don't swap to the hamburger icon at 1024, so tablets won't be able to use dropdown navigation well if we switch to this. Is there a way to recognize the device and use the hamburger for tablets and regular nav for PCs?	1	1	https://issues.jboss.org/browse/DEVELOPER-297
7628	DEVELOPER-281	[~lightguard] If you are running low on things to do, can you take this?	1	1	https://issues.jboss.org/browse/DEVELOPER-281
7629	DEVELOPER-281	[~crobson] where are we placing the stars for people to rate an item? I just went through the wireframes and I don't see a location for people to rate things. I see where we're displaying the average rating, but I'm not seeing where we expect people to rate them.	1	1	https://issues.jboss.org/browse/DEVELOPER-281
7630	DEVELOPER-281	[~lightguard] {{'Access-Control-Allow-Origin'}} header is added to the response from DCP when request contains {{'Origin'}} header. Check it is present there. It is also necessary to set {{withCredentials: true}} into XMLHttpRequest. I tested my tutorial code which uses http://dcp-stg.jboss.org just now, and it works as expected in both Firefox and Chrome. Which browser do you use?	1	1	https://issues.jboss.org/browse/DEVELOPER-281
7631	DEVELOPER-273	[~ray.ploski] Also, is it possible to get higher-res profile images from these guys?	1	1	https://issues.jboss.org/browse/DEVELOPER-273
7632	DEVELOPER-266	Is this issue still relevant?	0	0	https://issues.jboss.org/browse/DEVELOPER-266
7633	DEVELOPER-262	cool - what is the actual rss.url ? would like to check if jbds "just works" with it or need some yahoo pipes filtering in worst case ;)	0	0	https://issues.jboss.org/browse/DEVELOPER-262
7634	DEVELOPER-258	Where do I get the credentials from, and how do I set them in my environment? Once that is done, do I just do this: {code} rake clean preview[staging] {code} I think it would be useful to have something documented. I'm not suggesting you paste in the credentials, but knowing what env vars to set and who to ask for the values would be handy.	0	0	https://issues.jboss.org/browse/DEVELOPER-258
7635	DEVELOPER-252	Should this issue also cover having the buzz you generate be exposed as an RSS feed so we can show the same in jbds central ?	1	1	https://issues.jboss.org/browse/DEVELOPER-252
7636	DEVELOPER-244	[~pmuir] I look at Developer Materials on www-beta.jboss.org and I see no pagination, it seems to be neverending. Am I looking at the right place?	1	1	https://issues.jboss.org/browse/DEVELOPER-244
7637	DEVELOPER-243	How's this coming along [~vineet.reynolds]?	1	1	https://issues.jboss.org/browse/DEVELOPER-243
7638	DEVELOPER-238	Is there an example of what the correct URL string is supposed to be? For instance, they all carry this portion: "http://www.jboss.org/download-manager/f/", as in "http://www.jboss.org/download-manager/f/fuse/6.0.0.GA/(Full%20Zip)". I can find these projects at http://www.jboss.org/projects but is there some sort of download-manager component, as the URL indicates, that's not currently available?	0	0	https://issues.jboss.org/browse/DEVELOPER-238
7639	DEVELOPER-238	Is there by chance already login redirect functionality in the development site? And can it be incorporated into the same popup feature that was just put in for the event buttons?	0	0	https://issues.jboss.org/browse/DEVELOPER-238
7640	DEVELOPER-236	[~crobson] what type of popup should we use here?	1	1	https://issues.jboss.org/browse/DEVELOPER-236
7641	DEVELOPER-233	[~james.cobb] I noticed: https://github.com/jboss-developer/jboss-developer-site/commit/f348c98ecca2813a09bd51e4f7a9c84444c1939c. Can Cheyenne also wire up this page to the links on the videos in http://www-stg.jboss.org/products/eap/resources/, or do we need to do that bit?	1	1	https://issues.jboss.org/browse/DEVELOPER-233

#	Requirement	Questions	Answered	Successfully Answered	URL
7642	DEVELOPER-233	Jason, Can you make it so that when the user clicks on the video's title, it goes to the detail page? For example... 1) Go to: http://localhost:4242/products/eap/resources/ 2) Click on "London JBUG, April 2012: Developing Apps the JBoss Way" 3) Observe that it shows a jpg, not the detail page Also, the breadcrumb seems to still be broken. It shows: {code} Get Started > Developer Materials > Video Name {code}	0	0	https://issues.jboss.org/browse/DEVELOPER-233
7643	DEVELOPER-224	Okay I've implemented this. Can someone check the copy before it's merged? https://github.com/jboss-developer/www.jboss.org/pull/156 Currently setup like this: Subject: Why Subscribe to JBoss? Body: Here are a few reasons to subscribe to JBoss. You can view them all at http://jboss.org/why-subscribe More capability, less cost Delivered via subscription. So you pay for what's most important to you. Service Software that works. No vendor lock-in The development process is open. The technology is open. There are no secrets, no surprises, no lock-in. Award-winning support Ranked highest in customer satisfaction over Oracle, IBM, and BEA in an independent study 2 years running. Performance Just ask our customers why they trust Red Hat JBoss Middleware with their mission-critical applications. Modular architecture Robust enough to handle large, complex Java™ EE workloads, yet agile and flexible enough to run lighter ones.	1	1	https://issues.jboss.org/browse/DEVELOPER-224
7644	DEVELOPER-223	Ray, Is Dan in Jira? I don't seem to see him, so can't assign it to him.	0	0	https://issues.jboss.org/browse/DEVELOPER-223
7645	DEVELOPER-223	Are we still waiting for data for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-223
7646	DEVELOPER-222	I think we need to explore our options on where to store the image. Do we store a link to the original or a copy of the original?	1	1	https://issues.jboss.org/browse/DEVELOPER-222
7647	DEVELOPER-222	sys_description has all the html stripped, so we'd need to add in sys_content, which will increase the payload, possibly quite a bit depending on the size of the blogs we're pulling down. Pulling out images will certainly be easier if we parse the content as HTML, but that opens us up to possible XSS issues. We could attempt to use Regexp to find the image, but JavaScript doesn't support lookbehind assertions which may further complicate the regex we'd need to write :(I also went through some results from the DCP and many blogs look like the first image is a tracking image. I can craft up some javascript regex, just tried, to grab the sources and find images then take the first one that's an actual image. Shall we proceed with that attempt and see what comes out the other side?	1	1	https://issues.jboss.org/browse/DEVELOPER-222
7648	DEVELOPER-222	To be able to get "hero" image for blog posts would be very useful and is cool idea. Though, I have couple of questions: # If the blog post contains more than one image, how we select the best one? May be we should store also metadata of this image (size, number of colors, alt value?) which can help us use some heuristics to automate hero detection. A manually verified flag for hero image should be considered as well. # Where we will store the images? I am against storing binary image data into DCP (Lucene index). We should store it somewhere where each image can have a public URL (so browsers can load them quickly without need to interact with DCP at all). # How about copyrighted images? # Are we going to download (and override) all images again with every reindexing? What if the blog post becomes unavailable as time goes? Shall we still keep its images?	1	1	https://issues.jboss.org/browse/DEVELOPER-222
7649	DEVELOPER-222	I think we should think a little bit about making this more general (yet simple). I can imagine this can be useful for any kind of content (and since we want to introduce a new {{sys_}} field, it should be more generic). How about this. Let's introduce two new fields: # {{sys_featured_image_url}}: is expected to contain at most a single URL (or is empty or not provided at all) # {{sys_image_urls}}: is expected to contain any number of image URLs In your case you would check only {{sys_featured_image_url}} field and use it if provided. On the other hand this will also support other use cases, for example tweet can contain several image urls and it can make sense to display all of them while none of them is featured (as twitter does itself - [example here|https://twitter.com/AngLandiMoSobra/status/555334969241776129]). In compact search results it make sense to display only {{sys_featured_image_url}} and in more detailed search results it can make sense to display more images. This will also allow to disable individual featured image while all image URL links are still kept with the document. Moreover if we could rely on the order of images in {{sys_image_urls}} field then we could turn {{sys_featured_image_url}} field to a field called {{sys_featured_image_number}} which would only contain position (number) in the {{sys_image_urls}} list. But I am not sure about this option. I think Elasticsearch returns list items in consistent order but I am not sure if this is "officially" guarantied (thus might change in the future?). WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-222
7650	DEVELOPER-221	I see this is for 1.1, but I know something was done in the buzz section with filters, did we fix this one?	0	0	https://issues.jboss.org/browse/DEVELOPER-221
7651	DEVELOPER-219	I think this might be possible using histogram facet and might be already supported in Searchisko. But before I try to give you some query examples, do you think you can provide me more info about what type of chart/data do you want to present on the page? I am not sure I fully follow.	1	1	https://issues.jboss.org/browse/DEVELOPER-219
7652	DEVELOPER-219	One thing we should pay attention to in this case is how many XHR requests this page will need to execute to get all the data. Right now it seems it will be one or two requests per product forum. There are 11 forums on the page ATM. Firing that many requests may not be very good idea. Not only it will be probably slowed down by the browser (do they still allow only two concurrent requests per domain?) but also the risk is at our public webproxy which throttles number of requests a client can do per second. Unfortunately, I can not think of some quick workaround now but going forward situations like this should be avoided. I think this could be addressed by https://github.com/searchisko/searchisko/issues/123 (specifically by the multi_search).	1	1	https://issues.jboss.org/browse/DEVELOPER-219
7653	DEVELOPER-219	[~lightguard] did you see this query: http://dcp-stg.jboss.org/v1/rest/search?facet=activity_dates_histogram&project=jbosseap&sys_type=forumthread&size=0&activity_date_interval=month&activity_dates_histogram_interval=day Notice the 'date_histogram' section. It gives you the number of posts (count entry) on a particular date (time entry). It is returning too many results at the moment so we would need to filter those out. I think it also omits an entry for days with no posts created. Is this enough? {code} activity_dates_histogram: { _type: "date_histogram", entries: [{ time: 1381968000000, count: 1 }, { time: 1382054400000, count: 2 }, { time: 1382572800000, count: 1 }, ... {code}	1	1	https://issues.jboss.org/browse/DEVELOPER-219
7654	DEVELOPER-219	Wouldn't it be simpler to just add the project id to: https://github.com/jboss-developer/www.jboss.org/blob/master/_config/forums.yml? for example: {code} eap: name: 'JBoss EAP' id: jbosseap url: https://community.jboss.org/en/jbosseap description: 'An application platform for hosting your apps that provides an innovative modular, cloud-ready architecture, powerful management and automation, and world class developer productivity.' {code}	1	1	https://issues.jboss.org/browse/DEVELOPER-219
7655	DEVELOPER-219	Why do we have the wrong codes for some products?	1	1	https://issues.jboss.org/browse/DEVELOPER-219
7656	DEVELOPER-219	[~petemuir] What you mean? Think of it as hash ID. They could have been MD5 hash codes as well.	1	1	https://issues.jboss.org/browse/DEVELOPER-219
7657	DEVELOPER-219	[~petemuir] OK, good to know. We will need to change it. Do you know if there is anything similar for projects?	1	1	https://issues.jboss.org/browse/DEVELOPER-219

#	Requirement	Questions	Answered	Successfully Answered	URL
7658	DEVELOPER-219	[~james.cobb] can you review http://www-stg.jboss.org/pr/428/build/635/forums/ please	0	0	https://issues.jboss.org/browse/DEVELOPER-219
7659	DEVELOPER-218	We should use the same method for submitting the data as this form does: http://www.jboss.org/contact.html and use the same server to receive and process the data. One thing of concern is that the jboss.org contact form has AntiSpam control. We might need to work this into our form, in order to get our request accepted by the server. [~visuale], can you find out if the AntiSpam is needed by the server and report back.	1	1	https://issues.jboss.org/browse/DEVELOPER-218
7660	DEVELOPER-218	Do you by chance know if the regular JBoss website was built with Magnolia CMS? I see a lot of script that looks based in Magnolia – don't know if that means the whole site is built on that platform of if it's just some that some of the front-end UI borrows from Magnolia's front end UI. I've asked in the chat room and inquired with Pete Muir as well. If it is the case that JBoss uses Magnolia, then we'll need our own separate solution here, yes?	1	1	https://issues.jboss.org/browse/DEVELOPER-218
7661	DEVELOPER-218	[~visuale] Yes, the current jboss.org does use magnolia. Can't we just re-use the same back-end? I.e the service that the form submits to?	1	1	https://issues.jboss.org/browse/DEVELOPER-218
7662	DEVELOPER-218	I've tried out Magnolia in the past (a very little bit) but if memory serves and as this example shows (http://documentation.magnolia-cms.com/display/DOCS/Form+module#Formmodule-Installing), I believe that Magnolia's form processor would be entirely-self contained and not using some kind of service on in its host environment. This example runs though how to get the form data into a database: http://wiki.magnolia-cms.com/display/WIKI/Storing+form+data+in+the+Data+module There's also, of course, an example for getting the form data processed and sent as a regular email. I guess a question here would be what should the end recipient of the form be? I assume an email, of course, but will submissions also end up in a database?	1	1	https://issues.jboss.org/browse/DEVELOPER-218
7663	DEVELOPER-218	Are you still able to set this up, or shall I assign to someone else?	0	0	https://issues.jboss.org/browse/DEVELOPER-218
7664	DEVELOPER-218	[~newtonm] [~unibrew] what will the URL of this form be after the launch? The same?	1	1	https://issues.jboss.org/browse/DEVELOPER-218
7665	DEVELOPER-216	Do you know of any tools out there to help us with a creation of script like this?	1	1	https://issues.jboss.org/browse/DEVELOPER-216
7666	DEVELOPER-216	Is selenium the right tool for the job? We have over 600 pages in the site. Even doing it in parallel, selenium will take a very long time to go through the site.	1	1	https://issues.jboss.org/browse/DEVELOPER-216
7667	DEVELOPER-216	Sure. We want to make our own using the underlying tools, or do we want to use wraith in our solution? If we can use wraith with the patches you sent (assuming they take them) I'm all for not having to recreate the screenshot capture and diffing, but if not, it isn't the end of the world.	0	0	https://issues.jboss.org/browse/DEVELOPER-216
7668	DEVELOPER-213	There are no resources for EWS - we have not invested in this area for any sort of materials. Would we be able to remove some headers from the page if there is no content in them?	1	1	https://issues.jboss.org/browse/DEVELOPER-213
7669	DEVELOPER-210	I've committed some changes for this here: https://github.com/sbryzak/jboss-developer-site/commit/bfbd85efbb65680652a128081bb0d26292902967 Here's what I've done so far: 1) Implemented the script to close the popup (using JQuery) 2) Added an anchor for the direct download link 3) Added an anchor that links to the terms and conditions page 4) Used the abbreviated product name in the title of the popup (e.g. "Thank you for downloading JBoss EAP") A couple of questions: a) Should the agreement date for the terms and conditions (i.e. September 25, 2013) be hard coded? Or should it be a dynamic value (e.g. the current date)? b) Is the direct link URL correct? I used the same download URL that is used in Step 1 of the Instructions for Installation.	1	1	https://issues.jboss.org/browse/DEVELOPER-210
7670	DEVELOPER-210	I haven't sent a pull request with the changes yet, are you able to see my commit from the URL I pasted? It's possible that GitHub might restrict access since it's a private repo - if that's the case I can send a pull request to the master repo with my changes. On the other issues - a) how do we obtain the date on which the user signed the agreement? b) the download link seems to be the only one we have defined, hopefully it's correct	1	1	https://issues.jboss.org/browse/DEVELOPER-210
7671	DEVELOPER-210	It looks like it's going to be difficult to get the agreement date. I suggest we set the text to: {code} By downloading this product you have agreed with our terms and conditions. See your profile for the date you subscribed. └You will be notified again in 364 days after your subscription ends. {code} Where "your profile" links to the users profile on Jive. I think this is acheivable at the moment? Lets create another Jira issue for looking up the date from Jive. Does this work?	1	1	https://issues.jboss.org/browse/DEVELOPER-210
7672	DEVELOPER-210	Added DEVELOPER-226 for adding the date. [~sbryzak2] I think this issue is ready for merge now? Is the PR ready?	1	1	https://issues.jboss.org/browse/DEVELOPER-210
7673	DEVELOPER-207	What changes if an item is featured? We're looking for a short-term hack, yes?	1	1	https://issues.jboss.org/browse/DEVELOPER-207
7674	DEVELOPER-207	Will this always been a series of images and text? If it is we might be able to get away with a yaml file containing the image url and the text.	1	1	https://issues.jboss.org/browse/DEVELOPER-207
7675	DEVELOPER-202	[~pmuir] I seem to recall you having discussions with Lukas around contributor details in the DCP. Can you update this issue with what details we can get from the DCP?	1	1	https://issues.jboss.org/browse/DEVELOPER-202
7676	DEVELOPER-202	I can imagine some "contributor profile proxy" url on DCP side. This will be used when you can show user profile. You will simply pass identifier from sys_contributor there. When used, then DCP will search for relevant profile, and then redirect browser to final contributor profile in jive. WDYT?	1	1	https://issues.jboss.org/browse/DEVELOPER-202
7677	DEVELOPER-196	[~unibrew2] Can you make a Magnolia form that looks like this: http://design.jboss.org/jbossdeveloper/sitedesign/jbossdeveloper_1_0/concepts/get_involved/images/jbossdeveloper_getinvolved_events_addevent-1200px.png?	1	1	https://issues.jboss.org/browse/DEVELOPER-196
7678	DEVELOPER-194	is this still active?	0	0	https://issues.jboss.org/browse/DEVELOPER-194
7679	DEVELOPER-193	[~wesbos]: Can the "For Linux or Mac" bit be removed from what is displayed? That's now implied by the selected OS. It also adds a lot of cruft as it appears in many places on the page.	1	1	https://issues.jboss.org/browse/DEVELOPER-193
7680	DEVELOPER-179	[~lightguard] where did we get to with querying the dcp for URLs to profile pages on Jive?	1	1	https://issues.jboss.org/browse/DEVELOPER-179
7681	DEVELOPER-179	Do we still have a question where we're at here, do we want to document something?	1	1	https://issues.jboss.org/browse/DEVELOPER-179
7682	DEVELOPER-179	How much more of this is missing/outstanding?	1	1	https://issues.jboss.org/browse/DEVELOPER-179
7683	DEVELOPER-178	Catherine, could you check that you agree I've selected the right endpoints for the URLs?	1	1	https://issues.jboss.org/browse/DEVELOPER-178
7684	DEVELOPER-172	Would a js library help us here? for example highlight.js to give us syntax highlighting. Or maybe something more heavyweight providing line-numbers, a copy-to-clipboard button and a title?	0	0	https://issues.jboss.org/browse/DEVELOPER-172
7685	DEVELOPER-172	We could use a library, but the base font in red looks horrible. Highlight.js work work okay. Is that how we'd like to proceed?	0	0	https://issues.jboss.org/browse/DEVELOPER-172
7686	DEVELOPER-172	Do we need [~james.cobb] or [~crobson] to approve this? My concern would be that going with the OOTB http://highlightjs.org/ is going to look odd, without an appropriate L&F applied.	1	1	https://issues.jboss.org/browse/DEVELOPER-172

#	Requirement	Questions	Answered	Successfully Answered	URL
7687	DEVELOPER-161	Catherine, do you want me to sub this text in?	1	1	https://issues.jboss.org/browse/DEVELOPER-161
7688	DEVELOPER-159	Question for James - should we have "home" as a navigation option in the tray? On the website it is obvious that we don't have a home button on the navigation banner, so people know to click the logo, but on mobile format people might default to opening the tray and then have to close it to get back to the logo for the home page.	1	1	https://issues.jboss.org/browse/DEVELOPER-159
7689	DEVELOPER-155	[~james.cobb] The link to the mobile filtering wireframe is actually the desktop. Can you send me what it should look like on mobile?	1	1	https://issues.jboss.org/browse/DEVELOPER-155
7690	DEVELOPER-151	Are we done with this James?	0	0	https://issues.jboss.org/browse/DEVELOPER-151
7691	DEVELOPER-151	Shane, What is the status of the "Demo's" table?	1	1	https://issues.jboss.org/browse/DEVELOPER-151
7692	DEVELOPER-146	Can this be marked as resolved?	0	0	https://issues.jboss.org/browse/DEVELOPER-146
7693	DEVELOPER-143	Can this be marked as resolved?	1	1	https://issues.jboss.org/browse/DEVELOPER-143
7694	DEVELOPER-141	Does this mean we are killing the dropdowns? Or keeping them?	0	0	https://issues.jboss.org/browse/DEVELOPER-141
7695	DEVELOPER-141	Is this still an issue?	0	0	https://issues.jboss.org/browse/DEVELOPER-141
7696	DEVELOPER-140	Just about to push these changes to github. For the mobile hero, I think we need to look at skinny and tall heros rather than wide ones. We can supply different images for mobile / desktop. Just take the same content and make it stacked rather than wide?	1	1	https://issues.jboss.org/browse/DEVELOPER-140
7697	DEVELOPER-134	Current issues: -1) Time commitment needs removing- 2) Waiting for Searchisco for filtering of "rating, formats, publish date" 3) This needs removing from results somehow (as it's not a real quickstart): http://www-stg.jboss.org//quickstarts/eap/template/index.html. We might need to prevent this item from being indexed in the first place. -3) Remove the "??" time from each result- -4) Populate list with everything on first visit to the page. Currently the page is empty when first visited. Items start to appear once the criteria is changed.- -5) H2 topic is on a new-line. It should be on the same line, like the other items.- -6) Is pagination missing? I can only get a max of 10 results to display.This could simply be that we have only uploaded 10 items to the DCP. Alternatively, there could be some pagination going on, but with no links to navigate through the pages. This is to be investigated.- -7) After clicking on an item to see the details, hitting the 'back' button resets the filter. After hitting back, I think the same result-set should be displayed as was there before clicking on more details.- -8) Check what the skill level selector does. I think it should return just those quickstarts at the selected level. For example, "intermediate" would return just "intermediate" quickstarts, rather than "beginner" ones too. I think this is more obvious to the user.- -9) Every item has the text "VIDEO" next to a star rating. I assume it should say "Rating"?-	1	1	https://issues.jboss.org/browse/DEVELOPER-134
7698	DEVELOPER-130	This page is finished now. Can I close it?	1	1	https://issues.jboss.org/browse/DEVELOPER-130
7699	DEVELOPER-127	If we have authors for pages (quickstarts, bom, etc.) all that info is already in the page (the ruby object) otherwise... are we scraping the html for them, using contributors in git?	1	1	https://issues.jboss.org/browse/DEVELOPER-127
7700	DEVELOPER-127	[~velias], [~newtonm] indicated you are still able to help us out with some tasks. Can you take this one on please? We basically have some pages on the site that don't get indexed ([~lightguard] can you provide a list please?). We need to make sure they get indexed in some form on the dcp, so that when people search, they hit them. I also think [~newtonm] understands this task, so you can ask him for more info as well.	1	1	https://issues.jboss.org/browse/DEVELOPER-127
7701	DEVELOPER-127	[~wesbos] once we have the indexed data in the DCP, we need to be able to search for it on the site. Pete and I were thinking the simplest approach for now would be to use some form of "auto-complete" style search. Here the user would start typing in their criteria in the search box, and results would be suggested. Selecting a result would redirect to that page. The idea was to try and keep it as simple as possible for now. Is this something you can do, once the data is in the DCP?	1	1	https://issues.jboss.org/browse/DEVELOPER-127
7702	DEVELOPER-127	[~paul.robinson] absolutely! I think that approach would be best as well. Please assign me to this (or another issue) once the data is ready. Do you know approximately when it should be ready to go?	1	1	https://issues.jboss.org/browse/DEVELOPER-127
7703	DEVELOPER-127	[~velias] How are you getting on with this?	1	1	https://issues.jboss.org/browse/DEVELOPER-127
7704	DEVELOPER-122	What additional things do we need here? We have emails that GitHub and Travis send out. GitHub is also sending us irc notifications.	1	1	https://issues.jboss.org/browse/DEVELOPER-122
7705	DEVELOPER-105	Do you have a version with the media queries as well? This example isn't mobile ready but http://example.jboss.org/ is	1	1	https://issues.jboss.org/browse/DEVELOPER-105
7706	DEVELOPER-105	I had the blue separator but was asked to take it off. Should I put it back on?	0	0	https://issues.jboss.org/browse/DEVELOPER-105
7707	DEVELOPER-100	I've committed a rough first draft for this page here: https://github.com/sbryzak/jboss-developer-site/commit/32369ce0917ad9b75287a14888c783937ddd8feb There is a new file, _config/forums.yml which contains the data from which the page is generated. I have a few questions: 1) What should the text be at the top for the forum guidelines? 2) How/where can I obtain the posts count for each forum? 3) Where can I get the monthly activity graph for each forum? 4) WFK doesn't seem to have a forum page. 5) The descriptions were copied from http://www.jboss.org/products - are these too verbose? 6) The 'See all JBoss Forums' button links to http://www.jboss.org/projects, is this suitable? 7) Where should the 'Sticky' link in the first row link to?	0	0	https://issues.jboss.org/browse/DEVELOPER-100
7708	DEVELOPER-100	1) I'll look into if we already have something suitable for this, and if not will create some copy for you 2) We should talk to Mark Newton's team about if they have this type of data you can grab through an API or DCP 3) We should talk to Mark Newton's team about if they have this type of data you can grab through an API or DCP 4) WFK can be removed for now. We'll need to make sure the list of products forums we have maps to the products on the site in general at this point. 5) Maybe - but lets start here and we can modify them once we see how it looks on the site 6) No - it should link to: https://community.jboss.org 7) This forum doesn't exist yet - but it was something I thought we should consider creating. For now, lets remove it and we can add it if we get around to creating such a forum area.	0	0	https://issues.jboss.org/browse/DEVELOPER-100
7709	DEVELOPER-100	I agree, that the forum descriptions are too verbose. How about we use the descriptions off /products for the Beta? We can improve them for the GA if we need to. I think we can leave the post count and activity graph for the Beta. It's quite clear what that feature is going to look like. I see there is a discussion in-hand with Mark's team investigating this.	1	1	https://issues.jboss.org/browse/DEVELOPER-100
7710	DEVELOPER-100	Moved 1) and 2) into DEVELOPER-219.	0	0	https://issues.jboss.org/browse/DEVELOPER-100
7711	DEVELOPER-96	Reassigning to me for review of architecture, and will assign to Wes for JS impl as needed	0	0	https://issues.jboss.org/browse/DEVELOPER-96
7712	DEVELOPER-92	Wes implemented this page, but it is not in the right place in the navigation. What is shown today at http://www-stg.jboss.org/get-involved/contribute/ should be moved to http://www-stg.jboss.org/get-involved	1	1	https://issues.jboss.org/browse/DEVELOPER-92
7713	DEVELOPER-92	Jason - Can you move this page to the right place in the navigation? This should overwrite the current get-involved page.	1	1	https://issues.jboss.org/browse/DEVELOPER-92

#	Requirement	Questions	Answered	Successfully Answered	URL
7714	DEVELOPER-91	I've committed the first draft of this page here: https://github.com/sbryzak/jboss-developer-site/commits/DEVELOPER-91 I have a few questions: 1) I've pulled the images from DESIGN-479 just using Gimp to extract the icons - if this isn't acceptable then I'll need the actual icon resources or the PSD file so that I can extract the icons from their layer. 2) Why should happen when you click the "SEND THIS CONTENT TO YOUR BOSS" button? 3) What should the URLs be for all of the content in the side-panel? Specifically: a) Subscription Guide b) Support Satisfaction Report c) Case Studies d) Red Hat JBoss Middleware vs. IBM Websphere e) Red Hat JBoss Middleware vs. Oracle Fusion f) Datasheets and brochures g) Whitepapers h) Webinars 4) What should the URL be for the "Fill out an interest form" link? 5) What should the URL be for the "Find out how to contact sales for your country" link? 6) What should the URL be for the "Customer Portal" link? 7) What should the URL be for the "Contact us" link? (I'm assuming this should be a mailto: address)	1	1	https://issues.jboss.org/browse/DEVELOPER-91
7715	DEVELOPER-91	[~crobson] Also, can you answer Shane about the links? Or maybe you know who we should ask?	1	1	https://issues.jboss.org/browse/DEVELOPER-91
7716	DEVELOPER-91	3) What should the URLs be for all of the content in the side-panel? Specifically: a) Subscription Guide https://www.redhat.com/rhecm/rest-rhecm/jcr/repository/collaboration/jcr:system/jcr:versionStorage/c52b53bb0a05260220b4e3d889f155db/1/jcr:frozenNode/rh:resourceFile b) Support Satisfaction Report https://www.redhat.com/rhecm/rest-rhecm/jcr/repository/collaboration/jcr:system/jcr:versionStorage/5e6e3e7a7f0000011d40cd22e322196b/4/jcr:frozenNode/rh:pdfFile.pdf c) Case Studies https://www.redhat.com/resourcelibrary/results?portal:componentId=2e9e0c58-a437-4c88-b1a8-c492d711eca7&portal:type=action&portal:isSecure=false&navigationalstate=JBPNS_rO0ABXcaAAlTZWFyY2hLZXkAAAABAAAAB19fRU9GX18*&fName=Products&facetValueName=jbossenterprisemiddleware&actionType=addFacetsForProducts&SearchKey=&nestedSearch=false d) Red Hat JBoss Middleware vs. IBM Websphere https://engage.redhat.com/forms/jboss-v-websphere e) Red Hat JBoss Middleware vs. Oracle Fusion https://engage.redhat.com/forms/jb_oracle_fusion_competitive_sheet f) Datasheets and brochures https://www.redhat.com/resourcelibrary/results?portal:componentId=2e9e0c58-a437-4c88-b1a8-c492d711eca7&portal:type=action&portal:isSecure=false&navigationalstate=JBPNS_rO0ABXcaAAlTZWFyY2hLZXkAAAABAAAAB19fRU9GX18*&fName=DocumentType&facetValueName=datasheetsbrochures&actionType=addFacets&SearchKey=&nestedSearch=false g) Whitepapers https://www.redhat.com/resourcelibrary/whitepapers/?portal:componentId=220ad020-88ac-4404-a4e3-d2c19f3bdf22&portal:type=action&portal:isSecure=false&fName=Products&facetValueName=jbossenterprisemiddleware&actionType=addFacetsForProducts&SearchKey=&nestedSearch=false h) Webinars https://www.redhat.com/resourcelibrary/webinars/?portal:componentId=9f1f989c-f322-4db8-bb93-05211ed74970&portal:type=action&portal:isSecure=false&fName=Products&facetValueName=jbossenterprisemiddleware&actionType=addFacetsForProducts&SearchKey=&nestedSearch=false 4) What should the URL be for the "Fill out an interest form" link? http://www.redhat.com/contact/sales.html 5) What should the URL be for the "Find out how to contact sales for your country" link? I'm not sure where this one came from - I think it might have been from our contractor. Let's change the link name to "Find a Red Hat JBoss reseller" and link it here: http://redhat.force.com/finder/Partner_Finder 6) What should the URL be for the "Customer Portal" link? https://access.redhat.com/home 7) What should the URL be for the "Contact us" link? (I'm assuming this should be a mailto: address) Ray creating mailing list - will post back when we have it. Eventually this should be linking to a form on the website (http://design.jboss.org/jbossdeveloper/sitedesign/jbossdeveloper_1_0/concepts/utility/images/jbossdeveloper_contact-1200px.png)	1	1	https://issues.jboss.org/browse/DEVELOPER-91
7717	DEVELOPER-88	I just need a list of ISBNs. Can I get this from the DCP?	0	0	https://issues.jboss.org/browse/DEVELOPER-88
7718	DEVELOPER-85	[~alrubinger]: is this a dup of DEVELOPER-273?	1	1	https://issues.jboss.org/browse/DEVELOPER-85
7719	DEVELOPER-85	[~alrubinger] what's the plan with this? Will you have time to get this done by May 15th?	1	1	https://issues.jboss.org/browse/DEVELOPER-85
7720	DEVELOPER-82	Isn't this a duplicate of DEVELOPER-58?	0	0	https://issues.jboss.org/browse/DEVELOPER-82
7721	DEVELOPER-78	Isn't linking to the original blog post itself enough?	0	0	https://issues.jboss.org/browse/DEVELOPER-78
7722	DEVELOPER-77	Hi [~lightguard], There are some options that we can try to use right out of the box [More Like This|http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-more-like-this.html] (or [Suggesters|http://www.elasticsearch.org/guide/en/elasticsearch/reference/current/search-suggesters.html]) and it shouldn't be hard to expose such types of queries. But I would like to understand a little bit more context about what you mean by related materials first. Do you think you can give me some examples?	1	1	https://issues.jboss.org/browse/DEVELOPER-77
7723	DEVELOPER-77	More Like this is exactly what we need. Can we use that the dcpdev instance today?	1	1	https://issues.jboss.org/browse/DEVELOPER-77
7724	DEVELOPER-77	Just to be clear what does "listed field values" mean? I can guess, but ... ;-)	1	1	https://issues.jboss.org/browse/DEVELOPER-77
7725	DEVELOPER-77	To push this ticket forward let's talk about how you would like to make use of "similar" or "related" or "more like this" (or whatever we call it) on the developer material page. Do you think you can give me some idea about how it should be used on the website? For example do you want to display clickable "similar" link for each listed item? Or display "similar" link for each item only if the number of listed items drops down to just a few items? (meaning, user probably can not find what he is looking for). Trying to understand how you want to use/incorporate it on a real use case first.	1	1	https://issues.jboss.org/browse/DEVELOPER-77
7726	DEVELOPER-77	[~lvlcek] we want to use this on the right rail of the "detail" page for the developer material e.g. http://www.jboss.org/ticket-monster/BusinessLogic/index.html Here we would display items similar to the e.g. the tutorial chapter (but could also be for any developer material, such as a quickstart). Does this answer your question?	1	1	https://issues.jboss.org/browse/DEVELOPER-77
7727	DEVELOPER-76	There are a few of things to take into consideration here 1. Have we added content into the slim files, or kept them in asciidoc / markdown (quickstarts)? 2. How far do we want to go with i18n / l10n? Full site translation or just the content? 3. To my knowledge none of slim, kramdown (markdown processor), nor asciidoctor support i18n or l10n We'd certainly be creating our own solution at this point in time.	1	1	https://issues.jboss.org/browse/DEVELOPER-76
7728	DEVELOPER-71	Can this be resolved now?	1	1	https://issues.jboss.org/browse/DEVELOPER-71
7729	DEVELOPER-64	Catherine, could you do a wireframe of a standalone page for this?	1	1	https://issues.jboss.org/browse/DEVELOPER-64
7730	DEVELOPER-64	Also done, though I believe there's still some text that needs to be added. Pete, do you have that text?	0	0	https://issues.jboss.org/browse/DEVELOPER-64
7731	DEVELOPER-63	Catherine, could you do a wireframe of a standalone page for this?	0	0	https://issues.jboss.org/browse/DEVELOPER-63
7732	DEVELOPER-62	[~crobson], could you take a look at http://www.jboss.org/jdf/examples/get-started/ and propose a wireframe?	1	1	https://issues.jboss.org/browse/DEVELOPER-62
7733	DEVELOPER-58	Jason, can you take over this one?	0	0	https://issues.jboss.org/browse/DEVELOPER-58

#	Requirement	Questions	Answered	Successfully Answered	URL
7734	DEVELOPER-58	Outstanding issues: * Thumbnails missing. [~james.cobb] has the images. We need to decide where to store them and how to link to them. * "Show Me" and "Included In" drop-downs not being populated. What should they show? Data missing from DCP: * last update * download link * Additional resources (from view-more) * List of products it's contained in	1	1	https://issues.jboss.org/browse/DEVELOPER-58
7735	DEVELOPER-58	[~james.cobb], the white background seems to always be covering the default logo, even when the project logo is missing. See the attached screen-shot. I think this seems to be a problem with the current approach, as you can only hide the missing image. What I think you need is to also hide the white image as well?	0	0	https://issues.jboss.org/browse/DEVELOPER-58
7736	DEVELOPER-52	What's missing other than the webinar? I've only put in what was in the spreadsheet. BTW, there's a problem with the webinar that is being fixed. I'll chase that.	1	1	https://issues.jboss.org/browse/DEVELOPER-52
7737	DEVELOPER-49	[~paul.robinson], is this still outstanding?	1	1	https://issues.jboss.org/browse/DEVELOPER-49
7738	DEVELOPER-44	[~crobson] Also, I need a graphic for the title, like those on here: http://www-stg.jboss.org/products/. Do you have one or do we need to make it?	1	1	https://issues.jboss.org/browse/DEVELOPER-44
7739	DEVELOPER-34	Are we going to manually create this list for now until we can get it into the dcp?	0	0	https://issues.jboss.org/browse/DEVELOPER-34
7740	DEVELOPER-34	This will be up to us since there is no indexer out of the box, we'll have to create that as well. If the only ones using this are going to us, things we make and possibly other engineers I recommend a simple JSON (well, I guess that depends on where we want the extra library :) on the Java side creating the indexer or on the Ruby side writing XML). Either way works really. Something simple like: {code} book: authors: [name,name] title: title pub_date: date publisher: name topics: [topic, topic] tags: [tag, tag] {code} Is there anything else we need?	0	0	https://issues.jboss.org/browse/DEVELOPER-34
7741	DEVELOPER-34	Do we also want to link to some way of obtaining the book, or do we assume that they are all on Amazon and our audience can will just go and search there? Also, some may have a free eBook and maybe book-src that we can link to?	1	1	https://issues.jboss.org/browse/DEVELOPER-34
7742	DEVELOPER-30	Where has this been added, I'm not seeing it in developer?	0	0	https://issues.jboss.org/browse/DEVELOPER-30
7743	DEVELOPER-30	What was the outcome of this meeting talked about above?	0	0	https://issues.jboss.org/browse/DEVELOPER-30
7744	DEVELOPER-29	I talked about doing angular for this, do we think this might still be a good idea? I see there being a slight problem unless we create a backend for this in that we don't have all the options.	1	1	https://issues.jboss.org/browse/DEVELOPER-29
7745	DEVELOPER-29	Do you see it changing very much, or will it be pretty static once it's done?	1	1	https://issues.jboss.org/browse/DEVELOPER-29
7746	DEVELOPER-29	Do you mean the content, or the code for the finder?	1	1	https://issues.jboss.org/browse/DEVELOPER-29
7747	DEVELOPER-29	What results should we show by default? Just the first ones that show up from a blank search?	1	1	https://issues.jboss.org/browse/DEVELOPER-29
7748	DEVELOPER-29	I've made a start at this: https://github.com/paulrobinson/jboss-developer-site/commit/cdb8d7fd01772b4b6c1c185be88dac2b69426967 Essentially, I've cloned _layouts/product-developer-materials.html.slim and removed the product specific decorations from the page. See the attached screenshot for what it looks like. The result is working, but the formatting is wonky. I suspect it's because the current style expects the length of the page to be dictated by the size of the right-bar. This page has no right bar, so the footer ends up being too high. [~visuale] Can you take my commit and fix it? [~lightguard] can you check my approach is correct?	1	1	https://issues.jboss.org/browse/DEVELOPER-29
7749	DEVELOPER-29	Hi Paul, I'm having a bit of trouble cloning that specific commit. Is there a specific branch I should be working off of?	1	1	https://issues.jboss.org/browse/DEVELOPER-29
7750	DEVELOPER-28	Are we combining information from various sources, or will this be something the product people maintain?	1	1	https://issues.jboss.org/browse/DEVELOPER-28
7751	DEVELOPER-27	Will we be able to limit the buzz shown as to only those items pertinent to the product we are displaying it on?	1	1	https://issues.jboss.org/browse/DEVELOPER-27
7752	DEVELOPER-26	could stacks help here too?	1	1	https://issues.jboss.org/browse/DEVELOPER-26
7753	DEVELOPER-26	Each product has it's own xml file in this model, the uses portion allows a soft tie (put together via finding the actual project reference stored in the site within awestruct. I couldn't think of a good way to tie them together without creating a unified xml file, which I felt like would lead to confusion for those updating the files (I don't know what it uses, why is this extra info in there, there's too many options, etc). We can look at revising it if needed. The tie would be to use site.org(from the project or spec element).(project or spec)name. Each organization (Red Hat, Apache, Eclipse, etc) would have it's own file in _specs and _projects. Does that help?	1	1	https://issues.jboss.org/browse/DEVELOPER-26
7754	DEVELOPER-25	Is this just scraping JIRA and GitHub or do we need more than that?	1	1	https://issues.jboss.org/browse/DEVELOPER-25
7755	DEVELOPER-22	I've added a schema. [~lightguard] I'm wondering if we can reuse your XML parser? Could you look at that?	1	1	https://issues.jboss.org/browse/DEVELOPER-22
7756	HRCPP-364	Vittorio, can you create the same PR for the product please?	0	0	https://issues.jboss.org/browse/HRCPP-364
7757	HRCPP-25	The PR has been integrated, shouldn't you close this issue?	0	0	https://issues.jboss.org/browse/HRCPP-25
7758	HRCPP-20	Isn't this a duplicity for HRCPP-17 ?	0	0	https://issues.jboss.org/browse/HRCPP-20
7759	IMMUTANT-628	[~gantengx] - I have concerns about implementing this - all that will happen is that the call to {{stop}} will return immediately, it won't actually kill any long running jobs (unless you are killing an uberjar/{{lein run}} process with {{^C}}, etc, since that will kill all of the scheduler threads). How are you stopping your system currently? And are you running this as an uberjar or inside WildFly?	1	1	https://issues.jboss.org/browse/IMMUTANT-628
7760	IMMUTANT-628	[~tcrawley] Ah so there's no way to actually force killing the jobs? Currently I'm running this as an uberjar using init script. When stopping the system it's basically calling {{kill}} command	1	1	https://issues.jboss.org/browse/IMMUTANT-628
7761	IMMUTANT-580	This has been implemented, and is available in incremental build #682 and newer. Can you give it a try and confirm it works for you (see http://immutant.org/builds/2x/ for instructions on using incremental builds)? Relevant commits: https://github.com/projectodd/wunderboss/commit/0e6f802debb429d5ffb1537ecffae561fa72fc65 https://github.com/immutant/immutant/commit/be8e1009f51047dae432b7f5136a867871e23418	0	0	https://issues.jboss.org/browse/IMMUTANT-580
7762	IMMUTANT-568	Replying to the mailing list thread [1] here: > I think this option should be separate from `:singleton`, despite the > intuitive implication, and treat it as a local-only option. I can > imagine cases where you want a job running on only one machine, but > concurrently there. [~tcrawley] I'm not sure I understand correctly. Do you mean that setting `:allow-concurrent-exec?` to false would allow the job to execute 'in parallel' on another machine but not the same machine? Intuitively I would expect the option to behave differently by default: if set the job should not execute concurrently cluster wide, thus implying `:singleton` true. It's then up to the user to override by setting `:singleton` to true explicitly. ...just my two cents. [1] http://immutant.markmail.org/thread/n3wnofa6dg4dcgtg	1	1	https://issues.jboss.org/browse/IMMUTANT-568
7763	IMMUTANT-568	This is now available in [incremental build|http://immutant.org/builds/2x/] 591 and newer. Can you give that a try and confirm it works for you?	1	1	https://issues.jboss.org/browse/IMMUTANT-568
7764	IMMUTANT-562	Ah, that's a good idea. Would you mind sharing an example of how you are doing that?	1	0	https://issues.jboss.org/browse/IMMUTANT-562

#	Requirement	Questions	Answered	Successfully Answered	URL
7765	IMMUTANT-456	Immutant uses a custom configurator itself, to return a singleton Endpoint, a feature also described at that link. So if we took in a custom configurator, we would need to override its getEndpointInstance as we do now. That might lead to surprising behavior. What if we just go ahead and make Immutant's configurator put the HttpSession and cookies in the user properties?	1	1	https://issues.jboss.org/browse/IMMUTANT-456
7766	IMMUTANT-456	Another thought: how do you feel about passing a function instead of a custom configurator? Its signature would match that of the Configurator's modifyHandshake method. So we'd introduce another option key, :handshake, which would be a fn that should expect to be passed a ServletEndpointConfig, a HandshakeRequest and a HandshakeResponse. It's more limiting, in that we wouldn't expose the other Configurator methods, but possibly less surprising, and simpler to create a function than an object, I think.	1	1	https://issues.jboss.org/browse/IMMUTANT-456
7767	IMMUTANT-453	I've fixed Immutant to honor META-INF in resources (see IMMUTANT-455), so you no longer need to put he META-INF at the root, and it should work the same for .ima's or on-disk deployments. This will be in the latest incrementals (build #1206 or newer, #1206 is building now) - can you give it a try and confirm?	0	0	https://issues.jboss.org/browse/IMMUTANT-453
7768	IMMUTANT-436	Alex: The latest incremental build implements IMMUTANT-437 (which will load log4j.properties from resources/, or anywhere else on the classpath components that are within the app root) and IMMUTANT-443 (which restores log messages from Immutant namespaces when you use a custom logging configuration). Can you try the latest incremental build and make sure it works for you? You can get it via `lein immutant install LATEST`. I plan to get a 1.1.2 release out this week with these changes.	0	0	https://issues.jboss.org/browse/IMMUTANT-436
7769	IMMUTANT-436	Hmm, actually, since we now detect the presence of c.t.logging for IMMUTANT-443, we could force c.t.logging to use the log4j-factory. The question is: is that a bad idea? I suspect that most c.t.l users don't care which logging provider is being used unless they are providing custom logging configuration.	0	0	https://issues.jboss.org/browse/IMMUTANT-436
7770	IMMUTANT-415	I just pushed (via http://github.com/immutant/immutant/commit/047042ac0) a function that exposes this setting: {{(immutant.jobs/set-scheduler-options {:thread-count 10})}}. Can you see if this works for you? It will be available in incremental build #1166 (building now) or newer. Note that this function must be called before any jobs are scheduled, else it will throw an IllegalStateException.	1	1	https://issues.jboss.org/browse/IMMUTANT-415
7771	IMMUTANT-410	Kanwei: I tried to recreate this issue with https://github.com/tobias/immutant-410/blob/master/src/cookie_store_test/core.clj, but it's working for me. Can you share something that doesn't work, or guide me to make changes to my test app? The "Handler returned nil" error means that the ring handler fn returned nil instead of a response map. That error message could be more helpful. I'll take a look at that.	0	0	https://issues.jboss.org/browse/IMMUTANT-410
7772	IMMUTANT-405	Ian: Do you have any details about what header-based sessions might look like? Is it just a matter of storing a session-id in a header? How is that different than storing the session-id in a cookie header?	0	0	https://issues.jboss.org/browse/IMMUTANT-405
7773	IMMUTANT-390	Why is this ticket got closed? What does it mean "out of date"? Today I redeployed a simplest hello-world application war (built in clojure) on a standalone wildfly 8 (no clusters) and I was getting 404 during 10 seconds. If you think it's an upstream issue, could you link to it here?	1	0	https://issues.jboss.org/browse/IMMUTANT-390
7774	IMMUTANT-359	Have we marked them as deprecated somehow? A grep of the codebase says "nope".	0	0	https://issues.jboss.org/browse/IMMUTANT-359
7775	IMMUTANT-334	What would the workaround be for this? I added {{:lein-profiles [:default]}} to the {{:immutant}} descriptor in my {{project.clj}}, but e.g. {{target/classes}} still isn't on the classpath. (Though, {{target/user+dev+base/classes}} or whatever it is isn't being created anymore...)	1	1	https://issues.jboss.org/browse/IMMUTANT-334
7776	IMMUTANT-334	You shouldn't have seen {{target/user+dev+base/classes}} being generated I don't believe - were you? The issue is that since Immutant doesn't apply the default profile internally when resolving the project, but instead a subset, it expects the classes to be in {{target/user+dev+base/classes}}, where they actually are. I think the fix is to have Immutant always add {{target/classes}} to the effective classpath to cover this case, and the case where users are using a version of lein older than what Immutant uses internally.	1	1	https://issues.jboss.org/browse/IMMUTANT-334
7777	IMMUTANT-334	The latest incrementals are using 2.3.2, and I modified an integ to confirm that {{target/classes}} is on the ECP when no profiles are applied: https://github.com/immutant/immutant/commit/6f6be5cc6567f677b5cc90a99d2a0c3639159682. Chas: can you confirm?	0	0	https://issues.jboss.org/browse/IMMUTANT-334
7778	IMMUTANT-334	Mariano: can you confirm as well? (I forgot that you were the original reporter).	1	1	https://issues.jboss.org/browse/IMMUTANT-334
7779	IMMUTANT-329	I opted to just add the com.sun.net.httpserver.* packages instead of trying to determine a list of all 'public' com.sun packages, since the workaround is simple if more are discovered to be missing. This should be fixed by http://github.com/immutant/immutant/commit/87013ab53 and https://github.com/projectodd/wildfly/compare/56f59b98def1441727d9e797a7db57b65183724c...bd2a5fc891417c355690fe0ee436b6a73b542a46 Note that these changes will only affect the slim build of Immutant. The full build will not have these classes available, since we don't modify that build at all. Chas: can you confirm? These changes are already in the latest incremental builds.	1	1	https://issues.jboss.org/browse/IMMUTANT-329
7780	IMMUTANT-326	Paul, forgive my ignorance, can you give some more details around this notion of a passive listener, please? Are multiple clients supposed to see the same message, but only one ack's it? This sounds more like a topic to me. I'm not opposed to expanding the API, but it'd be nice if it had more application than just testing production traffic without affecting it (which seems kinda risky anyway). I'm not crazy about moving from jms to core, mainly because we may be switching to activemq in the distant future, and i'd like to support both it and hornetq. I would like to make sure that whatever core HornetQ abstractions you need (managers, factories, servers, etc) are within arm's reach of you and convenient, of course.	1	1	https://issues.jboss.org/browse/IMMUTANT-326
7781	IMMUTANT-326	So will it help to expose the ACK mode? It's not clear to me how setting the mode to CLIENT will achieve your goal, though. I don't think the broker will deliver a message to the "real" consumer just because the passive one doesn't ack it. There's no guarantee the passive consumer won't get the same message redelivered multiple times before the real consumer gets it once (though that's unlikely). Makes sense that you can't turn your production queues into topics just to get monitoring, and I hate for you to introduce another topic just for that. You can list the queue contents via the mgmt/jmx interfaces, but it's difficult to map them to actual consumers, if that's the level of granularity you need.	1	1	https://issues.jboss.org/browse/IMMUTANT-326
7782	IMMUTANT-317	I looked into what it would take to do this at deploy-time inside Immutant, and it's currently not straightforward, because: 1) we explicitly strip any Immutant deps from the project before passing them to {{resolve-dependencies}}. If we wanted to get the pedantic output for those deps, we would have to call {{resolve-dependencies}} twice, which is time consuming 2) most apps won't have any explicit Immutant deps, so we would need to add them before passing the project to {{resolve-dependencies}}. That would be possible for versioned releases, but wouldn't work for incremental builds, since the jar artifacts won't be available. 3) {{leiningen.core.classpath}} currently memoizes the pedantic output, which means an app that triggered warnings would not do so on a redeploy, nor would a second app that happens to have the same dependencies Putting this check in the plugin may be a better option, though we would still have the same issue with incremental builds. Maybe we could just support this for versioned builds? I'll give it some more thought.	0	0	https://issues.jboss.org/browse/IMMUTANT-317

#	Requirement	Questions	Answered	Successfully Answered	URL
7783	IMMUTANT-287	Is this still true (from web.org)? bq. NOTE: As of this writing, Pedestal route tables are not context-aware, so unless you deploy them to the root context, you must hard-code the context path in your routes.	0	0	https://issues.jboss.org/browse/IMMUTANT-287
7784	IMMUTANT-287	cache.org: the default locking mode is nil, but the description states: "One of :optimistic or :pessimistic". so what does nil mean?	0	0	https://issues.jboss.org/browse/IMMUTANT-287
7785	IMMUTANT-287	jboss.org: do we need to update the logging section?	0	0	https://issues.jboss.org/browse/IMMUTANT-287
7786	IMMUTANT-219	Should this be resolved as 'Done'? We didn't actually add scheduled delivery to publish.	0	0	https://issues.jboss.org/browse/IMMUTANT-219
7787	IMMUTANT-195	Pragmatically, I'm guessing that filing a bug with tomcat is not likely going to be profitable? (the code is very clearly wrong by design, to deal with a flaw in another part of the system, which is undoubtably even harder to fix) Not that I have any strong desire to do this, but I don't suppose it's possible to easily replace CoyoteAdapter with subclass that works?	1	1	https://issues.jboss.org/browse/IMMUTANT-195
7788	IMMUTANT-195	I've turned on ALLOW_BACKSLASH by default (https://github.com/immutant/immutant/commit/218b85bd4af809e22eb0d44136b03799459208fb). Can you confirm that the latest incremental works for you? If so, I'll resolve this issue.	0	0	https://issues.jboss.org/browse/IMMUTANT-195
7789	IMMUTANT-182	It might be worthwhile to do this by implementing one proper Chef/Pallet/Ansible/Puppet script; this is both practical and documents best practices for configuring a cloud node to be part of a cluster, would include init scripts for common init systems, also EC2-specific stuff for S3PING or AWSPING, logging, debugging, remote nREPL connectivity from a development system, etc. Ideally, if Ansible (my favorite), also shows how to do cluster-wide updates on EC2 through adding domain mode in Immutant or more likely a script to do rolling upgrades across the cluster. Other questions: - Is upgrading an app on an AMI better / worse than just doing a DNS swap across an entire cluster? - Maybe there is a good balance between power cycling an app, immutant, or an entire AMI/Node?	1	1	https://issues.jboss.org/browse/IMMUTANT-182
7790	IMMUTANT-107	Why? If we include deps in archives by default, why would we resolve them by default at deployment? It's too much of a pain to manually create the deployment descriptor just to turn :resolve-dependencies to false.	1	0	https://issues.jboss.org/browse/IMMUTANT-107
7791	IMMUTANT-96	Another reason to consider this change is that using '/' in queue/topic names will make it difficult to use HornetQ REST interface to JMS. See Section 3.1 of http://docs.jboss.org/resteasy/hornetq-rest/1.0-beta-1/userguide/html_single/index.html#d0e162 The URI for a topic/queue resourc would be /queues/jms.queue.{immutant_queue_name} Using '/' clearly doesn't play well in this scheme. Is there a strong reason to continue with '/', and not simply use '.' or some other separator?	1	1	https://issues.jboss.org/browse/IMMUTANT-96
7792	IMMUTANT-96	Right now, I believe the only naming requirement we have is that queues must start with /queue, and topics must start with /topic. We do that so we can recognize the type based on a naming convention. I'm fine with relaxing that, maybe to #"^.?queue". This would make queues starting with "/queue", ".queue", "queue", "xqueue" all legal. How does that sound?	1	1	https://issues.jboss.org/browse/IMMUTANT-96
7793	IMMUTANT-84	I pushed an update to lein-immutant today (0.8.2) that should resolve this. Can you try it and report back that it indeed works for you?	1	1	https://issues.jboss.org/browse/IMMUTANT-84
7794	IMMUTANT-78	Is this one still relevant?	1	0	https://issues.jboss.org/browse/IMMUTANT-78
7795	IMMUTANT-55	What does this even mean?	0	0	https://issues.jboss.org/browse/IMMUTANT-55
7796	IMMUTANT-27	Jim: Isn't this kind of done? Or are is this jira waiting on more formal Immutant support on OS?	1	1	https://issues.jboss.org/browse/IMMUTANT-27
7797	IMMUTANT-24	We already had this discussion in TB about 'at' API design. So I wrote a draft in Clojure. https://gist.github.com/184f7b2eb203df87727a Lemme know if I got the idea.	1	1	https://issues.jboss.org/browse/IMMUTANT-24
7798	IMMUTANT-21	Hey Jimbotron - is there anything else we need to do here?	1	1	https://issues.jboss.org/browse/IMMUTANT-21
7799	IMMUTANT-21	Hey Jimbotron - is there anything else we need to do here?	1	1	https://issues.jboss.org/browse/IMMUTANT-21
7800	ISPN-8004	I suppose sending the class name would solve the problem, but we want to send a numerical id instead to save some bandwidth. Could we use in internal cache as a replicated registry that maintains the mapping of encoder internal numerical id to class name (and the reverse)? Or what approach do you have in mind? I would like to avoid solutions that leave the assignment of ids to the users, like it happens with externalizers now.	0	0	https://issues.jboss.org/browse/ISPN-8004
7801	ISPN-7991	When I was working of REST refactoring (ISPN-7734) I noticed that we are not converting entries retrieved using REST interface. What we do instead is we take raw payload (as bytes) and store it [directly into the cache|https://github.com/infinispan/infinispan/blob/c8c632349a4899d8b6085329288f1615b0c2d33c/server/rest/src/main/java/org/infinispan/rest/operations/CacheOperations.java#L301-L301]. This is exactly the same thing the [old implementation was doing|https://github.com/infinispan/infinispan/blob/8.2.x/server/rest/src/main/scala/org/infinispan/rest/Server.scala#L364]. I'm not sure if this is the correct approach. But with JSON stuff coming soon I'm not sure if it matters. Do we want to put additional effort in fixing this bug or we maybe we should rely on transcoding?	1	1	https://issues.jboss.org/browse/ISPN-7991
7802	ISPN-7991	The current algorithm looks like the following: * We take the input data as bytes * [We create metadata for it|https://github.com/infinispan/infinispan/blob/c8c632349a4899d8b6085329288f1615b0c2d33c/server/rest/src/main/java/org/infinispan/rest/operations/CacheOperations.java#L314], * [We store it in cache with metadata|https://github.com/infinispan/infinispan/blob/c8c632349a4899d8b6085329288f1615b0c2d33c/server/rest/src/main/java/org/infinispan/rest/operations/CacheOperations.java#L325] By using metadata we know how to [convert it back]|https://github.com/infinispan/infinispan/blob/c8c632349a4899d8b6085329288f1615b0c2d33c/server/rest/src/main/java/org/infinispan/rest/operations/CacheOperations.java#L111-L112 when sending to the client. Now the biggest question that bothers me is what format needs to be used to make it working with compatibility? In case of String objects ({{content-type: text/plain}}), shall we put a String in cache? What about other types such as {{content-type: application/xml}}?	1	1	https://issues.jboss.org/browse/ISPN-7991
7803	ISPN-7991	And does it matter? I think the point is that no matter what you send via REST you need to pass a specific marshaller to the HotRod client to understand the data. That was the intended behaviour, IMO. But I'm not sure what you want to achieve exactly.	1	1	https://issues.jboss.org/browse/ISPN-7991

#	Requirement	Questions	Answered	Successfully Answered	URL
7804	ISPN-7990	bq. you won't know which indexes need to be cleaned up That's the issue this JIRA is trying to address. As I found out recently, there are some corner cases (considered rare, during state transfer) where it's not possible to know the type. bq. even cleaning up "all indexes" by id is troublesome as you might not have a full list to define "all indexes" I assume the QueryKnownClasses object would have all types stored? If not, doesn't the searchFactory know about all existent indexes? bq. without the type you don't know how to encode the key -> not matching the right key might lead to various mismatches, like delete the wrong data or not deleting it at all The key encoding is done by the KeyTransformationHandler which is cache-scoped; so I assume the id is consistent across all indexes. bq. what if the index was created by an instance having a different / out of date configuration How this could happen? I have a feeling if clusters have caches with outdated configurations indexing would not be the only problem	1	1	https://issues.jboss.org/browse/ISPN-7990
7805	ISPN-7990	[~sannegrinovero] Still waiting for your reply to my previous comments :) EDIT: Forgot to add, IMO operations across multiple search indices is not something too uncommon. Lucene has support for that and elasticsearch allows to use special placeholder "all" to delete from multiple indices or even use a regex to narrow them down [1]. And in Infinispan, we have the advantage of having unique document ids across all indices (it derives from the cache key). When you mention "unsafe", is it some internal limitation of the Hibernate Search engine, or the QueryKnownClasses (which is backed by a SYNC REPL cache) itself? [1] https://www.elastic.co/guide/en/elasticsearch/reference/current/multi-index.html	1	1	https://issues.jboss.org/browse/ISPN-7990
7806	ISPN-7990	bq. And in Infinispan, we have the advantage of having unique document ids across all indices (it derives from the cache key). When you mention "unsafe", is it some internal limitation of the Hibernate Search engine, or the QueryKnownClasses (which is backed by a SYNC REPL cache) itself? No I'm not referring to some obscure limitation in Hibernate Search. I was merely concerned about the reliability (especially over time and in exceptional cases) of {{QueryKnownClasses}}, but since Adrian clarified this is finally being taken care of I guess I have no more concerns. Incidentally you remind me of the terrible hack that Infinispan Query does to encode keys. We should see if we can improve on that, especially after we've dropped "auto-discovery" of index mapping.	1	1	https://issues.jboss.org/browse/ISPN-7990
7807	ISPN-7971	[~sebastian.laskawiec] What kind of warning did you have in mind? A dialog displayed once a user logs in for the first time?	1	1	https://issues.jboss.org/browse/ISPN-7971
7808	ISPN-7842	h5. High level questions # (?) I assume the intent that - for a given class type {{A}} - it's going to be mapped *either* via annotations *or* via explicit mapping? (I would suggest to expect that, at least for the first iterations, as we don't support integrating the two at the moment) # (/) (follow up on previous question) I hope you'll enforce (validate) against listing the same class in both ways # (flag) We will need to require the actual class definition to be "on classpath" during bootstrap of the Cache. I hope that's fine for the use case? h5. Property mapping # (-) Let's not expose index-time boost. We'll deprecate / remove this soon. # (?) I'm confused about the Spatial mapping example on _city_. Could you add the entity sources to clarify? i.e. you'll need matching {{@Latitude}} and {{@Longitude}} for each Spatial coordinates-set, unless these are embeddables but then you'll need nested mapping. # (i) _norms_, _term-vectors_, etc.. : in the annotations case we use enums. Let's implement the XML schema to match, i.e. don't accept any string. # Being able to support {{@FieldBridge}} is typically quite important. h5. Entity Properties vs Entity Fields Hibernate Search can map either properties and/or fields. We can read/write to Java fields, and this is supported by Infinispan Query as well. When having {code:xml} <property name="name" type="method"> <field ... {code} I guess _type_ refers to a getter? A "method" should typically match the full method name. Yet a "getter" is slang for *property* so the block already hinted about it. Should {{<property>}} be named differently, and have two different identifiers to get rid of the _type_ ? Should the inner _field_ be named _documentField_ or _indexField_ ? The {{@Field}} annotation is not ambiguous as one can see it's not a POJO field, but in this context it reads ambiguously. h5. Nested structures? What if my POJO contains a map/set of many (one/none/null) objects which need to be mapped ? In Hibernate Search we support the recursion into the related types, but the child element doesn't have to be {{@Indexed}} on its own. (!) Careful with polymorphic collections. h5. Analyzer definitions? I don't think we need to expose all of Hibernate Search features, but defining Analyzers is essential. h5. Mapping / converstion of Infinispan *keys* The concept of a separate key is not native in Hibernate Search, yet you need a way to define how key types being used are two-way converted to the index format. This is represented as a {{@ProvidedId}} in the Hibernate Search _internals_ but in this case you might want to expose them in a simpler form.	1	1	https://issues.jboss.org/browse/ISPN-7842

#	Requirement	Questions	Answered	Successfully Answered	URL
7809	ISPN-7842	don't want the index name to be customized, nor two entities to share the same index)? # I think support for class bridges is missing (I guess {{<property type="class">}}, without a property name, would work, but it's hardly elegant). See {{org.hibernate.search.annotations.ClassBridge}} and {{org.hibernate.search.cfg.EntityMapping.classBridge(Class<?>)}} in Hibernate Search. # You need some way to assign analyzers to fields (e.g. {{property name="foo" analyzer="myAnalyzer"}}). See {{org.hibernate.search.annotations.Field.analyzer()}} and {{org.hibernate.search.cfg.FieldMapping.analyzer(Class<?>)}} in Hibernate Search. # As Sanne mentioned, analyzer definitions are absolutely necessary. You may think they don't belong in entity mappings (and I could not agree more), but they should be _somewhere_, so that you can reference the definitions when defining index fields (see above). # As Sanne mentioned, there should probably be some support for embedding other objects in an entity's index, though you could probably add it later. And yes, embedded objects may be mapped to index fields without having their own index (therefore the "indexed-entity-mapping" tag probably wouldn't be a good fit for those). See {{org.hibernate.search.annotations.IndexedEmbedded}} and {{org.hibernate.search.cfg.PropertyMapping.indexEmbedded()}} as well as {{org.hibernate.search.annotations.ContainedIn}} and {{org.hibernate.search.cfg.PropertyMapping.containedIn()}} in Hibernate Search. # As Sanne mentioned, custom field bridges (and custom class bridges, but that's about the same) are an important feature and should probably be supported from the start. See {{org.hibernate.search.annotations.Field.bridge()}} and {{org.hibernate.search.cfg.PropertyMapping.bridge(Class<? extends FieldBridge>)}} in Hibernate Search. h4. Unadvisable features # As Sanne mentioned, it would be best not to include index-time boost, since it will probably disappear soon. On top of that, I seem to remember it's buggy: if you contribute to the same index field from multiple properties, each time with the same boost, I believe the boost will grow exponentially with the number of contributing properties. So... you've been warned. h4. Cosmetics # {{<property name="name" type="method">}} "type" may not be a very explicit name... maybe "access-mode" would be better? # Unless there's a specific reason for that, the "mode" attribute on the "spatial" tag should probably accept lowercase values instead of {{HASH}} and such, so it'll be consistent with other attributes. Also, it's called "spatial-mode" instead of just "mode" when used at a property level; you may want to use the same attribute name in both cases. # The values for the "term-vector" attribute in the "field" tag should probably be "*true/false*/with-offsets/with-positions/with-position-offsets" to be consistent with "norms", "store", etc. h4. Index-driven mapping Also, I'd like to point out that while annotation-based mapping is necessarily entity-driven, XML-based mapping could be index-driven, which may arguably be clearer when contributing to the same index field from multiple entities : {code} <indexed-entities> <!-- annotated entity --> <indexed-entity>org.infinispan.query.queries.faceting.Car</indexed-entity> <!-- non-annotated entity --> <index-entity-mapping index-name="people"> <!-- the FQN of the class to index --> <indexed-entity id="author">my.domain.model.Author</class> <indexed-entity id="fan">my.domain.model.Fan</class> <!-- optional --> <spatial name="place" mode="HASH"> <bridge entity="author"> <!-- class bridge --> </spatial> <!-- list of index fields --> <field name="name" store="true" index="true" analyze="true" norms="true" term-vector="true"> <bridge entity="author" property="name" access-mode="method" impl="my.package.MyBridge"> <bridge entity="fan" property="fullName" access-mode="method"> </field> {code} Of course, it would be a bit confusing for users switching over from annotation-based mapping, and it also may be a bit more challenging to implement. You may not want to do this for various reasons, but I wanted to mention it, since that's probably not something you'll be able to change	0	0	https://issues.jboss.org/browse/ISPN-7842
7810	ISPN-7426	We're reverting -the Hibernate Core- all upgrades since it breaks OSGI integration tests: {code} java.lang.Exception: Could not start bundle mvn:org.infinispan/infinispan-embedded/9.0.0-SNAPSHOT in feature(s) infinispan-embedded-9.0.0-SNAPSHOT: Unresolved constraint in bundle org.infinispan.embedded [46]: Unable to resolve 46.0: missing requirement [46.0] osgi.wiring.package; (&(osgi.wiring.package=javassist)(version>=3.20.0)(!(version>=4.0.0))) {code} The problem seems to be with entitymanager: {code} <lburgazzoli> 2017-02-07 18:15:53,882 | WARN | FelixStartLevel | FeaturesServiceImpl | 20 - org.apache.karaf.features.core - 3.0.4 | Error when installing feature hibernate: org.osgi.framework.BundleException: Jar is not a bundle, no Bundle-SymbolicName mvn:org.hibernate/hibernate-entitymanager/5.2.7.Final {code}	0	0	https://issues.jboss.org/browse/ISPN-7426
7811	ISPN-7418	Something to consider: * If the cache has a MIME type configured for key/value, should a cache write be able to override it? Maybe only if its the first write? Needs some thought...	1	0	https://issues.jboss.org/browse/ISPN-7418
7812	ISPN-7402	[~vblagojevic] The NPE you're getting now seems to indicate the service controller for the given cache-container is null, ie. the cache container is not running yet. Can you check ?	1	1	https://issues.jboss.org/browse/ISPN-7402
7813	ISPN-7359	[~ryanemerson] given that we can and we do use composite operations, should we perhaps close this issue?	1	0	https://issues.jboss.org/browse/ISPN-7359
7814	ISPN-7322	[~belaban] I disagree. Since we actually want to order only writes to each key (without respect to the other keys), full source ordering is excessive - one hiccup (lost/discarded message) delays all subsequent messages (until the problem is handled). About thread creation - do you mean that you deliver (to app) OOB messages received in one batch separately (one msg per thread)? Can we just switch this off and deliver the whole batch sequentially? The processing on backup should be non-blocking (not sure if there's something blocking left yet).	1	1	https://issues.jboss.org/browse/ISPN-7322
7815	ISPN-7276	[~ryanemerson] Why don't we just log out a user on server down event? That should take care of this and bunch of other potential transitory issues that will show up on client-server disconnect.	1	1	https://issues.jboss.org/browse/ISPN-7276
7816	ISPN-7224	hmmmm it's not that easy as you think guys... So let me try to explain my concerns a bit more... The {{Cache}} instance on which SpringCache support is baked can be re-used in many different places. There's nothing that prevents a user from grabbing an instance and locking the entire container for some reason. In that case {{Cache#get(Object, Callable<T>)}} would block forever. The second aspect is asynchronous replication. Imagine you cache is clustered and at the same time you supply object {{A}} on one node and object {{B}} on another one. In that case your clustered code will also invoke the supplier twice. The only thing we _could_ do is to provide "best effort" to synchronize on the key by locking it (with ability to disable this feature in configuration). This still doesn't give you any guarantees that your supplier won't be invoked twice and for me it's any better than {{computeIfAbsent}}. Would it solve your problem [~miken123]?	1	1	https://issues.jboss.org/browse/ISPN-7224
7817	ISPN-7120	bq. Reducing the number of tests may not be so easy... remember that we bq. need to test all versions of the protocol, not just the latest one. bq. And we still need to test stuff that's not explicitly in the protocol, bq. especially around state transfer/server crashes and around query bq. (which the protocol says almost nothing about). The above quote is from [~dan.berindei]. Should these also be included in the TCK?	1	0	https://issues.jboss.org/browse/ISPN-7120
7818	ISPN-7114	IspnPerfTest and RadarGun could use this after a run, to check if consistency is provided between cluster nodes. I'm thinking of a method in AdvancedCache that returns the number of inconsistent keys in a cluster. Perhaps it should return the keys directly (in a set)? Also, the (or an additional) method might return actual keys and values, for a reconciliation strategy to perform repair. Alternatively, as [~NadirX] mentioned, this could be done by shipping user code to the nodes to do repair... Can you propose some API and we'll discuss it on this JIRA?	0	0	https://issues.jboss.org/browse/ISPN-7114

#	Requirement	Questions	Answered	Successfully Answered	URL
7819	ISPN-7037	Here is a brief discussion about this topic, from the infinispan IRC channel, 22-SEP-2016 (EST): (11:09:36 AM) me@freenode: I know you are beyond busy so I don't really want to bug you about this stuff. However, do you recommend doing a unique quickstart based on the embedded and clustered quickstarts? the secure-embedded-cache quickstart only demonstrates granular security of cache nodes, and not really fully secured infinispan/jdg library mode clustering. (11:13:52 AM) ttarrant: mwigglesworth, what features would you like to highlight ? (11:14:30 AM) ttarrant: mwigglesworth, ideally such a quickstart should be able to run both with a SecurityManager and without one (while maintaining the same behaviour((11:15:25 AM) me@freenode: I would like to highlight the use of callbackhandlers to demonstrate how an end-user/developer would enable node-authentication, with infinispan/jdg. (11:16:45 AM) me@freenode: This would be more of a "secure clustering" example, however, since the secure-embedded-cache quickstart is appropriately labeled, and has part of the security enabled, I thought it may make sense to simply fork and update it. (which I have already started to do.0 (11:25:07 AM) jbossbot: new jira [ISPN-7037] Create quickstart that demonstrates library mode node security using callbackhandlers [New (Unresolved) Feature Request, Minor, Demos and Tutorials/Tasks, Unassigned] https://issues.jboss.org/browse/ISPN-7037 (11:34:34 AM) ttarrant: mwigglesworth, I'm not entirely familiar with JDF quickstart policies, but what about, instead of proliferating quickstarts, add to the existing one and enable the node auth via some switch / maven profile (11:35:13 AM) me@freenode: Sounds good.	0	0	https://issues.jboss.org/browse/ISPN-7037
7820	ISPN-6919	[~ksobolew] I'm afraid the functional API's commands can't run without exclusive access to the entries they modify, either. [~pruivo] numOwners == 1 should really be described separately. If we exclude that, the first case becomes "originator != primary" What happens when the backups install a new topology in which they're no longer owners, or if the backups crash? Also, why do we need the replication to the backups to be ordered? That won't help if the primary owner crashes, and another node becomes the primary owner. More general, we should really have some diagrams to discuss how the triangle is supposed to handle the various corner cases that appear when a node joins/leaves/crashes, or when there's a split+merge. I'm not saying this approach won't work with topology changes, it's just that I have no idea how it's supposed to work.	0	0	https://issues.jboss.org/browse/ISPN-6919
7821	ISPN-6797	Is this issue specific to HotRod, or do you plan to evaluate compression for RPCs/storeAsBinary/persistence as well?	1	0	https://issues.jboss.org/browse/ISPN-6797
7822	ISPN-6796	We don't have this information available. We can only access whatever DMR gives us. Closing it. If this is a requirement reopen another issue where this info is provided in DMR. If it is provided it will be automatically shown in error dialog	0	0	https://issues.jboss.org/browse/ISPN-6796
7823	ISPN-6706	Rereading what you wrote, I more understand your case. I would like to ask some more information though. You said the cluster is pausing due to the purge, this seems problematic of itself. The purge shouldn't stop other concurrent operations and in most cases it will be very heavily IO bound so it shouldn't be taking up CPU. How did the pause adversely affect your test? Can you also provide your configuration?	1	1	https://issues.jboss.org/browse/ISPN-6706
7824	ISPN-6706	Yeah that looks like the underlying issue. Can you try adding the index to the column and see if it fixes your issue? Either way it looks like we need a JIRA to add the index when creating the table.	1	1	https://issues.jboss.org/browse/ISPN-6706
7825	ISPN-6706	OK, I'm getting confused :) I too was under impression that the query to purge expired rows picks one node at a time by "manually" scanning the entire table, but I looked at the code more closely and the query indeed limits the rows selected to only those with timestamp in the past, and actually it is a DELETE statement. And indeed the CREATE TABLE statement doesn't create any indices beside the PK on the ID column. So it's much clearer now. Should I add another JIRA about the timestamp index?	1	0	https://issues.jboss.org/browse/ISPN-6706
7826	ISPN-6706	Thankfully I'm not :) So... it looks like these two should satisfy my use case (I hope at least ISPN-6710 is backported to 8.2.x!). Should this JIRA be closed then?	1	0	https://issues.jboss.org/browse/ISPN-6706
7827	ISPN-6677	Yes, retries are the only option. But I think we should implement reties in Core (this way we could use the same approach for both embedded and client server). WDYT? On Wed, Jan 18, 2017 at 10:27 AM, Gustavo Fernandes (JIRA) <issues@jboss.org	1	1	https://issues.jboss.org/browse/ISPN-6677
7828	ISPN-6676	Thoughts on design for this feature: * HTTP/2 [does not support upgrading to another protocol|https://http2.github.io/http2-spec/#rfc.section.8.1.2.2]. [TLS ALPN|https://http2.github.io/http2-spec/#rfc.section.3.1] should be used to ensure this type functionality. * HTTP/1.1 Upgrade header is supported to ensure backwards compatibility. In other words we should consider using [HTTP/1.1 Upgrade|https://http2.github.io/http2-spec/#rfc.section.3.2] or [ALPN|https://tools.ietf.org/html/rfc7301] to support this kind of functionality. * [HTTP/2 client with TLS support must use ALPN|https://http2.github.io/http2-spec/#rfc.section.3.4] * [HTTP/2 supports extensions, which can use new type of Frames, Settings and Error Codes|https://http2.github.io/http2-spec/#rfc.section.5.5]. * HTTP/2 supports [PING|https://http2.github.io/http2-spec/#rfc.section.6.7] and [CLOSE|https://http2.github.io/http2-spec/#rfc.section.6.8] operations out of the box. * [ALPN is considered as a protection for cross-protocol attacks|https://http2.github.io/http2-spec/#rfc.section.10.2] for HTTP/2	0	0	https://issues.jboss.org/browse/ISPN-6676

#	Requirement	Questions	Answered	Successfully Answered	URL
7829	ISPN-6673	The procedure for OpenShift looks like the following: # Start OpenShift cluster using: {code} oc cluster up {code} Note you are logged as a developer. # Create new Infinispan cluster using standard configuration. Later on I will use REST interface for playing with data, so turn on compatibility mode: {code} oc new-app slaskawi/infinispan-experiments {code} # Node that you should always be using labels for your clusters. I'll call my cluster=cluster-1 {code} oc env dc/infinispan-experiments OPENSHIFT_KUBE_PING_LABELS=cluster=cluster-1 oc env dc/infinispan-experiments OPENSHIFT_KUBE_PING_NAMESPACE=myproject oc label dc/infinispan-experiments cluster=cluster-1 {code} # Scale up the deployment {code} oc scale dc/infinispan-experiments --replicas=3 {code} # Create a route to the service {code} oc expose svc/infinispan-experiments {code} # Add some entries using REST {code} curl -X POST -H 'Content-type: text/plain' -d 'test' http://infinispan-experiments-myproject.192.168.0.17.xip.io/rest/default/1 {code} # Check if the entry is there {code} curl -X GET -H 'Content-type: text/plain' http://infinispan-experiments-myproject.192.168.0.17.xip.io/rest/default/1 {code} # Now we can spin up a new cluster. Again, it's very important to use labels which will avoid joining those two clusters together. The new Docker image with Infinispan Hot Rod server needs to contain the following Remote Cache Store definition: {code} 212,214d214 < <remote-store cache="default" hotrod-wrapping="false" read-only="true"> < <remote-server outbound-socket-binding="remote-store-hotrod-server" /> < </remote-store> 449,451c449 < <!-- If you have properly configured DNS, this could be a service name or even a Headless Service --> < <!-- However DNS configuration with local cluster might be tricky --> < <remote-destination host="172.30.14.112" port="11222"/> {code} # Spinning new cluster involves the following commands: {code} oc new-app slaskawi/infinispan-experiments-2 oc env dc/infinispan-experiments-2 OPENSHIFT_KUBE_PING_LABELS=cluster=cluster-2 oc env dc/infinispan-experiments-2 OPENSHIFT_KUBE_PING_NAMESPACE=myproject oc label dc/infinispan-experiments-2 cluster=cluster-2 oc expose svc infinispan-experiments-2 {code} # At this stage we have 2 clusters (the old one with selector {{cluster=cluster-1}} and the new one with selector {{cluster=cluster-2}}). Both should be up and running (check that with {{oc status -v}}). Cluster-2 has remote stores which point to Cluster-1. # Switch all the clients to the {{cluster=cluster-2}}. Depending on your configuration you probably want to create a new Route (if your clients connect to the cluster using Routes) or modify the Service. # Fetch all remaining keys from {{cluster=cluster-1}} {code} oc get pods --selector=deploymentconfig=infinispan-experiments-2 .. write down oc exec infinispan-experiments-2-3-pc7sg -- '/opt/jboss/infinispan-server/bin/ispn-cli.sh' '-c' '--controller=$(hostname -i):9990' '/subsystem=datagrid-infinispan/cache-container=clustered/distributed-cache=default:synchronize-data(migrator-name=hotrod)' {code} # Disconnect the source cluster {code} oc exec infinispan-experiments-2-3-pc7sg -- '/opt/jboss/infinispan-server/bin/ispn-cli.sh' '-c' '--controller=$(hostname -i):9990' '/subsystem=datagrid-infinispan/cache-container=clustered/distributed-cache=default:disconnect-source(migrator-name=hotrod)' {code}	0	0	https://issues.jboss.org/browse/ISPN-6673
7830	ISPN-6573	Answering to https://developer.jboss.org/message/955416#955416 ("Is this an issue only when cache store is involved, or are the functional operations broken in transactional mode completely?"): yes, it's broken in transactional mode completely. I'm attaching tests that illustrate this. FunctionalInMemoryNonTx passes, while FunctionalInMemoryTx fails.	1	0	https://issues.jboss.org/browse/ISPN-6573
7831	ISPN-6573	Okay. Couldn't you downgrade it to wishlist, though? :)	1	0	https://issues.jboss.org/browse/ISPN-6573
7832	ISPN-6537	Could you post your configuration? Are you using the modules distributed on http://downloads.jboss.org/infinispan/8.1.3.Final/infinispan-as-embedded-modules-8.1.3.Final.zip ?	1	1	https://issues.jboss.org/browse/ISPN-6537
7833	ISPN-6471	Isn't this a duplicate of ISPN-5083 ?	1	0	https://issues.jboss.org/browse/ISPN-6471
7834	ISPN-6457	I think the point of disabling authorization is to allow everything. So why disable access to these caches?	1	1	https://issues.jboss.org/browse/ISPN-6457
7835	ISPN-6433	Sorry Paul, Radoslav, but this slipped under the radar. By when do you need this?	1	1	https://issues.jboss.org/browse/ISPN-6433
7836	ISPN-6433	[~NadirX] It looks like it was erroneously closed - which is how it went unnoticed. Do you have a preferred process for flagging issues to backport to a specific branch?	0	0	https://issues.jboss.org/browse/ISPN-6433
7837	ISPN-6370	Ignored by hotrod connector?	0	0	https://issues.jboss.org/browse/ISPN-6370
7838	ISPN-6358	The entry interator tests are now actually testing the streams. Do we still have the same coverage after removing those tests?	1	1	https://issues.jboss.org/browse/ISPN-6358
7839	ISPN-6252	What do you mean by "manipulate it" and/or change the slot? Can this be done at runtime, on a per-deployment base? We can not patch the static definition of existing modules, that would break other applications. An alternative: as far as I remember the only module which needs both worlds in the same picture is the Infinispan DirectoryProvider for Hibernate Search: that's meant to link to the Hibernate Search instance included in WildFly/EAP, and therefore should be using its Lucene module (rather than the Infinispan/Lucene module). Considering that Infinispan Query has its own very specific needs on how it sets up the Directory, I'm wondering if we should rather not share the DirectoryProvider module for both Query and Hibernate Search. If you make the Infinispan DirectoryProvider compatible with Hibernate Search / WildFly only, and give Infinispan Query a similar alternative we should be good and the two projects would be able to have completely independent Lucene versions.	1	1	https://issues.jboss.org/browse/ISPN-6252
7840	ISPN-6252	> What do you mean by "manipulate it" and/or change the slot? Can this be done at runtime, on a per-deployment base? Infinispan query depends on Hibernate Search that depends on Lucene. Hibernate Search 5.6.0.Alpha1 (the version Infinispan points to now) is currently on Lucene 5.4. The only way to have Infinispan on 5.5 without a Hibernate search release would be to manipulate/hack the Hibernate Search modules zip on our side to include 5.5 instead of 5.4, but this is not very nice.	1	1	https://issues.jboss.org/browse/ISPN-6252
7841	ISPN-6252	> If you make the Infinispan DirectoryProvider compatible with Hibernate Search / WildFly only, > and give Infinispan Query a similar alternative we should be good and the two projects would > be able to have completely independent Lucene versions. That separation makes sense, but I think it'd would not allow Infinispan to have a Lucene version different from Hibernate Search: Infinispan Query using "ram" or "fs" directory will still depends on Hibernate Search that has its own exported Lucene version. I suppose it'd cause issues should Infinispan decides to have a dependency on another Lucene version. Am I right?	1	1	https://issues.jboss.org/browse/ISPN-6252
7842	ISPN-6252	HSearch alpha2 will use Lucene 5.5 (to be release really soo). So that mitigate the issue right? [~sannegrinovero] you lost me on your last comment though: no idea what you're talking about :)	1	0	https://issues.jboss.org/browse/ISPN-6252
7843	ISPN-6057	Migrate Cassandra cache store to the {color:#d04437}6.0{color} cachestore API Update Cassandra cache store to using Infinispan {color:#d04437}7/8{color} Is existed task only about 6x version?	1	0	https://issues.jboss.org/browse/ISPN-6057
7844	ISPN-6027	No! I deprecated this bundler and it will be removed in 4.0. Which perf test are you referring to?	1	1	https://issues.jboss.org/browse/ISPN-6027
7845	ISPN-6027	Was this with Radargun 3.0 which afair only focuses on latency? IMO, throughput should increase a lot by removing DONT_BUNDLE...	1	1	https://issues.jboss.org/browse/ISPN-6027

#	Requirement	Questions	Answered	Successfully Answered	URL
7846	ISPN-5962	That looks nice. But what does "acts much like a spin" mean exactly, will it still burn CPU time unnecessarily? We want this to lock pessimistically ASAP and free up the current thread.. actually encouraging context switching in this case.	1	1	https://issues.jboss.org/browse/ISPN-5962
7847	ISPN-5935	What's the new name you had in mind?	0	0	https://issues.jboss.org/browse/ISPN-5935
7848	ISPN-5807	[~pzapataf] What did we decide regarding this issue? Are you taking it over from me, are we postponing it or do you want me to address it for 8.2?	1	1	https://issues.jboss.org/browse/ISPN-5807
7849	ISPN-5704	Shall we provide order guarantees for Traversable returned from multi-key (evalMany methods) operations? Right now it's not guaranteed but it's implied.	0	0	https://issues.jboss.org/browse/ISPN-5704
7850	ISPN-5704	Use case: we want to add an element to a List and the list is stored as V in the functional map. You could try using this function: {code} <R> CompletableFuture<R> eval(K key, V value, BiFunction<V, ReadWriteEntryView<K, V>, R> f) {code} But the problem is that the element to add to the list can't be passed in {{value}} since V is List and not the individual element. Could we have a more flexible value type? Could we have this? {code} <R, T> CompletableFuture<R> eval(K key, T value, BiFunction<T, ReadWriteEntryView<K, V>, R> f) {code} The problem then would be that we could not call {{view.set(value)}}... It needs some thought.	0	0	https://issues.jboss.org/browse/ISPN-5704
7851	ISPN-5696	-Hi Tristan, I am still getting the same error i.e. Exception in thread "main" java.lang.NoClassDefFoundError: javax/transaction/TransactionManager I have tried with 'infinispan-7.2.5.Final-all' as I am using Java 1.7 Am i missing something ? Please help me out. Thanks in advance.- Working now. Thanks.	1	1	https://issues.jboss.org/browse/ISPN-5696
7852	ISPN-5545	is there any interest on this issue from other users ? actually on a large cluster (100 nearcache hotrod nodes) the nearcache floods the network	0	0	https://issues.jboss.org/browse/ISPN-5545
7853	ISPN-5535	Hi, could you please attach your wildfly configuration and you pom.xml?	0	0	https://issues.jboss.org/browse/ISPN-5535
7854	ISPN-5535	Can we see your MANIFEST.MF and jboss-deployment-structure.xml ? Infinispan jar files are not automatically exported to WildFly deployments and you need to explicitly add them as dependencies using the techniques described at https://docs.jboss.org/author/display/WFLY8/Class+Loading+in+WildFly	1	1	https://issues.jboss.org/browse/ISPN-5535
7855	ISPN-5520	Could you provide more information on circumstances, which should cause HR clients to use (be redirected to) backup site? Some ideas: 1. All nodes from the primary site go down 2. Administrator's intervention (outside of ISPN) - in fact this has the same effect as #1 as HR clients won't be able to connect to any server from the primary site 3. Status of primary site changes to "offline" - a possible way to avoid #2	1	1	https://issues.jboss.org/browse/ISPN-5520
7856	ISPN-5515	What guarantees do you have for this data to be valid? The definition of "first" seems rather subjective to me, and you have no idea how long ago that data was stored. Maybe this node has been down for the last month, while all other nodes worked and now this is being restarted.	1	1	https://issues.jboss.org/browse/ISPN-5515
7857	ISPN-5515	I think these are really tricky ideas which should be discussed on the mailing list, I noticed this JIRA by pure luck and find it concerning that such decisions are made without any wider discussion. What I find most tricky: - it's possible the new starting node starts while "thinking it's first", but then actually merge with a running cluster. The cluster detection protocols aren't foolproof, and you're relying on timeouts to be configured safely (when are they ever?). - it's unrealistic to push such a requirement to "admin's responsibility" especially but not least because node restarts might not be under their control - even with this design, the majority of cachestores are cleared so there is an assumption that "data loss is fine" for the user: so why even bother trying to keep a small portion of it at risk of consistency trouble? - this design seems to favour something else above correctness, and I'm not sure what "something else" you're aiming at.. why work hard to not wipe a single cachestore? I agree with you that this is an improvement over the current state, but I don't see why you would implement tricky code to provide a tricky solution when all what's needed is remove the preloading option from configuration. You'll be done in much less work and get a better reliable solution. Unless I'm missing the important reason to load this data?	1	1	https://issues.jboss.org/browse/ISPN-5515
7858	ISPN-5515	> I think these are really tricky ideas which should be discussed on the mailing list, I noticed this JIRA by pure luck and find it concerning that such decisions are made without any wider discussion. We already discussed this on the mailing list, and the conclusion was to implement [graceful restart|https://github.com/infinispan/infinispan/wiki/Graceful-shutdown-&-restore]. This issue is not really about implementing new functionality, it's about automating a recommendation we already have _for users who want it_. > it's possible the new starting node starts while "thinking it's first", but then actually merge with a running cluster. The cluster detection protocols aren't foolproof, and you're relying on timeouts to be configured safely (when are they ever?). If a node starts in a separate partition by itself, the behaviour with "purge on join" enabled will be exactly as it is now - not better, but not worse either. > it's unrealistic to push such a requirement to "admin's responsibility" especially but not least because node restarts might not be under their control A node restart will not affect nodes that are already running in any way. Also, this option will be disabled by default, and if the admin can't control the order in which nodes start, he should definitely not enable it. > even with this design, the majority of cachestores are cleared so there is an assumption that "data loss is fine" for the user: so why even bother trying to keep a small portion of it at risk of consistency trouble? In a replicated cache, it's not a small portion of the data, it's all the data. I agree that it makes a lot less sense in a distributed cache: in theory you could shut down the cluster such that all the state is transferred to a single node and all the data is preserved in that node's store, but it's definitely something you'd want to do on a regular basis. > this design seems to favour something else above correctness, and I'm not sure what "something else" you're aiming at.. why work hard to not wipe a single cachestore? These are my assumptions for using this option: * The cache is replicated and the number of nodes is small * Losing data is not fatal, as there is a backing store * Reading a stale value *is* fatal * Reading data from the canonical store is slow I realize these assumptions are quite narrow, and most users will not use it. But for applications who do fit these assumptions, I think this will help. And it would be back-portable to 7.2.x, unlike the graceful restart work. >I agree with you that this is an improvement over the current state, but I don't see why you would implement tricky code to provide a tricky solution when all what's needed is remove the preloading option from configuration. You'll be done in much less work and get a better reliable solution. I renamed the issue, since it's not really about preload - having preload does complicate the code, but stale values are possible with or without preload enabled.	1	0	https://issues.jboss.org/browse/ISPN-5515
7859	ISPN-5515	Thanks Dan, makes sense to consider replication too. But even then, why is it important for non-primary nodes to not load any entry? If any, I'd expect each node - even with REPL - to only load the entities which it is *primary owner* for. Consistency is an issue either way and I simply think one shouldn't trust any state from a shared CacheStore, but let's assume that some users really want this.. wouldn't you at least spread the load of loading operations from CacheStore among the nodes? I simply fail to see how picking any one node's CacheStore helps to improve on consistency.	0	0	https://issues.jboss.org/browse/ISPN-5515

#	Requirement	Questions	Answered	Successfully Answered	URL
7860	ISPN-5467	Is there a design document somewhere that describes this change ?	1	1	https://issues.jboss.org/browse/ISPN-5467
7861	ISPN-5467	[~belaban] it's probably worth a wiki page, but I'm going to add a brief description here :) I've just issued a new [PR|https://github.com/infinispan/infinispan/pull/4520], so that's what I'm going to describe. * Interceptors extending {{CommandInterceptor}} still work, and can use only {{invokeNext(ctx, command)}} to invoke the next interceptor synchronously. * The new base class is {{BaseAsyncInterceptor}}, which defines a single {{visitCommand(InvocationContext, VisitableCommand)}} entry point. * {{DDAsyncInterceptor}} extends {{BaseAsyncInterceptor}} to add the familiar {{visitXCommand(\[Tx\]InvocationContext, XCommand)}} entry points. * The return type of the {{visitXCommand}} methods is {{BasicInvocationStage}}, only allow you to extract a value or convert to a {{CompletableFuture}}. * {{BaseAsyncInterceptor}} defines some methods that return a {{BasicInvocationStage}}, or a more feature-full {{InvocationStage}} instance: ** {{invokeNext(ctx, command)}} works just like {{CommandInterceptor.invokeNext}}, except the invocation may be asynchronous. ** {{returnWith(value)}} returns a value, because we can't just return a value directly. ** {{goAsync(CompletableFuture<BasicInvocationStage>}} continues the invocation asynchronously. ** {{invokeNextAsync(ctx, command, future)}} is a shortcut for {{goAsync(f.thenApply(o -> invokeNext(ctx, command)))}}. ** {{returnWithAsync(value)}} is a shortcut for {{goAsync(f.thenApply(o -> returnWith(ctx, command)))}}. * {{InvocationStage}} extends {{BasicInvocationStage}} with a few more {{CompletionStage}}-like features. The main differences from {{CompletionStage}} are: ** All callbacks receive two extra parameters: the context and the command. In many interceptors, this allows us to store the callback in a field instead of creating a new instance on each invocation. ** Callbacks can throw {{Throwable}}, and received exceptions are not wrapped in `CompletionException`. I'm on the edge about this, but I'm sure a switch to `CompletionException` wrapping should be have a separate issue anyway. ** Not all {{CompletionStage}} methods are implemented, e.g. we don't support combining two invocation stages with {{thenCombine}}. ** We also have one extra method, {{compose((BasicInvocationStage, InvocationStage, VisitableCommand, Object, Throwable) -> BasicInvocationStage}}. It works like {{CompletionStage.thenCompose}}, but the callback gets called even if the invocation stage failed.	0	0	https://issues.jboss.org/browse/ISPN-5467
7862	ISPN-5414	Hi Adrain, A quick query on the scope of this implementation. Will we be avoiding parsing with this implementation. Will the parsed results be stored in query cache and named parameters be applied . We are working on similar modification and we can share our thoughts and design if its solving above two issues	1	1	https://issues.jboss.org/browse/ISPN-5414
7863	ISPN-5414	Can you please share the git url where we can view the code changes if coding is already done ?	1	1	https://issues.jboss.org/browse/ISPN-5414
7864	ISPN-5414	Thanks Adrian, We are presently implementing by keeping parsing result of antlr parser and storing it in Cache. Avoiding internals for AST Processor for now would give us easier way forward. Did some analysis around end to end speed after avoiding parsing we are able to achieve approx 4200 req/sec for 1 Million cache size which looks good for the moment. Our intention for now is not to make solution very generic which would be the final solution from your end but just to provide lookups from multiple keys stored in value part	0	0	https://issues.jboss.org/browse/ISPN-5414
7865	ISPN-5310	Is this another migration from hibernate-search codebase, more specifically [1]? [1] https://github.com/hibernate/hibernate-search/blob/master/modules/src/main/modules/search/infinispan-storage/module.xml	1	1	https://issues.jboss.org/browse/ISPN-5310
7866	ISPN-5310	Please take a look at the preview PR. The module is being produced solely on the migrated hibernate search directory provider and is called {{org.infinispan.hibernate-search.directory-provider}} slot {{ispn-7.2}}. Do you have a hard requirement on the name {{org.hibernate.search.infinispan-storage}} ?	0	0	https://issues.jboss.org/browse/ISPN-5310
7867	ISPN-5309	data consistency semantics in the absence of membership changes on the simplest of caches: no data store, no transactions, no passivation nor eviction. These are all refinements of the more basic model. > > I have a few comments/clarifications on the following issues: > - the expected data consistency for interleavings of gets and puts > - use of cache topology and CH in replicated mode and distributed mode > - some more detail on the processing of get and put operations in the context of a set of infinispan cache instances in distributed mode > - message order assumptions > > > data consistency > ---------- > In order to describe input sequences of get/put operations and the expected output values they return after processing by Infinispan, > I assume the following: > > - operations to be considered are ordered (or not) by the happens-before relation based on causality (as opposed to wall clock time) > - initial data consistency model for get and put operations: > - get(): operation that happens-after a get operation get(k) -> v sees the same value v (or the value of a later write) > - put(): operation that happens-after a write operation put(k=v) will see the updated value v (or the value of a later write) > > NB: this rules out any required data consistency between operations initiated at different nodes, even though one may occur after the other based on wall clock time This sounds good. Because a single client can connect to multiple nodes, in the future we should also model consistency for operations that happen on different but do not overlap (the first operation completely finishes before the other operation starts). > replicated mode vs distributed mode > ---------- > I'm a little unclear on the way in which replicated mode and distributed mode > make use of cache topology information (which we can assume > has been correctly distributed to all nodes): > > - do both of these modes make use the current cache topology (containing the CH) to lookup nodes corresponding to keys they need to access? Yes > - is the only difference in the way that the CH is set up (so that for cluster size N, there is one primary owner and N-1 backup owners)? There is one more difference: replicated mode always uses broadcasts instead of unicasts. It should only be relevant when the JGroups view changes, though. > - I assume only one CH is used in regular mode (as opposed to rehashing mode) Yes, we only have 2 CHs during state transfer. > distributed mode get()/put() design > ---------- > Here is a summary of the way in which get and puts are processed: > > (taken almost verbatim from section 4 of Consistency Guarantees in Infinispan, reorganized a little) > > get(): > - when an originator performs a get(K): > - if originator == owner, return the local value > - if originator != owner, originator sends a remote GET_REQ command to > all its owners in read-CH and uses the first successful non-null > response (GET_RSP) Correct. We do a retry when the topology id changes, but we can ignore that for now. > put:() > - when an originator performs a put(K,V): > - if originator == primary owner > - primary owner acquires the key lock, and replicates (PUT_REPL_REQ) > the write to the backup owner(s) > - backup owners confirm the write to the primary owner (PUT_REPL_RSP) > - primary owner updates the value locally, unlocks the key lock, and > replies to the originator > - if originator != primary owner > - originator will send a remote PUT_REQ to the primary owner of the > modified key (looked up from which CH?) > - primary owner acquires the key lock, and replicates (PUT_REPL_REQ) > the write to the backup owner(s) > - backup owners confirm the write to the primary owner (PUT_REPL_RSP) > - primary owner updates the value locally, unlocks the key lock, and > replies to the originator (PUT_RSP) Correct. > Are there any corrections to these descriptions that you would need to make? > Do you have the actual message types for these messages passed back and > forth? Remote get uses ClusteredGetCommand for replication, and the response is a SuccessfulResponse with the entry value (in some cases with an InternalCacheEntry, but we can ignore that for now). Remote put uses PutKeyValueCommand, both for the originator to talk to the primary and for the primary to talk to the backup owners. The responses are again SuccessfulResponses,	1	1	https://issues.jboss.org/browse/ISPN-5309
7868	ISPN-5298	Do we really want more than millisecond resolution? I doubt we can get millisecond precision with a remote distributed cache, anyway, as the write operation is very likely to take more than 1 millisecond.	1	1	https://issues.jboss.org/browse/ISPN-5298

#	Requirement	Questions	Answered	Successfully Answered	URL
7869	ISPN-5298	Ion, did you get a reply to the JSR issue? I'd wait to get an answer for that before making any immediate changes :| If we really need to implement this, I think Dan's approach is easiest. IOW, adding flags indicating the granularity of the expiration parameter. Seconds and Milliseconds would be enough. Remember what the protocol says: bq. Number of seconds that a entry during which the entry is allowed to life. If number of seconds is bigger than 30 days, this number of seconds is treated as UNIX time and so, represents the number of seconds since 1/1/1970. If set to 0, lifespan is unlimited. Seconds is enough to be able to represent any other expiration that's above it. Let's not complicate things beyond what's required :)	1	1	https://issues.jboss.org/browse/ISPN-5298
7870	ISPN-5298	The TCK issue opened: https://github.com/jsr107/jsr107tck/issues/76 . Based on the reply I think the tests might change to address the issue but not sure when.	1	1	https://issues.jboss.org/browse/ISPN-5298
7871	ISPN-5293	ISPN-5219 has two 'extra' needs when it comes for iterating: * Possibility to configure the batch size * Possibility to pass in the segments of interest * (Potentially) Pass in the server(s) of interest (shortcut for the above) Do you plan to cover these?	1	1	https://issues.jboss.org/browse/ISPN-5293
7872	ISPN-5246	Could you elaborate a bit about this change? (Or add a link to related -dev discussion, or link related JGroups JIRA)	1	1	https://issues.jboss.org/browse/ISPN-5246
7873	ISPN-5234	Isn't this already implemented in cluster-model.services.js, function getAvailability - we need to implement it though? It needs to actually check availability rather than simply returning true.	1	1	https://issues.jboss.org/browse/ISPN-5234
7874	ISPN-5234	Update: I don't understand exactly what functionality to imagine for these requests: - Filter by metrics (node view) - Filter by metrics (cache view) Can we elaborate and discuss it a liitle bit here? Meanwhile I -am working on- finished: - Search node (search as you type) - Search cache (search as you type) -I will issue PR with changes ONLY related to those 2 requests that is easy to understand and review. It should be done today (Tuesday) or tomorrow EOD at latest (Wednesday).- PR sent.	1	1	https://issues.jboss.org/browse/ISPN-5234
7875	ISPN-5167	Takayoshi, with ISPN 8 we have streams. Internally their API defines .count which returns a long. Is it sufficient for you to just call {code} cache.keySet().stream().count(){code} to get the size as a long?	0	0	https://issues.jboss.org/browse/ISPN-5167
7876	ISPN-5159	I can't remember precisely in which test this was causing trouble - probably in larger cluster tests, then. [~mcimbora], could you try to run a simple test without staggered start in LLNL, with e.g. 100 nodes?	1	0	https://issues.jboss.org/browse/ISPN-5159
7877	ISPN-5147	When we configure a new cache under the cache manager in JON, the cache is automatically started immediatelly once the configuration is finished. Just FYI. What we are still missing however is the possibility of changing configuration "on-fly" without restarting the server. Is that ^ what we are trying to achieve in this JIRA?	0	0	https://issues.jboss.org/browse/ISPN-5147
7878	ISPN-5131	{quote} Indeed you don't want to limit users deploying instances of AdvancedLoadWriteStore only. You want to support deploying AdvancedCacheLoader, CacheLoader, AdvancedCacheWriter, CacheWriter and ExternalStore as well. This should be tested to make sure it works as expected. Ideally, you'd make sure you can support all of these with the least amount of internal changes. It will require though supporting multiple META-INF/services/xxx definitions based on which interface the user decides to implement. {quote} I would like to propose slightly different approach. Let's create new interface called CacheStoreFactory: {code} public interface CacheStoreFacory { public CacheLoader createCacheLoader(); public CacheWriter createCacheWriter(); } {code} This interface will need to be implemented in deployable cache store. {{PersistenceManagerImpl}} would use such factories to obtain {{CacheLoader}} and {{CacheWriter}} instances. I think the biggest advantage of this approach is that we have only one service responsible for deployable Cache stores. The client has also the same level of flexibility (one might implement only one {{ExternalStore}} and return the same instance in both methods. The biggest drawback is slightly bigger refactoring in core (but on the other hand we guarded by a lot of unit/interaction tests there) Does it sound reasonable or would you still vote for multiple services?	1	1	https://issues.jboss.org/browse/ISPN-5131
7879	ISPN-5131	Not 100% sure about the factory idea. Although I like the simplification, what happens when you have both a cache loader and writer deployed? The bigger problem is how to tie a deployed cache loader/writer with the configuration itself. How are you going to say, "x" cache loader needs to be plugged into "y" cache? I really think you should write up your design in a design wiki (see https://github.com/infinispan/infinispan/wiki), where it's easier for others to collaborate and comment on the specifics. Also, once the wiki is complete, send it to the dev list to get everyone's input there.	1	1	https://issues.jboss.org/browse/ISPN-5131
7880	ISPN-5103	> Also: if we were to ban such "mixed types" from the same index we could use a straight forward Lucene Update command rather than applying the two commands Delete + Add Yes, that's what I described above :) But we don't necessary need to ban mixed types, it works nicely when the source data lives all in the same cache. The corner case (described above) is when entities with the same Id living in different caches share the same index. Is there any other scenario where delete by id is unsafe?	1	1	https://issues.jboss.org/browse/ISPN-5103
7881	ISPN-5083	Could you elaborate in what situations this should help? It seems to me that you only delegate the work to another thread pool, which will get exhausted soon. What do you intend to use as rejection policy?	1	0	https://issues.jboss.org/browse/ISPN-5083
7882	ISPN-5083	Maybe I realized what you might mean by this. Do you expect that a thread will read the request, issue asynchronous operation and the thread handling this async operation will later write the response to the socket? That would allow server-side multiplexing of several HotRod RPC calls over one connection, which is definitely a good thing! Though, in order to fully utilize such power, you'd need the client to multiplex the calls over this single socket, too. This is what SpyMemcached client does, IIRC, and our server implementation prohibits efficient use of socket connections due to the fact that the connection is owned by thread processing the HotRod call.	1	0	https://issues.jboss.org/browse/ISPN-5083
7883	ISPN-5073	[~NadirX] Isn't this issue obsolete now? It has to do more with the core then the admin console. Would you please update the status?	0	0	https://issues.jboss.org/browse/ISPN-5073
7884	ISPN-5042	Hey [~dan.berindei]! Is this issue still valid? I'm interested in this topic so I would be happy to implement some features here...	1	1	https://issues.jboss.org/browse/ISPN-5042
7885	ISPN-5035	I don't think put operations should be allowed to return arbitrary values, depending on the type of the value inserted into the cache. Couldn't you use {{DistributedExecutorService.submit(task, key)}} instead?	1	1	https://issues.jboss.org/browse/ISPN-5035
7886	ISPN-5035	I believe if the cache is not transactional DeltaAware isn't atomic either, unless the applying the Delta is an idempotent operation (as is the case with AtomicHashMap). Incrementing a counter is definitely not an idempotent operation... Can you expand a bit on how you expect {{cache.replace(key, deltaAwareExpectedValue, deltaAwareNewValue)}} to work?	1	1	https://issues.jboss.org/browse/ISPN-5035

#	Requirement	Questions	Answered	Successfully Answered	URL
7887	ISPN-5016	{quote}Overlapping operations can happen in any order{quote} What does 'happen' mean? I guess that eventually the state will be observed with those operations in certain order but does it mean that all threads will observe the writes in the same order? I think that JMM does not guarantee anything like that, and we should put emphasis on this if it's the case for Infinispan, too. {quote}Note that the transaction manager will hide the exception when the cache is registered as a synchronization (BATCH and NON_XA).{quote} Are you sure about this? I thought that write skew check will fire in prepare phase, and exceptions from prepare are not hidden neither in synchronization mode. {quote}maxIdle behaves similarly, except the reference timestamps are from the last write OR read operation.{quote} We could reference here ISPN-3202, expiration through idle time is not correctly maintained in cache store. {quote}In transactional caches, it also suspends the current transaction, and the next read in the transaction will not see the value.{quote} Sounds a bit like you had to resume the transaction manually after putForExternal read.	1	0	https://issues.jboss.org/browse/ISPN-5016
7888	ISPN-5016	>> Overlapping operations can happen in any order > What does 'happen' mean? I guess that eventually the state will be observed with those operations in certain order but does it mean that all threads will observe the writes in the same order? I think that JMM does not guarantee anything like that, and we should put emphasis on this if it's the case for Infinispan, too. I think JMM does guarantee that, if you consider each cache entry to be a volatile variable. >> Note that the transaction manager will hide the exception when the cache is registered as a synchronization (BATCH and NON_XA). > Are you sure about this? I thought that write skew check will fire in prepare phase, and exceptions from prepare are not hidden neither in synchronization mode. You're right, I'll remove it. >> maxIdle behaves similarly, except the reference timestamps are from the last write OR read operation. > We could reference here ISPN-3202, expiration through idle time is not correctly maintained in cache store. Sure. >> In transactional caches, it also suspends the current transaction, and the next read in the transaction will not see the value. > Sounds a bit like you had to resume the transaction manually after putForExternal read. I'll add "for the duration of the operation".	1	1	https://issues.jboss.org/browse/ISPN-5016
7889	ISPN-5016	{quote}Because data is stored in the shared store after the invalidation command was executed on all the nodes, it's also possible for a node to read the old value from the shared store after it has executed the invalidation.{quote} I don't understand the reason very well, could you elaborate a bit more? {quote}Invalidations triggered by write operations are synchronous and the operation will fail with a SuspectException if another node crashes. The entry will not be updated on the originator, but invalidation will still be performed on the non-crashed nodes.{quote} I thought SuspectExceptions are silently handled, aren't they? {quote}For multi-key operations like putAll(), the operation is sent to all the primary owners of the affected keys. When there is a single primary owner for all the affected keys, the operation succeeds and its effects are visible atomically. When there are multiple primary owners, the operation is not atomic: overlapping read operations will see random subsets of the updated values.{quote} As reads don't acquire any locks, I think that even update on single primary owner can be seen as non-atomic.	1	0	https://issues.jboss.org/browse/ISPN-5016
7890	ISPN-5016	{quote}>> Overlapping operations can happen in any order > What does 'happen' mean? I guess that eventually the state will be observed with those operations in certain order but does it mean that all threads will observe the writes in the same order? I think that JMM does not guarantee anything like that, and we should put emphasis on this if it's the case for Infinispan, too. I think JMM does guarantee that, if you consider each cache entry to be a volatile variable. {quote} OK, my bad, I was not inducing the ordering rules properly...	1	0	https://issues.jboss.org/browse/ISPN-5016
7891	ISPN-5016	{quote}If a node is suspected because of a Full GC, it might go from the initial JGroups view straight to the merge view. If that happens, its topology will be the largest one, and it will not be wiped, neither will it receive new data. Instead, it will keep the (possibly stale) entries it had before the Full GC.{quote} Sorry, this is not very clear to me. Why largest topology ID and why wouldn't that propagate its state to all other members (wiping their data)? {quote}If at least half of the nodes in the stable topology leave in quick succession{quote} What is quick succession? Does that depend on JGroups view installation, any timeout or the duration of rebalance? {quote}And if some of the nodes in the Available partition's consistent hash are not really accessible after the merge, the cache might stay Degraded.{quote} 'the cache'? We should be talking rather about partitions, not the whole cache. I understand the first part of the special case, but not the latter. {quote}If a node joins and becomes a backup owner after a write command was sent from the primary owner to the backups, but before the primary owner updates its own data container, it may not receive the value neither as a write command nor via state transfer.{quote} Sounds like a bug to me - is there a JIRA that could be linked? We could tolerate inconsistencies when the node crashes (if we can't fix it), but join or graceful leave should keep the cluster consistent. {quote}When a write to the store fails, it will fail the write operation.{quote} Does this hold for write-behind, too? {quote}With write-behind or asynchronous replication enabled, store write failures are hidden from the user (unless the originator is the primary owner, with async replication).{quote} When originator == primary and write-behind, how can the failure to write to the store be propagated to the user? I thought that user thread initiates both write to the store and async replication and returns. {quote}When the partitions merge back, there is no effort to replicate the values from one partition to another.{quote} Why is that different from non-tx mode, where the partitions with non-highest topology id are wiped? Moreover, in optimistic tx you write {quote}Same as with pessimistic and non-transactional caches.{quote} - what version, then? {quote}If the primary owners of the keys written by the transaction are all in the local transaction, {quote} local partition? {quote}If one partition stays available, its entries will replace all the other partitions' entries on merge, undoing partial commits in those partitions.{quote} Do I understand correctly that degraded partition may commit transaction, and this transaction will be later ignored (the data will later be overwritten to the previous values). Why is this behaviour desired? {quote}Transactions already prepared, however, will commit successfully even in minority partitions{quote} Is that true even if the originator is not in this minority partition? {quote}When a transaction needs to acquire more than one key lock with the same primary node, they are always acquired in the same order, so this will not cause a deadlock.{quote} If the keys have the same hashCodes, they can be locked in different order, though. ISPN-2491 {quote}The commit is always synchronous on the originator, so a transaction T3 started on node A after T1 finished will see T1's updates.{quote} Will it see all T1's updates, or just updates on those entries owned by A? {quote}The write to the attached cache store(s) is performed during the one-phase prepare command or the commit command, depending on the configuration.{quote} What configuration, exactly?	1	0	https://issues.jboss.org/browse/ISPN-5016

#	Requirement	Questions	Answered	Successfully Answered	URL
7892	ISPN-5016	{quote}TODO What happens if a node is a backup owner for two of the keys?{quote} Hmm? The primary owner sends the PutMapCommand to all nodes that are backup owners of at least one of the keys primary-owned by this node. But you're right that I am not sure whether this command makes the backup node write only {quote}When a node tries to read a key it does not own, it sends a remote get{quote} AFAIK the node checks if the entry is in data container, not whether it is owned. Regarding the situation that an entry is lost, I think that you should note somewhere that further attempts to read it will simply return null, not throwing any exceptions (the user could expect that). I lack explicit resolution on ISPN-4995, regarding reading and writing two (or more) entries and expectation whether any read is already visible. And in the end, I think that there should be a summarized table for the modes and ticks or crosses whether the cache will stay consistent after originator crash, primary owner crash, timeout exception/another exception. I think this is the best design document on Infinispan cache API so far. Thanks a lot, Dan, for the effort to put it together. I hope at least some users will be willing to read it (as it took me several hours just to go through it).	1	0	https://issues.jboss.org/browse/ISPN-5016
7893	ISPN-5016	shared store after it has executed the invalidation. > I don't understand the reason very well, could you elaborate a bit more? Rephrased. >> Invalidations triggered by write operations are synchronous and the operation will fail with a SuspectException if another node crashes. The entry will not be updated on the originator, but invalidation will still be performed on the non-crashed nodes. > I thought SuspectExceptions are silently handled, aren't they? Nope, they are silently handled only in some cases (e.g. L1 invalidation commands), but not in all cases (e.g. prepare commands, or regular invalidation commands). It does sound like a bug, I'll create one. >> For multi-key operations like putAll(), the operation is sent to all the primary owners of the affected keys. When there is a single primary owner for all the affected keys, the operation succeeds and its effects are visible atomically. When there are multiple primary owners, the operation is not atomic: overlapping read operations will see random subsets of the updated values. > As reads don't acquire any locks, I think that even update on single primary owner can be seen as non-atomic. Yeah, I was trying to say that it preserve the semantics of a local cache, but multi-key operations in a local cache aren't atomic either. >> If a node is suspected because of a Full GC, it might go from the initial JGroups view straight to the merge view. If that happens, its topology will be the largest one, and it will not be wiped, neither will it receive new data. Instead, it will keep the (possibly stale) entries it had before the Full GC. > Sorry, this is not very clear to me. Why largest topology ID and why wouldn't that propagate its state to all other members (wiping their data)? It's because the topology with the most nodes is elected as the merge topology. It seemed like a good idea when I implemented it, but perhaps it would have been better to use the one with the largest topology id instead. >> If at least half of the nodes in the stable topology leave in quick succession > What is quick succession? Does that depend on JGroups view installation, any timeout or the duration of rebalance? It means before the stable topology is updated again. In replicated mode there is no state transfer when a node leaves, that means the coordinator just needs to receive the rebalance confirmation from all the nodes. >> And if some of the nodes in the Available partition's consistent hash are not really accessible after the merge, the cache might stay Degraded. > 'the cache'? We should be talking rather about partitions, not the whole cache. I understand the first part of the special case, but not the latter. Changed to 'the merged partition'. >> If a node joins and becomes a backup owner after a write command was sent from the primary owner to the backups, but before the primary owner updates its own data container, it may not receive the value neither as a write command nor via state transfer. > Sounds like a bug to me - is there a JIRA that could be linked? We could tolerate inconsistencies when the node crashes (if we can't fix it), but join or graceful leave should keep the cluster consistent. I don't think it's any worse than the inconsistency in the previous paragraph. >> When a write to the store fails, it will fail the write operation. > Does this hold for write-behind, too? That's the next paragraph... >> With write-behind or asynchronous replication enabled, store write failures are hidden from the user (unless the originator is the primary owner, with async replication). > When originator == primary and write-behind, how can the failure to write to the store be propagated to the user? I thought that user thread initiates both write to the store and async replication and returns. Reworded, I meant with async replication but without write-behind. >> When the partitions merge back, there is no effort to replicate the values from one partition to another. > Why is that different from non-tx mode, where the partitions with non-highest topology id are wiped? Moreover, in optimistic tx you write >> Same as with pessimistic and non-transactional caches. > - what version, then? It's actually the same, I added that to the pessimistic mode description. >> If the primary owners of the keys written by the transaction are all in the local transaction, > local partition? Fixed. >> If one partition stays available, its entries will replace all the other partitions' entries on merge, undoing partial commits in those partitions. > Do I understand	0	0	https://issues.jboss.org/browse/ISPN-5016
7894	ISPN-4919	Does this suggest that cache xml attributes will no longer have default values? Let's say I define a template cache that defines <transaction isolation="REPEATABLE_READ"/> and a cache extension that defines <transaction mode="BATCH"/>. What then is the transaction isolation of the configured cache? The xsd says that transaction isolation defaults to READ_COMMITTED, meaning that <transaction isolation="READ_COMMITTED"/> and <transaction/> should be equivalent xml. However, I think this jira means to imply that the cache will use the "inherited" REPEATABLE_READ isolation. Thus, for the xsd to be correct, it will have to use 2 sets of attribute groups for each cache child resource, one that uses explicit default values (for cache templates) and one without defaults (meaning, the default value is implied by the configured template). That seems confusing to me. Consider too the child resources of a cache. A given resource (e.g. transaction), is fully defined, how would you determine which attributes meant to be "overrides" of the template and which are not, in particular when the template uses a non-default value, and the cache wants to override the value to match the original default?	1	1	https://issues.jboss.org/browse/ISPN-4919
7895	ISPN-4919	[~NadirX] I'm not arguing against the concept, just the usability/implementation. It's not just the schema that implies different behavior, but the Resources themselves. The cache resource and its children don't have nullable attributes. Every attribute of a cache and its children (e.g. transaction, locking, etc.) have default values that are applied to the model if no value was specified via a given operation. So, from the perspective of an add/write-attribute operation, there's no simple way to distinguish a value that is intended to be overridden from one that is intended to be inherited -- /unless/ you create parallel resource definitions for templates vs extending caches, where the resource attribute definitions for extending caches are completely nullable. In general though - the Configuration objects in 4.x did support the concept of overrides - where every property of the cache might be undefined, and thus take a default from some template or default cache. You guys changed this in 5.x-7.0, where Configuration objects are now completely defined. new ConfigurationBuilder().read(configuration) is semantically very different than the 4.x behavior, since a new configuration builder is already completely defined, and the read(Configuration) operation is effectively a means for cloning the entire configuration. Overriding is only every supported via programmatic calls to the resulting configuration builder - which is now a clone of the read configuration. It sounds like you guys want to revert the current Configuration behavior to be more like the 4.x configuration behavior. Is that a fair summation?	1	1	https://issues.jboss.org/browse/ISPN-4919
7896	ISPN-4754	Has the hibernate search slot "5.0" already been published?	1	1	https://issues.jboss.org/browse/ISPN-4754
7897	ISPN-4688	Hi Dan, it looks like this ISPN-2255 is not exclusively related to FGAHM but also to AHM which we use in EAP. Could you please check and possibly rebase?	1	1	https://issues.jboss.org/browse/ISPN-4688

#	Requirement	Questions	Answered	Successfully Answered	URL
7898	ISPN-4584	Maybe by rethinking how much of these settings the user should see and for how many use cases. Could we impose or default the names of the index data, metadata and locking caches to <owning-cache>_index_data, <owning-cache>_index_metadata, <owning-cache>_index_locking? Could we infer most of their settings, should they all be interlocked (same time out?, some are always REPL, some are always DIST?). Based on that you can infer things and end up with something like {code} <namedCache name="dist_lucene"> <transaction transactionMode="NON_TRANSACTIONAL"/> <clustering mode="DIST"> <hash numOwners="2" /> </clustering> <locking lockAcquisitionTimeout="3000" concurrencyLevel="1000" /> <indexing enabled="true" indexLocalOnly="true"> <properties> <!-- can we infer this, or make it a more compact setting before |properties| maybe an attribute of |indexing storage="infinispan"| <property name="default.indexmanager" value="org.infinispan.query.indexmanager.InfinispanIndexManager" /> <property name="default.exclusive_index_use" value="false" /> <!-- names inferred from the owning named cache, number of owners inferred from the owner's clustering.hash[numOwners], --> <! should we offer the ability to set replTimeout="30000" or can it be inferred? --> <property name="lucene_version" value="LUCENE_36" /> </properties> </indexing> </namedCache> {code} This approach does not replace the old model where all the guts and gory details are exposed. It just makes it much simpler for new comers that don't need to control and fine tune every quark of their body like Dr Manhattan and Sanne. The drawback is that if you need fine tuning, the complexity grows pretty quickly but that's the price of the backward compatibility. Note that most of this trick can probably be implemented in the configuration parser as a front end that spits out what the fully detailed configuration would have been.	0	0	https://issues.jboss.org/browse/ISPN-4584
7899	ISPN-4547	2 different settings? Wouldn't just the max queue size be sufficient? Or what are you thinking the other would be?	1	0	https://issues.jboss.org/browse/ISPN-4547
7900	ISPN-4534	I don't understand what you're suggesting. Do the embedded types have to be in persistence.xml? Do I have to somehow register those types?	0	0	https://issues.jboss.org/browse/ISPN-4534
7901	ISPN-4449	The PR's are integrated now. Can we close this or do you plan to add more?	1	1	https://issues.jboss.org/browse/ISPN-4449
7902	ISPN-4296	The limitation is restored successful only for REPL caches, as shown in [{{ClientAsymmetricClusterTest}}|https://github.com/infinispan/infinispan/blob/master/client/hotrod-client/src/test/java/org/infinispan/client/hotrod/ClientAsymmetricClusterTest.java] But for DIST caches, a nonexistent cache is started anyway in order to get hold of the {{ConsistentHash}} to write the topology update, as shown in [Encoder2x|https://github.com/infinispan/infinispan/blob/master/server/hotrod/src/main/scala/org/infinispan/server/hotrod/Encoder2x.scala#L100] If in the test we change from REPL to DIST, it fails, not throwing {{CacheNotFoundException}} anymore. Was this intentional? Shouldn't DIST behave the same way as REPL?	0	0	https://issues.jboss.org/browse/ISPN-4296
7903	ISPN-4229	If the question is: "Would it be enough to back port the fix for ISPN-3006?" I would say that I need to test on all supported dbs but I think that it could be enough if and only if we don't change the current default schema in the datasource configuration. For example on db like DB2, you can also use the same database and simply change the current default schema thanks to jdbc url of type jdbc:db2://MYHOST:MYPORT/MYDB:currentSchema=MYSCHEMA; On oracle it is also possible but it is a little bit more tricky https://community.oracle.com/thread/1038211?start=0&tstart=0	0	0	https://issues.jboss.org/browse/ISPN-4229
7904	ISPN-4144	Ok, I was recommend on our IRC chat that clear() would do what I wanted. There is a way to disable rebalancing but that call is only available through JMX which is ok I guess but it is weird to call system internal function through JMX. [~dan.berindei] how can we call LocalTopologyManagerImpl#setRebalancingEnabled as this method is not available through LocalTopologyManager interface or should we even do it to being with? What are you suggestions here?	1	1	https://issues.jboss.org/browse/ISPN-4144
7905	ISPN-4139	Is this preventing you to upgrade Search to Lucene 4.7? {quote}Might be able to create pull request, once I understand the whole Lucene 3 + 4 support thing ...{quote} Lol, that's the reason for me to work on including Search 5 in Infinispan asap: to get rid of that (remove the Lucene3 support) :-) I can take care of this, it's quite a minefield of backwards compatibility requirements.	1	1	https://issues.jboss.org/browse/ISPN-4139

#	Requirement	Questions	Answered	Successfully Answered	URL
7906	ISPN-4139	module. Here is an example: {noformat} Tests run: 1, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 0.228 sec <<< FAILURE! - in org.hibernate.search.infinispan.sharedIndex.SharedIndexTest testSingleResultFromDeviceIndex(org.hibernate.search.infinispan.sharedIndex.SharedIndexTest) Time elapsed: 0.228 sec <<< ERROR! java.lang.AbstractMethodError: org.apache.lucene.store.Lock.close()V at org.apache.lucene.index.IndexWriter.closeInternal(IndexWriter.java:1052) at org.apache.lucene.index.IndexWriter.close(IndexWriter.java:927) at org.apache.lucene.index.IndexWriter.close(IndexWriter.java:889) at org.hibernate.search.store.impl.DirectoryProviderHelper.initializeIndexIfNeeded(DirectoryProviderHelper.java:159) at org.hibernate.search.infinispan.impl.InfinispanDirectoryProvider.start(InfinispanDirectoryProvider.java:109) at org.hibernate.search.indexes.impl.DirectoryBasedIndexManager.initialize(DirectoryBasedIndexManager.java:104) at org.hibernate.search.indexes.impl.IndexManagerHolder.createIndexManager(IndexManagerHolder.java:272) at org.hibernate.search.indexes.impl.IndexManagerHolder.createIndexManager(IndexManagerHolder.java:561) at org.hibernate.search.indexes.impl.IndexManagerHolder.createIndexManagers(IndexManagerHolder.java:530) at org.hibernate.search.indexes.impl.IndexManagerHolder.buildEntityIndexBinding(IndexManagerHolder.java:107) at org.hibernate.search.spi.SearchFactoryBuilder.initDocumentBuilders(SearchFactoryBuilder.java:386) at org.hibernate.search.spi.SearchFactoryBuilder.buildNewSearchFactory(SearchFactoryBuilder.java:225) at org.hibernate.search.spi.SearchFactoryBuilder.buildSearchFactory(SearchFactoryBuilder.java:143) at org.hibernate.search.hcore.impl.HibernateSearchSessionFactoryObserver.sessionFactoryCreated(HibernateSearchSessionFactoryObserver.java:79) at org.hibernate.internal.SessionFactoryObserverChain.sessionFactoryCreated(SessionFactoryObserverChain.java:52) at org.hibernate.internal.SessionFactoryImpl.<init>(SessionFactoryImpl.java:588) at org.hibernate.cfg.Configuration.buildSessionFactory(Configuration.java:1857) at org.hibernate.search.test.util.FullTextSessionBuilder.build(FullTextSessionBuilder.java:181) at org.hibernate.search.infinispan.ClusterTestHelper.createClusterNode(ClusterTestHelper.java:97) at org.hibernate.search.infinispan.ClusterTestHelper.createClusterNode(ClusterTestHelper.java:61) at org.hibernate.search.infinispan.sharedIndex.SharedIndexTest.setUp(SharedIndexTest.java:100) {noformat} {quote} I can take care of this, it's quite a minefield of backwards compatibility requirements. {quote} If you want. I finally have Infinispan configured in my IDE though and I kind of see what I need to do. I guess I would need something similar to {{DirectoryBuilderImpl#create}} in {{BaseLockFactory}} (or somewhere in this care). Depending on the Lucene version I need to return an appropriate instance of {{Lock}}. I actually tried to implement {{#release()}} and {{#lock}} at the same time in {{BaseLuceneLock}}, but that did not work :-(The class is still not compatible then when run with Lucene 4.7. Just let me know if you fix this, then I leave it for now. {quote} BTW If you can't upgrade Search to 4.7 without an Infinispan release first, it's getting messy. We would need an Infinispan release, and the Infinispan patch would need to support both Lucene 3, and Lucene 4.6 and	1	1	https://issues.jboss.org/browse/ISPN-4139
7907	ISPN-4139	{quote} Lol, that's the reason for me to work on including Search 5 in Infinispan asap: to get rid of that (remove the Lucene3 support) {quote} How is this even related?	1	1	https://issues.jboss.org/browse/ISPN-4139
7908	ISPN-4132	"So at any time it might make sense to keep in memory some of the attributes of S, but not all of them." Not really. In WF/EAP, the sessions to be passivated are primarily expected to be abandoned sessions, that is, session that I expect to eventually expire. Our primary use case for passivation is *not* as a means to handle more active sessions that a given server can keep in memory - but rather to handle the abandoned, but not yet expired sessions. In this way, I would typically prefer to evict all entries of a presumably abandoned session over the attributes of an active session that have not been accessed recently. The likelihood that the former will be accessed before the latter is typically low. But no, there isn't anything in the JEE standard about how one specifies how sessions are passivated. Your point about this being too complex is certainly valid, so let's try a different tack... What if we had the ability to specify that a given cache entry was never to be passivated automatically (i.e. like an immortal cache entry, but for eviction purposes)? Infinispan would then passivate based on a max-size of "auto-evictable" cache entries (i.e. some subset of the total entries in memory). This way, I can leverage a passivation callback to manually passivate the non-auto-evictable session attributes for a given session and still be able to define my eviction settings based on a maximum number of sessions. WDYT?	0	0	https://issues.jboss.org/browse/ISPN-4132
7909	ISPN-4022	[~dan.berindei] and [~mircea.markus] I thought about this a bit further and I concluded there are additional benefits to this approach. I would call this enhancement "staggered combine". Just as Dan suggests we should invoke combine on certain thresholds (say 1K entries in a combiner) during map phase and move intermediate KOut/VOut values around cluster as these thresholds are reached. The benefit is that not only we will relieve memory pressure, we will also never run out of RAM storing map output in a Collector. In addition staggered migration of KOut/VOut to intermediate cache should alleviate some of the insertion stress we have observed in performance tests. If we are able to combine this feature with ISPN-3999 Sanne suggested this should be awesome! WDYT?	1	1	https://issues.jboss.org/browse/ISPN-4022
7910	ISPN-4021	Dan, we have two options here: a) we can pass intermediate cache directly to MapReduceTask and thus force correct typing b) pass the cache name only Option b) gives more flexibility in the future if things change internally while option a) minimizes potential errors on the behalf of user at the expense of flexibility b) has. What should we do?	1	1	https://issues.jboss.org/browse/ISPN-4021
7911	ISPN-4021	[~vblagojevic], can't we support both?	1	1	https://issues.jboss.org/browse/ISPN-4021
7912	ISPN-4011	[~mgencur] This issue has now become pretty much obsolete since the introduction of a proper mechanism to define protobuf schemas in a much nicer and friendlier way that does not require the use of JMX (see ISPN-4357). JMX support is still there but I think we should deprecate it in 7 and eventually remove it. Wdyt?	1	1	https://issues.jboss.org/browse/ISPN-4011
7913	ISPN-3954	From Hibernate docs: {quote} Collections are ignored by a stateless session {quote} Doesn't this mean that if I have entity with @ElementCollection, the collection won't get loaded/updated?	1	1	https://issues.jboss.org/browse/ISPN-3954
7914	ISPN-3953	I've noticed that when implementing the process operation. In my current implementation (see http://github.com/rvansa/infinispan-cachestore-jpa) I load only all keys and then load each entry in transaction. But this is far from optimal approach. Shouldn't we rename the cachestore to infinispan-persistence-hibernate instead, then, if we start using hibernate API? I have to write to infinispan-dev for the green.	1	1	https://issues.jboss.org/browse/ISPN-3953
7915	ISPN-3937	Setting the result set type and concurrency shouldn't be a problem. Are you worried about the lock on the table being held while we are processing the results? {{process()}} doesn't do that much work on the caller's thread, so our iteration over the result set should be only slightly slower than the driver's fetching of the entire result set.	0	0	https://issues.jboss.org/browse/ISPN-3937

#	Requirement	Questions	Answered	Successfully Answered	URL
7916	ISPN-3924	2,NonTxPutIfAbsentDuringLeaveStressTest-NodeD-3525:) [CommandAwareRpcDispatcher] ISPN000220: Problems un-marshalling remote command from byte buffer java.lang.NullPointerException 09:01:54,026 ERROR (testng-NonTxPutIfAbsentDuringLeaveStressTest:) [UnitTestTestNGListener] Test testNodeLeavingDuringPutIfAbsent(org.infinispan.distribution.rehash.NonTxPutIfAbsentDuringLeaveStressTest) failed. java.util.concurrent.ExecutionException: org.infinispan.remoting.RemoteException: ISPN000217: Received exception from NonTxPutIfAbsentDuringLeaveStressTest-NodeD-3525, see cause for remote stack trace at java.util.concurrent.FutureTask$Sync.innerGet(FutureTask.java:232) at java.util.concurrent.FutureTask.get(FutureTask.java:91) at org.infinispan.distribution.rehash.NonTxPutIfAbsentDuringLeaveStressTest.testNodeLeavingDuringPutIfAbsent(NonTxPutIfAbsentDuringLeaveStressTest.java:114) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.testng.internal.MethodInvocationHelper.invokeMethod(MethodInvocationHelper.java:80) at org.testng.internal.Invoker.invokeMethod(Invoker.java:714) at org.testng.internal.Invoker.invokeTestMethod(Invoker.java:901) at org.testng.internal.Invoker.invokeTestMethods(Invoker.java:1231) at org.testng.internal.TestMethodWorker.invokeTestMethods(TestMethodWorker.java:127) at org.testng.internal.TestMethodWorker.run(TestMethodWorker.java:111) at org.testng.TestRunner.privateRun(TestRunner.java:767) at org.testng.TestRunner.run(TestRunner.java:617) at org.testng.SuiteRunner.runTest(SuiteRunner.java:334) at org.testng.SuiteRunner.access$000(SuiteRunner.java:37) at org.testng.SuiteRunner$SuiteWorker.run(SuiteRunner.java:368) at org.testng.internal.thread.ThreadUtil$2.call(ThreadUtil.java:64) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) at java.util.concurrent.FutureTask.run(FutureTask.java:138) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:895) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:918) at java.lang.Thread.run(Thread.java:662) Caused by: org.infinispan.remoting.RemoteException: ISPN000217: Received exception from NonTxPutIfAbsentDuringLeaveStressTest-NodeD-3525, see cause for remote stack trace at org.infinispan.remoting.transport.AbstractTransport.checkResponse(AbstractTransport.java:41) at org.infinispan.remoting.transport.jgroups.CommandAwareRpcDispatcher.processSingleCall(CommandAwareRpcDispatcher.java:358) at org.infinispan.remoting.transport.jgroups.CommandAwareRpcDispatcher.invokeRemoteCommand(CommandAwareRpcDispatcher.java:167) at org.infinispan.remoting.transport.jgroups.JGroupsTransport.invokeRemotely(JGroupsTransport.java:521) at org.infinispan.remoting.rpc.RpcManagerImpl.invokeRemotely(RpcManagerImpl.java:280) at org.infinispan.interceptors.distribution.BaseDistributionInterceptor.handleNonTxWriteCommand(BaseDistributionInterceptor.java:238) at org.infinispan.interceptors.distribution.NonTxDistributionInterceptor.visitPutKeyValueCommand(NonTxDistributionInterceptor.java:72) at org.infinispan.commands.write.PutKeyValueCommand.acceptVisitor(PutKeyValueCommand.java:71) at org.infinispan.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:98) at org.infinispan.interceptors.EntryWrappingInterceptor.invokeNextAndApplyChanges(EntryWrappingInterceptor.java:326) at	1	1	https://issues.jboss.org/browse/ISPN-3924
7917	ISPN-3919	No need to verify, it won't work. This requires advanced externalizers which were left out of the Hot Rod client. You'd need to bring in the externalizer table and a lot of other JBoss Marshalling related components to make this work. Is it really needed?	1	0	https://issues.jboss.org/browse/ISPN-3919
7918	ISPN-3914	what events would you be interested in when it comes to groups? E.g. - entry added to the group - entry removed from the group Also there's no group-level locking, we can only lock on a per entry basis. Would be interesting to see more about your use case.	1	0	https://issues.jboss.org/browse/ISPN-3914
7919	ISPN-3905	Hi Dan. How is "reuse" of the string related to this issue?	1	1	https://issues.jboss.org/browse/ISPN-3905
7920	ISPN-3901	Deprecate in 9.0?	0	0	https://issues.jboss.org/browse/ISPN-3901
7921	ISPN-3900	Is this something we can do for 7.0?	0	0	https://issues.jboss.org/browse/ISPN-3900
7922	ISPN-3877	Hello Mircea, Thanks for your response. I have attached a small test program that can show the difference in the space usage. On my machine, a. with the original SingleFileStore version, the disk consumption reaches to around 10-11 MB and stays there. b. with the modified (patched version), the disk consumption reaches to around 10-11MB, but after the purge cycles kick-in, the file size remains below 500KB consistently. I am going to attach a pull request for this issue. So, please review the diff from there. After your review, can you also please let me know if this enhancement seems ok from your point of view? The overhead on the purge cycle because of our truncation logic is very minimal (around few tens of milli-seconds). Also, I am going to create another JIRA issue for another enhancement related to coalescing adjacent free entries. Thanks, Rajesh	1	1	https://issues.jboss.org/browse/ISPN-3877
7923	ISPN-3824	Make it "5.0.4.Final" ?	0	0	https://issues.jboss.org/browse/ISPN-3824
7924	ISPN-3808	Could we close this now and consider it resolved?	1	0	https://issues.jboss.org/browse/ISPN-3808
7925	ISPN-3807	I agree, this is a sensible step towards better performance. But there are some things we need to fix before this can work. 1. We need reliable support for parameters in the Hibernate Search HQL parser and not sure of its status now. Last time I tried using it (September) I was advised to avoid it as it was a bit too experimental. 2. Would be great if we had support for translating the 'order by'; right now it is recognized and ignored (there should be a HSEARCH jira but can't find it). Having 1,2 sorted will mean our query is going to be a fixed string + a variable map of parameter values instead of (currently) being a variable string + a list of sort criteria. Another related idea, would it be possible to cache the Lucene query object instead of caching the hql AST? Does lucene have params?	1	1	https://issues.jboss.org/browse/ISPN-3807
7926	ISPN-3792	PFER is an operation targeted for Hibernate's 2LC. Do you need this enhancement in this context or you're using it for something else? Also have you considered using the org.infinispan.Flag API on top of Cache.put?	1	1	https://issues.jboss.org/browse/ISPN-3792
7927	ISPN-3792	[~karlvr] like Mircea said, we only needed {{putForExternalRead}} in the context of Hibernate 2LC, and there it was enough to set the lifespan or maxIdle to the same values for all the entries in the cache. We're a bit reluctant to add more methods to the Cache interface, can you add some more details on why that won't work for you?	1	1	https://issues.jboss.org/browse/ISPN-3792
7928	ISPN-3792	You're right [~karlvr], an invalidation cache without {{putForExternalRead()}} wouldn't be very usable. The question in my mind was whether it makes sense in general to combine invalidation with time-based expiration, e.g. wouldn't a local cache without cluster-wide invalidation work better in your case?	1	1	https://issues.jboss.org/browse/ISPN-3792

#	Requirement	Questions	Answered	Successfully Answered	URL
7929	ISPN-3718	Yes, we need those options to bear more complex values than just a boolean. A little known aspect of protobuf options is they are very rich and can represent tuples and arrays also, even arrays of tuples or tuples of arrays. You get the idea :) So it would be perfectly doable even to completely emulate the whole HS indexing annotations. Here's a valid protobuf example of a field being annotated with an option: {code} message User { //... required string name = 3 [(infinispan.Field) = { indexed : true, analyzed : false, some_array : [111, 222, 333], /* yes! even arrays of values */ }]; //... } {code} The stuff between [ ] that follows a field declaration is the list of options attached to the field - its metadata/annotations. There can be multiple options, separated with comma. Unfortunately the protobuf syntax is not very nice and requires you to surround the option name in (). The value of the option is either a scalar value or a tuple surrounded by { }. The structure of the option has to be defined in advance (not shown here), so its usage can be validated by the protobuf schema compiler. Options values can have defaults, so the verbosity can be drastically reduced. So, would you consider this as a valid approach to define our indexing metadata?	1	1	https://issues.jboss.org/browse/ISPN-3718
7930	ISPN-3718	[~gustavonalle] The protostream library has just received support for parsing Java-like annotations from doc comments. Could you have a look please ? (https://github.com/infinispan/protostream/pull/9)	1	1	https://issues.jboss.org/browse/ISPN-3718
7931	ISPN-3668	One more question: org.infinispan.lucene.profiling.PerformanceCompareStressTest.profileInfinispanLocalDirectory() is present in both lucene-v3 and lucene-v4 Do we want the regression tests for both versions ?	0	0	https://issues.jboss.org/browse/ISPN-3668
7932	ISPN-3550	Are you using CQL or thrift in new implementation?	1	0	https://issues.jboss.org/browse/ISPN-3550
7933	ISPN-3549	What's the status of this cloud cache store implementation? Is there any particular reason why it hasn't been ported to Infinispan 6+? Is it only due to lack of interest/usage? I'd be surprised if that's the case, given today's incline to rely on the cloud for pretty much everything. Or is it related to some kind of performance issue where cloud storage in general doesn't lend itself to being used with Infinispan?? Anyway, there's [an Infinispan Github sub-repository|https://github.com/infinispan/infinispan-cachestore-cloud] that contains a 6.0.0-SNAPSHOT version (last update in September 2013). And there's also a 6.0.0.Alpha1 "infinispan-cachestore-cloud" Maven artifact available (uploaded in July 2013). So how much work would be necessary to update this implementation to at least Infinispan 6 (in case Infinispan 7 introduced more radical changes)? One thing to note is that Apache Jclouds, on which this cache store implementation is based, has also seen a lot of changes lately. From version 1.4.1 as referenced in the current WIP implementation, Jclouds is now at version 1.8.1. More important that the version number itself is the fact the Jclouds 1.8.1 has removed support for async operations (which are heavily used in the current cache store code). Jclouds 1.7.x is the last branch that contains support for async operations -- but then we loose new features like OpenStack Swift support in 1.8.x).	1	1	https://issues.jboss.org/browse/ISPN-3549
7934	ISPN-3549	[~dalbani]: great, your changes looks good to me. Would you mind me to cherry-pick your changes and continue with update to ISPN 7 and adding tests? Or you would like to do it yourself? Thanks!	1	1	https://issues.jboss.org/browse/ISPN-3549
7935	ISPN-3534	[~mlinhard], the old bundler was only a part of the issue. We also noticed that setting {{UFC.max_credits=200k}} slows things down (maybe more so with Infinispan 6.0 than 5.2). I see in the comparison you posted that the JDG 6.2 snapshots both ran with 200k, could you run the test again with 2m? And can you find the max_credits value in the JDG 6.1 run?	0	0	https://issues.jboss.org/browse/ISPN-3534
7936	ISPN-3533	Currently, HotRod client sets the protocol version and expects the server to understand it. Wouldn't it be better to include some negotiation: client connect with "I support 2.0, 2.1 and 2.2", server responds "I support 2.0 and 2.1", and then client knows he should use 2.1 instead of 2.2? Also, have you considered rather versioned "extensions" (such as querying, transactions, streaming, listeners...) to basic protocol than simple versions extending the previous one? Or is this overengineering?	1	1	https://issues.jboss.org/browse/ISPN-3533
7937	ISPN-3533	OK, I thought that adding a list of supported protocol versions/features into the ping response and client version auto-selection would be nice, rather than requiring to set the version in each client (the client does not need to send the list in ping request, in fact). Another topic: have you considered streaming data to and from server? It seems that it's a common request to be able to load the server via hotrod, and currently the user needs to spawn many threads in order to circumvent the put latency.	0	0	https://issues.jboss.org/browse/ISPN-3533
7938	ISPN-3480	Can protostream be released with the latest changes?	0	0	https://issues.jboss.org/browse/ISPN-3480
7939	ISPN-3455	Lukasz, i have just spent a week debugging the same issue. The issue is not due to Infinispan but to the usage of the Enum as keys. Infinispan to work need to be able for the key to calculate an hashcode that is consistent in the cluster, meaning if the key is Key1 then on node1,node2,node3 the hashcode produced need to be the same. When the key is an enum, that is not possible because the Java Enum generates an hashcode that is consistent only within one node, but when transported across JVM that is not guaranteed to be generated the same. You could think then to override the hashcode for the enum to generate a consistent one, sorry not possible, look at this SUN bug report: http://bugs.sun.com/view_bug.do?bug_id=6373406 In my case i end up creating a wrapper around the enum key and forcing for that wrapper a generation of a consistent hashkey. For code reference where the problem start in infinispan look at: https://github.com/infinispan/infinispan/blob/5.3.x/core/src/main/java/org/infinispan/interceptors/distribution/BaseDistributionInterceptor.java lines 127-130, there is were the troubles start. Bottom line as i said this is an issue caused by the Java ENUM design. For reference to a bugfix i have done to fix this issue in my code look at: https://git.opendaylight.org/gerrit/#/c/1292/ Hope it helps. Giovanni	1	1	https://issues.jboss.org/browse/ISPN-3455
7940	ISPN-3355	Few comments on the ClusterListeners: - Can the listener be made durable i.e. if the primary dies then the listener will be registered on the backup - Can the listener use a query api e.g. notify me of any Books having "title" like "%engine%" - Can the listener add the option to remove or modify the entry/entires its query matches.	1	1	https://issues.jboss.org/browse/ISPN-3355
7941	ISPN-3318	Interesting idea from Randall: {quote}Let's look at this from a user's perspective. Is it possible that the new file store **replaces** the old one, so that the user doesn't (really) need to know the difference? How similar are the configuration properties of each? Might it be possible to: rename and deprecate the existing one (e.g., "org.infinispan.loaders.file.LegacyFileCacheStore"), changing it to 'package' or 'protected' (iff the new one is stable enough that the old one should definitely not be used anymore), rename the new one to "org.infinispan.loaders.file.FileCacheStore" (so that existing configurations work without change), have some code in the new one detect the old file structure and automatically run a migration to convert from the old file structure to the new format (and possibly move the old files into a subdirectory) If this is possible, this is by far the best option for users since it provides a simple (perhaps even transparent) migration path with little if any changes required to configurations, but at the same time it allows the developers to get what they want: a new FCS that replaces the old one (which goes away).{quote}	1	1	https://issues.jboss.org/browse/ISPN-3318

#	Requirement	Questions	Answered	Successfully Answered	URL
7942	ISPN-3318	Hi Galder, your commit dbf63e2d4b21aaf8a7a5eda1dca1f88c72fc9070 concretely: git diff 43f82b7 dbf63e2 core/src/main/java/org/infinispan/util/logging/Log.java (on current master) made the infinispan-cachestore-rest build unfunctional http://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/jdg-qe-build-infinispan60-snapshot-auto/149/console {code} [ERROR] Failed to execute goal org.apache.maven.plugins:maven-compiler-plugin:3.0:compile (default-compile) on project infinispan-cachestore-rest: Compilation failure: Compilation failure: [ERROR] Only one message with the same format parameters is allowed. [ERROR] /home_local/workspace/jdg-qe-build-infinispan60-snapshot-auto/infinispan-cachestore-rest/src/main/java/org/infinispan/persistence/rest/logging/Log.java:[60,19] Only one message with the same format parameters is allowed. [ERROR] /home_local/workspace/jdg-qe-build-infinispan60-snapshot-auto/infinispan-cachestore-rest/src/main/java/org/infinispan/persistence/rest/logging/Log.java:[64,9] Only one message with the same format parameters is allowed. [ERROR] /home_local/workspace/jdg-qe-build-infinispan60-snapshot-auto/infinispan-cachestore-rest/src/main/java/org/infinispan/persistence/rest/logging/Log.java: Only one message with the same format parameters is allowed. {code} could you please fix this in the rest cachestore as well ? (it blocks building of infinispan-server for perf regression tests)	0	0	https://issues.jboss.org/browse/ISPN-3318
7943	ISPN-3285	[~sannegrinovero] These were already made optional in ISPN-3437. Is it worth removing them altogether?	0	0	https://issues.jboss.org/browse/ISPN-3285
7944	ISPN-3242	what's the reason for enforcing this?	1	1	https://issues.jboss.org/browse/ISPN-3242
7945	ISPN-3227	Shay Matasaro <smatasar@redhat.com> made a comment on [bug 973763|https://bugzilla.redhat.com/show_bug.cgi?id=973763] the complexity for this patch is really increased due to schema changes, i dont think its worth it how about using a store with pre-load configured but no persistence will that work as a workaround?	0	0	https://issues.jboss.org/browse/ISPN-3227
7946	ISPN-3227	Shay Matasaro <smatasar@redhat.com> made a comment on [bug 973763|https://bugzilla.redhat.com/show_bug.cgi?id=973763] i have not been able to check-in my code, here is a source patch file	0	0	https://issues.jboss.org/browse/ISPN-3227
7947	ISPN-3199	{code}put(K key, EntryVersion version);{code} <- that's already in 5.3, via put(K, V, Metadata) Shouldn't we take advantage of all the work around Metadata to include the rest of methods?	0	0	https://issues.jboss.org/browse/ISPN-3199
7948	ISPN-3194	[~NadirX] - there is a bit here (https://access.redhat.com/site/solutions/221103) about JDR but it doesn't really tell you what is being captured from the runtime. It seems to be more of a tool to zip up configs and logs. Do you have any more info on JDR?	0	0	https://issues.jboss.org/browse/ISPN-3194
7949	ISPN-3157	Shay Matasaro <smatasar@redhat.com> made a comment on [bug 973763|https://bugzilla.redhat.com/show_bug.cgi?id=973763] the complexity for this patch is really increased due to schema changes, i dont think its worth it how about using a store with pre-load configured but no persistence will that work as a workaround?	0	0	https://issues.jboss.org/browse/ISPN-3157
7950	ISPN-3157	Shay Matasaro <smatasar@redhat.com> made a comment on [bug 973763|https://bugzilla.redhat.com/show_bug.cgi?id=973763] i have not been able to check-in my code, here is a source patch file	0	0	https://issues.jboss.org/browse/ISPN-3157
7951	ISPN-3140	where would the data in this cluster be persisted during the shutdown? simpler with a shared cache store, each cache persisting locally would complicate things a bit.	0	0	https://issues.jboss.org/browse/ISPN-3140
7952	ISPN-3140	Pasted from Adrian's message ([http://markmail.org/message/ns7aojy7v7su2t7p]): {quote} 1. Add a JMX writable attribute (or operation?) to ClusterTopologyManager (name it suppressRehashing?) that is false by default but should also be configurable via API or xml. While this attribute is true the ClusterTopologyManager queues all join/leave/exclude(see below) requests and does not execute them on the spot as it would normally happen. [...] When it is set back to false all queued operations (except the ones that cancel eachother out) are executed. The setter should be synchronous so when setting is back to false it does not return until the queue is empty and all rehashing was processed. 2. We add a JMX operation excludeNodes(list of addresses) to ClusterTopologyManager. [...] This operation removes the node from the topology (almost as if it left) and forces a rebalance. The node is still present in the current CH but not in the pending CH. It's basically disowned by all its data which is now being transferred to other (not excluded) nodes. At the end of the rebalance the node is removed from topology for good and can be shut down without loosing data. Note that if suppressRehashing==true operation excludeNodes(..) just queues them for later removal. We can batch multiple such exclusions and then re-activate the rehashing. The parts that need to be implemented are written in italic above. Everything else is already there. excludeNodes is a way of achieving a soft shutdown and should be used only if we care about preserving data int the extreme case where the nodes are the last/single owners. We can just kill the node directly if we do not care about its data. suppressRehashing is a way of achieving some kind of batching of topology changes. This should speed up state transfer a lot because it avoids a lot of pointless reshuffling of data segments when we have many successive joiners/leavers. So what happens if the current coordinator dies for whatever reason? The new one will take control and will not have knowledge of the existing rehash queue or the previous status of suppressRehashing attribute so it will just get the current cache membership status from all members of current view and proceed with the rehashing as usual. If the user does not want this he can set a default value of true for suppressRehashing. The admin has to interact now via JMX with the new coordinator. But that's not as bad as the alternative where all the nodes are involved in this jmx scheme :) I think having only the coordinator involved in this is a plus. {quote} We're actually going to implement only point 1 now, and point 2 will be a separate issue (or perhaps as a part of ISPN-1394).	0	0	https://issues.jboss.org/browse/ISPN-3140
7953	ISPN-3128	How can we change in spring cache to use putforexternalread instead of put. I simply want to use putforexternalread instead of put. I am not using CachePut in any other places. So I don't have issues.	0	0	https://issues.jboss.org/browse/ISPN-3128
7954	ISPN-3126	Can you please close it once you confirmed that this is not a problem?	0	0	https://issues.jboss.org/browse/ISPN-3126
7955	ISPN-3102	I've tried running this test (with all logging set to WARN) on 5.3.x, 6.0 and on the exact revision when the test was added. The results were nearly identical for all runs: 1.5 mins. So this is not a regression. Still, indexing 2000 entries in 1.5 mins is very slow...	1	1	https://issues.jboss.org/browse/ISPN-3102
7956	ISPN-3102	[~sannegrinovero] How did you determine StateRequestCommand serialization (ISPN-3103) is a CPU usage hotspot here? This test uses waitForClusterToForm before anything else happens, so state transfer simply cannot have any influence on performance here. Anyway, I did try to profile and could not find any hotspot.	1	1	https://issues.jboss.org/browse/ISPN-3102
7957	ISPN-3076	Does Infinispan 5.2.5.Final use JGroups 3.3.x?	0	0	https://issues.jboss.org/browse/ISPN-3076

#	Requirement	Questions	Answered	Successfully Answered	URL
7958	ISPN-3076	Sorry my statement had a critical typo: I meant to say {quote}But I still don't get why Log4j*2* is of any interest to us.{quote} In other words, I think JBoss Logging should be our default logger implementation, and we also support log4j (v1) because it's a popular choice in other containers. Log4j V2 is not that popular yet to make this a wise move, especially until it has proven itself more efficient: I don't think we have data to back this up, for sure it defeats the purpose of even considering Log4J which should be "going with the popular version". Very interesting of course to explore alternatives, for debugging purposes for example we could have an in memory circular buffer which skips the String formatting step: just store the parameters as they are together with the log id, the buffer content could then be "rendered" in its string form at a second time when and only if needed.	1	1	https://issues.jboss.org/browse/ISPN-3076
7959	ISPN-3076	Any news about this subject ? I'm too struggling with jboss logging and Log4J2. My classpath is supposed to be correct (log4j2 api + core + bridges + jboss logging 3.2.0.Final - and no log4j 1 jar) and the only last logs not working are those ones with the exact same message : {quote} 2015-01-08 08:56:28,421 WARN The Logger org.infinispan.remoting.transport.jgroups.CommandAwareRpcDispatcher was created with the message factory org.apache.logging.log4j.message.ParameterizedMessageFactory@59aea098 and is now requested with the message factory org.apache.logging.log4j.message.StringFormatterMessageFactory@57b18927, which may create log events with unexpected formatting. {quote} Or is there a working workaround ?	0	0	https://issues.jboss.org/browse/ISPN-3076
7960	ISPN-3051	Exactly, like _localstorage=false_ in Oracle Coherence or event better like the Coherence *Extend feature (that also exclude the node from the distributed execution env). Check this [link|http://blog.tfanshteyn.com/2009/11/extending-data-grid-localstoragefalse.html] for more info. *USE CASE* For example, if you want to use infinispan as integration channel some embedded caches will be part of the channel keeping the data, but others will interact with the channel putting and getting data only (like a client). *This is exactly what you get using hotrod + _RemoteCache_ but it will be cool to be able to do this also with embedded caches.* Thanks in advance,	0	0	https://issues.jboss.org/browse/ISPN-3051
7961	ISPN-3045	Does this JIRA make any sense since the compatibility mode has been implemented?	0	0	https://issues.jboss.org/browse/ISPN-3045
7962	ISPN-2953	FWIW we're using the canonical ObjectName. As it happens though I've just worked around it by making a custom tree UI in hawtio for infinispan. So from my perspective I'm happy for you to close this issue now. Though I do think it'd be nice for the MBeans to appear a little nicer in all other JMX tools; but personally am less bothered now we've a nice hawtio plugin :)	1	1	https://issues.jboss.org/browse/ISPN-2953
7963	ISPN-2932	> what if a key that is not present in the cache is looked up? If preload=true they are all loaded in the cache, why should it not be there?	0	0	https://issues.jboss.org/browse/ISPN-2932
7964	ISPN-2926	I don't think it's a duplicate. The ability to access all the content of a cache is not what CLI provide at the moment (ISPN-1453) as far as I understand Can you reopen this?	1	0	https://issues.jboss.org/browse/ISPN-2926
7965	ISPN-2923	Infinispan doesn't allow inserting keys with null values, how did you end up with null values in the cache? I think this makes sense as a general requirement, e.g. stop a map/reduce task if enough values have been found.	0	0	https://issues.jboss.org/browse/ISPN-2923
7966	ISPN-2916	Is this something we can do for 7.0?	1	1	https://issues.jboss.org/browse/ISPN-2916
7967	ISPN-2822	FileStore.zip - contains the jprofiler snapshots when FileCacheStore was used; JDBC.zip - is for JDBC CacheStore use; No Store - is when the eviction is enabled, but no loaders are configured.	0	0	https://issues.jboss.org/browse/ISPN-2822
7968	ISPN-2808	[~pruivo]: will do, thanks	1	1	https://issues.jboss.org/browse/ISPN-2808
7969	ISPN-2792	{quote}There are some threads that don't follow this rule, e.g. the HotRod worker threads, but it should be useful in most cases.{quote} Would be nice to have these logging the test name as well. What advantage would this have over the grepping the logs for a single test?	1	1	https://issues.jboss.org/browse/ISPN-2792
7970	ISPN-2710	What kind of information do you need?	1	1	https://issues.jboss.org/browse/ISPN-2710
7971	ISPN-2689	I've attached a very hacky proof of concept. I basically just forked the module and patched the code to build with Lucene 4, then everything magically worked! So it seems some very trivial changes can in fact make what we have now at least work with Lucene 4.0 - though not in a backwards compatible way. Question now is, what's the best way to make both a Lucene 3 and 4 compatible version (or versions) available to users?	1	1	https://issues.jboss.org/browse/ISPN-2689
7972	ISPN-2689	Hi James, I'll be back from travels tomorrow. Had a quick preview already, but the patch doesn't apply cleanly; I'll have a better look on the planes. It would be great if you could use "unified format" for patches. Lucene 4.1 will be released very soon, possibly today. Shall we target that instead?	1	1	https://issues.jboss.org/browse/ISPN-2689
7973	ISPN-2675	"So I need for the 5.2 a way to provide custom create table syntax :-|" you are allowed to create the table by hand by setting "createOnStartup" to false. Isn't that good enough?	1	0	https://issues.jboss.org/browse/ISPN-2675
7974	ISPN-2669	[~saturnism], is this JIRA still relevant given what [~gabriel_mircea] said?	1	0	https://issues.jboss.org/browse/ISPN-2669

#	Requirement	Questions	Answered	Successfully Answered	URL
7975	ISPN-2665	{code}testBasicTargetRemoteDistributedCallableWithHighFutureAndLowTaskTimeout(org.infinispan.distexec.ReplSyncDistributedExecutorTest) Time elapsed: 0.416 sec <<< FAILURE! org.testng.TestException: Expected exception java.util.concurrent.TimeoutException but got java.util.concurrent.ExecutionException: org.infinispan.CacheException: org.jgroups.TimeoutException: timeout sending message to ReplSyncDistributedExecutorTest-NodeV-22475 at org.testng.internal.Invoker.handleInvocationResults(Invoker.java:1503) at org.testng.internal.Invoker.invokeMethod(Invoker.java:764) at org.testng.internal.Invoker.invokeTestMethod(Invoker.java:907) at org.testng.internal.Invoker.invokeTestMethods(Invoker.java:1237) at org.testng.internal.TestMethodWorker.invokeTestMethods(TestMethodWorker.java:127) at org.testng.internal.TestMethodWorker.run(TestMethodWorker.java:111) at org.testng.TestRunner.privateRun(TestRunner.java:767) at org.testng.TestRunner.run(TestRunner.java:617) at org.testng.SuiteRunner.runTest(SuiteRunner.java:334) at org.testng.SuiteRunner.access$000(SuiteRunner.java:37) at org.testng.SuiteRunner$SuiteWorker.run(SuiteRunner.java:368) at org.testng.internal.thread.ThreadUtil$2.call(ThreadUtil.java:64) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) at java.util.concurrent.FutureTask.run(FutureTask.java:138) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) at java.lang.Thread.run(Thread.java:680) Caused by: java.util.concurrent.ExecutionException: org.infinispan.CacheException: org.jgroups.TimeoutException: timeout sending message to ReplSyncDistributedExecutorTest-NodeV-22475 at java.util.concurrent.FutureTask$Sync.innerGet(FutureTask.java:232) at java.util.concurrent.FutureTask.get(FutureTask.java:91) at org.infinispan.distexec.DefaultExecutorService$DistributedTaskPart.innerGet(DefaultExecutorService.java:948) at org.infinispan.distexec.DefaultExecutorService$DistributedTaskPart.get(DefaultExecutorService.java:927) at org.infinispan.distexec.DistributedExecutorTest.testBasicTargetRemoteDistributedCallableWithHighFutureAndLowTaskTimeout(DistributedExecutorTest.java:159) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.testng.internal.MethodInvocationHelper.invokeMethod(MethodInvocationHelper.java:80) at org.testng.internal.Invoker.invokeMethod(Invoker.java:715) ... 15 more Caused by: org.infinispan.CacheException: org.jgroups.TimeoutException: timeout sending message to ReplSyncDistributedExecutorTest-NodeV-22475 at org.infinispan.util.Util.rewrapAsCacheException(Util.java:539) at org.infinispan.remoting.transport.jgroups.CommandAwareRpcDispatcher.invokeRemoteCommand(CommandAwareRpcDispatcher.java:184) at org.infinispan.remoting.transport.jgroups.JGroupsTransport.invokeRemotely(JGroupsTransport.java:515) at org.infinispan.remoting.rpc.RpcManagerImpl.invokeRemotely(RpcManagerImpl.java:179) at org.infinispan.remoting.rpc.RpcManagerImpl.invokeRemotely(RpcManagerImpl.java:201) at org.infinispan.remoting.rpc.RpcManagerImpl.invokeRemotely(RpcManagerImpl.java:258) at org.infinispan.remoting.rpc.RpcManagerImpl.access$000(RpcManagerImpl.java:81) at org.infinispan.remoting.rpc.RpcManagerImpl$1.call(RpcManagerImpl.java:292)	0	0	https://issues.jboss.org/browse/ISPN-2665
7976	ISPN-2657	[~saturnism] Could you go ahead and issue a pull request containing your patch, so I can review? Perhaps clean up the config defaults, etc., and Javadoc what can be configured?	1	1	https://issues.jboss.org/browse/ISPN-2657
7977	ISPN-2634	Hi Erik, has this ever seen the light of day ? :-)	0	0	https://issues.jboss.org/browse/ISPN-2634
7978	ISPN-2601	Is this about exposing the xsite configuration within the AS?	1	1	https://issues.jboss.org/browse/ISPN-2601
7979	ISPN-2596	Adrian, how far into this task are you ? Cause I want to verify ISPN-2574, and that basically means I'm gonna write an unit test, cause we have no stable scenario available where ISPN-2574 manifested.	1	1	https://issues.jboss.org/browse/ISPN-2596
7980	ISPN-2580	How many nodes make a cluster "large"? does this make the state transfer not to finish? (Bug 881070)	0	0	https://issues.jboss.org/browse/ISPN-2580
7981	ISPN-2573	Can you please elaborate a bit on the expiration flags bit? +1 re: DecoratedRemoteCache	0	0	https://issues.jboss.org/browse/ISPN-2573
7982	ISPN-2493	Are you documenting this on the wiki ?	0	0	https://issues.jboss.org/browse/ISPN-2493
7983	ISPN-2402	Sanne, what is the exception you get when you do a cache operation? I don't think the exception you posted can ever be thrown from a Cache method.	0	0	https://issues.jboss.org/browse/ISPN-2402
7984	ISPN-2402	The IllegalStateException that is thrown when the remote cache is in the process of shutting down should be handled similarly: {noformat} 15:01:27,140 ERROR (testng-NonTxPutIfAbsentDuringJoinStressTest:) [UnitTestTestNGListener] Test testNodeLeavingDuringPutIfAbsent(org.infinispan.distribution.rehash.NonTxPutIfAbsentDuringLeaveStressTest) failed. java.util.concurrent.ExecutionException: org.infinispan.remoting.RemoteException: ISPN000217: Received exception from NonTxPutIfAbsentDuringJoinStressTest-NodeH-31913, see cause for remote stack trace ... Caused by: org.infinispan.remoting.RemoteException: ISPN000217: Received exception from NonTxPutIfAbsentDuringJoinStressTest-NodeH-31913, see cause for remote stack trace ... Caused by: java.lang.IllegalStateException: Default cache is in 'STOPPING' state and this is an invocation not belonging to an on-going transaction, so it does not accept new invocations. Either restart it or recreate the cache container. at org.infinispan.interceptors.InvocationContextInterceptor.handleAll(InvocationContextInterceptor.java:93) at org.infinispan.interceptors.InvocationContextInterceptor.handleDefault(InvocationContextInterceptor.java:70) at org.infinispan.commands.AbstractVisitor.visitPutKeyValueCommand(AbstractVisitor.java:32) at org.infinispan.commands.write.PutKeyValueCommand.acceptVisitor(PutKeyValueCommand.java:62) at org.infinispan.interceptors.InterceptorChain.invoke(InterceptorChain.java:321) at org.infinispan.commands.remote.BaseRpcInvokingCommand.processVisitableCommand(BaseRpcInvokingCommand.java:39) at org.infinispan.commands.remote.SingleRpcCommand.perform(SingleRpcCommand.java:48) at org.infinispan.remoting.InboundInvocationHandlerImpl.handleInternal(InboundInvocationHandlerImpl.java:97) at org.infinispan.remoting.InboundInvocationHandlerImpl.access$000(InboundInvocationHandlerImpl.java:46) at org.infinispan.remoting.InboundInvocationHandlerImpl$2.run(InboundInvocationHandlerImpl.java:169) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) ... 1 more {noformat}	1	0	https://issues.jboss.org/browse/ISPN-2402
7985	ISPN-2397	If the c3p0.properties exists in the class path it is used for configuring the connection pool. You can add testConnectionOnCheckout/testConnectionOnChecking into that file, so I don't think this is an issue. Or am I not understanding your problem?	1	0	https://issues.jboss.org/browse/ISPN-2397
7986	ISPN-2379	Could you please explain in a few words what is this all about?	1	1	https://issues.jboss.org/browse/ISPN-2379

#	Requirement	Questions	Answered	Successfully Answered	URL
7987	ISPN-2366	I have a case in a functional test for ISPN-2362 that is very hard to create unless CacheManager.getCache() returns ASAP instead of waiting for NBST to perform the transfer. I'm not saying we should make getCache() always return immediately but we could make this configurable, some users might need it, most will probably prefer the current behaviour. So this whole issue can be solved by introducing a boolean config 'waitForInitialStateTransferToComplete' in StateTransferConfiguration (and friends) that is 'true' by default and updating StateTransferManagerImpl to skip the waiting if configuration says so. I already implemented this in my t_2362_m branch if you are interested. Would you agree to include this in 5.2.0?	1	1	https://issues.jboss.org/browse/ISPN-2366
7988	ISPN-2359	Is the status still "Coding in Progress" ?	0	0	https://issues.jboss.org/browse/ISPN-2359
7989	ISPN-2359	I was not able to reproduce the issue (linked PR with a test). Is there any particular situation it happens or any mailing list thread describing the issue?	1	1	https://issues.jboss.org/browse/ISPN-2359
7990	ISPN-2355	Can't you just use L1 for making sure data is available after writing? I'm not sure how such a flag can be implemented, as it would require (unless I'm missing something) the routing table to be aware of a particular key.	1	0	https://issues.jboss.org/browse/ISPN-2355
7991	ISPN-2355	I thought the L1 is only filled at reading data, not at writing. If the putted value is always available on the local node when L1 enabled, then this would be a proper solution. Exists there a way to avoid rehashing of this entry? (Since the entry is mostly not needed on other nodes. :-)	0	0	https://issues.jboss.org/browse/ISPN-2355
7992	ISPN-2336	Hmmm, read-only would need to be defined too. I.e. do you allow removals? Which put ops do you allow (you still need to be able to populate the cache, i.e. with putIfAbsent or PFER)? If an entry is evicted, how do you control no one is going to put a newer version of the entry?	0	0	https://issues.jboss.org/browse/ISPN-2336
7993	ISPN-2332	Why not both ? Allow backup strategies to be defined globally, and override them on a per-cache basis. An empty <backups/> or <sites/> means we don't back up a named cache... This would allow for the following: - If we want almost all caches to be backed up: define the strategies in the global section and override the named caches that should not get backed up with an empty <backups/> - If we have predominantly local, non-backed up caches, don't define <backups> in the global section, but instead on a per cache basis.	1	1	https://issues.jboss.org/browse/ISPN-2332
7994	ISPN-2328	Hmm, is this cos we're missing some configuration that was present in the old configuration? Or is it to provide support for ISPN-701 or similar functionality?	1	1	https://issues.jboss.org/browse/ISPN-2328
7995	ISPN-2327	Thomas, can you check again with 5.2.0.Beta1? If you do manage to reproduce it, please attach a log here (preferably with TRACE enabled for org.infinispan).	0	0	https://issues.jboss.org/browse/ISPN-2327
7996	ISPN-2319	I'm thinking for the configuration to look like this: {code:xml} <backup> <takeOffline afterFailures="10" minTimeToWait="30000"/> </backup> {code} Wdyt?	1	1	https://issues.jboss.org/browse/ISPN-2319
7997	ISPN-2319	So the time before a site-unreachable would be "failuresUntilOffline" * "timeout" ms? By consequence, would this also handle the case where an async remote transaction fails (say, due to a node failing on the remote site)? I had implemented a brute-force method similar to: <backup retryNum="5" retryBackoff="1000"/>	1	1	https://issues.jboss.org/browse/ISPN-2319
7998	ISPN-2318	Tomas Sykora <tsykora@redhat.com> made a comment on [bug 868832|https://bugzilla.redhat.com/show_bug.cgi?id=868832] Still no luck. Setting back to ON_DEV. Attaching new TRACE log from rack test - ER5. I don't know how can I help more now. Just let me know and I will do maximum.	0	0	https://issues.jboss.org/browse/ISPN-2318
7999	ISPN-2285	Hmmm, why is this really needed?	1	1	https://issues.jboss.org/browse/ISPN-2285
8000	ISPN-2284	[~vblagojevic] - do we have the new workings of map/reduce noted down somewhere for us to compare for docs?	1	0	https://issues.jboss.org/browse/ISPN-2284
8001	ISPN-2259	We've made some improvements recently to the async store, can you test with 5.2.0.Beta2 and see if you still such slow behaviour?	0	0	https://issues.jboss.org/browse/ISPN-2259
8002	ISPN-2204	Can't remember. I think one reason was, that the local executed callable got no keys (is filed as another ticket).	1	1	https://issues.jboss.org/browse/ISPN-2204
8003	ISPN-2198	Hmmm, odd. The key should be '8Az4Ia2V5NzYzNDI=', but the exception says that the key is 'PRIMARY'? @Radim, can you replicate this in a smaller scale and generate some logs with TRACE on for org.infinispan package? Tristan/Mircea, can either of you have a look to this?	0	0	https://issues.jboss.org/browse/ISPN-2198
8004	ISPN-2198	@Radim - the attached logs files were produced with DEBUG level enabled. This is not good for me, as it doesn't highlight individual key added to the cache. Can you please reproduce with TRACE level?	1	0	https://issues.jboss.org/browse/ISPN-2198
8005	ISPN-2198	Looking at the attached logs I can see that a put takes place on the node 4 *at the same* time when the other server(number 2) is shutdown (time 12:51:17,827). My understanding of the problem is that there's no (put) activity *during* and after the shutdown - otherwise the increasing number of entries might simply be explained by the addition of more entries to the system. Can you please confirm this? I also didn't see any size() being invoked on all the caches in the cluster (e.g on node 3) - how was the size of each individual cache obtained?	1	0	https://issues.jboss.org/browse/ISPN-2198
8006	ISPN-2192	Ok I see what you mean Randall! You basically want to execute certain code on all cache entries regardless of cache mode. We could make another class that implements DistributedExecutorService, say DefaultLocalExecutorService?	0	0	https://issues.jboss.org/browse/ISPN-2192
8007	ISPN-2181	Pete, a bit more guidance please. So for step 1) I have to do something similar to [how we implement a Bean for RemoteCache|https://github.com/infinispan/infinispan/blob/master/cdi/extension/src/main/java/org/infinispan/cdi/InfinispanExtension.java#L191]? I don't think I need to capture these injection points like we do for RemoteBean? When you mention extraction from cc you mean extraction from CreationalContext? Do I need to save Producer for Cache and use it to produce Cache similar to [RemoteCache|https://github.com/infinispan/infinispan/blob/master/cdi/extension/src/main/java/org/infinispan/cdi/InfinispanExtension.java#L198]? Step 3) So I put the cache into CreationContext in onPreExecute callback? I simply use the same [BeanManager|https://github.com/infinispan/infinispan/blob/master/cdi/extension/src/main/java/org/infinispan/cdi/CDIMapReduceTaskLifecycle.java#L22], get the CreationContext and call push with the specific Cache? Anything else? What about onPostExecute, do I have to do some cleanup? For step 0, lets call this qualifier @ImpliedInputTaskCache. Do I need to make a producer for it?	1	1	https://issues.jboss.org/browse/ISPN-2181

#	Requirement	Questions	Answered	Successfully Answered	URL
8008	ISPN-2181	{quote} So for step 1) I have to do something similar to how we implement a Bean for RemoteCache? I don't think I need to capture these injection points like we do for RemoteBean? {quote} You just need one Bean, you don't need to call it in a loop or anything. {quote} When you mention extraction from cc you mean extraction from CreationalContext? {quote} Yes. {quote} Do I need to save Producer for Cache and use it to produce Cache similar to RemoteCache? {quote} I would just implement the necessary custom logic in the ContextualLifecycle. I think you don't need any producer. There, you extract the right cache from the creational context. {quote} Step 3) So I put the cache into CreationContext in onPreExecute callback? {quote} Yes. {quote} I simply use the same BeanManager, get the CreationContext and call push with the specific Cache? {quote} Hmm, I guess you can't use the default creational context. You can probably just wrap CreationalContext and add some additional field to it, and then cast it. If this breaks Weld, it's a bug in Weld, and we can look at other options {quote} Anything else? What about onPostExecute, do I have to do some cleanup? {quote} Don't think so. {quote} For step 0, lets call this qualifier @ImpliedInputTaskCache. Do I need to make a producer for it? {quote} No. Just give the bean in (1) that qualifier.	1	1	https://issues.jboss.org/browse/ISPN-2181
8009	ISPN-2181	{quote}Hmm, I guess you can't use the default creational context. You can probably just wrap CreationalContext and add some additional field to it, and then cast it. If this breaks Weld, it's a bug in Weld, and we can look at other options{quote} What do you mean exactly? How can I do this after I create CreationalContext instance using BeanManager's createCreationalContext method? Do you mean to wrap that instance with a new instance of a class I write (which implements CreationalContext and has extra cache field) and that instance I cast then in create callback of ContextualLifecycle of the bean?	1	1	https://issues.jboss.org/browse/ISPN-2181
8010	ISPN-2177	Mircea, how about removing {{backupKeyLocks}} completely and using {{affectedKeys}}/{{clusterLockedKeys}} instead? Is there any place where having a false positive is really a problem?	0	0	https://issues.jboss.org/browse/ISPN-2177
8011	ISPN-2158	@Galder - are you reviewing this pull req?	1	1	https://issues.jboss.org/browse/ISPN-2158
8012	ISPN-2152	Michal Linhard <mlinhard@redhat.com> made a comment on [bug 841894|https://bugzilla.redhat.com/show_bug.cgi?id=841894] Tristan, I'm not sure what's to verify here, I've checked both distros: jboss-datagrid-library-6.0.1.ER1.zip jboss-datagrid-server-6.0.1.ER1.zip and yes, they contain JGroups 3.0.11.Final We've also ran hotrod client stress tests and everything seems to work fine. Do you propose to check anything else ?	0	0	https://issues.jboss.org/browse/ISPN-2152
8013	ISPN-2134	Fix version?	1	1	https://issues.jboss.org/browse/ISPN-2134
8014	ISPN-2133	Another issue that is making OGM a bit less efficient than it should. Should we create a tag SLOWING_OGM?	1	1	https://issues.jboss.org/browse/ISPN-2133
8015	ISPN-2129	Out of curiosity, how do you expect to call C APIs from Java (in a plattform independent manner)? :)	0	0	https://issues.jboss.org/browse/ISPN-2129
8016	ISPN-2119	Hmmm, why does CLI need sessions? :)	0	0	https://issues.jboss.org/browse/ISPN-2119
8017	ISPN-2104	We have these methods that receive a timeout value in millis in our config builder API. Do we want to enhance all of them? I suppose we keep the existing version taking a long value. org.infinispan.configuration.cache.AsyncConfigurationBuilder.replQueueInterval( ) org.infinispan.configuration.cache.AsyncLoaderConfigurationBuilder.flushLockTimeout( ) org.infinispan.configuration.cache.AsyncLoaderConfigurationBuilder.shutdownTimeout( ) org.infinispan.configuration.cache.DeadlockDetectionConfigurationBuilder.spinDuration( ) org.infinispan.configuration.cache.ExpirationConfigurationBuilder.lifespan( ) org.infinispan.configuration.cache.ExpirationConfigurationBuilder.maxIdle( ) org.infinispan.configuration.cache.ExpirationConfigurationBuilder.wakeUpInterval( ) org.infinispan.configuration.cache.FileCacheStoreConfigurationBuilder.fsyncInterval( ) org.infinispan.configuration.cache.FileCacheStoreConfigurationBuilder.lockAcquistionTimeout( ) org.infinispan.configuration.cache.HashConfigurationBuilder.rehashRpcTimeout( ) org.infinispan.configuration.cache.HashConfigurationBuilder.rehashWait( ) org.infinispan.configuration.cache.L1ConfigurationBuilder.lifespan( ) org.infinispan.configuration.cache.L1ConfigurationBuilder.cleanupTaskFrequency( ) org.infinispan.configuration.cache.LockingConfigurationBuilder.lockAcquisitionTimeout( ) org.infinispan.configuration.cache.SingletonStoreConfigurationBuilder.pushStateTimeout( ) org.infinispan.configuration.cache.StateTransferConfigurationBuilder.timeout( ) org.infinispan.configuration.cache.SyncConfigurationBuilder.replTimeout( ) org.infinispan.configuration.cache.TransactionConfigurationBuilder.cacheStopTimeout( )	1	1	https://issues.jboss.org/browse/ISPN-2104
8018	ISPN-2085	Has this test been disabled?	1	1	https://issues.jboss.org/browse/ISPN-2085
8019	ISPN-2085	No, should we disable it? sorry not sure what the final result of the massive thread was.	1	1	https://issues.jboss.org/browse/ISPN-2085
8020	ISPN-2044	I'd avoid wrapping, that takes more memory. Taking for example PutKeyValueCommand, the Flags are transmitted to other nodes as part of {code} Object[] getParameters(){code} still the command is extending org.infinispan.commands.AbstractFlagAffectedCommand would it not be more sensible to invoke the Externalizer explicitly rather than letting River figure out the type from the Object[] stored reference?	1	1	https://issues.jboss.org/browse/ISPN-2044
8021	ISPN-2038	Great. I was not aware that single thread performs all of the RPC for a given transaction. Does this mean that even though I have async replication, with sync listeners I will get all changes within a transaction in order ? Just by my curisity, If any JMS implementation does support the above functionality for async listeners, How would it know that all messages related to a transaction have already been received ? (Since replication mode is async and listener is async) Or is this functionality only applicable for sync replication and async notification ?	1	1	https://issues.jboss.org/browse/ISPN-2038

#	Requirement	Questions	Answered	Successfully Answered	URL
8022	ISPN-2037	one off quickly <sannegrinovero> vblagoje, just read your ISPN-2037 .. good timing ! let's talk about it? <jbossbot> jira [ISPN-2037] Map/Reduce tasks should process entries in the CacheLoader as well [Open (Unresolved) Feature Request, Critical, Vladimir Blagojevic] https://issues.jboss.org/browse/ISPN-2037 <vblagoje> ok sannegrinovero <vblagoje> sure, have time now <vblagoje> ? <sannegrinovero> yes <sannegrinovero> so the problem is that you have to load all entries from the cacheloader <sannegrinovero> but you don't actually know which keys exist <vblagoje> aha i see <vblagoje> side question, does keySet return stuff in cache loader? <sannegrinovero> so you have to load "Select * from database", which will very likely kill you by OOM <sannegrinovero> I don't know that <vblagoje> hmmm, how are these caches configured to begin with? any eviction, passivation, shared/single etc? <vblagoje> shared/single cache loader <vblagoje> sannegrinovero, who is most familiar with cahceloaders on our team? <sannegrinovero> mmarkus is the leader I think.. ttarrant accumulated some experience too. <sannegrinovero> sorry was away, back now. <sannegrinovero> Does it matter how caches are configured? <sannegrinovero> for sure if it's a single shared cacheloader, you'll have the problem that you only want to load entries for the current node, skipping all the rest. <sannegrinovero> otherwise you definitely kill the purpose of map reduce.. loading the whole database from a single slow store N times where N is the number of nodes in the cluster... very bad :D <sannegrinovero> vblagoje, ^ <vblagoje> so basically the problem is when map command arrives to node where it is executed and it needs to load all local keys, right now it does not load keys in cache loader - you say, is that right? <mmarkus> vblagoje: anything I can help with? <vblagoje> this is the case of map/reduce when input keys are not specified - basically do map/reduce on all keys <sannegrinovero> vblagoje, I'm not sure. maybe it works but I don't think so. we definitely need to have tests with Map/Reduce && cacheloader interactions. <sannegrinovero> yes, and for Query I never have the keys. I think even for map/reduce it's the most common case (to not have the set of keys) <sannegrinovero> on the other hand, with non-shared cacheloaders it's even trickier. <mmarkus> vblagoje: keySet ignores the cache loader interceptor <vblagoje> yeah ok <vblagoje> aha mmarkus, thanks dude; how do we get them? <mmarkus> vblagoje: only returns what's in memory <mmarkus> you need a way to get all the keys from a cache loader? <vblagoje> exactly <sannegrinovero> this is strongly related guys: https://community.jboss.org/wiki/CacheLoaderAndCacheStoreSPIRedesign#comment-9835 <mmarkus> vblagoje: that's a hard one man <sannegrinovero> I had sketched an API proposal on the mailing list <sannegrinovero> which manik had +1'ed .. searching for a link <sannegrinovero> vblagoje: http://lists.jboss.org/pipermail/infinispan-dev/2012-May/010760.html <vblagoje> mmarkus, sanne , maybe we can use loadAllKeys on CacheLoader interface and then slowly load values for each of those keys <mmarkus> vblagoje: thinkig about it from an jdbc perspective <sannegrinovero> vblagoje, you are the map/reduce API master :) the proposal I described has a similar API in concept, that the visitor "demands" collection of entries from the CacheLoader implementation, in blocks. <sannegrinovero> so the details of how the blocks of data are loaded are delegated to the implementation, which is important <vblagoje> you mean this Processor API sanne? <sannegrinovero> but the flow control is handled by your consumer <mmarkus> sannegrinovero: +1. I was about to say the same thing :) <sannegrinovero> mmarkus, nice :) <sannegrinovero> vblagoje, yes <mmarkus> sannegrinovero: indeed, it was ur idea :) <sannegrinovero> guys am I the only one wasting my dev days by writing on design on the ML :D ? <vblagoje> No, we all are hahaha <sannegrinovero> admittedly this was assigned to Manik ;) <mmarkus> sannegrinovero: you do have a fare share though :) <vblagoje> there are so many things to do noone has time to review anyone else's design <vblagoje> so Manik was going to do these Processor callbacks? <sannegrinovero> yea I realize that, I admit I was hooked to this as I need it. <sannegrinovero> no vblagoje I think I moved it to	1	1	https://issues.jboss.org/browse/ISPN-2037
8023	ISPN-2015	@Galder so the memory consumation looks ok without your decoder ?	1	1	https://issues.jboss.org/browse/ISPN-2015
8024	ISPN-2015	@Norman, which server did you test exactly? Still Memcached? And what kind of loads did you apply? Btw, I was testing JDG, which is the data grid server, which deploys the memcached, hotrod and rest endpoints. You can build it from https://github.com/infinispan/jdg. Shouldn't be a difference.	1	1	https://issues.jboss.org/browse/ISPN-2015
8025	ISPN-2015	Ok, can you try startMemcachedAgent.sh and start the Memcached server (-r memcached)? What load are you using? By default this is 400 ops per thread, with 100 threads.	1	1	https://issues.jboss.org/browse/ISPN-2015
8026	ISPN-2015	@Galder I fixed this and deployed a snapshot which contains the fix. Could you please try this out and let me know if it works ? You can find the snapshot here: https://oss.sonatype.org/content/repositories/snapshots/io/netty/netty/3.4.5.Final-SNAPSHOT/	1	1	https://issues.jboss.org/browse/ISPN-2015
8027	ISPN-2015	MP Swift <mp-swift@outlook.com> made a comment on [bug 899966|https://bugzilla.redhat.com/show_bug.cgi?id=899966] Hi, Anyone, can help me? Can I ask the org.jboss.netty.channel.Channel support getRemoteAddress support IPv6? Hope there is someone can assist me with that =) From Swift	0	0	https://issues.jboss.org/browse/ISPN-2015
8028	ISPN-1967	Now I'm using attribute "View" at jgroups:type=channel,cluster="default" So for example, would be nice to have something like "PendingView" and "CommitedView" at jboss.infinispan:type=Cache,name="testCache(dist_sync)",manager="default",component=DistributedStateTransferManager Is such thing feasible ?	0	0	https://issues.jboss.org/browse/ISPN-1967
8029	ISPN-1959	Manik, Sane, can I assign this to me?	1	1	https://issues.jboss.org/browse/ISPN-1959
8030	ISPN-1956	to get the goal/scope of this issue specified in a more precise manner: Forge may be setup to use INFINISPAN as persistence provider: {code} persistence setup --provider INFINISPAN --container JBOSS_AS7 {code} So what kind of inifinispan projects shall be bootstrapped?	1	1	https://issues.jboss.org/browse/ISPN-1956
8031	ISPN-1956	Have you had a look at the maven archetypes we have? I would think follow something similar.	0	0	https://issues.jboss.org/browse/ISPN-1956
8032	ISPN-1956	Looking at the archetype(s) and forge intensly I started creating a simple forge plugin that would be able to create 4 different examples that match what's currently the Application class that by the archetype (basicUse, lifespans, asyncOperations, registeringListeners). I guess that from there extending stuff would be pretty easy. A first enhancement could be some options that would tune the infinispan config that you get. Does this sound right to you? I'll push this stuff to github anytime soon.	1	1	https://issues.jboss.org/browse/ISPN-1956
8033	ISPN-1956	[~adietish], with the help of [~gastaldi] and [~lincolnthree] we created a prototype infinispan-forge plugin https://github.com/infinispan/infinispan-forge/ Can we use that to host your work ?	1	1	https://issues.jboss.org/browse/ISPN-1956
8034	ISPN-1956	sure, I'll make sure I move my current code into that stub if that's ok for you. Shall we proceed in usual git manner, me forking this repo or do you want to give me write premissions?	0	0	https://issues.jboss.org/browse/ISPN-1956
8035	ISPN-1956	I'm currently have a working base structure. I'm polishing, fixing stuff and will then file a pull request. The plugin is currently named *infinispan-forge*, shouldn't it be called forge-infinispan (as all others are? ex. forge-errai, forge-plugin-primefaces, etc.)? The base package is *org.jboss.infinispan.forge*. Shouldn't it be org.infinispan.forge (as the whole infinispan source itself)?	0	0	https://issues.jboss.org/browse/ISPN-1956
8036	ISPN-1956	[~aanderson] This is integrated, what about closing the JIRA and creating new ones for bugs/other features?	1	1	https://issues.jboss.org/browse/ISPN-1956
8037	ISPN-1935	Hello - do we have a solution for this? I am particularly interested in the callback implementation.	0	0	https://issues.jboss.org/browse/ISPN-1935
8038	ISPN-1901	Shouldn't this be in 5.1.x as well?	0	0	https://issues.jboss.org/browse/ISPN-1901
8039	ISPN-1863	Does it really matter if (http://download.oracle.com/maven/com/sleepycat/je/5.0.34/je-5.0.34.pom) is unstable, because when nexus has fetched the artifact, you already have a copy in your repo.	1	1	https://issues.jboss.org/browse/ISPN-1863

#	Requirement	Questions	Answered	Successfully Answered	URL
8040	ISPN-1861	Galder, could I assign this to you?	1	1	https://issues.jboss.org/browse/ISPN-1861
8041	ISPN-1836	Mircea, can you correct this? It says the same thing on the TransactionConfigurationBuilder, which is where I took the xsd version from.	0	0	https://issues.jboss.org/browse/ISPN-1836
8042	ISPN-1835	Is this out of date?	0	0	https://issues.jboss.org/browse/ISPN-1835
8043	ISPN-1812	Thanks Bela, I noticed something was not right but I did not know what! Should we wait for alpha2?	1	1	https://issues.jboss.org/browse/ISPN-1812
8044	ISPN-1812	What do you mean: what wasn't right ? My comment was just regarding a performance improvement, otherwise this should be fine !	0	0	https://issues.jboss.org/browse/ISPN-1812
8045	ISPN-1799	Can this be fixed so that it gets fixed in next ER build for JDG?	1	1	https://issues.jboss.org/browse/ISPN-1799
8046	ISPN-1799	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 765759|https://bugzilla.redhat.com/show_bug.cgi?id=765759] Prabhat Jha <prabhat.jha@jboss.com> made a comment on jira ISPN-1799 Can this be fixed so that it gets fixed in next ER build for JDG?	1	1	https://issues.jboss.org/browse/ISPN-1799
8047	ISPN-1799	Dan could you please try and get it in 5.1.4.CR1 or at least a patch for JDG	1	1	https://issues.jboss.org/browse/ISPN-1799
8048	ISPN-1799	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 765759|https://bugzilla.redhat.com/show_bug.cgi?id=765759] Tristan Tarrant <ttarrant@redhat.com> made a comment on jira ISPN-1799 Dan could you please try and get it in 5.1.4.CR1 ?	1	1	https://issues.jboss.org/browse/ISPN-1799
8049	ISPN-1799	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 765759|https://bugzilla.redhat.com/show_bug.cgi?id=765759] Tristan Tarrant <ttarrant@redhat.com> made a comment on jira ISPN-1799 Dan could you please try and get it in 5.1.4.CR1 or at least a patch for JDG	1	1	https://issues.jboss.org/browse/ISPN-1799
8050	ISPN-1799	we saw this error in our 3 node cluster in which the network connectivity to one of nodes is not reliable and has high latency. Does this exception have anything to do with the network connectivity and high latency?	1	1	https://issues.jboss.org/browse/ISPN-1799
8051	ISPN-1799	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 765759|https://bugzilla.redhat.com/show_bug.cgi?id=765759] dex chen <dex80526@yahoo.com> made a comment on jira ISPN-1799 we saw this error in our 3 node cluster in which the network connectivity to one of nodes is not reliable and has high latency. Does this exception have anything to do with the network connectivity and high latency?	0	0	https://issues.jboss.org/browse/ISPN-1799
8052	ISPN-1797	Great, thanks! Could you fork the project on GitHub, commit your changes and issue a pull request?	0	0	https://issues.jboss.org/browse/ISPN-1797
8053	ISPN-1797	Hi Andrew, did you make any progress on this? We could help finish it, as we need it soon and there are other volunteers willing to help on it.	1	1	https://issues.jboss.org/browse/ISPN-1797
8054	ISPN-1792	Ok, I can live with that. Closing the feature request?	1	0	https://issues.jboss.org/browse/ISPN-1792
8055	ISPN-1748	Mircea, did you try increasing the OOB thread pool?	1	0	https://issues.jboss.org/browse/ISPN-1748
8056	ISPN-1744	What is this exactly about?	1	0	https://issues.jboss.org/browse/ISPN-1744
8057	ISPN-1732	Why is this critical ? You're assuming that 3.0.2 has bad performance, which we haven't proven yet. Sanne's test (with Manik's Transactional test) actually showed no performance increase in 3.0.3 compared to 3.0.2...	1	1	https://issues.jboss.org/browse/ISPN-1732
8058	ISPN-1732	We where assuming there was a performance regression in JGroups 3.0.2, and we where assuming we could fix this in 3.0.3. But talking to Bela it seems that while it's true there is some contention in the RetransmissionBuffer, this isn't possibly a regression as it has been like this for a while. Shall I close this?	1	1	https://issues.jboss.org/browse/ISPN-1732
8059	ISPN-1636	Guys, why is this issue marked for 6.0? We should do it sooner! I made necessary changes on my fork https://github.com/vblagoje/infinispan/tree/t_1636_m and all we need to do is wire in CDI into SPI. Also issue ISPN-1634 depends on ISPN-1634 and it does not make sense to solve ISPN-1634 through a mechanism other than ISPN-1636.	1	1	https://issues.jboss.org/browse/ISPN-1636
8060	ISPN-1636	Why unrealistic Sanne?! I already did the majority of the work needed to enable the hookup! Kevin, I am moving this for 5.2! Let us know any estimates/conclusion after your initial investigation and assessment!	1	1	https://issues.jboss.org/browse/ISPN-1636
8061	ISPN-1636	Sorry to interfere but if this issue is expected for 5.2, may I still expect to have at least ISPN-1634 for 5.1?	1	1	https://issues.jboss.org/browse/ISPN-1636
8062	ISPN-1636	I've just checked what you did in the context of this issue and what is weird (functionally speaking) is the fact that in {{WordCountMapReduceCDITest}} you create a {{MapReduceTask}} on a cache called {{mapreducecache}} but what you actually inject in {{WordCountMapper}} and {{WordCountReducer}} is the default cache. I know that when we use the annotation {{Inject}}, we inject the default cache but in this context I personally found it very misleading, I believe that everybody expects to get the cache from which the {{MapReduceTask}} has been created, don't you agree?	1	1	https://issues.jboss.org/browse/ISPN-1636
8063	ISPN-1636	I do, it was an honest oversight on my part! Do we need any extra support to inject particular cache or is it already supported by our CDI extension?	1	0	https://issues.jboss.org/browse/ISPN-1636
8064	ISPN-1632	@Pete, is there any easy way to do what Manik suggests?	0	0	https://issues.jboss.org/browse/ISPN-1632
8065	ISPN-1616	It only reduces the number of RPCs if you have a pessimistically locked cache. Not otherwise.	1	1	https://issues.jboss.org/browse/ISPN-1616
8066	ISPN-1580	Isn't Sybase just a variant of MSSQL (or rather the other way round) ? So this should be fairly simple	1	1	https://issues.jboss.org/browse/ISPN-1580
8067	ISPN-1523	I don't understand - what is incorrect about this?	1	1	https://issues.jboss.org/browse/ISPN-1523
8068	ISPN-1492	maybe I'm missing something; is the thread needed for the entries inserted with a limited lifetime? I can of course disable the thread creation using reaperEnabled="false"	1	0	https://issues.jboss.org/browse/ISPN-1492
8069	ISPN-1453	We do have ispncon ... https://github.com/infinispan/ispncon Is that a valid starting point?	1	1	https://issues.jboss.org/browse/ISPN-1453
8070	ISPN-1430	Huh? This is on the current master and doesn't return any value: http://bit.ly/s7AWOU	1	1	https://issues.jboss.org/browse/ISPN-1430
8071	ISPN-1421	Guys, I've created https://docs.jboss.org/author/display/ISPN/Upgrading+from+5.0+to+5.1 with a few that came to my mind. Can you please update the document with any missing? Cheers :)	1	1	https://issues.jboss.org/browse/ISPN-1421
8072	ISPN-1416	[~manik] what kind of information would it need from the InvocationContext?	0	0	https://issues.jboss.org/browse/ISPN-1416
8073	ISPN-1413	We currently have: {code:java} Cache appCache = cacheManager.getCache("appCache").getAdvancedCache().withClassLoader(appClassLoader); //use the appCache instance which uses the specified class appClassLoader for marshalling, class loading etc. {code} Doesn't this cover the use case for this JIRA?	0	0	https://issues.jboss.org/browse/ISPN-1413
8074	ISPN-1407	So this functionality is needed in order to switch the HR clients connected to cluster A (running e.g. ISPN 5.2) to cluster B (running e.g. ISPN 5.3) when doing rolling upgrades. This involves an sys admin connecting to every client and manually updating this list. What I'm suggesting is a potentially more elegant approach to achieve this, by upgrading the HR protocol: - sys admin connects(JMX) to cluster A and marks the site as "upgraded to cluster B" - at this stage the HR endpoints on cluster A, when receiving requests from connected users, would also tell them(piggyback, same as we do with topology info) to connect to "cluster B" on next requests - this way the sys admin does not have to update the list of servers for all the users (there might be lots of them), but only update it once on the server side	1	1	https://issues.jboss.org/browse/ISPN-1407

#	Requirement	Questions	Answered	Successfully Answered	URL
8075	ISPN-1407	Why do we need HR protocol changes ? Can't we just fill the topology cache on cluster A with the addresses from cluster B and make sure all clients perform a GET to retrieve the new topology view ?	0	0	https://issues.jboss.org/browse/ISPN-1407
8076	ISPN-1394	Regarding the removal of nodes, why don't we tell the system which nodes we intend to remove beforehand, so it can actually work out if there might be inconsistencies (e.g. insufficient numOwners).	0	0	https://issues.jboss.org/browse/ISPN-1394
8077	ISPN-1391	Is this just a drop-in replacement?	0	0	https://issues.jboss.org/browse/ISPN-1391
8078	ISPN-1385	testStreamingAPIReusingStreams(org.infinispan.loaders.decorators.ChainingCacheLoaderTest) org.infinispan.loaders.CacheLoaderException: com.thoughtworks.xstream.converters.ConversionException: Cannot construct org.infinispan.container.entries.ImmortalCacheEntry as it does not have a no-args constructor : Cannot construct org.infinispan.container.entries.ImmortalCacheEntry as it does not have a no-args constructor ---- Debugging information ---- message : Cannot construct org.infinispan.container.entries.ImmortalCacheEntry as it does not have a no-args constructor cause-exception : com.thoughtworks.xstream.converters.reflection.ObjectAccessException cause-message : Cannot construct org.infinispan.container.entries.ImmortalCacheEntry as it does not have a no-args constructor class : org.infinispan.container.entries.ImmortalCacheEntry required-type : org.infinispan.container.entries.ImmortalCacheEntry path : /org.infinispan.container.entries.ImmortalCacheEntry line number : 1 ---------- at org.infinispan.loaders.dummy.DummyInMemoryCacheStore.fromStream(DummyInMemoryCacheStore.java:91) at org.infinispan.loaders.decorators.ChainingCacheStore.fromStream(ChainingCacheStore.java:73) at org.infinispan.loaders.BaseCacheStoreTest.testStreamingAPIReusingStreams(BaseCacheStoreTest.java:595) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.testng.internal.MethodHelper.invokeMethod(MethodHelper.java:644) at org.testng.internal.Invoker.invokeMethod(Invoker.java:546) at org.testng.internal.Invoker.invokeTestMethod(Invoker.java:700) at org.testng.internal.Invoker.invokeTestMethods(Invoker.java:1002) at org.testng.internal.TestMethodWorker.invokeTestMethods(TestMethodWorker.java:137) at org.testng.internal.TestMethodWorker.run(TestMethodWorker.java:121) at org.testng.TestRunner.runWorkers(TestRunner.java:909) at org.testng.TestRunner.privateRun(TestRunner.java:618) at org.testng.TestRunner.run(TestRunner.java:499) at org.testng.SuiteRunner.runTest(SuiteRunner.java:332) at org.testng.SuiteRunner.access$000(SuiteRunner.java:33) at org.testng.SuiteRunner$SuiteWorker.run(SuiteRunner.java:358) at org.testng.internal.thread.ThreadUtil$CountDownLatchedRunnable.run(ThreadUtil.java:142) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:603) at java.lang.Thread.run(Thread.java:722) Caused by: com.thoughtworks.xstream.converters.ConversionException: Cannot construct org.infinispan.container.entries.ImmortalCacheEntry as it does not have a no-args constructor : Cannot construct org.infinispan.container.entries.ImmortalCacheEntry as it does not have a no-args constructor ---- Debugging information ---- message : Cannot construct org.infinispan.container.entries.ImmortalCacheEntry as it does not have a no-args constructor cause-exception : com.thoughtworks.xstream.converters.reflection.ObjectAccessException cause-message : Cannot construct org.infinispan.container.entries.ImmortalCacheEntry as it does not have a no-args constructor class : org.infinispan.container.entries.ImmortalCacheEntry required-type : org.infinispan.container.entries.ImmortalCacheEntry path : /org.infinispan.container.entries.ImmortalCacheEntry line number : 1 ---------- at com.thoughtworks.xstream.core.TreeUnmarshaller.convert(TreeUnmarshaller.java:89) at	0	0	https://issues.jboss.org/browse/ISPN-1385
8079	ISPN-1367	And is this crucial for 5.1.0?	0	0	https://issues.jboss.org/browse/ISPN-1367
8080	ISPN-1362	Apologies, I just saw that you are using 5.0.0.Final. Just that your config file had the 4.2 XSD which is what threw me off. Can you confirm that you are in fact using 5.0.0.Final?	1	1	https://issues.jboss.org/browse/ISPN-1362
8081	ISPN-1354	Is this really the job of Infinispan? What's HornetQ there for?	0	0	https://issues.jboss.org/browse/ISPN-1354
8082	ISPN-1338	I can't reproduce the bug. Could you provide a test project?	0	0	https://issues.jboss.org/browse/ISPN-1338
8083	ISPN-1335	I now know that a cache store is made "read only" by setting ignoreModifications on <loader> element. What is not so clear is the use case - why would one use read only store.	0	0	https://issues.jboss.org/browse/ISPN-1335
8084	ISPN-1326	As the cacheName is now needed to create Command instances, the org.infinispan.commands.module.ModuleCommandFactory should pass the cacheName to the method {code} org.infinispan.commands.module.ModuleCommandFactory.fromStream(String, byte, Object[]){code} Or implementors won't be able to create any command. Unfortunately this is not covered by core's test suite, but you can have a look into the Query tests (currently doesn't compile because of the field change in _org.infinispan.commands.remote.BaseRpcCommand_	1	1	https://issues.jboss.org/browse/ISPN-1326
8085	ISPN-1308	Why do you need a separate project for this? Couldn't you include it within cdi/src/test... ? All dependencies would be marked as test scope, so you would not leak anything to production users.	1	1	https://issues.jboss.org/browse/ISPN-1308
8086	ISPN-1303	Rather than attaching src code, could you pls fork the project on GitHub, create a new topic branch on your fork and commit your changes there? That makes it easy for me to do a diff and pull in commits into the main project. :)	0	0	https://issues.jboss.org/browse/ISPN-1303
8087	ISPN-1303	Hi Robert, your patch is interesting, but it's pretty hard to understand which changes you made compared to the original one, as the FileCacheStore in Infinispan is today likely quite different than the one you started from. If you're not familiar or interested in Git, could you send a patch or diff file from the original file you edited, and add some comments? If you could apply your changes to the latest version of FileCacheStore that would be ideal.	1	1	https://issues.jboss.org/browse/ISPN-1303
8088	ISPN-1302	Can I assign this to you, Robert, since you spotted and reported the original JIRAs?	1	0	https://issues.jboss.org/browse/ISPN-1302
8089	ISPN-1300	Sorry for digging this one up but I tried setting the maximum number of buckets to 4096 (for an ext3 file system) and that made Infinispan infinitely slower. When this is configurable, could this be done in a way that's free of side-effects?	0	0	https://issues.jboss.org/browse/ISPN-1300
8090	ISPN-1297	It may not be enabled by default, but do we have the option to enable it?	1	1	https://issues.jboss.org/browse/ISPN-1297
8091	ISPN-1284	Should this be marked resolved? The linked pull request doesn't seem related.	0	0	https://issues.jboss.org/browse/ISPN-1284
8092	ISPN-1261	[~mircea.markus], [~galderz]: Could you please take a look at the following tasks: * Support all events - Done * Verify how cache and cache manager observers are registered - What exactly needs to be done? * Tests - Added for each type of event. * Doc - Added Java docs. Currently there is no information about event support in our documentation: http://infinispan.org/docs/7.0.x/user_guide/user_guide.html. Perhaps we should add it? * Remove deprecated event - Done, but not sure if we should do this... For backwards compatibility we might leave it?	0	0	https://issues.jboss.org/browse/ISPN-1261
8093	ISPN-1253	Is this still needed? IOW, would you guys still like this integration?	1	0	https://issues.jboss.org/browse/ISPN-1253
8094	ISPN-1239	Shouldn't this be closed as well?	0	0	https://issues.jboss.org/browse/ISPN-1239

#	Requirement	Questions	Answered	Successfully Answered	URL
8095	ISPN-1228	Infinispan already supports this. See this unit test that proves it: https://github.com/infinispan/infinispan/blob/master/core/src/test/java/org/infinispan/config/StringPropertyReplacementTest.java which makes use of this test config file: https://github.com/infinispan/infinispan/blob/master/core/src/test/resources/configs/string-property-replaced.xml Please try this out and reopen if it doesn't work for you, and we will investigate.	0	0	https://issues.jboss.org/browse/ISPN-1228
8096	ISPN-1212	Galder, is this issue not solved?	1	1	https://issues.jboss.org/browse/ISPN-1212
8097	ISPN-1168	Are we also going to ship with an impl that uses TransactionSyncRegistry? Or is this something that is going to be impl'd and shipped with JBoss AS?	0	0	https://issues.jboss.org/browse/ISPN-1168
8098	ISPN-1158	Manik, in the end this was implemented as throwing an exception rather than a warning. Exception seems a bit too much to me and couldn't it be a case that a user wants passivation and do manual Cache.evict() calls? WDYT? I think a WARN message would be better.	1	1	https://issues.jboss.org/browse/ISPN-1158
8099	ISPN-1149	@Jeff, thanks for this feature request. Have you discussed this on the user forums and/or developer mail list? One point of concern that I would have in auto-assigning ports on a server daemon process is how would you inform clients of this port?	1	0	https://issues.jboss.org/browse/ISPN-1149
8100	ISPN-1147	When can we expect this ?	1	0	https://issues.jboss.org/browse/ISPN-1147
8101	ISPN-1141	Not really a dupe of ISPN-871. ISPN-871 attempts to use a JNI-based heap in C. Magnus' approach above is a NIO direct buffer. Magnus, would you be interested in contributing a patch? :-)	0	0	https://issues.jboss.org/browse/ISPN-1141
8102	ISPN-1141	Can the DirectMemory project (https://github.com/raffaeleguidi/DirectMemory/wiki) be used for this?	1	1	https://issues.jboss.org/browse/ISPN-1141
8103	ISPN-1137	How is this related to ISPN-1136? Could you pls close this and use ISPN-1136 if they are duplicates? Make sure all of these notes are copied across.	1	1	https://issues.jboss.org/browse/ISPN-1137
8104	ISPN-1132	Suggested design. This feature is only available if *lockAtPrepareTime*, as defined within ISPN-1131, is enabled. It should be possible to control weather lock-reordering is enabled or not: e.g. there might be application where locks are being acquired in a ordered sequence already, so this optimization would not bring any benefit. The element that controls lock reordering is *reorderLockAcquisition*: {code:xml} <locking lockAtPrepareTime="true"> <reorderLockAcquisition enabled="true" keyComparator="extends.java.util.Comaparator"/> </locking> {code} *keyComparator* is optional. It is an {{java.util.Comaparator<K>}} implementation, with K from {{Cache<K,V>}}. If not specified defaults to a implementation that: * for each key calculates the CH value * compares the CH values of two keys in order to determine the order. Having *lockAtPrepareTime* disabled and *reorderLockAcquisition* enabled makes the system log a warning stating that lock reordering is not active. In order to support this feature {{TxInterceptor.visitPrepareCommand}} is modified as follows: * it would first invoke {{invokeNextInterceptor(ctx, command)}} and only then {code:java} if (!ctx.isOriginLocal()) { // replay modifications for (VisitableCommand modification : command.getModifications()) { VisitableCommand toReplay = getCommandToReplay(modification); if (toReplay != null) { invokeNextInterceptor(ctx, toReplay); } } } {code} * this is needed in order to have {{LockInterceptor.visitPrepareCommand}} acquire locks at once for all keys involved in the transaction Also {{PrepareTimeLockingInterceptor.visitPrepareCommand}} is enhanced as follows: * {code:java} for (WriteCommand wc : command.getModifications()) { acquireLock(wc); } {code} is replaced with {code:java} for (WriteCommand wc : reorder(command.getModifications())) { acquireLock(wc); } {code} * {{reorder(WriteCommand[] commandsToReorder)}} works as follows: ** if *reorderLockAcquisition* is disabled returns {{commandsToReorder}} ** if {{commandsToReorder}} contrains a {{ClearCommand}} it returns {{commandsToReorder}} (unmodified) ** breaks {{PutMapCommand}} in individual {{PutKeyValue}} commands ** after the step above we have a sequence of {{DataCommands}} - based on the original {{commandsToReorder}} sequence, with {{PutMapCommand}} "expanded" ** this sequence is the coalesced (a la ISPN-328), resulting in another sequence of {{DataCommands}} ** A List<K> is then determined from last {{DataCommands}} sequence by using {{DataCommands.getKey()}}. (K from Cache<K,V>) ** This list of keys is then ordered using the configured *keyComparator* ** locks are being acquired in the sequence resulted at the previous step	0	0	https://issues.jboss.org/browse/ISPN-1132
8105	ISPN-1123	Fixing ConsistentHashV1IntegrationTest too in the Hot Rod client. Why on earth was it transactional if we don't support transactions in Hot Rod yet? Also tightened the interceptor used to keep counter of calls.	1	1	https://issues.jboss.org/browse/ISPN-1123
8106	ISPN-1116	Manik, how could we solve this as it's marked depending on ISPN-1195, which is not solved ?	1	0	https://issues.jboss.org/browse/ISPN-1116
8107	ISPN-1115	Whys is applyDelta(K deltaAwareValueKey, Delta delta, Object... locksToAcquire) a method on the AdvancedCache? Is there any situation in which a user would call this method directly?	1	1	https://issues.jboss.org/browse/ISPN-1115
8108	ISPN-1115	I don't think so. How else can we invoke a method on cache from a fairly high-abstraction like AtomicHashMapProxy?	1	1	https://issues.jboss.org/browse/ISPN-1115
8109	ISPN-1115	Just thinking out loud, haven't read the design preview email on the dev list: Does the method need to be in the Cache? Couldn't the method be in a particular component which AtomicHashMapProxy gets injected?	1	1	https://issues.jboss.org/browse/ISPN-1115
8110	ISPN-1103	I completely agree that storing metadata as a JSON structure is a brilliant approach. I've added several features the list in the description. Storing the jsonObject as a string might work, especially if it is simply loaded and stored as an atomic unit. However, there are several advantages to storing it in a BSON (or BSON-like) representation, and all boil down to the fact that any schema-aware service will need to access the document contents for indexing, validation, computing differences (for DeltaAware functionality), and even applying changes (e.g., something similar to MongoDB's [atomic operations|http://www.mongodb.org/display/DOCS/Atomic+Operations]). Using a BSON representation will also make it more natural to handle binary values within the document. Now, if we did store the document internally as [BSON|http://www.bson.org], we actually don't need to store the metadata separately. For example, we can store a single document with this structure: {code} { "document" : /* user's document */ "metadata" : { "schema-ref" : blah-blah ... } {code} If represented as a [BSONObject|https://github.com/mongodb/mongo-java-driver/blob/master/src/main/org/bson/BSONObject.java], then the user's document is actually a nested BSON object stored under the "document" property name. This approach means that all the differencing, atomic operations, serialization, etc. functionality will work without having to distinguish between a user's document and a system metadata document. The resulting value class would be: {code} class SchematicValue { org.bson.BSONObject json; } {code} The downside of doing this is the increase in size of the {{SchematicValue}} class, from 2 references to a 1 reference plus an extra BSONObject implementation. The {{org.bson.BasicBSONObject}} extends {{LinkedHashMap<String,Object>}} (which is fairly substantial), but we could probably provide an optimized implementation for the top-level that still implemented the {{org.bson.BSONObject}} interface but more directly. Is it worth the apparent simplicity?	1	1	https://issues.jboss.org/browse/ISPN-1103

#	Requirement	Questions	Answered	Successfully Answered	URL
8111	ISPN-1103	What do you mean by map-reduce-based validation exactly? I guess this is not distributed at all but rather happening locallt?	1	0	https://issues.jboss.org/browse/ISPN-1103
8112	ISPN-1096	Fine by me, however what do you propose as a "correct solution" here?	0	0	https://issues.jboss.org/browse/ISPN-1096
8113	ISPN-1090	Good idea for an umbrella JIRA. Could you make sure the others are linked accordingly? These would include: * Eventing over Hot Rod * Transactions over Hot Rod * Map/Reduce over Hot Rod * Querying over Hot Rod We have JIRAs for at least some of the above. :)	1	0	https://issues.jboss.org/browse/ISPN-1090
8114	ISPN-1090	Which API would need to change for this? Just that the SuspectException would need to include the address, right? How is this API? Or are you referring to the HR protocol's error codes?	1	0	https://issues.jboss.org/browse/ISPN-1090
8115	ISPN-1090	To support virtual nodes more efficiently, the hash topology header must change: - -When virtual nodes are enabled, we're sending host:port information repeated for each of the hash ids belonging to a HotRod server. So, the protocol needs to be more efficient in v2 sending the host:port only once and then N hash ids, where N is the number of configured virtual nodes. If virtual nodes is disabled, the default is 1 virtual node, so N=1.- Actually, forget what I just said cos as Manik rightly pointed out in a chat we had earlier, it won't be atypical to configure with 500-1000 vnodes. So, even if we applied the suggested change, that'd still require, for a cluster of 10 nodes, shipping 10*1000*4bytes=~40kb which is not very optimal. During the chat it was suggested that the clients themselves generate the hash ids, but in order to do that, they need to have the same hash seed that the Infinispan server does, and currently that is org.infinispan.remoting.transport.Address which represents the cluster address. Clearly, this won't work for clients but what about if the server hashed on String UTF-8 version of org.infinispan.remoting.transport.Address? If that worked, then a hash topology update need only shipping: existing header params plus num virtual nodes, and tuples for host:port:clusterAddress(UTF-8). The key thing here is that the number of tuples to send back would be equal to the number of nodes in the cluster. The clients could then generate the hash codes for all virtual nodes themselves. At the moment, Infinispan cluster addresses are backed by JGroups address which are UUID driven, so its toString() representation seems deterministic. However, if there was ever a move back to IP based addresses that would break because toString() representations rely on DNS lookups which means the toString() representation could vary over time. @Pete, could you confirm whether this change on the key on which cluster addresses are hashed would work as expected? The hash topology header size could be further reduced if clients were assumed to be able to keep state in which case the server could just send information about nodes that have joined or left. Since nodes pick fixed positions in the hash wheel, this would in theory work cos the addition of a node would not require other nodes to recalculate hashes.	0	0	https://issues.jboss.org/browse/ISPN-1090
8116	ISPN-1084	This is a new feature - should be scheduled for 5.1.	1	1	https://issues.jboss.org/browse/ISPN-1084
8117	ISPN-1079	Dan, how much do you allocate for your RAM disk?	1	0	https://issues.jboss.org/browse/ISPN-1079
8118	ISPN-1068	Hey [~O.Bergner][~obergner] (sorry wasn't sure which account is active)! Could you tell me whether there is still an interest in this? I would love to help you with this but I'm not sure what should be the final outcome. Ping me please when you have a while (here or off-list if you prefer).	0	0	https://issues.jboss.org/browse/ISPN-1068
8119	ISPN-1065	Pete, when you write such parser, what do you do about validation of the XML? Do you delegate it to the schema? Or do you do that while parsing too? The type of validations I'm refering to is about the number of possible instances of a particular element, the allowed values for an XML attribute...etc. I see some benefit from having docs in schema and that's the fact that IDE or other tools can piggyback on that to offer information to the user while typing the XML, without having to go to the config doc.	0	0	https://issues.jboss.org/browse/ISPN-1065
8120	ISPN-1065	I'd love to be able to add structured configuration options to the Query engine, i.e. not to pass a list of properties but have a set of proper elements which users could validate via XSD. The tricky part is that this structure is defined by the query submodule, so I wonder if we can make this while still having the infinispan-query.jar optional: the parser should be able to self-register extensions for different parts, or ignore unknown elements. I don't expect the XSD to be dynamic too, so this might need to be built upfront from the aggregation of all extension modules. If not possible, what are the alternatives? shall we separate configuration files?	1	1	https://issues.jboss.org/browse/ISPN-1065
8121	ISPN-1059	Yes, I think so. It would be nice if someone with proper permissions could link to my document from the official Infinispan documentation. I can't. But otherwise, yes, this is done.	1	1	https://issues.jboss.org/browse/ISPN-1059
8122	ISPN-1014	Have you reported these issues to the RestEasy community?	1	0	https://issues.jboss.org/browse/ISPN-1014
8123	ISPN-1001	This needs to be tested once ISPN-1000 is closed. @Dan, could you pls coordinate with Bela on testing this issue once you are happy with ISPN-1000.	1	1	https://issues.jboss.org/browse/ISPN-1001
8124	ISPN-1000	Some design ideas: {code} # PUSH based rehashing The purpose behind this is to provide for a better rehashing scheme that is: * more robust * has fewer RPC messages * performs better The current rehashing scheme is RPC-heavy and attempts to be non-blocking, resulting in brittleness. ## Overview * Every node registers a view change listener * On view change, * determine if current node is affected and kick off a RehashTask * If related to a JOIN, RehashTask blocks on JOIN message to grab TopologyInfo (if config is topology-aware). * Enter READ_ONLY mode * if affected, loop through keyset in data container and non-shared cache store. For each key: * if LAST_OWNER on OLD_CH and key has a new owner in NEW_CH, add to push set * if LAST_OWNER is no longer an owner based on NEW_CH, add key to invalidation set (deals with JOINs and LEAVEs in the same way now!) * Push state to new owners * Apply NEW_CH to local node, exit READ_ONLY mode. * If successful, invalidate keys in invalidation set ### Determining if node is affected by ViewChange: * If JOIN: * If idx(JOINER) within + or - numOwners of idx(self) * Or if LEAVE: * and idx(LEAVER) within + or - (numOwners + 1) idx(self) _TODO_: HOw would this affect vnodes? LEAVER or JOINER will have multiple positions to check, and self will have multiple positions too. Can this be encapsulated in the ConsistentHash impl? ### READ_ONLY mode: I think attempting a non-blocking rehash at this stage is way too complex. We have it with the current scheme and we can see how brittle it is. IMO we should go for a blocking model for now and make sure rehashing is stable and robust, and consider non-blocking later. READ_ONLY mode is achieved by using the TransactionLogger. But we won't actually use this to *log* any transactions for now. Instead we can use it to block new transactions and non-transactional writes since it exists in all the appropriate places. ### Parallel joiners and leavers How will this work with parallel joiners and leavers? {code}	0	0	https://issues.jboss.org/browse/ISPN-1000
8125	ISPN-999	Is this still slated for Infinispan 7? I can see it's been bumped from 5.1 to 6 to 7 so far...	1	0	https://issues.jboss.org/browse/ISPN-999
8126	ISPN-996	I was thinking of a @CacheScoped new scope, so that objects may be injected by retrieving from the cache. The lifetime of an object would depend on the life of it on the cache. I implemented something alike in https://github.com/gastaldi/seam-cache WDYT ?	1	1	https://issues.jboss.org/browse/ISPN-996

#	Requirement	Questions	Answered	Successfully Answered	URL
8127	ISPN-996	Obvious issue here which is why this isn't in phase 1 is that Infinispan is not like JBoss POJO Cache and won't update the object in the cache when changes are made to the reference, the user needs to explicitly push the object back into the cache. So given {code} @CacheScoped class Foo { String name; //getters and setters } class Bar { @Inject void init(Foo foo) {foo.setName("George");} } {code} How does this updated copy of foo get back into the cache?	1	1	https://issues.jboss.org/browse/ISPN-996
8128	ISPN-994	"2nd most relevant digit" ? if you consider decimal representation, means that you won't support more then 10 nodes? You would likely need some VInt encoding to make sure you have as much bits as you need. How would you deal with nodes restarting, and so resetting their local owned counter?	1	1	https://issues.jboss.org/browse/ISPN-994
8129	ISPN-993	@Nicolas - can you wait for ISPN-1132?	1	1	https://issues.jboss.org/browse/ISPN-993
8130	ISPN-984	Galder Zamarreño added a comment on the mailing list: Hmmm, isn't that going opening a can of worms with clients expecting information that can be obtained remotely via JMX?	0	0	https://issues.jboss.org/browse/ISPN-984
8131	ISPN-961	Olaf, is this complete now?	1	1	https://issues.jboss.org/browse/ISPN-961
8132	ISPN-958	So whats the snapshot policy around ISPN? When does this get to a snapshot we could depend on (possibly)?	1	0	https://issues.jboss.org/browse/ISPN-958
8133	ISPN-950	Mircea, this has been mentioned before, see ISPN-509. I'd suggest changing ISPN-509 to focus on enabling configuration of listeners via XML, and indicating how Hot Rod servers would be a use case for it. WDYT?	0	0	https://issues.jboss.org/browse/ISPN-950
8134	ISPN-949	Shall we close this? This was an old idea of mine and I actually think now it's not very useful.	1	0	https://issues.jboss.org/browse/ISPN-949
8135	ISPN-949	Why is this not useful?	1	0	https://issues.jboss.org/browse/ISPN-949
8136	ISPN-949	you're right it might be useful in certain cases. Just not sure about a compelling enough use case vs. more API clustter. Ok let's keep it open, but I take it out 5.2	0	0	https://issues.jboss.org/browse/ISPN-949
8137	ISPN-924	Is there any chance that this will be implemented in near future? It would be a feature which certainly would be used by my company. At the moment, i implement a cluster-wide counter with the help of hard locks, i think your version will surely perform better.	1	1	https://issues.jboss.org/browse/ISPN-924
8138	ISPN-905	Patrick, could you please provide use cases for each of this operations so that we can understand better the meaning/aim of them? For example, why do you need cacheExists()/getCacheIfExists()? I assume that it'd return true only if a cache with that name has been started. If stopped it'd returned false. Whether the cache is defined or not in the config would have no bearing here. Wrt remove(), would not cache.stop() be enough for you? For example, stopping a cache removes all its contents. The only thing that cache.stop() could not do but CacheManager.remove() could, would be to update CacheManager's caches collection. Your comment wrt remove() being transaction safe is also something that should prob affect cache.stop() as well. In fact, I think this should be separated into a diff jira. FAO Manik, do you have any background information for this? For some reason it rings a bell that something related to this was discussed in dev/user forums but cannot find the link right now.	1	1	https://issues.jboss.org/browse/ISPN-905
8139	ISPN-905	I see the point for CacheManager.remove(String). There's two things that need some thought: First, how to make CacheManager.remove() wait for on-going txs to finish, within a time period of course (I think Cache.stop() should do as well. although as you said the target is different - I'll open a diff jira for this). The second is how to deal with concurrent CacheManager.remove() calls for the same cache, particularly coming from diff nodes. I don't understand the need for cacheExists()/getCacheIfExists() yet. What is the use case for this from a client's perspective? In other words, what would you do if a cache does not exist? What would you do different?	1	1	https://issues.jboss.org/browse/ISPN-905
8140	ISPN-905	The problem with listeners is that once you've closed a Cache, what is the other node gonna do? Close itself? You'll start to get some hard to deal reentrant @CacheStopped callbacks. Not nice. We avoid using listeners for internal stuff since they're quite restrictive. They're very useful for user code though. Most likely there will be some kind of RPC call. I've just finished with ISPN-915 which is key to making sure Cache.stop()/CacheManager.remove() are tx friendly. Patrick, I still don't get what cacheExists/getCacheIfExists give you even if you plan to use them as store. You're gonna have to give me a tangible example to convince me to add them. As a side note, I can see the point of getting a cache but not starting it (ISPN-627) since that would allow you to hook listeners to Cache before it's started and so receive callbacks related to rehashing/state transfer. So, to convince me, show me what you'd do with those methods, particularly, what would you do different if cacheExists() return false, or getCacheIfExists() returned null.	1	1	https://issues.jboss.org/browse/ISPN-905
8141	ISPN-899	Pete, do u expect to fix this for 4.2.0.CR2 or 5.0.0.ALPHA3?	1	1	https://issues.jboss.org/browse/ISPN-899
8142	ISPN-899	Galder, removing these methods broke the hotrod server. Can you comment if you need these methods there? (It looks to me naively like you do?).	1	1	https://issues.jboss.org/browse/ISPN-899
8143	ISPN-894	Do we wanna update 4.2.x as well?	1	1	https://issues.jboss.org/browse/ISPN-894
8144	ISPN-891	Mircea, is ClusterCacheLoaderConfig bean only used programmatically? Does it need properties?	1	1	https://issues.jboss.org/browse/ISPN-891
8145	ISPN-891	@Valdimir Not sure I understand. The config is as follows: <loaders preload="false"> <loader class="org.infinispan.loaders.cluster.ClusterCacheLoader"> <properties> <property name="remoteCallTimeout" value="20000"/> </properties> </loader> </loaders> How does this impact the ordering between singletonStoreConfig and asyncStoreConfig elements?	1	1	https://issues.jboss.org/browse/ISPN-891
8146	ISPN-872	I get the feeling that the issue here is that at read time, we don't differentiate between a pure byte[] which we shouldn't serialize as opposed one that we should: obj match { case ba: Array[Byte] => Response.ok.`type`("application/x-java-serialized-object").entity(streamIt(_.write(ba))).build case ser: Serializable => Response.ok.`type`("application/x-java-serialized-object").entity(streamIt(new ObjectOutputStream(_).writeObject(ser))).build If we send the byte[] as is, then this is not a java serialized object but it could be as shown by this example. So, if someone sends a byte[] as is, which does not come from having serialized an object, a different type is needed here, maybe "application/octet-stream" or some other pure binary MIME type? I must say that I thought that your test is a bit odd and I've wondered if it actually makes sense. If you have String, why convert it to byte array via getBytes() and then serialize it? You can just take the String and serialize it directly and the resulting object is a String on the server side. I guess what the getBytes() gives you is the confidence that it is being transformed into a particular encoding of your preference which in this case is plattform independent, i.e. UTF-8. Odd but I think deep down is a valid use case.	0	0	https://issues.jboss.org/browse/ISPN-872

#	Requirement	Questions	Answered	Successfully Answered	URL
8147	ISPN-872	Actually, the test is not correct. If you're already retrieving the byte[] from getBytes(), why pass it through ObjectOutputStream? That makes no sense at all. Even in the case that you pass getBytes() directly (which I'd add to the testsuite) currently you'd pass the byte[] as "application/x-java-serialized-object" which will probably fail cos the server side will attempt to deserialize which should fail cos it's not really a Java serialized object. So, the root cause is still present, a different type is needed for pure byte[] arrays that are not serialized such as getBytes() result.	0	0	https://issues.jboss.org/browse/ISPN-872
8148	ISPN-871	Was this ever done? If so, is it documented how to configure an off-heap Cache in a specific ISPN version? If not, I would recommend the consideration of Peter Lawrey's OpenHFT Chronocle net.openhft.collections.HugeHashMap implementation ... which is based on a off-heap MappedByteBuffer onto any/all of (/dev/shm, tmpfs, ext4) and is exceedingly high performant. https://github.com/OpenHFT/HugeCollections	1	1	https://issues.jboss.org/browse/ISPN-871
8149	ISPN-871	Hi, We want to get started w a POC build of doing this using Peter Lawreys HugeHashMap solution (details of the POC build ambition = http://infinispan-developer-list.980875.n3.nabble.com/infinispan-dev-Infinispan-embedded-off-heap-cache-td4026102.html). Any success (of course) to be immediately contributed back for ISPN's consideration to use. ROAD-BLOCK: Can't build! How do you build ISPN 7 from GIT? Tryng this: Assuming all I need is JDK 1.7 anf Maven 3 1. Fork https://github.com/infinispan/infinispan/ to my GIT repo 2. Clone from my GIT repo's ISPN origin to my desktop workspace 3. cd to the infinispan root pom.xml directory 4. Build from my workspace = `mvn clean compile` #this is the ROADBLOCK point - how do we resolve the below unresolved ' infinispan-extended-statistics'? See ben@52-b.liq,risk.ny.jpm.com> mvn clean compile <snip ....> [INFO] ---------- [INFO] BUILD FAILURE [INFO] ---------- [INFO] Total time: 48.156s [INFO] Finished at: Thu Feb 06 00:14:16 EST 2014 [INFO] Final Memory: 50M/361M [INFO] ---------- [ERROR] Failed to execute goal on project infinispan-extended-statistics: Could not resolve dependencies for project org.infinispan:infinispan-extended-statistics:jar:7.0.0-SNAPSHOT: Failure to find org.infinispan:infinispan-core:jar:tests:7.0.0-SNAPSHOT in https://repository.jboss.org/nexus/content/groups/public was cached in the local repository, resolution will not be reattempted until the update interval of jboss-public-repository has elapsed or updates are forced -> [Help 1] [ERROR] [ERROR] To see the full stack trace of the errors, re-run Maven with the -e switch. [ERROR] Re-run Maven using the -X switch to enable full debug logging. [ERROR] [ERROR] For more information about the errors and possible solutions, please read the following articles: [ERROR] [Help 1] http://cwiki.apache.org/confluence/display/MAVEN/DependencyResolutionException [ERROR] [ERROR] After correcting the problems, you can resume the build with the command [ERROR] mvn <goals> -rf :infinispan-extended-statistics ben@52-b.liq,risk.ny.jpm.com>	1	1	https://issues.jboss.org/browse/ISPN-871
8150	ISPN-870	Pete, I've integrated your commit. Do you plan to continue to work under the umbrella of this JIRA, or will you create a separate jira for CR1?	1	1	https://issues.jboss.org/browse/ISPN-870
8151	ISPN-863	Are you suggesting using Runtime.freememory() and Runtime.totalMemory()? This would be dependent on GC cycles though... but yeah, interesting potential. Only, how would you know when to *stop* evicting? Calling freememory() repeatedly in a loop is no help since memory freed by evicting would only be reported after a GC cycle.	1	0	https://issues.jboss.org/browse/ISPN-863
8152	ISPN-863	@Dave - what's the status with this? Are you still working with Vladimir on contributing this patch?	1	0	https://issues.jboss.org/browse/ISPN-863
8153	ISPN-863	I would agree with Robert; I think the best solution is the Instrumentation solution. The Instrumentation solution seems much less expensive than serialization and reflection to me. Although, if you have a write behind cache enabled that is already serializing, I could see then trying to leverage the size assuming the write behind has APIs to do so. Reflection seems way too expensive. It also seems that this wouldn't really need to "kick into gear" unless the JVM heap size is running "low". Maybe consider making the eviction policy something that kicks in when memory begins to consume a configurable portion of the total heap. One idea is to have the cache heap size calculated as a background non-blocking thread, where the heap does not affect "puts" to the cache unless the heap is already running low. If the heap is running low, we switch to "agressive eviction" mode, which blocks writes to the cache until the point that the heap-to-available memory threshold clears. Just thinking out loud. Does anyone know if getObjectSize actually already transverses sub-objects? http://docs.oracle.com/javase/7/docs/api/java/lang/instrument/Instrumentation.html#getObjectSize(java.lang.Object)	1	0	https://issues.jboss.org/browse/ISPN-863
8154	ISPN-863	I am entirely unfamiliar with the implementation details and internal architecture of Infinispan that might make an instrumentation-based eviction approach desirable. That said, if a typical objective is to make it so that caches utilize as much available heap as possible, while never provoking an {{OutOfMemoryError}}, that approach sounds incredibly complicated compared to putting entry values in e.g. {{SoftReference}}s. Is there a reason why such an approach is unworkable in Infinispan, or does it represent an orthogonal potential feature (perhaps because eviction strategies have no control over the object model around cache entries, etc)? Obviously, it doesn't address the notion of providing ways to allocate percentages of the max heap to caches, but I wonder if that's actually desirable given typical use cases and the complexity associated with implementing such controls.	0	0	https://issues.jboss.org/browse/ISPN-863
8155	ISPN-848	Erik, are you aware of http://community.jboss.org/docs/DOC-15604 ?	0	0	https://issues.jboss.org/browse/ISPN-848
8156	ISPN-825	Suggestion from Manik: numInitialRemoteGets (default 1) remoteGetTimeout (default 500ms? What's our average remote GET time anyway?) So with this scheme, we'd: * Randomly select 'numInitialRemoteGets' from dataOwners * Send them the remote get * After remoteGetTimeout, send a remote get to the next data owner, until we run out of data owners * Return the first valid response.	1	1	https://issues.jboss.org/browse/ISPN-825
8157	ISPN-808	Hmm, ConfigurationProperties is not something I want to necessarily expose to user code as API. Could you give me the use cases for something like this?	1	0	https://issues.jboss.org/browse/ISPN-808
8158	ISPN-800	Luca, any news on this? Is this finished already? Or do you need to do more work on it?	1	1	https://issues.jboss.org/browse/ISPN-800
8159	ISPN-800	Luca, I managed to talk to David and he's evaluating the right approach to this problem. Is this the last stumble, or are you waiting on other libraries too?	1	1	https://issues.jboss.org/browse/ISPN-800
8160	ISPN-800	What news are you looking for?	0	0	https://issues.jboss.org/browse/ISPN-800
8161	ISPN-800	Resolution/Fix Version still says resolution is 6.0 which I see is Oct 2013 for this capability. Is this still accurate?	1	1	https://issues.jboss.org/browse/ISPN-800
8162	ISPN-800	Hey [~galder.zamarreno], I'm more than willing to help with some of this. We'd like it for hibernate-infinispan as well (see https://hibernate.atlassian.net/browse/HHH-8214). There were a few ClassLoader issues that came up. What's the status here?	1	1	https://issues.jboss.org/browse/ISPN-800
8163	ISPN-800	[~bibryam], I'm pulling in your changes. Thanks! Here are a couple of questions: 1.) Why was it necessary to make commons a fragment attached to core, rather than ensuring the imports and exports are correct and running commons as a separate bundle? Just curious. 2.) Was allowing BND to automatically define imports causing issues, or do you simply prefer to explicitly list them?	1	1	https://issues.jboss.org/browse/ISPN-800

#	Requirement	Questions	Answered	Successfully Answered	URL
8164	ISPN-800	Would that also mean creating a similar bundle for the hotrod client ?	1	1	https://issues.jboss.org/browse/ISPN-800
8165	ISPN-800	Hey, can you guys take a look at https://github.com/infinispan/infinispan/pull/2457 and the questions within it?	1	1	https://issues.jboss.org/browse/ISPN-800
8166	ISPN-764	Trustin, when you try to upgrade, could you check if ISPN-390 now works?	1	1	https://issues.jboss.org/browse/ISPN-764
8167	ISPN-755	Filip, is the maven plugin finished? Have you tried to generate documentation for Infinispan's scala classes?	1	1	https://issues.jboss.org/browse/ISPN-755
8168	ISPN-721	Once again, similar to comment to what I said on ISPN-717. Why do this? Which 'current implementation' are you talking about? There're a fair few cache store impls out there... :). And again, IIRC, the loading is done by the CacheLoaderInterceptor (CLI) unless I'm mistaken. If a particular CacheStore impl is also loading again, that shouldn't be the case. I agree that it'd be a good optimization to be able to store something and get previous value in just one roundtrip to the cache store, but how many cache stores impls can do that? And how would that work with the current CLI and CSI architecture?	1	0	https://issues.jboss.org/browse/ISPN-721
8169	ISPN-715	thinking about it not sure how that would reduce number of deadlocks? [14:57] mmarkus: or cause deadlocks [14:58] an1310: ? [14:59] an1310: You would only acquire the one lock. If node A and node B both acquired a local lock, but we hashed to node C, wouldn't that be a deadlock? [15:01] mmarkus: so we have tx1 running on A, tx2 running on B and both work on a key that's hashed to C [15:01] mmarkus: is this your scenario? [15:03] an1310: Yes, unless there's some way that I could tell my app to always go to C (since that's where the data lives). [15:04] an1310: I'm colocating all the necessary caches involved in a tx on one data node and using that "eager lock single node" feature. [15:04] an1310: ...data on a cache; typing too fast. [15:04] mmarkus: returning to the example: how does that cause a deadlock? [15:05] mmarkus: i.e. the fact that I also acquire a local lock on A and B for key "k" which is hashed to C [15:06] an1310: I may be wrong about the deadlock part. [15:08] mmarkus: I don't see the benefit of it. Actually I think it would even cause performance faults: if another tx3 runs on A and tries to acquire lock on "k" which is held by tx1 then it would have to go to C and only then to timeout. Right now it would fail fast, without need to go to C	0	0	https://issues.jboss.org/browse/ISPN-715
8170	ISPN-710	FYI - AS6 is now prepped and ready for the upgrade to jgroups 2.11.0. Can we target this upgrade for 4.2.0.CR1?	0	0	https://issues.jboss.org/browse/ISPN-710
8171	ISPN-706	Should this be marked as superseded by ISPN-2281? This issue is still referenced from https://docs.jboss.org/author/display/ISPN/Interacting+With+Hot+Rod+Server+From+Within+Same+JVM	1	0	https://issues.jboss.org/browse/ISPN-706
8172	ISPN-699	I've run the test case attached to https://jira.jboss.org/secure/attachment/12337734/infinispan-hang.tar.gz and after running it a fair few times, I finally was able to reproduce it, i.e.: 2010-10-19 11:58:17,774 ERROR [Scheduled-eviction-thread-0] (FileCacheStore.java:184) - Error while reading from file: /Users/z/Dropbox/JBoss/infinispan/jira/697/infinispan-hang/./store/Users/-202784597 java.lang.NullPointerException at org.infinispan.marshall.jboss.ConstantObjectTable.readObject(ConstantObjectTable.java:253) At this point, I've made a clone of the Users store and tried to read that particular user from another program to verify whether the issue was due to the data that is actually stored being corrupt. The test run fine which indicates that there's no storage corruption. Instead, it seems like there's a race condition at read time that ends up with wrong byte being read. As a FYI, the byte returned should always be 42 which is the id for org.infinispan.loaders.bucket.Bucket. Now, we have to remember that reading/writing to/from a bucket is lock protected, so I can't see right now how this could be the result of concurrent read and write of same bucket. Could it be the result of two concurrent reads maybe? Hmmm, sounds odd. Anyway, I'm gonna add some extra logging and try to replicate the issue again.	1	1	https://issues.jboss.org/browse/ISPN-699
8173	ISPN-699	Hmm, this is puzzling me now. Once I added more information about the exception, I see: org.infinispan.CacheException: Type of data read is unknown. Byte read ('42') is not amongst known reader indexes. at org.infinispan.marshall.jboss.ConstantObjectTable.readObject(ConstantObjectTable.java:259) What the heck? static final byte BUCKET = 42; and @Marshallable(externalizer = Bucket.Externalizer.class, id = Ids.BUCKET) public final class Bucket So, either: - The HashMap contains a null for 42 index - Or the HashMap does not contain the 42 index at all - Or the ConstantObjectTable in use is not the one initially set up This is stored in a simple HashMap, nothing fancy. Adding more logging to this.	1	1	https://issues.jboss.org/browse/ISPN-699
8174	ISPN-694	The patch doesn't seem to notify the entries purged by CacheStore.purgeExpired(). Is it intended or just missed out? If the entries purged by CacheStore.purgeExpired() should also be notified, it will also require a change in CacheStore API, which means this issue needs to be resolved in 5.0. By the way, Instead of adding CacheNotifier parameter to purgeExpired(), I'd modify it to return the set of purged entries (i.e. Set<InternalCacheEntry>).	1	1	https://issues.jboss.org/browse/ISPN-694
8175	ISPN-694	From IRC: {code} Apr 29 14:47:51 <pferraro> mmarkus: whoops - I meant, also ISPN-694 Apr 29 14:47:52 <jbossbot> jira [ISPN-694] Create Expiration Notification [Reopened (Unresolved) Feature Request, Major, Unassigned] https://issues.jboss.org/browse/ISPN-694 Apr 29 14:49:04 <pferraro> mmarkus: I'm in the process of redesigning web session clustering for wildfly, and these 2 jiras stand in the way of leveraging infinispan's native expiration and eviction support. ... Apr 29 18:12:41 <mmarkus1> pferraro: in the scope of ISPN-694, we'd need to make the CacheStore trigger expiration notifications Apr 29 18:12:42 <jbossbot> jira [ISPN-694] Create Expiration Notification [Reopened (Unresolved) Feature Request, Major, Galder Zamarreño] https://issues.jboss.org/browse/ISPN-694 Apr 29 18:13:25 * manik has quit (Quit: Leaving.) Apr 29 18:13:39 <mmarkus1> pferraro: are you using a CustomCache store or one that exists in Infinispan? ... Apr 29 18:27:34 <pferraro> mmarkus1: For my case, if the cache entry was already passivated to a cache store, I only need the expiration notification if a subsequent cache get attempts to load it, but it should have expired. ... Apr 29 18:28:21 <pferraro> mmarkus1: so - for my case, it is enough that the in-memory entries are promptly expired - stuff in the cache store can receive expiration notifications as they are discovered Apr 29 18:28:38 <pferraro> mmarkus1: does that make sense? Apr 29 18:29:55 <pferraro> mmarkus1: So, currently, I'm not using a custom cache store, but rather an existing cache store that emits an expiration notification if it is ever fetched from the store Apr 29 18:34:28 <mmarkus1> pferraro_afk: yes, thanks! {code} So, for Paul's use case, expiration notifications for entries in memory is really what's needed. Give that as discussed in this JIRA, expiration of entries in cache store would require some careful thought, this part will be deferred until the CacheStore interfaces are redesigned in the next big major API review.	1	1	https://issues.jboss.org/browse/ISPN-694

#	Requirement	Questions	Answered	Successfully Answered	URL
8176	ISPN-675	There's a problem with this. Our rhq-plugin.xml is automatically generated as per http://anonsvn.jboss.org/repos/infinispan/trunk/tools/src/main/java/org/infinispan/tools/jopr/RhqPluginXmlGenerator.java The problem is to do with org.rhq.helpers.pluginGen component used to generate the XML file. It does not support generating XML like this: <runs-inside> <parent-resource-type name="JMX Server" plugin="JMX"/> <parent-resource-type name="JBossAS Server" plugin="JBossAS"/> <parent-resource-type name="JBossAS Server" plugin="JBossAS5"/> </runs-inside> Ian, could you guys improve the XML plugin generator so that it can generate these tags? We currently use: <version.rhq.plugingen>1.4.0.B01</version.rhq.plugingen> Finally, note that the RHQ version currently used by the JOPR plugin is: <version.rhq>1.2.0.GA</version.rhq> Ian, would working against this version cause any problems?	1	1	https://issues.jboss.org/browse/ISPN-675
8177	ISPN-675	Galder, Once you release v4.2.0.Final of the plugin, would you make sure to publish it to the JBoss Maven repo? That will allow me to update the EmbJopr trunk build to include the plugin.	1	1	https://issues.jboss.org/browse/ISPN-675
8178	ISPN-658	re: problem 1 Stop C1 on N1. Guess we would also need a new notification event, something like TopologyChanged perhaps? Currently we have ViewChanged and MergeView which extends ViewChanged. Perhaps ViewChanged should extends TopologyChange?	1	1	https://issues.jboss.org/browse/ISPN-658
8179	ISPN-658	currently support asymmetric clusters, such as: NodeA [c1, c2] NodeB [c1, c2, c3] NodeC [c2, c3] where each cache (cX) is created on a single CacheManager - which in turn starts a single JGroups channel and is represented in JGroups as a single member. Note that this only pertains to cache modes DIST and REPL. Other cache modes such as LOCAL and INVALIDATION are unaffected, the former for obvious reasons and the latter since invalidation doesn't perform any logic at startup. Invalidation messages are simply discarded by the receiving transport if the cache in question does not exist on the receiving node. So the issue really only has to do with joining and leaving. Normal runtime operation for both REPL and DIST will work fine (simply discard RPCs targeted to caches that don't exist). There are 3 specific areas which are challenging: 1) REPL uses a coordinator to determine where to pull state from. 2) DIST relies on a coordinator to coordinate a JOIN process, since JOINs are serialised (concurrent JOINs not currently supported) 3) DIST relies on view changes to determine whether nodes have joined or left To get around these, and to have minimal impact on the existing codebase, here is what I think we can do (very early designs, so may well be invalid!) * Create a new event: VirtualViewChange. A VVC is specific to a named cache. * It may be delivered by the CacheManager, which upon receiving a real ViewChange, would deliver a VVC to named caches to whom this may be relevant. * Note that this is only useful for leaves. Joins handled separately. * Cache.start() and Cache.stop() should emit a broadcast RPC to the cluster informing the cluster that it has started, including the cache name. * CacheManagers, on seeing remote caches start, may deliver VVCs to matching named caches. * RpcManager.getMembers() should be specific to the named cache. Transport.getMembers() returns the entire view, as per the JGroups channel, however RpcManager.getMembers() - which is named-cache specific - should prune the list from the transport to remove members that aren't relevant to the named cache. * Based on a cache-specific member list, we'd have a cache-specific coordinator as well. {quote} and {quote} In general, I think virtual views should work. However, here are some special cases we need to handle, too: #1 Views and virtual views - There needs to be a map of associations between views and virtual views, not sure what the keys will be. Maybe a BiMap... Yes. The correlation could also me maintained on ViewID (and possibly a corresponding VVID) #2 Broadcasts of CACHE_STARTED / CACHE_STOPPED - These (*instead* of view changes) are the notifications which trigger rebalancing - Essentially what Manik discussed below Yes, but these could be triggered by REAL ViewChanges. E.g., nodes crash, and existing nodes get a ViewChange. The CacheManager should then dispatch a VirtualViewChange to all relevant caches (which are affected by this ViewChange; not others) #3 Where are the consisten hash functions located ? - We probably cannot have CHs per DistributionManager anymore, but per cache (?) The CHs are not based on views, but on virtual views DistributionManager is per-cache anyway so the ConsistentHash can still be owned by the DistributionManager. #4 Handling of left members - When we get a view change excluding node P, everybody emits *local* notifications CACHE_STOPPED for all caches associated with P. This is where we need the hashmap of #1 Yes. See my comment on #2 above. #4.1 - As an alternative to #4, we could have only the *coordinator* broadcast CACHE_STOPPED events for P. This might eliminate the need for #5... Not really, since the coordinator may not know of certain cache instances. I think the CACHE_STOPPED notification should be delivered locally, once the CacheManager detects a ViewChange. #5 Handling of new members - When a new member starts, will it need the state (of #1) copied to it ? If we do #4, then probably yes. Not nice though... This becomes tricky. We have a few options (none of which are particularly palatable ...) A. On Cache.start(), a cache broadcasts CACHE_START and in *response*, gets the Addresses of nodes which have the same named cache running (or null if not). And use this + a real View to build a VirtualView. B. On Cache.start(), broadcast CACHE_START and	1	1	https://issues.jboss.org/browse/ISPN-658
8180	ISPN-658	Does this mean that programmatically creation of cache won't work, because the new cache can be created from any node? If this is true, I guess this is a show stopper for us, as we rely on dynamically creating many distributed caches at runtime. Predefined named cache has little value for us. Is there any work around for programmatically creation of caches (in jboss AS 6)?	1	1	https://issues.jboss.org/browse/ISPN-658
8181	ISPN-658	Agree with yangju as we are very much keen on programmatically creating the cache definitions. Infact, we dont keep anything on the Infinispan configuration files and all the named caches would be created during the startup of the cache also at the same on on all the nodes N1, N2 and N3. I guess the startup is working for us. Now coming to the runtime changes to the configurations on the different nodes, we are planning to do the same programmatic config changes simulataneously on all the nodes to avoid any assymmetry. For e.g., assume N1, N2, N3 initially had C1, C2 and during runtime C3 is added simulataneouly (almost) on all the nodes. Whats you opinion on this approach. But there can be some small differnce in the timing of this parallel config change. Does this cause any impact and resulting in the same issues as mentioned above in the thread ?	0	0	https://issues.jboss.org/browse/ISPN-658
8182	ISPN-639	I tested with Infinispan 4.2.0.ALPHA2 and am getting: 2010-09-24 09:16:58,563 ERROR [org.infinispan.remoting.rpc.RpcManagerImpl] (RMI TCP Connection(4)-127.0.0.2) unexpected error while replicating: java.lang.NullPointerException at org.infinispan.atomic.AtomicHashMapDelta.merge(AtomicHashMapDelta.java:58) [:4.2.0-SNAPSHOT] at org.infinispan.commands.write.PutKeyValueCommand.perform(PutKeyValueCommand.java:94) [:4.2.0-SNAPSHOT] at org.infinispan.interceptors.CallInterceptor.handleDefault(CallInterceptor.java:71) [:4.2.0-SNAPSHOT] . . . exception call stack is at http://pastebin.com/ftfm99vW Should this be a new jira?	0	0	https://issues.jboss.org/browse/ISPN-639
8183	ISPN-636	Hmmmm, how different would that be having a separate je.properties file? In both cases you'd need an external file. I think, and I might be mistaken here, that what Supin was looking for was to be able to define all je properties via infinispan.xml file cos having a separate file was impractical, i.e. <loader class="org.infinispan.loaders.bdbje.BdbjeCacheStore" fetchPersistentState="true" ignoreModifications="true" purgeOnStartup="true"> <properties> <property name="je.log.useODSYN" value="true"/> </properties> </loader>	1	0	https://issues.jboss.org/browse/ISPN-636
8184	ISPN-627	This issue has been migrating from 4.0.0 Final till 5.1.0 Final. (Carried over many final releases - 4.0.0, 4.1.0, 4.2.0, 4.2.1, 4.2.2, 5.0.0, 5.0.1, 5.0.2 and now 5.1.0) Is the impact too huge? Its going to be a very useful enhancement for Infinispan.	0	0	https://issues.jboss.org/browse/ISPN-627
8185	ISPN-626	Perhaps start a discussion on the dev mail list?	1	1	https://issues.jboss.org/browse/ISPN-626

#	Requirement	Questions	Answered	Successfully Answered	URL
8186	ISPN-608	Manik, What is the status of this one? Just wondering since I did not see the 07/Sep/10 comment and thought I'd get the changes I submitted in the 4.2.0.ALPHA1 release. Since I did not see those changes there, I came to this JIRA and saw your comment. Are changes planned for a future release number now?	1	1	https://issues.jboss.org/browse/ISPN-608
8187	ISPN-608	Or, flush() could return an int, being the number of entries flushed? Re: the out-of-order, could you pls create a separate JIRA and detail the issue there?	1	1	https://issues.jboss.org/browse/ISPN-608
8188	ISPN-573	I'm a bit puzzled about this. I always thought that async is great but only usable in some situations, depending on what the application is doing and what kind of consistency is required. So if async is an option, you should use it for both local and wan replication, but if it's not an option, how could this be used in wan only? Could anybody explain a use case in which this might make sense? One case occurs to me, which is having a local cluster on which changes are performed, and a remote cluster which is just a readonly view to the same data - but I assume this use case could be solved by setting up an async store and storing on the remote cluster?	0	0	https://issues.jboss.org/browse/ISPN-573
8189	ISPN-556	Is this for the batching use case?	1	1	https://issues.jboss.org/browse/ISPN-556
8190	ISPN-509	Re: first option "Enable listeners to be set via configuration XML." - can't we add this to the general cache configuration? I.e. not related to hotrod.	1	0	https://issues.jboss.org/browse/ISPN-509
8191	ISPN-509	Hmmmm, is this really needed for 5.1? There is a valid workaround available and if anything, I'd tackle it once Pete has changed the XML stuff (ISPN-1065), otherwise it'd be pointless.	1	1	https://issues.jboss.org/browse/ISPN-509
8192	ISPN-509	Yes this would need to happen after ISPN-1065. Isn't the workaround a bit involved? For folks using EDG, asking them to write their own startup code within JBoss AS 7 internals is non-trivial.	1	0	https://issues.jboss.org/browse/ISPN-509
8193	ISPN-508	Some notes: Protobuf requires a model to be defined in advance and classes to be generated, so this means that from a user perspective, whatever key/value that is passed to the Hot Rod client, it'd need to be casted to MessageLite to be able to call toByteArray method (http://code.google.com/apis/protocolbuffers/docs/reference/java/com/google/protobuf/MessageLite.html#toByteArray()). In a way this reminds me a bit of what users had to with their user classes to put them into PojoCache where they had to pre-compile and aspectize their classes so that they could be woven. This could potentially generate a lot of questions in forums based on my experience of users having to do these type of things themselves. Another thing that worries me about Protobufs would be how to do the reading part. According to http://code.google.com/apis/protocolbuffers/docs/javatutorial.html, the reading requires calling parseFrom(byte[]) static method on the object class. However, this parseFrom does not appear to belong to any class. The only thing I've found is http://code.google.com/apis/protocolbuffers/docs/reference/java/com/google/protobuf/DynamicMessage.html which can parseFrom a arbitrary type given the descriptor. How would the unmarshaller know either the descriptor or class to use to return? I'll look at Avro next...	0	0	https://issues.jboss.org/browse/ISPN-508
8194	ISPN-508	I've been playing around with Avro. Reflection based marshalling, which is the one that avoids clients having to precompile their classes, still requires some kind of sending around of the class name/type before hand in order to deserialize read data which would not be portable. Specific marshalling is the same that Protobuf and Thrift do. However, it seems to rely on a static singleton class that maintains a class cache in a CHM which would cause leaks on app server environments. Although the most interesting thing with Avro would be to allow clients to parse their own byte[] using reflection based marshalling, I seriously doubt this is portable in any way. IOW, how is a python client gonna know how to read a Java object that has been marshalled based on reflection where there's no schema? I'm attaching a sample project that uses reflection based marshalling but it's use is limited.	0	0	https://issues.jboss.org/browse/ISPN-508
8195	ISPN-506	Will we ever get this ? :)	1	0	https://issues.jboss.org/browse/ISPN-506
8196	ISPN-493	Is this expected to be fixed for 4.2.x?	0	0	https://issues.jboss.org/browse/ISPN-493

#	Requirement	Questions	Answered	Successfully Answered	URL
8197	ISPN-488	SEVERE: Could not connect to server: /127.0.0.1:11311 java.net.ConnectException: Connection refused at sun.nio.ch.Net.connect(Native Method) at sun.nio.ch.SocketChannelImpl.connect(SocketChannelImpl.java:507) at java.nio.channels.SocketChannel.open(SocketChannel.java:146) at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransport.<init>(TcpTransport.java:34) at org.infinispan.client.hotrod.impl.transport.tcp.TransportObjectFactory.makeObject(TransportObjectFactory.java:31) at org.apache.commons.pool.impl.GenericKeyedObjectPool.borrowObject(GenericKeyedObjectPool.java:1179) at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransportFactory.borrowTransportFromPool(TcpTransportFactory.java:164) at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransportFactory.getTransport(TcpTransportFactory.java:92) at org.infinispan.client.hotrod.impl.protocol.HotRodOperationsImpl.get(HotRodOperationsImpl.java:37) at org.infinispan.client.hotrod.impl.RemoteCacheImpl.get(RemoteCacheImpl.java:276) at replication.GetOnReadLoop$.main(GetOnReadLoop.scala:24) at replication.GetOnReadLoop.main(GetOnReadLoop.scala) Jun 9, 2010 1:32:29 PM org.infinispan.client.hotrod.impl.transport.tcp.TcpTransportFactory borrowTransportFromPool SEVERE: Could not fetch transport HotRodServerException{messageId=0, errorStatusCode=0} org.infinispan.client.hotrod.exceptions.TransportException: Could not connect to server: /127.0.0.1:11311 at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransport.<init>(TcpTransport.java:40) at org.infinispan.client.hotrod.impl.transport.tcp.TransportObjectFactory.makeObject(TransportObjectFactory.java:31) at org.apache.commons.pool.impl.GenericKeyedObjectPool.borrowObject(GenericKeyedObjectPool.java:1179) at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransportFactory.borrowTransportFromPool(TcpTransportFactory.java:164) at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransportFactory.getTransport(TcpTransportFactory.java:92) at org.infinispan.client.hotrod.impl.protocol.HotRodOperationsImpl.get(HotRodOperationsImpl.java:37) at org.infinispan.client.hotrod.impl.RemoteCacheImpl.get(RemoteCacheImpl.java:276) at replication.GetOnReadLoop$.main(GetOnReadLoop.scala:24) at replication.GetOnReadLoop.main(GetOnReadLoop.scala) Caused by: java.net.ConnectException: Connection refused at sun.nio.ch.Net.connect(Native Method) at sun.nio.ch.SocketChannelImpl.connect(SocketChannelImpl.java:507) at java.nio.channels.SocketChannel.open(SocketChannel.java:146) at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransport.<init>(TcpTransport.java:34) ... 8 more Exception in thread "main" HotRodServerException{messageId=0, errorStatusCode=0} org.infinispan.client.hotrod.exceptions.TransportException: Could not fetch transport at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransportFactory.borrowTransportFromPool(TcpTransportFactory.java:168) at org.infinispan.client.hotrod.impl.transport.tcp.TcpTransportFactory.getTransport(TcpTransportFactory.java:92) at org.infinispan.client.hotrod.impl.protocol.HotRodOperationsImpl.get(HotRodOperationsImpl.java:37) at org.infinispan.client.hotrod.impl.RemoteCacheImpl.get(RemoteCacheImpl.java:276) at replication.GetOnReadLoop$.main(GetOnReadLoop.scala:24) at	1	1	https://issues.jboss.org/browse/ISPN-488
8198	ISPN-465	I have published first version of prototype of Tomcat Session Manager using infinispan to distribute session state across multiple nodes. See https://github.com/zvrablik/tomcatInfinispanSessionManager Would anybody have time to discuss further work of that prototype? I am going to contact Tomcat developers too. Regards, Zdenek	1	1	https://issues.jboss.org/browse/ISPN-465
8199	ISPN-456	Which specific dependencies are you referring to? The Lucene modules seem pretty dependency-free... ?	0	0	https://issues.jboss.org/browse/ISPN-456
8200	ISPN-449	[~galderz] the title of this is a bit confusing. Is this intended as an extension of ISPN-78 over hotrod? If so mind updating the title?	1	1	https://issues.jboss.org/browse/ISPN-449
8201	ISPN-405	Just committed the first spike of this. Can I close this general issue now?	0	0	https://issues.jboss.org/browse/ISPN-405
8202	ISPN-390	Does it? Awesome!! I'll get going with this asap.	1	1	https://issues.jboss.org/browse/ISPN-390
8203	ISPN-390	Manik, how did you verify that the issue was fixed? I've tried: @Marshallable(externalizer = classOf[CacheValue.Externalizer], id = Ids.SERVER_CACHE_VALUE) class CacheValue(val data: Array[Byte], val version: Long) { And IntelliJ complains with: error: annotation argument needs to be a constant; found: Ids.SERVER_CACHE_VALUE @Marshallable(externalizer = classOf[CacheValue.Externalizer], id = Ids.SERVER_CACHE_VALUE) I've then run maven and it compiles fine. I've checked IntelliJ and is indeed compiling with Scala 2.8.1...	1	1	https://issues.jboss.org/browse/ISPN-390
8204	ISPN-375	When this will be released?	0	0	https://issues.jboss.org/browse/ISPN-375
8205	ISPN-360	hi mmarkus_ <galderz> wanted to clarify couple of things in http://community.jboss.org/message/528762 <mmarkus_> hi galderz <mmarkus_> sure <galderz> preloading is part of the startup of the cache loader, so when cache is started, the preloaded has finished, correct? <mmarkus_> yes <galderz> indeed: @Start(priority = 56) <galderz> but the state transfer is a diff thing <galderz> the cache is started already and it's having state applied to it? <mmarkus_> hmm, not sure which one comes first, i.e. state transfer or preload <galderz> http://community.jboss.org/wiki/JBossCacheNon-BlockingStateTransfer <galderz> assuming a node joining D, it's only marked as started when the state transfer has completed <galderz> not sure either about the order <mmarkus_> yes, a node is not started until state transfer is finished <galderz> @Start(priority = 55) -> @Start(priority = 56) <galderz> state trasnfer is first, the preload <galderz> then preload <mmarkus_> right, makes sense <galderz> now, what's the point of https://jira.jboss.org/jira/browse/ISPN-360 ? <mmarkus_> because you update the store and then, if you want, load it into memory <galderz> i can see maybe the use of a rehash started/completed notifications in running nodes affected by the start of a new node <mmarkus_> and that brings another thing to the forum post <galderz> and I think my last comment is what you're aiming at in 360? <galderz> the capability of other nodes to find out when a new node has joined, which would mean that any rehashing that might have happened in the running node is completed? <mmarkus_> I would like to be notified AFTER a node is started <mmarkus_> even when a remote node joined the cluster <galderz> i see ur point <mmarkus_> yes, exactly <galderz> but just bear in mind, in nodes that are affected by rehashing caused by the start of this node, that notification needs to happen after rehashing has finished, correct? <mmarkus_> this guy wants to start N nodes, and then tell the application that the cluster is started <galderz> assuming the app is running in each of these N nodes? <mmarkus_> yes. Basically the node that receives state will know that it finished receiving state, and should notify all other nodes that NodeJoined <mmarkus_> yes <mmarkus_> you've just raised another problem I wasn't aware of <mmarkus_> so if we have both preload and fetchInMemory state, then local cache store state before starting will be erased <galderz> yeah, i guess that assuming one of those nodes X, it should receive N-1 NodeJoined, but a NodeJoined from X-1, meaning that X-1 is next to X in the hash loop, X would receive NodeJoined when X-1 has started and X has done any reshahing <galderz> mmarkus_, not sure if it will erase it <galderz> it will override it <galderz> i'd assume that if you receive k=akey,v=1 via in memory and u had preloaded k=akey,v=2, v=1 would prevail <mmarkus_> yes, at the time X receives NodeJoioned from X-1, X(and all other nodes) finished rehashing	1	1	https://issues.jboss.org/browse/ISPN-360

#	Requirement	Questions	Answered	Successfully Answered	URL
8206	ISPN-359	Would it be possible to make this feature work with buddy replication in cooperation with mod_cluster? For example you have node1,node2,node3,node4. Session A is on node1 and node2. When node1 fails it would be best mod_cluster to hit deterministically node2 instead of the other nodes. I remember I have seen hints that Infinispan is looking into rehashing or rebalancing, but I am not sure how smart is the algorithm and is it working cooperatively with mod_cluster?	1	1	https://issues.jboss.org/browse/ISPN-359
8207	ISPN-359	By Buddy Replication I meant the DIST mode. I thought it is still called that way. Anyway, I will see what Brian says about mod_cluster. Other similar products use configurable affinity key independently in the cache and in the load balancers. In Mobicents we use pluggable algorithms for example and we will probably try to teach it to cooperate with this feature deterministically to cover some telco cases (some telco protocols dont have cookies or redirects to hunt down the right node). But what is still unclear on this topic is the rehashing/rebalancing. If a node fails is everything completely re-hashed to a new location (along with the huge network overhead of moving the chunks around over the network) or is there some effort to minimize the state transfer?	0	0	https://issues.jboss.org/browse/ISPN-359
8208	ISPN-328	Can we control that writes to DB is done only through one of instances and that instance might be chosen by cluster elements.	0	0	https://issues.jboss.org/browse/ISPN-328
8209	ISPN-318	Is this really necessary?	0	0	https://issues.jboss.org/browse/ISPN-318
8210	ISPN-317	Is ISPN-1539 a duplicate of this issue?	1	1	https://issues.jboss.org/browse/ISPN-317
8211	ISPN-313	[~belaban] would the new jgroups protocol configuration extension (graphic?) cover this as well?	1	0	https://issues.jboss.org/browse/ISPN-313
8212	ISPN-312	How would you handle the following case? Cache A uses a custom key class where the group is dynamically generated. Cache B's key is an unmodifiable String that needs to be grouped with the dynamic value of a key from Cache A. (Cache B's key can't be encapsulated in a custom key class since it needs to be queried from external entities that will not know the group context). The only way I can see doing this is to add a group parameter to a put() or replace() method.	1	1	https://issues.jboss.org/browse/ISPN-312
8213	ISPN-312	Erik, can you show me what you mean using a pseudo-API that you propose?	1	1	https://issues.jboss.org/browse/ISPN-312
8214	ISPN-312	I'm resolving this for 5.0 for tracking purposes. Erik can you please open a new issue for your feature request?	0	0	https://issues.jboss.org/browse/ISPN-312
8215	ISPN-303	Mazz: "Have you tried to create a "compatible group" containing your different infinspan servers? You could get some interesting aggregated metrics if you create a group containing all your servers and look at the group metrics."	0	0	https://issues.jboss.org/browse/ISPN-303
8216	ISPN-293	If I could get multiple values at once, that would be good enough for the use cases I met so far and AFAIR what Emmanuel described for OGM. Map<K,V> getAll(Key... keys) ? (assuming this simplifies the implementation)	0	0	https://issues.jboss.org/browse/ISPN-293
8217	ISPN-272	What happens right now in such a scenario? Are locks cleared? Or do we have stale locks?	1	1	https://issues.jboss.org/browse/ISPN-272
8218	ISPN-272	whilst at this also fix this todo - // TODO Couldn't affected keys be derived from modifications?	0	0	https://issues.jboss.org/browse/ISPN-272
8219	ISPN-256	Hmm; the problem with this is that it then makes infinispan a message bus. Completely not what it is designed to do, but I can see how it could aid other use cases. Anyway, what you would need is like you said an ExtensionCommand, which could contain an arbitrary payload of Serializable[], a target cache, and a mechanism to register a listener for such ExtensionCommands. And also a mechanism of dispatching such commands: dispatchExtensionCommand(ExtensionCommand ec, List<Address> recipients, boolean sync) ?	1	1	https://issues.jboss.org/browse/ISPN-256
8220	ISPN-256	I suppose this might be released before Final for community testing? If so, will it be included in Beta1?	1	1	https://issues.jboss.org/browse/ISPN-256
8221	ISPN-233	Galder: can this be closed?	1	0	https://issues.jboss.org/browse/ISPN-233
8222	ISPN-232	seems definitely useful, but if I have a traditional key, how will I be able to get() the value? should I keep my own map of originalKey->currentInfinispanKey or will the service be able to care for this? I'd like to be able to get() as usually and make use of this feature transparently. FYI this can potentially give an interesting performance boost in the way we manage the Lucene indexes.	0	0	https://issues.jboss.org/browse/ISPN-232
8223	ISPN-221	hello all, I'm going to need some more information, I have no context nor I know what the goal/purpose is. what is the meaning of "natural language" in this context? what is the object to be searched for? I don't consider JPA-QL to be natural language. Which module is this affecting? infinispan-query? why not Hibernate Search?	0	0	https://issues.jboss.org/browse/ISPN-221
8224	ISPN-193	Any suggestions for the names of the new XML elements? Where do we want to put them in schema? As a child of default configuration or any namedCache?	1	1	https://issues.jboss.org/browse/ISPN-193
8225	ISPN-162	Exclusion added, the following still needs doing: >> +1 to ensuring it does not get packaged with the test jar, but it should >> still be included in the distro (in the conf directory) as an example of >> how folks can enable logging. > > If we want to include it in the distro, it definitely must have a diff name to log4j.xml, otherwise it'll get picked up and users won't be able to get override it in any way. > > Currently we're not including a log4j.xml in the main distro at all. Hmm. Ok, then when it is copied to the conf directory of the distro it should be renamed to log4j-sample.xml, and there should be a README-logging.txt to explain how to enable this file, either by renaming it to log4j.xml or by using -Dlog4j.configuration. WDYT?	0	0	https://issues.jboss.org/browse/ISPN-162
8226	ISPN-143	Hi Manik - do we have any updates on this? Have any proposals been made regarding an infinispan security model. Like Aleksander, I work for an organisation where sensitive data will be held in the cache and securing access to it is a requirement. I'm not only talking about encrypted communication but perhaps more importantly role based api access control.	1	0	https://issues.jboss.org/browse/ISPN-143
8227	ISPN-139	I'm eager for news about this, did you publish some numbers or any thoughts to share?	1	0	https://issues.jboss.org/browse/ISPN-139

#	Requirement	Questions	Answered	Successfully Answered	URL
8228	ISPN-134	manik: This thread http://lists.jboss.org/pipermail/infinispan-dev/2009-July/000883.html also has info [1:36pm] manik: but basically [1:36pm] mmarkus_: also re:coherence migration, I think that will take about 3 days by itself, as I really have to understand coherence first [1:36pm] manik: Ah ok [1:36pm] manik: (the jopr stuff should take a few hrs IMO) [1:36pm] mmarkus_: I would do that together with a feature comparison with coherence (ISPN-46) [1:37pm] mmarkus_: kewl, I take your word for that [1:37pm] manik: well basically [1:37pm] manik: Heiko has already written the jopr plugin [1:37pm] manik: the only problem is that all of the JMX names and types are hard-coded in XML [1:37pm] manik: E.g., the jopr plugin refers to an XML file [1:37pm] manik: which tells it where all of the JMX attribs are [1:37pm] manik: now this sucks [1:37pm] mmarkus_: i c [1:37pm] manik: sunce every time we need to add a jmx attrib [1:38pm] manik: we need to update this XML as well [1:38pm] mmarkus_: where is the hand written xml now? [1:38pm] manik: so I think we can generate this XML [1:38pm] manik: http://fisheye.jboss.org/browse/Infinispan/trunk/jopr-plugin/src/main/resources/META-INF [1:38pm] manik: in trunk, under the jopr-plugin module [1:39pm] mmarkus_: i c [1:39pm] manik: basically most of this info is already available in our @MBean annotations [1:39pm] mmarkus_: and the code that generates it should be in jopr-plugin module as well, i think? [1:39pm] manik: well, the way I see it [1:40pm] manik: this should be generated an build-time [1:40pm] manik: and not at runtime [1:40pm] manik: I.e., when we build a release, the build process generates this XML and packages it with the jopr-plugin jar [1:40pm] manik: so we already have 3 psot-comple processes running in the 'tools' module [1:41pm] manik: 1) generate the config docs - http://infinispan.sourceforge.net/4.0/apidocs/config.html [1:41pm] manik: 2) generate the jmx docs - http://infinispan.sourceforge.net/4.0/apidocs/jmxComponents.html [1:41pm] manik: and 3) generate the XSD schema from the annotations on Configuration [1:41pm] mmarkus_: ah i c [1:41pm] manik: so I think a 4th step can be added here, to generate the jopr XML [1:42pm] mmarkus_: and how do these post-compile tools get called? [1:42pm] manik: maven has hooks for this [1:43pm] manik: e.g. [1:43pm] mmarkus_: I'm looking at pom.xml for tools [1:43pm] manik: http://fisheye.jboss.org/browse/Infinispan/trunk/core/pom.xml?r=728 [1:43pm] manik: line 88 [1:43pm] manik: Tools only builds the code that will do the generation [1:43pm] manik: tools doesnt have the build hooks [1:43pm] manik: the build hooks will be in whichever module that needs it [1:43pm] manik: e.g., core needs the XSD to be generated [1:44pm] manik: so it calls the XSD generator in core/pom.xml : 88 [1:44pm] manik: The Maven <process-classes> phase happens after <compile> [1:45pm] manik: So you could easily write something similar to generate the JOPR XML [1:45pm] manik: now here is one thing: [1:45pm] manik: the JOPR XML needs more info than we have in the @MBean annotations [1:45pm] mmarkus_: i c [1:45pm] manik: and hence my discussion in the email thread above [1:45pm] manik: do we enhance the annotations to include this extra info? Or (for now) can we just make assumptions? [1:46pm] manik: Ideally I would like to add the extra info in the annotations [1:46pm] mmarkus_: i'll have to add something similar to the pom.xml in jopr-module [1:46pm] manik: Yup [1:46pm] mmarkus_: I c [1:47pm] manik: well, have a look and a think and see what you feel - I think it should be pretty straightforward stuff. It is just reformatting metadata after all... [1:47pm] manik: metadata inannotation -> metadata in XML [1:47pm] mmarkus_: yep, much clearer to me now [1:47pm] mmarkus_: I'll update the jira with the conversation [1:47pm] manik: ok great [1:48pm] manik: also update with the link to the discssion thread on infinispan-dev [1:48pm] manik: since that has a lot of info as well [1:48pm] manik: To test the jopr plugin, you can follow these steps [1:48pm] manik: http://www.jboss.org/community/wiki/MonitoringInfinispanwithJopr [1:48pm] mmarkus_: mailing list, you mean? [1:48pm] manik: yes [1:49pm] mmarkus_: I'll paste the conversation, that includes the link [1:50pm] manik: ok	1	1	https://issues.jboss.org/browse/ISPN-134
8229	ISPN-134	<galderz> i do need a release that contains the updates to the plugin generator Oct 01 15:05:27 <galderz> so pilhuhn, we're discussing whether to add these annotations to ispn, or we should just work on the helper module on rhq? Oct 01 15:05:47 <galderz> i already need a release of plugin generator with the changes i'm implementing Oct 01 15:06:09 <galderz> so, how doable is to get a release of the annotation module at the same time? Oct 01 15:06:33 <galderz> if u could set up the module, i could start working on the annotations as I need them on ispn Oct 01 15:07:03 <galderz> if not doable, we'd have to add the annotations to ispn, but if we can start on the rhq module already, it'd be better Oct 01 15:07:59 <galderz> thoughts pilhuhn ? Oct 01 15:09:08 <pilhuhn> Howdy Oct 01 15:09:09 <pilhuhn> I think I have some uncommited code in it - but it is not in any official release anyway Oct 01 15:09:43 <galderz> what uncommited code? Oct 01 15:09:49 <galderz> annotations? Oct 01 15:10:04 <pilhuhn> I'd love to do it in RHQ, so that it gets re-usable Oct 01 15:10:15 <galderz> indeed Oct 01 15:10:17 <galderz> so? Oct 01 15:10:40 <pilhuhn> And yes, it makes sense to put them in a module within helpers, but outside the plugin generator, so that the jar file will stay small Oct 01 15:10:54 <pilhuhn> no, to discovery.ftl Oct 01 15:10:54 <galderz> +1 Oct 01 15:11:08 <galderz> the plugin generator is needed at build time by the jopr plugin Oct 01 15:11:18 <galderz> the annotations are needed by core and potentially other modules Oct 01 15:12:44 * joseph42 (n=joseph42@redhat/jboss/joseph42) has joined #infinispan Oct 01 15:12:49 <pilhuhn> yes Oct 01 15:13:04 <galderz> joseph42, summary: Oct 01 15:13:19 <galderz> i'd like to have a rhq/jopr release at some point that includes Oct 01 15:13:36 <galderz> a tiny module with jopr annotations Oct 01 15:13:46 <galderz> and an update to rhq plugin generator Oct 01 15:14:08 <galderz> these jopr annotations annotate metrics...etc in the infinispan project Oct 01 15:15:04 <pilhuhn> joseph42: Infinispan exports a lot of stuff via jmx - they use annotations, but that data is not enough to build the plugin descriptor from it. So the idea is to have a RHQ module with annotations, that projects can use - with some helper class to parse etc - to build fragments of the deployment descriptor and probably also of the component class facets Oct 01 15:16:24 <pilhuhn> So the idea is to have the annotations @RhqMetric, @RhqOperation in modules7helpers/annotations to build a very small jar and then have e.g. the plugin generator parse the classes and use the information contained Oct 01 15:16:53 <joseph42> so...that'll attempt to automate the process of building a plugin? Oct 01 15:17:19 <pilhuhn> help to automate Oct 01 15:17:26 <galderz> it will automatte the descriptor generation Oct 01 15:17:39 <pilhuhn> joseph42: e.g. in Infinispan a RHQ-resource is done via a bunch of Mbeans Oct 01 15:17:55 <pilhuhn> so the generic jmx-plugin does not apply Oct 01 15:17:57 <joseph42> but instead of a one-off solution, you want something formal that comes out of rhq/jopr Oct 01 15:18:06 <joseph42> as well as a generator youcan run over your codeto gen the plugin.xml Oct 01 15:18:27 <joseph42> did i capture that correctly? Oct 01 15:18:29 <pilhuhn> joseph42: They would do it in ispn, but I talked them into having that in RHQ, as others may want that stuff too Oct 01 15:18:37 <galderz> it's the plugin.xml that we're mostly interested in Oct 01 15:19:02 <joseph42> right, understood Oct 01 15:19:23 <joseph42> so if that module resides in ispn, you shouldn't have any deps on rhq/jopr,right? Oct 01 15:19:37 <joseph42> (assuming the xsd for plugin.xml is relatively static) Oct 01 15:19:42 <galderz> i do have a dep Oct 01 15:19:47 <galderz> the plugin generator module Oct 01 15:20:12 <galderz> i'm improving it to be able to pass in X number of metrics/ops Oct 01 15:20:24 <pilhuhn> joseph42: in an "ideal world" that module would live in RHQ for other projects to just use it Oct 01 15:20:36 <joseph42> hmm...so you want extenions to heiko's plugin generator to understand annotation-based enhancements in the ispn? Oct 01 15:21:12 <galderz> i need modifying the plugin generator to generate metric/ops dynamically Oct 01 15:21:35 <joseph42> galderz: completely understood Oct 01 15:21:46 <galderz> then, our own jopr plugin calls this, having checked the annotations, to generate the xml Oct 01 15:22:07 <galderz>	1	1	https://issues.jboss.org/browse/ISPN-134
8230	ISPN-133	----- "Philippe Van Dyck" <pvdyck at gmail dot com> wrote: > On Mon, Apr 19, 2010 at 12:36 PM, <galder at jboss dot org> wrote: > > > > Hi Philippe, > > > > So, to sum up, trying to fix > ISPN-133 won't solve your problem, > correct? Did you test this out by any chance? > > > > Yep, "Basically InputStream#available() is utterly broken.". > GzipIbputStream's implementation is : > public int available() throws IOException { > ensureOpen(); > if (reachEOF) { > return 0; > } else { > return 1; > } > } > > You definitely cannot use the return value to set your buffer size. Interesting. It seems like it's not only Apache's HTTP client library that did not care much about available() implementations, but also JDK classes as well. I'd rather read on 1k batches than 1 byte batches. See things like this, I'm even less inclined to urgently look into ISPN-133.	1	1	https://issues.jboss.org/browse/ISPN-133

#	Requirement	Questions	Answered	Successfully Answered	URL
8231	ISPN-105	Mircea - isn't this possible right now? If so, could you write a short wiki page on using/configuring the Jdbc Cache Store?	1	1	https://issues.jboss.org/browse/ISPN-105
8232	ISPN-98	So, we know that hudson show failures when using MPING but having run the testsuite in cluster10 (the slave used by Hudson) several times, the failures do not always happen. When they happen, you see stuff like this in the logs in spite of java.net.preferIpv4Stack = true: 2009-07-16 03:55:28,030 WARN [org.jgroups.protocols.MPING] (pool-1-thread-5) could not bind to /228.10.10.16 (IPv4 address); make sure your mcast_addr is of the same type as the IP stack (IPv4 or IPv6). Will ignore mcast_addr, but this may lead to cross talking (see http://www.jboss.com/wiki/Edit.jsp?page=CrossTalking for details). Exception was: java.net.BindException: Cannot assign requested address I'm trying to get TRACE on org.jgroups when it fails but no luck so far. Having compared good logs against bad logs generated in cluster10, I see the following appearing *only* when the tests fail: 2009-07-16 03:55:28,153 WARN [org.jgroups.protocols.TCP] (pool-1-thread-4) failed to join /224.0.75.75:7500 on eth2: java.net.SocketException: No such device Not sure if this is somehow cause or consequence though.	1	1	https://issues.jboss.org/browse/ISPN-98
8233	ISPN-95	@Alejandro is this complete now? Or is there more stuff you want to add to this tutorial?	1	1	https://issues.jboss.org/browse/ISPN-95
8234	ISPN-78	Would it be possible to make this a configuration option? For some applications it may make more sense to just save the object -- for example, if the object is not expected to be accessed much. T	0	0	https://issues.jboss.org/browse/ISPN-78
8235	ISPN-78	Wouldn't a non streaming API somehow defeat this feature's sole purpose? It is my understanding that it is driven by the desire to store objects the size of which exceeds that of every single JVM in the cluster, correct? If that is indeed the case, it would be *impossible* to store such an object without first fragmenting it into pieces each of which fits into a single node's heap. So we need a streaming API to read a large object from some place - think file - outside the current JVM's address space in order not to blow up that JVM's heap. It is out of the question that this approach considerably complicates matters, but I don't see a way around it. Ah, I see. Stupid me. Streaming to disk means writing a large object directly through to the cache store, without actually caching it in memory. Might make sense in some cases, though it could be argued if in this case Infinispan would really act as a "cache". We would still need a streaming API, though. Anyway, IMHO we should first concentrate on implementing the "fragmentation approach" and then check back whether streaming to disk might make sense. One step at a time.	1	1	https://issues.jboss.org/browse/ISPN-78
8236	ISPN-78	Hi Gabriel, it is my understanding that Reed-Solomon codes aide in detecting and correcting data corruption. In our scenario this equates to repairing chunks that have somehow been corrupted. I have thought about later implementing some means of doing so myself, so thanks for the hint. What I fail to see, though, is how this would help to address the problem of availability in the face of node failures. If two nodes storing the same chunk are down that chunk and thus the large object it belongs to are not available. Or would it be possible using RS codes to reconstruct a missing chunk from the information contained in the other chunks?	1	1	https://issues.jboss.org/browse/ISPN-78
8237	ISPN-75	Thanks for this. A few comments on the patch, before you check it in: 1. What is the purpose of the comparator? Exposing a comparator here implies that the entire cache is ordered - which it is not. I don't foresee this being useful until there is ordering applied to the entire cache. 2. addNewExpiry should store enough information to be able to purge on a lifespan expiry as well as an idle time expire (as per your comment in the patch). Perhaps this could be calculated to an absolute time based on lifespan + creation time or maxIdle + lastModified, whichever is sooner. 3. Why do you need to synchronize on expiryLock? From what I see, expiryTree is just another JDBM stucture and JDBM should handle concurrent access to this structure? Or have I missed something?	1	1	https://issues.jboss.org/browse/ISPN-75
8238	ISPN-58	How are we going to implement service behind this URI? Is it going to be a servlet that runs submitted Runnable/Callable?	0	0	https://issues.jboss.org/browse/ISPN-58
8239	ISPN-57	Some discussions on the subject: http://lists.jboss.org/pipermail/infinispan-dev/2009-November/001953.html and an email from Vladimir: "I think that we should give up, for the time being, trying to deploy Infinispan in GAE/j. After I have encountered the below mentioned blocker with our regular marshaling I have gone in direction recommended by Galder. I replaced river marshalling with java serialization. Although, java serialization does not really work, as Galder pointed, I wanted to jump over the first hurdle and see how far I can get with deployment. Didn't get very far. After I used java serialization instead of river I got this exception: java.security.AccessControlException: access denied (java.lang.RuntimePermission accessClassInPackage.sun.reflect) at java.security.AccessControlContext.checkPermission(AccessControlContext.java:323) at java.security.AccessController.checkPermission(AccessController.java:546) at java.lang.SecurityManager.checkPermission(SecurityManager.java:532) at com.google.appengine.tools.development.DevAppServerFactory$CustomSecurityManager.checkPermission(DevAppServerFactory.java:151) at java.lang.SecurityManager.checkPackageAccess(SecurityManager.java:1512) at java.lang.ClassLoader$3.run(ClassLoader.java:410) at java.security.AccessController.doPrivileged(Native Method) at java.lang.ClassLoader.checkPackageAccess(ClassLoader.java:408) at org.jboss.marshalling.reflect.SunReflectiveCreator$1.run(SunReflectiveCreator.java:40) at org.jboss.marshalling.reflect.SunReflectiveCreator$1.run(SunReflectiveCreator.java:38) at java.security.AccessController.doPrivileged(Native Method) at org.jboss.marshalling.reflect.SunReflectiveCreator.(SunReflectiveCreator.java:38) at org.infinispan.marshall.jboss.JBossMarshaller.start(JBossMarshaller.java:107) at org.infinispan.marshall.VersionAwareMarshaller.start(VersionAwareMarshaller.java:75) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at com.google.appengine.tools.development.agent.runtime.Runtime.invoke(Runtime.java:100) at org.infinispan.util.ReflectionUtil.invokeAccessibly(ReflectionUtil.java:170) at org.infinispan.factories.AbstractComponentRegistry$PrioritizedMethod.invoke(AbstractComponentRegistry.java:846) at org.infinispan.factories.AbstractComponentRegistry.internalStart(AbstractComponentRegistry.java:666) at org.infinispan.factories.AbstractComponentRegistry.start(AbstractComponentRegistry.java:572) at org.infinispan.factories.GlobalComponentRegistry.start(GlobalComponentRegistry.java:130) at org.infinispan.factories.ComponentRegistry.start(ComponentRegistry.java:134) at org.infinispan.CacheDelegate.start(CacheDelegate.java:303) at org.infinispan.manager.DefaultCacheManager.createCache(DefaultCacheManager.java:394) at org.infinispan.manager.DefaultCacheManager.getCache(DefaultCacheManager.java:358) at org.infinispan.manager.DefaultCacheManager.getCache(DefaultCacheManager.java:333) at guestbook.SignGuestbookServlet.doPost(SignGuestbookServlet.java:63) "	0	0	https://issues.jboss.org/browse/ISPN-57
8240	ISPN-49	Perhaps enlist David Lloyd's help in trying to make JBMAR faster? :)	1	1	https://issues.jboss.org/browse/ISPN-49
8241	ISPN-29	When tackling this, check what Alex has done in: http://www.jboss.org/index.html?module=bb&op=viewtopic&p=4232775#4232775	0	0	https://issues.jboss.org/browse/ISPN-29

#	Requirement	Questions	Answered	Successfully Answered	URL
8242	ISPN-29	* remember jmemcached is still pre 1.0, and is under active development, so nothing is written in stone. * I inherited the code base from another author, though it has changed radically since then. Many things about the design that I don't like are products of that. * it's a side project for me, and I do not currently have it in any production systems at my Day Job (an earlier version was used at my previous job), so it gets less love that it could. Other people are using it. It's being used a lot for integration tests. Still, it is under active development and getting better. Re #1 (locking): The choice was made to do locking further up the call chain because I had intentions at the time of potentially introducing the ability to chain or sequence cache storages together, which would require one lock for the entire operation against all storage engines. Further I had some notions of adding transactional (commit/rollback) style semantics at the protocol level. Should I do that, I felt there would be some rather explicit locking behaviours required, and it would be better off managed by Cache. But the most important reason that Cache performs the locking is that some operations (incr, decr, cas, prepend, and append) combine reads and writes into one atomic operation. So in this case there are 3 possible options: * Re-implement all these operations for every storage engine. * Add "lock/unlock" methods to the storage engine. * Lose atomicity on these operations. If you have other suggestions, I'm open to them. In any case, maybe your storage would sit at the level of Cache (which could be made into an interface or abstract), not a CacheStorage implementation. #2 (String/MCElement storage): The protocol really dictates this. Firstly, in the protocol all keys are just strings and can't be anything else. Secondly, every element can have storage flags stored along with it, as well as its timeout. This basically requires that either the value be encapsulated in something or these attributes be external to the value. It's just easier to do the former. #3 (asynchronous deletes): This is incorrect, only delayed removes use the delayed asynchronous queue. Deletes with no delay time attached are removed instantaneously. #4 (hardcoded to netty): It took me one day to switch from MINA to Netty; so while I am not sure I truly see the advantage of using raw NIO or XNIO, it certainly is possible and if you know NIO well, it wouldn't take long. That said, I'm not sure why you would do this? Do you have benchmarks to backup your performance claims? And why would a JBoss product have a problem with a Netty dependency? One of my goals has been to get the protocol decoding/encoding abstracted out enough that I could do unit testing of it (right now there's just integration tests), but it's difficult and time consuming. #5 (license): The license is Apache 2.0, there's nothing unclear about it. Also since your comment, the main site has switched now to a google code project. #6 (jgroups impl of the protocol): Looking at the jgroups implementation of the protocol... It's not complete or correct. There's more to the protocol than sets and gets. You can reinvent the wheel and write the protocols yourself, but ... why? #7 Enum comparisons: The big if/else chain isn't sexy, but switch statement in Java is braindead. In particular, GET, GETS, and a number of others all take variations of the same logic; switch would be ungainly here. You could argue (as someone else has) that this should be a polymorphic method call per operation, but I think this would be overkill. #8 slf4j: What's wrong with slf4j? I hope you're not suggesting commons-logging...	1	1	https://issues.jboss.org/browse/ISPN-29
8243	ISPN-29	Also is there sense in implementing memcached's binary protocol as well? Does anyone know what the status is with existing clients using this binary proto? http://www.elctech.com/articles/memcached-1-4-released-binary-protocol http://cloud.github.com/downloads/memcached/memcached/protocol-binary.txt	0	0	https://issues.jboss.org/browse/ISPN-29
8244	ISPN-29	Question to address: what happens when a connection breaks between client and server? What sort of keepalives, etc do we have, e.g., a ping-pong protocol?	0	0	https://issues.jboss.org/browse/ISPN-29
8245	ISPN-25	When is JGroups 2.8.0.CR2 due? Is there a JGroups JIRA we can link to?	0	0	https://issues.jboss.org/browse/ISPN-25
8246	ISPN-24	This is great news. What we're invisioning is a way to slot Infinispan between HIbernate and our app. So, Hibernate would be the Cachestore if possible and we'd use JPA-like constructs to query and persist objects. This would then enable a JPA-like distributed write-behind data grid using JPA objects. Good stuff.	1	1	https://issues.jboss.org/browse/ISPN-24
8247	ISPN-2	Indeed, thought it might related to something like that. Once you've figured out locally that the value matches and you need to delete, it's a go on the other nodes and don't need to double check. How about we mark value in some way so that this is obvious for the reader? transient?	1	1	https://issues.jboss.org/browse/ISPN-2
8248	IPAAS-325	Some general thoughts which come to mind while reading the story: * What's about error handling ? Would something like offering circuit breaking for API 2 sense (i.e. using default values when connection fails) ? * What with the response of API-3 ? Should it be recorded / logged ? Or is it up to the use to add an additional step after the call to API-3 to perform actions on this response ? * If I understand right, after the mapping from API-1 (resp) to API-2 (req), the message from API-1 (resp) should still remain in the exchange so that the call to API-3 can use mappings from API-1 (resp) and API-2 (resp) ? (which would require probably some fundamental changes in the data mapper)	1	0	https://issues.jboss.org/browse/IPAAS-325
8249	IPAAS-325	{quote}What's about error handling ? Would something like offering circuit breaking for API 2 sense (i.e. using default values when connection fails) ? {quote} I think this is an interesting use case as it starts to test the balance of what's functional vs. technical from a user perspective. Does it make more sense to have a developer add circuit breaker type functionality separately when needed or is there a citizen-level representation of that possible? {quote} What with the response of API-3 ? Should it be recorded / logged ? Or is it up to the use to add an additional step after the call to API-3 to perform actions on this response ? {quote} Good call - that's my mistake. I should note that API-1 is GET and API-2 and API-3 are POST. API-3 is putting data into another system, API-2 is pulling additional data based on what we got from API-1. I guess API-2 could do that via GET with id or via POST - let's discuss and choose the option that makes the most sense. {quote} If I understand right, after the mapping from API-1 (resp) to API-2 (req), the message from API-1 (resp) should still remain in the exchange so that the call to API-3 can use mappings from API-1 (resp) and API-2 (resp) ? (which would require probably some fundamental changes in the data mapper){quote} This is something we discussed back in the early design phases of iPaaS, so I would check in with Hiram to get the low down on the technical side. I believe the plan was to stash earlier messages as exchange properties or message headers to propagate them through the route/flow. The iPaaS UI would take care of passing the appropriate data slots to the mapper, which is designed to be able to map multiple source objects to a single target object.	0	0	https://issues.jboss.org/browse/IPAAS-325
8250	IPAAS-324	A WebHook typically also receives a payload (in addition to the WebHook's URL). The format is typically defined by the system calling the WebHook (e.g. GitHub). Should Syndesis also support evaluation for WebHook payload data which then can be used e.g. for data mapping ? Use Case: We add the webhook to a GitHub repo and receive a trigger each time a commit happens in GitHub. The data from the commit could be used in an integration mapping scenario, too. Security: How do we ensure that only authorized requests trigger the integration ? Do we generate unique random WebHook URLs for each integration ? (This would require to add / remove OpenShift `Routes` dynamically)	1	0	https://issues.jboss.org/browse/IPAAS-324

#	Requirement	Questions	Answered	Successfully Answered	URL
8251	IPAAS-323	[~kcbabo] could you please give a short example what this unique id could be in the case of Twitter ? Is it the twitter handle which is stuffed in one attribute of the salesforce object ? On which action should the unique id mapping happen ? Isn't it that this mapping would need two params, one from twitter, one from salesforce, so that this cannot be directly connected with a single action only (which is specific to a connector) ?	1	0	https://issues.jboss.org/browse/IPAAS-323
8252	IPAAS-319	Do we also want/need to display Exchange properties? How much do we want to expose to the user? For example, do we want to provide the ability to specify log name, etc?	1	1	https://issues.jboss.org/browse/IPAAS-319
8253	IPAAS-319	[~iocanel], the way the log UI design is now, it prompts the user to indicate whether or not to include the Log Level in the log. Is that desired? I expect that the Log Level is a separate and different option since it is an indication of what types of information you want to send to the log. I think the Logging UI needs two parts: *What kinds of information do you want to send to the log?* With a drop down for selecting one of ERROR, WARNING, INFO, DEBUG, TRACE. And they have to be in that order since each one includes the ones before it. Probably we will want some popup help to explain that. *What information do you want to send to the log?* And then list Message Header, Message Body, Custom Text.	1	1	https://issues.jboss.org/browse/IPAAS-319
8254	IPAAS-315	Looks good to me. Tested it today works fine. What is the plan for displaying status messages while a deployment is in progress? a popup on hover over the spinner?	1	1	https://issues.jboss.org/browse/IPAAS-315
8255	IPAAS-315	Right now the only granularity available from the API is really 'Activated' and 'Deactivated', we'd need a way of fetching any additional information we'd want to show in the UI. [~sjcox] I can't recall, do we have a hover/pop up in the design for integrations that are transitioning states?	1	1	https://issues.jboss.org/browse/IPAAS-315
8256	IPAAS-315	[~chirino] Ah cool I see that. Why don't I add a popup for now with the status message when it's in progress?	0	0	https://issues.jboss.org/browse/IPAAS-315
8257	IPAAS-306	[~kcbabo] [~chirino] can you provide more details about this. The address I assume is the hostname, eg which today is named host https://github.com/redhat-ipaas/connectors/blob/master/connectors/trade-insight-buy-connector/src/main/resources/camel-connector.json#L20 And the other parts are currently hardcoded if you see further below in that JSon file. But that API key fields. Are those some custom fields that the citizen user must configure. And I think Hiram said something about it was supposed to be a HTTP Header when calling the HTTP service. Are you going to use any of that for the demo, but regardless I assume maybe in the real world there may be some API keys the user must provide as credentials or something.	1	0	https://issues.jboss.org/browse/IPAAS-306
8258	IPAAS-299	This was done last week. Have you tested on BlueJeans to see how it looks? The gray will still be light (much darker on a "regular" screen), but shouldn't be white like before, so there is more of a contrast. I am going to do a quick pass-thru to see if I missed any sections of the site. The background color is custom and not in PatternFly, it's #E1E1E1. GitHub Issue: https://github.com/redhat-ipaas/ipaas-ui/issues/393	0	0	https://issues.jboss.org/browse/IPAAS-299
8259	IPAAS-295	What do you mean by an "API icon"? Don't you want to be able to use custom icons for APIs? Isn't it the service that's important, not that it's an API? Most things are APIs of course.	1	1	https://issues.jboss.org/browse/IPAAS-295
8260	IPAAS-295	Just to close off based on our IRC discussion. While most connectors are using an API under the covers, they offer a level of indirection - the end user is not interacting with the API directly. There will be cases where we have an API connector which just takes in a Swagger definition and allows you to invoke an API through that. I think we should use an API icon for that connector. Agree?	1	0	https://issues.jboss.org/browse/IPAAS-295
8261	IPAAS-294	Hi [~kurtstam], could you add these connectors to the deployment.json ?? They don't need to be usable, just show up in the UI.	1	1	https://issues.jboss.org/browse/IPAAS-294
8262	IPAAS-293	Kurt, could you make sure all 3 actions are in the deployments.json?	0	0	https://issues.jboss.org/browse/IPAAS-293
8263	IPAAS-284	[~kcbabo] so if you select trade insight as "start" connection, then what should trigger it? Should it just run 1 time, or should it run every X period or what?	1	1	https://issues.jboss.org/browse/IPAAS-284
8264	IPAAS-284	How do you map between suggestion -> trade from https://github.com/kcbabo/ipaas-demo/blob/master/trade-insight-api/src/main/java/org/example/ipaas/Suggestion.java to https://github.com/kcbabo/ipaas-demo/blob/master/day-trade-api/src/main/java/org/example/ipaas/Trade.java The fields do not seem to match-up so intuitively.	1	1	https://issues.jboss.org/browse/IPAAS-284
8265	IPAAS-250	Thanks, SJ. Icons look good to me. I'm on the fence with the HTTP one - I don't have an alternate suggestion, however, so I'm fine to stick with the current version. :-) A few questions: * Will there be a generic icon for steps? * Looks like some of our icons are color and some are black and white - do you think that will look OK? * We might need to add a few more icons to build out the connection catalog for Sprint 11. I'm certain we'll need another one for REST API. We might be able to get away with a non-specific icon with "API" underneath?	1	1	https://issues.jboss.org/browse/IPAAS-250
8266	IPAAS-250	[~kcbabo] yeah http was a tricky one for me because I wanted to avoid using text within the icon. I can explore further designs in the time being. To address your question, there will be a generic icon for steps so I'll provide that as well. I agree with you that we need to add additional icons to build out the connection list. I was thinking of adding facebook, instagram, linked in, google+ and github. What do you think of those? Also, that will add more colors instead of just blue and gray so hoping that will look ok. Let me know what you think!	1	1	https://issues.jboss.org/browse/IPAAS-250
8267	IPAAS-245	I dont see this in the narrative. Where is this being used?	1	1	https://issues.jboss.org/browse/IPAAS-245
8268	IPAAS-244	Where should it be stored then? And I bet you may want to show **** or something in the ipaas ui as placeholders. But still know if a field is in use or not. We cannot just show **** in all fields that is marked secret=true in the Camel connector.	1	1	https://issues.jboss.org/browse/IPAAS-244
8269	IPAAS-244	bq. Where should it be stored then? I should say that they can't be stored in plain text or read by untrusted parties in the git repository. If they are encrypted, that would work too. A consideration there is how these values are actually decrypted when used outside of iPaaS (e.g. on-premise runtime with FIS). IINM, Jimmi was thinking of using secrets to store passwords. bq. And I bet you may want to show **** or something in the ipaas ui as placeholders. But still know if a field is in use or not. We cannot just show **** in all fields that is marked secret=true in the Camel connector. Yes, good point.	0	0	https://issues.jboss.org/browse/IPAAS-244
8270	IPAAS-237	I want to see them all when we go to choose an action. Twitter Mention is the one that will actually be used. Just for my own curiosity, does it take a lot of effort to add these actions? Having that info will help us schedule/plan as we'll need to consider not only the connectors we'll support for PR1, GA, etc., but also the actions on those connectors.	0	0	https://issues.jboss.org/browse/IPAAS-237
8271	IPAAS-235	[~kcbabo], which actions will you actually use?	1	1	https://issues.jboss.org/browse/IPAAS-235

#	Requirement	Questions	Answered	Successfully Answered	URL
8272	IPAAS-232	[~kcbabo], [~dongniwang], [~crobson] Here is the first pass for the landing page design. I used the branding colors and patterns used for Fuse as well as Red Hat brand styling. This is what they would see when they are logged out of iPaaS. Consequently, if they are logged in, they would not be directed to this page, rather the dashboard of iPaaS. A couple questions for you [~kcbabo]: - Is Red Hat iPaaS the legal name for it? If so, this is what the product logo would look like. If not, then this design would have to change slightly. - Could you provide us with a short description of iPaaS to put on this page? Happy to hear any feedback you may have on this design, thanks! !iPaaS_LandingPage.png|thumbnail!	1	1	https://issues.jboss.org/browse/IPAAS-232
8273	IPAAS-203	[~jdyson] If I understand the story right, its about having the Rest API call, even when there is no real backend functionality. So some sort of simulation is ok (like blocking for some seconds and maybe throw some errors for certain requests). Is this correct ?	1	1	https://issues.jboss.org/browse/IPAAS-203
8274	IPAAS-203	[~jdyson] But do we really want to have the validation on an existing conection only and not _before_ creating the connection, too ? I mean from a UI perspective, you will enter the form, try a validation and only aftter this you create / save the connection. I believe we need both: Pre-creation validation (otherwise you would have to deal with nasty rollback issues, messing up history / auditing, ...) and validation for existing connections. Even when one tries to update a connection one does the validation step before persisting changes. Even more dragons: If you persist first, then validate but then you decide not to create the connection because you have forgotten your password, nobody forces you to push "cancel". You can simply close your browser. And who is then cleaning up the stale connection in the DB ?	1	1	https://issues.jboss.org/browse/IPAAS-203
8275	IPAAS-201	Does it actually get deleted or are we marking this as resolved because the UI elements to delete it are in place? Last I checked we don't make the actual call to the REST API (but that should be a quick fix, we just didn't have time to make it for the demo).	0	0	https://issues.jboss.org/browse/IPAAS-201
8276	IPAAS-176	[~jdyson] - Can I just call this Edit Integration? Or Integration Edit View? I think it's the same thing.	0	0	https://issues.jboss.org/browse/IPAAS-176
8277	IPAAS-159	This takes care of the initial create button. Can we add additional design links which show the user selecting the Start and Finish connections as well?	1	1	https://issues.jboss.org/browse/IPAAS-159
8278	IPAAS-157	Hey - with and without configuration is great. Do we have a version of these distinct from the integration editor? Both of these look like inline configuration of a connection in the integration editor.	1	1	https://issues.jboss.org/browse/IPAAS-157
8279	IPAAS-153	This is the epic for the whole sprint - do we really want to reject it? (holding back tears of rejection)	1	0	https://issues.jboss.org/browse/IPAAS-153
8280	IPAAS-146	tech debt.. moving to github	0	0	https://issues.jboss.org/browse/IPAAS-146
8281	IPAAS-116	FWIW I vote for option c). Trying to productionise iPaaS using the same approach as traditinal JBoss libraries sounds very wrong and hopefully whatever processes, document or checkboxes that originally needed this can be avoided. If Red Hat want to be able to deliver software faster to meet both customer and product demands then this really needs to be resolved. BTW did OpenShift do the same for Jenkins? I'd imagine not as there are a lot more plugins and dependencies I'd have thought, can we just use their approach?	1	0	https://issues.jboss.org/browse/IPAAS-116
8282	IPAAS-90	[~rajdavies]- Is this a duplicate of IPAAS-87? If so, can I close it?	1	1	https://issues.jboss.org/browse/IPAAS-90
8283	IPAAS-86	What's a project in this context? Is this the thing that holds the code for a connection or integration?	1	1	https://issues.jboss.org/browse/IPAAS-86
8284	IPAAS-72	Stan I think you did this already, right? Can this issue be closed or broken up into sub-tasks?	1	1	https://issues.jboss.org/browse/IPAAS-72
8285	IPAAS-69	Hi Hiram, I've re-opened this issue as I feel it is critical to the release and so did the program team as it was deemed an MVP. I appreciate it may not be the most exciting engineering work on the planet, however, iPaaS is all about the customer perception and we have 3 tiers. 1. Will it intrigue potential users enough to look at it ad say, wow, I can use something like this! 2. How easy it is to use and 3. How easy it is to get help. 2 and 3 go hand in hand and supportability can't be an afterthought. It needs to stay on the backlog and be addressed before GA. Susan	1	0	https://issues.jboss.org/browse/IPAAS-69
8286	IPAAS-52	[~joqian] - Do we need any designs for this, btw? Or maybe just a normal confirmation modal?	1	0	https://issues.jboss.org/browse/IPAAS-52
8287	IPAAS-47	[~kcbabo] [~stlewis] - Is this the same as the 'Create an Integration' page? The designs for which are here: https://redhat.invisionapp.com/share/TJ8KKMWGX#/screens/197805926	1	1	https://issues.jboss.org/browse/IPAAS-47
8288	IPAAS-46	Does this link into ConfigMap or is it just the variables are passed into the app container?	1	1	https://issues.jboss.org/browse/IPAAS-46
8289	IPAAS-46	[~joqian] [~kcbabo] - Do we need designs for this? Is there a page for this?	1	0	https://issues.jboss.org/browse/IPAAS-46
8290	IPAAS-44	[~davsclaus] Does this command work from any project or no project? If I deploy a connector jar to nexus so that ipaas-search-connectors can find it, I can't seem to get ipaas-connector-details to show any details on it. Also, can you get ipaas-connector-details to show up in the fabric8-console forge UI? Seems it's set to be hidden, ideally I kinda want to be able to see all of these commands from our test UIs for development purposes. Thanks!	1	1	https://issues.jboss.org/browse/IPAAS-44
8291	IPAAS-44	Thanks! Hey, should I be logging ipaas command related issues in that repo BTW?	1	1	https://issues.jboss.org/browse/IPAAS-44
8292	IPAAS-42	Agree on technical limitations of just checking connection details only. Despite that fact, we need to find a sensible way of making this work for iPaaS. Can we synthesize a route with the connection as the consumer endpoint and just route to a Mock endpoint to capture the response?	0	0	https://issues.jboss.org/browse/IPAAS-42
8293	IPAAS-41	Can we clarify on the actor here when we say "build a new version of the connector"? The expert user in this case is using the iPaaS UI to edit details of a connection. How does this user build a new connector through that interface?	1	1	https://issues.jboss.org/browse/IPAAS-41
8294	IPAAS-41	What's confusing here is that there are 2 distinct parts - server side and UI. Can we not reopen - but create a new UI specific issue instead?	1	1	https://issues.jboss.org/browse/IPAAS-41
8295	IPAAS-41	I agree that it's confusing, but I think it's also confusing if we mark the user story as done when it's not done. Can we create sub-tasks or related issues for the back end and UI work instead?	1	1	https://issues.jboss.org/browse/IPAAS-41
8296	IPAAS-41	An expert user really shouldn't be writing or updating code. The bulk of their interaction with iPaaS will be through the UI. Are we saying that changing things such as the name of a connection requires code-level tasks from a developer? That seems like a pretty steep drop off in usability and pulls in multiple personas to handle relatively basic tasks.	1	1	https://issues.jboss.org/browse/IPAAS-41
8297	IPAAS-41	Same comment as above w/r/t description, available fields, and defaults. Should we consider storing this type of information in a separate data store/repository vs. cooking it into the component code?	1	1	https://issues.jboss.org/browse/IPAAS-41
8298	IPAAS-41	[~davsclaus] Does this map to ipaas-edit-connector-options? If so, do we have a way of mapping a connection that's returned from ipaas-search-connectors to the project that builds the connector? Asking as I see that ipaas-edit-connector-options only seems to be available when I create a 'connector' project, from the connection library I'd then need to be able to map an entry in the library to the connector project I guess, unless there's a better forge command to use. Thanks!	1	1	https://issues.jboss.org/browse/IPAAS-41

#	Requirement	Questions	Answered	Successfully Answered	URL
8299	IPAAS-41	Yeah, maybe changing the name would be more a matter of duplicating the connector as a new project, might be a better question for [~kcbabo] to answer? Anyhoo, thanks for the additional commands! Follow on question, for these two items am I correct then in how these map to the relevant forge commands: * Changing the available fields (add or remove) ----> ipaas-choose-connector-options * Changing the default values for available fields ----> ipaas-edit-connector-options For either of the above it seems again like I need to reference the actual connector project when I fetch the command inputs. Are we going to have a way to associate a connector discovered via ipaas-search-connectors with the project source?	1	1	https://issues.jboss.org/browse/IPAAS-41
8300	IPAAS-40	From a user standpoint, I think there will be an expectation that I can add/update/remove labels at any time. If that means the component needs to be rebuilt, we need to think about what that means from a user experience standpoint (e.g. how is it rebuilt?).	0	0	https://issues.jboss.org/browse/IPAAS-40
8301	IPAAS-40	Can we document how this would be done through the iPaaS UI and how QE can test?	0	0	https://issues.jboss.org/browse/IPAAS-40
8302	IPAAS-39	Can we describe the remote API that can be used by QE and UI to satisfy this user story? Would prefer to keep these user stories open until they are fully satisfied.	0	0	https://issues.jboss.org/browse/IPAAS-39
8303	IPAAS-34	Will there be a UI component to this and, if so, should it come before or after the Users component is implemented? I'm asking because the Users component is now a priority for Sprint 4.	1	0	https://issues.jboss.org/browse/IPAAS-34
8304	IPAAS-33	Is this considered a user story for Sprint 4? It's written in a model-oriented manner (which I prefer, but I know we're using user stories like, "As a user I can login and logout" etc.). Either way I'll assign it to Sprint 4 since we need this for integration with the CD pipeline.	1	1	https://issues.jboss.org/browse/IPAAS-33
8305	IPAAS-33	[~crobson] [~joqian] - Do we have designs for user-related pages? Just asking as I'd like to have issues opened in GitHub to track them for the UI.	1	1	https://issues.jboss.org/browse/IPAAS-33
8306	IPAAS-25	Can links to the UI be added as comment or something so we/I can find information about this ticket?	0	0	https://issues.jboss.org/browse/IPAAS-25
8307	IPAAS-25	Hey Rachel - IPAAS-25 is just a top-level tracking issue for all of the connection-related user stories. It's fine if you want to include links here, but the main place we can use them from a tracking perspective are in the detailed users stories linked from this issue. That way, at the end of a sprint we can look at our user stories for that sprint and track whether we're done, in progress, how far, etc. Does that make sense?	1	0	https://issues.jboss.org/browse/IPAAS-25
8308	IPAAS-24	Is this for Sprint 4 or 5?	1	0	https://issues.jboss.org/browse/IPAAS-24
8309	IPAAS-24	Also [~crobson] [~joqian] - Do we have designs for this already?	1	0	https://issues.jboss.org/browse/IPAAS-24
8310	IPAAS-10	Are you running this using someone who has dedicated-admin role? https://docs.openshift.com/dedicated/3.2/getting_started/dedicated_administrators.html	1	0	https://issues.jboss.org/browse/IPAAS-10
8311	IPAAS-6	Where's the picture of [~kcbabo]?	0	0	https://issues.jboss.org/browse/IPAAS-6
8312	IPAAS-6	can we mark this as closed now?	0	0	https://issues.jboss.org/browse/IPAAS-6
8313	JBJCA-1210	Where can I find the commits for that issue? I searched here: https://github.com/ironjacamar/ironjacamar.org/commits What does setting the system property to true change? Thx.	0	0	https://issues.jboss.org/browse/JBJCA-1210
8314	JBJCA-989	Thank you for the advice ;-) What about would be this discussion in this case?	1	0	https://issues.jboss.org/browse/JBJCA-989
8315	JBJCA-989	What isn't working when building - there seems to be 2 issues here. But JIRA isn't a forum	0	0	https://issues.jboss.org/browse/JBJCA-989
8316	JBJCA-704	Good job! Where should ConnectionListener implementations be placed ?	1	0	https://issues.jboss.org/browse/JBJCA-704
8317	JBJCA-637	Will this automatically disable "our" pooling so that pooling can be delegating to the underlying <datasource-class>? If not, we might want to add another option for this because some vendors' javax.sql.DataSource implementations do their own pooling.	0	0	https://issues.jboss.org/browse/JBJCA-637
8318	JBJCA-637	Tom Fonteyne <tfonteyn@redhat.com> made a comment on [bug 957175|https://bugzilla.redhat.com/show_bug.cgi?id=957175] @John Doyle: "if a datasource class is defined first, then fallback to java.sql.Driver - the opposite of today." This bit is incorrect: "the opposite of today." The situation "today" is that the datasource class can *never* be used in EAP for the simple fact that all major vendors driver have the driver class preset. @Jesper Pedersen: " The "datasource-class" element is there in order to support JDBC like drivers which doesn't implement the java.sql.Driver class." Could you please provide an example ? I mean do you have such a non-driver that can be deployed to EAP so we can see the functionality working as you state ?	0	0	https://issues.jboss.org/browse/JBJCA-637
8319	JBJCA-571	Could you update this patch against trunk/ ? Also this mvn org.apache.maven.plugins:maven-deploy-plugin:2.6-SNAPSHOT:deploy-file looks incorrect.	1	1	https://issues.jboss.org/browse/JBJCA-571
8320	JBJCA-28	Is this done?	1	0	https://issues.jboss.org/browse/JBJCA-28

#	Requirement	Questions	Answered	Successfully Answered	URL
8321	JBJCA-12	ObjectName("jboss.jca:service=XaTxCM,name=" + jndiURL); // try to remove mbean if leftover... if (server.isRegistered(cmName)) { server.unregisterMBean(cmName); } Object cmBean = server.createMBean("org.jboss.resource.connectionmanager.TxConnectionManager", cmName); assert cmBean != null : "Server never created TxConnectionManager for: " + schema; assert server.isRegistered(cmName) : "Server never registered ConnectionManager for: " + schema; //ManagedConnection Factory ObjectName mcfName = new ObjectName("jboss.jca:service=ManagedConnectionFactory,name=" + jndiURL); // try to remove mbean if leftover... if (server.isRegistered(mcfName)) { server.unregisterMBean(mcfName); } Object mbean = server.createMBean("org.jboss.resource.connectionmanager.RARDeployment", mcfName); assert mbean != null : "Server never created ManagedConnectinoFactory for: " + schema; assert server.isRegistered(mcfName) : "Server never registered ManagedConnectionFactory for: " + schema; //ManagedConnectionPool ObjectName mcpName = new ObjectName("jboss.jca:service=ManagedConnectionPool,name=" + jndiURL); // try to remove mbean if leftover... if (server.isRegistered(mcpName)) { server.unregisterMBean(mcpName); } Object pool = server.createMBean("org.jboss.resource.connectionmanager.JBossManagedConnectionPool", mcpName); assert pool != null : "Server never created Connection Pool for: " + schema; assert server.isRegistered(mcpName) : "Server never registered Connection Pool for: " + schema; ObjectName serviceControllerName = new ObjectName("jboss.system:service=ServiceController"); // anon block so I don't rename var al { AttributeList al = new AttributeList(); al.add(new Attribute("JndiName", jndiURL)); al.add(new Attribute("ManagedConnectionPool", mcpName)); al.add(new Attribute("CachedConnectionManager", new ObjectName("jboss.jca:service=CachedConnectionManager"))); final ObjectName tmsName = new ObjectName("jboss:service=TransactionManager"); al.add(new Attribute("TransactionManagerService", tmsName)); //BUG? // al.add(new Attribute("LocalTransactions", Boolean.TRUE)); al.add(new Attribute("LocalTransactions", Boolean.FALSE)); al.add(new Attribute("TrackConnectionByTx", Boolean.TRUE)); // try to set the attributes on the bean try { server.setAttributes(cmName, al); } catch (UndeclaredThrowableException e) { final Throwable cause = e.getCause(); if (cause instanceof MBeanException) { LOGGER.error("Exception setting Attribute on Server handle: " + e.getMessage(), e); throw (MBeanException) cause; } //else throw new HRNException(cause, ApplicationConstants.BASE_RESOURCE, "common.error.db.connectionfailed", new Object[] { cause.getMessage() }); } } // RARDeployment { AttributeList al = new AttributeList(); al.add(new Attribute("ManagedConnectionFactoryClass", "org.jboss.resource.adapter.jdbc.xa.XAManagedConnectionFactory")); al.add(new Attribute("ConnectionDefinition", "javax.sql.DataSource")); // al.add(new Attribute("LogWriter", ?));//PrintWriter // try to set the attributes on the bean server.setAttributes(mcfName, al); } // anon block so I don't rename vars { AttributeList al = new AttributeList(); al.add(new Attribute("PoolJndiName", jndiURL)); al.add(new Attribute("ManagedConnectionFactoryName", mcfName)); al.add(new Attribute("MinSize", new Integer(0))); al.add(new Attribute("MaxSize", new Integer(50))); al.add(new Attribute("BlockingTimeoutMillis", new Long(5000))); al.add(new Attribute("IdleTimeoutMinutes", new Integer(7))); al.add(new Attribute("Criteria", "ByNothing")); // try to set the attributes on the bean server.setAttributes(mcpName, al); } //DataSource ObjectName dsName = new ObjectName("jboss.jca:service=DataSourceBinding,name=" + jndiURL);	1	0	https://issues.jboss.org/browse/JBJCA-12
8322	JIRA-521	[~paul.robinson] what type should be 'PDD Experience Link' field? One URL, multiple URLs, something else?	1	1	https://issues.jboss.org/browse/JIRA-521
8323	JIRA-520	Is it possible to add a 'Epic Type' field, with two options: "Business Requirement", "Technical Requirement"? The default would be "Business Requirement". This would appear after the "Epic Name" in the issue editor view.	1	1	https://issues.jboss.org/browse/JIRA-520
8324	JIRA-502	Can't RHMAP have have own workflow with changed default priority? Why to change default for all the projects because of one project :) ? I thing MAJOR is the default as long as I can remember ... 7 years.	1	1	https://issues.jboss.org/browse/JIRA-502
8325	JIRA-502	Rather than changing the default for every established project wouldn't it be better to reach out to Atlassian and see if they would consider a feature request? Alternatively if for a single project couldn't an alternative triage based workflow be developed that allows for untriaged issues to be separated from the triaged issues.	1	1	https://issues.jboss.org/browse/JIRA-502
8326	JIRA-502	I don't see strong enough benefits for such a wide change. How about adding a normal/medium severity level as teh new default? Currently Jira has 6 levels and major is already at the median level. So I don't see what this change will really achieve. I think this is what Bugzilla does (having 7 levels of severity with normal as the middle default).	1	1	https://issues.jboss.org/browse/JIRA-502
8327	JIRA-502	[~rrajesh] [~jfrizell] [~lgriffin] would you be so kind to explain the initial motivation from the RHMAP side please (with already ammended knowledge from Vlastimil & others)? Tried looking at the original Jira, but couldn't find deep reasoning explained, just saw the plain request without justification as of why should this change happen and how would that help RHMAP? After reading some comments from Leigh here in the Jira, I get the feeling, but still I don't see the big picture why should this help more than the other way around. Anyway a lot of interesting points brought there. Let me say one thing: * "Major" changed to "Minor" on triage / by component lead is better, than the other way around, when a potentially important "Major"+ issue will be left unnoticed as "Minor" as triage didn't happen. Think of triage as asynchronous process, going from minor by default & raising issue priority requires active participation and has far more significant possible negative concaveness effect [1]. Read only optionally for context of the term above: (optional)[1] Convexity, Robustness, and Model Error inside the Fourth Quadrant http://www.fooledbyrandomness.com/OxfordBTLecture.pdf	1	1	https://issues.jboss.org/browse/JIRA-502
8328	JIRA-502	[~vtunka] - It's really about the fact that Major should not be the default priority for all new tickets - it looses all meaning when 75% of all tickets [1] have that priority. In reality, it's less about changing it *to* Minor as it is about changing it *from* Major. I quite like [~cfang]'s idea of a new priority level that things default to - whether it's "Normal/Medium" or "Not Set" - personally, I think defaulting Priority to "Not Set" would be a HUGE benefit - it allows people to immediately see what tickets have not yet been triaged and allows "Major" to actually mean something again. Interestingly, based on the linked chart [1] it seems there is an option to not have a priority set (741 tickets have this). Even more interesting, this does not seem to be a historical thing - there were tickets created event *today* with the priority field not set - Filter for the list of these tickets: https://issues.jboss.org/issues/?filter=12330317. Is this to do with the idea of committers are the only ones who can set priority? Is there any way to leverage this (perhaps at a project level so it does not affect everyone) to prevent priority from ever being set when a ticket is created - i.e. have a custom workflow for RHMAP that simply hides the priority field at ticket creation time? [1] https://www.dropbox.com/s/uwyqttxojhnbkje/Screenshot%202016-12-08%2012.04.48.png?dl=0	0	0	https://issues.jboss.org/browse/JIRA-502
8329	JIRA-502	[~velias] Regarding second drawback, could you show Priority in Create form only to internal people? Regarding first drawback, if the fields were named User urgency and Development priority, that would lower the confusion factor, IMO. Personally, I've been much confused by shoelaces many years ago, I am hoping that resolving Priority/Severity confusion is also just a matter of time and practice. Although, if the problem could be resolved by more apt naming of the field, either what I suggested, or something else, that would be awesome.	1	1	https://issues.jboss.org/browse/JIRA-502
8330	JIRA-462	What do you think [~mfreer], would this plugin suffice instead?	1	1	https://issues.jboss.org/browse/JIRA-462

#	Requirement	Questions	Answered	Successfully Answered	URL
8331	JIRA-462	Sounds good to me Gliffy supports imports from *.vsdx and *.vdx Diagrams so its good enough for me thanks. if we could have Gliffy by default for our profiles [~velias] in JIRA that would be great. Could you give me some pointers as how to enable / use Gliffy as I am new to JIRA. Thanks	1	1	https://issues.jboss.org/browse/JIRA-462
8332	JIRA-406	Really cool plugin. How it works in different browsers?	1	1	https://issues.jboss.org/browse/JIRA-406
8333	JIRA-389	Hi, may we better define what means "participate" on issue. There are these possibilities: - Reporter - Assignee - Created comment - Attached file If we define this, I may create simple jira report which gives you email addresses of users who do something we defined for selected versions. And maybe I may add you list of issue keys where participated. You may then send "Thanks" emails by some standard way. Second option is create more sophisticated plugin which will perform complete mailing, but this seems rather heavy for now. So is report with list of emails enough for now?	1	1	https://issues.jboss.org/browse/JIRA-389
8334	JIRA-389	Vlastimil, where is this report? I'm excited to use it for the Seam 3 release.	0	0	https://issues.jboss.org/browse/JIRA-389
8335	JIRA-386	>> 2. there are some implications for text emails sent from jira - wiki markup stays here (is not rendered). May be problem if text emails are processed by some other system. Wouldn't this be a major problem while viewing JIRA email notifications in a mail client? I have intentionally configured text notifications from JIRA, instead of HTML. Would that mean, we would have to switch it to HTML or some other format for notifications?	0	0	https://issues.jboss.org/browse/JIRA-386
8336	JIRA-382	Looks like this plugin is not maintained for over a year. I could fork the project and make it work in JIRA 6, would that be enough to migrate?	1	1	https://issues.jboss.org/browse/JIRA-382
8337	JIRA-382	It looks like someone already did that: https://bitbucket.org/badazzhindu/jira-component-watcher-plugin/src Would you be willing to test it locally?	1	1	https://issues.jboss.org/browse/JIRA-382
8338	JIRA-382	Attempting to login in a mobile devices doesn't seem to work. Could it be a CAS issue?	1	1	https://issues.jboss.org/browse/JIRA-382
8339	JIRA-369	That seems to contain enough information, there's actually no need to download full build logs. How should develop this gadget ? Is this something that can be undertaken by the jboss.org team ?	1	1	https://issues.jboss.org/browse/JIRA-369
8340	JIRA-363	Hey Vlastimil, was there a reason the project was created as WINDUPRULE not WINDUPRULES? Thanks, Lincoln	1	1	https://issues.jboss.org/browse/JIRA-363
8341	JIRA-363	Also, could you set the permissions to be the same as WINDUP? thank you.	1	1	https://issues.jboss.org/browse/JIRA-363
8342	JIRA-363	Hi [~lincolnthree], jira key has length constraint to 10 characters, so WINDUPRULE is at maximal length. I set both permission scheme and project roles to same value as WINDUP during creation, so I believe permissions are same. Do you have any problem in this area?	1	1	https://issues.jboss.org/browse/JIRA-363
8343	JIRA-363	Also, could you please make the issue edit page be just one page, without the split into tabs?	1	1	https://issues.jboss.org/browse/JIRA-363
8344	JIRA-344	Simplified screen has been implemented in production JIRA and activated for * ARQGRA * ARQRUSH * MODE who requested it. I can activate it for other projects by request (I'm going to send announcement to the community to inform it)	0	0	https://issues.jboss.org/browse/JIRA-344
8345	JIRA-338	Which is real name of project to be used and who should be configured as Project Lead?	0	0	https://issues.jboss.org/browse/JIRA-338
8346	JIRA-338	[~ozizka] any news on this? Do you need this project to be created still or not? If yes then provide me necessary info please (see previous comment)	0	0	https://issues.jboss.org/browse/JIRA-338
8347	JIRA-323	Hi Aleksandar, do you have some examples of issues or comments where wiki syntax is a problem? I don't know yet about any example of this problem. On the other side I noticed some old descriptions and comments where wiki syntax helps and shows content nicely. But this definitely depends on content. I'm not sure if this kind of bulk change is necessary now, maybe better solution is only solve cases with problems.	1	1	https://issues.jboss.org/browse/JIRA-323
8348	JIRA-318	The http://issues.jboss.org/browse/ORG project is used to track releases of services maintained by the JBoss Community Development Team so using a component wouldn't make sense. I think using a separate project is the best way forward so that you can manage your own release cycles. Do you have a working name for the project or should we simply call it Build Tracker and use the key BUILDTRK?	0	0	https://issues.jboss.org/browse/JIRA-318
8349	JIRA-317	Copying some of the IRC discussion about this on #jbossorg: 12:17:02 IST) Jaikiran: anyone around for a quick chat about the jira/jboss.org email id consolidation? (12:27:33 IST) jchochol: velias: ^^^ (12:28:01 IST) velias: Jaikiran: hi (12:28:25 IST) Jaikiran: velias: hi velias (12:28:50 IST) Jaikiran: I saw a mail on jboss-dev yesterday that the email in jira will now be the same as the one in jboss.org account (12:28:57 IST) Jaikiran: is that correct? (12:29:46 IST) velias: yep (12:30:30 IST) Jaikiran: is there any way an individual can have separate emails for them? (12:30:48 IST) Jaikiran: right now, I have set jira to my work email so that i get those important notifications (12:31:11 IST) Jaikiran: and for my jboss.org account which is used in forums, mailing list etc, i have used a non-work email (12:32:15 IST) velias: there is not possibility to have separate emails for now, but I can imagine some solution for this if necessary (12:32:41 IST) Jaikiran: i would really like to have it that way. do you want me to add a feature request? (12:32:45 IST) velias: just for now we are trying to consolidate user profile management to one place and keep it as simple as possible (12:34:19 IST) Jaikiran: yeah, i understand that part. and have no issues with that. I however believe that letting individual users have multiple emails for different activities would be a good thing (12:34:47 IST) Jaikiran: one more reason why i want to keep it separate is because, for my work email, i will require VPN access (12:35:14 IST) Jaikiran: so i can't use it for most part of the community work because I won't be able to access it from any other machine whether vpn isn't setup (12:36:01 IST) velias: ok, I understand. Add feature request into jira component of "jboss.org development" project please (12:37:26 IST) Jaikiran: ok thanks, will do (12:38:08 IST) velias: Probably simplest solution is to create some "extended" part of user profile, which allow to set separate email for jira (and maybe for other systems). If not set, then jira use main email from account. (12:39:53 IST) Jaikiran: yep, that sounds good. (12:42:51 IST) velias: Jaikiran: ok, thanks for your suggestion	1	1	https://issues.jboss.org/browse/JIRA-317
8350	JIRA-315	Hi, JIRA has no function to change project key. One workaround is described here: https://jira.atlassian.com/browse/JRA-11140?focusedCommentId=64516&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-64516 We need to create new project with new key and name, set all settings to same values as on old one, create all components and versions here, and then bulk move all issues into new project. Issue keys are changed in this case, but URL's with old issue keys still work and are redirected to new one. So process should be: # I'll create new project with settings (schemes, users in project roles etc) # You will create all necessary components and versions here (with exact same names as in old projects to be mapped automatically during move) # You will move all issues into new project # I'll create redirect from https://issues.jboss.org/browse/ARQAJO to https://issues.jboss.org/browse/ARQGRA in our WebProxy Apache to have redirect for main project page too (moved issues are redirected automatically by JIRA) # I'll delete project ARQAJO Are you OK with this Lukas?	1	1	https://issues.jboss.org/browse/JIRA-315

#	Requirement	Questions	Answered	Successfully Answered	URL
8351	JIRA-315	Ok, I have finished steps 2 and 3. Could you please finish the steps 4 and 5?	1	1	https://issues.jboss.org/browse/JIRA-315
8352	JIRA-310	1) I think your mention is about Suggestimate functionality on "Issue detail" page. What do you mean about this info on "Create issue" page, do you seem it useful here? 2) I try to negotiate with plugin provider some change in showing this info on "Issue detail" page, I see these possibilities: - disable/enable it on per project basis - disable/enable it on per user basis - move this info from main working area into tab (beside other tabs: All, Comments, History, ...) As temporary improvement I decreased max. number of shown similar issues to 5.	1	1	https://issues.jboss.org/browse/JIRA-310
8353	JIRA-307	You would need admin user pass word for Hibernate Jira to create the bidirectionality. Is there really a difference if we each set up the linking on our own side?	1	1	https://issues.jboss.org/browse/JIRA-307
8354	JIRA-287	hi [~jnunes], I looked shortly at your request and I'm not sure if I understand why do you need new "excel like" plugin. I believe that for work with issues jira provides quite powerful search API. If you switch is into "List View" mode it provides excel like table view of issues where you can control which columns are visible. https://confluence.atlassian.com/jira/using-the-issue-navigator-185729658.html#UsingtheIssueNavigator-ListView You can also use two modes of search GUI - simplified which uses selectboxes for filtering, and advanced where you can use whole power of JQL (JIRA query language) for filtering. You can also store your searches as filters for later use, and use dashboards to provide current overviews around issues in some project (https://confluence.atlassian.com/jira/customizing-the-dashboard-185729498.html). Another way is to use JIRA Agile Reports - https://confluence.atlassian.com/agile/jira-agile-user-s-guide/using-a-board/using-reports Should you somehow describe why these tools do not suit your needs please, what is missing in them so why you need some "Excel like" plugin? I really prefer to use existing JIRA features than installing extra plugins. Then we can try to install plugins you proposed into our JIRA STAGING server where you can test them to be sure they match your needs.	1	1	https://issues.jboss.org/browse/JIRA-287
8355	JIRA-287	Any news on this, or can I close it as obsolete?	1	1	https://issues.jboss.org/browse/JIRA-287
8356	JIRA-272	Hi [~NadirX] We use status called {{New}} in some other workflows for this. Is it OK for you to use it there also not to confuse users too much?	1	1	https://issues.jboss.org/browse/JIRA-272
8357	JIRA-272	Hi [~NadirX], I prepared new workflow in JIRA, see attached image. Only change against original Git workflow is "New" status at the begining, which is used when issue is created. There are three transitions from this status: * "Hand Over to Development" to switch issue into Open status * "Resolve issue" as common from other states * "Close issue" as common from other states Once I switch ISPN project to use this WF then new issues start to flow into "New" status. Old issues will stay in current statuses. Do you have some preferred time when I should switch ISPN project to this new WF so you can test new WF immediately?	1	1	https://issues.jboss.org/browse/JIRA-272
8358	JIRA-269	Vlasta, could you please also enable Epics for RF, RFSBOX and RFPL projects?	1	1	https://issues.jboss.org/browse/JIRA-269
8359	JIRA-269	Hi, RF and RFSBOX has been switched. But RFPL do not use Git workflow currently. It uses "JBoss Platforms Support Case Workflow v1" scheme which defines some special workflow for "Support Patch" issue type and common jira workflow for other issue types. Are you sure we can switch it to Git workflow?	1	1	https://issues.jboss.org/browse/JIRA-269
8360	JIRA-269	Hi bq. I can't create Epic for RF, RFPL or RFSBOX. What do you do to create Epic? I think there was some change in GreenHopper in this area. Do you follow this process? https://confluence.atlassian.com/display/GH/Creating+an+Epic bq. On the other hand, I see "Support Patch" type of issue in ARQGRA. ARQGRA project uses "Default Issue Type Scheme" which contains all issue types available in JIRA. I can switch it to some other issue types scheme with less types. Here on bottom right I can see which types you use in ARQGRA - https://issues.jboss.org/browse/ARQGRA#selectedTab=com.atlassian.jira.plugin.system.project%3Aissues-panel There is not suitable type scheme for now, but I can create it for you. Are you OK with this?	1	1	https://issues.jboss.org/browse/JIRA-269
8361	JIRA-267	It's possible to create it, but which project need this? There are many Notification schemes in our jira, many of them are clones for distinct projects due emailing to lists. So add this to all projects? And which notifications send to this role, all, only issues related, or work and worklog related too?	1	1	https://issues.jboss.org/browse/JIRA-267
8362	JIRA-267	I had the JBBUILD project in mind specifically, but I think this should be available to all projects. Is there a default "subscribers" role that come with jira? This role should email members when new issues are created, new comments are added, issues are closed, projects are released, etc. Basically any changes to the issues of that project should send a notification to the members of that role.	1	1	https://issues.jboss.org/browse/JIRA-267
8363	JIRA-267	59 sounds like a lot of different notification schemes, I bet a lot of them are very similar, can they be consolidated? So how do I use the notification scheme in JBBUILD? I only see 4 possible roles to add myself to "Administrators", "Developers", "Private Reviewers", "Users". "Private Reviewers" should be removed because it doesn't seem to apply to this project. What is the difference between "Developers" and "Users"?	1	1	https://issues.jboss.org/browse/JIRA-267
8364	JIRA-261	Hi Xavier. Last build date is not part of 'project info' returned from jenkins server REST call. http://hudson.jboss.org/hudson/job/as7-master-build/api/json Only way to obtain it is to do subsequent REST call to the jenkins server to get build details. http://hudson.jboss.org/hudson/job/as7-master-build/16181/api/json Are you interested to add date/time only for "last build", or for "Last successful build" and "Last stable build" also?	1	1	https://issues.jboss.org/browse/JIRA-261
8365	JIRA-251	It is possible to update Target Release field if release ack flag is empty or "-", or always by user with project role "CDW PM". Solution we should implement is to reset Target Release (and all acks) field during clone, hopefully it will not collide with other clone use cases, I'm not sure if and how other teams use it. WDYT? [~jpallich] [~atangrin] any comments on this?	1	1	https://issues.jboss.org/browse/JIRA-251
8366	JIRA-220	Any chance of an update on this - is it going to be done at some stage / at all / never ?	1	1	https://issues.jboss.org/browse/JIRA-220
8367	JIRA-219	I agree that it is not a good idea to automatically apply that change to all projects. Thanks for your help. Heiko	1	1	https://issues.jboss.org/browse/JIRA-219
8368	JIRA-217	Hi Pete, should you elaborate a bit how "default QE assignee" should work? Assign "component QE lead" to the Tester field only when new issue is created with filled component? Or even later when component is set first time?	0	0	https://issues.jboss.org/browse/JIRA-217
8369	JIRA-209	Hi [~velias] / [~rrajesh] - can we get some feedback on this? What do we need to do to push this forward? Thanks.	1	1	https://issues.jboss.org/browse/JIRA-209
8370	JIRA-209	good points. can you elaborate the "see multiple sprints on the one screen" - you mean the backlog view ? since for now I just have a backlog that queries everything and thus I can see it all show up.	1	0	https://issues.jboss.org/browse/JIRA-209
8371	JIRA-207	+1 for changing CSS for all PR links (in the jira field and also in the body of all comments) to reflect status. Would it not be better to know which are merged (strikethrough) vs. closed (greyed out) vs. merged AND closed (strikethrough and greyed out)? What about when the PR's topic branch has been deleted?	0	0	https://issues.jboss.org/browse/JIRA-207

#	Requirement	Questions	Answered	Successfully Answered	URL
8372	JIRA-207	Is anybody able to help me prepare javascript function for this feature? I can embed it into JIRA then. Specification of JS is: * use jQuery or plain JS only * one javascript function which accepts {{id}} of hrml {{span}} element. Then it should: ** take whole content of this span ** check if it is an url of github pull request, if yes then ** do REST call to github API to get status of PR ** add necessary style to the input span to strike it if PR is closed already It will help me a lot to bring this feature to JIRA if anybody is able to prepare this JS. Thanks in advance.	0	0	https://issues.jboss.org/browse/JIRA-207
8373	JIRA-202	Hi Ondrej, some products are migrated to bugzilla for some time. Do you know if this kind of patterns is really used in jira, so is worth to implement this feature?	1	1	https://issues.jboss.org/browse/JIRA-202
8374	JIRA-201	"is related to by" sounds really clunky I also agree that the "was cloned *" is better as "cloned *" due to fewer words	1	1	https://issues.jboss.org/browse/JIRA-201
8375	JIRA-201	In terms of trimming the sentence length, I don't think you ever see the wording on a Jira issue exactly like "foo is blocked by bar". For example, there is a "Issue links" section in this particular Jira. As you can see ORG-2283 is not actually present within the "sentence". Also, I do agree that "relates to" does not imply directionality, therefore should it be allowed on either side as if you look at a bug and it says "A relates to B" a casual observer can't necessarily infer the direction from the wording present on the actual Jira. How about "relevant to"/"is relevant to"?	0	0	https://issues.jboss.org/browse/JIRA-201
8376	JIRA-198	There is one small problem in "Seam Workflow" workflow - "Pull Request Closed" step leaves "Resolution" on bad value "Unresolved". Compared to standard "Resolve Issue" step it has no screen assigned. So solution is to assign standard "Resolve Issue Screen" screen to this step. I tested it on STAGING and seems OK. What do you mean Dan, may I add this to PROD too? Second thing we need to solve if we want reuse this workflow is to map "Git Pull Request" field to other projects, because now seems to be mapped only in "Seam Field Configuration". So we need to enable it in "Default Field Configuration" too.	1	1	https://issues.jboss.org/browse/JIRA-198
8377	JIRA-198	bq. So solution is to assign standard "Resolve Issue Screen" screen to this step. I tested it on STAGING and seems OK. What do you mean Dan, may I add this to PROD too? Go for it. bq. So we need to enable it in "Default Field Configuration" too. And we also need to modify configuration for "Git Pull Request" custom field to allow it for new projects. I had the permissions limited in scope for the purpose of prototyping. Now that we are moving forward with it, you can open it up to other projects. bq. I also noticed that "Seam Workflow Scheme" assigns "Seam Workflow" to only small subset of issue types I excluded the other issue types because frankly, I don't know why we have them. They just confuse people. The less issue types there are, the better, IMHO. bq. Dan can migrate Seam projects to this updated workflows if want That sounds like the best approach.	1	1	https://issues.jboss.org/browse/JIRA-198
8378	JIRA-181	Problem with using issue status for "need feedback" is that you may need feedback/info from some person in nearly any phase of issue processing, which means you have to be able to switch to this status from nearly all other statuses (excluding resolved and closed). But there is a problem then how to switch issue back to the relevant status after feedback/info is provided. We have "NeedInfo" field present in our jira instance, which is 'user multiselect' type. It is not clearly described how to use it, but I believe it is used to name users you want info from. It is enabled only for some screens for now (only CDW screens), I can enable it for other screens/projects also if you think it may help you. WDYT?	1	1	https://issues.jboss.org/browse/JIRA-181
8379	JIRA-158	Hi David, I just found this plugin which can solve notification requirements partially https://marketplace.atlassian.com/plugins/com.riadalabs.jira.plugins.notificationassistant Limitations: - only JIRA administrator can configure notification emails sending, not Project Admin (which is organizational limitation only) - notification emails can be send to Assignee or defined email address only - I can try to communicate with plugin author to add more flexibility here What do you think?	1	1	https://issues.jboss.org/browse/JIRA-158
8380	JIRA-158	Hi David, notification assistant plugin is deployed on production jira server now. Should you specify which projects do you want to configure for "deadline" related notifications you requested by this issue please? I do not want to configure this notifications for all projects for now. I plan to announce this new feature to the community in near future to ask other projects leads if they want this or another type of notification for their projects.	1	1	https://issues.jboss.org/browse/JIRA-158
8381	JIRA-140	Hi Xavier, I'm not sure if I understand your request correctly. Do you want these fields to be added to issue detail page even if they are not filled in? It is not possible to configure this, JIRA simply shows only fields with some value at issue detail page. You have to use Edit page to add value, and field is shown on issue detail page then. If you have problem with too much fields on Edit Issue page then you can simply use "Configure Fields" > "Custom" option on top right of the Edit page and filter out fields you never use. We should try to somehow improve order of fields on Edit issue page, or we should try to organize them a bit better by use of Tabs, but it is a bit problematic as different projects and products have very different requirements in this area.	1	1	https://issues.jboss.org/browse/JIRA-140
8382	JIRA-140	we got Epic moved up in past, can't we just move Sprint up just under it ?	1	1	https://issues.jboss.org/browse/JIRA-140
8383	JIRA-135	How can I get notified on the 'Elastic Data Grid' component in the 'JBoss Enterprise Application Platform' (JBPAPP) project? I guess I have no permission to add myself to the recipient list.	1	1	https://issues.jboss.org/browse/JIRA-135
8384	JANDEX-23	Is it acceptable that it would be limited to declaring types? Methods/Classes/Fields etc	1	1	https://issues.jboss.org/browse/JANDEX-23
8385	JANDEX-23	If you had the full line number table is that better? What would you be able to do with it? For example, it could tell you that a local variable "i" was located at bytecode offset X, which was declared in method Foo. However, in the use case of error reporting, how would you know to look for "i"?	1	1	https://issues.jboss.org/browse/JANDEX-23
8386	JANDEX-22	[~jason.greene] Any update on this? If you give me any hint how to implement it reliably I can try to make some progress. Anyway, what's the Fix Version for this issue? 1.2.0.Beta2 obviously not...	1	1	https://issues.jboss.org/browse/JANDEX-22
8387	JANDEX-4	Thanks for the patch! I made some additional changes around close handling as a zip file is not always used. Can you retest the latest master and reopen if needed? Thanks	1	1	https://issues.jboss.org/browse/JANDEX-4
8388	JANDEX-1	Awesome, could you deploy a new release or at least a SNAPSHOT?	0	0	https://issues.jboss.org/browse/JANDEX-1

#	Requirement	Questions	Answered	Successfully Answered	URL
8389	JASSIST-257	Thanks for your reactivity ! It's better but now i get this : java.lang.UnsupportedClassVersionError: javassist/CtClass : Unsupported major.minor version 52.0 at java.lang.ClassLoader.defineClass1(Native Method) at java.lang.ClassLoader.defineClassCond(ClassLoader.java:631) at java.lang.ClassLoader.defineClass(ClassLoader.java:615) at java.security.SecureClassLoader.defineClass(SecureClassLoader.java:141) at java.net.URLClassLoader.defineClass(URLClassLoader.java:283) at java.net.URLClassLoader.access$000(URLClassLoader.java:58) at java.net.URLClassLoader$1.run(URLClassLoader.java:197) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:247) at com.meilleuregestion.instrumentation.Transformer.transform(Transformer.java:33) at sun.instrument.TransformerManager.transform(TransformerManager.java:169) at sun.instrument.InstrumentationImpl.transform(InstrumentationImpl.java:365) at java.lang.ClassLoader.defineClass1(Native Method) at java.lang.ClassLoader.defineClassCond(ClassLoader.java:631) at java.lang.ClassLoader.defineClass(ClassLoader.java:615) at java.security.SecureClassLoader.defineClass(SecureClassLoader.java:141) at org.apache.catalina.loader.WebappClassLoader.findClassInternal(WebappClassLoader.java:2895) at org.apache.catalina.loader.WebappClassLoader.findClass(WebappClassLoader.java:1173)	0	0	https://issues.jboss.org/browse/JASSIST-257
8390	JASSIST-198	The Exception is : Exception in thread "main" javassist.CannotCompileException: [source error] ; is missing Caused by: compile error: ; is missing	0	0	https://issues.jboss.org/browse/JASSIST-198
8391	JASSIST-197	It's not clear what you want. Can you show an example?	1	1	https://issues.jboss.org/browse/JASSIST-197
8392	JASSIST-197	Has a way of doing this been implemented yet?	0	0	https://issues.jboss.org/browse/JASSIST-197
8393	JASSIST-197	Is there any progress on this yet?	0	0	https://issues.jboss.org/browse/JASSIST-197
8394	JASSIST-163	I am happy to fix the above and submit a patch, but need a little help making sense of ProxyFactory.makeForwarder as I am not well-versed in bytecode directives. Ideally, could someone provide a plain java or pseudo-code translation of bytecode being generated in makeForwarder - i can't quite understand what that method is doing? Specifically, what is the value of Method[] methods when find2Methods is invoked (indirectly) from makeForwarder...	1	1	https://issues.jboss.org/browse/JASSIST-163
8395	JASSIST-163	Nikita, Your patch looks good but filling Method[] eagerly may cause unnecessary memory consumption and initialization overheads. I think the real problem is a jumbo lock. So what about modifying find2Methods() as following: {code} public static void find2Methods(Object self, String superMethod, String thisMethod, int index, String desc, java.lang.reflect.Method[] methods) { // synchronized (methods) { // if (methods[index] == null) { // methods[index + 1] = thisMethod == null ? null // : findMethod(self, thisMethod, desc); // methods[index] = findSuperMethod(self, superMethod, desc); // } // } if (methods[index] == null) { Method m1 = thisMethod == null ? null : findMethod(self, thisMethod, desc); Method m0 = findSuperMethod(self, superMethod, desc); synchronized (methods) { if (methods[index] == null) { methods[index + 1] = m1; methods[index] = m0; } } } } {code} Could you check this fix also removes performance hotspot? Thanks!	1	1	https://issues.jboss.org/browse/JASSIST-163
8396	JASSIST-163	The above change doesn't look safe, since other threads would be reading methods[index] without synchronizing (methods[index] might contain non-null but methods[index+1] could be null). Would it be easier to use a different datastructure instead for holding the pair of methods? Does the Methods[] methods array need to be longer than two elements or is it always sized for two?	1	1	https://issues.jboss.org/browse/JASSIST-163
8397	JASSIST-163	I agree that it's not easy to make sure of the correctness of the code above. In this sense, the code is not very good. However, since methods[even-number] is always set after methods[odd-number], if methods[index] is non null, methods[index + 1] is also non-null. (I suppose that the JVM's memory model guarantees this fact.) Furthermore, the values of m0 and m1 are always the same no matter which thread computes them. Am I correct?	1	1	https://issues.jboss.org/browse/JASSIST-163
8398	JASSIST-163	How are the elements in methods sort of final fields? I think that final fields have to be initialized during the referencing class constructor or class static initializer. Your change looks good to me. I'll paste the additional lines (from your change) below for others to also review. {code} /** * Finds two methods specified by the parameters and stores them * into the given array. * * @throws RuntimeException if the methods are not found. * @see javassist.util.proxy.ProxyFactory */ public static void find2Methods(Object self, String superMethod, String thisMethod, int index, String desc, java.lang.reflect.Method[] methods) { /* Once methods[index] and methods[index + 1] are set to non-null, * then their values never change. */ if (methods[index] == null || methods[index + 1] == null) { Method m1 = thisMethod == null ? null : findMethod(self, thisMethod, desc); Method m0 = findSuperMethod(self, superMethod, desc); synchronized (methods) { if (methods[index] == null) { methods[index + 1] = m1; methods[index] = m0; } } } } {code} If a another thread attempts to read methods[index] or methods[index+1], they will either see one of the fields set to null or both fields initialized. If either field is null, the thread will synchronize on methods and then only set both fields if they are still null.	1	1	https://issues.jboss.org/browse/JASSIST-163
8399	JASSIST-163	Thanks for your comment. As for "sort of final fields", I mean, the elements of methods never change once they are set. My interpretation of the new code is the same as what you wrote above. But I am not sure that the Method object contained in methods[index] is up-to-date when methods[index] is non null. Is this guaranteed by the memory model? Since a reader thread may not synchronize before reading methods[index], it may see that methods[index] is non null but the Method object in methods[index] contains older values than the ones visible to the writer thread, which sets methods[index]. Any advice?	1	1	https://issues.jboss.org/browse/JASSIST-163
8400	JASSIST-163	Hi Nikita, > as typically all properties of a proxied class are proxied, and, over the lifecycle of JVM, will be invoked. Is this true in general? Not only in your particular use-cases?	1	1	https://issues.jboss.org/browse/JASSIST-163

#	Requirement	Questions	Answered	Successfully Answered	URL
8401	JASSIST-163	{quote} My interpretation of the new code is the same as what you wrote above. But I am not sure that the Method object contained in methods[index] is up-to-date when methods[index] is non null. Is this guaranteed by the memory model? Since a reader thread may not synchronize before reading methods[index], it may see that methods[index] is non null but the Method object in methods[index] contains older values than the ones visible to the writer thread, which sets methods[index]. {quote} The Method object probably came from Class.getDeclaredMethods() or maybe Class.getDeclaredMethod(). Since those calls are thread safe, the entire Method object contents should be in main working memory. {quote} Count me biased, but I maintain the approach of eagerly initializing static methods[] - submitted above is most practical: {quote} Would it be possible to support both lazy and eager initialization?	1	1	https://issues.jboss.org/browse/JASSIST-163
8402	JASSIST-163	bq. The Method object probably came from Class.getDeclaredMethods() or maybe Class.getDeclaredMethod(). Since those calls are thread safe, the entire Method object contents should be in main working memory. Yes, but being thread safe means that a thread doing synchronization (or accessing a volatile field) can safely use the method according to the memory model of the JVM. The memory model does not guarantee anything about a thread doing no synchronization. Is this correct? We can expect that some thread safety is preserved on real hardware for such a thread that does no synchronization, though. bq. Would it be possible to support both lazy and eager initialization? This is a good idea.	1	1	https://issues.jboss.org/browse/JASSIST-163
8403	JASSIST-163	Good to hear that "eager initialization" is a good idea. :) Is anyone working on adding eager initialization?	1	1	https://issues.jboss.org/browse/JASSIST-163
8404	JASSIST-163	There were a couple of commits, prior to most of the discussion here. What was the resolution? Can this be re-opened?	1	1	https://issues.jboss.org/browse/JASSIST-163
8405	JASSIST-163	I planed the GA release early this week but a problem was found in JASSIST-160. The released date will be late this week or early next week. However, the current snapshot does not contain the code for "eager" initialization. Can I extend the release date to merge the code?	1	1	https://issues.jboss.org/browse/JASSIST-163
8406	JASSIST-163	I thought that you were going to leave lazy initialization in and also add eager initialization. So, that existing code could continue to run as it used to (with lazy initialization) and with code changes, applications could switch to eager initialization. Is that possible?	1	1	https://issues.jboss.org/browse/JASSIST-163
8407	JASSIST-163	This is a bug of org.hibernate.bytecode.internal.javassist.BulkAccessorFactory.java (Javassist is no guilty!) I wrote a fix but I haven't tested it yet. Can you check BulkAccessorFactory.java.diff?	1	1	https://issues.jboss.org/browse/JASSIST-163
8408	JASSIST-163	I have updated the diff file. Could you try BulkAccessorFactory.java.diff2?	1	1	https://issues.jboss.org/browse/JASSIST-163
8409	JASSIST-163	Would it be possible to also have a release that targets Java 6? The 3.17.0-GA release appears to target Java 7.	1	1	https://issues.jboss.org/browse/JASSIST-163
8410	JASSIST-163	I see. So, which kind of binary do you think we should upload to maven, Java 6 or 7? I'm not an expert of maven...	0	0	https://issues.jboss.org/browse/JASSIST-163
8411	JASSIST-163	Is there a reason that the source and target were changed to 1.7 in the following changeset?: https://source.jboss.org/changelog/Javassist?cs=678	0	0	https://issues.jboss.org/browse/JASSIST-163
8412	JASSIST-163	I'm not sure why, but only when we are using 3.17.1, we experience a new Hibernate test regression https://hibernate.onjira.com/browse/HHH-7850. The generated test class disassembled is here http://pastie.org/5494599 and failure exception is here http://pastie.org/5494609 It seems obvious to me that we need to fix this in the Hibernate [BulkAccessorFactory|https://github.com/hibernate/hibernate-orm/blob/4.1/hibernate-core/src/main/java/org/hibernate/bytecode/internal/javassist/BulkAccessorFactory.java] class by generating the StackMapTable (as mentioned in the comments of the above jira). Why do we need to change BulkAccessorFactory with Javassist 3.17.1 but not 3.15.0?	1	1	https://issues.jboss.org/browse/JASSIST-163
8413	JASSIST-150	This is a worst case. It will however illustrate the problem. Running this takes 32s per test on my machine. Running it will a small set of changes to the code and it takes only 0.32s.	1	1	https://issues.jboss.org/browse/JASSIST-150
8414	JASSIST-150	Oh, does your patch improve the speed?	1	1	https://issues.jboss.org/browse/JASSIST-150
8415	JASSIST-150	Because Bruce's patch changes the behavior of ClassPool and the change affects all the behavior of the whole Javassist, I have modified only the compiler with the same idea. Bruce, can you kindly check this fix with your own junit tests? I will release a new version within several days. Thanks.	1	1	https://issues.jboss.org/browse/JASSIST-150
8416	JASSIST-150	Hello, I just checked out trunk (rev 626) and 'mvn clean install' periodically fails testJIRA150b (seemingly 9 out of 10 times). Is that expected? {noformat} ---------- ---------- Test set: javassist.JvstTest ---------- ---------- Tests run: 295, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 6.902 sec <<< FAILURE! testJIRA150b(javassist.JvstTest4) Time elapsed: 1.807 sec <<< FAILURE! junit.framework.AssertionFailedError at junit.framework.Assert.fail(Assert.java:47) at junit.framework.Assert.assertTrue(Assert.java:20) at junit.framework.Assert.assertTrue(Assert.java:27) at javassist.JvstTest4.testJIRA150b(JvstTest4.java:688) {noformat}	1	0	https://issues.jboss.org/browse/JASSIST-150
8417	JASSIST-150	The results really depends on your machine. What machine are you using to run the test? Maybe I should remove this performance test?	0	0	https://issues.jboss.org/browse/JASSIST-150
8418	JASSIST-144	I think I found a solution. I introduced Proxy, a super interface of ProxyObject. Proxy only provides setHandler(). Then I modified Javassist so that a proxy object will implement the Proxy interface instead of ProxyObject *if* the original object (= the super class of the proxy) has getHandler(). If not, a proxy object will implement ProxyObject and getHandler() as before. This implies slight API incompatibility but its impacts would be minimized. Can anyone check Hibernate with the latest snapshot of Javassist?	1	1	https://issues.jboss.org/browse/JASSIST-144
8419	JASSIST-144	Shigeru, sorry I missed the earlier notifications... I am not following what the expected fix is here. As I understand it, the proxy object still requires a method setHandler to inject the Javassist handler object. If the class being proxied also defines a setHandler method we still have this problem. What am I missing?	1	0	https://issues.jboss.org/browse/JASSIST-144
8420	JASSIST-134	The releases from 3.9 to 3.12 were supposed to be copied from the JBoss repo over to the Sonatype repo and, from there, to Maven Central when we started publishing to Sonatype in August. However, this appears not to have happened. I'll check with the Sonatype admins to see if this was just forgotten about or if there was a problem which meant they could not be copied. In the latter case the only options are either to obtain them from the JBoss public repo or to upgrade to the versions which are in Maven Central (3.13 and 3.14). Note that the shift to Sonatype/Maven Central required changing the groupid from javassist to org.javassist.	0	0	https://issues.jboss.org/browse/JASSIST-134
8421	JASSIST-126	What is the status of this request?	0	0	https://issues.jboss.org/browse/JASSIST-126
8422	JASSIST-126	The bundle information is not in the manifest of the the 3.16.1-GA jar on maven central. Interestingly, the POM inside this jar has the maven-bundle-plugin configuration but it does not appear to have been executed? Is there another version I need to pick up with the OSGi bundle information? If not, can you roll a new release with this bundle information in the Manifest?	0	0	https://issues.jboss.org/browse/JASSIST-126
8423	JASSIST-126	Is there an ETA for 3.17.0? I'm running into the same issue. :(	1	0	https://issues.jboss.org/browse/JASSIST-126

#	Requirement	Questions	Answered	Successfully Answered	URL
8424	JASSIST-96	I am also happy to publish Javassist to Maven central. As for b), how do I do that? Is just sending an email to Pete enough?	1	1	https://issues.jboss.org/browse/JASSIST-96
8425	JASSIST-96	Andrew/Chiba, we've been asked to not use javassist as the group id, but instead use a qualfied name. For example, org.jboss.javassist. WDYT about that?	0	0	https://issues.jboss.org/browse/JASSIST-96
8426	JASSIST-64	You can choose either Javassist or ASM for Hibernate EntityManager. So if you choose Javassist (and turn ASM off) for Hibernate, you should be able to use ASM 3.1 for Jersey without any problems.	1	1	https://issues.jboss.org/browse/JASSIST-64
8427	JASSIST-62	Hi Chiba, Have you had a chance to think about this one yet? Will $type be available in insertBefore in the future? Best regards, Yanic	1	1	https://issues.jboss.org/browse/JASSIST-62
8428	JASSIST-62	I updated insertBefore() so that $type will be available. However, why do you need $type? CtMethod.getReturnType() returns the return type. You can use this to change the source code give to insertBefore().	1	1	https://issues.jboss.org/browse/JASSIST-62
8429	JASSIST-48	Do we really need such a fine grained type hierarchy?	0	0	https://issues.jboss.org/browse/JASSIST-48
8430	JASSIST-20	This problem is that a ObjectInputStream must recreate a proxy class when it deserializes a proxy object originally created by Javassist. ObjectInputStream deals with a proxy class by the Reflection API in a special way and I could not find a trick to add a similar thing for Javassist's proxies to ObjectInputStream. Is there anybody who knows how?	1	1	https://issues.jboss.org/browse/JASSIST-20
8431	JBERET-346	org.jberet.spi.JobOperatorContext.getJobOperatorContext(JobOperatorContext.java:47) at org.jberet.testapps.common.AbstractIT.before(AbstractIT.java:72) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:50) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:47) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:24) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:325) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:78) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:57) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:290) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:71) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:288) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:58) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:268) at org.junit.runners.ParentRunner.run(ParentRunner.java:363) at org.junit.runner.JUnitCore.run(JUnitCore.java:137) at com.intellij.junit4.JUnit4IdeaTestRunner.startRunnerWithArgs(JUnit4IdeaTestRunner.java:68) at com.intellij.rt.execution.junit.IdeaTestRunner$Repeater.startRunnerWithArgs(IdeaTestRunner.java:51) at com.intellij.rt.execution.junit.JUnitStarter.prepareStreamsAndStart(JUnitStarter.java:242) at com.intellij.rt.execution.junit.JUnitStarter.main(JUnitStarter.java:70) Caused by: org.infinispan.commons.CacheException: Unable to invoke method public void org.infinispan.persistence.manager.PersistenceManagerImpl.start() on object of type PersistenceManagerImpl at org.infinispan.commons.util.SecurityActions.lambda$invokeAccessibly$0(SecurityActions.java:95) at org.infinispan.commons.util.SecurityActions.doPrivileged(SecurityActions.java:83) at org.infinispan.commons.util.SecurityActions.invokeAccessibly(SecurityActions.java:88) at org.infinispan.commons.util.ReflectionUtil.invokeAccessibly(ReflectionUtil.java:165) at org.infinispan.factories.AbstractComponentRegistry$PrioritizedMethod.invoke(AbstractComponentRegistry.java:869) at org.infinispan.factories.AbstractComponentRegistry.invokeStartMethods(AbstractComponentRegistry.java:635) at org.infinispan.factories.AbstractComponentRegistry.internalStart(AbstractComponentRegistry.java:624) at org.infinispan.factories.AbstractComponentRegistry.start(AbstractComponentRegistry.java:549) at org.infinispan.factories.ComponentRegistry.start(ComponentRegistry.java:231) at org.infinispan.cache.impl.CacheImpl.start(CacheImpl.java:803) at org.infinispan.cache.impl.AbstractDelegatingCache.start(AbstractDelegatingCache.java:411) at org.infinispan.manager.DefaultCacheManager.wireAndStartCache(DefaultCacheManager.java:652) at	0	0	https://issues.jboss.org/browse/JBERET-346
8432	JBERET-274	Q: Some JDK APIs have disappeared… where did they go? There are two major factors which contribute to the presence of JDK APIs. The first factor is related to dependency resolution: if you are running a JDK module which does not depend on the API in question, it will not "see" it and will fail. The second factor is more complicated. Java 9 does not make all modules available by default. Certain modules, while bundled with the JDK, cannot be loaded unless a command-line argument is given to make the module available. The argument is "--add-modules module.name". The list of module names which must be added in this way is as follows: java.activation java.annotations.common java.corba java.transaction java.xml.bind java.xml.ws These modules may be upgraded with replacements by adding Jigsaw-style modular JARs to --upgrade-module-path. References http://mail.openjdk.java.net/pipermail/jigsaw-dev/2016-April/007209.html	1	1	https://issues.jboss.org/browse/JBERET-274
8433	JBERET-263	Thanks for your comments. I didn't test the CLI after the WFLY-7000 being fixed. I'll try tomorrow. I didn't totally understand the James Perkins' comment about ArchiveXmlLoader? Is there anything I can do?	1	1	https://issues.jboss.org/browse/JBERET-263
8434	JBERET-263	How do you run your integration tests? Do they start-up WildFly then just deploy the application and hit some endpoints? In your WAR, how do you start batch deployments? If you can, could you post your implementation of {{org.jberet.spi.JobXmlResolver}}? My test one was pretty basic, though maybe using JBoss Modules a better one could be implemented. {code:java} public class BatchXmlResolver implements JobXmlResolver { @Override public InputStream resolveJobXml(final String jobXml, final ClassLoader classLoader) throws IOException { return BatchXmlResolver.class.getResourceAsStream("/" + JobXmlResolver.DEFAULT_PATH + jobXml); } @Override public Collection<String> getJobXmlNames(final ClassLoader classLoader) { return Collections.singleton("simple-module.xml"); } @Override public String resolveJobName(final String jobXml, final ClassLoader classLoader) { return "simple-module"; } } {code}	1	1	https://issues.jboss.org/browse/JBERET-263

#	Requirement	Questions	Answered	Successfully Answered	URL
8435	JBERET-263	{quote} How do you run your integration tests? Do they start-up WildFly then just deploy the application and hit some endpoints? {quote} I've a Maven module called {{integration-test}} and all the integration tests related to WildFly are written in this Maven module. Yes, the [Maven Module "integrationtest" starts up the WildFly server | https://github.com/mincong-h/gsoc-hsearch/blob/JBERET-263/integrationtest/wildfly/pom.xml#L153-L160]. Once WF is started, arquillian deploys the integration test, e.g. [RestartIT.war | https://github.com/mincong-h/gsoc-hsearch/blob/JBERET-263/integrationtest/wildfly/src/test/java/org/hibernate/search/jsr352/RestartIT.java] to this managed container. Once every integration test is finished, the server will be stopped. {quote} In your WAR, how do you start batch deployments? {quote} The batch job {{massIndex}} is not part of the WAR deployment, it was installed previously by a Maven module called [{{module}} | https://github.com/mincong-h/gsoc-hsearch/tree/JBERET-263/modules]. Then the WAR deployment for integration test looks up this batch job {{massIndex}} inside the installed module. {quote} If you can, could you post your implementation of {{org.jberet.spi.JobXmlResolver}}? {quote} I don't have my own implementation of JobXmlResolver, instead, I used {{org.jberet.tools.MetaInfBatchJobsJobXmlResolver}} in my core module, you can check [here | https://github.com/mincong-h/gsoc-hsearch/blob/JBERET-263/core/src/main/resources/META-INF/services/org.jberet.spi.JobXmlResolver]: {code:title=core/src/main/resources/META-INF/services/org.jberet.spi.JobXmlResolver} org.jberet.tools.MetaInfBatchJobsJobXmlResolver {code} By the way, do I also need to implement the SPI for the integration tests? I didn't do that. I don't fully understand how your attachment {{ModuleBatchXmlResolver.java}}, but I got the idea. h4. Clarification about this issue Actually, my main question is how to start a batch job installed in a module using the CLI directly? There should have no requirement on deployment. Since the job belongs to a module but not a deployment, user should be able to start it any time using CLI or GUI without deployment. Is it possible with JSR 352? The scenario could be system maintenance plans (back up, clean up) or indexing recent inserted entities for search feature (what I'm doing). These can be scheduled at any off-hour but also be launched by the user manually using CLI / GUI.	1	1	https://issues.jboss.org/browse/JBERET-263
8436	JBERET-263	{quote} Actually, my main question is how to start a batch job installed in a module using the CLI directly? There should have no requirement on deployment. Since the job belongs to a module but not a deployment, user should be able to start it any time using CLI or GUI without deployment. Is it possible with JSR 352? The scenario could be system maintenance plans (back up, clean up) or indexing recent inserted entities for search feature (what I'm doing). These can be scheduled at any off-hour but also be launched by the user manually using CLI / GUI. {quote} The Java EE specifications are based around deployments. There is no generalized specification for CLI or management operations that would start a batch job. Modules are not scanned to see if they are some kind of EE module. In fact there is no concept of an EE module in terms of a JBoss Module. It's not only a JSR 352 issue. CDI won't generally work in a module neither with JPA. The resources can't be injected without having the context of a deployment. This means that batch jobs in a module will require a deployment just as CDI or JPA will. Consider an {{EntityManagerFactory}}. It is created explicitly for a deployment and can't be shared across other deployments.	1	1	https://issues.jboss.org/browse/JBERET-263
8437	JBERET-209	consumer config: http://kafka.apache.org/090/javadoc/org/apache/kafka/clients/consumer/ConsumerConfig.html producer config: http://kafka.apache.org/090/javadoc/org/apache/kafka/clients/producer/ProducerConfig.html both Kafka consumer and producer have long list of config properties. How to pass them to batch ItemReader and ItemWriter? We can consider the following options: * one config property corresponds to one @Inject @BatchProperty field in ItemReader or ItemWriter class. ** Pro: easy to see all config properties ** Con: when Kafka adds, removes, or changes config properties in the future, we will need to change ItemReader or ItemWriter classes ** Con: too many @BatchProperty injections * have a single @Inject @BatchProperty configPropertiesFile; and have the application create a properties file to hold all config properties. ** Pro: @BatchProperty can be the same for both ItemReader and ItemWriter classes; any future changes in Kafka config properties does not affect batch injection properties; ** Con: these Kafka properties reside in a separate file from the main job.xml file * use CDI to inject a java.util.Properties instance ** Pro: flexible ** Con: users need to write a CDI producer method * mix of the above 3 options.	0	0	https://issues.jboss.org/browse/JBERET-209
8438	JBERET-209	Looks like Kafka 0.8.2.2 server doesn't work with 0.9.0 client. Got the following exception: {code} java.lang.AssertionError: assertion failed: Version 2 is invalid for OffsetCommitRequest. Valid versions are 0 or 1. at scala.Predef$.assert(Predef.scala:179) at kafka.api.OffsetCommitRequest$.readFrom(OffsetCommitRequest.scala:35) at kafka.api.RequestKeys$$anonfun$9.apply(RequestKeys.scala:47) at kafka.api.RequestKeys$$anonfun$9.apply(RequestKeys.scala:47) at kafka.network.RequestChannel$Request.<init>(RequestChannel.scala:50) at kafka.network.Processor.read(SocketServer.scala:450) at kafka.network.Processor.run(SocketServer.scala:340) at java.lang.Thread.run(Thread.java:745) {code}	0	0	https://issues.jboss.org/browse/JBERET-209
8439	JBERET-173	Can you give more details about your use case, just to help us better understand your requirement? What is the error when trying to include batchee cdi scopes?	1	1	https://issues.jboss.org/browse/JBERET-173
8440	JBERET-139	Are there any vulnerable situation of loading job XMLs that created by a remote or malicious user? Of course I know security is important than gaining little convenience so probably we should go the opposite direction that stop setting XMLResolver and invoke xmlInputFactory.setProperty(XMLInputFactory.SUPPORT_DTD, false); . it's safest and reasonable way. I think the best solution is making configurable (enable/disable) that feature via subsystem configuration (standalone.xml or something) but it might be too much.	0	0	https://issues.jboss.org/browse/JBERET-139
8441	ENTMQ-1694	I pushed a unit test that demonstrates a way to get this to work. The idea is to configure a static bridge with dynamicOnly=false and use the BrokerAdminView (JMX) to create a durable sub on the local or forwarding broker. On a restart/recycle of the networkConnector, because dynamicOnly=false it will search in the message store for durable subs that match it "well known" clientId and sub prefix and will start forwarding them. Note: this bridge does no dynamic forwarding, so it does not listen for advisory messages at all. the staticBridge=true attribute. https://github.com/apache/activemq/commit/3c0a4d960ee99483c9160d0b6b80c27abd8f3319	0	0	https://issues.jboss.org/browse/ENTMQ-1694
8442	ENTMQ-1498	the activemq-osgi bundle has a javax.jms [1.1,3] import - so can work with either jms 1.1 or jms 2	0	0	https://issues.jboss.org/browse/ENTMQ-1498
8443	ENTMQ-1309	Is it documentation bug, or this attribute should work?	0	0	https://issues.jboss.org/browse/ENTMQ-1309
8444	ENTMQ-1093	Hi Gary, What do we need to set on configuration side? <virtualDestinations> <virtualTopic name=">" ... concurrentSend="true" Is " use implicit broker transaction for virtual topic fanout a" done automatically? Sorry if I've mis-read anything, Susan	1	1	https://issues.jboss.org/browse/ENTMQ-1093

#	Requirement	Questions	Answered	Successfully Answered	URL
8445	ENTMQ-1093	Hi Gary, Thanks for the information. I ran a few quick tests just to make sure everything was in order on my side this time: <virtualTopic name="vt.>" prefix="vt.perfTest.>." selectorAware="false" concurrentSend="true"/> Is invalid in build 19 [1]. Looking at your unit test - it seems to be in the right place? [2] [1] {quote} Caused by: org.springframework.beans.NotWritablePropertyException: Invalid property 'concurrentSend' of bean class [org.apache.activemq.broker.region.virtual.VirtualTopic]: Bean property 'concurrentSend' is not writable or has an invalid setter method. Does the parameter type of the setter match the return type of the getter? {quote} Apologies if I am missing something obvious ... [2] https://fisheye6.atlassian.com/changelog/activemq-git?cs=340728f2d1eb4c1aa74f4002c5b16bf7a70c57b7	0	0	https://issues.jboss.org/browse/ENTMQ-1093
8446	ENTMQ-1043	The following seem to be current options for message group <-> consumer mapping within the broker: SimpleMessageGroupMap CachedMessageGroupMap MessageGroupHashBucket Documentation doesn't seem to exist on these. When should someone use one and not the other? What is default if nothing specified? If someone has use-case where can't guarantee clients will always close groups, and plan to have lots and lots of (unbounded) number of dynamically created groups, what then?	0	0	https://issues.jboss.org/browse/ENTMQ-1043
8447	ENTMQ-948	Hi, I don't have an access to this KB article. Can someone attach it to this issue or look at the privileges for dejan.bosanac user?	1	1	https://issues.jboss.org/browse/ENTMQ-948
8448	ENTMQ-866	Hi Bill, Are you able to leverage producer flow control for this or perhaps implement a camel throttler?	1	1	https://issues.jboss.org/browse/ENTMQ-866
8449	ENTMQ-866	Hi Susan, Should there be a JBAMQPL jira for this issue?	1	0	https://issues.jboss.org/browse/ENTMQ-866
8450	ENTMQ-866	Hi Bill, Please re-open with additional information if you still need this as per last comments. thanks, Susan	1	1	https://issues.jboss.org/browse/ENTMQ-866
8451	ENTMQ-835	Hi Pat, I just changed the text of that admonition to: {quote} *Important:* The LevelDB store is supported only on the Red Hat Enterprise Linux (RHEL) platform (for the RHEL versions and configurations specified in Supported Configurations). If you want to store the LevelDB files on a distributed file system, the following configuration is also supported: the broker itself and the LevelDB adapter plugin run on a supported RHEL platform, while the LevelDB files are stored on an NFSv4 file system (the O/S of the server that hosts the NFSv4 file system is irrelevant). Note that the distributed file system must be NFSv4. For all other JBoss A-MQ-supported platforms, the LevelDB store is provided for technical preview only. {quote} How does that sound? Cheers, Fintan	1	1	https://issues.jboss.org/browse/ENTMQ-835
8452	ENTMQ-829	that makes sense. guess it could also be an info level message rather than warn?	0	0	https://issues.jboss.org/browse/ENTMQ-829
8453	ENTMQ-795	Hi Gary, A quick follow up on the above as we noted "So CACHE_SESSION or CACHE_CONNECTION will work fine." from 23/Dec/14 7:34 AM. Is CACHE_SESSION really ok? {code} static int CACHE_AUTO Constant that indicates automatic choice of an appropriate caching level (depending on the transaction management strategy). static int CACHE_CONNECTION Constant that indicates to cache a shared JMS Connection. static int CACHE_CONSUMER Constant that indicates to cache a shared JMS Connection and a JMS Session for each listener thread, as well as a JMS MessageConsumer for each listener thread. static int CACHE_NONE Constant that indicates to cache no JMS resources at all. static int CACHE_SESSION Constant that indicates to cache a shared JMS Connection and a JMS Session for each listener thread. {code} I just wanted to double check as CACHE_CONNECTION and CACHE_NONE make sense to me here but perhaps CACHE_SESSION has a nuance to it.	0	0	https://issues.jboss.org/browse/ENTMQ-795
8454	ENTMQ-789	From the Salesforce reference, it appears that this issue is actually for 6.1, not 6.0 (which makes more sense, because AMQP was tech preview, not supported, in 6.0). So, I am going to change the 'Affects version/s' field to 6.1.	0	0	https://issues.jboss.org/browse/ENTMQ-789
8455	ENTMQ-731	Hello [~tabish121], supplying bad credentials will result in "Not authorized to connect (5)" instead of bad username or password (4). In the fuse 6.2 ER3 log it is clear that bad credentials will raise java.lang.SecurityException instead of javax.security.auth.login.CredentialException. {code} org.apache.activemq.activemq-osgi - 5.11.0.redhat-620099 | Failed to add Connection ID:avano-ThinkPad-T430s-59300-1430387201283-6:11 due to java.lang.SecurityException: User name [admin] or password is invalid. {code} As far as I see, there is no CredentialException used in the code. Could you please take a look? Thanks	1	1	https://issues.jboss.org/browse/ENTMQ-731
8456	ENTMQ-673	that is not good. the master should shutdown within the keepAlivePeriod if the lock file is deleted. Is the lockKeepAlivePeriod specified in configuration? {code}<kahaDB ... lockKeepAlivePeriod="100" />{code}	1	1	https://issues.jboss.org/browse/ENTMQ-673
8457	ENTMQ-673	Jakub/Gary: Does this mean that it will not be fixed for GA?	0	0	https://issues.jboss.org/browse/ENTMQ-673
8458	ENTMQ-673	What about some timeout before start of slave? E.g. When lock is deleted and slave grabs the lock it could first wait for period of time T (it would be holding the lock but waiting - leaving time for previous master to shutown). After that period T the slave would become master. Would it be possible? If T will be bigger than the biggest keepAlivePeriod of all brokers. Then It would eliminate having two brokers running and keepAlivePeriod could be larger number.	0	0	https://issues.jboss.org/browse/ENTMQ-673

#	Requirement	Questions	Answered	Successfully Answered	URL
8459	ENTMQ-544	https://access.redhat.com/products/FuseSource/supportedPlatforms/ This should be changed to point to Red Hat JBoss Fuse - https://access.redhat.com/articles/310603 ------ Page 12 It's tedious, but shouldn't this text: Red Hat JBoss Fuse Tooling The JBoss Fuse Tooling provides a set of developer tools for developing Red Hat JBoss Fuse applications within Red Hat JBoss Development Studio Be changed to: Red Hat JBoss Fuse Tooling The Red Hat JBoss Fuse Tooling provides a set of developer tools for developing Red Hat JBoss Fuse applications within Red Hat JBoss Development Studio ------ Page 13 This text: Accessing dependencies through the Internet—Maven has the ability to download dependencies automatically through the Internet, by searching through known Maven repositories. This implies that you must have access to the Internet, when building with Maven. Should be changed to: Accessing dependencies through the Internet—Maven has the ability to download dependencies automatically through the Internet, by searching through known Maven repositories. This implies that you must either have access to the Internet, or local copies of the required Maven repositories, when building with Maven. ------ Page 14 This text: For larger Maven projects, there are two special kind of POM files that you might also need: Should be changed to: For larger Maven projects, there are two special kinds of POM files that you might also need: ------ Page 15 This text: Should be changed to: The following kinds of metadata are accessible to administrators, who can use them to customize or change the behavior of bundles at run time: ------ Page 19 This text: The dependency element declares a dependency on the Maven artifact with coordinates org.apache.camel:camel-blueprint:6.1.0.redhat-379. You can add as many dependency elements as you like inside the dependencies element. Should be changed to: The dependency element declares a dependency on the Maven artifact with coordinates (for example: org.apache.camel:camel-blueprint:6.1.0.redhat-379). You can add as many dependency elements as you like inside the dependencies element. ------ Pages 27-32 mvn archetype:generate -DarchetypeGroupId=org.apache.servicemix.tooling -DarchetypeArtifactId=servicemix-cxf-code-first-osgi-bundle -DarchetypeVersion=2013.01.0.redhat-610379 -DgroupId=org.fusesource.example -DartifactId=cxf-basic -Dversion=1.0-SNAPSHOT Fabric8:karaf@root> features:install cxf Refreshing bundles org.apache.cxf.cxf-rt-transports-http (132), org.apache.servicemix.bundles.spring-context-support (173), org.apache.cxf.cxf-rt-rs-extension-providers (135) Fabric8:karaf@root> features:install cxf-http-jetty Fabric8:karaf@root> features:install cxf-jaxws Fabric8:karaf@root> features:install cxf Fabric8:karaf@root> install mvn:org.fusesource.example/cxf-basic/1.0-SNAPSHOT Bundle ID: 297 Fabric8:karaf@root> start 297 Error executing command: Error starting bundles: Unable to start bundle 297: Unresolved constraint in bundle cxf-basic [297]: Unable to resolve 297.0: missing requirement [297.0] osgi.wiring.package; (&(osgi.wiring.package=javax.jws)(version>=2.0.0)(!(version>=3.0.0))) Starting the bundle fails^ ------ Page 38 Test the route with the WS client The cxf-basic project includes a simple WS client, which you can use to test the deployed route and Web service. In a command prompt, navigate to the cxf-basic directory and run the simple WS client as follows: cd ../cxf-basic mvn -Pclient -Dexec.args="http://localhost:8282/cxf/PersonServiceCF" The client fails with this exception: INFO: Creating Service {http://example.fusesource.org/}PersonService from class org.fusesource.example.Person Invoking getPerson... Sep 04, 2014 12:52:12 PM org.apache.cxf.phase.PhaseInterceptorChain doDefaultLogging WARNING: Interceptor for {http://example.fusesource.org/}PersonService#{http://example.fusesource.org/}GetPerson has thrown exception, unwinding now org.apache.cxf.binding.soap.SoapFault: Error reading XMLStreamReader. at	1	1	https://issues.jboss.org/browse/ENTMQ-544
8460	ENTMQ-544	Hi Jane - just finished the review - can you tell me where/from whom you received the examples? I may have found a couple of bugs in the DR2 build. ;-)	1	1	https://issues.jboss.org/browse/ENTMQ-544
8461	ENTMQ-527	Hi Fintan, The fix mentioned did not appear to make it into the docs. Could you have a look? Thanks Pat	1	1	https://issues.jboss.org/browse/ENTMQ-527
8462	ENTMQ-463	JBAMQPL-139 is created to track this request. Closing this issue.	0	0	https://issues.jboss.org/browse/ENTMQ-463
8463	ENTMQ-416	Do you plan to do another release at Apache before we ship A-MQ 6.1 so that this gets picked up in 6.1?	0	0	https://issues.jboss.org/browse/ENTMQ-416
8464	ENTMQ-406	Re-opening this case to further explore larger maxPageSizes. 1. Added some logic to AbstractStoreCursor to dump the time to doFillBatch in 5.8 patches branch: protected final synchronized void fillBatch() { //LOG.trace("{} - fillBatch", this); if (batchResetNeeded) { resetSize(); setMaxBatchSize(Math.min(regionDestination.getMaxPageSize(), size)); resetBatch(); this.batchResetNeeded = false; } if (this.batchList.isEmpty() && this.storeHasMessages && this.size >0) { try { long start = System.currentTimeMillis(); doFillBatch(); System.out.println(" **** doFillBatch completed in " + (System.currentTimeMillis()-start)); } catch (Exception e) { LOG.error(this + " - Failed to fill batch", e); throw new RuntimeException(e); } this.storeHasMessages = !this.batchList.isEmpty() || !hadSpace; } With a MPS of 1000, it's quick 300-500ms. With 10K MPS, it goes to about 18 seconds. My message size is 5KB. As these page loads are slower than smaller ones, can it block consumers from committing and working already loaded messages from the cursor? We can see with similar timestamps in consumer around session.commit that commit times degrade with larger MPS. Re-ran the original use case on 5.9 for this ticket and will add results shortly.	1	1	https://issues.jboss.org/browse/ENTMQ-406
8465	ENTMQ-364	This is an interesting request. A-MQ is a preconfigured instance of karaf. One way to approach this is to use a different karaf child instance for a different configuration - http://karaf.apache.org/manual/latest-2.3.x/users-guide/child-instances.html so replace etc/activemq.xml with custom xml for the child. An alternative, and closer to the request approach, may be to modify the precanned confit to use a system property to reference the xml config. currently in etc/org.fusesource.mq.fabric.server-default.cfg (the a-mq instance config file) we have:{code} broker-name=amq openwire-port=${activemq.port} standalone=true data=${karaf.data}/amq config=${karaf.base}/etc/activemq.xml{code} We could replace:{code}config=${karaf.base}/etc/activemq.xml{code} with: {code} config=${karaf.base}/etc/${amq-xml}{code} and have a default system property config for {code}amq-xml=activemq.xml{code} Then a user could do: $ cp ~/myCustomBroker.xml ${karaf.base}/etc $ amq -Damq-xml=myCustomBroker.xml How does something like that sit?	0	0	https://issues.jboss.org/browse/ENTMQ-364
8466	ENTMQ-347	Would your proposed workaround work with springbeans?	1	1	https://issues.jboss.org/browse/ENTMQ-347
8467	ENTMQ-311	Where can we get this documentation for JBoss-AM 6.1? Thanks.	1	0	https://issues.jboss.org/browse/ENTMQ-311
8468	ENTMQ-256	Is this still an issue for 6.1?	1	1	https://issues.jboss.org/browse/ENTMQ-256
8469	ENTMQ-251	Hi Freeman, in Fuse you should use mq-web-console (branded version of console) and I think it should use a proper admin/admin credentials by default. You can try using mq-web-console profile and see how it goes. I'd keep karaf/karaf in activemq-web-console as it's easier to install it in plain karaf. Can you try it out and let me know how it works.	0	0	https://issues.jboss.org/browse/ENTMQ-251
8470	ENTMQ-155	-015 has it in extras - may need to investigate reuse of the code in the mq-client bundle for examples - can the bundle be found when in extras?	0	0	https://issues.jboss.org/browse/ENTMQ-155
8471	ENTMQ-145	@Fintan - do you have something you can share?	0	0	https://issues.jboss.org/browse/ENTMQ-145
8472	ENTMQ-83	The enterprise product is not targeted at new users. There is the simple command line producer/consumer that allows a user to validate an install. see: https://fusesource.com/issues/browse/ENTMQ-27 Is there something particular that is missing or blocking a user from progressing?	0	0	https://issues.jboss.org/browse/ENTMQ-83

#	Requirement	Questions	Answered	Successfully Answered	URL
8473	ENTMQ-44	also add a warning not to close the shell unless they want to shut down the server?	1	0	https://issues.jboss.org/browse/ENTMQ-44
8474	ENTMQ-43	Update thread is a little eager, it can update the broker while it is starting, so causing a stop during start and ugly errors:{code} FuseMQ:karaf@root> java.lang.IllegalStateException: PageFile is not loaded at org.apache.kahadb.page.PageFile.assertLoaded(PageFile.java:809) at org.apache.kahadb.page.PageFile.tx(PageFile.java:303) at org.apache.activemq.store.kahadb.KahaDBStore.getDestinations(KahaDBStore.java:952) at org.apache.activemq.store.kahadb.KahaDBPersistenceAdapter.getDestinations(KahaDBPersistenceAdapter.java:125) at org.apache.activemq.broker.region.DestinationFactoryImpl.getDestinations(DestinationFactoryImpl.java:68) at org.apache.activemq.broker.BrokerService.startVirtualConsumerDestinations(BrokerService.java:2510) at org.apache.activemq.broker.BrokerService.startDestinations(BrokerService.java:2339) at org.apache.activemq.broker.BrokerService.doStartBroker(BrokerService.java:651) at org.apache.activemq.broker.BrokerService.startBroker(BrokerService.java:642) at org.apache.activemq.broker.BrokerService.start(BrokerService.java:578) at org.fusesource.mq.fabric.ActiveMQServiceFactory$ClusteredConfiguration$$anon$3.run(ActiveMQServiceFactory.scala:307) Exception in thread "Startup for ActiveMQ Broker: root" java.lang.NullPointerException at org.fusesource.mq.fabric.ActiveMQServiceFactory$ClusteredConfiguration$$anon$3.run(ActiveMQServiceFactory.scala:334){code}	0	0	https://issues.jboss.org/browse/ENTMQ-43
8475	ENTMQ-36	To support recovery of the broker, we need some additional xml configuration on the esb. I think this is preferable to doing it automatically in the broker pid service factory because the pool may want to be configured and different applications may wish to use different resource names. Currently the pid factory just exports a simple vm connection factory, no pooling of wrapping. It is also consistent with 4.x esb, save that the broker config is excluded. maybe in a file called: *xa.xml* (can it just be dropped into the etc directory?) {code} <!-- Licensed to the Apache Software Foundation (ASF) under one or more contributor license agreements. See the NOTICE file distributed with this work for additional information regarding copyright ownership. The ASF licenses this file to You under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at http://www.apache.org/licenses/LICENSE-2.0 Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License. --> <blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0" xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.0.0" xmlns:ext="http://aries.apache.org/blueprint/xmlns/blueprint-ext/v1.0.0" xmlns:blueprint="http://www.osgi.org/xmlns/blueprint/v1.0.0"> <manifest> Export-Service = javax.jms.ConnectionFactory </manifest> <!-- resolve local transaction manager and connection factory from osgi --> <reference id="recoverableTxManager" interface="org.apache.geronimo.transaction.manager.RecoverableTransactionManager" availability="mandatory" /> <reference id="activemqConnectionFactory" interface="javax.jms.ConnectionFactory" filter="(name=default)" availability="mandatory" /> <!-- configure a recoverable xa connection pool with a 'name' that matches the 'resourceName' of the resourceManager --> <bean id="pooledXAConnectionFactory" class="org.apache.activemq.pool.JcaPooledConnectionFactory"> <property name="maxConnections" value="1" /> <property name="connectionFactory" ref="activemqConnectionFactory" /> <property name="transactionManager" ref="recoverableTxManager" /> <property name="name" value="fusemq.default" /> </bean> <!-- register the connection factory with the transaction manager to ensure the 'resourceName' indoubt transactions are recovered on startup --> <bean id="resourceManager" class="org.apache.activemq.pool.ActiveMQResourceManager" init-method="recoverResource"> <property name="transactionManager" ref="recoverableTxManager" /> <property name="connectionFactory" ref="activemqConnectionFactory" /> <property name="resourceName" value="fusemq.default" /> </bean> <!-- publish the recoverable xa connection pool for use by your application --> <service ref="pooledXAConnectionFactory" interface="javax.jms.ConnectionFactory"> <service-properties> <entry key="name" value="localhost"/> <entry key="recoverable" value="true"/> </service-properties> </service> </blueprint>{code}	1	1	https://issues.jboss.org/browse/ENTMQ-36

#	Requirement	Questions	Answered	Successfully Answered	URL
8476	ENTMQ-16	camel Refreshing bundles org.springframework.core (91), org.springframework.context.support (97) org.apache.aries.blueprint[org.apache.aries.blueprint.container.ServiceRecipe] : Error retrieving service from ServiceRecipe[name='.component-1'] org.osgi.service.blueprint.container.ComponentDefinitionException: Unable to intialize bean registrationHandler at org.apache.aries.blueprint.container.BeanRecipe.runBeanProcInit(BeanRecipe.java:640) at org.apache.aries.blueprint.container.BeanRecipe.internalCreate(BeanRecipe.java:726) at org.apache.aries.blueprint.di.AbstractRecipe.create(AbstractRecipe.java:64) at org.apache.aries.blueprint.di.RefRecipe.internalCreate(RefRecipe.java:60) at org.apache.aries.blueprint.di.AbstractRecipe.create(AbstractRecipe.java:64) at org.apache.aries.blueprint.container.BlueprintRepository.createInstances(BlueprintRepository.java:219) at org.apache.aries.blueprint.container.BlueprintRepository.createInstance(BlueprintRepository.java:198) at org.apache.aries.blueprint.container.BlueprintRepository.create(BlueprintRepository.java:137) at org.apache.aries.blueprint.container.ServiceRecipe.createRecipe(ServiceRecipe.java:370) at org.apache.aries.blueprint.container.ServiceRecipe.createService(ServiceRecipe.java:278) at org.apache.aries.blueprint.container.ServiceRecipe.internalGetService(ServiceRecipe.java:248) at org.apache.aries.blueprint.container.ServiceRecipe.getService(ServiceRecipe.java:327) at org.apache.aries.blueprint.container.ServiceRecipe$TriggerServiceFactory.getService(ServiceRecipe.java:430) at org.apache.felix.framework.ServiceRegistrationImpl.getFactoryUnchecked(ServiceRegistrationImpl.java:308) at org.apache.felix.framework.ServiceRegistrationImpl.getService(ServiceRegistrationImpl.java:219) at org.apache.felix.framework.ServiceRegistry.getService(ServiceRegistry.java:310) at org.apache.felix.framework.Felix.getService(Felix.java:3497) at org.apache.felix.framework.BundleContextImpl.getService(BundleContextImpl.java:468) at org.osgi.util.tracker.ServiceTracker.addingService(ServiceTracker.java:442) at org.osgi.util.tracker.ServiceTracker$Tracked.customizerAdding(ServiceTracker.java:896) at org.osgi.util.tracker.AbstractTracked.trackAdding(AbstractTracked.java:261) at org.osgi.util.tracker.AbstractTracked.track(AbstractTracked.java:233) at org.osgi.util.tracker.ServiceTracker$Tracked.serviceChanged(ServiceTracker.java:840) at org.apache.felix.framework.util.EventDispatcher.invokeServiceListenerCallback(EventDispatcher.java:934) at org.apache.felix.framework.util.EventDispatcher.fireEventImmediately(EventDispatcher.java:795) at org.apache.felix.framework.util.EventDispatcher.fireServiceEvent(EventDispatcher.java:544) at org.apache.felix.framework.Felix.fireServiceEvent(Felix.java:4337) at org.apache.felix.framework.Felix.registerService(Felix.java:3352) at org.apache.felix.framework.BundleContextImpl.registerService(BundleContextImpl.java:346) at org.apache.aries.blueprint.container.BlueprintContainerImpl.registerService(BlueprintContainerImpl.java:410) at org.apache.aries.blueprint.container.ServiceRecipe.register(ServiceRecipe.java:184) at org.apache.aries.blueprint.container.BlueprintContainerImpl.registerServices(BlueprintContainerImpl.java:653) at	1	1	https://issues.jboss.org/browse/ENTMQ-16
8477	JBADMCON-175	Which version of the console does this refer to?	0	0	https://issues.jboss.org/browse/JBADMCON-175
8478	JBADMCON-170	Hi All, Can we have an update on this please. Thanks in Advance. Surendra	1	0	https://issues.jboss.org/browse/JBADMCON-170
8479	JBADMCON-168	Shouldn't this be logged in AS7 JIRA? I don't think this JIRA is used for the admin console in AS7.	0	0	https://issues.jboss.org/browse/JBADMCON-168
8480	JBADMCON-134	I ran into a problem updating the mbean attributes of the existing destinations. Taking queue/DLQ as an example. The mbean name defined in jbossmq-service.xml is: jboss.mq.destination:name=DLQ,service=Queue (name 1) but the name reported by jmx-console is: jboss.mq.destination:service=Queue,name=DLQ (name 2) In order to locate the mbean through DeploymentService, name 2 must be provided. But updating of the xml file failed because name 2 cannot be found in the xml file! The quick and easy fix is to change name 1in the xml file to name 2, but we'll have to do this for each mbean as we discover them. For now, the definition of all the default destinations must be fixed (alghough most of them are commented out). Is this an acceptable solution, or is there a better one? There're also problems updating the existing data sources. Unlike destinations, data sources are NOT defined within an <mbean>...</mbean> structure, and the way the attributes are defined is different from the mbean attributes. That means we cannot use the mbean-update template to perform the update. We may be able to come up with a different template for datasource update, but it is going to take some time. Do we want to do it now?	0	0	https://issues.jboss.org/browse/JBADMCON-134

#	Requirement	Questions	Answered	Successfully Answered	URL
8481	JBADMCON-129	My EAR contains an EJB3 archive, and a WAR with a Tapestry App. All Tapestry JARs are in the WEB-INF/lib directory. 2005-10-26 21:08:19,937 DEBUG [org.jboss.deployment.MainDeployer] create step for deployment file:/C:/java/jboss-5.0.0alpha/server/default/tmp/deploy/tmp22853techconf.ear-contents/techconf.war 2005-10-26 21:08:19,937 DEBUG [org.jboss.web.tomcat.tc5.Tomcat5] create, techconf.war 2005-10-26 21:08:19,937 DEBUG [org.jboss.system.ServiceController] Creating service jboss.web.deployment:war=techconf.war,id=1444246046 2005-10-26 21:08:19,937 DEBUG [org.jboss.system.ServiceController] adding depends in ServiceController.register: [] 2005-10-26 21:08:19,937 DEBUG [org.jboss.web.WebModule] Creating jboss.web.deployment:war=techconf.war,id=1444246046 2005-10-26 21:08:19,937 DEBUG [org.jboss.web.WebModule] Created jboss.web.deployment:war=techconf.war,id=1444246046 2005-10-26 21:08:19,937 DEBUG [org.jboss.system.ServiceController] Creating dependent components for: jboss.web.deployment:war=techconf.war,id=1444246046 dependents are: [] 2005-10-26 21:08:19,937 ERROR [org.jboss.deployment.MainDeployer] Could not create deployment: file:/C:/java/jboss-5.0.0alpha/server/default/tmp/deploy/tmp22853techconf.ear-contents/techconf.war java.lang.IllegalStateException: java.lang.ClassNotFoundException: org.apache.tapestry.ApplicationServlet at org.jboss.ws.server.WebServiceDeployerJSE.isWebserviceDeployment(WebServiceDeployerJSE.java:143) at org.jboss.ws.server.WebServiceDeployer.create(WebServiceDeployer.java:83) at org.jboss.ws.server.WebServiceDeployerJSE.create(WebServiceDeployerJSE.java:49) at org.jboss.deployment.SubDeployerInterceptorSupport$XMBeanInterceptor.create(SubDeployerInterceptorSupport.java:160) at org.jboss.deployment.SubDeployerInterceptor.invoke(SubDeployerInterceptor.java:77) at org.jboss.mx.server.Invocation.invoke(Invocation.java:76) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:247) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:653) at org.jboss.mx.util.MBeanProxyExt.invoke(MBeanProxyExt.java:176) at $Proxy30.create(Unknown Source) at org.jboss.deployment.MainDeployer.create(MainDeployer.java:933) at org.jboss.deployment.MainDeployer.create(MainDeployer.java:923) at org.jboss.deployment.MainDeployer.deploy(MainDeployer.java:787) at org.jboss.deployment.MainDeployer.deploy(MainDeployer.java:751) at sun.reflect.GeneratedMethodAccessor46.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:143) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:82) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:123) at org.jboss.mx.server.Invocation.invoke(Invocation.java:76) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:125) at org.jboss.mx.server.Invocation.invoke(Invocation.java:76) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:247) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:653) at org.jboss.mx.util.MBeanProxyExt.invoke(MBeanProxyExt.java:176) at $Proxy9.deploy(Unknown Source) at org.jboss.deployment.scanner.URLDeploymentScanner.deploy(URLDeploymentScanner.java:319) at org.jboss.deployment.scanner.URLDeploymentScanner.scan(URLDeploymentScanner.java:507) at	0	0	https://issues.jboss.org/browse/JBADMCON-129
8482	JBADMCON-78	The attached admin-console.war file contains the GUI (and some services) implementation of the general JBossMQ settings. Please take a look at it and give us your comments. Some notes/questions: . In the ?Settings? tab, the mechanism behind selecting a different data source is yet to be implemented, so don?t expect the sql properties to be changed when you select a different data source. If possible, I?d like to check-in the current code before starting this implementation. . In the ?More?? tab, the ?Has security manager? attribute is a write-only attribute. I?ve figured out a way to accept the settings without displaying it (similar to what jmx-console is doing). Does this look ok? Should I change the ?Make soft references? attribute to use radio buttons as well so the boolean attributes have the same look and feel? . In the ?Read-Only? tab, the ?CacheStore? and ?MessageCache? attributes are configurable attributes for the mbean, but we?ve decided to make it read-only on the console. Is this information of any value to the user? Should we even include them at all? . The ?NextInterceptor? is a configurable attribute of the SecurityManager mbean, but we?ve decided not to include it in the console display. Should I remove the ?Interceptor? attribute from the ?Read-Only? tab as well?	1	1	https://issues.jboss.org/browse/JBADMCON-78
8483	JBADMCON-78	Chi, I've successfully deployed to JBoss4.0.3sp1 and finally can look at the pages in question. Looks very good. A couple of comments: The tabs at the top are a little less conspicuous than I'd be comfortable with (on Windows with IE, anyway). Could we make them a tiny bit larger, more conspicuous, something? Or does the style sheet control thils? I notice that the readonly tab's text is normal, whereas all the other tabs are bold. Not a big deal, just a little odd-looking. Radio-buttons for booleans seems correct to me (as long as there are only two buttons :)) I've always found it a little bit silly to type the word TRUE, or FALSE to denote something that is one bit. I sometimes worry that the text will not be interpreted correctly because the case was wrong, whatever.	0	0	https://issues.jboss.org/browse/JBADMCON-78
8484	JBADMCON-54	Create that as a method on the DeploymentService API for now. Be able to copy from either a local or remote url: /** * Upload a new library to server lib dir. A different * filename may be specified, when writing the library. * * If the target filename exists, upload is not performed. * * @param src the source url to copy * @param filename the filename to use when copying (optional) * @return true if upload was succesful, false otherwise */ public boolean uploadLibrary(URL src, String filename);	1	1	https://issues.jboss.org/browse/JBADMCON-54
8485	JBAOP-686	Flavia, Can you please update the build to use this? Ståle is already using it for the aophelper documentation, but his build file seems different from the example provided by Scott. I also don't think that the aophelper docs get included in the release.	0	0	https://issues.jboss.org/browse/JBAOP-686
8486	JBAOP-674	should i remove the 1.5.x compatibility tests?	1	1	https://issues.jboss.org/browse/JBAOP-674
8487	JBAOP-674	I want to be able to continue running the tests using one command, i.e.: $./build.sh -f build-tests-jdk50.xml tests Rather than $./build.sh -f build-tests-jdk50.xml tests $cd somewhere $ant -f compatibiity-tests.xml $cd somewhere/else $ant -f compatibility-tests2.xml etc. Since then people are a lot more likely to forget to run them. If the conditions are getting too big, do something else. A few immediate ideas that come to mind are: *Create a new _run-compatibility target for the compatibility tests in the main build file and call that rather than the target currently used *Create external files as you have now, but include and call them from the main build file	1	1	https://issues.jboss.org/browse/JBAOP-674

#	Requirement	Questions	Answered	Successfully Answered	URL
8488	JBAOP-642	Apparently this fix was added to the wrong place. ---------- Kabir, I think that we inadvertently committed the changes for JBAOP-642 [1] to the wrong branch. Instead of committing to the AOP CP branch [2], we committed to the AOP CP02 tag [3]. I'm not sure it really matters since CP02 is soon to become obsolete, but just for clarity, I think we need to revert the changes to the CPO2 tag [3] and make them to the CP branch [2]. Is this something you could do? I feel sure that I can undo your changes and move them to the CP branch myself, but since I am not familiar with these changes, I would feel better if you did them. Of course, I can do this if necessary. Could you do this while referencing JBAOP-642 in your commit messages as well so that we have that recorded in reference to that issue? [1] https://jira.jboss.org/jira/browse/JBAOP-642 [2] https://svn.jboss.org/repos/jbossas/branches/JBoss_AOP_1_5_5_GA_CP/ [3] https://svn.jboss.org/repos/jbossas/tags/JBoss_AOP_1_5_5_GA_CP02/ Jimmy	1	0	https://issues.jboss.org/browse/JBAOP-642
8489	JBAOP-578	Staale, can you please look at this asap? It will be called LOADS once I integrate the aop-mc-int stuff into AS 5.	1	1	https://issues.jboss.org/browse/JBAOP-578
8490	JBAOP-560	The feature is already implemented. However, the ConcurrentHashMap instances, used for storing aspect instances, are not compatible with null values.. Not storing null aspects results in extra calls to factories. Besides representing an overhead, this might introduce a bug. Suppose that a factory returns non-null instances at every 2 times it is called. If we call this factory more than once for the same scope, it will result in an inconsistent behaviour. I suggest storing an aspect that represents null in those tables. What do you say, Kabir?	1	1	https://issues.jboss.org/browse/JBAOP-560
8491	JBAOP-461	Taking another look at the memory-test ant script, I see that I was mistaken when I said that those tests were being run with loadtime-weaving. They are being run with javaagent but, before that, aopc performs an instrumentation on the test classes. AFAIK, this will result in javaagent doing nothing, since all the AOP target classes will be instrumented at that time. Kabir, what am I missing here? I don't get why we are using aopc and javaagent.	1	0	https://issues.jboss.org/browse/JBAOP-461
8492	JBAOP-431	Flavia, as discussed on MSN I have got rid of the casting for Invocation, JoinpointInfo, and Target (for around advice). Would you mind seeing if it is possible to get rid of this casting for things like Caller, Arg etc?	1	1	https://issues.jboss.org/browse/JBAOP-431
8493	JBAOP-387	Is this a dupicate of JBAOP-385?	1	1	https://issues.jboss.org/browse/JBAOP-387
8494	JBAOP-387	What would be the exception type? We could either throw a RuntimeException with a message describing the cause of invalidation, or we could throw a RuntimeException of our own, like: public class InvalidAdviceException extends RuntimeException {...}	1	1	https://issues.jboss.org/browse/JBAOP-387
8495	JBAOP-373	I'm thinking on JoinPointBean. What I don't know, if we go with it, is: - will this crash with the fact that I rever to both Invocation and JoinPointInfo hierarchies as being joinpoint beans? - should we change the name of the subtypes we have agreed before (i.e., org.jboss.aop.joinpoint.FieldRead -> org.jboss.aop.joinpoint.FieldReadBean)?	1	1	https://issues.jboss.org/browse/JBAOP-373
8496	JBAOP-373	Hi Flavia, I noticed our naming is a bit off, but please read to the end before changing anything! Part 1: Construction.getConstructor() ConstructorCall.getConstructor() ConstructorExecution.getConstructor() MethodCall.getMethod() vs FieldAccess.getAdvisedField() (I think this should be renamed getField()) MethodExecution.getAdvisedMethod() (I think this should be renamed getMethod()) Part 2: MethodCalledByMethod.getCallingMethod() and ConstructorCalledByMethod.getCallingMethod() vs MethodCalledByConstructor.getCalling() and MethodCalledByConstructor.getCalling() (I think these should be renamed getCallingConstructor()) I realise I am probably the one who came up with these names, but since you are looking at it would you mind taking a look? If these cannot be renamed, in other words if these methods exist in 1.5.x and are being called, then I think we will need to rename the methods in the interfaces and create duplicate methods in the implementing classes. For example: interface MethodExecution { Method getMethod(); } class MethodInfo implements MethodExecution { ... //keep this public Method getAdvisedMethod() { return advisedMethod; } //add this public Method getMethod() { return advisedMethod; } } I am not 100% sure if this is public API, so this might not be necessary. let's sync up tomorrow when you get online	1	1	https://issues.jboss.org/browse/JBAOP-373
8497	JBAOP-373	Which tests do you mean? This stuff hasn't really been released so far, so changing the API should be ok? I had a look through 1.5.x and could not really find much usage of this outside of the "internal" code. Some things like EJB 3 do use the XXXInfo classes though, so I think we should add this. I will take a look at this now	0	0	https://issues.jboss.org/browse/JBAOP-373
8498	JBAOP-355	Kabir Khan wrote: >> <snip> >> >>>> aopclasspath -- Huh??? >>> A folder to the annotated classes >> >> How is that different from src or classpath? > > By default we do not browse for annotations (it is time-consuming to > look inside every file), so if you have annotations you want parsed > you need to specify them. > OK, that makes sense. When I reread the doc and actually followed the link to the "Annotation Bindings" chapter I figured that was it. ;) But this is confusing: "aopclasspath - This should mirror your class path and contain all JARs/directories..." The "mirror class path" thing threw me off, as to me "mirror" implies "duplicate". Perhaps "aopclasspath -- A subset of the JARs/directories in classpath or classpathref that includes all JARs/directories that may have annotated aspects (See Chapter "Annotation Bindings"). The AOPC compiler will browse each class file in this path to determine if any of them are annotated with @Aspect. Browsing for annotations is time consuming, so only classes in the JARs/directories listed in aopclasspath will be scanned." The next sentence "This gets used for the jboss.aop.class.path optional system property which is described in the "Command Line" section." is really confusing. Is it necessary? (No need to reply to that question. Or to any of this :)) > btw src, (which I agree is craply named, but created before my time) > specifies the files that you want to transform - there may be a LOT of > other things on the classpath which you do not want to transform... > Yeah, it's pretty clear what it is and how to use it (other than the docs should say "A directory containing class files to be transformed", not just "containing files". There's just a mental disconnect because of the name that I expect affects people looking at the PojoCache examples. > IIRC you can specify multiple src entries, but I agree using > includes would be nicer Yep, you can have multiple src entries. You've got includes as well; it's just that for some reason they are at the same level as src rather than being nested.	1	1	https://issues.jboss.org/browse/JBAOP-355
8499	JBAOP-326	Need to decide: if an advice method receives a @Target parameter, and the join point context is static (i.e., there is no target available), should we call the advice with a null target, or not call the advice at all?	0	0	https://issues.jboss.org/browse/JBAOP-326
8500	JBAOP-290	there are a few .jars thats needed to compile the source/tests which we dont include in the dist. should we start to package them into the dist or use ant/buildmagic/ivy/etc?	1	1	https://issues.jboss.org/browse/JBAOP-290

#	Requirement	Questions	Answered	Successfully Answered	URL
8501	JBAOP-288	I downloaded and installed JBoss 4.0.5.GA on 11-13-2006. Our applications worked fine. Today on 11-29-2006 I downloaded and installed another copy of JBoss 4.0.5.GA on a different machine. Endless lines of "ERROR [STDERR] [error] ClassLoader has been garbage collected.. Do verbose mode if you want full stack trace." appeared in the log for this copy of JBoss, and everything in our application having to do with AOP became inoperative. I noticed that the files in the jboss-aop-jdk50.deployer changed their date - from 9-01-2006 11:59 AM to 9-22-2006 4:08 PM and various other times, some which did not agree with those found in the download of the standalone copy of AOP 1.5.2.GA. After reverting my second install to the 9-01-2006 version of AOP which was bundled with JBoss 4.0.5.GA as of 11-13-2006, the errors disappear and the application work as expected. Is it possible there is an incompatibility with AOP dated 9-22-2006 currently bundled with JBoss 4.0.5.GA?	1	1	https://issues.jboss.org/browse/JBAOP-288
8502	JBAOP-288	That is odd. I just tried a fresh download of jboss 4.0.5.GA and deployed the jboss-aop.deployer from aop 1.5.2.GA in there and it works as expected. Can you try again, and provide some more info? (Please create a new JIRA issue, link to this, and assign to me)	1	1	https://issues.jboss.org/browse/JBAOP-288
8503	JBAOP-288	I used the JEMS installer on both of these, and picked the EJB 3.0 option. Is this how you obtained your new copy, or did you use the standard download?	1	1	https://issues.jboss.org/browse/JBAOP-288
8504	JBAOP-269	Would the fix be included in AS 4.0.5? I need to run my tests against a released AS version and cannot wait for AS 5.0. Can we do that? Thanks a lot!	1	1	https://issues.jboss.org/browse/JBAOP-269
8505	JBAOP-269	Did you perform the tests?	0	0	https://issues.jboss.org/browse/JBAOP-269
8506	JBAOP-237	Kabir, Can you look into this task. Thanks.	0	0	https://issues.jboss.org/browse/JBAOP-237
8507	JBAOP-215	What is the current status of this?	1	1	https://issues.jboss.org/browse/JBAOP-215
8508	JBAOP-142	This as now been added as the 'maxmemory' property, in v1.11, Looks like it will be in the JBoss_Aop_1_3_3 release? http://anoncvs.forge.jboss.com/viewrep/JBoss/jboss-aop/src/main/org/jboss/aop/ant/AopC.java	0	0	https://issues.jboss.org/browse/JBAOP-142
8509	JBAOP-56	Bill, I am re-opening this issue since I just realized perhaps nothing has been done according to Adrian's comment. I thought I have made myself clear and the comemnts in the forum thread as well. But I will re-iterate the problem here again. Please keep in mind this has nothing to do with JBossCache. Any aop user will run into this problem when doing instrumentation. 1. For import package statements that are not used anywhere in the code. True, a user needs to clean up. But it can happen that during compile time, the package library is in the classpath, but during runtime it is not. Aop/Javassist will spew out a huge exception stack that is diccifult to trace during load time. I propose we need to re-package the exception and make it more verbose. 2. For import package statement that is actually used in the code but is not needed in particular runtime (for example, I don't need JNDI lookup when I am running outside AS). In this case, Javassist will puke if the package is not around. Now this is a legitimate use case since, for example, we package JBossCache to be run inside and outside of AS with the same jar. Solution? Offer a flag maybe to allow Javassist to continue even if the package is not found. What do you think? -Ben	1	1	https://issues.jboss.org/browse/JBAOP-56
8510	JBAOP-54	Just commited support for declare-warning and declare-error via XML, and a "visual" compile time test (I cannot see any way to automate this completely right now. The test is a separate target: $ant compiletime-declare Rather than creating a joinpoint query thing I plugged this into the Instrumentor and CallerTransformer, since they are iterating over everything anyway. You can match on type expressions (i.e. does this class implement this method) or on method/constructor calls. TBD: -Define via annotations -Loadtime tests -Once "[JBAOP-45]callfield pointcut" has been done, this should be added to this. Created new task JBAOP-57. -Create tutorial	0	0	https://issues.jboss.org/browse/JBAOP-54
8511	JBAOP-20	Added support for JDK 5.0. We may skip JDK 1.4 because of time. JDK 5.0 still needs some packaging issues as well.	0	0	https://issues.jboss.org/browse/JBAOP-20
8512	JBASMP-69	I have not tested with the patch command, but you can execute CLI commands, https://docs.jboss.org/jbossas/7/plugins/maven/latest/examples/execute-commands-example.html. Would that be viable?	1	1	https://issues.jboss.org/browse/JBASMP-69
8513	JBASMP-62	I'd prefer to leave this to CLI to handle. Would something like this work https://community.jboss.org/wiki/If-elseControlFlow? Or possibly https://community.jboss.org/wiki/Try-catch-finallyControlFlow.	0	0	https://issues.jboss.org/browse/JBASMP-62
8514	JBASMP-53	This is the definition for the `--runtime-name`. Does this fit what you'd need? {quote} optional, the runtime name for the deployment. This will form the basis for such things as default Java EE application and module names. This would typically be the same as --name, and if not specified the value used for --name will be used. In some cases users may wish to have two deployments with the same 'runtime-name' (e.g. two versions of "example.war") both available in the management configuration, in which case the deployments would need to have distinct 'name' values but would have the same 'runtime-name'. Within an individual server, only one deployment with a given 'runtime-name' can deployed. However, multiple deployments with the same 'runtime-name' can exist in the configuration, so long as only one is enabled. {quote}	1	1	https://issues.jboss.org/browse/JBASMP-53
8515	JBASMP-53	This would be a great feature to add. [edited] I see org.wildfly.plugins:wildfly-maven-plugin:1.1.0.Alpha5-SNAPSHOT implements the ability to specify the runtime-name feature. It appears to work with JBoss EAP 6. Is it a good idea to use the wildfly plugin with JBoss EAP 6?	1	0	https://issues.jboss.org/browse/JBASMP-53
8516	JBASMP-45	Do you have the {{<configuration/>}} element defined on the plugin or just the execution? The plugin won't inherit the configuration from an execution definition.	1	1	https://issues.jboss.org/browse/JBASMP-45
8517	JBASMP-45	I see. Is it possible to tell mvn jboss-as:deploy-artifact which execution to run?	1	1	https://issues.jboss.org/browse/JBASMP-45
8518	JBASMP-41	It's great to see a prompt response! Thanks! For now, would it fail also if I use CLI scripts?	0	0	https://issues.jboss.org/browse/JBASMP-41
8519	JBASMP-25	Comment Reference: https://github.com/jbossas/jboss-as-maven-plugin/issues/23#issuecomment-9693313 Ok, I thought that the run goal would not block, instead return control to maven once the server is started. If it is supposed to block then how should JBoss be shut down so that Maven would continue? BTW, {{StandaloneServer.isRunning()}} will return false if the server state is STARTING so the loop would exit rather than block if JBoss doesn't start almost immediately: {code} final ModelNode result = client.execute(Operations.createReadAttributeOperation(Operations.SERVER_STATE)); isRunning = Operations.successful(result) && !STARTING.equals(Operations.readResultAsString(result)) && !STOPPING.equals(Operations.readResultAsString(result)); {code} The scenario we are looking to support is: 1. Start JBoss from Maven, configuring datasources, queues etc. in the pre-integration-tests phase 2. Run a suite of integration tests using failsafe in the integration-tests phase 3. Shut down JBoss with the :shutdown command in the post-integration-tests phase 4. Failsure verifies the test results in the verify phase For this to work we would need the Run goal to fork JBoss once it's started rather than block. If it is supposed to block by design then can a configuration option be added to support forking? {code} <fork>true</fork> {code}	1	1	https://issues.jboss.org/browse/JBASMP-25
8520	JBASMP-25	The scenario described above is what we aimed to implement as well. Will it be supported in the near future?	1	1	https://issues.jboss.org/browse/JBASMP-25

#	Requirement	Questions	Answered	Successfully Answered	URL
8521	JBASMP-25	Finally getting around to this. Is there a reason not to use Arquillian for in container testing rather than starting the server through the maven plugin?	1	1	https://issues.jboss.org/browse/JBASMP-25
8522	JBASMP-25	Good question. I should say that before the scenario above is performed we are adding our own modules to support jsr289 and updating the javaee.api module to include the SIP Servlet API etc to the JBoss/EAP. We do this by using several maven plugins and the result is an extended EAP ready to go. Do you think we can do the same with Arquillian?	1	1	https://issues.jboss.org/browse/JBASMP-25
8523	JBASMP-25	Sounds good. Do you have an estimate when these features will be released and in what version of the plugin?	1	1	https://issues.jboss.org/browse/JBASMP-25
8524	JBASMP-2	Another feature that should also be added is checking what packaging type the project is - for example a MojoFailureException if the packaging type is pom. The other thing I noticed is that the JbossAsControlCreateConfigurationMojo doesn't seem to perform a delta copy as ALR has mentioned in places, but instead copies the entire directory every time and overwrites existing files. Is this a feature worth pursuing? It took quite a long time (at least 60 seconds) copying the config directory on every invoke for me - i.e. it appears worth it to me. It also looks like it might be worth abstracting out the other common properties / functions from the mojo's into a super class, as I have done with JbossAbstractDeploymentMojo. Maven still appears to set all the properties etc correctly and it gets rid of a lot of duplicate code.	1	1	https://issues.jboss.org/browse/JBASMP-2
8525	ASPATCH-391	To concerns need to be addressed: - 1 - The bug is a race condition so I do not have either an automated or manual test case, however the synchronisation will be minimal so AS testsuite testing should still confirm no regressions. 2 - There is no upstream, JBoss AS 4.0.3.SP1 was the last release to use JBoss WS4EE, can we consider the CP branch the most upstream now?	1	1	https://issues.jboss.org/browse/ASPATCH-391
8526	ASPATCH-363	Malik Saheb said in his email sent on the 11/28/2008: ========== The customer seems satisfied by the provided patch. We could add it to the next CP. By the way, we have convinced this customer to migrate to an EAP version (4.3). He has planned to study the migration process next year, so he is asking if the xalan correction applied for 4.0.4 will be also available for JBoss EAP 4.3. ========== The patch has not been included in CP15. It would be most appreciated if it can be included in CP16. Can we also be certain that it will be applied to EAP 4.3?	1	1	https://issues.jboss.org/browse/ASPATCH-363
8527	ASPATCH-343	The following branches have been switched to the 5.5.23.patch04-brew release of Tomcat JBoss_4_0_4_GA_CP JBoss_4_0_5_GA_CP which contains all the following fixes: ASPATCH-234 and IT #168408 (ASPATCH-346), as well as fixes for CVE-2007-2449, CVE-2007-2450, CVE-2007-3386, CVE-2007-3382, CVE-2007-3385, 2007-5342, and CVE-2007-5461.	0	0	https://issues.jboss.org/browse/ASPATCH-343
8528	ASPATCH-338	Can you provide an ETA for this? Code freeze is today, however we can wait a couple of days.	1	1	https://issues.jboss.org/browse/ASPATCH-338
8529	ASPATCH-331	The problems I see with this change are the following: 1. It jumps from 5.5.9 to 5.5.23 2. If we don't fix the vulnerability in the Tomcat 5.5.9 patch branch, we'll have to upgrade 4.0.2 to Tomcat 5.5.23 as well. We have previously agreed to maintain 5.5.9 with security errata to support JBoss AS: 4.0.3SP1 cp's 4.0.2 cp's The point is, if we swallow this for 4.0.3.SP1, We'll also have to swallow it for the 4.0.2 cp branch. Are we ready to do this?	0	0	https://issues.jboss.org/browse/ASPATCH-331
8530	ASPATCH-331	How much work is it to backport this fix to Tomcat 5.5.9?	0	0	https://issues.jboss.org/browse/ASPATCH-331
8531	ASPATCH-322	Brian, would it be possible for you to backport your fix to 4.0.1 into the 4.0.0 CP branch?	1	1	https://issues.jboss.org/browse/ASPATCH-322
8532	ASPATCH-317	In that case are we Ok to go ahead and I will commit the fix?	1	1	https://issues.jboss.org/browse/ASPATCH-317
8533	ASPATCH-239	Can we make the legacy, broken behavior the default and provide the corrected behavior if a given option is checked? This will eliminate the impact on those who depend on the broken functionality, while allowing use of the corrected behavior if desired. This will also allow inclusion into the CP. So long as the patch can be done in a form which allows inclusion into the CP, I'm fine with producing a temporary one-off. The should know they are required to upgrade to the CP once it is released.	1	1	https://issues.jboss.org/browse/ASPATCH-239
8534	ASPATCH-231	Anil, what should we do?	0	0	https://issues.jboss.org/browse/ASPATCH-231
8535	ASPATCH-207	what is the revision id of your last commit?	0	0	https://issues.jboss.org/browse/ASPATCH-207
8536	ASPATCH-204	are we going to postpone CP01 until June?	1	1	https://issues.jboss.org/browse/ASPATCH-204
8537	ASPATCH-203	are we going to postpone CP01 until June?	0	0	https://issues.jboss.org/browse/ASPATCH-203
8538	ASPATCH-203	Vivek, what's the status? Code freeze is June 1st.	0	0	https://issues.jboss.org/browse/ASPATCH-203
8539	ASPATCH-196	I merged the changes from: https://svn.jboss.org/repos/jbossas/branches/JBoss_4_0_1_JBAS-2704/ because that is what was mentioned in the JBAS-2704 Now I see that there is also a: https://svn.jboss.org/repos/jbossas/branches/JBoss_4_0_1_JBAS-2704_1/ What is this second one? Which one should have been merged? Thanks	1	0	https://issues.jboss.org/browse/ASPATCH-196
8540	ASPATCH-196	Amit, can you please clarify which is the right patch for 4.0.1.SP1.CP ? Thanks	1	0	https://issues.jboss.org/browse/ASPATCH-196
8541	ASPATCH-194	Hi Alex, Have you determined if this needs to be merged to CP? If it is then please assign to me and I will do the merge. Vivek	1	1	https://issues.jboss.org/browse/ASPATCH-194
8542	ASPATCH-190	can this be closed?	0	0	https://issues.jboss.org/browse/ASPATCH-190
8543	ASPATCH-171	Jim, what do we do with this?	0	0	https://issues.jboss.org/browse/ASPATCH-171
8544	ASPATCH-164	While committing the changes, I made two errors in the last two commits in this branch. It should refer to ASPATCH-164 not JBAS-3988, do you know how to change them?	1	1	https://issues.jboss.org/browse/ASPATCH-164
8545	ASPATCH-164	Test now passes :) I need to change the patch instructions and get those two commit messages changed. I'll do that on Monday, does that fit with the CP02 release? AFAIK, it should be released the 2nd Tuesday of the month	1	1	https://issues.jboss.org/browse/ASPATCH-164
8546	ASPATCH-157	Is this going to be comitted before march 1st (tomorrow)?	1	1	https://issues.jboss.org/browse/ASPATCH-157
8547	ASPATCH-154	At the moment, I do not have any bandwith for this task. Would it be possible to reassign it to someone else? Thanks,	1	0	https://issues.jboss.org/browse/ASPATCH-154
8548	ASPATCH-138	OK, I can't compute the source diff because the JBossXB project has a branch JBossXB-1.0.0.CR6_JBAS-3882 but there isn't a JBossXB-1.0.0.CR6 tag (only CR7 and CR8 were tagged). The only remaining solution would be to blindly replace the jboss-xml-binding.jar with the one attached to http://jira.jboss.com/jira/browse/JBAS-3882 and creating a thirdpart repository version 1.0.0.CR6.patch01 for that. Imprtant: Is this patch included in XB RC8? Or the problem does not occur there?	1	1	https://issues.jboss.org/browse/ASPATCH-138

#	Requirement	Questions	Answered	Successfully Answered	URL
8549	ASPATCH-138	Hi Alexey, I've created a new version in thirdpart for jbossxb, namely1.0.0.CR6.patch01, that is supposed to replace the buggy 1.0.0.CR6 for jbossas 4.0.5_GA_CP. I've got the JAR that you have attached to JBAS-3882 which is supposedly a fixed 1.0.0.CR6, from the sources at: https://svn.jboss.org/repos/common/jbossxb/branches/JBossXB-1.0.0.CR6_JBAS-3882/ I am just wondering if that JAR was indeed crated from these sources and how pervasive a change this has been (remember, there is no CR6 tag, so I can't compare sources). The reason I ask is because the 1.0.0.CR6 JAR in thirdpart has: 346762 Jul 13 2006 jboss-xml-binding.jar while the one you attached (1.0.0.CR6_JBAS-3882) has: 265414 Feb 1 20:26 jboss-xml-binding.jar I did a comparison of the files and they do have the same classes exactly, but the byte count of every class is significantly different (thus the large difference in the file sizes). Is this really OK? Can we ship this new JAR in a CP release?	0	0	https://issues.jboss.org/browse/ASPATCH-138
8550	ASPATCH-129	Brian just edited JBAS-3808 and added this to the description: NOTE: THIS PATCH SHOULD NOT BE INCLUDED IN ANY CUMULATIVE PATCH RELEASE However, this patch has already been committed to 4.0.3SP1_CP (but not released, yet) Brian, could you explain why this should not be included in a CP?	0	0	https://issues.jboss.org/browse/ASPATCH-129
8551	ASPATCH-129	I think we should revert this patch from the CP branch. Vivek, can you do it today?	0	0	https://issues.jboss.org/browse/ASPATCH-129
8552	ASPATCH-129	Brian, we probably don't want to leave this customer using a one-off branch forever. In which newer version the new structure is used? Is in only in 5.0 or will that be the format for 4.2? (I take as it is incompatible that it is not in any of the 4.0.x releases) Regards, Fernando	1	1	https://issues.jboss.org/browse/ASPATCH-129
8553	ASPATCH-128	Thanks Dimitris! Merged into JBoss_4_0_4_GA_CP branch: 58752. Is it possible request QA on JBoss_4_0_4_GA_CP or does the entire task have to be completed first?	1	1	https://issues.jboss.org/browse/ASPATCH-128
8554	ASPATCH-128	Dimitris, can you comment?	0	0	https://issues.jboss.org/browse/ASPATCH-128
8555	ASPATCH-127	The patch.zip contains both server/all/deploy/http-invoker.sar/invoker.war/WEB-INF/web.xml server/all/deploy/httpha-invoker.sar/invoker.war/WEB-INF/web.xml Shouldn't the first one be removed?	0	0	https://issues.jboss.org/browse/ASPATCH-127
8556	ASPATCH-127	hey Ryan, did we add any kind of reminder to CP03 release notes?	1	0	https://issues.jboss.org/browse/ASPATCH-127
8557	ASPATCH-125	The one-off branch adds jboss/axis/lib/axis-ws4ee.jar. I am not familiar with the manner in which thirdparty jars were integrated in AS 3.2.7, I didnt see a build-thirdparty.xml or another file that could tell me if there is a change in versions. Nor is there anything in the sources to suggest a change in axis-ws4ee.jar is necessary. Is this an errant jar committed to the branch or is it actually deliberate?	1	1	https://issues.jboss.org/browse/ASPATCH-125
8558	ASPATCH-124	Vicky, diff between JBoss_4_0_3_JBAS-2638 and 4.0.3.SP1 results in tons of diffs. Can you help us merge this patch into the 4.0.3.SP1_CP03 branch?	1	1	https://issues.jboss.org/browse/ASPATCH-124
8559	ASPATCH-124	Vicky, which branch did you check this in to? I made a typo in my previous comment. Correct branch URL is https://svn.jboss.org/repos/jbossas/branches/JBoss_4_0_3_SP1_CP CP releases are tagged off the CP branch, so we don't have a separate branch for each incremental CP release.	1	1	https://issues.jboss.org/browse/ASPATCH-124
8560	ASPATCH-124	not sure i understand where it was checked in since there is no 4.0.3.SP1_CP03 branch (it's a tag). will you be able to finish this by the end of this week? code freeze is on 12/01	1	1	https://issues.jboss.org/browse/ASPATCH-124
8561	ASPATCH-122	Since there isnt a seperate one-off branch for this, I cant isolate the change to Branch_3_2. Was there an actual commit made to that branch for this? I am all 'svn log'ed out... Perhaps Scott could take a quick look...	0	0	https://issues.jboss.org/browse/ASPATCH-122
8562	ASPATCH-121	Ok, but don't know what you need me to do. The jira issue links are almost impossible for me to follow, but looks like the fix for JBAS-3316 was included in patch issued under ASPATCH-3. Are you saying that same code change in the patch for JBossAS-4.0.3.SP1_CP01 (ASPATCH-3) needs to be done for JBossAS-4.0.4.GA_CP03?	1	1	https://issues.jboss.org/browse/ASPATCH-121
8563	ASPATCH-118	Changes from ASPATCH-34 seem to overlap with changes in JBoss_4_0_3_SP1_JBAS-3650 branch: svn merge --dry-run --ignore-ancestry https://svn.jboss.org/repos/jbossas/tags/JBoss_4_0_3_SP1 https://svn.jboss.org/repos/jbossas/branches/JBoss_4_0_3_SP1_JBAS-3650 Skipped 'thirdparty' A thirdparty/.project U build/build-thirdparty.xml U security/src/main/org/jboss/test/LoginModulesTestCase.java C security/src/main/org/jboss/security/auth/spi/LdapExtLoginModule.java Needs developer comment.	0	0	https://issues.jboss.org/browse/ASPATCH-118
8564	ASPATCH-118	ASPATCH-34 incorporated JBAS-3645 which appears to have been addressed by JBAS-3650. Scott, could you confirm?	1	1	https://issues.jboss.org/browse/ASPATCH-118
8565	ASPATCH-117	Weston, could you take a look at the merge conflict? thanks	1	1	https://issues.jboss.org/browse/ASPATCH-117
8566	ASPATCH-117	do you mean this is a one-off patch that doesn't apply to other customers?	1	1	https://issues.jboss.org/browse/ASPATCH-117
8567	ASPATCH-77	Requesting triage due to above comments (custom branch, special config instructions). Can we require customers to move around config files? What else will this impact?	0	0	https://issues.jboss.org/browse/ASPATCH-77
8568	ASPATCH-70	Dimitris, i assigned this to you by mistake. sorry Fernando, just create a new JIRA and link it ("is related to") to the JBAS-2866	0	0	https://issues.jboss.org/browse/ASPATCH-70
8569	ASPATCH-61	Alexey, can you respond to this? Was the fix checked it into JBoss_4_0_4_GA_JBAS-3335?	1	1	https://issues.jboss.org/browse/ASPATCH-61
8570	ASPATCH-61	Should I properly fix it in the JBoss_4_0_4_GA_JBAS-3335?	1	1	https://issues.jboss.org/browse/ASPATCH-61
8571	ASPATCH-53	Yes, which branch should it be committed to?	0	0	https://issues.jboss.org/browse/ASPATCH-53
8572	ASPATCH-39	This patch will not be included because it relies on Java5. JBoss 4.0.x must be able to run on Java 1.4	1	1	https://issues.jboss.org/browse/ASPATCH-39
8573	ASPATCH-32	should we not include this into 4.0.3.SP1_CP branch?	1	1	https://issues.jboss.org/browse/ASPATCH-32
8574	JBASM-38	thanks for reporting. If I understand your report right, you actually run the shutdown batch a 2nd time (just after running it once already) and it showed you the above stack-trace? If that is the case then you filed a bug against jboss tools whereas this is a jboss as (and not tooling) related issue. Could you please indicate what AS you're trying to stop?	0	0	https://issues.jboss.org/browse/JBASM-38
8575	JBASM-23	Hallo, Anil. First thank you for your reply. I see that the "server" property is false by default. And setServer() method makes it true. So actually here the problem is not that we want single credentials for all servers but rather different credentials for the different servers. Would you be able to explain or point me at documentation how to make certain calls with certain credentials? For example I want to invoke an mbean operation with this username/password pair then operation on another server's mbean with different credentials. [code] doAs(credsServer1); invoke(...); doAs(credsServer2); invoke(...); [/code] Or something similar? To be clear it's not guaranteed invokes will be done by a different or the same thread. I can see that if I SecurityAssociation.setServer() then I can create a controlling thread for every server and pass operations through it. But I hope there will be some more lightweight solution to the problem. Thanks again!	1	1	https://issues.jboss.org/browse/JBASM-23
8576	JBASM-6	I'm planning on at least one more release before 5.0.GA. Should this be scheduled for the next release?	1	1	https://issues.jboss.org/browse/JBASM-6
8577	JBASM-3	So is there any conclusion to the topic how should the thread dump be generated? From what I read following Galder's link I can see that only "kill -QUIT" is reliable so should I aim at that or go for the best (which is it?) pure java solution?	1	1	https://issues.jboss.org/browse/JBASM-3

#	Requirement	Questions	Answered	Successfully Answered	URL
8578	JBASPECT-21	Andrew, what is the version of aspects security now?	0	0	https://issues.jboss.org/browse/JBASPECT-21
8579	JBASPECT-6	What do I need to do to see these compilation errors?	1	1	https://issues.jboss.org/browse/JBASPECT-6
8580	JBBENCH-140	Do you have more information to attach to the case? Like graph calls or something similar?	1	1	https://issues.jboss.org/browse/JBBENCH-140
8581	JBBENCH-76	Mark: I thought the log4j error was happening when executing driver.sh on Specj's driver. Can you please specify where/when this error is happening? I have seen this once, but I don't recall where this is happening.	1	1	https://issues.jboss.org/browse/JBBENCH-76
8582	JBBENCH-76	Clebert, I don't see the actual error listed in this jira anywhere. Am I missing it? We should have the whole error listed here..... Then I'll verify for you. Thanks, Deb	0	0	https://issues.jboss.org/browse/JBBENCH-76
8583	JBBENCH-55	We are waiting for Sasha's testcase for proceed on this. Did you get the testcase done Sasha? Also... did you get this working by changing Tomcat's properties at tomcat-service.xml?	0	0	https://issues.jboss.org/browse/JBBENCH-55
8584	JBBENCH-45	We can keep creating testcases for this task. I have already added some JRunit tests in JBossCache what is already a start. We could use the same framework for soak tests. Can please someone take a look on it?	1	1	https://issues.jboss.org/browse/JBBENCH-45
8585	JBBENCH-45	Clebert, could you follow this up with Ben and his work with JBento? Thanks!	1	1	https://issues.jboss.org/browse/JBBENCH-45
8586	JBBENCH-37	If we have the same application with three different jboss-web.xml, generating three different WAR applications. Is it valid, for general soak tests, if we execute the three versions simultaneously, and if something breaks execute them individually?	0	0	https://issues.jboss.org/browse/JBBENCH-37
8587	JBBENCH-33	Which one is the related bug in cache?	0	0	https://issues.jboss.org/browse/JBBENCH-33
8588	JBBENCH-27	Bela, what should we do with this issue? Performance results for 2.5 are clearly superior to INRIA results.	1	0	https://issues.jboss.org/browse/JBBENCH-27
8589	JBBENCH-27	How long would it take you to recreate their tests and publish new results ? The pain is that their paper still shows up in Google search...	0	0	https://issues.jboss.org/browse/JBBENCH-27
8590	JBBENCH-21	Hi Clebert, Is this the latest kit that your testing with then? Deb -- Deborah Beaman Hewlett - Packard Company TSG / Solution Alliances Engineering Linux Infrastructure Portfolios Engineering 110 Spit Brook Road ZK03-2X61 Nashua, NH 03062-2698 Tel: 603-884-1812 Fax: 603-884-0627 Email: deb.beaman@hp.com	1	1	https://issues.jboss.org/browse/JBBENCH-21
8591	JBBENCH-19	Do you want to work on this task? We can give you all the support for this activity if you want to!	0	0	https://issues.jboss.org/browse/JBBENCH-19
8592	JBBENCH-14	Do we have access the sourcecode for this benchmark? I guess if we use JBossSerialization here will have something better than KaRMI It would be nice to measure it.	0	0	https://issues.jboss.org/browse/JBBENCH-14
8593	JBBENCH-12	Specj is not automatable in terms of cron-tabs. Although Specj is well automated in terms of generating server configurations. The process will have to be manual, but mostly 80% automated. (Already part of Jboss specific files distribution)	0	0	https://issues.jboss.org/browse/JBBENCH-12
8594	JBBENCH-7	This should include running the benchmark against various JVMs including: JRockit (latest) sun JVM 1.3 sun JVM 1.4 sun JVM 1.5	1	1	https://issues.jboss.org/browse/JBBENCH-7
8595	JBBOOT-89	The attached patch is generating the start-up message relying on the concrete Server implementation's toString method to produce everything before the "Started in..." portion of the message. Do we want to include the (MicroContainer) portion or would the following be better? JBoss AS [1.2.3.VersionTag (build: SVNTag=VersionCVSTag date=BuildID)] Started in 27s:158m	1	1	https://issues.jboss.org/browse/JBBOOT-89
8596	JBBOOT-89	Should we also update the MC impl to also output a similar message? It doesn't appear to have Server toString method defined.	1	1	https://issues.jboss.org/browse/JBBOOT-89
8597	JBBOOT-89	Are there some nonstandard patch instructions? Typically we make via "svn di > /path/to/file" and patch in using "patch -p0 < /path/to/file". patching file impl-as/src/main/java/org/jboss/bootstrap/impl/as/server/AbstractJBossASServerBase.java can't find file to patch at input line 382 Perhaps you used the wrong -p or --strip option? The text leading up to this was: -------------------------- |Index: impl-base/src/main/java/org/jboss/bootstrap/impl/base/server/AbstractVersion.java |=== |--- impl-base/src/main/java/org/jboss/bootstrap/impl/base/server/AbstractVersion.java (revision 92321) |+++ impl-base/src/main/java/org/jboss/bootstrap/impl/base/server/AbstractVersion.java (working copy)	1	1	https://issues.jboss.org/browse/JBBOOT-89
8598	JBBOOT-86	The only downside of this is that you can't directly copy/paste the url, but not sure how many does that ? And if you use JBOSS_HOME as the base that should be printed full, allowing you to copy paste in two sets. Maybe print the full url if log level is TRACE and otherwise relative ?	0	0	https://issues.jboss.org/browse/JBBOOT-86
8599	JBBOOT-86	Forums are down ATM, so I'll do my peer review of the patch discussion here. ;) 1) The patch fails the TestSuite with NPEs. Check the development instructions: http://www.jboss.org/community/wiki/JBossBootstrapDevelopmentandContribution 2) This line: // Log full paths when log level Trace is enabled. if (log.isTraceEnabled()) { return url.toExternalForm(); } ...I don't like due to the conditional logic depending upon this log level. Essentially this creates a case where the test output (with logging set to TRACE) is different than in other environments, and the rest of the method doesn't get tested. So I'd skip this step. Let's take another look with those two issues addressed? Awesome to see your contribution Ray, thanks for getting involved.	1	1	https://issues.jboss.org/browse/JBBOOT-86
8600	JBBOOT-56	Where are we with this one, John?	1	0	https://issues.jboss.org/browse/JBBOOT-56
8601	JBBOOT-56	Sorry, didn't see your comment. I believe the patch I attached last is sufficient for this task. It may be a good idea to create a separate JIRA to centralize some of the test classes to avoid re-use. Would you like me to commit the attached patch?	0	0	https://issues.jboss.org/browse/JBBOOT-56
8602	JBBOOT-43	After a series of tests, it is a pretty straight forward change to get full URL support. The change just requires the AbstractMCServerBase start to process all nested <url> elements with the bootstrapUrl as the context and not the bootstrapHome. The only thing that does not work is having relative URLs what begin with a slash either in the bootstrapUrl or a nested <url>. This does not work in the existing implementation either. Should this be supported, or should it continue to throw an exception when starting with a slash? I am attaching a patch that includes some tests to show the failure when the slash is present.	1	1	https://issues.jboss.org/browse/JBBOOT-43
8603	JBBOOT-43	Thanks for looking it over. I have cleaned up 1 and 3 from above. I was attempting to test multiple different configurations of bootstrapUrl and bootstrapHome in a single test, which does not seem to work even with the shutdowServer lifecycle called. The server seems to leave the system property set for the bootstrapUrl after shutdown. This causes fun issues when you attempt to start a new server with just the bootstrapHome. As for the new exception type. Is there a standard place to put this, and do we just want to check if the URL starts with a slash and throw right away?	1	1	https://issues.jboss.org/browse/JBBOOT-43

#	Requirement	Questions	Answered	Successfully Answered	URL
8604	JBBOOT-43	After second thought, I am not able to remove the additional lifecycle method in MCServerBootstrapUrlTestCase.@After. Even with the shutdownServerAndClear lifecycle, the system properties are not being removed between tests. If the system properties are not removed between tests, subsequent tests will fail if the only the bootstrapHome is set. When each test runs through the start process it runs through the property overrides and sets system properties for each of the bootstrap(Home, Name, URL) properties. So if the bootstrapUrl is set in one test, it will still be set as a system property for the next test. If the subsequent test only sets the bootstrapHome and bootstrapName, it will get the bootstrapUrl applied from the last test as part of the override properties step, and will likely fail with the wrong bootstrapUrl. This seems like something that should get cleaned up when the server is shutdown. To test this behavior, remove the clearServerProperties method and rerun the tests. You will see a failure occur and if traced through it is bases on a lingering bootstrapUrl system property. Should these system properties be cleaned up at server shutdown?	1	1	https://issues.jboss.org/browse/JBBOOT-43
8605	JBBOOT-43	Seems like the server.shutdown shouldn't be necessary for tests that successfully startup, as it is in the @After lifecycle already. As you showed in the patch, the issue was with the exception cases the shutdown was not happening because it wasn't being started. ServerInitializer.cleanup resolves the issue. Is there a reason we still need to call the server.shutdown in the test body?	1	1	https://issues.jboss.org/browse/JBBOOT-43
8606	RHBPMS-4812	I understand that it's possible to achieve the goal with some workaround; nevertheless, I've seen this use case often and I'd like to achieve this goal with a more elegant and straightforward way. I was not aware of https://issues.jboss.org/browse/JBPM-5363, so it basically address the problem! Although, I think that the parametric signalRef is a way to workaround the limitations of the correlation key. In fact, there are many situations where the business key is created during the process or comes from an external system at the first interaction. Moreover, there could be many business keys that should be used to talk with different external parties. (e.g. CRM id, ERP id, etc). What do you think?	1	1	https://issues.jboss.org/browse/RHBPMS-4812
8607	RHBPMS-4812	I wouldn't call it a workaround. That is how it was designed to keep single correlation key for process instance to allow same look up mechanism as by process instance id though it is caller defined - aka business key. So how would you like to handle situation when we allow to set the business key within active instance - what should happen to the process instance if it attempts to use already used business key? The other thing you mentioned, many business keys for given process instance - that could work as the uniqueness should be on business key cannot be shared across instances but single instance can have many business key, that could be a valid extension. Though you can achieve that by using process instance variables and then whenever needed get hold of process instance by searching by its variables.	1	0	https://issues.jboss.org/browse/RHBPMS-4812
8608	RHBPMS-4812	so what you're suggesting is to have ability to assign more correlation keys into given process instance and pretty much at anytime of its life time, correct? I would agree but with restriction you can't change it but only add new one, does it make sense?	1	0	https://issues.jboss.org/browse/RHBPMS-4812
8609	RHBPMS-4803	Dear Engineering I have accidentally set the target release to 6.4.5, but unable to remove it. Can you please remove the content I entered. Thank you Toby	1	1	https://issues.jboss.org/browse/RHBPMS-4803
8610	RHBPMS-4673	In what sense is PostgreSQL supposed to be "included in a OpenShift process server image"? We do not bundle third-party components.	0	0	https://issues.jboss.org/browse/RHBPMS-4673
8611	RHBPMS-4480	[~tirelli] [~KrisVerlaenen] I need more time to finish work on this. Would estimate 2 to 3 days. Let me know if that's ok or should we move to next release?	1	1	https://issues.jboss.org/browse/RHBPMS-4480
8612	RHBPMS-4421	Hi [~tlivora] I'm thinking about suggesting container names like this: "org.jboss:integration:2.0.0-SNAPSHOT" (and if it's an invalid name, "org.jboss:integration:2.0.0-SNAPSHOT-2" and so on). What do you think?	1	1	https://issues.jboss.org/browse/RHBPMS-4421
8613	RHBPMS-4421	I tested this with BPM Suite 6.4.1 CR1 and the name is automatically generated. I just have one problem with inconsistency of the behavior: when I have several artifacts listed as a serach result, clicking some other artifacts sets it's GAV, but leaves the container name unchanged (courtesy of [this if|https://github.com/droolsjbpm/kie-wb-common/pull/614/files#diff-cd5ac2ebd6279017c3a017bfe3433840R244]). Are we ok with this inconsistency - wouldn't it be better to regenerate the container name as well to have the same behavior as for gav? [~karreiro] See video attached and notice how GAV changes when selecting different artifact but container name remains the same.	1	1	https://issues.jboss.org/browse/RHBPMS-4421
8614	RHBPMS-4198	can we consider this as limited solution on EAP7 and provide only support for non ssl JMS connections? [~tlivora] [~KrisVerlaenen]	1	1	https://issues.jboss.org/browse/RHBPMS-4198
8615	RHBPMS-4198	[~swiderski.maciej], as this use case (JMS over SSL for task operations) has been available only for EAP 6.x so far, I think we don't need to fix it for EAP 7.0. [~paradhya2], do you agree?	1	0	https://issues.jboss.org/browse/RHBPMS-4198
8616	RHBPMS-4150	Please explain what you mean by "build fails"? Can you provide a stack trace? There's no reason I can think of for the build to *fail* (if the project the copy was taken from built OK)	1	1	https://issues.jboss.org/browse/RHBPMS-4150
8617	RHBPMS-4150	[~kgaevski] Thanks for the explanation and video.. so the second build/deploy is expected to fail because a KJAR for the GAV has already been deployed. That said I agree it'd be nice if when copying a Project the User is either prompted for new GroupId, ArtifactID and Version or we automatically set the ArtifactId to the new Project name (as is the case when a new Project is created). Which is your preference? Automatically updating the ArtifactID is much simpler but may be the less preferred solution.	1	1	https://issues.jboss.org/browse/RHBPMS-4150
8618	RHBPMS-4150	Hi [~manstis], thank you for the fast reply. In my opinion the best case is when the User is prompted for new GroupId, ArtifactID and Version. But ArtifactID will be automatically set to new value (Project name) and Version will be set to default one (Version: 1.0). What do you think, is it possible to do in this way?	1	0	https://issues.jboss.org/browse/RHBPMS-4150
8619	RHBPMS-4123	[~lpetrovicky] [~kgaevski] , since this is a minor issue, given the effort that [~manstis] estimated above, and the short term nature of the change, I tend to agree with Michael. Can we postpone this for 7.x where a different solution will be used?	1	0	https://issues.jboss.org/browse/RHBPMS-4123
8620	RHBPMS-4079	[~swiderski.maciej] you are correct, it was not designed to do that. I think it should be fine to improve the end point to also start the scanner after the server is created, if the user sets the corresponding parameter. Wdyt?	1	1	https://issues.jboss.org/browse/RHBPMS-4079
8621	RHBPMS-3980	Maciej, can you please backport this into 6.4.x?	0	0	https://issues.jboss.org/browse/RHBPMS-3980
8622	RHBPMS-3832	Cherry-picked into 6.0.x: http://github.com/droolsjbpm/kie-wb-common/compare/2921103...c2f54f4	0	0	https://issues.jboss.org/browse/RHBPMS-3832
8623	RHBPMS-3729	Hi Ivo, Doc Text has been added. Can you please review the CCFR.	1	1	https://issues.jboss.org/browse/RHBPMS-3729
8624	RHBPMS-3610	This should be possible based on the work done for deployment descriptors (https://bugzilla.redhat.com/show_bug.cgi?id=1017327) that allows to configure persistence unit name for audit and runtime separately. So configuring the persistence unit names would be sufficient when using JPA mode of audit logs. Though for JMS based separate queues are needed to be created and then configuration file (jbpm.audit.jms.properties) would need to be added to kjar to instruct the runtime what queue it should use. Would this be enough or we need more advanced support?	0	0	https://issues.jboss.org/browse/RHBPMS-3610
8625	RHBPMS-3266	We can not make these changes, as they would affect all products ... are classes that are use in many places.	0	0	https://issues.jboss.org/browse/RHBPMS-3266

#	Requirement	Questions	Answered	Successfully Answered	URL
8626	RHBPMS-3199	Hi Sona, Doc Text has been added. Can you please review the CCFR.	1	1	https://issues.jboss.org/browse/RHBPMS-3199
8627	RHBPMS-3198	Marek, there should not be possibility to add listeners via services. Please note that they operate on deployment units and as such each deployment unit can have different configuration in that regard. So services are completely independent of that thus do not expose any way of registering listeners. When using deployment service to put kjar into runtime you can use deployment descriptor to define all listeners, handlers, etc you need and they will be available on runtime. Does that make sense?	1	0	https://issues.jboss.org/browse/RHBPMS-3198
8628	RHBPMS-3041	Carmine Talpa <carmine.talpa@gmail.com> made a comment on jira JBPM-3600 Hi Maciej, How can i specify time interval for work day ? e.g.: http://docs.jboss.com/jbpm/v3/userguide/businesscalendar.html weekday.monday= 9:00-12:00 & 12:30-17:00 ... This is an example of my businesscalendar properties file: business.hours.per.week = 6 business.hours.per.day = 36 business.start.hour = 8 business.end.hour = 15 business.weekend.days = 7 business.holiday.date.format = yyyy-DD-mm business.holidays = 2013-01-01,2013-05-01,2013-06-02,2013-06-13:2013-06-16,2013-08-15,2013-12-25 Could we have an example of a new "jbpm.business.calendar.properties" file Thanks a lot	0	0	https://issues.jboss.org/browse/RHBPMS-3041
8629	RHBPMS-3040	There have been several enhancements / fixes related to data assignments for 6.0.3 and 6.0.3 roll-up 1. Could we retest what still needs to be fixed? I would recommend opening individual issues for open issues, as that is much easier to keep track. Some additional feedback: 1) Is as expected, as there are two ways you could specify a value, either by giving it an explicit value or by mapping a variable to it. In most cases, the value is probably filled in, so by default the editor shows "data input" "is equal to" and the user can type a value in "To value". 2) Once a data output is defined, it should show in when creating a data output in the "From Object" dropdown. 3) I agree, an error message or something similar that some info (like "mapped to") is missing would be useful. Note that https://bugzilla.redhat.com/show_bug.cgi?id=1152826 is still an open issue related to data assignments as well.	1	1	https://issues.jboss.org/browse/RHBPMS-3040
8630	RHBPMS-3040	Tiho or Anton, please can you tell me what is a result of this bugzilla? There was mention a lot of problems but ER3 does not contain any fix. Problems: 1) it is not a bug. The current solution is expected. An user can assign a value. If he needs to map a value from a process variable then he has to rearrange the particular assignment. (Swap from and to object) 2) it is a real issue. This issue is not fixed by ER3. Predefined assignments should contains data inputs and data outputs together. (The assignment of "Result" is missing in ER3) 3) it is not a bug. The video sequence does not fill the type of the assignment. There was the additional request (comment #4) for the error notification. Is it the part of this bz? ER3 has similar behaviour as ER2. I assign this issue for 6.1.0 ER3 BPMS. Please, close this issue if you open new issue for each problem. In case that this bugzilla will be modified than specify what should be fixed.	0	0	https://issues.jboss.org/browse/RHBPMS-3040
8631	RHBPMS-2983	Could someone provide screenshots of the what this looks like for BPMS vs. Polymita?	0	0	https://issues.jboss.org/browse/RHBPMS-2983
8632	RHBPMS-2974	Toshiya, is this something for you to organize within the Zanata teams?	0	0	https://issues.jboss.org/browse/RHBPMS-2974
8633	RHBPMS-2974	Okay, I have fixed them in Zanata. You can pull them but Neus seems to be working on Zanata for other fixes and doesn't push those changes to github master yet. Could you check with Neus?	1	1	https://issues.jboss.org/browse/RHBPMS-2974
8634	RHBPMS-2974	Neus, After the outstanding i18n issues have been localized and pushed to Zanata; could you please pull the latest back into all modules from Zanata too. Thanks, Mike	0	0	https://issues.jboss.org/browse/RHBPMS-2974
8635	RHBPMS-2558	Is there any reason to remove the jbpm ddls in product or this is just a productisation mistake? jBPM ddls has never been existed in product. I didn't see a requirement in PRD which requires to include those ddls. or I am wrong? Ryan, creating DDLs was a requirement for the release. I remember having conversation with the team on this. I agree that it was not clearly defined on where and how it should be packaged, supplemental zip makes lot more sense. Given where we are with CR2 for 6.0, I would not push for changing the bits at this point. But we should get this in for 6.0.1 Alessandro, please do me a favor. File a bugzilla issue highlighting all the DDL scripts that needs to be included. There should be scripts for most of the DBs we certify. Lets track it closely for 6.0.1	1	1	https://issues.jboss.org/browse/RHBPMS-2558
8636	RHBPMS-2501	Hi Zuzana, Doc Text has been added. Can you please review the CCFR.	1	1	https://issues.jboss.org/browse/RHBPMS-2501
8637	RHBPMS-2454	I have checked with BPMS 6.1.0.DR3 and on page http://localhost:8080/business-central can be seen a help page listing basic REST API commands. However the reported page http://localhost:8080/business-central/rest still shows the same error as in bug description. New basic help screen in exec-server profile definitely improved the situation, but is there more work to be done? What is meant by swagger-ui? Thanks, Jiri	1	0	https://issues.jboss.org/browse/RHBPMS-2454
8638	RHBPMS-2400	Can't the existing BOMs be used for this purpose already (and if so, why not)?	1	0	https://issues.jboss.org/browse/RHBPMS-2400
8639	RHBPMS-2353	Jiri, are you saying you want to restrict the captured image to the size of the editor window? The bpmn2 editor delegates to an image capture/save dialog defined in the Graphiti framework. Maybe this enhancement request should be filed with the eclipse Graphiti project, otherwise I would have to duplicate a whole bunch of Graphiti internal code in the editor, just to override the current behavior.	0	0	https://issues.jboss.org/browse/RHBPMS-2353
8640	RHBPMS-2314	Marco, is this possible using the new QueryFilter?	1	0	https://issues.jboss.org/browse/RHBPMS-2314
8641	RHBPMS-2314	Hello Marco. Does "the new query methods underwater" mean org.jbpm.services.task.impl.TaskQueryServiceImpl.java ? If this is not the one, could you point out the class for me please? Then I'll guide to our customer that "firstResult and maxResults option" will be provided by another API from 6.1. Thx	1	1	https://issues.jboss.org/browse/RHBPMS-2314
8642	RHBPMS-2314	Hisao, what would you like to do with this BZ? We could either move this to product 6.2 and see it as a feature request, or close it.	1	0	https://issues.jboss.org/browse/RHBPMS-2314
8643	RHBPMS-2314	Hello. Thank you for your updates. Our customer wants to archive what I explained as "Expected results:" with BPMS 6.1. If user can specify "firstReuslt" and "maxResults" using "getTasksByVariousFields()" comes with BPMS 6.1, could you share with me how to do it please? If it's not possible, does customer need to use the new query methods? I might not correctly understand your suggestion. So could you explain what customer can do with BPMS 6.1 to set "firstResult" and "maxResults" with some simple sudo-code please? Thx	1	1	https://issues.jboss.org/browse/RHBPMS-2314
8644	RHBPMS-2314	Hisao, Does this help? String userId = "user"; InternalTaskService internalTaskService = (InternalTaskService) taskService; List<TaskSummary> taskSumList = internalTaskService.taskQuery(userId) .offset(10).maxResults(15) .orderBy(TaskQueryBuilder.OrderBy.processInstanceId) .buildQuery().getResultList();	1	0	https://issues.jboss.org/browse/RHBPMS-2314
8645	RHBPMS-2231	this property is currently not used in parser and is set to true by default on bpmn2 level. not user configurable at this point.	0	0	https://issues.jboss.org/browse/RHBPMS-2231
8646	RHBPMS-2219	Marco, can this one be closed based on what's there currently already?	0	0	https://issues.jboss.org/browse/RHBPMS-2219

#	Requirement	Questions	Answered	Successfully Answered	URL
8647	RHBPMS-2219	I've tested this with ER7 and written several tests for this feature, see: http://git.app.eng.bos.redhat.com/jbossqe/brms.git/tree/test-jbpm-integration/src/test/java/org/jboss/qa/bpms/jbpm/integration/remote/VariableHistoryTest.java However none of them work. 1.) I cannot work with variables correctly. String values are somehow converted to Long and I don't know why: parameters.put("var1", "30"); produces: java.lang.ClassCastException: java.lang.Long cannot be cast to java.lang.String Starting the process from business-central manually works well. 2.) In other case JaxbHistoryLogList contains empty internal historyLogList, no variable "myobject" found in the log. 3.) I've also encountered this error: org.jboss.resteasy.plugins.providers.jaxb.JAXBUnmarshalException: javax.xml.bind.UnmarshalException: unexpected element (uri:"", local:"process-instance-list"). // Expected elements are <{}log-instance-list>,<{}node-instance-log>,<{}process-instance-log>,<{}variable-instance-log> All errors are commented in corresponding test methods, which reproduce them. To reproduce issues listed above: 1.) Download BPMS 6 ER7 (eap6) and setup a db (in my case PostgreSQL). 2.) Add user ibek: ibek=9143dd65119c7b733e02c5076dcf8f84 (into application-users.properties) ibek=admin (into application-roles.properties) 3.) Run the server. 4.) Run the tests: git clone git://git.app.eng.bos.redhat.com/jbossqe/brms.git cd brms mvn clean test -Pjbpm-integration -Dtest=VariableHistoryTest (it will install test projects from bpms-assets.git automatically). I think this should be fixed, as this was requested by a customer.	0	0	https://issues.jboss.org/browse/RHBPMS-2219
8648	RHBPMS-2183	Hi Sal, what kind of custom authentication are you talking about? Can you elaborate a little bit more on how exactly the customer wants to use this?	1	0	https://issues.jboss.org/browse/RHBPMS-2183
8649	RHBPMS-2055	Verified in BPM Suite 6.1.0.ER4. QE test coverage: https://gitlab.mw.lab.eng.bos.redhat.com/bxms/brms/commit/d69e38edea2a503952f313f700773f03286e7744 Created also a pull request, because the current tests did not fail properly with strict option disabled: https://github.com/droolsjbpm/jbpm/pull/235 Can you please review and pull?	0	0	https://issues.jboss.org/browse/RHBPMS-2055
8650	RHBPMS-2011	I think all the various grains of control are what you strive for: - configure it per project - configure it per deployment (then have a way for the project to select the deployment to run in) - per server such as the default is always per_instance if not defined instead of SINGLETON as is now. I guess the solutions offered in previous posts would all be possible and may be desirable depending on the above configuration needs: - configure per project (in the kjar) - configure per deployment (via UI, some xml file on server to override defaults) - configure per server (same as deployments?) If you use kjars and allow dependencies, what level wins? project or deployment or server overrides? Just my thoughts after playing with it a bit today.	0	0	https://issues.jboss.org/browse/RHBPMS-2011
8651	RHBPMS-2003	currently simulation supports only XOR and AND gateways. For OR gateways there will be a need to introduce more simulation properties as more than one can be selected. Moreover distribution of the branches to be selected needs to somehow configurable as well as then set of outgoing flows from OR gateway should change per simulation runs and not be constant. This requires quite some development effort to be properly (and useful) implemented. Thus I would opt to move that to 6.0.1. Wdyt?	0	0	https://issues.jboss.org/browse/RHBPMS-2003
8652	RHBPMS-2003	Please, can you add some error message to the BPMS 6.0.1 that "OR gateway is not supported"? Please, give this alert into validator or end simulation engine with this error if a process contains "OR" gateway. Should I create a new bugzilla issue for this request?	1	1	https://issues.jboss.org/browse/RHBPMS-2003
8653	RHBPMS-2003	(In reply to Sona Mala from comment #6) > Please, can you add some error message to the BPMS 6.0.1 that "OR gateway is > not supported"? Please, give this alert into validator or end simulation > engine with this error if a process contains "OR" gateway. > > Should I create a new bugzilla issue for this request? Sona, can you create a new bz for better error feedback? We triaged this today and agreed the request as a Feature Request which cannot be done for 6.0.1.	0	0	https://issues.jboss.org/browse/RHBPMS-2003
8654	RHBPMS-1979	I am not totally sure that this is really an unexpected behaviour. I see the business-central as a collaborative tool. If I am working in a project and I save a process in a inconsistent state I think reasonable that every logged user working in the same BPM Suite instance (and with visibility for that process) see the error. If I am working in a rule of a project, where a different user is working in a process in the same project - and I or the other user save something with errors, we need to inform all users about that or the build will not succeed. Is not that expected ?	1	0	https://issues.jboss.org/browse/RHBPMS-1979
8655	RHBPMS-1943	[~jwest] We believe this was improved in the 6.3 release, can you check and close this JIRA if that is the case?	1	0	https://issues.jboss.org/browse/RHBPMS-1943
8656	RHBPMS-1936	Can we easily make job name a configurable column in the table?	0	0	https://issues.jboss.org/browse/RHBPMS-1936
8657	RHBPMS-1829	https://gitlab.mw.lab.eng.bos.redhat.com/bxms/brms/blob/2df19b26eea173e81e774dad3f60a9c97bb61c0c/test-jbpm-regression/src/test/java/com/bpms/customer/BZ1110206LdapUserGroupCallbackSearchScopeTest.java https://gitlab.mw.lab.eng.bos.redhat.com/bxms/brms/blob/2df19b26eea173e81e774dad3f60a9c97bb61c0c/test-jbpm-regression/src/test/java/com/bpms/customer/BZ1110206LdapUserInfoSearchScopeTest.java Maciej, great job! According to the results of the tests listed above both components are working as expected. However, I am not sure if the default search scope should be OBJECT_SCOPE when it is not possible to parse the value of 'ldap.search.scope' property. If this property is not set, ONELEVEL_SCOPE is used (as it was in the previous versions). I would recommend to unify this behavior and use ONELEVEL_SCOPE in both cases. What do you think?	0	0	https://issues.jboss.org/browse/RHBPMS-1829
8658	RHBPMS-1823	Could you give some idea about the response time you are getting? If it's quick, showing a loading page might not be required? How many processes do you have and how long is it taking to load these?	1	0	https://issues.jboss.org/browse/RHBPMS-1823
8659	RHBPMS-1823	(In reply to Kris Verlaenen from comment #2) > Could you give some idea about the response time you are getting? If it's > quick, showing a loading page might not be required? How many processes do > you have and how long is it taking to load these? There are 35 processes and load out after about 10 seconds. "No process definition found" always shown in the process list page.	0	0	https://issues.jboss.org/browse/RHBPMS-1823
8660	RHBPMS-1823	Mauricio, could we just show something else in the table, like "Loading ...", or show the "Loading ..." popup while the results haven't returned yet?	1	0	https://issues.jboss.org/browse/RHBPMS-1823
8661	RHBPMS-1808	Kris, I'm not sure just process instance id is sufficient. We still need to identify the user task node - think of a process definition with multiple user tasks, we need to track every task event to a single user task. And what about process definition with a user task in a loop ? (my assumption is that the listener erasing completed tasks is active). I guess the concept of runtime table & history table does not have any performance implications - it works for process instances, am I right? Maybe I just simplified it too much, what do you think about this solution? Thanks.	1	1	https://issues.jboss.org/browse/RHBPMS-1808
8662	RHBPMS-1807	To recap the changes: 1. Java File renamed to Data Object (in both DM and Project explorer). 2. List of types to select from when adding new field no longer contains [0..N] variant for each type, replaced by "[] List" checkbox. 3. Create button kept where it was. 4. Position attribute kept. 5. Extra attributes on hover - unknown.	0	0	https://issues.jboss.org/browse/RHBPMS-1807

#	Requirement	Questions	Answered	Successfully Answered	URL
8663	RHBPMS-1794	Would we be able to use a configurable shared ssh login, one that only background processes use? I wouldn't be against enabling the git daemon again as an implementation detail, but I'd rather the user not be directed to clone from it if they can't push. That means removing the options to copy the git address from the UI.	0	0	https://issues.jboss.org/browse/RHBPMS-1794
8664	RHBPMS-1794	I guess if we provide information for the user that git is read only and ssh is read/write that should provide enough information to choose the right one. wdyt? not sure I understand your question about ssh - behavior should not change for ssh at all.	0	0	https://issues.jboss.org/browse/RHBPMS-1794
8665	RHBPMS-1752	Hi Anton, The installer has been modified to configure multiple datasources now, using the same JDBC driver. For BPMS they can seperately install a datasource for each of dashbuilder and business-central. The datasources are independent of each other with the exception of the JDBC driver. Is this what you had in mind? Thanks, Mansi	1	0	https://issues.jboss.org/browse/RHBPMS-1752
8666	RHBPMS-1751	This ticket includes that home page enhancements that depend on BS3 upgrade, so it was postponed to product 6.3. During the face to face meeting QE also indicated this will have an impact on their test suite. Rajesh, can you please nack-it for 6.2? I don't have the permissions to do it myself. Thank you.	0	0	https://issues.jboss.org/browse/RHBPMS-1751
8667	RHBPMS-1733	Isn't this a database configuration issue? What happens if/when the underlying database implementation is changed? Does the jbpm plugin need to be changed again?	1	1	https://issues.jboss.org/browse/RHBPMS-1733
8668	RHBPMS-1700	Hello, Can this issue be marked as closed?	1	1	https://issues.jboss.org/browse/RHBPMS-1700
8669	RHBPMS-1670	I don't think this is a good idea at all... the intention of help icon ("?") is to provide some levle of help/support for users, move HOME link from left hand side to right hand side is not counter intuitive in my opinium, but this is also a misuse of the `help` component. Joe, what do you think?	1	0	https://issues.jboss.org/browse/RHBPMS-1670
8670	RHBPMS-1661	Hello Jeff, What do you think about the Marco's comment, does it address your request? Can you please describe how this issue should be tested? Do you have a reproducer/test for this issue? Thanks, Jiri	1	0	https://issues.jboss.org/browse/RHBPMS-1661
8671	RHBPMS-1661	Hello, Thank you Marco! In regards to QA ... i do not have a test case handy, sorry. My recommendation is to build a test script that invokes these modified human task query APIs using a variable number of human tasks persisted in a postgresql and/or mysql database. ie: test with 1 task persisted in the jbpm database . Then test with 100 tasks --> 1K --> 10K ---> 100K . How much is the performance degredation between tests ? thanks again! jeff	1	0	https://issues.jboss.org/browse/RHBPMS-1661
8672	RHBPMS-1620	OK - The issue is date granularity. Only units bigger than days are taken into consideration. Basically the command will not delete logs older than a day from today. Could we make this configurable so an admin is able to delete also logs older then a minute or second?	1	1	https://issues.jboss.org/browse/RHBPMS-1620
8673	RHBPMS-1607	Ivo, could you attach the stack trace? Also, if you have a link to the test in the jbossqe/brms repo, that would be great.	1	0	https://issues.jboss.org/browse/RHBPMS-1607
8674	RHBPMS-1599	[~KrisVerlaenen] Can we get some update on this? Or should it be closed and replaced with RFE?	0	0	https://issues.jboss.org/browse/RHBPMS-1599
8675	RHBPMS-1513	getVariable() currently does return the actual value. I'm not sure why you would expect this to be a copy? ksession contains a lot of similar operations (startProcess, insert, getObject, getProcessInstance) al behaving the same why, would you expect all these to be changed as well then? I agree a setVariable operation on ksession is a good idea.	1	0	https://issues.jboss.org/browse/RHBPMS-1513
8676	RHBPMS-1433	FYI: I believe that fix for this issue was already implemented as a part of fix for bug #1155480 Mauricio, what do you think?	1	1	https://issues.jboss.org/browse/RHBPMS-1433
8677	RHBPMS-1426	What do you mean with process description, could you give an example? Since you're looking at the process instance list, wouldn't all process instances of the same process show the same process description? How will this help differentiating process instances?	1	1	https://issues.jboss.org/browse/RHBPMS-1426
8678	RHBPMS-1426	(In reply to Kris Verlaenen from comment #2) > What do you mean with process description, could you give an example? Since > you're looking at the process instance list, wouldn't all process instances > of the same process show the same process description? How will this help > differentiating process instances? Hi, Kris Our requirement is just to make user can differentiate the process instances clealy. Add Process Description column is one workaround, do you have any other workaround? Our idea is: Add Process Description input on each process start form by default. Then the process description input can be shown in Process Description column to make user saw the difference by different process description. Any question please feel free to let we know.	1	0	https://issues.jboss.org/browse/RHBPMS-1426
8679	RHBPMS-1426	Adding a description column is good. Can I specify more columns? For example, for a purchase process, I would like to see: Id, desc, ... | who | item | price | number | total ---------- x | John | pen | $1 | 10 | $10 ... Gererally I want a customized query here by conditions of both fixed columns (such as start/end date) and process variables (such as who=john) and I can specify which colomns are displayed on the list.	0	0	https://issues.jboss.org/browse/RHBPMS-1426
8680	RHBPMS-1416	Thanks for quick implementation. Unfortunately, customer is running BRMS on Windows where is not allowed to install 3rd party software (e.g. git). Is it possible to enable this option without git command?	1	1	https://issues.jboss.org/browse/RHBPMS-1416
8681	RHBPMS-1413	Joe, since you own the UX, can you +1 this (or give feedback in case you have other suggestions)?	1	0	https://issues.jboss.org/browse/RHBPMS-1413
8682	RHBPMS-1387	Hi Marek, Doc Text has been added. Can you please review the CCFR.	0	0	https://issues.jboss.org/browse/RHBPMS-1387
8683	RHBPMS-1330	Ryan, could you add these migration scripts we have in the installer in the community to the supplementary tools?	1	1	https://issues.jboss.org/browse/RHBPMS-1330
8684	RHBPMS-1330	I have added the following into supplementary-tools: jboss-brms-bpmsuite-6.2-supplementary-tools/ upgrade-scripts/ ├── db2 | ├── bpmsuite-6.0-to-6.1.sql | └── bpmsuite-6.1-to-6.2.sql ├── h2 | ├── bpmsuite-6.0-to-6.1.sql | └── bpmsuite-6.1-to-6.2.sql ├── mysql5 | ├── bpmsuite-6.0-to-6.1.sql | └── bpmsuite-6.1-to-6.2.sql ├── oracle | ├── bpmsuite-6.0-to-6.1.sql | └── bpmsuite-6.1-to-6.2.sql ├── postgresql | ├── bpmsuite-6.0-to-6.1.sql | └── bpmsuite-6.1-to-6.2.sql ├── sqlserver | ├── bpmsuite-6.0-to-6.1.sql | └── bpmsuite-6.1-to-6.2.sql └ sybase ├── bpmsuite-6.0-to-6.1.sql └── bpmsuite-6.1-to-6.2.sql Is there anything missing?	1	0	https://issues.jboss.org/browse/RHBPMS-1330
8685	RHBPMS-1330	(In reply to Kris Verlaenen from comment #26) > Both pull requests related to the second and third issue mentioned in > comment 22 were merged to 6.3.x: > https://github.com/droolsjbpm/jbpm/commit/ > 751359c9bbca1b49e850723b9f527b8cd220f7d6 > https://github.com/droolsjbpm/jbpm/commit/ > b29c4fd348b33821623930c4631db8d282ba8e35 > > Keeping it assigned to Ryan for the mysql innodb fix. OK, I have added mysql innodb into the supplementary tools. @Tibo, However, I can't find the dashbuilder sql for mysql innodb. Shoud we assign this to dashbuilder owner to add the mysql innodb? Let me mark this as MODIFIED for test on CR1 first.	1	0	https://issues.jboss.org/browse/RHBPMS-1330

#	Requirement	Questions	Answered	Successfully Answered	URL
8686	RHBPMS-1253	This issue should take high priority for Openshift, Freemodeler and cloud deployments, since the loading of a simple process over a internet connection creates great delays, which greatly hurts usability. For example, loading a process involves about 223 different HTTP requests, summarized as follows (tested in 6.0.1 version) 223 requests 2 documents 10 stylesheets 121 images 12 scripts 78 SVG and others (Errai) Of course, this creates a big server overload and also a big delay in opening the process. Can you please confirm whether this has been targeted for 6.1 ?	1	1	https://issues.jboss.org/browse/RHBPMS-1253
8687	RHBPMS-1232	Can we mark this as modified now? I believe that what is now on 6.0.x meet the basic expectations.	1	1	https://issues.jboss.org/browse/RHBPMS-1232
8688	RHBPMS-1232	(In reply to Mauricio Salatino from comment #2) > Can we mark this as modified now? I believe that what is now on 6.0.x meet > the basic expectations. I'd like to wait for the more advanced table/grid widget before you mark this as modified (unless I understood wrong and the widget is already in the code, just not in the ER2 build).	0	0	https://issues.jboss.org/browse/RHBPMS-1232
8689	RHBPMS-1217	This looks to be complete to me!?! Can we set to MODIFIED?	0	0	https://issues.jboss.org/browse/RHBPMS-1217
8690	RHBPMS-1212	David, does the above make sense? Could you please try your use-case: - without removing the repository declaration from settings.xml, - with specifying <offline/> element in the settings.xml? Thanks!	0	0	https://issues.jboss.org/browse/RHBPMS-1212
8691	RHBPMS-1212	David, to help us think of the best solution, what are you trying to achieve? Are you setting up a machine which will have no internet connection when put into production? If so, the offline mode should help. Thanks!	0	0	https://issues.jboss.org/browse/RHBPMS-1212
8692	RHBPMS-1070	Adam, could you elaborate why you would need access to the WorkItemInfo object? This is considered an internal persistence Object and should only be accessed with caution (similar to for example ProcessInstanceInfo). Note that you could get a WorkItem using: ((org.drools.core.process.instance.WorkItemManager) ksession.getWorkItemManager()).getWorkItem(id)	1	0	https://issues.jboss.org/browse/RHBPMS-1070
8693	RHBPMS-1055	Anton, could you please elaborate more on this? how would you like to see that? We have tests in code base to cover this [1] though not sure what you'd like to have here... [1] https://github.com/droolsjbpm/jbpm/blob/master/jbpm-services/jbpm-kie-services/src/test/java/org/jbpm/kie/services/test/KModuleWithDependenciesDeploymentServiceTest.java	1	0	https://issues.jboss.org/browse/RHBPMS-1055
8694	RHBPMS-1000	Looking at the JaxbProcessDefinition class [1] I would like to gather opinion if we shall expose forms via that object? Main concern is that most of the forms will be created with form modeler then it will require form modeler to actually understand the form content and render it thus making it less useful for remote clients if at all. What do you think? [1] https://github.com/droolsjbpm/droolsjbpm-integration/blob/master/kie-remote/kie-remote-jaxb/src/main/java/org/kie/services/client/serialization/jaxb/impl/process/JaxbProcessDefinition.java	0	0	https://issues.jboss.org/browse/RHBPMS-1000
8695	RHBPMS-937	The jira mentioned in comment 6 will be included in 6.1 version so the remaining part is to configure handler to use BASIC authentication. My vote is to stick with what we have now and if needed reconfigure it using deployment descriptors per kjar/project. What do you think?	1	0	https://issues.jboss.org/browse/RHBPMS-937
8696	RHBPMS-937	Ivo, do you by any change use singleton session and both processes are in the same project/kjar?	1	0	https://issues.jboss.org/browse/RHBPMS-937
8697	RHBPMS-937	We have discovered a limitation of the current solution, see BZ1209479. Http authentication method of a REST task cannot be explicitly set. It can be set only indirectly like this new org.jbpm.process.workitem.rest.RESTWorkItemHandler("ibek","ibek1234;") see [1]. This means a limitation, because imagine a process definition which has several REST tasks, where each one uses a different remote endpoint and a different auth method. One configuration of one RestWorkItemHandler cannot handle that. Documentation team, can you please document this? [1] https://github.com/droolsjbpm/jbpm/blob/6.2.x/jbpm-workitems/src/main/java/org/jbpm/process/workitem/rest/RESTWorkItemHandler.java#L128	1	0	https://issues.jboss.org/browse/RHBPMS-937
8698	RHBPMS-937	Kris, have you read the linked bugzilla description? You can set "Username" and "Password" as you say, but the problem is elsewhere. You cannot set explicitly http auth type.	1	0	https://issues.jboss.org/browse/RHBPMS-937
8699	RHBPMS-914	I am not able to reproduce this issue, since new issue is present all the time. at step 8. Process model window is opened and designer is loading but never starts. There is printed error in log: 14:27:47,677 ERROR [stderr] (http-/127.0.0.1:8080-2) context [anonymous] 6597:12 attribute includedo isn't defined Isn't it possible that fix broke something else? It seems that there was same issue many time in history: https://developer.jboss.org/thread/238296?start=0&tstart=0 returning back to Assigned	1	1	https://issues.jboss.org/browse/RHBPMS-914
8700	RHBPMS-896	@Bob, please in what version of BPMN2 Modeler is this issue fixed? I have installed BPMN2 Modeler 1.1.3.201504171557 (Luna), but in properties of User Task element were no new tabs. Is needed some more configuration? For example Window > Preferences > BPMN2 > ...	1	1	https://issues.jboss.org/browse/RHBPMS-896

#	Requirement	Questions	Answered	Successfully Answered	URL
8701	RHBPMS-896	>(UserTask)-(EndEvent) 2. Double click on (UserTask) 3. In pop-up window named 'Edit User Task' click 'Add' in section Notifications 4. Fill all fields of new Notification 5. Click Ok 6. Save process 7. Repeat step 2. Pop-up window with Properties of (UserTask) should be displayed, but it is not. This Exception can be found in error log: Message: GFCommandStack.execute(Command) org.eclipse.core.runtime.AssertionFailedException: assertion failed: StackTrace: org.eclipse.core.runtime.AssertionFailedException: assertion failed: at org.eclipse.core.runtime.Assert.isTrue(Assert.java:110) at org.eclipse.core.runtime.Assert.isTrue(Assert.java:96) at org.eclipse.bpmn2.modeler.runtime.jboss.jbpm5.property.extensions.TaskNotification.fromString(TaskNotification.java:338) at org.eclipse.bpmn2.modeler.runtime.jboss.jbpm5.property.extensions.TaskNotificationList.add(TaskNotificationList.java:48) at org.eclipse.bpmn2.modeler.runtime.jboss.jbpm5.property.extensions.TaskNotificationsListComposite$1.inputChanged(TaskNotificationsListComposite.java:169) at org.eclipse.jface.viewers.ContentViewer.setInput(ContentViewer.java:288) at org.eclipse.jface.viewers.StructuredViewer.setInput(StructuredViewer.java:1675) at org.eclipse.bpmn2.modeler.core.merrimac.clad.AbstractListComposite.bindList(AbstractListComposite.java:510) at org.eclipse.bpmn2.modeler.runtime.jboss.jbpm5.property.JbpmUserTaskDetailComposite.createBindings(JbpmUserTaskDetailComposite.java:47) at org.eclipse.bpmn2.modeler.core.merrimac.clad.AbstractDetailComposite.setBusinessObject(AbstractDetailComposite.java:131) at org.eclipse.bpmn2.modeler.core.merrimac.clad.DefaultDialogComposite.setData(DefaultDialogComposite.java:212) at org.eclipse.bpmn2.modeler.core.merrimac.dialogs.AbstractObjectEditingDialog.aboutToOpen(AbstractObjectEditingDialog.java:186) at org.eclipse.bpmn2.modeler.core.merrimac.dialogs.AbstractObjectEditingDialog.open(AbstractObjectEditingDialog.java:254) at org.eclipse.bpmn2.modeler.core.merrimac.dialogs.ObjectEditingDialog.open(ObjectEditingDialog.java:72) at org.eclipse.bpmn2.modeler.core.features.ShowPropertiesFeature.execute(ShowPropertiesFeature.java:92) at org.eclipse.graphiti.features.custom.AbstractCustomFeature.execute(AbstractCustomFeature.java:83) at org.eclipse.graphiti.internal.command.GenericFeatureCommandWithContext.execute(GenericFeatureCommandWithContext.java:64) at org.eclipse.graphiti.ui.internal.command.GefCommandWrapper.execute(GefCommandWrapper.java:51) at org.eclipse.graphiti.ui.internal.editor.EmfOnGefCommand.execute(EmfOnGefCommand.java:54) at org.eclipse.graphiti.internal.command.GFPreparableCommand2.doExecute(GFPreparableCommand2.java:43) at org.eclipse.emf.transaction.RecordingCommand.execute(RecordingCommand.java:135) at org.eclipse.emf.workspace.EMFCommandOperation.doExecute(EMFCommandOperation.java:119) at org.eclipse.emf.workspace.AbstractEMFOperation.execute(AbstractEMFOperation.java:150) at org.eclipse.core.commands.operations.DefaultOperationHistory.execute(DefaultOperationHistory.java:513) at org.eclipse.emf.workspace.impl.WorkspaceCommandStackImpl.doExecute(WorkspaceCommandStackImpl.java:208) at	1	1	https://issues.jboss.org/browse/RHBPMS-896
8702	RHBPMS-873	I understand that we cannot add such a method to KieSession/TaskService. ksession.getKieBase().getProcess() may not be good for Remote API, either. Hmm, how about introducing a new interface class for Remote API?	0	0	https://issues.jboss.org/browse/RHBPMS-873
8703	RHBPMS-855	Hi Everyone, Is there any new progress on this? The localization team is currently working on the QA for 6.3 atm and the fact of having 14 different Zanata links makes it very time consuming to find the strings to be modified/corrected. Let me know if I can be of any help from my end. Thanks a lot! Angela García	1	0	https://issues.jboss.org/browse/RHBPMS-855
8704	RHBPMS-813	This BZ states that Documentation is not included in PDF after exporting a process... I can reproduce it on BPM Suite 6.0.3 and 6.1.0. At the bottom of the page [1], doc states: "Note that when exporting to PDF, the output PDF file will contain also the values of the Documentation properties of the Process and its elements." Are we missing something? [1] https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_BPM_Suite/6.1/html/User_Guide/Exporting_a_Process1.html	1	0	https://issues.jboss.org/browse/RHBPMS-813
8705	RHBPMS-745	Is it still make no sense for current release? Is it reasonable to close this issue?	1	0	https://issues.jboss.org/browse/RHBPMS-745
8706	RHBPMS-734	For now you can see "Unable to import provided BPMN2" after import incorrect process instead of "Successfully imported BPMN2" and it is good. But I guess it is better to go by way proposed by Marek: *** Expected results: Only 'Unable to import provided JSON' or 'Unable to import provided BPMN2' error message will be displayed, Key word is "Only". So user will not see the "Replace existing model?" dialogue and immediately see message "Unable to import provided BPMN2". What do you think?	0	0	https://issues.jboss.org/browse/RHBPMS-734
8707	RHBPMS-718	There is also a column 'createdby_id' - it seems confusing. What should be the difference between taskinitiator_id and createdby_id ?	1	0	https://issues.jboss.org/browse/RHBPMS-718
8708	RHBPMS-703	Good catch Anton, fixed on master jbpm master: https://github.com/droolsjbpm/jbpm/commit/f39782a9c36562bbfabfb16927caf8024bc56747 do we target that for 6.0.2 or 6.1.0 as I don't know if I shall mark this bz as modified or not?	1	1	https://issues.jboss.org/browse/RHBPMS-703
8709	RHBPMS-702	Maciej, will this make into the DR4 build?	0	0	https://issues.jboss.org/browse/RHBPMS-702
8710	RHBPMS-650	this will be actually partially available in 6.1 that will come as part of the asset management work. So when repository will be managed (meaning it will be under control of gunner for branching etc) it will have enhanced capabilities when it comes to build of the project: - build only will just verify if the assets can be build - build and install will build, package and install jar into maven repository (both local and the one hosted by business central) - build and deploy will build, package install to maven and deploy to runtime - this will be valid option only for BPMS so when doing tutoring only most of the time only two first options should be used and that will eliminate problems described in this bz. We could consider if we should still provide disable auto deploy option. So that are the thoughts on this?	1	1	https://issues.jboss.org/browse/RHBPMS-650
8711	RHBPMS-650	Justin, as part of other bz () we provided option to disable auto deploy via system property. So it will be possible run with no auto deploy but still possible to deploy via deployments panel if the property is set: -Dorg.kie.auto.deploy.enabled=false Would that cover this case or there is more to it?	1	1	https://issues.jboss.org/browse/RHBPMS-650
8712	RHBPMS-650	Thanks for the update Maciej - this flag would indeed cover our case. Will this be in product version 6.1.0?	1	0	https://issues.jboss.org/browse/RHBPMS-650

#	Requirement	Questions	Answered	Successfully Answered	URL
8713	RHBPMS-650	Jiri, this is expected behavior and works as designed. Meaning whatever has been deployed into runtime environment it will be brought up after server restart. Regardless if the version is SNAPSHOT or concrete. The difference for snapshots it that only latest snapshot of given version is deployed and not all available snapshots of given version (so called timestamped versions). So in your case you deployed 3 distinct deployment - different versions - so all of them will be deployed upon server startup. Next the property org.kie.auto.deploy.enabled does not alter already deployed versions but it allows to do not deploy when you click 'Build & Deploy' button for unmanaged repositories/projects. So if you build a project with that property set to false it will not deploy to runtime but only to maven repo (artifact repository). Whenever you don't want to have deployments present at runtime any more you much explicitly undeploy it, otherwise engine itself cannot make such a decision because it will impact running instances (if any). Does it make sense?	1	1	https://issues.jboss.org/browse/RHBPMS-650
8714	RHBPMS-551	Thanks Maciej. I had to add myself to a new kiemgmt role, why a new role? Why admin role is not enough to create a managed git repository? Unfortunately there is still no option to convert existing or cloned repositories to managed repositories. As a result the managed repository does not solve the problem. My existing or cloned repositories still have only Build&Deploy button without gui option to deploy using selected session strategy. Adding also Walter to CC. As a side effect of working on this bugzilla, I have filed a new installer BZ about kiemgmt role required to do full repository management: BZ1176597	1	0	https://issues.jboss.org/browse/RHBPMS-551
8715	RHBPMS-482	Ah, sorry, Marco. Please forget my stupid comment#3. My main point is that users may want to add their own classes/interfaces to implement internal logic, not as data models. However, DynamicJaxbContext.setupDeploymentJaxbContext() picks up *all* classes/interfaces in a project to JAXBContext.newInstance() so results in the Exception. If it's expected, we would need to document guidance like: "All Java sources in a project are considered as data models so need to follow JAXB conventions." Is it fine?	0	0	https://issues.jboss.org/browse/RHBPMS-482
8716	RHBPMS-459	Since we now have retry etc., are these checkboxes never used?	1	0	https://issues.jboss.org/browse/RHBPMS-459
8717	RHBPMS-421	How to reproduce: 0) I use JDK 7 to compile my project / start JBoss EAP/BPMS, so you need to use JDK 7 if you want to use my compiled jar... 1) first I installed EAP 6.4 + BPMS 6.1 in standalone mode (for quick setup I just use the default H2 datasource) 2) so I created a new project in business-central, say, "org.acme : myartifact : 1.0-SNAPSHOT" 3) now "New Item - Business Process", give a name "TestProcess", and import my "TestProcess.bpmn2" 4) open "TestProcess", notice the "Content" field value of the UserTask, I'm not sure how you setup this as you didn't show me your code (or any bpmn2 example...)... I guess you can setup via a process variable, or using TaskService to setup this programmatically, etc.... 5) ok now assuming we have the custom model, class CustomModel , in my maven project "model", you can find compiled jar in "target" folder in the project, you want to setup this jar as project dependency (upload the jar in "Authoring - Artifact repository" first), so that you can use this class to setup the Content field. 6) build the project in business-central 7) now import another maven project "brms61" into JBDS(I use 8.1), or Eclipse, whatever, this is the REST API client, check the server URL / USER / PASSWORD settings in the class Case01442738 8) compile this project and run Case01442738 as standalone Java program 9) it will connect to REST API and call 3 APIs: start a process of "TestProcess"; then get task list using the process id; then get the task content using task id. Now you should see this printing out the content retrieved from REST API, it's serialized object. In my testings I didn't even use "CustomModel" to set the Content field, I just use a simple string and looking into the src of Content class it saves the content object as Base64Binary format and that's basically what you'll get, and it seems you need to convert it into whatever you set , by yourself . (check my "// TODO" comments in class Case01442738) Checking the src of class "JaxbContentResponse" (response object returned by REST API "/task/{taskId}/content"), it does't seem to be JSON supported Comparing to the other "JaxbTaskSummaryListResponse" (response object returned by API "/task/query") , that one has "@JsonAutoDetect" and "@JsonTypeInfo" annotations and I've tried that it can be converted into JSON format (human-readable) So I was trying to setup simple string into Content field and couldn't get it printed out when requesting that via REST API. (though I think it should, by using something like Base64Binary decode utils.)	0	0	https://issues.jboss.org/browse/RHBPMS-421
8718	RHBPMS-374	This is already solved in master with the column picker component, that as far as I know it will not be migrated back to 6.0.x it will be in 6.1. Can I mark this as modified?	1	0	https://issues.jboss.org/browse/RHBPMS-374
8719	RHBPMS-315	Verified with BPMS-6.1.0.ER6 that using custom marshalling strategy configured via deployment descriptor inside kjar, variables can be persisted outside the process instance. Maciej, may I drop there a link to your example project you pointed me to? It illustrates the usage of marshalling strategies and deployment descriptors pretty well, Anton told me he could use an example like that.	1	1	https://issues.jboss.org/browse/RHBPMS-315
8720	RHBPMS-315	Maciej, what approach would you recommend if I wanted to retrieve variable value for a process instance which has *ended* already. For running process instance I am using: ================= Map<String,Object> params = new HashMap<String,Object>(); Person p = new Person(); p.setName("name"); params.put("processVar1", p); WorkflowProcessInstance w = (WorkflowProcessInstance) ksession.startProcess("com.sample.bpmn.hello",params); Person x = (Person)w.getVariable("processVar1"); System.out.println("var name:"+x.getName()); ================= And I don't need to care about JPAPlaceholderResolverStrategy at all. But what about completed? There is no relationship between entity Person and the process instance. So why would I ever wanted to use JPAPlaceholderResolverStrategy if for running process instance I can just use the approach above and for completed process instance it doesn't help me at all - of course unless I make sure that in my Entity class I also put variable name, and process instance id as attributes. And then I implement some custom querying logic, let's say in JPQL - but that seems like huge amount of work. I guess I am confused because I don't understand the value proposition of this JPAStrategy - mainly because the entity stored in database has absolutely no relationship with the Process Instance which persisted it and there is no out-of-the-box jBPM API which can help me to get this entity. Can you please give me some use cases / scenarios (I don't mean code) where we should recommend this option to customers? I understand what this strategy does - I don't understand what it's good for. Thanks, Anton	1	0	https://issues.jboss.org/browse/RHBPMS-315

#	Requirement	Questions	Answered	Successfully Answered	URL
8721	RHBPMS-315	Anton, main use case (in general regardless of the type of the strategy) is that it shall be used when process instance is not owner of the data. That means if data can be changed outside of the process instance then you should use variable persistence strategy. Moreover process instance id or variable name should not be part of the variable you persist externally - that simply breaks the idea of decoupled data. Imagine case where you have order entity that is stored in some order management db/system. Then why would you like to "corrupt" it with some internals of process engine/process definition that uses that data? So to recap, use variable persistence strategy when you have some other systems/components that are considered master system of that data and process instance must have access to always latest version of that data to properly process it. At the same time process instance can modify that data as that is primary responsibility of process instance - produce some results. Then don't expect that you will be able o find such data from within process instance same way you would for process variables owned by process instance, this is because strategy is all about persist and load operations, it does not provide any searching capabilities. If you need those you would have to extend the logic by custom listeners that can store (in any way) relationship between process instance and given data. Though not sure that makes much sense.... Hope this helps	0	0	https://issues.jboss.org/browse/RHBPMS-315
8722	RHBPMS-271	Yeah.. Interesting use case.. we can do that and then add an extra box in the assignments configurations for the task in the UI to add the excluded owners.	0	0	https://issues.jboss.org/browse/RHBPMS-271
8723	RHBPMS-157	Could you please elaborate little bit more on this case as I don't really understand what should happen? Do you mean next task owner as the task owner of currently active human task that process instance is waiting at?	1	1	https://issues.jboss.org/browse/RHBPMS-157
8724	RHBPMS-148	Every 'New'.. dialog has "Resource Name", change that to appropriate "XXXX Name" For instance for the new project dialog call it - "Project Name' etc.	1	1	https://issues.jboss.org/browse/RHBPMS-148
8725	RHBPMS-89	Hi Jeremy, there are two options how I can create a custom task in Designer: a) import it from some Service Repository, b) create a WID file in a project where I have to reference an icon path. Description of this issue: ---------- This issue is about a referencing of an icon from a WID file (WID can be place in various packages). There should be possibility to create an custom icon. It means that a user can add, delete or change icons in project. ========== Notes: ------ I mentioned 3 problems in comment #10: 1) WID editor contains a selector of icons - scope of this issue: If I use the selector, I want to insert the right path of the icon and Designer displays the right icon in Object repository - there should be an agreement between WID Editor and Designer 2) The folder 'global' does not allow to add an icon - scope of this issue: User can add, delete or change an icon 3) The empty project does not allow to select any icon (no predefined icons) - this is not important when an user can defined and use own icon. I put all these problems here because they can be solve in one system solution. If some problem lefts I will report new issues. Please, can you write me if something is still unclear?	1	1	https://issues.jboss.org/browse/RHBPMS-89
8726	RHBPMS-89	Hi Sona, Thanks for the clarification. I am wondering whether we should mark this BZ entry as modified now that the initial bug is (almost) fixed, and create a separate Feature Request for your notes in Comment 10 / Comment 14. In particular, your note in Comment 14 "2) The folder 'global' does not allow to add an icon..." requires a change to the "New -> Uploaded File" functionality which is a Business-Central issue rather than a Designer issue, so maybe it makes sense to make a new Feature Request BZ entry for this. If you agree that a new Feature Request should be created, will you create one? Thanks. - When I say the original bug is 'almost' fixed, there's an aspect of it that I'll work on today to complete it.	1	0	https://issues.jboss.org/browse/RHBPMS-89
8727	RHBPMS-89	Thanks Sona for your response. I've completed the fix for the initial bug so that relative and absolute icon paths work in WID files. For example: "icon" : "myicon.png" (icon in same folder as .wid file) "icon" : "com/test/myicon.png" (icon relative to .wid file location) "icon" : "../test/myicon.png" (icon relative to .wid file location) "icon" : "/HR/src/main/resources/myicon.png" (absolute path to icon) master: https://github.com/droolsjbpm/jbpm-designer/commit/68b31b993f011045b89c26747f6f9f29a47a0f56 6.2.x: https://github.com/droolsjbpm/jbpm-designer/commit/808124d608ef46b1d1955c2249dba4e5dfbe4df7	1	0	https://issues.jboss.org/browse/RHBPMS-89
8728	RHBPMS-86	(In reply to Kris Verlaenen from comment #2) > sourceRef of a data association should be a valid IDREF, so strictly > speaking, it can't be something like <sourceRef>myMap.var1</sourceRef>. A > transformation element (which is an expression inside the data association) > should be used to store the expression. > > So when using BPMN2 (both Designer and Eclipse), currently the user cannot > set expressions as source for data associations. As alternative, scripts > should be used to set a variable with the right value and then use a simple > data association to map that variable to the data input. > > We will have to consider adding support for transformation expressions in > the future. we already have support for transformation from the engine point of view. I talked to both Tihomir and Bob regarding support in the tooling though not sure it has been developed. I am quite sure eclipse modeler can show it when defined in source xml.	0	0	https://issues.jboss.org/browse/RHBPMS-86
8729	RHBPMS-7	[~ghjboss] Is this still necessary? If not, can you please close it?	0	0	https://issues.jboss.org/browse/RHBPMS-7
8730	RHBRMS-2858	Hi [~manstis], can you point me a doc that explain how to set up the workbench to load the list of possible values from an external data source. This is interesting but not fully address my requirement, since the focus of my request is dealing with a long list of values that cannot be handled in a static way (e,g, thousands of values). So the "auto complete" behaviour fits better the situation, so as far as the user fills some characters, he will get back a list of possible values in an interactive way. !img_001.png|thumbnail! A step further could be provide a complex form to filter the desired value. Maybe, this could be achieved with a mashup capability leveraging external html fragment. E.g. * address will pop up a form with 4 fields: country, city, address, zip code; starting the from to top to down the previous field narrow the actual field range. * Select a point on a map (embedding google map).	1	1	https://issues.jboss.org/browse/RHBRMS-2858
8731	RHBRMS-2810	Hi [~grgrzybek]. We need to cover more details about the instructions. See an example at: https://access.redhat.com/jbossnetwork/restricted/softwareDetail.html?softwareId=38733&product=bpm.suite&version=6.1.0&downloadType=patches Could you please use a similar approach for instruction ? Probably some steps will be similar. Thanks	1	1	https://issues.jboss.org/browse/RHBRMS-2810
8732	RHBRMS-2810	[~alessandrolt] [~grgrzybek] Would it be possible to rename one-off Maven GAVs to something like "drools-karaf-features-6.5.0.Final-redhat-7-RHBRMS-2810" and "drools-karaf-features-6.4.0.final-redhat-13-RHBRMS-2810"? This kind of suffix is often used for BxMS one-offs and I think it would work also for features files. It may be a little confusing to have two different artifacts with the same GAV. It is also problem for QE automation.	1	1	https://issues.jboss.org/browse/RHBRMS-2810
8733	RHBRMS-2750	Yes, we can close this one. Could you please close it as duplicated by RHBRMS-2743 ? Thanks Tibor.	1	1	https://issues.jboss.org/browse/RHBRMS-2750

#	Requirement	Questions	Answered	Successfully Answered	URL
8734	RHBRMS-2693	[~jsoltes] Hi, can you please advise what makes you consider this a regression; do you have a failing unit test? I only ask because we generate the DRL by iterating the data rows and construct a rule per row. If there's no data AFAIK we've never attempted to generate the static component. Can you also explain the expectation when we have (for example): {code} Person(age > 10) Address(city = {param1}) Loan(amount > 1000) {code} i.e. the first and third patterns are literals; and the second uses a template key. What about a more complex example:- {code} Person(age > 10, age <= {param1}) {code} If there are no data rows, do you expect the following to be generated? {code} Person(age > 10) {code}	1	0	https://issues.jboss.org/browse/RHBRMS-2693
8735	RHBRMS-2693	[~tirelli] This is an enhancement not a bug and definitely not a regression as originally reported. Are we doing enhancements in roll-up patches?	1	0	https://issues.jboss.org/browse/RHBRMS-2693
8736	RHBRMS-2693	[~jsoltes] As Michael mentioned, templates are supposed to be generated by iterating data. If there is no data to iterate, there is no reason to generate the source code. Having said that, the only thing I would change if possible [~manstis] is to disable the source code tab and the data tab if there is no "parameter" to interpolate in the template. What do you think? In any case, lets target 7.0 for this. There is no reason to change the behavior in a patch.	1	1	https://issues.jboss.org/browse/RHBRMS-2693
8737	RHBRMS-2693	[~tirelli] I am sorry, but I don't quite like the proposed change. Although, it would give a hint to the template user that he needs to have template key for the source generation to work, I am afraid this would put an inconsistency into how asset editors work (in means, this is the only editor, that would support this kind of behavior). I would rather prefer having no change done at all and just wait for the new Drools centric Web tooling proposals to happen (https://docs.google.com/document/d/1VYcwhoIidLlEssUHjZlGd4R7dnsTplzs698dpZdyXxA). I think there is a proposal to remove this editor, since guided data tables can now fully replace it. Or if we want to keep similar UI - another idea is to merge guided rule editor with guided template editor by just adding data table column into the guided editor... And I am ok with this happening in the 7.x series as originally proposed. Having said that, I still think that it should not be possible to add data row in guided templates when there is no template keys. My video shows, that it is possible (even though it does not generate any visible table). If something is to be disabled it is the add row button for this case. And I also think, this is a bug. What do you think?	1	1	https://issues.jboss.org/browse/RHBRMS-2693
8738	RHBRMS-2610	This JIRA lacks the necessary ACKs to apply the fix for 6.4.x Product. Can you please organise QE's ACK?	1	1	https://issues.jboss.org/browse/RHBRMS-2610
8739	RHBRMS-2413	Hi Francisco, IMHO the presence of those files does not necessarily imply that another BPMS installation has already been performed. The general idea was to allow a customer to use his own keystore. That brings in another question. What will the installer do if there are already those file present? If you think that the current behavior is sufficient and the above is a new feature or enhancement then I will not dispute that and I will change the issue type. Thank you for the feedback.	1	0	https://issues.jboss.org/browse/RHBRMS-2413
8740	RHBRMS-2413	Hi Marek, I see. I think that, if it's the case that we allow the customer to use their own keystore, we should make the following changes: 1) Make the generation of a client keystore optional through a selection on the panel. a) If the selection is to use their own keystore, we will require a path and a password to confirm access to the keystore, so the generated server.keystore certificate can be imported successfully. We will support only JKS / JCEKS type keystores here (unless you advise otherwise). b) If the selection is to have the installer generate them, we do exactly what we are doing now. In b), we can check for the existence of a keystore to avoid the failed installation (keytool fails upon file existence, so this really should be checked anyway.) How does this sound for you? Thanks, Tom	0	0	https://issues.jboss.org/browse/RHBRMS-2413
8741	RHBRMS-2413	Hi Tom! your proposal sounds great. When implemented, it will look really cool :) I'm also changing this issue to a Feature request as it looks as a far better match. Thanks! Marek	1	0	https://issues.jboss.org/browse/RHBRMS-2413
8742	RHBRMS-2413	Should this be included in 6.0.2?	1	1	https://issues.jboss.org/browse/RHBRMS-2413
8743	RHBRMS-2185	Update. Currently the A and B do not need to be added to import suggestions They are automatically picked up from the jar. In the future the goal should be to make import suggestions act as a mask for import options in the editors. Benefits: * If a jar with 1000 classes are added and only few of them are needed, the editor's import tab does not show 1000 items. * The project admin has the control to declare what facts can be used in the guided editors. An user that only has access to guided editor can not break the package by using the wrong Person fact if there are more than one available * This will unify the way imports are done. Currently imports from java.lang.* need to be added using import suggestions, in order for them to show up in guided editors. * In the future we might want to use aliases for the imports. org.test.Person could be masked as "User" or "Persona" if the users are spanish.	1	1	https://issues.jboss.org/browse/RHBRMS-2185
8744	RHBRMS-2000	(In reply to Petr Siroky from comment #0) > Description of problem: > The "-Dproductized" build is now creating distribution wars for containers > that are not supported for product, namely JBossAS 7.1.1 and WildFly 8.x. In > order to make the maintenance easier and the builds faster, I think it makes > sense to remove these assemblies from the productized build. > > I already discussed this with Mike Anstis and Kris and they agree that this > is a good idea. > > This should not involve any changes in prod. process as the resulting WARs > are not used anyway. @Ryan, can you please confirm this? > > > Version-Release number of selected component (if applicable): > 6.1.0.ER4 Yes, this is good idea. Product don't use the war for JBossAS 7.1.1, wildFly 8.x. So it's good to remove them from producitzed profile.	1	1	https://issues.jboss.org/browse/RHBRMS-2000
8745	RHBRMS-1937	See https://bugzilla.redhat.com/show_bug.cgi?id=1020925#c7 It was agreed to remove the ability to delete artifacts from the repository via the UI. Should a customer need to delete rogue artifacts it should be the role of a system administrator to delete the JAR from whatever places are necessary. The places depend upon how the customer has Maven configured/deployed.	0	0	https://issues.jboss.org/browse/RHBRMS-1937
8746	RHBRMS-1900	What is a solution for this?	0	0	https://issues.jboss.org/browse/RHBRMS-1900

#	Requirement	Questions	Answered	Successfully Answered	URL
8747	RHBRMS-1900	I'm using version 6.0.1.Final and facing this issue recently. Let me know the resolution pls. Below is my test case. Simply trying to fetch a release id of artifact deployed on WAS. ---------- String url = "http://localhost:8080/kie-drools-wb-distribution-wars-6.0.1.Final-jboss-as7.0/maven2/com/creditdecision/1.0/creditdecision-1.0.jar"; KieServices ks = KieServices.Factory.get(); KieRepository kr = ks.getRepository(); UrlResource urlResource = (UrlResource) ks.getResources().newUrlResource(url); urlResource.setUsername("*******"); urlResource.setPassword("*******"); urlResource.setBasicAuthentication("enabled"); InputStream is = urlResource.getInputStream(); KieModule kModule = kr.addKieModule(ks.getResources().newInputStreamResource(is)); KieContainer kContainer = ks.newKieContainer(kModule.getReleaseId()); System.out.println("ReleaseId loaded into local Rule Engine(KieContainer) : " +kContainer.getReleaseId()); ---------- Error: java.lang.NoSuchMethodError: com.ning.http.client.Realm$RealmBuilder.setEnconding(Ljava/lang/String;)Lcom/ning/http/client/Realm$RealmBuilder; at org.sonatype.maven.wagon.AhcWagon.openConnectionInternal(AhcWagon.java:84) at org.apache.maven.wagon.AbstractWagon.openConnection(AbstractWagon.java:116) at org.apache.maven.wagon.AbstractWagon.connect(AbstractWagon.java:218) at org.sonatype.aether.connector.wagon.WagonRepositoryConnector.connectWagon(WagonRepositoryConnector.java:345) at org.sonatype.aether.connector.wagon.WagonRepositoryConnector.pollWagon(WagonRepositoryConnector.java:385) at org.sonatype.aether.connector.wagon.WagonRepositoryConnector$GetTask.run(WagonRepositoryConnector.java:571) at org.sonatype.aether.util.concurrency.RunnableErrorForwarder$1.run(RunnableErrorForwarder.java:60) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at java.lang.Thread.run(Thread.java:745) ----------	0	0	https://issues.jboss.org/browse/RHBRMS-1900
8748	RHBRMS-1894	Pedro, could you please arrange for Form Builder's icon to be updated (and then re-assign to tsurdilo for Designer)? Thanks, Mike	0	0	https://issues.jboss.org/browse/RHBRMS-1894
8749	RHBRMS-1848	My apologies - What information do you need from me? If the question is weather we want to tighten the password restriction, then this question should be directed at PM (namely at Prakash Aradhya). Thank you.	1	1	https://issues.jboss.org/browse/RHBRMS-1848
8750	RHBRMS-1848	Has any more thought been given to this issue?	1	1	https://issues.jboss.org/browse/RHBRMS-1848
8751	RHBRMS-1848	I am OK to do it same as EAP. But lets think about it for a minute. EAP typically has Administrators as the users. For BRMS / BPMS we have 5 different kinds of users. Many of the users have simple roles like business user / analyst etc. The level of security imposed by EAP may be valid for EAP Admins. I am wondering if we really need to impose such restrictions for BxMS users of all types. Remember, users can still use passwords that are like EAP ones. How about making it configurable ? Have a setting in Userprofile to set using restrictive passwords like EAP. By default keep the current behaviour.	1	0	https://issues.jboss.org/browse/RHBRMS-1848
8752	RHBRMS-1848	Hi Marek, Can you clarify a little bit as to what is supposed to be done here? Did you want the password requirements to be the same as those in EAP? Thanks, Mansi	1	0	https://issues.jboss.org/browse/RHBRMS-1848
8753	RHBRMS-1848	Hi, I would like to give brief status and what is already included (6.2.0) about this BZ. Password restrictions looks same to me for both EAP admins and BxMS users (e.g. 8 chars long, one alphanumeric char, one digit and one non-aplhanumeric not including &). Second part is concerning passwords restrictions when using keystores, JMS etc. Here the only restriction is length more than 6. I can even use (& - char). So to sum it up. This RFE is to make password restrictions consistent across the installer including pwdvaults, ssl-keystores, etc. While EAP admin and BxMS users has this consistent, passwords for keystores and etc. doesn't. Thomas? Can you provide info about how much work will this need and if this is necessary to have, because user CAN use strong passwords, like in EAP, for keystores aswell, only difference is that if the password is weak Installer will let him use it. Do we need to enforce restrictons for these too?	1	0	https://issues.jboss.org/browse/RHBRMS-1848
8754	RHBRMS-1591	Ryan, did you remove .niogit folder from previous installations? As it might be it is still present there so kie-wb does not need to load it from remote and just uses what has already been there.	0	0	https://issues.jboss.org/browse/RHBRMS-1591
8755	RHBRMS-1454	This related to how metadata for files is stored in Guvnor for BRMS/BPMS6. Is metadata stored in an additional file on the file system (so it can be editor from Eclipse using for example EGit)?	1	1	https://issues.jboss.org/browse/RHBRMS-1454
8756	RHBRMS-1443	Hello, Can you please provide your complete DSL definition? Using a DSL with a single definition (see below) in the "mortgages" example repository the UI and generated source were OK (see below): *DSL* [when]There is an Applicant with the name of {name} whose CR is any of {rating:ENUM:Applicant.creditRating}=applicant:Applicant(name=="{name}", creditRating in {rating}) *Source code* > package org.mortgages; > > import java.lang.Number; > > rule "gre2" > dialect "mvel" > when > applicant:Applicant(name=="bananna", creditRating in ("AA","OK")) > then > end Thanks, Mike	1	0	https://issues.jboss.org/browse/RHBRMS-1443
8757	RHBRMS-1443	Mike, you are right, it works as you describe under normal conditions. The issue I encountered depends on specific DSL definition. The one I used is as follows: [when]There is a person with the name of {name} who likes {cheeseType:ENUM:Cheese.cheeseTypes}=person:Person(name=="{name}", cheese=="{cheeseType}") [when]There is a person with the name of {name} who likes any of {cheeseType:ENUM:Cheese.cheeseTypes}=person:Person(name=="{name}", cheese in {cheeseType}) I suspect the issue may be connected to the fact that the two sentences have a common prefix (I think there already was a bug related to sentences with common prefix). What to do now? We can a) revert the feature and postpone it to 6.2, b) accept it as is and provide further fixes later, c) include the additional fixes now. Depends on how severe the issue I hit is. I'm not sure ATM.	1	1	https://issues.jboss.org/browse/RHBRMS-1443
8758	RHBRMS-1422	Pedro, can you check if someone in your group can tackle this one? The core engine has a configuration option to define the severity of rule name duplication build result message: ERROR, WARN, INFO. Maybe all that is missing here is to properly catch that message and present to the user in the UI?	0	0	https://issues.jboss.org/browse/RHBRMS-1422
8759	RHBRMS-1422	Sorry, forgot to add: Did you check this in the 6 version? if so, could you please recheck this in a recent build and if the problem persists give a more detailed scenario to reproduce the defect?	0	0	https://issues.jboss.org/browse/RHBRMS-1422
8760	RHBRMS-1330	Maciej, The PR is for master, but we need to support for websphere in 6.0.2. Can you clarify what we need to do for 6.0.x? Is this PR supposed to be applied to both master and 6.0.x? Thanks	1	1	https://issues.jboss.org/browse/RHBRMS-1330
8761	RHBRMS-1330	Maciej, I applied the PR to master and cherry-picked it into 6.0.x as we discussed. Rajesh, This is required for the Websphere support in 6.0.2. Can you please review and set the acks? Thank you	1	1	https://issues.jboss.org/browse/RHBRMS-1330

#	Requirement	Questions	Answered	Successfully Answered	URL
8762	RHBRMS-1229	Adding Needinfo for Tomas, I have captured screen shots of each page of the GUI for advanced configuration but haven't been able to find any further information about configuring each thing specifically. In particular, SSL keystores, Business management LDAP, Security Domain, and JSSE. Is there community documentation I can use to develop a clear procedure for users wanting to enable the advanced configuration? Thanks, Gemma.	1	1	https://issues.jboss.org/browse/RHBRMS-1229
8763	RHBRMS-1199	Let's move this issue to 6.1 (and probably 6.2). There's no point in doing this analysis now if there's no bandwidth to implement it.	0	0	https://issues.jboss.org/browse/RHBRMS-1199
8764	RHBRMS-1061	Fixed all but #4. The contents of the pop-up can't be modified as this is chart library internals. (master) https://github.com/droolsjbpm/kie-wb-common/commit/f46d6713dbbfc19df914cec24bac399da1207782 (6.2.x) https://github.com/droolsjbpm/kie-wb-common/commit/a7d53beb6b027ba860531c4b32cbd0398c06e98c	1	1	https://issues.jboss.org/browse/RHBRMS-1061
8765	RHBRMS-988	Hi Jan Fixed in master: http://github.com/droolsjbpm/kie-wb-distributions/commit/716f25bab http://github.com/droolsjbpm/kie-wb-distributions/commit/0286d03cf Before also committing this to 6.2.x (this BZ is in the list of additional BZs that have to be tackled for 6.2.x), could you possibly check if the images are enough / acceptable, or changes are required? Thanks!	1	1	https://issues.jboss.org/browse/RHBRMS-988
8766	RHBRMS-857	Toshiya, I've pulled all translations for all of the BRMS related projects from Zanata. Some of these projects however lack any translations (e.g. https://translate.jboss.org/iteration/view/kie-wb-distributions/6.0.0). How can we get the Zanata community to complete translations? There is only so much that we, engineering, can do. Thanks, Mike	1	1	https://issues.jboss.org/browse/RHBRMS-857
8767	RHBRMS-857	Thanks. Could you please let me know (comment on this BZ) when translations are complete? I'll keep this BZ as "open" until we've merged the completed translations.	1	1	https://issues.jboss.org/browse/RHBRMS-857
8768	RHBRMS-857	Can you check if the latest translations have been committed for form modeler and BAM?	0	0	https://issues.jboss.org/browse/RHBRMS-857
8769	RHBRMS-857	Toshiya, we need your input here. Can we close this?	1	1	https://issues.jboss.org/browse/RHBRMS-857
8770	RHBRMS-701	Joe, could you please give a view on how to improve the popup? Thanks.	0	0	https://issues.jboss.org/browse/RHBRMS-701
8771	RHBRMS-658	Hi Sir, Can you please give an example, how to specify more than one patterns. I tried giving as below but its not classes are not imported: com.jcms.db.entity.** com.aml.**	0	0	https://issues.jboss.org/browse/RHBRMS-658
8772	RHBRMS-658	(In reply to Sona Mala from comment #0) > Describe the issue: > ---------- > Can you document the existance of a new file which name is > 'package-names-white-list' which was created by bz 1155371 see comment #9. > > The location of this file is the root directory: > + src > + package-names-white-list > + pom.xml > + project.imports > > This file is visible only in Repository view of Project Explorer in kie-wb > (of course, it is saved into git repository as well). The default content is > empty. > > An user defines facts which are visible in data model (authoring part). For > example, it filters a drop-down menu for importing a fact (Guided Rule > Editor, tab "Config", click on button "New item"). It may have an effect on > engine (for customers which build own GUI part). > > The language of file defines a packages which are allows to be displayed. > For example: > #1 org.jbpm > #2 org.jbpm.* > #3 org.jbpm.** > > > The first line allows these facts: > - org.jbpm.MyFact > - org.jbpm.MyFact2 > > The second line allows these facts: > - org.jbpm.anypackage.MyFact > - org.jbpm.anypackage2.TestFact > > The third line allows these facts: > - org.jbpm.MyFact > - org.jbpm.anypackage2.TestFact > - org.jbpm.firstpackage.secondpackage.MyTestFact > > Developers may set a group of facts which are visible for analyst (used in > rules). > > Section Number and Name: > ---------- > I am not sure which guide or section is the best for it. I suggest User > Guide, "3. Setting up a new project" chapter. Not sure, if this feature > should be present in Developer Guide as well. Hi Sir, Can you please give an example, how to specify more than one patterns. I tried giving as below but its not classes are not imported: com.jcms.db.entity.** com.aml.**	0	0	https://issues.jboss.org/browse/RHBRMS-658
8773	RHBRMS-650	Alexander, this BZ only lists commits into the master branch. Did this also land in 6.0.x, or shall I reopen?	1	1	https://issues.jboss.org/browse/RHBRMS-650
8774	RHBRMS-591	Joe, could you please give a view on how to improve the popup? Thanks.	1	0	https://issues.jboss.org/browse/RHBRMS-591
8775	RHBRMS-571	Mike, can you take a look at this issue and see if ti can be fixed in 6.0 ? The work around is a little frustrating.	1	0	https://issues.jboss.org/browse/RHBRMS-571
8776	RHBRMS-560	Zuzana, can you please be more specific about what needs to be done?	0	0	https://issues.jboss.org/browse/RHBRMS-560
8777	RHBRMS-464	(In reply to Kris Verlaenen from comment #2) > kie-tomcat-integration is already in the engine zip Not in ER5. Is this something new (as in after ER5 tag)?	1	0	https://issues.jboss.org/browse/RHBRMS-464
8778	RHBRMS-456	I have a few questions/remarks: 1. I believe that drools-karaf-features should be distributed also as the xml file in addition to the drools-karaf-features-6.2.0.Final-redhat-2.jar, i.e. we need to include the following artifact: <dependency> <groupId>org.drools</groupId> <artifactId>drools-karaf-features</artifactId> <version>${droolsjbpm.version}</version> <type>xml</type> <classifier>features</classifier> </dependency> 2. We can discuss if the drools-karaf-features-6.2.0.Final-redhat-2.jar should be there, since it contains only the drools-karaf-features xml file. 3. What about transitive dependencies, e.g. for kie-camel, should we distribute them in this zip? Now some of them are not available even in the Maven repo - please see BZ 1196257. 4. Regarding the requirement 3 from the Description: "3. Split the build script for said components to allow this bundle to have an independent release lifecycle. This means that the version number for theses components will be different from the rest of the components. This will allow us to release a new version of this "integration package" every time that one becomes necessary." We have moved drools-karaf-features.xml, which defines versions of Drools/jBPM components for deploying into Fuse, into the camel-integration package. If I understand correctly, the version of drools-karaf-features.xml might not be anymore aligned with BRMS/BPMS version which can cause problems when: - deploying BRMS/BPMS version which has no corresponding camel-integration pack release into Fuse - using integration pack with a different BRMS/BPMS core components version I would suggest excluding kie-camel (and related) features from drools-karaf-features.xml, which should remain distributed with BRMS/BPMS, and creating a new bxms-integration-pack-features.xml defining the integration pack features such as kie-camel. What do you think?	1	1	https://issues.jboss.org/browse/RHBRMS-456
8779	RHBRMS-456	Ryan, could you please put the transitive dependencies into lib subdirectory to be consistent with jboss-brms-6.1.0.GA-engine.zip and jboss-bpmsuite-6.1.0.GA-engine.zip? The zip also contains two versions of the same lib, probably because groupId was changed between those two versions: - stax-api-1.0.1.jar - stax-api-1.0-2.jar Thanks!	0	0	https://issues.jboss.org/browse/RHBRMS-456
8780	RHBRMS-432	Proposing for 6.4.0 to consider for documentation. [~tirelli] please recommend Engineering's proposal (can a warning message be shown in the UI?).	0	0	https://issues.jboss.org/browse/RHBRMS-432
8781	RHBRMS-397	Can we automatically import classes from java.lang? Or some similar solution?	1	1	https://issues.jboss.org/browse/RHBRMS-397

#	Requirement	Questions	Answered	Successfully Answered	URL
8782	RHBRMS-126	I've not heard anything to suggest a change in name to "Decision Logic Editor". Please let me know what decision is made and we can update the text in the UI accordingly. What comes of "Guided Rule Template" too? This is effectively the same editor but allows use of "template keys". Should this be renamed "Decision Logic Template Editor"? Await you advices before proceeding.	0	0	https://issues.jboss.org/browse/RHBRMS-126
8783	RHBRMS-126	This ticket did not make 6.0.0.GA cut. Will it remain as "Guided Rule Editor", or should we change this for 6.0.1.GA?	0	0	https://issues.jboss.org/browse/RHBRMS-126
8784	RHBRMS-126	Isn't decision logic editor the editor in designer to define constraints for gateways? This imho doesn't have anything to do with the guided rule editor. This editor is simply called "editor" in the UI tooling (when you try to edit a constraint for a gateway by selecting the sequence flow and editing the constraint property). If the name needs to be changed, this should imho only happen in the documentation.	1	0	https://issues.jboss.org/browse/RHBRMS-126
8785	RHBRMS-114	Pedro, is this something Juanma could do short-term? "The image slider, if it stays at all, should at least switch between different images - using the same one over and over looks strange. It might perhaps contain links to demos."	0	0	https://issues.jboss.org/browse/RHBRMS-114
8786	RHBRMS-77	I think appropriated document this feature to allow users tune their environments based in this new property. Could it be documented please ?	1	0	https://issues.jboss.org/browse/RHBRMS-77
8787	JBBUILD-720	Nick, Could I accomplish the same thing by just changing the job command in your first code block above to what is below (not sure of syntax)? Could I get away from having a custom script then? Thanks. The lines of yours I modified are 4, 8, 9. The new lines I added are 5, 6, 7, 8. 1. mkdir -p ${WORKSPACE}/build/publish 2. cd ${WORKSPACE}/build/publish 3. rm -f ${WORKSPACE}/build/publish/publish.sh 4. wget http://https://github.com/ModeShape/modeshape-tools/scripts/publish.sh 5. export RELEASETYPE="development" 6. export ECLIPSERELEASE="juno" 7. export PROJECTVERSION=3.3.0.Beta5" 8. export SITE_DIR="modeshape@filemgmt.jboss.org:/downloads_htdocs/modeshape/tools/updates/" 9. export DESTINATION="${SITE_DIR}/${RELEASETYPE}/${ECLIPSERELEASE}/${PROJECTVERSION}/" 10. . ${WORKSPACE}/build/publish/publish.sh tools/tmp	1	1	https://issues.jboss.org/browse/JBBUILD-720
8788	JBBUILD-720	[~elvisisking] Any reason this issue is still open? What else do you need from me?	1	1	https://issues.jboss.org/browse/JBBUILD-720
8789	JBBUILD-710	The re-indexing completed successfully. Geoffrey, can you verify that the search is working correctly?	1	1	https://issues.jboss.org/browse/JBBUILD-710
8790	JBBUILD-689	What groupIds will be served from this repository? I'm not able to browse it to see what is there. Does it include valid checksums?	1	1	https://issues.jboss.org/browse/JBBUILD-689
8791	JBBUILD-683	Hi Manik, Can you tell me which groupIds will come from this repo? Just org.springframework?	1	1	https://issues.jboss.org/browse/JBBUILD-683
8792	JBBUILD-683	Should we be using this URL instead? http://maven.springframework.org/milestone The most recent docs I found about the various spring repos are from this blog post: http://blog.springsource.com/2009/12/02/obtaining-spring-3-artifacts-with-maven/ An older blog post points to the amazon URL http://blog.springsource.com/2007/09/18/maven-artifacts-2/	0	0	https://issues.jboss.org/browse/JBBUILD-683
8793	JBBUILD-679	Can you provide more information on this? Are you asking for a new proxy repository, or do you just want to upload a few jars?	1	1	https://issues.jboss.org/browse/JBBUILD-679
8794	JBBUILD-672	Hi Paul, Theres any idea of when this will be finished? Thanks for your attention.	1	1	https://issues.jboss.org/browse/JBBUILD-672
8795	JBBUILD-667	It looks like this is a artifactory repository and I can't browse the files. Is there a login required for this repo? What groupIds will be coming from this repo?	1	1	https://issues.jboss.org/browse/JBBUILD-667
8796	JBBUILD-666	If this is only needed for testing, maybe it can be added to a QA repository? My concern with adding it to the main JBoss repo is that it will eventually be deployed to central, and could cause a conflict with our repo. It seems like we need some temporary location for testing. One thing I could do is deploy it to a staging repo, and then just never promote it. That way the staging repo can be used for testing it. Does that sound ok?	1	1	https://issues.jboss.org/browse/JBBUILD-666
8797	JBBUILD-664	Thanks Paul! Could I get you to add the groupIds org.fedorahosted.openprops and org.fedorahosted.tennera to the routing rule, please?	1	1	https://issues.jboss.org/browse/JBBUILD-664
8798	JBBUILD-664	I think there was a problem with the way the routing was configured. Sean, can you try again? It should be working now.	1	1	https://issues.jboss.org/browse/JBBUILD-664
8799	JBBUILD-664	Thank you Paul and Graeme; that's looking great. The groups all work, including Staging, except Public Thirdparty Releases Group. Was that one meant to include Flies? Either way, I don't think it will cause a problem. I see the indexing is a known problem area: https://jira.jboss.org/browse/JBBUILD-663	1	1	https://issues.jboss.org/browse/JBBUILD-664
8800	JBBUILD-639	detectJavaApiLink - Is the default not working correctly for this one? It should put a link matching the compile source version. doctitle - The default is "${project.name} ${project.version} API". Is there a reason to remove API from the title? bottom - The only change I see in this one compared to the default is removing the project inception date, is there a reason for this? I would prefer to leave the APIViz configuration up to the individual projects in case some projects don't want this.	1	1	https://issues.jboss.org/browse/JBBUILD-639
8801	JBBUILD-638	Is this something that should be run for all JBoss projects? It seems more like something that should be left for each project to configure. Also, is there some documentation available for this plugin?	1	1	https://issues.jboss.org/browse/JBBUILD-638
8802	JBBUILD-638	No, not run, but it would be nice to not have to do this bit of configuration in every project. Is that possible?	1	1	https://issues.jboss.org/browse/JBBUILD-638
8803	JBBUILD-635	Fixed in r500. And deployed a snapshot. Do you think we should also add the index generation to the source jars? http://fisheye.jboss.org/changelog/Maven/?cs=500	1	1	https://issues.jboss.org/browse/JBBUILD-635
8804	JBBUILD-630	Hi Stan, I would prefer not to add this repository since it's just a single project snapshot repo. Can you just add this to your POM or settings file?	1	1	https://issues.jboss.org/browse/JBBUILD-630
8805	JBBUILD-630	I take it that we wouldn't want to proxy all sourceforge snapshots? I need my project to build/run with Hudson. So using the settings file won't help. I already tried adding it to the main JSFUnit pom, but it didn't work. I figured that this had something to do with extending the latest jboss-parent pom. Can I still use jboss-parent and still go around Nexus to get the HtmlUnit snapshots?	1	0	https://issues.jboss.org/browse/JBBUILD-630
8806	JBBUILD-629	It seems that the staging system in Nexus doesn't produce events that include comments from staged repositories, so there's currently no way to gain access to that information for the purposes of audit logging. I talked to the Nexus folks, and they said we need to file a feature request against the commercial Nexus JIRA instance to get this done. They seemed interested in the idea of firing events on completion of different staging actions, but need the JIRA to work against. I don't have access to this JIRA system, so can you file the feature request and then we'll go from there?	1	0	https://issues.jboss.org/browse/JBBUILD-629
8807	JBBUILD-605	Does this repository work? I get a 404 when I try to access it.	1	0	https://issues.jboss.org/browse/JBBUILD-605
8808	JBBUILD-601	I have added the scala-tools, jclouds, and spy repository to nexus and they are now included in the group URL. The java.net repo was already available. I didn't add the oracle repo because it doesn't seem to be a valid Maven repository. Can you verify that this is the correct URL?	1	1	https://issues.jboss.org/browse/JBBUILD-601
8809	JBBUILD-601	Also I don't see the jclouds, scala, etc. artifacts in https://repository.jboss.org/nexus/content/groups/public/ as yet. Does it just need some time to sync?	1	1	https://issues.jboss.org/browse/JBBUILD-601

#	Requirement	Questions	Answered	Successfully Answered	URL
8810	JBBUILD-601	So what's the "correct" solution re: Oracle? Pull in the artifacts we need and deploy to our own third-party repo?	1	1	https://issues.jboss.org/browse/JBBUILD-601
8811	JBBUILD-507	Should there be a separate distribution for each subproject? Or a single project that creates a distribution zip that includes all the subprojects?	1	1	https://issues.jboss.org/browse/JBBUILD-507
8812	JBBUILD-507	Ok, so taking the main MC project as an example, there is already a source and dist assembly in the "build" directory (maybe this should be changed to "dist"?). The assembly config is currently in a profile, so you have to call "mvn -Passembly install" for it to run. Should this to run by default for each build? The src assembly looks ok to me as is, it's just a zip of the whole project minus the target directories. Should the dist zip include dependency jars in addition to the module jars?	0	0	https://issues.jboss.org/browse/JBBUILD-507
8813	JBBUILD-489	I removed the thirdparty directory of my Branch_5_x workspace. performed svn update. However, building via build.sh results in a ConnectionException. The attempt eventually times out but it is not clear initially that will happen. This is also being seen on Hudson runs. Is this a known issue? I don't typically clean thirdparty so not sure if it recently started. createthirdparty: [java] [INFO] Scanning for projects... [java] WAGON_VERSION: 1.0-beta-2 [java] [INFO] ---------- [java] [INFO] Building JBoss Application Server Thirdparty Build [java] [INFO] task-segment: [generate-resources] [java] [INFO] ---------- [java] [INFO] [enforcer:enforce {execution: ban-bad-dependencies}] [java] [INFO] [enforcer:enforce {execution: enforce-java-maven-versions}] [java] [INFO] [enforcer:enforce {execution: enforce-plugin-versions}] [java] [INFO] snapshot org.jboss.jbossas:jboss-as-component-matrix:5.1.0-SNAPS[java] [WARNING] [java] java.net.ConnectException: Connection timed out HOT: checking for updates from snapshots.jboss.org [java] [INFO] [enforcer:enforce {execution: ban-app-server-modules}] [java] [INFO] [jboss-license:download-licenses {execution: default}] [java] [INFO] Downloading licenses [java] [WARNING] [java] java.net.ConnectException: Connection timed out [java] at java.net.PlainSocketImpl.socketConnect(Native Method) [java] at java.net.PlainSocketImpl.doConnect(PlainSocketImpl.java:333) [java] at java.net.PlainSocketImpl.connectToAddress(PlainSocketImpl.java:195) [java] at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:182) [java] at java.net.Socket.connect(Socket.java:520) [java] at java.net.Socket.connect(Socket.java:470) [java] at sun.net.NetworkClient.doConnect(NetworkClient.java:157) [java] at sun.net.www.http.HttpClient.openServer(HttpClient.java:388) [java] at sun.net.www.http.HttpClient.openServer(HttpClient.java:523) [java] at sun.net.www.http.HttpClient.<init>(HttpClient.java:231) [java] at sun.net.www.http.HttpClient.New(HttpClient.java:304) [java] at sun.net.www.http.HttpClient.New(HttpClient.java:321) [java] at sun.net.www.protocol.http.HttpURLConnection.getNewHttpClient(HttpURLConnection.java:813) [java] at sun.net.www.protocol.http.HttpURLConnection.plainConnect(HttpURLConnection.java:765) [java] at sun.net.www.protocol.http.HttpURLConnection.connect(HttpURLConnection.java:690) [java] at sun.net.www.protocol.http.HttpURLConnection.getInputStream(HttpURLConnection.java:934) [java] at java.net.URL.openStream(URL.java:1007) [java] at org.jboss.maven.plugins.license.DownloadLicensesMojo.execute(DownloadLicensesMojo.java:134) <snip> [java] [WARNING] [java] java.net.ConnectException: Connection timed out	1	0	https://issues.jboss.org/browse/JBBUILD-489
8814	JBBUILD-447	Is the following failure somehow related to this issue? [ERROR] BUILD ERROR [INFO] ---------- [INFO] An Ant BuildException has occured: The following error occurred while executing this line: Target "dist" does not exist in the project "JBoss/Testsuite". and don't know wheter create new JIRA issue, or not?	1	1	https://issues.jboss.org/browse/JBBUILD-447
8815	JBBUILD-442	Ready for an on-going battle until the ant build is deleted ;) Still get some compile errors in varia. --- /home/pether/sandbox/workspaces/jboss/trunk/varia/src/main/org/jboss/varia/stats/DataSourceInterceptor.java:[67,7] org.jboss.varia.stats.DataSourceInterceptor is not abstract and does not override abstract method isWrapperFor(java.lang.Class<?>) in java.sql.Wrapper .. .. .. --- Will try to find out why later on. Have some changes in build/pom.xml for missing items in the dependency management section that showed up once i added embedded to trunk/pom.xml modules section again. Also created pom.xml for testsuite and spring-int that I will do more work on later..	1	1	https://issues.jboss.org/browse/JBBUILD-442
8816	JBBUILD-442	Pether, what else needs to be done on this issue? Are you able to generate the eclipse files now?	1	1	https://issues.jboss.org/browse/JBBUILD-442
8817	JBBUILD-437	I thought about adding this option when I added the groupId one. But it seemed like it could lead to confusion about where matching artifacts are. Can we just change the build scripts that depend on the old name? Or will that cause too many problems? It's not a big deal to add this, I just want to make sure it is the right behaviour.	0	0	https://issues.jboss.org/browse/JBBUILD-437
8818	JBBUILD-418	The maven issue http://jira.codehaus.org/browse/MNG-2045 has been applied, and will be released with 2.0.8. The app server maven builds will be dependent on the 2.0.8 release of maven.	0	0	https://issues.jboss.org/browse/JBBUILD-418
8819	JBBUILD-386	I think this issue is out of date now. Can I close this?	0	0	https://issues.jboss.org/browse/JBBUILD-386
8820	JBBUILD-316	Ruel, didn't you implement some snapshot related feature recently? Does that impact this request?	1	0	https://issues.jboss.org/browse/JBBUILD-316
8821	JBBUILD-316	Maven includes this feature in the release plugin. Does this still need to be added to the ant based build system?	0	0	https://issues.jboss.org/browse/JBBUILD-316
8822	JBBUILD-298	I've added it at: jboss/backport-concurrent/2.1.0.GA/lib/jboss-backport-concurrent.jar Renamed to jboss-backport-concurrent.jar, to distinguish, so this needs changing in jbossretro, too. I guess we can use the 3rd version digit to record our modifications (if any) on the base 2.1 library. Not sure though what do use in the component-info.xml for the licence field. Do we maintain the "creative commons public domain dedication" license?	0	0	https://issues.jboss.org/browse/JBBUILD-298
8823	JBBUILD-270	I created an ejb3 pom as part of issue JBBUILD-411. Is it ok to close this one?	1	0	https://issues.jboss.org/browse/JBBUILD-270
8824	JBBUILD-244	I think I need some help with reproducing the bug. Could anyone be so king ang give me the list of steps I should take to achieve this? Thanks!	1	1	https://issues.jboss.org/browse/JBBUILD-244
8825	JBBUILD-188	Are the remaining subtasks blockers for dealing with JBAS-2510 ?	0	0	https://issues.jboss.org/browse/JBBUILD-188
8826	JBBUILD-177	What is this issue about? It seems out of date.	1	0	https://issues.jboss.org/browse/JBBUILD-177
8827	JBBUILD-158	The maven assembly plugin can be used to package dependencies up with a project. It can extract and rejar the dependencies together with your project files. Is this something we want to do?	0	0	https://issues.jboss.org/browse/JBBUILD-158
8828	JBBUILD-113	Can you expand on this? How are we currently not satisyfing this feature? Apologies if we've discussed it, but this bug is rather brief.	1	0	https://issues.jboss.org/browse/JBBUILD-113
8829	JBBUILD-98	You cannot do that, unless there is mechanism to disable it (e.g. an unless="property"). What happens if somebody is working offline?	1	1	https://issues.jboss.org/browse/JBBUILD-98

#	Requirement	Questions	Answered	Successfully Answered	URL
8830	JBBUILD-95	When integrating this with the existing build I found some problems. Namely, when the build is started, this file is not created yet, so ANT will fail. This is due to the defintion of the import in the build/build.xml file. (<!ENTITY buildmagic SYSTEM "../tools/thirdparty/libraries.ent">). Scott's suggestion in the build forum is as follows: The libraries are only used by the main build when creating the dist. They have no purpose for building thirdparty or compiling the modules. Couldn't we simply break out the install target into a seperate build-dist.xml file that contains the install and related targets? Unless you do this, the initial build will be using the obsolete ../tools/buildmagic/etc/libraries.ent version and would fail anyway when the incorrect libraries were used. This is correct. The install target is actually many targets. The existing buildmagic system uses a process such that if a target exists such as : <target name="_module-j2ee-most">, then it will be executed during the most build. For most modules, targets such as these are defined. The caveat is that they must be present in the main build file, otherwise the buildmagic system will just skip over them. These targets are where the deployment operations take place (copying thirdparty files to certain places). I have implemented this locally and made it work. Changes are as follows: Build/build.xml 1) removal of reference to libraries.ent 2) properties for output structure moved to external file (they will need to be used by other xml files) 3) all module targets now call targets defined in file build-dist.xml (this removes the dependency on the libraries.ent file) Module build files 1) change the location of libraries.ent file build-distr.xml 1) this file contains all of the targets which copy thirdparty files to output directories 2) it is also where the libraries.ent location is defined install-structure.properties 1) contains properties which define the output structure, used by build.xml and build-distr.xml	0	0	https://issues.jboss.org/browse/JBBUILD-95
8831	JBBUILD-85	what's the difference between this and JBBUILD-79?	1	1	https://issues.jboss.org/browse/JBBUILD-85
8832	JBBUILD-76	I also had the same cvs issue, which the forge admins resolved. I would like to double-check what you wish to add to the thirdparty directory. The file you attached to this issue is a source structure with a number of files included. Is this what you wish to add, or are you wanting to include a single jar? The current thirdparty lib for jboss-cache currently has a file, JBossCache-1.2.3.zip, if the zip file is included, do you want it to be named the same as this?	1	1	https://issues.jboss.org/browse/JBBUILD-76
8833	JBBUILD-69	Can I close this issue? It seems to be out of date.	0	0	https://issues.jboss.org/browse/JBBUILD-69
8834	JBBUILD-60	I believe this is complete. Can I close this issue?	0	0	https://issues.jboss.org/browse/JBBUILD-60
8835	JBBUILD-59	Is there anything remaining to be done on this issue? I believe logging has been separated into the common project.	0	0	https://issues.jboss.org/browse/JBBUILD-59
8836	JBBUILD-58	Is there anything else that needs to be done for this?	0	0	https://issues.jboss.org/browse/JBBUILD-58
8837	JBBUILD-48	Has the circular dependency between the server and security modules been broken to do this? If not, it should be. All of the org.jboss.security.* classes should be in the security module.	0	0	https://issues.jboss.org/browse/JBBUILD-48
8838	JBBUILD-39	JMX creates 3 artifacts and puts them in the etc directory under test. Start.jar, target.mjar and MyMbeans.jar. Need to investigate why and or implement.	0	0	https://issues.jboss.org/browse/JBBUILD-39
8839	JBBUILD-28	Verified with 1.4 build. The jbossha.jar will probably need to be adjusted though. From Ben: Ruel, These are deprecated already with the new http session replcation that i did. Please remove them. Thanks, -Ben ---------- From: Ruel Loehr Sent: Monday, August 01, 2005 10:18 AM To: Ben Wang Subject: cluster build file Ben, I am working on a next-gen build system for jboss. In the cluster build I see the following: <!-- Build the ClusteredHttpSessionEB.jar --> <mkdir dir="${build.lib}/ClusteredHttpSessionEB.jar"/> <copy todir="${build.lib}/ClusteredHttpSessionEB.jar" filtering="no"> <fileset dir="${build.classes}"> <include name="org/jboss/ha/httpsession/beanimpl/interfaces/**"/> <include name="org/jboss/ha/httpsession/beanimpl/ejb/**"/> </fileset> <fileset dir="${source.resources}/ClusteredHTTPSessionBean"> <include name="META-INF/**"/> </fileset> </copy> <!-- Build the jbossha-httpsession.sar --> <mkdir dir="${build.lib}/jbossha-httpsession.sar"/> <copy todir="${build.lib}/jbossha-httpsession.sar" filtering="no"> <fileset dir="${build.lib}"> <include name="ClusteredHttpSessionEB.jar/**"/> </fileset> <fileset dir="${build.classes}"> <include name="org/jboss/ha/httpsession/beanimpl/interfaces/**"/> <include name="org/jboss/ha/httpsession/server/**"/> </fileset> <fileset dir="${source.resources}/ClusteredHTTPSessionService"> <include name="**"/> </fileset> </copy> 2 directories are created with a sar and jar in the name. How is this used? Is something dependent upon the names having the sar and jar in them?	1	0	https://issues.jboss.org/browse/JBBUILD-28
8840	JBBUILD-5	I believe this issue is out of date. Can I close this?	0	0	https://issues.jboss.org/browse/JBBUILD-5
8841	JBCACHE-1546	when runign with jgroups 2.6.12 GA: Tests run: 3772, Failures: 6, Errors: 0, Skipped: 0, Time elapsed: 587.558 sec <<< FAILURE! Results : Failed tests: testSTWithWritingTxThread(org.jboss.cache.statetransfer.NBSTCacheLoaderTest) testSTWithThirdWritingNonTxCache(org.jboss.cache.statetransfer.NBSTCacheLoaderTest) testConfigWarnings(org.jboss.cache.config.parsing.SampleConfigFilesCorrectnessTest) testStateTransfer(org.jboss.cache.statetransfer.TotalStateTransfertConcurrencyTest) testAbortBeforeCompletion(org.jboss.cache.transaction.pessimistic.AbortionTest) testSTWithWritingNonTxThread(org.jboss.cache.statetransfer.NonBlockingStateTransferTest) Tests run: 3772, Failures: 6, Errors: 0, Skipped: 0	0	0	https://issues.jboss.org/browse/JBCACHE-1546
8842	JBCACHE-1464	Marking this as high complexity since there is no API call in the cache loader interface for bulk gets at this time, nor can we "guess" which Fqns to retrieve. At best, we could implement a mechanism where we could enqueue get calls in the cache loader implementation and flush this queue as a single batch get every so often. This time interval would be user definable and also be pretty short to provide reasonable response times. Perhaps a second or so.	1	0	https://issues.jboss.org/browse/JBCACHE-1464
8843	JBCACHE-1416	Do you know if this affects 3.x as well? The startup/shutdown code in 3.x has changed.	1	1	https://issues.jboss.org/browse/JBCACHE-1416

#	Requirement	Questions	Answered	Successfully Answered	URL
8844	JBCACHE-1416	Hello, I don't know but it is easy to test it if the api is about the same : For example you can just launch the cache#stop and cache#destroy method in a separate thread that will starts after the webapp is undeployed (5s in this code). For example : public void destroy() { new Thread() { @Override public void run() { try { Thread.sleep(5000); } catch (InterruptedException e) { e.printStackTrace(); } cache.stop(); cache.destroy(); } }.start(); } generate this stacktrace in the log each time : Exception in thread "Thread-11" java.lang.NullPointerException at fr.arkhineo.ntp.controller.ApplicationLifeCycle$1.run(ApplicationLifeCycle.java:90) 29 sept. 2008 19:28:21 org.apache.catalina.loader.WebappClassLoader loadClass INFO: Illegal access: this web application instance has been stopped already. Could not load org.apache.log4j.spi.VectorWriter. The eventual following stack trace is caused by an error thrown for debugging purposes as well as to attempt to terminate the thread which caused the illegal access, and has no functional impact. java.lang.IllegalStateException at org.apache.catalina.loader.WebappClassLoader.loadClass(WebappClassLoader.java:1273) at org.apache.catalina.loader.WebappClassLoader.loadClass(WebappClassLoader.java:1233) at java.lang.ClassLoader.loadClassInternal(ClassLoader.java:319) at org.apache.log4j.spi.LoggingEvent.<init>(LoggingEvent.java:154) at org.apache.log4j.Category.forcedLog(Category.java:388) at org.apache.log4j.Category.log(Category.java:853) at org.apache.commons.logging.impl.Log4JLogger.error(Log4JLogger.java:193) at org.jgroups.protocols.Discovery$PingSenderTask$1.run(Discovery.java:389) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:441) at java.util.concurrent.FutureTask$Sync.innerRunAndReset(FutureTask.java:317) at java.util.concurrent.FutureTask.runAndReset(FutureTask.java:150) at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.access$101(ScheduledThreadPoolExecutor.java:98) at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.runPeriodic(ScheduledThreadPoolExecutor.java:181) at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.run(ScheduledThreadPoolExecutor.java:205) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:885) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:907) at java.lang.Thread.run(Thread.java:619)	1	0	https://issues.jboss.org/browse/JBCACHE-1416
8845	JBCACHE-1389	Is this irrelevant if OL is deprecated in 3.0?	1	0	https://issues.jboss.org/browse/JBCACHE-1389
8846	JBCACHE-1387	Can this be closed?	1	1	https://issues.jboss.org/browse/JBCACHE-1387
8847	JBCACHE-1286	If we were to include this in 2.2.0, could you add documentation and the patch/test to trunk, as well as any additional stuff like pooled HTTP connections? Couldn't you test this based on a stub that simulates the HTTP responses from S3, rather than an S3 account?	1	1	https://issues.jboss.org/browse/JBCACHE-1286
8848	JBCACHE-1250	Do you have a specific use case in mind for this? What sort of prio are we talking?	1	1	https://issues.jboss.org/browse/JBCACHE-1250
8849	JBCACHE-1250	? Didn't see anything specially relevant to my points above, regarding the cost of doing clearChildren() internally, and the risk of a concurrent insert if lockParent is not used. Did I miss something?	1	1	https://issues.jboss.org/browse/JBCACHE-1250
8850	JBCACHE-1218	If you feel like taking this on for 2.2, sure.	1	1	https://issues.jboss.org/browse/JBCACHE-1218
8851	JBCACHE-1197	And where does this belong? Current English docs reside in subversion alongside the code. Not sure if we want all translations here as well since translations may be out of sync or a couple of releases behind. E.g., the code is at version 2.1.0.Beta1, (and so is the English documentation) but given that we cut a public release ever couple of months, there may not be time for translations to catch up. What would you suggest here?	1	0	https://issues.jboss.org/browse/JBCACHE-1197
8852	JBCACHE-1197	Hi Manik, Sorry, there mate, can you assign this task to me? We are in process of recursing all the hibernate content into PO format and have used a new layout that keeps the translation in sync for the 4 AS books. This jira issue is a placeholder for me to ensure it stays in sync with your english. As soon as the tech review is done for AS422 I'll move onto this. http://jira.jboss.com/jira/browse/JBAS-4696 - Mike	0	0	https://issues.jboss.org/browse/JBCACHE-1197
8853	JBCACHE-1191	Has there been any progress on this?	1	1	https://issues.jboss.org/browse/JBCACHE-1191
8854	JBCACHE-1175	I coded it, but it's untested. It gets a bit complex in the OPTIMISTIC case, where if there's an implicit transaction I have to ensure that the async requirement is stored in the tx entry so the call gets replicated async when the implicit transaction commits. Use case is where they're trying to use a single cache for all 2nd level caching. Cache has to be REPL_SYNC for this to apply, as Hibernate will reject using INVALIDATION for the timestamps cache. So, currently a fairly narrow use case. But, it's possible I could eventually get INVALIDATION to work reasonably for timestamps, in which case the use case would be quite a bit broader.	1	1	https://issues.jboss.org/browse/JBCACHE-1175
8855	JBCACHE-1153	Identified by a boolean flag on the node?	1	0	https://issues.jboss.org/browse/JBCACHE-1153
8856	JBCACHE-1134	The difference between the classes extending DelegatingCacheLoader and the classes extending AbstractDelegatingCacheLoader is: Classes extending DelegatingCacheLoader are standalone classloaders, whereas the ones extending AbstractDelegatingCacheLoader are decorators of other class loaders. Maybe decorators, such as asynchronous or singleton store could one day become standalone cache loaders and stack/chaim them up/down other standard cache loaders? I'll open a discussion in the design forums soon.	1	1	https://issues.jboss.org/browse/JBCACHE-1134
8857	JBCACHE-1127	Isn't this the same as http://repository.jboss.com/maven2/jboss/jboss-j2ee/4.2.0.GA/ save for the name? 4.2.0.GA should be EE5 based, right?	1	1	https://issues.jboss.org/browse/JBCACHE-1127
8858	JBCACHE-1023	JBoss Forum reference should be a forum, not mail thread list. Anyway, why not have the Channel part of the RuntimeConfig object? This would result in a lot less number of changes throughout.	0	0	https://issues.jboss.org/browse/JBCACHE-1023
8859	JBCACHE-1023	Manik, I've been doing other stuff in this area and will wrap it up tomorrow, so why don't I do this at the same time?	1	1	https://issues.jboss.org/browse/JBCACHE-1023
8860	JBCACHE-992	This shouldn't affect the output of this JIRA, but if we prove that JBC builds and runs on JDK 6, bear in mind this: As far as AS is concerned, Java 6 is not currently supported, at least until the following issue has been resolved or the code is changed (I think this might happen in AS 5.0): Email to JBoss-Dev, Subject: Array class loading regression in jdk6 From: Scott M Stark <scott.stark@jboss.com> To: JBoss.org development list <jboss-development@lists.jboss.org> Reply-To: JBoss.org development list <jboss-development@lists.jboss.org> Date: 17 Dec 2006 - 6:05am An issue where ClassLoader.loadClass(String) is not resolving arrays is showing up as users start to use jdk6: http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=6434149 Exception in thread "main" java.lang.ClassNotFoundException: [Ljava.lang.String; at java.net.URLClassLoader$1.run(URLClassLoader.java:200) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:188) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:276) at java.lang.ClassLoader.loadClass(ClassLoader.java:251) at test.main(test.java:7) The workaround is to use the Class.forName(String, boolean, ClassLoader) method.	1	1	https://issues.jboss.org/browse/JBCACHE-992

#	Requirement	Questions	Answered	Successfully Answered	URL
8861	JBCACHE-991	Re: approach 1, what's the exact usage pattern for this, particularly when it comes to locking and return values? E.g., 1. What happens to read calls when executed within a batch? Ignored? 2. Write calls that return values? E.g., put()? 3. Replication - I'm guessing this is the primary reason for this JIRA - deferred till commitBatch()? 4. Persisting? Same as 3.? 5. Notifications? 6. Locks - acquired with each WRITE invocation and held for the duration of the batch? Or just acquired and released at commit time? Depending on the answers above, it may look very similar to a JTA transaction. Also, should batches participate in ongoing JTA transactions?	1	1	https://issues.jboss.org/browse/JBCACHE-991
8862	JBCACHE-971	For the record (Manik's recent target version change reminded me to update this): I have now developed a cache library which includes implementations of the following algorithms: * ARC * FRC (non-self-adjusting version of ARC; not sure if this is covered by the same patents as ARC) * 2Q * MQ (Multi-Queue Replacement, by Zhou, Philbin and Li -- see http://www.usenix.org/event/usenix01/full_papers/zhou/zhou_html/index.html -- I am not sure if there are any patent encumbrances here, and have contacted 2 of the authors to try and find out) This library has been developed without being jboss-specific; I have produced an adapter for JBoss Cache 1.3 (which we're still stuck using) and then one for JBoss Cache 2.0, which I would be happy to contribute (my employer has given me permission to do so, and is in the process of formalising this internally). HOWEVER, I have not heard anything back from IBM in the last 5 months in relation to the ARC patent issues, which makes me loathe to release this (of course!). Is there anyone at JBoss who would be in a better position to shake the IBM tree and try to extract an answer? I would be very keen to release my library and think JBoss Cache would benefit from it.	1	0	https://issues.jboss.org/browse/JBCACHE-971
8863	JBCACHE-971	In the meanwhile, could you contrib the patent-free implementations? I can't expect the legal permission route to be quick in any way. :-)	0	0	https://issues.jboss.org/browse/JBCACHE-971
8864	JBCACHE-964	Are you also writing a "Migrator" application to do the looping - traverse the tree, read each node and put it back in the cache?	1	1	https://issues.jboss.org/browse/JBCACHE-964
8865	JBCACHE-957	Any reason why this needs to be in 2.0.0? Can this be moved off to 2.1.0?	1	1	https://issues.jboss.org/browse/JBCACHE-957
8866	JBCACHE-956	Manik, we discussed my adding docs for this behavior to LockParentForChildInsertRemove. Can you do this as part of this JIRA, since it might be done later than JBCACHE-955. I don't want to doc an unimplemented feature.	1	1	https://issues.jboss.org/browse/JBCACHE-956
8867	JBCACHE-941	Sounds interesting. Want to start a forum thread to discuss this? I see a few congruent lines, one being PojoCache and the other being the JSR-107 compliant JCACHE interface. Either way though, this would probably be something for the 2.1.0/2.2.0 timeframe.	1	1	https://issues.jboss.org/browse/JBCACHE-941
8868	JBCACHE-932	Pointless; will be superceded with lock striping when we move to MVCC	0	0	https://issues.jboss.org/browse/JBCACHE-932
8869	JBCACHE-905	This is great stuff; thanks! :-) I notice there are a couple packages in there that only exist in the testsuite. Is that the way it's meant to be, i.e. test if we're executing our tests?	0	0	https://issues.jboss.org/browse/JBCACHE-905
8870	JBCACHE-905	Do you know if a patched EMMA plugin is due soon as well (I know the EMMA and Clover plugins share a significant portion of their codebase)? All of our Maven scripts are built around EMMA right now. We could switch to clover, but if an EMMA plugin is due soon then I'd rather wait for that.	1	1	https://issues.jboss.org/browse/JBCACHE-905
8871	JBCACHE-905	Manik, please, have a look at http://hudson.qa.jboss.com/hudson/view/JBoss Cache/job/jbosscache-core-trunk-code-coverage/17/ Is that OK?	1	1	https://issues.jboss.org/browse/JBCACHE-905
8872	JBCACHE-880	Could you pls link to the discussion thread in this task. Also add that this will require documentation and configuration changes.	1	1	https://issues.jboss.org/browse/JBCACHE-880
8873	JBCACHE-879	Functionality added but there's some TODOs: - VersionAwareMarshallerPerfTest: comparison failures based on results from different marshalling methods, i.e, if difference bigger than x, fail. Also add marshalling different types of objects. - Increasing VAM performance: using similar techniques as used in ExposedByteArrayOutputStream for VAM? Needs profiling to determine effect. NOTE: All profiling and changes related to performance will be part of BETA2	0	0	https://issues.jboss.org/browse/JBCACHE-879
8874	JBCACHE-839	Done. But there are two minor issues: 1. The performance impact needs to be evaluated. 2. During testing, how do we turn it on and off?	0	0	https://issues.jboss.org/browse/JBCACHE-839
8875	JBCACHE-829	Isn't the EULA only for binary distributions, as per the discussion on core? Or have I missed something out?	1	1	https://issues.jboss.org/browse/JBCACHE-829
8876	JBCACHE-815	I thought this following code would mean having a dependency with jboss-serialization.jar: public static ObjectInputStream createObjectInputStream(byte[] bytes) throws IOException { ByteArrayInputStream in = new ByteArrayInputStream(bytes); if (useJBossSerialization) { return new JBossObjectInputStreamOverride(in); } else { return new MarshalledValueInputStream(in); } } Is it possible to use an interface as I proposed, so you don't need to use reflection to create JbossSerialization? If the stuff is working (not loading jboss-serialization.jar) and if you don't like the idea of the interface, just close the issue again.	1	1	https://issues.jboss.org/browse/JBCACHE-815
8877	JBCACHE-800	have you got a unit test for this, showing how this leads to fatal problems?	1	1	https://issues.jboss.org/browse/JBCACHE-800
8878	JBCACHE-797	I don't see why not. I think we need a more sophisticated NodeFactory in the first place though. The current approach to creating nodes is mimicking nodes of parent type, and all this starts with the root. A better approach would be to have the node instance type injected into the factory when the cache starts, based on configuration settings (i.e., optimistic or pessimistic locking, whether this is a pojo cache or a plain cache, etc)	1	1	https://issues.jboss.org/browse/JBCACHE-797
8879	JBCACHE-752	Manik, will you be able to look at this, or maybe ask Galder? If not, assign back to me.	1	1	https://issues.jboss.org/browse/JBCACHE-752
8880	JBCACHE-738	Vladimir, can I pass this over to you since you're dealing with these issues?	1	1	https://issues.jboss.org/browse/JBCACHE-738
8881	JBCACHE-735	Does this need to be in every log output? Or just do it once? When we switch to multiplexer, is this still necessary?	1	1	https://issues.jboss.org/browse/JBCACHE-735
8882	JBCACHE-675	Hi Ben, I couldn't sleep and I wanted to see if you had additional thoughts on this. I guess you did. :) I saw your forum message as well. Thanks for pointing out these related issues. http://jira.jboss.com/jira/browse/JBCACHE-676 http://jira.jboss.com/jira/browse/JBCACHE-110 http://jira.jboss.com/jira/browse/JBCACHE-675 To resolve JBCACHE-675, should we conform to our new requirement that Map keys be serialiable? Or do you want me to revert the requirement that keys be serializable? Same question is true for JBCACHE-676, I'm assuming your going to want me to support non-serializable key but I'll wait to hear back from you. JBCACHE-110, I'm not sure about this one yet. Should we continue on the forum? Scott <9880004.1151635871320.JavaMail.jira@cloud.prod.atl2.jboss.com>, "Ben Wang (JIRA)" <jira-events@jboss.com> wrote:	1	1	https://issues.jboss.org/browse/JBCACHE-675
8883	JBCACHE-675	Scott, can you update the 1.4 doc to reflect this? Thanks!	1	1	https://issues.jboss.org/browse/JBCACHE-675
8884	JBCACHE-672	Brian, I have fixed it and checked it in. But I have no way verifying the fix now. Can you please test it out really quick under AS4.0? Thanks!	1	1	https://issues.jboss.org/browse/JBCACHE-672
8885	JBCACHE-670	Brian, can you go thru one more time to see if this makes it easier for the user per our discussion? If you are sastified with it, please go ahead and close it.	0	0	https://issues.jboss.org/browse/JBCACHE-670

#	Requirement	Questions	Answered	Successfully Answered	URL
8886	JBCACHE-666	Brian, can you go thru one more time to see if this makes it easier for the user per our discussion? If you are sastified with it, please go ahead and close it.	1	1	https://issues.jboss.org/browse/JBCACHE-666
8887	JBCACHE-658	Can you give a code snnipet to illustrate this problem? Let's say if I do, cache.putObject("/org/jboss/person", ben); then only person node is WL, while org/jboss should be RL. And this is only for putObject. After that, fine-grained modification should only lock the appropriate sub-nodes.	1	1	https://issues.jboss.org/browse/JBCACHE-658
8888	JBCACHE-658	I see. But what is your fqn though? I thin there is way to get around your problem. Let's say, you are filling: "/person/1", "person/2", ..., "/person/n" then to prevent WL on person, you can pre-create "/person/1", "person/2", ..., "/person/n" after that you do, pubObject("/person/1", pojo1); pubObject("/person/2", pojo1); ... should give you the concurrency that you are looking for. The other thing is regarding the overall performance. You can try out the latest 1.4.CR1 that has been optimized for PojoCache specifically.	0	0	https://issues.jboss.org/browse/JBCACHE-658
8889	JBCACHE-642	Does this have to be a config option on the file cache loader or on all cache loaders? What if someone uses a different cache loader impl? Also, doesn't this completely beat the purpose of a cache loader if it gets purged on startup? I suggest a forum thread to discuss this ...	0	0	https://issues.jboss.org/browse/JBCACHE-642
8890	JBCACHE-627	Should this be done? I'm not sure if this silent creation is a good thing.	1	1	https://issues.jboss.org/browse/JBCACHE-627
8891	JBCACHE-620	Ben, could you pls look at this one for me?	0	0	https://issues.jboss.org/browse/JBCACHE-620
8892	JBCACHE-620	We do have a field called MaxAgeSeconds in LRU that is doing what you described. Did you try it?	1	1	https://issues.jboss.org/browse/JBCACHE-620
8893	JBCACHE-620	What version was the MaxAgeSeconds attribute added?	1	0	https://issues.jboss.org/browse/JBCACHE-620
8894	JBCACHE-619	So as I understand it, without this bug fix it is not possible to use enum values in entity objects when using the EJB3 entity cache? Is there a workaround? I found this JIRA from this forum thread: http://www.jboss.com/index.html?module=bb&op=viewtopic&t=71311&view=previous	0	0	https://issues.jboss.org/browse/JBCACHE-619
8895	JBCACHE-595	Cool. One other small thing: for completeness the MapModifications struct should probably also have a "public Map unmodifiedEntries" field. Manik Surtani wrote: Ok, sounds cool, will add this to the JIRA task. How about a variation on Ben's idea of providing an enum to specify the type of change? Sample API follows. /** * Called before and after a {@link Node} is modified. Data passed in the * <code>modData</code> param can be used to determine what was modified. * <p/> * When called with <code>pre == true</code>, <code>modData</code> is the * initial state of the {@link Node} before modification. * <p/> * When called with <code>pre == false</code>, the contents of <code>modData</code> * depend on the value of <code>modTtype</code>: * * - PUT_KEY_VALUE: Map contains the single key/value pair that was added. * - REMOVE_KEY_VALUE: Map contains the single key/value pair that was removed. * - PUT_MAP: Map contains the new state of the {@link Node} following modification. * This map includes modified key/value pairs as well as any that were not affected. * * <p/> * Implementations interested in seeing the difference in the node data in the PUT_MAP case * can cache the <code>modData</code> map passed when <code>pre == true</code>, and then * when the <code>pre == false</code> callback is received, pass the cached map and the new * <code>modData</code> to {@link org.jboss.cache.util.Util.diffNodeData(Map pre, Map post)}. * * @param fqn Fqn of the node being modified * @param pre <code>true</code> if the modification is about to be applied, * <code>false</code> if it has already been applied * @param isLocal <code>true</code> if the modification originated locally, * <code>false</code> if it was replicated from another node * @param modType PUT_KEY_VALUE, REMOVE_KEY_VALUE or PUT_MAP * @param modData Unmodifiable {@link Map}; will not be <code>null</code>. See * description above. */ public void nodeModified(Fqn fqn, boolean pre, boolean isLocal, ModificationType modType, Map modData); public class Util { public static MapModifications diffNodeData(Map pre, Map post) { // do the diff here } public static class MapModifications { public Map addedEntries; public Map removedEntries; public Map modifiedEntries; } } Advantage here is JBC goes to little effort to support this, so overhead is low if the listener isn't interested. CacheListener impls that are interested in modifications, but where the app never calls put(Fqn, Map) also have very little overhead. I believe Jerry's DistributedState is such a case. The big overhead of doing the diff is only incurred by CacheListener impls that care, and we provide much of the code to do it for them.	1	1	https://issues.jboss.org/browse/JBCACHE-595
8896	JBCACHE-574	This is pre-1.4Alpha version (05/01/06) profiling data under JProfiler. The snapshot is attached. This is for repeatedly putObject with pojo of list size=100 (from the perf tester). The profiling is for 24 times of putObject and it results in: cache.get(), 2383x3 times, 38.2% cpu cache.put() - 2383x2 times, 33.7% cpu Further profiling of plain cache get/put operations shows the bottleneck mostly in lock and tx interceptor (even with local mode). Note that PojoCache results in 100 times more of get/put becuase of the List size of 100 (so we have to iterate thru the list 100 times). Obviously, becuase of plain cache local mode is still too expensive, it adds up in total. Here are couple of solutions: 1. Create a pojo level lock, then all get operation can bypass lock interceptor (since we don't need it anymore). What about put operation? We want to bypass the lock operation but not the tx interceptor. 2. This can combine with a different kind of pojo mapping. E.g., we do map it in fine-grained from cache1. Instead, we will send the pojo over to cache2 directly and then do the mapping from there. Furthermore, the operation will be transactional in one shot.	0	0	https://issues.jboss.org/browse/JBCACHE-574
8897	JBCACHE-562	Not quite as simple with the pojo-based configuration in 2.0, since the cacheloader class is an inherent property of the typed config object. The config object should probably have an "enabled" property. The "empty class" idea could still work for parsing 1.x-style configs, though. If class is empty, parser doesn't create the config pojo, and thus no cacheloader is created.	1	1	https://issues.jboss.org/browse/JBCACHE-562
8898	JBCACHE-562	Is this still important? I vote that we don't attempt to inject behaviour based on an empty class name. If you want to disable cache loaders, simply comment out the relevant block in XML or don't set a CacheLoaderConfig in the Configuration object.	0	0	https://issues.jboss.org/browse/JBCACHE-562
8899	JBCACHE-555	Are your 'after' timings after switching to caching the trace boolean?	1	1	https://issues.jboss.org/browse/JBCACHE-555
8900	JBCACHE-549	Should use something like GET_CONDITIONAL - JGRP-518?	0	0	https://issues.jboss.org/browse/JBCACHE-549
8901	JBCACHE-535	Which corresponding unit test contradicts the TreeCache.getKeys() contract?	1	1	https://issues.jboss.org/browse/JBCACHE-535
8902	JBCACHE-535	Yes, I think Node.getDataKeys() needs to be more consistent as well. Perhaps return data != null ? (data.keySet().isEmpty() ? null : data.keySet()) : null; Would need to document this in the Javadocs very clearly - both in Node.java as well as interfaces implemented.	1	1	https://issues.jboss.org/browse/JBCACHE-535
8903	JBCACHE-525	Elias, how robust is this code ? I need to solve the following problem: 2 different processes modify the same (shared) file, e.g. on an NFS volume, does JDBM use file locks of some sort to serialize concurrent access ?	1	0	https://issues.jboss.org/browse/JBCACHE-525
8904	JBCACHE-520	Minor correction. I have moved the payload object construction out of the loop in the tester and got more performance because of CPU cycle, I guess. The number is 1400 res/sec (9.8MB) for async and 650 req/sec (4.6MB) for sync.	1	1	https://issues.jboss.org/browse/JBCACHE-520

#	Requirement	Questions	Answered	Successfully Answered	URL
8905	JBCACHE-520	On what cluster did you get the 1400 res/sec number ? dev02, dev03 too ?	1	1	https://issues.jboss.org/browse/JBCACHE-520
8906	JBCACHE-519	Just tried on the lab cluster usig cluster01 and cluster02 (with transactions) Total time spent (ms) on thread id 0: 94997 for 400 loops Throughtput for this node with: threads = 50 loops = 400 sleep interval = 1 object list size = 100 tranasaction? true is: 209 requests/sec (no transactions) Total time spent (ms) on thread id 0: 82528 for 400 loops Throughtput for this node with: threads = 50 loops = 400 sleep interval = 1 object list size = 100 tranasaction? false is: 241 requests/sec This was with repl_sync. What replication mode did you use for your tests?	0	0	https://issues.jboss.org/browse/JBCACHE-519
8907	JBCACHE-509	Just tried this on the lab cluster with 3 nodes in the cluster, with the latest srcs from HEAD and didn't see any OOMEs. Have you seen this when testing on the lab cluster?	1	0	https://issues.jboss.org/browse/JBCACHE-509
8908	JBCACHE-490	Not for 1.3.0 or 1.4.0, I presume? And how would you want to do this in code - POJOCache extends TreeCacheAop, with no added features?	0	0	https://issues.jboss.org/browse/JBCACHE-490
8909	JBCACHE-490	Definitely in 1.4. But what do you think in 1.3? This will be just an extra class that mirrors TreeCacheAop and TreeCacheAopMBean for now. I have refactor-ed it before, both files use delegation now. So they are very short anyway.	0	0	https://issues.jboss.org/browse/JBCACHE-490
8910	JBCACHE-490	It has been considered actually. What's your rationale?	1	1	https://issues.jboss.org/browse/JBCACHE-490
8911	JBCACHE-488	I presume this is regarding cache content rather than any of our nternal objects, for which we should have our own marshaller (e.g., Fqn, Gtx, MethodId, etc). I also presume this will have a dependency on JBoss Serialization solving it's current issue with generating a large number of bytes for each object it serializes. Is there an open JIRA task for this in JBoss Serialization?	1	1	https://issues.jboss.org/browse/JBCACHE-488
8912	JBCACHE-484	Elias, could we discuss the patches you have here on the JBoss Cache Design forum? Could you pls go ahead and start a thread there and link it to this task? Thanks, Manik	1	1	https://issues.jboss.org/browse/JBCACHE-484
8913	JBCACHE-471	Should we just define an API callback and leave it to the user to implement it? We can provide a simple default policy implementation where the new coordinator will always win, e.g., copy its state to everybody again? I think the user can come up with more sphosticated algorithm to contribute back.	0	0	https://issues.jboss.org/browse/JBCACHE-471
8914	JBCACHE-466	JSP-250 common annotations for the Java platform. Not too much there. Generated, Reousrce, and Security stuffs. Do we have use of Reousrce? I don't think so now.	0	0	https://issues.jboss.org/browse/JBCACHE-466
8915	JBCACHE-464	Do we include 3rd party licenses in the AS distro at the moment?	1	1	https://issues.jboss.org/browse/JBCACHE-464
8916	JBCACHE-450	After some further thinking, I conclude that eviction is really not well suited in PojoCache. Again, it will have to run as part of passivation. Here is why. If we use eviction only to evict the pojo then: 1. When evicted, the cache content will be empty (i.e., no long available). What happens to the interceptor? Do we remove it as well or not? If we keep the interceptor, what happens when the pojo is accessed again, e.g., pojo.getAge()? It will return null since the cache content has been evicted already. 2. However, if decide to remove the corresponding iterceptor, then pojo.getAge() will return correctly. However, pojo.setAge(21) won't be intercepted by the cache interceptor. It will be just a regular pojo operation. This is a semantic change that user won't be aware of. 3. For pojo-based eviction, if a sub-object is refereshed, we will need to refresh the corresponding parent object as well. E.g., if address is refershed, then Person needs to be refreshed (i.e., no evicted). We need to take care of multiple references as well. We can do this in 2.0 since we keep track of like of parent object from the sub-object. Nonetheless, it is not as straightforward. 4. What happens at the remote node if it is a passive node (e.g., only failing over purpose). Then the pojo won't be evicted since officially there won't be one yet (untill the user do cache.find(id)) So I think the best way is go with passivation. If we use passivation, we can use the current cache node-level eviction policy. We really don't care about the pojo level since from PojoCache point of view, the fqn value is always available, e.g., when I do a pojo.getAge(), behind the scene, the cache loader will retrieve the value associated with the age item, e.g., cache.get(id, "age"). This way, we are totally transparent. Couple things to take notice though: 1. When it is passivated, if the pojo becomes orphan (i.e., no one is referencing it anymore), GC will kick in. Originally, the pojo reference is also stored under AOPInstance (as a strong reference) but now AOPInstance has been passivated already (and the reference field is a transient modifier), so GC can freely to collect the pojo instance. No problem in the vm. 2. The only problem though is what happens to the data in the cache loader store? Since the pojo has been GCed, most likely no one will use it (unless, I am using cache.find(id) again to resurrect it), we will need to come up a way to trim it or age out. We can have a separate timer thread to weed out the ageout items. 3. This is the scenario: cache.attach("pojo", pojo); pojo.setName("Ben"); ... // it is passivated pojo.getName(); // no problem here Pojo anotherPojo = cache.find("pojo"); then pojo != anotherPojo because getAopInstance() return nulls during find (becuase it is a transient field). Can we come up a way to serialize the pojo reference but not replicating it? If somehow AopInstance contains non-transient pojo reference but AopInstance is stored as a local cache mode (e.g., doesnt replicate), then it is serializable but not replicatable. As a result, case solved. If we can't, then everything will work fine actually, e.g., pojo and anotherPojo can do indepdent update and get. But only comparison of object identity (by reference) will fail.	1	1	https://issues.jboss.org/browse/JBCACHE-450
8917	JBCACHE-445	Brian, what is the status of this one? Should I wait until it is complete before releasing?	1	1	https://issues.jboss.org/browse/JBCACHE-445
8918	JBCACHE-427	Added manually into HEAD as well - will look into modifying the build system to automate this for the next release.	0	0	https://issues.jboss.org/browse/JBCACHE-427
8919	JBCACHE-413	what exactly do you want to do ? Create a build-aop.xml and import that from build.xml ?	0	0	https://issues.jboss.org/browse/JBCACHE-413
8920	JBCACHE-404	The isActive method is in the new Region interface. Only open question is whether RegionImpl should check it's active status before running the activation code, and ignore the call (log warning?) if already active.	1	1	https://issues.jboss.org/browse/JBCACHE-404
8921	JBCACHE-397	This is being tackled by Brian for 1.2.4SP2 and 1.3. Ben, Manik and Bela will meet Feb 23/24 2006 in Neuchatel to discuss binary compatibility going forward (1.4 and onwards).	1	1	https://issues.jboss.org/browse/JBCACHE-397
8922	JBCACHE-394	Whats the status of this, can this be closed?	1	1	https://issues.jboss.org/browse/JBCACHE-394
8923	JBCACHE-372	I will take a closer look. But by looking your test bench, it looks to me you are doing putObject with different student pojo. This is expected to be slow (question is is it that slow?) since this is the acid test for TreeCacheAop. We expect the usage is 1. you will re-use pojo repeatedly. That is, pojo has a relative long live time; 2. It is read mostly. E.g., upadte on pojo fields doesn't happen with every single call.	1	1	https://issues.jboss.org/browse/JBCACHE-372
8924	JBCACHE-356	Bela, How about JBoss Cache content on Labs? How does it look? If all looks good, please close all the sub tasks. Cheers, Damon	1	1	https://issues.jboss.org/browse/JBCACHE-356
8925	JBCACHE-347	Should we be doing this? Still debatable, as per forum discussion... Deferring for now anyway.	0	0	https://issues.jboss.org/browse/JBCACHE-347
8926	JBCACHE-344	He's your friend :-) Q: can we have this by 1.3, or should we push it into 1.4 ?	0	0	https://issues.jboss.org/browse/JBCACHE-344
8927	JBCACHE-344	Are we going ahead with this for 1.3? I think we should push this out to 1.4 ...	0	0	https://issues.jboss.org/browse/JBCACHE-344

#	Requirement	Questions	Answered	Successfully Answered	URL
8928	JBCACHE-344	Gurkan, could you target your next release of the IDE to support JBossCache 1.4.0? New features (as far as configs go) over 1.3.0 should only be buddy replication (JBCACHE-61) config elements, and a new cache loader (JBCACHE-525). The moment these 2 tasks are completed and are in CVS HEAD you could use a CVS snapshot of JBossCache to test the IDE. Naturally as I release DRs, alphas, etc., you'd use these as well.	1	1	https://issues.jboss.org/browse/JBCACHE-344
8929	JBCACHE-330	Hany, what's the status of this? Is it just pending the implementation of an in-memory DB for the unit tests? Could you pls add the forum discussion link to this JIRA task? Thanks, Manik	1	1	https://issues.jboss.org/browse/JBCACHE-330
8930	JBCACHE-326	Are we still doing this? This is the concept of a separate 'processor' thread for state events.	0	0	https://issues.jboss.org/browse/JBCACHE-326
8931	JBCACHE-307	Sounds good. Can we have this document in JBossCache/docs/design and refer to it from here ? I don't want to go JIRA huting to look up an old design decision, would be nice to have those decisions be part of the code base.	1	1	https://issues.jboss.org/browse/JBCACHE-307
8932	JBCACHE-304	Manik, if you are done, can you close this one?	0	0	https://issues.jboss.org/browse/JBCACHE-304
8933	JBCACHE-273	Bela, can the new getState go through the interceptor chain to have proper locking?	0	0	https://issues.jboss.org/browse/JBCACHE-273
8934	JBCACHE-272	Hany, you have 3 remaining minor tasks for 1.2.4. Can you find time to finish them by tomorrow? 1.2.4beta is coming out this weekend.	1	1	https://issues.jboss.org/browse/JBCACHE-272
8935	JBCACHE-270	I had a discussion with Bela today over a teleconference call about this issue which results in the following requirements for this interceptor: 1. Needs to handle 2 types of events. One results from method call invocation, for example, evictNodeMethodLocal invoked from TreeCacheListener. The other results from an interceptor action, for example, cache loader interceptor emitting nodeLoaded event. 2. Needs to notify listeners to TreeCache of pre and post events. 3. Needs to support transactional events. For example, if a put has been issued during a transaction and a rollback has been followed through an undo operation. The interceptor should handle it correctly. The following questions and ideas are still under investigation: 1. Where do we place the interceptor? Should we place it right before the call interceptor or at the bottom of the stack? 2. Do we need a queue to hold the various events in the interceptor and other interceptor write to it during execution to have their notification emitted? 3. Method calls in TreeCache call the notification methods directly. Adding this interceptor means that we would need to remove those calls and the flags associated with them, for example TreeCache._removeData() which is called either from its own method call invocation or from other method call invocation such as evict(). Doing so might have undesirable effect on how the cache behaves and might introduce new bugs in the release. That actually makes me want to ask the following question: should we do this now or wait until release 1.3 and implement it along with TreeCache Refactor issue? That's all I can think of now. Bela, Please add you thoughts and any other questions/ ideas that we discussed and I've might forgot.	1	1	https://issues.jboss.org/browse/JBCACHE-270
8936	JBCACHE-261	Based on the discussion we had during our meeting, Bela added to JGroups 2.2.9 "payload" notion in MethodCall put and get methods. What do we need to add to TreeCacheListener to support DSListener and DSListenerEx? Can you please elaborate on the following quote? <quote>To implement this, we might add a thread local storage variable to TreeCache and set it/clear it depending on whether we are making a local/remote change.</quote>	0	0	https://issues.jboss.org/browse/JBCACHE-261
8937	JBCACHE-261	Sorry Bela, why do we need to use 2.2.9 get/set payload notion in order for this to work? Can't we tell from the interceptor stack to see whether the replication is local or not? We can't rely on 2.2.9 since AS won't use 2.2.9 as I understand it for version incomptability reason. Or unless we lobby it.	0	0	https://issues.jboss.org/browse/JBCACHE-261
8938	JBCACHE-261	I can add a new event type in ExtendedTreeCacheListener for local, would that satisfy the requirement? Also, where this event should be emitted? Should be called from TreeCache.callRemoteMethods or else where?	1	1	https://issues.jboss.org/browse/JBCACHE-261
8939	JBCACHE-258	Daniel, while you are at eviction, can you also take a look at this one? thanks.	1	0	https://issues.jboss.org/browse/JBCACHE-258
8940	JBCACHE-254	Removal and subsequent re-creation under B doesn't do it ?	0	0	https://issues.jboss.org/browse/JBCACHE-254
8941	JBCACHE-249	3 zip files are now created when a dist is done: meg:~/Code/JBoss/JBossCache manik$ ls -lh dist/ total 30832 -rw-r--r-- 1 manik manik 584K Sep 1 13:16 bdbje-for-jbosscache.zip -rw-r--r-- 1 manik manik 9M Sep 1 13:16 jboss-cache-dist.zip -rw-r--r-- 1 manik manik 5M Sep 1 13:16 jboss-cache-minimal-dist.zip drwxr-xr-x 4 manik manik 136B Sep 1 11:14 lib meg:~/Code/JBoss/JBossCache manik$ The bdbje-for-jbosscache.zip only contains the sleepycat jar and some readme docs. Who do we need to contact to make sure automated builds pick up the new zip and get it published on SF for downloads?	1	1	https://issues.jboss.org/browse/JBCACHE-249
8942	JBCACHE-229	I agree with Bela here. I like typed event better and it is more explicit as well. How about we do this? 1. We create a sub-interface from TreeCacheListener, say, ExtendedTreeCacheListener that has these extra events. Let's make sure it will cover what we will need for a while. Then we will need to add code to differentiate when a user is registering an regularu TreeCacheListener or an extended one. We will emit the events accordingly. 2. We combine them in 1.3 afterwards.	1	1	https://issues.jboss.org/browse/JBCACHE-229
8943	JBCACHE-229	JBoss Portal is a client of JBoss Cache and use the implements the TreeCacheListener interface. From what I read it sounds like the interface is maintained and does not change however according to the docs/design/NewTreeCacheListener.txt document it seems that it will need migration : <quote>Classes that need refactoring: All classes that currently implement TreeCacheListener</quote> I am worried about that because that could break JBoss Portal with future versions of JBoss Cache. Can you clarify this ?	1	1	https://issues.jboss.org/browse/JBCACHE-229
8944	JBCACHE-214	I take it this has been sorted? http://docs.jboss.com/jbcache/current/TreeCache/html_single/ now contains the new stuff to do with Async Cache Loaders, for example... Ben, could you pls check if your changes are up here as well, and confirm? Thanks.	1	1	https://issues.jboss.org/browse/JBCACHE-214
8945	JBCACHE-197	Okay, if there is no problem with perf I'll close the case. Okay with you ?	0	0	https://issues.jboss.org/browse/JBCACHE-197
8946	JBCACHE-193	With these changes 2 lock unit tests fail: - NonBlockingWriterLockTest.testUpgradeWithExistingReader() - NonBlockingWriterLockTest.testUpgradeWithMultipleReaders() Can you look into why they fail ?	1	1	https://issues.jboss.org/browse/JBCACHE-193

#	Requirement	Questions	Answered	Successfully Answered	URL
8947	JBCACHE-169	UDP.setProperties(): the following properties are not recognized: 2005-05-24 07:19:42,012 INFO [STDOUT] -- listing properties -- 2005-05-24 07:19:42,012 INFO [STDOUT] bind_to_all_interfaces=true 2005-05-24 07:19:42,013 ERROR [org.jboss.cache.TreeCache] Starting failed jboss.cache:service=TomcatClusteringCache ChannelException: unable to setup the protocol stack at org.jgroups.JChannel.<init>(JChannel.java:217) at org.jgroups.JChannel.<init>(JChannel.java:189) at org.jboss.cache.TreeCache.startService(TreeCache.java:1104) at org.jboss.system.ServiceMBeanSupport.jbossInternalStart(ServiceMBeanSupport.java:272) at org.jboss.system.ServiceMBeanSupport.jbossInternalLifecycle(ServiceMBeanSupport.java:222) at sun.reflect.GeneratedMethodAccessor2.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:141) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:80) at org.jboss.mx.server.Invocation.invoke(Invocation.java:72) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:249) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:644) at org.jboss.system.ServiceController$ServiceProxy.invoke(ServiceController.java:897) at $Proxy0.start(Unknown Source) at org.jboss.system.ServiceController.start(ServiceController.java:418) at sun.reflect.GeneratedMethodAccessor9.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:141) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:80) at org.jboss.mx.server.Invocation.invoke(Invocation.java:72) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:249) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:644) at org.jboss.mx.util.MBeanProxyExt.invoke(MBeanProxyExt.java:176) at $Proxy4.start(Unknown Source) at org.jboss.deployment.SARDeployer.start(SARDeployer.java:273) at org.jboss.deployment.MainDeployer.start(MainDeployer.java:964) at org.jboss.deployment.MainDeployer.deploy(MainDeployer.java:775) at org.jboss.deployment.MainDeployer.deploy(MainDeployer.java:738) at sun.reflect.GeneratedMethodAccessor17.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:141) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:80) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:121) at org.jboss.mx.server.Invocation.invoke(Invocation.java:74) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:127) at org.jboss.mx.server.Invocation.invoke(Invocation.java:74) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:249) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:644) at org.jboss.mx.util.MBeanProxyExt.invoke(MBeanProxyExt.java:176) at $Proxy8.deploy(Unknown Source) at org.jboss.deployment.scanner.URLDeploymentScanner.deploy(URLDeploymentScanner.java:325) at org.jboss.deployment.scanner.URLDeploymentScanner.scan(URLDeploymentScanner.java:501) at org.jboss.deployment.scanner.AbstractDeploymentScanner$ScannerThread.doScan(AbstractDeploymentScanner.java:204) at org.jboss.deployment.scanner.AbstractDeploymentScanner.startService(AbstractDeploymentScanner.java:277) at	1	1	https://issues.jboss.org/browse/JBCACHE-169
8948	JBCACHE-139	Is this still relevant? Invalidation is in place and works fine without such a notification. Can this be closed?	1	1	https://issues.jboss.org/browse/JBCACHE-139
8949	JBCACHE-137	This is just used to read the xml config files. Write our own xml parser to read this and get rid of this dependency?	0	0	https://issues.jboss.org/browse/JBCACHE-137
8950	JBCACHE-109	Bela, is this still an issue?	0	0	https://issues.jboss.org/browse/JBCACHE-109
8951	JBCACHE-92	1. Are there any plans in the roadmap to support the kind of forceful eviction policy I need in JBoss Cache? 2. Are there any side effects with me implementing a new LRUPolicy that forcefully evicts a node and its children? 3. Any advice on how to do #2? Off the top of my head I would probably create a new CascadingLRUPolicy class that overrides the evict(Fqn) method to accomplish what I need to do. The new evict() implementation would be largely a copy/paste of TreeCache._evict() but thatt doesn't check for children. Here's a quick mock-up of what the new CascadingLRUPolicy.evict(Fqn) implemention might look like: public void _evict(Fqn fqn) throws Exception { if(!cache_.exists(fqn)) return; // node does not exist. Maybe it has been recursively removed. // use remove method now if there is a child node. Otherwise, it is removed boolean create_undo_ops = false; boolean sendNodeEvent = false; boolean eviction=true; _remove(null, fqn, create_undo_ops, sendNodeEvent, eviction); } Thanks in advance for any/all help. Regards, Brian Dueck	0	0	https://issues.jboss.org/browse/JBCACHE-92
8952	JBCACHE-86	Do you mean to say that nodeEvicted(Fqn) is going to be called whenever a node gets invalidated?	1	1	https://issues.jboss.org/browse/JBCACHE-86
8953	JBCACHE-66	Did you update the documentation as well (JBossCache/docs/TreeCache/en/master.xml) ?	1	1	https://issues.jboss.org/browse/JBCACHE-66
8954	JBCACHE-61	Ah, I see it now. :-) Here is previous comments: Did I miss it or how come I don't see the buddy replication Jira on 1.3? I have one feedback from Greg Hinkle that WL is currently doing of which we should implement as well (may not be the first cut though). In configuring BR policy, can we allow user to configure the buddy group to be the one that is not located in the same physical machine? For example, we have instance A, B, C, D. A and B are in machine 1 and C and D are in machine 2. Then the buddy group of A will have to be either C or D, and vice versa. This is desired in case the machine goes down.	1	1	https://issues.jboss.org/browse/JBCACHE-61
8955	JBCACHE-47	Is this documented ?	1	0	https://issues.jboss.org/browse/JBCACHE-47
8956	JBCACHE-42	I am starting to gather a use case for the TreeCacheAop white paper after the fine-grained http session replication is almost complete. I want your contribution for any idea if you have! The use case that I am looking for will be either running outside or inside of the web application. But here are the key traits that need to have: 1. Use of fine-grained replication for performance. Ideally to have a large collection (huge list) and each element has another complex object. When a singular field is modified, it will only replicate that field (instead of the whole list!). In addition, for chatty replication, it can also do batch processing as well for performance optimization. 2. Transparent management of object relationship after replication. Each object potentially can also have object graphs that can have multiple or recursive references. Regular cache replication has no way of managing this relationship (when failing over, for example, and as a result, demands developer intervention). But TreeCacheAop can manage this automatically. 3. (Optional) Use of CacheLoader +CacheAop to provide state persistency. Some state data may need to guarantee persistency so it will be less expensive to reconstruct it. This can be either an individual or shared cache loader. So far, an example that I can think of is in the line of PIM (Personal Information Management). Let's say I need to keep track of a list of company employees. Each employee has the following fields: - Company address - Home address - Personal relationship (spouse, relative, friend) - Work relationship (boss, subordinate) So if I only change Joe's wife's age (in this case, Mary in another company), for example, I have a fine-grained replication scenario inside a huge list and fair size object. In addition, when replicated, the relationship is perserved automatically. Meaning, you only update it one place, and the Joe/Mary spousal relationship is preserved on the replicated node (this is good, right? :-). Do you better ideas or can you enhance this?	0	0	https://issues.jboss.org/browse/JBCACHE-42

#	Requirement	Questions	Answered	Successfully Answered	URL
8957	FUSETOOLS-2442	when done, two restrictions on tested version can be removed in org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForCXFCodeFirstIT.testCXFCodeFirstSpringProjectCreation() and org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForAMQIT.testAMQBlueprintProjectCreation()	0	0	https://issues.jboss.org/browse/FUSETOOLS-2442
8958	FUSETOOLS-2417	[~davsclaus] Hey Claus, any chance you can answer a handful of questions about the bean implementation in camel and the method properties? We allow a user to browse a bean to select a method for the init-method, destroy-method, and factory-method. Should those be public methods? Or can it be any private or protected method as well? And should it only be methods without parameters?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2417
8959	FUSETOOLS-2410	{quote}Is it a "one time" generation? or the user will update the Swagger API regularly?{quote} I think the user would like to update the Swagger (OpenAPI) specification and either re-run the generation or see the created REST DSL auto updated. {quote}After generation, the user will need to modify the generated code a bit or not at all?{quote} The current code generation available via [CAMEL-11006|https://issues.apache.org/jira/browse/CAMEL-11006] allows only customization via [DestinationGenerator|https://github.com/apache/camel/blob/master/tooling/swagger-rest-dsl-generator/src/main/java/org/apache/camel/generator/swagger/DestinationGenerator.java], i.e. the {{to}} part of the route generated, and gathers as much as possible from the specification file. The UI could offer the choice between built-in ({{direct:Swagger Operation ID}}) or a custom implementation of {{DestinationGenerator}}. {quote}Where is generally the swagger file available? A file hosted online? is it a maven artefact? A file inside the fuse project? A file in a specific project?{quote} Can't tell, to me it makes sense to be in the Fuse project, as you're modelling the specification file in the same place you're implementing it.	1	1	https://issues.jboss.org/browse/FUSETOOLS-2410
8960	FUSETOOLS-2410	bq. {quote}Is it a "one time" generation? or the user will update the Swagger API regularly?{quote} bq. bq. I think the user would like to update the Swagger (OpenAPI) specification and either re-run the generation or see the created REST DSL auto updated. as it is not a one-way generation, what about generating it at runtime only? and use directly the swagger file as an import in Camel dsl?	0	0	https://issues.jboss.org/browse/FUSETOOLS-2410
8961	FUSETOOLS-2410	If I follow what you're saying, that functionality exists, i.e. you can use the same {{swagger-rest-dsl-generator}} to generate a {{RoutesBuilder}} implementation at runtime or to generate Java DSL source code. Is then the question should this be supported in the tooling? I think supporting this has merits for users that just use the tooling without writing Java code.	1	1	https://issues.jboss.org/browse/FUSETOOLS-2410
8962	FUSETOOLS-2410	{quote}If I follow what you're saying, that functionality exists, i.e. you can use the same swagger-rest-dsl-generator to generate a RoutesBuilder implementation at runtime or to generate Java DSL source code. Is then the question should this be supported in the tooling? I think supporting this has merits for users that just use the tooling without writing Java code.{quote} my question is more, do we really need to generate the code at design time or not? If the user is using the swagger file, doesn't he/she want to always use this one and not dealing with a Camel Rest Dsl?	0	0	https://issues.jboss.org/browse/FUSETOOLS-2410
8963	FUSETOOLS-2410	{quote}Yeah, that's an option when user is willing to use swagger-rest-dsl-generator from Java, but how does the user that doesn't do any Java coding, i.e. user that only uses the IDE's design interface, use Swagger with Camel? There is no support (currently) in the camel-swagger-java to create consumers from Swagger/OpenAPI specifications.{quote} and if this support is added, will it still make sense to generate some code at design time?	0	0	https://issues.jboss.org/browse/FUSETOOLS-2410
8964	FUSETOOLS-2410	{quote}Might be useful to step up a level and consider the context in which this might be used. First, I think it's safe to say that the team developing this API is following a design-first approach.{quote} I personally never met a team doing API first without making it evolving, looks alike a lot like an old Waterfall model in which it is possible but if you say the use case is this one, seems fine to have a one-time generation. {quote}So they've used a designer tool or text editor or copied an existing API definition - now they want to implement the API behind that design. Pretty unlikely that the user wants to write REST DSL on their own and then manually map that to the Swagger definition, so one-time generation fits the bill nicely here.{quote} If there is a better tool/API definition format than Camel Rest DSL, why do we need the Camel Rest DSL at all? Why Camel Runtime is not consuming directly this api? {quote} A slightly more advanced case is when the original Swagger definition changes - design team made a mistake, new version with additional stuff, etc. In this case, it should be fairly straightforward to note which resources/verbs in the Swagger definition are not covered by the DSL and vice versa. Providing that info to the user would be enough to enable them to go in and make their own changes to the REST DSL. {quote} maybe "fairly straightforward", I have no idea of what is a swagger definition, a rest dsl definition and what are the difference between them. So if it's easy, yes would be fine. The best place for this validation is inside a Language Server Protocol (when we will have it one day...) but can be done in Fuse Tooling first. {quote} Automagic updates are always nice, but maybe not necessary here?{quote} that was a part of my question :-)	1	0	https://issues.jboss.org/browse/FUSETOOLS-2410
8965	FUSETOOLS-2410	{quote}I personally never met a team doing API first without making it evolving, looks alike a lot like an old Waterfall model in which it is possible but if you say the use case is this one, seems fine to have a one-time generation.{quote} I think you may be connecting dots that don't exist here. I didn't say that API definitions don't evolve. In fact, I specifically referenced the case where the API definition changes. {quote}If there is a better tool/API definition format than Camel Rest DSL, why do we need the Camel Rest DSL at all? Why Camel Runtime is not consuming directly this api?{quote} Swagger / Open API are independent of the API implementation. REST DSL is implementation-specific for Camel. In my mind, this JIRA is proposing a mechanism for linking definition and implementation. Seems like a good match to me, but definitely open to other suggestions and superior approaches. {quote}maybe "fairly straightforward", I have no idea of what is a swagger definition, a rest dsl definition and what are the difference between them. {quote} OK, well would probably be good to get a bit of an intro/overview from Zoran or another Fuse engineer that's a bit closer to REST DSL.	0	0	https://issues.jboss.org/browse/FUSETOOLS-2410

#	Requirement	Questions	Answered	Successfully Answered	URL
8966	FUSETOOLS-2355	SocketCleaner" #88 daemon prio=8 os_prio=1 tid=0x0000000026378800 nid=0x2a68 in Object.wait() [0x0000000034c5f000] java.lang.Thread.State: WAITING (on object monitor) at java.lang.Object.wait(Native Method) at sun.net.www.http.KeepAliveStreamCleaner.run(KeepAliveStreamCleaner.java:101) - locked <0x00000006c948dc28> (a sun.net.www.http.KeepAliveStreamCleaner) at java.lang.Thread.run(Thread.java:745) "[ThreadPool Manager] - Idle Thread" #83 daemon prio=5 os_prio=0 tid=0x0000000026373000 nid=0x339c in Object.wait() [0x000000003475f000] java.lang.Thread.State: WAITING (on object monitor) at java.lang.Object.wait(Native Method) at java.lang.Object.wait(Object.java:502) at org.eclipse.equinox.internal.util.impl.tpt.threadpool.Executor.run(Executor.java:106) - locked <0x00000006ca285018> (a org.eclipse.equinox.internal.util.impl.tpt.threadpool.Executor) "Worker-10" #82 prio=5 os_prio=0 tid=0x0000000021b00000 nid=0x4724 in Object.wait() [0x000000003465e000] java.lang.Thread.State: TIMED_WAITING (on object monitor) at java.lang.Object.wait(Native Method) at org.eclipse.core.internal.jobs.Semaphore.acquire(Semaphore.java:39) - locked <0x00000006cb738230> (a org.eclipse.core.internal.jobs.Semaphore) at org.eclipse.core.internal.jobs.OrderedLock.doAcquire(OrderedLock.java:170) at org.eclipse.core.internal.jobs.OrderedLock.acquire(OrderedLock.java:106) at org.eclipse.core.internal.jobs.OrderedLock.acquire(OrderedLock.java:82) at org.eclipse.wst.common.componentcore.internal.util.ComponentImplManager.createComponent(ComponentImplManager.java:210) at org.eclipse.wst.common.componentcore.ComponentCore.createComponent(ComponentCore.java:64) at org.eclipse.wst.common.componentcore.internal.builder.DependencyGraphImpl$GraphUpdateJob$1.run(DependencyGraphImpl.java:527) at org.eclipse.core.runtime.SafeRunner.run(SafeRunner.java:42) at org.eclipse.wst.common.componentcore.internal.builder.DependencyGraphImpl$GraphUpdateJob.run(DependencyGraphImpl.java:462) at org.eclipse.core.internal.jobs.Worker.run(Worker.java:55) "Worker-9" #81 prio=5 os_prio=0 tid=0x0000000021b04000 nid=0x191c in Object.wait() [0x000000003455e000] java.lang.Thread.State: TIMED_WAITING (on object monitor) at java.lang.Object.wait(Native Method) at org.eclipse.core.internal.jobs.WorkerPool.sleep(WorkerPool.java:188) - locked <0x00000006c5fe5a68> (a org.eclipse.core.internal.jobs.WorkerPool) at org.eclipse.core.internal.jobs.WorkerPool.startJob(WorkerPool.java:220) at org.eclipse.core.internal.jobs.Worker.run(Worker.java:52) "Worker-8" #78 prio=5 os_prio=0 tid=0x0000000021afe000 nid=0x1ec8 in Object.wait() [0x000000003337e000] java.lang.Thread.State: WAITING (on object monitor) at java.lang.Object.wait(Native Method) at java.lang.Object.wait(Object.java:502) at org.eclipse.core.internal.jobs.ThreadJob.waitForRun(ThreadJob.java:270) - locked <0x00000006cb676858> (a java.lang.Object) at org.eclipse.core.internal.jobs.ThreadJob.joinRun(ThreadJob.java:197) at org.eclipse.core.internal.jobs.ImplicitJobs.begin(ImplicitJobs.java:92) at org.eclipse.core.internal.jobs.JobManager.beginRule(JobManager.java:307) at org.eclipse.core.internal.resources.WorkManager.checkIn(WorkManager.java:121) at org.eclipse.core.internal.resources.Workspace.prepareOperation(Workspace.java:2188) at org.eclipse.core.internal.resources.Workspace.run(Workspace.java:2235) at org.eclipse.core.internal.resources.Workspace.run(Workspace.java:2267) at	0	0	https://issues.jboss.org/browse/FUSETOOLS-2355
8967	FUSETOOLS-2355	java.lang.Thread.State: TIMED_WAITING (on object monitor) at java.lang.Object.wait(Native Method) at org.eclipse.core.internal.jobs.Semaphore.acquire(Semaphore.java:39) - locked <0x00000006cb738230> (a org.eclipse.core.internal.jobs.Semaphore) at org.eclipse.core.internal.jobs.OrderedLock.doAcquire(OrderedLock.java:170) at org.eclipse.core.internal.jobs.OrderedLock.acquire(OrderedLock.java:106) at org.eclipse.core.internal.jobs.OrderedLock.acquire(OrderedLock.java:82) at org.eclipse.wst.common.componentcore.internal.util.ComponentImplManager.createComponent(ComponentImplManager.java:210) at org.eclipse.wst.common.componentcore.ComponentCore.createComponent(ComponentCore.java:64) at org.eclipse.wst.common.componentcore.internal.builder.DependencyGraphImpl$GraphUpdateJob$1.run(DependencyGraphImpl.java:527) at org.eclipse.core.runtime.SafeRunner.run(SafeRunner.java:42) at org.eclipse.wst.common.componentcore.internal.builder.DependencyGraphImpl$GraphUpdateJob.run(DependencyGraphImpl.java:462) at org.eclipse.core.internal.jobs.Worker.run(Worker.java:55) {noformat} {noformat} "Worker-8" #78 prio=5 os_prio=0 tid=0x0000000021afe000 nid=0x1ec8 in Object.wait() [0x000000003337e000] java.lang.Thread.State: WAITING (on object monitor) at java.lang.Object.wait(Native Method) at java.lang.Object.wait(Object.java:502) at org.eclipse.core.internal.jobs.ThreadJob.waitForRun(ThreadJob.java:270) - locked <0x00000006cb676858> (a java.lang.Object) at org.eclipse.core.internal.jobs.ThreadJob.joinRun(ThreadJob.java:197) at org.eclipse.core.internal.jobs.ImplicitJobs.begin(ImplicitJobs.java:92) at org.eclipse.core.internal.jobs.JobManager.beginRule(JobManager.java:307) at org.eclipse.core.internal.resources.WorkManager.checkIn(WorkManager.java:121) at org.eclipse.core.internal.resources.Workspace.prepareOperation(Workspace.java:2188) at org.eclipse.core.internal.resources.Workspace.run(Workspace.java:2235) at org.eclipse.core.internal.resources.Workspace.run(Workspace.java:2267) at org.eclipse.wst.common.project.facet.core.internal.FacetedProject.modify(FacetedProject.java:339) at org.eclipse.wst.common.project.facet.core.internal.FacetedProject.installProjectFacet(FacetedProject.java:255) at org.eclipse.wst.jsdt.web.core.internal.project.ConvertJob.runInWorkspace(ConvertJob.java:57) at org.eclipse.core.internal.resources.InternalWorkspaceJob.run(InternalWorkspaceJob.java:39) at org.eclipse.core.internal.jobs.Worker.run(Worker.java:55) {noformat} {noformat} "Worker-7" #77 prio=5 os_prio=0 tid=0x0000000021afd000 nid=0x2154 waiting for monitor entry [0x000000003327e000] java.lang.Thread.State: BLOCKED (on object monitor) at org.eclipse.wst.common.componentcore.internal.util.ComponentImplManager.createComponent(ComponentImplManager.java:223) - waiting to lock <0x00000006c8a0b1e8> (a org.eclipse.wst.common.componentcore.internal.util.ComponentImplManager) at org.eclipse.wst.common.componentcore.internal.util.ComponentImplManager.createComponent(ComponentImplManager.java:212) at org.eclipse.wst.common.componentcore.ComponentCore.createComponent(ComponentCore.java:64) at	0	0	https://issues.jboss.org/browse/FUSETOOLS-2355
8968	FUSETOOLS-2268	My opinion is that this is a bad approach. I don't see any benefit of using the jgit timestamps. I see only the mess of versions :(It's quite easy to check whether the code was modified since there are source bundles. And users really don't need to know which features were changed. Nice example is BPMN2 modeller - one feature is 1.3.1 and the other one is 1.3.2. What's the real version of the BPMN2 modeller?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2268
8969	FUSETOOLS-2268	From my point of view jgit timestamps make perfect sense. Here is an example of jbosstoos.as plugins from a users point of view (Eclipse -> Installation details etc): !as_jgit.png! You can see that all plugins are in version 3.2.3 but if you look more closely you'll see that e.g. as.core is timestamped based on commit from 2017-01-18, however as.ui is timestamped based on commit from 2017-02-15. [~apodhrad] not sure why is this a blocker for you. Can you maybe elaborate a bit more?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2268

#	Requirement	Questions	Answered	Successfully Answered	URL
8970	FUSETOOLS-2265	what is the status of this issue? 9.2 branch has been created but it is based on Neon. Should we rename it and create a 9.3 branch based on Oxygen? I think that we decide to not do that and use master with 10.0 for oxygen.	0	0	https://issues.jboss.org/browse/FUSETOOLS-2265
8971	FUSETOOLS-2263	Question: Is the Camel Catalog always available? I see the following issues: # user is working in an offline environment (probably just a small chance but its existing) # no catalog shipped with the runtime # runtime is remote and we need to grab catalog from a remote file system That brings up another question: Do we want to use the camel catalog or do we parse each component on the classpath. That would also affect the palette contents. If going that way we definitely need a way to add new components to the classpath in a visual way.	1	1	https://issues.jboss.org/browse/FUSETOOLS-2263
8972	FUSETOOLS-2263	Does this really block https://issues.jboss.org/browse/JBIDE-24243 ?	0	0	https://issues.jboss.org/browse/FUSETOOLS-2263
8973	FUSETOOLS-2257	I need to have a deeper look because it is trange that it is failing on a single slave RHEL-7.2-m1.xlarge-10.8.185.31 currently: {noformat} org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForAMQIT.testAMQSpringProjectCreation[2.17.3] 10 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForAMQIT.testAMQSpringProjectCreation[2.15.1.redhat-621117] 10 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForCBRIT.testCBRSpringCreation[2.17.3] 10 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForCBRIT.testCBRSpringCreation[2.15.1.redhat-621117] 9.6 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForCXFCodeFirstIT.testCXFCodeFirstSpringProjectCreation[2.17.3] 14 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForCXFCodeFirstIT.testCXFCodeFirstSpringProjectCreation[2.15.1.redhat-621117] 12 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForSimpleIT.testEmptySpringProjectCreation[2.17.3] 9.6 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForSimpleIT.testEmptySpringProjectCreationOnLocationOutsideWorkspace[2.17.3] 8.7 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForSimpleIT.testEmptySpringProjectCreation[2.15.1.redhat-621117] 9.9 sec 1 org.fusesource.ide.projecttemplates.tests.integration.wizards.FuseIntegrationProjectCreatorRunnableForSimpleIT.testEmptySpringProjectCreationOnLocationOutsideWorkspace[2.15.1.redhat-621117] 8.2 sec 1 {noformat} {noformat} Error Message Expecting empty but was:<[{"severity"=2, "sourceId"="JDT", "Creation time: "=1488288879705L, "resource affected"="/home/jenkins/workspace/Fuse-Tooling-9.2.x-Update-Site/sources/editor/tests/org.fusesource.ide.projecttemplates.tests.integration/target/work sp@cé/FuseIntegrationProjectCreatorRunnableForAMQIT-AMQSpringProject-2.17.3", "message"="The project cannot be built until build path errors are resolved", "categoryId"=10, "type: "="org.eclipse.jdt.core.problem"}, {"severity"=2, "sourceId"="JDT", "Creation time: "=1488288877301L, "classpathFileFormat"="false", "message"="Archive for required library: '/home/jenkins/.m2/repository/org/apache/camel/camel-test-spring/2.17.3/camel-test-spring-2.17.3.jar' in project 'FuseIntegrationProjectCreatorRunnableForAMQIT-AMQSpringProject-2.17.3' cannot be read or is not a valid ZIP file", "cycleDetected"="false", "type: "="org.eclipse.jdt.core.buildpath_problem", "outputOverlappingSource"="false", "location"="Build path", "arguments"="0:", "resource affected"="/home/jenkins/workspace/Fuse-Tooling-9.2.x-Update-Site/sources/editor/tests/org.fusesource.ide.projecttemplates.tests.integration/target/work sp@cé/FuseIntegrationProjectCreatorRunnableForAMQIT-AMQSpringProject-2.17.3", "id"=964, "categoryId"=10}]> {noformat}	1	1	https://issues.jboss.org/browse/FUSETOOLS-2257
8974	FUSETOOLS-2229	[~chirino] [~kcbabo] maybe you have a better idea? Should we remove some of them? Should we add some others?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2229
8975	FUSETOOLS-2223	[~chirino] What is the URl of the FIS catalog? Is it somthing in? https://repository.jboss.org/nexus/content/groups/ea/io/fabric8/archetypes/archetypes-catalog/ ? How can we determine the exact version? Same as some FIS bom?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2223
8976	FUSETOOLS-2223	Yes thats it. Note that the repo and version suffix change when we GA. Ga will use a .redhat- version instead of a .fuse suffix. The latest DR build is at: https://repository.jboss.org/nexus/content/groups/ea/io/fabric8/archetypes/archetypes-catalog/2.2.180.fuse-000008/archetypes-catalog-2.2.180.fuse-000008-archetype-catalog.xml If you want inspiration/code on how to do it, see: https://github.com/jboss-fuse/fuse-forge/blob/master/src/main/java/org/jboss/fuse/forge/addon/ui/FuseProjectSetupStep.java#L270	0	0	https://issues.jboss.org/browse/FUSETOOLS-2223
8977	FUSETOOLS-2223	I created https://issues.jboss.org/browse/JBIDE-24283 to have the Red Hat catalog with all the FIS archetypes available. If it get accepted I will change target version to 10.0 or maybe close it we might not even still want this issue, do we really want to provide several catalogs for each specific version of FIS?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2223
8978	FUSETOOLS-2197	How can we choose a template for the user? Starting from a blank project seems to be the more common use case	1	1	https://issues.jboss.org/browse/FUSETOOLS-2197
8979	FUSETOOLS-2191	The issue will be with matching that approach, since the Tabbed Properties view in Eclipse really doesn't have a way to do a decorator label built in. We can easily do a tooltip and error decorator as we've done with SwitchYard. The image below is from the binding properties for a scheduling binding. You'll notice the little red X on the left side of the field and the tooltip with the error message. !switchyard-scheduling-binding-properties.png|thumbnail! I would like to see other examples since we have many Fuse components that have a ton of properties and we're not really going to want to space them out by putting an error label between them. What would the error look like if it was on a higher field? Where does the error message go (the one at the bottom)? What's with that label on top? Some of the Fuse components are very property dense. !fuse-activemq-properties.png|thumbnail! Adding the X (error) and I (information) decorators to the left of the field is great and I agree that the red background on the field makes it unreadable - but I don't see how we can make this quite match what they're attempting in the PatternFly approach.	0	0	https://issues.jboss.org/browse/FUSETOOLS-2191
8980	FUSETOOLS-2180	Is that solved now or not? [~jmurphey]	1	1	https://issues.jboss.org/browse/FUSETOOLS-2180
8981	FUSETOOLS-2176	I believe this belongs here, to JBTIS, right? Can somebody please set the fix version and component and resolve again? Thanks!	0	0	https://issues.jboss.org/browse/FUSETOOLS-2176
8982	FUSETOOLS-2172	Can you double check? I think this issue is no longer occuring.	0	0	https://issues.jboss.org/browse/FUSETOOLS-2172
8983	FUSETOOLS-2160	[~rafaelcba] Even if it has not yet been validated by QE you can try this feature by using the latest version: http://download.jboss.org/jbosstools/neon/integration/updates/integration-stack/fuse-tooling/9.1.0.CI/ If you have time to test it and tell us if it fulfills the requirements for the workshop/trainings/demo that you are doing it will be nice. Do you have other blocking points for workshop/training/demo?	1	0	https://issues.jboss.org/browse/FUSETOOLS-2160

#	Requirement	Questions	Answered	Successfully Answered	URL
8984	FUSETOOLS-2154	[~chirino] To determine if we need to use the SpringBoot catalog or not, I based the condition of having the "camel-spring-boot-starter" artifact as dependency. Can you confirm that it is the good condition? If not, what is a good condition for it?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2154
8985	FUSETOOLS-2149	[~chirino] is detecting class extending org.apache.camel.spring.boot.FatJarRouter enough? Do we need more precise constraints? Do we have other cases?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2149
8986	FUSETOOLS-2149	there was some discussions to get rid of Maven dependency and to be able to support Camel project based on Gradle for instance. So we were trying to be based on File detection. I'm not sure where we are for that, looking for maven dependency can be easier. [~lhein] Do you have more information? What do you think?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2149
8987	FUSETOOLS-2142	What kind of alignment needs to happen here?	0	0	https://issues.jboss.org/browse/FUSETOOLS-2142
8988	FUSETOOLS-2140	{noformat} Failed tests: CamelIOHandlerIT.testLoadAndSaveOfSimpleModel:70->assertModelRoundTrip:96 Expecting: <"<?xml version="1.0" encoding="UTF-8"?> <beans xmlns="http://www.springframework.org/schema/beans" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://www.springframework.org/schema/beans http://www.springframework.org/schema/beans/spring-beans-2.5.xsd http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-spring.xsd"> <!-- some beans... --> <bean class="org.fusesource.camel.rider.sample.MyBean" id="devBean1"/> <bean class="org.fusesource.camel.rider.sample.AnotherBean" id="devBean2"/> <!-- some camel route stuff now... --> <camelContext autoStartup="true" id="context1" lazyLoadTypeConverters="false" streamCache="false" xmlns="http://camel.apache.org/schema/spring"> <endpoint id="something" uri="seda:myConfiguredEndpoint"/> <route> <description>edit file!</description> <from uri="seda:someWhere"/> <filter> <xpath>/foo/bar</xpath> <to uri="seda:anotherPlace2"/> </filter> </route> </camelContext> <!-- some other stuff here... --> </beans> "> to be equal to: <"<?xml version="1.0" encoding="UTF-8"?> <beans xmlns="http://www.springframework.org/schema/beans" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://www.springframework.org/schema/beans http://www.springframework.org/schema/beans/spring-beans-2.5.xsd http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-spring.xsd"> <!-- some beans... --> <bean class="org.fusesource.camel.rider.sample.MyBean" id="devBean1"/> <bean class="org.fusesource.camel.rider.sample.AnotherBean" id="devBean2"/> <!-- some camel route stuff now... --> <camelContext autoStartup="true" id="context1" lazyLoadTypeConverters="false" streamCache="false" xmlns="http://camel.apache.org/schema/spring"> <endpoint id="something" singleton="false" uri="seda:myConfiguredEndpoint"/> <route> <description>edit file!</description> <from uri="seda:someWhere"/> <filter> <xpath>/foo/bar</xpath> <to uri="seda:anotherPlace2"/> </filter> </route> </camelContext> <!-- some other stuff here... --> </beans> "> but was not. CamelIOHandlerIT.testLoadAndSaveOfSimpleModel:70->assertModelRoundTrip:96 Expecting: <"<?xml version="1.0" encoding="UTF-8"?> <beans xmlns="http://www.springframework.org/schema/beans" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://www.springframework.org/schema/beans http://www.springframework.org/schema/beans/spring-beans-2.5.xsd http://camel.apache.org/schema/spring http://camel.apache.org/schema/spring/camel-spring.xsd"> <!-- some beans... --> <bean class="org.fusesource.camel.rider.sample.MyBean" id="devBean1"/> <bean class="org.fusesource.camel.rider.sample.AnotherBean" id="devBean2"/> <!-- some camel route stuff now... --> <camelContext autoStartup="true" id="context1" lazyLoadTypeConverters="false" streamCache="false" xmlns="http://camel.apache.org/schema/spring"> <endpoint id="something" uri="seda:myConfiguredEndpoint"/> <route> <description>edit file!</description> <from uri="seda:someWhere"/> <filter> <xpath>/foo/bar</xpath> <to uri="seda:anotherPlace2"/> </filter> </route> </camelContext> <!-- some other stuff here... --> </beans> "> to be equal to: <"<?xml version="1.0" encoding="UTF-8"?> <beans xmlns="http://www.springframework.org/schema/beans" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://www.springframework.org/schema/beans	0	0	https://issues.jboss.org/browse/FUSETOOLS-2140
8989	FUSETOOLS-2140	this test org.fusesource.ide.camel.editor.integration.CamelEditorIT.createASecondRoute() detected that we can't add anymore a route. I reproduced manually. I think that it comes from the fact that now everything is declared in EIP and so the org.fusesource.ide.camel.editor.features.create.ext.CreateFigureFeature.isAttemptToCreateWiredFigure(AbstractCamelModelElement, ICreateContext) is returning true although before it was returning false --> the issue comes from teh catalog whcih contains some null values for input/output The JSON retrieved from Camel Catalog is missing the input/output attributes for the ones that have a null value: {noformat} { "model": { "kind": "model", "name": "endpoint", "title": "Endpoint", "description": "Camel endpoint configuration", "javaType": "org.apache.camel.spring.CamelEndpointFactoryBean", "label": "spring,configuration,endpoint" }, "properties": { "singleton": { "kind": "attribute", "required": "false", "type": "boolean", "javaType": "java.lang.Boolean", "deprecated": "true", "secret": "false", "defaultValue": "false", "description": "Not in use" }, "uri": { "kind": "attribute", "required": "true", "type": "string", "javaType": "java.lang.String", "deprecated": "false", "secret": "false", "description": "Sets the URI to use to resolve the endpoint. Notice that additional options can be configured using a series of property." }, "pattern": { "kind": "attribute", "required": "false", "type": "string", "javaType": "org.apache.camel.ExchangePattern", "enum": ["InOnly", "InOptionalOut", "InOut", "OutIn", "OutOnly", "OutOptionalIn", "RobustInOnly", "RobustOutOnly"], "deprecated": "true", "secret": "false", "description": "Sets the exchange pattern of the endpoint" }, "properties": { "kind": "element", "required": "true", "type": "array", "javaType": "java.util.List<org.apache.camel.model.PropertyDefinition>", "oneOf": ["property"], "deprecated": "false", "secret": "false", "description": "To configure additional endpoint options using a XML style which is similar as configuring Spring or Blueprint beans." }, "camelContextId": { "kind": "attribute", "required": "false", "type": "string", "javaType": "java.lang.String", "deprecated": "false", "secret": "false", "description": "Id of CamelContext to use if there are multiple CamelContexts in the same JVM" }, "id": { "kind": "attribute", "required": "false", "type": "string", "javaType": "java.lang.String", "deprecated": "false", "secret": "false", "description": "Sets the id of this node" } } } {noformat} answer from Claus is that it is normal that there is no input/output attributes. I kept it as is for now, so we have some null string values for these components. As all of them are new, I needed to do a special case of camelContext to eb able to add a route.	0	0	https://issues.jboss.org/browse/FUSETOOLS-2140
8990	FUSETOOLS-2136	I'm not sure what we is the best way to handle it. Tycho assumes that there is always a translation files and the default is plugin.properties. Should we create empty ones? Is there a way to say that there is no translation files?	0	0	https://issues.jboss.org/browse/FUSETOOLS-2136
8991	FUSETOOLS-2120	Seems the jar is too large to email or to attach. what to do?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2120
8992	FUSETOOLS-2120	[~tsedmik] Did you get the jar file and are you testing it?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2120

#	Requirement	Questions	Answered	Successfully Answered	URL
8993	FUSETOOLS-2119	{noformat} <janstey_mtg> would be good to add both: https://maven.repository.redhat.com/ga and https://maven.repository.redhat.com/earlyaccess/all <janstey_mtg> we do in the container as well <janstey_mtg> put ga first in list because that is the main one ... <lhein> what is https://access.redhat.com/maven-repository ? <janstey> it's just an informational web page <lhein> General availability repository: https://maven.repository.redhat.com/ga/ <lhein> Early-access repository: https://maven.repository.redhat.com/earlyaccess/all/ <lhein> You can choose between the general availability repository (GA releases) and the early access repository (Alpha and Beta builds). ... <lhein> so these are basically the ones we need <lhein> right? <janstey> yep {noformat}	1	1	https://issues.jboss.org/browse/FUSETOOLS-2119
8994	FUSETOOLS-2118	As further step it would be cool if * it could auto complete bean ID's - for instance when in camel route {{<process ref="...">}} it could list available beans implementing {{Processor}} * currently it lists bean IDs in Designer - ref fields, but it lists all beans, not just "matching" ones	0	0	https://issues.jboss.org/browse/FUSETOOLS-2118
8995	FUSETOOLS-2099	Hello [~lhein] [~apodhrad]. Would it be possible to incorporate this one any time soon? I believe it would be convenient just double click on route entry and see it's properties.	1	1	https://issues.jboss.org/browse/FUSETOOLS-2099
8996	FUSETOOLS-2095	Where do we define the checkstyle for the Fuse tooling? Or do we?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2095
8997	FUSETOOLS-2095	So [~aurelien.pupier] do we want to simply disable it for the SY tooling? Or what are you asking here?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2095
8998	FUSETOOLS-2095	I would prefer to deactivate it for now. If you activate the Eclipse checkstyle checker, you will notice that it reports a lot of violations even if these ones are not breaking the build. So to follow them is very complicated and they even are currently not compatible with Fuse Tooling (tabs vs spaces) i don't really want to enter in a long debate for this. i think the best is to deactivate it. Do you agree?	0	0	https://issues.jboss.org/browse/FUSETOOLS-2095
8999	FUSETOOLS-2075	How can I reproduce that? Usually the progress dialog can be hidden.	1	1	https://issues.jboss.org/browse/FUSETOOLS-2075
9000	FUSETOOLS-2074	Is that done?	0	0	https://issues.jboss.org/browse/FUSETOOLS-2074
9001	FUSETOOLS-2064	What will be the location for the item "No Runtime Selected" at the beginning? or according to order?	1	1	https://issues.jboss.org/browse/FUSETOOLS-2064
9002	FUSETOOLS-1983	Added a first shot at defining the template but can't get my categories to show up in the wizard. https://github.com/fusesource/fuseide/pull/644 What am I doing wrong? [~lhein] - I'm guessing it's something brain dead.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1983
9003	FUSETOOLS-1983	I merged the template. Are you planning to do a BP or Java DSL one too or is that not working on EAP?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1983
9004	FUSETOOLS-1979	is it ok to do an API break on ICustomPaletteEntry? (as we are moving to 8.0.0) Or do we prefer to modify the extension point and putting an optional attribute for another class responsible for the update? Edit: i'm wondering if using a second would be cleaner and simpler for user in fact (even without considering API break disadvantage of first proposal)	0	0	https://issues.jboss.org/browse/FUSETOOLS-1979
9005	FUSETOOLS-1979	Could you provide more details how this issue can be verified, please? I'm not sure which camel version is changed (the property camel.version or a version defined for an artifact?)	1	1	https://issues.jboss.org/browse/FUSETOOLS-1979
9006	FUSETOOLS-1977	[~lhein] but if it doesn't work as local camel context, we introduce an inconsistency and introduce a big usability issue. How can we, at least, proposed to not "launch as camel context"? Perhaps if no camel maven plugin is not defined in pom.xml?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1977
9007	FUSETOOLS-1976	Do we have a specific feedback when we can drop or not on another node? Do we have a specific feedback when we drop on a connection? To be complete, I feel that we miss only: - to be able to "add before"/"Prepend" when an element has outgoing transition but no incoming. Then it will allow to place a new element from the palette to every place on the flow. - to handle containers, how to Prepend/Append for them?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1976
9008	FUSETOOLS-1976	for 1. if you drop on the timer:start, this is normal behavior. To insert between timer and bean, you need to drop on the connection for 2. What is your issue here? when you drop a Log on teh connection between bean and direct, it doesn't work?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1976
9009	FUSETOOLS-1957	I started to work on it. Several issues still. Test are emphasising some of them. You can check it on this branch: https://github.com/apupier/fuseide/tree/FUSETOOLS-1957-CXF-ContractFirstTemplate	0	0	https://issues.jboss.org/browse/FUSETOOLS-1957
9010	FUSETOOLS-1940	I feel that with the time that we have we need to do: # introduce microversions in facets (currently we only have major.minor) --> so which one do we create? 2.15.1 and 2.17.0 only? or should we create 2.15.1, 2.15.1-redhat-621084, 2.17.0 and 2.17.0.redhat-630077? # adapt camel configurator to install the correct facet --> the idea is to default to a generated catalog based on closest version? the 2 other solutions seems far more complicated: # in the model we should check if we have a supported service for that version and if not we then generate one using the json catalog we then grab via maven? # (or alternatively refactor the model to rely only on the catalog jar dependency and download missing ones on the fly if possible) --> I would say that if the version doesn't exist in our embedded catalogs, we then use the closer previous version.	0	0	https://issues.jboss.org/browse/FUSETOOLS-1940
9011	FUSETOOLS-1932	The org.jboss.tools.forge plugins are available as a standard part for JBT core - they can be found here: http://download.jboss.org/jbosstools/neon/staging/updates/core/4.4.1.AM2/plugins/ The forge addon plugins are also from JBoss but they're not part of core or org.jboss.tools. Their dependency pom fragment looks like: {code} <dependency> <groupId>io.fabric8.forge</groupId> <artifactId>camel</artifactId> <classifier>forge-addon</classifier> <version>${version}</version> </dependency> {code} So - they're their own animal. [~nickboldt] - would we just create a new dependency in the IS TP with the forge-addon repo? Do we need to mirror it?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1932
9012	FUSETOOLS-1932	"Just" 17Mb. Can the addon be reviewed if it can be smaller? cc: [~davsclaus] ?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1932
9013	FUSETOOLS-1932	I'd need to understand a bit more. Some immediate questions: 1. +17MB. What's the size of JBDS Forge today? (+17MB on a 3MB distro is a big deal. +17MB on a 200MB distro is not.) 2. If this addon is added to the distro, how do we keep them properly in sync? e.g., if the camel commands addon is serviced/updated, does that force us to make a new release of JDBS Forge? I think I have others, but let's start with those... :)	0	0	https://issues.jboss.org/browse/FUSETOOLS-1932
9014	FUSETOOLS-1932	Having a way to easily browse and install would be a convenient feature for Forge in general and could save space for JBoss Tools. For Fuse Tooling itself, I think that it will be better to include the specific Camel addon as it is relevant. I don't undestand very well what you said [~gastaldi]: "The Forge plugin provides an extension point that allows you to specify additional addon repositories. AFAIK this is how the Windup team do to load their addons in the embedded runtime. If the camel team already provides an Eclipse plugin, it could implement that extension point. If not, perhaps we can consider another strategy." So it means that we can implement an extension in our Fuse Tooling project to include the addon? What do you mean by "camel team already provides an Eclipse plugin"?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1932

#	Requirement	Questions	Answered	Successfully Answered	URL
9015	FUSETOOLS-1932	I believe I need to create a new eclipse plugin that bundles these addons and implements the Forge plugin extension point. What repository is recommended to contain that?	1	0	https://issues.jboss.org/browse/FUSETOOLS-1932
9016	FUSETOOLS-1932	I don't have much experience in creating Eclipse plugins (I only maintain the jbosstools-forge one but I am not sure what files are required so Eclipse can recognize it). Can someone create the basic structure in there so I can add the business code? This is an example of what should be implemented in the plugin so the JBoss Tools Forge plugin can call it: https://github.com/windup/windup-eclipse-plugin/blob/master/plugins/org.jboss.tools.windup.runtime/src/org/jboss/tools/windup/runtime/internal/FurnaceRepositoryProvider.java	1	0	https://issues.jboss.org/browse/FUSETOOLS-1932
9017	FUSETOOLS-1908	[~jmurphey] Can you update the documentation with a note to run Java DSL based samples? in this case the "Goals" field should be exec:java https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Fuse/6.0/html/Using_the_Fuse_Developer_Tools/files/RiderRunningMaven.html	1	1	https://issues.jboss.org/browse/FUSETOOLS-1908
9018	FUSETOOLS-1896	Can u explain the usecase a bit more? When using the import wizard for SAP it installs the SAP binaries as plugin into your eclipse. That makes the bundles available for all projects if they reference them. Why should I need to declare them in the POM? What is the use case?	1	0	https://issues.jboss.org/browse/FUSETOOLS-1896
9019	FUSETOOLS-1889	Would need to see what it's like at moment to compare - is this image the merged version?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1889
9020	FUSETOOLS-1865	I tested this issue with JBDSIS 9.0.1.CR3. It seems to be still broken (see attached video). The Progress bar disappears but project creation is not completed. The UI freeze is not convenient. Can we move project creation process into background (_Run in background_ button)?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1865
9021	FUSETOOLS-1837	Remember you are using the Maven M2E Project import wizard. How can this wizard know about the nature of an imported project or set a specific project nature at all. That doesn't sound very logical to me.	0	0	https://issues.jboss.org/browse/FUSETOOLS-1837
9022	FUSETOOLS-1798	yes I saw this. Thanks for opening the feature request. I'm currently on several other topics so not in my top list. What is the priority for you? Is it a blocking point? If yes, can you link the feature that is blocked by this one please? We might need to discuss about it with [~lhein] in this case.	1	0	https://issues.jboss.org/browse/FUSETOOLS-1798
9023	FUSETOOLS-1784	This works only for the first time you add a destination/server. But not for this use case: 1. Add a destination 2. Select tab 'Pool' 3. Add another destintion 4. The tab 'Pool' is selected for this new destination Tested with JBDS-IS 9.0.1.CR2 (JBoss Fuse SAP Tool Suite 8.0.0.Beta3-v20160713-1930-B63)	1	1	https://issues.jboss.org/browse/FUSETOOLS-1784
9024	FUSETOOLS-1740	Hello [~lhein], would it be possible to use Camel's capability to dump routes from camel context as XML file ? One can see this functionality in hawtio web console - there is tab for route diagram and also for "Source". Internally it calls [ManagedCamelContext.dumpRoutesAsXml|https://github.com/apache/camel/blob/master/camel-core/src/main/java/org/apache/camel/management/mbean/ManagedCamelContext.java#L408] I tried this small example and it seems to work ([^java-visualise-demo.zip]) !java-dsl-to-xml.png|thumbnail! {code:java} // get camel context with required route CamelContext context = new DefaultCamelContext(); CamelRoute r = new CamelRoute(); r.addRoutesToCamelContext(context); // get routes definition from context List<RouteDefinition> definitions = context.getRouteDefinitions(); RoutesDefinition routes = new RoutesDefinition(); routes.setRoutes(definitions); // dump given routes to xml String xmlContent = ModelHelper.dumpModelAsXml(context, routes); System.out.println(xmlContent); // try to open generated xml in camel editor - it should visualise routes Files.write(Paths.get("src/main/resources/routes-dump.xml"), xmlContent.getBytes(Charset.forName("utf-8"))); {code} There are definitely cases where this simple approach would fail, like this example with CDI - it creates one camel route for each entry injected in {{periods}} {code:java} @Inject private List<Period> periods; @Override public void configure() throws Exception { // Main route (for each period) for (Period period : periods) { final String periodNameWithDash = period.toString().replaceAll(" ", "-"); from(String.format("quartz2://timedRoutes/%s?cron=%s", period.name(), period.cronCamelFormat())) .id(Constants.PREFIX + "poll-" + periodNameWithDash) {code} What do you think about this approach?	1	0	https://issues.jboss.org/browse/FUSETOOLS-1740
9025	FUSETOOLS-1700	Aurelien, can you take care of this one? I think there is already a method to cut the text. See https://github.com/fusesource/fuseide/blob/7.3.x/editor/plugins/org.fusesource.ide.camel.editor/src/org/fusesource/ide/camel/editor/features/add/AddNodeFeature.java#L95	1	1	https://issues.jboss.org/browse/FUSETOOLS-1700
9026	FUSETOOLS-1665	Do we need to remove also the select box to choose the project in the transformation wizard? !screenshot-1.png!	0	0	https://issues.jboss.org/browse/FUSETOOLS-1665
9027	FUSETOOLS-1650	What happens if we disconnect from the JMX server...shouldn't that also close the editor then? And what if I want to look at it after disconnecting? I am rather unsure if its really a good idea to close it right away.	1	0	https://issues.jboss.org/browse/FUSETOOLS-1650
9028	FUSETOOLS-1643	I'd say the Ugly picture is actually the more clean and straight forward here. Wdyt?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1643
9029	FUSETOOLS-1603	What is the goal of having the version defined in the project? isn't the version in the pom the right place to determine it? If version in the pom and in the project not matching, we should provide a warning validation?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1603
9030	FUSETOOLS-1602	Can the status of producer/consumer can be modified during design time? If yes, we need to think about the validation. If a property field is not valid but not shown due to its producer/consumer facet, what should we do? Is it working at deployments? How to remove/update these properties?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1602
9031	FUSETOOLS-1585	Spring IDE is providing useful Bean Editor but only for Spring beans. For Blueprint beans, that have a different syntax, I didn't find any editor. Is it worthy to propose Beans editor only for Spring in a first time?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1585
9032	FUSETOOLS-1585	[~aurelien.pupier] or [~apodhrad] Aurelien, do you by chance have a copy of the example project you were using for your screen shots? Or can either of you point me at a good example somewhere on the internet? I'm just trying to find a handful of simple examples so I can start working through this. It appears that you basically are exposing the bean declaration and some properties? Something along the lines of: {code} <bean id="myBean" class="org.testspringbeans.MyBean"> <property name="somename" value="Bob Hope" /> <property name="somevalue" value="30" /> </bean> {code}	1	1	https://issues.jboss.org/browse/FUSETOOLS-1585
9033	FUSETOOLS-1567	[~bfitzpat] suggested in the mailing list to have an enhancement to suppor thi suse case: "At this point, the Fuse Transformation Editor is not meant as a general purpose dozer configuration editor. It's not a use case we have tested. At this point it only works with dozer configurations inside Camel projects. " What is the cost of the enhancement? Is it possible in a first time to not allow user to open the editor when the format is not supported?	1	0	https://issues.jboss.org/browse/FUSETOOLS-1567
9034	FUSETOOLS-1566	Aurelien, can you check and resolve if all is done?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1566

#	Requirement	Questions	Answered	Successfully Answered	URL
9035	FUSETOOLS-1566	I started work of externalizing String and it is just a matter of time to finish it. for the point 4, did you have some specific place to check? for the point 6, there is no use of CamelIOHandler, do you mean to get rid of CamelConfigBuilder?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1566
9036	FUSETOOLS-1560	when it tries to add one. I'm getting this. It looks like it's trying to insert the node before a text node in line 147. Even if I add a "<dataFormats/>" node by hand into the XML it's failing. I'm also unable to add anything via the Configuration tab of the editor, so am wondering about that as well. {code} org.w3c.dom.DOMException: NOT_FOUND_ERR: An attempt is made to reference a node in a context where it does not exist. at org.apache.xerces.dom.ParentNode.internalInsertBefore(Unknown Source) at org.apache.xerces.dom.ParentNode.insertBefore(Unknown Source) at org.fusesource.ide.camel.model.service.core.model.CamelContextElement.addDataFormat(CamelContextElement.java:147) at org.jboss.tools.fuse.transformation.core.camel.CamelConfigBuilder.createJsonDataFormat(CamelConfigBuilder.java:200) at org.jboss.tools.fuse.transformation.core.camel.CamelConfigBuilder.createDataFormat(CamelConfigBuilder.java:66) at org.jboss.tools.fuse.transformation.editor.wizards.NewTransformationWizard$1.run(NewTransformationWizard.java:137) at org.eclipse.jface.operation.ModalContext.runInCurrentThread(ModalContext.java:463) at org.eclipse.jface.operation.ModalContext.run(ModalContext.java:371) at org.eclipse.jface.wizard.WizardDialog.run(WizardDialog.java:1002) at org.jboss.tools.fuse.transformation.editor.wizards.NewTransformationWizard.performFinish(NewTransformationWizard.java:206) at org.eclipse.jface.wizard.WizardDialog.finishPressed(WizardDialog.java:799) at org.eclipse.jface.wizard.WizardDialog.buttonPressed(WizardDialog.java:429) at org.eclipse.jface.dialogs.Dialog$2.widgetSelected(Dialog.java:619) at org.eclipse.swt.widgets.TypedListener.handleEvent(TypedListener.java:248) at org.eclipse.swt.widgets.EventTable.sendEvent(EventTable.java:84) at org.eclipse.swt.widgets.Display.sendEvent(Display.java:4362) at org.eclipse.swt.widgets.Widget.sendEvent(Widget.java:1113) at org.eclipse.swt.widgets.Display.runDeferredEvents(Display.java:4180) at org.eclipse.swt.widgets.Display.readAndDispatch(Display.java:3769) at org.eclipse.jface.window.Window.runEventLoop(Window.java:827) at org.eclipse.jface.window.Window.open(Window.java:803) at org.jboss.tools.fuse.transformation.extensions.DataMapperEndpointFigureFeature.createNode(DataMapperEndpointFigureFeature.java:82) at org.fusesource.ide.camel.editor.features.create.ext.CreateFigureFeature.create(CreateFigureFeature.java:192) at org.fusesource.ide.camel.editor.features.create.ext.CreateEndpointFigureFeature.create(CreateEndpointFigureFeature.java:99) at org.jboss.tools.fuse.transformation.extensions.DataMapperEndpointFigureFeature.create(DataMapperEndpointFigureFeature.java:57) at org.eclipse.graphiti.features.impl.AbstractCreateFeature.execute(AbstractCreateFeature.java:100) at org.eclipse.graphiti.internal.command.GenericFeatureCommandWithContext.execute(GenericFeatureCommandWithContext.java:64) at org.eclipse.graphiti.internal.command.GFPreparableCommand.doExecute(GFPreparableCommand.java:37) at org.eclipse.emf.transaction.RecordingCommand.execute(RecordingCommand.java:135) at org.eclipse.graphiti.ui.internal.editor.GFWorkspaceCommandStackImpl.execute(GFWorkspaceCommandStackImpl.java:121) at org.eclipse.emf.transaction.impl.AbstractTransactionalCommandStack.execute(AbstractTransactionalCommandStack.java:219) at	1	1	https://issues.jboss.org/browse/FUSETOOLS-1560
9037	FUSETOOLS-1560	exception: com/sun/xml/bind/api/ErrorListener !STACK 0 java.lang.NoClassDefFoundError: com/sun/xml/bind/api/ErrorListener at java.lang.ClassLoader.defineClass1(Native Method) at java.lang.ClassLoader.defineClass(ClassLoader.java:760) at org.eclipse.osgi.internal.loader.ModuleClassLoader.defineClass(ModuleClassLoader.java:272) at org.eclipse.osgi.internal.loader.classpath.ClasspathManager.defineClass(ClasspathManager.java:632) at org.eclipse.osgi.internal.loader.classpath.ClasspathManager.findClassImpl(ClasspathManager.java:588) at org.eclipse.osgi.internal.loader.classpath.ClasspathManager.findLocalClassImpl(ClasspathManager.java:540) at org.eclipse.osgi.internal.loader.classpath.ClasspathManager.findLocalClass(ClasspathManager.java:527) at org.eclipse.osgi.internal.loader.ModuleClassLoader.findLocalClass(ModuleClassLoader.java:324) at org.eclipse.osgi.internal.loader.BundleLoader.findLocalClass(BundleLoader.java:327) at org.eclipse.osgi.internal.loader.sources.SingleSourcePackage.loadClass(SingleSourcePackage.java:36) at org.eclipse.osgi.internal.loader.BundleLoader.findClassInternal(BundleLoader.java:398) at org.eclipse.osgi.internal.loader.BundleLoader.findClass(BundleLoader.java:352) at org.eclipse.osgi.internal.loader.BundleLoader.findClass(BundleLoader.java:344) at org.eclipse.osgi.internal.loader.ModuleClassLoader.loadClass(ModuleClassLoader.java:160) at java.lang.ClassLoader.loadClass(ClassLoader.java:357) at org.jboss.tools.fuse.transformation.editor.internal.wizards.XMLPage$6.handleValueChange(XMLPage.java:422) at org.eclipse.core.databinding.observable.value.ValueChangeEvent.dispatch(ValueChangeEvent.java:63) at org.eclipse.core.databinding.observable.ChangeManager.fireEvent(ChangeManager.java:119) at org.eclipse.core.databinding.observable.value.DecoratingObservableValue.fireValueChange(DecoratingObservableValue.java:57) at org.eclipse.core.databinding.observable.value.DecoratingObservableValue.handleValueChange(DecoratingObservableValue.java:99) at org.eclipse.core.databinding.observable.value.DecoratingObservableValue$1.handleValueChange(DecoratingObservableValue.java:72) at org.eclipse.core.databinding.observable.value.ValueChangeEvent.dispatch(ValueChangeEvent.java:63) at org.eclipse.core.databinding.observable.ChangeManager.fireEvent(ChangeManager.java:119) at org.eclipse.core.databinding.observable.value.AbstractObservableValue.fireValueChange(AbstractObservableValue.java:74) at org.eclipse.core.internal.databinding.property.value.SimplePropertyObservableValue.notifyIfChanged(SimplePropertyObservableValue.java:127) at org.eclipse.core.internal.databinding.property.value.SimplePropertyObservableValue.access$2(SimplePropertyObservableValue.java:119) at org.eclipse.core.internal.databinding.property.value.SimplePropertyObservableValue$1$1.run(SimplePropertyObservableValue.java:69) at org.eclipse.core.databinding.observable.Realm$1.run(Realm.java:150) at org.eclipse.core.runtime.SafeRunner.run(SafeRunner.java:42) at org.eclipse.core.databinding.observable.Realm.safeRun(Realm.java:154) at org.eclipse.core.databinding.observable.Realm.exec(Realm.java:172) at	1	1	https://issues.jboss.org/browse/FUSETOOLS-1560

#	Requirement	Questions	Answered	Successfully Answered	URL
9038	FUSETOOLS-1560	my local environment. That leaves just the "Other" support in the tooling. I'm working through this now: {code} !ENTRY org.eclipse.ui 4 0 2016-01-21 09:18:36.310 !MESSAGE Unhandled event loop exception !STACK 0 java.lang.NullPointerException at org.jboss.tools.fuse.transformation.editor.internal.wizards.OtherPage.initialize(OtherPage.java:269) at org.jboss.tools.fuse.transformation.editor.internal.wizards.OtherPage.setVisible(OtherPage.java:320) at org.eclipse.jface.wizard.WizardDialog.updateForPage(WizardDialog.java:1229) at org.eclipse.jface.wizard.WizardDialog.access$4(WizardDialog.java:1208) at org.eclipse.jface.wizard.WizardDialog$8.run(WizardDialog.java:1197) at org.eclipse.swt.custom.BusyIndicator.showWhile(BusyIndicator.java:70) at org.eclipse.jface.wizard.WizardDialog.showPage(WizardDialog.java:1194) at org.eclipse.jface.wizard.WizardDialog.nextPressed(WizardDialog.java:885) at org.eclipse.jface.wizard.WizardDialog.buttonPressed(WizardDialog.java:425) at org.eclipse.jface.dialogs.Dialog$2.widgetSelected(Dialog.java:619) at org.eclipse.swt.widgets.TypedListener.handleEvent(TypedListener.java:248) at org.eclipse.swt.widgets.EventTable.sendEvent(EventTable.java:84) at org.eclipse.swt.widgets.Display.sendEvent(Display.java:4362) at org.eclipse.swt.widgets.Widget.sendEvent(Widget.java:1113) at org.eclipse.swt.widgets.Display.runDeferredEvents(Display.java:4180) at org.eclipse.swt.widgets.Display.readAndDispatch(Display.java:3769) at org.eclipse.jface.window.Window.runEventLoop(Window.java:827) at org.eclipse.jface.window.Window.open(Window.java:803) at org.jboss.tools.fuse.transformation.extensions.DataMapperEndpointFigureFeature.createNode(DataMapperEndpointFigureFeature.java:82) at org.fusesource.ide.camel.editor.features.create.ext.CreateFigureFeature.create(CreateFigureFeature.java:192) at org.fusesource.ide.camel.editor.features.create.ext.CreateEndpointFigureFeature.create(CreateEndpointFigureFeature.java:99) at org.jboss.tools.fuse.transformation.extensions.DataMapperEndpointFigureFeature.create(DataMapperEndpointFigureFeature.java:57) at org.eclipse.graphiti.features.impl.AbstractCreateFeature.execute(AbstractCreateFeature.java:100) at org.eclipse.graphiti.internal.command.GenericFeatureCommandWithContext.execute(GenericFeatureCommandWithContext.java:64) at org.eclipse.graphiti.internal.command.GFPreparableCommand.doExecute(GFPreparableCommand.java:37) at org.eclipse.emf.transaction.RecordingCommand.execute(RecordingCommand.java:135) at org.eclipse.graphiti.ui.internal.editor.GFWorkspaceCommandStackImpl.execute(GFWorkspaceCommandStackImpl.java:121) at org.eclipse.emf.transaction.impl.AbstractTransactionalCommandStack.execute(AbstractTransactionalCommandStack.java:219) at org.eclipse.graphiti.internal.command.CommandExec.executeCommand(CommandExec.java:85) at org.eclipse.graphiti.ui.internal.command.CreateModelObjectCommand.execute(CreateModelObjectCommand.java:49) at org.eclipse.graphiti.ui.internal.editor.EmfOnGefCommand.execute(EmfOnGefCommand.java:54) at org.eclipse.graphiti.internal.command.GFPreparableCommand2.doExecute(GFPreparableCommand2.java:43) at org.eclipse.emf.transaction.RecordingCommand.execute(RecordingCommand.java:135) at	1	1	https://issues.jboss.org/browse/FUSETOOLS-1560
9039	FUSETOOLS-1555	I think that this issue should be transformed in "providing a Template using REST DSL" What do you think?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1555
9040	FUSETOOLS-1549	Hiram, isn't that more a runtime task for the camel cdi component instead of the tooling? Or can you elaborate how tooling is involved here?	1	0	https://issues.jboss.org/browse/FUSETOOLS-1549
9041	FUSETOOLS-1521	Can you take a look at that Jon?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1521
9042	FUSETOOLS-1510	I tried to reproduce this issue with the latest Fuse Tooling (_7.3.2.FINAL_, part of _JBDSIS 8.0.7.GA_) for _JBDS 8.1.0.GA_ and the following runtimes: - JBoss Fuse 6.1.1.GA - JBoss Fuse 6.2.1.GA I can create and start a runtime. Also I can browse the runtime in _JMX Navigator_ view. --> *Cannot reproduce Bug* Did you set a user for JBoss Fuse? see https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Fuse/6.2.1/html/Installation_on_Apache_Karaf/InstallAddUser.html#topic-34584 Did you use an appropriate server type? In attached screenshot I see _Apache Karaf_ instead of _JBoss Fuse_.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1510
9043	FUSETOOLS-1488	[~jpav]there is a merged PR attached to this issue. Is it already fixed?	1	0	https://issues.jboss.org/browse/FUSETOOLS-1488
9044	FUSETOOLS-1434	I'm still confused about this issue. The linked pull request has no actual effect on the Fuse Tooling (commented code) and HL7 support via HAPI java object models and _Other_ type in Data Transformation is there from the very beginning. I could not find any related commit in fuseide repository. Am I wrong? Can you edit name of this issue to be more descriptive? I'm still for reopening this issue and set the fix version to 9.0.0.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1434
9045	FUSETOOLS-1434	isn't the priority for HL7 has been lowered? Can we put it to major or low instead of critical?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1434
9046	FUSETOOLS-1430	[~jpav] and I talked about this on Tuesday a bit at our weekly meeting... but I want to document something here. My initial concern holds true. Right now we don't have a solid way of interpreting which file the user used to create their model in the case of JSON or XML. For Java and Other, the model is used -as is-, so it is automatically loaded each time the transformation editor is opened. We talked about stashing the path in the project metadata, but that seems like an incomplete solution for a few reasons now that I've been thinking about it. There's a few potential user paths to consider. 1. User creates new Transformation in a Camel project. No transformation has existed before. We have the opportunity stash the source and target for the models in the project metadata. 2. User imports a quickstart with existing models and transformation file. Right now that doesn't exist in the quickstarts, so we would need to add the metadata for proper importing and I'm not sure if m2e will properly process it. 3. User creates project in a previous version of the editor and wants to update the model. Same issues as #2. What do we do in those cases? I think it may be possible simply to look for the class name as a file name (either as *.xml, *.xsd, or *.json)... It would at least give us a guess. Still working on this.	0	0	https://issues.jboss.org/browse/FUSETOOLS-1430

#	Requirement	Questions	Answered	Successfully Answered	URL
9047	FUSETOOLS-1422	Transformation Editor fails to create the part's controls *_<camelContext>_ element is not present in a Camel XML file *Steps:* 1. Create a new SwitchYard Project. Specify no runtime (select "<None>" as the Target Runtime in the SwitchYard Version Details) and type "2.1.0.redhat-630133" as the Library Version 2. Drag a Camel (XML) Implementation onto the composite (box with light blue background in the SwitchYard editor). 3. In the wizard, create a new service interface (whatever you want to call it) and hit Finish. 4. Open the generated "route.xml" file in the Camel editor. 5. Drag and drop the Transformation->Data Transformation item from the palette onto the canvas. 6. Finish the wizard (use xml2xml transformation from the following XML example) {code} <?xml version="1.0"?> <order> <customer id="A0001"> <name>Antwerp Zoo</name> <city>Antwerp</city> <country>BE</country> </customer> <date>2012-03-01</date> <orderlines> <orderline> <article id="A0001"> <description>Aardvark</description> </article> <quantity>1</quantity> </orderline> <orderline> <article id="A0011"> <description>Alpaca</description> </article> <quantity>10</quantity> </orderline> </orderlines> </order> {code} *Error* {code} eclipse.buildId=9.1.0.GA-v20160414-0124-B497 java.version=1.8.0_91 java.vendor=Oracle Corporation BootLoader constants: OS=linux, ARCH=x86_64, WS=gtk, NL=en_US Framework arguments: -product com.jboss.devstudio.core.product -pluginCustomization /home/tsedmik/devel/git/jbds-installer/jbds-9.1.0.GA_jbdsis-9.0.1.CR2/resources/pluginCustomization.ini Command-line arguments: -os linux -ws gtk -arch x86_64 -product com.jboss.devstudio.core.product -pluginCustomization /home/tsedmik/devel/git/jbds-installer/jbds-9.1.0.GA_jbdsis-9.0.1.CR2/resources/pluginCustomization.ini -data /home/tsedmik/devel/git/jbds-installer/jbds-9.1.0.GA_jbdsis-9.0.1.CR2/target/tmp -clean org.eclipse.e4.ui.workbench Error Fri Jul 29 13:36:45 CEST 2016 java.lang.RuntimeException: java.lang.ClassNotFoundException: generated_1469792188714.Order at org.jboss.tools.fuse.transformation.core.dozer.DozerMapperConfiguration.loadModel(DozerMapperConfiguration.java:433) at org.jboss.tools.fuse.transformation.core.dozer.DozerMapperConfiguration.getSourceModel(DozerMapperConfiguration.java:380) at org.jboss.tools.fuse.transformation.editor.internal.util.TransformationManager.rootSourceModel(TransformationManager.java:195) at org.jboss.tools.fuse.transformation.editor.internal.SourceTabFolder.<init>(SourceTabFolder.java:44) at org.jboss.tools.fuse.transformation.editor.TransformationEditor.createPartControl(TransformationEditor.java:231) at org.eclipse.ui.internal.e4.compatibility.CompatibilityPart.createPartControl(CompatibilityPart.java:151) at org.eclipse.ui.internal.e4.compatibility.CompatibilityEditor.createPartControl(CompatibilityEditor.java:99) at org.eclipse.ui.internal.e4.compatibility.CompatibilityPart.create(CompatibilityPart.java:341) at sun.reflect.GeneratedMethodAccessor216.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.eclipse.e4.core.internal.di.MethodRequestor.execute(MethodRequestor.java:56) at org.eclipse.e4.core.internal.di.InjectorImpl.processAnnotated(InjectorImpl.java:898) at org.eclipse.e4.core.internal.di.InjectorImpl.processAnnotated(InjectorImpl.java:879) at org.eclipse.e4.core.internal.di.InjectorImpl.inject(InjectorImpl.java:121) at org.eclipse.e4.core.internal.di.InjectorImpl.internalMake(InjectorImpl.java:345) at org.eclipse.e4.core.internal.di.InjectorImpl.make(InjectorImpl.java:264) at	1	1	https://issues.jboss.org/browse/FUSETOOLS-1422
9048	FUSETOOLS-1422	java.lang.InterruptedException: sleep interrupted at java.lang.Thread.sleep(Native Method) at org.eclipse.core.internal.jobs.JobManager.join(JobManager.java:1002) at org.jboss.tools.fuse.transformation.editor.TransformationEditor.waitJavaBuild(TransformationEditor.java:389) at org.jboss.tools.fuse.transformation.editor.TransformationEditor.init(TransformationEditor.java:376) at org.eclipse.ui.internal.EditorReference.initialize(EditorReference.java:361) at org.eclipse.ui.internal.e4.compatibility.CompatibilityPart.create(CompatibilityPart.java:319) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.eclipse.e4.core.internal.di.MethodRequestor.execute(MethodRequestor.java:56) at org.eclipse.e4.core.internal.di.InjectorImpl.processAnnotated(InjectorImpl.java:898) at org.eclipse.e4.core.internal.di.InjectorImpl.processAnnotated(InjectorImpl.java:879) at org.eclipse.e4.core.internal.di.InjectorImpl.inject(InjectorImpl.java:121) at org.eclipse.e4.core.internal.di.InjectorImpl.internalMake(InjectorImpl.java:345) at org.eclipse.e4.core.internal.di.InjectorImpl.make(InjectorImpl.java:264) at org.eclipse.e4.core.contexts.ContextInjectionFactory.make(ContextInjectionFactory.java:162) at org.eclipse.e4.ui.internal.workbench.ReflectionContributionFactory.createFromBundle(ReflectionContributionFactory.java:104) at org.eclipse.e4.ui.internal.workbench.ReflectionContributionFactory.doCreate(ReflectionContributionFactory.java:73) at org.eclipse.e4.ui.internal.workbench.ReflectionContributionFactory.create(ReflectionContributionFactory.java:55) at org.eclipse.e4.ui.workbench.renderers.swt.ContributedPartRenderer.createWidget(ContributedPartRenderer.java:129) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.createWidget(PartRenderingEngine.java:971) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.safeCreateGui(PartRenderingEngine.java:640) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.safeCreateGui(PartRenderingEngine.java:746) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.access$0(PartRenderingEngine.java:717) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine$2.run(PartRenderingEngine.java:711) at org.eclipse.core.runtime.SafeRunner.run(SafeRunner.java:42) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.createGui(PartRenderingEngine.java:695) at org.eclipse.e4.ui.internal.workbench.PartServiceImpl$1.handleEvent(PartServiceImpl.java:99) at org.eclipse.e4.ui.services.internal.events.UIEventHandler$1.run(UIEventHandler.java:40) at org.eclipse.swt.widgets.Synchronizer.syncExec(Synchronizer.java:186) at org.eclipse.ui.internal.UISynchronizer.syncExec(UISynchronizer.java:145) at org.eclipse.swt.widgets.Display.syncExec(Display.java:4761) at org.eclipse.e4.ui.internal.workbench.swt.E4Application$1.syncExec(E4Application.java:211) at org.eclipse.e4.ui.services.internal.events.UIEventHandler.handleEvent(UIEventHandler.java:36) at org.eclipse.equinox.internal.event.EventHandlerWrapper.handleEvent(EventHandlerWrapper.java:197) at	1	1	https://issues.jboss.org/browse/FUSETOOLS-1422

#	Requirement	Questions	Answered	Successfully Answered	URL
9049	FUSETOOLS-1422	Spring namespace in the route file. But I also don't see the error you have reported, even though I do see the Sleep interrupted exception from the transformation editor again. That seems to be the same issue reported with FUSETOOLS-1507, but maybe it's back again? {code} java.lang.InterruptedException: sleep interrupted at java.lang.Thread.sleep(Native Method) at org.eclipse.core.internal.jobs.JobManager.join(JobManager.java:1002) at org.jboss.tools.fuse.transformation.editor.TransformationEditor.waitJavaBuild(TransformationEditor.java:389) at org.jboss.tools.fuse.transformation.editor.TransformationEditor.init(TransformationEditor.java:376) at org.eclipse.ui.internal.EditorReference.initialize(EditorReference.java:361) at org.eclipse.ui.internal.e4.compatibility.CompatibilityPart.create(CompatibilityPart.java:319) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.eclipse.e4.core.internal.di.MethodRequestor.execute(MethodRequestor.java:56) at org.eclipse.e4.core.internal.di.InjectorImpl.processAnnotated(InjectorImpl.java:898) at org.eclipse.e4.core.internal.di.InjectorImpl.processAnnotated(InjectorImpl.java:879) at org.eclipse.e4.core.internal.di.InjectorImpl.inject(InjectorImpl.java:121) at org.eclipse.e4.core.internal.di.InjectorImpl.internalMake(InjectorImpl.java:345) at org.eclipse.e4.core.internal.di.InjectorImpl.make(InjectorImpl.java:264) at org.eclipse.e4.core.contexts.ContextInjectionFactory.make(ContextInjectionFactory.java:162) at org.eclipse.e4.ui.internal.workbench.ReflectionContributionFactory.createFromBundle(ReflectionContributionFactory.java:104) at org.eclipse.e4.ui.internal.workbench.ReflectionContributionFactory.doCreate(ReflectionContributionFactory.java:73) at org.eclipse.e4.ui.internal.workbench.ReflectionContributionFactory.create(ReflectionContributionFactory.java:55) at org.eclipse.e4.ui.workbench.renderers.swt.ContributedPartRenderer.createWidget(ContributedPartRenderer.java:129) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.createWidget(PartRenderingEngine.java:971) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.safeCreateGui(PartRenderingEngine.java:640) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.safeCreateGui(PartRenderingEngine.java:746) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.access$0(PartRenderingEngine.java:717) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine$2.run(PartRenderingEngine.java:711) at org.eclipse.core.runtime.SafeRunner.run(SafeRunner.java:42) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.createGui(PartRenderingEngine.java:695) at org.eclipse.e4.ui.internal.workbench.PartServiceImpl$1.handleEvent(PartServiceImpl.java:99) at org.eclipse.e4.ui.services.internal.events.UIEventHandler$1.run(UIEventHandler.java:40) at org.eclipse.swt.widgets.Synchronizer.syncExec(Synchronizer.java:186) at org.eclipse.ui.internal.UISynchronizer.syncExec(UISynchronizer.java:145) at org.eclipse.swt.widgets.Display.syncExec(Display.java:4761) at org.eclipse.e4.ui.internal.workbench.swt.E4Application$1.syncExec(E4Application.java:211) at	1	1	https://issues.jboss.org/browse/FUSETOOLS-1422
9050	FUSETOOLS-1411	Length Data Format. This gets long, so my apologies... We're getting all the way to the Bindy marshaller, which is choking. The header row wasn't set on the exchange header here - https://github.com/apache/camel/blob/master/components/camel-bindy/src/main/java/org/apache/camel/dataformat/bindy/fixed/BindyFixedLengthDataFormat.java#L90 So it doesn't process anything for the header or footer, and then when it tries to unbind the model into the text file it appears that it's attempting to process the orderHeader, orderFooter, and items fields as fields - which obviously doesn't work. It's doing that in here: https://github.com/apache/camel/blob/master/components/camel-bindy/src/main/java/org/apache/camel/dataformat/bindy/BindyFixedLengthFactory.java#L327 If I don't wrap the items using the same rough structure we use for the json side (items/item) we have to specify an index in the transformation.xml file and you can get a single item (albeit without the header & footer). That doesn't fix the header/footer issue but at least you get a body row out of the marshaller. Unfortunately after looking at beanio I'm not sure it'll do much better. It has a similar approach to annotations. Any ideas? The transformer is doing just fine - the model is generated beautifully. It's just not being marshalled into the output file appropriately. Are there any ways to somehow intercept the model and populate the exchange header and footer properties? (Sorry, still learning Camel!) I'm attaching the two versions of the project (the one that spits out one line and the one that spits out nothing at all). (Here's the debug output.) {code} [main] TransformationTest INFO ********** [main] TransformationTest INFO Testing: transform(org.example.TransformationTest) [main] TransformationTest INFO ********** [main] CamelNamespaceHandler DEBUG Using org.apache.camel.spring.CamelContextFactoryBean as CamelContextBeanDefinitionParser [main] CamelNamespaceHandler DEBUG Registered default: org.apache.camel.spring.CamelProducerTemplateFactoryBean with id: template on camel context: camel-1 [main] CamelNamespaceHandler DEBUG Registered default: org.apache.camel.spring.CamelConsumerTemplateFactoryBean with id: consumerTemplate on camel context: camel-1 [main] SpringCamelContext DEBUG Set the application context classloader to: sun.misc.Launcher$AppClassLoader@54bedef2 [main] bstractCamelContextFactoryBean DEBUG Setting up routes [main] bstractCamelContextFactoryBean DEBUG Found JAXB created routes: [Route[[From[file:/tmp/inbox?fileName=abc-order.json]] -> [To[ref:json2fixed], To[file:/tmp/outbox?fileName=xyz-order.fixed]]]] [main] SpringCamelContext DEBUG onApplicationEvent: org.springframework.context.event.ContextRefreshedEvent[source=org.springframework.context.support.ClassPathXmlApplicationContext@246ae04d: startup date [Thu Sep 17 16:46:41 MDT 2015]; root of context hierarchy] [main] SharedProducerServicePool DEBUG Starting service pool: org.apache.camel.impl.SharedProducerServicePool@d41f816 [main] SpringCamelContext DEBUG Using ComponentResolver: org.apache.camel.impl.DefaultComponentResolver@8576fa0 to resolve component with name: properties [main] DefaultComponentResolver DEBUG Found component: properties in registry: null [main] DefaultComponentResolver DEBUG Found component: properties via type: org.apache.camel.component.properties.PropertiesComponent via: META-INF/services/org/apache/camel/component/properties [main]	1	1	https://issues.jboss.org/browse/FUSETOOLS-1411
9051	FUSETOOLS-1353	I have put "date", "Date", and "java.util.Date" as the field type in the json schema and it either comes across in the generated classes as a string or a java Object. {code} "date": { "type": "Date" }, {code} Becomes {code} @JsonProperty("date") private Object date; {code} Working through this from point A to point B, how do we expect to get "java.util.Date" into the generated POJO from org.jsonschema2pojo? Does the user have to tweak the generated code?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1353

#	Requirement	Questions	Answered	Successfully Answered	URL
9052	FUSETOOLS-1353	Ok, so now that I have a working example, I'm starting to look at how to implement this. As the data format has to be on the source, not the target, in the dozer config, I've added a new menu item to the drop-down on the source (left) side of the mapping detail pane. If the target is set to "java.lang.Date" as the type, the user will see a new menu - "Set date format" - which opens up the following dialog: !data-format-dialog-draft.png! The drop-down has a list of a few formats, but allows the user to type one (will probably need to add some kind of validation)... * "yyyy-MM-dd HH:mm:ss" * "yyyy-MM-dd HH:mm:ss.SSS" * "yyyy-MM-dd HH:mm:ss.SSS Z" * "MM-dd-yyyy HH:mm:ss" * "MM/dd/yyyy HH:mm" If the user types something or selects something, they can hit OK. The selected format is then set as the "data-format" on the source. I will need to add a "data-format" attribute to the org.jboss.tools.fuse.transformation.model.Model class that handles the property in the configuration. Should we also include this menu on the target side of the equation (which then would set it on the source)?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1353
9053	FUSETOOLS-1352	[~kcbabo] Can you provide a good example of how this should look in the transformation file? Right now I'm working with this, which assumes that the file "simple.txt" is in the src/main/resources directory, which is on the classpath: {code} <mapping> <class-a>org.apache.camel.component.dozer.ExpressionMapper</class-a> <class-b>xyzorderschema.XyzOrderSchema</class-b> <field custom-converter-id="_expressionMapping" custom-converter-param="simple:resource:classpath:simple.txt"> <a>expression origin </field> </mapping> {code} My other questions are: 1) Will this work with the custom converter in the Camel dozer component? 2) The three schemes supported are: classpath, file, and http? Where are these documented? I can't find them listed in http://camel.apache.org/language.html Thanks.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9054	FUSETOOLS-1352	Next question... With three "schemes" available for importing a script file, what do we want to present the user? In the http case, we're just pasting in a web URL. I'm assuming it's something like: {code} from("direct:start") // the script will be loaded on each message, as we disabled cache .to("language:simple:http://someurl.com/myscript.txt?contentCache=false") .to("mock:result"); {code} In the file case, are we browsing anywhere on the file system? What format should the file be in? The lone example I have found is: {code} from("direct:start") // the script will be loaded on each message, as we disabled cache .to("language:simple:file:target/script/myscript.txt?contentCache=false") .to("mock:result"); {code} Where is the file based? What context does that entail? And classpath... {code} from("direct:start") // load the script from the classpath .to("language:simple:classpath:org/apache/camel/component/language/mysimplescript.txt") .to("mock:result"); {code} Is this the classpath of the project? If so, then we should be able to limit it to a similar classpath file browse dialog as I've made a few other places. Sorry for all the questions. Documentation (and even the Camel in Action book) have not been very forthcoming with answers.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9055	FUSETOOLS-1352	I suspected that was going to be the route you wanted to go. The problem is we have a dialog on top of a dialog on top of an editor, so each level of UI we continue to add gets hairier. We can do it, but it won't be very pretty. I will see what I can come up with. Do we also want to take this opportunity to expand the expression text box so it's scrollable? How much text should we allow in an embedded expression before we suggest to the user that they really should import a file? I also wonder at which languages can have loaded expression files- every language? Or just some languages?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9056	FUSETOOLS-1352	Yeah, just like that! Personally, I prefer that interface. What do you and John think?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9057	FUSETOOLS-1352	It'll work for me. Should I just chug forward?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9058	FUSETOOLS-1352	[~janstey] or [~davsclaus], with [~kcbabo] stepping away from the Transformation project I'm hoping you guys might be able to help with what may be a bug in the camel-dozer component? [~lhein] Pointed me in your direction. :) What we're trying to do is have a way of passing a language object with a resource script to be loaded, but it seems to be being ignored when we run it through Camel. It would be one thing if it was just a path problem, but even if I provide a bogus file it's not being processed for some reason. Can you point me at the code where language expressions are handled in camel-dozer or any unit tests where we try to leverage that sort of "Loading scripts from resources" approach (as detailed here: http://camel.apache.org/language.html)? Thanks in advance!	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9059	FUSETOOLS-1352	[~davsclaus] Thanks again for the code pointer. Next question... Do you have any good examples for how the classpath, file, and http URIs are formed in this case? I found examples of classpath in several places in the docs, but the rest are more of a mystery. I have been Googling with no joy to find test cases to use when banging on the UI side (obviously they don't work right now, but will work when we get this issue fixed in the camel component). Do we have a test somewhere that puts this through its paces? Or a quickstart somewhere?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9060	FUSETOOLS-1352	[~davsclaus] [~kcbabo] I have created a PR against Camel master - https://github.com/apache/camel/pull/548 Can you take a look? I added code to handle the language script loading and tests to verify it's working. Would be great if this could also find its way into the 2.15.x branch. I don't have a way to test the http functionality at this point, so if you have suggestions please let me know. But it seems to work fine now. Loads from the classpath or the file scheme just fine.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9061	FUSETOOLS-1352	Hey Tomas, 1. The empty string in the Source combo is there for a reason. The model-driven combo box won't accept a null selection, so we must add an empty element to keep it happy in the Data Binding framework. 2. Source: File is definitely a bug. We should open a File dialog there, you are correct. Workaround right now is to type the path. 3. I'm unable to reproduce your case where the browse button is always enabled. Can you do a video? 4. We can definitely reset the Details after the user changes the Language, effectively clearing any options the user has set up to that point. 5. [~jpav] Any ideas on the scalar to vector/indexing issue when combined with expressions? [~tsedmik], can you create a couple of new JIRAs - #2, #3, and #4 can be combined into one. And #5 should be another. Thanks!	1	1	https://issues.jboss.org/browse/FUSETOOLS-1352
9062	FUSETOOLS-1348	If I understood well, the current solution requires to update the camel-dozer dependency. Is it in Fuse Runtime? In Camel itself? I think that we can't have it for 9.1.0 then as we will still be based on Fuse Runtime 6.3.0 or Camel 2.18.0. Which version can we target on Camel/Fuse? And last question, are we able to handle hiding this feature when user is based on a previous version which doesn't have the updated Dozer version?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1348

#	Requirement	Questions	Answered	Successfully Answered	URL
9063	FUSETOOLS-1347	We already have a mechanism for specifying Custom Operations in the Transformation Editor. In a mapping, click the drop-down menu arrow on the Source detail and select "Add Custom Function." This opens the "Add Custom Operation" dialog [1], which allows you to select from a pre-existing list of Classes or create a new one. Right now we are limited to methods with a single parameter. If we change the name of this to simply be "Add Function" (which still also allows the creation and use of user-created functions), we can easily add a library of String functions to match the requested functions out of the box with a single parameter (i.e. the original string). The class could be included in the camel-dozer component and exposed to the user through the current dialog. My logic here is to keep things simple. A function is a function, whether brought in from the user's project or from the upstream camel-dozer component, so having a complete list in one place would be handy since the selection dialog has a search area already. We could also expose them via a cascading menu (Add Function->[functionName]) in the same drop-down list as "Add Custom Function" shows up today if we want to keep them separate. Either way, the simple String functions (single parameter) we could expose initially would include: * toLowerCase (String original) * toUpperCase (String original) * trim (String original) Unfortunately, as we are only afforded a single parameter (the original string), we are limited in what we can accomplish here without some significant changes to org.apache.camel.component.dozer.CustomMapper. The "mapCustom" method would need to handle an arbitrary number of input parameters and the org.jboss.tools.fuse.transformation.editor.internal.AddCustomFunctionDialog would need to offer a way for users to provide inputs. Once we have that capability, we could offer all sorts of mapping functions out of the box: * concat (String original, String strToAdd) * replace (String original, char oldchar, char newchar) * replaceAll (String original, String regex, String replacement) * replaceFirst (String original, String regex, String replacement) * split (String original, String regex, int index) * substring (String original, int beginIndex) * substring (String original, int beginindex, int endIndex) And obviously the expanded parameter option would enhance the custom user-created operations as well. [~jpav] and [~kcbabo] - Do we have time in this sprint to push further Camel changes upstream to the camel-dozer component? These would be rather significant I suspect and we're down to the wire with time. [1] Note: We should probably settle on the text for this menu item and make the dialog text match.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1347
9064	FUSETOOLS-1347	[~jpav] I've been looking at your initial implementation of the functionality in the "Add Function" menu and associated dialog. Here's my feedback in no particular order and not trying to be overly critical, but seems to go contrary to other dialogs in Eclipse: * It opens without anything selected, so the Arguments group is empty and leaves you guessing as to what "applicable arguments" you may be referring to. Perhaps put a tooltip or some label in the group box if nothing is selected to hint (in addition to the text at the top of the dialog) that they need to select a function. * The list of functions is in a listbox with no vertical scrollbar. Need to add that * There should be a thin black line at the bottom of the "title" area of the dialog - not sure why that's missing. Looks odd. If you look at other title area dialogs, it's usually there. * I would change the dark black splitter and make it a darker gray or remove it entirely. We don't need it in this case since the user can resize the entire dialog to see whatever they wish to see. * Please use an actual table for the arguments or change the black column headers to be more like a traditional table in Eclipse. We don't need to custom build every control. TableViewer has built in functionality we can use for column editors and the like. * The arguments are a bit strange to me. In all cases there's just the one input parameter which is always the string to be converted. Can the user change the value? is the name really 'string'? Unless there are arguments that the user has to fill in (which there aren't any yet) I'd get rid of Arguments part of the dialog completely for this release. I think I get what you're after here in the future (when the user can actually set input for some of these arguments) but for now it's more confusing than anything else. Let's get some feedback from [~kcbabo] on this as well.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1347
9065	FUSETOOLS-1346	We now have: * csv-to-xml * json-to-json * json-to-xml * xml-to-csv * xml-to-json * xml-to-xml Still need: * csv-to-fixed (other to other) * java-to-java * java-to-json (may not be needed) * java-to-xml (may not be needed) * xml-to-java (may not be needed) * json-to-java (may not be needed) I'm erring on the side of completeness here. Other-to-Other and the Other-to-? and ?-to-Other examples essentially handle the Java-to-Java cases. We will for sure add the Other-to-Other (csv-to-fixed probably, since we've had some questions about fixed length messages) and a Java-to-Java example before I'm done. Also created FUSETOOLS-1429 so we get these examples moved back to the 7.3.x stream as well.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1346
9066	FUSETOOLS-1346	[~bfitzpat] The status is "coding in progress" but last comment is from last September. I suppose that it is not the case. But should it be resolved or set back in To do?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1346
9067	FUSETOOLS-1282	Currently we generate the components drawer entries in the palette from what we get from Camel meta data. Unfortunately this meta data doesn't provide any label for a component. So it would need a change on Camel side to get this fixed.	0	0	https://issues.jboss.org/browse/FUSETOOLS-1282
9068	FUSETOOLS-1279	Created an initial version of the extension point in the editor plugin. To use it you have to create a class which implements ICustomDblClickHandler interface. see https://github.com/fusesource/fuseide/blob/master/editor/plugins/org.fusesource.ide.camel.editor/src/org/fusesource/ide/camel/editor/provider/ext/ICustomDblClick Handler.java The method *canHandle* decides if you want to handle a dbl click event on a specific figure or not. In the case of the data transformation you could check if its a ref and then lookup the ref and see if its a dozer endpoint. The method *handleDoubleClick* is called if *canHandle* returns true. This method should contain the actual action you want to do like opening an editor etc. I only tried if the handling itself works. You have to check if you can access the model to lookup the referenced EP in a REF endpoint. If thats not possible let me know and we extend the parameters of the methods a bit.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1279
9069	FUSETOOLS-1272	[~lhein] the SAP Tooling Beta 3 has been tagged? Where has it been pushed? Do we need to modify something in JBDSIS? [~tsedmik] Can you tell which version of SAP plugins are installed?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1272
9070	FUSETOOLS-1263	+1 for looking at doing the git move. -0 on changing plugin/features too much. We'll need to be careful about the renaming/refactoring since it changes API's and upgrade paths. I think we can for most things add p2 metadata to handle bundle id renames, but how about preferences or other properties/settings ? are their keys using the bundle id's? how many settings are there in Fuse tooling ? how would it affects users if we do the rename?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1263
9071	FUSETOOLS-1262	Whats the state here? I can't see anything in the new repo.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1262
9072	FUSETOOLS-1259	How about creating a tooling tutorial supplement repo in https://github.com/fusebyexample?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1259
9073	FUSETOOLS-1259	Ok. What would I need to do to get them there? And how easy will it be for Fuse Tooling users to get them out and installed in the correct folder in their CBRroute project?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1259

#	Requirement	Questions	Answered	Successfully Answered	URL
9074	FUSETOOLS-1259	"Why not zip them up and put them up on the Customer Portal with the book?" Yes, that's a possibility and might be a good solution in this case. You can put arbitrary files under the {{files/}} directory in Publican (though I would have to check that our {{sibin}} preprocessor script honours that). Having a place in {{github.com/FuseByExample}} to put doc demos would also be very useful. I've been thinking about this recently and it is something we could probably use for other doc demos (but I haven't figured out the details yet). Maybe we can have a chat about this sometime, Jane.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1259
9075	FUSETOOLS-1258	Modified the main project POM to add a targetplatform profile and created Jenkins build and publish projects to establish an independent SAP update site: https://github.com/punkhorn/jboss-fuse-sap-tool-suite/commit/16ec15aac8b0ccd405a60d6bd0d56eb57f052a84 mvn -Ptargetplatform clean install # mirror the JBTIS base-ea and is-stack-and-data-layer.target into a combined repo https://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/jbosstools-fuse-sap-tooling https://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/jbosstools-fuse-sap-tooling-publish The resulting update site is: http://download.jboss.org/jbosstools/updates/integration/luna/integration-stack/jboss-fuse-sap-tool-suite/7.3.0.Beta1/ Modified the ide-config.properties to create a property that may be used to identify the IS extras URL: https://github.com/jbosstools/jbosstools-download.jboss.org/commit/71b6294966f31ccce728634855c3baa7d40728c8 https://github.com/jbosstools/jbosstools-download.jboss.org/commit/26c879653786989884536cbce2fa3dc73495754d Also established IS composites to hold the "extras" URLs (currently only component is SAP tooling): http://download.jboss.org/jbosstools/updates/development/luna/integration-stack/extras/compositeArtifacts.xml http://download.jboss.org/jbosstools/updates/development/luna/integration-stack/extras/compositeContent.xml http://download.jboss.org/jbosstools/updates/stable/luna/integration-stack/extras/compositeArtifacts.xml http://download.jboss.org/jbosstools/updates/stable/luna/integration-stack/extras/compositeContent.xml What needs to be done? The wizard Java code which looks for the update site should pull the SAP tooling update site URL from ide-config.properties like this: PropertiesHelper.getPropertiesProvider().getValue("jboss.discovery.site.integration-stack-extras.url"); See this code for an example: https://github.com/jbosstools/jbosstools-central/blob/master/discovery/plugins/org.jboss.tools.discovery.core/src/org/jboss/tools/discovery/core/internal/DiscoveryActivator.java#L131	1	1	https://issues.jboss.org/browse/FUSETOOLS-1258
9076	FUSETOOLS-1258	looks like a good start. could we please avoid publishing non-stable ocntent into stable urls ? i.e. http://download.jboss.org/jbosstools/updates/stable/luna/integration-stack/extras/compositeArtifacts.xml shuold IMO just be empty. Also, shouldn't the actual published site be an aggregate instead of a composite pointing to dynamic changing sap site content ?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1258
9077	FUSETOOLS-1258	Just to be clear, the plugin should be made so it can survive without the property set. i.e. have a default value in case a mistake is made in the properties upload or whatnot. also should we make it a comma separated list to allow for freedom to extend it ? (i'm fine with staying just at one but then we just need to always be okey that the sap tooling is *only* one feature)	1	1	https://issues.jboss.org/browse/FUSETOOLS-1258
9078	FUSETOOLS-1258	I guess I'm confused now. We currently have this defined in ide-config.properties: jboss.discovery.site.integration-stack-extras.url|jbosstools|4.2=http://download.jboss.org/jbosstools/updates/development/luna/integration-stack/extras/ This is a generic property value for identifying anything the IS defines as 'extras' - currently that's a composite site with just one entry: <child location="http://download.jboss.org/jbosstools/updates/development/luna/integration-stack/jboss-fuse-sap-tool-suite/7.3.0.Beta1/all/repo/"/> There could be more. You're only interested in jboss-fuse-sap-tool-suite which is a real repo. Could you XML parse (in Java) the compositeArtifacts.xml and find the SAP tooling suite you're looking for? What's the alternative - creating a SAP-specific property value in ide-config.properties? This should be the SAP feature name? Please more details.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1258
9079	FUSETOOLS-1258	Looking at the actual live file: I see http://download.jboss.org/jbosstools/configuration/ide-config.properties {{{ jboss.discovery.site.integration-stack-sap.url|jbosstools|4.2=http://download.jboss.org/jbosstools/updates/development/luna/integration-stack/jboss-fuse-sap-tool-suite/7.3.0.Beta1/all/repo/ jboss.discovery.site.integration-stack-sap.feature|4.2=org.fusesource.ide.sap.feature.feature.group }}} The thing I notice is that the latter does not have "jbosstools" specified despite it has the version number - so I think that should be fixed the |4.2 removed for the sap.feature line if you want it to be globally available. btw. is it really intended/expected the feature is named "feature.feature.group" ?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1258
9080	FUSETOOLS-1258	Thanks - looks like it doing everything in a nice straight line but one thing that I'm missing is how *updates* are going to be handled ? This code removes the generate updatesite meaning when next release comes out of IS with SAP tooling updates these will not be picked up. How are users going to get updates ? By running the install *again* ?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1258
9081	FUSETOOLS-1258	if you kept the updatesite created in a less temporary location and kept it in place then updates the SAP tools plugin will just work. I assume we are not locked into too specific SAP client lib versions ? With respect to the sap libraries if we actually put proper versions onto the generated plugins and you use proper version ranges then if you are just fixing bugs in the tools then updates will just become available and only if we actually have some strict incompatbility with sap client libs will you bump the version range and users will need to run the SAP wizard install again to get proper versions in place. makes sense ? could it work ?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1258
9082	FUSETOOLS-1225	[~ceposta] on the page that you pointed, all configuratiosn are based on beans (so should be covered by https://issues.jboss.org/browse/FUSETOOLS-1585). is it really the case? or do you mean something else? if yes, can you open an enhancement request linked to this one containing a sample of the xml code that you would like to manage please?	1	0	https://issues.jboss.org/browse/FUSETOOLS-1225
9083	FUSETOOLS-1200	[~lhein] Hi, I am new around these parts so not sure if JIRA is the place for questions. I have a few questions on this ticket if this enhancement is still in scope. * Does the screenshot !screenshot-1.png|thumbnail! depict what you have in mind? * I implemented this only on the 'Project Explorer' since 'Navigator' & 'Package explorer' doesn't follow [CNF|http://help.eclipse.org/mars/index.jsp?topic=%2Forg.eclipse.platform.doc.isv%2Fguide%2Fcnf.htm] so it doesn't seem straightforward to implement custom trees for those views. * Only <route/> child nodes are shown currently in expanded state (ref: issue description "show the contained routes as child elements"). When you double click a route it will take you to that route in the diagram editor (i.e., it behaves similar to the outline view) * I have a fork [here|https://github.com/2point718/fuseide] if you'd prefer to have this discussion on a pull request	1	1	https://issues.jboss.org/browse/FUSETOOLS-1200

#	Requirement	Questions	Answered	Successfully Answered	URL
9084	FUSETOOLS-1181	Jane, good to hear that. Can you open a PR for that? (I can't find one in our git repo)	0	0	https://issues.jboss.org/browse/FUSETOOLS-1181
9085	FUSETOOLS-1181	Jane, this issue is already resolved, isn't it? _(according to https://github.com/fusesource/fuseide/commit/dd8f66d230c961a763e95453bbe85244c317e6a4)_ Can you resolve the issue here?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1181
9086	FUSETOOLS-1177	I think Max's idea is great. Customer can create his project based on many quickstarts like the ones in $Fuse61_HOME/quickstart and moreover, camel-archetype-activemq is the old archetype, creating a project from it can not run in the fuse. about camel-archetype-spring-dm, is spring still the suggestion ? I suppose we should recommend blueprint instead.	0	0	https://issues.jboss.org/browse/FUSETOOLS-1177
9087	FUSETOOLS-1168	This seems to be a bug in Camel's dumpAsXML JMX method. Claus, is that the same we already reported to Camel?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1168
9088	FUSETOOLS-1168	Huh? I am using the camel naming factory to create those ids. Don't you check on Camel side for uniqueness when assigning IDs? Can't believe that.	0	0	https://issues.jboss.org/browse/FUSETOOLS-1168
9089	FUSETOOLS-1151	Hey James, could you add more details what you want to do? From what the headline says you want to check out some files from a fabric8 git repo. But what I am missing is details like what files, how can I reach that git repo, where do the files belong to? (guess a server?) If its a server thing I would recommend we add a configuration node below the server in the servers view and put the files under that node in some to be defined ordering / grouping. Also of interest is what those files contain and if we need custom editors for those files. Thanks in advance Lars	1	0	https://issues.jboss.org/browse/FUSETOOLS-1151
9090	FUSETOOLS-1120	Hi Lars, does 'publishing' means deploing to a server? Because as far as I know we do not support deployment of JAR projects. Am I correct?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1120
9091	FUSETOOLS-1106	[~pleacu]: I am not seeing that issue when running Fuse Tooling only. What kind of blocking dialog is showing up? As far as I remember we do not define special keys at all.	0	0	https://issues.jboss.org/browse/FUSETOOLS-1106
9092	FUSETOOLS-1104	Will it support both Java 7 and Java 8 and what is the timeframe for JBDS 8?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1104
9093	FUSETOOLS-1104	wait wait - build and test with JDK 8 is two very differnt things. We are *not* building our eclipse plugins with JDK 8. For JBDS 7 it is Java 6, for JBDS it is Java 7. We are *Testing* running our eclipse plugins with Java 6+7 in JBDS 7, and JAva 7+8 for JBDS 8. *Support* for *using* Java 8 is possible on kepler but only enabled on jbosstools as commnuity testing. JBDS 8 is where the first full support is for Java 8. If you need to target *users* apps on Java 8 - you currently need to move to JBDS 8/JBoss Tools Luna. Which one is this jira about ?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1104
9094	FUSETOOLS-1104	what are "some servers" ? some specific versions of Fuse or ?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1104
9095	FUSETOOLS-1103	Basic functionality like Start / Stop / Restart Server will be needed. For the publishing we need additional information: 1. What can be published / deployed to Fabric8? 2. How does the deployment work? 3. Do we deploy to local Fabric8 instances only or is remote publishing possible too? 4. If remote publish, then how can it be done? 5. How to remove published stuff in a clean way (local, remote)? ...	0	0	https://issues.jboss.org/browse/FUSETOOLS-1103
9096	FUSETOOLS-1103	[~jastrachan], could you answer above questions?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1103
9097	FUSETOOLS-1103	1. jars, wars, ears, bundles 2. see http://fabric8.io/#/site/book/doc/index.md?chapter=mavenPlugin_md. basically the deployment artefact is uploaded to fabric8's maven proxy: http://fabric8.io/#/site/book/doc/index.md?chapter=mavenProxy_md then a profile is lazily created/updated 3. see above docs; it works with any fabric; you can specify the URL to it. 4. there's only one way wherever the fabric 5. there's no way to do that; just like there's no clean way to do that with any maven repository/proxy. Why do you want to?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1103
9098	FUSETOOLS-1102	Can that plugin be also used for deployments to Apache Karaf and Apache ServiceMix?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1102
9099	FUSETOOLS-1079	Can we close that one?	0	0	https://issues.jboss.org/browse/FUSETOOLS-1079
9100	FUSETOOLS-1066	So now the gui has changed such that Fuse Fabric perspective --> Fabric8 perspective, and Fabric Explorer --> Fabric8 Explorer? (The last gui I saw -- 7.2.0.v20140312-2041-H37-CR3 -- Fabric Navigator had changed back to Fabric Explorer.) I will need new screen captures as well. When can I get a new build?	1	1	https://issues.jboss.org/browse/FUSETOOLS-1066
9101	FUSETOOLS-1066	Jane, can you try with a clean workspace? If that doesn't help you could also try to reset the perspective.	1	1	https://issues.jboss.org/browse/FUSETOOLS-1066
9102	FUSETOOLS-1031	Yes, in my tests I had no selection active. But besides that you can always create a route from the main menu and also inside the editor...why should I use the outline view for that. We don't even open that view in our perspectives by default.	1	0	https://issues.jboss.org/browse/FUSETOOLS-1031
9103	FUSETOOLS-955	It would appear that it's possible to co-install SpringIDE 3.3 and FuseIDE 7.2 into JBDS 7.0.0.GA: [^springide33_fuseide72_coinstalled_in_JBDS700GA.png] I'm just not sure if having both 4.0.M1 and 3.1.4.RELEASE plugins running at the same time will cause problems. Has anyone tried to USE both SpringIDE 3.3 and FuseIDE 7.2 at the same time from within the same JBDS 7 install? Give me some tests to run and I can try it out...	1	0	https://issues.jboss.org/browse/FUSETOOLS-955
9104	FUSETOOLS-955	Nick, what is the state here? Can we close that issue?	1	0	https://issues.jboss.org/browse/FUSETOOLS-955
9105	FUSETOOLS-955	State was "Give me some tests to run and I can try it out..." but no one did, so no testing was done. Have you tried to install FuseIDE 7.2.CR1 and Spring 3.3 in the same JBDS?	0	0	https://issues.jboss.org/browse/FUSETOOLS-955
9106	FUSETOOLS-948	Can we defer until after 6.1?	0	0	https://issues.jboss.org/browse/FUSETOOLS-948
9107	FUSETOOLS-913	Can this be deferred til after 6.1?	0	0	https://issues.jboss.org/browse/FUSETOOLS-913
9108	FUSETOOLS-907	Jane, could you also document, that using the context menu "Maven / Update Project" is only available on the project folder and not on the pom.xml file? I know its unrelated to this issue, but would be nice if you also put that in place. Thanks, Lars	1	1	https://issues.jboss.org/browse/FUSETOOLS-907
9109	FUSETOOLS-907	Jane, did u document it somehow? If yes, please resolve the issue.	0	0	https://issues.jboss.org/browse/FUSETOOLS-907
9110	FUSETOOLS-907	No, I didn't put it in the user guide because you said it needs to be fixed. So I was treating it as a known issue, which should go into the release notes. Is this a known issue that will be fixed? If not, I will need to add it as part of the procedure to view other items in the Fuse JMX Navigator. This isn't the only place where this occurs. I've also experienced it when selecting a message trace in Messages View and had to unpin Properties view to see the Message Details page correlated with the message trace I selected.	0	0	https://issues.jboss.org/browse/FUSETOOLS-907
9111	FUSETOOLS-868	Reviewed that patch from Paul and it gives an error dialog if the zookeeper password is wrong. Though the message inside that dialog is not really useful. It says that there might be no fabric and we should create or join one. But the cause was a wrong zookeeper password. Another strange thing for me: You can enter whatever credentials as Username / password as long as the zookeeper password is fine. Is that correct?	1	1	https://issues.jboss.org/browse/FUSETOOLS-868
9112	FUSETOOLS-868	fixed. Jason, can you verify the fix?	1	1	https://issues.jboss.org/browse/FUSETOOLS-868
9113	FUSETOOLS-832	Is that still an issue with the latest version?	1	1	https://issues.jboss.org/browse/FUSETOOLS-832
9114	FUSETOOLS-707	Assigning to you Luis - can we have a new Route icon with a trace modifier; so folks can visually see the routes being traced?	1	1	https://issues.jboss.org/browse/FUSETOOLS-707

#	Requirement	Questions	Answered	Successfully Answered	URL
9115	FUSETOOLS-658	the parameter "-Dcamel.blueprint=true" is now set if the context is a blueprint one. Not yet tested because we still use an old version of the camel:run plugin. Claus, could u test it by time?	1	1	https://issues.jboss.org/browse/FUSETOOLS-658
9116	FUSETOOLS-623	can we close that one, Charles?	1	1	https://issues.jboss.org/browse/FUSETOOLS-623
9117	FUSETOOLS-613	Oh and I just noticed in the server launch configuration, it has options for Apache Derby in the VM arguments (see screenshot 3). Is that really needed?	1	0	https://issues.jboss.org/browse/FUSETOOLS-613
9118	FUSETOOLS-565	thought that this is no longer there? could u dbl check and close this issue if it away?	0	0	https://issues.jboss.org/browse/FUSETOOLS-565
9119	FUSETOOLS-565	Tried the same here. When hitting the stop button the whole jmx process vanished from the JMX tree. I tried to get the context menu of the diagram view but it had no entry for dynamic instance.	1	1	https://issues.jboss.org/browse/FUSETOOLS-565
9120	FUSETOOLS-537	Can't reproduce that problem. When using add (tried both adding a standalone unconnected node and a connected node by using add on another node) context menu the new node gets aligned as wanted by the layout. If you add it to an existing route it is added to the right of the node otherwise it is aligned to the left as no connection is available. Also the newly created figure is selected automatically. Guess you had some buggy environment? Please recheck and reopen this issue if you still run into this problem.	0	0	https://issues.jboss.org/browse/FUSETOOLS-537
9121	FUSETOOLS-481	Don't the doc team do release notes now?	1	1	https://issues.jboss.org/browse/FUSETOOLS-481
9122	FUSETOOLS-477	Which perspective are you talking about?	1	1	https://issues.jboss.org/browse/FUSETOOLS-477
9123	FUSETOOLS-475	I haven't been able to reproduce this. What I did: 1. open the diagram 2. selected a figure --> the properties page was filled with properties of that figure 3. selected the Camel Context File from Project Explorer --> the Properties view was filled with the properties of that file 4. selected Context Menu "Run as / Local Camel Context Without Tests" The console was filled with texts about startup of the local camel context. The properties view was still showing the properties of the selected Camel Context file. After selecting the EIP figure again, the properties view was filled with the values of that figure again. Looks all fine to me. We have to remember that the properties view always tries to show the properties of the selected element. In the above case the last selected element was the Camel Context File and not the EIP figure inside the diagram.	1	0	https://issues.jboss.org/browse/FUSETOOLS-475
9124	FUSETOOLS-454	Was that new endpoint connected to something else or was it standalone? Just to clarify things. I think it may got dismissed because it was not valid at that moment.	1	1	https://issues.jboss.org/browse/FUSETOOLS-454
9125	FUSETOOLS-451	Claus, I can't find that context entry any longer. Can you check if its still there for you?	1	1	https://issues.jboss.org/browse/FUSETOOLS-451
9126	FUSETOOLS-420	Thanks Eric. Sorry I missed it! Instead of "Registering Fuse IDE" maybe we can specifically call it "Licensing Fuse ID"? Also, can we specifically say that if you are interested in a trial version to contact us? It's really the non-subscribers trying it out that seem to run into difficulties. Thanks, Susan	1	1	https://issues.jboss.org/browse/FUSETOOLS-420
9127	FUSETOOLS-383	wonder if this helps? http://stackoverflow.com/questions/6286701/an-eclipse-console-view-that-respects-ansi-color-codes	0	0	https://issues.jboss.org/browse/FUSETOOLS-383
9128	FUSETOOLS-325	With ECLIPSE-707, we have distinguished the icon for the camel route in the tree to show that tracing is on. Do we think we want to change the text still, to say "Fuse Message Tracing"?	1	0	https://issues.jboss.org/browse/FUSETOOLS-325
9129	FUSETOOLS-322	I think this is now resolved as we re-fresh connections every second or so - could you double check sometime if you can reproduce?	1	1	https://issues.jboss.org/browse/FUSETOOLS-322
9130	FUSETOOLS-309	Claus, if we would add such a bypass for the unit test, would you ever fix the broken tests in future? I am not sure if we should add such features as those unit tests should ensure that the application does what it was designed for. Bypassing unit tests is a risk and if used on production systems it is really dangerous.	0	0	https://issues.jboss.org/browse/FUSETOOLS-309
9131	FUSETOOLS-303	> Why does this occur when the following was entered into the eclipse.ini file (as per the Fuse > IDE for Camel installation guide, version 1.1 - Chapter 1 Prerequisites section): > > -vm > C:/Program Files/Java/jdk1.6.0_26/bin/javaw.exe > > Can you please also update this Fuse IDE for Camel Installation Guide document to ensure a clean Eclipse configuration on MS Windows? I've tried this setting on my WinXP box and it works fine for me.	0	0	https://issues.jboss.org/browse/FUSETOOLS-303
9132	FUSETOOLS-288	{quote} Calling a feature "Deploy To..." implies that artifacts are going to be packaged for and deployed into a run time environment of some sort. It also implies that some feedback about the success or failure of the deployment will be given. {quote} It does and it is. Though the output goes to the Console. The build may take a little while so it avoids having a modal dialog appear or anything like that. {quote} The Fuse IDE "Deploy To..." menu option does not do anything of the sort. It simply builds the project and copies the result to a local folder. {quote} How is that any different to your first paragraph? I'm confused. {quote} The only feedback presented to the developer is if the build succeeds or fails. {quote} The console shows the file being copied and that the build succeeded. What kind of output did you have in mind? {quote} The easiest fix for this would be to call the option "Export to...". {quote} Am not 100% sure Deploy is the right word, but am less sure Export is. Developers talk about deploying to containers. The directory is called "deploy" in Karaf for example; so folks are used to the term "Deploy" meaning put build artifacts into the deploy directory to deploy your application. {quote} Ideally, this feature would be coupled with a mechanism for configuring a connection to a run time environment, ensure that the artifacts are properly packaged (at least through a warning if the package property in the POM is wrong), and provide some feedback about the status of the generated artifact once it is deployed. {quote} So we could use a modal dialog with OK / failed or some kind of progress meter or something. Whether a runtime connection is used or not depends; currently the Deploy To mechanism deploys to a hot deploy directory; the "Servers -> Add Karaf / ServiceMix Server" option is about adding a 'runtime connection' to a server that is run inside Eclipse and so can grok what kinds of builds are suitable and so forth.	0	0	https://issues.jboss.org/browse/FUSETOOLS-288
9133	FUSETOOLS-285	Eric do you have suggestions for descriptions / better descriptions? If its the description in the maven pom.xml of the archetype, then we can adjust those. Here is for example the camel-archetype-activemq description: <description>Creates a new Camel project that configures and interacts with ActiveMQ.</description> If you supply better descriptions, then I can adjust them in the Fuse MR codebase Since we use the same descriptions as at Apache you can see them here: https://svn.apache.org/repos/asf/camel/trunk/tooling/archetypes/ eg its the pom.xml in the first directory. such as https://svn.apache.org/repos/asf/camel/trunk/tooling/archetypes/camel-archetype-activemq/pom.xml	0	0	https://issues.jboss.org/browse/FUSETOOLS-285
9134	FUSETOOLS-285	I can look through the Camel ones. Do the other projects have similar archetype structures in their repos and have the descriptions located in the archetype's POM?	0	0	https://issues.jboss.org/browse/FUSETOOLS-285
9135	FUSETOOLS-285	I can add descriptions to the ServiceMix ones, but I will need someone to commit them. Should I make a JIRA and a patch? I already fixed up the CXF one at Apache.	0	0	https://issues.jboss.org/browse/FUSETOOLS-285

#	Requirement	Questions	Answered	Successfully Answered	URL
9136	FUSETOOLS-281	All projects we create at Apache, Fuse Forge and our Fuse distributions are Maven based. As are all of our training projects and all projects our customers have created using our stuff to my knowledge. So I fail to see how the lack of the redundant "Maven" word in the menu is a bug. Folks who use Fuse kinda assume Maven is in there. Since there is no non-maven projects, I fail to see how mentioning Maven is that big a deal. BTW we are not creating a new "Fuse IDE Type" nor have any immediate plans to do so; so far we're using maven as our build system and don't provide any facade that hides setting up a class loader, compile time dependencies or the huge complexities involved of making WARs for web apps, using Spring or using OSGi with or without spring-dm or blueprint - never mind choosing some permutation of camel / cxf / activemq too. We've simply created a helper wizard for creating new Fuse based projects (via Fuse IDE). We're using "Fuse IDE Project" specifically to differentiate from other stuff outside of Fuse IDE; e.g. if you're using Fuse IDE with Spring STS one day. So this isn't really a bug - its more arguing for better text on a menu item right? I'd even be tempted to just use "Fuse Project" instead of "Fuse IDE Project" as all projects are going to be Maven for the foreseeable future. Then in the tooltip / description of the wizard talk about Maven. If we ever provide any other kind of project one day, we'd make using a maven or non-maven one an option/page in the same wizard anyway rather than having one menu item for "Fuse Maven Project" and "Fuse not-Maven Project" etc.	1	1	https://issues.jboss.org/browse/FUSETOOLS-281
9137	FUSETOOLS-280	I don't follow - could you be more specific which bit of IDE you're talking about?	0	0	https://issues.jboss.org/browse/FUSETOOLS-280
9138	FUSETOOLS-258	James I can see the icon is from the auto generated model. /* (non-Javadoc) * @see com.fusesource.rider.model.AbstractNode#getIconName() */ @Override public String getIconName() { return "endpoint.png"; } Can you change so idempotent consumer uses the filter icon? return "filter.png"	0	0	https://issues.jboss.org/browse/FUSETOOLS-258
9139	FUSETOOLS-220	What about Debug>Debug Configurations and Debug>Deploy To Container? Do they also need documenting, and are they still works in progress? The Fix Version of the dependency bug 213 is set to 2.2.	0	0	https://issues.jboss.org/browse/FUSETOOLS-220
9140	FUSETOOLS-206	Lets just not support it when its removed in Camel. We will remove it in Camel 3.0, or maybe earlier.	0	0	https://issues.jboss.org/browse/FUSETOOLS-206
9141	FUSETOOLS-204	Eric do you mind fixing the icon?	0	0	https://issues.jboss.org/browse/FUSETOOLS-204
9142	FUSETOOLS-202	I wonder if the real fix for this is a Camel one? i.e. remove the "ref" XML attribute and add a new XML element <dataFormatRef ref="foo"/> which would be a child of <marshal>? I guess we can try invent some hack in Fuse IDE though...	1	1	https://issues.jboss.org/browse/FUSETOOLS-202
9143	FUSETOOLS-194	Eric do you do icons? The icon in com.fusesource.rider.editor/icons/dynamicRouter24.png has its top black border chopped off when in the rider. Maybe we can redraw the icon so the top line is there?	0	0	https://issues.jboss.org/browse/FUSETOOLS-194
9144	FUSETOOLS-169	I've just used the latest build to create this wire tap... {code} <wireTap uri="file:something"> <body> <language language="xpath">/foo/bar</language> </body> </wireTap> {code} is that enough? Is there something missing?	1	1	https://issues.jboss.org/browse/FUSETOOLS-169
9145	FUSETOOLS-163	Jane, could you document that new flag?	1	1	https://issues.jboss.org/browse/FUSETOOLS-163
9146	FUSETOOLS-64	Any chance you could take a stab at allowing reordering of the tree (e.g. CBR when clauses) using drag and drop in the outline view?	1	0	https://issues.jboss.org/browse/FUSETOOLS-64
9147	FUSETOOLS-52	all of the patterns now have corresponding reference pages that describe the properties settable in the editor. Can this be closed?	0	0	https://issues.jboss.org/browse/FUSETOOLS-52
9148	FUSETOOLS-46	The documentation link on the welcome page opens the Fuse IDE help locally as suggested. Can this be closed?	0	0	https://issues.jboss.org/browse/FUSETOOLS-46
9149	FUSETOOLS-45	I've added Mac specific instructions to the install guide. Can this be confirmed and closed?	1	1	https://issues.jboss.org/browse/FUSETOOLS-45
9150	FUSETOOLS-44	This issue was fixed when the need to add the FuseSource archetypes was removed from the IDE. The offending link was removed from the install guide. Can this be confirmed and closed?	0	0	https://issues.jboss.org/browse/FUSETOOLS-44
9151	FUSETOOLS-43	The help plug-in shows content. Can this be confirmed and closed?	0	0	https://issues.jboss.org/browse/FUSETOOLS-43
9152	FUSETOOLS-30	I'm assuming you get this error when entering the archetype catalog as per the instructions at http://fusesource.com/docs/developer/camel/beta/install_guide/RiderInstallOther.html In m2eclipse you need to make sure "http://" precedes the URL, otherwise you will get the error. So the instructions are correct as it lists http://repo.fusesource.com/nexus/content/groups/public/archetype-catalog.xml Did this occur even with the http:// included?	0	0	https://issues.jboss.org/browse/FUSETOOLS-30
9153	FUSETOOLS-30	This issue seems to have been solved, but I'm getting the same errors...Eric mentioned it's not a problem because m2e is installed with fuse catalog installed... i assume that means when you install FUSE IDE, the catalog will be installed automatically... and it is... but it's empty. trying to add manually gives the errors originally reported by James. Using Fuse IDE 2.0.26 on Eclipse Indigo Service Release 2 (build 20120216-1857).... which I believe is Eclipse 3.7.x... is that version not supported? I'll try 3.6 and see if that helps...	0	0	https://issues.jboss.org/browse/FUSETOOLS-30
9154	FUSETOOLS-30	Same thing on eclipse 3.6. The arch list shows up empty. Should I open a new JIRA for my issue?	0	0	https://issues.jboss.org/browse/FUSETOOLS-30
9155	FUSETOOLS-28	Can we get html versions of the FDFC docs?	1	1	https://issues.jboss.org/browse/FUSETOOLS-28
9156	FUSETOOLS-18	The install guide does use the suggested pattern for adding an update site. Can this be confirmed and closed?	0	0	https://issues.jboss.org/browse/FUSETOOLS-18
9157	FUSETOOLS-7	A link to the help is shown on the welcome page now. Can this be confirmed and closed?	0	0	https://issues.jboss.org/browse/FUSETOOLS-7
9158	JBCL-185	We have a customer facing this issue currently in a migration effort. I whipped up a potential fix for them to test that would allow them to cache locally to the BCL to return from BCL.loadClass() calls before trying to go through synchronization and the domain. They confirmed this change resolved the performance degradation experienced moving from 4.3 to 5. Attached is the diff file and source file encompassing my changes for this on 2.0.9.GAs BCL. Would anyone more well versed mind reviewing these changes? Can we implement these in the JBoss Classloader? Or would there be a better way to get these loadClass calls to return from cache before synchronizing and bottlenecking?	1	1	https://issues.jboss.org/browse/JBCL-185
9159	JBCL-185	Did you try this on v2.0.11? * http://anonsvn.jboss.org/repos/jbossas/projects/jboss-cl/tags/2.0.11.GA/	1	1	https://issues.jboss.org/browse/JBCL-185
9160	JBCL-185	Could you clarify what you were wondering if I've tried? The performance testing exposing the issue on 2.0.11.GA? Or implementing this change on 2.0.11.GA? Currently, neither. Were you wondering if 2.0.11.GA would already have addressed this? If so, I don't believe so. Comparing 2.0.11.GA and 2.0.9.GA so far, I'm not really seeing any notable differences in BaseClassLoader or BaseClassLoaderDomain that would appear to provide any different caching/locking behavior seen here on 2.0.9.	1	1	https://issues.jboss.org/browse/JBCL-185
9161	JBCL-131	Can we please keep discussions in the forums. To avoid future confusion, jboss-cl has the exactly the same semantics. Optional dependencies can't trigger a new "wire" getting constructed once the deployment has passed the classloader created stage. Dynamic imports get resolved at runtime so they can pick up the results from lazy resolution (or plain old non-lazy resolution).	1	1	https://issues.jboss.org/browse/JBCL-131

#	Requirement	Questions	Answered	Successfully Answered	URL
9162	JBCL-131	This issue is not resolved in AFAICS. Could you please have another look? mvn -Dtest=DynamicImportPackageTestCase test Running org.jboss.test.osgi.classloader.DynamicImportPackageTestCase Tests run: 3, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 3.457 sec <<< FAILURE! testDynamicImportWithWildcard(org.jboss.test.osgi.classloader.DynamicImportPackageTestCase) Time elapsed: 0.43 sec <<< FAILURE! java.lang.AssertionError: Unexpected ClassNotFoundException for: dynamic-wildcard-a loads org.jboss.test.osgi.classloader.support.c.C at org.junit.Assert.fail(Assert.java:92) at org.jboss.osgi.testing.OSGiTestHelper.assertLoadClass(OSGiTestHelper.java:279) at org.jboss.osgi.testing.OSGiTest.assertLoadClass(OSGiTest.java:213) at org.jboss.osgi.testing.OSGiFrameworkTest.assertLoadClass(OSGiFrameworkTest.java:135) at org.jboss.test.osgi.classloader.DynamicImportPackageTestCase.testDynamicImportWithWildcard(DynamicImportPackageTestCase.java:118)	0	0	https://issues.jboss.org/browse/JBCL-131
9163	JBCLUSTER-220	Brian, did you know that RpcDispatcher already has a method callRemoteMethods() which has an 'oob' parameter ? I don't think this is exposed at the callMethodOnCluster() method though. If a recursive synchronous method call had an OOB flag, the call would not block !	1	1	https://issues.jboss.org/browse/JBCLUSTER-220
9164	JBCLUSTER-220	In 2.6.7, only the request is OOB, *not* the response. In 2.7 and later, the response is *always* OOB. I changed that because I removed the deadlock detection code in 2.7 and later. Should I backport this change to 2.6.x (2.6.8) ?	1	1	https://issues.jboss.org/browse/JBCLUSTER-220
9165	JBCLUSTER-220	Are there chances that this fix will be applicable for JBoss AS 4.2.x which uses JGroups 2.4.1-SP4?	1	1	https://issues.jboss.org/browse/JBCLUSTER-220
9166	JBCLUSTER-220	2.7 does not work with JBoss 4.2 because of conflicts in class definitions. 2.6.7 works though I see lots of warnings: [UDP] down_thread was deprecated and is ignored [UDP] up_thread was deprecated and is ignored [UDP] Attribute "use_outgoing_packet_handler" has been deprecated and is ignored [PING] down_thread was deprecated and is ignored [PING] up_thread was deprecated and is ignored [MERGE2] down_thread was deprecated and is ignored [MERGE2] up_thread was deprecated and is ignored [FD_SOCK] down_thread was deprecated and is ignored [FD_SOCK] up_thread was deprecated and is ignored [FD] down_thread was deprecated and is ignored [FD] up_thread was deprecated and is ignored [VERIFY_SUSPECT] down_thread was deprecated and is ignored [VERIFY_SUSPECT] up_thread was deprecated and is ignored [NAKACK] down_thread was deprecated and is ignored [NAKACK] up_thread was deprecated and is ignored [NAKACK] max_xmit_size was deprecated in 2.6 and will be ignored [UNICAST] down_thread was deprecated and is ignored [UNICAST] up_thread was deprecated and is ignored [STABLE] down_thread was deprecated and is ignored [STABLE] up_thread was deprecated and is ignored [GMS] down_thread was deprecated and is ignored [GMS] up_thread was deprecated and is ignored [GMS] join_retry_timeout has been deprecated and its value will be ignored [FC] down_thread was deprecated and is ignored [FC] up_thread was deprecated and is ignored [FRAG2] down_thread was deprecated and is ignored [FRAG2] up_thread was deprecated and is ignored [STATE_TRANSFER] down_thread was deprecated and is ignored [STATE_TRANSFER] up_thread was deprecated and is ignored [STATE_TRANSFER] use_flush has been deprecated and its value will be ignored Is this normal to ignore them?	1	1	https://issues.jboss.org/browse/JBCLUSTER-220
9167	JBCLUSTER-206	If this new library is already in SVN, can you please add the link? It would also need to be in jboss repository as well so that portal can easily bundle it.	1	1	https://issues.jboss.org/browse/JBCLUSTER-206
9168	JBCLUSTER-153	Any progress on this? When can we expect timer clustering working? Please, give some light.	0	0	https://issues.jboss.org/browse/JBCLUSTER-153
9169	JBCLUSTER-117	JBCLUSTER-87 was incorporated by this fix. How does the profile service support the previous netboot functionality? If the profile service variant of netboot is already supported in JBoss AS 5.0.0, where can we find documentation about the configuration? Thanks in advance, Dirk Daems	1	0	https://issues.jboss.org/browse/JBCLUSTER-117
9170	JBCLUSTER-115	Does JBCLUSTER49 have to be done for 5.0.0 to avoid compatibility problems?	0	0	https://issues.jboss.org/browse/JBCLUSTER-115
9171	JBCLUSTER-115	Does JBCLUSTER-68 have to be done in 5.0.0 to avoid compatibility problems?	0	0	https://issues.jboss.org/browse/JBCLUSTER-115
9172	JBCLUSTER-99	#5 Ran it with 2.2.9.1 (btw, why is it not on the jboss download page?). No error: 463 tps. Looks like this is what we should be recommending? Maybe more tests are needed.	1	1	https://issues.jboss.org/browse/JBCLUSTER-99
9173	JBCLUSTER-99	>>2006-03-20 16:05:18,669 ERROR [org.jgroups.protocols.pbcast.NAKACK] (requester=1 72.16.130.57:41887, local_addr=172.16.129.57:49802) >>message with seqno=3280 not found in sent_msgs ! sent_msgs=[1477 - 3503] I am having the same problem as above mentioned with JGroups 4rc2(cannot as yet update version to higher). How can i fix it ? Tips on how you resolved the above would be welcome.	0	0	https://issues.jboss.org/browse/JBCLUSTER-99
9174	JBCLUSTER-94	Jerry, whenever you get a chance to look at this, could you open a JBAS issue for it and close this one as a duplicate? I need to figure out how to keep people from opening issues in the JBCLUSTER project; it's meant to be an area for high-level planning issues. (No blame meant here for the people who open issues; it's the logical place if you don't know otherwise).	1	1	https://issues.jboss.org/browse/JBCLUSTER-94
9175	JBCLUSTER-80	Ben, can we close this one?	0	0	https://issues.jboss.org/browse/JBCLUSTER-80
9176	JBCLUSTER-72	Ben, Is there anything I can do to help you get your content certified on Labs? Anything in the instructions vague? Let me know. I'm here to help. Cheers, Damon	1	1	https://issues.jboss.org/browse/JBCLUSTER-72
9177	JBCLUSTER-68	Does JBCLUSTER-68 have to be done in 5.0.0 to avoid compatibility problems?	0	0	https://issues.jboss.org/browse/JBCLUSTER-68
9178	JBCLUSTER-67	Can you comment on this., Scott ?	0	0	https://issues.jboss.org/browse/JBCLUSTER-67
9179	JBCLUSTER-53	Michael, this issue belongs to JBossCache. Can you go over there and open up a Jira and assign it to 1.2.4 release. We will address it. thanks!	0	0	https://issues.jboss.org/browse/JBCLUSTER-53
9180	JBCLUSTER-16	How about failover testing and performance testing? Current JBento does not provide failover testing.	1	1	https://issues.jboss.org/browse/JBCLUSTER-16
9181	JBCLUSTER-16	Do we need to provide buddy replcation node in the JSESSIONID?	0	0	https://issues.jboss.org/browse/JBCLUSTER-16
9182	JBCLUSTER-13	So the standalone Tomcat clustering has no dependency on AS, except the some common code base and packaging details. The other issue is how do we do regular testing?	1	0	https://issues.jboss.org/browse/JBCLUSTER-13
9183	JBCLUSTER-9	1. Do we already have integration test for entity bean in ejb3 with clustering?	0	0	https://issues.jboss.org/browse/JBCLUSTER-9

#	Requirement	Questions	Answered	Successfully Answered	URL
9184	JBCOMMON-130	Bootstrap Environment JBOSS_HOME: /opt/jboss-eap-5.0/jboss-as JAVA: java JAVA_OPTS: -Dprogram.name=run.sh -server -Xms256m -Xmx512m -XX:MaxPermSize=256m -Dorg.jboss.resolver.warning=true -Dsun.rmi.dgc.client.gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000 -Dsun.lang.ClassLoader.allowArraySyntax=true -Djava.net.preferIPv4Stack=true CLASSPATH: /opt/jboss-eap-5.0/jboss-as/bin/run.jar ========== 08:50:15,463 INFO [ServerImpl] Starting JBoss (Microcontainer)... 08:50:15,464 INFO [ServerImpl] Release ID: JBoss [EAP] 5.0.0.GA (build: SVNTag=JBPAPP_5_0_0_GA date=200910202128) 08:50:15,464 INFO [ServerImpl] Bootstrap URL: null 08:50:15,464 INFO [ServerImpl] Home Dir: /opt/jboss-eap-5.0/jboss-as 08:50:15,464 INFO [ServerImpl] Home URL: file:/opt/jboss-eap-5.0/jboss-as/ 08:50:15,465 INFO [ServerImpl] Library URL: file:/opt/jboss-eap-5.0/jboss-as/lib/ 08:50:15,475 INFO [ServerImpl] Patch URL: null 08:50:15,475 INFO [ServerImpl] Common Base URL: file:/opt/jboss-eap-5.0/jboss-as/common/ 08:50:15,476 INFO [ServerImpl] Common Library URL: file:/opt/jboss-eap-5.0/jboss-as/common/lib/ 08:50:15,476 INFO [ServerImpl] Server Name: jb336 08:50:15,476 INFO [ServerImpl] Server Base Dir: /opt/jboss-eap-5.0/jboss-as/server 08:50:15,476 INFO [ServerImpl] Server Base URL: file:/opt/jboss-eap-5.0/jboss-as/server/ 08:50:15,476 INFO [ServerImpl] Server Config URL: file:/opt/jboss-eap-5.0/jboss-as/server/jb336/conf/ 08:50:15,477 INFO [ServerImpl] Server Home Dir: /opt/jboss-eap-5.0/jboss-as/server/jb336 08:50:15,477 INFO [ServerImpl] Server Home URL: file:/opt/jboss-eap-5.0/jboss-as/server/jb336/ 08:50:15,477 INFO [ServerImpl] Server Data Dir: /opt/jboss-eap-5.0/jboss-as/server/jb336/data 08:50:15,477 INFO [ServerImpl] Server Library URL: file:/opt/jboss-eap-5.0/jboss-as/server/jb336/lib/ 08:50:15,477 INFO [ServerImpl] Server Log Dir: /opt/jboss-eap-5.0/jboss-as/server/jb336/log 08:50:15,477 INFO [ServerImpl] Server Native Dir: /opt/jboss-eap-5.0/jboss-as/server/jb336/tmp/native 08:50:15,477 INFO [ServerImpl] Server Temp Dir: /opt/jboss-eap-5.0/jboss-as/server/jb336/tmp 08:50:15,478 INFO [ServerImpl] Server Temp Deploy Dir: /opt/jboss-eap-5.0/jboss-as/server/jb336/tmp/deploy 08:50:17,377 INFO [ServerImpl] Starting Microcontainer, bootstrapURL=file:/opt/jboss-eap-5.0/jboss-as/server/jb336/conf/bootstrap.xml 08:50:18,639 INFO [VFSCacheFactory] Initializing VFSCache [org.jboss.virtual.plugins.cache.CombinedVFSCache] 08:50:18,642 INFO [VFSCacheFactory] Using VFSCache [CombinedVFSCache[real-cache: null]] 08:50:19,252 INFO [CopyMechanism] VFS temp dir: /opt/jboss-eap-5.0/jboss-as/server/jb336/tmp 08:50:19,255 INFO [ZipEntryContext] VFS force nested jars copy-mode is enabled. 08:50:21,898 INFO [ServerInfo] Java version: 1.6.0_17,Sun Microsystems Inc. 08:50:21,898 INFO [ServerInfo] Java Runtime: OpenJDK Runtime Environment (build 1.6.0_17-b17) 08:50:21,898 INFO [ServerInfo] Java VM: OpenJDK 64-Bit Server VM 14.0-b16,Sun Microsystems Inc. 08:50:21,898 INFO [ServerInfo] OS-System: Linux 2.6.32-71.el6.x86_64,amd64 08:50:21,907 INFO [ServerInfo] VM arguments: -Dprogram.name=run.sh -Xms256m -Xmx512m -XX:MaxPermSize=256m -Dorg.jboss.resolver.warning=true -Dsun.rmi.dgc.client.gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000 -Dsun.lang.ClassLoader.allowArraySyntax=true -Djava.net.preferIPv4Stack=true -Djava.endorsed.dirs=/opt/jboss-eap-5.0/jboss-as/lib/endorsed 08:50:21,956 INFO [JMXKernel] Legacy JMX core initialized 08:50:31,297 INFO [WebService] Using RMI server codebase: http://localhost.localdomain:8083/ 08:50:44,095 INFO [NativeServerConfig] JBoss Web Services - Stack Native Core 08:50:44,096 INFO [NativeServerConfig] 3.1.2.SP3 08:50:46,223 INFO [LogNotificationListener] Adding notification listener for logging mbean "jboss.system:service=Logging,type=Log4jService" to server org.jboss.mx.server.MBeanServerImpl@7f1fc4b8[defaultDomain='jboss'] 08:51:02,661 WARN [MappedReferenceMetaDataResolverDeployer] Unresolved references exist in	1	0	https://issues.jboss.org/browse/JBCOMMON-130
9185	JBCOMMON-85	I assigned both above mentioned issues to me. I'll take care of the soon. I just need to know, which component new EntityResolver should reside in. Or should I create new JBAS subcomponent?	1	1	https://issues.jboss.org/browse/JBCOMMON-85
9186	JBCOMMON-85	Can JBossEntityResolver be refactored to do that? JBossEntityResolver is used directly from XB (at least) as the default resolver (it can be changed though via API).	1	1	https://issues.jboss.org/browse/JBCOMMON-85
9187	JBCOMMON-39	Can you explain again what the problem is? Is there a forum thread discussing it?	1	1	https://issues.jboss.org/browse/JBCOMMON-39
9188	JBCOMMON-37	Looks like this issue fell through the cracks, sorry about that. The jboss-common 1.2.1 binaries are in the maven repository here: http://repository.jboss.org/maven2/jboss/common/jboss-common/	0	0	https://issues.jboss.org/browse/JBCOMMON-37
9189	JBCOMMON-14	Is there a date in mind for this release? If possible, I would like to work on the next release of JBoss Common to test the maven release process. I have made a couple of modifications to the build scripts, and I will probably need to make a few more changes to have the release fully automated by maven.	1	1	https://issues.jboss.org/browse/JBCOMMON-14
9190	JDG-995	The text looks good, but I think that it should refer to the methods {{removeWithVersion}} and {{replaceWithVersion}}, instead of saying "conditional replace" or "conditional remove". Aren't those the correct Hot Rod API calls [~galder.zamarreno] or [~NadirX]?	0	0	https://issues.jboss.org/browse/JDG-995

#	Requirement	Questions	Answered	Successfully Answered	URL
9191	JDG-941	[Custom Externalizers|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#customize_externalizers] have been updated to use the Full Product Name. ---- 23. [About Synchronous and Asynchronous Notifications|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#about_synchronous_and_asynchronous_notifications] has had the typo corrected, and has included `XML` in the description. ---- 24. [Cache Entry Modified Listener Configuration|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#cache_entry_modified_listener_configuration] has been updated to include a link to the API Documentation, and the classes have been highlighted. Should `Cache` be highlighted as well? ---- 25. I have corrected the link in [The Cache Listener API|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#the_cache_listener_api]. You were correct - this should have linked to the About Synchronous and Asynchronous Notifications section. ---- 26. The code in [Clustered Listener Example|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#clustered_listener_example] has been updated. ---- 27. This appears to be the same as 26? ---- 28. I have included the Maven example in [Cassandra Cache Store Programmatic Configuration|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#cassandra_cache_store_sample_programmatic_configuration]. ---- 29. I have changed the version to `8.4.0.Final-redhat-2` in [Remote Event Client Listener Example|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#remote_event_client_listener_example], and adjusted the text to be the Full Product Version. ---- 30. I have included the missing link in [Filtering Remote Events|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#filtering_remote_events]. ---- 31. I have corrected the text in [Customizing Remote Events|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#customizing_remote_events] to be the full product name. ---- 32. `KeyValueFilter` and `Converter` have been emphasized in [Event Marshalling|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#event_marshalling]. ---- 33. I have changed the text to use the full product version in [JSR-107 (JCache) API|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#jsr_107_jcache_api_2]. ---- 34. Beta references have been removed per #1, and the version has been updated in [Option 2: Adding the necessary files to the classpath|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#option_2_adding_the_necessary_files_to_the_classpath]. ---- 35. Text in [Clustering JCache Instances|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#clustering_jcache_instances] and [The REST Interface|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#the_rest_interface] has been changed to use the full product name. ---- 36. The example in [Java Client Code|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#java_client_code] now defines one import per line, and the text has been updated to use the full product name. ---- 37. The text in [About Hot Rod|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/#about_hot_rod] has been updated to use the full product name. ---- 38 and 39 don't exist.	1	1	https://issues.jboss.org/browse/JDG-941
9192	JDG-941	[~chuffman] I probably missed something, but where should I see the changes? When I look at the https://access.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/developer_guide/ , I don't see any of the changes you made. I'm probably using wrong link :)	1	1	https://issues.jboss.org/browse/JDG-941
9193	JDG-920	I've made the following changes: * In 3.2, [Required Maven Repositories|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/getting_started_guide/#required_maven_repositories], I've mentioned that this should be installed, and included a link to the below section. I have also changed the tech-preview repo to GA. * In 4.1, [Using JBoss Data Grid with Supported Containers|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/getting_started_guide/#using_jboss_data_grid_with_supported_containers_2], I have corrected the link, and added a link to the JDG 7 release notes. This link should be live once the 7.1 Release Notes are published. * In 4.7.4, [Installing Red Hat JBoss Data Grid in Karaf (Library Mode)|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/getting_started_guide/#installing_red_hat_jboss_data_grid_in_karaf_library_mode] I have corrected the typo, and the formatting is correct. * In 5.2, [The camel-jbossdatagrid component|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/getting_started_guide/#camel_jbossdatagrid_component], I have adjusted the tables here. Please let me know your thoughts on this new presentation. There are currently several conversations inside CCS regarding table formatting. * In 5.8.1, [Installing camel-jbossdatagrid for Red Hat JBoss Enterprise Application Platform|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/getting_started_guide/#installing_camel_jbossdatagrid_for_red_hat_jboss_enterprise_application_platform], I have included two links to the EAP Installation Guide, and updated the versions for EAP. One question - In section 5.2, [The camel-jbossdatagrid component|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_data_grid/7.1/html-single/getting_started_guide/#camel_jbossdatagrid_component], the version is listed as matching JDG. In the maven repo this is 2.17.0.Final-redhat-2 for the camel-jbossdatagrid component in CR3. I'm assuming this should indicate 2.17.0 instead of the JDG version?	1	1	https://issues.jboss.org/browse/JDG-920
9194	JDG-903	Hi Wolf, Why was this closed? Do you have definitive response from engineering that it will not be a supported feature.	1	0	https://issues.jboss.org/browse/JDG-903
9195	JDG-903	Wolf, Further would you know which version will support this feature fully in production?	1	0	https://issues.jboss.org/browse/JDG-903
9196	JDG-814	Hi [~pzapataf], what exactly is the reason for upgrade? I'm not very familiar with Vert.x . Is this a nice-to-have or really necessary?	1	1	https://issues.jboss.org/browse/JDG-814
9197	JDG-814	[~jsenko], will netty be upgraded to 4.1.8 in CR1 as well? I ask as I noticed the Fix Version hasn't been set.	1	1	https://issues.jboss.org/browse/JDG-814

#	Requirement	Questions	Answered	Successfully Answered	URL
9198	JDG-751	Hi, [~anistor] while verifying this issue I have noticed the following: 1. Performing: ./cli.sh -c '/subsystem=datagrid-infinispan/cache-container=clustered:register-proto-schemas(file-names=[test3.proto],file-contents=[kaboom])' Getting: {"outcome" => "success"} 2. Performing: ./cli.sh -c '/subsystem=datagrid-infinispan/cache-container=clustered:get-proto-schema-names' { "outcome" => "success", "result" => ["test3.proto"] } 3. Performing: ./cli.sh -c '/subsystem=datagrid-infinispan/cache-container=clustered:get-proto-schema-errors(file-name=test3.proto)' { "outcome" => "success", "result" => "java.lang.IllegalStateException: Syntax error in test3.proto at 1:7: unexpected label: kaboom" } 4. Performing: ./cli.sh -c '/subsystem=datagrid-infinispan/cache-container=clustered:get-proto-schemas-with-errors' { "outcome" => "success", "result" => ["test3.proto"] } 5. Then I am adding another .proto file without errors: ./cli.sh -c '/subsystem=datagrid-infinispan/cache-container=clustered:register-proto-schemas(file-names=[test2.proto],file-contents=[package test;])' {"outcome" => "success"} 6. Performing: ./cli.sh -c '/subsystem=datagrid-infinispan/cache-container=clustered:get-proto-schemas-with-errors' { "outcome" => "success", "result" => [] } 7. Performing: ./cli.sh -c '/subsystem=datagrid-infinispan/cache-container=clustered:get-proto-schema-errors(file-name=test3.proto)' { "outcome" => "success", "result" => "java.lang.IllegalStateException: Syntax error in test3.proto at 1:7: unexpected label: kaboom" } As you can see the result of step 6 doesn't contain the test3.proto, even though the step 7 contains it. My question is: is this a normal behavior? Shouldn't the outcome of step 6 still contain the test3.proto file's name? Thanks in advance, Anna. P.S. I have tried another way of operations: if on step 5 I am adding another .proto file which also contains syntax errors, then the step 6 shows only the new created .proto file's name and doesn't include the 'test3.proto' .	1	1	https://issues.jboss.org/browse/JDG-751
9199	JDG-751	[~amanukyan] Those two PRs are in now. Could you please retry the validation? Thanks	1	1	https://issues.jboss.org/browse/JDG-751
9200	JDG-739	Hi, [~gustavonalle] I have just checked the code, and seems that after changing the hibernate-search version to 5.6.0.CR1 another change was done to make it 5.6.0.Final. Is this ok? The tests are passing and seems that the change doesn't create any problems, but I wanted to recheck it with you. Thanks in advance, Anna.	1	1	https://issues.jboss.org/browse/JDG-739
9201	JDG-683	I've done the following: * Removed Tech notice from [Hot Rod C++ Client Encryption|https://access.qa.redhat.com/documentation/en/jboss-data-grid/7.1/single/developer-guide/#hot_rod_c_client_encryption] * Included section on [Hot Rod C++ Working with Sites|https://access.qa.redhat.com/documentation/en/jboss-data-grid/7.1/single/developer-guide/#hot_rod_c_client_working_with_sites] * Adjusted [Hot Rod C++ Client Encryption|https://access.qa.redhat.com/documentation/en/jboss-data-grid/7.1/single/developer-guide/#hot_rod_c_client_encryption] to include SNI function. * Included the additional Async methods in [Hot Rod C++ Client Asynchronous API|https://access.qa.redhat.com/documentation/en/jboss-data-grid/7.1/single/developer-guide/#hot_rod_c_client_asynchronous_api]: - clearAsync - putIfAbsentAsync - removeAsync - removeWithVersionAsync - replaceAsync * Added [Hot Rod C++ Client Remote Event Listeners|https://access.qa.redhat.com/documentation/en/jboss-data-grid/7.1/single/developer-guide/#hot_rod_c_client_working_with_sites] Am I missing any of the new features that were included in JDG 7.1? And should any of these be marked Tech Preview?	0	0	https://issues.jboss.org/browse/JDG-683
9202	JDG-639	Seems that {{SkipOnOs.OS.SOLARIS}} doesn't work as well, because the JniLevelDBCacheStoreTest.* and JniLevelDBStoreFunctionalTest.* tests are still running for Solaris and fail with exception mentioned by Martin. Another question is the integration tests. E.g. {{org.infinispan.server.test.cs.leveldb.LevelDBCacheStoreIT}} doesn't have the SkipOn annotation and fails on Solaris. Should the annotation is added here as well, or it only refers to functional tests?	1	1	https://issues.jboss.org/browse/JDG-639
9203	JDG-639	[~jsenko] Why did you change the state of this issue (and reassigned it to yourself and back to me)? If this was done by a mistake, please change it back to pull request sent.	1	1	https://issues.jboss.org/browse/JDG-639
9204	JDG-639	[~sebastian.laskawiec] The PR has been open for three weeks and looks invalid, so I wanted to move it back to coding in progress. Didn't know JIRA would assign it to me automatically. Does the code compile for you?	1	1	https://issues.jboss.org/browse/JDG-639
9205	JDG-523	Here is a brief discussion about this topic, from the infinispan IRC channel, 22-SEP-2016 (EST): (11:09:36 AM) me@freenode: I know you are beyond busy so I don't really want to bug you about this stuff. However, do you recommend doing a unique quickstart based on the embedded and clustered quickstarts? the secure-embedded-cache quickstart only demonstrates granular security of cache nodes, and not really fully secured infinispan/jdg library mode clustering. (11:13:52 AM) ttarrant: mwigglesworth, what features would you like to highlight ? (11:14:30 AM) ttarrant: mwigglesworth, ideally such a quickstart should be able to run both with a SecurityManager and without one (while maintaining the same behaviour((11:15:25 AM) me@freenode: I would like to highlight the use of callbackhandlers to demonstrate how an end-user/developer would enable node-authentication, with infinispan/jdg. (11:16:45 AM) me@freenode: This would be more of a "secure clustering" example, however, since the secure-embedded-cache quickstart is appropriately labeled, and has part of the security enabled, I thought it may make sense to simply fork and update it. (which I have already started to do.0 (11:25:07 AM) jbossbot: new jira ISPN-7037 Create quickstart that demonstrates library mode node security using callbackhandlers [New (Unresolved) Feature Request, Minor, Demos and Tutorials/Tasks, Unassigned] https://issues.jboss.org/browse/ISPN-7037 (11:34:34 AM) ttarrant: mwigglesworth, I'm not entirely familiar with JDF quickstart policies, but what about, instead of proliferating quickstarts, add to the existing one and enable the node auth via some switch / maven profile (11:35:13 AM) me@freenode: Sounds good.	0	0	https://issues.jboss.org/browse/JDG-523
9206	JDG-457	Looks like this was pushed to jdg-7.0.x by [~pruivo]. Is this approved for ER8 [~pzapataf] [~mgencur]?	1	0	https://issues.jboss.org/browse/JDG-457
9207	JDG-441	Hi Hiram, is this issue still valid? Thanks	0	0	https://issues.jboss.org/browse/JDG-441
9208	JDG-311	Is this issue backported to JDG?	1	1	https://issues.jboss.org/browse/JDG-311
9209	JDG-229	Thanks for explanation [~william.burns]! As I'm not aware about any JIRA for REST, I'll keep it in "reopend" state to track also REST logs. Could you please also elaborate more on console log? It's allowed to setup Console handler for HR logger, so I'd expect it works and logs into console. If this is not the case, IMHO it shouldn't be possible to set it up, otherwise user would be confused as I am :-) Thanks!	1	1	https://issues.jboss.org/browse/JDG-229
9210	JBDEPLOY-245	Is this still an issue with new CL Shutdown policy == GC?	1	0	https://issues.jboss.org/browse/JBDEPLOY-245
9211	JBDEPLOY-191	Anil, can you check if this is already fixed, via UnmarshallerFactoryDeployer::start. If not, can you give more instructions or do it yourself?	1	0	https://issues.jboss.org/browse/JBDEPLOY-191
9212	JBDEPLOY-171	This is done. There was one change required which is to make the "defaultCodeBaseRoot" URL a property of the SAR deployer rather than using the ServiceConfig which is JBossAS specific. This will need changing in the JBossAS config to add that property from the system property.	0	0	https://issues.jboss.org/browse/JBDEPLOY-171
9213	JBDEPLOY-79	This has been merged to the 2.0 and trunk.	0	0	https://issues.jboss.org/browse/JBDEPLOY-79

#	Requirement	Questions	Answered	Successfully Answered	URL
9214	JDF-842	[~pgier]: I see a pretty specific error: Tests in error: org.jboss.as.quickstarts.bean_validation_custom_constraint.MyPersonTest: The server is already running! Managed containers do not support connecting to running server instances due to the possible harmful effect of connecting to the wrong server. Please stop server before running or change to another type of container. Don't you see that? The 'bean-validation-custom-contraint' quickstart is a clone of the 'bean-validation' quickstart. It's interesting that 'bean-validation' doesn't have the same issue. I will have to investigate.	1	1	https://issues.jboss.org/browse/JDF-842
9215	JDF-819	Hi [~lhein]. Thanks for the PR. Can you please, also add Fuse to majorReleases (6) and minorReleases (6.0 and 6.1) ? You can use JPP 6 as and example.	1	1	https://issues.jboss.org/browse/JDF-819
9216	JDF-818	I just checked an the carmart quickstart is not in the JBoss EAP quickstarts: https://github.com/jboss-developer/jboss-eap-quickstarts [~fbricon]: Do you know where the quickstarts with the unusual quickstarts originate?	1	1	https://issues.jboss.org/browse/JDF-818
9217	JDF-818	It looks like this is a JDG quickstart. Adding the quickstart owners, Tristan and Martin, so they are aware of this issue. I will modify the Contributing Guide to talk about the required directory structure. [~rafabene]: Is it possible for the qstools Maven plug-in to test for this too?	1	1	https://issues.jboss.org/browse/JDF-818
9218	JDF-818	[~fbricon] : I'm not really clear what the problem is. The top level quickstarts are all in a directory one level below the top parent level. Some quickstarts, like `inter-app`, that are composed of multiple applications that talk to each other, are another level deep. Can you point to a specific problem? This is a shortened, simplified tree structure for the JBoss EAP quickstarts (https://github.com/jboss-developer/jboss-eap-quickstarts/) current directory = jboss-eap-quickstarts . ├── app-client │ ├── client-simple │ │ ├── pom.xml │ │ └── src │ │ └── ... │ ├── ear │ │ └── pom.xml │ ├── ejb │ │ ├── pom.xml │ │ └── src │ │ └── ... │ ├── pom.xml │ └── README.md ├── ejb-in-ear │ ├── ear │ │ └── pom.xml │ ├── ejb │ │ ├── pom.xml │ │ └── src │ │ └── ... │ ├── pom.xml │ ├── README.md │ └── web │ ├── pom.xml │ └── src │ └── ... │ ├── ejb-in-war │ ├── pom.xml │ ├── README.md │ └── src │ ├── ... ├── inter-app │ ├── appA │ │ ├── pom.xml │ │ └── src │ │ └── ... │ ├── appB │ │ ├── pom.xml │ │ └── src │ │ └── ... │ ├── pom.xml │ ├── README.md │ └── shared │ ├── pom.xml │ └── src │ └── ...	1	1	https://issues.jboss.org/browse/JDF-818
9219	JDF-818	[~fbricon]: : I am still not clear what the problem is. Can you point to a specific problem?	1	1	https://issues.jboss.org/browse/JDF-818
9220	JDF-818	Didn't I do that here: https://github.com/jboss-developer/jboss-developer-shared-resources/blob/master/guides/CONTRIBUTING.md#general-guidelines ?	1	1	https://issues.jboss.org/browse/JDF-818
9221	JDF-792	[~rafabene] or [~hugofirth]: Where does this issue stand and why does it belong to me?	1	1	https://issues.jboss.org/browse/JDF-792
9222	JDF-771	I am not qualified to make this code change. I thought I had added some developers to this previously, but apparently I did not, so I am doing that now. Can someone review this proposal?	1	0	https://issues.jboss.org/browse/JDF-771
9223	JDF-771	[~pmuir], yes, I can. But before I take this on, should we stop at kitchensink-angularjs? Pretty much every other quickstart (excepting contacts-mobile and it's derivatives) needs to be reviewed for the same; we could in fact copy over the patterns used in that quickstart. Also, should this be completed for WFK 2.7?	1	0	https://issues.jboss.org/browse/JDF-771
9224	JDF-771	[~pmuir], [~vineet.reynolds]: I think [~hugofirth] was offering to contribute the changes. I was looking for confirmation of the proposal and someone to do a code review when it's complete. [~hugofirth], am I reading your proposal correctly?	1	0	https://issues.jboss.org/browse/JDF-771
9225	JDF-771	[~sgilda], [~vineet.reynolds], [~pmuir], [~jowilson], [~paul.robinson]: I am happy to contribute some changes to kitchensink-angularjs, but might not have time to percolate those changes throughout the other quickstarts. [~jowilson] is right in thinking that [~paul.robinson] and I originally wanted these fixes to support a university coursework we're designing, however we reverted to building a derivative of the contacts-mobile QS and side stepped any issues. In general these changes (the corrected HTTP error codes specifically) would be a good idea long term IMO. Does that clear things up for everyone? :)	1	0	https://issues.jboss.org/browse/JDF-771
9226	JDF-769	[~rafabene]: Can we do this after we branch for the next release of JBoss EAP?	0	0	https://issues.jboss.org/browse/JDF-769
9227	JDF-740	[~lightguard], Can you elaborate on the new rules?	1	1	https://issues.jboss.org/browse/JDF-740
9228	JDF-740	[~lightguard] : I have another question. Are the meta data tags mapped to full descriptions? In other words, if I search for quickstarts that implement 'Bean Validation', will they appear in the result list?	0	0	https://issues.jboss.org/browse/JDF-740
9229	JDF-732	Hey [~manaRH]. Can you please send the branch that we work together to upstream? Thanks	1	1	https://issues.jboss.org/browse/JDF-732
9230	JDF-728	What problem are you running into here? Did you replace all jQueryMobile 1.0.x references with 1.4.x (including the JS, CSS and images)? And did you follow this guide: http://jquerymobile.com/upgrade-guide/1.4/ ?	1	0	https://issues.jboss.org/browse/JDF-728
9231	JDF-728	Is it related to ticket-monster ?	0	0	https://issues.jboss.org/browse/JDF-728
9232	JDF-719	Hi Ray, The quickstarts that use the PostgreSQL database have instructions to install the Postgres driver as a core module: https://github.com/jboss-developer/jboss-developer-shared-resources/blob/master/guides/CONFIGURE_POSTGRESQL.md#configure-the-postgresql-database-for-use-with-the-quickstarts Would it make sense to create a wrapper quickstart that makes it more obvious?	1	1	https://issues.jboss.org/browse/JDF-719
9233	JDF-706	Matej, Could you take this issue next when you have some time? Thank you! ~Lincoln	1	1	https://issues.jboss.org/browse/JDF-706
9234	JDF-706	Sande, could you post the relevant source locations for this project? http://www.jboss.org/jdf/quickstarts/jboss-as-quickstart/forge-from-scratch/ Thank you, Lincoln	1	1	https://issues.jboss.org/browse/JDF-706
9235	JDF-706	[~lincolnthree] and [~mbriskar] Are JBoss Developer Studio 8 and Forge 2 out in final release now? If so, do you want to rebase this on the 6.4.x-develop branch and submit the pull again?	0	0	https://issues.jboss.org/browse/JDF-706
9236	JDF-694	If you are looking for more things to validate, you could validate the name. It should only have letters and numbers, don't you think?	1	0	https://issues.jboss.org/browse/JDF-694
9237	JDF-682	I created general instructions in the JBoss developer shared repository here: https://github.com/jboss-developer/jboss-developer-shared-resources/blob/master/guides/DEPLOY_TO_OPENSHIFT.md Would it be okay to just point to them?	1	1	https://issues.jboss.org/browse/JDF-682
9238	JDF-673	[~fbricon]: Which repositories and branches should I be checking? I verified the following repositories and branches do not contain 'http://red.ht/jbeap-6-docs'. EAP quickstarts <https://github.com/jboss-developer/jboss-eap-quickstarts> * 6.3.x * 6.3.x-develop * 6.2.x EAP Archetypes <https://github.com/jboss-developer/jboss-eap-archetypes> * 6.2.x-develop WFK quickstarts <https://github.com/jboss-developer/jboss-wfk-quickstarts> * 2.5.x * 2.5.x-develop * 2.4.x WFK Archetypes <https://github.com/jboss-developer/jboss-wfk-archetypes> * 2.4.x * 2.5.x-develop Did I miss anything?	1	1	https://issues.jboss.org/browse/JDF-673
9239	JDF-670	I found unused imports in: kitchensink-ear/web/src/main/java/org/jboss/as/quickstarts/kitchensink_ear/util/WebResources.java Are there any others?	0	0	https://issues.jboss.org/browse/JDF-670

#	Requirement	Questions	Answered	Successfully Answered	URL
9240	JDF-649	Rafael, thanks for prompt fix, however I have just noticed a regression in 2.0.1 was found which we will need to address (ARQGRA-422). Could you please hold off this patch?	1	1	https://issues.jboss.org/browse/JDF-649
9241	JDF-607	[~rafabene] what is the problem? I can see it there http://search.maven.org/#artifactdetails|org.jboss.arquillian.extension|arquillian-drone-selenium-depchain|1.2.3.Final|pom Or bom-check is confused by pom type packaging?	1	1	https://issues.jboss.org/browse/JDF-607
9242	JDF-607	Can't this be a problem for customers that only have maven.repository.redhat.com and Maven Central? I don't think we instruct customers to include JBoss Nexus on their settings.xml	0	0	https://issues.jboss.org/browse/JDF-607
9243	JDF-607	So giving we have a bug on this update it should keep it opened, right? Should we reverte the change also ?	0	0	https://issues.jboss.org/browse/JDF-607
9244	JDF-592	[~bleathem] I was thinking about how to implement it and I the most close solution that I thought (that can work for multiple project) is the following: - Create a removePomProperties property - So you can remove any properties declaration from the original quickstart: version.richfaces, version.jboss.bom.eap, version.jboss.bom.wfk, etc - Create a insertPomPropertyFile property - A property pointing to a file that holds the full #if #else #end clause. What do you think about this proposal ? Any other idea on how to implement this solution ?	1	1	https://issues.jboss.org/browse/JDF-592
9245	JDF-592	Something that I was not considering is that this conditional expression is not only for Pom properties. Maybe a <applyPatch/> property to QSTools should work better. WDYT ?	1	1	https://issues.jboss.org/browse/JDF-592
9246	JDF-592	I was thinking of a comment syntax we could use that qstools could uncomment, however this would complicate the pom of the quickstart which defeats the usefulness of the quickstarts as a learning tool. {{<applyPatch/>}} would be brittle, with the patch file having to be adjusted as changes are made to the pom. However if we used a custom patch syntax, where the substitution was made accoding to commented markers in the pom then it would be more robust. Something like: {code:title=pom.xml} <!-- #enterprise versions of the BOM dependencies --> <version.richfaces>${richfaces-enterprise-version}</version.richfaces> <version.jboss.bom.eap>${jboss-bom-enterprise-version-eap}</version.jboss.bom.eap> <version.jboss.bom.wfk>${jboss-bom-enterprise-version-wfk}</version.jboss.bom.wfk> <!-- /#enterprise versions of the BOM dependencies --> {code} {code:title=substitutions.xml} <enterprise> #if ($enterprise == "true" || $enterprise == "y" || $enterprise == "yes") <version.richfaces>${richfaces-enterprise-version}</version.richfaces> <version.jboss.bom.eap>${jboss-bom-enterprise-version-eap}</version.jboss.bom.eap> <version.jboss.bom.wfk>${jboss-bom-enterprise-version-wfk}</version.jboss.bom.wfk> #else <version.richfaces>${richfaces-version}</version.richfaces> <version.jboss.bom>${jboss-bom-version}</version.jboss.bom> #end </enterprise> {code} In this proposal, the <!-- #enterprise --> ... <!-- /#enterprise --> comments mark the boundary of the substitution with the id _enterprise_. The substitutions file then provides the replacement code identified by the substitution code. WDYT?	1	1	https://issues.jboss.org/browse/JDF-592
9247	JDF-592	Can you include a usage sample in this jira issue?	1	0	https://issues.jboss.org/browse/JDF-592
9248	JDF-588	Wonderful. I'll probably head to Webapp. Is BRMS capable of running Webapps ? I tried to but it didn't work. I can't remember the exception at the moment.	1	1	https://issues.jboss.org/browse/JDF-588
9249	JDF-588	[~ge0ffrey] and [~sgilda] Can you review it when you have time, please ? Thanks	1	1	https://issues.jboss.org/browse/JDF-588
9250	JDF-560	[~pgier] As discussed with [~manaRH] we decided to remove the fail property since it could hide the "errors" and add an option to ignore specific dependencies only. ok ?	0	0	https://issues.jboss.org/browse/JDF-560
9251	JDF-543	Does this needs to be included on EAP BOMs also ?	1	1	https://issues.jboss.org/browse/JDF-543
9252	JDF-543	[~manaRH] [~aslak] I couldn't find the right Arquillian Persistence GAV .Final to be included. Can you provided me this info ?	1	1	https://issues.jboss.org/browse/JDF-543
9253	JDF-539	Thanks Rafael It does not work by specifying the maven settings xml file using "-s" option: mvn -s ~/building/eap6-quickstart-settings.xml -DeapMavenRepoUrl=file:///home/lgao/tmp/eap-repo/jboss-eap-6.2.0.GA-maven-repository/ -U org.jboss.maven.plugins:maven-qstools-plugin:bom-check But it works after I copied the ~/building/eap6-quickstart-settings.xml to ~/.m2/settings.xml. Is it possible to checks the resolution based on the maven settings xml file transferred using the '-s' option? Jenkins job does not have permission to write to ~/.m2/settings.xml.	1	1	https://issues.jboss.org/browse/JDF-539
9254	JDF-498	Brian, save the attached .cheatsheet.xml file in the root directory of the kitchensink-rf quickstart. It should open automatically when you import the project into JDBS. This is just a template. You will need to change wordings and descriptions because I have no idea what I'm talking about. :-) Please feel free to ask me for help if you need it! Please refer to the kitchensink quickstart for a complete example: https://github.com/jboss-developer/jboss-eap-quickstarts/blob/master/kitchensink/.cheatsheet.xml	0	0	https://issues.jboss.org/browse/JDF-498
9255	JDF-480	[~glamperi]: I have a number of questions on this: * You mentioned Fuse. Does this quickstart require Fuse? If so, you need to contribute it to the Fuse team. The quickstarts are now separated into separate repositories by product. * If it does not require Fuse, does it work on both AS 7.1 and EAP 6.1? * Could you provide an overview of the things you had to change from the original to get it to work? As a quick test, I tried to build this quickstart using both Maven Central and the EAP 6.1.1 repository. It failed because it could not find artifact `gnu-getopt:getopt:jar:1.0.13`.	1	0	https://issues.jboss.org/browse/JDF-480

#	Requirement	Questions	Answered	Successfully Answered	URL
9256	JDF-480	More information from Gary: Hi Sande, Here are answers to your earlier questions: 1) You mentioned Fuse. Does this quickstart require Fuse? If so, you need to contribute it to the Fuse team. The quickstarts are now separated into separate repositories by product. Yes, I am working with someone on the Fuse Team on this. The Fuse issue is perhaps a little worse, so the EAP Team shouldn't feel that bad The Security reference points you to non-existent references in the community. 2) If not, does this work on both AS 7.1 and EAP 6.1? My idea is that this should all be self-dependent in the helloworld-ws-secure example. I was in a hurry to get this to work and my primary objective was to show Fuse Web Service Proxy with security. When I started to run into issues with the amount of code needed to get this working and the lack of similar examples, I thought it was best to throw this out there. I would like this to all work with EAP 6.1 Maven repos without any outside, unsupported community repos. I needed to bring this repo in (http://repository.jboss.org/nexus/content/groups/public/) in addition to the downloaded EAP 6.1 repo. 3) Could you provide an overview of the things you had to change from the original to get it to work? A) Bring in the community repo mentioned in #2 into the settings.xml B) 1 additional dependency needs to be added to the pom.xml: <dependency> <groupId>org.jboss.ws.cxf</groupId> <artifactId>jbossws-cxf-client</artifactId> <version>4.1.3.Final</version> <!--<version>4.0.4.GA-redhat-1</version>--> <scope>provided</scope> </dependency> C) I renamed the Client class to ClientWS, so that I could bring in org.apache.cxf.endpoint.Client. I could have used fully qualified class names, but I was in a hurry. D) The ClientWS class was modified to accept all server certs: Client client = ClientProxy.getClient(helloWorldService); HTTPConduit conduit = (HTTPConduit)client.getConduit(); TLSClientParameters tcp = new TLSClientParameters(); tcp.setTrustManagers(new TrustManager[]{ new TrustAllX509TrustManager() }); conduit.setTlsClientParameters(tcp); AuthorizationPolicy auth = conduit.getAuthorization(); if (null == auth) { auth = new AuthorizationPolicy(); } auth.setUserName("local"); auth.setPassword("local"); E) added TrustAllX509TrustManager as an inner class F) Changed ClientTest.java to use a local wsdl and the pom.xml has a reference to the Fuse proxy location that needs to be replaced. This 2 steps should be changed for the real quickstart TBD, add ability to use client authentication as well as server only authentication. Regards, Gary	1	0	https://issues.jboss.org/browse/JDF-480
9257	JDF-455	Should we close this as compiler target is needed to specify? Otherwise the compiler target is derived from default version of maven-compiler-plugin, which is still 1.5 AFAIK.	1	1	https://issues.jboss.org/browse/JDF-455
9258	JDF-435	This change concerns me. I tried building from root after adding this module and the build fails with the following. Do the JPP quickstarts require a specific repository? $ mvn clean install '-Pdefault,!complex-dependencies'[INFO] Scanning for projects... Downloading: http://repo.maven.apache.org/maven2/org/jboss/bom/gatein-3.6-bom/1.0.1.Final-redhat-1/gatein-3.6-bom-1.0.1.Final-redhat-1.pom [ERROR] The build could not read 1 project -> [Help 1] [ERROR] [ERROR] The project org.gatein.portal.quickstarts:gatein-portal-quickstarts-parent:3.6.0.Final-qs-9-redhat-1 (/home/sgilda/GitRepos/quickstart-jdf/jpp-quickstarts/pom.xml) has 1 error [ERROR] Non-resolvable import POM: Could not find artifact org.jboss.bom:gatein-3.6-bom:pom:1.0.1.Final-redhat-1 in central (http://repo.maven.apache.org/maven2) @ line 190, column 25 -> [Help 2] [ERROR] [ERROR] To see the full stack trace of the errors, re-run Maven with the -e switch. [ERROR] Re-run Maven using the -X switch to enable full debug logging. [ERROR] [ERROR] For more information about the errors and possible solutions, please read the following articles: [ERROR] [Help 1] http://cwiki.apache.org/confluence/display/MAVEN/ProjectBuildingException [ERROR] [Help 2] http://cwiki.apache.org/confluence/display/MAVEN/UnresolvableModelException	1	1	https://issues.jboss.org/browse/JDF-435
9259	JDF-435	[~rafabene]: I spoke with Peter on IRC and he confirmed that the JPP quickstarts require that you configure your Maven settings to use their repository. The rest of the quickstart use Maven Central so it is quick and easy to do the root build at this time. I am reluctant to modify the root POM to include these quickstarts at this time. What do you think?	1	1	https://issues.jboss.org/browse/JDF-435
9260	JDF-435	I agree with you. But what about using "complex-dependencies" profile? As the description says: "All the quickstarts that have weird requirements around deployment that Maven can't handle" I'm facing the same problem on BRMS quickstarts and the decision around jpp quickstarts will be carried to brms quickstarts.	1	1	https://issues.jboss.org/browse/JDF-435
9261	JDF-426	Is Yaml the best format for this? Would JSON be better?	1	1	https://issues.jboss.org/browse/JDF-426
9262	JDF-407	Replace Errai and GWT with a poller based HTML5 solution. New services were written for obtaining the performance status, and to modify the bot state. New Backbone.js views and model were created. The polling strategy used is not a long-polling solution. It is timed to occur at the same interval for the Bot (3 seconds). What needs to be completed: * Updates to the TicketMonster guide * Consideration for a long-polling solution * Review of the JAX-RS resources created. The Bot resource is not RESTful since it is a direct port of the earlier RPC endpoint.	0	0	https://issues.jboss.org/browse/JDF-407
9263	JDF-388	Does this mean a quickstart which would show DS JSF module? There's a handful of DS examples here: https://github.com/rsmeral/deltaspike-examples expense-tracker, tickets and task-management use the DS JSF module. They are not necessarily quickstart quality, but could surely be adapted that way.	1	1	https://issues.jboss.org/browse/JDF-388
9264	JDF-388	+1. You can add it at two places: - https://deltaspike.apache.org/external.html (Please, just send me a PR and I'll merge it for you) - http://www.jboss.org/developer-materials/#!sys_type=demo (Take a look at: https://github.com/jboss-developer/jboss-developer-demos/blob/master/CONTRIBUTING.md and feel free to call me if you need any help) For this Jira, I have some ideas for kitchensink-deltaSpike: - Use JsfMessage - Use Pages (type-safe view config) What do you think ?	1	1	https://issues.jboss.org/browse/JDF-388
9265	JDF-358	I have done some research and it appears you use `:with_toc_data` to add HTML anchors to each header in the output HTML so you can link to each section. I checked the dist/github-flavored-markdown.rb script, and it appears to already use this option. However, when I look at the generated HTML for the README.md files, I can see that it generates an `id`: <h2 id="toc_1">What is it?</h2> But I don't see the anchors. I found some examples (https://github.com/vmg/redcarpet, http://dev.af83.com/2012/02/27/howto-extend-the-redcarpet2-markdown-lib.html) that used `:with_toc_data => true`. I tried it but it didn't seem to help. I see someone else has this issue: http://stackoverflow.com/questions/14673807/redcarpet-anchors-to-headers [~rafabene]: If you have any hints, let me know. I will continue to investigate.	1	1	https://issues.jboss.org/browse/JDF-358
9266	JDF-345	Can we at least update the lab guide with instructions in the PDF when using the PDF that this cut and pastes badly, and maybe link to the HTML source? To make this easier to understand, and note it is kinda messy. Maybe link to this jira also, so people know we are working it. Jim	1	1	https://issues.jboss.org/browse/JDF-345
9267	JDF-313	We have some test cases that show how to create and configure custom identity stores. Would that be enough ? Or the idea is provide a very simple store (with only a few methos implemented) with focus on the configuration ?	1	1	https://issues.jboss.org/browse/JDF-313

#	Requirement	Questions	Answered	Successfully Answered	URL
9268	JDF-312	This quickstart is complete, however depends on a currently unreleased version of PicketLink. Will send a pull request once PicketLink 2.5.0.Beta4 is released. Branch containing the quickstart is here: https://github.com/sbryzak/jboss-as-quickstart/tree/picketlink-permission-resolver	0	0	https://issues.jboss.org/browse/JDF-312
9269	JDF-282	about file size - why would you not want it there ? it is to provide an idea to the user how big the download will be before we start the download. yes jbosstools does not *need* it, but users would like to know if they are downloading 1 mb or 500 mb :)	1	1	https://issues.jboss.org/browse/JDF-282
9270	JDF-282	Why don't we put this in a new file e.g. legacy-stacks.yml. This could be used only by JBDS.	1	1	https://issues.jboss.org/browse/JDF-282
9271	JDF-282	I'm not against the idea of multiple stacks files, but, I don't think we should create a second file without being clear on a few things. By declaring that it'd "only be used by JBDS" seems to make its scope "JBDS responsibility" and removes your hands from keeping it up to date and accurate. If that's what we wanted, we would just maintain our own stacks.yaml and not even involve you. So we need a better 'scope' of the document than just 'only used by jbds'. 1) What is the scope of each document. This should not even mention consumers such as JBDS, but be more open to who's consuming it while being more exact on scope. 2) When do items become 'legacy', ie, when would they need to be moved from the primary stacks to the 'legacy' file? Take for example, seam. Currently seam 2.3.0 is in this patch. If a 2.3.1 comes out, do we remove 2.3.0 from the primary stacks and move it to the 'legacy'? There needs to be some rules here as to what we're planning on doing. Otherwise it'll turn into a big mess.	1	1	https://issues.jboss.org/browse/JDF-282
9272	JDF-282	If that's what you guys wish, then can we make it happen? Our client would then need to fetch stacks from two separate files, and, I assume, simply group them into a larger array. Which should be fine as long as there's no conflicting information between them, correct?	1	1	https://issues.jboss.org/browse/JDF-282
9273	JDF-282	We could also do this via the existing stacks file. How soon do you need this?	1	1	https://issues.jboss.org/browse/JDF-282
9274	JDF-269	Of course. I'll let you know. Can you guess archetype version right now? Will it be 7.1.3.CR6?	1	1	https://issues.jboss.org/browse/JDF-269
9275	JDF-264	I think I released with the latest hibernate-core .Final version at the moment. Maybe I'm confusing but I can't find the 4.3.1.Final version on maven repositories. Do you mind me send the link to any repository and GAV which contains the 4.3.1.Final ?	1	1	https://issues.jboss.org/browse/JDF-264
9276	JDF-254	In general, please file one issue per issue and don't compound issues, it makes it hard to track them, as not all may be fixable at the same time. 1) This is clearly documented in the tutorial, and is also a dupe of JDF-6 2) This is a bug in Forge 1.1.x, it's not reproducible in 1.2.x 3) This is a covered in https://issues.jboss.org/browse/JDF-257 which I've spun out of this issue (because it's covering too much).	1	1	https://issues.jboss.org/browse/JDF-254
9277	JDF-254	Karel Piwko <kpiwko@redhat.com> made a comment on [bug 929102|https://bugzilla.redhat.com/show_bug.cgi?id=929102] Tomas, can you please verify that 2) cannot be reproduced with 1.2.x? It seems that Pete and you are experiencing different behavior.	1	1	https://issues.jboss.org/browse/JDF-254
9278	JDF-254	Tomas Repel <trepel@redhat.com> made a comment on [bug 929102|https://bugzilla.redhat.com/show_bug.cgi?id=929102] (In reply to comment #8) > Tomas, can you please verify that 2) cannot be reproduced with 1.2.x? It > seems that Pete and you are experiencing different behavior. I have tried using: Ticekt Monster 2.0.7.Final-redhat-1 (the one from WFK 2.2.0.CR1 jboss-wfk-2.2.0-ticketmonster-demo.zip) Forge 1.2.2.Final The richfaces-bom can't be found for me.	1	1	https://issues.jboss.org/browse/JDF-254
9279	JDF-254	Karel Piwko <kpiwko@redhat.com> made a comment on [bug 929102|https://bugzilla.redhat.com/show_bug.cgi?id=929102] Do you guys use different settings.xml by any chance?	1	1	https://issues.jboss.org/browse/JDF-254
9280	JDF-231	Karel, how do we do this? https://github.com/jboss-jdf/jboss-bom/blob/master/jboss-javaee-6.0-with-tools/pom.xml - it's not in the BOM, it's coming in transitively.	0	0	https://issues.jboss.org/browse/JDF-231
9281	JDF-216	To be accurate, the dialog allows me to create the project anyway. But I'm just wondering why it says the server is missing. Maybe this is something in JBDS/JBT itself? But then what is the effect of the supported runtimes definition in the yaml file?	1	1	https://issues.jboss.org/browse/JDF-216
9282	JDF-209	How about: Check for a `<finalName>${project.artifactId}</finalName>` in the file to prevent appending version numbers to archives.	1	1	https://issues.jboss.org/browse/JDF-209
9283	JDF-209	[~rafabene]: Does this tool use the .m2/settings.xml file? The reason I ask is because when I test it against the EAP 6.1 quickstarts, I get an error: [ERROR] [ERROR] Non-resolvable import POM: Could not find artifact org.jboss.bom:jboss-javaee-6.0-with-tools:pom:1.0.4.Final-redhat-3 in central (http://repo.maven.apache.org/maven2) @ line 76, column 25 So it appears to be going to Maven Central instead of the repository I configured in my settings. It's not a big deal. I believe I can work around it by copying the repository into .m2/ instead of using the settings.xml file.	1	1	https://issues.jboss.org/browse/JDF-209
9284	JDF-208	Hi [~ppalaga] Is this Jira valid ?	0	0	https://issues.jboss.org/browse/JDF-208

#	Requirement	Questions	Answered	Successfully Answered	URL
9285	JDF-192	comment to the readme file? {noformat} ^C12:19:12,437 INFO [org.apache.coyote.http11] (MSC service thread 1-4) JBWEB003075: Coyote HTTP/1.1 pausing on: http-/127.0.0.1:8080 12:19:12,438 INFO [org.apache.coyote.http11] (MSC service thread 1-4) JBWEB003077: Coyote HTTP/1.1 stopping on : http-/127.0.0.1:8080 12:19:12,440 INFO [org.apache.catalina.core] (MSC service thread 1-4) JBWEB001079: Container org.apache.catalina.core.ContainerBase.[jboss.web].[default-host].[/] has not been started 12:19:12,439 INFO [org.jboss.as.quickstarts.cluster.hasingleton.service.ejb.HATimerService] (MSC service thread 1-6) Stop HASingleton timer service 'org.jboss.as.quickstarts.cluster.hasingleton.service.ejb.HATimerService' 12:19:12,449 INFO [org.apache.tomcat.util] (http-/127.0.0.1:8080-Acceptor) JBWEB003022: Channel processing failed 12:19:12,450 INFO [org.jboss.as.connector.subsystems.datasources] (MSC service thread 1-4) JBAS010409: Unbound data source [java:jboss/datasources/ExampleDS] 12:19:12,459 INFO [org.apache.coyote.ajp] (MSC service thread 1-2) JBWEB003048: Pausing Coyote AJP/1.3 on ajp-/127.0.0.1:8009 12:19:12,459 INFO [org.apache.coyote.ajp] (MSC service thread 1-2) JBWEB003051: Stopping Coyote AJP/1.3 on ajp-/127.0.0.1:8009 12:19:12,473 ERROR [org.jboss.as.ejb3.invocation] (MSC service thread 1-6) JBAS014134: EJB Invocation failed on component SchedulerBean for method public abstract void org.jboss.as.quickstarts.cluster.hasingleton.service.ejb.Scheduler.stop(): javax.ejb.EJBException: JBAS014559: Invocation cannot proceed as component is shutting down at org.jboss.as.ejb3.component.interceptors.ShutDownInterceptorFactory$1.processInvocation(ShutDownInterceptorFactory.java:59) [jboss-as-ejb3-7.2.0.Alpha1-SNAPSHOT.jar:7.2.0.Alpha1-SNAPSHOT] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) [jboss-invocation-1.1.1.Final.jar:1.1.1.Final] at org.jboss.as.ejb3.component.interceptors.LoggingInterceptor.processInvocation(LoggingInterceptor.java:59) [jboss-as-ejb3-7.2.0.Alpha1-SNAPSHOT.jar:7.2.0.Alpha1-SNAPSHOT] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) [jboss-invocation-1.1.1.Final.jar:1.1.1.Final] at org.jboss.as.ee.component.NamespaceContextInterceptor.processInvocation(NamespaceContextInterceptor.java:50) [jboss-as-ee-7.2.0.Alpha1-SNAPSHOT.jar:7.2.0.Alpha1-SNAPSHOT] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) [jboss-invocation-1.1.1.Final.jar:1.1.1.Final] at org.jboss.as.ejb3.component.interceptors.AdditionalSetupInterceptor.processInvocation(AdditionalSetupInterceptor.java:54) [jboss-as-ejb3-7.2.0.Alpha1-SNAPSHOT.jar:7.2.0.Alpha1-SNAPSHOT] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) [jboss-invocation-1.1.1.Final.jar:1.1.1.Final] at org.jboss.as.ee.component.TCCLInterceptor.processInvocation(TCCLInterceptor.java:45) [jboss-as-ee-7.2.0.Alpha1-SNAPSHOT.jar:7.2.0.Alpha1-SNAPSHOT] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) [jboss-invocation-1.1.1.Final.jar:1.1.1.Final] at org.jboss.invocation.ChainedInterceptor.processInvocation(ChainedInterceptor.java:61) [jboss-invocation-1.1.1.Final.jar:1.1.1.Final] at org.jboss.as.ee.component.ViewService$View.invoke(ViewService.java:165) [jboss-as-ee-7.2.0.Alpha1-SNAPSHOT.jar:7.2.0.Alpha1-SNAPSHOT] at org.jboss.as.ee.component.ViewDescription$1.processInvocation(ViewDescription.java:182) [jboss-as-ee-7.2.0.Alpha1-SNAPSHOT.jar:7.2.0.Alpha1-SNAPSHOT] at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) [jboss-invocation-1.1.1.Final.jar:1.1.1.Final] at org.jboss.invocation.ChainedInterceptor.processInvocation(ChainedInterceptor.java:61) [jboss-invocation-1.1.1.Final.jar:1.1.1.Final] at org.jboss.as.ee.component.ProxyInvocationHandler.invoke(ProxyInvocationHandler.java:72) [jboss-as-ee-7.2.0.Alpha1-SNAPSHOT.jar:7.2.0.Alpha1-SNAPSHOT] at org.jboss.as.quickstarts.cluster.hasingleton.service.ejb.Scheduler$$$view1.stop(Unknown Source) [jboss-as-cluster-ha-singleton-service.jar:] at	0	0	https://issues.jboss.org/browse/JDF-192
9286	JDF-176	[~pmuir]: Vikram is asking if this is even possible yet? Does the full release version of RESTEasy support asynch processing?	1	1	https://issues.jboss.org/browse/JDF-176
9287	JDF-172	Using JBDS to publish ticket-monster is pretty easy (I can do a screencast if that helps): # ASSERT: have ticketmonster repo cloned to your machine # EXEC: Launch *OpenShift Application* wizard from central and create a jbossas-7/eap-6 application. # ASSERT: you get the application imported into a project in your workspace, and a server adapter is created for it. # EXEC: *File->Import->Maven->Existing Maven Project* and point it to the demo folder in the ticket-monster repo. # EXEC: drag and drop the ticket-monster project (you just imported in the former step) to your Openshift server adapter (*Servers* view). When it asks you if it shall publish, answer *No* # EXEC: in the openshift project: go to the deplyoments folder and rename the ticket-monster.war - that was created when you dragged'n'dropped it - to ROOT.war # EXEC: open the *Navigator* view, browse the imported OpenShift project and get to the *.openshift/markers* folder. Create an empty file called *skip_maven_build*. # EXEC: select .openshift/markers/skip_maven_build and pick *Team->Add to Index* from context-menu. Your imported Openshift project and the server adapter should now look like this: !ticket-monster.png! # EXEC: Tell your OpenShift server adapter to *Publish* To now look at your ticket-monster application, wait until the publishing progress disappeared from the status bar and point your browser to the OpenShift application: *Show in -> Web Browser* in the context menu of your OpenShift server adapter.	1	1	https://issues.jboss.org/browse/JDF-172
9288	JDF-172	Andre, can you modify your process for binary deployment?	1	1	https://issues.jboss.org/browse/JDF-172
9289	JDF-172	Not sure I follow you completely. In OpenShift source and binary deployments are done via pushing the source or pushing a war in the deployments folder. Those 2 are not mutually exclusive, you have both at the same time in cmd-line and JBDS. There's no difference btw the 2 in this matter. The only additional feature in JBDS is that it'll war any (WTP) project you drag and drop to the OpenShift server adapter (to the deployments folder). When publishing, the server adapter git pushes the wohle repo to openshift (exactly as a git push does), pushing the source and wars within the deployments folder. I'm thus not sure what workflow isn't clear, what added value there is with a checkbox. JBDS completely matches the workflow you have when you git push on the cmd-line. In the above scenario we rename the war (that JBDS built when we drag'n dropped the ticket-monster to the adapter) from ticket-monster.war to ROOT.war so that it'll appear at /. Since source and binary deployments both happen, the source would get built, deployed to ROOT.war, and override our prebuilt ROOT.war. We avoid that by skipping the maven build with the skip_maven_build marker. Do I miss something?	1	1	https://issues.jboss.org/browse/JDF-172
9290	JDF-172	sure, no prob with discussing, I just dont see any obvious improvement right now. The only improvement I spot with checkboxes is that one would have the "skip_maven_build" marker set. IMHO we should then provide some ceckbox "skip source build". Missing something? What feels unnatural in the above usecase is that you need the OpenShift repo imported as a project whereas you'd only want to work with the ticket-monster app/repo. Agree?	1	1	https://issues.jboss.org/browse/JDF-172
9291	JDF-172	+1 for having a 'binary deployment only' option which will go disable the the build marker...but what do we do about the existing project content - just leave it in place I would say. wdyt?	0	0	https://issues.jboss.org/browse/JDF-172
9292	JDF-159	The Deactivateable feature was added to cdi-portable-extension and cdi-veto quickstarts. But I'm now wondering if we should keep just on cdi-portable-extension to avoid any confusion that could exist mixing a veto with deactivateable. What you think?	0	0	https://issues.jboss.org/browse/JDF-159
9293	JDF-159	[~lightguard] wrote: I would say if your extension is going in a library it should, if it's for your application, I don't see the point. Why would you create an extension only to disable it? I just don't want to muddy the waters and have people wonder what's going on or think it's required to create an extension.	0	0	https://issues.jboss.org/browse/JDF-159
9294	JDF-159	In this case I think that for JDF-159 we shouldn't deactivate any of the existing (relevant) quickstarts but create another quickstart that disable some of the existing DeltaSpike extensions. Do you all agree ?	0	0	https://issues.jboss.org/browse/JDF-159

#	Requirement	Questions	Answered	Successfully Answered	URL
9295	JDF-153	I've done some ARQ tests using the RESTEasy Client, didn't have any problems, but as you say doing this from within an application may different. What difficulties have you seen?	0	0	https://issues.jboss.org/browse/JDF-153
9296	JDF-153	wasn't this met by jax-rs-client quickstart ?	1	1	https://issues.jboss.org/browse/JDF-153
9297	JDF-152	Hi Jeff, Thanks for volunteering to do this one! I put it in my name to take it off the table. Could you send me your email address? Thanks, Sande	1	1	https://issues.jboss.org/browse/JDF-152
9298	JDF-151	[~bleathem] Hi Brian, The tests said that: Can't resolve Archetype [org.richfaces.archetypes:richfaces-archetype-kitchensink:4.2.3.Final-1] I looked at https://repository.jboss.org/nexus/content/groups/public/org/richfaces/archetypes/richfaces-archetype-kitchensink/ and I didn't see this version: 4.2.3.Final-1 that you added in this commit: https://github.com/bleathem/jdf-stack/commit/6b222608338c121a8d0b3fc33f0a0d3ce672df2b Could you check the release of this version ? Thanks.	1	1	https://issues.jboss.org/browse/JDF-151
9299	JDF-151	[~bleathem] Brian, This BOM (org.jboss.spec:jboss-javaee-6.0:pom:3.0.0.Final-redhat-1) is available at the EAP6 maven repository. I'm not sure with we should merge the PR before the staging is finished since this can break anything. Shouldn't we wait the RF KS archetype to be released ? Most of stacks tools can use an alternate stacks repo. So you can use this feature to test your changes like [~fbricon] commented on JBIDE-13059: The archetype can be tested in Central by starting eclipse/jbds using the system property : -Dorg.jboss.examples.stacks.url=https://raw.github.com/bleathem/jdf-stack/a4d996b89e5a311759c2f8c9445cfc75a70cc49f/stacks.yaml	1	1	https://issues.jboss.org/browse/JDF-151
9300	JDF-147	Does JBoss Logging allow us to actually get the value of a message rather than log it?	0	0	https://issues.jboss.org/browse/JDF-147
9301	JDF-134	For jboss-bom release [~pgier] and [~kpiwko] asked me to not report by email, but just submit a JBPAPP jira instead as this example: https://issues.jboss.org/browse/JBPAPP-10284 . Should we keep notifying them also by email ?	1	1	https://issues.jboss.org/browse/JDF-134
9302	JDF-134	[~rafabene] this looks great, apart from my comment on removing EAP references :-) Also, I wondered if we want to add a question to the script "Do you want to send release notifications?" as there might be some situations where the user doesn't. Or perhaps a flag --do-not-notify to the command line...	1	1	https://issues.jboss.org/browse/JDF-134
9303	JDF-132	I can see it. I am using Chrome Version 22.0.1229.94. I can also see it on Firefox 16.0.1. Pete, what are you using?	1	1	https://issues.jboss.org/browse/JDF-132
9304	JDF-132	* Should open in a new tab * Would be nice to have it like the other ones, where it has the modal dialog that explains how to use it - show how to import to Eclipse maybe? Ok. I'll do the modal dialog instead of opening a new tab. Do you know if is there a DOC ready that explains that ? * What is the icon really? It seems a lego build block. I used because build makes me associate it with Maven :)	1	1	https://issues.jboss.org/browse/JDF-132
9305	JDF-129	The patches do *not* modify SectionAllocationBean. This is a Forge-generated class which apparently isn't generated correctly. Can you attach the generated source, Hardy? Does the project compile if you don't apply the patches?	1	1	https://issues.jboss.org/browse/JDF-129
9306	JDF-129	Hardy, Would you mind attaching the project files as generated by Forge as well? Trying to figure out what the error actually is (and the copy/pasted code is not too much of a help - I believe also due to the reformatting). Thanks, Marius	1	1	https://issues.jboss.org/browse/JDF-129
9307	JDF-128	No, Forge does not deliver the complete functionality we require for this example. The patches are required for: - changing the look and feel of the application; - including functionality that is not provided by Forge (e.g. handling of images as described in the tutorial); - fixing critical issues in Forge (we probably could do away with those by testing against a later release where they are presumably fixed) Now, I haven't been able to reproduce the above with JBoss Forge 1.1.1 (on a Mac OS X). What is your OS, Hardy?	1	1	https://issues.jboss.org/browse/JDF-128
9308	JDF-128	Hi Cuneyt, Can you try tag 2.0.4.Final or master? Thanks, Marius	1	1	https://issues.jboss.org/browse/JDF-128
9309	JDF-127	Hardy, can you list an exact list of points where the readme is not in sync, or is this simply missing the Forge installation bits? Because the Building and Running on OpenShift bits do not seem out of sync. Given the scenario when the repository is forked by the user, having the readme file is a must in order for it being consistent with the other example we provide.	0	0	https://issues.jboss.org/browse/JDF-127
9310	JDF-108	When you finish the quickstarts in the jdf repository you can assign to me again so I'll update the site	1	1	https://issues.jboss.org/browse/JDF-108
9311	JDF-108	Karel, in the comment above, you say kitchensink-html5-mobile is targeted for WFK, however, it appears in the master eap6 branch of https://github.com/jboss-eap/quickstart. 1. Should it be in the WFK instead? 2. If so, what about helloworld-html5?	0	0	https://issues.jboss.org/browse/JDF-108
9312	JDF-101	Jiri, Can you clarify why is it invalid ? The suggest BOM is here: http://mvnrepository.com/artifact/org.jboss.spec/jboss-javaee-6.0/3.0.0.Final	0	0	https://issues.jboss.org/browse/JDF-101
9313	JDF-76	Have we actually seen the cdn fail?	1	1	https://issues.jboss.org/browse/JDF-76
9314	JDF-70	Jason, would you be interested in coordinating this? My thought was to ask some of the seam module leads to write one or two, to show off what they thought was an interesting problem.	1	1	https://issues.jboss.org/browse/JDF-70
9315	JDF-70	We have the cdi-portable-extension quickstart now, however, it only covers some of the SPI. Shall we do one or two more to cover other parts of the SPI?	1	1	https://issues.jboss.org/browse/JDF-70
9316	JDF-51	Burr, can you do this?	1	1	https://issues.jboss.org/browse/JDF-51
9317	JDF-46	Hi [~sgilda]. Could you check please if this Jira can be closed ? Thanks	1	1	https://issues.jboss.org/browse/JDF-46
9318	JDF-19	Marius, did you document Postgres?	1	1	https://issues.jboss.org/browse/JDF-19
9319	JDF-8	That link is 404 now. What is it that we're supposed to do for this ticket?	1	0	https://issues.jboss.org/browse/JDF-8
9320	JDF-1	I thinks that the separation between quickstarts and demos solved this. Can you confirm that, [~pmuir] ? Thanks	0	0	https://issues.jboss.org/browse/JDF-1
9321	EMB-61	Newly-attached patch resolves Endorsed issues (though you must first add endorsed dirs properly, not through Surefire config). Error is now: org.jboss.ws.WSException: Cannot find configFile: META-INF/standard-jaxws-client-config.xml at org.jboss.ws.metadata.config.JBossWSConfigFactory.filenameToURL(JBossWSConfigFactory.java:196) at org.jboss.ws.metadata.config.JBossWSConfigFactory.getConfig(JBossWSConfigFactory.java:150) at org.jboss.ws.metadata.umdm.EndpointMetaData.initEndpointConfigMetaData(EndpointMetaData.java:882) at org.jboss.ws.metadata.umdm.EndpointMetaData.initEndpointConfig(EndpointMetaData.java:858) at org.jboss.ws.metadata.builder.jaxws.JAXWSClientMetaDataBuilder.rebuildEndpointMetaData(JAXWSClientMetaDataBuilder.java:295) at org.jboss.ws.core.jaxws.spi.ServiceDelegateImpl.getPortInternal(ServiceDelegateImpl.java:281) at org.jboss.ws.core.jaxws.spi.ServiceDelegateImpl.getPort(ServiceDelegateImpl.java:208) at javax.xml.ws.Service.getPort(Service.java:92) at org.jboss.embedded.testsuite.wsclient.Embedded.getEmbeddedWsPort(Embedded.java:63) at org.jboss.embedded.testsuite.ServerTestCase.testWs(ServerTestCase.java:487)	1	1	https://issues.jboss.org/browse/EMB-61

#	Requirement	Questions	Answered	Successfully Answered	URL
9322	EMB-42	Good. So as a prereq for this, what do we need to install into AS? Is there a single JAR distro we can depend upon in Maven, and deploy directly into the server? Or some more manual patching first?	1	0	https://issues.jboss.org/browse/EMB-42
9323	EMB-42	WB is integrated into JBoss 5.1.0 and above, so probably nothing much to update I think. In terms of tests, I would recommend copying a few out of the WB testsuite as a starting point. Or are you thinking to actually reuse these tests?	0	0	https://issues.jboss.org/browse/EMB-42
9324	EMB-37	Why are the slf4j classes added in Seam, rather than in embedded legacy?	1	1	https://issues.jboss.org/browse/EMB-37
9325	EMB-37	because of misunderstanding - anyway in Seam there was a requirement for that and therefore I understood it differently. So I have added slf4j-* to the distribution, but there is a question if it shouldn't be version beta3.SP10? WDYT?	1	1	https://issues.jboss.org/browse/EMB-37
9326	EMB-28	Might this later be moved to something in bootstrap?	0	0	https://issues.jboss.org/browse/EMB-28
9327	EMB-12	What does "legacy" version mean? Is this done?	1	0	https://issues.jboss.org/browse/EMB-12
9328	EMB-11	Pete, the old "embedded" module won't have these targets, but the new projects will. Is that enough to close this? If not, please provide for some more description.	1	0	https://issues.jboss.org/browse/EMB-11
9329	JBEAP-11780	??AFAICR we've only done so once?? Unfortunately, there has been made another change, see [my comment (JBEAP-11781) | https://issues.jboss.org/browse/JBEAP-11781?focusedCommentId=13426540&page=com.atlassian.jira.plugin.system.issuetabpanels%3Acomment-tabpanel#comment-13426540] in cloned 7.0.z JIRA please (I should have probably posted it in here). Not to bringing confusion to topic, I'm digging more into patching and I've just found out that the configuration files are excluded from patches completely [1]. Why is that so? Un-ignoring these would leads to patching conflict, which is IMO more acceptable than missing piece of configuration, without even knowing it. Moreover, CLI (not sure about web-console) has features for dealing with conflicts. -- [1] https://github.com/jbossas/patch-gen/blob/master/patch-gen/src/main/java/org/jboss/as/patching/generator/DistributionStructureImpl.java#L60	0	0	https://issues.jboss.org/browse/JBEAP-11780
9330	JBEAP-11764	-Setting Blocker priority. The text is confusing and misleading.- I am sorry, I should have read the description more carefully. I expected that this issue is about the first Note, not about the first Important. Reverting priority back. The text in the description is confusing -- the cited note does not describe requirements on the feature -- the other note does -- and the other note needs the improvement. [~rhn-support-pnag], could you clarify the scope of this issue please?	0	0	https://issues.jboss.org/browse/JBEAP-11764
9331	JBEAP-11670	It states: For information about how to deploy and run the quickstarts and to run the Arquillian tests in JBoss Developer Studio, see Run the Quickstarts in JBoss Developer Studio in the Development Guide for JBoss EAP. If you follow the link to _+Run the Quickstarts in JBoss Developer Studio+_ [1], you'll see a subsection titled *Run the bean-validation Quickstart* that describes how to run the Arquillian tests. Is that too obscure? [1] https://access.redhat.com/documentation/en-us/red_hat_jboss_enterprise_application_platform/7.0/html-single/development_guide/#run_the_quickstarts_in_red_hat_jboss_developer_studio	1	1	https://issues.jboss.org/browse/JBEAP-11670
9332	JBEAP-11670	[~sgilda], I think the link to devel guide is fine. However, I agree with [~mjurc] that these instructions could be included also directly in the shared resources [1]. Similar to what we have for "Quickstarts Containing an EAR" or "Quickstarts Containing a Server and Java Client", there could be "Quickstarts Containing Arquillian tests". What do you think? [1] https://github.com/jboss-developer/jboss-developer-shared-resources/blob/master/guides/USE_JBDS.md#use-red-hat-jboss-developer-studio-or-eclipse-to-run-the-quickstarts	1	0	https://issues.jboss.org/browse/JBEAP-11670
9333	JBEAP-11606	Why is this needed? Is it essential for ER2? CC [~psakar]	0	0	https://issues.jboss.org/browse/JBEAP-11606
9334	JBEAP-11565	How do you plan to handle domain? Do you set MALLOC_ARENA_MAX for DC, HC and servers? Or only for servers?	0	0	https://issues.jboss.org/browse/JBEAP-11565
9335	JBEAP-11391	I get a 404 error on that link - is there an alternative link?	0	0	https://issues.jboss.org/browse/JBEAP-11391
9336	JBEAP-11167	Reopen on revision e7cfdf2f [~nids1409]: I'm not sure, what do you mean by "`logger` attribute" [here|https://gitlab.cee.redhat.com/red-hat-jboss-enterprise-application-platform-documentation/eap-documentation/merge_requests/3116/diffs#08b3c1ebfa9fec2bfb153588a31f783869b67055_0_4]. Can you please change "`logger` attribute" to "logging"? So the whole sentence would be: * "RESTEasy logs the used providers and interceptors in the `DEBUG` level of logging."	1	1	https://issues.jboss.org/browse/JBEAP-11167
9337	JBEAP-11167	[~nids1409]: Your JIRAs are somewhat confusing because even though you mark them as resolved, you have no comment indicating where to preview the fix. It takes extra effort to flip the JIRAs when we push to stage. Could you please remember to add the preview URL when you mark the JIRA resolved?	1	0	https://issues.jboss.org/browse/JBEAP-11167
9338	JBEAP-11129	I've qa_acked this, but I have some notes. After a cursory look at the PR, I think this fully implements EAP7-708. [~pferraro] is that correct, or is there some more integration work needed? [~pferraro] do you have some XML fragments or CLI commands that show how to configure the new AUTH protocol and new token types? If you don't have anything at hand, I'll eventually figure it out, but any resources would be appreciated. Thanks. Since security-related RFEs require that legacy configuration still works, I'll assume that AUTH legacy configuration also still works (e.g. configuration with no Elytron references required), in a similar way encrypt protocols do. [~pferraro] are there some notable issues with legacy configuration?	1	1	https://issues.jboss.org/browse/JBEAP-11129
9339	JBEAP-11036	[~rsvoboda] [~jmartisk] [~simkam] Can you ack please? Apparently this will fix (or be needed at least) JBEAP-10651.	1	1	https://issues.jboss.org/browse/JBEAP-11036
9340	JBEAP-11000	[~kabirkhan] [~jmesnil] could you move incorporated jiras to ready fo qa? Thanks. Edit: There are 11 jiras [1] for JMS and ActiveMQ component in resolved state and fix version 7.1.0.DR19. Can they be moved as well? [1] https://issues.jboss.org/issues/?jql=project%20%3D%20JBEAP%20AND%20fixVersion%20%3D%207.1.0.DR19%20AND%20component%20in%20(ActiveMQ%2C%20JMS)	1	0	https://issues.jboss.org/browse/JBEAP-11000
9341	JBEAP-10941	[~bmaxwell], do we need https://issues.jboss.org/browse/JBEAP-10156 for 7.0.7?	1	1	https://issues.jboss.org/browse/JBEAP-10941

#	Requirement	Questions	Answered	Successfully Answered	URL
9342	JBEAP-10853	Reopen on EAP 7.1.0.DR19. Module add operation can't detect recursive symlinks. [~jdenise]: Can "module add" operation detect recursive symlinks, skip this symlinks and print warning message about this recursive symlinks? Example: {noformat} cd mkdir test cd test mkdir test1 mkdir test2 echo "test1" > test1/a echo "test2" > test2/b cd test1 ln -s /home/mkopecky/test/test2 cd .. cd test2 ln -s /home/mkopecky/test/test1 module add --name=org.my.example --resources=/home/mkopecky/test/test1/ {noformat} *Results:* \[disconnected /\] module add --name=org.my.example --resources=/home/mkopecky/test/test1/ Module not added. An error occurred while copying directory /home/mkopecky/test/test1 to /home/mkopecky/playground/eap/7.1.0.DR19/jboss-eap-7.1/modules/org/my/example/main/test1 :java.nio.file.FileSystemException: /home/mkopecky/playground/eap/7.1.0.DR19/jboss-eap-7.1/modules/org/my/example/main/test1/test2/test1/test2/test1/test2/test1/test2/test1/test2/test1/test2/test1/test2/test1/test2/test1/...../test2/test1: File name too long \[disconnected /\] This situation is quite safe, because size of files included in test folder is small. But if test1 folder contains for example big zip file, module add operation causes no free space on disk.	1	1	https://issues.jboss.org/browse/JBEAP-10853
9343	JBEAP-10853	Reopen on EAP 7.1.0.ER1. [~jdenise]: Can you please also add check for adding EAP folder itself as module? Example: {noformat} ./jboss-cli.sh --command="module add --name=org.my.example --resources=../" {noformat}	1	1	https://issues.jboss.org/browse/JBEAP-10853
9344	JBEAP-10853	[~mkopecky], copying folder can be a dangerous weapon. EAP folder is a case, but root directory or user home are some others. Do we really want to have a long list of cases? I am not sure.	1	1	https://issues.jboss.org/browse/JBEAP-10853
9345	JBEAP-10826	[~pkremens] I think it's a test issue, IBM JDK needs a different KeyManagerFactory (IbmX509). I need to pull the name from somewhere dynamically or differentiate based on JDK. Could you file a new jira and assign it to me? I'll have a look at it after I'm done with ER1 testing. Thanks! edit: Or reopen this, ok. Thanks.	1	1	https://issues.jboss.org/browse/JBEAP-10826
9346	JBEAP-10796	[~jmesnil] Can you provide the devel-ack? [~zhfeng] What does this upgrade fix/why is it needed?	1	1	https://issues.jboss.org/browse/JBEAP-10796
9347	JBEAP-10759	[~jason.greene], I filed JBEAP-10859 for including the artifact and BOM in the product Maven repo. I don't see any legacy EJB client JARs in the EAP distribution. Should there be a module containing it, or perhaps some kind of {{bin/client/jboss-client-legacy.jar}} with it? Or are users supposed to import it only via Maven?	0	0	https://issues.jboss.org/browse/JBEAP-10759
9348	JBEAP-10723	[~jmesnil] can you check if all issues solved by this upgrade are linked here? I think JBEAP-6606 should be linked here as well. Can you add it please? Thanks.	0	0	https://issues.jboss.org/browse/JBEAP-10723
9349	JBEAP-10621	Martin Svehla asked: {quote} Hey Kabir, Seems that JIRAs that have labels eap7.1.0-defer or eap7.1.0-to-prd, while having dev ack-, cause issues in some of our filters. Can you please make a mass update and set dev ack to '+' or '?' on all JIRAs with that label? {quote} Since this issue is one of these, these labels implicitly mean it is a valid issue, so setting devel-ack to +	1	0	https://issues.jboss.org/browse/JBEAP-10621
9350	JBEAP-10578	[~msvehla] How do we go about converting this to a documentation issue as you mentioned on the call?	1	1	https://issues.jboss.org/browse/JBEAP-10578
9351	JBEAP-10578	[~kanovotn] What is the process for these docs issues? It shows up in the security team's list of blockers, but it has been done really, and is not a dev task	1	1	https://issues.jboss.org/browse/JBEAP-10578
9352	JBEAP-10552	The scope of this task should be more than just aligning the full-ha profile. It should be an overall review of CONSOLE handler use in our config files. 1) In a managed domain, why is the full-ha profile different? Is any domain assumed to be "production" environment, thus deserving the higher perf config setting we use for standalone-full-ha.xml? 2) If we consider there's no good reason for full-ha to be different, are we ok with turning off all server logging via the HC console in EAP? Is that an acceptable change in an EAP minor? 3) If we consider there's no good reason for full-ha to be different in domain.xml, is there a good reason for it to be different in the standalone.xml variants? If not, is turning off console logging in those an acceptable change in an EAP minor? FWIW, I have no real strong opinions on these questions; I just want the task to be to consider the whole picture. My 2 cents (nothing more): I have a *very mild* dislike for changing only full-ha because it just doesn't make a lot of logical sense and that bothers my logical side. And dropping console logging in all configs including standalone variants feels too extreme; like something we'd regret when we realize some unforeseen impact. I'm indifferent between the other permutations, including doing nothing at all.	1	1	https://issues.jboss.org/browse/JBEAP-10552
9353	JBEAP-10527	Has the server knowledge of the AuthenticationContext defined by the client or is this object only stored on the client? On the server-side any authentication and authorization for JMS is performed by https://github.com/wildfly/wildfly/blob/master/messaging-activemq/src/main/java/org/wildfly/extension/messaging/activemq/ElytronSecurityManager.java using the ServerAuthenticationContext.	0	0	https://issues.jboss.org/browse/JBEAP-10527
9354	JBEAP-10527	Is it really blocker? I would think it is not: we were not offering that feature with PicketBox and it was my understanding that providing the same feature set with Elytron in the messaging subsystem was sufficient? From a technical perspective, If we can get the client's AuthenticationContext info on the server-side using ServerAuthenticationContext, that *new feature* can be done for EAP 7.1. Otherwise, it would require changes to Artemis to integrate it with SASL and I don't think it is reasonable to expect it for EAP 7.1	0	0	https://issues.jboss.org/browse/JBEAP-10527
9355	JBEAP-10261	how are you testing this? executing this scripts from cmd or from powershell? powershell "fixes" the path and is correctly set. .bat scripts needs to be executed in cmd.	0	0	https://issues.jboss.org/browse/JBEAP-10261
9356	JBEAP-10061	I updated the community [https://docs.jboss.org/author/display/WFLY/JPA+Reference+Guide] to refer to "Migrating from" the non-Hibernate persistence providers. Is that clear enough to also change in the EAP 7.1 documentation? I also updated the *jboss.as.jpa.classtransformer* description to: {quote} set to false to disable class transformers for the persistence unit. Set to true, to allow entity class enhancing/rewriting. {quote} For the EAP 7.1 description of *jboss.as.jpa.classtransformer*, we can change to: {quote} jboss.as.jpa.classtransformer |Set to `false` to disable entity class transformers for the persistence unit. Set to `true` to enable entity class tranforming, or set one of { hibernate.enhancer.enableDirtyTracking, hibernate.enhancer.enableLazyInitialization, hibernate.enhancer.enableAssociationManagement, hibernate.ejb.use_class_enhancer} properties to `true`. {quote}	0	0	https://issues.jboss.org/browse/JBEAP-10061
9357	JBEAP-10052	[~istraka] Could you please review the topics added here: https://gitlab.cee.redhat.com/nchaudha/eap-documentation/commit/8181db6c4d9f2d467daeb3bca41b7c20ab529940	1	1	https://issues.jboss.org/browse/JBEAP-10052
9358	JBEAP-9901	Hi [~ochaloup], I'm almost ready to start writing the transaction manager content for the Performance Tuning Guide (to be tracked in this sub-task JIRA). [You mentioned in the parent JIRA|https://issues.jboss.org/browse/JBEAP-7233?focusedCommentId=13345191&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13345191] that for XA transactions, using a journal store for logging would be better for performance. Is all that is required to enable the {{use-journal-store}} setting, or is any other configuration necessary? {code} /subsystem=transactions:write-attribute(name=use-journal-store,value=true) {code}	1	1	https://issues.jboss.org/browse/JBEAP-9901

#	Requirement	Questions	Answered	Successfully Answered	URL
9359	JBEAP-9746	Does this mean it should a configuration attribute on the logging subsystem? If so that will not work. It has to be done via a system property because loggers are generally static in WildFly/EAP and they will be created before the attribute can be parsed.	0	0	https://issues.jboss.org/browse/JBEAP-9746
9360	JBEAP-9746	Cannot the configuration be made as a system property in the standalone.xml?	1	1	https://issues.jboss.org/browse/JBEAP-9746
9361	JBEAP-9746	[~Takis] You mentioned "And it is a feature that comes with the next component upgrade to 3.3.1." Will this component upgrade make 7.1.0 GA?	0	0	https://issues.jboss.org/browse/JBEAP-9746
9362	JBEAP-9716	[~trepel]: It seems like someone may have already fixed this as I am not able to reproduce this issue now. Can you give it a shot with the latest source? If it is fixed, I will remove those warnings from the README file.	1	1	https://issues.jboss.org/browse/JBEAP-9716
9363	JBEAP-9716	[~jeffmaury] I can see the warnings even if using https://github.com/jbossas/eap-quickstarts/blob/7.x/kitchensink-jsp/pom.xml#L60 As you can see it uses taglibs-standard-impl and the scope is compile (default scope if nothing specified). I tried adding jsp-api to the pom.xm as you suggested but only one of the two warnings disappeared. I wonder [~sgilda] , could you please try it as well, it's really weird. There are other quickstarts that use jsp (e.g. spring-greeter). It seems that warnings are not shown for these quickstarts because of {noformat} <dependency> <groupId>org.jboss.spec.javax.servlet.jsp</groupId> <artifactId>jboss-jsp-api_2.3_spec</artifactId> </dependency> {noformat} in their pom.xml However when I tried adding this in kitchensink-jsp quickstart, again only one of the two warnings disappeared.	1	1	https://issues.jboss.org/browse/JBEAP-9716
9364	JBEAP-9716	[~ctomc]: This issue is beyond me. I've tried the suggestions, but they don't really work. Is there someone else who could take a look at this?	1	1	https://issues.jboss.org/browse/JBEAP-9716
9365	JBEAP-9712	@[~rjanik], do you have any suggestions for this chapter? [~johara] suggested mentioning simple caches for local caches that don't need transactions or persistence, but it seems the EAP implementation already enables that option automatically (according to JBEAP-1905 and https://github.com/jbossas/jboss-eap7/pull/26/). Is my understanding correct?	1	1	https://issues.jboss.org/browse/JBEAP-9712
9366	JBEAP-9603	What is the setup you are using to test websockets where you see this behavior?	1	1	https://issues.jboss.org/browse/JBEAP-9603
9367	JBEAP-9590	[~dlofthouse], [~jkalina] can you please switch linked JBEAP jiras to proper state before I ack this? Some of them are still in NEW state, should be in PR sent.	1	1	https://issues.jboss.org/browse/JBEAP-9590
9368	JBEAP-9324	[~emmartins]: Are you asking [~smarlow]?	0	0	https://issues.jboss.org/browse/JBEAP-9324
9369	JBEAP-9243	Changing the icon just requires swapping the JFrameIcon found in izpack. I was thinking of just using shadowman, or should it be more wildfly specific?	0	0	https://issues.jboss.org/browse/JBEAP-9243
9370	JBEAP-9193	https://github.com/jbossas/jboss-eap7/pull/1441 is to move to 3.0.0.Beta7 but that's not expected to be the last core upgrade for DR14.	1	1	https://issues.jboss.org/browse/JBEAP-9193
9371	JBEAP-8950	[~dmichael], could you please include the link of the JIRA that was included in the RN?	1	1	https://issues.jboss.org/browse/JBEAP-8950
9372	JBEAP-8944	This issue is fixed, can it be moved to fixed issues?	0	0	https://issues.jboss.org/browse/JBEAP-8944
9373	JBEAP-8941	1. Should the EAP 7.1 release notes refer to [https://docs.jboss.org/author/display/WFLY/WildFly+Elytron+Security#WildFlyElytronSecurity-GetStartedusingtheElytronSubsystem], which contents may change to reflect future WildFly versions? Same question for [https://docs.jboss.org/author/display/WFLY/Using+the+Elytron+Subsystem] + [https://docs.jboss.org/author/display/WFLY/WildFly+Elytron+Security] + [https://docs.jboss.org/author/display/WFLY/WildFly+Elytron+Security#WildFlyElytronSecurity-UsetheElytronandSecuritySubsystemsinParallel]. 2. The link for "Migrate from Hibernate ORM 5.0 to Hibernate ORM 5.1", doesn't work for me, I get a http 404 error. [https://access.redhat.com/documentation/en-us/red_hat_jboss_enterprise_application_platform/7.1.alpha/html-single/migration_guide/#migrating_from_hibernate_5_0_to_5_1]	0	0	https://issues.jboss.org/browse/JBEAP-8941
9374	JBEAP-8941	Yes, the EAP 7.1 Alpha Release Notes intentionally links to upstream Elytron docs, because we had the Elytron writer help draft upstream Elytron content for Alpha use. For EAP 7.1 Beta & GA, the Elytron info in the Release Notes will be modified to link to relevant Elytron documentation in the official EAP product documentation and the Release Notes will no longer point to upstream WildFly docs.	1	1	https://issues.jboss.org/browse/JBEAP-8941
9375	JBEAP-8853	[~martyn-taylor] The Artemis upgrade incorporates [commit|https://github.com/rh-messaging/jboss-activemq-artemis/commit/3d765ae4cf1b122328ce7bfa8f8db671ad758b92] which likely fixes JBEAP-7454 and JBEAP-8523. Is it true? If yes these JIRAs should be linked here and consequently marked as Resolved. [~kabirkhan] it also affects RFEs EAP7-591 and EAP7-660.	1	1	https://issues.jboss.org/browse/JBEAP-8853
9376	JBEAP-8853	[~kabirkhan] There are some reopen issues related to this upgrade. Is there general policy for upgrade jiras which incorporates reopened issues? Should I reopen this jira as well or verify with comment that something is wrong with it? Thanks	1	1	https://issues.jboss.org/browse/JBEAP-8853
9377	JBEAP-8665	[~wolfc] Could you comment on this issue please?	1	0	https://issues.jboss.org/browse/JBEAP-8665
9378	JBEAP-8665	[~wolfc] [~kabirkhan] how can we push this issue forward for review/updates/action?	1	1	https://issues.jboss.org/browse/JBEAP-8665
9379	JBEAP-8659	Martin Svehla asked: {quote} Hey Kabir, Seems that JIRAs that have labels eap7.1.0-defer or eap7.1.0-to-prd, while having dev ack-, cause issues in some of our filters. Can you please make a mass update and set dev ack to '+' or '?' on all JIRAs with that label? {quote} Since this issue is one of these, these labels implicitly mean it is a valid issue, so setting devel-ack to +	1	0	https://issues.jboss.org/browse/JBEAP-8659
9380	JBEAP-8599	Hi, can you please include JDG-708? It's critical for us. It was fixed in 3.0.20.Final in upstream: https://issues.jboss.org/browse/RESTEASY-1205	1	1	https://issues.jboss.org/browse/JBEAP-8599
9381	JBEAP-8599	[~asoldano] hi, i've created and linked the issue. We need the fix for 7.0.5.GA actually, because of JDG release schedule, 7.0.6.GA would be too late. Is it possible to release SP2 in 7.0.5.GA (CR2)? I'll talk to my PM and escalate this if needed.	1	1	https://issues.jboss.org/browse/JBEAP-8599
9382	JBEAP-8593	[~pjelinek], is this a big concern given that the passwords are shown plainly in CLI? Changing it in the given form is trivial but the passwords would still show up in the Management Model browser (as do other passwords).	1	1	https://issues.jboss.org/browse/JBEAP-8593
9383	JBEAP-8532	This requires some changes that I don't want to make this late in the 7.1 cycle - can we move this to 7.2 and I'll take another look then?	1	0	https://issues.jboss.org/browse/JBEAP-8532
9384	JBEAP-8489	[~asoldano]: Can you please set dev_nack and close this jira?	0	0	https://issues.jboss.org/browse/JBEAP-8489
9385	JBEAP-8437	So basically only "name" attribute is check during XML validation and others are skipped and checked by model? Is there a plan to extend it the validation as information which attribute is required is in xsd?	1	1	https://issues.jboss.org/browse/JBEAP-8437
9386	JBEAP-8437	Is there an accepted requirement anywhere that all configuration mistakes are caught via xml parsers? I don't recall agreeing to one, although it's been a lot of years now. :) BTW, my comments are not meant as a rejection of this issue. I just want behavior improvements separately categorized.	1	1	https://issues.jboss.org/browse/JBEAP-8437
9387	JBEAP-8434	Hi [~jpallich]! What do you think about this issue becoming a RFE?	0	0	https://issues.jboss.org/browse/JBEAP-8434
9388	JBEAP-8434	Hi [~bilge]! What do you think about this issue becoming a RFE?	0	0	https://issues.jboss.org/browse/JBEAP-8434

#	Requirement	Questions	Answered	Successfully Answered	URL
9389	JBEAP-8371	[~jamezp]: Thank you for explanation. I see two choices: * add description of this use case to documentation and all similar use cases too. or * logging load all properties at first. Writer will be opened correctly after all properties are loaded. Anything else would be done after all properties are loaded. This is probably not a bug, it is enhancement. Component is also "User experience", because this could be really annoying for users although users could set logging correctly ... WDYT?	1	1	https://issues.jboss.org/browse/JBEAP-8371
9390	JBEAP-8308	[~kabirkhan] [~ataylor] [~clebert] I'm not happy with the "wildfly" in the suffix. Can we please use SP instead of wildfly ?	1	1	https://issues.jboss.org/browse/JBEAP-8308
9391	JBEAP-8287	Hello [~jamezp] I would like to ask you about supported languages. According to EAP7-30 we support zh_CN. On the other hand, 3.2.3. Change the Language of the Management Console [here|https://doc-stage.usersys.redhat.com/documentation/en-us/red_hat_jboss_enterprise_application_platform/7.1.alpha/html-single/configuration_guide/#management_console_overview] is said, that we support zh_Hans in management console. Is EAP7-30 valid?	1	0	https://issues.jboss.org/browse/JBEAP-8287
9392	JBEAP-8287	One more question. These files can be found in eap sources: {code} ./appclient/src/main/resources/org/jboss/as/appclient/logging/AppClientLogger.i18n_fr.properties ./appclient/src/main/resources/org/jboss/as/appclient/logging/AppClientLogger.i18n_zh_Hans.properties ./appclient/src/main/resources/org/jboss/as/appclient/logging/AppClientLogger.i18n_pt_BR.properties ./appclient/src/main/resources/org/jboss/as/appclient/logging/AppClientLogger.i18n_de.properties ./appclient/src/main/resources/org/jboss/as/appclient/logging/AppClientLogger.i18n_es.properties ./appclient/src/main/resources/org/jboss/as/appclient/logging/AppClientLogger.i18n_zh_CN.properties ./appclient/src/main/resources/org/jboss/as/appclient/logging/AppClientLogger.i18n_ja.properties {code} zh_Hans can not be set as language to be used (starting eap with java opts -Duser.language=zh -Duser.country=Hans - nothing is translated). Is there any reason why AppClientLogger.i18n_zh_Hans.properties is there? Note that this is only example. There are many zh_Hans property files.	1	0	https://issues.jboss.org/browse/JBEAP-8287
9393	JBEAP-8280	bq. The instance pool max-pool-size value is the maximum number of EJB instances that can exist in the pool. The default max-pool-size value is 20, which means that for a particular EJB there can only be 20 instances of that EJB at one time. Actually the default is that the pool size is derived from the size of worker pools: {noformat} [domain@localhost:9990 /] /subsystem=ejb3/strict-max-bean-instance-pool=slsb-strict-max-pool:read-resource { "outcome" => "success", "result" => { "derive-size" => "from-worker-pools", "max-pool-size" => 20, "timeout" => 5L, "timeout-unit" => "MINUTES" } } {noformat} If the {{derive-size}} attribute is defined, it overrides the value from {{max-pool-size}}. Please add some discussion about the {{derive-size}} attribute (and both of its possible values - {{from-worker-pools}} and {{from-cpu-count}}) and the possibility to undefine it (in which case the explicit {{max-pool-size}} is applied). This affects the {{5.1.1.1. Creating a Bean Instance Pool}} section where both possibilities should be described. bq. If more than the maximum number of threads try to use a particular EJB, the client will block and wait for an instance to become available. If a client does not get an instance within the time set in instance-acquisition-timeout, an exception is thrown. - I only see {{timeout}} and {{timeout-unit}} attributes, no {{instance-acquisition-timeout}}. Am I missing something? - Also the wording "more than the maximum number of threads" is a bit confusing, because obtaining an instance from the pool depends on whether the instance is free or not, it has nothing to do with number of threads (either someone is using it or not). bq. {{5.1.1.2. Specifying the Instance Pool a Bean Should Use}} - IMO it is worth mentioning that you can do this using the {{@org.jboss.ejb3.annotation.Pool}} annotation as well, not just using the deployment descriptor. - {{By default, EJBs use the default instance pool.}} - that sentence sounds a bit silly and tautological :) and it repeats what was already said at the beginning of {{5.1.1. Bean Instance Pools}} I think we could remove it. bq. Bean instance pools are used for Stateless Session Beans (SLSBs) and Message Driven Beans (MDBs). By default, beans use the default instance pool specified by default-slsb-instance-pool. The {{default-slsb-instance-pool}} is for SLSBs only, for MDBs, it is {{default-mdb-instance-pool}} bq. By default, a bean thread pool named default is used for remote EJB requests, asynchronous EJB calls, and EJB timers. This is no longer true for remote EJB requests, in EAP 7 these are by default handled in the worker defined in the IO subsystem. It is possible to use an EJB thread pool for remote requests, but you have to set the value of {{execute-in-worker}} of the {{/subsystem=ejb3/service=remote}} resource to {{false}}	1	1	https://issues.jboss.org/browse/JBEAP-8280
9394	JBEAP-8280	Verified in revision 018262f. Looks good to me now, thanks!	0	0	https://issues.jboss.org/browse/JBEAP-8280
9395	JBEAP-8054	Verification failed on Performance guide, revision 8288495. ---- {quote}CMS is the default garbage collector in Java 8{quote} Default garbage collector in Java 8 is "Parallel GC", you can check this by -XX:+PrintCommandLineFlags option: \[mkopecky@dhcp-10-40-5-170 bin\]$ java -XX:+PrintCommandLineFlags -version -XX:InitialHeapSize=257824576 -XX:MaxHeapSize=4125193216 -XX:+PrintCommandLineFlags -XX:+UseCompressedClassPointers -XX:+UseCompressedOops -XX:+UseParallelGC ... \[mkopecky@dhcp-10-40-5-170 bin\]$ ---- {quote} To explicitly enable large page ...: -XX:+UseLargePages To explicitly disable large page ...: -XX:-UseLargePages {quote} Large pages should be disabled by different attribute then is used for enabling of them, right? ---- Information about getting/setting large pages information from OS should be in documentation too. It could be enough to add link to [oracle documentation|http://www.oracle.com/technetwork/java/javase/tech/largememory-jsp-137182.html] ---- {quote}These operating system options do the same thing as the large pages JVM option, but are instead applied system-wide. For more information, see the Memory chapter in the Red Hat Enterprise Linux Performance Tuning Guide. {quote} Memory chapter is quite huge. Can you be more specific and add link to some subchapter? ---- There could be mentioned some basic informations about ulimit tool. On the [article on Customer portal|https://access.redhat.com/solutions/61334] is almost nothing about ulimit tool. There should be at least: * how to get soft limit values: ulimit -Sa * how to get hard limit values: ulimit -Ha * how to set some value: for example this: ulimit -Hn 4090 * difference between soft limit and hard limit	1	1	https://issues.jboss.org/browse/JBEAP-8054
9396	JBEAP-7969	[~rhusar]: Should I assign this one to you?	1	1	https://issues.jboss.org/browse/JBEAP-7969
9397	JBEAP-7877	I can't reproduce this. How are you actually configuring the proxy?	0	0	https://issues.jboss.org/browse/JBEAP-7877
9398	JBEAP-7875	[~kabirkhan] [~jmesnil] All the linked jira are in state PR sent or resolved but this upgrade jira is ready for QA. Could you update the state of jiras, please?	1	0	https://issues.jboss.org/browse/JBEAP-7875
9399	JBEAP-7867	I don't think this is a VDX concern. VDX has no knowledge of the namespaces for which parsers are registered. Not for 7.1; we're past the point where we can consider non-trivial enhancements.	0	0	https://issues.jboss.org/browse/JBEAP-7867

#	Requirement	Questions	Answered	Successfully Answered	URL
9400	JBEAP-7796	Hi Gunnar, could you give a few more details about the changes included. You mentioned that this upgrade fixes a regression and two bugs. Are there any new features? Or just the regression fix + bug fixes?	0	0	https://issues.jboss.org/browse/JBEAP-7796
9401	JBEAP-7790	[~rsvoboda] I presume we already test this using the property? Can we add it to the management model at this stage?	0	0	https://issues.jboss.org/browse/JBEAP-7790
9402	JBEAP-7780	[~rsvoboda] This one has been merged, can you ack it?	0	0	https://issues.jboss.org/browse/JBEAP-7780
9403	JBEAP-7762	Martin, I can see this ticket is assigned to me. Please can you be more specific on how the localization team can help you here? Thanks.	1	1	https://issues.jboss.org/browse/JBEAP-7762
9404	JBEAP-7721	[~mstyk] If I understand correctly, the guide should add the following line to the Core and Jms examples found at the end of the thread management chapter[1]: {{myFactory.allowCoreThreadTimeOut(true);}} Is there a CLI equivalent for allow-time-out (or whatever it may be called) like there is for use-global-pools? I'm unable to find the attribute if it exists. Reply when you can. Thank you! [1]https://access.redhat.com/documentation/en/red-hat-jboss-enterprise-application-platform/7.0/single/configuring-messaging/#client-side_thread_management	1	1	https://issues.jboss.org/browse/JBEAP-7721
9405	JBEAP-7720	[~pjelinek] or [~mnovak]: Could you please take a look at this merge request and let me know if I am understanding this one correctly?	1	1	https://issues.jboss.org/browse/JBEAP-7720
9406	JBEAP-7719	[~jpallich]: I spoke with [~rhn-support-pfestoso] and we both think this JIRA can be closed. Can you confirm?	1	1	https://issues.jboss.org/browse/JBEAP-7719
9407	JBEAP-7719	[~nziakova] or [~dudae]: Since this is now targeted for 7.2, should it still be a blocking issue?	1	1	https://issues.jboss.org/browse/JBEAP-7719
9408	JBEAP-7705	[~rsvoboda] since the general undertow+elytron use case is covered here: https://issues.jboss.org/browse/JBEAP-7228 are there any other usecases/scenarios I'm missing? If not can we close this JIRA?	1	1	https://issues.jboss.org/browse/JBEAP-7705
9409	JBEAP-7698	Verification failed in revision 864568d5. Chapter CHAPTER 2. HOW RED HAT JBOSS ENTERPRISE APPLICATION PLATFORM 7 HANDLES SECURITY OUT OF THE BOX looks good to me. There are only some 'major' issues - JBEAP-11399, JBEAP-11410, JBEAP-11411, JBEAP-11412. What needs to be improved is following - related Analysis document says: _From a documentation perspective we should add commands to show how to convert the default resources to use Elytron security, We should also show commands to remove the legacy security once it is no longer referenced._ There is JBEAP-8472 which will provide CLI script {{enable-elytron.cli}} in examples directory which will enable Elytron. By mention this script in documentation will be covered mentioned above requirement (otherwise there will be needed documentation for how to enable every particular resource which is integrated with Elytron). Please add information about this script once JBEAP-8472 will be resolved - there should be mentioned: * where this CLI script can be found * some basic information that it can be used as example for enabling Elytron in related subsystems and management interfaces	0	0	https://issues.jboss.org/browse/JBEAP-7698
9410	JBEAP-7694	Failing verification. What is missing: * common place with set of included mechanisms ** Anonymous ** Digest ** Entity ** External ** Gs2 ** Gssapi ** LocalUser ** OAuth2 ** OTP ** Plain ** Scram * information how and when to use each of them * list of parameters which can be used for the mechanisms configuration and their description: ** wildfly.sasl.local-user.quiet-auth / jboss.sasl.local-user.quiet-auth ** wildfly.sasl.gssapi.client.delegate-credential ** wildfly.sasl.otp.alternate-dictionary ** ... * how to configure set of mechanisms on server side (filtering, ordering, parameters etc). * how to configure set of mechanisms on client side (filtering, ordering, parameters etc).	0	0	https://issues.jboss.org/browse/JBEAP-7694
9411	JBEAP-7690	[~ahoffer] It needs to be listed as a newly supported feature and I think it also needs some documentation on the configuration of distributed workmanager. That should go in chapter 16 (JCA) of Configuration Guide. In addition to the usual stuff (what is a distributed workmanager, what attributes does it take, etc.) there are some specifics that I think need to be documented, such as that the thread pool for a distributed workmanager needs to be of the same name as the workmanager. [~maeste] I see the Jirban Community Docs flag on EAP7-495 has been set to Done/Not Relevant, but couldn't find the documentation in the WildFly docs *. Could you point [~ahoffer] to the docs, or provide some, if there are none? Thanks. * https://docs.jboss.org/author/display/WFLY10/Documentation	0	0	https://issues.jboss.org/browse/JBEAP-7690
9412	JBEAP-7689	[~nziakova]: I am a little confused by two of your comments as they contradict each other: 1. In the first bullet, you say: "The artifacts are not downloaded from JBoss EAP Maven repository, but from Red Hat Maven repository..." 2. In the last bullet you say it "Should be the JBoss EAP Maven repository" Should it be the "Red Hat Maven repository" or the "JBoss EAP Maven repository"? We refer to it extensively throughout our documentation as the "JBoss EAP Maven repository". It only contains artifacts for JBoss EAP, so it might be better to keep referring to it that way. Also, the downloaded version of the repository is branded with the JBoss EAP name. Your thoughts?	1	1	https://issues.jboss.org/browse/JBEAP-7689
9413	JBEAP-7679	[~jpallich]: We do not document Shrinkwrap resolvers. Is this JIRA to update the quickstarts that use it? If so, the component should be changed to "Quickstarts".	0	0	https://issues.jboss.org/browse/JBEAP-7679
9414	JBEAP-7672	[~jpallich]: We do not document AngularJS. Is this JIRA to update the quickstarts that use it? If so, the component should be changed to "Quickstarts".	0	0	https://issues.jboss.org/browse/JBEAP-7672
9415	JBEAP-7671	[~jpallich]: We do not document JQuery or JQuery Mobile. Is this JIRA to update the quickstarts that use it? If so, the component should be changed to "Quickstarts".	0	0	https://issues.jboss.org/browse/JBEAP-7671
9416	JBEAP-7670	[~jpallich]: We do not document Arquillian. Is this JIRA to update the quickstarts that use Arquillian? If so, the component should be changed to "Quickstarts".	0	0	https://issues.jboss.org/browse/JBEAP-7670
9417	JBEAP-7669	[~jpallich]: Should this component be changed to "Quickstarts"?	0	0	https://issues.jboss.org/browse/JBEAP-7669
9418	JBEAP-7665	- Please add a note about the {{org.wildfly.extension.batch.jberet.deployment.BatchPermission}} which is required to allow authenticated users deal with Batch jobs. The supported targets of this permissions are {{start}} for starting jobs, {{stop}} for stopping jobs, {{restart}} for restarting jobs, {{abandon}} for marking executions as ABANDONED, {{read}} for reading metadata about job executions using the JobOperator and management interfaces. bq. For example, a secured RESTful endpoint is created to initiate batch jobs and the batch-jberet subsystem. - This sentence is not very senseful, did you mean for example "using" instead of "and" the batch-jberet subsystem? bq. By default, the batch-jberet subsystem is configured to use the ApplicationDomain Elytron security domain. - I think that at least in 7.1.0.DR11 the security domain is undefined by default. [~jamezp] are we planning to keep it this way?	1	1	https://issues.jboss.org/browse/JBEAP-7665
9419	JBEAP-7615	Can you confirm that this is not default in EAP7? Can you point me to what system property it is? If above is true, maybe it would make sense to extend that property or add a new one to incorporate removing the file content entirely? Current approach doesn't fit well in Docker images, these are tested afterwards and we release only working images, so no rollback would be required.	1	0	https://issues.jboss.org/browse/JBEAP-7615

#	Requirement	Questions	Answered	Successfully Answered	URL
9420	JBEAP-7615	We are talking here about different things. We use the {{ageout-history}} command. It removes all the unnecessary, previous patches, which let us go down from 1.4 GB to around 280 MB. But there are *still left scrapped jar files* that were updated while patching under {{modules/system/layers/base}} directory. Example? {code} [root@d12b18135037 ~]# jar tf /opt/eap/modules/system/layers/base/org/jboss/log4j/logmanager/main/log4j-jboss-logmanager-1.1.1.Final-redhat-1.jar Picked up JAVA_TOOL_OPTIONS: -Duser.home=/home/jboss -Duser.name=jboss java.util.zip.ZipException: error in opening zip file at java.util.zip.ZipFile.open(Native Method) at java.util.zip.ZipFile.<init>(ZipFile.java:219) at java.util.zip.ZipFile.<init>(ZipFile.java:149) at java.util.zip.ZipFile.<init>(ZipFile.java:120) at sun.tools.jar.Main.list(Main.java:1115) at sun.tools.jar.Main.run(Main.java:293) at sun.tools.jar.Main.main(Main.java:1288) {code} When we remove these unusable jar files we save over 100 MB with 6.4.12.GA CP.	1	0	https://issues.jboss.org/browse/JBEAP-7615
9421	JBEAP-7470	This will be done in a later release, as it requires a non compatible Undertow API change that is not suitable for the 1.4.x branch.	0	0	https://issues.jboss.org/browse/JBEAP-7470
9422	JBEAP-7469	[~jdenise]: When do you expect to merge code for this jira? Before EAP 7.1.0.DR11?	1	0	https://issues.jboss.org/browse/JBEAP-7469
9423	JBEAP-7469	@Marek, although we seem to have a stable integration, I am wandering if we should really rush to integrate in in 7.1 We could wait and only push it upstream when we have identified a branch for 7.1 So should come in 7.2 with part of CLI.next.	1	1	https://issues.jboss.org/browse/JBEAP-7469
9424	JBEAP-7469	[~jdenise]: Well, I prefer to wait and push it to upstream (I mean wf-core, not wf-core-eap) when we have identified a branch for 7.1. Community can fix some potential initial issues ... When will you have Aesh-readline ready? DR10? Do you run WildFly TS with Aesh.readline? (I mean WF, not WF-core) So you plan to add CLI.next to EAP 7.2, right? Btw, I'm unable to build you branch locally: {noformat} [ERROR] Failed to execute goal org.apache.maven.plugins:maven-compiler-plugin:3.5.1:testCompile (default-testCompile) on project wildfly-core-testsuite-standalone: Compilation failure: Compilation failure: [ERROR] /home/mkopecky/git/jbossqe-eap-testsuite-eap7-core/testsuite/standalone/src/test/java/org/jboss/as/test/integration/management/cli/FileInputTestCase.java:[24,27] package com.sun.javafx.util does not exist [ERROR] /home/mkopecky/git/jbossqe-eap-testsuite-eap7-core/testsuite/standalone/src/test/java/org/jboss/as/test/integration/management/cli/FileInputTestCase.java:[61,82] cannot find symbol {noformat}	1	0	https://issues.jboss.org/browse/JBEAP-7469
9425	JBEAP-7469	[~jdenise]: What do you plan with this jira? Can we change "Target Release" to 7.2.0.GA? I don't have permission to change it, but I can ask Jiri Pallich. Jiri has those permissions.	1	0	https://issues.jboss.org/browse/JBEAP-7469
9426	JBEAP-7469	[~jpallich]: Can you please change target release of this jira to 7.2.0.GA?	0	0	https://issues.jboss.org/browse/JBEAP-7469
9427	JBEAP-7438	Verified on EAP 7.1.0 DR9 tag https://github.com/jbossas/jboss-eap7/blob/EAP_7.1.0.DR9-dev/webservices/tests-integration/pom.xml#L138 [~kabirkhan] [~asoldano], is this temporary thing till netty & artemis are synced or this will stay in for safety reasons in the future?	0	0	https://issues.jboss.org/browse/JBEAP-7438
9428	JBEAP-7362	[~soul2zimate]: When do you plan to fix this issue?	1	1	https://issues.jboss.org/browse/JBEAP-7362
9429	JBEAP-7362	Thank you. I updated status of this jira. [~soul2zimate]: Can you please set dev_ack?	0	0	https://issues.jboss.org/browse/JBEAP-7362
9430	JBEAP-7235	The originating JIRA is closed and I believe this one, which also targets 7.0,is complete. [~jpallich]: Can we close this issue also?	0	0	https://issues.jboss.org/browse/JBEAP-7235
9431	JBEAP-7233	Thanks a lot for the suggestions, [~ochaloup] and [~mkopecky]! @[~mkopecky]: I am writing the JVM chapter at the moment (tracked in sub-task jira JBEAP-8054), and I will be sure to mention ulimits as suggested. @[~ochaloup]: I have made a note in the planning doc to cover the transaction subsystem log store. When I get up to writing that section, I will probably contact you for some more details.	1	1	https://issues.jboss.org/browse/JBEAP-7233
9432	JBEAP-7228	Hey [~rhn-support-pnag], I have added comments to the MR. When do you expect you will add documentation for the other areas of this RFE?	0	0	https://issues.jboss.org/browse/JBEAP-7228
9433	JBEAP-7224	[~jpallich]: This seems to be a duplicate of https://issues.jboss.org/browse/JBEAP-1493. Or am I missing something here?	0	0	https://issues.jboss.org/browse/JBEAP-7224
9434	JBEAP-7221	Verification failed in revision n_2419454_how-to-configure-server-security_version_7.1.Alpha_edition_1.0_release_0-revision_8279980 (docs-bulk-12) The documentation for this RFE is insufficient. I'm missing a couple of important pieces of information: * How is it possible to create a custom manager * Where can be custom trust manager implementations referred from (e.g. server modules, in deployments, ...) * What must be changed in server's configuration in order to use a custom trust manager - including CLI commands with a brief reason and description * Explicitly say that usage of a custom trust manager will not be possible when running a JVM in FIPS mode (+ describe a way how to achieve similar behaviour as stated in analysis doc [1]) * An example of use case in which a custom manager could be used [1] https://docs.google.com/a/redhat.com/document/d/1oMPYgHgmYNG0psiQDx8p3D_A5hu5IZ_NCy3UMyR9cms/edit?usp=sharing	1	1	https://issues.jboss.org/browse/JBEAP-7221
9435	JBEAP-7218	> So for #3, is there any way that starting a job through the console would actually work? Yes, I use some batch examples from here https://github.com/javaee-samples/javaee7-samples/tree/master/batch Their job applications doesn't requires parameters. > Or perhaps should there be a way to allow the user to provide the necessary parameters when starting the job instead of just failing every time? Yes, it is possible, it would require some changed in the UI to allow the user to specify the parameters, but I would like to deter this implementation to a later release as the time is tight for the EAP 7.1. > Could you also verify the wording in my merge request to make sure it is accurate for the console behavior? I think the text reflects the actual UI, nice job.	1	1	https://issues.jboss.org/browse/JBEAP-7218
9436	JBEAP-7212	This should probably be treated as a review of all pom.xml examples related to EJB clients and suggesting this particular strategy for handling dependencies - using the {{wildfly-ejb-client-bom}} for versions and then just adding the {{jboss-ejb-client}} library. {noformat} <dependencyManagement> <dependencies> <dependency> <groupId>org.jboss.eap</groupId> <artifactId>wildfly-ejb-client-bom</artifactId> <version>EJB_CLIENT_BOM_VERSION</version> <type>pom</type> <scope>import</scope> </dependency> </dependencies> </dependencyManagement> <dependencies> <dependency> <groupId>org.jboss</groupId> <artifactId>jboss-ejb-client</artifactId> </dependency> </dependencies> {noformat} We should also mention where to find the right EJB_CLIENT_BOM_VERSION - probably tell the users to take a look into the Maven repo, which version is present there? Any better idea?	1	1	https://issues.jboss.org/browse/JBEAP-7212

#	Requirement	Questions	Answered	Successfully Answered	URL
9437	JBEAP-7212	I would say that the chapter {{Project Dependencies for Remote EJB Clients}} as a whole needs a bit of restructuralization, it doesn't make much sense anymore. The pom.xml examples in the new chapter {{Single artifactID for jboss-ejb-client Dependencies}} basically just introduces a new way of achieving the same result as the two pom.xml snippets above (in {{Maven Dependencies for Remote EJB Clients}}) - the result is the same. But in the docs, these two sections are too separated and it looks like they are something different. Can we add an explanation that these are two different ways to achieve the same? Also, this list of added dependencies {quote} jboss-ejb-client xnio-nio jboss-marshalling-river jboss-sasl jboss-ejb-api_3.1_spec jboss-transaction-api_1.1_spec {quote} is not complete, and probably not necessary at all anyway. Do we really need to list all transitive dependencies? When the developer uses one of the two ways to declare an EJB client dependency, these transitive dependencies will be automatically added, it's not really necessary to list them in the documentation, IMO. Btw, please note that I also filed related JBEAP-11505 which is to remove a similar list of transitive dependencies from the beginning of the {{Maven Dependencies for Remote EJB Clients}} section - for the same reasoning.	0	0	https://issues.jboss.org/browse/JBEAP-7212
9438	JBEAP-7192	Hi [~dmichael]! It currently looks like there will no documentation apart RN. Do you plan to add rest of it to relevant chapters in Configuration Guide?	0	0	https://issues.jboss.org/browse/JBEAP-7192
9439	JBEAP-7178	reopening in n_2419413_security-architecture_version_7.1.Alpha_edition_1.0_release_0-revision_8279994 with these objections: * "Elytron uses SASL to implement HTTP authentication" - This is new information for me. Could you double ches it is true? Does this somehow impacts user? As I didn't realized that until now, I think this is internal implementation detail, isn't it? Is it worth mentioning then? * "...provides a complete set of HTTP authentication mechanisms including BASIC, FORM, DIGEST, and CLIENT-CERT." I am missing SPNEGO in the list * Document how to setup custom http authentication mechanism.	0	0	https://issues.jboss.org/browse/JBEAP-7178
9440	JBEAP-7178	[~zrhoads] [~rhn-support-pnag], we have agreed with Zach configuring custom component wil be tracked separately. Does this issue already exists? Could you link it? Or it should be created?	0	0	https://issues.jboss.org/browse/JBEAP-7178
9441	JBEAP-7174	@[~hsvabek], I think that details on how to create a custom implementation in the Development Guide is a different enough use case and should be tracked in a new JIRA. I don't think it should hold up verification of this one. What do you think?	1	1	https://issues.jboss.org/browse/JBEAP-7174
9442	JBEAP-7173	Verification failed with How to configure server security revision 6bff93b8. *3.3. Credential Stores* * IMPORTANT paragraph There is change in behaviour if you fill both - credential-reference and password attribute. Here is jira issue https://issues.jboss.org/browse/JBEAP-8764 According this jira issue please remove "IMPORTANT" paragraph. It is no longer necessary. *3.3.2. Supply a Master Password to a Credential Store* * Example Credential Store Created with a Password from Another Credential Store In example is "alias==master-pw" which contain double =, please remove one. *MISSING* * There is missing how to explicitly define type of Credential store Some explanation of /subsystem=elytron/credential-store=credStore:add(type= I can see a lot of texts in MR which is out of scope this jira issue. Can you give me a link to given issue?	1	1	https://issues.jboss.org/browse/JBEAP-7173
9443	JBEAP-7173	Documentation should give answer for these questions: - What is "uri" attribute and what it is used for? -- what are possible parameters in uri? -- and give their description - what is happened if I replace storage file on filesystem when EAP server runs? - Credential store storage file can be removed after server starts up, passwords stay loaded in memory. -- what happens if you add/remove/update new alias in this credential store -- what happens if you lost permission to write folder - When is loaded credential store to memory from file? start, reload, restart, creation of new credential store, ... - What is relation between storage file and credential store loaded to memory? - How to define absolute path to credential store storage file? - Table A.27. credential-store Attributes - wrong URI example, there isn't "vault://" -- keep on mind that URI parameter can be different for default implementation and custom credential store implementation	1	1	https://issues.jboss.org/browse/JBEAP-7173
9444	JBEAP-7173	Hi [~hsvabek], I'll be taking over docs for this feature. I have a few questions: * I am trying to go through the process of creating and using a credential store, but the commands currently documented don't work. I assume the implementation of this feature has changed? For example: {code} /subsystem=elytron/credential-store=exampleCS:add(\ > uri="cr-store://exampleCS?create=true", \ > credential-reference={clear-text=cs-secret} \ >) 'uri' is not found among the supported properties: [credential-reference, relative-to, case-sensitive, create, implementation-properties, location, modifiable, other-providers, provider-name, providers, type] {code} In looking at the management model in the CLI, it looks like there is now {{location}}, {{relative-to}}, and {{create}} attributes for the adding a new store. It looks like something like the following works for creating a new store (and having it stored in the {{data}} directory of the server): {code} /subsystem=elytron/credential-store=test_store01:add(location="cred_stores/test_store01.jceks", relative-to=jboss.server.data.dir, credential-reference={clear-text=pass123},create=true) {code} Can you confirm if this is correct? I assume the URI of {{cr-store://}} etc isn't required any more. * I think we should provide some advice on why a user would use a credential store vs using a password vault. Am I correct in assuming that we would recommend that users should prefer using a credential store because it is more easier to use and reference inside EAP (e.g. using the management CLI), rather than managing a vault externally using {{vault.sh}}? * For your suggestions in your previous comment re memory, permissions, etc: I'm not sure that we need to go in to that much detail on how the memory loading of the store file works. I think an admonition like the following would be sufficient? {quote}IMPORTANT: JBoss EAP reads the credential store file into memory and writes changes to it at varying times. You must ensure that the user running the JBoss EAP process has permissions to the store file, and that you do not externally modify or remove the store file while JBoss EAP is running.{quote}	1	1	https://issues.jboss.org/browse/JBEAP-7173
9445	JBEAP-7173	After doing a bit of googling, I've found these great blog posts by [~pskopek]: http://undertheelytron.blogspot.com.au/2017/04/basics-credential-store-in-wildfly.html http://undertheelytron.blogspot.com.au/2017/05/credential-references.html I think his example of using a credential reference for a datasource is a much better practical example of using a credential from a credential store (rather than using one store to create another credential store). @[~pskopek]: Do you mind if I take some content from your blog posts and adapt it for the EAP docs?	1	1	https://issues.jboss.org/browse/JBEAP-7173

#	Requirement	Questions	Answered	Successfully Answered	URL
9446	JBEAP-7173	Hi [~lcosti.redhat], You are right, 'uri' attribute was removed as this jira issue JBEAP-8450. Besides new attributes 'location', 'relative-to' and 'create', there are also 'modifiable' and case-sensitive attributes (maybe will be renamed, check JBEAP-8871 for details). And of course your example how to create credential store is correct. It would be nice go through whole documentation and update parts where is 'uri' attribute is used. Why prefer using credential store rather than vaults? One reason is easier usage, but it is not main reason for it. Main reason is better security how to get access to secure storage. Please [~pskopek] give here detailed answer, thanks. About loading storage to memory - there is reload() operation for credential store, to reload it from filesystem to memory. Now it doesn't work good JBEAP-10650. Your IMPORTANT paragraph looks good, please add there information about reload, keep on mind how reload works: *Scenario 1* # you replace storage file for credential store on filesystem # process reload() # you have in memory new content of storage file *Scenario 2* # you replace storage file for credential store on filesystem # through CLI you make remove/add operation for aliases -> memory is persisted to filesystem and your replaced storage file from first step is replaced. # process reload() # you have in memory same content as before reload() You must be very careful with reload() operation.	1	1	https://issues.jboss.org/browse/JBEAP-7173
9447	JBEAP-7168	Hi [~hsvabek], like the other credential store docs issues, I'll be taking over this one too. :) [~hsvabek] [~pskopek]: I'm a bit confused on the use cases for this feature. I understand that it is for Elytron Client, but is it expected that credential stores will only be available to clients deployed on EAP, [as this existing doc implies|https://access.qa.redhat.com/documentation/en-us/red_hat_jboss_enterprise_application_platform/7.1.beta/html-single/how_to_configure_identity_management/#using_elytron_client_deployed]? Or, are clients that are not deployed on EAP expected to be able to create and use credential stores? If this is the case, how is it supported for users to create a credential store? Would using the elytron tool be supported? I can't find many details on how to refer to created stores for Elytron Client. I've found the [test cases here|https://github.com/wildfly-security/wildfly-elytron/blob/master/src/test/java/org/wildfly/security/auth/client/XmlConfigurationTest.java#L120], but while it looks like it uses custom-config.xml, I'm not sure about the EAP implementation of this. Those tests looks like it creates a credential store, but I can't see how it adds credentials to the store? Some assistance in helping me understand how this works would be very much appreciated. :)	1	1	https://issues.jboss.org/browse/JBEAP-7168
9448	JBEAP-7167	[~zrhoads], could you provide correct link? Custom component -- OK, I will check {{custom-realm}} documentation when available.	1	1	https://issues.jboss.org/browse/JBEAP-7167
9449	JBEAP-7167	Hi [~okotek], I'll be taking over this docs issue. I'm not sure how to document this change. As I understand it (please correct me if I am wrong), if a user's roles is changed in LDAP (or other modifiable realm), that change will not be propagated automatically to a relevant Elytron caching-realm. What should we advise a customer to do? Should we advise them that on every LDAP update that affects users' roles, that they should be clearing affected Elytron caching-realms?	1	1	https://issues.jboss.org/browse/JBEAP-7167
9450	JBEAP-7165	Verification failed in How to Configure Server Security in revision 6bff93b8. This documentation Jira should provide introduction to the new feature described in Analysis document in Jira description. Provided documentation misses some important information about configuration and understanding of this new feature. What is missing or incorrect: * Section 3.1.4. Mapping Identity for Authenticated Management Users includes sentence "Once you have added the identity resource, the identity of an authenticated user will appear when accessing the management interfaces.", but it is incorrect. In seems that in case when no access=identity is configured then identity from security domain used for authentication appears. * One of requirement for this RFE mentioned in Analysis documents says: "It will be possible to configure the SecurityDomain to be used by domain management to obtain the current SecurityIdentity which will be used for the invocation of the current management request." This requirement should be reflect in documentation. * Analysis document includes table which describes which identity is used in some cases (note, that scenarios 5 a 6 in table are out of the scope of that RFE and should not be documented). This table can be used for improving documentation of this RFE. * Documentation should provide answers for following questions: ** What is and how work inflow identity for Management security? (i.e. one security domain is used for authentication and results identity is inflowed into security domain used in access=identity) ** Provide some information about trusted-security-domains attribute from security domain - when and why it is needed. *** What happens if security domain from access=identity does not trust security domain from authentication. *** What happens if security domain from access=identity does not trust security domain from authentication, but both of them using the same security realm. *** There should be mentioned that trusted-security-domains is not transitive. ** How are mapped roles/groups to the final identity if access=identity is used (i.e. which domain maps it). In case when some information are needed for documenting mentioned above notes please ask developers. I also reported some issues related to this Jira - I linked them to this Jira.	0	0	https://issues.jboss.org/browse/JBEAP-7165
9451	JBEAP-7153	[~zrhoads]: Should this be moved to "Ready for QA"?	1	1	https://issues.jboss.org/browse/JBEAP-7153
9452	JBEAP-7152	[~ahoffer] Opps, sorry pasted the wrong link. Yeah, I meant the EAP7.0 release notes. I was talking about the community docs yeah for EAP 7.0, but I thought you had done some EAP docs for that at some point (I may not be remembering correctly.) WRT how we talk about it, the ignore-unused-configuration is generally used on the thing being managed, so in this case the slaves (The domain controller generally doesn't want to ignore anything, as it has to be able to distribute the entire domain model and deployments etc to a slave that is not ignoring or running --backup etc). FWIW, we mentioned this in the 7.1 alpha notes here: https://access.redhat.com/articles/2662761 (see: Slave Host Controllers Ignore Irrelevant Parts of the Domain Configuration and Content Repository). I'll double check what was available and what the default was in 7.0 vs 7.1 and post an update here. I don't remember exactly when the change went in.	0	0	https://issues.jboss.org/browse/JBEAP-7152
9453	JBEAP-7150	Hi [~pjelinek], the doc work to use the new subsystem is tracked and already done in this JIRA https://issues.jboss.org/browse/JBEAP-10180. Could you move this JIRA back to verified? Thanks!	1	1	https://issues.jboss.org/browse/JBEAP-7150

#	Requirement	Questions	Answered	Successfully Answered	URL
9454	JBEAP-7149	Hi [~zrhoads], thank you for the notification. I've briefly read that section and here are my notes: # In the first paragraph of {{17.11. CONFIGURING HTTP/2}}: ** There is link to HTTP2 RFC draft. Although that link is automatically redirected to the final specification I believe we should provide direct link to final spec instead: {{https://datatracker.ietf.org/doc/rfc7540/}} ** Regarding to the link to the SPDY protocol in the very same paragraph, maybe we might prefer to actually provide link to IETF RFC draft {{https://datatracker.ietf.org/doc/draft-mbelshe-httpbis-spdy/}} instead of the current {{www.chromium.org/spdy}} link? # In the {{WARNING}} section, there is {quote} HTTP/2 requires a TLS stack that supports ALPN... {quote} I would emphasize that this is true only when using HTTP2 over a secured connection (TLS). When you use HTTP/2 in cleartext, the TLS is not utilized at all. Thus I propose something like: {quote} When using HTTP/2 over secured TLS connection the TLS stack that supports ALPN TLS protocol extension is required as it is not provided by the default installation of Java 8... {quote} Also I should note that using HTTP2 over secured connection is probably 99% of the time it will be actually used as all modern browsers don't support HTTP2 in cleartext. # Section {{17.11.1. Configuring Undertow to use HTTP/2}}: ** Well, I think that all necessary information is already included in the description but I feel uncomfortable with the structure a little bit. Basically all you need to do to enable HTTP2 is to set 'enable-http2' attribute on particular http(s)-listener. In case of https-listener the HTTPS setup must be prepared also. Maybe we can separate the steps for the http-listener (with some warning note that most modern browsers enforce HTTP/2 over secured TLS connection - 'h2' - and does not support HTTP/2 in clear text - 'h2c') and https-listener. # I think that we should also provide information how to setup and use OpenSSL libs downloaded from JBCS in EAP somewhere in the documentation. This configuration might slightly differ depending on whether EAP is installed from .zip archive vs RPM packages. Should you have any additional questions or something is not clear, just let me know.	1	1	https://issues.jboss.org/browse/JBEAP-7149
9455	JBEAP-7149	Hi [~rhn-support-pnag], according to the latest discussion, we should also explicitelly mention that there is actually also some ALPN implementation put into the EAP itself (although we still prefer our customers to use OpenSSL based ALPN implementation instead). Thus I propose following - in the Configuration guide, section {{17.11. CONFIGURING HTTP/2}}, there is a following paragraph in the warning column: {quote} When using HTTP/2 over a secured TLS connection, a TLS stack that supports ALPN TLS protocol extension is required as it is not provided by the default installation of Java 8. To get ALPN TLS protocol extension support, utilize the OpenSSL provider in JBoss EAP 7.1 with OpenSSL libraries that implement ALPN capability. The supported OpenSSL libraries are the ones from JBoss Core Services, installed and configured. {quote} Proposed updated paragraph text (feel free to alternate/modify): {quote} When using HTTP/2 over a secured TLS connection, a TLS stack that supports ALPN TLS protocol extension is required. As ALPN is not available in Java 8, there was introduced implementation of it right into the JBoss EAP 7.1, which depends on Java internals. As such it works only with Oracle and OpenJDK (and doesn't work with IBM Java) and there is a risk (even though it is rather low) that it can break with some Java update. Because of this risk it is strongly recommended to utilize ALPN TLS protocol extension support from the OpenSSL provider in JBoss EAP 7.1 with OpenSSL libraries that implement ALPN capability. The supported OpenSSL libraries are the ones from JBoss Core Services, installed and configured. {quote} Also I thought that maybe we should reorder some paragraphs in the warning column so they are more logically ordered: # The paragraph starting with sentence '{{JBoss EAP will automatically try to search for the OpenSSL libraries on the system and use them.}}' should be moved right after the example of how to configure {{openssl}} provider via legacy security-based configuration. # The paragraph '{{If OpenSSL is loaded properly, you will see a message in the server.log during JBoss EAP startup, similar to}}' inluding the log message example should be moved right after the previously mentioned one - that means after the '{{JBoss EAP will automatically try to search for the OpenSSL libraries on the system and use them.}}' I think that this change should make information more clear regarding to the order of the steps customer would probably do. What do you think?	0	0	https://issues.jboss.org/browse/JBEAP-7149
9456	JBEAP-7146	* bq. If your application contains a persistence.xml file or the code uses the @PersistenceContext or @PersistenceUnit annotations, JBoss EAP 7 detects this during deployment and assumes the application uses JPA. It implicitly adds the Hibernate ORM 5 libraries plus a few other dependencies to your application class path and defaults to using these libraries. ** This has worked this way in EAP 6.x as well, I think. Is this really relevant for the migration guide then? * The {{Oracle Byte Array Mapping Changes}} speaks about the {{hibernate.dialect.oracle.prefer_longvarbinary}} property, but the property is correctly named {{hibernate.dialect.oracle.prefer_long_raw}} - yes, it's wrong in the official migration guide document, I'll report that. FYI this is where it is defined: https://github.com/hibernate/hibernate-orm/commit/177bbfe5e371796a5af523e0d1a70839ff0bdb71#diff-1e7261dc64d13954f60f7b0c87758dabR28 * Also, are we planning to include information about new features in Hibernate 5.1.x in addition to just the migration related stuff? I mean, the stuff which is not covered by RFEs, that means without Envers to-one associations	1	1	https://issues.jboss.org/browse/JBEAP-7146
9457	JBEAP-7140	Is this documentation for the Configuration Guide, the Migration Guide, or both? From what I can see, it seems like it might belong here: https://access.redhat.com/documentation/en-us/red_hat_jboss_enterprise_application_platform/7.1.Beta/html-single/configuration_guide/#manage_a_transaction If so, should the new operation be mentioned in the Migration Guide as well?	0	0	https://issues.jboss.org/browse/JBEAP-7140
9458	JBEAP-7139	What is the point of this issue?	0	0	https://issues.jboss.org/browse/JBEAP-7139
9459	JBEAP-7130	Hi [~rhn-support-pfestoso]! Is this a duplicate of JBEAP-7718? If so, can you please close it?	1	1	https://issues.jboss.org/browse/JBEAP-7130
9460	JBEAP-7044	[~brian.stansberry]: I agree with keeping type. But what about to rename "input-stream-index" to "input"? This would be still correct (for both "tooling" and "direct input by users"), but it would be also intuitive for customers. We can still keep "input-stream-index" (as something like alias) for backward compatibility and make it deprecated. And I really recommend to update description of this attribute.	1	0	https://issues.jboss.org/browse/JBEAP-7044

#	Requirement	Questions	Answered	Successfully Answered	URL
9461	JBEAP-7044	[~mjurc][~mkopecky] My apologies; in my comments above I think I didn't understand properly your perspective when you filed this. I forgot the CLI added a trick where it would check some metadata on the param and based on that let the user provide a file name as the attribute value. That's useful but it does have the impact of meaning the r-r-d output, which is a server management API description and is not directly meant as CLI help text, is made inconsistent with what the CLI user experiences. I'm not willing to change the server side management API description away from being a description of the server side management API and turn it into CLI help text. We need a proper server-side management API description. But still, we should try to improve. Some things to do: 1) Look into the param names and descriptions as you suggest here. That should be the focus of this particular JIRA. Text descriptions we of course have some flexibility, and param names we can tweak. (Note that in some cases the param is already in a .Final API and a change means a new param not a rename.) I personally prefer 'input-stream' to 'input' as it better hints at what is going on and is a better description of the server management API. But the '-index' part isn't really necessary. [~ehugonnet] can you have a look at this? Note that we should not change the type from INT, as that would be changing the description from a server API description to a CLI description. 2) [~jfdenise] Perhaps when the CLI feature you've worked on that generates CLI help text from r-r-d output comes along, it could factor this scenario into what is generated? Add some text or something to the CLI help output that explains what the CLI allows with the input? 3) CLI high level commands. One of the main points of CLI high level commands is to hide the less friendly aspects of the server side API from end users. So any place where we are using this CLI input trick is a candidate for a high level command. (We have high level commands around deployment and deployment overlays already.)	0	0	https://issues.jboss.org/browse/JBEAP-7044
9462	JBEAP-7007	[~psakar][~kabirkhan] This upgrade happened in WildFly Core 3.0.0.Alpha13, so it should be linked to JBEAP-7028, not JBEAP-7572, and moved to On QA, Fix Version DR9. I'll make a bold decision and decide I'm right and do all that now. :)	1	1	https://issues.jboss.org/browse/JBEAP-7007
9463	JBEAP-7007	[~jdenise]: Can you please close this as "Out of date"?	0	0	https://issues.jboss.org/browse/JBEAP-7007
9464	JBEAP-6868	[~asoldano] I can take this over, unless you want to keep it.	1	1	https://issues.jboss.org/browse/JBEAP-6868
9465	JBEAP-6868	We've found it works most easily to do major API version bumps in upstream (which in this case would be ported back to 7.x). At some point the EAP code base and the upstream one will fork more than they are now, and doing major bumps in upstream ensures there is a minor stream available for future changes in EAP's branch of the fork. Right now the two branches are the same so you could use 2.1 and then bump upstream to 3.0 in the next upstream release where there is a change. But then we have to remember it's being done this way, and guard against the upstream change being to 2.2, which would be the intuitive thing for someone seeing 2.1 as the current version to do. But that would prevent the EAP branch from producing a 2.2 if it didn't want everything that was in the upstream 2.2. So, a 2.1 would be ok if you prefer but just to follow the KISS principle and stay consistent across subsystems we say always bump the upstream version by a major increment.	0	0	https://issues.jboss.org/browse/JBEAP-6868
9466	JBEAP-6858	Do I understand it correctly? It is not possible to do it because of limitation of current implementation. I noticed you changed the type of Jira from Bug to Enhancement. Do you plan to address this issue in the future? If no we should close it.	1	1	https://issues.jboss.org/browse/JBEAP-6858
9467	JBEAP-6853	[~tomekadamski]: This is a bad user experience. Is there any underlying technical reasoning for nacking this?	1	1	https://issues.jboss.org/browse/JBEAP-6853
9468	JBEAP-6842	I added a blue focus box to the Jtree in treepackspanel. The focus box currently spans the height of the rows. Is this okay or should the focus box span the jtrees true size?	1	1	https://issues.jboss.org/browse/JBEAP-6842
9469	JBEAP-6842	What about input field focus proposal, would it be processed? !input_proposed_focus.png|thumbnail!	1	1	https://issues.jboss.org/browse/JBEAP-6842
9470	JBEAP-6842	Hi Jan, I'm unsure if an input field focus indicator is required, since the cursor appears on the focused text field. Additionally, this is quite a pain to implement since it requires drawing an indicator beyond the components bounds and any solution would be quite hacky.. I can do it but I believe it will take much more than it is worth. What do you think?	1	1	https://issues.jboss.org/browse/JBEAP-6842
9471	JBEAP-6842	Well, it was only a proposal. There are probably more important issues to work on. Still, can ve avoid drawing focus indicator outside component by wraping JTextField in JPanel and draw the border there so that it doesn't interfere with other components?	1	1	https://issues.jboss.org/browse/JBEAP-6842
9472	JBEAP-6842	Hi Daniel, I checked that and I am glad you make it possible, it looks good. However, the margins could be played with a little more to find the best visual effect. Currently the space around the field is not equal in every direction (bottom + right larger): !textfield-focus-current.png|thumbnail! And I think there is no need to have more than 1px space between filed and border. With these sizes for inset in JFocusPanel I get the desired effect: {code:java} gbc.insets = new Insets(1, 1, 0, 0); {code} !textfield-focus-altered-size.png|thumbnail! I am not quite sure why I see equally distributed space with inset sizes not the same, I was just trying different values. This way also the "blank" space between different fields get a bit smaller, but still bigger than the original without field focus. Can these margins can be altered too to get as close as possible to the margins in DR11 installer?	1	1	https://issues.jboss.org/browse/JBEAP-6842
9473	JBEAP-6834	The need to reload e.g. many servers in production domain just to enable transactional statistics is very user unfriendly behaviour. If this needs to add additional way how to inform Narayana about the variable value change this new way should be implemented. Do you need RFE for this?	1	1	https://issues.jboss.org/browse/JBEAP-6834
9474	JBEAP-6834	Hi [~tomjenkinson], I changed this ticket from bug to enhancement. Could you please take a look at the PR which is linked to the WFLY-7534. Thank you.	1	1	https://issues.jboss.org/browse/JBEAP-6834
9475	JBEAP-6702	Setting qa_ack. [~jmesnil] There is dependency for <module name="org.apache.activemq.artemis" /> in org.apache.activemq.artemis.journal module. I'm not sure if it's correct as everyone else should depend on "journal" module and "journal" module should not load "artemis" main module. wdyt?	1	1	https://issues.jboss.org/browse/JBEAP-6702
9476	JBEAP-6702	[~jmesnil] As JDBC store functionality was added to EAP 7.1.0.DR8, there is new jar {{artemis-jdbc-store-1.5.0.redhat-1.jar}} in artemis main module. Is it correct location. Shouldn't it be move to journal module as well?	0	0	https://issues.jboss.org/browse/JBEAP-6702
9477	JBEAP-6702	[~jmesnil] Looking at the code {{JDBCJournalImpl}} from artemis-jdbc-store.jar is implementing interface {{Journal}} from artemis-journal.jar which in "journal" module. It seems that artemis-jdbc-store.jar should be located in "journal" module as well. wdyt?	0	0	https://issues.jboss.org/browse/JBEAP-6702
9478	JBEAP-6694	Hey [~baranowb], what is the current state of this one? I see that all related Jiras are in "Pull Request Sent" for quite a long time. Just trying to clean things a little bit. Thank you.	1	1	https://issues.jboss.org/browse/JBEAP-6694
9479	JBEAP-6591	Below taken from email thread. Will update here with reply from Jimmy/Rich per below... Hi Bilge, what do you think? Can we remove the JBoss EAP 5 Legacy Clients section from the EAP 7 documentation? Erich ---------- I am ok removing it from the documentation. I would like to get confirmation from Jimmy and Rich as well. (Now cc-ed)	1	1	https://issues.jboss.org/browse/JBEAP-6591

#	Requirement	Questions	Answered	Successfully Answered	URL
9480	JBEAP-6556	Component upgrade contains two re-opened issues: JBEAP-5965, JBEAP-6622 . Re-opening.	1	1	https://issues.jboss.org/browse/JBEAP-6556
9481	JBEAP-6484	[~jmesnil] Hi Jeff, there is already jira for this - JBEAP-5830. Can this be closed as duplicate?	0	0	https://issues.jboss.org/browse/JBEAP-6484
9482	JBEAP-6355	[~rhn-support-pnag]: Was this merged to stage? Should this be moved to 'Ready for QA'?	0	0	https://issues.jboss.org/browse/JBEAP-6355
9483	JBEAP-6344	[~sgilda]: I'd suggest to drop the "such a loser!" part of the message. It looks a bit unprofessional to me. What do you think?	1	1	https://issues.jboss.org/browse/JBEAP-6344
9484	JBEAP-6290	Since there is proper fix now in jackson, what about to revert the https://github.com/wildfly/wildfly/pull/9173 in wildfly?	1	1	https://issues.jboss.org/browse/JBEAP-6290
9485	JBEAP-6253	[~rsvoboda] This is in, can you qa-ack?	1	1	https://issues.jboss.org/browse/JBEAP-6253
9486	JBEAP-6244	[~spyrkob]: When do you plan to fix this issue?	0	0	https://issues.jboss.org/browse/JBEAP-6244
9487	JBEAP-6019	CXF 3.1.8 is out yet updated and I think it should be pulled in via jbossws-cxf update once CXF 3.1.8 is available and pass TCK ;).	0	0	https://issues.jboss.org/browse/JBEAP-6019
9488	JBEAP-5919	[~sguilhen] Thanks. Is there any progress with jiras for [2] and [3]?	1	1	https://issues.jboss.org/browse/JBEAP-5919
9489	JBEAP-5830	This should be upgraded to tag created from maintenance branch: https://github.com/rh-messaging/jboss-activemq-artemis/tree/jboss-1.4.0-x [~ataylor] [~jmesnil] Any plans when tag could be created? It would be great to have it in EAP 7.1.0.DR7 (dead line is in ~ 1 week) [~kabirkhan] All jiras are in new state but most of them (maybe all) has already PR in jboss-1.4.0-x. Do you know who moves jiras to correct state?	1	1	https://issues.jboss.org/browse/JBEAP-5830
9490	JBEAP-5830	[~ataylor] [~jmesnil] Please see Miroslav's comments above. Andy, can you relink accordingly? For unlinked issues, please unresolve them and clear the fix version	1	0	https://issues.jboss.org/browse/JBEAP-5830
9491	JBEAP-5830	[~jmesnil] [~ataylor] is there any update on this? Do we have Artemis upgrade jira for EAP 7.1.0.DR8?	1	1	https://issues.jboss.org/browse/JBEAP-5830
9492	JBEAP-5784	Hi Martin. I'm working on the solution to this, and I just wanted to clarify with you what you believe the ideal solution to be. I've opted for the option '2' of the ones you've suggested; however, I am not sure which screen to return to. Do you mean the advanced options screen? Or the JDBC vendor screen? I'm unable to exit cleanly from the vendor screen as well. Perhaps it's best to add the 'press 1 to continue ...' dialogue in between the JDBC vendor selection and the jar entry?	1	1	https://issues.jboss.org/browse/JBEAP-5784
9493	JBEAP-5784	Hi Thomas, I've tested your changes. I found missing text and one regression. I think the new functionality is otherwise good. *Missing text*: When there are no JARs, pressing 1 results in: {code:java} Driver JAR locations: /home/mschvarc/lib/postgresql.jar [/home/mschvarc/lib/postgresql.jar] Enter 1 to accept, 2 Add additional jar, 3 Remove jar, 4 Redisplay JDBC vendor 3 [ ] Enter 1 to accept, 2 Add additional jar, 3 Remove jar, 4 Redisplay JDBC vendor 1 dynamicComponents.jars.empty //missing resource press 1 to continue, 2 to quit, 3 to redisplay. {code} "dynamicComponents.jars.empty" is not found anywhere in the 3 git repositories. From the behavior, pressing 1 sends the user back to JDBC Driver Vendor selection. Pressing 3 also sends the user to the same screen. Either way the user can now *quit* the installation or re-display JDBC Driver Vendor selection, which is what this JIRA is about. *Bugs*: IBM DB2 Driver selected Input Postgres or any other JAR (has to be an actual jar file) Remove the JAR (press '3 to remove' repeatedly) {code:java} You cannot remove anymore elements. [ ] Enter 1 to accept, 2 Add additional jar, 3 Remove jar, 4 Redisplay JDBC vendor 1 dynamicComponents.jars.empty press 1 to continue, 2 to quit, 3 to redisplay. 1 None of the provided jars contained the chosen driver's Driver class: "com/ibm/db2/jcc/DB2Driver.class". Do you want to continue? (y/n) [n]: {code} I should not be able to proceed with *zero* JARs present (regression against DR4).	1	1	https://issues.jboss.org/browse/JBEAP-5784
9494	JBEAP-5784	Hi Thomas, your above solution looks good. However, as of 7.1.0.DR5 and when building the installer manually from master, this is what I get: {code:noformat} JDBC Driver Setup .... Driver vendor: 0 [x] IBM DB2 1 [ ] Microsoft SQL Server 2 [ ] MySQL .... Input Selection: 1 Driver JAR locations: /home/mschvarc/ [/home/mschvarc/] [Starting processing] ... Starting process Advanced configuration - JDBC (5/25) Failed to copy jars. /home/mschvarc (Is a directory) Starting process Onfail Server Shutdown Job (6/25) 11:17:45,756 ERROR [org.jboss.as.protocol] (Thread-39) WFLYPRT0003: Failed to close resource org.jboss.as.cli.embedded.ThreadContextsModelControllerClient@474ab1e9: java.lang.IllegalStateException: WFLYEMB0023: The embedded server is stopping and invocations on the ModelControllerClient are not available {code} The installer selects the file without any validation and begins/advances the installation, resulting in crash if file is invalid or a directory. "Validation error:: At least one jar is required for the JDBC driver." Where does this warning take you afterwards?	1	1	https://issues.jboss.org/browse/JBEAP-5784
9495	JBEAP-5751	[~mnovak] Thanks. I'm actually fixing this bug based on your comments in JBEAP-4363. (updated today with a question for you. please take a look when you can:) Close this one as dupe?	1	1	https://issues.jboss.org/browse/JBEAP-5751
9496	JBEAP-5576	AFAIK, JBoss is open source, even EAP. So the source code should be published. So where is the source code of wildfly-core 2.1.6-redhat-1 ?	1	0	https://issues.jboss.org/browse/JBEAP-5576
9497	JBEAP-5576	That's fine. But why don't you want to put the source code in the RedHat GA repository (https://maven.repository.redhat.com/ga/), just like version 2.1.2.Final-redhat-1, or the other artefacts ?	0	0	https://issues.jboss.org/browse/JBEAP-5576
9498	JBEAP-5518	Hi, It started logging in output system log. Can we make it customised to log into specific folder. under log dir. Thanks Jivendra	1	0	https://issues.jboss.org/browse/JBEAP-5518
9499	JBEAP-5480	[~treblereel] [~fjuma] I am trying to clean up some states of issues, and it isn't clear to me if this one is in EAP 7.x or not. Can you please let me know so I can adjust accordingly?	1	1	https://issues.jboss.org/browse/JBEAP-5480
9500	JBEAP-5480	[~treblereel] Would you be able to submit PRs against wildfly and jboss-eap7 to get this upgrade in?	1	1	https://issues.jboss.org/browse/JBEAP-5480
9501	JBEAP-5443	[~wolfc], what's the reason for the subject change? This was a jira for moving from 3.0.18 to 3.0.19 in 7.0.x, similar to JBEAP-5438 which did the same on 7.1. I've also opened up a PR one month ago.	1	1	https://issues.jboss.org/browse/JBEAP-5443
9502	JBEAP-5429	[~msvehla] JBEAP-5112 follows up this one to get fixed, so would you please provide qa_ack? Thank you.	1	1	https://issues.jboss.org/browse/JBEAP-5429
9503	JBEAP-5400	Singleton MDBs (if that's what the 1st item in this issue is about?) are described at https://access.redhat.com/documentation/en/red-hat-jboss-enterprise-application-platform/version-7.0/developing-ejb-applications/#clustered_singleton_mdbs. Singleton services are described at https://access.redhat.com/documentation/en/red-hat-jboss-enterprise-application-platform/version-7.0/development-guide/#clustered_ha_singleton_service (there's also a chapter in the EJB Dev Guide, but that's going away as a part of JBEAP-3816). Singleton deployments are described at https://access.redhat.com/documentation/en/red-hat-jboss-enterprise-application-platform/version-7.0/development-guide/#ha_singleton_deployments. So the docs already have it all. Is the issue about adding links into the Config Guide that would point to the Dev Guides?	0	0	https://issues.jboss.org/browse/JBEAP-5400
9504	JBEAP-5349	Upstream pull request: - https://github.com/wildfly/wildfly/pull/9082 Tag: - https://github.com/wildfly-security/jboss-negotiation/releases/tag/3.0.3.Final Are SET Ok to take over with the EAP updates?	0	0	https://issues.jboss.org/browse/JBEAP-5349
9505	JBEAP-5331	[~manaRH]: I updated the merge request as you suggested. Could you please review again? Thanks for the help!	1	1	https://issues.jboss.org/browse/JBEAP-5331

#	Requirement	Questions	Answered	Successfully Answered	URL
9506	JBEAP-5331	[~mnovotn], the last changes look good(*) to me. Do you want to have a look at them or can assume they are alright? (*) There's "Red Hat JBoss Web Framework" instead of "Red Hat JBoss Web Framework Kit". I don't know if we care about that. I've also gone through the previous changes and found that there's an almost identical duplicate paragraph in the beginning of chapter 7.2 and 7.3. I don't think these are big issues, but I'll reopen the Jira.	1	1	https://issues.jboss.org/browse/JBEAP-5331
9507	JBEAP-5310	How is this related to JBEAP-5253 ? Duplicate?	1	1	https://issues.jboss.org/browse/JBEAP-5310
9508	JBEAP-5275	[~treblereel], [~brian.stansberry], JBEAP-5310 has been verified, but console logging in standalone-full-ha.xml is still enabled in DR4. When can we expect this fixed?	0	0	https://issues.jboss.org/browse/JBEAP-5275
9509	JBEAP-5256	@Jeff what is the next planned version in 7.1.0?	0	0	https://issues.jboss.org/browse/JBEAP-5256
9510	JBEAP-5243	[~nziakova]: I just noticed the outdated year in the license headers today. Should I merge this update? Waiting for QE to acknowledge this is an issue before I merge the 1200+ file changes.	0	0	https://issues.jboss.org/browse/JBEAP-5243
9511	JBEAP-5126	[~tomjenkinson] yep, ok, I do understand now. I haven't read the mailing list properly. With the new annotation it will be possible to forbid transaction propagation to happen. And what about JTS and IIOP? It behaves currently in the same fashion as ejb remoting. Is this by design behavior which will remain?	1	1	https://issues.jboss.org/browse/JBEAP-5126
9512	JBEAP-5057	[~lthon], [~sgilda] -- Does this JIRA need to be CLOSED manually? Looks like the work is completed. Thanks!	1	1	https://issues.jboss.org/browse/JBEAP-5057
9513	JBEAP-5041	Hello [~mjurc], can you be more specific please ? Or can you provide an example ? Thank you.	1	1	https://issues.jboss.org/browse/JBEAP-5041
9514	JBEAP-5041	[~mstefank]: Would you kindly provide this JIRA with ack+ seeing this has already been merged in upstream?	1	1	https://issues.jboss.org/browse/JBEAP-5041
9515	JBEAP-4980	Isn't this an EAP change on how it implements JBossWS?	0	0	https://issues.jboss.org/browse/JBEAP-4980
9516	JBEAP-4957	Is this still an issue?	1	1	https://issues.jboss.org/browse/JBEAP-4957
9517	JBEAP-4957	Hello Romain, Could you check this PR?	1	1	https://issues.jboss.org/browse/JBEAP-4957
9518	JBEAP-4943	I just checked with aesh-0.66.11 on Mac OS, I can skip words using Alt +Arrow. And this is not that bad already. =,/,(are delimiters. It seems that the issue is for Windows only. Would you confirm that this is Windows only? Thanks. JF.	1	1	https://issues.jboss.org/browse/JBEAP-4943
9519	JBEAP-4930	Each tree node displays the resource description retrieved from r-r-d, so the "Management Model" view is fairly descriptive, because is allows the user navigates in the domain model and see all attributes, operations and its descriptions. It there should be an enhancement, can you elaborate on which description is necessary ?	0	0	https://issues.jboss.org/browse/JBEAP-4930
9520	JBEAP-4886	[~mstyk] Hi Martin, Can you attach your configuration file and examples?	1	0	https://issues.jboss.org/browse/JBEAP-4886
9521	JBEAP-4886	[~mstyk] Hi Martin,I have made a modify according to your example, when i do the operation ??/subsystem=messaging-activemq/server=default:list-connection-ids?? ,the connection is built on two servers .And when i change the "RandomConnectionLoadBalancingPolicy" into "FirstElementConnectionLoadBalancingPolicy",all of the connection are built on one server,this is right according to the documentation .So I think the policy is useful,what do you think about it?	1	0	https://issues.jboss.org/browse/JBEAP-4886
9522	JBEAP-4886	How is the connection factory defined on the server? Is it the same on every server? I assume the connection factory is retrieved via JNDI. If so, how is the JNDI lookup defined? Is every client looking up the connection factory on the same server? In any event, no client-side solution is going to ultimately solve the problem you've outlined in your use-case as there's still a chance all the clients would end up connecting to the same server. The real solution would require a server-side feature that balanced connections smartly across the cluster.	1	1	https://issues.jboss.org/browse/JBEAP-4886
9523	JBEAP-4886	How are the clients doing their JNDI lookup? Is every client looking up the connection factory on the same server? If so, this behavior is not surprising. There are a few ways to address this with configuration: * Add connectors for the other (maybe even all) nodes in the cluster to "RemoteConnectionFactory." This, of course, is a bit tedious and prone to error, but it is nonetheless effective. * Remove the "connectors" attribute from the "RemoteConnectionFactory" and use "discovery-group" instead (referencing "dg-group1"). This is the recommended approach. * Ensure somehow that the clients' JNDI lookups are balanced among the cluster nodes.	1	1	https://issues.jboss.org/browse/JBEAP-4886
9524	JBEAP-4886	How can you update the cluster topology on the client side before the first connection is created? Where would you get the cluster topology information from? As I see it, the client would need to connect to its target, download the cluster topology, and then potentially connect to a different node in the cluster. This turns 100 connection events into potentially 200 connection events and would significantly increase the time it took the client to connect to the cluster. I think this solution would also get a bit messy because the client would want to maintain its original connection to the cluster while it attempted to connect to a different node so in case the topology changed it could receive those updates. In the end you could still end up with many (or even all) clients connected to the same cluster node since there's nothing "smart" about this approach. As I mentioned before, the real solution to this problem would require a server-side feature that balanced connections smartly across the cluster. This has been on our radar for some time, but we haven't gotten to it yet. See https://issues.apache.org/jira/browse/ARTEMIS-23.	1	1	https://issues.jboss.org/browse/JBEAP-4886
9525	JBEAP-4715	Upstream pull request: - https://github.com/wildfly/wildfly/pull/9082 Tag: - https://github.com/wildfly-security/jboss-negotiation/releases/tag/3.0.3.Final Are SET Ok to take over with the EAP updates?	1	1	https://issues.jboss.org/browse/JBEAP-4715
9526	JBEAP-4673	Hi [~asoldano], do you mind if I take this one?	1	0	https://issues.jboss.org/browse/JBEAP-4673
9527	JBEAP-4673	Thinking about this, I believe this can't work the way it's implemented. System properties in WildFly, IIRC, are only available at boot time. If we somehow tie this to a model change on WildFly, then I believe we can have this work on-the-fly. What do you think?	1	0	https://issues.jboss.org/browse/JBEAP-4673
9528	JBEAP-4673	Any update on this ? how hard would it be to achieve it ?	1	0	https://issues.jboss.org/browse/JBEAP-4673
9529	JBEAP-4673	Hi [~rsvoboda], the easiest and more manageable solution would be if the deployed web services already use logging interceptors. Then we would set a new property on WildFly that would trigger the on-the-fly logging functionality implemented in CXF. Would that be good enough?	0	0	https://issues.jboss.org/browse/JBEAP-4673
9530	JBEAP-4609	Why exactly is this an issue? This is really just about coding conventions. Either you use static imports, or you don't, or you mix it. Besides, even if we unified it now, we would have to introduce it as a new rule for developers to adhere to this convention. We would probably have to add it as a new checkstyle rule to enforce it. Just fixing it once and then leaving it would make absolutely no sense. Also I don't think that coding conventions fall under QA responsibility. Was this discussed with developers?	1	0	https://issues.jboss.org/browse/JBEAP-4609
9531	JBEAP-4600	[~thofman]: Has this jira correct description? There is mentioned "3.0.15.Final", but current version of RESTEasy in EAP is 3.0.16.	1	1	https://issues.jboss.org/browse/JBEAP-4600

#	Requirement	Questions	Answered	Successfully Answered	URL
9532	JBEAP-4547	[~jamezp]: That sounds like a good plan at this time. At some point in the future, we can revisit this. Since you created the JIRA, would you like to close it? Or should I close it was Won't Fix? ;-)	1	0	https://issues.jboss.org/browse/JBEAP-4547
9533	JBEAP-4485	[~manaRH]: I believe I addressed all of the JIRAs in this merge request. Would you please review and let me know if it needs anything? Also make sure I didn't restate anything improperly :-).	1	1	https://issues.jboss.org/browse/JBEAP-4485
9534	JBEAP-4485	I got an info that EAP7 docs are now frozen, but we need this for Windup EAP6->7 rule migrations to point to new added parts of Migration guide which I requested. [~nziakova], could you do something to ack this changes?	1	1	https://issues.jboss.org/browse/JBEAP-4485
9535	JBEAP-4293	[~jpallich] Just out of curiosity which Weld issue is a blocker for JDG 7? Is it WELD-2110?	1	1	https://issues.jboss.org/browse/JBEAP-4293
9536	JBEAP-4292	@MapKeyColumn refers to a column mapped in embeddable map value [HHH-8305] - Envers auditing OneToOne inserts null in Aud JoinTable in a not stablished relationship [HHH-8507] - SQLServer2005LimitHandler throws "string out of range -1" error [HHH-8999] - NullPointerException when updating or deleting multiple entities of same type with non-comparable IDs [HHH-9322] - Updating multiple entities with the same name in the same flush throws a ClassCastException when using a non-Comparable user type as the entity ID [HHH-9676] - TypeSafeActivator.applyDDL doesn't process composing constraints when @NotNull is present [HHH-10230] - Using boolean or null literals results in a parser error [HHH-10245] - SelfDirtinessTracker not found [HHH-10277] - AttributeConverter not applied to attributes of an embeddable used as collection element [HHH-10287] - Cache keys no longer include the entity type [HHH-10290] - org.hibernate.test.cfg.BatchVersionedDataConfigTest does not work properly when hibernate.jdbc.batch_versioned_data is set in environment [HHH-10299] - After adding "@Audited" to an entity field the warning "Use of DOM4J entity-mode is considered deprecated" appears [HHH-10334] - @formula annotation javadoc contains database dependent sample usage [HHH-10341] - Specification violation in Implementation of TupleElement [HHH-10352] - HBM2DDL does not use described name for foreign key in join inheritance [HHH-10360] - Natural ID columns are nullable by default (hbm mappings only) [HHH-10361] - Child module osgi does not exist in hibernate-tutorials [HHH-10373] - Sequence generator for idbag ignores generator parameters [HHH-10375] - adding an entity at index to list with @OrderColumn in detached Entity [HHH-10413] - byte[] as the version attribute broken [HHH-10427] - ServiceRegistry creates multiple service instances and returns uninitialized services [HHH-10476] - Getting an entity with composite ID containing a detached entity fails if PersistenceContext contains a managed copy [HHH-10511] - PluralAttribute.getPersistentAttributeType() wrong value with enhanced classes [HHH-10515] - Stored procedure execution fails to find column [HHH-10522] - Invalid reference to naming chapter in user guide [HHH-10524] - Invalid reference wrt composite identifiers [HHH-10525] - Invalid reference to bytecode enhancement in User Guide [HHH-10535] - Documentation: hibernate.order_updates configuration property [HHH-10540] - Incorrect bytecode instrumentation when there are multiple @Embedded attributes [HHH-10551] - NonIdentifierAttribute.isLazy() returns false when @Basic(fetch=FetchType.LAZY) [HHH-10562] - ManagedEntity linked list broken when mutable, non-associated enhanced entity is evicted [HHH-10583] - BytecodeEnhancement's AssociationManagement breaks with entity containing two ToOne relations with the same mappedBy value. [HHH-10587] - Test NationalizedIgnoreCaseTest.testIgnoreCaseCriteria fails on some databases that don't support nString type [HHH-10588] - ConnectionsReleaseTest should use H2Dialect [HHH-10596] - Some tests fail on oracle12c because of missing dialect-scope [HHH-10598] - Oracle JDBC driver can't handle entities with a one-character name in some cases [HHH-10608] - Avoid HibernateException with Hibernate.initialize and ENABLE_LAZY_LOAD_NO_TRANS [HHH-10614] - Auto evict collection cache throws exception with null values [HHH-10635] - EntityManager.isOpen() should respect status of Session [HHH-10640] - SchemaCreationTest fails on DB2 due to wrong expected statement [HHH-10641] - Some tests fail on Oracle due to ORA-00972: identifier is too long [HHH-10642] - BytecodeEnhancements PersistentAttributesEnhancer bytecode error INVOKESPECIAL [HHH-10644] - Enhancement cause Exception if property field not present [HHH-10646] - Lazy fetch not properly handled when annotation in @MappedSuperclass [HHH-10648] - No MultipleBagFetchException is thrown when fetching two bags eagerly [HHH-10678] - hibernate-mapping's schema is ignored [HHH-10679] - Subselect fetching is not used for collections if their entity owners are initialized in a batch [HHH-10708] - Accessing a lazy collection in an enhanced class deletes it	0	0	https://issues.jboss.org/browse/JBEAP-4292
9537	JBEAP-4281	[~ataylor] Hi ,Andy,about this issue, in order to disable slow client reconnecting with KILL slow client policy,i just want to add a new attribute in ServerConsumer named "isKill" to record the client whether or not was killed, when the slow client is killed,we can set the attribute is true,then when the slow client want to reconnect the server,we can judge according to the attribute ,if the attribute is true we will ignore the reconnect attempt.This is my resolution to fix this issue,what do you think about it ? expected your reply,thanks!	1	0	https://issues.jboss.org/browse/JBEAP-4281
9538	JBEAP-4281	[~jbertram] Hi,Justin,yeah I have done it,when detected the slow consumer with the policy KILL,I will send a packet to the client,but i think it is not necessary for the client to response it.Even though I have done only work for core client,I do not understand why this not work for any other kind of client,can you tell me that?	1	0	https://issues.jboss.org/browse/JBEAP-4281
9539	JBEAP-4275	I would expect that if a server to which a connection in the pool is connected dies then that connection's ExceptionListener (i.e. {{org.apache.activemq.artemis.ra.ActiveMQRAManagedConnection#onException}}) would be invoked and the connection would then be removed from the pool (assuming reconnect-attempts isn't -1 and also assuming there is no backup configured). Then later if pool utilization was high enough to force the creation of a new connection to replace the one that was lost then a new connection would be added to the pool and would be created according to the connection load balancing policy configured on the pool and the updated cluster topology. That said, my hunch is that the main use-case in mind here is when a node is _added_ to the remote cluster. However, if a new node was added to remote cluster then a similar process could occur. If the pool's utilization was high enough so that a new connection was necessary (and there was room in the pool for another connection) then a new connection would be created according to the connection load balancing policy configured on the pool and the updated cluster topology. Therefore, connections could eventually go to the new node. Of course, this doesn't help in the situation where the pool is full. Also, I'm surprised that a complete server restart would be required to reset the pool. Isn't there a way to simply flush the connection pool so that the connections are destroyed and then recreated as they're needed? This was possible in previous versions of EAP. This may not be exposed on the pooled-connection-factory, but I'm fairly certain that the JCA subsystem supports ways to time-out idle connections which would be a way to get connections onto a new node. There are plans to implement server-side client connection re-balancing when cluster topology changes so I think development effort should probably go there rather than this particular feature since there ostensibly are ways to deal with this administratively in the mean-time. Lastly, was this based on a customer request? Typically EAP feature requests come from customers through GSS and product management for approval.	1	0	https://issues.jboss.org/browse/JBEAP-4275

#	Requirement	Questions	Answered	Successfully Answered	URL
9540	JBEAP-4260	management layer submits a task to handle some of the work of processing the CLI request: 11:43:52,363 INFO [org.jboss.as] (management-handler-thread - 1) Created thread: java.lang.Exception at org.jboss.as.server.ServerService$ServerExecutorService$1.newThread(ServerService.java:444) at java.util.concurrent.ThreadPoolExecutor$Worker.<init>(ThreadPoolExecutor.java:612) at java.util.concurrent.ThreadPoolExecutor.addWorker(ThreadPoolExecutor.java:925) at java.util.concurrent.ThreadPoolExecutor.execute(ThreadPoolExecutor.java:1368) at org.jboss.as.protocol.mgmt.AbstractMessageHandler$ManagementRequestContextImpl.executeAsync(AbstractMessageHandler.java:403) at org.jboss.as.protocol.mgmt.AbstractMessageHandler$ManagementRequestContextImpl.executeAsync(AbstractMessageHandler.java:392) at org.jboss.as.controller.remote.ModelControllerClientOperationHandler$CompletedCallback.sendResponse(ModelControllerClientOperationHandler.java:287) at org.jboss.as.controller.remote.ModelControllerClientOperationHandler$ExecuteRequestHandler.doExecute(ModelControllerClientOperationHandler.java:224) at org.jboss.as.controller.remote.ModelControllerClientOperationHandler$ExecuteRequestHandler.access$300(ModelControllerClientOperationHandler.java:130) at org.jboss.as.controller.remote.ModelControllerClientOperationHandler$ExecuteRequestHandler$1$1.run(ModelControllerClientOperationHandler.java:152) at org.jboss.as.controller.remote.ModelControllerClientOperationHandler$ExecuteRequestHandler$1$1.run(ModelControllerClientOperationHandler.java:148) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.Subject.doAs(Subject.java:422) at org.jboss.as.controller.AccessAuditContext.doAs(AccessAuditContext.java:149) at org.jboss.as.controller.remote.ModelControllerClientOperationHandler$ExecuteRequestHandler$1.execute(ModelControllerClientOperationHandler.java:148) at org.jboss.as.protocol.mgmt.AbstractMessageHandler$ManagementRequestContextImpl$1.doExecute(AbstractMessageHandler.java:363) at org.jboss.as.protocol.mgmt.AbstractMessageHandler$AsyncTaskRunner.run(AbstractMessageHandler.java:472) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) at org.jboss.threads.JBossThread.run(JBossThread.java:320) 2) Clustering submits a task to offload potentially blocking work from the MSC threads: 11:32:43,073 INFO [org.jboss.as] (MSC service thread 1-2) Created thread: java.lang.Exception at org.jboss.as.server.ServerService$ServerExecutorService$1.newThread(ServerService.java:444) at java.util.concurrent.ThreadPoolExecutor$Worker.<init>(ThreadPoolExecutor.java:612) at java.util.concurrent.ThreadPoolExecutor.addWorker(ThreadPoolExecutor.java:925) at java.util.concurrent.ThreadPoolExecutor.execute(ThreadPoolExecutor.java:1368) at org.wildfly.clustering.service.AsynchronousServiceBuilder.start(AsynchronousServiceBuilder.java:112) at org.jboss.msc.service.ServiceControllerImpl$StartTask.startService(ServiceControllerImpl.java:1948) at org.jboss.msc.service.ServiceControllerImpl$StartTask.run(ServiceControllerImpl.java:1881) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at	0	0	https://issues.jboss.org/browse/JBEAP-4260
9541	JBEAP-4232	Here's what standalone.sh does: {code} if ["x$CLIENT_SET" = "x" -a "x$SERVER_SET" = "x"]; then # neither -client nor -server is specified if $darwin && ["$JVM_OPTVERSION" = "-d32"]; then # Prefer client for Macs, since they are primarily used for development PREPEND_JAVA_OPTS="$PREPEND_JAVA_OPTS -client" else PREPEND_JAVA_OPTS="$PREPEND_JAVA_OPTS -server" fi fi {code} So, for macs running 32 bit, we append "-client", else we append "-server", appending nothing of course if the user configured one or the other. So, is the proposal here to 1) ignore the mac stuff and just append "-server", or 2) not append anything at all? If the latter, that implies that you consider "-server" by itself to be meaningless, completely redundant to the default tiered compilation behavior, and that the only reason anyone would meaningfully use -server is in conjunction with turning off tiered compilation (-server -XX:-TieredCompilation). Re: the code cache size, it sounds like the 240MB is the default with -server, so the only people not already using that amount are OSX users running 32 bit. BTW, AFAICT the lines I pasted above plus one other line are all that would come out of standalone.sh because of this change. The rest of the nearby logic is also used for putting OSX users on a 32 bit VM. I'm not saying getting rid of just those lines isn't a valid thing to consider, just that doing this doesn't mean the large chunk of script surrounding them can all come out.	1	1	https://issues.jboss.org/browse/JBEAP-4232
9542	JBEAP-4232	[~brian.stansberry], can you check what it shows on Mac OS ?	1	0	https://issues.jboss.org/browse/JBEAP-4232
9543	JBEAP-4232	Why haven't we resolved this issue already? Nothing needs to be changed.	1	0	https://issues.jboss.org/browse/JBEAP-4232
9544	JBEAP-4227	Hi [~brian.stansberry], [~anrobert] - this JIRA is marked as Affecting the RN. Is there any additional information needed in order to have this RN'ed correctly?	1	0	https://issues.jboss.org/browse/JBEAP-4227
9545	JBEAP-4133	[~dlofthouse], will elytron have new Kerberos/SPNEGO implementation? Is it opportunity to address long-lived kerberos bugs?	1	1	https://issues.jboss.org/browse/JBEAP-4133
9546	JBEAP-4133	Preliminary testing with elytron and eap 7.1 shows in this case "Continuation required" roundtrip doesn't occure. [~dlofthouse], could you elaborate more how is this achieved? Is important here the fact, that in {{kerberos-security-facotry}} only "Kerberos V5" and "SPNEGO" can be actually configured as {{mech-oids}}. When trying to configure "Kerberos V5 Legacy" I get {{java.security.PrivilegedActionException: GSSException: Unsupported mechanism requested: 1.2.840.48018.1.2.2}} [~paramjindal] , could you link customer case, please?	1	1	https://issues.jboss.org/browse/JBEAP-4133
9547	JBEAP-4133	[~darranl] , is there any progress on this issue ?	0	0	https://issues.jboss.org/browse/JBEAP-4133
9548	JBEAP-4004	Yes, you are right it uses a public API currently (updating description now) but the main reason here is to avoid doing an EE lookup from MSC service. So I think the strategy should be the following: # remove the QS as is in the existing (now closed) PR as this is not the recommended way to do similar thing # add a QS reaching the same goal, but the right way by using singleton deployments instead of EE lookup from MSC # add a QS showing singleton service, but doing something else than using EE lookup from MSC Does that sounds good now?	1	1	https://issues.jboss.org/browse/JBEAP-4004
9549	JBEAP-4004	I still do not think we should remove the quickstart. I would prefer to get someone to fix the code you have issues with. These are my concerns about removing it, especially this late in the process: * This will break the documentation as we have references to it. * It may also break the jboss.org developer site as they provide information and details for all the quickstarts. * Anyone that implemented code using this quickstart in the previous release will be looking for a way to migrate it to this release. As [~lthon] mentioned in the other JIRA [1], the migration concerns are real. * The code is actually a valid example of a timer service. Can you point out which class and lines of code you are concerned about? I am not clear what you mean when you say it does an "EE lookup from MSC service". [~wdfink] is out on PTO for a while, so he won't be able to update the offending code, but maybe I can find someone else to do it. [1] https://issues.jboss.org/browse/JBEAP-3270?focusedCommentId=13163319&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13163319	1	1	https://issues.jboss.org/browse/JBEAP-4004
9550	JBEAP-4004	Ha, did I change the _Assignee_ to myself? I must have clicked somewhere... Definitely wasn't intentional, sorry about that!	0	0	https://issues.jboss.org/browse/JBEAP-4004

#	Requirement	Questions	Answered	Successfully Answered	URL
9551	JBEAP-3984	[~jpallich] I have released 7.0.0.CR1 in Jira. This issue is still open, but we went with Paul's suggestion: {code} [~/sourcecontrol/jboss-eap7/jboss-eap7] $git grep jgroups -- pom.xml pom.xml: <version.org.jgroups>3.6.8.Final-redhat-2</version.org.jgroups> {code} CR1 contains 3.6.8.Final-redhat-2 with the cherry-pick. What to do with this issue?	1	1	https://issues.jboss.org/browse/JBEAP-3984
9552	JBEAP-3967	[~ataylor] this kind of improvement could be done at the Artemis level (when the broker is started) or at the WildFly level (when the configuration is read)? Do you have something similar in Artemis backlog?	1	0	https://issues.jboss.org/browse/JBEAP-3967
9553	JBEAP-3907	[~psakar] updated the groupId in the 7.0.x branch of the quickstarts. This JIRA is a reminder to update and synchronize the 7.0.x-development branch after GA. After GA, cherry-pick this commit back to 7.0.x-devel: https://github.com/jboss-developer/jboss-eap-quickstarts/commit/24df1d89c9a7d8a9b01dcf2f304e367425a637be	1	1	https://issues.jboss.org/browse/JBEAP-3907
9554	JBEAP-3907	artifacts in the public Maven repository use `7.0.0.GA`. [~pgier]: Do you know how I should proceed with this JIRA? {code} [INFO] Scanning for projects... [ERROR] The build could not read 15 projects -> [Help 1] [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-app-client-ejb:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/app-client/ejb/pom.xml) has 1 error [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ org.jboss.quickstarts.eap:jboss-app-client-ejb:[unknown-version], /home/sgilda/GitRepos/jboss-eap-quickstarts/app-client/ejb/pom.xml, line 40, column 21 [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-app-client-client-simple:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/app-client/client-simple/pom.xml) has 1 error [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ org.jboss.quickstarts.eap:jboss-app-client-client-simple:[unknown-version], /home/sgilda/GitRepos/jboss-eap-quickstarts/app-client/client-simple/pom.xml, line 41, column 21 [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-ejb-asynchronous-client:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-asynchronous/client/pom.xml) has 1 error [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ org.jboss.quickstarts.eap:jboss-ejb-asynchronous-client:[unknown-version], /home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-asynchronous/client/pom.xml, line 55, column 21 [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-ejb-security-interceptors:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-security-interceptors/pom.xml) has 2 errors [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-security-api:jar is missing. @ line 136, column 21 [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ line 147, column 21 [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-one-ejb:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-one/ejb/pom.xml) has 1 error [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-one-ejb:[unknown-version], /home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-one/ejb/pom.xml, line 102, column 21 [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-one-ear:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-one/ear/pom.xml) has 1 error [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-one-ear:[unknown-version], /home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-one/ear/pom.xml, line 102, column 21 [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-two-ejb:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-two/ejb/pom.xml) has 1 error [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-two-ejb:[unknown-version], /home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-two/ejb/pom.xml, line 102, column 21 [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-two-ear:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-two/ear/pom.xml) has 1 error [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-two-ear:[unknown-version], /home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-two/ear/pom.xml, line 102, column 21 [ERROR] [ERROR] The project org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-main-ejb:7.0.0-SNAPSHOT (/home/sgilda/GitRepos/jboss-eap-quickstarts/ejb-multi-server/app-main/ejb/pom.xml) has 1 error [ERROR] 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. @ org.jboss.quickstarts.eap:jboss-ejb-multi-server-app-main-ejb:[unknown-version],	1	1	https://issues.jboss.org/browse/JBEAP-3907
9555	JBEAP-3907	Added the following comment here: https://issues.jboss.org/browse/JBEAP-2584?focusedCommentId=13251125&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13251125 This change [1] was applied to the product version of the quickstarts [2]. I originally thought the change had to be back-ported to the development branch of the quickstarts [3] after the 7.0.0.GA Maven repository was delivered, but I now question whether this is correct. When I cherry-pick the commit back to this branch, the quickstarts don't build because the BOMs in the development repository do not contain the correct version of the artifacts. [~pgier] and I discussed this and we both think that maybe this does not need to be ported back to the development quickstarts. Can someone please clarify whether the development quickstarts should be using the `org.jboss.eap` group ID? [1] https://github.com/jboss-developer/jboss-eap-quickstarts/commit/24df1d89c9a7d8a9b01dcf2f304e367425a637be [2] https://github.com/jboss-developer/jboss-eap-quickstarts/tree/7.0.x [3] https://github.com/jboss-developer/jboss-eap-quickstarts/tree/7.1.x-develop cc @paul gier, [~rafabene], [~rsvoboda], [~nziakova]	1	1	https://issues.jboss.org/browse/JBEAP-3907
9556	JBEAP-3901	[~jpallich] Hi JIri, could you pm_ack this jira, please?	0	0	https://issues.jboss.org/browse/JBEAP-3901
9557	JBEAP-3772	[~jason.greene], do you see any risk of including this upgrade for 7.0.0.GA?	0	0	https://issues.jboss.org/browse/JBEAP-3772
9558	JBEAP-3750	[~stalep] [~andy.miller] Can you please ack/nack/add labels as outlined in http://post-office.corp.redhat.com/archives/jboss-as/2017-March/msg00001.html	1	0	https://issues.jboss.org/browse/JBEAP-3750
9559	JBEAP-3698	[~ahoffer]: It looks like this MR was merged 2 months ago. Should this be handed over to QE?	1	1	https://issues.jboss.org/browse/JBEAP-3698
9560	JBEAP-3664	[~rob.stryker] The product name was changed to follow official product naming guidelines. Please note that while EAP 7.0 product name was changed to JBoss EAP, EAP 6.4.x product name remains "EAP". Can we close the issue ?	1	1	https://issues.jboss.org/browse/JBEAP-3664
9561	JBEAP-3664	The point of this issue is to alert you guys that changes such as this will break consuming tools, and you should telegraph your intentions to the teams who may depend on those flags BEFORE you make such a change. > Can we close this issue? I'm not sure. Have you updated your process documentation to ensure that 'alert the tools team' is now part of the process for any time you want to change the product name?	1	1	https://issues.jboss.org/browse/JBEAP-3664
9562	JBEAP-3664	Hey all: It's been over 6 months. When will these simple tests be added? Can we get it assigned to someone for the next sprint or whatever?	1	1	https://issues.jboss.org/browse/JBEAP-3664
9563	JBEAP-3664	Hi Brian: Any idea when this can get done? Do you guys need any more info from me?	0	0	https://issues.jboss.org/browse/JBEAP-3664
9564	JBEAP-3664	Hi Brian: Any idea when this can get done? Do you guys need any more info from me?	0	0	https://issues.jboss.org/browse/JBEAP-3664
9565	JBEAP-3662	[~valery_sorokin]: Please, can you fill "Affects Version" entry in this jira?	0	0	https://issues.jboss.org/browse/JBEAP-3662
9566	JBEAP-3651	I thought we talked about just setting an offset? Why does it need to be random?	1	1	https://issues.jboss.org/browse/JBEAP-3651

#	Requirement	Questions	Answered	Successfully Answered	URL
9567	JBEAP-3632	Understood, it is whatever the user sets. But if they use that upstream service script for something that sets and maybe expects WILDFLY_HOME. Wouldn't this risk causing some confusion? P.S.: Customers know EAP7 is based on Wildfly. GSS and PM have recently created some text to reply to the usual question "what is the Wildfly version that EAP 7 is based on?" So they will not be surprised to find env variables with a WILDFLY_ prefix.	0	0	https://issues.jboss.org/browse/JBEAP-3632
9568	JBEAP-3593	Is there an ETA for this version?	1	1	https://issues.jboss.org/browse/JBEAP-3593
9569	JBEAP-3590	[~dimitris] [~jason.greene] - Do we need to include this to 7.0.0?	1	0	https://issues.jboss.org/browse/JBEAP-3590
9570	JBEAP-3572	[~sgilda] How would you like to do it if you would like to support multiple versions concurrently ?	1	0	https://issues.jboss.org/browse/JBEAP-3572
9571	JBEAP-3572	Maybe the new QS can be added to WildFly Quickstarts, until it can be merged into product. What do you think, [~sgilda]?	1	0	https://issues.jboss.org/browse/JBEAP-3572
9572	JBEAP-3572	[~nziakova]: I was thinking the same thing! [~ozizka]: Do you want to modify and submit your pull to the WildFly quickstart repository?	0	0	https://issues.jboss.org/browse/JBEAP-3572
9573	JBEAP-3571	The only linked issue is JBEAP-3364 which is *MINOR*. Update from JGroups-3.6.7.Final to JGroups-3.6.8.Final contains *38 commits*, 4 from them are "Merge pull request". So it is quite a lot of changes, for details see https://github.com/belaban/JGroups/compare/JGroups-3.6.7.Final...JGroups-3.6.8.Final We are in *BLOCKERS only phase*, PR https://github.com/jbossas/jboss-eap7/pull/369 is already merged. How did this happen ? [~n1hility] [~jpallich]	0	0	https://issues.jboss.org/browse/JBEAP-3571
9574	JBEAP-3442	Hi [~ahoffer], sorry for unclear description. There is a small difference between the current text and text which is asked to be added by this jira. The information which is currently in doc talks about commit phase of two phase commit. This should be clarified in some way like change {{The JDBC driver returns XAER_RMERR when a network failure occurs during the commit phase.}} to something like {{The JDBC driver returns XAER_RMERR when a network failure occurs on commit phase of 2PC.}} This jira asks to add one more information about limitation of JDBC drivers as mentioned in description. This limitation lays in returning wrong error code when one phase commit optimization of 2PC is used. The description of the limitation should be clarify. My idea would be {{When connection failure happens when one phase commit optimization is used JDBC driver should return XAException error code XAER_RMFAIL but the driver returns XAER_RMERR.}}. So basically it's about adding more limitation info under the section. Now concerning one phase commit optimization for various jdbc drivers. Is my intention understandable or should I add some more info? Thank you Ondra	1	1	https://issues.jboss.org/browse/JBEAP-3442
9575	JBEAP-3433	[~lcosti.redhat]: The links for the JBCS docs are not available on the QA preview server yet, are they? I tried the following URL and didn't find anything: {code} :LinkJBCShttpdInstall: /documentation/en/red-hat-jboss-core-services/version-1/apache-http-server-installation-guide/ {code}	1	1	https://issues.jboss.org/browse/JBEAP-3433
9576	JBEAP-3426	Thanks for the review, [~jstefl]. Do you know what the download is going to be named on the Customer Portal? [~mbabacek], do you know where we can find this out?	1	1	https://issues.jboss.org/browse/JBEAP-3426
9577	JBEAP-3389	Hi [~hhovsepy]: Does this capture what you are trying to say? The Postgres Plus JDBC driver returns XIDs for all prepared transactions that exist in the Postgres Plus database, so there is no way to determine the database to which the XID belongs. This issue is limited to cases where you define more than one data source for the same database in JBoss EAP.	1	1	https://issues.jboss.org/browse/JBEAP-3389
9578	JBEAP-3389	Hi [~sgilda], I will answer for Hayk as he moved to other department and I'm responsible for transaction component in EAP. [~dmichael] It's so. The problem, caused by the behavior, is visible for a user when there are more datasources configured under EAP pointing to the same database server. I would rephrase the sentence a bit, would that give you a sense? {code} The Postgres Plus JDBC driver returns XIDs for all prepared transactions that exist in the Postgres Plus Server, so there is no way to determine the database to which the XID belongs. This could cause a trouble during recovery in case that JBoss EAP configures more than one datasource pointing to the same Postgres Plus Server. As that in-doubt transaction recovery could be run under wrong account which causes fail of the recovery attempts. {code}	1	1	https://issues.jboss.org/browse/JBEAP-3389
9579	JBEAP-3389	[~ochaloup]: Thanks for the explanation. I might rephrase it just a bit more. How is this: The Postgres Plus JDBC driver returns XIDs for all prepared transactions that exist in the Postgres Plus Server, so there is no way to determine the database to which the XID belongs. If you define more than one data source for the same database in JBoss EAP, an in-doubt transaction recovery attempt could be run under wrong account, which causes the recovery to fail.	1	1	https://issues.jboss.org/browse/JBEAP-3389
9580	JBEAP-3374	[~tremes]: Can you please review the content in my merge request and let me know if it looks okay? Thanks!	1	1	https://issues.jboss.org/browse/JBEAP-3374
9581	JBEAP-3357	Hm ok....I overlooked this. Any suggestions how to rewrite this ? E.g "Weld should scan the archive to perform bean discovery" ?	1	1	https://issues.jboss.org/browse/JBEAP-3357
9582	JBEAP-3346	When will 5.0.8 be released?	1	1	https://issues.jboss.org/browse/JBEAP-3346
9583	JBEAP-3341	[~stalep] [~andy.miller] Can you please ack/nack/add labels as outlined in http://post-office.corp.redhat.com/archives/jboss-as/2017-March/msg00001.html	0	0	https://issues.jboss.org/browse/JBEAP-3341
9584	JBEAP-3306	[~ochaloup]: Where did you get the above code example? I am just curious if it has been tested. Some of the code seems incomplete. Here are a couple of examples. There is no definition for the `orb` variable referenced in this line: orb = com.arjuna.orbportability.ORB.getInstance("ClientSide"); This line does not assign the new object to a variable: new ContextPropagationManager();	1	1	https://issues.jboss.org/browse/JBEAP-3306
9585	JBEAP-3305	[~ochaloup]: Could you please review the procedure here and let me know your feedback: https://docs.google.com/document/d/1wCgjL-4Iqr4T4FmHEtd-L7khJo8uH8Wb-R7jPXNTOlY/edit#	1	1	https://issues.jboss.org/browse/JBEAP-3305
9586	JBEAP-3305	[~rhn-support-pnag] and [~ochaloup]: I commented in the merge request but it has already been merged and I don't want to lose the comment. I think it's the wording that is bothering me. It's isn't really narrowing the home interface, is it? Isn't it narrowing the reference to the home interface? For example: . Use the context lookup to narrow the reference to the `IIOPBeanHome` home interface. Or am I being too picky here? Also, if this change is made in another merge request, I would suggest changing the variable name for the Properties from "prope" to "props". That's the more common variable name used for properties. Thanks!	1	1	https://issues.jboss.org/browse/JBEAP-3305
9587	JBEAP-3305	[~rhn-support-pnag]: Sorry I didn't mention the "prope" variable name before. I had meant to, but I was too bothered by the phrasing in the "narrow" statement. Based on the reply by [~ochaloup], would you mind submitting the fixes to the 'narrowing' phrase and 'prope' I suggested above in another merge request?	1	1	https://issues.jboss.org/browse/JBEAP-3305

#	Requirement	Questions	Answered	Successfully Answered	URL
9588	JBEAP-3193	[~anrobert]: Do you think we should clone this issue so it can be tracked more easily? [~jstefl]: Can you clarify the following statement: * Connectors (mod_cluster, mod_jk, isapi, nsapi): mod_cluster,mod_jk are part of Apache HTTP Server delivered in separate product JBoss Core Services, isapi and nsapi are delivered with JBoss Core Serices as well. Can I just combine them into one statement and say the following: * The mod_cluster, mod_jk, isapi, and nsapi connectors are delivered with the product in a separate "JBoss Core Services" download.	1	1	https://issues.jboss.org/browse/JBEAP-3193
9589	JBEAP-3193	[~jstefl]: I'm not sure I'm clear what you are saying, particularly the 3rd bullet. Is the following okay? {code} Package JBoss EAP 6 JBoss EAP 7 AIO Natives for Messaging Delivered with the product in a The new ActiveMQ Artemis separate "Native Utilities" subsystem is included in the download JBoss EAP download. Note: The libraries are located in the JBoss EAP 7 org.apache.activemq.artemis module. {code}	1	1	https://issues.jboss.org/browse/JBEAP-3193
9590	JBEAP-3193	[~jstefl]: 1. The only reference I could find for JBoss Core Services was the link you provided here: https://access.qa.redhat.com/documentation/en/red-hat-jboss-core-services/version-1/apache-http-server-installation-guide/#about_red_hat_jboss_core_services. I don't see an equivalent product category on https://access.redhat.com/documentation/en/red-hat-jboss-core-services/. Do you know of a published link? 2. I thought you wanted me to mention the module location. I will simplify it. 3. Done!	1	1	https://issues.jboss.org/browse/JBEAP-3193
9591	JBEAP-3134	Let me summarize current situation: 1. For starting EAP as service via JSVC on RHEL or Solaris the "commons-daemon.jar" (name is only illustational) has to be delivered. 2. All combination of JSVC was not delivered in jbcs DR2.1, see JCSP-19 3. There was problem with naming of delivered bits, see JCSP-16 4. I expect JSVC (jar) will be delivered as part of jbcs in zip and rpm distro because jsvc was delivered via zip (it was in native utils) and rpm in EAP 6 as well. [~mladen.turk]: Could you please help us with construction of commands starting EAP as service with JSVC in EAP/standalone and domain mode? Any product specific scripts should be delivered with specific product. For example see handling of tooling for handling [EAP as Windows service|https://issues.jboss.org/browse/JBEAP-2735?focusedCommentId=13152301&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13152301].	1	1	https://issues.jboss.org/browse/JBEAP-3134
9592	JBEAP-3134	I am not sure the JBCS is the place to add this JAR, and we were actually told not to include it. Other consumers of JSVC (it is not only EAP) have their own version of commons-daemon.jar and if we ship our own it may cause class loader issues. Also, where would we instal it, as all consumers have different layouts (e.r. Wildfly and Tomcat)? AFAIK the Tomcats on JWS already have commons-daemon, so maybe EAP7 should ship as an optional RPM the commons-daemon to be used when JSVC is used. Not sure how to do that for the ZIP, maybe some optional ZIP as well. Or just include it on EAP from the starts. We need PMs input on this (JBCS, JWS, EAP).	1	1	https://issues.jboss.org/browse/JBEAP-3134
9593	JBEAP-3134	We'll create a jbcs-httpd24-apache-commons-daemon for that then. EAP7 customers will need to get that from the JBCS channel if they are using RPMs. And stick the .jar with the executables for the ZIP. A JBCS JIRA will be created for that. I still wonder... EAP6 and EWS are not the only consumers of JBCS. What happens if the product already has a commons-daemons.jar in the classpath?	0	0	https://issues.jboss.org/browse/JBEAP-3134
9594	JBEAP-3079	"Component Upgrade"? What gets upgraded? :-)	1	1	https://issues.jboss.org/browse/JBEAP-3079
9595	JBEAP-3079	[~baranowb], [~pgier] can you please comment on the state of the JIRA?	1	1	https://issues.jboss.org/browse/JBEAP-3079
9596	JBEAP-3067	[~pgier]: I was told this change made it into the ER4 release, so I updated the quickstart POM files to use the new GAV: <dependency> <groupId>org.jboss.eap</groupId> <artifactId>wildfly-ejb-client-bom</artifactId> <type>pom</type> <scope>provided</scope> </dependency> <dependency> <groupId>org.jboss.eap</groupId> <artifactId>wildfly-jms-client-bom</artifactId> <type>pom</type> </dependency> However, when I try to build the quickstarts, I get errors: 'dependencies.dependency.version' for org.jboss.eap:wildfly-ejb-client-bom:pom is missing. 'dependencies.dependency.version' for org.jboss.eap:wildfly-jms-client-bom:pom is missing. Have either the GA [1] or Staging [2] repository been updated with this change yet? [1] https://maven.repository.redhat.com/ga/ [2] http://jboss-developer.github.io/temp-maven-repo/	1	1	https://issues.jboss.org/browse/JBEAP-3067
9597	JBEAP-3067	[~pgier]: I thought that was the purpose of the staging repository. We update that repository with changes needed for the public version of the quickstarts [1]. [~rafabene]: Is that your understanding of the purpose of that repository? I don't think I can make this change until that is done. And I usually cherry-pick the changes to the product version. Is there a reason we can't update the staging repository?	1	1	https://issues.jboss.org/browse/JBEAP-3067
9598	JBEAP-3061	I tried to set auto-create-jms-queues to true and auto-delete-jms-queues to false. When I try to send message to non-existent queue then queue is created and consumer is able to receive this message. Problem occurs after the consumer was closed. Queue was deleted despite the fact that auto-delete-jms-queues is set to false. I'm not sure if it's problem in Artemis or integration. Providing reproducer below. How to run the test locally: {code} git clone git://git.app.eng.bos.redhat.com/jbossqe/eap-tests-hornetq.git cd eap-tests-hornetq/scripts/ git checkout refactoring_modules groovy -DEAP_VERSION=7.0.0.ER5 PrepareServers7.groovy export WORKSPACE=$PWD export JBOSS_HOME_1=$WORKSPACE/server1/jboss-eap export JBOSS_HOME_2=$WORKSPACE/server2/jboss-eap export JBOSS_HOME_3=$WORKSPACE/server3/jboss-eap export JBOSS_HOME_4=$WORKSPACE/server4/jboss-eap cd ../jboss-hornetq-testsuite/ mvn clean test -Dtest=AutoCreateQueueTestCase#testAutoDeleteQueueDisabled -DfailIfNoTests=false -Deap=7x | tee log {code}	0	0	https://issues.jboss.org/browse/JBEAP-3061
9599	JBEAP-3061	[~jbertram] Could you please have a look at this issue please? I quickly looked at https://github.com/rh-messaging/jboss-activemq-artemis/commit/e293d80f08de58e4ce2bafd9d63c1a4185e079a9 and artemis code base and I see no use for AddressSettings#isAutoDeleteJmsQueues(). Is the deletion of auto-created jms queue implemented?	1	1	https://issues.jboss.org/browse/JBEAP-3061
9600	JBEAP-3050	[~jstefl], I just wanted to make sure: I was previously told that Apache HTTP Server from JBoss Web Server 3 was also going to be supported for use with EAP 7. Is this still correct?	1	1	https://issues.jboss.org/browse/JBEAP-3050
9601	JBEAP-3050	[~dmichael]: I'm not clear which topic you mean. Are you referring to the topic shared between the Release Notes and the Migration Guide or a different topic? Either way, if you understand what needs to be done, please submit a merge request and I will be happy to review it.	1	1	https://issues.jboss.org/browse/JBEAP-3050
9602	JBEAP-3050	Hi Jan, How does this sound: {code} * The Apache HTTP server from JBoss Core Services does not include support HP-UX. If you need to configure load balancing for JBoss EAP 7 instances running on HP-UX servers, you have two choices: 1. You can configure the JBoss EAP 7 instance as a load balancer on the HP-UX server. For details, see Configuring {ProductShortName} as a Front-end Load Balancer. 2. You can deploy Apache HTTP server on a machine that is supported and certified and then run the load balancer on that machine. {code}	1	1	https://issues.jboss.org/browse/JBEAP-3050
9603	JBEAP-3028	Does this tag make it to be in EAP 7.0.0.ER5?	0	0	https://issues.jboss.org/browse/JBEAP-3028

#	Requirement	Questions	Answered	Successfully Answered	URL
9604	JBEAP-2998	[~gaol] I don't get it :) The PR linked here only adjusts a testcase, but where's the code which added the {{--connection-properties}} parameter? Am I blind?	1	1	https://issues.jboss.org/browse/JBEAP-2998
9605	JBEAP-2990	[~clebert] I'm willing to take a look at that for you, but could you point me to the right direction ? Maybe I am wrong, but given your description, it sounds like you have clear idea on where, in the code, the changes needs to be done, and how those unwanted calls could happen...	1	1	https://issues.jboss.org/browse/JBEAP-2990
9606	JBEAP-2931	Hi Mansi, can you please add one additional check? Installer should not allow user to add more modules in case that pool of Codes for given element is empty. Testing this on authentication element option, trying to add another module once the all code were selected by previous modules Installer simply skip to JSSE panel. {noformat} Operation qwe=asd [PermitAll, Required, asd=qwe] [DenyAll, Required, asd=zxc] [Delegating, Required, asd=qwe] [Web, Required, asd=qwe] [JACC, Required, qwe=asd] [XACML, Required, qwe=asd] Enter 1 to accept, 2 Add Module, 3 Remove Module 2 Code JSSE Configuration {noformat} Apart from this, all changes looks good to me. Thanks, Petr	1	1	https://issues.jboss.org/browse/JBEAP-2931
9607	JBEAP-2931	Hi Petr, Is it alright if when all codes are selected, it simply prints an error if the user tries to add another module? Something like "You cannot add more modules." This is what it would look like {noformat} ... Enter 1 to accept, 2 Add Module, 3 Remove Module 2 You cannot add more modules. [DenyAll, Required, sdfsd=sdf] [PermitAll, Required, sdf=sdf] [Delegating, Required, sdf=sdf] [Web, Required, sdf=sdf] [JACC, Required, sdf=sdf] [XACML, Required, asd=asd] Enter 1 to accept, 2 Add Module, 3 Remove Module {noformat}	1	0	https://issues.jboss.org/browse/JBEAP-2931
9608	JBEAP-2929	[~kabirkhan] or [~ppenicka]: Can you correct the Target Release on this to 7.1.0? [~jpallich] has probably PTO ...	0	0	https://issues.jboss.org/browse/JBEAP-2929
9609	JBEAP-2897	[~johnathonlee] [~jolee] [~gaol] The PR is for 7.1.0. Should the target release of this Jira be changed?	1	1	https://issues.jboss.org/browse/JBEAP-2897
9610	JBEAP-2897	[~rsvoboda] The PR is for 7.1.0. Should the target release of this Jira be changed?	1	1	https://issues.jboss.org/browse/JBEAP-2897
9611	JBEAP-2858	[~lcosti.redhat], [~dmichael], Thank you for this JIRA. There are some other reported issues in Release Notes. I think we should track all of them here. What do you think? If you agree, I will link all open RN issues to this one.	0	0	https://issues.jboss.org/browse/JBEAP-2858
9612	JBEAP-2789	[~emmartins]: I created a topic for migration of the `threads` subsystem, but can not find anything on the updates to Infinispan 4.0. Do you know who can help me with this one?	1	1	https://issues.jboss.org/browse/JBEAP-2789
9613	JBEAP-2789	[~pferraro] or [~emmartins]: Is there any information available about the important changes to the Infinispan subsystem? I don't see anything in the community docs here: https://docs.jboss.org/author/display/WFLY10/Admin+Guide	1	1	https://issues.jboss.org/browse/JBEAP-2789
9614	JBEAP-2786	{quote} I believe this belongs in the Configuration Guide, not the Migration Guide. {quote} Yes, the correct place for that is Configuration Guide, but the stack is now configured on a {{channel}} element instead of {{subsystem}}. I see the commit contains a link to Configuration Guide, so that's fine by me. {quote} The Migration Guide contains a section on threads subsystem changes here: {quote} In that case, there's no change needed for JGroups topic. {quote} Where do you see that thread-factory is deprecated? {quote} It is in the output of the following CLI command when EAP is started with HA capabilities ({{./standalone.sh -c standalone-ha.xml}}): {code} /subsystem=jgroups/stack=udp/transport=UDP:read-resource-description(recursive=true) {code} The output contains the following: {code} "thread-factory" => { "type" => STRING, "description" => "The thread factory to use for handling asynchronous transport-specific tasks.", "expressions-allowed" => false, "nillable" => true, "min-length" => 1L, "max-length" => 2147483647L, "deprecated" => { "since" => "3.0.0", "reason" => "Deprecated. Supports EAP 6.x slaves." }, "access-type" => "read-write", "storage" => "configuration", "restart-required" => "no-services" }, {code} I think the project documentation hasn't been updated since this changed. I'll reopen the Jira due to this, but the commit looks good otherwise.	0	0	https://issues.jboss.org/browse/JBEAP-2786
9615	JBEAP-2766	According to discussion in https://gitlab.cee.redhat.com/red-hat-jboss-enterprise-application-platform-documentation/eap-documentation/merge_requests/930 , section Unsupported and Deprecated functionality will be created. Currently, the note about unsupported org.jboss.genericjms module is placed under Pruning heading, but it would be much better to have it in section Unsupported and Deprecated functionality. Could you please move it there? Question about section https://access.qa.redhat.com/documentation/en/red-hat-jboss-enterprise-application-platform/version-7.0/configuration-guide/#configure-generic-jms-resource-adapter is still open. Is it OK to keep section in and add note that it is unsupported?	0	0	https://issues.jboss.org/browse/JBEAP-2766
9616	JBEAP-2758	MR looks good, I only have one comment. Saying "Consider how to approach the migration of large managed domains that include hundreds of hosts" sounds weird for people that don't operate such domains. I'd reword like this: "If you operate a large managed domain (say 10s of hosts or more), consider how to approach gradual migration." We should have a dedicated chapter on this ourselves, but for now, a sentence like this will have to suffice :-) Also, we should say something about migrating a cluster. As I commented above, there's a topic about this in the Patching Guide that we can reuse. Here's how I'd like to proceed: let's fix the few things I mentioned above and then consider this JIRA resolved. For additional topics that require more detail, I'll file other JIRAs. I already filed JBEAP-3715 for gradual migration of large domain. What other topics do you think we need to cover and need help from other parties (developers, consultants, ...)? The recommended practices?	1	1	https://issues.jboss.org/browse/JBEAP-2758
9617	JBEAP-2758	[~lthon]: We do not use parenthesis in running text in our documentation. That is why I modified your suggestions of "If you operate a large managed domain (say 10s of hosts or more), consider " to say "If you are currently running a large managed domain consisting of 10 or more hosts, consider..." They both state pretty much the same thing. If we don't want a hard limit, we should not include it in the text at all. I would suggest we just say: "If you are currently running a large managed domain, consider a gradual migration approach." Does that work for you?	1	1	https://issues.jboss.org/browse/JBEAP-2758
9618	JBEAP-2584	How many/which artifacts does this affect ? If this 'leak' into artifacts used in examples and quickstarts we are looking at a rather massive impact and a direct opposition for the assumption used for the rules/guidelines for Project Wolf/Enterprise Maven repo.	1	1	https://issues.jboss.org/browse/JBEAP-2584
9619	JBEAP-2584	Do you have a matching PR for all the quickstarts to illustrate the impact this will have for users, documentation, examples and tools ?	1	1	https://issues.jboss.org/browse/JBEAP-2584
9620	JBEAP-2584	What are they being renamed from? I presume this breaks maven's conflict detection? If a wildfly and eap dependency that contain the same packages/classes are in the same dependency tree, then they end up both on the classpath (no conflict detection), so arbitrarily one will win. On BRMS/BPMS we make each war in a way that it deploys to both WildFly 10 and EAP 7 (or even EAP 6). Our current approach is very brittle. Is there a documented, recommended approach for that? If this issue means we need different dependency sets to support both WildFly 10 and EAP 7, especially different groupId/artifactId's of those dependencies (so this leaks beyond dependencyManagement), then this can make it very difficult for us to keep supporting both platforms - and therefore test against both platforms.	0	0	https://issues.jboss.org/browse/JBEAP-2584

#	Requirement	Questions	Answered	Successfully Answered	URL
9621	JBEAP-2584	This change [1] was applied to the product version of the quickstarts [2]. I originally thought the change had to be back-ported to the development branch of the quickstarts [3] after the 7.0.0.GA Maven repository was delivered, but I now question whether this is correct. When I cherry-pick the commit back to this branch, the quickstarts don't build because the BOMs in the development repository do not contain the correct version of the artifacts. [~pgier] and I discussed this and we both think that maybe this does not need to be ported back to the development quickstarts. Can someone please clarify whether the development quickstarts should be using the `org.jboss.eap` group ID? [1] https://github.com/jboss-developer/jboss-eap-quickstarts/commit/24df1d89c9a7d8a9b01dcf2f304e367425a637be [2] https://github.com/jboss-developer/jboss-eap-quickstarts/tree/7.0.x [3] https://github.com/jboss-developer/jboss-eap-quickstarts/tree/7.1.x-develop cc @paul gier, [~rafabene], [~rsvoboda], [~nziakova]	1	1	https://issues.jboss.org/browse/JBEAP-2584
9622	JBEAP-2529	I am not sure they should be stored in the history. Does any other shell add commands from executed scripts to the command history?	0	0	https://issues.jboss.org/browse/JBEAP-2529
9623	JBEAP-2529	[~jpallich] I would like to set QA_NACK to this jira, because this jira is "Feature Request". And feature requests shouldn't be filled in JPEAt jira project. But "CDW Engine" remove my QA_NACK. Do you know what I need to do for setting QA_NACK?	1	1	https://issues.jboss.org/browse/JBEAP-2529
9624	JBEAP-2529	[~baranowb]: Ok, can you please set "target release" and "Affects Version"? I recommend to move this jira to EAP7/PRODMGT project. AFAIK, JBEAP jira project shouldn't contain feature requests.	1	1	https://issues.jboss.org/browse/JBEAP-2529
9625	JBEAP-2516	[~dmlloyd] the PR seems to have been merged without attention to your comment https://github.com/jbossas/jboss-eap7/pull/143/files#r48203353 -- is it still a valid concern?	1	1	https://issues.jboss.org/browse/JBEAP-2516
9626	JBEAP-2516	Ok if you are saying that the line-swap is fine, I don't object. I just didn't want to close this until we resolve concerns about the correctness, but I don't understand that code enough to be able to determine it myself. Could someone please provide the devel_ack? Then we'll move it forward, it is already in ER4.	0	0	https://issues.jboss.org/browse/JBEAP-2516
9627	JBEAP-2463	Since EAP 7.0.0.ER5 already contains Weld 2.3.2.Final (or rather 2.3.2.Final-redhat-1.jar), shouldn't this issue be resolved?	1	1	https://issues.jboss.org/browse/JBEAP-2463
9628	JBEAP-2456	Is this still valid ? The text appear correctly for me (Chrome and Firefox), see runtime_tx.png attached. I would like to close as Will not fix.	0	0	https://issues.jboss.org/browse/JBEAP-2456
9629	JBEAP-2443	[~jmesnil] Hi, is fix version for this upgrade 7.0.0.ER4?	1	1	https://issues.jboss.org/browse/JBEAP-2443
9630	JBEAP-2418	Regression definition = the user configures semaphores to be unfair, but that doesn't take any effect and the semaphores stay fair, therefore the performance benefit will not occur. We need to be sure that this won't happen. Can performance team ensure this?	0	0	https://issues.jboss.org/browse/JBEAP-2418
9631	JBEAP-2351	I assume that pressing the "OK" will continue the installation, correct? In that case how about we add something like "Do you wish to use the password anyway?" to the dialogue?	1	1	https://issues.jboss.org/browse/JBEAP-2351
9632	JBEAP-2283	Have you tried the operation /subsystem=messaging-activemq/server=default:import-journal ? It should be able to take a xml dump from HornetQ (EAP6) and import it into Artemis (EAP7)	1	0	https://issues.jboss.org/browse/JBEAP-2283
9633	JBEAP-2240	When could we expect this upgrade please?	0	0	https://issues.jboss.org/browse/JBEAP-2240
9634	JBEAP-2224	Thanks [~bilge] and [~jesper.pedersen] for the acks. Anyone any idea why the issue moved to New/Unresolved once all the acks were in? /cc [~jason.greene]	1	1	https://issues.jboss.org/browse/JBEAP-2224
9635	JBEAP-2214	[~okalman] Please note that the forward-when-no-consumers attribute is no longer used and is replaced by the message-load-balancing-type (see ARTEMIS-19). The current setup uses the default for message-load-balancing-type (ON_DEMAND) which is the same behaviour than forward-when-no-consumers=false I'm not sure if the configuration is correct. Why does the server #4 (which is the target of the bridges) define broadcast-group/discovery-group and a cluster-connection? From the test description, the cluster is formed of server #1 & #3 but not #4.	1	1	https://issues.jboss.org/browse/JBEAP-2214
9636	JBEAP-2158	You need to explain a little how these numbers have been captured. What exact process has been monitored (i.e. selenium process or actual browser)? How did you record the heap useage (tools, etc)?	1	1	https://issues.jboss.org/browse/JBEAP-2158
9637	JBEAP-2158	Did you compare it to other browsers in such an environment?	1	1	https://issues.jboss.org/browse/JBEAP-2158
9638	JBEAP-2064	Why would you split them further? Note I've organized the current module structure essentially in two categories: # the ones the end user code needs directly (contains API) # the ones which shouldn't be used directly (for example the backwards compatibility codecs) but need to be picked up by lucene-core and Hibernate Search Please keep in mind this module structure is maintained for both Hibernate Search (upstream and EAP) and Infinispan (upstream and JDG product). If given suggestions I'm happy to make changes but would prefer to do them upstream first.	1	1	https://issues.jboss.org/browse/JBEAP-2064
9639	JBEAP-2064	Note that we are talking about the way it is delivered in .rpm files. (I've adjusted the Summary to reflect that now, sorry). "Why would you split them further?" Because QE doesn't like the unused JARs to be installed anywhere in the system when the Server does not use them. We can't avoid this completely, but when we split the JARs between separate .rpm files we can avoid shipping some of these. "BTW We removed dependencies to Solr a long time ago." That is another JAR we need to have in a separate sub package. The way it was done now was not to include it in the RPM which is a bad way as prevents another product to re-use the RPM if it needs lucene-solr.	1	1	https://issues.jboss.org/browse/JBEAP-2064
9640	JBEAP-2020	Hi, please, what is status of this jira? It seems that Narayana 5.2.8.Final is part of the EAP7 ER3 but the fix version here or status to 'Ready for QA' is not changed. Thanks Ondra	1	1	https://issues.jboss.org/browse/JBEAP-2020
9641	JBEAP-1916	I checked Chapter 3. Configuring Destinations and I have few notes: * in section 3.3 there is no explanation for attributes queue-address, entries, durable, legacy-entries and selector. What are entries, why can we have more entries, what do entries mean for application developer, etc * which differences are between entries in the example (queue/myQueue jms/queue/myQueue java:jboss/exported/jms/queue/myQueue), why do we have various prefixes, especially java:jboss/exported	0	0	https://issues.jboss.org/browse/JBEAP-1916
9642	JBEAP-1915	First, I don't get why the content about configuration options should be placed in development guide? Why would we want this duplicated in configuration and in development guide? I think that simply something similar to this will be enough for development guide: {code:java} There are several configuration options that can help setup your application or help in your development process (ie. logging). For configuration options of webservices subsystem please consult Configuration guide, chapter xxx. {code} Second, even if it is acked to be placed in Development guide for webservices, is should be distinguished JAX-WS related things from JAX-RS related things, ie. configuring webservices subsystem should be moved within JAX-WS related chapters, as mentioned in JBEAP-2723	0	0	https://issues.jboss.org/browse/JBEAP-1915
9643	JBEAP-1905	[~rhusar], [~pferraro], how can this JIRA be resolved when it lacks the qa_ack+ ? The "workaround / hack" of closing and reopening the bug is meant for situation when you want to start to work on an issue before all the acks are provided. Resolved == MODIFIED in Bugzilla. It is not possible to move bugs to Resolved status without qa_ack. Reopening.	1	1	https://issues.jboss.org/browse/JBEAP-1905

#	Requirement	Questions	Answered	Successfully Answered	URL
9644	JBEAP-1905	This is already fixed and merged. The QA ack was automatic when moving from WFLY to JBEAP. What is your goal here?	0	0	https://issues.jboss.org/browse/JBEAP-1905
9645	JBEAP-1856	[~wolfc] - is the fixVersion intentional? I think it needs a target release instead?	0	0	https://issues.jboss.org/browse/JBEAP-1856
9646	JBEAP-1856	[~jmesnil] This jira seems to be out-of-date. Can it be closed?	1	1	https://issues.jboss.org/browse/JBEAP-1856
9647	JBEAP-1829	[~anrobert]: I'm not clear what content needs updating for the QA automated testing? [~nziakova]: Do you know what is needed here?	1	1	https://issues.jboss.org/browse/JBEAP-1829
9648	JBEAP-1758	[~mnovak] what's the status of this issue? the underlying issue in Artemis has been fixed some times ago. Is it still an issue or did the resolution slips through the cracks?	1	1	https://issues.jboss.org/browse/JBEAP-1758
9649	JBEAP-1667	[~bmaxwell]: I've updated the implicit dependency list via [JBEAP-2361|https://issues.jboss.org/browse/JBEAP-2361] with details from the spreadsheet you'd shared: [EAP 6.4.x Implicit Dependencies|https://docs.google.com/spreadsheets/d/1qx_HM9JZDeratNdU7KmB8iZzc5N2lv0zwsaPiWC6cnc/edit#gid=0] Could you please confirm if CMP or any other module/subsystem in the spreadsheet is not included in EAP 7? https://docs.google.com/spreadsheets/d/1qx_HM9JZDeratNdU7KmB8iZzc5N2lv0zwsaPiWC6cnc/edit#gid=0	1	1	https://issues.jboss.org/browse/JBEAP-1667
9650	JBEAP-1667	[~bmaxwell]: I've updated the 'Implicit Module Dependencies' table for EAP 7. Following are the changes made: * Removed the subsystems that are no longer there: CMP; JacORB; JAXR; Messaging; Threads; Web * Updated the subsystems that have been changed: Messaging to Messaging-ActiveMQ; Web to Undertow * Added newly added subsystems: Batch-JBeret; IIOP-OpenJDK * Added dependencies for the newly added subsystems with inputs from the source. Could you please review the changes below and confirm? Thanks! https://gitlab.cee.redhat.com/nchaudha/eap-documentation/blob/9b50b46b8f71d563d30dd1a9e50bbeedb12d36c0/docs/topics/4353.adoc	1	1	https://issues.jboss.org/browse/JBEAP-1667
9651	JBEAP-1667	[~bmaxwell] Also, could you please confirm if the '{{org.jboss.as.jacorb}}' dependency under EJB3 subsystem needs to be updated/removed? I checked in the source and it's still listed there even though JacORB is no longer supported.	1	1	https://issues.jboss.org/browse/JBEAP-1667
9652	JBEAP-1667	[~bmaxwell]: Let me take a stab at it. {code} JSF (Java Server Faces) Subsystem org.jboss.as.jsf jsf-api jsf-impl The package dependencies listed above are added if the WAR `web.xml` does not bundle JSF. The module slot added for the package dependencies listed below depends on the `web.xml`. javax.faces.api com.sun.jsf-impl {code} Does that sound right?	1	1	https://issues.jboss.org/browse/JBEAP-1667
9653	JBEAP-1667	[~nids1409]: Did you my comment in your merge request and [~bmaxwell]'s comments above? At the minimum, WAR and JSF need to be capitalized. This may also need some rewording. I will try to get to it today.	1	1	https://issues.jboss.org/browse/JBEAP-1667
9654	JBEAP-1667	[~nids1409] Looking at my notes now, it seems that I messed some dependencies in my previous comment. I'm sorry about that. in EE subsystem I forgot to mention that dependency {{org.hibernate.validator}} should be removed too. also I suggested to remove {{org.jboss.resteasy.resteasy-atom-provider}} from JAX-RS, but it is valid dependency, could you please re-add it? Again I'm sorry about these inaccuracies.	1	1	https://issues.jboss.org/browse/JBEAP-1667
9655	JBEAP-1604	Sounds to me like your asking for certain HHH fixes to be included in EAP 7. I think that the individual (JBEAP-1596, JBEAP-1460, JBEAP-1595) jiras should be assigned to someone that is doing the back porting of fixes for EAP 7. [~gbadner] should this jira be assigned to you? Or directly to David Stahl?	1	1	https://issues.jboss.org/browse/JBEAP-1604
9656	JBEAP-1604	All of the "incorporates" links have already been fixed in Hibernate ORM 5.0.2, so there is nothing listed that needs to be backported; they will already be fixed when updated to Hibernate ORM 5.0.3.Final. Scott Marlow, I have not been involved in getting new versions incorporated. Should I be? Let's discuss offline	0	0	https://issues.jboss.org/browse/JBEAP-1604
9657	JBEAP-1360	[~martyn-taylor] You're right. This was issue with the test. I've fixed it in the test suite. There is a question: I've tried to remove backup connector from configuration of core bridge so there was just connector to live server. If core bridge is configured with ha=true, should it receive live/backup pair topology and failover to backup without the need to have static connector to backup configured? Decreasing priority as reported issue is not valid.	1	1	https://issues.jboss.org/browse/JBEAP-1360
9658	JBEAP-1360	[~jmesnil] Jeff. When the connectors are set up, the ActiveMQServerService is doing an address resolution for loopback interfaces. This is causing problems when configuring static connectors. In order for Artemis to assume 2 TransportConfigurations are the same, the params must be equal. Since, WildFly is replacing the host="localhost.localdomain" when the server bind address is 127.0.0.1. Configuring a static connector with 127.0.0.1 means that the resulting TransportConfigurations are different. See above. This is causing failover to stop working on bridges with single static connector configured. Is there any reason for doing this? I am not sure if this is a bug in the WildFly code or if there is some policy to always use "localhost.localdomain"? {code:title=ActiveMQServerService.java|borderStyle=solid} // resolve the host name of the address only if a loopback address has been set if (sa.getAddress().isLoopbackAddress()) { host = NetworkUtils.canonize(sa.getAddress().getHostName()); } else { host = NetworkUtils.canonize(sa.getAddress().getHostAddress()); } {code} There is an additional problem here, I think either through the management interface or the tests where an additonal property is being set on the static connector "localAddress".	1	1	https://issues.jboss.org/browse/JBEAP-1360
9659	JBEAP-1360	[~jmesnil][~martyn-taylor]] Thanks for looking at this! I understand this way is correct. My concern is usability as it's not intuitive. I've never tried to even set {{local-destination-outbound-socket-binding}} with EAP 6/7 and did not know its existence. Is there a way to make it simple and use just socket-binding for remote-connector in live server?	1	1	https://issues.jboss.org/browse/JBEAP-1360
9660	JBEAP-1360	[~jmesnil] What do you think of the changes? It's usability problem.	1	1	https://issues.jboss.org/browse/JBEAP-1360
9661	JBEAP-1360	[~jmesnil] Hi Jeff, I can see that this was changed from bug to feature request. From usability point of view this seems to be bug. wdyt?	0	0	https://issues.jboss.org/browse/JBEAP-1360
9662	JBEAP-1291	The DR13 installer still contains checks for services running on conflicting ports, even though this enhancement obsoletes the checks. Possibly create a new JIRA?	0	0	https://issues.jboss.org/browse/JBEAP-1291
9663	JBEAP-1281	Hi Petr, Which VM are you testing this on? I'm testing on W10 and in some cases, I do see the path reselection; however, complications with aesh make it difficult to test. I believe there are some conditions on windows in which we will see this issue reproduced, but not all. We need to narrow down the reproduction procedure. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBEAP-1281

#	Requirement	Questions	Answered	Successfully Answered	URL
9664	JBEAP-1281	Seems there is still something wrong, in some cases installer doesn't offer reselecting target and instead ends with {{XML not written successfully}} Checked inputs: PASS 1) Unwritable location provided: G:\auto.xml or /root/auto.xml PASS 2) Directory provided: C:\tmp or /home/jbliznak PASS 3) Relative location: auto.xml PASS 4a) Existing file on local filesystem: C:\test\existing.xml or /home/jbliznak/existing.xml FAIL 4b) Existing file (not writable) on local filesystem C:\test\existing-nowritable.xml or /home/jbliznak/existing-nowritable.xml 5) Writing to network drive (windows): FAIL 5a) new file H:\home\hudson\jbliznak\users-tmp\auto.xml PASS 5b) Existing file H:\home\hudson\jbliznak\users-tmp\existing.xml FAIL 5c) Existing file (not writable) H:\home\hudson\jbliznak\users-tmp\existing-nowritable.xml ad 4b) this issue is for both, Windows and Linux platforms. 5c) seems to be the same problem - missing check after user select to overwrite. {code:java} [jbliznak@erebos-rh installer]$ sudo touch /home/jbliznak/existing-nowritable.xml [jbliznak@erebos-rh installer]$ ll /home/jbliznak/existing-nowritable.xml -rw-r--r--. 1 root root 0 29. bře 15.55 /home/jbliznak/existing-nowritable.xml ..... Would you like to generate an automatic installation script and properties file? (y/n) [n]: y Select path for the automatic installation script: [/home/jbliznak/test/EAP-7.0.0/auto.xml] /home/jbliznak/existing-nowritable.xml File already exists, would you like to overwrite it? (y/n) [n]: y XML not written successfully. [Console installation done] [jbliznak@erebos-rh installer]$ cat /home/jbliznak/existing-nowritable.xml [jbliznak@erebos-rh installer]$ {code} {code:java} Would you like to generate an automatic installation script and properties file? (y/n) [n]: y Select path for the automatic installation script: [C:\Users\Administrator\Deskt op\.\test-console\EAP-7.0.0\auto.xml] H:\home\hudson\users-tmp\jbliznak\existing-nowritable.xml File already exists, would you like to overwrite it? (y/n) [n]: y XML not written successfully. [Console installation done] {code} ad 5a) [~thauser] you wrote in JBEAP-3580, that you were unable to write the file even when checks say "OK to write", [~pkremens] on the other hand was able to write it. I am using the very same command he provided and I am also unable to write the file {code:java} Select path for the automatic installation script: [C:\Users\Administrator\Desktop\.\test-console\EAP-7.0.0\auto.xml] H:\home\hudson\users-tmp\jbliznak\auto.xml Cannot write to path, please choose another location. {code} This is something strange, I am able to create/edit files/dirs on that network drive, but installer says no in this case (although it is ok with 5b). In any case, would it be possible to extend the check and offer the reselection if the file was unable to write (although other checks claimed it should be writable)?	1	1	https://issues.jboss.org/browse/JBEAP-1281
9665	JBEAP-1093	This functionality works correctly in standalone.sh and standalone.ps1. This functionality don't work on bat script. Is it correct behaviour? If yes, this should be mentioned in documentation (documentation jira should be created now).	1	1	https://issues.jboss.org/browse/JBEAP-1093
9666	JBEAP-1093	[~jpallich] or [~ppenicka]: Can you correct the Target Release on this to 7.1.0?	0	0	https://issues.jboss.org/browse/JBEAP-1093
9667	JBEAP-1006	Yes that's it. Thanks [~jmesnil]!	0	0	https://issues.jboss.org/browse/JBEAP-1006
9668	JBEAP-760	Hi Pavel, I couldn't reproduce the problem, can you elaborate on how to reproduce it ?	1	0	https://issues.jboss.org/browse/JBEAP-760
9669	JBEAP-760	Can this issue be closed ?	0	0	https://issues.jboss.org/browse/JBEAP-760
9670	JBEAP-732	Can you guys do some proper clean up? Which version it affects, which version it's fixed, etc.	0	0	https://issues.jboss.org/browse/JBEAP-732
9671	JBEAP-523	[~rafabene]: Is this an issue with the QS tools Maven plugin? I see that Darran replaced the private API with the public API in this quickstart on Mar 13, 2014 with this commit: https://github.com/jboss-developer/jboss-eap-quickstarts/commit/9be73458bfc1cd909a19f5f1c175fbd0fe5b8160	1	1	https://issues.jboss.org/browse/JBEAP-523
9672	JBEAP-523	[~dlofthouse] or [~darranl]: See Rafael's comment above. You fixed this quickstart over a year ago. Do you know why these classes are still marked as private?	1	1	https://issues.jboss.org/browse/JBEAP-523
9673	JBEAP-523	[~bdaw]: There is a long history related to the use of private APIs in this quickstart and I thought it was resolved for JBoss EAP 6.3.1. These were the issues for this quickstart and as you can see, they are marked resolved. * https://issues.jboss.org/browse/EAP6-69 * https://issues.jboss.org/browse/EAP6-136 See [~rafabene]'s comment above. Do you know why these classes are still marked private?	1	1	https://issues.jboss.org/browse/JBEAP-523
9674	JBEAP-523	I took a look and I saw that `org.jboss.as.core.security` is private on the module `org.jboss.as.core-security` but `org.jboss.as.core.security.api` from `org.jboss.as.core-security-api` module is marked as "deprecated". [~sgilda]. Does QSTools still complaining about those imports? From my perspective it looks like org.jboss.as.security.api.* is a valid import	1	1	https://issues.jboss.org/browse/JBEAP-523
9675	JBEAP-459	Shouldn't those be returned by /subsystem=messaging-activemq/server=default:read-children-names(child-type=ha-policy) ? Which is what CLI tab-completion relies upon in this case.	0	0	https://issues.jboss.org/browse/JBEAP-459
9676	JBEAP-385	Petr, can you close this one as a duplicate? It's older than JBEAP-2529 but the other one has more data on it.	1	0	https://issues.jboss.org/browse/JBEAP-385
9677	JBEAP-235	Hi [~swd847], I have successfully verified support for options C, L, N and S (their respective form in Undertow predicate language). Alhough I am in trouble to verify options E and NE. Regarding to E: I have tried to create rewrite rule where I first define some env variable and then use that variable to determine what kind of rewrite should be used. My particular expression in expression-filter looks like this: {code} set(attribute='%{r,myvar}',value='myval'); equals('${ENV:myvar}',myval)->rewrite(/jboss-helloworld/second.html) else rewrite(/jboss-helloworld/third.html) {code} Output of this should be rewrite to 'second.html' page but unfortunatelly it seems that 'myvar' is not set. I think I am missing something here - did I misunderstand how this should work? Regarding to NE: does that mean that user can specify whether not to use URI escaping rules only when using rewrite handler, not in expression in predicate logic? Could you please provide more info for me? Thanks	1	1	https://issues.jboss.org/browse/JBEAP-235
9678	JBEAP-235	Thank you [~swd847] for your answer. Regarding to E: unfortunatelly I am still not able to make it working. When I try with {code}${r, myvar}{code} I get following error: {code} [standalone@localhost:9990 /] /subsystem=undertow/configuration=filter/expression-filter=envvar:write-attribute(name=expression,value="set(attribute='%{r, myvar}', value='myval'); equals('${r, myvar}','myval')->rewrite('/jboss-helloworld/second.html') else rewrite('/jboss-helloworld/third.html')") { "outcome" => "failed", "failure-description" => "WFLYCTL0211: Cannot resolve expression 'set(attribute='%{r, myvar}', value='myval'); equals('${r, myvar}','myval')->rewrite('/jboss-helloworld/second.html') else rewrite('/jboss-helloworld/third.html')'", "rolled-back" => true } {code} Well, when I use {{%}} instead of {{$}} then command succeed (although guess this is not correct input): {code} /subsystem=undertow/configuration=filter/expression-filter=envvar:write-attribute(name=expression,value="set(attribute='%{r, myvar}', value='myval'); equals('%{r, myvar}','myval')->rewrite('/jboss-helloworld/second.html') else rewrite('/jboss-helloworld/third.html')") {code} - after reload this still is not working -> still redirected to {{third.html}}. Could you please tell me what else am I doing wrong? Thanks. Regarding to NE: ok. I have checked that now URL encoding does not apply when using in rewrite predicate so user can do it manually.	1	1	https://issues.jboss.org/browse/JBEAP-235

#	Requirement	Questions	Answered	Successfully Answered	URL
9679	JBEAP-235	Hi [~swd847], I've just realized I have marked different Stuart Douglas in previous comment. Therefore now hopefully this will ping you - could you please look and see what am I doing wrong in previous comment regarding to setting own servlet request attribute? Thank you.	1	1	https://issues.jboss.org/browse/JBEAP-235
9680	JBEAP-176	[~jmesnil] After the long discussion in the cloned WFLY jira, I assume the only way to fix would be adding javax.jms.api to global modules. Is this correct, and is this a desired thing in Wildfly? Should this be discussed with someone else from Wildfly/EAP team before fixing or deciding to keep it in the backlog?	1	1	https://issues.jboss.org/browse/JBEAP-176
9681	JBEAP-143	What's the best component for this? Productization? RPM?	0	0	https://issues.jboss.org/browse/JBEAP-143
9682	JBEAP-142	What's the best component for this? Productization? RPM?	0	0	https://issues.jboss.org/browse/JBEAP-142
9683	JBEAP-117	This is reading like a feature request so as Brian says this is the wrong place - can we get this one closed?	0	0	https://issues.jboss.org/browse/JBEAP-117
9684	JBEAP-62	This was intended for JBoss EAP 6.1 as the comment above states. Can it be moved?	1	1	https://issues.jboss.org/browse/JBEAP-62
9685	JBEAP-52	Is it any workaround for this issue? We are migrating to jboss 7.2 and we have a big number of unicast remote objects, so it's really critical for us.	1	1	https://issues.jboss.org/browse/JBEAP-52
9686	JBEAP-31	How is this class packaged in the project? Are you sure it is not packaged twice?	1	1	https://issues.jboss.org/browse/JBEAP-31
9687	JBEAP-31	Do you have a reproducer for this issue?	0	0	https://issues.jboss.org/browse/JBEAP-31
9688	JBPAPP-11240	[~raggz] I can see that this is not built as roll-up patch as we usually do for HornetQ. Could you built it as roll-up patch as we did for JBPAPP-11226, please?	0	0	https://issues.jboss.org/browse/JBPAPP-11240
9689	JBPAPP-11214	multithreaded application. I tried running the reproducer against EAP 5.1.2 and it worked correctly there. I tried "patching" 5.0.1 by using jbosscx-client.jar and jboss-common-jdbc-wrapper.jar from a 5.1.2.GA distribution and it worked correctly too. I have a heavy suspicion that https://issues.jboss.org/browse/JBPAPP-3970 is the cause -- it needs to be backported to this patch too. The description of JBPAPP-3970/JBAS-7811 fits my testing logs very well: From the server log - two client threads get assigned connections with the same hashCode, so it is stored only once in the WrapperDataSourceService. {noformat} 10:01:51,723 DEBUG [WrapperDataSourceService] Created Connection proxy for invoker=jboss:service=invoker,type=jrmp, targetName=jboss.jca:service=DataSourceBinding,name=GlobalDS, cacheID=912805409 10:01:51,723 DEBUG [WrapperDataSourceService] Created Connection proxy for invoker=jboss:service=invoker,type=jrmp, targetName=jboss.jca:service=DataSourceBinding,name=GlobalDS, cacheID=912805409 {noformat} after that, one client uses the connection, then he closes it. The other client then tries to use it, but the connection with that ID is already closed. {noformat} Exception in thread "Thread-1" java.lang.RuntimeException: java.lang.reflect.UndeclaredThrowableException at AccessTest$1.run(AccessTest.java:48) Caused by: java.lang.reflect.UndeclaredThrowableException at com.sun.proxy.$Proxy1.close(Unknown Source) at AccessTest$1.run(AccessTest.java:41) Caused by: java.lang.IllegalAccessException: Failed to find connection: 912805409 at org.jboss.resource.adapter.jdbc.remote.WrapperDataSourceService.invoke(WrapperDataSourceService.java:213) at sun.reflect.GeneratedMethodAccessor256.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:855) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:422) at sun.reflect.GeneratedMethodAccessor255.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:303) at sun.rmi.transport.Transport$1.run(Transport.java:159) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:155) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:535) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run0(TCPTransport.java:790) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:649) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:895) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:918) at java.lang.Thread.run(Thread.java:662) at sun.rmi.transport.StreamRemoteCall.exceptionReceivedFromServer(StreamRemoteCall.java:255) at sun.rmi.transport.StreamRemoteCall.executeCall(StreamRemoteCall.java:233) at sun.rmi.server.UnicastRef.invoke(UnicastRef.java:142) at org.jboss.invocation.jrmp.server.JRMPInvoker_Stub.invoke(Unknown Source) at	0	0	https://issues.jboss.org/browse/JBPAPP-11214
9690	JBPAPP-11127	What version of postgres plus are we talking about? Wasn't support for that first added in EAP 6?	0	0	https://issues.jboss.org/browse/JBPAPP-11127
9691	JBPAPP-11127	Guessing off your HHH-8349 link: If you're talking about pgplus 9.2, support for that wasn't added until ORM 4.2 and EAP 6 (HHH-8175). Can this be closed?	0	0	https://issues.jboss.org/browse/JBPAPP-11127
9692	JBPAPP-11073	I have deployed the tag HornetQ_2_2_28_Final I tried to upload it.. but for some reason the poms are not being accepted. This is an ancient project now... using ant... Would productization have better poms to this? or do I need to fix this?	0	0	https://issues.jboss.org/browse/JBPAPP-11073
9693	JBPAPP-11073	Could you be more specific on the issue you're meeting ? What error message do you get ?	0	0	https://issues.jboss.org/browse/JBPAPP-11073
9694	JBPAPP-10981	[~huwang] [~pferraro] I've tagged it. https://svn.jboss.org/repos/jbosscache/core/tags/3.2.12.GA/ Do you guys require maven artifacts to be uploaded as well or you're building it yourself in brew? If later I'd rather not release it in maven as I don't want people to see any activity around it (it is not maintained anymore and users are encouraged to use Infinispan)	0	0	https://issues.jboss.org/browse/JBPAPP-10981
9695	JBPAPP-10959	So far I built the following: jboss-seam2-2.2.7.EAP5-2.ep5.el6 jboss-seam2-2.2.7.EAP5-1.ep5.el6 rhel4 build is in WIP.	0	0	https://issues.jboss.org/browse/JBPAPP-10959
9696	JBPAPP-10916	Hi Marek, how about the 3.3.1.SP4 tag?	1	1	https://issues.jboss.org/browse/JBPAPP-10916
9697	JBPAPP-10916	Hi Brian, thanks for the SP4 tag. But I fond a problem, there is SP3 zip file in SP4 tag: http://anonsvn.jboss.org/repos/richfaces/tags/3.3.1.SP4/docs/ which is same in the SP3 tag: http://anonsvn.jboss.org/repos/richfaces/tags/3.3.1.SP3/docs/ Is that OK?	0	0	https://issues.jboss.org/browse/JBPAPP-10916
9698	JBPAPP-10911	[~navssurtani] thanks for details, I see commits 18199 and 18187 linked with this JIRA. Will you provide documentation for customers how to enable this feature ?	1	1	https://issues.jboss.org/browse/JBPAPP-10911
9699	JBPAPP-10893	So, based on the start script mentioned, the way to run EAP as a service on Solaris (without using the scipt itself) is something like: * Define JBOSS_HOME, SERVER_NAME, JBOSS_OPTIONS, JBOSS_USER, JBOSS_CONSOLE AND JAVA_HOME. * Ensure the user has write permission in these locations. * Stop any already running EAP processes * Start EAP with: ${JBOSS_HOME}/bin/run.sh -c ${SERVER_NAME} ${JBOSS_OPTIONS} Is there more to it than that?	1	1	https://issues.jboss.org/browse/JBPAPP-10893

#	Requirement	Questions	Answered	Successfully Answered	URL
9700	JBPAPP-10893	Note to self: Topic 3758 has good instructions for doing this in EAP 6. Would they apply to 5?	0	0	https://issues.jboss.org/browse/JBPAPP-10893
9701	JBPAPP-10893	Ivo; I have no such script in my EAP 5.2 zip installation (From CSP: JBEAP-5.2.0/jboss-eap-5.2.0.zip). Is it contained in one of the other distribution archives?	0	0	https://issues.jboss.org/browse/JBPAPP-10893
9702	JBPAPP-10889	Hi Bartosz, Could you elaborate? If you're talking about communicating between Remoting 2 on EAP 5 and Remoting 3 on EAP 6, then that's not likely ever to work. If you're talking about using RMI, then it's not a Remoting issue. Thanks, Ron	1	0	https://issues.jboss.org/browse/JBPAPP-10889
9703	JBPAPP-10847	Dennis, is it possible to get some info on how this race issue was fixed for inclusion in the 5.3.0 Release Notes? Also, is the text I provided in the Release Notes Text field accurate?	1	0	https://issues.jboss.org/browse/JBPAPP-10847
9704	JBPAPP-10839	Hi Tom, there is missing link to git/svn from where one-off patch was built. I need it so I can run HornetQ test suite and check that all fixes are present. Based on the naming convention described in [1] could you remove attachment JBPAPP-10839-signed.zip, please? Return patch to me as soon as it's done. Thanks, Mirek [1] https://docspace.corp.redhat.com/docs/DOC-151974	0	0	https://issues.jboss.org/browse/JBPAPP-10839
9705	JBPAPP-10816	Can you provide the stacktrace and/or a sample app to reproduce the issue?	1	0	https://issues.jboss.org/browse/JBPAPP-10816
9706	JBPAPP-10720	How do you launch your build so this issue happens ?	0	0	https://issues.jboss.org/browse/JBPAPP-10720
9707	JBPAPP-10710	Will this be included in EAP 5.3 release? also, will a new SOA-P 5 version be release along with EAP 5.3?	1	0	https://issues.jboss.org/browse/JBPAPP-10710
9708	JBPAPP-10666	Tom, it doesn't look like this issue is past EAP QE yet. Do we have more work to do on it? Thanks, Rick	1	1	https://issues.jboss.org/browse/JBPAPP-10666
9709	JBPAPP-10666	[~raggz]: This one has been sitting for a long time, and it only recently got turned over to QA. When I check the case, it's been closed with no apparent need for a patch. Can you confirm?	0	0	https://issues.jboss.org/browse/JBPAPP-10666
9710	JBPAPP-10645	Would you create 2.2.3.GA tag to include the patch?	1	1	https://issues.jboss.org/browse/JBPAPP-10645
9711	JBPAPP-10623	I am returning this one-off because of following issues: 1. Incorrect names of jars - it should be just jboss-messaging.jar and jboss-messaging-client.jar. If you would like to highlight that these jars are signed, try zipping them both into JBPAPP-10623-signed.zip or so. 2. Incorrect MANUAL INSTALL INSTRUCTIONS - the path should be without 'server' part, just $JBOSS_HOME/common/lib and $JBOSS_HOME/client. Could you please fix these issues? Thank you, Nikoleta	1	1	https://issues.jboss.org/browse/JBPAPP-10623
9712	JBPAPP-10623	It looks good now, thanks. Could you please hand over the jira back to QA?	1	1	https://issues.jboss.org/browse/JBPAPP-10623
9713	JBPAPP-10610	can this be uploaded soon?	1	1	https://issues.jboss.org/browse/JBPAPP-10610
9714	JBPAPP-10600	Does anyone watching this issue have some raw documentation with which we (ECS) can make a start? Without a starting point, this simply can't be delivered.	0	0	https://issues.jboss.org/browse/JBPAPP-10600
9715	JBPAPP-10596	testStatefulBank21() is the only test failure on this TestCase. testRemoteBindingInterceptorStack() passes locally and has been passing on Boo as well. Testcase: testStatefulBank21 took 0.046 sec FAILED null expected:<[_CREATED]> but was:<[]> junit.framework.ComparisonFailure: null expected:<[_CREATED]> but was:<[]> at org.jboss.ejb3.test.bank.unit.BankDeploymentDescriptorTestCase.testStatefulBank21(BankDeploymentDescriptorTestCase.java:248) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:23) at junit.extensions.TestSetup.run(TestSetup.java:27)	1	1	https://issues.jboss.org/browse/JBPAPP-10596
9716	JBPAPP-10557	Why do we have one-off patch for the same issue for both EAP 600 and EAP 601? I think we should push to use latest EAP and so only one one-off patch should be processed - the one for EAP 601.	0	0	https://issues.jboss.org/browse/JBPAPP-10557
9717	JBPAPP-10494	Could you please change MANUAL INSTALL INSTRUCTIONS: ? This step also remove hibernate-infinispan-4.1.3.Final-redhat-1.jar, but shouldn't: {noformat} Backup and remove the following files: $JBOSS_HOME/modules/org/hibernate/main/hibernate*.jar {noformat} Thanks	1	1	https://issues.jboss.org/browse/JBPAPP-10494
9718	JBPAPP-10444	There is another commit in Branch_2_0 of the microcontainer which relates to JBPAPP-6020 which is still open. Should this commit be included in 2.0.11 upgrade or not?	1	1	https://issues.jboss.org/browse/JBPAPP-10444
9719	JBPAPP-10444	Hi Ivo, The new source tag depends on google-guice-1.0, but google-guice-2.0 has been included in our EAP 5.2.0 release and compare to 1.0, 2.0 has great change, I have tried to patch sources but failed. Is it possible for you to tag new source which depends on google-guice-2.0? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-10444
9720	JBPAPP-10418	Attention: Miroslav The installation of libaio is already documented in the EAP Installation Guide, as can be confirmed here at [1]. I'm open to suggestions as to how these instructions or the HornetQ User Guide might be improved with reference to libaio. On re-reading the Installation Guide it seems that there is no direct mention of the native libraries in the HornetQ installation section. Is this what you want clarified? [1] https://access.redhat.com/knowledge/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/Installation_Guide/index.html#Install_Native_Components	1	1	https://issues.jboss.org/browse/JBPAPP-10418

#	Requirement	Questions	Answered	Successfully Answered	URL
9721	JBPAPP-10349	Test requires: java.net.preferIPv4Stack=true. Not sure why it failed in big run, since this option seems to be included in setup. In case this is not included: {noformat} ERROR [org.jgroups.protocols.UDP] failed sending message to null (53 bytes) java.lang.Exception: dest=/233.54.54.54:54545 (56 bytes) at org.jgroups.protocols.UDP._send(UDP.java:374) at org.jgroups.protocols.UDP.sendToAllMembers(UDP.java:302) at org.jgroups.protocols.TP.doSend(TP.java:1478) at org.jgroups.protocols.TP.send(TP.java:1468) at org.jgroups.protocols.TP.down(TP.java:1186) at org.jgroups.protocols.PING.sendMcastDiscoveryRequest(PING.java:278) at org.jgroups.protocols.PING.sendGetMembersRequest(PING.java:259) at org.jgroups.protocols.Discovery$PingSenderTask$1.run(Discovery.java:407) at org.jgroups.util.TimeScheduler$RobustRunnable.run(TimeScheduler.java:196) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:471) at java.util.concurrent.FutureTask$Sync.innerRunAndReset(FutureTask.java:351) at java.util.concurrent.FutureTask.runAndReset(FutureTask.java:178) at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.access$301(ScheduledThreadPoolExecutor.java:178) at java.util.concurrent.ScheduledThreadPoolExecutor$ScheduledFutureTask.run(ScheduledThreadPoolExecutor.java:293) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:603) at java.lang.Thread.run(Thread.java:722) Caused by: java.io.IOException: Invalid argument at java.net.PlainDatagramSocketImpl.send(Native Method) at java.net.DatagramSocket.send(DatagramSocket.java:676) at org.jgroups.protocols.UDP._send(UDP.java:358) ... 16 more {noformat} Im going to resolve this issue since test does not fail anymore.	1	1	https://issues.jboss.org/browse/JBPAPP-10349
9722	JBPAPP-10294	Is this a Documentation issue? If so please add that as the Component. Also, if it is being marked as Done then it must have a Fix Version....	1	1	https://issues.jboss.org/browse/JBPAPP-10294
9723	JBPAPP-10293	The tag is here: - https://svn.jboss.org/repos/jbossas/projects/security/security-negotiation/tags/security-negotiation-2.1.2.GA Who is next for this Jira to handle the build?	0	0	https://issues.jboss.org/browse/JBPAPP-10293
9724	JBPAPP-10221	I just noticed that yesterday was the deadline for component upgrades. Can I get a new Remoting 2 in?	0	0	https://issues.jboss.org/browse/JBPAPP-10221
9725	JBPAPP-10211	1. .classpath has some ear ref - is this required, used in test? 2. .classpath has refs to local copy of jboss, this should be changed to $JBOSS_HOME or something 3. class names are wrong binary(EJB3FailoverServices-1.0.0-SNAPSHOT.jar) bean refers to HelloServiceRemote, however there is no such class, there is HelloServiceRemote2(EJB3FailoverPublic2-1.0.0-SNAPSHOT.jar) 4. instructions in Readme.txt state that EJB3FailoverPublic2-1.0.0-SNAPSHOT.jar should be deployed. Why interface is deployed, not bean?(or both :)) 5. IPs in Unit test class are kind of wrong.	0	0	https://issues.jboss.org/browse/JBPAPP-10211
9726	JBPAPP-10211	What is more important - I wonder how other places which have the "call-by-value" flags handle this.	0	0	https://issues.jboss.org/browse/JBPAPP-10211
9727	JBPAPP-10211	Noted. Initially CTR has been "reinitialized" after each lookup, but I wasnt sure if this is correct - hence the check for "success". WRT to issue, Im still unable to replicate locally failure which may be caused by this fix: {noformat} org.jboss.test.cluster.defaultcfg.ejb3.ustxsticky.test.UserTransactionStickyUnitTestCase.testMoreBeansInTransaction (from org.jboss.test.cluster.defaultcfg.ejb3.ustxsticky.test.UserTransactionStickyUnitTestCase(Default-udp)) Failing for the past 3 builds (Since #27) Took 28 ms. Error Message Could not dereference object Stacktrace javax.naming.NamingException: Could not dereference object [Root exception is java.lang.RuntimeException: Unreachable?: Service unavailable.] at org.jnp.interfaces.NamingContext.getObjectInstanceWrapFailure(NamingContext.java:1551) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:869) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:713) at javax.naming.InitialContext.lookup(InitialContext.java:422) at org.jboss.test.cluster.defaultcfg.ejb3.ustxsticky.test.UserTransactionStickyUnitTestCase.test(UserTransactionStickyUnitTestCase.java:83) at org.jboss.test.cluster.defaultcfg.ejb3.ustxsticky.test.UserTransactionStickyUnitTestCase.testMoreBeansInTransaction(UserTransactionStickyUnitTestCase.java:67) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) Caused by: java.lang.RuntimeException: Unreachable?: Service unavailable. at org.jboss.aspects.remoting.ClusterChooserInterceptor.invoke(ClusterChooserInterceptor.java:241) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.ejb3.proxy.impl.remoting.IsLocalProxyFactoryInterceptor.invoke(IsLocalProxyFactoryInterceptor.java:104) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.ClusteredPojiProxy.invoke(ClusteredPojiProxy.java:79) at $Proxy2.createProxyBusiness(Unknown Source) at org.jboss.ejb3.proxy.impl.objectfactory.session.SessionProxyObjectFactory.createProxy(SessionProxyObjectFactory.java:129) at org.jboss.ejb3.proxy.clustered.objectfactory.session.SessionClusteredProxyObjectFactory.getProxy(SessionClusteredProxyObjectFactory.java:89) at org.jboss.ejb3.proxy.impl.objectfactory.ProxyObjectFactory.getObjectInstance(ProxyObjectFactory.java:161) at javax.naming.spi.NamingManager.getObjectInstance(NamingManager.java:332) at org.jnp.interfaces.NamingContext.getObjectInstance(NamingContext.java:1526) at org.jnp.interfaces.NamingContext.getObjectInstanceWrapFailure(NamingContext.java:1543) {noformat}	0	0	https://issues.jboss.org/browse/JBPAPP-10211
9728	JBPAPP-10211	What is the status of this - Is the patch suitable for testing or are there still issues ?	0	0	https://issues.jboss.org/browse/JBPAPP-10211
9729	JBPAPP-10187	The tag: JBossMessaging_1_4_8_SP9 does not exist in http://anonsvn.jboss.org/repos/messaging/tags/, could you please check it, Howard?	0	0	https://issues.jboss.org/browse/JBPAPP-10187
9730	JBPAPP-10102	Dominik Pospisil <dpospisi@redhat.com> made a comment on [bug 900989|https://bugzilla.redhat.com/show_bug.cgi?id=900989] That would work. But as I said this need to be copied before other tests are executed so I guess this needs to be hacked into testsuite POM file which is something I would not prefer to do. I guess you are much more familiar with the testsuite POMs. Can you do that Ondro?	1	1	https://issues.jboss.org/browse/JBPAPP-10102
9731	JBPAPP-10102	Rostislav Svoboda <rsvoboda@redhat.com> made a comment on [bug 900989|https://bugzilla.redhat.com/show_bug.cgi?id=900989] Is the workaround sufficient ? If yes and you don't want changes in TS, please close this BZ. Btw, how many tests are affected by this BZ, I can't access https://jenkins-jbossqa.rhcloud.com.	1	1	https://issues.jboss.org/browse/JBPAPP-10102

#	Requirement	Questions	Answered	Successfully Answered	URL
9732	JBPAPP-10017	I've tried replicating this situation on my localhost, using the following steps: 1. checkout sources from SVN, using the checkout.log for guidance on revId, etc. 2. apply VMan patch from dist-git, using patches.log for guidance 3. created proxy to jb-eap-6-rhel-6-build MEAD Maven repository on my local repository manager, disabling local caching to make it a "pass-through" proxy 4. setup variant settings.xml to <mirror/> to the pass-through proxy, so all I can access are artifacts in the jb-eap-6-rhel-6-build MEAD repository 5. ran `mvn -s ~/.m2/settings-jbeap6.xml dependency:resolve-plugins` which is the command that failed according to the root.log Nothing failed. So, I ran: `mvn -s ~/.m2/settings-jbeap6.xml org.apache.maven.plugins:maven-help-plugin:effective-pom -Doutput=effective.pom.xml` which calculates what the POM looks like after processing inheritance, interpolation, BOM imports, etc. and looked for the version of the project-info-reports plugin in that effective POM. Here's what I found: {noformat} [...] <pluginManagement> [...] <plugin> <artifactId>maven-project-info-reports-plugin</artifactId> <version>2.2-redhat-1</version> </plugin> [...] </pluginMangement> <plugins> [...] <plugin> <artifactId>maven-site-plugin</artifactId> <version>3.0-beta-3-redhat-1</version> <executions> <execution> <id>default-site</id> <phase>site</phase> <goals> <goal>site</goal> </goals> <configuration> <outputDirectory>/Users/jdcasey/workspace/redhat/temp/neethi-3.0.2/target/site</outputDirectory> <reportPlugins> <reportPlugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-project-info-reports-plugin</artifactId> </reportPlugin> </reportPlugins> </configuration> </execution> <execution> <id>default-deploy</id> <phase>site-deploy</phase> <goals> <goal>deploy</goal> </goals> <configuration> <outputDirectory>/Users/jdcasey/workspace/redhat/temp/neethi-3.0.2/target/site</outputDirectory> <reportPlugins> <reportPlugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-project-info-reports-plugin</artifactId> </reportPlugin> </reportPlugins> </configuration> </execution> </executions> <configuration> <outputDirectory>/Users/jdcasey/workspace/redhat/temp/neethi-3.0.2/target/site</outputDirectory> <reportPlugins> <reportPlugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-project-info-reports-plugin</artifactId> </reportPlugin> </reportPlugins> </configuration> </plugin> [...] </plugins> {noformat} In short, the only version for the project-info-reports plugin I could find was 2.2-redhat-1, and I've verified that this doesn't depend on joda-time. Nor does the plugin declaration mention any additional dependency on joda-time. I'm not sure how this project ever came to depend on a different version of the project-info-reports plugin...is it possible the jboss-parent was rebuilt and had this plugin added?	1	1	https://issues.jboss.org/browse/JBPAPP-10017
9733	JBPAPP-9923	Sep 13 17:24:26 EDT 2012, how are you?> but was:<null> at org.jboss.test.ee5client.unit.AppClientUnitTestCase.test1(AppClientUnitTestCase.java:72) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) Standard Error java.lang.reflect.InvocationTargetException at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.jboss.ejb3.client.ClientContainer.invokeMain(ClientContainer.java:269) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.jboss.ejb3.client.ClientLauncher.launch(ClientLauncher.java:301) at org.jboss.ejb3.client.ClientLauncher.launch(ClientLauncher.java:174) at org.jboss.test.ee5client.unit.AppClientUnitTestCase.test1(AppClientUnitTestCase.java:64) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at junit.framework.TestCase.runTest(TestCase.java:164) at junit.framework.TestCase.runBare(TestCase.java:130) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:120) at junit.framework.TestSuite.runTest(TestSuite.java:230) at junit.framework.TestSuite.run(TestSuite.java:225) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.extensions.TestSetup.run(TestSetup.java:25) at junit.framework.TestSuite.runTest(TestSuite.java:230) at junit.framework.TestSuite.run(TestSuite.java:225) at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.run(JUnitTestRunner.java:420) at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.launch(JUnitTestRunner.java:911) at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.main(JUnitTestRunner.java:768) Caused by: java.lang.IllegalStateException: Container jboss.j2ee:ear=ee5client-test.ear,jar=ee5client-test.jar,name=HelloWorldBean,service=EJB3,VMID=1f556e0744b10064:28a024c7:139c178713d:-7899 is not registered at org.jboss.ejb3.Ejb3Registry.getContainer(Ejb3Registry.java:167) at org.jboss.ejb3.remoting.IsLocalInterceptor.invoke(IsLocalInterceptor.java:73) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.PojiProxy.invoke(PojiProxy.java:62) at $Proxy16.invoke(Unknown Source) at org.jboss.ejb3.proxy.impl.handler.session.SessionProxyInvocationHandlerBase.invoke(SessionProxyInvocationHandlerBase.java:188) at $Proxy15.sayHelloTo(Unknown Source) at org.jboss.test.ee5client.client.HelloWorldClient.main(HelloWorldClient.java:85) ... 33 more {noformat}	0	0	https://issues.jboss.org/browse/JBPAPP-9923
9734	JBPAPP-9905	This seems like a duplicate issue JBPAPP-7888 Was there any code change in the tests due to this jira? I still see these tests failing when the proper crypto files are not updated https://jenkins.mw.lab.eng.bos.redhat.com/hudson/view/EAP5/view/AS-TestSuite/job/eap5.2Winjdk6/15/jdk=java16_default,label=dev101-W2k8r2-x86_64/testReport/ I am reopening this jira ...if this jira was not associated with any code change/improvement please resolve it as duplicate for ease of tracking.	1	1	https://issues.jboss.org/browse/JBPAPP-9905
9735	JBPAPP-9889	Can you sign the jars and upload it?	1	1	https://issues.jboss.org/browse/JBPAPP-9889
9736	JBPAPP-9889	Mirek, it was fixed already.. I just didn't resolve it.. sorry about that. https://github.com/hornetq/hornetq/commit/1641d180c9639d503cf889c6b02303584dd5c170 still not resolved.. as I'm not sure what's the fix-version it should be in... can you resolve it for me at the proper fix please?	1	1	https://issues.jboss.org/browse/JBPAPP-9889
9737	JBPAPP-9889	Hi Jimmy, Tom, could you sign last jars from Clebert so I can start testing. Thanks, Mirek	1	1	https://issues.jboss.org/browse/JBPAPP-9889
9738	JBPAPP-9889	is https://issues.jboss.org/browse/HORNETQ-1011 in this? https://issues.jboss.org/browse/HORNETQ-1011 seems to suggest it should be but it is not listed above.	1	1	https://issues.jboss.org/browse/JBPAPP-9889

#	Requirement	Questions	Answered	Successfully Answered	URL
9739	JBPAPP-9771	Tom, could you prepare a tag for CP13? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-9771
9740	JBPAPP-9771	Martha, could you please build the JBossTS from the tag mentioned above? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-9771
9741	JBPAPP-9754	{noformat} org.jboss.test.security.test.EJB3SpecUnitTestCase.testUnauthenticatedMethodAccess (from org.jboss.test.securitymgr.test.EJB3SpecUnitTestCase(tests-security-manager)) Failing for the past 1 build (Since #966) Took 0.2 sec. add description Error Message Container jboss.j2ee:jar=security-ejb3.jar,name=SimpleStatelessSessionBean,service=EJB3,VMID=1f556e0744b10064:-3e4b8920:139c43cf0d3:-7f3a is not registered Stacktrace java.lang.IllegalStateException: Container jboss.j2ee:jar=security-ejb3.jar,name=SimpleStatelessSessionBean,service=EJB3,VMID=1f556e0744b10064:-3e4b8920:139c43cf0d3:-7f3a is not registered at org.jboss.ejb3.Ejb3Registry.getContainer(Ejb3Registry.java:167) at org.jboss.ejb3.remoting.IsLocalInterceptor.invoke(IsLocalInterceptor.java:73) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.PojiProxy.invoke(PojiProxy.java:62) at $Proxy5.invoke(Unknown Source) at org.jboss.ejb3.proxy.impl.handler.session.SessionProxyInvocationHandlerBase.invoke(SessionProxyInvocationHandlerBase.java:188) at $Proxy4.invokeUnprotectedMethod(Unknown Source) at org.jboss.test.security.test.EJB3SpecUnitTestCase.assertMethodAccessResults(EJB3SpecUnitTestCase.java:264) at org.jboss.test.security.test.EJB3SpecUnitTestCase.testUnauthenticatedMethodAccess(EJB3SpecUnitTestCase.java:163) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) {noformat}	0	0	https://issues.jboss.org/browse/JBPAPP-9754
9742	JBPAPP-9754	This is last issue which remains open for JBPAPP-9712, is there any resolution or conclusion whether its supposed to be fixed or ignored??	0	0	https://issues.jboss.org/browse/JBPAPP-9754
9743	JBPAPP-9754	Is there a similar log like the one attached, that we can get from a successful run?	1	1	https://issues.jboss.org/browse/JBPAPP-9754
9744	JBPAPP-9736	However Ive managed to fetch log from Jenkins, since its full log and deployment unit is reused its a bit hard to guess which log is correct(probably the second one): {noformat} 2012-08-22 05:08:16,559 DEBUG [org.jboss.system.ServiceCreator] (RMI TCP Connection(453)-127.0.0.1) Created mbean: jboss.jacc:service=jacc,id="vfszip:/var/lib/jenkins/jobs/EAP-5.x-JDK7-latest/workspace/JBPAPP_5_1/testsuite/output/lib/invoker-adaptor-test.ear/dummy.war/",parent="invoker-adaptor-test.ear" 2012-08-22 05:08:16,560 DEBUG [org.jboss.system.ServiceController] (RMI TCP Connection(453)-127.0.0.1) Creating service jboss.jacc:service=jacc,id="vfszip:/var/lib/jenkins/jobs/EAP-5.x-JDK7-latest/workspace/JBPAPP_5_1/testsuite/output/lib/invoker-adaptor-test.ear/dummy.war/",parent="invoker-adaptor-test.ear" 2012-08-22 05:08:16,560 DEBUG [org.jboss.web.WebPermissionMapping] (RMI TCP Connection(453)-127.0.0.1) Qualified url patterns: {/=PatternInfo[pattern=/,type=3,isOverriden=false,qualifiers=[]]} 2012-08-22 05:08:16,561 DEBUG [org.jboss.system.ServiceController] (RMI TCP Connection(453)-127.0.0.1) starting service jboss.jacc:service=jacc,id="vfszip:/var/lib/jenkins/jobs/EAP-5.x-JDK7-latest/workspace/JBPAPP_5_1/testsuite/output/lib/invoker-adaptor-test.ear/dummy.war/",parent="invoker-adaptor-test.ear" 2012-08-22 05:08:16,564 DEBUG [org.jboss.deployers.plugins.deployers.DeployersImpl] (RMI TCP Connection(453)-127.0.0.1) Fully Deployed vfszip:/var/lib/jenkins/jobs/EAP-5.x-JDK7-latest/workspace/JBPAPP_5_1/testsuite/output/lib/invoker-adaptor-test.ear/ 2012-08-22 05:08:16,805 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=0 2012-08-22 05:08:17,811 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=1 2012-08-22 05:08:18,807 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=2 2012-08-22 05:08:19,811 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=3 2012-08-22 05:08:20,811 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=4 2012-08-22 05:08:21,811 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=5 2012-08-22 05:08:22,811 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=6 2012-08-22 05:08:23,807 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=7 2012-08-22 05:08:24,811 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=8 2012-08-22 05:08:25,811 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Sending notification on timer, count=9 2012-08-22 05:08:25,811 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-8) Cancelled timer 2012-08-22 05:08:25,884 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-9) Sending notification on timer, count=0 2012-08-22 05:08:26,887 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-9) Sending notification on timer, count=1 2012-08-22 05:08:27,887 INFO [org.jboss.test.jmx.invoker.InvokerTest] (Timer-9) Sending notification on timer, count=2 2012-08-22 05:08:27,901 DEBUG [org.jboss.jmx.connector.invoker.InvokerAdaptorService] (ThreadPool(4)-1) Failed to notify client, unregistering listener java.rmi.UnmarshalException: Error unmarshaling return header; nested exception is: java.net.SocketTimeoutException: Read timed out at sun.rmi.transport.StreamRemoteCall.executeCall(StreamRemoteCall.java:227) at sun.rmi.server.UnicastRef.invoke(UnicastRef.java:160) at org.jboss.test.jmx.invoker.RMIBadListener_Stub.handleNotification(Unknown Source) at org.jboss.jmx.connector.invoker.InvokerAdaptorService$NotificationListenerDelegate.handleNotification(InvokerAdaptorService.java:350) at sun.reflect.GeneratedMethodAccessor219.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at	1	1	https://issues.jboss.org/browse/JBPAPP-9736

#	Requirement	Questions	Answered	Successfully Answered	URL
9745	JBPAPP-9736	TestSuite/job/eap5.2Winjdk6/3/jdk=java16_default,label=w2k8/testReport/ Stacktrace: Error Message jboss.test:service=InvokerTest is not registered. Stacktrace javax.management.InstanceNotFoundException: jboss.test:service=InvokerTest is not registered. at org.jboss.mx.server.registry.BasicMBeanRegistry.get(BasicMBeanRegistry.java:526) at org.jboss.mx.server.MBeanServerImpl.getAttribute(MBeanServerImpl.java:559) at org.jboss.jmx.connector.invoker.InvokerAdaptorService.invoke(InvokerAdaptorService.java:266) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:138) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:140) at org.jboss.jmx.connector.invoker.SerializableInterceptor.invoke(SerializableInterceptor.java:74) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.invocation.jrmp.server.JRMPProxyFactory.invoke(JRMPProxyFactory.java:180) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:855) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:422) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:303) at sun.rmi.transport.Transport$1.run(Transport.java:159) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:155) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:535) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run0(TCPTransport.java:790) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:649) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) at java.lang.Thread.run(Thread.java:662) Stactrace for SecureJMXInvokerUnitTestCase: Error Message jboss.test:service=InvokerTest,secured=true is not registered. Stacktrace javax.management.InstanceNotFoundException: jboss.test:service=InvokerTest,secured=true is not registered. at org.jboss.mx.server.registry.BasicMBeanRegistry.get(BasicMBeanRegistry.java:526) at org.jboss.mx.server.MBeanServerImpl.getAttribute(MBeanServerImpl.java:559) at org.jboss.jmx.connector.invoker.InvokerAdaptorService.invoke(InvokerAdaptorService.java:266) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:138) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at	1	1	https://issues.jboss.org/browse/JBPAPP-9736
9746	JBPAPP-9736	TestSuite/job/eap5.2Winjdk6/3/jdk=java16_default,label=w2k8/testReport/ Stacktrace: Error Message jboss.test:service=InvokerTest is not registered. Stacktrace javax.management.InstanceNotFoundException: jboss.test:service=InvokerTest is not registered. at org.jboss.mx.server.registry.BasicMBeanRegistry.get(BasicMBeanRegistry.java:526) at org.jboss.mx.server.MBeanServerImpl.getAttribute(MBeanServerImpl.java:559) at org.jboss.jmx.connector.invoker.InvokerAdaptorService.invoke(InvokerAdaptorService.java:266) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:138) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:140) at org.jboss.jmx.connector.invoker.SerializableInterceptor.invoke(SerializableInterceptor.java:74) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.invocation.jrmp.server.JRMPProxyFactory.invoke(JRMPProxyFactory.java:180) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:264) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.invocation.jrmp.server.JRMPInvoker$MBeanServerAction.invoke(JRMPInvoker.java:855) at org.jboss.invocation.jrmp.server.JRMPInvoker.invoke(JRMPInvoker.java:422) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:303) at sun.rmi.transport.Transport$1.run(Transport.java:159) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:155) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:535) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run0(TCPTransport.java:790) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:649) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) at java.lang.Thread.run(Thread.java:662) Stactrace for SecureJMXInvokerUnitTestCase: Error Message jboss.test:service=InvokerTest,secured=true is not registered. Stacktrace javax.management.InstanceNotFoundException: jboss.test:service=InvokerTest,secured=true is not registered. at org.jboss.mx.server.registry.BasicMBeanRegistry.get(BasicMBeanRegistry.java:526) at org.jboss.mx.server.MBeanServerImpl.getAttribute(MBeanServerImpl.java:559) at org.jboss.jmx.connector.invoker.InvokerAdaptorService.invoke(InvokerAdaptorService.java:266) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:157) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:96) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:138) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at	1	1	https://issues.jboss.org/browse/JBPAPP-9736
9747	JBPAPP-9712	Is there a JIRA tracking these two test failures? They have been failing consistently for a while now. org.jboss.test.cluster.multicfg.test.HAJndiAutoDiscoveryTestCase.testAutoDiscovery org.jboss.test.cluster.multicfg.test.HAJndiAutoDiscoveryTestCase.testAutoDiscovery	1	0	https://issues.jboss.org/browse/JBPAPP-9712
9748	JBPAPP-9616	Hi [~jbertram] Would it be possible to get a bit more information about what was causing this exception and how it was fixed? Both this and the original HornetQ JIRA are a little light on for explanatory details. Cheers.	1	1	https://issues.jboss.org/browse/JBPAPP-9616

#	Requirement	Questions	Answered	Successfully Answered	URL
9749	JBPAPP-9572	jboss-cache-pojo-3.0.1-1.ep5.el6 built in Brew from the tag 3.0.1.GA, unsigned jars from repolib uploaded to r.j.o. http://repository.jboss.org/jboss/cache/jbosscache-pojo/3.0.1.GA-brew Committed revision 36971. Local build was successfull with this new version, the JBPAPP_5_1 branch updated to use this new version. Committed revision 113362. Marking as resolved.	1	1	https://issues.jboss.org/browse/JBPAPP-9572
9750	JBPAPP-9503	org.jboss.profileservice.management.KernelBusRuntimeComponentDispatcher.get(KernelBusRuntimeComponentDispatcher.java:119) at org.jboss.profileservice.management.DelegatingComponentDispatcherImpl.get(DelegatingComponentDispatcherImpl.java:83) at org.jboss.aop.Dispatcher.invoke(Dispatcher.java:121) at org.jboss.aspects.remoting.AOPRemotingInvocationHandler.invoke(AOPRemotingInvocationHandler.java:82) at org.jboss.profileservice.remoting.ProfileServiceInvocationHandler.invoke(ProfileServiceInvocationHandler.java:99) at org.jboss.remoting.ServerInvoker.invoke(ServerInvoker.java:967) at org.jboss.remoting.transport.local.LocalClientInvoker.invoke(LocalClientInvoker.java:106) at org.jboss.remoting.Client.invoke(Client.java:2070) at org.jboss.remoting.Client.invoke(Client.java:879) at org.jboss.aspects.remoting.InvokeRemoteInterceptor.invoke(InvokeRemoteInterceptor.java:60) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.remoting.MergeMetaDataInterceptor.invoke(MergeMetaDataInterceptor.java:74) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.security.SecurityClientInterceptor.invoke(SecurityClientInterceptor.java:65) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aop.generatedproxies.AOPProxy$0.get(AOPProxy$0.java) at org.jboss.profileservice.management.client.ManagedPropertyDelegate.getValue(ManagedPropertyDelegate.java:151) at org.jboss.profileservice.management.ManagementViewImpl.updateComponent(ManagementViewImpl.java:1135) at org.jboss.profileservice.remoting.SecurityContainer.invoke(SecurityContainer.java:81) at $Proxy119.updateComponent(Unknown Source) at org.jboss.aop.Dispatcher.invoke(Dispatcher.java:121) at org.jboss.aspects.remoting.AOPRemotingInvocationHandler.invoke(AOPRemotingInvocationHandler.java:82) at org.jboss.profileservice.remoting.ProfileServiceInvocationHandler.invoke(ProfileServiceInvocationHandler.java:99) at org.jboss.remoting.ServerInvoker.invoke(ServerInvoker.java:967) at org.jboss.remoting.transport.socket.ServerThread.completeInvocation(ServerThread.java:791) at org.jboss.remoting.transport.socket.ServerThread.processInvocation(ServerThread.java:744) at org.jboss.remoting.transport.socket.ServerThread.dorun(ServerThread.java:548) at org.jboss.remoting.transport.socket.ServerThread.run(ServerThread.java:234) Caused by: java.lang.NullPointerException at org.jboss.ha.framework.server.managed.OpenChannelsMapper.createMetaValue(OpenChannelsMapper.java:200) at org.jboss.ha.framework.server.managed.OpenChannelsMapper.createMetaValue(OpenChannelsMapper.java:59) at org.jboss.profileservice.management.KernelBusRuntimeComponentDispatcher.get(KernelBusRuntimeComponentDispatcher.java:98) at org.jboss.profileservice.management.DelegatingComponentDispatcherImpl.get(DelegatingComponentDispatcherImpl.java:83) at org.jboss.aop.Dispatcher.invoke(Dispatcher.java:121) at org.jboss.aspects.remoting.AOPRemotingInvocationHandler.invoke(AOPRemotingInvocationHandler.java:82) at org.jboss.profileservice.remoting.ProfileServiceInvocationHandler.invoke(ProfileServiceInvocationHandler.java:99) at org.jboss.remoting.ServerInvoker.invoke(ServerInvoker.java:967) at org.jboss.remoting.transport.local.LocalClientInvoker.invoke(LocalClientInvoker.java:106) at	1	1	https://issues.jboss.org/browse/JBPAPP-9503
9751	JBPAPP-9503	After fixing some problems with a closed channel: {noformat} junit.framework.AssertionFailedError: Channel DefaultPartition-pess-entity has clusterName at org.jboss.test.cluster.channelfactory.managed.ManagedObjectTestUtil.validateOpenChannelValue(ManagedObjectTestUtil.java:160) at org.jboss.test.cluster.channelfactory.managed.ManagedObjectTestUtil.validateOpenChannels(ManagedObjectTestUtil.java:82) at org.jboss.test.cluster.defaultcfg.test.ChannelFactoryManagedObjectsTestCase.testChannelFactory(ChannelFactoryManagedObjectsTestCase.java:173) {noformat} Why is {{DefaultPartition-pess-entity}} closed anyway?	1	1	https://issues.jboss.org/browse/JBPAPP-9503
9752	JBPAPP-9502	Unfortunately, I've got this again in EAP-5.2.0-ER2 build when I've tried to run RESTEasy module jboss-integration-testing against running EAp instance. How to reproduce: svn co https://resteasy.svn.sourceforge.net/svnroot/resteasy/tags/RESTEASY_JAXRS_1_2_1_GA_CP02 resteasy cd resteasy/jboss-integration-testing <EAP_DIR>/bin/run.sh mvn2 installI've got: {code} [ERROR] BUILD ERROR [INFO] -- [INFO] Mojo error occurred: Server returned HTTP response code: 500 for URL: http://localhost:8080/jmx-console/HtmlAdaptor?action=invokeOpByName&name=jboss.system:service%3DMainDeployer&methodName=deploy&argType=java.net.URL&arg0=file:/home/pjanouse/devel/zdroje/RESTEasy-1.x/RESTEASY_JAXRS_1_2_1_GA_CP02/6/jboss-integration-testing/basic-integration-test/target/basic-integration-test.war {code} And on EAP part {code} 10:45:07,600 WARN [HtmlAdaptorServlet] Received request for invokeOpByName with invalid or missing CSRFToken token. This might indicate that someone is trying CSRF attack on your server. 10:45:07,601 ERROR [[HtmlAdaptor]] Servlet.service() for servlet HtmlAdaptor threw exception javax.servlet.ServletException: Operation invoke prohibited. at org.jboss.jmx.adaptor.html.HtmlAdaptorServlet.invokeOpByName(HtmlAdaptorServlet.java:319) . . {code}	1	1	https://issues.jboss.org/browse/JBPAPP-9502

#	Requirement	Questions	Answered	Successfully Answered	URL
9753	JBPAPP-9497	Result of Petr's review: ---------- * What sections we should remove, instead referring to the Migration Guide; Chapter 2, 3 and 4 - suitable candidates for move to the Migration Guide, if the Migration Guide is going to be published at all. If there will be no Migration Guide in the end, I'd suggest moving this to the end of this guide, so that New installation (currently chapter 5) is right after Introduction - better positioning of the information because a new installation is, imo, what users will most typically do when referring to this guide. ---------- * How we might restructure the remaining content. 5.1.2. Configuring Your Java Environment - points to Appendix C, but Appendix D is also relevant (maybe merge the two appendices as suggested further below) 6.1. Web Services Stack, 6.2. PicketLink Federation - these two don't go together with 6.3 Installation Methods; I'd divide this chapter into two, resulting in 6.3 being in its own chapter, and maybe move 6.1 and 6.2 to a further location in the guide not to interrupt the flow of information Chapter 14 - divide into two sections based on the type of installation (graphical (uninstall script) vs. others (rm -r [EAP root]), and indicate the installation types in the section heading Appendix A - Disabling Authentication - is there a reason for having this in the Installation Guide? I'd suggest moving this to the Security Guide or/and the Administration and Configuration Guide instead Appendix C - Installing JDK on RHEL - I'd merge this with Appendix D - Installing Sun JDK on Windows - info about JDK installation on RHEL6 is completely missing - we need to add it ---------- * other notes 1 - "Installing JBoss Enterprise Web Platform is simple and easy. You can have it installed and running in no time. This guide will teach you to install and uninstall JBoss." - referring to Enterprise Web Platform - should be Enterprise Application Platform - + a bit too marketing-ish, not suitable for technical docs - I'd suggest a re-write or even leaving this paragraph out 1.1 Other manuals: the link here is pointing to the global docs overview, shouldn't we rather be pointing to the EAP docs overview: http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/index.html? if not, I'd at least modify the text a bit to specify what products' docs (all products in this case) can be found at the link Chapter 9. Installation using the Graphical Installer - running the installer as root - causes problems on RHEL-6.3 (errors, unable to run, no shortcuts in apps menu), getting expected results when launching as a regular user - need to verify with SMEs about the need to run as root 12.4.2. Run the Enterprise Application Platform as a Service on Red Hat Enterprise Linux - I'd expect the same information here as in 12.4.1 - need to verify with an SME on how to start the service manually on linux, and if it is possible unless installed from an RPM - second note: missing end comma + I'd turn "Administration and Configuration Guide" into a link leading to the guide Appendix E - Installing Apache Ant Procedure E.2: "Do this on Red Hat Enterprise Linux by adding the following line" - this procedure is about installing "on Other Operating Systems", not on RHEL - shouldn't we change it to "on Unix/Linux systems"	1	1	https://issues.jboss.org/browse/JBPAPP-9497
9754	JBPAPP-9497	Petr - My comments are prefixed $$$ ---------- What sections we should remove, instead referring to the Migration Guide; Chapter 2, 3 and 4 - suitable candidates for move to the Migration Guide, if the Migration Guide is going to be published at all. If there will be no Migration Guide in the end, I'd suggest moving this to the end of this guide, so that New installation (currently chapter 5) is right after Introduction - better positioning of the information because a new installation is, IMHO, what users will most typically do when referring to this guide. It seems we won't have the MG ready for 5.1.3 but I agree that the first chapter after the introdcution should be Installation. ---------- How we might restructure the remaining content. 5.1.2. Configuring Your Java Environment points to Appendix C, but Appendix D is also relevant (maybe merge the two appendices as suggested further below) $$$ Agree - please merge these two appendices. 6.1. Web Services Stack, 6.2. PicketLink Federation - these two don't go together with 6.3 Installation Methods; I'd divide this chapter into two, resulting in 6.3 being in its own chapter, and maybe move 6.1 and 6.2 to a further location in the guide not to interrupt the flow of information $$$ Agree - please divide into two and move 6.1 and 6.2. Chapter 14 - divide into two sections based on the type of installation (graphical (uninstall script) vs. others (rm -r [EAP root]), and indicate the installation types in the section heading $$$ Agree. Appendix A - Disabling Authentication is there a reason for having this in the Installation Guide? I'd suggest moving this to the Security Guide or/and the Administration and Configuration Guide instead $$$ Agree. Please move this to the A&C Guide. IMHO the Security Guide should not give steps which actually reduce the platform's security. Appendix C - Installing JDK on RHEL I'd merge this with Appendix D - Installing Sun JDK on Windows info about JDK installation on RHEL6 is completely missing - we need to add it $$$ Agree. ---------- other notes 1 - "Installing JBoss Enterprise Web Platform is simple and easy. You can have it installed and running in no time. This guide will teach you to install and uninstall JBoss." referring to Enterprise Web Platform - should be Enterprise Application Platform $$$ Wow - please fix the product name. + a bit too marketing-ish, not suitable for technical docs - I'd suggest a re-write or even leaving this paragraph out $$$ Agree - please remove this paragraph. 1.1 Other manuals: the link here is pointing to the global docs overview, shouldn't we rather be pointing to the EAP docs overview: http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/index.html? if not, I'd at least modify the text a bit to specify what products' docs (all products in this case) can be found at the link $$$ Agree - change so that it points to the EAP docs. Chapter 9. Installation using the Graphical Installer running the installer as root - causes problems on RHEL-6.3 (errors, unable to run, no shortcuts in apps menu), getting expected results when launching as a regular user - need to verify with SMEs about the need to run as root $$$ Agree - I have seen a JIRA recently about the issue of running the installer as root. If I can find it I'll provide you the link. 12.4.2. Run the Enterprise Application Platform as a Service on Red Hat Enterprise Linux I'd expect the same information here as in 12.4.1 - need to verify with an SME on how to start the service manually on linux, and if it is possible unless installed from an RPM $$$ Agree. second note: missing end comma + I'd turn "Administration and Configuration Guide" into a link leading to the guide $$$ Please fix comma. The Red Hat style guide states that we DON'T link directly from one book to another. Appendix E - Installing Apache Ant Procedure E.2: "Do this on Red Hat Enterprise Linux by adding the following line" this procedure is about installing "on Other Operating Systems", not on RHEL - shouldn't we change it to "on Unix/Linux systems" $$$ Agree - change to "on Unix/Linux..." I also noticed the word "Repository" and this should be deleted.	1	1	https://issues.jboss.org/browse/JBPAPP-9497

#	Requirement	Questions	Answered	Successfully Answered	URL
9755	JBPAPP-9497	should remove, instead referring to the Migration Guide; Chapter 2, 3 and 4 - suitable candidates for move to the Migration Guide, if the Migration Guide is going to be published at all. If there will be no Migration Guide in the end, I'd suggest moving this to the end of this guide, so that New installation (currently chapter 5) is right after Introduction - better positioning of the information because a new installation is, IMHO, what users will most typically do when referring to this guide. $$$ It seems we won't have the MG ready for 5.1.3 but I agree that the first chapter after the introdcution should be Installation. €€€ Moved chapters 2, 3 and 4 to a further location, now they are chapters 10, 11 and 12. ---------- How we might restructure the remaining content. 5.1.2. Configuring Your Java Environment points to Appendix C, but Appendix D is also relevant (maybe merge the two appendices as suggested further below) $$$ Agree - please merge these two appendices. €€€ Appendices merged, currently it is appendix B. 6.1. Web Services Stack, 6.2. PicketLink Federation - these two don't go together with 6.3 Installation Methods; I'd divide this chapter into two, resulting in 6.3 being in its own chapter, and maybe move 6.1 and 6.2 to a further location in the guide not to interrupt the flow of information $$$ Agree - please divide into two and move 6.1 and 6.2. €€€ Moved the Installation methods under Chapter 2. New Installation. Moved WSS and PicketLink under Chapter 9. Optional Components. Chapter 14 - divide into two sections based on the type of installation (graphical (uninstall script) vs. others (rm -r [EAP root]), and indicate the installation types in the section heading $$$ Agree. €€€ Divided into 2 sections (currently 13.1 and 13.2). Added 13.3 Uninstallation from RPM installation. Appendix A - Disabling Authentication is there a reason for having this in the Installation Guide? I'd suggest moving this to the Security Guide or/and the Administration and Configuration Guide instead $$$ Agree. Please move this to the A&C Guide. IMHO the Security Guide should not give steps which actually reduce the platform's security. €€€ Completely removed from the Installation Guide. Added to Administration and Configuration Guide. Appendix C - Installing JDK on RHEL I'd merge this with Appendix D - Installing Sun JDK on Windows info about JDK installation on RHEL6 is completely missing - we need to add it $$$ Agree. €€€ Merged, currently Appendix B. ---------- other notes 1 - "Installing JBoss Enterprise Web Platform is simple and easy. You can have it installed and running in no time. This guide will teach you to install and uninstall JBoss." referring to Enterprise Web Platform - should be Enterprise Application Platform $$$ Wow - please fix the product name. €€€ Paragraph completely removed, see below. + a bit too marketing-ish, not suitable for technical docs - I'd suggest a re-write or even leaving this paragraph out $$$ Agree - please remove this paragraph. €€€ Paragraph removed. 1.1 Other manuals: the link here is pointing to the global docs overview, shouldn't we rather be pointing to the EAP docs overview: http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/index.html? if not, I'd at least modify the text a bit to specify what products' docs (all products in this case) can be found at the link $$$ Agree - change so that it points to the EAP docs. €€€ Changed, now it is pointing to EAP docs. Chapter 9. Installation using the Graphical Installer running the installer as root - causes problems on RHEL-6.3 (errors, unable to run, no shortcuts in apps menu), getting expected results when launching as a regular user - need to verify with SMEs about the need to run as root $$$ Agree - I have seen a JIRA recently about the issue of running the installer as root. If I can find it I'll provide you the link. €€€ Verified, no need to run the installer as root, information removed. 12.4.2. Run the Enterprise Application Platform as a Service on Red Hat Enterprise Linux I'd expect the same information here as in 12.4.1 - need to verify with an SME on how to start the service manually on linux, and if it is possible unless installed from an RPM $$$ Agree. €€€ Verified, only RPM installation causes EAP to run as a service. second note: missing end comma + I'd turn	0	0	https://issues.jboss.org/browse/JBPAPP-9497
9756	JBPAPP-9495	Hui, I did the RHEL-6 bits for you. Can you please try replicating it for RHEL-4/5? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-9495
9757	JBPAPP-9493	Hi [~lakagwu] and [~jfclere]. I'm wondering if I can get some more information about exactly what was causing this issue and how it was resolved. Was a CPING/CPONG implemented in the end?	1	0	https://issues.jboss.org/browse/JBPAPP-9493
9758	JBPAPP-9493	Thanks for the info [~jfclere]. Can you confirm that the draft Release Note above accurately describes the bug and its fix? If not, can you let me know where I misunderstood. Cheers.	1	0	https://issues.jboss.org/browse/JBPAPP-9493
9759	JBPAPP-9493	I am not sure about the status of the this JIRA. Is there a one-off patch ? If yes, do we have a link to share with customers ?	0	0	https://issues.jboss.org/browse/JBPAPP-9493
9760	JBPAPP-9493	A customer urgently needs a one-off patch. Do I created a another JIRA for this ? Environment: * EAP 5.1.1 * mod_cluster 1.0.10.GA_CP02 * mod_cluster Native-Libs + a patch mod_cluster-native-1.0.10-5.1.1.GA_CP02_patch02.ep5.el6.x86_64.rpm (https://issues.jboss.org/browse/JBPAPP-8957) * Apache Webserver 2.2.15 * RHEL 6.2	0	0	https://issues.jboss.org/browse/JBPAPP-9493
9761	JBPAPP-9493	Can we close this one as the fix will be in EAP 5.2.0 ER5 as indicated in JBPAPP-10409 and this one doesn't have the right workflow?	0	0	https://issues.jboss.org/browse/JBPAPP-9493
9762	JBPAPP-9466	(10:15:02 PM) ozizka-ntb: the .awssecret can't be in ~ (10:15:07 PM) ozizka-ntb: think jenkins (10:15:20 PM) ozizka-ntb: can I specify different location? (10:16:17 PM) akostadinov: will need to hack aws: (10:16:18 PM) akostadinov: $secrets_file \|\|= "$home/.s3cfg" if !-e "$home/.awssecret" && -e "$home/.s3cfg"; (10:16:18 PM) akostadinov: $secrets_file \|\|= "$home/.awssecret";	1	1	https://issues.jboss.org/browse/JBPAPP-9466
9763	JBPAPP-9454	[~smarlow], should we reject this one? Since we had several issues with Javassist 3.17.1 in Hibernate 4 in general, as well as AS7, I'd be worried about attempting to backport it into EAP 5.	1	1	https://issues.jboss.org/browse/JBPAPP-9454
9764	JBPAPP-9454	IMO, It would be good to defer JBPAPP-9454 until the known issues are resolved and we get more community testing of the Javassist 3.17.x.x changes. HHH-7884/JASSIST-189 is still open.	1	1	https://issues.jboss.org/browse/JBPAPP-9454
9765	JBPAPP-9454	When the known issues are resolved and enough testing has occurred, I think the next question is can the fix be back ported. The Java 7 verify error fixes would not be included in the back port, which might make the change less risky. Lets discuss further when I'm back.	1	1	https://issues.jboss.org/browse/JBPAPP-9454
9766	JBPAPP-9454	We have all ack's for including in EAP 6.2.0 (https://bugzilla.redhat.com/show_bug.cgi?id=900690). However, who should we ping about this in EAP 5? In my opinion, with the regressions we hit before, it's a bit of a risky move.	1	1	https://issues.jboss.org/browse/JBPAPP-9454
9767	JBPAPP-9393	Hi [~jbertram] Could you please supply a bit more information about what was causing HornetQ to fail in this scenario and what was done to correct the problem? Both this and the original HQ JIRA seem to be light on details. I'll do up a draft release note as soon as I have a bit more to go on. Cheers.	1	1	https://issues.jboss.org/browse/JBPAPP-9393
9768	JBPAPP-9348	@Paul, can you create a tag so we can bring ha-server-api upgrade into EAP 5.1.3?	0	0	https://issues.jboss.org/browse/JBPAPP-9348
9769	JBPAPP-9337	claprun@redhat.com made a comment on [bug 831445|https://bugzilla.redhat.com/show_bug.cgi?id=831445] Kyle, Why is the standard XML XSD resolved from the wsdl director? Also, there's a jax-ws-catalog.xml file in $JBOSS_HOME/server/$PROFILE/conf which contains mappings from Public/System Ids to local files, couldn't we piggy back on this file instead of adding a new file? I tried adding references to the WSRP files there but it didn't work so maybe this why EAP doesn't support catalog files properly? I would also argue that the directory that contains the WSRP-specific schemas should be named differently (probably wsrp). Finally, if the current solution is the only one that would work with minimal effort at this time (i.e. targeting EAP 5.1.2), could you also do the same changes so that the WSRP 2 schemas are also resolved locally, please?	1	1	https://issues.jboss.org/browse/JBPAPP-9337

#	Requirement	Questions	Answered	Successfully Answered	URL
9770	JBPAPP-9337	Kyle Lape <klape@redhat.com> made a comment on [bug 831445|https://bugzilla.redhat.com/show_bug.cgi?id=831445] > Why is the standard XML XSD resolved from the wsdl director? This can be changed easily. Changing the dir name to 'wsrp' should be fine. > What about conf/jax-ws-catalog.xml? I'm not sure why this file is there for you, but it's only used by CXF. JBossWS Native doesn't have any code to parse jax-ws-catalog.xml and add the information to the deployment metadata. Adding this code would be a major change to code, and this is why I am opting for the simpler (code-change-wise) solution. I see that file is in EPP 5.2.1, but I'm not sure why. JBossWS Native doesn't use it (confirmed by Jim Ma). > Can you add WSRP 2 schemas? Sure, I can do this. Would you mind to give me the information needed to add these and test them?	1	1	https://issues.jboss.org/browse/JBPAPP-9337
9771	JBPAPP-9337	claprun@redhat.com made a comment on [bug 831445|https://bugzilla.redhat.com/show_bug.cgi?id=831445] Kyle, Did you get a chance to add the WSRP 2 files?	1	1	https://issues.jboss.org/browse/JBPAPP-9337
9772	JBPAPP-9337	Hi Kyle, I've added the ability to resolve the jbossws-entities.properties under WEB-INF/classes , and committed the fix to jbossws-native-3.1.2 branch. Can you please help try if that works in WSRP2 ? Thanks.	1	0	https://issues.jboss.org/browse/JBPAPP-9337
9773	JBPAPP-9336	Instrumentor failed to transform org.jboss.test.aop.bean.POJO javassist.CannotCompileException: by javassist.CannotCompileException: [source error] cannot find [Ljava/lang/Object; at org.jboss.aop.instrument.GeneratedAdvisorFieldAccessTransformer.doBuildFieldWrappers(GeneratedAdvisorFieldAccessTransformer.java:84) at org.jboss.aop.instrument.FieldAccessTransformer.buildFieldWrappers(FieldAccessTransformer.java:103) at org.jboss.aop.instrument.Instrumentor.transform(Instrumentor.java:778) at org.jboss.aop.instrument.GeneratedAdvisorInstrumentor.transform(GeneratedAdvisorInstrumentor.java:117) at org.jboss.aop.SuperClassesFirstWeavingStrategy.instrumentClass(SuperClassesFirstWeavingStrategy.java:202) at org.jboss.aop.SuperClassesFirstWeavingStrategy.translate(SuperClassesFirstWeavingStrategy.java:69) at org.jboss.aop.AspectManager.translate(AspectManager.java:1079) at org.jboss.aop.AspectManager.transform(AspectManager.java:1023) at org.jboss.aop.asintegration.jboss5.DefaultTranslator.transform(DefaultTranslator.java:47) at org.jboss.classloader.spi.translator.TranslatorUtils.applyTranslatorsOnTransform(TranslatorUtils.java:71) at org.jboss.classloader.spi.ClassLoaderSystem.transform(ClassLoaderSystem.java:564) at org.jboss.classloader.spi.base.BaseClassLoaderDomain.transform(BaseClassLoaderDomain.java:266) at org.jboss.classloader.spi.base.BaseClassLoaderPolicy.transform(BaseClassLoaderPolicy.java:159) at org.jboss.classloader.spi.base.BaseClassLoader$2.run(BaseClassLoader.java:615) at org.jboss.classloader.spi.base.BaseClassLoader$2.run(BaseClassLoader.java:596) at java.security.AccessController.doPrivileged(Native Method) at org.jboss.classloader.spi.base.BaseClassLoader.loadClassLocally(BaseClassLoader.java:595) at org.jboss.classloader.spi.base.BaseClassLoader.loadClassLocally(BaseClassLoader.java:572) at org.jboss.classloader.spi.base.BaseDelegateLoader.loadClass(BaseDelegateLoader.java:135) at org.jboss.classloader.spi.filter.FilteredDelegateLoader.loadClass(FilteredDelegateLoader.java:131) at org.jboss.classloader.spi.base.ClassLoadingTask$ThreadTask.run(ClassLoadingTask.java:455) at org.jboss.classloader.spi.base.ClassLoaderManager.nextTask(ClassLoaderManager.java:267) at org.jboss.classloader.spi.base.ClassLoaderManager.process(ClassLoaderManager.java:166) at org.jboss.classloader.spi.base.BaseClassLoaderDomain.loadClass(BaseClassLoaderDomain.java:304) at org.jboss.classloader.spi.base.BaseClassLoaderDomain.loadClass(BaseClassLoaderDomain.java:1261) at org.jboss.classloader.spi.base.BaseClassLoader.loadClassFromDomain(BaseClassLoader.java:914) at org.jboss.classloader.spi.base.BaseClassLoader.doLoadClass(BaseClassLoader.java:506) at org.jboss.classloader.spi.base.BaseClassLoader.loadClass(BaseClassLoader.java:451) at java.lang.ClassLoader.loadClass(ClassLoader.java:356) at java.lang.Class.getDeclaredConstructors0(Native Method) at java.lang.Class.privateGetDeclaredConstructors(Class.java:2404) at	1	1	https://issues.jboss.org/browse/JBPAPP-9336
9774	JBPAPP-9336	Carlo, could you confirm? I got this error in my machine. [aopc] Exception in thread "main" java.lang.ClassNotFoundException: org.jboss.test.cluster.web.aop.Address [aopc] at java.net.URLClassLoader$1.run(URLClassLoader.java:202) [aopc] at java.security.AccessController.doPrivileged(Native Method) [aopc] at java.net.URLClassLoader.findClass(URLClassLoader.java:190) [aopc] at java.lang.ClassLoader.loadClass(ClassLoader.java:306) [aopc] at java.lang.ClassLoader.loadClass(ClassLoader.java:247) [aopc] at org.jboss.aop.standalone.Compiler.loadFile(Compiler.java:349) [aopc] at org.jboss.aop.standalone.Compiler.compile(Compiler.java:208) [aopc] at org.jboss.aop.standalone.Compiler.main(Compiler.java:88)	0	0	https://issues.jboss.org/browse/JBPAPP-9336
9775	JBPAPP-9316	A draft Release Note text has been supplied above, however confirmation is requested from the assignee as to whether a Release Note is required for this issue. Does the issue and it's resolution have an end-user impact? Might they need to know this to use EAP 5.1.3?	1	1	https://issues.jboss.org/browse/JBPAPP-9316
9776	JBPAPP-9299	According to JIRA the JBoss Web SME is Remy Maucherat - https://issues.jboss.org/secure/ViewProfile.jspa?name=rmaucher. How much work is likely to be involved in adding documentation of the Filter element? If it's more than a day or so I'd like to schedule it for a later date.	0	0	https://issues.jboss.org/browse/JBPAPP-9299
9777	JBPAPP-9295	After performing some research, I found that a link to the supported configurations page for EAP (the same as in the KB article) is already included in: - EAP Release Notes (as of 5.1.0): http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/Release_Notes_5.1.2/index.html#id2300403 - EAP Installation Guide: http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/Installation_Guide/index.html#Getting_Started-Pre_Requisites - here it is mentioned as "Supported Operating Systems", I suggest to change it to "Supported Configurations" and modify the text accordingly From the other guides, only Administration and Configuration Guide seems as a potential candidate for inclusion of this info. I have, however, not been able to identify a suitable chapter to include it in. The guide also seems focused on EAP itself, not on the environment, so I'm really not sure about including it here. Russel: 1) Can I go ahead and modify the linked section of Installation Guide? 2) What is your opinion on inclusion in Administration and Configuration Guide? Do we want this type of information in it?	1	1	https://issues.jboss.org/browse/JBPAPP-9295

#	Requirement	Questions	Answered	Successfully Answered	URL
9778	JBPAPP-9294	I found out that we only have the "2.1 Do You Need Help?" section in the Preface of each book. It directs users to the Red Hat Customer Portal, where they can open a support case. However, not a word about what information should be provided. If anywhere, I would expect the information from the KB article here. I have two ideas: 1) Extend the "2.1 Do You Need Help" section with the required information. This would, however, require changing it in the Publican brands used for JBoss docs, which I'm not sure we really want to do. What do you think Russel? 2) I guess it might be much more effective to have the information included directly on the page where support cases are submitted: https://access.redhat.com/support/cases/new I guess this is out of our scope, but I reckon this might really help.	1	1	https://issues.jboss.org/browse/JBPAPP-9294
9779	JBPAPP-9258	This JIRA has been marked as affecting Release Notes however nothing in it contains any matter that would do so. The original ticket that was cloned from (HORNETQ-940) highlights an issue, however it contains no reference to how the problem was solved. More information is required before this issue can be included in the EAP 5.1.3 Release Notes: What was causing the issue? How was the issue resolved?	1	1	https://issues.jboss.org/browse/JBPAPP-9258
9780	JBPAPP-9163	I fixed JBPAPP-9507 + renamed oracle11RAC to oracle11gRAC to reflect product name and message about supported DBS which mentions Oracle 11gRAC. (Commit #113210) I'd like to bring up JDBC 3 and JDBC 4 topic. Currently this tool forces JDBC 4 drivers. Would you agree to change name.properties to be less specific and allow jdbc3 drivers?	1	1	https://issues.jboss.org/browse/JBPAPP-9163
9781	JBPAPP-9112	Wouldn't an enforcer rule be more appropriate here? Or something similar to that? It should be obvious when a dependency or plugin diverges from the parent POM or BOM(s): it will have a version specified. Right?	1	0	https://issues.jboss.org/browse/JBPAPP-9112
9782	JBPAPP-9002	Marking this issue as "Needs more info" as, from a user perspective, does this change actually need a release note? Does the user need to know that these error messages will no longer be logged every five seconds?	1	1	https://issues.jboss.org/browse/JBPAPP-9002
9783	JBPAPP-8986	Dmaven.test.failure.ignore=true {code} testMessagingClient(org.jboss.as.test.smoke.messaging.client.messaging.MessagingClientTestCase): Timed out waiting to receive cluster topology. Group:null testMessagingClient(org.jboss.as.test.smoke.messaging.client.jms.JmsClientTestCase): Failed to create session factory That's expected, JMS is not supported in EC2. Failed tests: testSendWithNoTx(org.jboss.as.test.smoke.stilts.SimpleStompletTestCase): Bundle started testPropertyBasedEnvEntry(org.jboss.as.test.smoke.property.EjbDDWithPropertyTestCase): expected:<[foo/bar]> but was:<[${from.jboss.props1}${file.separator}${from.jboss.props2}]> testPropertyBasedEnvEntryWithOverride(org.jboss.as.test.smoke.property.EjbDDWithPropertyTestCase): expected:<[foo-\|-bar]> but was:<[${from.jboss.props1}${path.separator}${from.jboss.props2}]> Props replacing issue, looking into it. testClientDeploymentAsArchive(org.jboss.as.test.smoke.osgi.SimpleRunAsClientTestCase): expected:<[INSTALLED]> but was:<[ACTIVE]> This seems like timing issue. testListAvailableModules(org.jboss.as.test.smoke.jsr88.JSR88DeploymentTestCase) testDeploymentStreamApi(org.jboss.as.test.smoke.mgmt.servermodule.ServerInModuleDeploymentTestCase) testDeploymentFileApi(org.jboss.as.test.smoke.mgmt.servermodule.ServerInModuleDeploymentTestCase): expected:<1> but was:<0> Tests in error: testSendWithTxCommit(org.jboss.as.test.smoke.stilts.SimpleStompletTestCase): Connection timed out. testFindPackageCapability(org.jboss.as.test.smoke.osgi.SimpleModuleRegistrationTestCase): org/osgi/resource/Requirement testResolveModule(org.jboss.as.test.smoke.osgi.SimpleModuleRegistrationTestCase): org/osgi/resource/Requirement testRepositoryService(org.jboss.as.test.smoke.osgi.SimpleRepositoryTestCase): Could not serialize object: TestResult[status=FAILED,time=95ms] Timing issue probably. testStandaloneXml(org.jboss.as.test.smoke.subsystem.parse.ParseAndMarshalModelsTestCase): org/jboss/as/controller/services/path/PathManagerService testStandaloneHAXml(org.jboss.as.test.smoke.subsystem.parse.ParseAndMarshalModelsTestCase): org/jboss/as/controller/services/path/PathManagerService testStandaloneFullXml(org.jboss.as.test.smoke.subsystem.parse.ParseAndMarshalModelsTestCase): org/jboss/as/controller/services/path/PathManagerService testStandaloneFullHAXml(org.jboss.as.test.smoke.subsystem.parse.ParseAndMarshalModelsTestCase): org/jboss/as/controller/services/path/PathManagerService testHostXml(org.jboss.as.test.smoke.subsystem.parse.ParseAndMarshalModelsTestCase): org/jboss/as/controller/services/path/PathManagerService testDomainXml(org.jboss.as.test.smoke.subsystem.parse.ParseAndMarshalModelsTestCase): org/jboss/as/controller/services/path/PathManagerService testRegistryConfiguration(org.jboss.as.test.smoke.deployment.rar.tests.inflow.InflowTestCase): org/jboss/as/connector/util/ConnectorServices testMetadataConfiguration(org.jboss.as.test.smoke.deployment.rar.tests.inflow.InflowTestCase): org/jboss/as/connector/util/ConnectorServices testRegistryConfiguration(org.jboss.as.test.smoke.deployment.rar.tests.pure.PureTestCase): org/jboss/as/connector/util/ConnectorServices testMetadataConfiguration(org.jboss.as.test.smoke.deployment.rar.tests.pure.PureTestCase): org/jboss/as/connector/util/ConnectorServices I think this is fixed in GA. Note: Tested with ER4 which is in AMI.	1	1	https://issues.jboss.org/browse/JBPAPP-8986
9784	JBPAPP-8986	Dominik Pospisil <dpospisi@redhat.com> made a comment on [bug 900989|https://bugzilla.redhat.com/show_bug.cgi?id=900989] That would work. But as I said this need to be copied before other tests are executed so I guess this needs to be hacked into testsuite POM file which is something I would not prefer to do. I guess you are much more familiar with the testsuite POMs. Can you do that Ondro?	1	1	https://issues.jboss.org/browse/JBPAPP-8986
9785	JBPAPP-8986	Rostislav Svoboda <rsvoboda@redhat.com> made a comment on [bug 900989|https://bugzilla.redhat.com/show_bug.cgi?id=900989] Is the workaround sufficient ? If yes and you don't want changes in TS, please close this BZ. Btw, how many tests are affected by this BZ, I can't access https://jenkins-jbossqa.rhcloud.com.	1	1	https://issues.jboss.org/browse/JBPAPP-8986
9786	JBPAPP-8891	Do we need to build with Java 7? The requirement, as I understood it, was to only run on Java 7.	1	1	https://issues.jboss.org/browse/JBPAPP-8891
9787	JBPAPP-8891	I hate subversion: {noformat} svn cp varia/src/main/org/jboss/varia/stats/DataSourceInterceptorJDK6.java varia/src/main/org/jboss/varia/stats/AbstractDataSourceInterceptorJDBC4.java svn cp varia/src/main/org/jboss/varia/stats/DataSourceInterceptorJDK6.java varia/src/main/org/jboss/varia/stats/DataSourceInterceptorJDK7.java cd connector/src/main/org/jboss/resource/adapter/jdbc mkdir jdbc4 svn add jdbc4 svn cp jdk6/* jdbc4 for i in jdbc4/*; do svn mv $i `echo $i | sed "s/JDK6/JDBC4/"`; done cd - prep-svn-patch /tmp/JBPAPP-8891.patch patch -p0 </tmp/JBPAPP-8891.patch svn add varia/src/main/org/jboss/varia/stats/AbstractDataSourceInterceptorJDBC4_1.java cd connector/src/main/org/jboss/resource/adapter/jdbc svn add jdbc4_1/ jdk7/ cd - {noformat}	0	0	https://issues.jboss.org/browse/JBPAPP-8891
9788	JBPAPP-8786	Justin, is it possible to get more information about the fix that resolved this issue? Also, is that release notes draft text even approaching accurate?	1	1	https://issues.jboss.org/browse/JBPAPP-8786
9789	JBPAPP-8731	Justin, could you please confirm whether shutting down the main server or the backup server creates the problem? (refer to the <remark> tag in the Release Notes draft above for context) Also, did I identify the fix correctly from your and Clebert's comments on HORNETQ-911?	1	1	https://issues.jboss.org/browse/JBPAPP-8731
9790	JBPAPP-8683	from the code, i can see the fix provided from HHH-5043 is already in the EAP 5.1.2 release, would you provide a test case?	1	1	https://issues.jboss.org/browse/JBPAPP-8683

#	Requirement	Questions	Answered	Successfully Answered	URL
9791	JBPAPP-8683	Hello Liu, We used the same test case in the project JIRA. Do you really need another test case or the one attached in the other JIRA is enough?	1	0	https://issues.jboss.org/browse/JBPAPP-8683
9792	JBPAPP-8617	Do you see these OOM errors in upstream Jenkins builds? For this job [1] we're building against as console master branch (with -Peap), I'm just wondering if you are seeing the same errors. [1]https://hudson.qa.jboss.com/hudson/job/JBoss-EAP-6.0.x-Community/	1	1	https://issues.jboss.org/browse/JBPAPP-8617
9793	JBPAPP-8617	paul, did you figure it out?	1	1	https://issues.jboss.org/browse/JBPAPP-8617
9794	JBPAPP-8377	Can someone please explain just what it is that has been resolved in this JIRA issue? I presume JBTM = JBoss Transaction Manager? It's not at all clear in this issue nor linked issues just what the Cause, Consequence, Fix and Result are.	1	1	https://issues.jboss.org/browse/JBPAPP-8377
9795	JBPAPP-8363	Dimitris, can you create a tag to include 2.2.20 upgrade in EAP 5.1.3? Is there someone else this should be assigned to?	0	0	https://issues.jboss.org/browse/JBPAPP-8363
9796	JBPAPP-8168	Hi Justin, we're having difficulties to verify this jira. I can see that there were some changes to ClientCrashTest.java but it seems to be unstable: https://jenkins.mw.lab.eng.bos.redhat.com/hudson/view/EAP5/view/HornetQ/job/_hornetq-RHEL5-x86_64-SunJDK1.6-NIO/lastCompletedBuild/testReport/ Could you help with reproducing this issue, please? Thank you, Mirek	0	0	https://issues.jboss.org/browse/JBPAPP-8168
9797	JBPAPP-8125	Is this still an issue? can we close it? [~bershath27]?	0	0	https://issues.jboss.org/browse/JBPAPP-8125
9798	JBPAPP-8119	Hi Justin, it seems that work on this is done. Is it possible to commit it to upstream? Thanks, Mirek	1	1	https://issues.jboss.org/browse/JBPAPP-8119
9799	JBPAPP-8073	Since this issue is marked "Resolved", shouldn't the Fix Version/s field have been updated to "EAP 5.1.3"?	0	0	https://issues.jboss.org/browse/JBPAPP-8073
9800	JBPAPP-8073	Sorry, I didn't quite understand the initial request. That said, the part of the commit I backported should be fine as the synchronization will now log the stack trace. It looks to me like with both parts of the fix you would get the same stack trace twice? If you think it is safest, I can put the other part of the commit in if you want? I not too sure of the protocol in that case here, should I create a new tag and raise a Jira in JBPAPP to get the new version included?	1	1	https://issues.jboss.org/browse/JBPAPP-8073
9801	JBPAPP-8038	Which JBoss version uses EAP 5.1.3? As it is not fixed AS 7.1.1, I'd like know which version this bug fix for.	1	0	https://issues.jboss.org/browse/JBPAPP-8038
9802	JBPAPP-7881	Customer Answer: I think you don't understand my use case. We have a legacy app and JPA is a persistence option we are trying to bring (the app is using a proprietary ORM framework). We try to be as clean as possible (support the two through configuration with a clear migration path). Assume this app has 3 modules A, B and C. Each of them is using persistence (the proprietary one). B and C depends on A. I don't need to step back because we thought about that issue *a lot*. We're introducing module D which gives *RUNTIME* support for JPA through Hibernate. A, B and C can still perfectly work with the current system and support JPA if it's available. I don't want to bring hibernate as a dependency if we're not using it. Relying on JPA API and Hibernate annotations is fine though. In D, I am already registering the custom types we had to define to support custom data types. It's fine if D heavily relies on hibernate-core. Right now A still has hibernate core dependency because we have a @GenericGenerator on the base class and this refers to the generator class. Being able to register that through configuration in D is the feature I am looking for. Thanks for that. Regarding annotations split, it's the same. Why on earth would I have to bring the whole hibernate engine to configure a JPA entity with some advanced options? Again, a entity in B or C or A is using a custom type. I must use @Type. In order to do that, I have to bring the whole hibernate-core jar and its dependencies. Sorry, that does not make any sense to me.	1	0	https://issues.jboss.org/browse/JBPAPP-7881
9803	JBPAPP-7813	be the inclusion of the facelets-1.1.15.B1.patch, which doesn't work with properly with the JBoss Embedded version present in the seam distribution in EAP5.0.0, which is used to run the integration tests. See also JBPAPP-5456, which discusses a previous attempt to include facelets-1.1.15.B1.patch on EAP5.1. and JBPAPP-7057 JBPAPP-7058 which fixes the inclusion of the facelets-1.1.15.B1.patch in EAP5.1.2 by updating jboss embedded. The question is whether it is really necessary to include that part of the patch (which introduces the VFSClasspath to tothe jsf-facelets.jar) {noformat} java.lang.AssertionError: java.lang.IllegalStateException: File cannot contain children: FileHandler@1496691980[path= context=file:/home/maschmid/testing/one-offs/JBPAPP-7813/jboss-eap-5.0/seam/examples/excel/test-build/META-INF/MANIFEST.MF real=file:/home/maschmid/testing/one-offs/JBPAPP-7813/jboss-eap-5.0/seam/examples/excel/test-build/META-INF/MANIFEST.MF] at org.jboss.seam.mock.AbstractSeamTest$Request.onException(AbstractSeamTest.java:455) at org.jboss.seam.mock.AbstractSeamTest$Request$2.doFilter(AbstractSeamTest.java:504) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:83) at org.jboss.seam.web.IdentityFilter.doFilter(IdentityFilter.java:40) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.MultipartFilter.doFilter(MultipartFilter.java:90) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.ExceptionFilter.doFilter(ExceptionFilter.java:64) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.RedirectFilter.doFilter(RedirectFilter.java:45) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.ajax4jsf.webapp.BaseXMLFilter.doXmlFilter(BaseXMLFilter.java:178) at org.ajax4jsf.webapp.BaseFilter.handleRequest(BaseFilter.java:290) at org.ajax4jsf.webapp.BaseFilter.processUploadsAndHandleRequest(BaseFilter.java:368) at org.ajax4jsf.webapp.BaseFilter.doFilter(BaseFilter.java:495) at org.jboss.seam.web.Ajax4jsfFilter.doFilter(Ajax4jsfFilter.java:56) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.LoggingFilter.doFilter(LoggingFilter.java:60) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.HotDeployFilter.doFilter(HotDeployFilter.java:53) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.servlet.SeamFilter.doFilter(SeamFilter.java:158) at org.jboss.seam.mock.AbstractSeamTest$Request.run(AbstractSeamTest.java:491) at org.jboss.seam.mock.AbstractSeamTest$FacesRequest.run(AbstractSeamTest.java:862) at org.jboss.seam.example.excel.test.RenderTest.testSimple(RenderTest.java:27) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.testng.internal.MethodHelper.invokeMethod(MethodHelper.java:580) at org.testng.internal.Invoker.invokeMethod(Invoker.java:478) at org.testng.internal.Invoker.invokeTestMethod(Invoker.java:617) at org.testng.internal.Invoker.invokeTestMethods(Invoker.java:885) at org.testng.internal.TestMethodWorker.invokeTestMethods(TestMethodWorker.java:126) at org.testng.internal.TestMethodWorker.run(TestMethodWorker.java:110) at org.testng.TestRunner.runWorkers(TestRunner.java:712) at org.testng.TestRunner.privateRun(TestRunner.java:582) at	1	1	https://issues.jboss.org/browse/JBPAPP-7813
9804	JBPAPP-7729	[~stliu] Can you please complete the Release Notes text with the nominated information about this bug. Unfortunately I wasn't able to grok the nature of the problem from the comments on this ticket (or the linked HornetQ ticket). Cheers	1	1	https://issues.jboss.org/browse/JBPAPP-7729

#	Requirement	Questions	Answered	Successfully Answered	URL
9805	JBPAPP-7671	121) 17. Complete Working Example Proxy Server example has an opening < that probably shouldn't be there 122) Appendix: workers.properties Reference What exactly does a range of _ to / comprise? "The structure of a Worker Property is worker.worker_name.directive" <-- missing final full stop 123) Appendix: workers.properties Reference – workers.properties Connection Directives "For example, a correct directive flag set is worker.node1.ping_mode=CI" <-- missing full stop, and should probably describe what it is attempting to set for clarity's sake 124) Appendix: workers.properties Reference – workers.properties Load Balancing Directives "...three times the load than other workers..." <-- three times the load OF other workers >>> Fixed. 125) Appendix: Java Properties Reference – B.1. Proxy Configuration itemizedlist is incorrectly punctuated Proxy Configuration Values <-- not values, parameters or attributes The default values are displayed next to the parameter name here, but nowhere else in the document. >>> Fixed. 126) Appendix: Java Properties Reference – B.1. Proxy Configuration "an usable worker" <-- a usable worker >>> Fixed. 127) Appendix: Java Properties Reference – B.1. Proxy Configuration cping and cpong or CPing and CPong? Consistency required. >>> Fixed.	1	1	https://issues.jboss.org/browse/JBPAPP-7671
9806	JBPAPP-7671	Jan, I have made the following edits as per your feedback. Revision 20 is the latest available edition of this book. I have rebuilt it and it should soon be available at [1]. 109) 15.2.1 Configuring HttpSession Passivation number of sessions managed by the the session manager <-- remove a the >>> "the" was not removed Checked and confirmed as removed. Note that this text is now contained in section 17.2.1 (previously 15.2.1). 117) 16.1 Configuration SSO's <-- NO. SSOs. Not fixed Checked and confirmed as corrected. Note that this text is now contained in section 17.1 (previously 16.1). 125) Appendix: Java Properties Reference – B.1. Proxy Configuration Proxy Configuration Values <-- not values, parameters or attributes >>> not fixed Can you please clarify exactly what's to be changed here? 127) Appendix: Java Properties Reference – B.1. Proxy Configuration cping and cpong or CPing and CPong? Consistency required. >>> not fixed >>> Could you please do fulltext search on entire document and make this consistent? Thanks. I have done a full text search on the book's content and changed all references to CPIng and CPon, with the exception of references in log files. [1] http://documentation-devel.engineering.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/HTTP_Connectors_Load_Balancing_Guide/index.html	1	1	https://issues.jboss.org/browse/JBPAPP-7671
9807	JBPAPP-7670	90) 13. What Is Netscape Server API Question mark? lower case "is" "to gain a basic understand about" --> "to gain a basic understanding of" web server or HTTP Server? pick a term >>> Fixed. Question mark added, "Is" converted to lower case, "understand" changed to "understanding", HTTP changed to "web" 91) 13. What Is Netscape Server API Itemizedlist is not punctuated correctly – closing full stop with no commas or semi-colons >>> Fixed. Semi-colons added. 92) 14. Configuring the NSAPI Connector on Solaris worker node and master node use different cases here than in other places in the document >>> Invalid. I searched all the files which make up this book and the only instances I could find of capitalisation were correct - i.e. in titles.	1	1	https://issues.jboss.org/browse/JBPAPP-7670
9808	JBPAPP-7670	93) Task: Configure Initial Clustering Behaviour What is non-initial clustering behaviour? >_> Also, Step 1 – perhaps say "disable the mappings ..... by commenting them out, as shown...". Basic, but helps clarify the task for the unfamiliar. >>> Fixed 94) Task: Configure a Basic Cluster with NSAPI "The procedure specifies [that] worker01..." >>> Fixed. Also removed the unnecessary text "Use the following procedure".	1	1	https://issues.jboss.org/browse/JBPAPP-7670
9809	JBPAPP-7669	63) 12.1 Prerequisites and Configuration Assumptions Why is this a massive section, but most preprequisites in the first half of the book are as tiny little formalparas? Consistency please. >>> Fixed. I changed the layout and text so that it was more consistent with the rest of the text. 64) 12.1 Prerequisites and Configuration Assumptions The IMPORTANT note should probably be integrated into the steps rather than at the beginning of the chapter. (Which I think ought to be Step 4 of 12.2.) >>> Fixed. 65) 12.2 Configure Server Instance as a Worker Node – Prerequisites Step 1 – specifies use of the default profile. Is this the only profile supported? If so, remove "you want to configure as a worker node". If it's not the only supported profile, then the directive to copy default profile specifically should be removed. You can use default as an example after the main directive. Same again for subsequent steps – I'd make a main directive that summarizes what they are doing and why, followed by a specific example. >>> Fixed. The default profile is used as it already contains the necessary components. 66) Task: Define ISAPI Filter Not immediately clear what the filter will do or why you are defining it. Screenshots would be nice, if anyone has licensed products around. >>> Fixed. I added a sentence which explains briefly why the filter is created. 67) Task: Define ISAPI Filter Master Server or Master server? Be consistent. >>> Fixed. Changed "Master Server" to "master server".	1	1	https://issues.jboss.org/browse/JBPAPP-7669
9810	JBPAPP-7669	68) Task: Define ISAPI Filter – Step 5 Is the location of the executable always C:\connectors\jboss-ep-5.1\native\sbin\isapi_redirect.dll? >>> Fixed. I changed this so that it uses a path convention instead of an explicit path. 69) Task: Define ISAPI Virtual Directory Not immediately clear why this needs to be done. Perhaps a summary of the configuration required at the beginning of Ch. 12 would be useful to orient people. >>> Fixed. I added a brief description of what the filter will do. 70) Task: Define ISAPI Virtual Directory – Step 1 Missing closing full stop. >>> Fixed. 71) Task: Define ISAPI Virtual Directory – Step 2 Again, is this ALWAYS the path? >>> Fixed. I used a path convention here instead of the explicit path. 72) Task: Define ISAPI Virtual Directory – Step 4 Missing full-stop. >>> Fixed.	1	1	https://issues.jboss.org/browse/JBPAPP-7669
9811	JBPAPP-7669	73) Task: Define ISAPI Web Service Extension Bleh @ difference in cases between a button and a window – I'm assuming this is present in the software, but still. >>> I couldn't find where "window" and "button" are mentioned in this section so I've skipped this point. 74) Task: Define ISAPI Web Service Extension – Step 2 Missing "exactly as written" used in previous procedures. Also, is this always the path? >>> Fixed. Added "exactly as written" and changed path. 75) 12.4 Microsoft IIS 7 Initial Clustering Configuration – Terms URL at the end of the definition is not a hyperlink. Also, why does only this chapter have a Terms section? >>> Made text into a hyperlink. This chapter has only 1 entry in the Terms section and it's there instead of in an appendix because it's important the reader understands this term when reading the rest of the chapter. 76) Task: Define a JBoss Native ISAPI Restriction – Step 1 "Start -> Run, then type inetmgr." <- I think you also need to hit Enter or click Run afterward. >>> Fixed. 77) IMPORTANT: Task: Define a JBoss Native ISAPI Restriction – Step 2 Ends abruptly. Opens in the what? >>> Fixed. Deleted "in the what" 78) Task: Define a JBoss Native ISAPI Restriction – Step 5 Is this always the path? Also, the spacing around the "jboss" value is too wide (spaces inside and outside the tag?). >>> Fixed. Changed path. The spacing is a result of the styling applied. 79) Task: Define an JBoss Native Virtual Directory s/an/a >>> Fixed. 80) Task: Define an JBoss Native Virtual Directory – Step 2 Missing full stop after (exactly as written). Is this path always the same? >>> Fixed. Added full stop. Changed path. 81) Task: Enable the ISAPI-dll Handler – Step 1 Ends abruptly. >>> Fixed. 82) Task: Enable the ISAPI-dll Handler – Step 4 Missing full stop. >>> Fixed.	1	1	https://issues.jboss.org/browse/JBPAPP-7669

#	Requirement	Questions	Answered	Successfully Answered	URL
9812	JBPAPP-7669	83) Task: Configure ISAPI to serve a Basic Cluster – Step 1 Double spacing around the directory. Append to, not append into. >>> Fixed "append". Spacing is a matter of the JBoss styles and is not significant enough to change. 84) Task: Configure ISAPI to serve a Basic Cluster – Step 2 "... has an abstract directory name containing images" --> "has an abstract directory name, and contains images" >>> Fixed. 85) Task: Configure ISAPI to serve a Basic Cluster – Step 3 & 4 Append to, not into. >>> Fixed. 86) Task: Configure ISAPI to serve a Basic Cluster – Step 6 Is there a default location for the log_file property? We should tell people this. Also, tell them what exactly they are meant to be "verifying" about the logs. >>> Fixed. The default location is specified in the example configuration file. Re: checking logs, readers are to check for events of Warning and/or Error. 87) Task: Configure ISAPI to serve a Load Balancing Cluster Load-Balancing >>> Fixed. 88) Task: Configure ISAPI to serve a Load Balancing Cluster – Step 1 Directory has extra spaces around it >>> Spacing is a matter of the JBoss styles and is not significant enough to change. 89) Task: Configure ISAPI to serve a Load Balancing Cluster – Step 3 & 4 Append to, not into. >>> Fixed.	1	1	https://issues.jboss.org/browse/JBPAPP-7669
9813	JBPAPP-7668	29) 5.1.1. mod_manager.so "uses this location for generated keys" <-- uses this location to store generated keys, perhaps,to avoid for-for confusion >>> Fixed 30) 5.1.1. mod_manager.so Are there any limits on Maxcontext, Maxnode, Maxhost, Maxsessionid parameters? >>> Don't know but I'll try to find out. A quick Internet search didn't reveal any limits. 31) 5.1.1. mod_manager.so images/mod_status.png <--could stand to go inside a <figure> tag with a title, since it is all white it just kinda makes the spacing in the document look weird. >>> Fixed 32) 5.1.2. mod_proxy_cluster.so Again, keep an eye on case in your module names. "to change load balancing behaviour" --> load-balancing >>> Fixed 33) 5.1.3. mod_advertise.so case in module names "If port is not specified, the default port specified is 23364." --> the default port used, not specified >>> Fixed 34) 5.2 Install Proxy Server Components "on a JBoss Enterprise Web Server." <-- on a server instance, or on JBoss etc. not on a JBoss etc. >>> Fixed 35) 5.2 Install Proxy Server Components "The Native components are Operating System and processor architecture specific." --> The native components are specific to particular operating systems and processor architectures. (Either that or you need to hyphenate in a complicated way.) Also, check caps on Native Components. >>> Fixed 36) Task: Install Proxy Server Components "Edit the JBoss Enterprise Web Server Apache configuration file" --> Edit the [Apache] configuration file in your JBEWS instance: JBOSS_EWS_DIST/httpd/conf/httpd.conf Instruction for commenting out the line could be clearer, too. (Question for the devs: is there anything the mod_proxy_balancer module is compatible with, and if not, why is it here?) >>> Fixed 37) 6. Configure Basic Proxy Server "to configure a JBoss Enterprise Web Server" --> again, a JBEWS instance, or just JBEWS. Also, "JBoss HTTP connector mod_cluster" is poorly marked up and capitalized in the opening here. >>> Fixed	0	0	https://issues.jboss.org/browse/JBPAPP-7668
9814	JBPAPP-7668	45) Task: Install and Configure a JBoss Enterprise Web Server Worker Node Appendix B, Java Properties Reference is only mentioned in the EWS worker note config section, but the opening of this Appendix suggests that it refers to JBEAP, not JBEWS. Which is it? And should this appear in the previous task's prerequisites as well? >>> Fixed. Added to previous tasks's pre-reqs, also appendix opening changed so that EWS and EAP are mentioned. 46) 8. Further Server Configuration feels like a bit of a "misc" section. Can this content go elsewhere? >>> Fixed. In fact this material was somehow duplicated so I have deleted what was chapter 8. 47) Task: Configure Application Platform Worker Node with Static Proxy List comma separated --> comma-separated >>> Fixed 48) 9.2 Clustered Node Operation "The JBoss HTTP Connector can operate in non-clustered or clustered mode." --> in either clustered or non-clustered mode >>> Fixed 49) IMPORTANT: 9.2 Clustered Node Operation Does not include details on configuring clustered/non-clustered operation. >>> Still pending action. 50) 10. Load Balancing Demonstration jboss-eap-5.1/mod_cluster/demo directory. Only in 5.1? I imagine this will need to be edited in future; may as well make it multi-version compliant now. >>> Fixed. 51) 10. Load Balancing Demonstration I think QE likes us to avoid letting applications possess things (load-demo.war's) >>> Fixed. "load-demo.war's" change to "the application's" 52) 10. Load Balancing Demonstration "Load-demo app" <-- decide how you will refer to the app and use it consistently; avoid swapping between this and load-demo.war Similarly, decide whether to refer to it as the Demonstration or the Demo >>> Fixed. "Demo" changed to demonstration. 53) 10.1 Set up the Demonstration Step 2 Note – text and admonition weight (note vs important) should match the admonition in the previous section, if this note is here at all. >>> Fixed. 54) 10.1 Set up the Demonstration Step 3 - just prefix the step name with "Tomcat 6 only" or something, remove the admonition >>> Fixed. Admonition removed and prefix added instead.	0	0	https://issues.jboss.org/browse/JBPAPP-7668
9815	JBPAPP-7668	55) 10.1 Set up the Demonstration Step 3 - "<web-app>element," <-- missing space, and element not marked up as I would expect Also, is it actually "appending" (that is, is </web-app> the last thing in the file before you add this)? >>> Yet to be addressed 56) Task: Configure Client Control Tab Fields "...to learn the values you must supply in the Client Control tab of the Load Balancing Demonstration." <-- 'complete this task to learn' is a vague task goal; what does the task actually accomplish? I would recommend making Step 2 "Supply values for <required fields> and, optionally, <optional fields>." Move your list of references outside the task - either beforehand with "the demo client requires a number of values, defined here" or after with "the fields in step two expect values as defined in the following list". >>> Fixed. 57) Task: Interact with the Demonstration Task Prerequesites <-- differs from the straight "Prerequisites" in all other sections, and is misspelled >>> Fixed 58) Task: Interact with the Demonstration I'm not sure "experiment" should be a task/procedure. Additionally, screenshots couldn't go amiss. >>> Fixed. I have reworded the description to avoid "experiment". For the sake of the document's structure I have left this as a task. Regarding screenshots, I disagree and believe they are not necessary. The only steps which might benefit from screenshots are instructions to click on a tab in a GUI interface, a task so simple that I don't believe screenshots would help.	0	0	https://issues.jboss.org/browse/JBPAPP-7668
9816	JBPAPP-7668	Request for help with item #49 was made via email to Samuel Mendenhall (at James Livingston's suggestion): "Samuel, James Livingston recommend I contact you for help regarding documentation of HTPP Connectors. I am working on a JIRA issue in which some identified that the HTTP Connectors Load Balancing Guide states (in section 8.2 - "The JBoss HTTP Connector can operate in either clustered or non-clustered mode." but that is all. There is no further explanation of whether or not any configuration is required to use the HTTP Connector in a clustered configuration. Can you please give me some hints?"	0	0	https://issues.jboss.org/browse/JBPAPP-7668

#	Requirement	Questions	Answered	Successfully Answered	URL
9817	JBPAPP-7667	7) 1. Overview "takes the load of the disabled node" <-- reads like a typo for 'takes the load off'; awkward >>> Reworded 8) 1. Overview "Two distinct functions must be performed" <-- by...? Try "This involves two distinct operations:" >>> Reworded 9) IMPORTANT: 1. Overview "Sessions state replication" --> "Session state replication" "is handled by JBoss at the application level" <-- if it's handled at the application level, it's not handled by JBoss. Perhaps this should say "on an application by application basis"? Or does JBoss delegate responsibility to the application? >>> Reworded 10) 1. Overview "A cost effective way of managing load-balancing is..." <-- hyphenation is wrong. cost-effective and load balancing, please >>> Fixed 11) 1. Overview "Apache httpd " <-- described previously as Apache httpd Server, pick one and stick with it >>> "Apache HTTP Server" chosen as this is how it's described on the project's home page. Fixed. 12) 2. Download and Install "Apache httpd is included..." <-- described previously as Apache httpd Server, pick one and stick with it >>> Fixed 13) 3. Configure load balancing using Apache and mod_jk "Apache" <-- described previously as both Apache httpd Server and Apache httpd, pick one and stick with it, or say "hereafter referred to as Apache" (Last time I will mention this unless it develops a new reference; it's running rife throughout the document.) >>> Fixed 14) STYLE: 3. Configure load balancing using Apache and mod_jk – Prerequisites The weighting of the Task header and the Prerequisites header is the same, which makes it look like the steps actually belong to the Prerequisites. >>> Fixed	1	1	https://issues.jboss.org/browse/JBPAPP-7667
9818	JBPAPP-7666	OK, my fault but I must reopen - none of mentioned points are on https://access.redhat.com/knowledge/docs/en-US/JBoss_Enterprise_Application_Platform/5/html/HTTP_Connectors_Load_Balancing_Guide/index.html Am I missing something?	1	1	https://issues.jboss.org/browse/JBPAPP-7666
9819	JBPAPP-7658	I have added the requested information to the EAP 5.1.2 Release Notes. The book has now being re-Brew-ed and the modified version is available on DocStage: http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/Release_Notes_5.1.2/index.html#id3501816	0	0	https://issues.jboss.org/browse/JBPAPP-7658
9820	JBPAPP-7610	Ales, this upgrade has been identified for the EAP 5.1.3 release. Is it possible to create a tag including the linked issue to be incorporated?	0	0	https://issues.jboss.org/browse/JBPAPP-7610
9821	JBPAPP-7610	When do you need this? I can do it tomorrow. (hopefully it will work w/o problems, as it's been a while since I released JBCL :-))	0	0	https://issues.jboss.org/browse/JBPAPP-7610
9822	JBPAPP-7603	Clebert, Is there any info about this available anywhere? Cheers	1	1	https://issues.jboss.org/browse/JBPAPP-7603
9823	JBPAPP-7603	Thanks Justin The content from the upstream docs has been added to the enterprise Guide and is available for review at: http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/HornetQ_User_Guide/index.html#id4021952 Could you please confirm that it is technically accurate and suitable for publication. Cheers	1	1	https://issues.jboss.org/browse/JBPAPP-7603
9824	JBPAPP-7603	Thanks for the feedback Justin. I had misunderstood your info, but have it straight now. Can you please confirm the changes to be sure they're clear and accurate. http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html/HornetQ_User_Guide/ch29s07.html Thanks again	1	1	https://issues.jboss.org/browse/JBPAPP-7603
9825	JBPAPP-7556	Can I please double-check what's being requested? I have read the latest comments on JIRA JBPAPP-7521 and I believe the problem was that the redelivery-delay parameter was *NOT* supposed to have been changed from 6000 to 0 but instead this was a mistake. If I understand correctly, you're asking that (1) the Release Notes entry for this JIRA issue be deleted, and (2) the change to the HornetQ documentation be undone - i.e. the value "6000" be changed back to "0" in the appendix. I don't understand just why in JBPAPP 7521, Clebert and Rostislav seem to confirm that having "redelivery-delay" set to a value greater than 0 is a problem, yet a later comment from Miroslav states that if this parameter is not greater than 0 that failover doesn't work. Can someone please clarify the situation and what changes are required to documentation (including Release Notes)?	1	1	https://issues.jboss.org/browse/JBPAPP-7556
9826	JBPAPP-7556	Can someone please confirm just what the status of this issue is? As far as I can tell the change was documented correctly but some comments indicate that the documentation changes need to be undone. If I haven't received a reply by 3 February 2012 I will resolve this JIRA issue.	0	0	https://issues.jboss.org/browse/JBPAPP-7556
9827	JBPAPP-7494	Martha, this file is needed in the main zip, I've updated the JBPAPP_5_1 branch (Committed revision 113599.) with the fix, and local build now includes that file under docs/examples/jms/, can you please merge that into your 5.2.0 tag? Also, we should have this file in the rpm and zip now. Please update those as well if needed. Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-7494
9828	JBPAPP-7480	Misty, Can you move JBPAPP-7480 to bz and then cross reference the bz number in this Jira and close this one? Thanks, Mike	1	0	https://issues.jboss.org/browse/JBPAPP-7480
9829	JBPAPP-7455	Hi Clebert, I am reopening this issue because I have not found any documentation updates related to this new feature. There is no description in new documentation for EAP 5.1.2 [1] and there is no description in project documentation [2]. I have tried to add min-large-message-size to connection factory and this works fine. Core-bridge takes this configuration value. But I want to verify new configuration properties (well same property but at different locations) according discussion at help desk ticket. Could you add documentation for this new feature, please? [1] http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/HornetQ_User_Guide/index.html [2] http://anonsvn.jboss.org/repos/hornetq/branches/Branch_2_2_EAP	1	1	https://issues.jboss.org/browse/JBPAPP-7455
9830	JBPAPP-7455	Russe, Can you add something like the diff.patch attached here? I have committed this to the open source's doc.	1	1	https://issues.jboss.org/browse/JBPAPP-7455
9831	JBPAPP-7455	Attention: Clebert I am happy to add to the documentation mention of the "min-large-message-size" configuration option but before I do I need to check I understand correctly. (1) When I wrote the Release Notes entry for this issue I thought the configuration option already existed but that it wasn't recognised. If in fact this is a *new* configuration option, I need to correct the Release Notes text; (2) The patch you attached to this issue also contains documentation of two other configuration options. Is the documentation of those options covered by separate JIRA issues? I'll wait on your feedback before making any changes. Regards, Russell	1	1	https://issues.jboss.org/browse/JBPAPP-7455
9832	JBPAPP-7455	Attention: Clebert Sorry, but I'm now even more confused. :) "Russel: there should be two parameters to be added on the cluster-connection: min-large-message-size and max-retries" I understand from what you've written that these parameters already existed but previously *only* existed on ClientSessionFactory, not the cluster connnection. I understand that part. What I don't understand is the sudden mention of the "max-retries" parameter. It's not mentioned anywhere in this JIRA issue, nor in the patch you attached to this issue. The patch documents two parameters: min-large-message-size and reconnect-attempts. Can you please explain the difference between what you have said about the two parameters and the documentation patch? On an side note, I notice in the patch that you have a probably unwanted " character between the <literal> tags of the "reconnect-attempts" parameter. I'm sorry that we're having to play JIRA comment tag but I want to be sure I understand before I edit the documentation. Is there a better method by which we could communicate, perhaps IRC? I am rdickens and always in #docs, #brisbane, #jboss-docs when I am at work. Regards, Russell	1	1	https://issues.jboss.org/browse/JBPAPP-7455

#	Requirement	Questions	Answered	Successfully Answered	URL
9833	JBPAPP-7363	These issues require SME input to allow development of this document to continue: * Is it possible to get an example of how the SAML11SPRedirectFormAuthenticator method could be implemented? Config files, usecase? * What does Artifact Resolution Protocol offer end users, in a practical sense? * What happens if GLO participant numbers are too high? * There's tables in the source content for SAML2IssuerTrustHandler and SAML2LogOutHandler.java configuration parameters that have nothing in them. Is there content available or should they be removed? * Are the'authors' under BinaryTokenHandler (and the ensuing sections) the author of the code or the documentation about it? * How much of the content in the PicketLink STS Login Modules page is required in the enterprise docs?(http://community.jboss.org/wiki/PicketLinkSTSLoginModules). Part of ECS mandate is to avoid linking to community documents as much as possible. Is it important enough to port the entire page into the doc? * Is this an appropriate trace category for org.jboss.security? The documentation says to add one, but not what it should look like. <!-- <category name=" org.jboss.security"> <priority value="TRACE"/> </category> --> These questions are also in the PLINK doc itself (highlighted in yellow) http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/PicketLink_Reference_Guide/index.html	1	1	https://issues.jboss.org/browse/JBPAPP-7363
9834	JBPAPP-7325	Anil, can this change be added to the requested file without impacting the current customers?	1	1	https://issues.jboss.org/browse/JBPAPP-7325
9835	JBPAPP-7325	We can do this at the EAP level. Can we do this as part of the EAP build?	1	1	https://issues.jboss.org/browse/JBPAPP-7325
9836	JBPAPP-7275	Attention: Permaine We are past the docs freeze date for EAP 5.1.2 but I am willing to make this correction to the documentation if it's minor. The description and your comments on this issue don't seem to match so I'm confused. In the Description you say "But there isn't such package jbossas-web-tp-licences..." yet in the comments "We now have jbossas-{,web}tp-licences...". Even more confusing is your latest comment which states that this page should be removed from the yum command you quoted. If the package exists, why is it to be removed from the yum command? I look forward to your reply.	1	1	https://issues.jboss.org/browse/JBPAPP-7275
9837	JBPAPP-7275	Unfortunately, I haven't seen correction on Stage Docs server yet - licence package is still mentioned on every RHEL (EAP/EWP). Isn't the problem that IG for EWP will be not updated (PM decision)? Although I've seen this package also in EAP docs. Update: Note from Permaine is correct, licence package is deployed to system in the first phase automatically as dependency for jbossas.	1	1	https://issues.jboss.org/browse/JBPAPP-7275
9838	JBPAPP-7275	Attention: Pavel (and Rajesh) (1) The latest change noted on 22 November 2011 was to the *EAP* Installation Guide, not the *EWP* Installation Guide as this issue requests, so I'll need to correct my error ASAP. I wonder, though, if this change is also relevant to EAP? In other words, if this change affects EAP *and* EWP, should I leave my changes as they are? (2) Pavel, you're quite correct that in a PM meeting some time ago it was decided that EWP documentation was no longer to be maintained. So my original changes to the EWP Installation Guide should not have been done. I'll leave them as they are, though, since it's not worth undoing them.	1	1	https://issues.jboss.org/browse/JBPAPP-7275
9839	JBPAPP-7221	I'm trying to resolve old open issues filed against the Installer component, and it would seem that the [EAP 6 Installation Guide|https://access.redhat.com/knowledge/docs/en-US/JBoss_Enterprise_Application_Platform/6/html/Installation_Guide/Install_JBoss_Enterprise_Application_Platform_6_Using_the_Installer_Method.html] now contains documentation about this feature. Does this issue need to remain open?	1	1	https://issues.jboss.org/browse/JBPAPP-7221
9840	JBPAPP-7221	I echo Alex' question from 2012; do we need this to remain open now that documentation exists?	1	1	https://issues.jboss.org/browse/JBPAPP-7221
9841	JBPAPP-7203	The customer has indicated that their production date is October 28, 2011. Can we make this date?	0	0	https://issues.jboss.org/browse/JBPAPP-7203
9842	JBPAPP-7046	Can this issue be gone? EAP branch deploys and runs.	1	1	https://issues.jboss.org/browse/JBPAPP-7046
9843	JBPAPP-7014	Brew build: jboss-cluster-ha-client-1.1.3-1.1.ep5.el6 Repo location: http://download.lab.bos.redhat.com/brewroot/repos/jb-eap-5-rhel-6-build/latest/maven/org/jboss/cluster/jboss-ha-client/1.1.3.Final/	0	0	https://issues.jboss.org/browse/JBPAPP-7014
9844	JBPAPP-7014	Brew Build: jboss-cluster-ha-client-1.1.4-1.1.ep5.el6 Repo Location: http://download.lab.bos.redhat.com/brewroot/repos/jb-eap-5-rhel-6-build/latest/maven/org/jboss/cluster/jboss-ha-client/1.1.4.Final/	1	1	https://issues.jboss.org/browse/JBPAPP-7014
9845	JBPAPP-6996	Attention: Derek Horton Can you please expand on the draft Release Notes entry? Standard release notes are based on four pieces of information: Cause, Consequence, Fix and Result. I have looked over the issue's description and commits etc but simply don't understand the issue well enough to complete the Release Notes field with confidence. Thankyou, Russell Dickenson Content Author	1	0	https://issues.jboss.org/browse/JBPAPP-6996
9846	JBPAPP-6990	The contents of the attached patch have been incorporated into setcion 15.2. Datasource Parameters of the Administration and Configuration Guide.	1	1	https://issues.jboss.org/browse/JBPAPP-6990
9847	JBPAPP-6958	Sent email to Clebert and Tomas (with Pavel in cc) Hi Clebert and Tomas Pavel raised https://issues.jboss.org/browse/JBPAPP-6958 to capture the need for GFS2/SAN Fencing information for HornetQ. Can you assist with the correct information so we can incorporate this info before docs freeze on September 18? I would need the basic procedure with configuration options. The end result of the procedure would be the customer correctly configuring Fencing on their system. Because customer server implementations vary, if we can give the customer configuration info for the most common, or top two most common scenarios (like we did for Colocated Symmetrical Clustering), that would be ideal. Jared Morgan	1	1	https://issues.jboss.org/browse/JBPAPP-6958
9848	JBPAPP-6923	This jira description is too brief for me to verify it. Please add some explanatory usecase or description so as we know what we should verify. Also, is there a relationship to JBAPP-6954? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-6923
9849	JBPAPP-6893	Does this need to be mentioned in the Release Notes for EAP 5.1.2, considering it has a helpdesk ticket reference. The policy is that every GSS issue has CCFR provided in the Release Note Text field.	0	0	https://issues.jboss.org/browse/JBPAPP-6893
9850	JBPAPP-6892	We need the change in both the zip and rpm, is that correct? Just want to make sure before I make the changes. Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-6892
9851	JBPAPP-6846	Hi [~rhatlapa] The instance of 'JBoss Server' in the Release Notes had already been corrected in the source code when I went to fix it. A new version will be staged today. Nice catch. Regarding the Security Guide; I'm having trouble finding the instance of 'JBoss Server' you're referring to. A search of a html-single build of the book returns only the reference in my own Revision History entry, which says that the term had been removed. It's possible the latest version of the doc had not propagated to the stage by the time you checked it (there can be some lengthy delays sometimes). If the instance is still present, could you grab a screenshot and attach it to this JIRA. If not, and you're satisfied that this fix is verified, please close this ticket. Thanks for your diligence on this issue.	1	1	https://issues.jboss.org/browse/JBPAPP-6846
9852	JBPAPP-6794	I've spoke to colleagues about testing this. Without changing the actual code to put in a sleep to stop the synchronization and try and make it break it's not possible. The customer is happy to test the patch for us but neither the upstream or apache fix has a test case for this either. Can we just do some base testing. Any questions let me know please.	1	1	https://issues.jboss.org/browse/JBPAPP-6794

#	Requirement	Questions	Answered	Successfully Answered	URL
9853	JBPAPP-6752	jbossas-tp-licenses-5.1.2-1.ep5.el4 jbossas-tp-licenses-5.1.2-1.ep5.el5 jbossas-tp-licenses-5.1.2-1.ep5.el6 jbossas-web-tp-licenses-5.1.2-1.ep5.el4 jbossas-web-tp-licenses-5.1.2-1.ep5.el5 jbossas-web-tp-licenses-5.1.2-1.ep5.el6 Built in Brew. Martha, please let me know when you have a jbossas{,-web} built, I can update the build to R: -tp-licenses and remove the packaing of the licenses file in the jbossas{,-web} itself.	1	1	https://issues.jboss.org/browse/JBPAPP-6752
9854	JBPAPP-6743	I can confirm the EAP 5 documentation is fixed, but EWS 2 doc still need this fix: http://documentation-devel.engineering.redhat.com/docs/en-US/JBoss_Enterprise_Web_Server/2/html-single/Hibernate_Entity_Manager/index.html#setup-configuration-bootstrapping Could you please fix it? Thank you	0	0	https://issues.jboss.org/browse/JBPAPP-6743
9855	JBPAPP-6726	Stefan, this is supposed to be fixed in the JBPAPP_5_1 branch since December but I could not find the SVN commit with that JIRA in the subject in the svn log. Can you confirm that is fixed and find the svn rev number for us please?	1	1	https://issues.jboss.org/browse/JBPAPP-6726
9856	JBPAPP-6724	To my mind, the problem only occurs when requireReauthentication is set to True (I checked it), because of this : After looking at the code of WebAuthentication, I saw that it sets AUTH_TYPE field with private static final java.lang.String AUTH_TYPE = "PROGRAMMATIC_WEB_LOGIN"; In org.apache.catalina.authenticator.SingleSignOnEntry.updateCredentials(Principal, String, String, String), the method checks if authtype is FORM or BASIC to sets the canReauthenticate boolean to true, this one will allow to reauthenticate later. So as far as I understand, org.jboss.web.tomcat.security.login.WebAuthentication.AUTH_TYPE is incompatible with org.apache.catalina.authenticator.SingleSignOn.requireReauthentication. What do you think ?	0	0	https://issues.jboss.org/browse/JBPAPP-6724
9857	JBPAPP-6706	Seems like something that a customer could do with knowing about. Flagged for release note at a minimum. Could there also be a docs impact with this fix?	1	1	https://issues.jboss.org/browse/JBPAPP-6706
9858	JBPAPP-6680	Hi, I'm not sure, how does a typical usecase looks like. Where is this resources.xml file located? Could you provide me with some example? Thanks	1	1	https://issues.jboss.org/browse/JBPAPP-6680
9859	JBPAPP-6680	I'm sorry, I'm not familiar with LzPack and I'm apparently missing something. This splash screen should be displayed when the installer is run? And where in the installer jar directory structure is eap-config directory located? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-6680
9860	JBPAPP-6654	Hi Howard, I've identified some tests which fails regularly due to too strict timeout. Could you please increase them 10x for receive(...) methods and 2x for thread sleep. Thank you, this will help me a lot with futher testing. Mirek org.jboss.test.messaging.jms.ConnectionFactoryTest(Remote-bisocket).testAdministrativelyConfiguredOrderingGroup - line 829 FailoverTest(Clustering).testSimpleFailoverWithRemotingListenerEnabled - line 1429 FailoverTest.simpleFailover(FailoverTest.java:2321) JMSExpirationHeaderTest.testNoExpirationOnTimeoutReceive(JMSExpirationHeaderTest.java:90) DistributedQueueTest.clusteredQueue(DistributedQueueTest.java:611) HATest.testFailoverWithUnackedMessagesClientAcknowledge(HATest.java:416) ChangeFailoverNodeTest.killAllToOneAndBackAgain - increase timeout in all Thread.sleep(8000) method to 20000 seconds	0	0	https://issues.jboss.org/browse/JBPAPP-6654
9861	JBPAPP-6649	do you have a test replicating this?	1	1	https://issues.jboss.org/browse/JBPAPP-6649
9862	JBPAPP-6649	Hi, we have tried to re-test this issue with EAP 5.1.2 ER2 and we are able to reproduce this issue with this build. {code} java.lang.IndexOutOfBoundsException at org.jboss.netty.buffer.AbstractChannelBuffer.checkReadableBytes(AbstractChannelBuffer.java:657) at org.jboss.netty.buffer.AbstractChannelBuffer.readBytes(AbstractChannelBuffer.java:337) at org.jboss.netty.buffer.AbstractChannelBuffer.readBytes(AbstractChannelBuffer.java:343) at org.hornetq.core.buffers.impl.ChannelBufferWrapper.readSimpleStringInternal(ChannelBufferWrapper.java:84) at org.hornetq.core.buffers.impl.ChannelBufferWrapper.readNullableSimpleString(ChannelBufferWrapper.java:58) at org.hornetq.jms.client.HornetQTextMessage.doBeforeReceive(HornetQTextMessage.java:143) at org.hornetq.jms.client.HornetQMessageConsumer.getMessage(HornetQMessageConsumer.java:236) at org.hornetq.jms.client.HornetQMessageConsumer.receive(HornetQMessageConsumer.java:133) at org.jboss.qa.hornetq.TestConsumerSimple.main(TestConsumerSimple.java:54) {code} How to reproduce it? - setup EAP 5.1.2 ER2 with HornetQ, use PAGE address full policy and durable test queue - use standalone JMS client and push 1000 large text messages into queue. Every message include 1MB of real text - use standalone JMS client and try to receive all messages (no transaction, auto ack). Disconnect network when client consumes messages (just pull out network cable from client machine) and restore this connection. Do it several times during the test and do not put cable back after last disconnection. After some time you will see IndexOutOfBoundsException on the client side (you will have to wait several minutes). We have also noticed another interesting behavior. Client is configured to receive messages with 30s timeout but client will not end with null message. Try to pull out cable and wait more than 30s. You should see it. {code} while ((message = consumer.receive(30000)) != null) {code}	1	1	https://issues.jboss.org/browse/JBPAPP-6649
9863	JBPAPP-6649	Hi Clebert, I agree but is it possible to provide better exception here? IndexOutOfBoundsException does not look good. Could you add your comment about described issue with consumer.receive(30000)? It seems that client does not return null message in this case. Is it expected behavior? Thank you.	1	1	https://issues.jboss.org/browse/JBPAPP-6649
9864	JBPAPP-6618	This is due to the nature of AUTO_ACK mode. The user should be aware that using a small timeout this would happen. I think it is a normal behavior. Can we mark it as 'won't fix'? Howard	1	1	https://issues.jboss.org/browse/JBPAPP-6618
9865	JBPAPP-6618	Hi Mirek, If a user doesn't want to loose a message, he should use a transactional ack to guarantee this. The behavior of AUTO_ACK conforms to JMS spec. Can we postpone this jira to later releases until we have a better idea? Howard	1	1	https://issues.jboss.org/browse/JBPAPP-6618
9866	JBPAPP-6589	Hi Clebert, Roger, severity of log message was decreased to "debug" level. Is it ok? Jira says to decrease it to "either WARNING or INFO". If it's ok then jira can be closed. I'll wait for feedback. Thank you, Mirek	0	0	https://issues.jboss.org/browse/JBPAPP-6589
9867	JBPAPP-6577	how does this relate to HHH-2389?	1	0	https://issues.jboss.org/browse/JBPAPP-6577

#	Requirement	Questions	Answered	Successfully Answered	URL
9868	JBPAPP-6577	Strong Lui, Have you ever tried to use the private modifier for version/id setter? For sure not! If you have tried it, you would notice that, because the proxy extends the entity class, these setters need to be, at least, protected. About the use case not being a good practice. Why this is not a good practice? In 2006 (when this version of hibernate was released) we did not have conversation scope in JavaEE, remember? As far as I know, conversation scope was recently add to Java EE 6 (which was proposed by Seam Framework). Even though, conversation scope, in my opinion, is not a good option for all kinds of system. About the rejection. It is not a "Feature Request". It is a "BUG", for sure. Think about the below cases: 1) Entity is detached: hibernate receives the version and understand it. 2) Entity is attached: hibernate receives the version, DOES NOT understand it and simply SAY NOTHING to the programmer. So, it is a BUG because we have only two options: 1) hibernate implements the same semantics in both attached or detached (that is the better solution and used by other JPA frameworks, in my opinion) OR 2) hibernate throws an Exception when entity is attached saying it does not understand the new value. Ignoring a value just set is, for sure, A BUG. Of course, the second option can not be implemented at this point because it will break many systems around the world. So, if you do not want to fix this. Please, say that. But do not say it is not a bug because IT IS. I know a lot of people that have the same problem, even with newer versions of hibernate. If you don't want to fix this in this old version, please, fix this in the last version of hibernate. So we will not have the same issue forever.	1	0	https://issues.jboss.org/browse/JBPAPP-6577
9869	JBPAPP-6577	{quote} Have you ever tried to use the private modifier for version/id setter? For sure not! If you have tried it, you would notice that, because the proxy extends the entity class, these setters need to be, at least, protected. {quote} trust me, i know what i was saying. http://docs.jboss.org/hibernate/core/3.6/reference/en-US/html/tutorial.html#tutorial-firstapp-firstclass {quote} About the use case not being a good practice. Why this is not a good practice? In 2006 (when this version of hibernate was released) we did not have conversation scope in JavaEE, remember? As far as I know, conversation scope was recently add to Java EE 6 (which was proposed by Seam Framework). Even though, conversation scope, in my opinion, is not a good option for all kinds of system. {quote} did i say conversation in JAVA EE 6? i say "session-per-conversation", it is a well known hibernate pattern. http://docs.jboss.org/hibernate/core/3.6/reference/en-US/html/transactions.html#transactions-basics-apptx {quote} About the rejection. It is not a "Feature Request". It is a "BUG", for sure. Think about the below cases: 1) Entity is detached: hibernate receives the version and understand it. 2) Entity is attached: hibernate receives the version, DOES NOT understand it and simply SAY NOTHING to the programmer. So, it is a BUG because we have only two options: 1) hibernate implements the same semantics in both attached or detached (that is the better solution and used by other JPA frameworks, in my opinion) OR 2) hibernate throws an Exception when entity is attached saying it does not understand the new value. Ignoring a value just set is, for sure, A BUG. Of course, the second option can not be implemented at this point because it will break many systems around the world. So, if you do not want to fix this. Please, say that. But do not say it is not a bug because IT IS. I know a lot of people that have the same problem, even with newer versions of hibernate. If you don't want to fix this in this old version, please, fix this in the last version of hibernate. So we will not have the same issue forever. {quote} i'm saying it is not a bug because of we never said in our document hibernate supports that, and, take a look of JPA spec, "11.1.50 Version Annotation" -- "In general, fields or properties that are specified with the Version annotation should not be updated by the application.[112]"	0	0	https://issues.jboss.org/browse/JBPAPP-6577
9870	JBPAPP-6514	Email sent to Thomas seeking further information: "Thomas, I need your help with this JIRA issue as I need to write an appropriate Release Notes entry but don't have enough information. Since this is a few feature, we'll need to provide instructions to customers as to how it's to be used - both in creating the installation template file and using it to complete an automated installation. Does this documentation already exist and if so, where? If not, when might it be created? Of course all I need is a rough draft, from which I can write more complete instructions and add these to the installation guide. Since we're writing draft Release Notes now, please ACK this message by Monday, 5 September 2011 and provide a response by Wednesday, 7 September 2011. https://issues.jboss.org/browse/JBPAPP-6514 Thanks, -- Russell Dickenson Engineering Content Services Red Hat, Brisbane Australia"	1	1	https://issues.jboss.org/browse/JBPAPP-6514
9871	JBPAPP-6497	Tyronne, Are these requests still valid? To my knowledge, no changes have been made to the docs as a result of this JIRA. 1: What is the JCA adapter renaming you referred to? 3: There's many references to 'transactions' in the HornetQ Guide, which section (number and title) did you wish to see separated into a new chapter?	1	1	https://issues.jboss.org/browse/JBPAPP-6497
9872	JBPAPP-6497	Hi, >>1: What is the JCA adapter renaming you referred to? This has already been changed and looks fine. >>3: There's many references to 'transactions' in the HornetQ Guide, which section (number and title) did you wish to see separated into a new chapter? I doubt it's necessary at this stage. With the latest changes in JMS RAR, you don't truly need to perform any manual configurations on the JBossTS to enlist HQ resources for recover. This has been automated since JBoss-EAP-5.1.2 This Jira can now be closed.	1	1	https://issues.jboss.org/browse/JBPAPP-6497
9873	JBPAPP-6483	Hi Jared, item: 7.4 Standalone HornetQ JNDI configuration needed to be removed Invalid configuration and comment is still here: “In the default configuration,….” + including Note. Could you check it, please? http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/HornetQ_User_Guide/index.html#id3379235	0	0	https://issues.jboss.org/browse/JBPAPP-6483
9874	JBPAPP-6462	Jimmy, Can you please provide Cause, Consequence, Fix, Result information for this issue? I have looked over the patch attached above and the linked JIRA issues but still don't understand what the issue was and how it was resolved. I need this information for EAP 5.1.2's release notes so I'd appreciate an ACK by Thursday, 25 August 2011 and a response by Monday, 29 August 2011.	1	1	https://issues.jboss.org/browse/JBPAPP-6462
9875	JBPAPP-6462	Jimmy, Do you know if I will have confirmation on this issue's draft Release Notes field by the end of this week? If not I'll need to exclude it from the Release Notes.	1	1	https://issues.jboss.org/browse/JBPAPP-6462
9876	JBPAPP-6462	Attention: Jimmy Sorry for marking this "Not Required" but I was pressed for time to get Release Notes finalised. Other JIRA issues are have put me back a little and since you've corrected the text I thought it was best to include it. I have added a little information to the Release Notes about what was done to fix the issue. I read the SVN commit diff but am not sure if I have described the fix accurately. Could you please look over it and tell me if the description is wrong? Thanks	1	1	https://issues.jboss.org/browse/JBPAPP-6462
9877	JBPAPP-6444	Hi Rajesh, I believe that certification for Oracle 11gR2 11.2.0.2 [1] is done. We could close this jira. What do you think? Thank you, Mirek [1] https://docspace.corp.redhat.com/docs/DOC-30279#Oracle_112020_mm184	0	0	https://issues.jboss.org/browse/JBPAPP-6444
9878	JBPAPP-6418	I mean.. if we don't at least log an error, the user will never fix his application.	0	0	https://issues.jboss.org/browse/JBPAPP-6418

#	Requirement	Questions	Answered	Successfully Answered	URL
9879	JBPAPP-6418	Another reason for not to do it: Say, you mistyped an expression. Say: venorID=id, when you had a vendorID property on your messages. If I don't throw an exception, you will never know you didn't have an exception. If you want an improvement here.. the maximum we could do is log a warn with PROPERTY=<property> non existent. But that would be major change, with low priority to be developed since I don't see much gain compared to other issues we have to invest our time.	0	0	https://issues.jboss.org/browse/JBPAPP-6418
9880	JBPAPP-6368	Will the .bat appear eventually, or should we close?	0	0	https://issues.jboss.org/browse/JBPAPP-6368
9881	JBPAPP-6368	Hi Clebert, I've attached switch.bat file. I'm not sure how to get it to EAP distribution. Could you do it please? Thank you, Mirek	1	1	https://issues.jboss.org/browse/JBPAPP-6368
9882	JBPAPP-6368	Is there any other mechanism (like ant) to install HQ on the EAP distribution to replace JBM? If no this is a blocker.	0	0	https://issues.jboss.org/browse/JBPAPP-6368
9883	JBPAPP-6333	Rajesh, could you reassign this issue to some EAP QE? Or, just close it as won't fix if you don't mind EAP 5 logging this annoying INFO.	0	0	https://issues.jboss.org/browse/JBPAPP-6333
9884	JBPAPP-6303	I am reopening this JIRA because there is a new configuration parameter for ra (see committed changes). This parameter is not documented in project documentation and it is not documented in EAP documentation. Could you update project and EAP documentation, please?	1	1	https://issues.jboss.org/browse/JBPAPP-6303
9885	JBPAPP-6302	Builds updated to use hornetq-eap-exmples, resteasy-examples, jboss-messaging-examples and a svn export of the ws examples: rh-ewp-docs-5.1.1-3.ep5.el4 rh-ewp-docs-5.1.1-3.ep5.el5 rh-ewp-docs-5.1.1-3.ep5.el6 resteasy and ws examples updated in : rh-eap-docs-5.1.1-3.ep5.el4 rh-eap-docs-5.1.1-3.ep5.el5 rh-eap-docs-5.1.1-3.ep5.el6 Marking this as resolved, will be in ER2.	1	1	https://issues.jboss.org/browse/JBPAPP-6302
9886	JBPAPP-6260	Larry O'Leary <loleary@redhat.com> made a comment on [bug 710230|https://bugzilla.redhat.com/show_bug.cgi?id=710230] (In reply to comment #3) > Fixed in master - commit 5f9ddb7. I do not see 5f9ddb7, did you mean e431028f? Also, please note that this doesn't truly resolve this BZ. It resolves the example use-case provided but the BZ is more general then that. It seems that this confusion can exist with any operation that would have a boolean result. Instead of using the "Success" status on the operation status, maybe we should be looking at: In Progress --> [Failed | Complete] Or even going as far as displaying the operations result as the state. If a boolean, "True|False", if a simple object "<object>.toString()", if a complex object, "Object Name" or "See Result" link. The idea is that we want to convey to the user, via the UI, that what they are looking at is not the result of the operation but a link that will take you to the result of the operation.	1	1	https://issues.jboss.org/browse/JBPAPP-6260
9887	JBPAPP-6174	Fernando, could we get this component upgrade to the ER2?	0	0	https://issues.jboss.org/browse/JBPAPP-6174
9888	JBPAPP-6166	Prabhat, can you assign this patch to somebody on Richfaces QE for testing?	0	0	https://issues.jboss.org/browse/JBPAPP-6166
9889	JBPAPP-6157	Martha, we need to bump the release, to at least 7, of the RHEL 6 build to be newer than the base one, antlr-2.7.7-6.5.el6. Do you mind rebuilding it? Thanks!	1	0	https://issues.jboss.org/browse/JBPAPP-6157
9890	JBPAPP-6156	Martha, we need to bump the release of the RHEL 6 build, to at least 9, to be newer than the base one, hsqldb-1.8.0.10-8.el6. Do you mind rebuilding it? Thanks!	1	0	https://issues.jboss.org/browse/JBPAPP-6156
9891	JBPAPP-6145	Is this JIRA about a documentation section describing how nad why S3_PING (and possibly the JBM id script) use Amazon S3 to enable clustering?	1	0	https://issues.jboss.org/browse/JBPAPP-6145
9892	JBPAPP-6145	Is this information really relevant? We ship our cluster-ec2 and mod_cluster-ec2 profiles pre-configured to use S3_PING and tcp (although UDP unicast might still work we have tcp in config). I guess we need to document how S3_PING works and what configuration user needs to apply in the amazon AWS console but I don't see why should we document advanced jgroups configuration details.	1	0	https://issues.jboss.org/browse/JBPAPP-6145
9893	JBPAPP-6138	Hi Jared, could you add a note to documentation about this issue. Thank you, Mirek	0	0	https://issues.jboss.org/browse/JBPAPP-6138
9894	JBPAPP-6138	Attention: Jared Are you happy to run with this one? I thought about re-assigning it to myself but I don't think it's worth it.	1	1	https://issues.jboss.org/browse/JBPAPP-6138
9895	JBPAPP-6138	Hi Jared, problem is little bit different. I think that documentation update is not correct here. Problem: You can use HA activation property on your MDB. Problem is that MDB with this activation property cannot be used in default EAP configuration profile because this configuration profile does not have cluster=true in HornetQ configuration. Documentation should contain information that this activation property cannot be used with HornetQ without cluster enabled. If you use this MDB with the HornetQ without cluster=true, you will get the following exception: {code} 09:51:07,042 WARN [HornetQActivation] Failure in HornetQ activation org.hornetq.ra.inflow.HornetQActivationSpec(ra=org.hornetq.ra.HornetQResourceAdapter@f1abb1c destination=/queue/DLQ destinationType=javax.jms.Queue ack=Auto-acknowledge durable=false clientID=null user=null maxSession=15) HornetQException[errorCode=3 message=Timed out waiting to receive cluster topology] at org.hornetq.core.client.impl.ServerLocatorImpl.createSessionFactory(ServerLocatorImpl.java:723) at org.hornetq.ra.inflow.HornetQActivation.setupSession(HornetQActivation.java:352) at org.hornetq.ra.inflow.HornetQActivation.setup(HornetQActivation.java:290) at org.hornetq.ra.inflow.HornetQActivation$SetupActivation.run(HornetQActivation.java:587) at org.jboss.resource.work.WorkWrapper.execute(WorkWrapper.java:205) at org.jboss.util.threadpool.BasicTaskWrapper.run(BasicTaskWrapper.java:260) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:603) at java.lang.Thread.run(Thread.java:679) {code} Could you contact Clebert for more detailed information, please? Thank you.	1	1	https://issues.jboss.org/browse/JBPAPP-6138
9896	JBPAPP-6138	Hi Clebert Can you provide more information regarding Pavel's comment? Cheers Jared	1	1	https://issues.jboss.org/browse/JBPAPP-6138
9897	JBPAPP-6107	Is this upgrade being done for EAP 5.1.1 or 5.1.2? I would prefer this be updated in EAP 5.1.1 as we want to target JDK 1.6 U24 certification with EAP 5.1.1. We already have all our lab systems updated to this version.	0	0	https://issues.jboss.org/browse/JBPAPP-6107
9898	JBPAPP-6098	Mike, let's leave this around for the next release. However, Would it be ok to not update this JIRA until after EAP 5.1.1 and EWS 1.0.2 releases as it's causing some confusion. Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-6098
9899	JBPAPP-6098	Do we need the 1.0.10? Its just EC2 stuff - is it supported in EWS anyway? If yes could we make it to CR2?	1	1	https://issues.jboss.org/browse/JBPAPP-6098
9900	JBPAPP-6098	jboss-eap5-native built with this version of mod_cluster. Mladen, can you please update the EAP 5.1.2 natives to include this version of mod_cluster? Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-6098
9901	JBPAPP-6098	Updated. Do we need a -patch packages and what name that would be?	1	1	https://issues.jboss.org/browse/JBPAPP-6098
9902	JBPAPP-6098	Mladen, we don't need a -patch package, this is for EAP 5.1.2, can you please update the name of the zips to reflect it is for 5.1.2? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-6098
9903	JBPAPP-6098	Can somebody please fill in Affects and Fix versions? Now I don't know which release this is for.	0	0	https://issues.jboss.org/browse/JBPAPP-6098
9904	JBPAPP-5989	Juca, can this task be resolved?	1	1	https://issues.jboss.org/browse/JBPAPP-5989

#	Requirement	Questions	Answered	Successfully Answered	URL
9905	JBPAPP-5924	Hi Martha, can you build the patch-02 when Stefan sends us the new properties file please?	1	1	https://issues.jboss.org/browse/JBPAPP-5924
9906	JBPAPP-5893	Would it make sense to set the value inside the comment to {{jboss.messaging.ServerPeerID}} by default? i.e. instead of {code} <!-- property name="com.arjuna.ats.jta.xaRecoveryNode" value="CHANGE_ME" --> {code} put {code}<!-- property name="com.arjuna.ats.jta.xaRecoveryNode" value="${jboss.messaging.ServerPeerID}" --> {code} If the server is being used in a clustered environment, users would have to set {{jboss.messaging.ServerPeerID}} anyway.	1	1	https://issues.jboss.org/browse/JBPAPP-5893
9907	JBPAPP-5889	Hui, can you try and do this on Monday? The new tag is http://anonsvn.jboss.org/repos/jbossws/thirdparty/cxf/tags/cxf-2.2.12-patch-01/	0	0	https://issues.jboss.org/browse/JBPAPP-5889
9908	JBPAPP-5876	Hi Ivo I need more information before I can resolve this issue: For A. I will need a reference point with descriptions of these elements so I can add them to the table. For B. There is no mention of connection-factory.ha, connection-factory-ha, or any other permutation. Can you link to where you saw this element mentioned in the Enterprise HornetQ User Guide. Cheers Jared	1	1	https://issues.jboss.org/browse/JBPAPP-5876
9909	JBPAPP-5876	Hi Andy Can you help with the element descriptions? All I need is a link to a wiki or txt file with adequate info and I can take it from there. I've had no luck with my google searches for the missing element definitions. Cheers J	1	1	https://issues.jboss.org/browse/JBPAPP-5876
9910	JBPAPP-5875	Hi Ivo I've added descriptions for each of the elements for B. to the HQ User guide. I trawled the HQ user guide for the descriptions, and put them in with links to the most relevant section of the user guide. More info required to resolve this issue completely however: 1. For A, does the .xsd have descriptions in the form of comments for the missing configuration elements/parameters. I can list the elements as specified, but they will be meaningless without explanatory text. 2. Currently, the table seems to be largely unsorted, and quite chaotic. Can you see a reason why we might have the elements in the order they are? My only thought would be that it follows the nesting in each config file, but that seems a bit daft. I would propose a reorder (which would be after the Putnam release due to time to do this).	1	1	https://issues.jboss.org/browse/JBPAPP-5875
9911	JBPAPP-5875	Hi Andy Can you help with the element descriptions? All I need is a link to a wiki or txt file with adequate info and I can take it from there. I've had no luck with my google searches for the missing element definitions. Cheers J	0	0	https://issues.jboss.org/browse/JBPAPP-5875
9912	JBPAPP-5875	The following information has been provided by Clebert Suconic (many thanks). Some of these parameters need further context before Clebert can confirm their purpose etc. "I have embedded my answer, as well as adding some questions as I'm missing some context here: - initial-wait-timeout? In what context is that? cluster-connection there's a initial lookup for the cluster that has a separate timeout. log-delegate-factory-class-name... it's only for standalone usage. it should never be touched on the EAP context live-connector-ref... I"m not sure about it.... where is that definition? hornetq-configuration / root? log-journal-write-rate, perf-blast-pages, run-sync-speed-test is left undocumented on purpose. we can eventually enable to run some performance tests on the server and verify what's the write rate the system is capable off. I've been thinking of removing it anyway. connector-services? I'm not sure what's that? I'm missing a context here as well bridges.ha.. if it's using ha to connect to a server through the bridge. it should be under the birdge, right? bridge.user/password/static-connectors... that's used to define the connectivity on the birdge cluster-connection.static-connectors.. that's used to define static clusters message-counter-history-day-limit.. that's to define a limit on the MessageCounters.. which you can have a history on it. there's a chapter about it somewhere else."	0	0	https://issues.jboss.org/browse/JBPAPP-5875
9913	JBPAPP-5844	Dennis, can you please try and verify this fix? This pull the plug test requires a Windows host and on the very old XP host I have, it takes 23 minutes to start JBoss. The XP host executes so slowly that timeouts are causing members to be suspected, triggering merges and other problems which complicate verification. If you can't verify it, i'm going to have to pass it back to Rajesh so he can find someone with at least one moderately fast Windows platform from which a plug can be pulled.	1	1	https://issues.jboss.org/browse/JBPAPP-5844
9914	JBPAPP-5844	Good thinking, Ondrej. How about trying this out with the ER2 release? You'll need to start up two instances of EAP on two separate interfaces, let them cluster together, and then shut down / restart one of the interfaces as above. You could try to reproduce the error first using an older version of EAP.	0	0	https://issues.jboss.org/browse/JBPAPP-5844
9915	JBPAPP-5813	So should I take out the workaround then? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-5813
9916	JBPAPP-5813	Can anybody please fill in Affects Version/s and update Fix Version/s ?	0	0	https://issues.jboss.org/browse/JBPAPP-5813
9917	JBPAPP-5786	Why is tag called GA when version is SP1?	0	0	https://issues.jboss.org/browse/JBPAPP-5786
9918	JBPAPP-5775	Resolved with the fix for JBTM-821 in 4.6.1 CP09 via JBPAPP-5186.	1	1	https://issues.jboss.org/browse/JBPAPP-5775
9919	JBPAPP-5775	Jonathan, take a look at the Release Notes Text please. I think there should be TransactionManager or JBossTS instead of Hibernate. Could you help Misty with proper release notes wording, pls?	1	1	https://issues.jboss.org/browse/JBPAPP-5775
9920	JBPAPP-5757	Verified code in jboss-jpa-deployers.jar packed in EAP-5.1.2.ER1 distribution. MANIFESt.MF shows version 1.0.0 and PersistenceUnitInfoImpl still contains public ClassLoader getNewTempClassLoader() { return null; } Was the component upgraded in 5.1.2 as commented earlier?	0	0	https://issues.jboss.org/browse/JBPAPP-5757
9921	JBPAPP-5711	Anne-Louise, should we change the name of the ZIPs from what was shipped in EAP 5.1.0.GA or should we try and keep them consistent?	0	0	https://issues.jboss.org/browse/JBPAPP-5711
9922	JBPAPP-5702	We should probably move listConsumersAsJSON to JMS. Can you open a feature request on that case? (low priority) We would deprecate the method and create a new one for compatibility purposes.	1	1	https://issues.jboss.org/browse/JBPAPP-5702
9923	JBPAPP-5699	Is this on purpose, Fernando? If so we can close this out as Won't Fix.	0	0	https://issues.jboss.org/browse/JBPAPP-5699
9924	JBPAPP-5641	John, since resource recovery is now handled automatically, what exactly are our documentation requirements for this?	1	1	https://issues.jboss.org/browse/JBPAPP-5641
9925	JBPAPP-5641	Ivo, can you give your recommendation on this ticket? Reopen if necessary.	0	0	https://issues.jboss.org/browse/JBPAPP-5641
9926	JBPAPP-5616	Hi Jared, could you add a comment to HornetQ User Guide, please? Thank you, Mirek	0	0	https://issues.jboss.org/browse/JBPAPP-5616

#	Requirement	Questions	Answered	Successfully Answered	URL
9927	JBPAPP-5610	<jmorgan@redhat.com> CC: James 'Doc' Livingston <jlivings@redhat.com>, Laura Bailey <lbailey@redhat.com>, Jean-Frederic Clere <jclere@redhat.com> In the short term, if that's not possible, we could take the Enterprise configuration section that is there and publish it as an article; then refer readers to the existing community doc for further reference. On 12/15/2010 12:49 PM, Joshua Wulf wrote: > > Currently it's here: > > https://svn.devel.redhat.com/repos/ecs/JBoss_Enterprise_Application_Platform/trunk/5.x/HTTP_Connector_Guide > > > > > > I've built and tagged a version to the stage, but it hasn't shown up yet. > > > > In the svn repo is a file > > https://svn.devel.redhat.com/repos/ecs/JBoss_Enterprise_Application_Platform/trunk/5.x/HTTP_Connector_Guide/Content%20Map > > > > This file has the proposed TOC of the book, which is essentially a > > reorganisation of the existing mod_cluster 1.0.0 doc. > > > > It's about 40-50% complete. All the Enterprise-specific installation and > > configuration stuff for supportability has been covered. What remains > > now is to rewrite the existing guide information on specific configurations. > > > > The weaknesses of the existing community mod_cluster guide that I'm > > trying to address are: > > > > 1. Implementation-driven documentation rather than task- or feature-based > > > > Information is largely organised based on which implementation component > > provides the functionality, rather than around the end-user tasks. I've > > created a structure which groups end-user tasks, and presents the > > relevant information there. > > > > 2. Insufficient explanation of features. > > > > The example of that that I am currently stuck on is "CreateBalancers". > > The relationship between Balancers and VirtualHosts is not explained > > conceptually prior to this. From googling and reading forums it's still > > not clear to me how the two relate. My angle on this was to spin up some > > VMs (see the ticket referenced in the Content Map) and test it out. > > > > Once I got the basic configuration up and running and did some > > investigation I planned to ask Jean-Frederic some intelligent questions, > > and write it all up. > > > > I think to get this guide over the finish line we need to put someone > > exclusively on it for a couple of weeks. Jean-Frederic is more than > > happy to help; and I'm happy to help someone for an hour or two a day to > > bring them up to speed and help them along - we just need someone who > > can give it focused attention. > > > > - Josh > > > > On 12/14/2010 10:36 AM, Jared Morgan wrote: >> >> Hi Josh >> >> >> >> What needs to be done for this guide? >> >> >> >> Jared Morgan >> >> Content Author >> >> Red Hat Asia Pacific >> >> 1/193 North Quay >> >> BRISBANE QLD 4000 >> >> >> >> P: +61 7 3514 8242 >> >> M: +61 413 005 479 >> >> >> >> ---------- >> >> *From: *"Laura Bailey" <lbailey@redhat.com> >> >> *To: *"James 'Doc' Livingston" <jlivings@redhat.com> >> >> *Cc: *"Jared Morgan" <jmorgan@redhat.com>, "Joshua Wulf" <jwulf@redhat.com> >> >> *Sent: *Tuesday, December 14, 2010 7:51:07 AM >> >> *Subject: *Re: mod_cluster docs >> >> >> >> This seems to have fallen through the cracks because the post-release >> >> stuff was cancelled, but we do need to have a mod_cluster guide online. >> >> >> >> The HTTP Connectors Guide is missing ISAPI information, but is >> >> reasonably complete otherwise afaik. >> >> >> >> Josh, I think you were responsible for completing this guide - can you >> >> and Jared check on its progress, and figure out what we need to do to >> >> get it online (including tee-up with Andrew and Rajesh for QE)? >> >> >> >> ISAPI configuration details can be left until 5.1.1 release, but if the >> >> rest is more or less ready to go, we should get it out there. >> >> >> >> If need be, we can just publish the mod_cluster config details as an >> >> article and add it back to the HTTP Connectors Guide for 5.1.1. >> >> >> >> Cheers, >> >> Laura >> >> >> >> >> >> >> >> >> >> >> >> ----- Original Message ----- >> >> From: "James 'Doc' Livingston" <jlivings@redhat.com> >> >> To: "Laura Bailey" <lbailey@redhat.com>, "Jared Morgan" <jmorgan@redhat.com> >> >> Cc: "Sumit Bhat" <sbhat@redhat.com> >> >> Sent: Monday, 13 December,	1	1	https://issues.jboss.org/browse/JBPAPP-5610
9928	JBPAPP-5593	This should be fixed already. Can you check it?	1	1	https://issues.jboss.org/browse/JBPAPP-5593
9929	JBPAPP-5593	Hi Clebert, examples included in official distribution do not work. I have created new JIRA for it, see JBPAPP-6659. Could you check if examples were prepared correctly for distribution by productization, please? E.g. file readme [1] in root folder of examples should not be included in distribution. [1] http://source.jboss.org/browse/HornetQ/branches/Branch_2_2_EAP/examples-eap/readme?r=10177&r=10177&r=10177	0	0	https://issues.jboss.org/browse/JBPAPP-5593
9930	JBPAPP-5591	Here's the thread I started when I noticed it went missing... Jimmy -------- Original Message -------- Subject: IPv6 in JBoss Documentation Date: Mon, 11 Oct 2010 16:32:12 -0500 From: Jimmy Wilson <jawilson@redhat.com> To: Laura Bailey <lbailey@redhat.com> I have a customer asking about how to configure JBoss for use with IPv6. I have the information contained in the release notes [1], but I happen to know that bin/run.sh needs to be modified as well. This was previously documented in the Installation Guide. I know this because if I search for "ipv6 jboss" using the documentation search [2] it returns this: "7.3. IPv6 Configuration ... JBoss Enterprise Application Platform starts with an IPv4 stack by default. To use IPv6 instead, you will need to make some changes to the run scripts. ... www.redhat.com/.../4.3.0.cp08/html/Installation_Guide/IPv6.html - Cached" Clicking either the "www.redhat.com..." or "Cached" links doesn't work due to the docs changes, and looking at 4.3 CP08's Installation Guide, section 7.3 isn't about IPv6 anymore. Can you help me find the referenced page from the search? I've web searched the full text of *every* piece of 4.3 documentation that I thought relevant, and I've found nothing. I know what the documentation was about, so that's not the problem, but I'd like to reference it to the customer if it still exists. Should I open a JIRA to re-add it? [1] http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/4.3/html-single/Release_Notes_CP08/index.html#id2834368	1	1	https://issues.jboss.org/browse/JBPAPP-5591
9931	JBPAPP-5591	"""Networking JBPAPP-5591 IPv6 is not supported in version 5 of the Enterprise Platform. """ Shouldn't this be "Enterprise Application Platform"? Or does it refer generally to E*P? Reopening.	0	0	https://issues.jboss.org/browse/JBPAPP-5591
9932	JBPAPP-5571	Removed Resolution text from release notes text field. Anil, would the new attribute need to be documented for all customers in the Sec Guide. Would customers other than Mastercard require this attribute to use it for other applications?	1	0	https://issues.jboss.org/browse/JBPAPP-5571
9933	JBPAPP-5568	Anil, would the new attribute need to be documented for all customers in the Sec Guide. Would customers other than Mastercard require this attribute to use it for other applications?	1	1	https://issues.jboss.org/browse/JBPAPP-5568
9934	JBPAPP-5567	Removed Resolution text from release notes text field. Anil, would the new attribute need to be documented for all customers in the Sec Guide. Would customers other than Mastercard require this attribute to use it for other applications?	1	0	https://issues.jboss.org/browse/JBPAPP-5567
9935	JBPAPP-5536	Hi David I can't access the Issue Tracker issue that relates to this problem. I tried to search for the included number, but I can't find any trace of it. Is the Description the same as the issue tracker description. I will assign this to Anne-Louise Tangring so she can forward to the correct area for developmental action. Anne-Louise, JBPAPP-3798 is also currently unassigned, so should we put that in someone's queue as well. Cheers Jared	1	0	https://issues.jboss.org/browse/JBPAPP-5536
9936	JBPAPP-5516	found random regression in ejb3 testsuite run with EAP 5.1.2ER2.SendMessagesTestCase is failing at assertNotNull(msg); as message received is null. Does not fail always! https://hudson.qa.jboss.com/hudson/view/EAP5/job/eap5-ejb3/80/testReport/org.jboss.ejb3.test.jbas7922.unit/SendMessagesTestCase/testShutdown/	0	0	https://issues.jboss.org/browse/JBPAPP-5516

#	Requirement	Questions	Answered	Successfully Answered	URL
9937	JBPAPP-5516	For some reason Arjuna decides to abort the transaction with the first message "await". {noformat} 2011-11-16 13:41:51,348 WARN [com.arjuna.ats.arjuna.logging.arjLoggerI18N] (Thread-16) [com.arjuna.ats.arjuna.coordinator.TransactionReaper_18] - TransactionReaper::check timeout for TX a1058b2:ac0b:4ec3fc80:3bc4c in state RUN 2011-11-16 13:41:51,348 WARN [com.arjuna.ats.arjuna.logging.arjLoggerI18N] (Thread-17) [com.arjuna.ats.arjuna.coordinator.BasicAction_58] - Abort of action id a1058b2:ac0b:4ec3fc80:3bc4c invoked while multiple threads active within it. 2011-11-16 13:41:51,349 WARN [com.arjuna.ats.arjuna.logging.arjLoggerI18N] (Thread-17) [com.arjuna.ats.arjuna.coordinator.CheckedAction_2] - CheckedAction::check - atomic action a1058b2:ac0b:4ec3fc80:3bc4c aborting with 1 threads active! 2011-11-16 13:41:51,349 WARN [com.arjuna.ats.arjuna.logging.arjLoggerI18N] (Thread-17) [com.arjuna.ats.arjuna.coordinator.TransactionReaper_7] - TransactionReaper::doCancellations worker Thread[Thread-17,5,jboss] successfully canceled TX a1058b2:ac0b:4ec3fc80:3bc4c 2011-11-16 13:41:51,442 WARN [com.arjuna.ats.arjuna.logging.arjLoggerI18N] (WorkManager(2)-51) [com.arjuna.ats.arjuna.coordinator.BasicAction_40] - Abort called on already aborted atomic action a1058b2:ac0b:4ec3fc80:3bc4c 2011-11-16 13:41:51,457 ERROR [org.jboss.resource.adapter.jms.inflow.JmsServerSession] (WorkManager(2)-51) Unexpected error delivering message delegator->JBossMessage[5412736656638262]:PERSISTENT, deliveryId=0 java.lang.reflect.UndeclaredThrowableException {noformat} The await message is then redelivered and the MDB will go into the wait state again, which times out at the end.	1	1	https://issues.jboss.org/browse/JBPAPP-5516
9938	JBPAPP-5516	Jonathan, can you comment on the behaviour of Arjuna as noticed by Carlo?	0	0	https://issues.jboss.org/browse/JBPAPP-5516
9939	JBPAPP-5496	Hi Marek, In order to proceed with release notes, can I please have more information about how this affected the customer? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-5496
9940	JBPAPP-5463	There are still two dependencies containing -brew in their versions in root.pom.xml. Is this allowable? <dependency> <groupId>quartz</groupId> <artifactId>quartz</artifactId> <version>1.5.2-brew</version> </dependency> <dependency> <groupId>dom4j</groupId> <artifactId>dom4j</artifactId> <version>1.6.1-brew</version> <exclusions> <exclusion> <groupId>xml-apis</groupId> <artifactId>xml-apis</artifactId> </exclusion> </exclusions> </dependency>	1	1	https://issues.jboss.org/browse/JBPAPP-5463
9941	JBPAPP-5461	Tyronne, Could you please advise as to where information on the missing configuration options you mentioned when you reopened this JIRA?	0	0	https://issues.jboss.org/browse/JBPAPP-5461
9942	JBPAPP-5456	Should this be an "Upgrade JSF to 1.1.15.B1-patch01" JIRA following the upgrade rules (docspace, parent JIRA)? For me it souds like it should. Thanks for the patch BTW	1	1	https://issues.jboss.org/browse/JBPAPP-5456
9943	JBPAPP-5456	Why was the license on Classpath.java changed from CDDL to ASL 2.0? Was this intentional (since the newer jsf2 code made a switch to ASL 2.0)? There is also an ASL 2.0 license on the new file VFSClasspath.java, and I assume that is a compatible license, although it does not share the same license as the rest of the files. I noticed that our RPM said 'License: ASL 2.0', but this doesn't seem to be the case. All of the code appears to be under the CDDL License. According to <https://fedoraproject.org/wiki/Packaging/LicensingGuidelines>, the package must list both licenses, plus include a comment with the license breakdown, e.g.: # The entire source code is CDDL except Classpath.java and VFSClasspath.java which are ASL 2.0 License: CDDL and ASL 2.0	1	1	https://issues.jboss.org/browse/JBPAPP-5456
9944	JBPAPP-5456	java.lang.AssertionError: java.lang.RuntimeException: java.net.MalformedURLException: no !/ in spec at org.jboss.seam.mock.AbstractSeamTest$Request.onException(AbstractSeamTest.java:455) at org.jboss.seam.mock.AbstractSeamTest$Request$2.doFilter(AbstractSeamTest.java:504) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:83) at org.jboss.seam.web.IdentityFilter.doFilter(IdentityFilter.java:40) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.MultipartFilter.doFilter(MultipartFilter.java:90) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.ExceptionFilter.doFilter(ExceptionFilter.java:64) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.RedirectFilter.doFilter(RedirectFilter.java:45) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.ajax4jsf.webapp.BaseXMLFilter.doXmlFilter(BaseXMLFilter.java:178) at org.ajax4jsf.webapp.BaseFilter.handleRequest(BaseFilter.java:295) at org.ajax4jsf.webapp.BaseFilter.processUploadsAndHandleRequest(BaseFilter.java:373) at org.ajax4jsf.webapp.BaseFilter.doFilter(BaseFilter.java:500) at org.jboss.seam.web.Ajax4jsfFilter.doFilter(Ajax4jsfFilter.java:56) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.LoggingFilter.doFilter(LoggingFilter.java:60) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.web.HotDeployFilter.doFilter(HotDeployFilter.java:53) at org.jboss.seam.servlet.SeamFilter$FilterChainImpl.doFilter(SeamFilter.java:69) at org.jboss.seam.servlet.SeamFilter.doFilter(SeamFilter.java:158) at org.jboss.seam.mock.AbstractSeamTest$Request.run(AbstractSeamTest.java:491) at org.jboss.seam.mock.AbstractSeamTest$FacesRequest.run(AbstractSeamTest.java:862) at org.jboss.seam.example.excel.test.RenderTest.testSimple(RenderTest.java:27) Caused by: java.lang.RuntimeException: java.net.MalformedURLException: no !/ in spec at org.jboss.virtual.plugins.context.AbstractURLHandler.initCacheLastModified(AbstractURLHandler.java:76) at org.jboss.virtual.plugins.context.AbstractURLHandler.(AbstractURLHandler.java:65) at org.jboss.virtual.plugins.context.jar.SynthenticDirEntryHandler.(SynthenticDirEntryHandler.java:72) at org.jboss.virtual.plugins.context.jar.AbstractJarHandler.buildParents(AbstractJarHandler.java:244) at org.jboss.virtual.plugins.context.jar.AbstractJarHandler.initJarFile(AbstractJarHandler.java:179) at org.jboss.virtual.plugins.context.jar.JarHandler.(JarHandler.java:63) at org.jboss.virtual.plugins.context.jar.JarContext.createVirtualFileHandler(JarContext.java:85) at org.jboss.virtual.plugins.context.jar.JarContext.(JarContext.java:58) at org.jboss.virtual.plugins.context.jar.JarContextFactory.getVFS(JarContextFactory.java:51) at org.jboss.virtual.VFS.getVFS(VFS.java:132) at org.jboss.virtual.VFS.getRoot(VFS.java:146) at com.sun.facelets.util.VFSClasspath.search(VFSClasspath.java:48) at com.sun.facelets.util.Classpath.search(Classpath.java:63) at com.sun.facelets.compiler.TagLibraryConfig.loadImplicit(TagLibraryConfig.java:428) at com.sun.facelets.compiler.Compiler.initialize(Compiler.java:87) at com.sun.facelets.compiler.Compiler.compile(Compiler.java:104) at com.sun.facelets.impl.DefaultFaceletFactory.createFacelet(DefaultFaceletFactory.java:218) at	1	0	https://issues.jboss.org/browse/JBPAPP-5456
9945	JBPAPP-5456	How exactly does the full url look like? For some reason the right protocol Handler didn't kick in. e.g. I suspect there is a "jar" prefix present	1	1	https://issues.jboss.org/browse/JBPAPP-5456
9946	JBPAPP-5456	Hi guys, Thanks for all of your hard work. The customer is requesting an update on this patch. Do we have any?	1	0	https://issues.jboss.org/browse/JBPAPP-5456
9947	JBPAPP-5456	How exactly does the url look like? Its protocol is "jar", but then it's missing "!/", as per the spec.	1	1	https://issues.jboss.org/browse/JBPAPP-5456
9948	JBPAPP-5456	>> I don't have a way how to go into VFS sources to inspect what JAR URL could fail in it Debuging? Trace logging? >> there is bundled jboss-vfs 2.0.0.Beta6 in it. Is there a way to replace this ancient version?	1	1	https://issues.jboss.org/browse/JBPAPP-5456

#	Requirement	Questions	Answered	Successfully Answered	URL
9949	JBPAPP-5456	Ales, trace logging is already attached in test-log.zip, see my previous comment. JBoss Embedded dependencies can't be easily changed. I can only add newer version to classpath, what version of jboss-vfs works with your jsf-facelets fix?	1	0	https://issues.jboss.org/browse/JBPAPP-5456
9950	JBPAPP-5456	>> Ales, trace logging is already attached in test-log.zip, see my previous comment. And it has no info about which URL is failing? If not, what about remote debugging? >> what version of jboss-vfs works with your jsf-facelets fix? Take the latest 2.2.x. There are some API changes in between, specially with some util classes, so I hope it should work, otherwise drop to 2.1.x or even 2.0.0.GA.	1	1	https://issues.jboss.org/browse/JBPAPP-5456
9951	JBPAPP-5456	Do we have a update on this? The customer is still waiting for the patch. thanks. Gary	1	0	https://issues.jboss.org/browse/JBPAPP-5456
9952	JBPAPP-5434	This has no component. Where should it go in the Release Notes?	0	0	https://issues.jboss.org/browse/JBPAPP-5434
9953	JBPAPP-5418	Hi Yang, can you build and upload to maven2-brew for us please?	0	0	https://issues.jboss.org/browse/JBPAPP-5418
9954	JBPAPP-5376	Verified on EAP 4.3.0_CP09 CR1. See for example http://hudson.qa.jboss.com/hudson/view/EAP4.3-jbossws/job/eap-4.3-jbossws-sun6/jdk=java16_default,label=RHEL5_x86/lastCompletedBuild/testReport/org.jboss.test.ws.jaxws.jbws2637/	1	1	https://issues.jboss.org/browse/JBPAPP-5376
9955	JBPAPP-5342	Please provide a link to download EAP 4.3 CP10 binary	1	1	https://issues.jboss.org/browse/JBPAPP-5342
9956	JBPAPP-5334	Carlo, apparently there is also an upgrade needed for jboss-ejb3-proxy-impl Is this intentional or a mistake in the 1.0.2 bom? If it is intentional, what is the tag for it?	0	0	https://issues.jboss.org/browse/JBPAPP-5334
9957	JBPAPP-5330	Kabir, can you create the 2.1.6 CP02 tag?	0	0	https://issues.jboss.org/browse/JBPAPP-5330
9958	JBPAPP-5293	The issue here is that the listener is being added after the notification thread has started running. Thus, the map of listeners is modified while the notification thread is iterating over it. SSLSocketImpl does not protect the map from concurrent access and so we hit this exception. Can the listener not be added before the thread is started?	0	0	https://issues.jboss.org/browse/JBPAPP-5293
9959	JBPAPP-5293	I have verified this, doesn't happen in EAP 5.1.1. Is the documentation work done? If it is, please confirm that and/or close.	0	0	https://issues.jboss.org/browse/JBPAPP-5293
9960	JBPAPP-5293	I am experiencing this exception in jboss 6.1.0-FINAL. In the documentation (http://docs.jboss.org/jbossweb/2.1.x/printer/changelog.html) I see that bug JBPAPP-5293 was fixed in Jboss Web 2.1.11.GA, but I can't find this version anywhere. On http://www.jboss.org/jbossweb/downloads/jboss-web the latest available is 2.1.9. Where can I download jbossweb 2.1.11 or higher?	1	0	https://issues.jboss.org/browse/JBPAPP-5293
9961	JBPAPP-5259	Hi Weinan, can you take care of this one?	0	0	https://issues.jboss.org/browse/JBPAPP-5259
9962	JBPAPP-5228	Justin I'm looking at $JBOSS_DIST/docs/examples/binding-manager/sample-bindings.xml, which is the only file called sample-bindings.xml in the distributiuon. What do you mean by a missing tag? The file sizes for sample-bindings.xml in CP08 and CP09 are the same.	0	0	https://issues.jboss.org/browse/JBPAPP-5228
9963	JBPAPP-5227	Hi Mike Just doing the triage on this for 5.1.1. Do you remember what guide you saw this in? I found "apr-utils" in the Installation Guide for EAP 5 so I will assume this bug relates to this. http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Web_Platform/5/html/Installation_Guide/Install_Native_Components.html	0	0	https://issues.jboss.org/browse/JBPAPP-5227
9964	JBPAPP-5214	David, how is this now? How did we set it up for MEAD loading?	0	0	https://issues.jboss.org/browse/JBPAPP-5214
9965	JBPAPP-5193	Isn't this fix to be included in the next EAP 5 release? I don't see any JIRA linked to this one nor any "Upgrade JBss TS to <new version>" filed. Without that this customer may experience a regression when upgrading.	1	0	https://issues.jboss.org/browse/JBPAPP-5193
9966	JBPAPP-5161	Can you check if this has been resolved for EAP 4.3.CP09? The linked issues is marked resolved with a fix version of 4.3.CP09. If that's the case we can resolve this JIRA too.	1	1	https://issues.jboss.org/browse/JBPAPP-5161
9967	JBPAPP-5161	Verified on EAP 4.3.0_CP09 CR1 jboss jaxr in version 1.2.0.SP2-brew https://anonsvn.jboss.org/repos/jbossws/stack/native/tags/jbossws-native-2.0.1.SP2_CP09/version.properties http://anonsvn.jboss.org/repos/jbossas/tags/JBPAPP_4_3_0_GA_CP09/build/build-thirdparty.xml	0	0	https://issues.jboss.org/browse/JBPAPP-5161
9968	JBPAPP-5150	Im still struggling with functional tests for this. I don't know how to correctly reproduce the scenario that throws in unpatched and doesn't throw in patched version i'm attaching sample project where i'm testing this contains two methods: store and store2, both can be called via their respective urls http://localhost:8080/hibernate-sample-war/eventmanager/store http://localhost:8080/hibernate-sample-war/eventmanager/store2 store calls session.close() in transaction which means before autoclose store2 calls session.close() out of transaction, after autoclose in scenario with autoclose=false it shouldnt throw "Session was already closed" when I'm using JTA the autoclose is true regardless of the HibernateMBean setting. It's because of org.hibernate.context.JTASessionContext.buildOrObtainSession() that disregards this setting. do you have any code that I can use to functionally test this ?	0	0	https://issues.jboss.org/browse/JBPAPP-5150
9969	JBPAPP-5150	yeah, that was what I forgot to mention. When I deploy to the unpatched appserver, I comment out the AutoCloseSessionEnabled attribute in the jboss-service.xml. but I have still a problem when I do it as you propose (no AutoCloseSessionEnabled attribute) unpatched: 1. start an unpatched EAP 4.3.0.GA_CP08 2. deploy hibernate-sample.ear 3. visit http://localhost:8080/hibernate-sample-war/eventmanager/store2 4. I get "Session was already closed" exception patched: 1. start an patched EAP 4.3.0.GA_CP08 2. deploy hibernate-sample.ear 3. visit http://localhost:8080/hibernate-sample-war/eventmanager/store2 4. I also get "Session was already closed" exception with http://localhost:8080/hibernate-sample-war/eventmanager/store there's no problem in either patched and unpatched, which is expected because of the logic of the store() method - it closes the session within the transaction and then when after the transaction when the framework tries to autoclose it doesn't complain because it closes only when not already closed. EAP 4.3.0.GA_CP08 seems to be packaged with Hibernate 3.2.4.sp1.cp10 (http://anonsvn.jboss.org/repos/hibernate/core/tags/JBOSS_EAP_3_2_4_SP1_CP10) I can see that by removing the line settings.setProperty(Environment.AUTO_CLOSE_SESSION, "true"); in org.jboss.hibernate.jmx.Hibernate (line 483): the patch effectively changes the default value to false because in org.hibernate.cfg.SettingsFactory (line 146) boolean autoCloseSession = PropertiesHelper.getBoolean(Environment.AUTO_CLOSE_SESSION, properties); the getBoolean method returns false if the property isn't set. but as I mentioned org.hibernate.context.JTASessionContext.buildOrObtainSession() successfully ignores this see lines 117 and 128 which happens also in this example project and that's why I get "Session was already closed" also in patched version. or does it behave differently in your environment ?	0	0	https://issues.jboss.org/browse/JBPAPP-5150
9970	JBPAPP-5148	Well, the fix is supposedly in common-core 2.2.17.GA which is what the EAP 5.1.1 has been updated to. This is from the MEAD build information: Parameters SCM URL: svn+http://anonsvn.jboss.org/repos/common/common-core/tags?2.2.17.GA#3877 So you guys are saying that the 2.2.17.GA version of commons-core did not solve the problem? That was the only reason to switch to that version. The branch has been upgraded to it and the JAR has been available from the repo from quite some time (Feb 11), didn't anybody try it before now? Or maybe the problem requires further investigation?	0	0	https://issues.jboss.org/browse/JBPAPP-5148
9971	JBPAPP-5148	Hi Jean-Frederic, would the upgrade of commons-core to 2.1.17 alone fix this issue?	0	0	https://issues.jboss.org/browse/JBPAPP-5148

#	Requirement	Questions	Answered	Successfully Answered	URL
9972	JBPAPP-5148	Per my last comment, can you make sure the build references the correct commons tag so this is in the EAP 5.1.2 release?	0	0	https://issues.jboss.org/browse/JBPAPP-5148
9973	JBPAPP-5143	Hi Ivo, can you give us a patch or a new tag? We are in a hurry to rebuild for ER2. It is not a blocker, so if you cannot do this now it is OK, but if it is a quick thing we will rebuild the plugin with the changes. Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-5143
9974	JBPAPP-5139	What should be there, this is not clear for me, if there are still symlink to real location? "But all of these %{_javadir}/sun-jaxws links should point to the real external locations."	1	1	https://issues.jboss.org/browse/JBPAPP-5139
9975	JBPAPP-5139	Marek, is this fixed? Did you find the right packages?	0	0	https://issues.jboss.org/browse/JBPAPP-5139
9976	JBPAPP-5139	David didn't respond to my question. What are correct packages to use. jboss-seam2 packages uses glassfish-jaxws for jsr173-api.jar, jsr181-api.jar and jboss-ejb-api for jsr250-api.jar. So again my question, Are they correct?	1	1	https://issues.jboss.org/browse/JBPAPP-5139
9977	JBPAPP-5139	I think it is resolved, because tatletale report was reviewed and nobody filed the jira against differences between these 3 jars in JBoss AS and Seam distribution. Do you agree?	1	1	https://issues.jboss.org/browse/JBPAPP-5139
9978	JBPAPP-5094	Hi Elspeth Could you please make the changes specified in this ticket this morning. The task estimation is 1hr, so because you've never done it before I'll expect the task to be finished in two. Cheers J	1	1	https://issues.jboss.org/browse/JBPAPP-5094
9979	JBPAPP-5094	Hey Elspeth, is this issue resolved now? From a quick glance it looks like you've done all the base roll-ins. If so, just Resolve the issue and assign it to Rajesh R.	1	1	https://issues.jboss.org/browse/JBPAPP-5094
9980	JBPAPP-5094	Hi Elspeth Can I roll in the other changes you've still got outstanding while you're doing the diff task? Or will that clash with the rebase you're currently working on?	1	1	https://issues.jboss.org/browse/JBPAPP-5094
9981	JBPAPP-5092	>Section 38.5.2 >in the Note change "HornetQ 2.0.0" to "HornetQ 2.1.x" Changed to HornetQ 2.2 <-- is this correct, or does the issue only relate to HQ 2.1.x	1	1	https://issues.jboss.org/browse/JBPAPP-5092
9982	JBPAPP-5092	Hi Jared, I have found new issues. I am reopening this JIRA. Thank you. Section 3.1 Starting The Server This section should be changed. We have automated installer which replaces JBoss Messaging with HornetQ. User do not have to create profiles. And “default-with-hornetq” is not correct profile. Section 5. Core Architecture “This allows HornetQ to be instantiated in any dependency injection framework such as JBoss Microcontainer, or Spring.” We do not support standalone HornetQ and we should remove such as sentences. Section 6. Using the Server “For the remainder of this chapter ...” We do not support standalone server, remove this sentence it (e.g. JNDI service is provided by EAP)? Section 6.1 Configuration files “... in the HornetQ directory structure ...” HornetQ is installed into EAP dir. structure, this sentence is valid only for standalone server. 7.4 JNDI Configuration We should remove text between "In the default ..." to the end of the section. JNDI is provided by EAP. Section 10.1 and 10.2 We should remove sentence “If the client runs inside a Java 5 virtual machine, use instead hornetq-core-client-java5.jar.” We do not ship special jar for JDK5. Chapter 32. Application Server Integration and Java EE This chapter uses hornetq-ra.rar on several spots. HornetQ is installed as full replacement of JBoss Messaging and it uses jms-ra.rar. But explicit declaration of @ResourceAdapter("jms-ra.rar") is not required. Section 32.5.1 “can be found inside the hornetq-ra.rar archive” EAP 5 does not include this archive file. Section 32.5.2 Configuring JBoss 5 MDB should be renamed, this name does not seem correct to me.	0	0	https://issues.jboss.org/browse/JBPAPP-5092
9983	JBPAPP-5092	{quote}Hi Jared, I have found new issues. I am reopening this JIRA. Thank you.{quote} Thanks for these observations. Comments inline: {quote}Section 3.1 Starting The Server This section should be changed. We have automated installer which replaces JBoss Messaging with HornetQ. User do not have to create profiles. And “default-with-hornetq” is not correct profile.{quote} This section was removed completely based on feedback in the Bugzilla issues. {quote}Section 5. Core Architecture “This allows HornetQ to be instantiated in any dependency injection framework such as JBoss Microcontainer, or Spring.” We do not support standalone HornetQ and we should remove such as sentences.{quote} Thanks for this pick-up. It seems that as soon as Tyronne and I think we have removed all info about "standalone" another one sticks it's head up to bite us. {quote}Section 6. Using the Server “For the remainder of this chapter ...” We do not support standalone server, remove this sentence it (e.g. JNDI service is provided by EAP)?{quote} And another one... Removed. {quote}Section 6.1 Configuration files “... in the HornetQ directory structure ...” HornetQ is installed into EAP dir. structure, this sentence is valid only for standalone server.{quote} I have changed "in the hornetq directory structure" to "in the JBoss Enterprise Application Platform directory structure". {quote}7.4 JNDI Configuration We should remove text between "In the default ..." to the end of the section. JNDI is provided by EAP.{quote} Yep, the same observation was made by Tyronne and Andrew Ross in the Bugzilla HQ issues. It was removed as part of that feedback. {quote}Section 10.1 and 10.2 We should remove sentence “If the client runs inside a Java 5 virtual machine, use instead hornetq-core-client-java5.jar.” We do not ship special jar for JDK5.{quote} This section was removed as part of BZ feedback. {quote}Chapter 32. Application Server Integration and Java EE This chapter uses hornetq-ra.rar on several spots. HornetQ is installed as full replacement of JBoss Messaging and it uses jms-ra.rar. But explicit declaration of @ResourceAdapter("jms-ra.rar") is not required.{quote} I think the same observation was made in a number of BZ issues. All the references I can see in this chapter now refer to the jms rar {quote}Section 32.5.1 “can be found inside the hornetq-ra.rar archive” EAP 5 does not include this archive file.{quote} I can't locate this section in the new guide. Since this JIRA was created, I have removed some sections completely based on Bugzilla feedback. I suspect this happened in this case. {quote}Section 32.5.2 Configuring JBoss 5 MDB should be renamed, this name does not seem correct to me.{quote} This name has either been changed or the section is no longer present in the document.	1	1	https://issues.jboss.org/browse/JBPAPP-5092
9984	JBPAPP-5092	Hi Jared, some of theese issues are still not fixed so I am reopening this JIRA. Could you look at it please? Thank you, Nikoleta ========== Section 2.5 Migrating JMS-administered Objects and Bridges correct Table 2.5 - "text overflow" ========== Section 32.1.1 the paragraph: " The TransactionManagement annotation tells the container to manage the transaction. The TransactionAttribute annotation tells the container that a JTA transaction is required for this MDB. Note that the only other valid value for this is TransactionAttributeType.NOT_SUPPORTED which tells the container that this MDB does not support JTA transactions and one should not be created. " – it seems that it has a bad aligning but now it is Section 30.1.1. ========== Section 32.4.2 "</connection-definition>" in the configuration example is a bad aligning - move to right now Section 30.4.2. ========== Section 33.3 "Different implementations can be provided with the org.hornetq.jms.bridge.DestinationFactory interface. " - seems to have a bad aligning now Section 31.3. ========== Section 39.1.1.1.1 in the both configuration examples is a bad aligning "</factory-class>" now in sections 5.3 and 43.1.5 ==========	1	1	https://issues.jboss.org/browse/JBPAPP-5092

#	Requirement	Questions	Answered	Successfully Answered	URL
9985	JBPAPP-5090	Peter, the content I have in there already is about Remote Method Invocation of EJB3s - is this information for something different? It talks about configuring servlets. How does this relate to the existing content, about configuring the server for EJB3 RMI over SSL?	1	1	https://issues.jboss.org/browse/JBPAPP-5090
9986	JBPAPP-5090	For the webconnector there is no keypair value specified. What happens when you use a keystore that has more than one key pair in it? Can we use keyAlias on the connector like this: https://jira.jboss.org/browse/JBPAPP-538 ?	0	0	https://issues.jboss.org/browse/JBPAPP-5090
9987	JBPAPP-5089	Peter, this section has moved to here: http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/Security_Guide/index.html#id3130166 The book name has been changed to "Security Guide", and the chapter has moved to chapter 19. To this specific JIRA issue: What about if we change it from "java -cp" to "java -jar " - does that remove the need to specify the classpath? I'd like to keep the example as simple as possible.	0	0	https://issues.jboss.org/browse/JBPAPP-5089
9988	JBPAPP-5088	Wow, that's important - thanks for catching that! On my running EAP 5.1 instance production profile https://server_name:8443/invoker/JNDIFactory and http://server_name:8443/invoker/JNDIFactory both do not connect to a running service. Are you saying that the CN of the certificate has to be the same as the SERVER_NAME that a client uses to connect with the server? If so, where is this configured on the client?	1	1	https://issues.jboss.org/browse/JBPAPP-5088
9989	JBPAPP-5086	Marek, On this issue there seems to be some confusion as to which JARs have definitely been removed and which have not. Can you please confirm this? I am currently writing and compiling Release Notes for EAP 5.1.2 and I need this information to be able to provide the customer with accurate information. Please ACK this message by Thursday, 1 September 2011 and provide a response by Monday, 5 September 2011. Thankyou, Russell Dickenson (Content Author)	1	1	https://issues.jboss.org/browse/JBPAPP-5086
9990	JBPAPP-5061	Email sent to Marcus Moyses asking for help: "Marcus, I need help to make progress on a JIRA issue and hope that you might be able to guide me. The specific issue is https://issues.jboss.org/browse/JBPAPP-5061. I have looked at the commit diffs of the changes made in linked JIRA issues (3583 for example) but still don't understand what needs to be done with respect to documentation. In the EAP Security Guide (http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/Security_Guide/index.html#sect-JaasSecurityDomain_MBean) is mention of KeyStorePass and TrustStoreType but I'm not sure if something needs to be added to their description. Are you able to help? If not, can you suggest someone else I might contact? -- Russell Dickenson Engineering Content Services Red Hat Brisbane, Australia"	1	1	https://issues.jboss.org/browse/JBPAPP-5061
9991	JBPAPP-5061	Reply from Marcus: "Hi Russel, I think I added some descriptions of those options in http://community.jboss.org/wiki/JaasSecurityDomain Let me know if there is anything else you need. Regards On 10/25/2011 01:59 AM, Russell Dickenson wrote: > Marcus, > > I need help to make progress on a JIRA issue and hope that you might be able to guide me. The specific issue is https://issues.jboss.org/browse/JBPAPP-5061. > > I have looked at the commit diffs of the changes made in linked JIRA issues (3583 for example) but still don't understand what needs to be done with respect to documentation. In the EAP Security Guide (http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/Security_Guide/index.html#sect-JaasSecurityDomain_MBean) is mention of KeyStorePass and TrustStoreType but I'm not sure if something needs to be added to their description. > > Are you able to help? If not, can you suggest someone else I might contact? > > -- Marcus Moyses Senior JBoss Core Developer JBoss by Red Hat"	1	1	https://issues.jboss.org/browse/JBPAPP-5061
9992	JBPAPP-5020	Alessio, is there any harm in leaving the saaj-api.jar around when CXF is being used?	0	0	https://issues.jboss.org/browse/JBPAPP-5020
9993	JBPAPP-5016	Remy, what will those defaults be? I currently have a section in the performance tuning guide. It doesn't really cover everything, just the things I have tested that seem to show improvement, but I do talk about what the defaults are, so depending on what you are changing here, I will have to update the guide. If you are changing things that don't overlap than no problem. So, if you could shoot me a note that would be good. Also, I would like to test those defaults in my environment to see if they help or hurt my own tests. Thanks.	1	0	https://issues.jboss.org/browse/JBPAPP-5016
9994	JBPAPP-5009	Are you expecting to back port this to EAP 4.3 and 5.1? If so, please clone the JIRA and have one in TBD_EAP4 and one in TBD_EAP5. Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-5009
9995	JBPAPP-5006	Martin, can you verify the instructions?	1	1	https://issues.jboss.org/browse/JBPAPP-5006
9996	JBPAPP-4975	So Step 3 currently reads: <title>Deploy or package the deployment descriptor</title> <para>Move the deployment descriptor file to the <filename>deploy</filename> directory of the required server profile in your installation. </para> <para>Alternatively, package the deployment descriptor in the <filename>META-INF</filename> directory of the EJB Jar, or <filename>WEB-INF/classes</filename> of your web application (WAR).</para> If I understand correctly, you want me to change the Step to the following: 3. Deploy or package the deployment descriptor Deploy or package the deployment descriptor in a WAR file, and ensure the WAR is deployed to the WEB-INF folder. Could you please confirm I understand correctly?	1	1	https://issues.jboss.org/browse/JBPAPP-4975
9997	JBPAPP-4970	Marek, I still see it there: jboss-eap-5.1.1.ER1/seam/lib/hsqldb.jar jboss-eap-5.1.1.ER1/seam/bootstrap/deploy/messaging/hsqldb-persistence-service.xml jboss-eap-5.1.1.ER1/seam/bootstrap/deploy/hsqldb-ds.xml Is this a problem?	0	0	https://issues.jboss.org/browse/JBPAPP-4970
9998	JBPAPP-4970	Attention: Marek Novotny Should this issue's Release Notes text remain as it is or not? Since the issue has been reopened I expect perhaps it should *not* be included in the EAP 5.1.2 Release Notes document.	1	1	https://issues.jboss.org/browse/JBPAPP-4970
9999	JBPAPP-4946	/mnt/hudson_workspace/workspace/JBoss-EAP-5.1.x-Testsuite-sun16/JBossEAP_5_1_x/cluster/src/main/org/jboss/ha/framework/server/ClusterPartition.java:1742: cannot find symbol symbol : variable cvs location: class org.jgroups.Version return Version.description + "(" + Version.cvs + ")"; This rings a bell. Didn't we fix for Support a JAR that did not have the Version class?	1	1	https://issues.jboss.org/browse/JBPAPP-4946
10000	JBPAPP-4945	Hi Fernando, jettison-1.2-4.ep5.el5 and jettison-1.2-4.ep4.el5 are built with pom and repolib. I saw you've already built jettison 1.2 in rhel6, thank you for help! BTW, which hudson task I should check in future to ensure my build won't break anything? Could you please provide an url? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-4945

#	Requirement	Questions	Answered	Successfully Answered	URL
10001	JBPAPP-4933	Rebecca, Can you add these comments to the release notes? - Incorrect Numbering: 2.1. Consolidated Trust Store should actually be 2.1.1 and the others following it. Features introduced in 5.1 and features from previous releases are the two major sections. - Section 2.1.1 - This is the first time we are shipping the NSAPI connector,we should mention it. NSAPI is a connector used to connect Sun Java Web server to JBoss Enterprise Application Platform. - Important: Section 2.3 component versions for 5.1 are not in sync with what we are shipping. Some of the components had updates in some of the internal builds Cr2 and CR3. The latest can be found here. https://docspace.corp.redhat.com/docs/DOC-37998 Most of the current versions correspond to CR1 versions documented initially. - Section 2.4.4 Source zip should be ftp://ftp.redhat.com/pub/redhat/jbeap/5.1.0/en/source/jboss-eap-src-5.1.0.zip - Section 2, features introduced should talk about JDBC4 support which is offered from this release onwards Hibernate offers support for JDBC4 from this release onwards. https://jira.jboss.org/browse/JBPAPP-4180 Support for the connector component had existed from earlier. We should check with Jesper Pederson (dev) or Rich Sharples on what we want to call out here. - Section 2 should talk more about performance improvments made in this release.	1	1	https://issues.jboss.org/browse/JBPAPP-4933
10002	JBPAPP-4933	The above issues have been corrected, with the exception of: - JDBC4 support - I've emailed Rich and Jesper for assistance with this. - Performance improvements - I've emailed Andy Miller for assistance, but if you can suggest an alternate source that would be great. - Consolidated Trust Store - this section has been commented out as not applying to EAP 5.1. Is this correct?	1	1	https://issues.jboss.org/browse/JBPAPP-4933
10003	JBPAPP-4933	Hi Laura. I have found that known issue JBPAPP-2571 has been fixed. Related JBM issue (JBMESSAGING-1767) has been fixed in 1.4.7.GA which is included in EAP 5.1. Please could you remove note about this issue from section "Known Issues with this release"? Thank you.	1	1	https://issues.jboss.org/browse/JBPAPP-4933
10004	JBPAPP-4933	Moved JBPAPP-2571 from known to fixed in EAP 5.1 Release Notes. Renamed jboss-as to jboss-as-web three times in EWP 5.1 Release Notes. Removed JBPAPP-5014 and JBPAPP-2571 from EWP 5.1 Release Notes. Corrected JBoss Messaging and JBoss Web Services CXF versions in section 2.6 of RN as required. However, I was under the impression that a CXF tech preview of 3.1.2.SP6 had gone out alongside/post-release 5.0.1. I've removed the previous version in case I was incorrect, but can someone confirm/deny?	1	1	https://issues.jboss.org/browse/JBPAPP-4933
10005	JBPAPP-4933	Thanks! Looks good to me now. I am closing but how are we going to resolve the structure problem? Do you want a new Jira for it? "The issue with there being only one chapter is due to this book actually being an article - the form is different. I will fix this for the next release, but for now I am leaving it as it is. Every section should now have a heading, though. "	1	0	https://issues.jboss.org/browse/JBPAPP-4933
10006	JBPAPP-4920	Could I please get a description of this issue for the release notes? I need to know what was causing the issue, what effect that had, and how it was fixed. Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-4920
10007	JBPAPP-4902	Hi Hui, can you update this one for us and upload to maven2-brew? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-4902
10008	JBPAPP-4866	Hi, is there any test for this issue? Thanks Mirek	1	1	https://issues.jboss.org/browse/JBPAPP-4866
10009	JBPAPP-4866	How to run JBM 1.4 tests in QA Lab: cp /qa/tools/opt/jdbc-drivers/mysql-connector-java-5.0.4-bin.jar lib/jdbc-drivers/ echo '<container><database>mysql</database><database-configurations><database-configuration name="mysql"><url>jdbc:mysql://vmg08.mw.lab.eng.bos.redhat.com:3306/msgcc</url><driver>com.mysql.jdbc.Driver</driver><isolation>TRANSACTION_READ_COMMITTED</isolation><username>msgcc</username><password>msgcc</password></database-configuration></database-configurations><remoting-transport>bisocket</remoting-transport></container>' > etc/container.xml ant tests -Dintegration.base=EAP4 -Dtest.bind.address=`hostname` -DfakeOption=fakejms-mysql -Djboss.messaging.groupname=JBM_GRP4 -Djboss.messaging.datachanneludpport=45573 -Djboss.messaging.controlchanneludpport=45574 -Djboss.messaging.datachanneludpaddress=228.6.6.9 -Djboss.messaging.controlchanneludpaddress=228.7.7.10	1	1	https://issues.jboss.org/browse/JBPAPP-4866
10010	JBPAPP-4860	Is EWS able to support the JBoss Cache-based mechanism described in this "HTTP Connectors Load Balancing Guide"? I don't think it can, since AIUI it's just Tomcat with none of the integration hooks EAP provides for running JBoss Cache. (I'm phrasing this as a question since I don't follow this product closely, but I'd be shocked if EWS supports JBC-based session replication.) I don't think Clustered SSO would work on EWS either, for the same reason. The approach being documented is an alternative to the JBoss Cache based session replication discussed in http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html/HTTP_Connectors_Load_Balancing_Guide/clustering-http-state.html. It most logically fits along with those. I don't know how best to resolve the issue that much of the content there doesn't apply to EWS.	1	1	https://issues.jboss.org/browse/JBPAPP-4860
10011	JBPAPP-4860	I agree that this is the right spot: http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html/HTTP_Connectors_Load_Balancing_Guide/clustering-http-state.html An option would be to add an admonition to the chapter that this is not support on EWS. Would that be ok?	1	1	https://issues.jboss.org/browse/JBPAPP-4860
10012	JBPAPP-4860	Hi Paul, Could you, please, let me know, who is the person that can review the fix? Thanks	0	0	https://issues.jboss.org/browse/JBPAPP-4860
10013	JBPAPP-4835	In which case, can it be moved into fix version TBD EAP 5, and cloned for fix version TBD EAP 6 so that we can triage it for release notes properly? Thanks.	1	0	https://issues.jboss.org/browse/JBPAPP-4835
10014	JBPAPP-4828	Hi Fernando Can you help with the info for the upgrade procedure? If not, please flick the JIRA to someone who can assist. Cheers Jared	1	1	https://issues.jboss.org/browse/JBPAPP-4828
10015	JBPAPP-4828	Fernando responded by email: > Is there currently an upgrade path from 5.1 to 5.1.x for RPMs? I know you could probably just run up2date (RHEL 4) or yum update (RHEL 5) in most cases, but are there any customer gotchas that need to be specified in the upgrade instructions of the Installation Guide? No, just the usual recommendation, backup your stuff, look for .rpmsave or .rpmnew files if you have modified config files and merge you changes manually, Ahh, this is also standard for servers, but always good to emphasize, "stop the server before upgrading", restart afterwards. Another caveat: it is not supported to change the option made on the initial install (EAP 5.1.0) of CXF for the WS stack, the upgrade will keep that option and attempts to change it will have unpredictable results. The same is valif for the new option, HormetQ for JMS, it requires a new installation, not an upgrade. Note, we haven't yet finished the set of RPMs, never tried ourselves an upgrade from 5.1.1, we may find things we don't know about. Cheers, Fernando	0	0	https://issues.jboss.org/browse/JBPAPP-4828
10016	JBPAPP-4828	On QA we can't verify update instructions from 5.1.x to 5.1.2 (= added section to IG), because some blockers occurred (e.g. [JBPAPP-7272|https://issues.jboss.org/browse/JBPAPP-7272]) in ER1 build. So this verification should be postponed to ER2 build. Anne-Louise, could you please agree with my proposal and modify Fix version of this JIRA?	0	0	https://issues.jboss.org/browse/JBPAPP-4828
10017	JBPAPP-4828	Why is this reopened and assigned to me? If this is about instructions for how to install RPMs, it needs to be assigned to Fernando or someone on his team.	0	0	https://issues.jboss.org/browse/JBPAPP-4828

#	Requirement	Questions	Answered	Successfully Answered	URL
10018	JBPAPP-4796	Hey Laura, re: "release notes flag" - not sure what you mean. Are you talking to me or the person who will edit the build-thirdparty.xml file? -Ron	1	1	https://issues.jboss.org/browse/JBPAPP-4796
10019	JBPAPP-4795	Hi Marcus, is the new tag available?	0	0	https://issues.jboss.org/browse/JBPAPP-4795
10020	JBPAPP-4795	I don't think so. Can you ping the project lead (Alexey Loubyansky) please?	0	0	https://issues.jboss.org/browse/JBPAPP-4795
10021	JBPAPP-4795	No Fernando, this *is* the Jira for tracking that update but you changed its summary to '1.0.5' on 20/Aug/'10. The original description says it clearly [1]. Please try to keep Jira clean, makes it impossible for QE to verify the Jiras. [1] "Because of JBPAPP-4786 we need a new release of JBoss Metada (1.0.6?) and modify EAP 5.1 component matrix to use this version."	0	0	https://issues.jboss.org/browse/JBPAPP-4795
10022	JBPAPP-4795	Alexey, can you provide a tag for metadata 1.0.6 including a fix for the related issue?	1	1	https://issues.jboss.org/browse/JBPAPP-4795
10023	JBPAPP-4778	Carlo, what's a FIRO queue? Did you mean FILO?	1	0	https://issues.jboss.org/browse/JBPAPP-4778
10024	JBPAPP-4778	Carlo, what documentation is this patch for? Which book?	0	0	https://issues.jboss.org/browse/JBPAPP-4778
10025	JBPAPP-4778	One other thought. Why are we using jboss-beans.xml to try to accomplish this, instead of changing the code directly? That also seems like the wrong thing to do to me.	1	1	https://issues.jboss.org/browse/JBPAPP-4778
10026	JBPAPP-4778	Yeah, I'm aware of the pool size cannot be infinite, in terms of still getting the best possible throughput. I'm not so sure that the pool size as set is right though. The pool size for the HTTP connector is only 200 based on the document above. The best throughput is actually achieved using 600, not 200. So, I would argue its too small. The difference being that throttling the workload with a smaller thread pool just allows it to continue to run at higher number of clients, so our disagreement on this is really over what's most important. Supporting a larger number of clients, or getting the best possible throughput? When customers do their own benchmarks, they typically look for the best possible throughput. You answered Carlo's mistaken claim about the pool size not actually changing through the tests, just the number of sessions, or virtual users. The real question here is what do we do with this? I have been recommending that we close this issue, and not push this. The fact that it also breaks the metrics on the thread pool (admin-console and JON), make it a non-started as well. Can we agree not to apply this? Certainly it cannot be applied as is anyway since it breaks the admin console and JON.	0	0	https://issues.jboss.org/browse/JBPAPP-4778
10027	JBPAPP-4776	Verified the note about supported browsers is present in http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html/Administration_Console_Quick_Start_Guide/Administration_Console_Quick_Start_Guide-Accessing_the_Console.html . Closing the issue.	0	0	https://issues.jboss.org/browse/JBPAPP-4776
10028	JBPAPP-4765	Ondrej, can you verify this in the CR3.5 build? I think this has been marked closed by Dev.	0	0	https://issues.jboss.org/browse/JBPAPP-4765
10029	JBPAPP-4763	Ondrej, can you continue on the investigation reported here?	1	0	https://issues.jboss.org/browse/JBPAPP-4763
10030	JBPAPP-4763	Can you set the component? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-4763
10031	JBPAPP-4762	Marking it as a blocker since it is a performance issue which needs to be investigated for the EAP 5.1 release. Bruno, could you assign this to somebody in dev?	0	0	https://issues.jboss.org/browse/JBPAPP-4762
10032	JBPAPP-4759	Hi Andy, the final date for this document to go to QE is August 18. We will need it a few days before that to check it for style, etc. Josh tells me you are writing the content - can you give me a date you will be able to hand the doc off to us before QE?	0	0	https://issues.jboss.org/browse/JBPAPP-4759
10033	JBPAPP-4759	Andy, I had created a perf guide for EPP 4.3 but could not get a chance to create one for EPP5. I would like to create one for EPP 5.1 because couple of customers have already asked for it. IIUC, perf guide for EAP 5.1 is slated to be created post release. Do you have any ETA on this? It will help me immensely.	1	1	https://issues.jboss.org/browse/JBPAPP-4759
10034	JBPAPP-4759	Having all these issues added to this container ticket will also help with Docs QE and Engineering QE. Is this guide going to make 5.1.1 Post GA? Should it be realigned with the normal 5.1.1 release?	1	1	https://issues.jboss.org/browse/JBPAPP-4759
10035	JBPAPP-4715	If this needs to go in the Release Notes / Installation Guide, can I get an official summary of the problem, its effects on customers, and any workarounds?	1	1	https://issues.jboss.org/browse/JBPAPP-4715
10036	JBPAPP-4715	Rich, is it also valid for clients using older JDKs, like JBoss EAP 4.3 client using JBoss EAP 5.x server?	0	0	https://issues.jboss.org/browse/JBPAPP-4715
10037	JBPAPP-4715	HI guys, can I echo Laura in asking for a summary, effects on customers, and any workarounds, for the Release Notes, please? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-4715
10038	JBPAPP-4691	Why was "PTG" added to the issue summary? I am not familiar with this abbreviation and I am linking to this Jira in a report and this is adding unnecessary confusion. If its something internal, tags should be used instead.	0	0	https://issues.jboss.org/browse/JBPAPP-4691
10039	JBPAPP-4691	Attention: Andy Has the information requested in this issue been included in the Performance Tuning Guide? If not, where would you recommend it be added? Maybe the "Stateless Session Beans" section?	1	0	https://issues.jboss.org/browse/JBPAPP-4691
10040	JBPAPP-4691	I have added the following text to the PTG as section 4.1.4: " <title>Local EJB3 Calls</title> <para> When an application uses a remote EJB3 interface the call is serialized and deserialized by default because this is a requirement of network communication. If the EJB is running in the same JVM as the client calling it, all communications are local so there's no need for serialization or deserialization. To avoid the CPU load involved and so improve the application's responsiveness, enable the system property org.jboss.ejb3.IsLocalInterceptor.passByRef to true, for example: add -Dorg.jboss.ejb3.remoting.IsLocalInterceptor.passByRef=true to JAVA_OPTS in the server's <filename>run.conf</filename> configuration file. </para>" Does this accurately reflect what you intended, Andy?	0	0	https://issues.jboss.org/browse/JBPAPP-4691
10041	JBPAPP-4684	Ivo, can you verify this in CR3.5?	1	0	https://issues.jboss.org/browse/JBPAPP-4684
10042	JBPAPP-4653	Will this change?	1	0	https://issues.jboss.org/browse/JBPAPP-4653
10043	JBPAPP-4611	Peter, can you verify this in CR3?	1	0	https://issues.jboss.org/browse/JBPAPP-4611
10044	JBPAPP-4611	There are two jars in the EAP 5.1.0.CR3 containing other jars: ./server/default/deploy/admin-console.war/plugins/rhq-jmx-plugin.jar [1] ./server/default/deploy/admin-console.war/plugins/jopr-jboss-as-5-plugin.jar [2] The first [1] jar contains only signed jars which is OK. The second [2] jar contains two jars which are not signed: ./lib/rhq-jboss-as-common-3.0.0.EmbJopr2.jar ./lib/jbpm-3.1.1.jar From Common Criteria or Java Security Manager policy point of view I am fine with those two not signed jars, because policy can reference them by location and since we have just two of them it is not going to make policy unreadable. The other argument is that those jars are part of JBoss Admin Console which code cannot be used by other (user's) applications. I would reopen the issue and move fix version to the next release. Anil, what do you think?	1	1	https://issues.jboss.org/browse/JBPAPP-4611
10045	JBPAPP-4611	Peter, can you check in CR3.5 if this issue still holds? The above mentioned jars are now built in brew.	1	0	https://issues.jboss.org/browse/JBPAPP-4611
10046	JBPAPP-4602	Should it be built as part of the WS build? Maybe as a separate build? Alessio, any suggestions?	1	0	https://issues.jboss.org/browse/JBPAPP-4602

#	Requirement	Questions	Answered	Successfully Answered	URL
10047	JBPAPP-4601	Although this is a bit older JIRA, I can't find a related info in "Transactions Development Quick Start Guide" nor "Transactions Administrators Guide". Is current state of docs about JBossXTS as supposed to be?	1	1	https://issues.jboss.org/browse/JBPAPP-4601
10048	JBPAPP-4589	Martha, I assume this will need to be documented - can you or Yiqiao create a documentation JIRA with instructions, including any channel subscription/JDK requirements?	0	0	https://issues.jboss.org/browse/JBPAPP-4589
10049	JBPAPP-4551	Chris, can you tell me where users will be able to download the HornetQ and CXF zips and what they are likely to be named? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-4551
10050	JBPAPP-4551	Could you put the changes to the stage, please?	1	1	https://issues.jboss.org/browse/JBPAPP-4551
10051	JBPAPP-4551	I was not able to find the fixes on the stage. Now I can see there are two versions for EAP on the stage, 5.0 and 5. So now there are also two versions of the Installation Guide, one under the '5.0' section and the other one under the '5' section. Which of these versions is final and should be QE'd?	1	1	https://issues.jboss.org/browse/JBPAPP-4551
10052	JBPAPP-4546	password_tool.bat is in the $JBOSS_HOME/bin, but after trying to specify keystore (option 1) you get ClassNotFoundException (see at the bottom). password_tool.bat is missing following line: set JBOSSPASS_CLASSPATH=%JBOSSPASS_CLASSPATH%;%JBOSS_HOME%\client\jbosssx-client.jar Can you add the line, please? Exception in thread "main" java.lang.NoClassDefFoundError: org/jboss/security/plugins/FilePassword at org.jboss.security.integration.password.PasswordMaskManagement.ensureKeyStore(PasswordMaskManagement.java:297) at org.jboss.security.integration.password.PasswordMaskManagement.setKeyStoreDetails(PasswordMaskManagement.java:83) at org.jboss.security.integration.password.PasswordTool.main(PasswordTool.java:121) Caused by: java.lang.ClassNotFoundException: org.jboss.security.plugins.FilePassword at java.net.URLClassLoader$1.run(URLClassLoader.java:202) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:307) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:248) ... 3 more	0	0	https://issues.jboss.org/browse/JBPAPP-4546
10053	JBPAPP-4546	At this point it is not clear if a new build will be done for EAP 5.1 which would include this fix. The jira is resolved with a fix version of EAP 5.1, is this critical to include in EAP 5.1?	0	0	https://issues.jboss.org/browse/JBPAPP-4546
10054	JBPAPP-4546	Do we need to re-open this jira and put in TBD EAP 5 or should we add a new JIRA for the missing jar to be included in future release? Also, do we need the release notes to be updated with info about the missing jar and on where to get it?	0	0	https://issues.jboss.org/browse/JBPAPP-4546
10055	JBPAPP-4546	Anil, can you please commit this one line change to the tip of the JBPAPP_5_1 branch? I will try and backport it as a patch if we rebuild.	0	0	https://issues.jboss.org/browse/JBPAPP-4546
10056	JBPAPP-4539	Please see commit link https://issues.jboss.org/browse/JBAS-8344?page=com.atlassian.jirafisheyeplugin:fisheye-issuepanel#issue-tabs When committed, a type [JBAPP] make source link not work	1	1	https://issues.jboss.org/browse/JBPAPP-4539
10057	JBPAPP-4489	Andy, this is another change needed when switching to hornetq, can you please update the script in the JBPAPP_5_1 branch to include this change? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-4489
10058	JBPAPP-4489	Howard, Would it be possible to create a patch or fix it in a tag for EAP 5.1 CR2? We cannot include a patch which causes exception to come up even if it is harmless. We will start building CR2 next Monday, so we'll need something today or tomorrow. Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-4489
10059	JBPAPP-4483	Hi Michael Apologies for introducing the incorrect information back into the section again. It's what happens when I'm trying to close off issues close to the release. Don't look deep enough. Regarding [1] "Instead the users can base their application on the mavenized template WAR project "resteasy-jaxrs-war" that comes with the RESTEasy source distribution. " Would customers have to download that from the community? For enterprise customers, we try to avoid this scenario. Is there a project included in the enterprise binary a customer could access? Regarding [2], I have fixed the XML comments. The reason why the comments were formatted strangely here is because you can't comment out a section in XML that contains Comments. The author would have stripped out the comment markup so they could comment out the section.	1	1	https://issues.jboss.org/browse/JBPAPP-4483
10060	JBPAPP-4483	Hi again Michal That's exactly the information I was looking (and hoping) for. I'll direct the user to the location of the examples in the RESTEasy binary. Are you happy for me to close this issue after incorporating this change? Cheers Jared	1	1	https://issues.jboss.org/browse/JBPAPP-4483
10061	JBPAPP-4464	If I understand correctly the "fix version" will not include 4.3.0 CP09. Can you confirm.	1	1	https://issues.jboss.org/browse/JBPAPP-4464
10062	JBPAPP-4451	Dennis JGRP-1207 has status Rejected. In what sense is this JIRA issue resolved? Am I to understand that no changes were made?	1	1	https://issues.jboss.org/browse/JBPAPP-4451
10063	JBPAPP-4443	Clebert, can you please comment on the above?	0	0	https://issues.jboss.org/browse/JBPAPP-4443
10064	JBPAPP-4443	JORAM Tests are also consumed by AS6. Why don't you add an ignore on the libraries used by the JORAM tests.	1	1	https://issues.jboss.org/browse/JBPAPP-4443
10065	JBPAPP-4440	I have removed the '"-resteasy - a Technology Preview of RESTEasy - a portable implementation of JSR-311 JAX-RS Specification' reference from the Admin and Config guide. Does this resolve the issue?	0	0	https://issues.jboss.org/browse/JBPAPP-4440
10066	JBPAPP-4434	Hi Mladen, you know Solaris better, what do you think? I believe this script comes from the community AS, can this change be incorporated to the trunk if James clone this into a JBAS JIRA?	0	0	https://issues.jboss.org/browse/JBPAPP-4434
10067	JBPAPP-4434	The patch doesn't work. == is not allowed on solaris sh, and -c cannot be first param to grep. Has to be escaped (\-c) Here is the correct one: if ["x${isJBossRunningPid}" != "x"]; then args=`pargs -a ${JBossPid}` test "x${args}" = "x" && return 0 inst=`echo ${args} | grep "\-c ${instance}" 2>/dev/null` test "x${inst}" = "x" && return 0 fi Do you wish me to commit that?	1	1	https://issues.jboss.org/browse/JBPAPP-4434
10068	JBPAPP-4409	Can this please be cloned for EAP 4.x?	1	1	https://issues.jboss.org/browse/JBPAPP-4409

#	Requirement	Questions	Answered	Successfully Answered	URL
10069	JBPAPP-4403	Here are some comments after checking the docs: Chapter 2: Installation References to tech preview to be removed. Also the jbossws-cxf version is wrong. Chapter 3: WS Policy Framework Should be moved just before current Chapter 11 (keep all the general policy stuff together) Chapter 10: Messaging tutorial The chapter title should be something like "WS-Reliable Messaging tutorial" Chapter 12: How it works A reference to ws-policy or policies in general in the chapter title would be good Section 14.3 There's a reference to the community version JBossWS-CXF 3.1.1 that should be removed. The same in 14.4.1. Moreover, if you want to refer stack specific samples, those would need to be actually shipped with the EAP distribution (afaik this is currently not the case). The sample mentioned there is in the community version of jbossws-cxf stack specific testsuite. Otherwise you need to point users to some code online (using fisheye / viewcv) Section 14.4.3 There's a reference to the community jbossws-native stack userguide. Is that allowed? Chapters 14 and 15 should be swapped; as you did for WS-RM and WS-Addressing, first you talk about the topic leveraging the doc comning from Apache, then you provide a sample of integration into JBoss. Chapter 16: SOAP Messaging Logging The cxf-extensions-jbossws.xml file is actually not available on EAP 5.1, so what's written in the note applies; the whole chapter could then be reviewed with that in mind. Chapter 17: JMS Transport tutorial Remove this, it's actually not available in the version included in EAP 5.1 Chapeter 18: CXF configuration Remove this; most of what's there does not apply to the jbossws-cxf integration. Chapter 19: Logging All the logging related stuff (so everything except 19.4) either do not apply to jbossws-cxf integration or is already covered in Chapter 16, hence should be removed.	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10070	JBPAPP-4403	Besides that, is another section of the documentation going to be available for basic ws doc, jaxws, etc like we used to have for native stack?	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10071	JBPAPP-4403	Alessio, which books are you referring to in your last comment? Where would this information come from, and is there a reason it wasn't mentioned when the documentation plan for the CXF docs was put forward?	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10072	JBPAPP-4403	Thanks Alessio. That section will still be available for customers in this release. To avoid duplicating content, I think we should reference that chapter of the Admin and Config Guide. Can you tell us which sections in particular apply to CXF, and where the references to the Admin and Config Guide should be made?	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10073	JBPAPP-4403	Hi Alessio, I have some questions about your comments. 1. We'd rather avoid referencing community documentation, but I don't know where else we could reference for this information. If you could provide information about configuring the Bouncy Castle provider, or suggest another reference point, that would be great 2. When you say, 'there's a reference to the community version JBossWS-CXF 3.1.1, are you simply referring to the version number? 3. In your first comment, you say that the jbossws-cxf version is wrong; can you tell me where this information comes from (is it the downloadable zip file?), and what version we should be using? 4. Could you please clarify what a sample reference is, and what you would like us to do with it?	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10074	JBPAPP-4403	> Thanks Alessio. That section will still be available for customers in this release. To avoid duplicating content, I think we should reference that chapter of the Admin and Config Guide. Can you tell us which sections in particular apply to CXF, and where the references to the Admin and Config Guide should be made? Here is the list of what applies to *JBossWS-CXF* too: - Chapters 5.1 (just remove the sentence "Even though JAXRPC..." as JAX-RPC is NOT supported with the CXF stack in EAP51), 5.2, 5.3, 5.4, 5.5, 5.6, 5.8, 5.9, 5.10, 5.11, 5.12, 5.13, 5.14, 5.15, 5.16, 5.17, 5.18, 5.19, 5.20, 5.22 (except section 5.22.1.1 Endpoint Config) - Section 5.21.3 As to where put references... well, basically all the stuff above comes logically before everything that's in the current EAP 5.1 CXF guide, as the latter is basically about jbossws-cxf stack specific things.	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10075	JBPAPP-4403	> 1. We'd rather avoid referencing community documentation, but I don't know where else we could reference for this information. If you could provide information about configuring the Bouncy Castle provider, or suggest another reference point, that would be great http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html/Administration_And_Configuration_Guide/ch05s21s02s04.html > 2. When you say, 'there's a reference to the community version JBossWS-CXF 3.1.1, are you simply referring to the version number? Yes, given the version in EAP is 3.1.2.SPx, there's no need to provide and reference to a community version. The feature mentioned there is available and the customer most probably doesn't need to know in which community version of the stack that became available. > 3. In your first comment, you say that the jbossws-cxf version is wrong; can you tell me where this information comes from (is it the downloadable zip file?), and what version we should be using? jbossws-cxf-3.1.2.SP3-3.SP3.6.ep5.el4-tp-installer.zip is the name of the downloadable zip file for EAP 5.0. CXF TP. I don't know which name it will have in EAP 5.1, but for sure that will not have 3.1.2.SP3 in it, given the jbossws-cxf version EAP 5.1 is based on is 3.1.2.SP6. You probably need to get in touch with Fernando's team for the exact file name. > 4. Could you please clarify what a sample reference is, and what you would like us to do with it? A sample reference is a reference to a jbossws class being a testsuite sample. For instance things like "you can check the sources of org.jboss.test.ws.jaxws.samples.SecurityTestCase" As said to Laura, I think we should avoid directly mentioning sources that are not actually shipped in the EAP distribution. Either changes are done for shipping them, or we do not reference them.	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10076	JBPAPP-4403	Hi Alessio, I have made the changes you suggested and staged the book on docstage again. Aside from the section you recommended for removal, are there any other sections that are not directly related to something we support?	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10077	JBPAPP-4403	Hi Alessio, the email you sent was missing two of the documents: EAP51-Rebecca.txt and EAP51-CXF-Doc.txt. Could you send these to me again, please? :) Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10078	JBPAPP-4403	Alessio, can you get back to Rebecca on this? We are waiting on the missing documents mentioned in the previous comment.	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10079	JBPAPP-4403	Hi Alessio, have implemented the changes in the diffs. Can you please check here http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.1/html/CXF_User_Guide/index.html and let me know what you think? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10080	JBPAPP-4403	Alessio, the source document is part of an internal repo for the product docs, so probably not. Bec, can you reapply the missing changes?	0	0	https://issues.jboss.org/browse/JBPAPP-4403
10081	JBPAPP-4403	Hi Alessio, I have NO IDEA what happened to the changes before, but they should be implemented now: http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html/JBoss_WS_CXF_User_Guide/index.html I'm not sure if I've removed too much of the Spring XML Configuration; could you please let me know if I have? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-4403
10082	JBPAPP-4385	How should it look like the patch? Only changed file/s in special branch for the issue?	0	0	https://issues.jboss.org/browse/JBPAPP-4385
10083	JBPAPP-4378	Should this be documented in the release notes? Can I get a summary of the problem and any available workarounds?	1	1	https://issues.jboss.org/browse/JBPAPP-4378

#	Requirement	Questions	Answered	Successfully Answered	URL
10084	JBPAPP-4359	Hi Rostislav, could you please try locally updating to the version there to the one used in EAP 5.1? The maven testsuite builds the classpath using maven dependencies. Assuming this can be fixed just updating the version in the testsuite pom, we need to evaluate whether we want to cut another release for this single change or a simple patch can be applied when running the QA tests.	1	1	https://issues.jboss.org/browse/JBPAPP-4359
10085	JBPAPP-4359	According to http://neptune.usersys.redhat.com:8080/jboss-support/ CR3 has JBoss WS Framework 3.1.2.SP8 and HornetQ 2.1.2 Do we still need this note at all??? Sorry, I don't have time to come up to speed on testing this. Can someone in QE please test this and let us know.	1	1	https://issues.jboss.org/browse/JBPAPP-4359
10086	JBPAPP-4359	Rostislav, I've updated that version number. A couple of questions before I rebuild: 1. We are giving users an instruction to update a pom.xml file - what is the path of that file? 2. We say that tests fail - what tests are these? Is there a file path associated with them?	1	1	https://issues.jboss.org/browse/JBPAPP-4359
10087	JBPAPP-4359	Hi Joshua. This issue applies to EAP 5.1.0 with standard JBossWS Native stack + installed HornetQ messaging, so I suggest this change: The following tests fail with a java.lang.SecurityException error: --> For EAP 5.1.0 with standard JBossWS Native stack + HornetQ messaging the following tests fail with a java.lang.SecurityException error: 1. We are giving users an instruction to update a pom.xml file - what is the path of that file? It's about pom.xml for WS Native testsuite. Relative path is modules/testsuite/pom.xml, SVN repository path is http://anonsvn.jboss.org/repos/jbossws/stack/native/tags/jbossws-native-3.1.2.SP7/modules/testsuite/pom.xml 2. We say that tests fail - what tests are these? Is there a file path associated with them? It's enough to have these test listed as it's now. These tests use JMS as transport protocol.	1	1	https://issues.jboss.org/browse/JBPAPP-4359
10088	JBPAPP-4359	Rostislav, I cannot find the WS Native testsuite in my jboss-eap-5.1 installation. Is this something that is part of the EAP binary install? Is this situation something that users can encounter on their machine?	1	1	https://issues.jboss.org/browse/JBPAPP-4359
10089	JBPAPP-4359	Do we need to tell customers about this in the release note then? It sounds like an internal QE tweak...	1	1	https://issues.jboss.org/browse/JBPAPP-4359
10090	JBPAPP-4356	Jesper, thanks for the email on this - can you create a subtask docs JIRA with the exact information that needs to change in the Admin and Config Guide? Thanks!	1	0	https://issues.jboss.org/browse/JBPAPP-4356
10091	JBPAPP-4353	I looked at this... I don't normally review documentation, but nothing lept out at me. Is this supposed to be assigned to me?	0	0	https://issues.jboss.org/browse/JBPAPP-4353
10092	JBPAPP-4353	Chris, you're listed as rel-eng lead, so I assigned it to you to review the Installation section - I assumed you would know where the files would end up and whether the download instructions were correct. If not, who is the best person to review this? I understand we do not do RPMs for WFK 1.1, so that whole section should be removed?	0	0	https://issues.jboss.org/browse/JBPAPP-4353
10093	JBPAPP-4353	Chapter 3.2: I would hope that : "None of the information in this section seemed relevant given that we no longer ship Struts, Spring or GWT. David, Ondrej, please provide additional information if required, or assign to someone who can." would not be included in the final, and that these sections have been updated? Chapter 6 - Is there a link to this doc download - jboss-wfk-docs-1.1.0.zip? The other link only shows a page with a copy of the release notes - http://www.redhat.com/docs/en-US/JBoss_Web_Framework_Kit/ Chapter 7 & 8 "YOUR INPUT HERE" - assuming this will be updated base on actual issues fixed, open.... If this is what is needed - I believe David would be able to compile this for you.	0	0	https://issues.jboss.org/browse/JBPAPP-4353
10094	JBPAPP-4353	Maybe boiler-plate we need, but don't see that the XFS trademark is ever used in the document outside the Legal Notices. Should it be removed?	0	0	https://issues.jboss.org/browse/JBPAPP-4353
10095	JBPAPP-4339	Hi Alexey, this a CP release (despite the number change 5.0.1 -> 5.1.0). Is t ok to upgrade metadata from 1.0.2 to 1.0.5 in a CP release? Fully backwards compatible, no side-effects? Anee-Louise (PM) and Laura (Docs) will need a list of issues fixed between 1.0.2 and 1.0.5. Can you add that as a comment to this JIRA please? (assuming the upgrade is ok for a CP) Many thanks, Fernando	1	1	https://issues.jboss.org/browse/JBPAPP-4339
10096	JBPAPP-4339	Many thanks Ales. Carlo, can you comment about JBMETA-277 and possible effects to EAP 5.0.1 customers upgrading to EAP 5.1.0 please?	0	0	https://issues.jboss.org/browse/JBPAPP-4339
10097	JBPAPP-4335	Marius, you seem to have removed all my formatting and language changes in making your modifications... can you revert that and add only the additional lines please? Or just let me know what material has to be added. Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-4335
10098	JBPAPP-4292	If we are changing the version of rhq (and EMS) then we need a 1.4.3.SP1 tag for the embjopr right? Or will you simply re-tag 1.4.3.GA? It is OK to re-tag, we haven't given that to anyone else yet....	1	1	https://issues.jboss.org/browse/JBPAPP-4292
10099	JBPAPP-4292	Hi Simeon, do you have an update on the status of this?	1	1	https://issues.jboss.org/browse/JBPAPP-4292
10100	JBPAPP-4287	Please verify in EAP 5.1 documentation on stage.	1	1	https://issues.jboss.org/browse/JBPAPP-4287
10101	JBPAPP-4287	Guys, I had spent quite some time trying to look up the docs in question and I failed :-D There is just a general installation note in the corp preview and on the knowledge base web page... Can you point me to the archive?	0	0	https://issues.jboss.org/browse/JBPAPP-4287
10102	JBPAPP-4280	Could you verify this for accuracy, please, Strong? :)	1	1	https://issues.jboss.org/browse/JBPAPP-4280
10103	JBPAPP-4280	Hi Juraci, I think I have fixed the tabbing/rendering issue. Could you please verify here on stage? http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html/Hibernate_Search_Reference_Guide/search-mapping.html#search-mapping-entity	1	1	https://issues.jboss.org/browse/JBPAPP-4280
10104	JBPAPP-4278	Ales, transferring this content request from JBPAPP-2635. Can you write some example code for this section? Ideally your material would be inserted immediately following "Another way to limit scanning is to provide the jboss-classloading.xml file." http://java.dzone.com/articles/jboss-microcontainer-classloading goes into more detail with the jboss-classloading.xml files, but I'm no coder. :)	0	0	https://issues.jboss.org/browse/JBPAPP-4278

#	Requirement	Questions	Answered	Successfully Answered	URL
10105	JBPAPP-4278	Below is an email thread in which a JBoss Subject Matter Expert confirms that in fact this information should *NOT* be moved to the Performance Tuning Guide when the "Performance Tuning" chapter of the Admin & Config Guide is removed: ----- Original Message ----- > From: "Russell Dickenson" <rdickens@redhat.com> > To: "Andrig Miller" <anmiller@redhat.com> > Sent: Sunday, January 8, 2012 6:12:58 PM > Subject: Question about JIRA issue for EAP Performance Tuning Guide > > Andy, > > I need your help with a JIRA issue which was raised a long time ago: > https://issues.jboss.org/browse/JBPAPP-4278 > > The gist of the JIRA issue is that someone noted that with the > creation of the PTG the "Performance Tuning" section of the EAP > Admin & Config Guide could be removed. There was a further note > regarding this chapter's removal though: "If this chapter is > removed, a reference to the Classloading Layer chapter of the > Microcontainer User Guide should be added in the Performance Tuning > Guide's equivalent of the VFS Cache Tuning section." I have looked > over the PTG and can't find any reference to VFS, tuning or > otherwise. > > Assuming I'm correct and VFS is not mentioned in the PTG, should this > be added specifically so that this reference can be added, or have > circumstances changed perhaps and there's no longer a need to add > this reference? > I personally don't think that this should be added to the performance tuning guide. The actually changes only affect deployment performance, not actual runtime performance, and while you may increase the deployment time, the extra memory that the VFS will use, could actually hurt runtime throughput. So, I would not add this information to the performance tuning guide, since its focused on runtime throughput of the customers application, and the VFS tuning has nothing to do with that, and could actually hurt throughput. Also, there are no metrics where they show how much it actually improves deployment time. Most customers aren't going to care if there deployments take 6 seconds versus 5.5, or something like that. Andy > > -- > Regards, > > Russell Dickenson > Content Author, Engineering and Operations > > DID: +61 (0)7 3514 8159 > Email: rdickens@redhat.com > > Red Hat, Asia-Pacific Pty Ltd > Level 1, 193 North Quay > Brisbane 4000 > Office: +61 7 3514 8100 > Fax: +61 7 3514 8199 > Website: www.redhat.com <http://www.redhat.com/>	1	1	https://issues.jboss.org/browse/JBPAPP-4278
10106	JBPAPP-4251	Martha, is this the same as the JBoss XB component we report in the release notes? That version is currently listed as 2.0.1.GA... should that change for the Beta release?	0	0	https://issues.jboss.org/browse/JBPAPP-4251
10107	JBPAPP-4247	Rosta, can this jira be closed?	1	1	https://issues.jboss.org/browse/JBPAPP-4247
10108	JBPAPP-4229	On Thu, 2010-01-14 at 14:59 -0300, Hardy Ferentschik wrote: > Speaking of plugins, what does the maven-test-ext-plugin do for us and why > is it needed in > the annotations module? The build seems to work fine even w/o. Its there to mimic behavior JVZ promised me he would add to Maven but never has (in 3 years). Basically it allows you to *dynamically* alter the test classpath. We don't use it per se because its not really intended for us. Its intended for people running the hibernate testsuite(s) against "other databases". The build works ok for us without it there, but it needs to be there. -- Steve Ebersole <st...@hibernate.org> Hibernate.org	1	1	https://issues.jboss.org/browse/JBPAPP-4229
10109	JBPAPP-4229	Strong, on Hudson, for testing, it is OK to grab the plugins (all of maven) from central, so I vote we close this as WONTFIX. In Brew you can just patch out the reference, depmap it to nothing or something. Maybe David can suggest how to handle that but a small patch with some '-' lines will do it. Can we close this as WONTFIX?	1	1	https://issues.jboss.org/browse/JBPAPP-4229
10110	JBPAPP-4229	so David, would you please import it? I tried to work on jpackage, but don't know how to .. (can't remember the details, I remember i send you and Fernando a mail to ask this)	1	1	https://issues.jboss.org/browse/JBPAPP-4229
10111	JBPAPP-4229	David, would you please build this on jdk5? you only build this on jdk6, but hibernate can not be built on jdk6.. we need jdk5 thanks	1	1	https://issues.jboss.org/browse/JBPAPP-4229
10112	JBPAPP-4229	Strong, did you say that you no longer need a JDK5 build? If so, then this issue can be marked as Done.	1	1	https://issues.jboss.org/browse/JBPAPP-4229
10113	JBPAPP-4227	Fernando, can you take this from here?	1	1	https://issues.jboss.org/browse/JBPAPP-4227
10114	JBPAPP-4226	Manik, is this being included? If so, some mention should be made in the release notes about this Tech Preview. Can you let me know?	1	0	https://issues.jboss.org/browse/JBPAPP-4226
10115	JBPAPP-4226	Adding the release notes flag. Manik, Rich, someone will need to give me information on how to connect Infinispan to EAP - what do customers have to do to get Infinispan, and get it working with EAP?	1	0	https://issues.jboss.org/browse/JBPAPP-4226
10116	JBPAPP-4226	This is a significant undertaking. We will have to dedicate a person to it full-time. They will need to study Infinispan and then do the productization of the docs. What resources are QE dedicating to this?	0	0	https://issues.jboss.org/browse/JBPAPP-4226
10117	JBPAPP-4217	What is the target version? Please update: https://docspace.corp.redhat.com/docs/DOC-37998	0	0	https://issues.jboss.org/browse/JBPAPP-4217
10118	JBPAPP-4192	I cannot see the fixes on stage. Could you update it on stage please?	1	1	https://issues.jboss.org/browse/JBPAPP-4192
10119	JBPAPP-4191	Ales, do you think we should have an admin guide for the MC as part of EAP doc bundle? Do you see the administrator configuring the microcontainer in any ways?	0	0	https://issues.jboss.org/browse/JBPAPP-4191
10120	JBPAPP-4191	What is usually in these admin guides?	0	0	https://issues.jboss.org/browse/JBPAPP-4191
10121	JBPAPP-4189	This issue is a clone of JBPAPP-3591 to specifically track EAP5. The last comment on the source JIRA was: Fernando Nasser - 27/Apr/10 03:23 PM Can we move this to 5.0.2 for tracking purposes now?	1	1	https://issues.jboss.org/browse/JBPAPP-4189
10122	JBPAPP-4157	Hi Ruchi, I'm not sure I understand you fully. For example, JBossQueue is not a final class, you can extend it whether your sub class is abstract or not. Also JBossQueue implements javax.jmx.Queue, so it can be casted to javax.jmx.Queue. Am I mising anything there? Howard	1	0	https://issues.jboss.org/browse/JBPAPP-4157
10123	JBPAPP-4155	Mike, this chapter is not included in the community version of the AS5 documentation either. Are you sure it is still required? If so, please let us know where in the current Admin and Config Guide it should go. Thankyou. :)	0	0	https://issues.jboss.org/browse/JBPAPP-4155
10124	JBPAPP-4155	Mike, I've attached the naming chapter file from 4.2. Can you please review the contents for use with EAP 5.1 and let me know what needs to change before I add it to the Admin and Config Guide for 5.1? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-4155
10125	JBPAPP-4155	Hi Mike, I have added the Naming chapter into the Admin and Config guide in chapter 17, after Clustered JNDI Services, but I wasn't quite sure where to slot it in. Would you be able to look over it and see if there's a more appropriate place, please? Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-4155
10126	JBPAPP-4143	Ales, which version is expected? Please update https://docspace.corp.redhat.com/docs/DOC-37998	0	0	https://issues.jboss.org/browse/JBPAPP-4143
10127	JBPAPP-4143	This used to be top level issue for all potential MC upgrades -- VFS, Reflect, MDR, Man, Kernel, CL, Deployers, ... Hence there is no version. Or is there no need to top level MC component upgrade issue?	1	1	https://issues.jboss.org/browse/JBPAPP-4143
10128	JBPAPP-4143	Hi Ales, I don't see a subtask for JBoss Man but in DOC-37998 it is listed as having to be upgraded to 2.1.1.SP2 Is it no longer to be upgraded and can stay at 2.1.1.GA? Thanks, Fernando	1	1	https://issues.jboss.org/browse/JBPAPP-4143
10129	JBPAPP-4140	If the issue is resolved for 4.3.CP09, then we should clone off this issue and make one for EAP 4.2 so this issue can be closed off against 4.3.CP10. Or could we just remove the 4.2 Affects Version and close this ticket off since 4.2 is no longer under active development?	1	0	https://issues.jboss.org/browse/JBPAPP-4140

#	Requirement	Questions	Answered	Successfully Answered	URL
10130	JBPAPP-4139	Hi Rajesh, Could you please just verify that this section has been removed and is now correct? Thanks. http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html/Seam_Reference_Guide/mail.html	1	1	https://issues.jboss.org/browse/JBPAPP-4139
10131	JBPAPP-4131	Please verify in EAP 5.1 documentation on stage.	0	0	https://issues.jboss.org/browse/JBPAPP-4131
10132	JBPAPP-4111	removed jboss-cache 1.4.1 and pojo cache	1	1	https://issues.jboss.org/browse/JBPAPP-4111
10133	JBPAPP-4111	I have made the changes as specified in the FishEye. Rajesh, could you please verify this for 5.1? Available here: http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html-single/Seam_Reference_Guide/index.html#cache	0	0	https://issues.jboss.org/browse/JBPAPP-4111
10134	JBPAPP-4111	Hi. Just a quick comment before goes to archiving... What I encountered when trying the supported version of 3.2.x was that whenever you used the <s:cache> tag, an exception occurred. So we had to revert back to using the old 1.4.1 version of jboss cache. Is this now fixed?	1	1	https://issues.jboss.org/browse/JBPAPP-4111
10135	JBPAPP-4111	I am not aware of any exception in usage of s:cache component. Could you specify more details about it?	1	1	https://issues.jboss.org/browse/JBPAPP-4111
10136	JBPAPP-4089	Marek, where would I find this information? Would you suggest this forms part of the Seam Reference Guide, or part of the release notes?	0	0	https://issues.jboss.org/browse/JBPAPP-4089
10137	JBPAPP-4089	Hi Rajesh, I have made the changes as available in the fisheye. Could you please verify these for 5.1? Available here: http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html/Seam_Reference_Guide/index.html	1	1	https://issues.jboss.org/browse/JBPAPP-4089
10138	JBPAPP-4058	Docmented as a known issue in the EAP 4.3.8 release notes:<classname>CacheEntry</classname> resources are not cleaned when <classname>ApplicationContext</classname>s are shut down, causing a memory leak. Is there a workaround available for this?	1	0	https://issues.jboss.org/browse/JBPAPP-4058
10139	JBPAPP-4053	Tyronne, this has been set to Coding in Progress since April. Can you give an update on status and need for this? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-4053
10140	JBPAPP-4052	I merged the Transactions JTA, JTS, and XTS Programmers Guides into a single Transactions Development Guide for EAP 5.1. Here are the final EAP Transactions deliverables for EAP 5.1: 1. Transactions Administration Guide 2. Transactions Development Quick Start Guide 3. Transactions Development Guide	0	0	https://issues.jboss.org/browse/JBPAPP-4052
10141	JBPAPP-3961	Brian, since Infinispan is only a tech preview this time, can we really say Cache is deprecated? Doesn't deprecated imply 'no longer fully supported'?	1	1	https://issues.jboss.org/browse/JBPAPP-3961
10142	JBPAPP-3961	Attention: Jared Since this is marked with a Fix Version/s of TBD EAP 5, should it be targeted for EAP 5.1.2's Release Notes? It's a Documentation-specific issue.	0	0	https://issues.jboss.org/browse/JBPAPP-3961
10143	JBPAPP-3961	Eva - My memory of what you asked in our team meeting is not very clear. IIRC you asked if we were to *definitely* remove from documentation features which had been deprecated but were still supported. I believe we should leave documentation of supported features in place, removing only documentation of features which are no longer supported. Does this give you enough guidance? I'd honestly be ver surprised if any deprecated feature has since been declared unsupported, so there should be very little to do with this issue. I would expect a Release Notes entry to be written but expect that that's all.	0	0	https://issues.jboss.org/browse/JBPAPP-3961
10144	JBPAPP-3950	Yes some customers needlessly put jars in that shouldn't be added, however the bigger case is when a customer deploys a 3rd party product, I've seen several that include api classes for jms and such which should be ignored. When the customer has a 3rd party application, they have to go back to the vendor and get them to make a version with the classes removed so that they can maintain their support agreement. By delegation mode, do you mean to have it check the parent first? In that case it would not allow the customer to override other valid classes.	0	0	https://issues.jboss.org/browse/JBPAPP-3950
10145	JBPAPP-3950	Does this apply to EAP 4.x or 5.x or both? If both, this issue needs to be cloned and the one for 4.x put in TBD EAP 4 and the other one in TBD EAP 5. Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-3950
10146	JBPAPP-3935	Did you pass it the build.properties and with an empty version.tag ? This branch is set for a production build so the behavior without a version.properties is flaky.	0	0	https://issues.jboss.org/browse/JBPAPP-3935
10147	JBPAPP-3879	Considering that EAP 5.0.1 wasn't released yet, is it OK to run the regression/compatibility tests on 5.0.1 CR1?	1	1	https://issues.jboss.org/browse/JBPAPP-3879
10148	JBPAPP-3845	Can I get a description of this for the release notes, please?	1	1	https://issues.jboss.org/browse/JBPAPP-3845
10149	JBPAPP-3842	Laura, can you add the suggested changes?	0	0	https://issues.jboss.org/browse/JBPAPP-3842
10150	JBPAPP-3829	Why don't you use the full jbossweb.jar, which is in the repository? http://repository.jboss.org/maven2-brew/jboss/web/jbossweb/2.1.7.GA/jbossweb-2.1.7.GA.jar That is what the dependency is looking for and what the build in Brew uses.	1	1	https://issues.jboss.org/browse/JBPAPP-3829
10151	JBPAPP-3829	Checked the latest build of tomcat6 and jbossweb in Brew, can NOT find CometEvent in jbossweb jars, but in tomcat6 catalina.jar, CometEvent is there. CometEvent comes to tomcat since version 6, so jbossweb should use different catalina, a low version and it doesn't have CometEvent yet. So, we have to upload tomcat6 or upgrade the catalina jbossweb uses, Fernando, how do you think of it?	1	1	https://issues.jboss.org/browse/JBPAPP-3829
10152	JBPAPP-3828	Anil, I have looked over the project and I have some questions about the scope. There are multiple deliverables in the community repo: 1. User Guide - This appears to be a guided tutorial for creating a specific type of application using Microcontainer and POJOs 2. Getting Started Guide - This appears to be an overview for installing and configuring the Microcontainer 3. Reference Guide What level of detail are we looking at? I have been working from the User Guide, but I think the Getting Started Guide may be more the direction you were headed.	0	0	https://issues.jboss.org/browse/JBPAPP-3828
10153	JBPAPP-3817	Hi Fernando, In Brew, mojo-maven2-plugins-jaxb2 has a maven-jaxb2-plugin subpackage, but that maven-jaxb2-plugin has different version and groupId with the needed one, I'd like to add repolib to mojo-maven2-plugins and upload to maven2-brew, then liweinan can create a patch for the pom which has this dependency, this may resolve this problem. Do you have any advice?	1	1	https://issues.jboss.org/browse/JBPAPP-3817
10154	JBPAPP-3817	First of all, which version and groupId has the mojo-maven2-plugins-jaxb2 subpackage? Is this really the same plugin or some other plugin that just happens to be also related to JAXB2? How did we manage to build the resteasy RPM in Brew? Is there a depmap? I am not even sure if depmaps apply to plugins... CHeck the resteasy Brew root.log and brew.log and find out where the maven-jaxb2-plugin came from	1	1	https://issues.jboss.org/browse/JBPAPP-3817
10155	JBPAPP-3798	Hi Anne-Louise Should this issue be placed into a developer's queue, rather than being left unassigned? Cheers J	0	0	https://issues.jboss.org/browse/JBPAPP-3798
10156	JBPAPP-3797	This seems a reasonable RFE, other customers would likely benefit as well and is in-line with our focus on operational eas of use. Needs to be assessed by eng. to ensure it can be done without changing existing behaviour, eg. what happens if the run-as user doesn't have permission to create a pid file ?	0	0	https://issues.jboss.org/browse/JBPAPP-3797
10157	JBPAPP-3797	Lami, can you please be more specific. Which one of the http://anonsvn.jboss.org/repos/jbossas/branches/ this needs to go in	1	1	https://issues.jboss.org/browse/JBPAPP-3797
10158	JBPAPP-3764	If the test application still reports uncollected garbage, why is this issue resolved?	0	0	https://issues.jboss.org/browse/JBPAPP-3764
10159	JBPAPP-3754	I've updated the RHEL builds to 1.2.30: mod_jk-1.2.30-2.ep5.el4 mod_jk-1.2.30-1.1.jdk6.ep5.el5 Mladen, Can you please update the hudson build for the other platforms? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-3754

#	Requirement	Questions	Answered	Successfully Answered	URL
10160	JBPAPP-3732	Rich, as far as I can see, Oracle 9i isn't listed as a supported configuration for EAP 5.0.1, just 4.2 and 4.3: http://www.jboss.com/products/platforms/application/supportedconfigurations/ Can you advise on this? We don't have any information on support in the documentation itself, so would it be better to ensure the jboss.com page remained up to date?	1	1	https://issues.jboss.org/browse/JBPAPP-3732
10161	JBPAPP-3731	Rajesh, I've documented this as part of JBPAPP-3732 in the "Database Certifications" section of the release notes. Is the following still valid if this is TBD? JBoss Enterprise Application Platform 5.0.1 is now certified for the Oracle 11g R2 and Oracle 11g RAC databases when used with Oracle JDBC Driver version 11.2.0.1.0. It is also certified for Sybase ASE 15.0.3 when used with JConnect driver version 6.0.5.	1	1	https://issues.jboss.org/browse/JBPAPP-3731
10162	JBPAPP-3731	Rajesh, changed to "Oracle 11g R2 RAC". This is currently documented for 4.2 as well, is this correct?	1	1	https://issues.jboss.org/browse/JBPAPP-3731
10163	JBPAPP-3719	Here is where it appears in the community docs: http://labs.jboss.com/file-access/default/members/jbossmessaging/freezone/docs/userguide-1.4.0.SP3/html_single/index.html#inst.remoteclient Is this book in the EAP set? If so, is this section in the book? According to Tyrone the information cannot simply be cut-and-pasted. He knows how it needs to be done in EAP.	0	0	https://issues.jboss.org/browse/JBPAPP-3719
10164	JBPAPP-3719	Hi Tyronne, Can you provide the relevant information? Thanks. Ben	1	1	https://issues.jboss.org/browse/JBPAPP-3719
10165	JBPAPP-3719	Hi, i tried to find a description in JBM User Guide for EAP 4.3 which would give a list of needed jars for a JMS client to run, but without success. Can you give me a hint where the description was added? I expected something as in http://docs.jboss.org/jbossmessaging/docs/userguide-1.4.0.SP3/html_single/index.html#inst.remoteclient in section 4.4. Thank you	1	1	https://issues.jboss.org/browse/JBPAPP-3719
10166	JBPAPP-3719	Laura, could you help me to find the changes? Thank you	1	1	https://issues.jboss.org/browse/JBPAPP-3719
10167	JBPAPP-3719	Hi Howard I need a little help with this question raised by Tyronne. Would you be able to check it out for me, Howard? Thanks for any help you can give me here. Cheers Jared	1	1	https://issues.jboss.org/browse/JBPAPP-3719
10168	JBPAPP-3703	Bill, can you comment on the possible backward incompatibilities that may be encountered when upgrading RESTEasy from 1.1 to 1.2.1 ? This is a CP type of release, where we no longer support the previous version, so we must ensure absolute backwards compatibility.	0	0	https://issues.jboss.org/browse/JBPAPP-3703
10169	JBPAPP-3703	Will you upload the build result into http://repository.jboss.org/maven2-brew/org/jboss/resteasy/ ?	1	0	https://issues.jboss.org/browse/JBPAPP-3703
10170	JBPAPP-3677	I have a customer in support case [1]. Seems he is experiencing this issue. However the error message he is getting is slightly different. EAP 5.0.1 Oracle Driver Implementation-Version: 11.2.0.1.0 Oracle Server 11.2.0.1 They are requesting an patch for EAP 5.0.1. Any chance of a patch? If so, a timeline would be much appreciated. [1] https://c.na7.visual.force.com/apex/Case_View?id=500A0000003iovWIAQ [2] org.jboss.resource.JBossResourceException: Could not create connection; - nested throwable: (java.sql.SQLException: ORA-01017: invalid username/password; login denied)	0	0	https://issues.jboss.org/browse/JBPAPP-3677
10171	JBPAPP-3634	why we need this? if we do it, so, which one should users use in which case? what's the difference between http://hibernate.sourceforge.net/hibernate-mapping-3.0.dtd and "http://hibernate.sourceforge.net/hibernate-mapping-3.1.dtd" ?? take a look of http://fisheye.jboss.org/browse/Hibernate/core/trunk/core/src/main/resources/org/hibernate/hibernate-mapping-3.0.dtd?r1=15365&r2=17821&u=3&ignore=&k= this link shows the difference between these two dtd files, the one in hibernate-core jar in eap5 and another in source forge site as you can see the differences came with two project jira http://opensource.atlassian.com/projects/hibernate/browse/HHH-3275 http://opensource.atlassian.com/projects/hibernate/browse/HHH-4440 they are new features in hibernate trunk, and they are not available in eap5 for now we have carefully kept the compatibility between every change of that dtd file or if you want to use "hibernate-mapping-3.1.dtd", whatever it is, you can put it in one of your package, e.g. org.aaa.bbb and declare it in hbm files like this: <!DOCTYPE hibernate-mapping PUBLIC "-//Hibernate/Hibernate Mapping DTD 3.0//EN" "classpath://aaa/bbb/hibernate-mapping-3.1.dtd"> I leave this open in case others want to discuss more	1	0	https://issues.jboss.org/browse/JBPAPP-3634
10172	JBPAPP-3634	Strong, does this need to be documented as a known issue?	0	0	https://issues.jboss.org/browse/JBPAPP-3634
10173	JBPAPP-3620	Will it be before Aug 18? If not, let me know and we can reassign this to TBD EAP 5.	1	1	https://issues.jboss.org/browse/JBPAPP-3620
10174	JBPAPP-3620	Hi Marek, are you able to give us any information on this? E.g what changes actually need to be incorporated? Isaac, do you know who in JBDS we could contact for information on this? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-3620
10175	JBPAPP-3620	Isaac, was the intent behind these two JIRAs similar?	0	0	https://issues.jboss.org/browse/JBPAPP-3620
10176	JBPAPP-3618	The EWS release will include Tomcat Native 1.1.19. Do we need to be aligned on that version?	0	0	https://issues.jboss.org/browse/JBPAPP-3618
10177	JBPAPP-3591	Does this need to be documented for the EAP 4.2.9 and 4.3.8 release notes?	1	1	https://issues.jboss.org/browse/JBPAPP-3591
10178	JBPAPP-3591	Laura, the information about this certification is correctly there, but we need to mention the two issues linked to this one. I changed both to include the "Affects Version" for 4.x series, as they were missing (and weren't included in Release notes probably because of that). So, could you please include them in the Release Notes and assign this issue back to me for verification? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-3591
10179	JBPAPP-3591	Can we move this to 5.0.2 for tracking purposes now?	1	1	https://issues.jboss.org/browse/JBPAPP-3591
10180	JBPAPP-3529	Hi Remy, can we avoid making this copy? They may get out of sync again.	0	0	https://issues.jboss.org/browse/JBPAPP-3529
10181	JBPAPP-3529	Is fixing this really so important ? Not attempting to configure checkInterval would make it default to 0, which would disable background compilation if development is set to false. Moving to the next CR for now.	0	0	https://issues.jboss.org/browse/JBPAPP-3529
10182	JBPAPP-3499	Manik, can you let me know any details? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-3499
10183	JBPAPP-3498	Manik, I need to know if there are any changes required to this document for inclusion in the Enterprise Web Platform. EWP is not meant to expose any remote connections except the web server. Can you tell me which sections need to be removed so that this is the case? Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-3498
10184	JBPAPP-3498	"Can I run JBoss Cache outside of the JBoss Enterprise Web Platform?" - surely we don't want this? "Can JBoss Cache run as a second level cache inside Hibernate?" - does EWP support hibernate in the first place? If not all of the HIbernate related FAQs can be removed. "Which transaction managers are supported by JBoss Cache?" - does EWP support JTA? Again, these could be removed if unnecessary... HTH Manik	1	1	https://issues.jboss.org/browse/JBPAPP-3498
10185	JBPAPP-3498	Will remove the first. Brian, can you answer the support questions so that I know what to remove?	1	1	https://issues.jboss.org/browse/JBPAPP-3498
10186	JBPAPP-3498	"Can JBoss Cache run as a second level cache inside Hibernate?" -- yes, clustered 2nd level caching is supported "Which transaction managers are supported by JBoss Cache?" -- yes, JTA is supported, (although JTS is not)	1	1	https://issues.jboss.org/browse/JBPAPP-3498

#	Requirement	Questions	Answered	Successfully Answered	URL
10187	JBPAPP-3498	I've removed "Can I run JBoss Cache outside of the JBoss Enterprise Web Platform?" Brian, thanks for the quick response. The original content is: Which transaction managers are supported by JBoss Cache? JBoss Cache supports any TransactionManager that is JTA compliant, such as JBoss Transactions. While JBoss Cache does ships with a dummy transaction manager (org.jboss.cache.transaction.DummyTransactionManager), we do not recommend using this for production. It is not thread safe, and is intended for internal testing only. I have added the following: JTS is not supported Java Transaction API (JTA) is supported for use with Enterprise Web Platform. However, Java Transaction Service (JTS) is not supported. Keep this in mind when selecting a transaction manager. Please confirm that this is appropriate.	1	1	https://issues.jboss.org/browse/JBPAPP-3498
10188	JBPAPP-3489	after a manual searching and google site search, javax.transaction:jta:jar:1.0.1B could not be found in: http://repository.jboss.com/maven2-brew/ and http://repository.jboss.com/maven2/ could fnasser kindly indicate me where can I find it or is there any better searching tools I can use? Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-3489
10189	JBPAPP-3485	so, can i resolve this as won't fix?	1	1	https://issues.jboss.org/browse/JBPAPP-3485
10190	JBPAPP-3484	You're right, I do not like this. Though it has nothing to do with the DTD change. What is the use case for this?	1	0	https://issues.jboss.org/browse/JBPAPP-3484
10191	JBPAPP-3484	And have the measured the effect of causing the size of their query plan caches to explode?	1	0	https://issues.jboss.org/browse/JBPAPP-3484
10192	JBPAPP-3484	@sebastian I am more concerned with the why, not so much if it is possible (its programming, essentially anything is possible). I hear you saying that the db2 optimizer somehow treats "... where x = ?" differently from "... where x = 'SOME_VALUE'". Is that really what you are saying? What specifically is different? It's able to use an index in one case but not the other? Generally speaking a query optimizer is not about [optimizing results], so I am not really sure what "the optimizer is not able to properly optimize the result for next queries" means. Do you mean that you are asking it hold results for each query and here the "optimizer" is not able to distinguish the "cached results" based on bound parameter values? I am totally guessing here as I am not a db2 user (I have been lucky enough to avoid db2 my whole career) and the database with which I am familiar have a much different usage of the phrase "optimizer" in regards to queries.	1	0	https://issues.jboss.org/browse/JBPAPP-3484
10193	JBPAPP-3484	Hi Steve, no DB2 in your whole career? Really nice! :) Here's problem with the optimizer: Let's say we have a database table with 500.000 records. 200 records have an index 'a', the rest has an index 'b'. The table is partitioned with this index -> 2 Partitions (a = 200 recs, b=499800 recs). Now you query with "where index='?'", ?=a. The optimizer uses a table scan cause there are only 200 records. The acces plan is cached. Now you query with "where index='?'". ?=b The optimizer should now create a new access plan cause it's better to use the index tree if there are so much records. But there is a cached plan present for query ".. index=?" and db2 uses the cached one - and a full table scan is performed. Very, very slow :(If the index-field isn't a parameter-marker db2 uses a new plan every time this query is performed. There's no query plan cache used for db2, thats a bit bad. But the performance improves and costs went down.	1	1	https://issues.jboss.org/browse/JBPAPP-3484
10194	JBPAPP-3484	I am not sure how this a DB2 or any database specific issue. It is known that with any database there are 'cases' where using a literal value is better than a parameter marker for e.g. ... 'by substituting the actual predicate with a bind variable, you're removing the ability for the optimiser to compare your value with the data distribution in the column, which might make it opt for a full table scan or an index when this isn't appropriate' So the use-case above is one that fits into this category. The simplest solution for these cases is for the application to use a literal value in the query (HQL, Criteria whatever) instead of parameter marker. What is preventing the application from doing this... SQL injection ? I cant see how the index value choice is coming from user input data but even if it is, the app can duly sanitize it. The proposed solution seems a classic example of over engineering to me	1	1	https://issues.jboss.org/browse/JBPAPP-3484
10195	JBPAPP-3484	@Sebastian. As mentioned above, i don't think it matters whether this is Oracle or DB2 and also what's preventing the application from using a literal in the HQL or Criteria itself ?	1	1	https://issues.jboss.org/browse/JBPAPP-3484
10196	JBPAPP-3484	Mario, lets take a step back... Would you agree or disagree that this is a database issue? To me it boogles my mind quite honestly that the database does not account for these bind parameters (their values) when determining which execution plan to use, especially when linked to indexed columns, let alone columns used in partitioning.	1	0	https://issues.jboss.org/browse/JBPAPP-3484
10197	JBPAPP-3483	Documented as a known issue in the EAP 5.0.1 release notes.	0	0	https://issues.jboss.org/browse/JBPAPP-3483
10198	JBPAPP-3477	Howard, can you give me details on how to migrate from JBoss Messaging 1.4.6 to HornetQ and vice versa? This is a requirement on the EAP 5.1 ERD.	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10199	JBPAPP-3477	Laura, I'll try on this. By the way, I don't think migrate from HornetQ to JBM make any sense. Can you explain why EAP 5.1 need this? Howard	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10200	JBPAPP-3477	The guide looks very comprehensive, thanks Howard. Is the source code in a repo somewhere?	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10201	JBPAPP-3477	Howard, the HornetQ User Guide and HornetQ Quickstart Guides have been combined for EAP 5.1. Can you assign this to someone for technical review for EAP 5.1? I need to know what is not supported or does not apply to HornetQ in EAP 5.1. http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html-single/HornetQ_User_Guide/index.html	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10202	JBPAPP-3477	Hi Tim, Can you help review the HornetQ's doc for EAP 5.1? http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html-single/HornetQ_User_Guide/index.html Thanks Howard	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10203	JBPAPP-3477	I am in the process of adding the migration PDF to HornetQ, however, I feel the following points could use clarification. Tim, if you could provide details that would be great. 1) Sections 4 and 10 in the PDF no longer seem entirely relevant to EAP, given that migration entails uninstalling and reinstalling the entire platform. Are these still valid? http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html-single/Installation_Guide/index.html#alt-tech <-- The Alternative Technologies section details our current process. Are you able to help me synthesise these into one process? Where would those steps fit into this migration document? 2) "As HornetQ has a different transport layer and cluster design, you need to map the parameters in [Remoting and JGroups] configuration files to their HornetQ equivalents" - what are the equivalents? Can we list them here, or is there another section in the HQ user guide that covers this that we can link to? 3) The JMS Bridge parameters table is non-specific about the name of parameters after TargetDestinationFactory - is it possible to get the specific parameter name for the rest of the attributes, or is "Source user name parameter" as close as we can get? The JMS Bridges chapter does not specify these either and I think it would help to be specific. 4) Section 7 says: "You need to consult corresponding documents for a complete description of the parameters in Remoting and JGroups in order to get a proper settings for HornetQ." EAP does not ship with this documentation. What do HornetQ users need to know? 5) Section 8 says "After you have migrated all the JMS destinations on to HornetQ, you can " and does not finish. I assume the end of that sentence is " migrate existing messages from your JBoss Messaging server to the HornetQ server"?	1	1	https://issues.jboss.org/browse/JBPAPP-3477

#	Requirement	Questions	Answered	Successfully Answered	URL
10204	JBPAPP-3477	Laura - can you tell me which svn revision from the HornetQ repository you used when taking the HornetQ documentation to convert to EAP documentation? We will need to track this so we can do svn diffs to see any changes in the HornetQ docs that need to be also applied to the EAP docs.	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10205	JBPAPP-3477	Hi Laura, here is my answer. However, I think Tim's opinion is still important to have. :) 1) Sections 4 and 10 in the PDF no longer seem entirely relevant to EAP, given that migration entails uninstalling and reinstalling the entire platform. Are these still valid? http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html-single/Installation_Guide/index.html#alt-tech <-- The Alternative Technologies section details our current process. Are you able to help me synthesise these into one process? Where would those steps fit into this migration document? --- I think you can omit those sections as you think fit for the EAP. 2) "As HornetQ has a different transport layer and cluster design, you need to map the parameters in [Remoting and JGroups] configuration files to their HornetQ equivalents" - what are the equivalents? Can we list them here, or is there another section in the HQ user guide that covers this that we can link to? --- I think the parameters of remoting and jgroups (in hornetq it has it own parameters) are wildly different because their architectures (jbm vs hq) is different. It's hard to find a one-to-one mapping between the two system. So probably we can point to the document of remoting and jgroups for references. 3) The JMS Bridge parameters table is non-specific about the name of parameters after TargetDestinationFactory - is it possible to get the specific parameter name for the rest of the attributes, or is "Source user name parameter" as close as we can get? The JMS Bridges chapter does not specify these either and I think it would help to be specific. --- HornetQ's JMS bridge doesn't have parameter names, they are passed in by order. 4) Section 7 says: "You need to consult corresponding documents for a complete description of the parameters in Remoting and JGroups in order to get a proper settings for HornetQ." EAP does not ship with this documentation. What do HornetQ users need to know? --- They need to understand JBM's remoting and jgroups configuration parameters to find a proper hornetQ config. I assume the users should have that knowledge as they are existing JBM users. 5) Section 8 says "After you have migrated all the JMS destinations on to HornetQ, you can " and does not finish. I assume the end of that sentence is " migrate existing messages from your JBoss Messaging server to the HornetQ server"? --- Sorry, my brain must have gone astray. :) Yes you are right, the rest of the sentence is what you have indicated.	0	0	https://issues.jboss.org/browse/JBPAPP-3477
10206	JBPAPP-3477	Tim, I will do the diff you have suggested and merge changes that way. Are you able to do a full tech review before August 13 so that we have the chance to include your feedback before CR3?	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10207	JBPAPP-3477	1) Because that page was not present in the community HornetQ documentation as far as I could tell, and we received no response to our email queries on this subject. Can you provide us with a list of contributors, or direct us to a contributor page? We definitely want to give the bulk of the credit to the people who actually wrote the document! 2) Can you give examples? I can see AS5 references, which I will remove, but I can't see installation references. http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html-single/HornetQ_User_Guide/index.html	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10208	JBPAPP-3477	Chapter 3 - Probably remove the concept of Standalone.... (I hate to let it go though) On running examples, we should write that the JMS examples are starting a small standalone server that is suitable to run the server only, and not supported for production. Also, add a paragraph on examples that it's required to set JBOSS_HOME in order to have the examples running Chapter 6 - 6.1, 6.2, 6.3 needs to go away. LIBAIO (6.4) needs to be contextualized on how the .so is installed with EAP. Which I'm not sure ATM. 6.6, remove the line about run and run.sh. Also, remove log4j.xml and logging.properties paragraphs, since logging is part of jboss logging when hornetq is integrated 6.7 - needs to go 6.8 - needs to go Chapter 7 - 7.3 - needs to go Chapter 40 - I would keep it as is.. say if someone is running EAP on RHEL Mainframe. Chapter 43 - I would also keep it. Someone could do fancy things on the client, while bridging to a main server. (like running a Messaging server on a ATM or small store) Chapter 46.5 (Tuning the VM). That's a generic thing.. Maybe it needs to away, as long as the EAP also have similar suggestions	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10209	JBPAPP-3477	>>Chapter 3 - Probably remove the concept of Standalone.... (I hate to let it go though) >>On running examples, we should write that the JMS examples are starting a small standalone server that is suitable to run the server only, and not supported for production. jmorgan: Added <important> to Section 3.2.1 "The JMS Examples" >>Also, add a paragraph on examples that it's required to set JBOSS_HOME in order to have the examples running jmorgan: Added Section 3.2.2.1 "Option: Add JBOSS_HOME Environment Variable" >>Chapter 6 - >>6.1, 6.2, 6.3 needs to go away. jmorgan: Removed all sections as directed. Remaining comments have been renumbered to reflect the new numbering in the guide. >>LIBAIO (6.4) needs to be contextualized on how the .so is installed with EAP. Which I'm not sure ATM. jmorgan: NEW SECTION NUMBER 6.1 - >>6.6, remove the line about run and run.sh. >>Also, remove log4j.xml and logging.properties paragraphs, since logging is part of jboss logging when hornetq is integrated jmorgan: NEW SECTON NUMBER 6.3 - removed the line about the run scripts and the loging properties. I've also taken a look at what files are added when messaging is converted' to HornetQ, and added those into the section. I've broken the file locations into varlists to chunk them a bit. >>6.7 - needs to go >>6.8 - needs to go jmorgan: NEW SECTION NUMBERS 6.4 AND 6.5 - deleted as prescribed >>Chapter 7 - 7.3 - needs to go jmorgan: deleted 7.3 >>Chapter 40 - I would keep it as is.. say if someone is running EAP on RHEL Mainframe. >>Chapter 43 - I would also keep it. Someone could do fancy things on the client, while bridging to a main server. (like running a Messaging server on a ATM or small store) >Chapter 46.5 (Tuning the VM). That's a generic thing.. Maybe it needs to away, as long as the EAP also have similar suggestions jmorgan: Commented out the VM section.	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10210	JBPAPP-3477	Hi Clebert Regarding the LIBAIO question in your review comments, there seems to be information on that in Chapter 40 "Libaio Native Libraries" of the guide. Does this information nullify the comment you made in the document? Cheers Jared	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10211	JBPAPP-3477	Clebert, how much of Chapter 40. "Libaio Native Libraries" applies to EAP? I assume we do not need to compile, so it seems like only the paragraphs just before 40.1 starts are useful. I would then mention that LibAIO can be obtained from the Native zips, correct? Would the following be more useful as part of the Persistence chapter, rather than in a chapter on its own? (Say, around the "Installing AIO" note? [1]) ---------- HornetQ uses a native library, <filename>libHornetQAIO.so</filename>, in order to implement Linux Asynchronous IO as part of its high-performance journaling. Download the appropriate Native ZIP for your operating system and architecture to obtain the <filename>libHornetQAIO.so</filename> library. More information about the Native ZIPs is available in the <citetitle>Installation Guide</citetitle>. [1] http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5/html-single/HornetQ_User_Guide/index.html#persistence	1	1	https://issues.jboss.org/browse/JBPAPP-3477

#	Requirement	Questions	Answered	Successfully Answered	URL
10212	JBPAPP-3477	Hi Clebert {quote} Say if you are running EAP on Linux in a non common hardware chipset (such as a Mainframe, the experimental super fast PS3 computer.. etc). On that case you will have a proper JDK, a proper RHEL machine and everything else, but you won't have the HornetQ libaio. {quote} I guess the question here is whether it is supported under EAP product support. We don't make a habit of documenting cases in enterprise docs for those environments that are not under the SLA. {quote} So, if someone tell us we won't ever support such practice, then the compilations instructions can go away from the EAP doc. Otherwise it needs to stay. {We would need to offer the library for every supported platform also) {quote} Are you saying that we need to offer native compilation instructions because we don't provide libaio packages for all supported software environments? Cheers Jared	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10213	JBPAPP-3477	Bump. NEEDINFO (email sent 20111021) Hi [~clebert.suconic] Can you help with the last question I asked? This is holding up this resolution for 5.1.2. If you're not the best person to ask, feel free to flick the issue over to someone else. Cheers mate Jared	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10214	JBPAPP-3477	[~atangrin], can you confirm whether the situation Clebert raised in this issue is required from a docs perspective?	1	1	https://issues.jboss.org/browse/JBPAPP-3477
10215	JBPAPP-3364	Will something roll up to EAP 5.0.1?	1	1	https://issues.jboss.org/browse/JBPAPP-3364
10216	JBPAPP-3359	Built in RHEL-4-EP-5 after commons-build-plugin was imported, https://brewweb.devel.redhat.com/buildinfo?buildID=122823 Hi Yiqiao Wang, I saw you were working on RHEL-5-EP-5, so could you please update it then?	0	0	https://issues.jboss.org/browse/JBPAPP-3359
10217	JBPAPP-3344	Is this included in JBoss Web 2.1.7?	0	0	https://issues.jboss.org/browse/JBPAPP-3344
10218	JBPAPP-3319	Hi Peter I've incorporated everything in your comments, however I can't find the place you want this info to go. - Another JBoss specific permissions: -- org.jboss.naming.JndiPermission with actions: bind, rebind, unbind, lookup, list, listBindings, createSubcontext, * (stands for all actions). More info in javadoc of the class. -- org.jboss.security.srp.SRPPermission = A custom permission class for protecting access to sensitive SRP information like the private session key and private key -- org.hibernate.secure.HibernatePermission -- org.jboss.metadata.spi.stack.MetaDataStackPermission -- org.jboss.config.spi.ConfigurationPermission -- org.jboss.kernel.KernelPermission -- org.jboss.kernel.plugins.util.KernelLocatorPermission Could you let me know where this info needs to be inserted. Cheers Jared	1	1	https://issues.jboss.org/browse/JBPAPP-3319
10219	JBPAPP-3287	Shall I include this info in the release notes, section 2.5 "The Native Package"? http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0.0/html-single/Release_Notes/ The EAP 5.0 release notes do not currently list supported operating systems. These operating systems are listed in the EWS release notes - is support the same on EAP? * Red Hat Enterprise Linux 4, latest update (x86, x86_64) * Red Hat Enterprise Linux 5, latest update (x86, x86_64) * Solaris 9 (x86, SPARC32, SPARC64) * Solaris 10 (x86, x86_64, SPARC64) * Windows Server 2003 SP2 (x86, x86_64) * Windows Server 2008 R2 (x86, x86_64)	0	0	https://issues.jboss.org/browse/JBPAPP-3287
10220	JBPAPP-3279	***JBoss Web Console*** The Web Console gets a shout out in the Admin and Config Guide 24.5: "The JBoss Enterprise Application Platform web console http://localhost:8080/web-console/ provides you with monitoring tools starting with the JVM Hardware environment statistics on the default page and access to monitoring tools and snapshots." Configuring security for access is described in Section 7.3 of the Installation Guide: http://www.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0.0/html-single/Installation_Guide/index.html#Post-Installation *** Tomcat status*** Tomcat is the default servlet container used by the JBoss Application Server. Questions to answer: Why would you need to look at this status? What if you replace the default servlet container with another one?	0	0	https://issues.jboss.org/browse/JBPAPP-3279
10221	JBPAPP-3279	This JIRA is asking for clarification and detail on the following links, available on the main page of the EAP: Admin Console JMX Console JBoss Web Console Tomcat Status We need to answer the following questions: * What are the use-cases for each of these links? Why will users need them? * Is the Admin Console implemented as a Seam app? * What does the JMX console do, now that the App Server uses a microcontainer instead of a microkernel? Does it still provide a view into the microcontainer? * The JMX page lists a number of functions, but why would they be used? Examples? * What does the new microcontainer do? * What is the role of the microkernel now, and how does it relate to the microcontainer?	0	0	https://issues.jboss.org/browse/JBPAPP-3279
10222	JBPAPP-3279	Jason, we want to include an overview of the use cases of each of the links visible to users at http://localhost:8080/ when users start up their EAP instance. Could you give me an overview why each of the following would be used, please? - Admin Console - JMX Console - JBoss Web Console - Tomcat status Are you also able to clarify what the shift from microkernel to microcontainer means for the JMX Console (if anything)? If you are unable to answer these questions, could you suggest someone else to ask, please? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-3279
10223	JBPAPP-3279	The admin console is written using seam. The primary interface for management is Intended to be the admin console. However, we do not yet offer management of all services that offer a JMX interface. So the JMX console remains for this purpose. All it really is, is a generic JMX client. So it can be used for other things like operating on user service (provided they expose a JMX interface). The difference between the jmx microkernel and the microcontainer is really just that everything had to be done in jmx with the mc you don't need it unless you want to use it to expose some kind of simple management view	0	0	https://issues.jboss.org/browse/JBPAPP-3279
10224	JBPAPP-3273	The "1.7 Current Limitations" section needs updating: "* Dynamic load balance factor calculation support is currently limited to the JBoss AS integration." This is no longer true - nor is integration with other servers relevant to this document" "* Clustered-mode support is currently limited to the JBoss AS integration" This is not relevant (same reason). "* Checksum verification of proxy advertisements (if advertise is enabled) will never fail." This is the only current limitation - and will be fixed in the next point release (1.0.3). See https://jira.jboss.org/jira/browse/MODCLUSTER-104	1	1	https://issues.jboss.org/browse/JBPAPP-3273
10225	JBPAPP-3273	On http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/0.1/html/mod_cluster/chap-mod_cluster-Overview.html, the platforms supported are: RHEL4,5 x86, x64 Solaris x86, x64 Windows x86, x64 On http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/0.1/html/mod_cluster/sect-mod_cluster-Overview-Requirements.html, httpd is *not* included in the bundle. On http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/0.1/html/mod_cluster/sect-mod_cluster-Quick_Start_Guide-Install_the_httpd_binary.html, since we don't include httpd in our bundle, we don't need the instructions for 'Install the entire httpd'. In the section for 'Install only the modules', remove the reference for mod_proxy.so and mod_proxy_ajp.so . Same for 2.3, please remove the section that is for the extracted httpd. For 2.4 and 2.7, do we need the instructions for JBossweb and AS 4.2/3? In 3.4, are there text to follow this line: The HAModClusterService bean defines: ?	1	1	https://issues.jboss.org/browse/JBPAPP-3273
10226	JBPAPP-3273	Radoslav, I've made both changes - you should be able to see them at the link. HTTPD is an acronym, though. Why lower case? :(	1	1	https://issues.jboss.org/browse/JBPAPP-3273

#	Requirement	Questions	Answered	Successfully Answered	URL
10227	JBPAPP-3273	Paul, Can you please review the mod_cluster documentation? Please note that the review needs to be completed by cob 03 Feb in order to allow Content Services to amend and prepare the document for GA release. Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-3273
10228	JBPAPP-3273	Thanks for including previous changes, more needed though: Httpd conf ========== Content of 2.1. Download mod-cluster components should not be the chapter which was missing, that one belongs to 2.3. Configure httpd section - move it. In that section, all parts that start with "mod_" prefix are names of modules and currectly they are the same style as configuration values - change all of those to smaller headlines (h3). Furthermore, the rest of conf value listings are in a table format - why not here as well? Demo ==== there is one extra accidental - mod_proxy.so - remove (2.9, bullet 1 end) Two lines of code: [bes@besdev ~]$ cd /home/bes/mod_cluster/demo/client [bes@besdev mod_cluster]$ ./run-demo.sh change to: [user@machine ~]$ cd /home/user/mod_cluster/demo/client [user@machine client]$ ./run-demo.sh bullet 9 - remove extra number '2' from sessions2 Client Driver Configuration Options -> formatting problem, should be a headline remove sentence "(In the future this will likely be changed to a percentage input, so X% can be invalidated, the rest abandoned.) " Load Generation Scenarios -> should be title again Chapter 4 ========= there is extra italics in $JBOSS_HOME/server/$PROFILE/deploy/jbossweb.sar/server.xml - fix also in jvmRoute there is extra italics, 3 times Note mod_cluster currently works with JBoss AS 5.0. -> remove in 4.4. there is * The state of each proxy will be kept in sync on each node in the cluster. but missing second bullet, which should be: * Proxies will be proactively notified of view changes - specifically crashed members, allowing the proxy to gracefully reconfigure itself potentially avoiding failover processing. in Proxy Configuration, default for "stickySession" is in bold (tag), remove bold. Migration ========= Where is this part? Seems to be missing. http://www.jboss.org/mod_cluster/native/mod_proxy.html	0	0	https://issues.jboss.org/browse/JBPAPP-3273
10229	JBPAPP-3273	Hi Gary, assigning to you as I went through it, looks good to me. Lets be done with it soon :-) I found just few issues (they don't need review afterwards): * 1.7. Current Limitations this section is basically useless, unless this is our standard, shouldn't we remove this 'placeholder'? * In 3.2. Migration from mod_proxy This sentence seems to be there by an accident; remove: "To migrate from mod_proxy to mod_cluster, refer to http://www.jboss.org/mod_cluster/native/mod_proxy.html." * 4.1. Common Configuration this is "$JBOSS_HOME/server/$PROFILE/deploy/jbossweb.sar/server.xml" also in unnecessary italics	0	0	https://issues.jboss.org/browse/JBPAPP-3273
10230	JBPAPP-3257	Martha, I assume this change has also been made for EAP 4.3.8? Documented in the 4.2.9/4.3.8 release notes.	1	1	https://issues.jboss.org/browse/JBPAPP-3257
10231	JBPAPP-3232	Does it need to be mentioned at? https://docspace.corp.redhat.com/docs/DOC-37998	1	0	https://issues.jboss.org/browse/JBPAPP-3232
10232	JBPAPP-3216	Prakash, assigning to you for any/all of the following information: * Requirements as listed in WFK 1.1 ERD: o Package contents o Installation instructions o Configuration instructions: + How to use WFK with EAP 5.x + How to use WFK with EWP 5.x partially committed + How to use WFK with EWS 1.x If you cannot provide it yourself, please let me know who can. :) Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-3216
10233	JBPAPP-3212	Jason, can you exclude these for us? Or we should not have seen these?	1	1	https://issues.jboss.org/browse/JBPAPP-3212
10234	JBPAPP-3191	Chris, I see you have already signed the jar, but we should not rename the jar as this jar is directly passed on to the customer. Can you rename it and put it in an archive to indicate it's signed or through some other means without changing the jar name? Please "Hand Over to QE" and assign to "JBoss QA Lead" after that.	0	0	https://issues.jboss.org/browse/JBPAPP-3191
10235	JBPAPP-3157	Is this fixed and planned for EAP 5.0.1 and 4.3.0.GA_CP08? Can this case be closed?	0	0	https://issues.jboss.org/browse/JBPAPP-3157
10236	JBPAPP-3131	Hi guys, JBM 1.x is in maintenance mode now, so normally there will be no new features added to it. Users who seek features beyond JBM should turn to HornetQ. Is there any good reason to add this feature to JBM 1.4? Thanks, Howard	1	1	https://issues.jboss.org/browse/JBPAPP-3131
10237	JBPAPP-3131	Hi Tomohisa, Thanks for the help. I have already implemented the basics of this feature. The code already checked in to Branch_1_4. The listPreparedTransactions() method now lists only the Transaction Ids. Please see the details at https://jira.jboss.org/jira/browse/JBMESSAGING-1772. The main classes changed are ServerPeer and TransactionRepository. However I didn't implement the functions you listed, namely: 1. sort by creation time --- so far in JBM doesn't record the transaction creation time. To implement this, we need to add it to the DB. 2. show target destination name -- I am not clear what is the target destination name, can you explain? 3. show message header and body in summary -- can you explain what should be in the summary? Thanks Howard	1	1	https://issues.jboss.org/browse/JBPAPP-3131
10238	JBPAPP-3131	Hello Howard, Thanks for your reply! ---- 1. sort by creation time --- so far in JBM doesn't record the transaction creation time. To implement this, we need to add it to the DB. ---- jbm_msg table already have timestamp. I think sorting by this timestamp may be helpful for many systems, although I know this timestamp is not reliable because it depends on client clock... ---- 2. show target destination name -- I am not clear what is the target destination name, can you explain? ---- Just showing JMSDestination. I want to know which destination is the transaction about to send to or receive from. ---- 3. show message header and body in summary -- can you explain what should be in the summary? ---- For example, JMSMessageID, JMSTimestamp(above 1.), body, and etc.. Showing above information will helpful when operator use these features.Operator may want to check which data is abandoned in prepared and about to be commit/rollback. Thanks Tomohisa	1	1	https://issues.jboss.org/browse/JBPAPP-3131
10239	JBPAPP-3131	Hi Tomohisa, Sorry for the delay. I tried out the implementation bit. It looks nice. I have a couple of questions: 1. the table returned by listPreparedTransactions() is perhaps too wide. If possible, can we reduce the number of columns? For example we may retain the "transaction_id", "branch_qual", "format_id", "global_txid", "queue_name", "message_id", and consider to discard the rest of columns as they are mostly properties of each message, not the properties of each transaction. 2. How the user can commit/rollback the transactions? One way is we provide a servlet (like you did) and deployed it. The user then input transaction ids and decision (commit or rollback) and invoke the mbean service (JmsXAController). This must involve a bit of extra work. I would suggest we implement the functionality using xmbean services so that we can use jmx-console to generate the GUI for us. It could be easy to integrate the JmsXAController operations to ServerPeer. And this way we don't need to provide a separate servlet. What do you think?	1	1	https://issues.jboss.org/browse/JBPAPP-3131
10240	JBPAPP-3119	Documented in the EAP 4.2.CP08 and 4.3.CP07 release notes.	1	1	https://issues.jboss.org/browse/JBPAPP-3119
10241	JBPAPP-3088	In EAP 5, this feature is a bit extensive and uses a <SessionCookie> element nested in <Context>, and it can configure domain, path, comment, httpOnly, secure, maxAge and name. I designed it mirroring the Servlet 3.0 session cookie configuration features. Obviously it is a large backport (but adding httpOnly is a significant patch by itself, as Servlet 2.5 cookies do not allow this flag). This feature request should not be creating incompatible configuration, so <Context useHttpOnly="true"> is not acceptable as it is not going to be supported in EAP 5+.	1	1	https://issues.jboss.org/browse/JBPAPP-3088
10242	JBPAPP-3083	Jimmy, can you explain why this occurred and what you did to fix it, for the release notes?	0	0	https://issues.jboss.org/browse/JBPAPP-3083

#	Requirement	Questions	Answered	Successfully Answered	URL
10243	JBPAPP-3055	(10:18:53 AM) bstansberry: wolfc: jason was saying something about you having an problem w/ hibernate-jbosscache2 ? (10:19:32 AM) wolfc: bstansberry: hehe, yes EAP 5 uses it AS 5.2 does not. So I have an issue in ejb3/testsuite. (10:19:43 AM) wolfc: bstansberry: where are you btw? what's your time? (10:19:58 AM) bstansberry: wolfc: beijing. 5:20 PM (10:20:27 AM) bstansberry: wolfc: ic. would moving the clustered entity tests out of ejb solve your problem? (10:20:53 AM) wolfc: bstansberry: yes, what does it actually test? (10:21:13 AM) wolfc: bstansberry: I'm also looking at entityoptimisticlocking (10:21:45 AM) bstansberry: wolfc: it's all tests of Hibernate second level caching, except via JPA (10:22:14 AM) wolfc: bstansberry: that's what I thought. So it should actually be somewhere in jpa-hibernate-int, not in ejb3 (10:22:40 AM) bstansberry: wolfc: those tests require an appserver (10:23:38 AM) wolfc: bstansberry: no, these tests require a runtime :-) (10:23:55 AM) wolfc: bstansberry: using AS is the excuse :-) (10:24:10 AM) bstansberry: wolfc: ok, fair enough. but there are already a lot of tests of this stuff in Hibernate itself (10:25:14 AM) wolfc: bstansberry: yeah, we got a lot of superfluous tests in ejb3 (10:25:21 AM) bstansberry: wolfc: using AS is the lazy way to test that the stuff that's already tested actually works in the real runtime :) (10:25:23 AM) wolfc: bstansberry: but I can't determine whether this is one of them (10:25:48 AM) bstansberry: wolfc: well it has nothing to do with session beans or mdbs or service beans (10:27:23 AM) wolfc: bstansberry: so we can safely say this test is relatively bogus from the ejb3 p.o.v. (10:27:59 AM) bstansberry: wolfc: yes. it's there from the old AS 4 days when jpa integration was in the ejb3 module (10:28:11 AM) wolfc: bstansberry: I'm checking whether removing the actual hibernate.cache.region.factory_class setting will make stuff pass (10:28:39 AM) bstansberry: wolfc: oh, is that that problem? (10:28:53 AM) bstansberry: wolfc: use Hibernate 3.5 Beta2 (10:29:03 AM) wolfc: bstansberry: org.hibernate.cache.jbc2.JndiMultiplexedJBossCacheRegionFactory does not exist anymore (10:29:10 AM) bstansberry: wolfc: ^^^ (10:29:14 AM) wolfc: bstansberry: but how will this affect EAP 5? (10:29:21 AM) jbossfox is now known as fox_brb (10:29:31 AM) bstansberry: wolfc: steve refactored hibernate; i restored the old class (10:29:48 AM) bstansberry: wolfc: EAP 5 does not use Hibernate 3.5 (10:30:50 AM) bstansberry: wolfc: should say: steve refactored hibernate for 3.5 beta1, i restored the old class for backward compatiblity in beta2 (10:30:53 AM) wolfc: http://repository.jboss.org/maven2/org/hibernate/hibernate-jbosscache2/ (10:31:01 AM) wolfc: bstansberry: there is no 3.5b2 (10:32:40 AM) bstansberry: wolfc: http://repository.jboss.org/maven2/org/hibernate/hibernate-jbosscache/ (10:32:54 AM) bstansberry: wolfc: he made the whole "2" notion go away (10:33:10 AM) bstansberry: wolfc: since this stuff actually uses JBC 3 (10:33:49 AM) wolfc: bstansberry: ah I already found it weird that changing all settings to jbc instead of jbc2 would start trigger 'need jbc2 settings' errors (10:35:22 AM) bstansberry: wolfc: do you import org/hibernate/hibernate-jbosscache2 for any reason other than testsuite? (10:35:32 AM) wolfc: bstansberry: no, only there (10:36:19 AM) wolfc: building... (10:36:23 AM) bstansberry: wolfc: ok; i gotta go in a bit. can you open a JIRA to get these tests out of ejb3, assign to me? I'll do it tomorrow (10:37:05 AM) wolfc: bstansberry: this means porting them over to EAP 5, AS 5.2 & trunk (10:38:01 AM) bstansberry: wolfc: ok (10:38:13 AM) wolfc: bstansberry: got yourself a deal :-) (10:42:39 AM) wolfc: bstansberry: one more question: do we actually need those tests on EAP 5? (10:43:00 AM) bstansberry: wolfc: yes, i think so (10:43:05 AM) wolfc: rgr	1	1	https://issues.jboss.org/browse/JBPAPP-3055
10244	JBPAPP-3051	Strong, this looks like it's just a change in the way the code works - listing it as a fixed issue rather than a known issue. Will this change make it into EWP5?	0	0	https://issues.jboss.org/browse/JBPAPP-3051
10245	JBPAPP-3047	What is your opinion Jason: add to clean or create a clean-all target?	1	0	https://issues.jboss.org/browse/JBPAPP-3047
10246	JBPAPP-3047	Deferring to 5.0.2. Re: your question Fernando, if the artifacts in maven-thirdparty are being pulled from the local maven repo (i.e. ~/.m2/repository) then I think including this as a step in ./build.sh clean makes sense. Anyone know the timing difference to do a build following a ./build.sh clean w/ maven-thirdparty empty vs. full? Assuming that all artificats are in ~/.m2/repository i.e. there's no need to pull them from a remote repo.	0	0	https://issues.jboss.org/browse/JBPAPP-3047
10247	JBPAPP-3047	Didn't we have a clobber target that was supposed to be the 'super'-clean in the AS build that also trashed thirdparty downloads?	0	0	https://issues.jboss.org/browse/JBPAPP-3047
10248	JBPAPP-3042	13.1.2. Launching a JBoss Enterprise Application Platform Cluster ========== "for node1 isn't technically necessary, since ports-default is the ... default. " -> remove 3 dots in "@Cache (usage=CacheConcurrencyStrategy.TRANSACTIONAL)" what is our coding standard? i would remove the extra space after @Cache "public class Account implements Serializable" we probably want to keep the coding standards and show its the actual implementation, so it should include the curly brace "public class Account implements Serializable {"	1	1	https://issues.jboss.org/browse/JBPAPP-3042
10249	JBPAPP-3042	Brian, would you agree with the above? Reassign to Laura if you do.	1	1	https://issues.jboss.org/browse/JBPAPP-3042
10250	JBPAPP-3038	should the fix version include FP release?	1	1	https://issues.jboss.org/browse/JBPAPP-3038
10251	JBPAPP-3023	Rajesh, is this issue fixed for 5.0.1?	1	1	https://issues.jboss.org/browse/JBPAPP-3023
10252	JBPAPP-3023	Pavel, can you verify in the staged doc and close this issue?	0	0	https://issues.jboss.org/browse/JBPAPP-3023
10253	JBPAPP-3023	Pavel, I did not receive any comments when you re-opened the issue. Can you add them again and reassign to me if any changes are required?	1	1	https://issues.jboss.org/browse/JBPAPP-3023
10254	JBPAPP-3023	Howard, can you verify Pavel's comment re. queue/app1/email? This was not my understanding from the above discussion.	1	1	https://issues.jboss.org/browse/JBPAPP-3023
10255	JBPAPP-3018	Richard - Is EAP 5.1.1 going to need any modifications besides the run.sh to run with IPv6? Besides Transactions, are there any other caveats? I am trying to address JBPAPP-5555 for EAP 5.1.1 and Jimmy Wilson pointed me to this JIRA.	0	0	https://issues.jboss.org/browse/JBPAPP-3018

#	Requirement	Questions	Answered	Successfully Answered	URL
10256	JBPAPP-2997	NOTE ON HSQLDB and IPv6: A number of test cases use an HSQLDB instance during their execution. The database instance is started in the TestSeup.setUp() and torn down in TestSetup.tearDown(). A helper class DBSetup.java is used in EJB3 tests and clustering tests to do the initialization. The database is started up with no -address specified (and so by default binds to whatever localhost returns: java -cp .:hsqldb.jar org.hsqldb.server.Server -database cif-db -port 1701 -silent true -trace true -no_system_exit true The database is shutdown by using JDBC to connet to the URL: String dbURL = "jdbc:hsqldb:hsql://" + System.getProperty("jbosstest.server.host.url", "localhost") + ":1701"; The problem is that if by chance the system property jbosstest.server.host.url is not set oin the ant target for the test case, we use localhost instead. This was happening when running the AS testsuite against IPv6 and so IPv4 connections and hostnames were being used for the database accesses. The EJB3 testsuite fared better and revealed that HSQLDB 1.8.0 is not IPv6 compatible, and tests fail with the exception: SOO10 Invalid argument in JDBC call java.sql.SQLException: SOO10 Invalid argument in JDBC call at org.hsqldb.jdbcDriver.getConnection(Unknown Source) at org.hsqldb.jdbcDriver.connect(Unknown Source) at java.sql.DriverManager.getConnection(DriverManager.java:525) at java.sql.DriverManager.getConnection(DriverManager.java:171) at org.jboss.ejb3.test.DBSetup.tearDown(DBSetup.java:89) at junit.extensions.TestSetup$1.protect(TestSetup.java:20) at junit.extensions.TestSetup.run(TestSetup.java:23) This is what happens when we try to connect using a URL containing an IPv6 literal. Thus, the testsuite was not truly running against IPv6. It appears as though IPv6 support was intriduced in HSQLDB 1.9.0, but this is still in RC stage and no release has been made available. Furthermore, it has some package name chanhges which means that some clients will have to be rewritten. Hopefully only that to start the DB.	1	1	https://issues.jboss.org/browse/JBPAPP-2997
10257	JBPAPP-2970	What is the fix for this issue?	1	1	https://issues.jboss.org/browse/JBPAPP-2970
10258	JBPAPP-2968	Ondrej, can you pass this on for testing?	0	0	https://issues.jboss.org/browse/JBPAPP-2968
10259	JBPAPP-2954	How can I test this?Can you please provide further information?	1	1	https://issues.jboss.org/browse/JBPAPP-2954
10260	JBPAPP-2950	Could you put fixed version of the Installation Guide to the doc stage, please? I can still see the old version in [1]. [1] http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/4.3/html-single/Installation_Guide/index.html	0	0	https://issues.jboss.org/browse/JBPAPP-2950
10261	JBPAPP-2948	Is this security loophole of having the jmx_remoting.sar file in /deploy directory resolved in JBoss EAP 6 or AS 7.1.2? Is there a plan to work on a more secure JSR 160 implementation?	0	0	https://issues.jboss.org/browse/JBPAPP-2948
10262	JBPAPP-2938	Martha, Can you please build, upload and then switch the EAP 4 branch to use this? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2938
10263	JBPAPP-2935	Hi, will this JIRA be fixed for this EAP 4 CP cycle?	0	0	https://issues.jboss.org/browse/JBPAPP-2935
10264	JBPAPP-2888	Juca, what is the status on this?	0	0	https://issues.jboss.org/browse/JBPAPP-2888
10265	JBPAPP-2888	Juca, that patch looks okay, I'd think you can apply it. but i only see you added "org.hibernate.tool.instrument.cglib.InstrumentTask" to testsuite/imports/sections/hibernate.xml, no javassist instrumentTask. I kind of remember that you told me that should be changed to javassist manually, is it true?	1	0	https://issues.jboss.org/browse/JBPAPP-2888
10266	JBPAPP-2848	Assigning to you, Gary - note that this issue affects both the 4.2 and 4.3 CP releases.	0	0	https://issues.jboss.org/browse/JBPAPP-2848
10267	JBPAPP-2848	Amended Paragraph 6.3 of 4.2CP08 and 4.3CP07 docs as requested. I will reassign this issue for verification when the docs are available for review on the doc-stage.	0	0	https://issues.jboss.org/browse/JBPAPP-2848
10268	JBPAPP-2837	Assigning to Martha as she takes care of the EAP 4 CPs now. Is JBossWS the only one that needs upgrading, any other of the jbossws components or JAXB?	0	0	https://issues.jboss.org/browse/JBPAPP-2837
10269	JBPAPP-2836	of:org.jboss.aop.asintegration.jboss5.VFSClassLoaderScopingPolicy java.lang.Error: Error determining protection domain for ZipEntryHandler@12893236[path=org/jboss/aop/asintegration/jboss5/VFSClassLoaderScopingPolicy.class context=file:/qa/hudson_ws/workspace/JBoss-EAP-5.0.x-Testsuite-sun16/JBossEAP_5_0_x/build/output/jboss-5.0.0.Branch/lib/jboss-aop-asintegration-mc.jar real=file:/qa/hudson_ws/workspace/JBoss-EAP-5.0.x-Testsuite-sun16/JBossEAP_5_0_x/build/output/jboss-5.0.0.Branch/lib/jboss-aop-asintegration-mc.jar/org/jboss/aop/asintegration/jboss5/VFSClassLoaderScopingPolicy.class] at org.jboss.classloading.spi.vfs.policy.VFSClassLoaderPolicy.getProtectionDomain(VFSClassLoaderPolicy.java:664) at org.jboss.classloader.spi.base.BaseClassLoader$2.run(BaseClassLoader.java:550) at org.jboss.classloader.spi.base.BaseClassLoader$2.run(BaseClassLoader.java:535) at java.security.AccessController.doPrivileged(Native Method) at org.jboss.classloader.spi.base.BaseClassLoader.loadClassLocally(BaseClassLoader.java:533) at org.jboss.classloader.spi.base.BaseClassLoader.loadClassLocally(BaseClassLoader.java:510) at org.jboss.classloader.spi.base.BaseDelegateLoader.loadClass(BaseDelegateLoader.java:134) at org.jboss.classloader.spi.filter.FilteredDelegateLoader.loadClass(FilteredDelegateLoader.java:131) at org.jboss.classloader.spi.base.ClassLoadingTask$ThreadTask.run(ClassLoadingTask.java:452) at org.jboss.classloader.spi.base.ClassLoaderManager.nextTask(ClassLoaderManager.java:251) at org.jboss.classloader.spi.base.ClassLoaderManager.process(ClassLoaderManager.java:150) at org.jboss.classloader.spi.base.BaseClassLoaderDomain.loadClass(BaseClassLoaderDomain.java:265) at org.jboss.classloader.spi.base.BaseClassLoaderDomain.loadClass(BaseClassLoaderDomain.java:1119) at org.jboss.classloader.spi.base.BaseClassLoader.loadClassFromDomain(BaseClassLoader.java:801) at org.jboss.classloader.spi.base.BaseClassLoader.loadClass(BaseClassLoader.java:444) at java.lang.ClassLoader.loadClass(ClassLoader.java:252) at java.lang.ClassLoader.loadClassInternal(ClassLoader.java:320) at java.lang.Class.forName0(Native Method) at java.lang.Class.forName(Class.java:247) at org.jboss.reflect.plugins.introspection.IntrospectionTypeInfoFactoryImpl.resolveComplexTypeInfo(IntrospectionTypeInfoFactoryImpl.java:419) at org.jboss.reflect.plugins.introspection.IntrospectionTypeInfoFactoryImpl.getTypeInfo(IntrospectionTypeInfoFactoryImpl.java:388) at org.jboss.reflect.plugins.introspection.IntrospectionTypeInfoFactory.getTypeInfo(IntrospectionTypeInfoFactory.java:54) at org.jboss.classadapter.plugins.BasicClassAdapterFactory.getClassAdapter(BasicClassAdapterFactory.java:61) at org.jboss.config.plugins.AbstractConfiguration.getBeanInfo(AbstractConfiguration.java:86) at org.jboss.kernel.plugins.config.AbstractKernelConfig.getBeanInfo(AbstractKernelConfig.java:80) at org.jboss.kernel.plugins.config.AbstractKernelConfigurator.getBeanInfo(AbstractKernelConfigurator.java:77) at org.jboss.kernel.plugins.config.AbstractKernelConfigurator.getBeanInfo(AbstractKernelConfigurator.java:96) at org.jboss.kernel.plugins.dependency.PreInstallAction.installActionInternal(PreInstallAction.java:90) at	1	1	https://issues.jboss.org/browse/JBPAPP-2836
10270	JBPAPP-2814	Martha, can you please also build, upload and flip the EAP 4 branch to use this build?	0	0	https://issues.jboss.org/browse/JBPAPP-2814
10271	JBPAPP-2809	Gary, assigning to you. Please note that this affects both 4.2 and 4.3 CP releases.	1	1	https://issues.jboss.org/browse/JBPAPP-2809

#	Requirement	Questions	Answered	Successfully Answered	URL
10272	JBPAPP-2809	These changes have been applied to 4.2_CP08, 4.3_CP07. I will reassign this issue for verification when the docs are available for review on the doc-stage.	0	0	https://issues.jboss.org/browse/JBPAPP-2809
10273	JBPAPP-2809	Pavel, I have corrected this and will stage the docs so you can verify and close this afternoon. This fix will only be in the online documentation unless we do a CR3 build; I will mention the error in the bundled docs in the release notes.	1	1	https://issues.jboss.org/browse/JBPAPP-2809
10274	JBPAPP-2808	There are other IPv6 related issues which need to be added to this issues list of dependencies. Will do as time goes on and before the JGroups release 2.4.7.GA is tagged.	1	1	https://issues.jboss.org/browse/JBPAPP-2808
10275	JBPAPP-2808	Not yet, but it can be in relatively short notice. We wanted to double check some IPv6 related issues. What is the latest reasonable possible date to get it to you? (I'm stuck with EAP 5 testing at the moment, as well as IP_v6 issues for EAP 4.2./4.3 CPs)	1	1	https://issues.jboss.org/browse/JBPAPP-2808
10276	JBPAPP-2808	We're already starting to upgrade and build the components for this CP, would it be possible to tag this by the end of this week? Or do we need to get the IPv6 fixes into this tag as well?	1	1	https://issues.jboss.org/browse/JBPAPP-2808
10277	JBPAPP-2808	Hi Richard, Is the tag for JGroups 2.4.7 ready? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2808
10278	JBPAPP-2808	Martha, When the tag is ready, can you please build in brew, upload to r.j.o and switch the JBPAPP_4_2_0_GA_CP branch to use the new build? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2808
10279	JBPAPP-2795	Hi Ron, We've built Remoting 2.2.3.SP1 and it's up on http://repository.jboss.org/jboss/remoting/2.2.3.SP1-brew/, the EAP 4 branch has also been switched to use this version. The Hudson runs look good so far, it runs ant tests in the testsuite directory, are those testcases part of the regular Hudson run?	1	1	https://issues.jboss.org/browse/JBPAPP-2795
10280	JBPAPP-2795	Hi Permaine, I'm not sure which testcases you're referring to. I think there might be some Remoting integration tests in the testsuite directory, though I'm not sure if they get run in Hudson. Are those the tests you're talking about? The Remoting test suite itself gets run separately in Hudson, though Hudson is configured to run the tests from a newer branch than the one that yields the 2.2.x releases. I run the 2.2.x tests myself. -Ron	1	1	https://issues.jboss.org/browse/JBPAPP-2795
10281	JBPAPP-2795	Hi Ron, I just confirmed with Rajesh that those test failures listed are part of the nightly Hudson run, so, they are passing for both EAP 4.2 and 4.3. Now that we have 2.2.3.SP1 built in brew and uploaded to r.j.o, can this be resolved? Cheers, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2795
10282	JBPAPP-2795	Hi Permaine, Re: "I'm not sure which testcases you're referring to" Could I possibly be any more oblivious? The test cases you were referring to were the ones I mentioned in my previous post. Doh! Re: "can this be resolved?" I rarely run the AS test suite, so when tests fail, I assume the problem is one of timing, configuration, or something like that, local to my own environment. The important thing is that these tests have no apparent relationship to Remoting, so I didn't cause any regressions in release 2.2.3.SP1. I think the fact that these tests are passing in hudson is all the resolution that's needed. Let me know if you disagree. -Ron	1	1	https://issues.jboss.org/browse/JBPAPP-2795
10283	JBPAPP-2781	Martha, can you please build this in brew and upload to r.j.o? Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-2781
10284	JBPAPP-2780	Hi David, Can you please build this in brew and upload to r.j.o? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2780
10285	JBPAPP-2779	Wasn't this a problem on JBoss Common 1.x? EAP 5 uses JBoss Common 2.x, so does it apply? Also, if we needed a new JBoss Common we have not received it by component freeze date, there is no "Upgrade JBoss Comon to ...." JIRA opened, so I guess this must be moved to 5.0.2.	0	0	https://issues.jboss.org/browse/JBPAPP-2779
10286	JBPAPP-2779	Why is this issue still open? It looks like we have JBoss Commons 2.2.17 in EAP 5.1.1. Please close as Resolved if in EAP 5.1.1, if not, close as out of date. Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-2779
10287	JBPAPP-2726	Do we have a 2.2.16 tag yet?	0	0	https://issues.jboss.org/browse/JBPAPP-2726
10288	JBPAPP-2649	Found several minor issues, reopening: * Chapter 23. JBoss Cache Configuration and Deployment "JBoss Cache documentation available at http://www.jboss.org/jbosscache. " -> make the link clickable - but shouldnt we navigate users to the productized eap5 version of this document and not the community one? we need to make sure we are adding value * 23.1. Key JBoss Cache Configuration Options "...encouraged to read the JBoss Cache documentation available at http://www.jboss.org/jbosscache. " -> same comment as above "seen in the JBoss Cache javadocs. " -> capitalize J in javadocs * 23.1.2. Cache Mode 2 phase-commit -> change to "two-phase commit" * 23.1.4. Concurrent Access multi-versioned concurrency control -> remove the extra "ed", its "multi-version concurrency control" * 23.1.6. Eviction "see the Eviction chapter in the "Using JBoss Cache as a Hibernate Second Level Cache" guide at http://www.jboss.org/jbossclustering/docs/hibernate-jbosscache-guide-3.pdf." -> make the link clickable and possibly change it ot the productized version of that document * 23.1.8. Buddy Replication in a section starting with " The main things you would be likely to configure are: ", change 3 times two hyphen "--" to a single dash * 23.2.2. Deployment Via a -service.xml File JBoss 4 -> change to 4.x since there are more version of JBoss 4 JBoss 4, JBoss 5 -> four times, change to the product name, probably Enterprise App... (reviewed docs at: http://www.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0.0/html-single/Administration_And_Configuration_Guide/index.html#jbosscache.chapt)	0	0	https://issues.jboss.org/browse/JBPAPP-2649
10289	JBPAPP-2649	Some of the minor changes are still not reflected. "seen in the JBoss Cache javadocs. " -> capitalize J in javadocs 2 phase-commit -> change to "two-phase commit" multi-versioned concurrency control -> remove the extra "ed", its "multi-version concurrency control" JBoss 4, JBoss 5 -> change to the product name, probably Enterprise App... ps: is there no productized version of "Using JBoss Cache as a Hibernate Second Level Cache"?	1	1	https://issues.jboss.org/browse/JBPAPP-2649
10290	JBPAPP-2647	Hi Howard, how is this coming along?	1	1	https://issues.jboss.org/browse/JBPAPP-2647

#	Requirement	Questions	Answered	Successfully Answered	URL
10291	JBPAPP-2646	Issues found. Reopening. * Chapter 19. HTTP Services "Load-balancing of incoming invocations" - would somebody call a HTTP request an invocation? i would use a different term automtically -> automatically http://httpd.apache.org/ -> make it clicable - a link http://www.apache.org/dist/jakarta/tomcat-connectors/jk/binaries/ -> make it clicable - a link "Next, download mod_jk binaries. Several versions of mod_jk exist as well. We strongly advise you to use mod_jk 1.2.x, as both mod_jk and mod_jk2 are deprecated," -> what did the author mean? we recommend mod_jk and its deprecated? that is absurd "worker.node2.host= node2.mydomain.com" -> remove the space after `equals` sign JBoss AS instances -> JBoss Enterprise App.... http://jakarta.apache.org/tomcat/connectors-doc/config/workers.html -> wrong link! where do we want to point them? "The only thing you must change is the worker.loadbalancer.balanced_workers line: it must list all workers previously defined in the same file: load-balancing will happen over these workers. " -> there are 2 colons within the same sentence, rewrite somehow * 19.1.4. Configuring JBoss to work with mod_jk remove one group of 3 dots and leave only 1 group of 3 dots: ... ... </Engine> JBoss AS -> JBoss Enterprise App.... remove the two plus signs from here: <!-- Define an AJP 1.3 Connector on port 8009 --> +<Connector protocol="AJP/1.3" port="8009" address="${jboss.bind.address}" + redirectPort="8443" /> rewrite "Apache+mod_jk" to "Apache with mod_jk" http://www.jboss.org/community/wiki/UsingModjk12WithJBoss -> make it a link "with AS 5." -> with Enterprise App.... 5 JBoss AS 4 -> Enterprise App.... JBoss AS 5.0.0.GA -> Enterprise App... "In JBoss AS 5.1.0," -> talking about AS realease, is this desired? * 19.2.2. HttpSession Passivation and Activation AS 5 -> Enterprise App... * 19.3. Using FIELD level replication title says "FIELD level", the first pararaph starts with hyphenated "FIELD-level" -> sync to the best one @@org.jboss.cache.pojo.annotation.Replicable -> remove extra @, there should be only one JBoss AS 5 -> expand to eap " public class Address " -> remove extra space before word `public' /** * Represents a street address. * @@org.jboss.cache.pojo.annotation.Replicable */ in Xdoclet -> is this correct or only one @? // this will update and replicate the postal code. -> remove the last period in the comment * 19.4. Using Clustered Single Sign On title is not hyphenated but "single sign-on, " is -> synchronize and add the abbreviation (SSO) afterwards so that its introduced for later use * 19.4.2. SSO Behavior "as he accesses " -> "as he or she accesses "	0	0	https://issues.jboss.org/browse/JBPAPP-2646
10292	JBPAPP-2646	There are still many references to "AS" which should be changed to EAP. And one more forgotten: "worker.node2.host= node2.mydomain.com" -> remove the space after '=' (equals) sign	0	0	https://issues.jboss.org/browse/JBPAPP-2646
10293	JBPAPP-2645	This chapter is in very good shape, here are the issues: * 16 -> paths in this chapter are all relative, can be confusing it its some global conf dir or only profile's dir [what is our guide standard?] * 16.1. How it works in Note #1, remove the extra period/dot, expression in parenthesis is just part of the previous sentence "different "target" bound under the same name. (JNDI on node 1 will have a binding" * 16.2.1.3. How can I tell if things are being bound into HA-JNDI that shouldn't be? missing period at the end of the sentence: "Please visit the following link for an example: Problem with removing a Node from Cluster" * 16.2.2. For clients running outside the Enterprise Application Platform make "initialising" -> US spelling initializing "specifies the TTL (time-to-live) f or autodiscovery IP multicast packets" -> remove extra space after "f" to form the word "for" " Besides the above dependency injected services, the available configuration attributes for the HAJNDI bean are as follows: " -> make HAJNDI hyphenated HA-JNDI (watch out! not in the xml schema!) "Backlog specifies the maximum queue length" -> lower-case letter "B" in "backlog" (it refers to the property name which is all in lower-case) missing space in "LoadBalancePolicyimplementation" -> 2 words: LoadBalancePolicy implementation missing <a> link "See the Introduction and Quick Start chapter for details. "	0	0	https://issues.jboss.org/browse/JBPAPP-2645
10294	JBPAPP-2642	Radoslav, can you clarify what you meant in the section quoted in the previous comment?	1	1	https://issues.jboss.org/browse/JBPAPP-2642
10295	JBPAPP-2642	Hm, I dont know. Its probably fixed already, thanks for getting back to me. Where can I see the current version of the doc?	1	1	https://issues.jboss.org/browse/JBPAPP-2642
10296	JBPAPP-2642	Hey Laura, I just found what I meant. In section "21.1.1.2. POJO deployments using HASingletonController", there is one big code block with Java code. However within the same block there are 2 different public classes defined, which is not possible in Java. This is just cosmetic, but split the code block into two (so that there is a visual 'break' between those 2 classes). Move the fix version if its not doable. Otherwise verified. (at http://downtown.englab.bne.redhat.com/drafts/JBoss_Enterprise_Application_Platform/5.0.1/Administration_And_Configuration_Guide/html/deployment.chapt.html)	0	0	https://issues.jboss.org/browse/JBPAPP-2642
10297	JBPAPP-2641	Reopening (please talk to Brian/Paul for issues i didn't provide solution for.) * Chapter 18. Clustered Entity EJBs lonely "e.g.", make it part of the sentence above it, "Configuration of a the second-level cache is done via your EJB3 deployment's persistence.xml. " 2 times: in java code: make the quotes the `java quotes` or what to call it, like this " @Cache(usage = CacheConcurrencyStrategy.TRANSACTIONAL, region = "Account") "region named {region_prefix}/org/hibernate/cache/StandardQueryCache." -> for expressions in path we use probably these characters: <, >. missing space 'of the' in "region to use instead ofthe default" in java code: @QueryHint(name = "org.hibernate.cacheRegion, value = "Queries") rewrite quotes and add missing quote to: @QueryHint(name = "org.hibernate.cacheRegion", value = "Queries") can somebody please QA the sanity/up-to-dateness of the hibernate stuff? (Juca?) * 18.2. Entity Bean in EJB 2.x typo special -> "fit into the sepecial case situation" synched-> synchronized link is broken! remove? "Seppuku pattern http://dima.dhs.org/misc/readOnlyUpdates.html." in note: broken link again http://www.jboss.org/jbossgroup/partners.jsp	1	1	https://issues.jboss.org/browse/JBPAPP-2641
10298	JBPAPP-2641	Juca, can you please look at these 3 sections with Hibernate stuff and assign back to Laura? Thanks! 18.1.1. Configure the distributed cache 18.1.2. Configure the entity beans for cache 18.1.3. Query result caching http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html/Administration_And_Configuration_Guide/clustering-entity.html#clustering-entity-30-cache	1	1	https://issues.jboss.org/browse/JBPAPP-2641

#	Requirement	Questions	Answered	Successfully Answered	URL
10299	JBPAPP-2640	* Chapter 17. Clustered Session EJBs note that this chapter often refers to "JBoss" with JBoss AS/EAP in mind, please synchronize with the rest of the guide or is "JBoss" (not in the context of the name of the company) acceptable? * 17.2.1. The EJB application configuration "(irregardless of whether maxSize beans are cached.)" -> move period outside of the parentheses like). * 17.2.3. CacheManager service configuration (aka buddies) -> sync, missing periods "a.k.a." * 17.4.2. Optimize state replication remove "This feature is available on JBoss Enterprise Application Platform 3.0.1+ " * 17.4.3. The HASessionStateService configuration thought: maybe make path "all/deploy/cluster/ha-legacy-jboss-beans.xml" generic to include all clustered profiles not only all? * 17.4.4. Handling Cluster Restart missing link: "We have covered the HA smart client architecture in the section called "Client-side interceptor architecture"." failover -> fail-over (sync; or the other way in chapter 16) HAJNDI -> with hyphen HA-JNDI The 3.2.7+/4.0.2+ releases contain a RetryInterceptor that can be -> should this be here? is this *only* for EAP5? then make it "RetryInterceptor can be..." * 17.4.5. JNDI Lookup Process JVM -> VM all the java code references are in standard paragraph style, should they be made <code> ? missing <a> link in "See the section on "ClusteredJNDI Services"" some weird references to RC and CR (4.0.4.RC1), what with that, ask Brian/Paul	0	0	https://issues.jboss.org/browse/JBPAPP-2640
10300	JBPAPP-2638	possible system property setting we can use to configure VFS behavior. <property key> - <options>; <default> jboss.vfs.forceCopy - true|false; false This defines how nested jars should be handled. If forceCopy equals true, we create a temporary copy of the nested jar, and re-wire VFS accordingly. If forceCopy equals false, we handle nested jars in-memory, which doesn't create temporary copy, but is more memory consuming. Currently JBossAS forces temporary copy by default. If this query part - "useCopyJarHandler" - is used as part of URI query, you can configure force-copy at runtime, per URI root (if it doesn't already exist). jboss.vfs.forceVfsJar - true|false; false By using this setting, you can fall back to old jar handling. Note that old jar handling was deprecated in favor of new zip handling code. jboss.vfs.forceNoReaper - true|false; false To gain a bit on performance, we close jar files asynchronously - via separate reaper thread. If you want to close jars files synchronously, you can force no usage of reaper thread. This can also be defined via URI query - "noReaper" query part. jboss.vfs.forceCaseSensitive - true|false; false Do we require case sensitive file path matching? With this enabled you can force differentiation between lower and upper cased file paths. jboss.vfs.optimizeForMemory - true|false; false With this enabled we re-order in-memory jar handling a bit, to gain on memory consumption. jboss.vfs.cache - full class name of cache implementation; org.jboss.virtual.spi.cache.helpers.NoopVFSCache In order to re-use existing temporary files (in order not to re-do all unpacking and wiring), you can define a VFS cache implementation which the VFS registry will use to keep its existing VFS roots. Every VirtualFile lookup from VFS class uses this singleton cache instance to check for existing matching cache entry. By matching we also consider any existing "ancestor" from which you can get a hold of exact VirtualFile instance. *VFS Cache Tuning As described before, VFS cache holds VFS roots, from which any VirtualFile lookup can access existing VirtualFile instances. This is a sepcially useful in the case of temporary files (created from nested jars), meaning you don't have to do multiple unpackings for nested jar file related resources. By default in VFS there is no caching involved - org.jboss.virtual.spi.cache.helpers.NoopVFSCache is used. But you can provide your own implementation or choose from existing VFS implementations. Cache implementations from "org.jboss.virtual.plugins.cache" package: - SoftRefVFSCache; uses soft reference as map's entry value - WeakRefVFSCache; uses weak reference as map's entry value - TimedVFSCache; evicts cache entries after defaultLifetime - LRUVFSCache; evicts cache entries based on LRU, keeping min and max entries - CombinedVFSCache; holds few permanent roots, any other new root is cached in its realCache property In JBossAS we use CombinedVFSCache as we know which are our permanent roots to watch and keep. This is how it's configured in MC's bean configuration file. <bean name="VFSCache"> <constructor factoryClass="org.jboss.virtual.spi.cache.VFSCacheFactory" factoryMethod="getInstance"> <!-- Use the CombinedVFSCache implementation --> <parameter>org.jboss.virtual.plugins.cache.CombinedVFSCache</parameter> </constructor> <start ignored="true"/> <property name="permanentRoots"> <map keyClass="java.net.URL" valueClass="org.jboss.virtual.spi.ExceptionHandler"> <entry> <key>${jboss.lib.url}</key> <value><null/></value> </entry> <entry> <key>${jboss.common.lib.url}</key> <value><inject bean="VfsNamesExceptionHandler"/></value> </entry> <entry> <key>${jboss.server.lib.url}</key> <value><inject bean="VfsNamesExceptionHandler"/></value> </entry> <entry> <key>${jboss.server.home.url}deploy</key> <value><inject bean="VfsNamesExceptionHandler"/></value> </entry> </map> </property> <property name="realCache"> <bean class="org.jboss.virtual.plugins.cache.IterableTimedVFSCache"/> </property> </bean> Any new custom VFS root (e.g. additional deploy/ directory) should be added to this configuration. *Annotation Scanning Tuning There are currently three ways to limit resources scanning. First one is to provide	0	0	https://issues.jboss.org/browse/JBPAPP-2638
10301	JBPAPP-2638	Can this issue be assigned to the docs team to be added to the Admin and Config Guide?	1	1	https://issues.jboss.org/browse/JBPAPP-2638
10302	JBPAPP-2635	Reopening, mostly typos: * Chapter 24. JBoss Enterprise Application Platform 5 Performance Tuning * 24.3. Operating System Performance Tuning unix based operating systems -> capitalize word Unix java based applications -> hyphenate and capitalize Java, "Java-based applications" "java heap" -> capitalize Java bench-marking -> remove the hyphen, its "benchmarking" "JVM Hardware" -> i dont understand, probably a missing period or lowercase hardware but that doesnt make much sense to me. "Note the process id for the "org.jboss.Main" application" -> make ID capitalized; 2 times * 24.6.1. Memory usage "and that is not part of the heap per se." -> i would suggest changing "per se" to something easily comprehensible to non-native English users. e.g.: itself change "-XX:MaxPermSize?=256m" to "-XX:MaxPermSize=256m"; two times (i.e. remove the extra question mark) "The jstat tool (http://java.sun.com/javase/6/docs/technotes/tools/share/jstat.html)" -> make it a clickable link * 24.6.1.1. VFS Tuning can we explain the VFS abbreviation here too? its only at the beginning of the guide (Virtual File System) "More information about this can be found in the Class Loader documentation section as it covers a lot of other details as well, not just scanning." -> can we link to it as well? * 24.6.2. Database Connection "use of your hardware (Hard disk," -> lowecase "hard disk" "database read's and writes. " -> remove extra apostrophy after read, so that its "reads" in note, change "Postgres" to PostgreSQL (possibly add a link too http://www.postgresql.org/) * 24.6.3. Clustering Tuning "particularly if if" -> remove extra if * 24.6.3.3. JGroups Message Bundling "TCP socket -- small " -> fix double hyphen to a dash	0	0	https://issues.jboss.org/browse/JBPAPP-2635
10303	JBPAPP-2635	Ales, would you be able to write a piece of example code for this section? http://community.jboss.org/wiki/JBoss5custommetadatafiles describes a simple jboss-classloading.xml file: <classloading xmlns="urn:jboss:classloading:1.0" export-all="NON_EMPTY"> <capabilities> <module name="test5"/> </capabilities> <requirements> <module name="test2"/> </requirements> </classloading> I'm not sure this is specific to limiting scanning. Please look at 24.6.1.1.2. "Annotation Scanning Tuning" in http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html-single/Administration_And_Configuration_Guide_CP01/#Memory_Usage_Tuning Ideally your material would be inserted immediately following "Another way to limit scanning is to provide the jboss-classloading.xml file."	0	0	https://issues.jboss.org/browse/JBPAPP-2635
10304	JBPAPP-2624	Howard, how's this going?	0	0	https://issues.jboss.org/browse/JBPAPP-2624
10305	JBPAPP-2623	Howard, are there any instructions for JBoss Messaging 1.4.6 that should be included in this chapter, or can I comment out the entire section about HornetQ until 5.1?	1	1	https://issues.jboss.org/browse/JBPAPP-2623

#	Requirement	Questions	Answered	Successfully Answered	URL
10306	JBPAPP-2622	There are just 2 typos in the chapter 7.3 Transactional Resources on page 73. 1) the first line of the first paragraph: "The transaction manger coordinates.." -> "The transaction manager coordinates.." 2) the sixth line of the second paragraph: "XA transactions, it it also feasible.." -> "XA transactions, it is also feasible.."	1	1	https://issues.jboss.org/browse/JBPAPP-2622
10307	JBPAPP-2621	Hi Isaac, I have found several issues in this chapter (see my comments). Please, could you fix it and reassign this issue back to me? Thank you, Pavel	1	1	https://issues.jboss.org/browse/JBPAPP-2621
10308	JBPAPP-2619	I know JBossPortal does something with MC on top of Tomcat. I hacked something long ago on top of GF. We need to update this diagram 4.1, as it's very out-dated. e.g. container module no longer exists - for almost 2years now; broken into smaller pieces - reflect, mdr, ... Reliance module could be left off - it's still in a sandbox phase. "where <server_configuration> represents the name of the server profile, for example, all, default or minimal" What about web? "profile-repository.xml" - I don't think this name is used anymore. "It's an extension of always looking to the filesystem server/name/conf/jboss-service.xml and server/name/deploy to load deployments." What about deployers/ directory? "JBoss5VirtualFileSystem" - we can drop the JBoss5 prefix. "this bean translates deployer-beans.xml into" - should be "-jboss-beans.xml" suffix "It handles the legacy jboss-service.xml style of mbean deployment descriptors" - it can handle any -service.xml file "More information on the JBoss Microcontainer project can be obtained from http://labs.jboss.com/jbossmc/." - the link should be www.jboss.org/jbossmc Who picked "most important" beans? How should we proceed with fixing/updating this issues?	1	1	https://issues.jboss.org/browse/JBPAPP-2619
10309	JBPAPP-2619	Ales, are you finished with this? If so, please reassign it to me so I can keep track of it. If not, when do you expect to have completed the review? Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-2619
10310	JBPAPP-2617	Chapter 2. JBoss Enterprise Application Platform 5 architecture - The diagram is about the overall jboss portfolio, a super-set of EAP. EAP would really be the Application Server box from the picture. - The directory structure needs checking. OTOH maybe it would be better to discuss the content/purpose of individual dirs because it often changes. E.g. docs/examples/binding-manager - is gone docs/examples/netboot - obsolete server/default/conf/bootstrap - no profile-repository.xml server/default/conf/ - no bootstrap-norepo.xml server/default/xmbean - some descriptors are gone ... Overall Part I - Chapter 2 looks rather poor in content explaining the overall architecture.	1	1	https://issues.jboss.org/browse/JBPAPP-2617
10311	JBPAPP-2617	Dimitris, - Ch. 1 - changes applied. Thanks. :) - Ch. 2 - re. the diagram, would you be able to email me a mockup of a more appropriate diagram so that I can have one made? - re. the directory structure, until we can put together a chapter discussing the purpose of individual directories, can you point me to a more up-to-date directory structure? - Ch. 3 - I assume this means I can remove the bullet point referring to *-jboss-bean.xml? - Ch. 6 Does this mean we can remove Ch. 6 altogether? If not, there is a substantial amount of work required to render that massive diagram into something that will a) fit and b) be easily understood by readers. Please get back to me on this. Thanks, Laura	1	1	https://issues.jboss.org/browse/JBPAPP-2617
10312	JBPAPP-2606	where can I see customer support portal?Is this issue resolved there?	0	0	https://issues.jboss.org/browse/JBPAPP-2606
10313	JBPAPP-2601	Fernando, What do you want me to do with this? Javadocs are no longer stored on redhat.com/docs, and ECS doesn't handle javadocs. Let me know. :)	1	1	https://issues.jboss.org/browse/JBPAPP-2601
10314	JBPAPP-2600	Hey Permaine, please correct me if I'm wrong ;) since QA needs the javadocs for the distribution then once you know which ones you are including, would you be able to attach in zip format, the APIs that need to be included on redhat.com/docs? If the ones online need to be different API docs, can you please explain that a bit to me. I know there was an email about internal and external API docs that I didn't quite understand. Thanks, Isaac	1	1	https://issues.jboss.org/browse/JBPAPP-2600
10315	JBPAPP-2599	Permain, can you take this from here?	1	1	https://issues.jboss.org/browse/JBPAPP-2599
10316	JBPAPP-2595	Hi, is this resolved?	1	1	https://issues.jboss.org/browse/JBPAPP-2595
10317	JBPAPP-2579	I believe the way it is is the right way and we have been using it anyway.	0	0	https://issues.jboss.org/browse/JBPAPP-2579
10318	JBPAPP-2576	Ivo, do we know what the issues are? Is it documented somewhere, in a jira?	0	0	https://issues.jboss.org/browse/JBPAPP-2576
10319	JBPAPP-2571	Hi Isaac, we have found problem in JBoss Messaging on MS SQL Server with MS JDBC drivers version 2.0. But fix of this problem has been postponed (see JBPAPP-2595). We have workaround for this issue by adding parameter to JDBC URL connection string. Problem is related to Adaptive buffering. Adaptive buffering is introduced starting with the Microsoft SQL Server 2005 JDBC Driver version 1.2. But starting with the JDBC driver version 2.0 is this adaptive buffering selected as default behavior for the driver. When is Adaptive buffering active, large values can be read only once by using the get<Type>Stream methods. The second attempt causes exception "The data has been accessed and is not available for this column or parameter". Current JBoss messaging implementation causes this issue. It is necessary to reconfigure JDBC driver to use full buffering behavior. JDBC driver can be configured by parameter in JDBC URL string, please see comments on this issue. See http://msdn.microsoft.com/en-us/library/bb879937.aspx Please, could you add note about this workaround into the EAP5 documentation for MS SQL Server and MS JDBC drivers 2.0? Thank you.	0	0	https://issues.jboss.org/browse/JBPAPP-2571
10320	JBPAPP-2538	Hi David, can you give a status on this one please?	1	1	https://issues.jboss.org/browse/JBPAPP-2538
10321	JBPAPP-2520	Charles, can you assign this to the right person?	0	0	https://issues.jboss.org/browse/JBPAPP-2520
10322	JBPAPP-2488	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2488
10323	JBPAPP-2488	According to https://jira.jboss.org/jira/browse/JBPAPP-2328 RPMs for EAP5 should be available now, but I cannot find them anywhere. Does anyone know where they are?	1	1	https://issues.jboss.org/browse/JBPAPP-2488
10324	JBPAPP-2488	Is the RPM Spec (or maven-rpm-plugin configuration) available in Subversion somewhere?	0	0	https://issues.jboss.org/browse/JBPAPP-2488
10325	JBPAPP-2478	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2478
10326	JBPAPP-2458	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2458
10327	JBPAPP-2457	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	1	1	https://issues.jboss.org/browse/JBPAPP-2457
10328	JBPAPP-2453	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2453
10329	JBPAPP-2447	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2447
10330	JBPAPP-2447	Sean can you take a look at these and see how they turn out translated? Let me know about font type/size.	1	1	https://issues.jboss.org/browse/JBPAPP-2447
10331	JBPAPP-2447	Hi, could someone close this issue? I have no way to check weather those SVG images are in the correct place. Thanks	0	0	https://issues.jboss.org/browse/JBPAPP-2447
10332	JBPAPP-2446	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2446

#	Requirement	Questions	Answered	Successfully Answered	URL
10333	JBPAPP-2445	Isaac, Sorry, after further investigation, the web-console*properties under deploy/management should be used. jboss-as/server/<config>/deploy/management/console-mgr.sar/web-console.war/WEB-INF/classes/web-console-{users,roles}.properties Can we include those instructions? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2445
10334	JBPAPP-2445	Are you sure? :P I'll change it back again now.	0	0	https://issues.jboss.org/browse/JBPAPP-2445
10335	JBPAPP-2445	Sorry if it wasn't clear what I meant. What I see there now is this: (taken from http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0/html/Installation_Guide/Post-Installation.html) "$JBOSS_HOME/server/$CONFIG/deploy/management/console-mgr.sar/web-console.war/WEB-INF/classes/web-console-users.properties, where $JBOSS_HOME is the install directory and the $CONFIG is the server configuration being used. When a server configuration other than those mentioned, is in use, the $JBOSS_HOME/server/$CONFIG/conf/deploy/management/console-mgr.sar/web-console.war/WEB-INF/classes/web-console-{users,roles}.properties file should be modified." What there should be is this: "$JBOSS_HOME/server/$CONFIG/deploy/management/console-mgr.sar/web-console.war/WEB-INF/classes/web-console-users.properties, where $JBOSS_HOME is the install directory and the $CONFIG is the server configuration being used."	1	1	https://issues.jboss.org/browse/JBPAPP-2445
10336	JBPAPP-2441	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2441
10337	JBPAPP-2389	shared-udp --------- - use_incoming_packet_handler="true": remove (only option in 2.8) - use_concurrent_stack="true": remove (default=true) and deprecated in 2.8 - oob_thread_pool.queue_max_size="100": remove udp --- - <pbcast.NAKACK use_mcast_xmit="false": I'd set it to true. There is no additional cost when using UDP as transport, and more than 1 receiver could have lost a message so all of them receive the xmit udp-async --------- - oob_thread_pool.keep_alive_time="500": is a bit short, I'd set it to 1000 - oob_thread_pool.queue_max_size="100": remove as the queue is disabled tcp --- - Why is 'loopback' false here and true for the UDP based stacks ? - tcp_nodelay is true but enable_bundling is true, too: this is inconsistent - use_send_queues: I found this valuable to prevent threads getting blocked in the TCP socket write phase, and therefore blocking other threads too. If you enable this, also set send_queue_size General ------- - RPCs which run on a transport with message bunling enabled can override bundling on a per message basis by setting the OOB flag. Maybe this should be added as a comment to the "udp/tcp" stack - We use STATE_TRANSFER instead of STREAMING_STATE_TRANSFER, should we change this ? TDB - TBD: FLUSH: really needed for all stacks ? - There's a lot of repetition in this XML file, maybe we should look into using a templating mechanism and only define the diffs from a base stack (not just for shared-udp) - If bundling is disabled, I'd remove max_bundle_size="64000" and max_bundle_timeout="30"	1	1	https://issues.jboss.org/browse/JBPAPP-2389
10338	JBPAPP-2370	Hi Sean Which repo could we put those .PO files in? Currently all docs are in https://svn.jboss.org/repos/jbossas/projects/docs/enterprise/5.0. Cheers, Simon	1	1	https://issues.jboss.org/browse/JBPAPP-2370
10339	JBPAPP-2370	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2370
10340	JBPAPP-2370	Hi, which language/font pack I have to install on my computer in order to have correctly displayed chinese and japanese characters? I am using Fedora 11 and have installed several language/font packs, but chinese and japanese characters are still not correctly dispalyed.	1	0	https://issues.jboss.org/browse/JBPAPP-2370
10341	JBPAPP-2365	Is https://jira.jboss.org/jira/browse/JBAS-6722 marked private? I cannot read it.	1	1	https://issues.jboss.org/browse/JBPAPP-2365
10342	JBPAPP-2365	What is the impact of removing the two lines? I can see that the two lines are present in current EAP5 bits: <!-- This just checks metadata locations --> <bean name="MetaDataStructureModificationChecker" class="org.jboss.deployers.vfs.spi.structure.modified.MetaDataStructureModificationChecker"> <constructor> <parameter><inject bean="MainDeployer" /></parameter> </constructor> <property name="cache"><inject bean="StructureModCache" /></property> <property name="filter"><bean class="org.jboss.system.server.profile.basic.XmlIncludeVirtualFileFilter" /></property> </bean> <!-- We just visit wars, but exclude classes, lib, dev dirs --> <bean name="WebVisitorAttributes" class="org.jboss.system.server.profile.basic.IncludeExcludeVisitorAttributes"> <constructor> <parameter> <set> <value>.war</value> </set> </parameter> <parameter> <set> <value>WEB-INF/classes</value> <value>WEB-INF/lib</value> <value>WEB-INF/dev</value> </set> </parameter> </constructor> </bean>	1	1	https://issues.jboss.org/browse/JBPAPP-2365
10343	JBPAPP-2365	Is there still time to instead of just deleting those two lines, implement proper filter? I can do that tomorrow, first thing in the morning.	1	1	https://issues.jboss.org/browse/JBPAPP-2365
10344	JBPAPP-2365	Ales, what will the proper filter do? Which components will the filter impact (JBossWeb + Seam or only Seam?)	1	1	https://issues.jboss.org/browse/JBPAPP-2365
10345	JBPAPP-2363	The attached file should be all you need, but I'm not sure about the styling of the header. If it doesn't work, can you grab the XML of a similar page? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-2363
10346	JBPAPP-2363	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2363
10347	JBPAPP-2363	Is there a task for someone to create localized release notes too? The English release notes are up at http://www.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/5.0.0.BETA/html-single/Release_Notes/index.html , but the localized links in the readmes are still giving 404s.	1	0	https://issues.jboss.org/browse/JBPAPP-2363
10348	JBPAPP-2348	Martin, Have these changes been applied to the current source zips of CR1 by some means? I'm still getting security related failures in jca tests (see http://hudson.qa.jboss.com/hudson/view/EAP5/job/eap5-as-solaris9-jdk6/16/ and the test case org.jboss.test.jca.test.LocalTransactionTidyupUnitTestCase).	1	1	https://issues.jboss.org/browse/JBPAPP-2348
10349	JBPAPP-2348	Is there any plan for making these changes available to the Hudson builds? I'm getting hundreds of failures due to this issue at present and there really isn't much point in me running testsuites or looking at testsuite results until it is fixed. Am I to prepare my own custom builds to work around this issue? Is there a plan to fix the CR1 build? Information would be appreciated.	0	0	https://issues.jboss.org/browse/JBPAPP-2348

#	Requirement	Questions	Answered	Successfully Answered	URL
10350	JBPAPP-2348	Martin I've had a look at this issue last night and I think that the way to solve it is sketched out below. It's based on noting that the only classpath which needs to be changed is tests.classpath which provides an execution classpath for junit, which is where the unsigned matching jars are getting picked up. In theory, it should be possible to simply redefine this classpath to pick up the right jars (those from the distribution), leaving everything else the same. For test case invocations which make use of run-junit, and so make use of the classpath, <classpath> <pathelement location="${build.classes}"/> <pathelement location="${build.resources}"/> <path refid="tests.classpath"/> </classpath> we should be OK. But additionally, there are a number of ant targets which start junit executions not using run-junit, and these define their own classpaths which point back into the build. These targets need to be updated to point to the distribution jars too. Basic recipe: (i) leave all the .ent files as they are - they are needed for the build and the tests.classpath can be constructed from scratch without them. (ii) replace the classpath tests.classpath <path id="tests.classpath"> <fileset dir="${jboss.dist}/client"> <include name="**/*.jar"/> <exclude name="jaxws-rt.jar"/> <exclude name="jaxws-tools.jar"/> </fileset> <pathelement path="${classpath}"/> <pathelement path="${local.classpath}"/> <pathelement path="${project.tools}/lib/ant.jar"/> <pathelement path="${project.tools}/lib/ant-junit.jar"/> <path refid="jboss.ejb3.classpath"/> <path refid="jboss.jmx.classpath"/> <path refid="jboss.jmx-remoting.classpath"/> <path refid="jboss.test.classpath"/> <path refid="thirdparty.classpath"/> </path> with <path id="tests.classpath"> <fileset dir="${jboss.dist}/client"> <include name="**/*.jar"/> <exclude name="jaxws-rt.jar"/> <exclude name="jaxws-tools.jar"/> </fileset> <pathelement path="${classpath}"/> <pathelement path="${local.classpath}"/> <pathelement path="${project.tools}/lib/ant.jar"/> <pathelement path="${project.tools}/lib/ant-junit.jar"/> <path refid="jboss.dist.modules.classpath"/> <path refid="jboss.dist.thirdparty.classpath"/> </path> where jboss.dist.thirdparty.classpath and jboss.dist.modules.classpath are defined to include all the modules jars and thirdparty jars from the distribution. These may not be the right definitions, and we may want to use jboss.dist.lib.classpth, jboss.dist.common.classpath, etc. and combine them as required. (iii) fix a number of test execution macros in the testsuite (e.g. tests-compatability) which add in pointers to the build jars (e.g. varia) in their classpath. These pointers need to be updated to point to the distribution. All this stuff should be turned on if the user passes a non-null value for jboss.dist. What do you think? Can you see any problems, errors?	0	0	https://issues.jboss.org/browse/JBPAPP-2348
10351	JBPAPP-2348	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2348
10352	JBPAPP-2348	If they are only under thirdparty and not in the distribution, where is the conflict from signed vs unsigned coming from?	1	1	https://issues.jboss.org/browse/JBPAPP-2348
10353	JBPAPP-2346	Ivo, can you link the JBTM bugs that were fixed in the CP branch to this JIRA?	0	0	https://issues.jboss.org/browse/JBPAPP-2346
10354	JBPAPP-2328	Are these RPMs available to license holders? I cannot find them.	1	1	https://issues.jboss.org/browse/JBPAPP-2328
10355	JBPAPP-2325	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	1	0	https://issues.jboss.org/browse/JBPAPP-2325
10356	JBPAPP-2322	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2322
10357	JBPAPP-2322	Shelly, where can I verify this? Should this be included in the Migrating to Enterprise Application Platform 5.0 Chapter of Installation Guide? Or this will be part of the new Migration Guide (JBPAPP-2511) which is not ready yet? I suppose the latter. In that case I will have to wait for the guide.	0	0	https://issues.jboss.org/browse/JBPAPP-2322
10358	JBPAPP-2316	Hi Strong. If this is being made available in 5.1.0 only, should we remove the 4.x tags from the affects version. And did this issue miss the 5.1.0 release? There doesn't seem to be evidence of a release note for 5.1.0, even if it is just a known issue. If this doesn't require a release note, could you please set the Release Note Docs Status flag to "Not Required". Cheers Jared	1	0	https://issues.jboss.org/browse/JBPAPP-2316
10359	JBPAPP-2315	This fix broke 3.4.2.GA mapping to Oracle. See here: https://forum.hibernate.org/viewtopic.php?f=1&t=1014208 What is the workaround for this break in compatibility?	0	0	https://issues.jboss.org/browse/JBPAPP-2315
10360	JBPAPP-2306	Hi Stelios, I've attached an updated jboss_init_hpux.sh, I don't have a HP machine around, do you mind trying it out before I commit this back to the branch? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2306
10361	JBPAPP-2306	Martin, can you verify the patch attached to the JIRA?	0	0	https://issues.jboss.org/browse/JBPAPP-2306
10362	JBPAPP-2306	One other question: why didn't you take the approach for passing options to org.jboss.Main as was done in the script jboss_init_solaris.sh (i.e. allowing a variable JBOSS_OPTIONS?). These "init" scripts are starting to have more differences than similarities - they should work in the same way from a users point of view, as they are basically all doing the same thing.	1	1	https://issues.jboss.org/browse/JBPAPP-2306
10363	JBPAPP-2306	Richard, you're right, I forgot to add some information about the new parameter to the header for EAP 4.x - I did it for EAP 5. I'll add it there. But otherwise I verified and committed the patch as requested. I just tested it again and it works for me. I didn't know about JBPAPP-3036, but I cannot reproduce it anyway. And I don't think I'm using ksh. What are you referring to when you talk about passing options to org.jboss.Main? I don't see anything like that in the jboss_init_solaris.sh script. This how the AS is started there: CMD_START="${JBOSS_HOME[i]}/bin/run.sh -c ${SERVER_NAME[i]}" nohup su - ${JBOSS_USER[i]} -c "sh $CMD_START" >${JBOSS_CONSOLE[i]} 2>&1 &	1	0	https://issues.jboss.org/browse/JBPAPP-2306
10364	JBPAPP-2306	1. Shouldn't this thing needs to be tested under root, as the instructions in the header file explicitly state that. 2 Which shell is this command *supposed* to use? In the script, it is marked as #!/bin/sh, which defaults on HPUX their implementation of the POSIX shell. root on HPUX in our lab has bash defined as its shell in the password file. 3. I have JAVA_HOME defined and exported in the .bash_profile for root, as well as the .profile 4. If I set the shebang (shell indicator) in the file as #!/bin/sh or #!/usr/bin/bash or #!/usr/bin/ksh, I get the same results: ++ isUserOK ++ id root ++ ((0!=0)) ++ su root -c 'test -w /tmp/jboss-eap-4.2/jboss-as/server/default' ++ ((0!=0)) +++ whence java /sbin/init.d/jboss: line 139: whence: command not found +++ [[-d /opt/java1.5]] ++ su - root -c ++ ((2!=0)) ++ echo 'ERROR: The user "root" does not have $JAVA_HOME defined; either define $JAVA_HOME for this user or ensure that a path to "java" exists.' ERROR: The user "root" does not have $JAVA_HOME defined; either define $JAVA_HOME for this user or ensure that a path to "java" exists. But this is a non-POSIX shell complaining, as when I start /usr/bin/ksh by hand and execute whence java, it doesn't complain. I really don't understand what whence does - from the manual: whence -pv name: For each name, indicate how it would be interpreted as a command name. The -v produces a more detailed report. 5. If I replace su - ${JBOSS_USER[i]} -c `whence java | | [[ -d $JAVA_HOME ]]` >/dev/null 2>&1 with su - ${JBOSS_USER[i]} -c "test -d $JAVA_HOME " >/dev/null 2>&1 it works fine. Don't know if this is equivalent.	1	1	https://issues.jboss.org/browse/JBPAPP-2306

#	Requirement	Questions	Answered	Successfully Answered	URL
10365	JBPAPP-2306	I'm still having problems with this script. I think you need to get someone else to test this. Here is why: 1. Here is my /etc/rc.config.d file: -bash-3.1# more /etc/rc.config.d/jboss JBOSS_HOME[0]=/tmp/jboss-eap-4.2/jboss-as JBOSS_USER[0]=rachmatowicz SERVER_NAME[0]=default JBOSS_START[0]=1 JBOSS_BIND_ADDR[0]=10.16.94.186 and here is the AS distribution -bash-3.1# ls -la /tmp/jboss-eap-4.2/ total 48 drwxr-xr-x 4 rachmatowicrachmatowic 96 Nov 12 14:25 . drwxrwxrwt 62 root root 8192 Nov 20 11:09 .. drwxr-xr-x 7 rachmatowicrachmatowic 8192 Nov 12 01:02 jboss-as drwxr-xr-x 10 rachmatowicrachmatowic 8192 Nov 12 05:35 seam and here is the user using the system: -bash-3.1$ who rachmato pts/0 Nov 20 10:52 Would you like to tell me what is my login id on this system? I login as rachmatowicz: alias dev24='ssh -l rachmatowicz dev24.qa.atl2.redhat.com' I remember the sysadmin guys telling me that my username was too long and required special treatment. Perhaps this is part of the problems i'm seeing? If I can be given another non-root userid to use, i'll gladly continue the testing... I'm still getting the same errors I saw yesterday.	1	1	https://issues.jboss.org/browse/JBPAPP-2306
10366	JBPAPP-2306	Richard, I managed to make it work under your account. What I needed to do is change your shell in /etc/passwd from /bin/bash to /bin/sh. Now the question is: Why doesn't it work with bash? Anyway, feel free to use my account for testing if you need it.	0	0	https://issues.jboss.org/browse/JBPAPP-2306
10367	JBPAPP-2306	Very odd stuff going on here, and i'm not sure if it is an HPUX issue or an issue with our testing environment (i.e. the use of LDAP and other login/validation/password file issies). With a use with a /usr/bin/sh shell (HPUX POSIX shell), the script works. Wiith bash, odd things happen. For example, pskopek has a /etc/passwd entry indicating /bin/bash as a default shell. 1. If I execute (i.e. start a non-login shell under pskopek and execute the env command): # su pskopek # env I see that .bashrc is called as expected 2. If I execute (i.e. start a login shell under pskopek and execute the env command separately): # su - pskopek # env I see that .bash_profile and .bashrc are called as expected 3. If I execute (i.e. start a login shell under pskopek and execute the env command): # su - pskopek -c 'env' I see that neither .bashrc nor .bash_profile is called This latter case is why the script is not working with a user who has a bash shell. However: 4. if I execite (i.e. start a interactive login shell under pskopek and execute the env command): # su - pskopek -i -c 'env' I see that both .bashrc and .bash_profile are called Adding the extra -i to make the command interafctive makes the init script work with my userid. I haven't quite been able to explain what is going on here, nor have I tested it with a user who has a POSIX default shell.	0	0	https://issues.jboss.org/browse/JBPAPP-2306
10368	JBPAPP-2298	Martha, Can you please build a 1.1.0.jboss.patch01-brew with this patch, upload and switch the EAP 4 branch to use this? The branch to use is jakarta-commons-logging-jboss/RHEL-4-EP-V1. Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-2298
10369	JBPAPP-2270	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2270
10370	JBPAPP-2264	Martha, can you please add the patch, build, upload and then switch the EAP 4 branch to use this? The branch to use for this is: jacorb/RHEL-4-EP-V1 Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2264
10371	JBPAPP-2263	OK, I will clone this JIRA for the EAP 5. With regards to the upstream case, please add the link in a comment when you have one. Can you please add the IT link to the "Support Case Reference" box?	0	0	https://issues.jboss.org/browse/JBPAPP-2263
10372	JBPAPP-2186	Alessio, do you know what has to be done? Will this be OK for SOA-P?	0	0	https://issues.jboss.org/browse/JBPAPP-2186
10373	JBPAPP-2182	Hi Vicky, will the tag for JBoss Common be available soon?	1	1	https://issues.jboss.org/browse/JBPAPP-2182
10374	JBPAPP-2182	Martha, Dimitris has created the tag for jboss-common this for us: https://svn.jboss.org/repos/common/common-old/tags/JBossCommon_1_2_1_GA_CP01/ Can you please build, upload, then switch the EAP 4 branch to use this? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2182
10375	JBPAPP-2181	I will update the online Installation Guide for 4.2.0.GA_CP07 and 4.3.0.GA_CP05 but I will leave this issue open and set the fix to 4.2.0.GA_CP08 and 4.3.0.GA_CP06 so it can be recorded in those release notes for when this fix is included in the documentation bundle that is shipped with the software.	0	0	https://issues.jboss.org/browse/JBPAPP-2181
10376	JBPAPP-2160	Martha, Can you please work on this build? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-2160
10377	JBPAPP-2144	Hibernate integration code for non-supported second-level caches and connection pools will continue to be included in EAP 5 because we didn't deprecate it in EAP 4. How far in advance does functionality need to be deprecated before being removed? Would it need to be deprecated in EAP 5.0.GA to remove it in 5.1? Does it make sense to deprecate all of the integration code for non-supported components?	1	1	https://issues.jboss.org/browse/JBPAPP-2144
10378	JBPAPP-2144	These components have not been marked "@deprecated" in the code, so it will not be documented as deprecated in Javadoc. Is it sufficient, for now, to simply document that they are deprecated in the release notes? Also, as Steve mentioned in a discussion at https://docspace.corp.redhat.com/clearspace/message/12692#12692 , we may want to deprecate usage of a component specifically in EAP, but perhaps not in a standalone release. This would complicate how we mark source code as @deprecated.	1	1	https://issues.jboss.org/browse/JBPAPP-2144
10379	JBPAPP-2144	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	1	0	https://issues.jboss.org/browse/JBPAPP-2144
10380	JBPAPP-2142	I've always seen that to disable chunking encoding is always related to client configuration. How can it be disabled from server side? What configuration file can be changed in JBoss or JBossWS to do that?	1	0	https://issues.jboss.org/browse/JBPAPP-2142
10381	JBPAPP-2142	I'd like to but https://community.jboss.org/nfm/redirectToCSForum.jspa?module=bb&op=viewforum&f=247 and https://community.jboss.org/nfm/redirectToCSForum.jspa?module=bb&op=viewforum&f=242 are out of service. Do you know if is there another link to forum group? Thanks	1	0	https://issues.jboss.org/browse/JBPAPP-2142
10382	JBPAPP-2134	Shelly, have you talked to Jason about this already? Have you switched the Branch_5_0 to use these already?	1	1	https://issues.jboss.org/browse/JBPAPP-2134
10383	JBPAPP-2130	This fix has already been committed to the 2.4.x branch by Vladimir.	0	0	https://issues.jboss.org/browse/JBPAPP-2130
10384	JBPAPP-2129	Do you need me to find out, Permaine, or were you asking Fernando?	0	0	https://issues.jboss.org/browse/JBPAPP-2129
10385	JBPAPP-2075	Isaac, can you document the above?	1	1	https://issues.jboss.org/browse/JBPAPP-2075
10386	JBPAPP-2059	Ben, can you take care of this for me? Just go through the JBoss Messaging 1.4.3 documentation and remove/comment out information relating to 4.2 and 4.3, as above. :)	1	1	https://issues.jboss.org/browse/JBPAPP-2059

#	Requirement	Questions	Answered	Successfully Answered	URL
10387	JBPAPP-2059	Completed for JBossMessaging_1_4_0_SP3_CP10/ (EAP 4.3.8) except for the following, which should be added to the JBM documentation (or the Getting Started Guide) if required: <section id="properties"> <title>System Properties used by JBoss Messaging</title> <section id="properties.supportBytesId"> <title>support.bytesId</title> <para>This system property controls the default behavior when constructing a JBossMessage object from a foreign message object. If set to true, the JBossMessage constructor will try to extract the native byte[] correlation ID from the foreign message headers. If set to false, it will use the normal string type JMSCorrelationID. If this system property is absent or is given some value other than 'true' and 'false', it will defaults to 'true'.</para> </section> <section id="properties.retainOldxabehaviour"> <title>retain.oldxabehaviour</title> <para>This system property controls what kind of exception a JMS XAResource throws when the prepare is called after the connection is broken. If this property is not defined, an XAException with XA_RBCOMMFAIL error code will be thrown. If this propertie is defined, an XAException with XA_RETRY error code will be thrown instead. JBM by default doesn't define this property.</para> </section> <section id="properties.mapmessagesupportnullobject"> <title>mapmessage.support.null.object</title> <para>This system property, once defined, allows null values in MapMessage.setObject(String key, Object value). JBM by default doesn't define this property, which means passing null values to the setObject() method will cause a MessageFormatException to be thrown.</para> </section> </section> Howard, can you tell me whether this still applies to the product disributed with EAP and where the best place for this information would be? The installation chapter of the JBM guide is removed from the guide shipped with EAP.	1	1	https://issues.jboss.org/browse/JBPAPP-2059
10388	JBPAPP-2046	Misty, do you agree with the assessment in the JIRA? Can you give us advice on how to proceed for EAP 4.3.9?	1	1	https://issues.jboss.org/browse/JBPAPP-2046
10389	JBPAPP-2046	I cannot see the fix on the stage. Could you point me where it is, please? http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/4.3/html-single/Server_Configuration_Guide/index.html	0	0	https://issues.jboss.org/browse/JBPAPP-2046
10390	JBPAPP-2043	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-2043
10391	JBPAPP-2016	Phillip, is this resolved for this CP? If not, please move this to the next CP.	0	0	https://issues.jboss.org/browse/JBPAPP-2016
10392	JBPAPP-2002	JGroups-2.4.6.GA.src/build.properties has: # use true for IPv6 and false for IPv4 jgroups.useIPv6=false Is this the way we should build?	0	0	https://issues.jboss.org/browse/JBPAPP-2002
10393	JBPAPP-1996	Release Notes - JBoss Remoting - Version 2.2.3 - Text format Bug * [JBREM-1081] - Fix NPE in ServerInvokerCallbackHandler * [JBREM-1088] - MicroSocketClientInvoker(InvokerLocator locator, Map configuration) ctor - not propagating exceptions (only message is wrapped) * [JBREM-1099] - Make MulticastDetector detection message send buffer size configurable * [JBREM-1109] - Eliminate race in MicroRemoteClientInvoker.getDataType() * [JBREM-1111] - CLONE [JBREM-851] - In LeasePinger replace Timer if it has shut down * [JBREM-1112] - Potential race between ConnectionValidator and ConnectionListener upon connection failure * [JBREM-1119] - CLONE [JBREM-1113] - ServerInvokerCallbackHandlers leak when client doesn't shut down * [JBREM-1121] - Client SocketFactory should be configurable by InvokerLocator * [JBREM-1125] - Test for IllegalStateException when calling Timer.schedule() * [JBREM-1127] - ClassCastExceptions arising from Unmarsharller/Classloader caching in MicroRemoteClientInvoker * [JBREM-1129] - Eliminate nondeterminism in Lease updates Feature Request * [JBREM-1082] - Allow ConnectionValidator to access InvokerLocator parameters * [JBREM-1084] - Allow CallbackPoller to access Client and InvokerLocator parameters * [JBREM-1100] - Link ServerInvokerServlet instances to Connectors via MBean names rather than locator URLs * [JBREM-1102] - Make configuration map available to MarshalFactory * [JBREM-1128] - Introduce connection identity concept Release * [JBREM-1131] - Release 2.2.3 Task * [JBREM-139] - need automated test for servlet server invoker * [JBREM-1078] - Remove reference to "numberOfRetries" parameter from Remoting Guide * [JBREM-1079] - Port fixes to ServletServerInvoker from branch 2.x to branch 2.2 * [JBREM-1085] - Reduce log level of ServerSocketWrapper.close() log messages * [JBREM-1103] - Correct javadoc for Client.invokeOneway() * [JBREM-1104] - Identity.get() should create a more meaningful RuntimeException message * [JBREM-1108] - Warn against making ConnectionValidator.validatorPingPeriod shorter than ConnectionValidator.validatorPingTimeout * [JBREM-1110] - InvokerLocator.getParameters() should not return null * [JBREM-1130] - Assure version compatibility with earlier versions of Remoting	0	0	https://issues.jboss.org/browse/JBPAPP-1996
10394	JBPAPP-1996	Martha, Can you please build this in brew? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBPAPP-1996
10395	JBPAPP-1971	jbossws-jboss42 is the container integration artifact. Afaics there has been no changes there for 1.2.1.GA_CP05, so I did not cut any new release for that. Do you need a new release nevertheless?	1	1	https://issues.jboss.org/browse/JBPAPP-1971
10396	JBPAPP-1949	Hi Fernando, The related JBMESSAGING-1608 has been fixed into JBM 1.4.0.SP3-CP08. It should go in to EAP 4.3.0 CP05. What do you need me to do for the Porting? or shall I close this jira?	1	1	https://issues.jboss.org/browse/JBPAPP-1949
10397	JBPAPP-1948	So the RPM layout doesn't matter at all for the zip? For example, helpinfo is part of the main RPM and not a doc (-javadoc, -manual) RPM. Do you want the apidocs and helpinfo separated in the docs zip? Right now, everything from each package is being copied to the same directory.	1	1	https://issues.jboss.org/browse/JBPAPP-1948
10398	JBPAPP-1946	Is this included in CR3 build of WFK? How is this done? Reopening this issue just to keep track.	0	0	https://issues.jboss.org/browse/JBPAPP-1946
10399	JBPAPP-1935	Great. Not sure to understand what the technical preview contains: can this be proxified for example by our internal nexus instance, or does it contain only a subset of what the downloadable version offers? Thanks.	1	0	https://issues.jboss.org/browse/JBPAPP-1935
10400	JBPAPP-1898	Why does Richfaces need jhighlight? I can't find it being used in the community release	0	0	https://issues.jboss.org/browse/JBPAPP-1898
10401	JBPAPP-1898	Jay, can you please comment on why this is included for QE?	0	0	https://issues.jboss.org/browse/JBPAPP-1898
10402	JBPAPP-1898	I think he wants to know what jhighlight is for---why is it added to WFK, but not in the community (upstream) release? It is included in the list of jars in JBPAPP-1946.	1	1	https://issues.jboss.org/browse/JBPAPP-1898
10403	JBPAPP-1893	Jay, can you please read over the attached file and check that the information is correct, or assign it to someone who can? Thankyou. :) I think the only thing that needs to be added to this chapter is the note mentioned by Peter: When an application is deployed as .war file, it is expanded in a folder under ${jboss.server.home.dir}/tmp/deploy directory name as "tmp-exp.war". To configure a filter to archive app1.war, the correct jboss-log4j.xml should be : <param name="ConversionPattern" value="%d %-5p %c{1} %m%n"/> Please confirm that this is correct.	0	0	https://issues.jboss.org/browse/JBPAPP-1893
10404	JBPAPP-1877	The native libraries (at least as I defined them) were removed in CR3. But you are correct that there are tarballs of the mozilla libraries, as well as libraries for webkit (mac) in gwt/tools/redist. In that case, gwt/tools/sdk should be removed as well as it also contains mozilla- and webkit-related files. Aside from this, there are many jars here that come from upstream. Is this also a problem?	1	1	https://issues.jboss.org/browse/JBPAPP-1877
10405	JBPAPP-1876	Sorry Shaun, would you be able to clarify which book this is in? There are quite a few and it will save time if you can specify. :)	1	1	https://issues.jboss.org/browse/JBPAPP-1876

#	Requirement	Questions	Answered	Successfully Answered	URL
10406	JBPAPP-1873	Hey Mike, having a look at this issue in the JBoss Web project and even in the svn commits, it doesn't give any details on exactly how this issue was fixed but just that it was. Would you be able to give any details on the percise details of how this bug was fixed so I can document it correctly in the Release Notes? Thanks, Isaac	1	1	https://issues.jboss.org/browse/JBPAPP-1873
10407	JBPAPP-1868	Jared, can you get back to Anne-Louise about whether this is doable for 4.3.9?	0	0	https://issues.jboss.org/browse/JBPAPP-1868
10408	JBPAPP-1868	Hi Jared, there is comment "Implemented Eva's No. 2 Suggestion." but I am not able to find it at http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/4.3/html-single/Server_Configuration_Guide/index.html#clustering-session-sfsb30. Could you point me to the updated part, please?	1	1	https://issues.jboss.org/browse/JBPAPP-1868
10409	JBPAPP-1864	Hi Jason, are we really supposed to do that? Why? Is that someone who can help with the change in the JBPAPP_5_0 branch? Thansks in advance. --Fernando	1	1	https://issues.jboss.org/browse/JBPAPP-1864
10410	JBPAPP-1819	Can online docs for CP06 be updated?	0	0	https://issues.jboss.org/browse/JBPAPP-1819
10411	JBPAPP-1819	This change has been added to the EAP 4.2.CP08 and 4.3.CP07 release notes. Rajesh, please verify on http://documentation-stage.bne.redhat.com/docs/en-US/ and close this issue if you believe the information is sufficient.	1	1	https://issues.jboss.org/browse/JBPAPP-1819
10412	JBPAPP-1749	Darran, could you please paste here your build order again?	1	1	https://issues.jboss.org/browse/JBPAPP-1749
10413	JBPAPP-1749	One more thing Alessio, the jbossws-spi build uses ant <get> to grab thirdparty/jboss-logging-log4j.jar but does not add it to the classpath (see ant-import/build-thirdparty.xml). Is this JAR really needed for something in this context? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-1749
10414	JBPAPP-1740	I think a one-off patch is going to be necessary for this issue. Beyond saying that here, what else does support need to do in order to make this happen?	0	0	https://issues.jboss.org/browse/JBPAPP-1740
10415	JBPAPP-1733	Gail, is it a one-time test, or should I prepare some scripts for an automated test?	0	0	https://issues.jboss.org/browse/JBPAPP-1733
10416	JBPAPP-1733	Emmanuel, if annotations and entity manager are upgraded to 3.3.x in EAP 4.2/4.3, will the non-FP version of seam work without validator? Is validator only a build dependency?	0	0	https://issues.jboss.org/browse/JBPAPP-1733
10417	JBPAPP-1731	I've found the email thread discussing this. The issues mentioned there are: http://opensource.atlassian.com/projects/hibernate/browse/ANN-551 http://opensource.atlassian.com/projects/hibernate/browse/ANN-683 ANN-551 was fixed in 4.2.0.GA_CP01 by a patch that was applied when building with brew (https://jira.jboss.org/jira/browse/JBPAPP-313) It was backported to annotations v3_2_1_GA_CP and fixed (w/o the patch) starting with 4.2.0.GA_CP02. ANN-653 was not backported to v3_2_1_GA_CP. There is no unit test for this fix and I'm not familiar enough with the code that was fixed to document the impact of fixing this bug. Emmnauel, will fixing ANN-653 cause a change in behavior that could be regarded as a backward compatibility issue? If so, please describe the change in behavior.	1	1	https://issues.jboss.org/browse/JBPAPP-1731
10418	JBPAPP-1731	Andrig Miller wrote: > Gail, > > Emmanuel told us about this issue during the process on deciding what to do in the FP. It was my considered opinion then, and it still is now, that this sort order difference is fine for a CP and the FP release. I would have never given the go ahead for the FP to begin with, if I didn't think that was true. Yes, the behavior is different, but the new behavior is correct, so I think of it as a bug. > > We should close this issue, and update the components in the base, with a release note on it saying that the fix for the original bug (whatever that was), changes the sort order. > > Andy > > ----- "Gail Badner" <gbadner@redhat.com> wrote: > > There is a backward-compatibility issue upgrading Hibernate annotations > > in 4.2.0.GA_CP07 / 4.3.0.GA_CP05 to be the same version used in > > 4.3.0.GA_CP05_FP01. Details are documented in > > https://jira.jboss.org/jira/browse/JBPAPP-1731. > > > > I'm sorry this issue is coming up so close to the release. I found out > > late last week that it is not acceptable to require applications to be > > re-compiled when upgrading to a new CP. > > > > Fernando suggested adding a config parameter and making the default > > behavior be the same as in the previous 4.2/4.3 CPs. If the same build > > is used for 4.3.0.GA_CP05_FP01, then using that same default behavior > > will cause a backwards-compatibility issue for the FP. > > > > IMHO, it makes more sense to make the default be the 3.3.x behavior and > > deal with the change in 4.2.0.GA_CP07 / 4.3.0.GA_CP05, than make the > > default be the 3.2.x behavior and deal with the change in > > 4.3.0.GA_CP05_FP01. > > > > Here are some alternatives for dealing with the change in 4.2.0.GA_CP07 > > / 4.3.0.GA_CP05: > > 1) In the release notes, tell customers that if they are affected they > > should do one of the following: > > - add hibernate.foreignkey.hashcode_from_counter=false to > > persistence.xml or hibernate.cfg.xml > > - startup JBoss using "run.sh > > -Dhibernate.foreignkey.hashcode_from_counter=false" > > 2) Change run.sh to pass -Dhibernate.foreignkey.hashcode_from_counter=false > > > > Option 2) will not work if 4.3.0.GA_CP05 and 4.3.0.GA_CP05_FP01 use the > > same run.sh. > > > > Opinions? Other ideas? > > > > Thanks, > > Gail > > > > > > -- > Andrig (Andy) Miller > VP of Engineering for JBoss Middle-ware > Red Hat, Inc.	1	0	https://issues.jboss.org/browse/JBPAPP-1731
10419	JBPAPP-1731	Rich Sharples wrote: > Well I don't take the support calls - but if everyone thinks this has limited scope to impact customer's existing apps. I'm OK - is there a way we can pre-empt the problem and actively tell people about the --preserve-existing-but-wrong-behaviour flag ? >> Because we're trying not to have different behavior (i.e. add a parameter in run.conf) for CP and CP_FP...and also the fact that the broken behavior *is* broken... :) >> >> -CB >> >> Rich Sharples wrote: >>> If we think the impact is low then it's academic - but why wouldn't we make the broken but existing behaviour the default. >>>> Rich Sharples wrote: >>>> >>>>> The question is - how many existing applications will need to be recompiled / redeployed - if it's 0.1% that's OK - if it 50% then it's not OK. >>>> >>>> We don't know for sure. It seems like a very small number, if any. Also, I believe Gail implied that she was going to provide a switch to fall back to the broken (non-FP) behavior, but not make it the default. >>>> >>>> Are you cool with that? >>>> >>>> -CB	1	0	https://issues.jboss.org/browse/JBPAPP-1731
10420	JBPAPP-1729	I attached the JDiff reports for annotations and entitymanager. Is there a list somewhere on what we consider as public API? If so, I will cross the list with this report, to make sure we have no changes on it.	1	1	https://issues.jboss.org/browse/JBPAPP-1729
10421	JBPAPP-1729	For annotations hey Juca, Can you run a JDiff between ANN 3_2_1_GA_CP03 and Hibernate Validator (the one we include). Cause some of those APIs are public. The rest is fine for annotations	1	1	https://issues.jboss.org/browse/JBPAPP-1729
10422	JBPAPP-1729	Hardy/Juca, what's the final word on this? Are there any backward-compatibility issues that require that we stay with 3.2.x in EAP 4.2/4.3 CPs?	1	1	https://issues.jboss.org/browse/JBPAPP-1729
10423	JBPAPP-1728	The only fixes for annotations and entity manager that I see are for unit tests. Is it worth upgrading annotations and entity manager to include the unit test fixes, since they are going to be upgraded in the next CP anyhow?	0	0	https://issues.jboss.org/browse/JBPAPP-1728

#	Requirement	Questions	Answered	Successfully Answered	URL
10424	JBPAPP-1701	When I build Javadoc outside of brew, the HREFs are there. When I build with brew, they are gone. I'm not sure why this is happening, but I noticed the same thing happening for jboss-cache and jboss-messaging: - http://porkchop.devel.redhat.com/brewroot/packages/jboss-cache/1.4.1/6.SP13.1.ep1.el4/noarch/jboss-cache-javadoc-1.4.1-6.SP13.1.ep1.el4.noarch.rpm - http://porkchop.devel.redhat.com/brewroot/packages/jboss-messaging/1.4.0/2.SP3_CP08.1.ep1.el4/noarch/jboss-messaging-javadoc-1.4.0-2.SP3_CP08.1.ep1.el4.noarch.rpm Is this a problem with brew or the spec files?	1	1	https://issues.jboss.org/browse/JBPAPP-1701
10425	JBPAPP-1701	The problem is that the brew builds are done offline. When ant builds javadoc offline, it needs the "linkoffline" attribute specified with a location for the Java package list. Details about the ant javadoc task are at: http://ant.apache.org/manual/CoreTasks/javadoc.html Is the Java API package list already in the brew environment? If so, where will it be? If not, is there another package that should be added as a build requirement that has Java API package list? This is a minor issue and will be put off until the next CP.	1	1	https://issues.jboss.org/browse/JBPAPP-1701
10426	JBPAPP-1701	response from scm-tool mail list ---------- I checked with the person who packaged java-1.5.0-sun and she indicated that the javadocs (and hence the package-lists file) was not also packaged. However, you should be able to simply use a copy of the file http://java.sun.com/j2se/1.5.0/docs/api/package-list for instance, and instruct javadoc that this is the package list it should link to http://java.sun.com/j2se/1.5.0/docs/api/ - Chris On 08/05/2009 05:58 AM, Strong Liu wrote: > Hi all, > I'm trying to fix this issue : > Add HREFs for JDK methods to Hibernate Javadoc -- https://jira.jboss.org/jira/browse/JBPAPP-1701 > When I build Javadoc outside of brew, the HREFs are there, but in brew, it is missing, the problem is that the brew builds are done offline, when ant builds javadoc offline, it needs the "linkoffline" attribute specified with a location for the Java package list. > Details about the ant javadoc task are at: > http://ant.apache.org/manual/CoreTasks/javadoc.html > > Is the Java API package list already in the brew environment? If so, where will it be? > > If not, is there another package that should be added as a build requirement that has Java API package list?	1	1	https://issues.jboss.org/browse/JBPAPP-1701
10427	JBPAPP-1691	Rajesh, in previous CP releases, we used to deploy the *-jdbc2-service.xml files to server/all/deploy-hasingleton/jms, but now it seems that it doesn't works. Deploying to server/all/deploy/jms makes the test pass. Can you confirm this behaviour change?	0	0	https://issues.jboss.org/browse/JBPAPP-1691
10428	JBPAPP-1685	Perhaps we can call this "new" version 2.1.jdk15-brew as the 2.1-brew was still built with JDK 1.4 (JPackage has a patch to build with JDK 1.5 in their latest release of this RPM).	0	0	https://issues.jboss.org/browse/JBPAPP-1685
10429	JBPAPP-1679	What's the timeframe for this? I have this on my todo list, but currently I am quite busy with Bean Validation.	1	1	https://issues.jboss.org/browse/JBPAPP-1679
10430	JBPAPP-1679	In order to backport these things - what is the branch for hibernate core for these modules?	1	1	https://issues.jboss.org/browse/JBPAPP-1679
10431	JBPAPP-1648	Martin, I'm testing the patches, but is there a test case which I should pay special attention?	1	1	https://issues.jboss.org/browse/JBPAPP-1648
10432	JBPAPP-1648	Juraci, could you please run the tests again with the modified patch? Thanks!	1	1	https://issues.jboss.org/browse/JBPAPP-1648
10433	JBPAPP-1643	The dependency JBPAPP-1961 is still open, can this bug still be resolved? If yes, then should JBPAPP-1961 still be documented as fixed in this update?	1	0	https://issues.jboss.org/browse/JBPAPP-1643
10434	JBPAPP-1612	Clebert, could you help us with that change?	0	0	https://issues.jboss.org/browse/JBPAPP-1612
10435	JBPAPP-1612	I saw the solution in: http://fisheye.jboss.org/browse/JBossAS/branches/JBPAPP_4_2_0_GA_CP/varia/src/resources/services/binding/sample-bindings.xml?r1=82905&r2=82968 You probably haven't read my description. This same sample-bindinds.xml is used to build both EAP 4.2 and EAP 4.3. The way you "solve" it will only work for EAP 4.3. You need to either: a) Duplicate the binding for both 'jboss.jms' & 'jboss.mq' b) Leave it as it was and apply a patch when building 4.3 with the changed that you've made.	1	1	https://issues.jboss.org/browse/JBPAPP-1612
10436	JBPAPP-1612	If it's not too much of a hassle can you do it the other way, i.e. create a sample-bindings-4.3.xml for the 4.3 build so we can keep sample-bindings.xml synched with AS Branch_4_2? Thanks.	1	1	https://issues.jboss.org/browse/JBPAPP-1612
10437	JBPAPP-1598	Done?	0	0	https://issues.jboss.org/browse/JBPAPP-1598
10438	JBPAPP-1593	Done?	1	1	https://issues.jboss.org/browse/JBPAPP-1593
10439	JBPAPP-1553	Since the FP ships with Seam 2.0.2 whereas the EAP 4.3 CP ships with 1.2.1, this fix requires inclusion in the FP release notes as well -> however it should be stated that the FP fix is for the different version of Seam so there is no confusion for the customer when reading both the CP and FP release notes.	1	1	https://issues.jboss.org/browse/JBPAPP-1553
10440	JBPAPP-1550	I saw this jira notification and thought although error manifests in portal, it's caused by Webservices meaning that if a webservice sar/war is deployed in EAP, it would cause the same error. Rajesh, could you have some EAP folks look into it? If not, I will handle it.	0	0	https://issues.jboss.org/browse/JBPAPP-1550
10441	JBPAPP-1533	The subtasks JIRAs associated with this JIRA represent this information. Is there a reason we have to list them in an additional JIRA comment as well?	1	0	https://issues.jboss.org/browse/JBPAPP-1533
10442	JBPAPP-1532	The subtask JIRAs associated with this JIRA represent this information. Is there a reason we have to list them in an additional JIRA comment as well?	1	0	https://issues.jboss.org/browse/JBPAPP-1532
10443	JBPAPP-1522	Dominik, can you port Scott's fix from Branch_4_2 to the EAP branch and verify in a nightly build if the issue has been fixed?	0	0	https://issues.jboss.org/browse/JBPAPP-1522
10444	JBPAPP-1522	Has this fix been committed? Can this issue be resolved?	1	1	https://issues.jboss.org/browse/JBPAPP-1522
10445	JBPAPP-1512	Hi Chris, how is the investigation about signing JARs built in Brew automatically going? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-1512
10446	JBPAPP-1512	Peter, can you verify this on the CR1 build? Test cases are in testopia.	1	0	https://issues.jboss.org/browse/JBPAPP-1512
10447	JBPAPP-1479	Permaine, can you please check with Clebert if this is something we can get removed when the soa switch is specified?	0	0	https://issues.jboss.org/browse/JBPAPP-1479
10448	JBPAPP-1460	> 2) What do you mean? It is there since Nov 25. Where you looking under maven2? AS 4.2 is pre-maven, uses a buildmagic repo (top level of r.j.o.). Here is the URL: > http://repository.jboss.org/glassfish/jsf/1.2_10-brew/ I see now. I was looking under http://repository.jboss.org/sun-jsf/. Under that I find some old JSF brew builds. So I thought that's where 1.2_10-brew would go.	1	1	https://issues.jboss.org/browse/JBPAPP-1460
10449	JBPAPP-1439	A further question, does this only apply to EAP 4.3 or 4.2 as well? Cheers, Isaac	1	0	https://issues.jboss.org/browse/JBPAPP-1439
10450	JBPAPP-1433	Is there any specific test case I should pay more attention regarding this item?	1	1	https://issues.jboss.org/browse/JBPAPP-1433

#	Requirement	Questions	Answered	Successfully Answered	URL
10451	JBPAPP-1428	runs against the 'all' configuration: I am running: - testsuite checked out from JBPAPP_4_2_0_GA_CP04_JBPAPP-1180_JBPAPP-1268_JBPAPP-1428 - jboss 4.2.0.GA.CP04 with jars patched [rachmatowicz@dev91 reports]$ more TEST-org.jboss.test.cache.bean.MBeanUnitTestCase.txt Testsuite: org.jboss.test.cache.bean.MBeanUnitTestCase Tests run: 4, Failures: 3, Errors: 0, Time elapsed: 3.007 sec Testcase: testSetup took 0.142 sec Testcase: testPutTx took 0.362 sec FAILED javax.transaction.SystemException: java.lang.IllegalArgumentException: null object name junit.framework.AssertionFailedError: javax.transaction.SystemException: java.lang.IllegalArgumentException: null object name at org.jboss.test.cache.bean.MBeanUnitTestCase.testPutTx(MBeanUnitTestCase.java:112) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) Testcase: testRollbackTx took 0.069 sec FAILED javax.transaction.SystemException: java.lang.IllegalArgumentException: null object name junit.framework.AssertionFailedError: javax.transaction.SystemException: java.lang.IllegalArgumentException: null object name at org.jboss.test.cache.bean.MBeanUnitTestCase.testRollbackTx(MBeanUnitTestCase.java:149) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) Testcase: testReplicatedPutTx took 0.189 sec FAILED javax.transaction.SystemException: java.lang.IllegalArgumentException: null object name junit.framework.AssertionFailedError: javax.transaction.SystemException: java.lang.IllegalArgumentException: null object name at org.jboss.test.cache.bean.MBeanUnitTestCase.testReplicatedPutTx(MBeanUnitTestCase.java:190) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:24) at junit.extensions.TestSetup$1.protect(TestSetup.java:21) at junit.extensions.TestSetup.run(TestSetup.java:25) In the case of the second test, testRollbackTx(), I put in a Throwable.printStackTrace() and got the following: ---------- Standard Error ---------- javax.transaction.SystemException: java.lang.IllegalArgumentException: null object name at org.jboss.tm.usertx.client.ClientUserTransaction.begin(ClientUserTransaction.java:128) at org.jboss.test.cache.bean.MBeanUnitTestCase.testRollbackTx(MBeanUnitTestCase.java:138) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at junit.framework.TestCase.runTest(TestCase.java:164) at junit.framework.TestCase.runBare(TestCase.java:130) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at	1	1	https://issues.jboss.org/browse/JBPAPP-1428
10452	JBPAPP-1428	Galder I run the test org.jboss.cache.bean.MBeanUnitTestCase against the patched version of EAP 4.2.0.GA.CP04 and it fails with the above errors. I run the same test against the version of AS built from the tagged version JBPAPP_4_2_0_GA_CP04_JBPAPP-1180_JBPAPP-1268_JBPAPP-1428, after having unsecured it, and the test passes. Can you confirm the patch instructions? In the mean time, i'll try and follow up on possible differences.	1	1	https://issues.jboss.org/browse/JBPAPP-1428
10453	JBPAPP-1428	http://anonsvn.jboss.org/repos/jbossas/tags/JBPAPP_4_2_0_GA_CP04 --new http://anonsvn.jboss.org/repos/jbossas/branches/JBPAPP_4_2_0_GA_CP04_JBPAPP-1180_JBPAPP-1268_JBPAPP-1428 2. Used awk to remove the extra information and get a list of files changed: cluster/src/main/org/jboss/proxy/AbstractTransactionStickyInterceptor.java cluster/src/main/org/jboss/proxy/ClientUserTransactionStickyInterceptor.java cluster/src/main/org/jboss/proxy/TransactionStickyInterceptor.java cluster/src/main/org/jboss/proxy/generic/ProxyFactoryHA.java cluster/src/main/org/jboss/ha/framework/interfaces/TransactionSticky.java cluster/src/main/org/jboss/invocation/unified/interfaces/UnifiedInvokerHAProxy.java cluster/src/main/org/jboss/invocation/jrmp/interfaces/JRMPInvokerProxyHA.java cluster/build.xml transaction/src/main/org/jboss/tm/TransactionPropagationContextUtil.java testsuite/local.properties testsuite/src/main/org/jboss/test/cluster/invokerha/InvokerHaFailureType.java testsuite/src/main/org/jboss/test/cluster/invokerha/ejb/UserTransactionStickinessVerifierInterceptor.java testsuite/src/main/org/jboss/test/cluster/invokerha/ejb/test/UserTransactionStickyUnitTestCase.java testsuite/src/main/org/jboss/test/cluster/invokerha/ejb/UserTransactionStickyBean.java testsuite/src/main/org/jboss/test/cluster/invokerha/ejb/UserTransactionStickyHome.java testsuite/src/main/org/jboss/test/cluster/invokerha/ejb/UserTransactionStickyRemote.java testsuite/src/main/org/jboss/test/cluster/invokerha/UnifiedInvokerHaTransactionStickyUnitTestCase.java testsuite/src/main/org/jboss/test/cluster/invokerha/InvokerHaTransactionalMockUtils.java testsuite/src/main/org/jboss/test/cluster/invokerha/UnifiedInvokerHaMockUtils.java testsuite/src/main/org/jboss/test/cluster/invokerha/UnifiedInvokerHaUnitTestCase.java testsuite/src/main/org/jboss/test/cluster/invokerha/AbstractInvokerHaTransactionSticky.java testsuite/src/main/org/jboss/test/cluster/invokerha/InvokerHaInfrastructure.java testsuite/src/main/org/jboss/test/cluster/invokerha/JRMPInvokerHaTransactionStickyUnitTestCase.java testsuite/src/main/org/jboss/test/cluster/invokerha/AbstractInvokerHa.java testsuite/src/main/org/jboss/test/cluster/invokerha/JRMPInvokerHaMockUtils.java testsuite/src/main/org/jboss/test/cluster/invokerha/JRMPInvokerHaUnitTestCase.java testsuite/src/resources/cluster/invokerha/ejb/META-INF/jboss.xml testsuite/src/resources/cluster/invokerha/ejb/META-INF/ejb-jar.xml testsuite/src/resources/test-configs/ustxsticky/conf/jboss-service.xml testsuite/imports/sections/cluster.xml testsuite/imports/server-config.xml testsuite/build.xml server/src/main/org/jboss/tm/usertx/server/ClientUserTransactionService.java server/src/main/org/jboss/tm/usertx/server/UserTransactionSessionImpl.java server/src/main/org/jboss/tm/usertx/client/ClientUserTransactionObjectFactory.java server/src/main/org/jboss/tm/usertx/client/ClientUserTransaction.java server/src/main/org/jboss/tm/usertx/client/ServerVMClientUserTransaction.java server/src/main/org/jboss/proxy/ejb/RetryInterceptor.java	1	1	https://issues.jboss.org/browse/JBPAPP-1428

#	Requirement	Questions	Answered	Successfully Answered	URL
10454	JBPAPP-1428	Jimmy I reran the testsuite adding in the jboss.jar file to server/all/lib and server/production/lib. The testsuite ran very clean, with one error only (org.jboss.test.tm.test.TxTimeoutIUnitTestCase) and this test passed successfully on retry several times. I was still getting the following WARN messages appearing on the console: 13:17:59,301 WARN [loggerI18N] [com.arjuna.ats.internal.jta.resources.arjunacore.norecoveryxa] [com.arjuna.ats.internal.jta.resources.arjunacore.norecoveryxa] Could not find new XAResource to use for recovering non-serializable XAResource < 131075, 28, 26, 1-a1058a0:c685:4935a90e:7d4ba1058a0:c685:4935a90e:7d51 > 13:17:59,308 WARN [loggerI18N] [com.arjuna.ats.internal.jta.resources.arjunacore.norecoveryxa] [com.arjuna.ats.internal.jta.resources.arjunacore.norecoveryxa] Could not find new XAResource to use for recovering non-serializable XAResource < 131075, 28, 26, 1-a1058a0:c685:4935a90e:7d55a1058a0:c685:4935a90e:7d5b > 13:17:59,315 WARN [loggerI18N] [com.arjuna.ats.internal.jta.resources.arjunacore.norecoveryxa] [com.arjuna.ats.internal.jta.resources.arjunacore.norecoveryxa] Could not find new XAResource to use for recovering non-serializable XAResource < 131075, 28, 26, 1-a1058a0:c685:4935a90e:7d5fa1058a0:c685:4935a90e:7d65 > These messages appear after the testsuite run and when restarting the server after the test run. Deleting the data directory and tmp directory and restarting the server causes the messages not to reappear, so I assume they are due to some test artifacts left over from the test run. I'm going to ask Jonathan Halliday if he knows why they might arise. Jonathan, do these messages cause any concern for you?	1	1	https://issues.jboss.org/browse/JBPAPP-1428
10455	JBPAPP-1393	Isaac, does this make it to this CP? if not, can you move it to the next one? Thanks, Permaine	1	0	https://issues.jboss.org/browse/JBPAPP-1393
10456	JBPAPP-1393	Hi Tihomir, Can you advise where I can find information relating to the deployment of a foo-ds.xml application? I have found the following material online: http://www.jboss.org/file-access/default/members/jbossweb/freezone/docs/latest/deployer-howto.html foo-ds.xml refernced in the Seam Reference Guide. http://docs.jboss.org/seam/2.1.2/reference/en-US/html/testing.html ....but cannot find a single definitive guide. Alternatively, can you provide a description? Thanks, Ben	1	0	https://issues.jboss.org/browse/JBPAPP-1393
10457	JBPAPP-1393	Laura, this was assigned to Tiho. Is it because we need more information from him? After it was assigned to him, you moved the CP fix version, so I want to be sure what we need.	1	0	https://issues.jboss.org/browse/JBPAPP-1393
10458	JBPAPP-1382	Done?	1	1	https://issues.jboss.org/browse/JBPAPP-1382
10459	JBPAPP-1381	Most of it is already done right?	0	0	https://issues.jboss.org/browse/JBPAPP-1381
10460	JBPAPP-1380	Done?	0	0	https://issues.jboss.org/browse/JBPAPP-1380
10461	JBPAPP-1368	I don't see where you made the change so that embedded is built with the new hibernate. I see lots of file changes for embedded version, but I must be missing the one for the hibernate version change. Can you point me to it?	1	1	https://issues.jboss.org/browse/JBPAPP-1368
10462	JBPAPP-1362	Hey Justin, I'm trying to work out where to fit this information into the Server Configuration Guide, is this part of the JBoss JAXR Configuration?	1	1	https://issues.jboss.org/browse/JBPAPP-1362
10463	JBPAPP-1341	I have been able to find information on the BaseCertLoginModule configuration in the EAP Server Config Guide: http://www.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/4.3.0.cp02/html-single/Server_Configuration_Guide/index.html 9.5.3.7. BaseCertLoginModule However I have not been able to find anything for the others, does anyone have any information they can point me to that may be able to assist?	1	1	https://issues.jboss.org/browse/JBPAPP-1341
10464	JBPAPP-1334	Jean-Frederic, can you point Isaac to a list of the changes, please? Thanks in advance.	0	0	https://issues.jboss.org/browse/JBPAPP-1334
10465	JBPAPP-1333	A few discrepancies to note while reviewing the attachments: The RHEL 4 rpms at location: http://porkchop.devel.redhat.com/devel/candidates/JBEAP4.3.0/latest/rpm/JBEAP4.3.0-re20081028.0/i386/i386/RedHat/RPMS/ show the documentation rpm as: rh-eap-docs-4.3.0-4.GA_CP03.ep1.1.el4.noarch.rpm Also with the jbossas* rpms: From above link, the RHEL 5 rpms for jbossas* are shown as jbossas-4.3.0-3.GA_CP03.6.ep1.el5.noarch.rpm 27-Oct-2008 15:48 13M jbossas-client-4.3.0-3.GA_CP03.6.ep1.el5.noarch.rpm 27-Oct-2008 15:49 6.0M Shown in your RHEL 5 attachment: jbossas-4.3.0-3.GA_CP03.6.2.ep1.el5.noarch.rpm jbossas-client-4.3.0-3.GA_CP03.6.2.ep1.el5.noarch.rpm The former issues regarding the doc rpms seems to be that the doc updates were not made available for RHEL 5 [or to the RHEL 5 link] but can't explain the difference in the version of the jbossas rpms noted above. Can you comment?	0	0	https://issues.jboss.org/browse/JBPAPP-1333
10466	JBPAPP-1289	Darran, do you know how can we get the information Isaac needs? Thanks in advance.	0	0	https://issues.jboss.org/browse/JBPAPP-1289
10467	JBPAPP-1276	I see two tests failing. The first one I've no idea why, but wrt the second one, I wonder why is it trying to use MQ classes? <testcase classname="org.jboss.test.ws.jaxws.jbws1809.JBWS1809TestCase" name=" testWSDLAccess" time="0.415"> <failure message="JAXBIntros should turn the 'docRequest.name' pro perty into a XML attribute" type="junit.framework.AssertionFailedError">junit.fr amework.AssertionFailedError: JAXBIntros should turn the 'docRequest.name&a pos; property into a XML attribute at org.jboss.test.ws.jaxws.jbws1809.JBWS1809TestCase.testWSDLAccess(JBWS 1809TestCase.java:89) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) </failure> </testcase> <testcase classname="org.jboss.test.ws.jaxrpc.samples.jmstransport.JMSTranspor tTestCase" name="testSOAPMessageToEndpointQueue" time="0.042"> <error message="org/jboss/aop/pointcut/PointcutStats" type="java.lang.NoClas sDefFoundError">java.lang.NoClassDefFoundError: org/jboss/aop/pointcut/PointcutS tats at org.jboss.jms.client.delegate.ClientConnectionFactoryDelegate.<cli nit>(ClientConnectionFactoryDelegate.java) at sun.misc.Unsafe.ensureClassInitialized(Native Method) ... On Stderr: ... Caused by: org.jboss.deployment.DeploymentException: No ClassLoaders found for: org.jboss.mq.server.jmx.Queue; - nested throwable: (java.lang.ClassNotFoundExcep tion: No ClassLoaders found for: org.jboss.mq.server.jmx.Queue) ...	0	0	https://issues.jboss.org/browse/JBPAPP-1276
10468	JBPAPP-1276	Since we are talking about EAP 4.3 GA, JBPAPP-491 seem to be exactly the issues I'm seeing. Though I wonder why JBPAPP-678 has not fixed the MQ issue. Please advice on how should I proceed. All other tests pass so here are the MD5 sums in case we're going to ignore: da465b8d24d16a3dc8ac30a1545bd7eb jbossws-client.jar b96d726b4b88d0f9b91381cb83b1ca45 jbossws-common.jar b51ad8a99ba7c3ed925caf1b0145d457 jbossws-core.jar 1d077383a5930bade2c2e1f828a95c06 jbossws-native.jar	0	0	https://issues.jboss.org/browse/JBPAPP-1276
10469	JBPAPP-1276	Please ignore the second failure, this is against JBoss EAP 4.3.0.GA and JBPAPP-678 did not make it in until JBoss EAP 4.3.0.GA_CP01 so that test is expected to fail. Do you have any further information on the first test that is failing, that is not failing locally for me. When run on it's own does that test still fail? Can you attach the junit files and test.log and server.log for that failure so I can see what is happening?	0	0	https://issues.jboss.org/browse/JBPAPP-1276

#	Requirement	Questions	Answered	Successfully Answered	URL
10470	JBPAPP-1276	Just before I start testing again. Last time I've set -Dorg.jboss.ws.eagerInitializeJAXBContextCache=true when starting the AS server. Do I need to set that when running the test suite as well (ant command line)?	0	0	https://issues.jboss.org/browse/JBPAPP-1276
10471	JBPAPP-1272	Hi Isaac, The notes referenced in the pdf (e.g. JBPAPP-916 [1]) don't have any explanation in the footer. Can you check that? Also, I found a typo which has been there since the previous version: JavaSE 6 includes includes support... should be: Java SE 6 includes support... Besides that, it looks alright. Thanks, Martin	1	1	https://issues.jboss.org/browse/JBPAPP-1272
10472	JBPAPP-1268	Ond?ej Žižka, First of all, not sure how your name is spelt, so taking the spelling JIRA did. Would it be possible to stop the QA process? There's a small bottleneck that we need to resolve in the fix provided. I'll hand it back to you once it's been fixed.	1	1	https://issues.jboss.org/browse/JBPAPP-1268
10473	JBPAPP-1268	Ond?ej Žižka, Assigning it back to you as customer is happy to go with the original fix for the moment. Could you also please confirm when you expect QA for this patch to finish and for it to be uploaded? In theory, it should take 5 biz days from the moment I pass it on to you. So, should be uploaded by 19th/20th November. Can you please confirm?	1	1	https://issues.jboss.org/browse/JBPAPP-1268
10474	JBPAPP-1268	Galder, we have rerun the testsuite last night and should be able to give you a result in an hour or so. There seem to be a few problems, but let us see if we can give you a repeatable error. For the WaekHashMap -> Map change, couldn't we update the testcases in the branch?	1	1	https://issues.jboss.org/browse/JBPAPP-1268
10475	JBPAPP-1268	Richard, I was just having a chat with Ondrej and he said the test pass Ok now. What exactly is it going on? Re: For the WaekHashMap -> Map change, couldn't we update the testcases in the branch? Well, I have already! At the beginning the testsuite from EAP 4.2.0.GA_CP04 was being run where weak hash map was still used but of course I updated the test in the branch. Re: The QA lab has just been taken down for 63 hours... What exactly does that mean? Ondrej was fine just a few minutes ago...	0	0	https://issues.jboss.org/browse/JBPAPP-1268
10476	JBPAPP-1262	Dominik, is this for EAP 4.x or for EAP 5? Can you please fill out the Affects and Fix Versions.	0	0	https://issues.jboss.org/browse/JBPAPP-1262
10477	JBPAPP-1261	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	1	1	https://issues.jboss.org/browse/JBPAPP-1261
10478	JBPAPP-1260	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-1260
10479	JBPAPP-1260	Rado, can you file bugs and link it to this JIRA for any issue that showed up on the failover and soak testing on the CR2 build?	0	0	https://issues.jboss.org/browse/JBPAPP-1260
10480	JBPAPP-1260	Dominik, were you able to run the failover tests with the fixes from the linked JIRA's? Can you comment on the latest status? Are there any new issues?	0	0	https://issues.jboss.org/browse/JBPAPP-1260
10481	JBPAPP-1227	Patch applied JBPAPP_4_2_0_GA_CP branch. Rev 78994.	0	0	https://issues.jboss.org/browse/JBPAPP-1227
10482	JBPAPP-1222	Has anyone gone through the DB2-specific JIRA issues for the community version of Hibernate? If the JBPAPP issues for DB2 were created just from unit test failures, then there may be more issues in the community version that are not unit tested yet. It would be a good idea to check if there is anything important identified in the community version and document it as known for this CP, then I can plan to fix it for the next CP.	0	0	https://issues.jboss.org/browse/JBPAPP-1222
10483	JBPAPP-1222	There are some open issues, but very few with test cases attached. So, it would be good to select some of the issues and build a test case for them. I can build them, but I'd need the selected issues. This JIRA query can be helpful: http://opensource.atlassian.com/projects/hibernate/secure/IssueNavigator.jspa?reset=true&&type=1&type=5&query=db2&summary=true&description=true&body=true&environment=true&resolution=-1&status=1&status=4&status=10000&sorter/field=updated&sorter/order=DESC&sorter/field=votes&sorter/order=DESC (you may want to add the "Votes" to your view) These are the most voted issues, but I saw in the comments that one of them won't be fixed and another doesn't looks like an issue. Maybe it is time to close them then? HHH-2763 HHH-817 HHH-1348 HHH-2160	1	1	https://issues.jboss.org/browse/JBPAPP-1222
10484	JBPAPP-1214	Is this tag ready? Can you please list which fixes are included in it since the last CP? Cheers, Fernando	1	1	https://issues.jboss.org/browse/JBPAPP-1214
10485	JBPAPP-1214	No, it hasn't been created yet. Should I do this before finishing unresolved issues, or wait until everything is done?	1	1	https://issues.jboss.org/browse/JBPAPP-1214
10486	JBPAPP-1199	JBAOP-633 is closed, can you confirm this fix is in the 1.5.5 CP branch and that it will be in the 1.5.5.cp03 release?	1	1	https://issues.jboss.org/browse/JBPAPP-1199
10487	JBPAPP-1180	How do you revert back the workflow status?	1	1	https://issues.jboss.org/browse/JBPAPP-1180
10488	JBPAPP-1180	Martin, UnifiedInvokerHaTransactionStickyUnitTestCase and UnifiedInvokerHaUnitTestCase are basic unit test cases and I've just run them from Eclipse successfully. How are you running these tests? What were the failures? Could you attach the unit test reports for these two classes? In the mean time, I'm specting the issue reported by the customer when testing this patch.	0	0	https://issues.jboss.org/browse/JBPAPP-1180
10489	JBPAPP-1168	Why did Seam migrate to xsd files? What advantage did that give the project?	0	0	https://issues.jboss.org/browse/JBPAPP-1168
10490	JBPAPP-1131	Which version o Hibernate Core are you using? Gail did a port of getPostCollectionRecreateEventListeners() into the production version to ensure Hibernate Search work with it.	0	0	https://issues.jboss.org/browse/JBPAPP-1131
10491	JBPAPP-1029	We do have jbossws "common", "framework" and "spi", besides the main jbossws build, which also produces jbossws-native42. We don't have any "integration" module. What is the repository.jboss.org module for it? Is this a new thing?	1	1	https://issues.jboss.org/browse/JBPAPP-1029
10492	JBPAPP-1029	And this separation happened between 2.0.1.SP2_CP03 and 2.0.1.SP2_CP04? We did not need any "integration" to build CP03. This kind of refactoring does not seem appropriate for a CP branch. Are you sure this is not produced by the main WS build (which is also producing native42 for us)? You can also check that http://repository.jboss.org/jboss has no jbossws-integration module, nor is one mentioned in the component-info.xml of jbossws-native42: http://repository.jboss.org/jboss/jbossws-native42/2.0.1.SP2_CP03-brew/component-info.xml AFAIK the integration code was separated for JBoss AS 5, but that is a different version of the WS.	1	1	https://issues.jboss.org/browse/JBPAPP-1029
10493	JBPAPP-1029	We don't have this step in what Heiko gave us (Deepak). All we do is to build JBoss WS core which also produces the native42 bits: svn export http://anonsvn.jboss.org/repos/jbossws/stack/native/tags/jbossws-native-2.0.1.SP2_CP03 jbossws # Copy resources cp %{SOURCE3} jbossws-jboss42-resources.zip popd mv ant.properties.example ant.properties # Build jbossws ant -Djboss42.home=$JBOSS_HOME -Djboss.server.instance=all -Dchecksum.ok=true jars We do get these "resources" from somewhere. Deepak, the instructions to get the Source3 are missing in the spec file. Can you fill that in please?	1	1	https://issues.jboss.org/browse/JBPAPP-1029
10494	JBPAPP-1000	Pavel, can you determine the minimum hardware requirements to run EAP 4.x. Production server config would be a good candidate, minimum would be server startup and running a sample app like the jmx-console would be sufficient.	1	1	https://issues.jboss.org/browse/JBPAPP-1000
10495	JBPAPP-1000	Isaac, can you add the minimum hardware requirements to the installation guide?	1	1	https://issues.jboss.org/browse/JBPAPP-1000
10496	JBPAPP-1000	Pavel, a couple of questions: 1) What CPUs does 'modern' refer to? All 32bit and 64bit processors? We need to be specific for this section. 2) What is the configuration of a 'simple application'? Is a simple application the most common usage for users? If not, perhaps we should mention specifications for the most common setup s well in order to not lead customers astray. Thanks, Isaac	1	1	https://issues.jboss.org/browse/JBPAPP-1000
10497	JBPAPP-983	Rajesh, was the EAP 4.2 certified as well?	1	0	https://issues.jboss.org/browse/JBPAPP-983

#	Requirement	Questions	Answered	Successfully Answered	URL
10498	JBPAPP-978	FYI on javadoc problem This javadoc problem was seen some time ago and is related to a bug in the JDK. Following are are my notes from when I looked into it. Looks as though I updated all versions exce[pt Branch_2_4. I've now added in the change to Branch_2_4 so it won't happen next time. Notes: There is a JDK bug (http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=5082928) which presents itself as a java.util.Missing ResourceException when you try to compile the javadoc for JGroups using jdk1.5.0_06 or jdk1.5.0_11. I've included an example of the exception at the end of this mail. In a nutshell, versions of Javadoc bundled with JDK 1.5.0 have problems handling Javadoc text containg the characters "<A" in succession - they try to interpret it as an HTML anchor and then can't find the appropriate error message text and throw the exception. This bug is fixed in JDK 6. There is a workaround for the bug: replace "<A" with the equivalent <A. By playing around a bit, I found that this workaround: (i) prevents javadoc from throwing the exception, but does not generate correct javadoc for the element containing the <A with jdk1.5.0_06 (ii) prevents javadoc from throwing the exception, and does generate correct javadoc for the element containing the <A with jdk1.5.0_11 Here is a list of JGroups sourcefiles where this problem arises (I got these error messages using another trick mentioned in the bug report): src/org/jgroups/blocks/RequestCorrelator.java:703: warning - Malformed link: Stack<Address>. src/org/jgroups/util/Digest.java:34: warning - Malformed link: Map<Address, Entry> I've replaced the "<A" with <A in all cases, rebuilt the Javadoc and looked at it - everything seems OK with jdk1.5.0_11. Should I update the codebase with these changes? Richard	1	1	https://issues.jboss.org/browse/JBPAPP-978
10499	JBPAPP-978	Please do. Do you know if the resulting javadocs without your fix are generated correctly or if we should incorporate your patch and rebuild 2.4.4 to get proper javadocs?	1	1	https://issues.jboss.org/browse/JBPAPP-978
10500	JBPAPP-966	This is issue isn't fixed in Jboss 5.1.0GA (as downloaded from jboss site) with jbossws-native-3.1.2.GA Adding link to jboss-common-client.jar to classpath of wsrunclient does not make any changes. Stacktrace is (watch that string numbers in Provider and ProviderLoader changed): Exception in thread "main" java.lang.IllegalStateException: Failed to load javax.xml.ws.spi.Provider: org.jboss.ws.core.jaxws.spi.ProviderImpl at javax.xml.ws.spi.ProviderLoader.loadProvider(ProviderLoader.java:94) at javax.xml.ws.spi.Provider.provider(Provider.java:87) at javax.xml.ws.Service.<init>(Service.java:79) at org.jbia.ws.HelloService.<init>(HelloService.java:45) at org.jbia.ws.HelloWsClient.main(HelloWsClient.java:18) Caused by: java.lang.ClassNotFoundException: org.jboss.ws.core.jaxws.spi.ProviderImpl at java.net.URLClassLoader$1.run(URLClassLoader.java:200) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:188) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:276) at java.lang.ClassLoader.loadClass(ClassLoader.java:251) at javax.xml.ws.spi.ProviderLoader.loadProvider(ProviderLoader.java:88) ... 4 more Could you explain step by step how did you fix it? What does it mean "resolved in WS build" if it doesn't work in distributable jboss-5.1.0.GA ?	0	0	https://issues.jboss.org/browse/JBPAPP-966
10501	JBPAPP-943	While I begin work on this, do you know if the EAP 4.2 version and fixed issues be ready for the docs freeze on the 9th of July?	1	1	https://issues.jboss.org/browse/JBPAPP-943
10502	JBPAPP-932	Do we need to upgrade to JBossXB 1.0.0 SP3? What about jaxr?	0	0	https://issues.jboss.org/browse/JBPAPP-932
10503	JBPAPP-931	Hibernate entity manager and annotations will continue to be patched for EAP 4.2 CP04 / 4.3 CP02. The patches will be backported in the next release.	0	0	https://issues.jboss.org/browse/JBPAPP-931
10504	JBPAPP-918	It looks like the version for Entity Manager needs to be something different from 3.2.1.GA_CP02 to distinguish it from the previous version which was incorrectly set to 3.2.1.GA_CP02 for the code tagged v3_2_1_GA_CP01. The code has already been tagged as v3_2_1_GA_CP02 with the same version (3.2.1.GA_CP02) as v3_2_1_GA_CP01. Should I re-tag the branch as v3_2_1_GA_CP03 with version 3.2.1.GA_CP03 to keep the versions distinct? The code freeze is today so please let me know as soon as possible.	0	0	https://issues.jboss.org/browse/JBPAPP-918
10505	JBPAPP-910	Hi Clebert, is this done? Do we need a new JB Mesaging (we have 1.4.0.SP3_CP02)?	0	0	https://issues.jboss.org/browse/JBPAPP-910
10506	JBPAPP-910	Resolving this, the db2 persistence file has been tested and committed to the JBoss Messaging CP branch. All Messaging tests pass against DB2 9.1 Clebert, can you make sure to upload the db2 persistence file to the repository when you cut off the next JBM CP release? It will then make its way into the EAP example docs	0	0	https://issues.jboss.org/browse/JBPAPP-910
10507	JBPAPP-887	I am sorry, I do not understand the reason for the question. What problem are you trying to solve?	1	1	https://issues.jboss.org/browse/JBPAPP-887
10508	JBPAPP-887	In most cases, we see a jboss-<component>.jar up on repository.jboss.org, it's just that this one is different. Is the juddi-saaj.jar the only jar needed?	1	1	https://issues.jboss.org/browse/JBPAPP-887
10509	JBPAPP-886	This came too late, can you still include the fix? If not, please move it to the next CP versions, please.	1	0	https://issues.jboss.org/browse/JBPAPP-886
10510	JBPAPP-885	This came too late in the cycle, can you still include the fix by freeze date? If not, please move it to the next CP cycle please.	1	1	https://issues.jboss.org/browse/JBPAPP-885
10511	JBPAPP-878	TCK4 test should be run before releasing. (Changes in the API signatures).	0	0	https://issues.jboss.org/browse/JBPAPP-878
10512	JBPAPP-871	Clebert, can you take at this issue? This test case was added as part of the feature in JBAS-4945 Does org.jboss.ejb.plugins.jms.DLQHander handle DLQ's in jboss messaging as well? Please take a look at the analysis below from Martin. He doesn't see org.jboss.ejb.plugins.jms.DLQHander in the logs for EAP 4.3 From Martin: I was looking at the ResendMDBUnitTestCase failure and I'd like to give you a quick overview of what it's about. I first thought it was only testing JBoss Messaging and therefore used only in EAP 4.3, but that's not correct. In fact, it is being run on both releases (and it was first developed for Branch_4_2). Apparently it runs fine on 4.2. So how it works: The testcase first deploys 3 Message Driven Beans. The first one (MDBA) is connected to queue/A and whenever it reads a message, it does setRollbackOnly(). By default, the message is resent 10 times and then it should go to DLQ (Dead Letter Queue), which is set to be queue/B for this MDB. The second MDB (MDBB) is connected to queue/B and and whenever it gets a message, it resends it to queue/C. The third MDB (MDBC) is connected to queue/C and it does the same as the first MDB - when it receives a message, it just rolls it back. And then finally the message should go to DLQ, which is set to queue/D. The testcase sends a message to queue/A and then waits for it on queue/D as it should get there eventually. It works fine with JBoss MQ, but for some reason not with JBoss Messaging. All I can see in the log is that the message is 10 times rejected by MDBA, but it seems it never gets to the DLQ afterwards. For example JBoss MQ writes this after the 10 attempts: WARN [DLQHandler] Message resent too many times; sending it to DLQ; message id=ID:7-12136131402121 Whereas JBoss Messaging doesn't write anything and not even in the DEBUG level could I find any information of what was happening after those 10 attempts. But definitely the message never makes it to the second MDB (there is a log.info)	1	1	https://issues.jboss.org/browse/JBPAPP-871
10513	JBPAPP-871	Clebert, did you get a chance to look at the above comment?	1	1	https://issues.jboss.org/browse/JBPAPP-871
10514	JBPAPP-871	Isaac, can you please get in touch with Clebert to obtain a draft of the documentation bit? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-871

#	Requirement	Questions	Answered	Successfully Answered	URL
10515	JBPAPP-867	This already exists in AS5 branch - does it need to go in 4.2.x community before EAP?	1	1	https://issues.jboss.org/browse/JBPAPP-867
10516	JBPAPP-867	No, if you can backport it into the EAP CP branch. Is the change to the AS code only? P.S.: Freeze time approaching quicly. ;-)	0	0	https://issues.jboss.org/browse/JBPAPP-867
10517	JBPAPP-867	Chris, is there an HHH issue for this? Also, this issue is marked as a feature request. Should this be removed from the CP branch?	1	0	https://issues.jboss.org/browse/JBPAPP-867
10518	JBPAPP-865	Galder, is there any more information that you can provide about this issue and how it was fixed for release notes? Thanks, Isaac	1	0	https://issues.jboss.org/browse/JBPAPP-865
10519	JBPAPP-858	Rajesh, this -bin will be in this CP. Can you please add running one installation from -bin in addition to the regular RPM installation to the test plan? One from 4.2 and one from 4.3 would be what would make more sense. Note: You cannot install just the -bin, you need a regular jbossas installed on the server first. The -bin is for an _additional_ installation only, it cannot be used in place of a regular RPM jbossas installation. If you only have one version of the AS, it will have to be the regular RPM installation (or the old ZIP or installer ones). Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP-858
10520	JBPAPP-851	Darran, will this fix happen in the AS code or in WS code?	1	1	https://issues.jboss.org/browse/JBPAPP-851
10521	JBPAPP-823	Ok, thanks for checking this. W.r.t. jbossws-context.war I've got another idea: Can we leave this as an archive? Then we would patch it like any other jar file. And in the next CP04 release we would make it an archive again in the full distributions to keep it consistent between all distros. Or is there a reason why this is now an exploded war?	1	1	https://issues.jboss.org/browse/JBPAPP-823
10522	JBPAPP-823	I've slightly modified dist-diff to pick up jbossws-context.war and verified it works well. So I vote for making this dir an archive again and do the same for the other distributions in the next CP release. If we don't do so, we won't have consistent patch instructions. Sure we can remove that file and copy a whole dir when patching from GA, or CP01 or CP02 to CP03. But what do we say to a customer who patches from GA to CP04 and to another customer who patches from CP03 to CP04? The second one would already have a directory (from patching to CP03) so he wouldn't need to remove the jbossws-context.war file, the first one would still stick with that war archive and would need to perform this step.	1	1	https://issues.jboss.org/browse/JBPAPP-823
10523	JBPAPP-823	Due to JBPAPP-733 we cannot provide an archive :(Chris, we need to do something similar we do for Seam. So after you have built the new CP, you provide the Seam node maually in install-process.xml, right? <node name="155"> <action class="com.jboss.jbossnetwork.product.jbpm.handlers.NotificationActionHandler" config-type="bean"> <notification>To patch Seam to 1.2.1.EAP_4.2_CP03 rename the old directory #{jbossHomeDir}/../seam to #{jbossHomeDir}/../seam-1.2.0.AP and copy the new Seam directory from #{patchFolder}/seam to #{jbossHomeDir}/../seam </notification> </action> <transition name="success" to="156"/> </node> And now we also need another node for the jbossws-context.war, please add something like this: <node name="156"> <action class="com.jboss.jbossnetwork.product.jbpm.handlers.NotificationActionHandler" config-type="bean"> <notification>To patch jbossws-context.war rename the old file or directory #{jbossServerHomeDir}/deploy/jbossws.sar/jbossws-context.war to #{jbossServerHomeDir}/deploy/jbossws.sar/jbossws-context.war.replaced_by_CP03 and copy the new directory from #{patchFolder}/jboss-as/server/all/deploy/jbossws.sar/jbossws-context.war to #{jbossServerHomeDir}/deploy/jbossws.sar/jbossws-context.war </notification> </action> <transition name="success" to="157"/> </node> and don't forget to increase the last node's number, in this case 157, but that may change since you have added the 2 init scripts I mentioned above to patches.build.xml. And same way like you add the Seam directory to the CP, please also copy the jbossws-context.war directory from EAP-4.2_CP03 to patch_folder/jboss-as/server/all/deploy/jbossws.sar and zip the whole CP again.	1	1	https://issues.jboss.org/browse/JBPAPP-823
10524	JBPAPP-812	Who decided that this is going to be part of EAP5?	0	0	https://issues.jboss.org/browse/JBPAPP-812
10525	JBPAPP-812	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	0	0	https://issues.jboss.org/browse/JBPAPP-812
10526	JBPAPP-785	The RC3 build has the jar in the following locations: junit: Source distribution: ./jboss-as/thirdparty/junit/lib/junit.jar (this location is populated during build, downloads from thirdparty repo) ./jboss-as/ejb3/docs/examples/xpetstore/lib/ant/junit/junit.jar ./jboss-as/tools/lib/junit.jar ./jboss-as/tools/lib/ant-junit.jar (not sure if ant-junit.jar falls into the same issue) None in the binary zip: Docs bundle packages junit.jar /doc/examples/server-config-guide/lib/junit.jar JBPM: Binary Zip includes it at: jboss-eap-4.3/seam/lib/jbpm-3.1.4.jar Src zip has it at: jboss-eap-4.3-src/seam/lib/jbpm-3.1.4.jar If these have been removed from the rpm composes, how did they end up in the zip? (zip's built from rpm's) The EAP build downloads jbpm and bpel in its thirdparty while building. Is that allowed? ./jboss-as/thirdparty/jbpm ./jboss-as/thirdparty/jbpm/jpdl ./jboss-as/thirdparty/jbpm/jpdl/component-info.xml ./jboss-as/thirdparty/jbpm/jpdl/lib ./jboss-as/thirdparty/jbpm/jpdl/lib/jbpm-identity.jar ./jboss-as/thirdparty/jbpm/jpdl/lib/jbpm-console.war ./jboss-as/thirdparty/jbpm/jpdl/lib/jbpm-jpdl.jar ./jboss-as/thirdparty/jbpm/jpdl/lib/jbpm-jmx.jar ./jboss-as/thirdparty/jbpm/bpel ./jboss-as/thirdparty/jbpm/bpel/component-info.xml ./jboss-as/thirdparty/jbpm/bpel/lib ./jboss-as/thirdparty/jbpm/bpel/lib/jbpm-bpel-jbossws10.jar ./jboss-as/thirdparty/jbpm/bpel/lib/jbpm-bpel.jar ./jboss-as/thirdparty/jbpm/bpel/lib/jbpm-bpel-jbossws12.jar ./jboss-as/thirdparty/jbpm/bpel/lib/jbpm-bpel.sar ./jboss-as/thirdparty/jbpm/bpel/lib/jbpm-bpel-jbossws4ee.jar ./jboss-as/thirdparty/jbpm/bpel/lib/jbpm-bpel-config.jar ./jboss-as/thirdparty/jbpm/bpel/lib/jbpm-bpel-tools.jar ./jboss-as/thirdparty/jbpm/bpel/lib/jbpm-bpel.war	0	0	https://issues.jboss.org/browse/JBPAPP-785
10527	JBPAPP-757	This probably is a CSS issue. Can we fix this in the next CP?	0	0	https://issues.jboss.org/browse/JBPAPP-757
10528	JBPAPP-757	Jay can you provide the content to be added?	1	1	https://issues.jboss.org/browse/JBPAPP-757
10529	JBPAPP-757	Can you please rewrite this section (refer to http://www.redhat.com/docs/manuals/jboss/jboss-soa-4.2/html/JBPM_Users_Guide/jpdlgraphicalprocessdesigner.html and https://engineering.redhat.com/docbot/en-US/JBoss_SOA_Platform/4.2/html/SOA_ESB_JBPM_Integration_Guide/ch01s04.html for examples).	0	0	https://issues.jboss.org/browse/JBPAPP-757
10530	JBPAPP-756	does this also apply to 4.2 branch?	1	1	https://issues.jboss.org/browse/JBPAPP-756
10531	JBPAPP-756	Hey Jay, I have attached am image of what I believe is the issue you are talking about. Is the issue indeed about how the red round dots with the numbers do not seem to line up correctly with the text they are meant to be associated with? If this is the problem then I suggest either removing the use of these image bullet points in favour of the usual listing system (as I haven't seen images used as bullet points in other docs before), or it would be a CSS issue with spacing as the issue i believe is caused from the size of the images. If I am going down the totally wrong track then please let me know :) Cheers, Isaac	1	1	https://issues.jboss.org/browse/JBPAPP-756
10532	JBPAPP-756	Isaac, was this resolved in the last doc drop you provided? I assume it was not since jira issue is still open. Deferring this until next CP since it's minor. Please update this once resolved.	0	0	https://issues.jboss.org/browse/JBPAPP-756
10533	JBPAPP-755	Jay, Are we talking about one of the files included in the docs folder? Which readme.txt is this?	1	1	https://issues.jboss.org/browse/JBPAPP-755
10534	JBPAPP-754	Do we need to update the Release Notes? Or changelog.txt? Where is this changelog.txt?	0	0	https://issues.jboss.org/browse/JBPAPP-754

#	Requirement	Questions	Answered	Successfully Answered	URL
10535	JBPAPP-735	We are using 2.2.2.SP7. Do we have a 2.2.2.SP8? What are the fixes in it?	0	0	https://issues.jboss.org/browse/JBPAPP-735
10536	JBPAPP-723	According to https://enterprise.redhat.com/issue-tracker/?module=issues&action=view&tid=184222. The examples still don't work. Could some-one from product team investigate?	0	0	https://issues.jboss.org/browse/JBPAPP-723
10537	JBPAPP-723	Can we have some details on what is failing?	1	1	https://issues.jboss.org/browse/JBPAPP-723
10538	JBPAPP-683	But who fixes it? Jimmy or myself?	0	0	https://issues.jboss.org/browse/JBPAPP-683
10539	JBPAPP-679	What version is affected? i think that only 1.2.x. Am I right? Pete, please update the affected version(s). Thanks i advance.	0	0	https://issues.jboss.org/browse/JBPAPP-679
10540	JBPAPP-676	What version is affected? i think that only 1.2.x. Am I right? Pete, please update the affected version(s). Thanks i advance.	0	0	https://issues.jboss.org/browse/JBPAPP-676
10541	JBPAPP-665	I was unware of the presence of the JBPAPP branch until I saw the email accompanying this JIRA issue, so this patch has not been added there. What is the next step?	1	1	https://issues.jboss.org/browse/JBPAPP-665
10542	JBPAPP-661	Jay, Could you please confirm whether this has been fixed in the JBPAPP-4_2_0_GA_CP branch? Please commit the changes in this branch and resolve it. Thanks	1	1	https://issues.jboss.org/browse/JBPAPP-661
10543	JBPAPP-660	Chris, what's the plan in this one?	1	1	https://issues.jboss.org/browse/JBPAPP-660
10544	JBPAPP-655	Are we set for the upgrade?	0	0	https://issues.jboss.org/browse/JBPAPP-655
10545	JBPAPP-640	Hi Mike, what's the status of this?	1	1	https://issues.jboss.org/browse/JBPAPP-640
10546	JBPAPP-624	At the recent developers meeting, Andy Miller expressed the opinion that component projects shouldn't have to maintain separate community and EAP branches. Concretely, that means that Remoting 2.2.2.SP5 can serve in place of 2.2.2.SP4_CP01. It follows that this issue is superseded by JBPAPP-655 "Update Remoting to 2.2.2.SP5" (fix version 4.3.0.GA_CP01). Fernando is OK with this plan. Is everyone else OK?	1	1	https://issues.jboss.org/browse/JBPAPP-624
10547	JBPAPP-622	Does this apply to EAP as well? Is all	0	0	https://issues.jboss.org/browse/JBPAPP-622
10548	JBPAPP-620	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	1	1	https://issues.jboss.org/browse/JBPAPP-620
10549	JBPAPP-614	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	1	1	https://issues.jboss.org/browse/JBPAPP-614
10550	JBPAPP-613	Is this Issue a "Blocker" or "Critical" priority for EAP GA? If yes. please adjust the priority to indicate that the Jira should be considered for CR3.	1	1	https://issues.jboss.org/browse/JBPAPP-613
10551	JBPAPP-568	OK, fair enough. However, I know I would find it useful for the EAP source distro to include the source for every JBoss component included in the EAP binary. It seems a bit daft to have to download more source distros after having grabbed 193MB worth of EAP source.	1	1	https://issues.jboss.org/browse/JBPAPP-568
10552	JBPAPP-556	Hi Dimitris, is this applicable to AS 5? Should we open a JBAS for it?	0	0	https://issues.jboss.org/browse/JBPAPP-556
10553	JBPAPP-556	Can you please create and link to a forum thread to discuss the feature there, first?	1	0	https://issues.jboss.org/browse/JBPAPP-556
10554	JBPAPP-549	What's not clear? Try clicking on the linked task and you will see it was resolved in the JBM 1.4.0 SP3_CP02 release. ;)	1	1	https://issues.jboss.org/browse/JBPAPP-549
10555	JBPAPP-546	The changes for this has introduced a serialVersionUID mismatch on org.jboss.invocation.jrmp.server.JRMPInvokerHA serialVersionUID error for org.jboss.invocation.jrmp.server.JRMPInvokerHA, 405 -7657305823982668529, current: -2270982798654158900 Does this have an impact?	1	1	https://issues.jboss.org/browse/JBPAPP-546
10556	JBPAPP-534	Bill, can this patch be merged to CP branch by 01/22?	1	1	https://issues.jboss.org/browse/JBPAPP-534
10557	JBPAPP-532	Is there a reason why we need to add this test in a cumulative patch?	1	1	https://issues.jboss.org/browse/JBPAPP-532
10558	JBPAPP-527	This is a QA work, why am I assigned this?	0	0	https://issues.jboss.org/browse/JBPAPP-527
10559	JBPAPP-527	Juca, Can you check the version of the JDBC driver we test and address this JIRA?	1	1	https://issues.jboss.org/browse/JBPAPP-527
10560	JBPAPP-527	We test on this driver: Specification-Title: Oracle JDBC driver classes for use with JDK14 Sealed: true Created-By: 1.4.2_08 (Sun Microsystems Inc.) Implementation-Title: ojdbc14.jar Specification-Vendor: Oracle Corporation Specification-Version: Oracle JDBC Driver version - "10.2.0.2.0" Implementation-Version: Oracle JDBC Driver version - "10.2.0.2.0" Implementation-Vendor: Oracle Corporation Implementation-Time: Tue Jan 24 08:55:21 2006 There's also this other available in the lab, which is the latest as per Oracle's website: Specification-Title: Oracle JDBC driver classes for use with JDK14 Created-By: 1.4.2_14 (Sun Microsystems Inc.) Implementation-Title: ojdbc14_g.jar Specification-Vendor: Oracle Corporation Specification-Version: Oracle JDBC Driver version - "10.2.0.4.0" Implementation-Version: Oracle JDBC Driver version - "10.2.0.4.0" Implementation-Vendor: Oracle Corporation Implementation-Time: Sat Feb 2 11:53:00 2008 Can you please confirm which version should we use in the tests from now on? A more up-to-date version from the lab or 10.2.0.3 ?	1	1	https://issues.jboss.org/browse/JBPAPP-527
10561	JBPAPP-502	what's the status?	0	0	https://issues.jboss.org/browse/JBPAPP-502
10562	JBPAPP-502	Jimmy, is this change important enough to get done in CPs? Or can it wait till 5.0?	0	0	https://issues.jboss.org/browse/JBPAPP-502
10563	JBPAPP-488	When is this expected?	0	0	https://issues.jboss.org/browse/JBPAPP-488
10564	JBPAPP-413	Clebert, right. The HSQL clustered will be used for testing purposes. There will be instructions in the installation guide to configure a different db before using EAP. Coming to that, Clebert, can you point me to instructions for setting up a different db for clustered config? I saw the basic instructions on the messaging user guide. Is it the same for the clustered config as well?	1	1	https://issues.jboss.org/browse/JBPAPP-413
10565	JBPAPP-363	I am interested in configuring the JSessionID cookie name. However, I did not find the appropricate tag I can use from the jboss-web.xml. Could you tell me what I should use?	1	0	https://issues.jboss.org/browse/JBPAPP-363
10566	JBPAPP-360	Brian, the config files are not changed. Can you take a look at it? I am glad atleast this was brought to your attention now. The jgroups upgrade is already in place in branch JBPAPP_4_2. If you can change the config files, it will automatically be picked up in the next build for testing.	1	1	https://issues.jboss.org/browse/JBPAPP-360
10567	JBPAPP-360	Verified that bundling is enabled for all services except messaging. Is this intended? */deploy/jboss-messaging.sar/hsqldb-persistence-service.xml still has bundling disabled for both data and control channels.	1	1	https://issues.jboss.org/browse/JBPAPP-360

#	Requirement	Questions	Answered	Successfully Answered	URL
10568	JBPAPP-355	Yep Adrian wisely fixed the code and added a new version in the repository(1.3.4-jboss). The previous version was 1.3.4-brew. Oswego is not releasing any more builds because they are in maintenance. Matter of fact, these fixes weren't released as a package, but were available as individual files on the site. 1.3.4 is the last version that will ever be released from them. The original problem that Adrian fixed is that values, keySet, and entrySet toArray() methods were not thread safe. Adrian put a fix in there that used an iterator to build a secondary collection that we then call toArray with. This was the bug that was fixed. JBAS-4698. It doesn't look like that fix was ever moved into the cp. Actually there were a couple fixes in the oswego libraries that Adrian fixed, that didn't look like they were in EAP 4.2 CP/4.3. The version of the oswego libaries that was released for this fix(it went into 4.2.2) was 1.3.4-jboss(https://svn.jboss.org/repos/repository.jboss.org/oswego-concurrent/1.3.4-jboss/) The problem with the 1.3.4-jboss fix is that if an entry is removed, then it leaves a null in the map. When iterating through the collections the null would be propagated to the temporary collection being used to convert to an array. The end result was that when toArray() is called, it throws a null pointer exception. So the new version 1.3.4-jboss-update1 has that fix in there. So I put a check for a null in there for the values, keyset, and entry collections. This version of the oswego libraries that was just released is 1.3.4-jboss-update1.(https://svn.jboss.org/repos/repository.jboss.org/oswego-concurrent/1.3.4-jboss-update1). My plan is to change the build-thirdparty.xml file to reflect the new version. I also had to change the component-info.xml file in jbossxb version 1.0.0SP1-brew to reflect the compatibility with the new version of concurrent.jar.(http://repository.jboss.com/jboss/jbossxb/1.0.0.SP1-brew/component-info.xml) I had assumed that the same build process for EAP as AS, is that an incorrect assumption. The build-thirdparty.xml file holds the versions that are to be looked up in the repository. The version is then fetched from the repository and build. I have put a new version up in the repository. The build seems to run fine with the new entry in build-thirdpary. Does EAP build do the same thing? I had made the assumption that it did. Jay:)	1	1	https://issues.jboss.org/browse/JBPAPP-355
10569	JBPAPP-355	I provided the patch files to make a brew build. I also included the original source. They are at http://repository.jboss.com/oswego-concurrent/1.3.4-jboss-update1/src/ Why aren't the patch files enough. Do I have to attach this to this defect also? Cheers, Jay:)	1	1	https://issues.jboss.org/browse/JBPAPP-355
10570	JBPAPP-353	Ron, i see that the version of remoting has been upgraded to 2.2.2.SP1.patch01-brew for 4.2.0.GA_CP02 and 4.3.0.GA EAP. Is that the version you mean by "2.2.2.GA_CP01" above?	0	0	https://issues.jboss.org/browse/JBPAPP-353
10571	JBPAPP-353	By the way, what are the current freeze dates for EAP 4.3.0.GA and 4.2.0.GA_CP02?	1	1	https://issues.jboss.org/browse/JBPAPP-353
10572	JBPAPP-349	If a feature is needed, it must be discused in the forums with Adrian, first. Where is the discussion thread?	0	0	https://issues.jboss.org/browse/JBPAPP-349
10573	JBPAPP-349	is this really a blocker for the 4.3 release?	1	1	https://issues.jboss.org/browse/JBPAPP-349
10574	JBPAPP-343	Patrick, what do you think? This is for the RPM sets.	1	1	https://issues.jboss.org/browse/JBPAPP-343
10575	JBPAPP-334	Just create a 1.2.0.GA_CP01 release out of this branch. https://svn.jboss.org/repos/common/common-old/branches/JBossCommon_1_2_0_GA_CP/ The fix is already there.	1	1	https://issues.jboss.org/browse/JBPAPP-334
10576	JBPAPP-334	Should we upgrade to 1.2.1.GA for the EAP 4.3 BETA? Dimitris says it contains only this fix.	0	0	https://issues.jboss.org/browse/JBPAPP-334
10577	JBPAPP-332	So we need a JBossWS 1.2.1.GA_CP02 for EAP 4.2 CP01 and EAP 4.3 CP01 Could you please coordinate with the WS team so they incorporate the change you want?	1	1	https://issues.jboss.org/browse/JBPAPP-332
10578	JBPAPP-332	Darran, is JBossWS 1.2.1.GA_CP branch ready to be tagged? We need WS build integrated into EAP 4.2.0_CP02 by 01/22.	0	0	https://issues.jboss.org/browse/JBPAPP-332
10579	JBPAPP-330	I didn't get the problem, because Seam 1.2.1 uses the commons-beanutils-1.7.0.jar without the commons-collections classes in it. So what's the problem?	1	1	https://issues.jboss.org/browse/JBPAPP-330
10580	JBPAPP-330	Hi Marek, here is my proposal: We fix the <componentref name="apache-beanutils" version="1.7.0-brew"/> so it does not include the collections classes. They should not be thre in the first palce (i.e. we should be uploading commons-beanutils-bean-collections-1.7.0.jar not commons-beanutils-1.7.0.jar, even if we don't use the funky Apache naming of that JAR). This will not affect the AS as it explicitly includes <componentref name="apache-collections" version="3.1-brew"/> so it won't miss the removed classes. What I am concerned now is that the spurious collection classes in beanutils _in theory_ are the same ones from our collections 3.1-brew, so Seam should be seeing the same malfunction. It doesn't matter if the classes come from the same JAR as beanutils or from a separate JAR. Which common-collections does Seam use after all? Is there a special one that does not have the problem that is different from the one in the AS? How do you ensure these are used and not the ones already loaded by the container? Regards, Fernando	1	1	https://issues.jboss.org/browse/JBPAPP-330
10581	JBPAPP-330	Jon, can you flip the EAP branch of JBoss AS to use your new clean beanutils? Make sure it works first, of course. Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP-330
10582	JBPAPP-318	:(Don't forget http://jira.jboss.com/jira/browse/JBSEAM-1059 then too. Who is going to do this back port ? I really hope you checked that you can actually backport it safely without all the other fixes in Seam 1.2.1.	0	0	https://issues.jboss.org/browse/JBPAPP-318

#	Requirement	Questions	Answered	Successfully Answered	URL
10583	JBPAPP-314	1. Most of the linked Remoting issues are related somehow to JBossMessaging, in the sense that they were related to JBM test failures, or were simply pointed out by Tim Fox or Clebert Suconic. These are: JBREM-641 "re-implement the callback polling for http transport to reduce latency" JBREM-743 "For polling callback handler, org.jboss.remoting.Client.addListener() should create only one CallbackPoller per InvokerCallbackHandler? JBREM-747 "org.jboss.remoting.transport.Connector should unregister server invoker from MBeanServer" JBREM-751 "Eliminate unnecessary "Unable to process control connection:" message from BisocketServerInvoker: JBREM-754 "Reset timeout on each use of HttpURLConnection" JBREM-755 "Make ConnectorValidator parameters configurable" JBREM-756 "CallbackPoller should shut down if too many errors occur." JBREM-757 "Implement quick Client.removeListener() for polled callbacks." JBREM-761 "NPE in BisocketServerInvoker$ControlConnectionThread" JBREM-765 "Add a separate timeout parameter for callback clients" JBREM-766 "Guard against "spurious wakeup" from Thread.sleep()" JBREM-767 "Avoid deadlock in callback BisocketClientInvoker when timeout == 0" JBREM-774 "BisocketClientInvoker.replaceControlSocket() and handleDisconnect() should close control socket" JBREM-775 "MicroSocketClientInvoker.initPool() should omit pool from log message" JBREM-778 "BisocketServerInvoker.start() creates a new static Timer each time" JBREM-782 "Remove network i/o from synch block in ServerInvokerCallbackHandler.getCallbackHandler()" JBREM-783 "Remove network i/o from synch blocks that establish and terminate LeasePingers" JBREM-784 "Use separate maps for control sockets and ordinary sockets in BisocketClientInvoker" JBREM-785 "BisocketClientInvoker.transport() inadvertently uses listenerId member variable" JBREM-787 "Move network i/o in BisocketClientInvoker constructor to handleConnect()" JBREM-788 "Access to BisocketClientInvoker static maps should be synchronized in handleDisconnect()" JBREM-790 "NPE in BisocketClientInvoker$PingTimerTask" JBREM-793 "Lease should synchronize access to client map" 2. Two of the issues come from support cases: JBREM-781 "Socket transport needs to provide to the client local address of a TCP/IP connection, as seen from the server" JBREM-791 "In HTTPClientInvoker remove newlines and carriage returns from Base64 encoded user names and passwords" 3. One issue I found in some other context, but since it involves a possible socket leak, it's probably worth fixing: JBREM-771 "MicroSocketClientInvoker can experience socket leaks"	1	1	https://issues.jboss.org/browse/JBPAPP-314
10584	JBPAPP-311	The one off patch JBPAPP-264 appears to have been incorrectly noted as fixed in EAP 4.2 CP01. This JIRA was for incorporating the one-off patch into EAP 4.2 CP01 (based on the description). JBPAPP-264 should probably be removed from the EAP 4.2 CP01 change log and this one added. Based on the CP-related comments from Shelly McGowain in the one-off patch JIRA (this was apparently tested in CP01) and the fact that JBPAPP-314 was included (JBoss Remoting 2.2.2 GA contained this fix), I'd say we can close/resolve this issue after we get it in the right place. I'll mark JBPAPP-314 as having incorporated this change. I'll also resolve/close it. What shall we do for the fix version of JBPAPP-264 given that it was a one-off patch?	0	0	https://issues.jboss.org/browse/JBPAPP-311
10585	JBPAPP-304	Ok, sounds great. So besides from 1.2.1.GA already being available in SVN what/who else needs to do anything ?	0	0	https://issues.jboss.org/browse/JBPAPP-304
10586	JBPAPP-303	Move to CP02, as per Remy's complexity analisys?	1	1	https://issues.jboss.org/browse/JBPAPP-303
10587	JBPAPP-297	no need - just need to get them online	0	0	https://issues.jboss.org/browse/JBPAPP-297
10588	JBPAPP-278	comments from the duplicate ASPATCH issue that has been deleted: Darran Lofthouse [15/Aug/07 05:08 AM - Visible to JBoss Employee] Assigning to patch triage as this is technically a feature request. A number of customers have issues when they are attempting to use WSDL to Java with a WSDL that contains bound headers defined. In JBossWS 1.2.1.GA wstools just ignores these headers so the generated SEI does not reference the header type and the header types themselves are not generated. In the opposite direction Java to WSDL does support bound headers. The inclusion of this enhancement does not alter other scenarios for WSDLs that did not work they now start working. Non inclusion of this enhancement may be a barrier to some customers and could stop them upgrading to the platform, especially as we do not support JAX-WS. Darran Lofthouse [15/Aug/07 05:08 AM - Visible to JBoss Employee] Assigning to patch triage as this is technically a feature request. A number of customers have issues when they are attempting to use WSDL to Java with a WSDL that contains bound headers defined. In JBossWS 1.2.1.GA wstools just ignores these headers so the generated SEI does not reference the header type and the header types themselves are not generated. In the opposite direction Java to WSDL does support bound headers. The inclusion of this enhancement does not alter other scenarios for WSDLs that did not work they now start working. Non inclusion of this enhancement may be a barrier to some customers and could stop them upgrading to the platform, especially as we do not support JAX-WS. Ryan Campbell [16/Aug/07 05:50 PM] Are there test cases which will validate this new functionality? [Show »] Ryan Campbell [16/Aug/07 05:50 PM] Are there test cases which will validate this new functionality? Darran Lofthouse [16/Aug/07 05:54 PM] Every issue logged will be committed upstream and will either have it's own test case implemented specifically for it the only exception is if the issue can be reproduced with an existing test case or if an existing test case can be enhanced. So yes there is a set of new tests for this, there are a couple of minor dependent feature requests needed for this that also have their own test cases. [Show »] Darran Lofthouse [16/Aug/07 05:54 PM] Every issue logged will be committed upstream and will either have it's own test case implemented specifically for it the only exception is if the issue can be reproduced with an existing test case or if an existing test case can be enhanced. So yes there is a set of new tests for this, there are a couple of minor dependent feature requests needed for this that also have their own test cases. Ryan Campbell [17/Aug/07 03:43 PM] Since it sounds well-covered by tests, QE doesn't have any issue with this.	1	0	https://issues.jboss.org/browse/JBPAPP-278
10589	JBPAPP-266	We would need a jbossws 1.2.1.GA_CP01 uploaded to the repository.jboss.com so we could switch EAP AS 4.2.0.GA_CP01 to use it. Thomas, when could you get a build uploaded? The EAP AS CP01 freeze date is Aug 24th. Regards.	1	1	https://issues.jboss.org/browse/JBPAPP-266
10590	JBPAPP-266	Darran, What's ETA when all of these WS fixes will be ported to the CP branch? Some of the related JBWS issues have to reference to salesforce cases.	0	0	https://issues.jboss.org/browse/JBPAPP-266
10591	JBPAPP-266	thanks. can we get ETA for all 13 remaining WS issues? Is Aug 24th still feasible?	0	0	https://issues.jboss.org/browse/JBPAPP-266
10592	JBPAPP-251	so who would know ? :)	0	0	https://issues.jboss.org/browse/JBPAPP-251
10593	JBPAPP-233	can we close this issue now?	1	1	https://issues.jboss.org/browse/JBPAPP-233
10594	JBPAPP-233	are we going to be having HTML versions of the Arjuna docs?	1	1	https://issues.jboss.org/browse/JBPAPP-233
10595	JBPAPP-230	We will use either https://svn.jboss.org/repos/repository.jboss.org/jboss/remoting/2.2.0.SP4_CP02-brew/ or a Remoting 2.2.2.GA release, both which include these fixes. Ron, the code freeze for EAP AS is next week, how soon can you have a 2.2.2.GA (which also would support the JBM release) ready? P.S.: If possible we would like to propose that 2.2.2.GA be used in the EAP AS 4.2 CP01, so we would apply for the patch triage process to authorize that, if that could be out in time.	0	0	https://issues.jboss.org/browse/JBPAPP-230
10596	JBPAPP-230	who is getting these into remoting 2.2.2?	1	1	https://issues.jboss.org/browse/JBPAPP-230

#	Requirement	Questions	Answered	Successfully Answered	URL
10597	JBPAPP-229	Hi Carlo, who is going to produce a patch for the CP01?	1	1	https://issues.jboss.org/browse/JBPAPP-229
10598	JBPAPP-227	Jimmy, can we close this one?	0	0	https://issues.jboss.org/browse/JBPAPP-227
10599	JBPAPP-225	Should the appender reference be ASYNC like it is for the root element, or FILE as it is by default?	0	0	https://issues.jboss.org/browse/JBPAPP-225
10600	JBPAPP-225	Can you attach the latest jboss-log4j.xml so that I can have a look? I am fairly certain that org.jboss.ha should not have any appenders referenced (it gets them based on inheriting root), but I want to have a look to be sure we're on the same page.	0	0	https://issues.jboss.org/browse/JBPAPP-225
10601	JBPAPP-223	What is the impact of this appender not being commented? Does it affect performance or something else?	0	0	https://issues.jboss.org/browse/JBPAPP-223
10602	JBPAPP-223	Why was this reopened? The root category now uses ASYNC, as per Andy's comment....	1	1	https://issues.jboss.org/browse/JBPAPP-223
10603	JBPAPP-213	oh ? so the installer and the zip generated different content ? bummer. Which version should we (RHDS) use ? (note, currently we use the non-secured zip to allow the IDE to more easily start and run the server)	1	0	https://issues.jboss.org/browse/JBPAPP-213
10604	JBPAPP-199	JBossTS 4.2.3.SP4 is passed the JBossTS QA and available on the repository. I'm seeing some JBossAS testsuite failures with this version, but they may be unrelated. Is there a server from which I can get baseline testsuite results to compare? If not I'll have to run again on my (slow) machine.	1	1	https://issues.jboss.org/browse/JBPAPP-199
10605	JBPAPP-199	There are 2 known failures. What are you getting?	1	1	https://issues.jboss.org/browse/JBPAPP-199
10606	JBPAPP-196	why is this better to move *off* of a public mailing list where others can contribute information? ;) the problem is that sybase apparently makes no distinction between permanent and temporary tables. Why should one not be able to create temporary tables within a transaction? How does such a restriction make sense?	1	0	https://issues.jboss.org/browse/JBPAPP-196
10607	JBPAPP-196	Well I didn't want to talk about platform on a public mailing list. About sense and sybase... you know ;) Please tell us if it is possible to make hibernate not use ddl in transactions even create temporary tables. If it is too hard I guess it's late for the platform.	0	0	https://issues.jboss.org/browse/JBPAPP-196
10608	JBPAPP-196	Steve, the code freeze for AS 4.2 CP01 is Augist 24th. Do you have an ETA for this fix?	0	0	https://issues.jboss.org/browse/JBPAPP-196
10609	JBPAPP-196	Does this need release notes? If so, can you provide something in the format: Cause: Consequence: Fix: Result: Thanks!	1	0	https://issues.jboss.org/browse/JBPAPP-196
10610	JBPAPP-196	Strong, is this for EAP 5.x? If so, please move it to TBD EAP 5. Thanks.	1	0	https://issues.jboss.org/browse/JBPAPP-196
10611	JBPAPP-196	Anne, this issue was resolved as won't fix, Misty reopened it just to correct the release notes. since this issue affects to both eap4 & 5, i'd think set it to tbd would better, what do you think?	0	0	https://issues.jboss.org/browse/JBPAPP-196
10612	JBPAPP-195	Josh, is this being tracked in BZ? Was this addressed in time for RC6?	0	0	https://issues.jboss.org/browse/JBPAPP-195
10613	JBPAPP-195	Is this meant to go into the docs? I understood that this was just being jotted down for future reference internally. Do you want me to put this into the release notes for the platform?	0	0	https://issues.jboss.org/browse/JBPAPP-195
10614	JBPAPP-195	Well these are configuration options set for the test suite. Most of them are just optional. I think that having notes about options that users could have to change will improve user experience. . On https://wiki.corp.jboss.com/bin/view/IT/SystemServiceQaLab are described the database installations though this particular information is not always easy to extract. Did you included info from that url or just the options mentioned in the issue description?	0	0	https://issues.jboss.org/browse/JBPAPP-195
10615	JBPAPP-187	i.e. JBoss EAP 4.2.0.RC5 for 05/22 build Ram, are you ok with proposed version string?	0	0	https://issues.jboss.org/browse/JBPAPP-187
10616	JBPAPP-178	I added the following comments and left the old type-mappings in place instead of duplicating them or changing the type-mapping names which might break the current applications in place. Updated the following: 1 & 2) <type-mapping> <name>Oracle9i</name> <!-- | This type-mapping applies both to Oracle 9i and Oracle 10g | Make sure that you have the latest Oracle 10g version of ojdbc14.jar --> 3) <type-mapping> <name>MS SQLSERVER2000</name> <!-- | This type-mapping applies both to MS SQLSERVER 2000 and MS SQLSERVER 2005 --> 4.) <type-mapping> <name>DB2</name> <!-- | This type-mapping applies to DB2 V7.2 and V8.0 --> 7.) <type-mapping> <name>PostgreSQL 8.0</name> <!-- | This type-mapping applies to PostgreSQL V8.0 and V8.1 -->	0	0	https://issues.jboss.org/browse/JBPAPP-178
10617	JBPAPP-121	What's the status of this?	0	0	https://issues.jboss.org/browse/JBPAPP-121
10618	JBPAPP-120	Rajesh, can you verify in Beta1.	1	0	https://issues.jboss.org/browse/JBPAPP-120
10619	JBPAPP-120	I fixed this in Branch_4_2. Is there a separate source tree for Beta1?	0	0	https://issues.jboss.org/browse/JBPAPP-120
10620	JBPAPP-114	Can't the classpath for VersionRelease.java be modified to point to server/all/lib/*.jar as well?	0	0	https://issues.jboss.org/browse/JBPAPP-114
10621	JBPAPP-114	update on alpha6: svn vs. zip Following files still differ: server/{all | default}/deploy/jmx-console.war/WEB-INF/web.xml svn: <param-value>UDP(ip_mcast=true;ip_ttl=16;loopback=false;mcast_addr=${jboss.partition.udpGroup:228.1.2.3};mcast_port=45566): org.jboss.jmx.adaptor.control.FindView zip: <param-value>UDP(ip_mcast=true;ip_ttl=16;loopback=false;mcast_addr=${jboss.partition.udpGroup:228.1.2.3};mcast_port=${jboss.partition.udpPort:45566}): org.jboss.jmx.adaptor.control.FindView bin/run.conf: only in zip distro: if ["`uname`" = "Linux"]; then JAVA_OPTS="$JAVA_OPTS -Djava.net.preferIPv4Stack=true" fi Who should apply these changes in the svn branch? varia/src/resources/jmx/jmx-console-web.xml system/src/bin/run.conf and Branch_4_2/system/src/bin/run.conf ---------- This file differs: build/build-thirdparty.xml only in svn: <visit-componentref-graph componentVisitor="org.jboss.ant.util.graph.ComponentRefGraphLicenseVisitor"/> ---------- No differences in jar files.	0	0	https://issues.jboss.org/browse/JBPAPP-114
10622	JBPAPP-114	regarding svn vs src: Difference in */ServerConfig.class is OK and intentional. It is needed to make Difference in _ddsModule.class seems to come from rebuilt jacorb, which is okay. By RC6, the jacorb versions will be re-synched. For the xml/sh/properties files, all are okay except for the last 2: varia/src/resources/services/invoker/http/web.xml varia/src/resources/services/invoker/httpha/web.xml What differs in these? regarding source zip vs .zip: All is OK except for these: build/output/jbossas/docs/examples/jms/mssql-jdbc2-service.xml build/output/jbossas/docs/examples/jms/standalone/build.xml I don't understand what the "?" is doing there in the first place, and what is adding the other 2 question marks. Does anyone know what purpose they serve?	1	1	https://issues.jboss.org/browse/JBPAPP-114
10623	JBPAPP-114	How is "SOURCE-ZIP-distro" being built? For the zip/vs rpm: bin/wsconsume.sh bin/wsrunclient.sh bin/wsprovide.sh have been added to rpm	1	1	https://issues.jboss.org/browse/JBPAPP-114
10624	JBPAPP-114	Are these differences even after the as release checklist is followed to build from source zip?	1	1	https://issues.jboss.org/browse/JBPAPP-114

#	Requirement	Questions	Answered	Successfully Answered	URL
10625	JBPAPP-114	RC7: svn vs. src-zip no unexpected diffs in jars: lib/jboss-system.jar at org/jboss/system/server/ServerConfig.class (expected) client/jboss-system-client.jar at org/jboss/system/server/ServerConfig.class (expected) client/jbossall-client.jar at org/jboss/system/server/ServerConfig.class (expected) bin/run.jar at org/jboss/system/server/ServerConfig.class (expected) server/all/deploy/management/console-mgr.sar/web-console.war/applet.jar at org/jboss/system/server/ServerConfig.class (expected) server/default/deploy/management/console-mgr.sar/web-console.war/applet.jar at org/jboss/system/server/ServerConfig.class (expected) Have these been clarified? diff svn/system/src/bin/run.bat src-zip/system/src/bin/run.bat 82,84c82,83 < rem Add -server to the JVM options, if supported < "%JAVA%" -version 2>&1 | findstr /I hotspot > nul < if not errorlevel == 1 (set JAVA_OPTS=%JAVA_OPTS% -server) --- > rem Add -server to the JVM options > set JAVA_OPTS=%JAVA_OPTS% -server diff svn/system/src/bin/run.sh src-zip/build/jboss-as/system/src/bin/run.sh 13a14,30 > # Set conf if specified, else set to default > JBOSSCONF="production" > > arg_count=1 > eval SWITCH=\$$arg_count > while [! -z $SWITCH] > do > > if ["$SWITCH" = "-c"]; then > eval JBOSSCONF=\$`expr $arg_count + 1` > break > fi > > arg_count=`expr $arg_count + 1` > eval SWITCH=\$$arg_count > done > 52c69,74 < RUN_CONF="$DIRNAME/run.conf" --- > > if [! -z "$JBOSSCONF"] && [-f "$DIRNAME/../server/$JBOSSCONF/run.conf"]; then > RUN_CONF="$DIRNAME/../server/$JBOSSCONF/run.conf" > else > RUN_CONF="$DIRNAME/run.conf" > fi 150d171 < if ["x$HAS_HOTSPOT" != "x"]; then 155d175 < fi	0	0	https://issues.jboss.org/browse/JBPAPP-114
10626	JBPAPP-114	RC8: Differences between SOURCE-ZIP-distro and ZIP-distro (with-auth): diff in jar: client/jbossall-client.jar - the with-auth zip distro has follwoing empty directories: Only in org/jboss/iiop/rmi: ir Only in org/jboss/invocation: local Only in org/jboss/security: jacc Only in org/jboss/security: jce Only in org/jboss/security: propertyeditor Other diffs: Only in src-zip server/lib server/deploy server/deploy-ha Shouldn't this go into the production dir? The production dir isn't present in the src-zip. Only in source zip: seam/* same as in RC7	0	0	https://issues.jboss.org/browse/JBPAPP-114
10627	JBPAPP-114	Deepak, are we all set for a code freeze (CP01) on Aug 24th? Do we still have any issues BTW, or is this just a placeholder so we remember to check this at each release?	1	0	https://issues.jboss.org/browse/JBPAPP-114
10628	JBPAPP6-1741	Rostislav Svoboda <rsvoboda@redhat.com> made a comment on [bug 901281|https://bugzilla.redhat.com/show_bug.cgi?id=901281] Jim, can you share some details about changes, mainly motivation for adding BindingProvider properties? Tank you.	0	0	https://issues.jboss.org/browse/JBPAPP6-1741
10629	JBPAPP6-1692	Permaine Cheung <pcheung@redhat.com> made a comment on [bug 901242|https://bugzilla.redhat.com/show_bug.cgi?id=901242] Brian, can you please look into this? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP6-1692
10630	JBPAPP6-1680	Jimmy Wilson <jawilson@redhat.com> made a comment on [bug 901231|https://bugzilla.redhat.com/show_bug.cgi?id=901231] @Stuart: If this is the only way to address this, is a new property necessary? There are standard JVM properties that we could use to identify an IBM JVM and then act accordingly? Are there any IBM JVMs without this issue? My understanding is that all of them have it.	0	0	https://issues.jboss.org/browse/JBPAPP6-1680
10631	JBPAPP6-1680	John Doyle <jdoyle@redhat.com> made a comment on [bug 901231|https://bugzilla.redhat.com/show_bug.cgi?id=901231] We need a plan to close this BZ. We have lists of the the occurrences of the finalize in methods in the EAP 6, can we make that list more accessible and enable people to work on the occurrences as time permits, and reduce the list that way?	0	0	https://issues.jboss.org/browse/JBPAPP6-1680
10632	JBPAPP6-1626	@David, can you comment on this.	0	0	https://issues.jboss.org/browse/JBPAPP6-1626
10633	JBPAPP6-1615	Rosta, what's the reason to remove this from the Mead repository? Is it causing some problem?	0	0	https://issues.jboss.org/browse/JBPAPP6-1615
10634	JBPAPP6-1612	Clebert Suconic <csuconic@redhat.com> made a comment on [bug 900080|https://bugzilla.redhat.com/show_bug.cgi?id=900080] Do we really need that? We already have a flow control property per consumer.	0	0	https://issues.jboss.org/browse/JBPAPP6-1612
10635	JBPAPP6-1595	Hi Fernando, {quote} we do not modify upstream builds, so if they produce -sources.jar we include them if not we don't. {quote} that is basically what I am asking for :) We don't need to rely on the upstream build process. When we already have the sources available, why don't we just zip them, and upload as an artifact? 3 lines in bash, isn't it?	1	0	https://issues.jboss.org/browse/JBPAPP6-1595
10636	JBPAPP6-1595	Unfortunately is not usually that simple, these weird builds manipulate sources in crazy ways, they are not like well behaved maven builds that follow a convention. But we can try adding that for a few selected ones and see if we succeed (we tried in the past and failed). However, I inspected the 55 missing messages. There are several false positives for things like resources, test artifacts, console export etc. There are duplicates as well, some of the commons appear more than once I don't know why. All the rest are for Apache old projects, things like commons bits, dom4j, jdom, veocity, xom for which I am not sure we should be distributing APIs, maybe Jacorb (which is in there) but that probably requires separating APIs that we should and that we should not ship, which ain't going to be easy. Can you take a look and paste here a list of what you really need? And the smaller the list the easier if becomes for us to try some hack in the ant builds. P.S.: The list of things public and private is a bit fishy, I still cannot understand (nor does Rich Sharples) why we don't want to ship the RESTEasy javadocs. And I myself have serious doubts about us shipping APIs of these old Apache libraries. This sounds like trouble.	1	1	https://issues.jboss.org/browse/JBPAPP6-1595
10637	JBPAPP6-1595	{quote}But we can try adding that for a few selected ones and see if we succeed (we tried in the past and failed).{quote} Sounds good. And I will (try to) prepare a list of those which are missing to generate the javadoc for EAP. {quote}However, I inspected the 55 missing messages. There are several false positives for things like resources, test artifacts, console export etc. There are duplicates as well, some of the commons appear more than once I don't know why.{quote} Good, could you please consult the details with the maintainer of the job, Rostislav Svoboda - rsvoboda@redhat.com? {quote} All the rest are for Apache old projects, things like commons bits, dom4j, jdom, veocity, xom for which I am not sure we should be distributing APIs, maybe Jacorb (which is in there) but that probably requires separating APIs that we should and that we should not ship, which ain't going to be easy. {quote} Yep, some of them are not needed for javadoc. But for other purposes, like debugging, having them is, say, convenient, at least. Resteasy - No idea. Apache libs and such - +1, we should consider wisely what to include as public.	1	0	https://issues.jboss.org/browse/JBPAPP6-1595
10638	JBPAPP6-1595	Hmm, I can imagine how the jar's are assembled, I've seen e.g. Tomcat 5.x build. The problem is that JavaDoc plugin fails if some -sources.jar is not present for any dependency. IIRC this can't be overcame, but I will check once we have the JBoss bits. Could we perhaps create the source jars just for the EAP release? Not for CVEs etc. Just for the sake of JavaDoc. Probably not worth the special effort, just asking if it could be done once per minor release, along the way with your team's normal work.	1	0	https://issues.jboss.org/browse/JBPAPP6-1595
10639	JBPAPP6-1586	Has this fix been applied to both EAP and community branches?	0	0	https://issues.jboss.org/browse/JBPAPP6-1586
10640	JBPAPP6-1586	I'm not sure how this got into the EAP6 ER4 bin. Moving this to TBD until it is clear what is required. Tom, do you want this fix in EAP6 as well? The code changes for this issue show commits only on the EAP5 branch.	1	0	https://issues.jboss.org/browse/JBPAPP6-1586

#	Requirement	Questions	Answered	Successfully Answered	URL
10641	JBPAPP6-1585	You know, this whole discussion is telling. Once upon a time JBoss was all about being the best opensource appserver available. Now it's "we properly implement open specifications". Is that all you aspire to? Sad.	1	0	https://issues.jboss.org/browse/JBPAPP6-1585
10642	JBPAPP6-1585	My frustration with this issue has nothing to do with Jon. Furthermore, you can be sure that many other longtime JBoss users are feeling this same frustration. This bug has been in place for five years. Regardless of what the spec says, this behavior is fundamentally broken and seriously impedes the ability to use J2EE container managed security to develop applications. For example, the piece of software we're using *right now* - JIRA - could not use container managed security because of this bug prevents the most stupidly fundamental concepts on the web: * A page that renders differently depending on whether you are logged in or not. Yes, there are probably other issues as well, but this is one of the more serious ones. So, five years ago I was frustrated with a deficiency that is _glaring_ and _obvious_ to anyone who actually uses this technology to develop end-user websites. There were (and probably still are) numerous complaints about it in the forums. It (and a zillion other problems with JBoss of varying degrees of severity) sat neglected for five years. I gave up. Despite building a big chunk of my professional career since 2002 on JBoss, I threw away all my hard-fought knowledge (accumulated months of stepping through JBoss code in the debugger) and moved on to other platforms (Resin and Appengine). I'm bitter about this and I have a right to be; as an advocate of JBoss (don't think it was easy to convince EA or Kink.com) it sucks to shrug my shoulders and say "sorry guys, this platform is stagnant and you need to move on". So when I suddenly find an email in my box that says "We now have an official excuse for this bug: EE6 does not require us to fix our broken shit", it pulls the scab off the wound. You're getting both barrels because you said it. I've earned the right to be pissed with my blood and the subscription dues from two major corporations. Want to reverse JBoss' slow slide into obscurity? I hate to sound like the 15-pieces-of-flare guy from Office Space, but do you really just want to do the bare minimum? It's one thing to say "we can't fix this because of performance problems" or "we don't have the resources to fix this" but to close the bug as invalid just because the spec doesn't absolutely require that it be fixed is LAME.	0	0	https://issues.jboss.org/browse/JBPAPP6-1585
10643	JBPAPP6-1577	Strong, Does the hibernate team consider this to be a bug or a feature request?	0	0	https://issues.jboss.org/browse/JBPAPP6-1577
10644	JBPAPP6-1575	To be safe and consistent, can we just change the version suffix to .redhat-1 across the board, instead of -redhat-1?	0	0	https://issues.jboss.org/browse/JBPAPP6-1575
10645	JBPAPP6-1575	Unfortunately that would also create invalid versions because OSGi does not allow "." in the qualifier. 1.0.0.Final -> 1.0.0.Final.redhat-1 (invalid OSGi version)	1	1	https://issues.jboss.org/browse/JBPAPP6-1575
10646	JBPAPP6-1574	{quote} I think that's a great idea! I would love to see a quickstart that does the following: * Provides 2 login roles: administrator and regular user * If you log in as administrator, it takes you to a "Welcome Administrator" page. * If you log in as quickstartUser/quickstartPassword (this is the id used in the other quickstarts), it takes you to a "Welcome User" page. * It should also prevent you from accessing the Admin page when you log in as the regular user. * It should also provide a logout button so you can log in again as the other user. Are you volunteering? ;-) Thanks, Sande {quote}	1	1	https://issues.jboss.org/browse/JBPAPP6-1574
10647	JBPAPP6-1570	baranowb <bbaranow@redhat.com> made a comment on [bug 900560|https://bugzilla.redhat.com/show_bug.cgi?id=900560] Anne this seems to be fixed in 6.1. Is there any urge to have this fixed in 6.0.1 ?	0	0	https://issues.jboss.org/browse/JBPAPP6-1570
10648	JBPAPP6-1568	Misha H. Ali <mhusnain@redhat.com> made a comment on [bug 745912|https://bugzilla.redhat.com/show_bug.cgi?id=745912] Is this marked for Release Notes? If so, it does not need to be assigned to the Documentation component, the requires_release_notes? flag is sufficient.	0	0	https://issues.jboss.org/browse/JBPAPP6-1568
10649	JBPAPP6-1479	Rostislav Svoboda <rsvoboda@redhat.com> made a comment on [bug 900403|https://bugzilla.redhat.com/show_bug.cgi?id=900403] Peter, what's current status of this item? Is it incorporated in EAP 6.1.0 DRx ?	0	0	https://issues.jboss.org/browse/JBPAPP6-1479
10650	JBPAPP6-1479	Peter Skopek <pskopek@redhat.com> made a comment on [bug 900403|https://bugzilla.redhat.com/show_bug.cgi?id=900403] (In reply to comment #10) > Peter, what's current status of this item? Is it incorporated in EAP 6.1.0 > DRx ? It already there. run valut.sh --help to see all options. DR4 contains the change. There is also Java class which could be used from inside tests, so no need to call script from OS level. Quick example: I have tried to use API on my Vault and it works fine: public class Test { public static void main(String[] args) throws Exception { VaultSession s = new VaultSession("/home/pskopek/dev/as7/vault.jks", "password1234", "/home/pskopek/dev/as7/vault", "12345678", 23); s.startVaultSession("vault"); String attribute = s.addSecuredAttribute("DS", "password", "fero".toCharArray()); System.out.println(attribute); System.out.println(s.getKeystoreMaskedPassword()); } } Output: VAULT::DS::password::ODY4MmU5ZWEtNWNmYi00ZjUwLWFhNzgtYzJhMTdjOTZiOWM5TElORV9CUkVBS3ZhdWx0 MASK-3wLlsUw/wDupvXfU6WrXcZ CLI: [standalone@localhost:9999 /] /core-service=vault:add(vault-options=[("KEYSTORE_URL"=>"/home/pskopek/dev/as7/vault.jks"), ("KEYSTORE_PASSWORD" => "MASK-3wLlsUw/wDupvXfU6WrXcZ"), ("KEYSTORE_ALIAS" => "vault"), ("SALT" => "12345678"), ("ITERATION_COUNT" => "23"), ("ENC_FILE_DIR" => "/home/pskopek/dev/as7/vault/")]) {"outcome" => "success"}	0	0	https://issues.jboss.org/browse/JBPAPP6-1479
10651	JBPAPP6-1456	2. Who is the customer behind the request? Account name: Adam Halbardier Customer segment: 3 TAM/SRM customer: no Strategic Customer yes/no: no 3. What is the nature and description of the request? The customer is searching for the way to make JBoss EAP 6 FIPS 140-2 compliant. To use the NSS library is the only way which he found [1] and which is applicable for EAP 5. He tried it on EAP 6, but failed. The NSS library requires a PKCS#11 keystore, by contrast, a "keystore-type" attibute of JBoss EAP 6 configuration can not be "PKCS11", but "PKCS12" or "JKS", and JBoss EAP 6 checks whether a keystore would be a valid PKCS#12 keystore or a valid JKS keystore and recognizes PKCS#11 keystore as an invalid keystore. Therefore, the customer can not configure JBoss EAP 6 to be FIPS 140-2 compliant. [1] https://access.redhat.com/knowledge/solutions/42301 4. Why does the customer need this? (List the business requirements here) Because it will not be possible for the customer to procure and use JBoss EAP in production without an acceptable solution for FIPS 140-2 compliance. 5. How would the customer like to achieve this? (List the functional requirements here) It is a straightforward way to make EAP 6 to handle PKCS#11 keystore and use NSS library as EAP 5 can. Also, the other way is acceptable for the customer, if it exists. 6. For each functional requirement listed in question 5, specify how Red Hat and the customer can test to confirm the requirement is successfully implemented. If you take the way that the NSS library can be enabled, The steps to test is as follow:- 1. confirm if SSL/TSL works 2. confirm if a PKCS#11 keystore is used 3. confirm if the NSS library is used If you take the other way to fix this issue, need some discussions. 8. Does the customer have any specific timeline dependencies? The customer is currently evaluating JBoss EAP 6 in anticipation of migrating to the platform in the next few months.	0	0	https://issues.jboss.org/browse/JBPAPP6-1456

#	Requirement	Questions	Answered	Successfully Answered	URL
10652	JBPAPP6-1455	Peter Skopek <pskopek@redhat.com> made a comment on [bug 900887|https://bugzilla.redhat.com/show_bug.cgi?id=900887] Shay, was EAP6 web subsystem configured to use SSO? More info could be found at https://community.jboss.org/wiki/HowToUsePicketLinkIDPInClusteredEnvironment	0	0	https://issues.jboss.org/browse/JBPAPP6-1455
10653	JBPAPP6-1452	Matteo Mortari <matteo.mortari@gmail.com> made a comment on [bug 901070|https://bugzilla.redhat.com/show_bug.cgi?id=901070] Hello this was raised as Type: Feature Request and now CLOSED NOTABUG without a motivation. Can you kindly explain why, please? Thanks Ciao	0	0	https://issues.jboss.org/browse/JBPAPP6-1452
10654	JBPAPP6-1451	Guys how can you see this as a feature request? It has been available in the past (EAP4, EAP5), it is a common use case, it is not in EAP6 it was simply forgotten to put in...	0	0	https://issues.jboss.org/browse/JBPAPP6-1451
10655	JBPAPP6-1445	Rostislav Svoboda <rsvoboda@redhat.com> made a comment on [bug 900788|https://bugzilla.redhat.com/show_bug.cgi?id=900788] Is it still valid or can it be closed ?	0	0	https://issues.jboss.org/browse/JBPAPP6-1445
10656	JBPAPP6-1441	Heiko Braun <hbraun@redhat.com> made a comment on [bug 900675|https://bugzilla.redhat.com/show_bug.cgi?id=900675] Maybe some further clarification: try to answer this question yourself: What's are the drivers associated with selected profile?	0	0	https://issues.jboss.org/browse/JBPAPP6-1441
10657	JBPAPP6-1441	Bruno Machado <bmachado@redhat.com> made a comment on [bug 900675|https://bugzilla.redhat.com/show_bug.cgi?id=900675] JDBC Driver is a module, thus it's available to every profile. Sometimes the user doesn't want to start the server to configure a datasource. For instance, during a management schedule. User usually stop the server, change what he needs and start it again. With the current implementation, the JDBC driver will not be available if the server with the associated profile with the driver metadata is not up. Maybe it's not a UI issue, maybe it's the way AS7/EAP6 deals with datasources and drivers since its configurations are associated to a subsystem in a profile. It could be associated to a server settings in host.xml or to a server-group, not to a specific profile. Does it make more sense?	0	0	https://issues.jboss.org/browse/JBPAPP6-1441
10658	JBPAPP6-1337	Rostislav Svoboda <rsvoboda@redhat.com> made a comment on [bug 901079|https://bugzilla.redhat.com/show_bug.cgi?id=901079] > dist-git has been reverted to 1.3.1 should it be still open ?	0	0	https://issues.jboss.org/browse/JBPAPP6-1337
10659	JBPAPP6-1336	Is this a regression from EAP 6.0.0.GA?	0	0	https://issues.jboss.org/browse/JBPAPP6-1336
10660	JBPAPP6-1333	The remoting upgrade is being tested upstream (in the 7.1 and master branches).	0	0	https://issues.jboss.org/browse/JBPAPP6-1333
10661	JBPAPP6-1333	Oh, I had set remoting to be upgraded to 3.2.9 for ER2 in http://ett.usersys.redhat.com/brew_tags/jb-eap-6-rhel-6-wolf/packages Should I cancel that?	0	0	https://issues.jboss.org/browse/JBPAPP6-1333
10662	JBPAPP6-1299	Hi Rotislav, I can't reproduce this locally (tried multiple times) with OpenJDK 1.7.0-b146, current AS7 master and current JBossWS-CXF 4.0.x EAP6 branch. Can you please check if we can have fresher results from hudson (using latest EAP6 & ws stack)?	0	0	https://issues.jboss.org/browse/JBPAPP6-1299
10663	JBPAPP6-1299	Jim Ma <ema@redhat.com> made a comment on [bug 899644|https://bugzilla.redhat.com/show_bug.cgi?id=899644] Rosta, can you please try my lastest fix ? https://source.jboss.org/changelog/JBossWS?cs=17389 We wrongly define a non-exsiting portName, it will be loaded transiently with some versionjdk6 and jdk7.	1	1	https://issues.jboss.org/browse/JBPAPP6-1299
10664	JBPAPP6-1238	Brian: Could you look at the screenshot in [this subtask|https://issues.jboss.org/browse/JBPAPP-9108]? What do you think of the design? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP6-1238
10665	JBPAPP6-1236	Rostislav Svoboda <rsvoboda@redhat.com> made a comment on [bug 901066|https://bugzilla.redhat.com/show_bug.cgi?id=901066] CXF-4561 was fixed in 2.6.4, EAP 6.1.0 contains 2.6.6 Does it mean it's resolved ? If yes, please set proper Status to ON_QA	0	0	https://issues.jboss.org/browse/JBPAPP6-1236
10666	JBPAPP6-1195	Madhumita Sadhukhan <msadhukh@redhat.com> made a comment on [bug 900649|https://bugzilla.redhat.com/show_bug.cgi?id=900649] I tried to enter remote location for existing file http://maven.repository.redhat.com/techpreview/all/example-settings.xml However it always gives me error message that settings.xml does not exist. However specifying a maven repository and settings.xml available on local works. Is this fix in 6.1.0.ER1?	0	0	https://issues.jboss.org/browse/JBPAPP6-1195
10667	JBPAPP6-1195	Thomas Hauser <thauser@redhat.com> made a comment on [bug 900649|https://bugzilla.redhat.com/show_bug.cgi?id=900649] Madhumita, What this part of the installer does is add an active profile to the given (local) settings.xml with the location of the maven repository configured (the locaiton of the maven repo can be remote). This is to simplify the building of the quickstarts. One cannot specify a remote location for the settings.xml (the installer is modifying the file). Scenarios: a) Local repo location, local settings.xml - Useful if the user has downloaded the JBoss maven repo, can specify the local location b) Remote repo location, local settings.xml - Useful if the user hosts their own maven repository, or for using the public JBoss repos c) Local repo location, no settings.xml exists - Basic settings.xml will be created at the given local path d) Remote repo location, no settings.xml exists - Same as above So it seems like the issue you've described, we should prevent by making sure the user is not trying to specify a remote settings.xml location. Does that sound about right? Thanks, Tom	0	0	https://issues.jboss.org/browse/JBPAPP6-1195
10668	JBPAPP6-1101	Hi Alessio, when will the tag of CXF&jbossws-cxf be ready?	1	1	https://issues.jboss.org/browse/JBPAPP6-1101
10669	JBPAPP6-1078	Hi Rio, Can you please change the EAP branch for JBoss Native so it has the correct (for EAP 6) version and groupId for netty? We are using Netty 3.2.6.Final and the groupId is still org.jboss.netty Thanks P.S.: I worked arounf and I am building native 4.0.6.GA so that request is for a future 4.0.7.GA release.	1	1	https://issues.jboss.org/browse/JBPAPP6-1078
10670	JBPAPP6-1064	What I think Misty wanted was something more in the lines of: If the version of Remoting in the client is not the same as the one the Server was just updated to you may encounter this problem. Right Misty?	1	1	https://issues.jboss.org/browse/JBPAPP6-1064
10671	JBPAPP6-1044	Can this be altered to adhere to the ISO/IEC 19770-2 standard?	0	0	https://issues.jboss.org/browse/JBPAPP6-1044
10672	JBPAPP6-1044	Paul, as you create the zips for EAP 6, can you please handle this? Thanks!	0	0	https://issues.jboss.org/browse/JBPAPP6-1044
10673	JBPAPP6-1043	Jason, can you make this change? At least incorporate Paul's patch?	0	0	https://issues.jboss.org/browse/JBPAPP6-1043
10674	JBPAPP6-1041	Verified during 6.0.1-ER2 test cycle. Javadoc is now shipped in ZIP and RPM distribution (the same content), remains issues are still tracked on ZIP side, hence marking this issue as resolved and closing...	1	1	https://issues.jboss.org/browse/JBPAPP6-1041
10675	JBPAPP6-1040	Alex, which one is that? Agent at i-e9138494 can't be launched, and I don't see i-83900dfe in the list of nodes. Is that the one?	0	0	https://issues.jboss.org/browse/JBPAPP6-1040
10676	JBPAPP6-1040	Ondra, I think empty junit reports happen on timeouts and other bad failures where a report file coulnd't be created. Are you sure these reported failures are really harmless?	0	0	https://issues.jboss.org/browse/JBPAPP6-1040
10677	JBPAPP6-1040	Alex, what do you mean? Are you referring to https://jenkins-jbossqa.rhcloud.com/job/EAP6-EC2/106/#showFailuresLink ?	0	0	https://issues.jboss.org/browse/JBPAPP6-1040
10678	JBPAPP6-1040	ah, I misread the error as an empty junit report. Anyways, why can't these tests find that path? {code}Path not found: /mnt/workspace/EAP6-EC2/jboss-as/testsuite/integration/manualmode/../../../build/target{code}	0	0	https://issues.jboss.org/browse/JBPAPP6-1040

#	Requirement	Questions	Answered	Successfully Answered	URL
10679	JBPAPP6-987	{quote}Note that these do need adding, we should not replace the existing BOMs.{quote} Do we need to include both versions of the poms in the EAP 6.0.1 repo? There won't be anything in the 6.0.1 repo that depends on the older versions of the boms, but maybe some of the 6.0.1 quickstarts or other stuff would need the older versions.	1	1	https://issues.jboss.org/browse/JBPAPP6-987
10680	JBPAPP6-987	Karel, ok, so if I understand correctly, removing the arquillian-drone-bom dependency from the -with-tools bom was the correct solution, right?	1	0	https://issues.jboss.org/browse/JBPAPP6-987
10681	JBPAPP6-985	I can't find Release Notes for this issue in http://documentation-devel.engineering.redhat.com/docs/en-US/JBoss_Enterprise_Application_Platform/6/html-single/Release_Notes_6.0.1/ could you please look at it? Thanks	1	1	https://issues.jboss.org/browse/JBPAPP6-985
10682	JBPAPP6-982	The real question is the probability of the customer hitting this issue. Mirek confirmed that the server doesn't have to be killed multiple times to hit this issue. I don't understand the correlation to small message sizes here, when id-cache-size deals with message count. Is it likely to send more than 2000 messages during a single crash/restart when the message size is small? <mnovak> but when messages are small (means less then 500 bytes) then confirmation-window-size can contain more then 2000 messages (2000 is default value) and when duplicated message comes again after kill and restart case then it's not discarded because it's not in the cache <mnovak> so oldest message is always rewritten be the newest one <rajesh> Does the use case have to be killing/restarting multiple times to reproduce this? <mnovak> no <mnovak> it can be hit with just one kill	1	1	https://issues.jboss.org/browse/JBPAPP6-982
10683	JBPAPP6-982	What will be a better value? we could increase it to 20K, that will add performance impact to people who don't need that ammount of duplicate detection. You will have 20K strings in memory on a hashSet per address.	1	1	https://issues.jboss.org/browse/JBPAPP6-982
10684	JBPAPP6-982	I can increase the size... do you want to hold a release because of that?	0	0	https://issues.jboss.org/browse/JBPAPP6-982
10685	JBPAPP6-981	Unfortunately this seems to be a limitation of the jar. Only icons that go under /usr/share/icons seem to properly show up at proper resolution. The ones inside the jar (regardless of resolution) show up stretched/pixelated. I tried using one from the izpack itself : http://s01.yyz.redhat.com/dmondega/pastebin/files/gnomeicon.png As you can see its perhaps a bit better. In the menu bar (top-left) reason for cut-off is because we provided one resolution image, in order to accomodate for that we'd have to provide different resolution ones. Do you think its worth it to pursue this goal with the end result looking somewhat like what i posted ? Let me know, if so i can file in a request to get proper graphics done and use those instead.	1	0	https://issues.jboss.org/browse/JBPAPP6-981
10686	JBPAPP6-981	Do we really need a release notes text for this jira? Thomas, IMO its best to mark this as Release Notes text not required as officially we do not support Fedora. This is not included in release notes in docs devel yet.	0	0	https://issues.jboss.org/browse/JBPAPP6-981
10687	JBPAPP6-976	Hi Rotislav, I can't reproduce this locally (tried multiple times) with OpenJDK 1.7.0-b146, current AS7 master and current JBossWS-CXF 4.0.x EAP6 branch. Can you please check if we have fresher results from hudson?	0	0	https://issues.jboss.org/browse/JBPAPP6-976
10688	JBPAPP6-852	Does this mean GA will go out with a custom mvn build for customers to use !? How is the mvn that is being installed supposed to be used by users ? or is it to be ignored?	1	1	https://issues.jboss.org/browse/JBPAPP6-852
10689	JBPAPP6-852	@Max, F, J, and D have been addressed for CR1. We provided you with a CR1 Installer to verify. Can you?	0	0	https://issues.jboss.org/browse/JBPAPP6-852
10690	JBPAPP6-852	Can that JIRA be marked as resolved?	0	0	https://issues.jboss.org/browse/JBPAPP6-852
10691	JBPAPP6-852	Does this need a release note for EAP 6.0.1? If so, the issue needs to be re-opened and the release notes flags set (Affects Release Notes, Not Yet Documented). If not, the issue needs to be re-opened and the release notes flag set to "Release notes not required".	0	0	https://issues.jboss.org/browse/JBPAPP6-852
10692	JBPAPP6-850	I'm not sure if xsom is a common library that changes every release. Can we do a rebuild in mead for EAP 6.0.1 even if there is no code change to resolve these differences?	0	0	https://issues.jboss.org/browse/JBPAPP6-850
10693	JBPAPP6-849	David Walluck <dwalluck@redhat.com> made a comment on [bug 900762|https://bugzilla.redhat.com/show_bug.cgi?id=900762] How is this different from bug #957679?	0	0	https://issues.jboss.org/browse/JBPAPP6-849
10694	JBPAPP6-848	513 Utf8 Starting Apache Velocity v1.6.3 (compiled: 2012-04-20 14:27:08) Utf8 Starting Apache Velocity v1.6.3 (compiled: 2012-06-14 00:07:08) Is the only difference found.	0	0	https://issues.jboss.org/browse/JBPAPP6-848
10695	JBPAPP6-844	Hi String, you've rebuilt the hibernate4 RPMs after the ZIP was created for these fixes: * Tue Jun 26 2012 Strong Liu <stliu@redhat.com> 4.1.3-2.9.Final - Change to R antlr-eap6, javassist-eap6 - Remove hibernate-infinispan from zip This should not cause any class differences but QE is seeing differences in the org.hibernate.internal* packages between the ZIP (based on your previous build) and this latest RPM you've built. Of course, they are comparing the RHEL-6 previous build (the ZIP is created from the RHEL-6 build) and the RHEL-5 build of this last one. Are the differences somehow related to the different OSs? Or maybe you have a different patch in the two branches? Can you investigate please?	1	1	https://issues.jboss.org/browse/JBPAPP6-844
10696	JBPAPP6-844	that change was requested by EWS team (Permaine and Weinan). there is not other change except the two dependencies are changed to the -eap6 ones. I think the differences are caused by these dependencies change. but not sure. normally, it wouldn't. have they compared the jar from zip with the RHEL-5 build before this change applied? do they see any differences? and how the difference are compared?	1	0	https://issues.jboss.org/browse/JBPAPP6-844
10697	JBPAPP6-844	Hi Strong, you are abolutely right. The ZIP was created from the RHEL-6 build: And the RHEL-5 RPM is hibernate4-4.1.3-2.8.Final.ep6.el5 so it is from before the changes I mentioned above. hibernate4-4.1.3-2.8.Final.ep6.el6 So it should be a plain rebuild on RHEL-5 of your RHEL-6 one. But still, QE is seeing some differences in the classes above. They just want to make sure the code is equivalent. If you look at the classes in both 2.8 builds for RHEL-6 and RHEL-5 and on what they differ can you confirm they are just compiler (or some other tool) differences due to the different OS and JDK release? I don't know if you can access this: http://s01.yyz.redhat.com/jtripath/report/modules/org/hibernate/main/hibernate-core.jar/JarReport.html but even if you can't I can tell you the differences are all in classes related to JAXB. The RHEL-6 build was built with: glassfish-jaxb-2.2.5-4_redhat_1.ep6.el6 And, because of the mass rebuild build order and regen delay, the RHEL-5 one was built with: glassfish-jaxb-2.1.12-10_patch_02.ep5.el5 instead of glassfish-jaxb-2.2.5-4_redhat_1.ep6.el5 Maybe this explains the difference. Now the $1,000,000 question: At run time what will be used is jaxb 2.2.5 regardless of what we compiled hibernate-core with. Is there any problem using the code compiled with jaxb 2.1.12? BTW, your build for EWS (the 2.9 one) for RHEL-5 was also built with jaxb 2.1.12. If you think there is a problem, can you please rebuild the hibernate4 package on RHEL-5? I believe both EAP 6 6.0.0.GA and EWS can use the same builds as you say the changes do not affect code. Thanks for looking into this.	1	1	https://issues.jboss.org/browse/JBPAPP6-844

#	Requirement	Questions	Answered	Successfully Answered	URL
10698	JBPAPP6-844	the difference of jaxb generated source with different glassfish-jaxb do not matter, since they are only used in the new metamodel. this new metamodel is not finished/supported in EAP 6 (which we're currently working on and it is targeted to Hibernate 5), the _Configuration_ API is the only supported way in EAP 6, as same as before, no change. do you think this is still a problem that requires a rebuild?	1	0	https://issues.jboss.org/browse/JBPAPP6-844
10699	JBPAPP6-760	ER7 release is May 8. Is the resolution of this JIRA on track for that? Not sure if there is a QE element to this but we only have one build scheduled after if any adjustments in the build environment need to be made.	0	0	https://issues.jboss.org/browse/JBPAPP6-760
10700	JBPAPP6-760	There are projects in the src.zip that won't build. So far, what I'm finding is arquillian builds, which I'm sure are only used for running tests. I'll log the specific missing deps here, but it brings up a few questions: 1. Should we include sources for things we know won't build with our repository, in our src.zip file? 2. Should we be including test-time dependencies in the src.zip? 3. When we say "allow them to rebuild" are we talking about running the unit tests? If not, shouldn't we be looking at disabling tests? (we could possibly do this via an embedded settings.xml that specified maven.test.skip == true as a property in a profile or something.) I'm not clear on the rules of the road here. I'm supposed to verify that we can rebuild EAP6 from source, and document how to do it...but what does "rebuild EAP6 from source" mean, exactly?? Obviously, there are nuances to the phrase.	0	0	https://issues.jboss.org/browse/JBPAPP6-760
10701	JBPAPP6-760	Regarding building and running tests - I don't really have an opinion on this, our process now is mainly designed for convenience. For the short term, I think we can just focus on being able to reproduce the binaries without running tests as a requirement. Regarding jboss-osgi-logging - I see version 1.0.0-redhat-1 in the repository, doesn't this fall into the range [1.0.0.Final,1.0.0.Final-redhat-1]? If not, I'm not sure how the original build worked. Regarding openjpa - We don't ship this one, so we don't build it as part of EAP. Regarding BOMs - Users should not directly import the EAP BOMs into their builds, for the reasons you mentioned. For this use case, we have several user boms documented in the quickstarts and elsewhere. These main EAP BOMs purposely capture a superset of what is shipped in EAP because we use them to manage all the builds that go into EAP and many of these build dependencies are not shipped. Just as an example, the hibernate 3 stuff is use to verify the hibernate 3 integration, but this is not shipped because a user would be required to add the necessary components to their EAP instance to run a hibernate 3 based project.	0	0	https://issues.jboss.org/browse/JBPAPP6-760
10702	JBPAPP6-760	RE: jboss-osgi-logging, If I read the maven source code for version parsing, it *seems* as though 1.0.0 == 1.0.0.Final, which means it would fall in the stated range. However, something must have changed, or else that equivalence isn't tested. The bottom line is that I was using Maven 3.0.4 to run these builds, and it failed to find a valid version. I've verified that the metadata containing this version (1.0.0-redhat-1) is available to the build, yet it doesn't seem to satisfy the range. Is it possible that the BOM got changed after the build/pom.xml or dist/pom.xml ran? I thought we were using a different MEAD build to create the EAP distribution. If that's not the case, I'm at a loss to explain how it worked, but I can show the environment and logs that led to my result. RE: what to build (not building tests), should we exclude things from the src.zip if we're not intending them to be built? Isn't inclusion in the .zip implying rebuild-ability?	0	0	https://issues.jboss.org/browse/JBPAPP6-760
10703	JBPAPP6-757	Vladimir, can you create a list of the poms currently in the Maven repository which are still referring to jboss parent version 6? If possible put the list into a format that can be read by a script, maybe one line for each GAV that contains a reference to the older parent pom.	0	0	https://issues.jboss.org/browse/JBPAPP6-757
10704	JBPAPP6-692	The installer looks great! Here are a few things that came to mind as I was trying it out: Minor typo on the quickstart installation screen: - The title currently is "QuickStarts Installation" but it should just be "Quickstarts Installation" (since the docs don't seem to capitalize the "s"). - Similarly, "Do you want to install QuickStarts?" should be "Do you want to install quickstarts?" Minor typo on the maven repository screen: - "...containing Redhat artifacts,..." should be "...containing Red Hat artifacts,..." "Socket Binding Setup" screen: - It might be helpful to have a summary at the top of the screen that explains the options on the screen, as most of the other screens have something like this. "Socket Binding (Domain)" screen: - If a user chooses to only configure ports for domain.xml, when the "Socket Binding (Domain)" screen comes up, all fields in the first column are blank and so it might not be obvious to the user what kinds of values are allowed in that first column. Thus, it might be helpful to add names for the two columns.	1	1	https://issues.jboss.org/browse/JBPAPP6-692
10705	JBPAPP6-691	Is this the same as JBPAPP-7931 ?	0	0	https://issues.jboss.org/browse/JBPAPP6-691
10706	JBPAPP6-687	Rich, can you comment if this is a requirement for EAP6?	0	0	https://issues.jboss.org/browse/JBPAPP6-687
10707	JBPAPP6-686	@Tom, this is a roll over from EAP4/5. Has this been adderssed in 6 or planning to be?	0	0	https://issues.jboss.org/browse/JBPAPP6-686
10708	JBPAPP6-686	Hi, I read through the linked gss discussion. Here is what currently happens on timeout: com.arjuna.ats.arjuna.coordinator.TransactionReaper::doCancellations:349 tsLogger.logger.trace("Reaper Worker " + Thread.currentThread() + " attempting to cancel " + e._control.get_uid()); Are you saying that we need to bump this up to info from trace? Here are the normal transaction state change messages: com.arjuna.ats.arjuna.coordinator.BasicAction::Begin:1342 tsLogger.logger.trace("BasicAction::Begin() for action-id " + get_uid()); com.arjuna.ats.arjuna.coordinator.BasicAction::End:1342 tsLogger.logger.trace("BasicAction::End() for action-id " + get_uid()); com.arjuna.ats.arjuna.coordinator.BasicAction::Abort:1575 tsLogger.logger.trace("BasicAction::Abort() for action-id " + get_uid()); Are you saying we should bump these from trace to debug? To clarify, are we saying that the messages themselves are OK but you want me to bump their log levels? Tom	1	1	https://issues.jboss.org/browse/JBPAPP6-686
10709	JBPAPP6-686	Hi Taro, Thanks for the feedback. Assuming we have discrete log messages for these actions, is it not fair to say that the specific wording of the message (assuming it is on topic) is a matter of education for the user? Those messages are well established and to alter them for specific versions could actually introduce confusion for earlier users. I also can confirm that during regular transaction completion you will also get the following messages preceding those messages from BasicAction: com.arjuna.ats.internal.jta.transaction.arjunacore.TransactionImple::commit:172 if (jtaLogger.logger.isTraceEnabled()) { jtaLogger.logger.trace("TransactionImple.commit"); } com.arjuna.ats.internal.jta.transaction.arjunacore.TransactionImple::rollback:231 if (jtaLogger.logger.isTraceEnabled()) { jtaLogger.logger.trace("TransactionImple.rollback"); } Due to the way that the system is architected, those BasicAction messages I provided earlier are descriptive and apply to the transactions underlying "BasicAction". Tom	1	1	https://issues.jboss.org/browse/JBPAPP6-686

#	Requirement	Questions	Answered	Successfully Answered	URL
10710	JBPAPP6-686	Thanks Tom. As you read the original discussion, this request was originally one of change requests for new major version of EAP, 6.0, than for a particular version of JBossTS and from EAP perspective I am assuming incompatibility of trace messages between EAP 5 and 6 would be acceptable for customers. Also in general educating the user is important for gss to reduce support case load to explain what those messages mean and how to track those. The four messages in the description should be the typical ones the customers asking to gss and making those messages user-friendly from EAP6 release would be great for gss and customers. I understand that current messages are well established and BasicAction messages are descriptive from code perspective as you mentioned, however, it would be great that this enhancement is acceptable for EAP 6.x release than for a particular version of JBossTS. If it's not acceptable, can you please consider to add more details around those messages in JBossTS(and EAP) documentations? Taro	1	0	https://issues.jboss.org/browse/JBPAPP6-686
10711	JBPAPP6-682	Isn't this the same as JBPAPP-8931?	0	0	https://issues.jboss.org/browse/JBPAPP6-682
10712	JBPAPP6-677	Kyle, AS 7.1.2.Final already includes JAXB 2.2.5. Not sure which version of EAP 6 is picking that up.	0	0	https://issues.jboss.org/browse/JBPAPP6-677
10713	JBPAPP6-589	Karel, can you create a Mead build which contains only an Arquillian BOM which contains all the versions of arquillian components which should be used with EAP 6? Then we can ship this with the Maven repo, and use this for the build of the jboss-bom project.	1	1	https://issues.jboss.org/browse/JBPAPP6-589
10714	JBPAPP6-589	Karel, what about the javaee-with-tools bom? I can't build that one because it still points to BOMs in central, can this be changed to use the wfk-arquillian bom instead?	1	1	https://issues.jboss.org/browse/JBPAPP6-589
10715	JBPAPP6-503	Gao, could you please update the jbossws common and native ones?	0	0	https://issues.jboss.org/browse/JBPAPP6-503
10716	JBPAPP6-428	Can't this be pushed to AS 8 then..or into EAP 6.....it should be worked on at some point...	0	0	https://issues.jboss.org/browse/JBPAPP6-428
10717	JBPAPP6-380	Shelly, AFAIR in Beta1 logging was restored also for production profiles, does this address this issue as well?	0	0	https://issues.jboss.org/browse/JBPAPP6-380
10718	JBPAPP6-380	Shelly, any reason why don't we ship/enable async logging for production env?	1	1	https://issues.jboss.org/browse/JBPAPP6-380
10719	JBPAPP6-380	I know the heap sizes (just for standalone) were put back into the ER4 build. Did every other item targeted in JBPAPP-6738 make it into the ER4? If I am not wrong the commit made for JBPAPP-6738 was reverted completely? Is there another JIRA tracking any remaining work here?	0	0	https://issues.jboss.org/browse/JBPAPP6-380
10720	JBPAPP6-379	Fernando, would you mind I take this jira? Rajesh, Fernando, do you think this build has to come from brew or a local build is fine?	1	1	https://issues.jboss.org/browse/JBPAPP6-379
10721	JBPAPP6-379	Strong, I don't mind if the build is done locally. The official build is part of EAP 5.1.3 or an EWS deliverable. A local build is fine for EAP6 testing. Scheduling this for ER4, can you create a build and coordinate handoff with Shelly?	0	0	https://issues.jboss.org/browse/JBPAPP6-379
10722	JBPAPP6-337	What is the connection of this jira to DESIGN-209 ?	0	0	https://issues.jboss.org/browse/JBPAPP6-337
10723	JBPAPP6-304	thomas, can you elaborate on this?	0	0	https://issues.jboss.org/browse/JBPAPP6-304
10724	JBPAPP6-304	Ok - here's a start. Test OSGi Bundle deployment and runtime bundle management: On the 'Runtime' page: * Go to 'Runtime Operations->OSGi' * The console should ask whether the OSGi subsystem needs to be activated, select yes. Then, after a few seconds, the list of bundles will appear. * Deploy an OSGi Bundle (e.g. test TestOSGiBundle attached to this bug) via 'Deployments->Manage Deployments' * Enable the deployment ** You should see 'Hello from an OSGi Bundle!' appear on the AS7 console. * Go to 'Runtime Operations->OSGi' Initially the new deployment is not visible as the table only has 8 rows and our testbundle is the 9th. * Reverse the sorting on Bundle ID by clicking on the header. You'll see bundle with the highest ID is our TestOSGiBundle. * Click on the TestOSGiBundle row in the list and make sure that the Details pane is correctly filled with information. It should show: {noformat}Version: 1.0.0 Start Level: 1 SymbolicName: TestOSGiBundle Type: bundle State: ACTIVE{noformat} * Click the 'Stop' button on the bundle row. ** You should see 'Goodbye from an OSGi Bundle!' in the console ** The State column should be empty now and the state field in the details pane should say 'RESOLVED'. * Go back to 'Deployments->Manage Deployments' and remove the TestOSGiBundle deployment. Test capabilities configuration On the 'Profile' page: * Select 'OSGi->Framework' and go to the 'Capabilities' tab. * Click 'Add' and enter org.projectodd.stilts with start level 2. * Should appear now in the list. Go back to the Runtime page: * The stilts-stomplet-server-bundle should now be visible. Clicking on it should show Start Level 2. State should be INSTALLED. * Now move the framework to start level 2 by clicking the Framework tab. This causes the stilts bundle to be started, but unfortunately there is a resolution error. Thomas I attempted to add the org.jboss.netty module via capabilities, but that caused another error. Do you know what the issue is with netty? * Open the standalone.xml file and see that the following line was added to the configuration: {noformat}<capability name="org.projectodd.stilts" startlevel="2"/>{noformat} * Remove the stilts bundle again by going back to the Capabilities tab and clicking remove. ** Console should say that it needs to be restarted, currently it doesn't (JBPAPP-8358). Test framework properties On the 'Profile' page: * Select OSGi -> Framework and go to the 'Properties' tab. * Remove the 'org.osgi.framework.startlevel.beginning' and add it again now with value 3. (BTW we used to have an edit button here, it would be more user-friendly to allow editing instead of delete-add) ** Console should say that it needs to be restarted, currently it doesn't (JBPAPP-8358). * Check in the standalone.xml that it has the following configuration: {noformat} <properties> <property name="org.osgi.framework.startlevel.beginning"> 3 </property> </properties>{noformat} * Remove the check from the Lazy Activation configuration checkbox. * Check in the standalone.xml that it has the following configuration: {noformat}<subsystem xmlns="urn:jboss:domain:osgi:1.2" activation="eager">{noformat} * Restart the server and reopen the console. On the 'Runtime' page * Go to the 'Runtime Operations->OSGi' tab. Since the OSGi subsystem is eagerly started it should not ask to activate it. The list of bundles should appear instantly. * Select the Framework tab. It should show start level 3 as the active start level. Config Admin Service On the 'Profile' page select 'Core->Config Admin Service'. * Click 'Add' on top of the screen. Add a PID of org.foo.bar with mykey and myvalue as configuration values. Add another config value: xkey, yval * Check in the standalone.xml that it has the following configuration: {noformat} <subsystem xmlns="urn:jboss:domain:configadmin:1.0"> <configuration pid="org.foo.bar"> <property name="mykey" value="myvalue"/> <property name="xkey" value="yval"/> </configuration> </subsystem>{noformat}	0	0	https://issues.jboss.org/browse/JBPAPP6-304
10725	JBPAPP6-303	Why is this being rejected? If there is no documentation (yet) this issue should have a dependency on an issue that creates such docs from the jbosgi project docs. The Fix Version should reflect the EAP version that this is targeted for.	1	0	https://issues.jboss.org/browse/JBPAPP6-303
10726	JBPAPP6-302	Why would there be any tests in EAP that are not also available upstream?	0	0	https://issues.jboss.org/browse/JBPAPP6-302
10727	JBPAPP6-302	What does this mean? Is there anything I need to take care of?	1	0	https://issues.jboss.org/browse/JBPAPP6-302
10728	JBPAPP6-301	Can you please take a look at [4], the set of supported functionality? I believe we need to remove all the Jetty stuff from AS7 anyway: Webconsole, provided by Felix (requires HttpService) HttpService, provided by PaxWeb (runs on Jetty) Webapp support, provided by PaxWeb (runs on Jetty) or we should just get rid of all non-core OSGi bundles?	0	0	https://issues.jboss.org/browse/JBPAPP6-301
10729	JBPAPP6-301	The management console exposes OSGi as a runtime operation. If OSGi is removed, will there be an impact to the console functionality or is the availability of this subsystem dynamically detected?	0	0	https://issues.jboss.org/browse/JBPAPP6-301

#	Requirement	Questions	Answered	Successfully Answered	URL
10730	JBPAPP6-299	Isn't that the maven repo zip that contains the AS artifacts? Why do we need another zip, what's the main problem with rebuilding from source? In fact we should use the sources from the source distribution for running the AS testsuite. That validates the source distribution we ship as well.	1	1	https://issues.jboss.org/browse/JBPAPP6-299
10731	JBPAPP6-294	Mike, can you just write down in a wiki somewhere the steps you follow when creating the Mead repo snapshots? Then you can close this issue.	1	1	https://issues.jboss.org/browse/JBPAPP6-294
10732	JBPAPP6-274	This is the wrong package to use for xml-resolver. We are using some internal package-renamed version of xml-resolver that was being bundled with glassfish (which is certainly bad practice on both fronts). Can we use the correct GAV: xml-resolver:xml-resolver:1.2 instead?	0	0	https://issues.jboss.org/browse/JBPAPP6-274
10733	JBPAPP6-267	Actually is that command supposed to be in the DRx version? Maybe the devs forgot to include that command for the enterprise edition?	0	0	https://issues.jboss.org/browse/JBPAPP6-267
10734	JBPAPP6-264	Copied comment from AS-1338: "Hi all, I'm just adding some brief info from John about newly implemented remote jndi support and how to use it. As I understood everything what should be reachable in jndi from outside the server must be registered under "java:jboss/exported/" context. As an example I used it in standalone jms client to get connection factory and queue from jndi: // jms client code // initialize context final Properties env = new Properties(); env.put(Context.INITIAL_CONTEXT_FACTORY, "org.jboss.naming.remote.client.InitialContextFactory"); env.put(Context.PROVIDER_URL, "remote://localhost:4447"); Context context = new InitialContext(env); // do lookups ConnectionFactory cf = (ConnectionFactory) context.lookup("jms/RemoteConnectionFactory"); Queue queue = (Queue) context.lookup("jms/queue/test"); // do whatever you want How are jms objects registered to jndi in messaging subsystem in standalone-ha.xml: <jms-connection-factories> ... <!-- configuration for connection factory --> <connection-factory name="RemoteConnectionFactory"> <connectors> <connector-ref connector-name="netty"/> </connectors> <entries> <entry name="java:jboss/exported/jms/RemoteConnectionFactory"/> <!- this is the important part --> </entries> </connection-factory> ... </jms-connection-factories> <jms-destinations> ... <jms-queue name="testQueue"> <entry name="java:jboss/exported/jms/queue/test"/> <!- this is the important part --> </jms-queue> ... </jms-destinations> Used libraries (in my simple client): hornetq-core-2.2.11.Final.jar hornetq-jms-2.2.11.Final.jar jboss-as-naming-7.1.0.Final-SNAPSHOT.jar jboss-jms-api_1.1_spec-1.0.0.Final.jar jboss-logging-3.1.0.GA.jar jboss-marshalling-1.3.6.GA.jar jboss-marshalling-river-1.3.6.GA.jar jboss-modules.jar jboss-msc-1.0.2.GA.jar jboss-remote-naming-1.0.0.Beta4.jar jboss-remoting-3.2.0.CR9.jar jboss-sasl-1.0.0.Beta9.jar netty-3.2.6.Final.jar xnio-api-3.0.0.GA.jar xnio-nio-3.0.0.GA.jar Note: I had to add static block to my class so I could connect to server through remoting: static { Security.addProvider(new JBossSaslProvider()); } Hopefully I didn't forget anything and this info will help. Thanks for this cool feature, Mirek"	1	1	https://issues.jboss.org/browse/JBPAPP6-264
10735	JBPAPP6-263	Peter, can you link any upstream issue that details the issues?	0	0	https://issues.jboss.org/browse/JBPAPP6-263
10736	JBPAPP6-262	Hi Jonathan, I assume your comment on : https://jira.jboss.org/browse/PRODMGT-30 "JBossTS has no plans to support tx inflow with lazy interposition or non-interposed configurations. Changing the ejb interceptor chain to optimize out the tx import may be an alternative." applies here too. Can you confirm. Thank you. Bruno	1	1	https://issues.jboss.org/browse/JBPAPP6-262
10737	JBPAPP6-262	Does this need to be documented as a Known Issue for EAP 5.1.1? Thanks for your help!	0	0	https://issues.jboss.org/browse/JBPAPP6-262
10738	JBPAPP6-261	Anil, can you mark these dependencies as "optional"? Then you can close this issue.	0	0	https://issues.jboss.org/browse/JBPAPP6-261
10739	JBPAPP6-257	Why did you upgrade it only to 2.1.3.1? AS 7.1.3.Beta1 which is the base of EAP 6.0.1 ER1 has picketlink 2.1.4 in the dependency tree.	0	0	https://issues.jboss.org/browse/JBPAPP6-257
10740	JBPAPP6-208	It's both there: the landing pages and the console updates. Are you sure you are looking at the right instance? Maybe you need to clear your browser cache?	0	0	https://issues.jboss.org/browse/JBPAPP6-208
10741	JBPAPP6-204	It looks like there is a problem with the new maven bundle plugin 2.5 monolith. Can you take a look at this John? {noformat} [ERROR] Failed to execute goal org.apache.felix:maven-bundle-plugin:2.3.5-redhat-1:bundle (default-bundle) on project jboss-as-osgi-configadmin: Execution default-bundle of goal org.apache.felix:maven-bundle-plugin:2.3.5-redhat-1:bundle failed: Unable to load the mojo 'bundle' in the plugin 'org.apache.felix:maven-bundle-plugin:2.3.5-redhat-1'. A required class is missing: org/codehaus/plexus/archiver/UnArchiver [ERROR] ---------- [ERROR] realm = plugin>org.apache.felix:maven-bundle-plugin:2.3.5-redhat-1 [ERROR] strategy = org.codehaus.plexus.classworlds.strategy.SelfFirstStrategy [ERROR] urls[0] = file:/home/pgier/projects/jboss-eap/local-repo/org/apache/felix/maven-bundle-plugin/2.3.5-redhat-1/maven-bundle-plugin-2.3.5-redhat-1.jar [ERROR] urls[1] = file:/home/pgier/projects/jboss-eap/local-repo/org/codehaus/plexus/plexus-utils/2.1-redhat-1/plexus-utils-2.1-redhat-1.jar [ERROR] Number of foreign imports: 1 [ERROR] import: Entry[import from realm ClassRealm[project>org.jboss.as:jboss-as-osgi-configadmin:7.1.0.Beta1-redhat-1, parent: ClassRealm[maven.api, parent: null]]] [ERROR] [ERROR] ---------- ----------: org.codehaus.plexus.archiver.UnArchiver [ERROR] -> [Help 1] {noformat}	1	1	https://issues.jboss.org/browse/JBPAPP6-204
10742	JBPAPP6-195	Paul, will the JBossEAP-Release-Version be set only for artifacts that come from the AS source tree? External components for instance Hornetq, which is also built in mead, only has the community release version (2.2.6) in its manifest, does not have a corresponding EAP version. Typically components like this will have their own product branch and create releases from them, for instance say, 2.2.6.CP01. Should their manifests still have the corresponding EAP version (6.0.x).	0	0	https://issues.jboss.org/browse/JBPAPP6-195
10743	JBPAPP6-194	Vladmir, can you run the script against the -latest tag repository, and then upload the script?	1	1	https://issues.jboss.org/browse/JBPAPP6-194
10744	JBPAPP6-191	is this a permanent change or something that is done manually for every build?	0	0	https://issues.jboss.org/browse/JBPAPP6-191
10745	JBPAPP6-191	I mean, is there a manual process to remove openjpa for each build/release or is the removal automated?	0	0	https://issues.jboss.org/browse/JBPAPP6-191
10746	JBPAPP6-186	Could you please provide details how to import JBossPublicKey.RSA into keystore.	0	0	https://issues.jboss.org/browse/JBPAPP6-186
10747	JBPAPP6-179	What is the difference between -sources.zip and -sources.jar ? asking since maybe those sources.zip are needed but shuold be called sources.jar ? find on .tar shows the maven repo contains jboss-as-dist-7.1.0* zip - this is in .org mainly used for downloading/installing AS via arquillian afaik ? Is this a usecase we need to cover for EAP too or just remove this?	0	0	https://issues.jboss.org/browse/JBPAPP6-179
10748	JBPAPP6-169	Heiko, here is a patch with some minor updates to the content based on what we previously had in EAP 6, can you apply this to the ballroom git repo?	0	0	https://issues.jboss.org/browse/JBPAPP6-169

#	Requirement	Questions	Answered	Successfully Answered	URL
10749	JBPAPP6-165	If -redhat identifier represents EAP, should the version be 6.0.0-redhat-1 or 7.0.1-redhat-1/7.1.0-redhat-1? It relates to my comment in JBPAPP-6831 . When EAP 6.0.1 comes out, what will this version be? 7.1.1.-redhat-1? Actually incrementing the AS version like this in EAP releases was Andy's original recommendation, which you were opposed to. BTW, did we come to an agreement if the AS version in server startup messages would stay the same or increment with every EAP release. Basically the same question as above.	0	0	https://issues.jboss.org/browse/JBPAPP6-165
10750	JBPAPP6-164	Here are two versions of the EAP favicon based on the look and feel I designed for the Console. I've included the .png version and the .ico	0	0	https://issues.jboss.org/browse/JBPAPP6-164
10751	JBPAPP6-164	With DR11 build, I see the favicon only on the landing page, but not for the console pages. Can somebody check to verify if it's not an issue with my browser?	0	0	https://issues.jboss.org/browse/JBPAPP6-164
10752	JBPAPP6-154	btw. is http://lists.jboss.org/pipermail/jboss-as7-dev/2012-January/005017.html the one we should use going forward ? look for bin/product.conf for the "type" (eap, soa-p, etc.) and then modules/org/jboss/as/product/<product.conf>/dir/META-INF/MANIFEST.MF/JBoss-Product-Release-Version for the version ?	1	1	https://issues.jboss.org/browse/JBPAPP6-154
10753	JBPAPP6-154	Rob, I was asked to verify this. I just did - Max' steps work. Can you confirm that this is fine with you and close the JIRA? Thanks.	0	0	https://issues.jboss.org/browse/JBPAPP6-154
10754	BRMS-599	I've compared this request to our Support-Patch guidelines* and find the request reasonable. (We don't exactly follow the Support-Patch guidelines for BRMS at this point, but it can serve as a guideline.) The Support Patch guidelines are not particularly clear, so I will take it as a take-away to make sure they are clear when we make the move to BZ (perhaps this month.) Alessandro, sorry the doc failed us this time. Would you please make allowances for Tiho's request, and try to use this case as an example until we re-publish for BZ. Thanks, Rick *https://docspace.corp.redhat.com/docs/DOC-51939	1	1	https://issues.jboss.org/browse/BRMS-599
10755	BRMS-596	kylin <ksoong@redhat.com> made a comment on [bug 724753|https://bugzilla.redhat.com/show_bug.cgi?id=724753] As [1] hints the issue is fixed in drools 5.3.0.Final, but I test via BRMS 5.3.0 the issue still existed [1] https://issues.jboss.org/browse/JBRULES-3063	0	0	https://issues.jboss.org/browse/BRMS-596
10756	BRMS-596	Edson Tirelli <etirelli@redhat.com> made a comment on [bug 724753|https://bugzilla.redhat.com/show_bug.cgi?id=724753] Can you please clarify how you tested? By default, BRMS does allow rule replacement... it will only raise errors when the engine is configured to raise errors for rule duplication.	0	0	https://issues.jboss.org/browse/BRMS-596
10757	BRMS-555	The drools parent pom has been building the test-jar for a couple of months already: https://github.com/droolsjbpm/droolsjbpm-build-bootstrap/blob/master/pom.xml#L377 <plugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-jar-plugin</artifactId> <version>2.3.1</version> <executions> <execution> <goals> <goal>test-jar</goal> </goals> </execution> </executions> The test-jars should also be deployed to nexus. More info at http://maven.apache.org/guides/mini/guide-attached-tests.html So I don't understand the problem. One possible corner case: I don't know how that configuration affects war packaging, such as guvnor-webapp (formerly known as drools-guvnor). Of what artifact are you missing the test jars?	0	0	https://issues.jboss.org/browse/BRMS-555
10758	BRMS-428	Kris Verlaenen <kverlaen@redhat.com> made a comment on [bug 958408|https://bugzilla.redhat.com/show_bug.cgi?id=958408] This related to how metadata for files is stored in Guvnor for BRMS/BPMS6. Is metadata stored in an additional file on the file system (so it can be editor from Eclipse using for example EGit)?	0	0	https://issues.jboss.org/browse/BRMS-428
10759	BRMS-427	Lucas, Trevor is right. It should not be any different. We still expect users to deploy the deployable components on either EAP, EWP, SOA-P and EWS. So, they should use the same properties. I assume Docs would address this clearly. Btw, what does EWS use for the same ? Also, note, many of our customers use BRMS on other containers like Websphere and Weblogic. We dont need to get into those details and not necessarily document those scenarios. But we need to be congnizant of that fact. I am sure there are similar login module options in those products. My question is what does BRM (Guvnor) look for in that case.	0	0	https://issues.jboss.org/browse/BRMS-427
10760	BRMS-415	What else needs to be done with this? Can it be marked as 'Resolved'?	1	1	https://issues.jboss.org/browse/BRMS-415
10761	BRMS-415	This should now be a doc issue. Can we explain when one option is preferred over another? And then hand over to the docs group, please.	0	0	https://issues.jboss.org/browse/BRMS-415
10762	BRMS-415	IIUC, users don't have to change or decide on any of these options because the ModeShape service is now being configured correctly when used within BRMS. If that is the case, does anything need to be documented?	1	1	https://issues.jboss.org/browse/BRMS-415
10763	BRMS-415	Can this be closed? What else needs to be done? Does it need to be reassigned?	0	0	https://issues.jboss.org/browse/BRMS-415
10764	BRMS-397	Randall can I think 1.0 is out, I think it makes sense to upgrade MS to that version, rather then using an release candidate?	1	1	https://issues.jboss.org/browse/BRMS-397
10765	BRMS-393	Toshiya, checked this with trunk and I don't see the problem (unless am not running the right test). Can you show again here how to reproduce this issue in Guvnor by describing your model, rule and a test scenario?	0	0	https://issues.jboss.org/browse/BRMS-393
10766	BRMS-368	Should the fix version of this not also include 5.1.0 ? Idk.	0	0	https://issues.jboss.org/browse/BRMS-368
10767	BRMS-362	questions: 1 - Are these all the supported scenarios for BRMS 5.x ? 2 - what configuration will be provided out of the box in the EWP bundle ? 3 - how will deployment & configuration differ for tomcat in EWS ? 4 - what will be involved with BRE deployment ? Also what is it? And what is it for ?	0	0	https://issues.jboss.org/browse/BRMS-362
10768	BRMS-343	Note that all references to modeshape in this doc need to be prefaced with the fact that it is only supported in 5.1 and later versions of BRMS	1	1	https://issues.jboss.org/browse/BRMS-343
10769	BRMS-297	Document is here: http://downtown.englab.bne.redhat.com/drafts/JBoss_Enterprise_BRMS_Platform/5.0.2/Release_Notes/ Darrin - is the doc ready for review?	0	0	https://issues.jboss.org/browse/BRMS-297
10770	BRMS-297	Comments: Section: 4. Supported Deployment Platforms Supported Java Virtual Machines (JVM) * Sun JDK 1.5 * Sun JDK 1.6 Comment: This list may change. Waiting on a PM decision as RHEL dropped Sun 1.5 on March 31 ---------- Supported Browsers * Firefox 3 * Internet Explorer 7 * Internet Explorer 6 * Safari 3 Comment: This list needs to be expanded to also include IE8 and Safari 4 ---------- Section: 7. Possible Migration Issue Need to change the wording: Client applications built using the JBoss Rules engine libraries from the 5.0.1 and later releases are unable to user resources being server from a JBoss Enterprise BRMS Platform 5.0.0 server. Should read "use" not "user." ---------- Section: 8. Issues resolved in the 5.0.2 release The sequence of the JIRAs listed in this section seem to be random. Can the JIRAs be listed in numerical order?	0	0	https://issues.jboss.org/browse/BRMS-297
10771	BRMS-297	be handy if I actually set this issue to resolved wouldn't it ? ;-)	1	1	https://issues.jboss.org/browse/BRMS-297

#	Requirement	Questions	Answered	Successfully Answered	URL
10772	BRMS-285	Comments: ========== Figure 4.1. Main feature areas of the BRMS Platform The above picture shows the main feature areas of the BRMS Platform. • Info: This is the initial screen, with links to resources. • Rules: This is the category and business user perspective. • Package: This is where packages are configured and managed. • Deployment: this is where deployment snapshots are managed. • Admin: Administrative functions (categories, statuses, import and export) Comment: The image in figure 4-1 is truncated and the bulleted items do not match the screen display ========== Chapter 5 - The note that refers to: JBoss only provides support for this functionality in the JBoss Developer Studio 2 product. Should refer to: JBoss provides support for this functionality in the JBoss Developer Studio 2 and 3 products. ========== The first reference to Guvnor in the book is on page 4: Note These instructions tell you to use the URL http://<server>:8080/jboss-brms/ to access the BRMS application. However, that URL is just a convenience and it actually redirects your browser to the full URL, http://<server>:8080/jboss-brms/ org.drools.guvnor.Guvnor/Guvnor.html. You can use the full URL if you wish. But - the explanation of the relationship between Guvnor and BRMS is not explained until page 53: Important Guvnor is the name of the open-source project on which the JBoss Enterprise BRMS Platform has been built. The Eclipse Guvnor Tools were developed with it in mind. Any place where EGT refers to "Guvnor" it can be considered equivalent to "BRMS Platform". To avoid confusion the remainder of this chapter will use term "Guvnor". Comment: We should add this note to page 4: Important Guvnor is the name of the open-source project on which the JBoss Enterprise BRMS Platform has been built. ========== I think that it is worth defining DRL (Drools Rule Language - the "native" Drools rule language) and DSL (Domain Specific Languages - a way of extending the rule language to your problem domain) fo rth efirst time user/reader. ========== 4.6.1. The Rule Agent The Rule Agent is a component which is embedded in the core run-time of the Rules Engine. In order to use it, you do not need any extra components. In fact, if you are using the BRMS Platform, your application should only need to include the drools-core dependencies in its classpath (drools and mveljars only). There are no other rule-specific dependencies. Comment: We should refer the user to the "README_DEPENDENCIES.txt" file that is included in jboss-brms-engine.zip Note that there is also a drools-ant ant task... Comment: drools-ant-5.0.1.BRMS.jar is also included in: jboss-brms-engine.zip ========== 4.6.1. The Rule Agent The code fragment should also include: import org.drools.agent.RuleAgent; import org.drools.RuleBase; ========== Task to think about - not in the context of a point release - re-organizing the guide. Chapter 4 is named the "Quickstart Guide," but it is over 25 pages long. I'm thinking that we should redesign the guide and convert this long chapter into a set of shorter chapters where each chapter covers one task or use case.	1	1	https://issues.jboss.org/browse/BRMS-285
10773	BRMS-285	Items 1-7: 1 - already covered by BRMS-233 2 - JBDS 3 - will update 3 - will move these comments to the introduction of the book 4 - will add definitions to first usages 5 - will investigate, needs some clarification 6 - will update 7 - that is something that needs to occur upstream. The first section of the chapter is also called Quickstart Guide and actually *is* a quickstart guide. The chapter just needs renaming to something more appropriate	0	0	https://issues.jboss.org/browse/BRMS-285
10774	BRMS-284	Found the note - it looks good - one change to the wording is needed: NOTE: Quick Installation The JBoss Enterprise BRMS Platform WAR file can be quickly deployed with no configuration. This is only suitable for development and demonstration scenarios. To do this, skip three of the installation procedure and enable one or more users in jboss-as/server/production/conf/props/jmx-console-users.properties. This should read: "skip step 3"	0	0	https://issues.jboss.org/browse/BRMS-284
10775	BRMS-281	So, a user who is package admin and can access only rules of a particular category gets a permission error when he tries to access an asset that is assigned to a category (in this case no category) that this user should not have access - sounds good to me - or am I missing something? Let me summarize the problem: I have a user called "pawc" (packageadminwithcategory) which is - package admin for package "p" - analyst for category "c" and a user called "panc" (packageadminnocategory) which is package admin for package "p" I have an asset called a1 in "p" with category "c" I have an asset called a2 in "p" without any categories assigned User "panc" can access both assets, user "pawc" can only access a1 and cannot access a2. IMHO that works as expected, right? User "pawc" is limited to the package "p" but is also limited to a particular group of rules which are assigned to the "c" category.	0	0	https://issues.jboss.org/browse/BRMS-281
10776	BRMS-281	Which package admin? "073678" ? IMHO that's fine, because this is an admin who is limited with the analyst role to particular rules - so as a package admin he still can perform all the tasks a package admin is allowed to do (build pkg, create snapshots, etc) - that's what the package admin role is for. However he can only access rules which are assigned to the UWReporting, UWReporting/UWRMapping and UWReporting/UWRProviderSection categories. Rules assigned to other categories and also not assigned to any categories are out of sight for this guy, even he has package admin rights... these right are limited because of the additional analyst role.	0	0	https://issues.jboss.org/browse/BRMS-281
10777	BRMS-262	Is this content accurate and adequate ? Any suggestions ?	0	0	https://issues.jboss.org/browse/BRMS-262
10778	BRMS-257	DRL rules already supports import *, so this is a purely GUI issue. My impression is that this may not be a good idea, as the GUI needs to be eager and loads all type up for drop downs. This may make the GUI very inefficient. This is different to the import on demand in Java. In the import on demand case, java can be lazy and only resolves and loads classes that are needed during compiling. Having said that, we may still support import * if there is a valid use case and this support improves usability. As far as I know, the "Imported types" list is automatically populated with class types when a jar is uploaded. In this case, use does not need to edit the imported types manually, thus he/she does not need to use a import * which may save his/her time. Is there any other use cases when users may need to use import *?	1	1	https://issues.jboss.org/browse/BRMS-257
10779	BRMS-254	I think agents following multiple sources is rather complicated. What does it happen if the BRMS instance that is up does not have the latest package version? Instead of it, it would be better that all BRMS instances are aware they are clusterized. When a user builds a new version of a package, the BRMS should send messages to all the other BRMS instances asking to update that package. The fail over between BRMS instances could be provided by any well known HA alternative (like mod_jk).	0	0	https://issues.jboss.org/browse/BRMS-254
10780	BRMS-227	The doc looks good - can we get the URLs added to the new section? THAT will finish the book.	1	1	https://issues.jboss.org/browse/BRMS-227
10781	BRMS-219	In section 2.2. Installation, we tell the user to install the BRMS war as: http://localhost:8080/JBoss-BRMS/ instead of http://localhost:8080/jboss-brms/ (as was the case in the 5.0 GA release) as a result many other edits to the document are needed. (See the above comment.) Is there a need to change the URL on 5.0.1? If not, then we should restore the text to: http://localhost:8080/jboss-brms/	1	1	https://issues.jboss.org/browse/BRMS-219
10782	BRMS-219	Len, I removed the error message. I can't find any part of Section 2.2 or anywhere else taht referes to http://localhost:8080/JBoss-BRMS/ . They all seem to be lower case now. Can you please clarify where you saw this? Many thanks.	0	0	https://issues.jboss.org/browse/BRMS-219

#	Requirement	Questions	Answered	Successfully Answered	URL
10783	BRMS-219	Some more comments Why is JBoss *everywhere* in bold? page 3: Microsoft in bold https://support.redhat.com/JBossnetwork/ lowercase JB in JBoss page 4: If a new registry was created you will also ... it's a repository, not a registry page 7 and 8: Ensure that the BRMS Platform is not running, either by un-deploying it or halting your JBoss EAP server. You have to stop the EAP before you do step #2 (copying classes into the lib direcotry, otherwise you end up with a ClassNotFoundException page 8: Eg. ${INSTALL}/server/production/lib/ jboss-as dir is missing Eg. ${INSTALL}/jboss-as/server/production/lib/ page 11: org.JBoss.security.auth.spi.UsersRolesLoginModule should be lowercase JB directory ${INSTALL}/JBoss-as/server lowercase JB (several times on this page) step #3 and #4: missing sample xml snippets page 13, step #3: missing sample xml snippets page 14: This panel can be found be going to the Navigate panel 'be' should be 'by' page 15: The permissions assigned to a use can be updated 'use' should be 'user' page 18: from the BRMSWAR file no space between BRMS and WAR page 23: feature areas mentioned below figure 4.1 not in sync with this figure page 26: is a still a single category one 'a' needs to be removed page 27: and this action this will create a new item second 'this' needs to be removed does not effect its execution 'effect' should be 'affect' page 30: figure 4.6, 4.7, 4.8, 4.9, 4.10, 4.17, 4.18, 4.23 from old jbrms page 40: The picture bove shows 'bove' should be 'above' page 44: shows the alternativee way of exploring 'alternativee' should be 'alternative'	0	0	https://issues.jboss.org/browse/BRMS-219
10784	BRMS-219	page 47: those dervied from using 'dervied' should be 'derived' page 48: IMPORTANT: ... This section should be in orange I guess. page 49: see the "Package URI" 'Package URI' should be 'Deployment URL' then replace "NewSnapshot" with "LATEST". 'NewSnapshot' should be 'TEST' page 50: There are also many third-party WebDav clients are available double 'are' http://10.64.5.95:8080/jboss-brms localhost, <server> and <yourserver> is used everywhere else page 53: that are be the center doesn't sound correct... that *is* the center? *they* are the center? page 57: one the tool-bar these are 'one' should be 'on' page 58: There are two method of should be 'methods' page 62: Not that the wizard will not allow 'Not' should be 'Note'	0	0	https://issues.jboss.org/browse/BRMS-219
10785	BRMS-211	I think support for this is there, have a look at the screen shot (it's based on trunk), however it's not working if I set a name... still fact0, fact1, ... are used. Could you please create a corresponding GUVNOR jira for this issue and link them?	0	0	https://issues.jboss.org/browse/BRMS-211
10786	BRMS-134	Is the problem here the choice of domain for the example ?	0	0	https://issues.jboss.org/browse/BRMS-134
10787	BRMS-109	Hi Len, can you be more specific on this issue please? Not sure what you mean by "DB configuration process" in BRMS.	0	0	https://issues.jboss.org/browse/BRMS-109
10788	BRMS-84	Jarek - can you take the first look at this? Thanks!	1	1	https://issues.jboss.org/browse/BRMS-84
10789	BRMS-84	KnowledgeBuilder to modified. strange wording, maybe this sounds better: This allows the behavior of the KnowledgeBuilder to *be* modified. or: This allows to modify the KnowledgeBuilder's behavior. ---------- strange wording: The first parameter is for Properties and is optional and may be left null for the default options will be used. ---------- ======== PAGE 26 ======== First 6 lines of example 3.6 should be reomoved: KnowledgeBuilder kbuilder = KnowledgeBuilderFactory.newKnowledgeBuilder(); if(kbuilder.hasErrors()) { System.out.println(kbuilder.getErrors()); return; } ======== PAGE 29 ======== KnowledgePackages - the "s" at the end is bold, but shouldn't be ======== PAGE 30 ======== KnowledgeBaseFactory.createKnowledgeBaseConfiguration(null, cl); should be: KnowledgeBaseFactory.newKnowledgeBaseConfiguration(null, cl); ======== PAGE 31 ======== sample 3.15: the ">" and "<" chars have been replaced with < and > Also @SuppressWarnings is missing and it wouldn't compile anyway - KnowledgePackage and not KnowledgePackages and missing a cast Collection<KnowledgePackages> kpkgs = in.readObject( ); should be: @SuppressWarnings("unchecked") Collection<KnowledgePackage> kpkgs = (Collection<KnowledgePackage>)in.readObject( ); ======== PAGE 32 ======== The KnowlegeBuilder is created using the KnowledgeBuilderFactory. should be: The KnowlegeAgent is created using the KnowledgeAgentFactory. ======== PAGE 33 ======== remove one "for" This is also true for for notification. ======== PAGE 34 ======== KnowledgeBaseFactory.createKnowledgeBaseConfiguration(null, cl); should be: KnowledgeBaseFactory.newKnowledgeBaseConfiguration(null, cl); KnowledgeBase kbase KnowledgeBaseFactory.newKnowledgeBase(kbaseConf); should be: KnowledgeBase kbase = KnowledgeBaseFactory.newKnowledgeBase(kbaseConf); Example 3.21. ChangeSet XML which adds a directories contents I think this should be: Example 3.21. ChangeSet XML which adds a directorie's contents ======== PAGE 35 ======== KnowledgeAgentConfiguration kaconf = KnowledgeAgentFactory.newKnowledgeAgentConfiguation(); typo in newKnowledgeAgentConfiguation, missing "r", should be: KnowledgeAgentConfiguration kaconf = KnowledgeAgentFactory.newKnowledgeAgentConfiguration(); Example 3.23. Creating a new KnowledgeBuilder should be: Example 3.23. Creating a new KnowledgeBase ======== PAGE 36 ======== Example 3.24. Add KnowledgePackages to a KnowledgeBase should be: Example 3.24. Creating a new StatefulKnowledgeSession this provides a unified API to work with process and rules. should be process*es*: this provides a unified API to work with processes and rules. ======== PAGE 37 ======== It is common to hear the two terms used interchangably; so expect to hear the two used interchangeably. I would remove the second part ;) and fix the typo in interchangably (should be interchangeably) --------- As mentioned in the KnowledgeBase section a WorkingMemory may operate in one of two assertion modes: equality or identity. That hasn't been mentioned before, so I would change it to: A WorkingMemory may operate in one of two assertion modes: equality or identity. ======== PAGE 38 ======== The Rule Engine must be notified of modified Facts, so that they can be reprocessed. should be lower case "f" in Facts: The Rule Engine must be notified of modified facts, so that they can be reprocessed. -------- Within a rule the modify keyword is supported and provides block setters. I would rephrase it: Within a rule the modify statement is supported and provides calls to objects' setters. -------- The WorkingMemory provides access to the Agenda, query executions as well getting access to named EntryPoints. I would reword it: The WorkingMemory provides access to the Agenda, query execution results as well as named EntryPoints. ======== PAGE 40 ======== The KnowledgeRuntime provides further methods that are applicableto both rules and processes. "applicableto" - missing space ======== PAGE 43 ======== when calling FireAllRules(). lowercase "f", should be: when calling fireAllRules(). ======== PAGE 50 ======== kbuilder.add(ResourceFactory.newFileSystemResource(fileName), ResourceType.DRL); should be:	1	1	https://issues.jboss.org/browse/BRMS-84

#	Requirement	Questions	Answered	Successfully Answered	URL
10790	BRMS-83	When to use a BRMS Text: 1. You need to manage versions/deployment of rules Comment: This might read better: You need to manage versioning/deployment of rules ---------- Page 3, Section 2.1. Supported and recommended platforms Comment: This section should list Java 1.5 and 1.6 in the set of supported platforms. ---------- Page 3, Section 2.2. Installation Text: 2. Extract Files The JBoss Enterprise BRMS Platform package contains two files, jboss-brms.war and jboss-brms-engine.zip. jboss-brms.war is a Java web archive containing the BRMS application. jboss-brms-engine.zip contains the jar files for the JBoss Rules 5 rules engine. These files are a subset of jboss-brms.war. Comment: The jboss-brms.CR2.zip file contains these (2) files: jboss-brms-engine.zip jboss-brms-manager.zip The jboss-brms-manager.zip file contains the jboss-brms.war file only. ---------- Text: 3. "explode" the web archive Create a directory called jboss-brms.war and use a file archive tool to unpack the web archive into this directory. Comment: We should add this text as a warning: Be careful to name the directory "jboss-brms.war" - the ".war" is needed or the BRMS will not be properly deployed by the EAP server. ---------- Page 5, Figure 2.2. BRMS first screen after login Comment: This screenshot shows an IP address of 10.64.5.95. I'm guessing that this is a system in Brisbane. It would be safer to use localhost if we have time to replace the screenshot. ---------- Page 7, Procedure 2.2. Changing the repository location Text: 5. Move existing data store if required A new data store will be created by BRMS at the new location if there isn't one there already. If you wish to keep your existing data store then you can copy your existing files to the new location. Comment: Add this at the end of the above paragraph: "before you restart your server." ---------- Page 8, Section 2.3.2. BRMS Database Configuration Text: The means by which the BRMS Platform manages access to these databases is handled by a Persistence Manager. There is a generic Persistence Manager for use with any JDBC compliant database as well as several specific to particular database implementations, such as Oracle or MySQL. Comment: We should remove the reference to MYSQL here as we do not yet support that DB. ---------- Page 8, Section 2.3.2. BRMS Database Configuration Warning If you already have assets stored using the existing configuration that you wish to keep then you should export those before performing these steps. You can then import them into your new database afterwards. Comment: We should add a note to the warning that exporting and then importing a Rules repository is not supported as a backup and restore mechanism. Users should be told to set up a regular database backup procedure for their Worskpace and Versions databases. ---------- Page 18, 2.6.1. Backups Comment: We should repeat the note from 2.3.2 - instructing the users to not depend on export/import as a backup mechanism. ---------- Page 18, 2.6.2. Asset list customization Comment: This section refers to a AssetListTable.properties file. That file is not included in the .war: [ldimaggi@ldimaggi jboss-brms.war]$ pwd /jboss/local/EAP/jboss-eap-4.3/jboss-as/server/all/deploy/jboss-brms.war [ldimaggi@ldimaggi jboss-brms.war]$ ll total 56 -rw-rw-r-- 1 ldimaggi ldimaggi 3825 Apr 17 17:23 drools_logo.gif -rw-rw-r-- 1 ldimaggi ldimaggi 131 Apr 17 17:23 index.jsp drwxrwxr-x 3 ldimaggi ldimaggi 4096 Apr 17 19:57 META-INF -rw-rw-r-- 1 ldimaggi ldimaggi 130 Apr 17 17:23 NOTE.txt drwxrwxr-x 5 ldimaggi ldimaggi 4096 May 8 11:04 org.drools.guvnor.Guvnor drwxrwxr-x 2 ldimaggi ldimaggi 4096 Apr 17 19:57 org.drools.guvnor.Guvnor-aux drwxrwxr-x 4 ldimaggi ldimaggi 4096 May 12 13:41 WEB-INF [ldimaggi@ldimaggi jboss-brms.war]$ find . -name AssetListTable.properties -print ---------- Page 18, 2.6.3. Customized selectors for package building Comment: Selectors should be marked as tech preview for this release. ---------- Page 19, Section 2.6.4. Customizing the BRMS Platform User Interface Text: The most common branding change is to replace the images jbossrules_hdrbkg_blue.gif and drools.gif. Respectively these images are the logo at the top of the	1	0	https://issues.jboss.org/browse/BRMS-83
10791	BRMS-83	More info from Michael - thanks Michael! Yes - it would not be a sys admin but someone with specific rules expertise. Len DiMaggio wrote: > Hey Michael, > > Do we want the admins to configure the packages? Or maybe just users who have a better understanding of the customer's app? I vote for the second option - what do you think? > > > --Thanks, Len > > > > > > Page 38, 4.2.7. Package management > > Text: > > Configuring packages is generally something that is done once, and by someone with some > experience with rules/models. Generally speaking, very few people will need to configure packages, > and once they are setup, they can be copied over and over if needed. Package configuration is most > definitely a technical task that requires the appropriate expertise. > > Comment: > > I will ask Michael Neale about this one - I would like for us to be able tell readers which type of user should perform this task. >	1	0	https://issues.jboss.org/browse/BRMS-83
10792	BRMS-83	1.5 and 1.6 in the set of supported platforms. ---------- Page 3, Section 2.2. Installation Text: 2. Extract Files The JBoss Enterprise BRMS Platform package contains two files, jboss-brms.war and jboss-brms-engine.zip. jboss-brms.war is a Java web archive containing the BRMS application. jboss-brms-engine.zip contains the jar files for the JBoss Rules 5 rules engine. These files are a subset of jboss-brms.war. Comment: The jboss-brms.CR2.zip file contains these (2) files: jboss-brms-engine.zip jboss-brms-manager.zip The jboss-brms-manager.zip file contains the jboss-brms.war file only. ---------- Text: 3. "explode" the web archive Create a directory called jboss-brms.war and use a file archive tool to unpack the web archive into this directory. Comment: We should add this text as a warning: Be careful to name the directory "jboss-brms.war" - the ".war" is needed or the BRMS will not be properly deployed by the EAP server. ---------- Page 5, Figure 2.2. BRMS first screen after login Comment: This screenshot shows an IP address of 10.64.5.95. I'm guessing that this is a system in Brisbane. It would be safer to use localhost if we have time to replace the screenshot. ---------- Page 7, Procedure 2.2. Changing the repository location Text: 5. Move existing data store if required A new data store will be created by BRMS at the new location if there isn't one there already. If you wish to keep your existing data store then you can copy your existing files to the new location. Comment: Add this at the end of the above paragraph: "before you restart your server." ---------- Page 8, Section 2.3.2. BRMS Database Configuration Text: The means by which the BRMS Platform manages access to these databases is handled by a Persistence Manager. There is a generic Persistence Manager for use with any JDBC compliant database as well as several specific to particular database implementations, such as Oracle or MySQL. Comment: We should remove the reference to MYSQL here as we do not yet support that DB. ---------- Page 8, Section 2.3.2. BRMS Database Configuration Warning If you already have assets stored using the existing configuration that you wish to keep then you should export those before performing these steps. You can then import them into your new database afterwards. Comment: We should add a note to the warning that exporting and then importing a Rules repository is not supported as a backup and restore mechanism. Users should be told to set up a regular database backup procedure for their Worskpace and Versions databases. ---------- Page 18, 2.6.1. Backups Comment: We should repeat the note from 2.3.2 - instructing the users to not depend on export/import as a backup mechanism. ---------- Page 18, 2.6.2. Asset list customization Comment: This section refers to a AssetListTable.properties file. That file is not included in the .war: [ldimaggi@ldimaggi jboss-brms.war]$ pwd /jboss/local/EAP/jboss-eap-4.3/jboss-as/server/all/deploy/jboss-brms.war [ldimaggi@ldimaggi jboss-brms.war]$ ll total 56 -rw-rw-r-- 1 ldimaggi ldimaggi 3825 Apr 17 17:23 drools_logo.gif -rw-rw-r-- 1 ldimaggi ldimaggi 131 Apr 17 17:23 index.jsp drwxrwxr-x 3 ldimaggi ldimaggi 4096 Apr 17 19:57 META-INF -rw-rw-r-- 1 ldimaggi ldimaggi 130 Apr 17 17:23 NOTE.txt drwxrwxr-x 5 ldimaggi ldimaggi 4096 May 8 11:04 org.drools.guvnor.Guvnor drwxrwxr-x 2 ldimaggi ldimaggi 4096 Apr 17 19:57 org.drools.guvnor.Guvnor-aux drwxrwxr-x 4 ldimaggi ldimaggi 4096 May 12 13:41 WEB-INF [ldimaggi@ldimaggi jboss-brms.war]$ find . -name AssetListTable.properties -print ---------- Page 18, 2.6.3. Customized selectors for package building Comment: Selectors should be marked as tech preview for this release. ---------- Page 19, Section 2.6.4. Customizing the BRMS Platform User Interface Text: The most common branding change is to replace the images jbossrules_hdrbkg_blue.gif and drools.gif. Respectively these images are the logo at the top of the interface and the site "favorites icon" Comment: The jbossrules_hdrbkg_blue.gif is displayed multiple times on the top of the UI display. ---------- Page 20, 2.6.5. Import and Export Text: This should not be used as a substitute for the backup system provided by	1	0	https://issues.jboss.org/browse/BRMS-83

#	Requirement	Questions	Answered	Successfully Answered	URL
10793	BRMS-83	of rules Comment: This might read better: You need to manage versioning/deployment of rules ---------- Page 25, 4.2.1. Rules are assets Text: As the BRMS can manage many different types of rules (and more), they are all classed as "assets". An asset is anything that can be stored as a version in the repository. This includes decision tables, models, DSLs and more. Sometimes the word "rule" will be used to really mean "asset" (i.e. the things you can do also apply to the other asset types). You can think of asset as a lot like a file in a folder. Assets are grouped together for viewing, or to make a package for deployment etc. Comment: Can we reword this to: As the BRMS can manage many different types of rules and other rules-related data, which are referred to as "assets". An asset is anything that can be stored as a version in the repository. This includes decision tables, models, DSLs and more. In this guide and in the UI, the word "rule" is sometimes used mean "asset" (i.e. the actions you can perform on rules can also be performed on the other asset types). Assets are grouped together for viewing, or to make a package for deployment. ---------- Page 31, 4.2.4.1.3. A more complex example: Text: In the above example, you can see it is using a mixture of literal values, and formulas. The second constraint on the "Person" fact, is a formula (in this case it is doing a silly calculation on the persons age, and checking something against their name - both "age" and "name" are fields of the Person fact in this case. In the 3rd line (which says "age is less than .." - it is also using a formula, although, in this case the formula does a calculation and returns a value (which is used in the comparison) - in the former case, it had to return True or False (in this case, its a value). Obvious formulas are basically pieces of code - so this is for experienced users only. Comment: Can we reword the last line to: "Formulas are pieces of executable code - so this is for experienced users only." ---------- Page 34, Section 4.2.4.5. Technical rules (drl) Text: Technical (drl) rules are stored as text - they can be managed in the BRMS. A DRL can either be a whole chunk of rules, or an individual rule. if its an individual rule, no package statement or imports are required (in fact, you can skip the "rule" statement altogether, just use "when" and "then" to mark the condition and action sections respectively). Normally you would use the IDE to edit raw DRL files, since it has all the advanced tooling and content assistance and debugging, however there are times when a rule may have to deal with something fairly technical. In any typical package of rules, you generally have a been for some "technical rules" - you can mix and match all the rule types together of course. Comment: Can we reword this as: Technical (drl) rules are stored as text and can be managed in the BRMS. A DRL can either be an individual rule or multiple rules. If it's an individual rule, no package statement or imports are required. You can skip the "rule" statement altogether and just use "when" and "then" to mark the condition and action sections respectively. Normally you would use an IDE to edit raw DRL files, since it has all the advanced tooling to support editing, content assistance and debugging. Also - we have to ask Michael what this means - how does a rule being "technical" require that it be edited in the BRMS? However there are times when a rule may have to deal with something fairly technical. In any typical package of rules, you generally have a been for some "technical rules" - you can mix and match all the rule types together of course. ---------- Page 35, Section 4.2.4.7. Data enumerations (drop down list configurations) Text: Data enumerations are an optional asset type that technical folk can configure to provide drop down lists for the guided editor. These are stored and edited just like any other asset, and apply to the package that they belong to. Comment: Can we reword the paragraph to read: Data enumerations are an optional asset type that can be configured to provide drop down lists for the guided editor. These are stored and edited just like any other asset, and apply to the package that they belong to. ---------- Page 37, 4.2.4.8. Advanced enumeration concepts Text: Mode advanced enumerations: In the above cases, the values in the lists are calculated up front. This is fine for relatively static data, or	1	0	https://issues.jboss.org/browse/BRMS-83
10794	BRMS-83	From Michael - thanks! Try this: Technical (drl) rules are stored as text - they can be managed in the BRMS. A DRL can either be a whole chunk of rules, or an individual rule. if its an individual rule, no package statement or imports are required (in fact, you can skip the "rule" statement altogether, just use "when" and "then" to mark the condition and action sections respectively). Normally you would use the IDE to edit raw DRL files, since it has all the advanced tooling and content assistance and debugging. However, there are times when a rule may have to deal with something fairly technical in a package in Guvnor. In any typical package of rules, you generally have a need for some "technical rules" - you can mix and match all the rule types together of course. Len DiMaggio wrote: > Hey Michael, > > Question on something in the user guide - can you add to this? > > > Page 34, Section 4.2.4.5. Technical rules (drl) > > Text: > > Technical (drl) rules are stored as text - they can be managed in the BRMS. A DRL can either be a > whole chunk of rules, or an individual rule. if its an individual rule, no package statement or imports > are required (in fact, you can skip the "rule" statement altogether, just use "when" and "then" to mark > the condition and action sections respectively). Normally you would use the IDE to edit raw DRL files, > since it has all the advanced tooling and content assistance and debugging, however there are times > when a rule may have to deal with something fairly technical. In any typical package of rules, you > generally have a been for some "technical rules" - you can mix and match all the rule types together of > course. > > Comment: > > ------> How does a rule being "technical" require that it be edited in the BRMS? Maybe this needs rewording? > > However there are times when a rule may have to deal with something fairly technical. In any typical package of rules, you > generally have a been for some "technical rules" - you can mix and match all the rule types together of > course.	1	0	https://issues.jboss.org/browse/BRMS-83
10795	BRMS-83	Document has been updated with the above feedback. Feedback above has been incorporated with the exception of a couple of "formatting issues" (no time sorry) and the items below which need more information: Page 33, Section 4.2.4.3. Spreadsheet decision tables Text: To use a spreadsheet, you upload an xls (and can download the current version, as per the picture above). To create a new decision table, when you launch the rule wizard, you will get an option to create one (after that point, you can upload the xls file). Comment: We should also state that BRMS currently supports uploading from Microsoft Excel(tm) spreadsheets. ---------- Is that to say we should specify that when we say xls we specifically mean Microsoft Excel(tm) spreadsheets? Page 36, 4.2.4.8. Advanced enumeration concepts Comment: The references to ULP and PULP types of fuel may be lost on readers who are not in Australia. Maybe we can add a footnote: http://www.fuelwatch.wa.gov.au/info/dsp_fuel_types.cfm ---------- Examples in less specific domains would be preferred but adding further clarification is not really necessary. The description of the problem domain is fairly clear. Page 20, 2.6.5. Import and Export We've seen failures on importing large repos. We may need to add some additional warnings to this section. Let's talk to Darrin about how to describe this. ---------- what sort of errors? are there any plans to address this? JIRA? should I include it as a known issue in the release notes?	1	0	https://issues.jboss.org/browse/BRMS-83
10796	BRMS-82	updated, committed & sent to soa-pm-list aswell. Can you clarify this statement from the comments above? "The only JVMs that we should list are Sun 1.5, Sun 1.6, and Apple 1.5 - note that this is for the BRMS UI. " I'm a bit confused here, does this mean the BRMS UI has a client side java dependency (applets or something) ? And that we are supporting JDK1.5 on OSX for that, but only the Sun JDKs on the server ?	0	0	https://issues.jboss.org/browse/BRMS-82
10797	BRMS-8	Are you sure that's not a cache issue? If it's the same in CR1 can you attach a screenshot	0	0	https://issues.jboss.org/browse/BRMS-8
10798	BRMS-2	What's the status?	0	0	https://issues.jboss.org/browse/BRMS-2
10799	EDG-122	Rado, as you said this was executed with CR3 and precisely for CR4 I fixed messages and providing clearer exceptions. Closing now. Please reopen if no satisfied with CR4.	0	0	https://issues.jboss.org/browse/EDG-122
10800	EDG-121	Michal, what was the outcome of this? Seems like we're going with 2.2.x, so Infinispan would need to upgraded to REST 2.2.x, right?	0	0	https://issues.jboss.org/browse/EDG-121

#	Requirement	Questions	Answered	Successfully Answered	URL
10801	EDG-31	Why not Netty? Bill Burke had some interest in running RestEasy on top of Netty and that makes sense in my opinion. Less dependency, too.	0	0	https://issues.jboss.org/browse/EDG-31
10802	EDG-28	What should the official bundle look like? Does it look different from the usual tar/zip distribution? Any pointers would be appreciated.	0	0	https://issues.jboss.org/browse/EDG-28
10803	EDG-28	What I thought by official build is basically a somehow stabler tag - the build process and format of the distribution is not any different. I'd like to see it based on tagged AS 7.0.0.CR1 that's about to be released a tagged version of infinispan, that means probably 5.0.0.CR6 also I'd like to establish a nomemclature for later tags let's say ALPHA1, ALPHA2 etc. so that we can report against a concrete version in JIRA (set proper affectsVersion, fixVersion) so that when we take ALPHAx and find serious blockers in it, new ALPHA(x+1) is released and every change and bug report is trackable. the customer will get one of the later ALPHA tags, that would be the result of our QA process wdyt ?	0	0	https://issues.jboss.org/browse/EDG-28
10804	EDG-28	So, what you mean is tag Alpha1 real soon and keep fixing the known issues and keep tagging, and then provide the latest alpha (e.g. alpha5) to the customer when it's ready?	0	0	https://issues.jboss.org/browse/EDG-28
10805	EDG-23	ISPN-1267 has been resolved again. Michal, can you verify it looks good on your side now and close if necessary?	1	1	https://issues.jboss.org/browse/EDG-23
10806	EDG-22	Is this still an issue with 5.0.0.CR6?	0	0	https://issues.jboss.org/browse/EDG-22
10807	EDG-20	I think this issue can be better phrased as 'REST server cannot access the entries written by other protocol servers'. Can the REST server access the entries written via Hot Rod or direct calls?	0	0	https://issues.jboss.org/browse/EDG-20
10808	EDG-20	No, all server modules are basically isolated. HotRod is isolated from others by means of different key "encoding". All keys are stored as ByteArrayKey. In memcached and REST they are stored as plain String. It would be cool if we were able to share data between server modules, but there is one fundamental problem. All three server modules store different metadata with the entry. We would have to unite the way how the metadata is stored for each module and define "conversions" and default values for metadata that isn't written by one but read by the other. E.g. REST module writes mime type of the entry on PUT operation. HotRod doesn't. What kind of mime would be supplied to a value written by hotrod and read by REST ? I'm all for starting a design document on this, but so far the consensus I think was to isolate the server modules (via isolating their keyspaces).	0	0	https://issues.jboss.org/browse/EDG-20
10809	EDG-20	What does isolating keyspace really mean? Is it having a different EmbeddedCacheManager? Then does that mean we have to create three cache manager in a single VM?	0	0	https://issues.jboss.org/browse/EDG-20
10810	EDG-16	Michal, should this be closed then?	1	1	https://issues.jboss.org/browse/EDG-16
10811	EDG-15	Michal, can you check that with CR8 this is gone?	1	1	https://issues.jboss.org/browse/EDG-15
10812	EDG-14	What would be the desired behavior? Do you want Netty not to rename threads at all?	1	0	https://issues.jboss.org/browse/EDG-14
10813	EDG-12	1) Understood. See here for domain schema requirements: http://community.jboss.org/wiki/DomainRequirements 2) It covers *most* of it. The modifications can be largely summarized as: * More concise * Prune stuff that the AS handles internally. e.g. everything in <global>. Things like channel, thread pools, mbean server, transaction manager, are all injected from their respective subsystems. * Prune obscure configuration options (e.g. multiple cache loaders) * Only expose configuration options that are relevant for a given cache mode. e.g. don't expose rehash options for non-distributed caches * Provides sensible per cache mode custom default values * Combine multiple dependent options wherever possible (e.g. combine multiple booleans into single enum) The schema can be found here: https://github.com/jbossas/jboss-as/blob/master/clustering/src/main/resources/schema/jboss-infinispan.xsd Are there configuration options that you see missing? While the AS7 and EDG release schedules will vary, the EDG/EAP6 release schedules probably won't vary much. One of the domain schema requirements is that the domain schema should be *stable*. So, there *shouldn't* be any schema changes across minor releases. Changes across major releases are handled by versioning the schema itself.	0	0	https://issues.jboss.org/browse/EDG-12
10814	JBEPP-1468	The requested content has been included in the EPP Reference Guide. The new version (5.2.1-5) has been pushed to the Docs Stage.	1	1	https://issues.jboss.org/browse/JBEPP-1468
10815	JBEPP-1451	Matt, can you go ahead and commit for English language ?	1	0	https://issues.jboss.org/browse/JBEPP-1451
10816	JBEPP-1441	Scott, can you please include this in the documentation?	1	1	https://issues.jboss.org/browse/JBEPP-1441
10817	JBEPP-1430	Can you be more specific about what you would add to: http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Portal_Platform/5.1/html-single/Installation_Guide/index.html#sect-Reference_Guide-Database_Configuration ? Beside a warning note that the 2 DB must be different ?	0	0	https://issues.jboss.org/browse/JBEPP-1430
10818	JBEPP-1365	From Release Note text on the dev issue: {quote} The fix introduces JBossClusteredSSOValveFilter, which removes the patching and workarounds customers had to implement in earlier versions of the product, and increases overall platform security.{quote} Hi Marek The new ValveFilter sounds like a good thing to document, and replace the current workaround documented in the Reference Guide (which uses BASIC authentication, currently the subject of that worm). Can you link me to any info about the Filter, and what parameters it accepts? Even better, would it be possible to get a real-world configuration example of the ValveFilter in a directive file? Cheers Jared	1	0	https://issues.jboss.org/browse/JBEPP-1365
10819	JBEPP-1361	Hi Jared, This jira is not something, which needs to be secure. I set it with security level RH+EXO probably by accident, sorry for that:-/ Now I made it public. I returned release notes status back to "None". Is some more info needed from me? Because "Release notes text" already have some value, which can be used to identify purpose of this jira.	1	0	https://issues.jboss.org/browse/JBEPP-1361
10820	JBEPP-1350	Thank Marek, I've commented the points you mentioned, but would it be possible for you to give me some guidance on removing references to Tomcat altogether? (if EPP doesn't run on it, it probably shouldn't be in the documentation) There's some code snippets that reference tomcat, some URLs that seem Tomcat specific, notes on changing Tomcat port numbers and instructions to download and extract Tomcat. Which of those can be changed and what should they be changed to? Thanks	1	1	https://issues.jboss.org/browse/JBEPP-1350
10821	JBEPP-1290	I think we would benefit from it once we'll be able to consume EAP 5.1.2 (That should be in EPP 5.2.1), is that ok to do that or do we really have to override the versions shipped with Seam ? Thomas	1	1	https://issues.jboss.org/browse/JBEPP-1290
10822	JBEPP-1117	Thomas, not sure if this is a problem, can you please check this?	0	0	https://issues.jboss.org/browse/JBEPP-1117
10823	JBEPP-1098	Is this new info bar functionality documented anywhere that can be linked to in a release note?	1	0	https://issues.jboss.org/browse/JBEPP-1098
10824	JBEPP-1098	Based on Theute's comment, the Release Notes text has been redrafted. Old text: In previous versions of JBoss Enterprise Portal Platform, new portlets added to a page did not, by default, display the 'info bar' which usually surrounds the portlet. In this release, behaviour to have the 'info bar' displayed by default is available under the portal config settings. New text: This release of JBoss Enterprise Portal Platform adds a configurable setting to control whether the 'info bar' which surrounds new portlets is displayed by default. Is the revised text technically accurate?	1	0	https://issues.jboss.org/browse/JBEPP-1098

#	Requirement	Questions	Answered	Successfully Answered	URL
10825	JBEPP-1016	Busted, I did that too quickly ;) Matt can you replace the text ? The previous one: "Add priority for skin module css loading" was not very detailed.	0	0	https://issues.jboss.org/browse/JBEPP-1016
10826	JBEPP-1016	Is the new release text ok? Or does it need more details?	1	0	https://issues.jboss.org/browse/JBEPP-1016
10827	JBEPP-1013	If the underlying redirection handling has not been corrected, will this issue not just present again if a user points the iFrame to a site that does use redirection?	1	1	https://issues.jboss.org/browse/JBEPP-1013
10828	JBEPP-952	Can we also get the work done for GTNPORTAL-1706 in, please? (r6010 and r6136)	0	0	https://issues.jboss.org/browse/JBEPP-952
10829	JBEPP-894	I've checked it and everything seems to be ok now. Thanks Scott! Should I resolve this jira or you will?	1	1	https://issues.jboss.org/browse/JBEPP-894
10830	JBEPP-800	Scott, can you please add a comment at Chapter 8 (Portlets and Gadgets Administration) about JBEPP-927? It can be a Note or Warning that gadgets which are not used on Dashboard do not properly store preferences. It's about gadgets which you can add on portal pages in Gadget Wrapper Portlet. See comment from Thomas as well. This can be also added in Rel.notes as a known issue. Thanks	1	1	https://issues.jboss.org/browse/JBEPP-800
10831	JBEPP-780	Should this still be marked as a Known Issue? The problem causing the issue seems to have been corrected and the JIRA re-opened for a different reason (if I'm understanding the comments correctly)	0	0	https://issues.jboss.org/browse/JBEPP-780
10832	JBEPP-779	I am reopening since the main purpose of this Jira (setting multiplexer channel for MOPSessionManager) is not done. Shouldn't we create JIRA in eXo kernel for multiplexer channel support ?	0	0	https://issues.jboss.org/browse/JBEPP-779
10833	JBEPP-766	Download link in the page [1] fails with 'An Unexpected Error Has Occurred. ' Also the patch is not listed in [2]. Could you check it please? [1] https://support.redhat.com/jbossnetwork/restricted/softwareDetail.html?softwareId=5183 [2] https://access.redhat.com/jbossnetwork/restricted/listSoftware.html?downloadType=patches&product=jbportal&version=5.1.0	1	1	https://issues.jboss.org/browse/JBEPP-766
10834	JBEPP-735	Which jgroups configuration files ? (there are multiple ones for multiple channels) ? In general, you can define where to pickup the configuration file, so that you can always leave the one that we package intact in its jar and just specify a new location for a custom one.	0	0	https://issues.jboss.org/browse/JBEPP-735
10835	JBEPP-627	This table has been separated into EPP 5.0 and 5.1 specific sections. Entries for Oracle 11g R2 and Oracle 11g R2 RAC have been added to the 5.1 section. These new entries require version numbers.	1	1	https://issues.jboss.org/browse/JBEPP-627
10836	JBEPP-618	Thomas, would the instruction still be relevant for users on EPP 5.0?	1	1	https://issues.jboss.org/browse/JBEPP-618
10837	JBEPP-618	But I verified and it was already in EAP 5.0 as well, and should be removed from both versions of the doc	0	0	https://issues.jboss.org/browse/JBEPP-618
10838	JBEPP-616	Can you attach or list out all queries generated during CRUD of user and role?	0	0	https://issues.jboss.org/browse/JBEPP-616
10839	JBEPP-616	I have been logging sql queries for simple CRUD use cases (attached in a zip file). Then i've been counting the select queries with grep 'select this' sql.log | wc -l for different jboss-cache eviction algorithm set. there was no difference between algorithms, the eviction seems to be fully satysfying since it keeps all the data cached as long as possible. Anyway, there seems to be a place for improvement, since i.e. the query (in create.log) 2011-02-08 15:17:15,489 DEBUG [org.hibernate.SQL] org.jboss.logging.Logger select this_.ID as ID4_2_, this_.IDENTITY_TYPE as IDENTITY2_4_2_, this_.NAME as NAME4_2_, this_.REALM as REALM4_2_, type1_.ID as ID11_0_, type1_.NAME as NAME11_0_, rm2_.ID as ID0_1_, rm2_.NAME as NAME0_1_ from jbid_io this_ inner join jbid_io_type type1_ on this_.IDENTITY_TYPE=type1_.ID inner join jbid_realm rm2_ on this_.REALM=rm2_.ID where this_.NAME=? and type1_.NAME=? and rm2_.NAME=? is called several times with the same parameters. I will debug to understand the process better. my use case: start of startup.log 1. start clean EPP, after start 2. open "/portal/public/classic/" end of startup.log 3. click on login as root icon (bottom left) 4. click on Group -> Organization -> New staff start of create.log 5. fill in all fields, click save 6. click OK on modal confirmation window end of create.log 7. click on Group -> Organization -> Users and groups management start of retrieve.log 8. fill in the previously created username into the search field 9. click on search link end of retrieve.log 10. click on the edit button for that user start of update.log 11. change his first name 12. click save 13. click OK on modal confirmation window end of update.log 14. fill in the previously created username into the search field 15. click on search link start of remove.log 16. click on the delete button for that user 17. click ok on both modal confirmation windows end of remove.log	1	1	https://issues.jboss.org/browse/JBEPP-616
10840	JBEPP-592	Sorry, there is one change. If you change Template value it take effect, but after signing out. So this should be changed immediately. And what about some validation of Template input? Is it possible so that it could validate template path?	0	0	https://issues.jboss.org/browse/JBEPP-592
10841	JBEPP-592	One more comment: - if you change value of footer portlet from par:/groovy/groovy/webui/component/UIFooterPortlet.gtmpl to "text in footer" - it is saved properly, finish Edit Layout - sign out Error in log: 08:36:29,140 ERROR [FutureCache] Computing of resource org.exoplatform.resolver.ResourceKey@d2ace7e5 threw an exception java.lang.NullPointerException: No null input stream accepted at org.exoplatform.commons.utils.IOUtil.getStreamContentAsBytes(IOUtil.java:153) at org.exoplatform.groovyscript.text.TemplateService$1.retrieve(TemplateService.java:76) at org.exoplatform.groovyscript.text.TemplateService$1.retrieve(TemplateService.java:69) at org.exoplatform.commons.cache.future.FutureCache$1.call(FutureCache.java:83) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) at java.util.concurrent.FutureTask.run(FutureTask.java:138) at org.exoplatform.commons.cache.future.FutureCache.get(FutureCache.java:117) at org.exoplatform.groovyscript.text.TemplateService.getTemplate(TemplateService.java:154) ... 08:36:29,142 ERROR [portal:Lifecycle] Template: text in footer not found.	0	0	https://issues.jboss.org/browse/JBEPP-592
10842	JBEPP-569	Patch applied to the EPP INstallation Guide configuration section and edited for doc consistency. svn -r5051 Is there further documentation needed for HTTPS and Gadgets? Or does this process correct that issue?	0	0	https://issues.jboss.org/browse/JBEPP-569
10843	JBEPP-560	JIRAs created: JBPAPP-5313 - One-off to upgrade JBC to version 3.2.6 in EAP 5.1.0 JBPAPP-5323 - Upgrade JBC in EAP 5.1.1 to version 3.2.6 (or higher if available)	0	0	https://issues.jboss.org/browse/JBEPP-560
10844	JBEPP-545	There looks to be enough content in the PL IDM Reference Guide to create a standalone, albeit slim, document. Are we able to apply the JBEPP name? I.E: "JBoss Enterprise Portal Platform PicketLink IDM Reference Guide" Based on ECS metric estimates, this conversion should take about two weeks to complete (including revision/editing).	1	1	https://issues.jboss.org/browse/JBEPP-545
10845	JBEPP-541	Is there info on how to change the IDM configuration anywhere?	1	1	https://issues.jboss.org/browse/JBEPP-541
10846	JBEPP-541	Could you describe a usecase?	1	0	https://issues.jboss.org/browse/JBEPP-541

#	Requirement	Questions	Answered	Successfully Answered	URL
10847	JBEPP-541	smumford@redhat.com made a comment on [bug 793458|https://bugzilla.redhat.com/show_bug.cgi?id=793458] Martin, I'm not sure I see the connection between the request for instructions on how to implement a custom DB schema and the Kbase article you linked which is about LastLoginTime. Can you confirm that the URL [1] is correct, and if so, please explain how the two topics are related. Thanks 1: https://access.redhat.com/knowledge/solutions/74953	0	0	https://issues.jboss.org/browse/JBEPP-541
10848	JBEPP-541	smumford@redhat.com made a comment on [bug 793458|https://bugzilla.redhat.com/show_bug.cgi?id=793458] Thanks Martin, I've drafted some simple content based on the KBase article, trying to be generic while using the kbase usecase as an example. Can you please review it and recommend any changes in a comment. I won't commit it to the document until it's verified by an SME. Thanks ######### Draft Content ######################## Overview: In order to implement a custom class, the following steps must be taken: * Edit the 02portal.war/WEB-INF/conf/organization/idm-configuration.xml file to use the custom class. * Create the custom implementation class. The easiest way to do this is to copy and edit one of the java files found at: http://anonsvn.jboss.org/repos/gatein/epp/portal/tags/EPP_5_1_1_GA/component/identity/src/main/java/org/exoplatform/services/organization/idm/ DOCS NOTE: Is there any way for a user to get these from their local EPP instance? Or are these downloads the only way? * Edit or remove any methods in the class that may interfere with the customized behavior you are implementing. DOCS NOTE; A 'for further information about possible changes please refer to ...' would be good here. Example: The following is a step-by-step example of the changes outlined above. This example disables writing a user's last login time to the database. 1: Change the value for org.exoplatform.services.organization.OrganizationService in idm-configuration.xml: <component> <key>org.exoplatform.services.organization.OrganizationService</key> <!--type>original.ExampleServiceImpl</type--> <type>custom.ExampleServiceImpl</type> <init-params> Obtain a copy of http://anonsvn.jboss.org/repos/gatein/epp/portal/tags/EPP_5_1_1_GA/component/identity/src/main/java/org/exoplatform/services/organization/idm/PicketLinkIDMOrganizationServiceImpl.java 2: Change the constructor of the custom PicketLinkIDMOrganizationServiceImpl to use custom DAO implementations: public PicketLinkIDMOrganizationServiceImpl(InitParams params, PicketLinkIDMService idmService) throws Exception { groupDAO_ = new GroupDAOImpl(this, idmService); userDAO_ = new UserDAOImpl(this, idmService); userProfileDAO_ = new UserProfileDAOImpl(this, idmService); membershipDAO_ = new MembershipDAOImpl(this, idmService); membershipTypeDAO_ = new MembershipTypeDAOImpl(this, idmService); 3: Remove the following code from the authenticate(..) method in the custom UserDAOImpl (as it sets the last login time): if (authenticated) { UserImpl userImpl = (UserImpl)user; userImpl.setLastLoginTime(Calendar.getInstance().getTime()); saveUser(userImpl, false); }	0	0	https://issues.jboss.org/browse/JBEPP-541
10849	JBEPP-530	Is there more information on how this fix was implemented? More info on the workaround ported from earlier JB Portal?	1	0	https://issues.jboss.org/browse/JBEPP-530
10850	JBEPP-519	* Should SP Installation be documented; 1) in it's own document, 2) in the SP User Guide, 3) in an SP specific section of the EPP Installation guide or 4) a combination (ie; in both the SP User Guide and the EPP Install Guide)? * How many installation scenarios are there? Is there an EPP+SP unit that has SP pre-deployed, or is the only method to obtain both EPP and SP (separately) and install EPP then deploy SP into it?	1	1	https://issues.jboss.org/browse/JBEPP-519
10851	JBEPP-519	Doc Note: Is Site Publisher 'separate' enough from EPP to require a new book for installation, or does it fall under the same product, and hence the "collapsed docs" regime (ie, it would go into EPP Installation Guide under 'Site Publisher' section)?	1	1	https://issues.jboss.org/browse/JBEPP-519
10852	JBEPP-517	Section 2.7.3.1: Second sentence: Why is "webapplication" one word and bold font? I think that sentence should read: "This icon can be changed by adding an image to the following directory of the <bold>web</bold> application:" What is even more confusing about this sentence is that the directory shown does not exist. And then the next sentence shows an example of an image placed in a different directory. This makes me think the above-mentioned directory should be: skin/DefaultSkin/portletIcons/portlet_name.png where "portlet_name" is the name of the portlet in portlet.xml. That would make the example consistent with the explanation.	1	1	https://issues.jboss.org/browse/JBEPP-517
10853	JBEPP-517	Rich, would you mind combining few comments together so that we get less emails? I currently get one email for every comment you make.	0	0	https://issues.jboss.org/browse/JBEPP-517
10854	JBEPP-517	Pg 124: Inside the "Assumptions" note: Number 2 needs to be updated to WSRP 2.0.0 (I think?). Number 3 mentions EPP Version 5.0-GA. Does that need to be updated to 5.1? Section 15.4: Third paragraph: First sentence: "...is used to *accomplish* this." Page 126: Last paragraph: Missing a space between "the" and "org.gatein.pc.remotable". Page 130: Inside the "Values" note: Second sentence ends with a couple extra periods. Page 130: Paragraph that proceeds the "Values" note: First word should be "Enter". Section 15.6.3 has the same couple of paragraphs as Section 15.6.2.2 (exact copies of each other). These two sections need to be merged into one. Section 15.7.1.1: Second paragraph: Second sentence: "...a valid email and the *previous* email address is not..." Section 15.8.1: Second paragraph: First sentence: "...can be modified with *repsect* to whether..." Section 15.8.2: First sentence; "...register with it before providing access *to* its services." Page 141: Step 4: I think the "A" in "Add property" needs to be bold. Page 142: Inside "Note": Second sentence: "...deploy the policy implementation as *a* JAR file in the..."	1	1	https://issues.jboss.org/browse/JBEPP-517

#	Requirement	Questions	Answered	Successfully Answered	URL
10855	JBEPP-517	Page 145: Last sentence: "Exo" should be "eXo". (Never mind - this might take a while to change since it appears numerous times in this chapter.) Section 16.3.2: Fourth paragraph: "...type, and a method on *the* target component to use..." Section 16.4: Fourth paragraph: Sentence should end "...can be left empty." Page 150: Second paragraph: First sentence: "...defines one or more <property> *elements*, each of which..." Page 150: Paragraph following the first code sample: "...requires one or more <value> *elements*." Page 151: First paragraph: "...which is used for plain *Java* style object..." Page 153: Middle of the page: "The Unified servlet context and unified classloader are..." (No comma after *context*) Page 153: Following paragraph: "...archives, configured with *the* dependencies property..." Section 16.6: Second paragraph: Remove the comma after "...the original portal .war archives by adding a .war..". A comma there creates a pause that makes the sentence confusing. Section 16.7: End of fifth paragraph has some extra spacing (I think). The term "org.exoplatform.container.ExoContainerContext" is a single word. Page 155: End of third paragraph - same story: should be "org.exoplatform.container.web.AbstractFilter" Page 155: Paragraph following the "Note": First sentence: "...may also need access to portal information..." (remove the first *to*) Page 155: The following sentence has a semi-colon in the middle of it that should be removed. Section 17.2.1: Question: Why does the Portal Reference Guide discuss using eXo JCR *outside* of Portal, i.e., in standalone mode? (I'm only curious- there may be a good reason for it.) Page 160: Next to the last paragraph: Remove the comma after *with*. Same sentence: Replace "J2EE" with "JavaEE". Section 17.2.2: Same story as before: From what I can tell, this entire section is for the standalone mode of JCR. The config file that takes up pages 162 and 163 should be replaced with the EPP version of the file. (I can send that to you if you need it.) Page 168: Third paragraph: Should be "...with the following *Relational* Database Management...". (It's kind of funny how it reads now though!) Section 17.3.2: From my understanding, this entire section refers to the standalone mode of JCR only, not EPP. I could be wrong though. The config file shown on pages 169-170 looks nothing like the EPP one. Page 171 and 172: Can you define your own JCR workspaces in EPP? I am either lost or confused at this point. If this information stays in the Reference Guide, then we need a use-case to explain why we would do this. Section 17.3.3: I don't know if this is EPP either. If it is, I can't find these config files so we should clarify where they are. Page 190: Within the "Note" is a reference to the /webapps folder, which is Tomcat. Just take off "/webapps". Section 17.10.1.1: Section title - "Jboss" should be "JBoss"	1	1	https://issues.jboss.org/browse/JBEPP-517
10856	JBEPP-517	The above issues have been addressed in version 1-1.19 of the Reference Guide. One issue requires more information to resolve: Re Marek's comment: "5.2 - Reference to "access.control.workspace.roles" , which is probably does not correct as well." NeedInfo: I don't know what the correct code is here. "access.control.workspace.roles" appears in 5.0.0 portal-configuration.xml, but not as shown in code example. "access.control.workspace.roles" does not appear in 5.1(ER04) file at all. Can anyone let me know what the correct code sample(s) would be (I need to cover both 5.0 AND 5.1.)? See the attached text file for info about the resolution of other issues. https://jira.jboss.org/secure/attachment/12339012/Feedback-action.txt	1	1	https://issues.jboss.org/browse/JBEPP-517
10857	JBEPP-517	VIlliam; I can't locate the typos you outlined (I think the page numbers are different in the version you looked at). p93. typo in the "You" word - Y is missing p145. typo in <field> - "f" is missing Can you give me the section title and possibly some surrounding text to make searching easier please? Thanks.	1	1	https://issues.jboss.org/browse/JBEPP-517
10858	JBEPP-517	Villiam: Are you sure about changing all instances of kernel_1_0.xsd to kernel_1_1.xsd? I grepped the jboss-epp-5.1 instance I have locally and numerous hits for kernel_1_0.xsd were found in the 02portal.war files	1	1	https://issues.jboss.org/browse/JBEPP-517
10859	JBEPP-517	VIlliam: RE: p143-144 - confgis should be updated from 02portal.war/WEB-INF/conf/common/common.configuration.xml The code sample on pages 143 and 144 of the PDF I have refers to web.xml and doesn't resemble the code in common-configuration.xml at all. Can you provide a section heading and some surrounding text to ensure I'm seeing the same code sample as you, please.	1	1	https://issues.jboss.org/browse/JBEPP-517
10860	JBEPP-517	On the WSRP part of the doc: - Why is section 15.3.2 on a different page as the rest of the 15.3 section? - 15.5: the URLs there are actually incorrect. Both URLs only work in EPP 5.1. EPP 5.0 only has access to WSRP 1 but the URL is different: in EPP 5.0: http://{hostname}:{port}/wsrp-producer/MarkupService?wsdl. (WSRP 1 only) in EPP 5.1: WSRP 1: http://{hostname}:{port}/wsrp-producer/v1/MarkupService?wsdl. WSRP 2: http://{hostname}:{port}/wsrp-producer/v2/MarkupService?wsdl. - 15.6.2.1.1: it doesn't matter that the screenshots are not updated as the example is not specific to the self consumer. Plus, it's still correct in EPP 5.0 (even though, again, the point of the screenshots here are not to show the self consumer, but rather how the portlet behaves if the given registration information is invalid). - Why is section 15.6.3.2 on its own page instead of the same page as the rest of section 15.6.3? - 15.6.4: self producer -> selfv1 and selfv2 consumers - Procedure 15.3: screenshots are barely legible... @Michal: you should add your documentation on how to remove the WSRP admin portlet...	1	1	https://issues.jboss.org/browse/JBEPP-517
10861	JBEPP-517	The above comments have been actioned. @Michal; I followed Chris's comments that, although the screenshots are slightly outdated, they serve the purpose they're intended for. Regarding Section "3.3. The Command Servlet" :I updated the para to read as follows: The CommandServlet is called by the portlet container for requests to particular portlets. It also includes some init code when the portal is launched. This servlet (org.gatein.wci.command.CommandServlet) is automatically added during the deployment of each portlet application and mapped to /gateinservlet. Is that more technically accurate? @Chris: Re the page placements. Were you looking at a chunked html version? This has to do with Publican's default chunking of sections. I can adjust that if needed, but it doesn't seem to be too much of a problem this close to release.	1	1	https://issues.jboss.org/browse/JBEPP-517

#	Requirement	Questions	Answered	Successfully Answered	URL
10862	JBEPP-516	Hi Scott, there are my comments for User Guide verification. Thanks! - First page + title + second page title - 5.0 is mentioned, should be 5.1 2.4 Navigation - Dashboard editor should be bold as well, it is twice in paragraph. - Site + Group editors are available for admin, but Dashboard editor is available for all users. 3.3.1 - There should be screen of Dashboard workspace and not calendar 3.3.2 - Screen 3. is misunderstanding, it could be removed or there could be only screen of Add gadgets link 5.1.2 (?) - "..that allows them to; add a new page..." -> "..that allows them to add new page..." 5.2.1 - There is not latest screen of point 2. Portal templates tab was added. - + there should be some comment (+screen) about portal templates as well 5.2.2 - It seems that ideas are not completed here. - First sentence ends with :, but nothing follows. - An administrator can change many aspects of a portal: - there should be some list of what can be changed I think - And I think this chapter is not about Edit Portal, because it is 5.2.5. 5.2.2 is only about displaying present portals. 5.2.3 - There should be mentioned ending of edit layout action: by Finish button on Edit Inline composer or with Abort (do not save changes) 5.2.4 - There is too big screenshot of Navigation management 5.2.7 - point 2. Screen could be updated with deployed SimpleSkin (to have list with at least 2 items) - Screen for permament skin change is not correct - Properties tab is missing, and Vista + Mac skins do not exists at gatein or epp - permanent, point 3. "Click Save then Finish" -> "Click Save (to finish)" 5.3.3 - point 3., it should be something like: Right-click on node where you want to paste this node - and screen is not correct, you should not paste in system nodes 5.3.4 - same as for 5.3.3 - remove last sentence: "This new pages will appear in the page lists:" 5.3.5 - same as for 5.3.3 5.3.6 - screen is not correct, you can't delete system nodes 5.3.7 - screen is not correct, system nodes should not have sub-nodes 5.4.1.1 - point 6., "or click the icon to see more templates" -> "or click the select box in right pane to see more templates" 5.4.1.2 - screen in point 3. is duplicated, second one doesn't bring anything new - Owner id for user owner type is not described, describe separately for portal, group, user. Remove sentence "The name of the page owner" 5.4.4 - point 1., "Go to Site Editor in..." -> "Mouse over Site Editor in..." 5.5.1 remove sentence: "You can not users to a group but you can remove them from the group". - it's not true, you can add users to groups as well General: - on some places there is Mouse over and somewhere it is Mouse-over, it should be unified 5.5.1.1 - under screen of point 2., "Account Settings" -> Account Setting 5.6.2.1 - should be numbered 1.-6. and not 1.-3.+1.3. 5.6.2.2 - same as for 5.6.2.1 (via page management numbering)	1	1	https://issues.jboss.org/browse/JBEPP-516
10863	JBEPP-516	Scott, there is wrong name of Preface page in EPP User Guide: pref-JBoss_Portlet_Bridge_Reference_Guide-Preface.html Can you edit page name? It doesn't look well in url :) Thanks.	1	1	https://issues.jboss.org/browse/JBEPP-516
10864	JBEPP-516	Thanks Michel. Fixed. How did you even catch that?!	1	1	https://issues.jboss.org/browse/JBEPP-516
10865	JBEPP-515	Prabhat, The section 2.1 about not upgrading has been removed and the Tomcat section had already been removed. Sorry, but I'm not sure what you mean in your last comment. Can you explain a little more? Should the command <mysql> CREATE DATABASE gateinjcr;> in step two of the "Creating a MySQL Database" procedure be repeated as <mysql> CREATE DATABASE gateinidm;> That way the line "Do the same with gateinidm." would be replaced with an example of the step. Is there more to it that should be documented? Should the whole procedure be repeated, and if so, what needs to be different?	1	1	https://issues.jboss.org/browse/JBEPP-515
10866	JBEPP-487	More information required w.r.t Release Notes: Thomas; does this issue require a release note?	0	0	https://issues.jboss.org/browse/JBEPP-487
10867	JBEPP-482	Adam Kovari <akovari@redhat.com> made a comment on [bug 793400|https://bugzilla.redhat.com/show_bug.cgi?id=793400] Hello, in what version is this scheduled to be fixed? In JPP 6 beta it seems to be still reproducable: Step to reproduce : - Edit the pages.xml of the classic portal - Change the page title with #{portal.classic.home} - Start the server with a clean jcr Result: - #{portal.classic.home} is display on page title instead of the "Home" label define in classic.properties ressource bundle.	0	0	https://issues.jboss.org/browse/JBEPP-482
10868	JBEPP-406	Who's in charge ?	1	0	https://issues.jboss.org/browse/JBEPP-406
10869	JBEPP-277	Wesley Hales said: >Error on bridge build > >[INFO] ---------- >[ERROR] FATAL ERROR >[INFO] ---------- >[INFO] basedir en-US does not exist >[INFO] ---------- >[INFO] Trace >java.lang.IllegalStateException: basedir en-US does not exist >Can you take a look? Scott Mumford said: Hi Wesley I updated my local copy of the PB Ref Guide and it built ok (using maven, which I assume you use). (also tried checking out the whole thing again and then again from the folder level above) I can't recreate the error. Can you give me some details of your circumstances; what you've check out of where, version number, build system, build commands, java version, etc. Thanks	1	1	https://issues.jboss.org/browse/JBEPP-277
10870	JBEPP-276	In attachement are notes for reference guide - Part 2 (Chapter 11, Chapter 13, Chapter 14). So till now, We went through almost whole reference guide and remaining chapters are only: Chapter 12 - Building JSF Portlet Chapter 16 - Foundations Chapter 17 - eXoJCR We will provide the rest as soon as possible. But I have question: Is Part 3 finished and ready for our review or you are still working on it? Because today Viliam noticed really big number of typos and typographic inconsistencies, especially in chapter 17.	1	1	https://issues.jboss.org/browse/JBEPP-276
10871	JBEPP-276	Section "11.2.2.4. JSF example using the JBoss Portlet Bridge" 1.) PortletBridge examples are places in "portletbridge/examples" folder instead of "examples" folder. 2.) No example "JSFHelloUser" is bundled with EPP5 (I don't know nothing about this example). Do we really have this example at all?	0	0	https://issues.jboss.org/browse/JBEPP-276
10872	JBEPP-276	I'm recommending triage of sections 17.12 (Transactions), 17.13 (TransactionManagerLookup) and possibly 17.14 (eXo JCR Statistics) due to inadequate content. Does anyone have any comments or objections?	1	1	https://issues.jboss.org/browse/JBEPP-276

#	Requirement	Questions	Answered	Successfully Answered	URL
10873	JBEPP-276	The latest version of the Reference Guide has been committed. Find it at the Hudson URL above. This revision completes the first pass edits and included as much of the QE feedback as ECS is able to complete. Outstanding issues (from various feedback sources) are listed below: _____ *Apply changes from Chris's commit no. 31061 3107!!!! (Not sure what this means - SM) *7.2.5 Debugging resources??? (Not sure what this means - SM) *this discuss about eXo JCR standalone configuration. is it the subject of EPP5 reference guide? (I can comment refernces to this out if nobody deems it necessary -SM) *Villiam mentioned the variance in file paths are formatted. I found the following examples: 02portal.war/WEB-INF/conf/portal/portal/classic/portal.xml WEB-INF/gatein-resources.xml 01eXoResources.war/WEB-INF/gatein-resources.xml $JBOSS_HOME/bin/run.sh eXoGadgets.war/WEB-INF/ web.war/WEB-INF/classes/locale/portlet/portal WEB-INF/classes/locale/portlet/gadget/GadgetPortlet. /web/skin/portal/webui/component/UIFooterPortlet/DefaultStylesheet-rt.css /GateInResources/skin/DefaultSkin/webui/component/UITabSystem/UITabs/background/NormalTabStyle-rt.gif /portletbridge/examples/ We definitely need a consistent style. We just need someone to choose one. I'm in favour of starting them all with /$PORTAL_HOME/blah/blah/ _____ Anything I've missed? Major additions and changes to the document are marked with comments in the XML code for reference.	1	1	https://issues.jboss.org/browse/JBEPP-276
10874	JBEPP-276	Latest Reference Guide edits committed. Solved XML layout problem in included config files. Adjusted column widths in Section 17.5 to correct overflow issues. Various minor edits. Outstanding issues: Section 11.2.2.3 - Sentence "My name is GateIn Portal. What's yours?" In example is used: "My name is JBoss Portal". But this is not issue in documentation, because "My name is JBoss Portal" should not be used (SM - Unable to get the example portlet running so cannot update screenshot if portlet has been updated to correct wording. Section 7.5.2 Debugging resources - magic locale??? (I have never seen this before and I don't know if it's even possible, it is great and there should be better description how to do it) (SM - Further information unavailable at the moment. Section commented out and marked for next release) Section 14.3.1. Configuration of org.exoplatform.services.organization.idm.PicketLinkIDMOrganizationServiceImpl - This service has one more option "useJTA" which is not mentioned in reference guide. (SM - Added to useJTA to relevant list. No information regarding option included here, so content inferred from context referencing cacheconfig options)	1	1	https://issues.jboss.org/browse/JBEPP-276
10875	JBEPP-163	Can we resolve this JIRA? The Site Publisher JIRA can be reassigned to JBEPP-372 (5.0.1 tracker) if required.	1	1	https://issues.jboss.org/browse/JBEPP-163
10876	JBEPP-158	The note mentioned in the first comment has been removed. WRT to Mike's comment, is there an argument for overhauling the section in question? Or can it be removed entirely? For the moment the section has been commented out pending decision and/or new content. A CP05 version of the Reference Guide has been committed to: http://svn.jboss.org/repos/portal/docs/enterprise/trunk/4.3.cp05_Reference_Guide/	1	1	https://issues.jboss.org/browse/JBEPP-158
10877	JBEPP-97	Could we change it in the next non-CP release? I would think most customers would prefer security first and as you said, they'll likely be modifying login.jsp anyhow so can take away extra security then if desired.	1	0	https://issues.jboss.org/browse/JBEPP-97
10878	JBEPP-82	What do you mean by API change? I'd argue that the changes to CustomizationManager are fair as CustomizationManager was never part of the supported API...	0	0	https://issues.jboss.org/browse/JBEPP-82
10879	JBEPP-26	what about licences?	0	0	https://issues.jboss.org/browse/JBEPP-26
10880	JBEPP-23	Which version of the EAP installation guide should this be based on?	1	1	https://issues.jboss.org/browse/JBEPP-23
10881	JBEPP-23	> 3. Add database part. What content relating to database setup is needed?	1	1	https://issues.jboss.org/browse/JBEPP-23
10882	JBEPP-23	I made some changes. Please synch up your code with repository and see how everything looks. On another note, can you add DB descriptor parts for other databases? I think they were present in old installation guide that I pointed on my previous comment.	1	1	https://issues.jboss.org/browse/JBEPP-23
10883	JBEPP-23	For section 2.1.2, if that's what it says for EAP then it's true for EPP as well. I did not understand what you mean by "And further to that, are other databases?". What I meant is similar instruction for Oracle, Mssql, Postgresql to what you have for MySQL.	1	1	https://issues.jboss.org/browse/JBEPP-23
10884	JBEPP-23	Prabhat, Here are a few other things I've noticed in this version: 1.1 Feedback The JIRA information should be for JBEPP project and component "Documentation". Version is unset (currently says "<version>") **Should be 4.3.0.GA** Chapter 2 2.1.1 FYI: The supported configurations page linked to has up to date information for JBEAP, but no mention of JBEPP. 2.2 FYI: The components version page linked to has information about the level of the JBEAP that is included in JBEPP, but does not include the JBEPP component. Chapter 3 3.2 Installing. This section currently has inconsistent references to file names - some references are jboss-eap, one is jboss-epp Chapter 4 4.1 Should the directory name change from EnterprisePlatform to something like EnterprisePortalPlatform?	0	0	https://issues.jboss.org/browse/JBEPP-23
10885	JBEPP-23	Prabhat, I can not find the documentation for the descriptors for Oracle, Postgres, MSSQL etc. - it doesn't seem to be in the guide you previously linked.. Can you please direct me to where it is so I can import it into this guide?	1	1	https://issues.jboss.org/browse/JBEPP-23
10886	JBEPP-20	Do you have machine to do this? Can you use one from the lab? Otherwise, I will do it.	1	1	https://issues.jboss.org/browse/JBEPP-20
10887	JBEPP-17	I have found the difference between IE7 and Firefox 2 on wherethe page Test and TestWithAjax are displayed. On IE, it's displayed towards the right of the window and on FF, it's on left. Does this need to be fixed?	0	0	https://issues.jboss.org/browse/JBEPP-17

#	Requirement	Questions	Answered	Successfully Answered	URL
10888	JBEPP-4	Viliam, This has priority and if you get overwhelmed, I will take care of some of other tasks that I have assigned to you. Please start on it once you are done with selenium builds which I think you are very close. Currently, for JBoss Portal 2.7 release process can be summarized as: 1. svn checkout portal code 2. svn checkout portal docs 3. build portal and docs 4. unzip JBoss AS 5. copy portal stuff to AS for bundled distro 5. copy portal docs 6. zip it and actual commands are: $ mkdir release $ cd release release $ svn export https://svn.jboss.org/repos/portal/tags/JBoss_Portal_2_7_0 project release $ svn export https://svn.jboss.org/repos/portal/docs/tags/JBoss_Portal_2_7_0 docs release $ svn export https://svn.jboss.org/repos/portal/examples/tags/JBoss_Portal_2_7_0_BETA1 examples release $ unzip jboss-4.2.3.GA.zip release $ cd docs release/docs $ mvn package release/docs $ cd ../project/build release/project/build $ sh build.sh release/project/build $ export JBOSS_HOME=$PWD/../../jboss-4.2.3.GA release/project/build $ ant -f distrib.xml Please try this out and see what you get in projects/build/output. These are the output of above steps. Now this needs to be modified for EPP. So, please leave distrib.xml as it is and create a new one (epp-release.xml) taking it as a baseline. What new script should do is: 1. Read different version name (Portal, EAP, location of EAP.zip) from a properties file so that all one needs to do is update this file for a new release. 2. Build JBoss Portal source 3. Copy jboss-portal.sar to default configuration 4. Copy jboss-portal-ha.sar to production and all configuration 5. Build portal docs and copy them to doc folders in EAP 6. Produce one zip out of it. Ultimately, the output if EAP + Portal = EPP. I understand you have many questions so please feel free to ask.	0	0	https://issues.jboss.org/browse/JBEPP-4
10889	JBEPP-4	in distro.xml there are two patches applied : run.conf patch and log4j patch... should i apply them too?	0	0	https://issues.jboss.org/browse/JBEPP-4
10890	JBEPP-2	Dependencies such as <componentref name="sun-jaf" version="1.1"/> <componentref name="sun-javamail" version="1.3.1"/> <componentref name="sun-jsf" version="1.2_08-brew"/> have moved to glassfish such as <componentref name="glassfish/javamail" version="1.4.0-brew"/> <componentref name="glassfish/jstl" version="1.2.0-brew"/> This of course requires changing changing some build.xml. Should I go ahead and change them if there is no conflict with other jar version?	0	0	https://issues.jboss.org/browse/JBEPP-2
10891	SOA-3550	What depends on ant: - EDS - HornetQ - CXF - seam - schema tool - quick starts	1	0	https://issues.jboss.org/browse/SOA-3550
10892	SOA-3464	The following are some initial suggestions at the beginning of the documnt, but section 3.1 Data Services Examples will really change, how would you prefer I submit the changes? I can attach a doc with the rest of the suggested changes or I can cut and paste them here. Your choice. ---------- 5.2 Getting Started Guide Review General [note] suggestion: remove the version (i.e., 5.1.0) when JBoss Enterprise Data Services Platform is referenced, provides no additional context other than what the version of the document is referring to in the Title. If not, change the 5.10 to 5.2.0. Chapter 1 1.1 - reference to http://www.jboss.com/products/platforms/soa/components/ , is this right link? Cause this page doesn't contain Data Services or anything about its components. Or maybe that page needs to be changed to point to: http://www.jboss.com/products/platforms/dataservices/components/. Chapter 2 2.1 - JBoss SOA Platform version should be changed to 5.2.0, as well as, Standalone. - Link to tested configurations should be to: http://www.jboss.com/products/platforms/dataservices/supportedconfigs/ 2.3 - (my observation) The WARNING about CXF is no longer an issue in ER5, cause you can only deploy EDS to all, production or default. Since CXF and EDS has to be deployed to the same configuration, and CXF can only be deployed to those 3 because there is no option to install elsewhere, this shouldn't occur, since EDS is now changed to not allow other profiles. - in step 4 - "Run installation script", can the sentence "Apache CXF must ..." or the word "must" be emphasized to stress its need if using webservices? - in step 5 - "Configure", can we rephrase to: However, user accounts will need to be created to enable data access and other features. Refer.... Chapter 3 - sentence: VDB's can be queried ..... add ODBC: VDB's .... JDBC, ODBC and web-services.	0	0	https://issues.jboss.org/browse/SOA-3464
10893	SOA-3264	The client exception I get (from squirrel) is: Remote org.teiid.client.security.LogonException: VDB "bqt" version "latest" does not exist. This is using 7.4. The project indicates there is nothing to be fixed and we don't want to list what the available vdb's are for security purposes.	1	0	https://issues.jboss.org/browse/SOA-3264
10894	SOA-3154	Kurt, can you please respond to all of the open items here: https://docspace.corp.redhat.com/docs/DOC-70182 ? Thanks.	0	0	https://issues.jboss.org/browse/SOA-3154
10895	SOA-3146	We can't and are not going to upgrade to Apache CXF 2.4.x in JBossWS 3.4.x or lesser. The move to CXF 2.4.x requires integration changes that are in JBossWS 4.x only. As this has been isolated into an Apache CXF issue, we need to backport the fix for that issue into our maintenance fork of the CXF version that's consumed by EAP / SOA-P as we did for other issues.	0	0	https://issues.jboss.org/browse/SOA-3146
10896	SOA-3146	The fix has been ported in EAP 5.1.1 branch . Gary, do you know if the SOA 5.2 will consume EAP 5.1.1 ? Thanks.	1	1	https://issues.jboss.org/browse/SOA-3146

#	Requirement	Questions	Answered	Successfully Answered	URL
10897	SOA-3117	examples without references on appropriate quickstarts (although they exist). This leads to low information value from reading the documentation h2.Problematic parts of the guide Following things I consider particulary bad: Inconsistent/wrong description of Administration console [where] Chapter 1, (especially 1.1.2 Metrics) [now] In documentation there is a list of metrics that can be "viewed" [notes] - The list contains actions: Terminate Session, Terminate Requests, Terminate Transaction. These actions are not present on Metrics tab and are now on "Control" tab in actual Console. Next problem with these actions is that they are not "viewable" it seems that they were put to this chapter (1.1.2) by mistake or the heading "Control" is just missing (Control). - In my opinion this documentation is somehow very obsolete because not all actions available in console are listed here and Control tab is missing description in documentation. Typography [where] 2.1 Getting started (Note box) [now] use theadminHelp() [suggest] use the adminHelp() Wrong method name [where] 2.3 The Log and Recorded Script Files [now] connectionAsAdmin() [suggest] connectAsAdmin() Bad english [where] 3 Deploying Virtual Databases [now] A VDB is the primary means to define a virtual database [note] "VDB are" or "VDB is primary mean". Although I don't know what is meant by VDB. I thought VDB is "virutal database" in which case the sentence wouldn't make sense. Typo [where] 3.2. Deploying VDB Dependencies, second paragraph [now] TO give Typo [where] 3.2. Deploying VDB Dependencies, last paragraph [now] Once you have deplored ... Invalid path [where] Example 3.4 WS-Security [now] ${jboss.server.home.dir}/server/default/conf/xxx-jbossws-cxf.xml [suggest] ${jboss.server.home.dir}/conf/xxx-jbossws-cxf.xml Type in example [where] Example 3.4 WS-Security [now] missing ending tag for <map> Incomplete or wrong information [where] 3.2.3.3 Logging [now] Logging, when enabled, will be performed at an INFO level to the org.apache.cxf.interceptor context [info] As I understood this sentence, after turing this feature ON, in my server log I should have logged requests and responses under mentioned logger. I made sure to have correctly configured log4j in my SOA platform and although I see logged request and response (encrypted) I don't see it coming out from org.apache.cxf.interceptor logger. [suggest] Add more information to this paragraph about where user can expect the output of such logger and add INFO box that he should have logging set up accordingly (as default settings of SOA P won't allow him to see the INFO outputs) Typo causing non-deployable DS [where] Example 3.6. Example Disabling Host-name Verification [now] ...rs disableCNcheck="true" /> [suggest] ...rs disableCNCheck="true"/> [note] When you supply examples, please validate at least against available xml schemas... Wrong chapter numbering [where] 3.2.5 CXF Configuration [now] 3.2.5 [suggest] 3.2.4.1 or consolidate the paragraph with 3.4.2 [note] It seems that CXF configuration is logically connected to 3.2.4 because notes about salesforce is there. Therefore it certainly shouldn't be numbered as 3.2.5 Confusing and duplicit information [where] 3.2.4, 3.2.5 [now] Information about how to configure salesforce connector with CXF configuration [note] The information is duplicit and is essentially copy-paste of 3.2.3.1. The documentation is confusing mainly becase of ambiguous titles of headings [suggest] Delete 3.2.5 and just refer to the fact that teiid-connector-salesforce.rar uses same configuration as CXF web services for Bus configuration. Also remove second Note box for "bus info". Cubersome language [where] 3.3 [now] If a specific version is set, then you can only connect to that VDB. If no version is set, the deployed VDBs are searched until the appropriate version is found. This feature helps support more fluid migration scenarios. [suggest] Specifying version is optional. If the version is not specified then deployer will decide which version of the VDB should be used. This helps... Not following conventions [where] 3.3 VDB Versioning [now] ... the deployment file (such as vdbname.version.vdb or marketdata.2.vdb... [suggest] use italic for "vdbname.version". Also restructure to something like: "... the deployment file such as	0	0	https://issues.jboss.org/browse/SOA-3117
10898	SOA-3115	Chapter 5) - all list items --> highlight terms being exaplained where appropriate - whole document --> english quotation marks - Install the schema in your IDE to avail yourself of its auto-completion functionality. --> rewrite - The Javabean Cartridge sets the following conditions to Java beans: --> following item list -- all end with a period or all without - ((1), (2), (3) - all accepts (4)) --> too many parenthesis -- (1, 2, 3 - all accepts 4) -- what does it mean all accepts? -- the hyphen is written as minus (rewrite) - all description of the above configuration contain too many parenthesis -- quite unreadable with the doubles -- ((- all terms in the above configuration should be emphasized -- for example "execOnElement" - (see the beanIdRef attribute values and how the reference the beanId values defined on (2) and (3)) --> see ... and how the reference the beanId? -- rewrite - Configuration (2) creates the com.acme.Header bean instance. --> missing bullet - ontoor the Header.customerNumber --> onto? - (3) Configuration (3) creates the ListOrderItem bean instance for holding the OrderItem instances. --> too much "3" - (4) Configuration (4) --> too much "4" - (one example being that the DateDecoder [decoder="Date"]). --> that? - The decodeParam iformat s based on the ISO 8601 standard for date formatting. --> probrably "format is" - Here is an Integer example: --> following note is much the same as the previous one, rewrite not to be so redundant -- readability - (for instance, the numeric value 876592.00 might be represented as "876_592!00"jb:bean element for some reason). --> rewrite, insert some space probably - In order to decode this value as (for isntalce) --> for instance? - which is a simple expression that you can applied to the Sting value before decoding it. --> you can applied? - that operates on the value String --> that operates on the String value - (The 'instanceMethod part is optional. --> ""? - There are three ways in which you can declare which definition language you want to use: --> rewrite these three ways, details follow: - For example. --> For example: <example style>....</example style> - Note:If you have a definition with your default language that includes a ':' then you must prefix that definition with 'default:' else you will run into an Exception. --> put into note style - MVEL as factory definition language --> is it a heading or what? - If MVEL where the default language then we couldn’t do anything to improve --> where? -- probably was or were - Binding Key Value Pairs into Maps --> heading? - On the config: --> label the examples better - as part of a model driven transform --> transformation - Merging Multiple Data Entities Into a Single Binding --> heading? - The ValueBinder class is the visitor that does the value binding. --> heading? - The value binding XML configuration is part of the JavaBean schema from Smooks 1.3 on: http://www.milyn.org/xsd/smooks/javabean-1.4.xsd. The element for the value binding is value. --> following item list -- end all items with a period or without - default : The default value for if the selected data is null or an empty string. --> rewrite - "Example" following above list item --> emphasize somehow - "Message" --> convert to example description - "The Configuration" --> convert to example description - The value binder can be programmatic configured using the org.milyn.javabean.Value Object. --> programmatically, "Object" = "object" - "The Message:" --> convert to example description - "The Java Model (not including getters/setters):" --> convert to example description - "The Configuration Code:" --> convert to example description... do the same rules as for headings apply here? -- "Code"="code"? - "The Execution Code:" --> convert to example description... do the same rules as for headings apply here? -- "Code"="code"? - The optional "-p" properties file parameter allows specification of additional config parameters: --> next list item is nested, so the bullets have to be indented more - The Html Reporting tool --> HTML? - your Source data --> "Source"="source"? - Note the "$TODO$" tokens. --> emphasize $TODO$ in other way - This will mean that some of the JavaSource bean instances may be visible in the JavaResult. --> meant?	0	0	https://issues.jboss.org/browse/SOA-3115

#	Requirement	Questions	Answered	Successfully Answered	URL
10899	SOA-3114	Hi guys, Could we please add to the Smooks_User_Guide (referenced above) something similar to what we find at [1]? (I've just notice we reference that link in the SOA-P Programmer's Guide.) The below referenced wiki article contains a couple of important points that I don't think are in the User Guide above: 1- From an ESB programmer's perspective, where do I get the thing I want to transform, and where do I set it (on the message)? 2- What's the relationship between my config file for Smooks and the ESB deployment descriptor (jboss-esb.xml)? I'd suggest putting this information near the front of the Smooks_User_Guide for SOA-P. If the user starts with this guide, they'll be wondering those things for quite a while. Another tip we might put in the user guide-- if they are seeking advice on how to use the various Smooks components available, give them a hint to check the schemas (XSDs) referenced in the quickstarts. There are comments in-line, these are helpful but I don't think we ever suggest that the user look at them. Thanks! Rick [1] http://community.jboss.org/wiki/MessageTransformation	1	1	https://issues.jboss.org/browse/SOA-3114
10900	SOA-3114	OK... sorry for not commenting on this earlier. @Rick... wrt #1 above... what do you mean by "get" and "set"? Do you mean "get-payload-location" and "set-payload-location". That's standard across all ESB actions. wrt #2 above... what do you mean by "relationship"?	1	1	https://issues.jboss.org/browse/SOA-3114
10901	SOA-3114	That all looks great to me Rick... how about the docs guys just add that then?	0	0	https://issues.jboss.org/browse/SOA-3114
10902	SOA-3108	3.1 first jboss-esb.xml sample, according to xsd, there cannot be invmScope attribute in the root element. I did not find any example configured like this. Move the attribute to the service element. 3 several occurances "end-point" -> "endpoint", the rest of the document uses "endpoint" 3.2 second sentence "the reason being that this will encourage loose coupling." I'm not a native speaker but this sentence seems incomplete to me. 3.2 "The message structure can be also >>be<< depicted in UML" 3.2 "Version 5.0 of the JBoss Enterprise Service Bus will be the first to support user-enhanced contexts." - are we going to have version 5.0? Do we add new features to ESB? IMHO this won't happen since there is SwitchYard now. 3.2 Second occurence of the same warning box "In JBossESB, Attachments and Properties are not treated differently from the Body. The general concepts they embody are currently being re-evaluated and may change significantly in future releases. As such, we recommend developers do not use Attachments." It should be removed since it do not have any relations to the surrounding text. 3.6 "unto" -> "into" 3.7 This section seems uncomplete. There is nothing regarding default value or behavior of fault-to header. I do not understand the reference to LogicalEPR in the information box in this section. 3.10 "Every implementation of the org.jboss.soa.esb.message.Message interface can >>be<< obtained from the", "be" is missing in the original doc. 3.10 information box at the end of the section, "JBOSS_XML . To" - additional space is present before the dot. 3.11 "Message Formats, There are two serialized message formats, JBOSS_XML and JBOSS_SERIALIZED. Read this section to learn about them." - isn't it obvious what is this chapter about? Is the last sentence in the first paragraph necessary? 3.11 last paragraph, "When it is packing objects in XML, The", capital T should be in lowercase	1	1	https://issues.jboss.org/browse/SOA-3108
10903	SOA-3108	This comment is completely on section 4.1.3 Notifiers IMHO this section need a major rewrite. 1) First paragraph reads "Use notifiers to pass on success and error messages to ESB-unaware end-points. ... This may mean that one cannot have ESB-aware and -unaware end-points listening to the same channel. If this is the case, consider using either couriers or the ServiceInvoker...". It is not clear what is exactly meant by "If this is the case". What is the case - having ESB-aware and -unaware endpoint on the same channel or the opposite? 2) Second paragraph starts with "Not every ESB-aware transport is supported by notifiers". What is the reason for mentioning this? In the first paragraph we learned that none ESB-aware endpoint is supported by a notifier. 3) Third paragraph (starting with "As outlined above, it is a listener's responsibilities...") and the following list is very confusing, inaccurate and full of errors. Fisrt, description of listeners configuration should go to section 4.1.1 Listeners. Second, several concepts are mixed together - providers configuration, services configuration and listeners configuration. Listener does not supply information for registry, according to XSD of jboss-esb.xml, there is no listenerClass tag or attribute, endpoint reference is configured using providers not in the listener, action pipeline is configured in the service at the same level as listener (not within the listener). There are no tags named service-category, service-name, service-description, EPR-description, instead there are attributes of the service tag called category, name and description. Terms 'tag' and 'attribute' are often confused in this part. 4) "This example configuration instantiates a listener object (the jms-listener attribute), which will wait for those incoming ESB messages that are serialized within a javax.jms.ObjectMessage." It configures two listener objects. 'jms-listener' is a tag, not an attribute. Messages can have various formats (various classes can be used), for instance javax.jms.TextMessage, not only ObjectMessage. I suggest specifying interface insted. It is necessary to mention why there are two listeners and clarify the difference between gateway listener (it would be sufficient that this is the ESB-unaware listener and that its role is to encapsulate the JMS message in the ESB message used throughout ESB). 5) In the example with "public class MyJMSListenerAction", the keyword extends and the correct predecessor (probably ActionProcessingPipeline). Starting with this example, there should be one more section called "Action Classes". 6) "Use action classes to customised the framework for specific needs. specific needs." Remove duplicate 'specific needs'. Fix 'customized' -> 'customize'. What framework is customized? There is no framework mentioned previously. I would prefer something like "Use action classes to provide additional functionality to ESB services that are not implemented in the out-of-the-box actions." 7) "ActionProcessingPipeline class expects any other action class to provide, at the minimum, this information:". I believe that this sentence is just a complicated and cryptic way of sayin "Each descendant of ActionProcessingPipeline must, at minimum, provide:".	1	1	https://issues.jboss.org/browse/SOA-3108
10904	SOA-3108	7.1 "advantage of thse features", "thse" -> "these" 7.1.1 "The service will then receive the JMSEPR use it to send the message." I don't understand this sentence. What should it mean? 7.1.1 "service (an application service.)" -> "service (an application service)." (note the dot position change) 7.1.6 terms should be in bold: ServiceInvoker, EPR, DeadLetterService, DLQ, MessageStore, DeadLetterService, RDLVR, jbossesb.esb 7.2.5 and subsections, dtto, I won't list all the terms here, it is obvious what should go in bold 7.2.5.1 code sample is missing syntax highlight 7.2.5.2 numbered list, "3. One of:", following list should be sub-list of the numbered one (now it is at the same level as the numbered list) 7.2.5.2 code sample is missing syntax highlight 7.2.5.3 "start by implementing a javax.jms.MessagListener:", 'start' -> 'Start' 7.2.5.3 all code samples are missing syntax highlight 7.2.5.3 "You could also used the ServiceInvoker", 'used' -> 'use' Given the shape of the docs, it would be more effective if I could fix the document directly...	1	1	https://issues.jboss.org/browse/SOA-3108

#	Requirement	Questions	Answered	Successfully Answered	URL
10905	SOA-3108	10.3 "JBossESB 4.8 has the following MimeDecoder implementations.", currently we have JBossESB 4.10 in SOA-P 5.2. Ommiting the version number seems to be the most elegant solution. 11.1 "In order to learn more about these components, one should refer to the Services Guide. It explains, in detail, how to configure them.", this text references to SOAPProcessor and SOAPClient. These are described later in this guide (Programmers Guide). There is no reference to them in Services Guide. It should be mentioned that there are several versions of these components at the moment, specifically I mean WISE. 12. Classes, methods and types should be in bold. See all subsections for this. 12.1.1 First table in this section, "byte[]" should be in bold 12.1.9 "The SmooksTransformer action will be deprecated" -> "The >>former<< SmooksTransformer action will be deprecated", I suggest adding the word 'former' here to explicitely note that this is the action mentioned in the previous section. For the first time reader, the terms may coincide. 12.1.9 "The easiest way to avoid this issue (in JBoss ESB 4.5.GA+) is", we do not have any GA releases of projects. This should refer to a plain version or SOA-P release. 12.1.9 "you should use the http://www.milyn.org/xsd/smooks/javabean-1.2.xsd configuration", this might be a clickable link 12.1.9 "See the Smooks v1.1+ documentation", we have Smooks 1.4 in SOA-P 5.2, are we sure that there are no changes? Shouldn't we rather refer to 1.4 docs? 12.1.9.1 "For more information on for Smooks Execution Report refer to http://www.smooks.org/mediawikiindex.php?title=V1.2:Smooks_v1.2_User_Guide#Checking_the_Smooks_Execution_Process.", what about updating the link to point to v1.4? 12.1.9.3 "12.1.9.3. XML, EDI, CSV etc Input Payloads" -> "12.1.9.3. XML, EDI, CSV and Other Input Payloads" 12.1.10 First table, "Message" should be in bold. 12.2.1 First table, "Message" should be in bold. 12.5.3 Table 12.8, missing http-client-properties description. A sample usage would be great as well: <action name="HttpRouter" class="org.jboss.soa.esb.actions.routing.http.HttpRouter"> <property name="endpointUrl" value="http://@HOSTNAME@:8080/Plain_ws_ejb/HelloWorldWS"/> <property name="http-client-properties" > <http-client-property name="connection.http.connection.stalecheck" value="false" /> <http-client-property name="max-total-connections" value="200" /> <http-client-property name="max-connections-per-host" value="200" /> </property> ... 12.7.1 "12.7.1. SOAPProcessor JBoss Webservices SOAP Processor", what is the line after section caption good for? 12.7.1.3 Bad formatting of the Caution Box, the text "Caution" is over the previous paragraph 12.7.1.4 "a .esb deployment" -> "an .esb deployment" 12.7.1.4 "This however means that this action can only be used when your .esb deployment is installed on the JBoss Application Server i.e. It is not supported on the JBossESB Server. ", "It is" should have lowercase 'i', the sentence is nonsense - you cannot deploy an .esb archive on JBoss AS without JBossESB. Doesn't it refer to standalone JBossESB usage without the AS? 12.7.2 "12.7.2. SOAPClient" -> "12.7.2. WISE SOAPClient" 12.7.2 "If there is a SOA Fault and HTTP 500", 'SOA' -> 'SOAP' 12.7.2.2 "See the "webservice_bpel" quickstart.", there is no webservice_bpel quickstart. Why there is a reference to this particular QS? 12.7.3 "12.7.3. SOAPClient" -> "12.7.3. soapUI SOAPClient" 12.7.3 Delete the first sentence in this section ("SOAP Client action processor.")	1	1	https://issues.jboss.org/browse/SOA-3108
10906	SOA-3107	Configuration [now] SOA_ROOT/server/CONFIG/deployers/esb.deployers/jbossesb-properties.xml [suggestion] SOA_ROOT/server/CONFIG/deployers/esb.deployer/jbossesb-properties.xml Incomplete information [where] Chapter 1.3.1 Configuring a Java Message Service Provider [now] Docs suggest user should set destination-type, destination-name in jboss-esb.xml [note] After brief tryouts I wasn't able to find out how these two parameters should be set. I am used to the fact that when I want to set destination for JMS provider i will just use jms-message-filter element and dest-type/dest-name attributes. However usage of destination-type and destination-name is not even in any quickstart that would help and neither is there any reference of these two parameters in XSD schema (they are probably set somewhere via string property of some xml element). As it is there is no way to easily follow what docs want to say for somebody who didn't use these params before. [suggest] Add these configuration into some hello world Quickstart with configuration pointing to implicit JBoss Messaging and reference this quickstart in docs. Also explain why there is need for these params when user can achieve this via dest-type/dest-name under jms-bus element. Maybe not complete information [where] 1.3.17.2. Message Tracing [now] By default, every interaction between components and messages is logged. [note] I am using 5.0.2.ER3 build and by default (with log4j config set to INFO) I dont see any of these audit logs when running quickstarts. [suggest] please supply more information what should be done to se these logging messages Incomplete information [where] Chapter 1.3.17.2 Message tracing [now] The TraceFilter will only log information if you set the org.jboss.soa.esb.messagetrace [note] Where should this property be set? Broken links [where] 1.3.17.5. Installing and...; Procedure 1.2. Deploying the OpenSSO to Tomcat, step 4 [now] broken links: http://java.net/projects/opensso/public/use/docs/fampdf/index.html https://opensso.dev.java.net/public/use/index.html Because of this i couldn't review rest of the chapter 1 Cubersome language [where] Chapter 2 The Registry [now] when services run, they usually add an end-point reference to the registry. It is through this that they are contacted. If services have been programmed well, they automatically remove these end-point references when they terminate. However, if they are poorly programmed or crash, theese entries will remain there. These "orphan" entries will prevent subsequent deployments from executing correctly, so remove them manually. [suggest] After service starts up, it usually adds an end-point reference (EPR) to the registry. EPR can be further used to contact the service. \par Well programmed service removes its corresponding EPR upon termination. However poorly programmed service can leave EPR in the case of service crash, this kind of EPR is reffered to as "orphan entry". Because orphan entry will prevent subsequent redeployments of the service it must be removed manually. Typo [where] Chapter 4 Default [now] word "prefixed" is used instaned of "suffixed" Incomplete information [where] 6.1 The Monitoring and ... [now] For instance, it displays message counts by service, action, and node. It also displays such useful information as: processing times,.... [note] Where are those information in the console located? Simplistic information about hotdeploy [where] Chapter 7. Hot Deployment [now] Very limited information about Hot deployment is explained in suprisingly long paragraphs. [suggest] Add information about hdscanner-jboss-beans.xml and limitations of Hot Deploy Cubersome language [where] 11.2. Message Store Interface [now] can ignore this field if need be. [suggest] can ignore this field if needed. Missing security settings [where] Table 12.1. jbossesb-properties.xml Security Settings [now] table is missing <property name="org.jboss.soa.esb.services.security.publicKeyTransformation" ... [suggest] add it Typo [where] 12.1.1. Configuring Security on Services [now] ... Her is a code ... [suggest] ... Here is a code... h2. Minor comments Confusing [where] Preface 1.1 Key combinations can be distinguished from keycaps by the hyphen	0	0	https://issues.jboss.org/browse/SOA-3107

#	Requirement	Questions	Answered	Successfully Answered	URL
10907	SOA-3093	Hi David a few notes to get thinks started 1.2.2. Logging in. True for riftsaw installation on pure JBoss AS. Not true for SOA-P though . default login:password is admin:admin which is defined in soa-user.properties, soa-roles.properties. Login policy is definied in login-config.xml and can be changed in deploy/bpel-console/gwt-console-server.war/WEB-INF/jboss-web.xml in the <security-domain/> element. 1.2.5. What about process archives which do not have a version definition? e.g. process archive name Quickstart_bpel_simple_invoke.jar. These are handeled as version 0. So Quickstart_bpel_simple_invoke-1.jar may be deployed as the next process version. 2.2. Example 2.1. Sample Configuration: IMHO this should include a sample build file from SOA-P not Riftsaw. 2.3. consider to include: BPEL processes can be tested directly using JBDS. No ANT scripts are necessary. right click the process WSDL -> Web Services -> Test with Web Service Explorer -> ... 6.2. Mentioning the farm directory is good but we don't want to make the impression that it's suited for enterprise use. The process archives should be copied separately to each server.	1	1	https://issues.jboss.org/browse/SOA-3093
10908	SOA-3093	I have made most of these changes thanks, Marek. :-) The only thing I haven't altered is the sample build file - can an SME please provide a relevant SOA one? Cheers.	1	1	https://issues.jboss.org/browse/SOA-3093
10909	SOA-3093	Could you please add to guide steps to configure riftsaw clustering? In default It's disabled and old steps are for deploy riftsaw with clustering. It would be nice to have some template files in docs/examples as jboss-beans.xml and riftsaw-cache-manager-jboss-beans.xml (these files are ready in qa/clusterinstall2/resources/riftsaw/META-INF where is 1 configurable part @database@ in jboss-beans.xml). Step 1 would be set @database@ in jboss-beans.xml example file to some of these mysql, postgre, db2, sqlserver, oracle or sybase Step 2 would be copy jboss-beans.xml to [server-config]/deploy/riftsaw.sar/META-INF/ Step 3 would be replace [server-config]/deploy/cluster/jboss-cache-manager.sar/META-INF/jboss-cache-manager-jboss-beans.xml with riftsaw-cache-manager-jboss-beans.xml	1	1	https://issues.jboss.org/browse/SOA-3093
10910	SOA-3093	Hi David, can you have a look at the following please? 1.2 Quickstart Examples - Hello World - wrong path - should be SOA_ROOT/samples/quickstarts/bpel_hello_world - the user can send a message to the process as: ant runtest 1.3 ESB Examples. - wrong path - should be SOA_ROOT/samples/quickstarts/bpel_esb_helloworld 2.3 BPEL properties bpel.uddi.* property is located in the table header. I suppose this is a typo. 3.4 Changing End-Point Configuration Properties Everything in this section refering to Apache ODE should be removed since it's not configurable in Riftsaw. Riftsaw only supports process specific config. Also for completion the only property supported at this time is mex.timeout (as stated in the guide). 5.2 The note about the ClerkManager includes also Table 5.1. Is this a mistake or intention? 7.1 Installation Steps provided above by ibek should be in this section. Attempting to install another Riftsaw deployment can break the integration.	0	0	https://issues.jboss.org/browse/SOA-3093
10911	SOA-3086	Again, I'm not on-board with making this a blocker unless we are specific about the requirements. We're already at the end of the development window, but we don't fully understand the customer needs. What can be done in the near-term is to just reintroduce changes to have pluggable role validation, but that still leaves open issues of: -how many calls are being made to the custom implementation (1 per user query?) currently we make at least 1 call for each query/subquery/function call. -do they need the ability to modify the incoming query (this is the most problematic) -do they need to specific mechanism to conveys denials (they had indicated a potential conversational approach to refine the submitted query) or will the exception text of a SQLException be sufficient. There are possibly more.	0	0	https://issues.jboss.org/browse/SOA-3086
10912	SOA-3049	one off my head is the data instance cleanup configuration: http://ode.apache.org/instance-data-cleanup.html	1	0	https://issues.jboss.org/browse/SOA-3049
10913	SOA-3038	David - it's not there yet - can you check on this? Thx!	0	0	https://issues.jboss.org/browse/SOA-3038
10914	SOA-3012	Per DOC-53797: "The goal of this process is to provide a quick and clear response to the customer - we should respond to the request withing two weeks. The response has to be unambgious - if we have rejected the enhancement a reason needs to be provided; if it is accepted then a release version has to be commited to." Ken, can we please make a decision on this ASAP?	0	0	https://issues.jboss.org/browse/SOA-3012
10915	SOA-3009	Review comments: Comments on new 4.3 CP05 Getting Started Guide ---------- Page 1, Overview Text: Chapter 3 contains instructions on how to launch, shutdown and troubleshooting the JBoss Enterprise SOA Platform product. Comment: Grammar - should be "troubleshoot" ---------- Page 1, Section 1.1 Text: JBoss Enterprise SOA Platform is also the supported deployment platform for JBoss Enterprise Data Services. Refer to http://www.jboss.com/products/platforms/dataservices/ for more information about this product. Comment: Remove this paragraph - EDS is SOA 5.1 only. ---------- Page 4, Section 2.3 Text: Important This document does not cover the installation of the JBoss Enterprise Data Services Platform. Refer to the EDS Getting Started Guide for information for installation instructions other details for that product. Comment: Remove this paragraph - EDS is SOA 5.1 only. ---------- Page 10, Section 2.2.4 Text: 2.4.4. Managing jUDDI Console Users Read this section to learn about managing user accounts for publishers in the jUDDI Console. Only the jUDDI publisher root, or another publisher with administration rights, can grant administration rights to another publisher. Do this in the jUDDI Console by enabling the isAdmin checkbox in the Publisher Details screen. To enable root log-on, create the user root using the procedure described in Section 2.4.1, “Creating Server Console User Accounts”. The role of user must be assigned to this user. Comment: Remove this section. The jUDDI console is added in SOA-P 5.x ---------- Page 13, Section 3.1 Text: Important The startup time for the server using the production server profile is approximately two minutes. This is slightly longer than previous releases due to the high number of new features in this version. Comment: Remove this paragraph. ---------- Page 17, Section 4.1.1 Text: 4.1.1. The "Helloworld" Quick-Start Comment: All references to "Helloworld" should be changed to "helloworld" ---------- Page 19, Section 4.1.1 Comment: We should add a more descriptive definition of "ESB-aware" to this section of the guide: In describing application clients and services, the term "ESB-aware" means that they can understand the message format and transport protocols used by the JBossESB in the SOA Platform. Gateways are used to convert messages in ESB-unaware formats into a form that can be processed and routed by the ESB. ---------- Page 21, section A2 Comment: We should refer to the Oracle JDK as the "Oracle JDK" not the "Sun Microsystems" JDK. (The following appendix does this already.) ---------- Page 27, Appendix C. Verifying Downloaded Files Comment: The screenshot and example text in this section refer to SOA-P 5 builds, not 4.3. ---------- Page 29, Appendix D. Installing Apache Ant Comment: References to ant 1.8 should be changed to ant 1.7	1	1	https://issues.jboss.org/browse/SOA-3009

#	Requirement	Questions	Answered	Successfully Answered	URL
10916	SOA-3008	Hi Bob: Have we changed the facet name in OUR svn? Or only in the eclipse branch? Currently I have noticed the following files that might be affected: - VCFUtil (one line patch), very easy fix. - Large amounts of org.jboss.tools.bpel.runtimes/plugin.xml The second is the more worrisome version here. I need to know whether this plugin or some other integration plugin like it is going to continue to be distributed at all times. Before I can move forward with this, I need to understand the migration path in general, when we will switch from using our svn version to eclipse-released versions, and whether our svn will continue to hold an integration plugin of some sort. Ready to help, though. =]	0	0	https://issues.jboss.org/browse/SOA-3008
10917	SOA-3008	OK, after thinking about this some more, I've come to the conclusion that we should leave our jbt facet alone since we'll be pushing this stuff back into eclipse anyway. I'm thinking we leave the "bpel.facet.core" that's currently defined in the org.eclipse.bpel.apache.ode.runtime plugin (in the eclipse.org repo) as is, and we leave the "jbt.bpel.facet.core" defined in org.jboss.tools.bpel.runtimes as is, but rename it to something like "JBoss Riftsaw BPEL 2.0 Project Facet". Project facets appear to be bound to specific runtimes anyway, so this kind of makes sense. The two adapter plugins that define these facets could then become separate, installable features; so we'd have the features: 1. BPEL Editor 2. JBoss Riftsaw Runtime Adapter 3. Apache ODE Runtime Adapter Then you (the user) can switch between the two facets if you want to deploy to different runtimes. Is this reasonable?	0	0	https://issues.jboss.org/browse/SOA-3008
10918	SOA-3008	Just an afterthought: the ODE runtime requires that the service address include "ode/processes" in its URL, whereas Riftsaw does not. So for example, a "hello_world" process would be addressed as: http://localhost:8080/ode/processes/hello_world when deployed to ODE and as: http://localhost:8080/hello_world when deployed to Riftsaw. This means that the WSDL (where the service binding is defined) needs to be changed depending on which runtime the service is deployed to. Should this service URL change become part of the facet installation (esp. considering that each runtime may have defined different host and port#s) ?	0	0	https://issues.jboss.org/browse/SOA-3008
10919	SOA-2991	Tom, can you add this to the jbpm config in ESB?	0	0	https://issues.jboss.org/browse/SOA-2991
10920	SOA-2984	Zip file containing signed jar built from: https://svn.devel.redhat.com/repos/jboss-soa/branches/5.0.2-SOA-2984/build-tools/ Revision 8758 which uses: http://anonsvn.jboss.org/repos/labs/labs/jbossesb/branches/JBESB_4_7_CP2_JBESB_3584/ Revision 36860 README.txt: The included file: jbossesb-soap.jar should be copied over the following files: jboss-soa-p.5.0.0/jboss-as/server/<profile>/deploy/soap.esb/jbossesb-soap.jar where "<profile>" is "all", "default" and "production". The included file should also be copied to any additional profiles that have been created.	0	0	https://issues.jboss.org/browse/SOA-2984
10921	SOA-2978	Per denial note, we now have a support case to justify this request. I'll re-open it. https://c.na7.visual.force.com/apex/Case_View?id=500A0000008UO8NIAW&sfdc.override=1 SF case 00561240 Thanks, Rick	1	1	https://issues.jboss.org/browse/SOA-2978
10922	SOA-2974	3.1.2 SP9 - is included in ER6 - for example: ./jboss-as/jbossws-cxf-installer/deploy/lib/jbossws-cxf-server.jar ./jboss-as/jbossws-cxf-installer/deploy/lib/jbossws-cxf-client.jar	1	0	https://issues.jboss.org/browse/SOA-2974
10923	SOA-2925	Do we get this for "free" when we do the MRG-M cert?	0	0	https://issues.jboss.org/browse/SOA-2925
10924	SOA-2924	Kev, Great to have positive feedback from GSS. Can you please add the sample code and possibly the diagram to the project doc and I will pull it down into the product guide? Thanks.	1	0	https://issues.jboss.org/browse/SOA-2924
10925	SOA-2922	Need a little direction. In which document(s) does this appear? (What are migration notes"?)	0	0	https://issues.jboss.org/browse/SOA-2922
10926	SOA-2919	(Per Steve Hawkins) This seems overly general. Teiid exceptions generally contain detailed messages and a detailed exception chain. Can you provide examples when this is occurring and what you would like to see instead?	1	1	https://issues.jboss.org/browse/SOA-2919
10927	SOA-2916	Martin - is this still an issue?	1	0	https://issues.jboss.org/browse/SOA-2916
10928	SOA-2900	The current message in 5.1 that will display is: EDS cannot be deployed to profile ${profile}, it doesnt have all required dependencies. Is there different/improved message that you wish to display?	1	1	https://issues.jboss.org/browse/SOA-2900
10929	SOA-2688	It should already be possible to encrypt the passwords in jbossesb-properties.xml, can you give us an example of where you think this fails?	0	0	https://issues.jboss.org/browse/SOA-2688
10930	SOA-2630	If there are parts missing that you think should be included in the docs then please list them. As for where it should reside, are you intending to move it?	0	0	https://issues.jboss.org/browse/SOA-2630
10931	SOA-2626	Hi Van, I think you said that at least one of the GSG examples would include this. Can you please comment here on which one ? Thanks	0	0	https://issues.jboss.org/browse/SOA-2626
10932	SOA-2616	When using jUDDI v3, it makes much more sense to use user defined keys, and have part of the binding key be the host and port name. This way the existing binding keys are simply reused rather then a new set being created when the servers come back up. Also the new jUDDI v3.1 client ships with a service cache, which is updated by a subscription (when a service or endpoints change or are removed the client caches are informed). This is using standard UDDI v3 functionality so should work with any v3 registry and makes the registry much more dynamic. This reduces the problem to removing dead endpoint information of nodes that are really gone. I would think that a 3 'strikes you're out' policy in the ServiceInvoker should do the trick, that with an optional periodic query ("am I in the registry") - if not, reregistration to bring back services that could not be seen due to for example temporary network issues. "The real solution is to have a 'service repository' that reacts dynamically to cluster members appearing/disappearing" if you have this info then you can do registration/unregistration right there.. My 2 cents. --Kurt	1	0	https://issues.jboss.org/browse/SOA-2616
10933	SOA-2591	Lee Carlon <lcarlon@redhat.com> made a comment on [bug 724407|https://bugzilla.redhat.com/show_bug.cgi?id=724407] Pasted from email: > This needs to be fixed. Lee, why don't we sync up early next week > review this doc ? > Sounds good, Prakash. I'm in the process of having the document added to the BRMS docs on the internal Doc stage. I've marked off the chapters that I believe shouldn't be included for 5.2 as 'Not Supported in 5.2'. I'll let you know as soon as the document is up and we can review from there. - Lee	0	0	https://issues.jboss.org/browse/SOA-2591

#	Requirement	Questions	Answered	Successfully Answered	URL
10934	SOA-2376	doc updated & pushed to stage outstanding items emailed for feedback: Section 1.1 Can be a JCA Connector or any other custom connection provider. Teiid recommends and uses JCA as its specification defines how you can write, package, and configure access to EIS system in consistent manner. Open issue – We recommend that customers use the vendor-supplied JCA adapters for using JMS with JCA. Will we do the same for using Teiid with JCA? Section 1.2.1, 1.3.1 For sample Translator code, refer to the Teiid source code athttp://anonsvn.jboss.org/repos/teiid/trunk/connectors/ Where will the Teiid source be available from ? Will it be a standalone ZIP or included with the SOA source? And (assuming it will be available for download as a ZIP) what directory of the repo will be the root of the ZIP ? eg, if the docs refer to http://anonsvn.jboss.org/repos/teiid/trunk/connectors/ how will that be available in the ZIP - noting that the product docs will really be referring to the branch source although the upstream docs generally refer to the trunk. Section 3.1.7. Metadata Override and implement the method getMetadata(), if you want to expose the metadata about the source for use in Dynamic VDBs. This defines the tables, column names, procedures, parameters, etc. for use in the query engine. This method is not yet used by Designer tooling. Open issue: Need to check with the Dev team about “This method is not yet used by Designer tooling.” What are the user impacts of this?	0	0	https://issues.jboss.org/browse/SOA-2376
10935	SOA-2370	Toshiya - Do not set the fix version when opening Jira issues, leave the fix version unassigned so it can be picked up to be triaged. Looking at the linked Jira issue this is already resolved for 5.1 and is a duplicate? Or is there further work for this issue?	0	0	https://issues.jboss.org/browse/SOA-2370
10936	SOA-2309	AdminShell Bootstrap Environment TEIID_HOME = /jboss/local/51_ER2/jboss-soa-p-5/eds/teiid/adminshell CLASSPATH = /jboss/local/51_ER2/jboss-soa-p-5/eds/teiid/adminshell/lib/patches/*:/jboss/local/51_ER2/jboss-soa-p-5/eds/teiid/adminshell/lib/teiid-adminshell-7.1.1.jar:/jboss/local/51_ER2/jboss-soa-p-5/eds/teiid/adminshell/lib/* JAVA = /usr/lib/jvm/java-1.6.0-openjdk-1.6.0.0//bin/java ========== ===> [import static org.teiid.adminshell.AdminShell.*; import static org.teiid.adminshell.GroovySqlExtensions.*; import org.teiid.adminapi.*;] ===> [import static org.teiid.adminshell.AdminShell.*; import static org.teiid.adminshell.GroovySqlExtensions.*; import org.teiid.adminapi.*;, import org.testng.annotations.*] ===> [import static org.teiid.adminshell.AdminShell.*; import static org.teiid.adminshell.GroovySqlExtensions.*; import org.teiid.adminapi.*;, import org.testng.annotations.*, import org.testng.TestNG] ===> [import static org.teiid.adminshell.AdminShell.*; import static org.teiid.adminshell.GroovySqlExtensions.*; import org.teiid.adminapi.*;, import org.testng.annotations.*, import org.testng.TestNG, import org.testng.TestListenerAdapter] FATAL: org.codehaus.groovy.control.MultipleCompilationErrorsException: startup failed: groovysh_parse: 6: unexpected token: @ @ line 6, column 5. @BeforeClass ^ 1 error org.codehaus.groovy.control.MultipleCompilationErrorsException: startup failed: groovysh_parse: 6: unexpected token: @ @ line 6, column 5. @BeforeClass ^ 1 error at org.codehaus.groovy.control.ErrorCollector.failIfErrors(ErrorCollector.java:296) at org.codehaus.groovy.control.ErrorCollector.addError(ErrorCollector.java:103) at org.codehaus.groovy.control.ErrorCollector.addFatalError(ErrorCollector.java:142) at org.codehaus.groovy.control.ErrorCollector.addError(ErrorCollector.java:113) at org.codehaus.groovy.control.ErrorCollector.addError(ErrorCollector.java:125) at org.codehaus.groovy.control.SourceUnit.addError(SourceUnit.java:353) at org.codehaus.groovy.antlr.AntlrParserPlugin.transformCSTIntoAST(AntlrParserPlugin.java:104) at org.codehaus.groovy.antlr.AntlrParserPlugin.parseCST(AntlrParserPlugin.java:71) at org.codehaus.groovy.control.SourceUnit.parse(SourceUnit.java:249) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:616) at org.codehaus.groovy.runtime.callsite.PojoMetaMethodSite$PojoCachedMethodSiteNoUnwrapNoCoerce.invoke(PojoMetaMethodSite.java:229) at org.codehaus.groovy.runtime.callsite.PojoMetaMethodSite.call(PojoMetaMethodSite.java:52) at org.codehaus.groovy.runtime.callsite.AbstractCallSite.call(AbstractCallSite.java:121) at org.codehaus.groovy.tools.shell.RigidParser.parse(Parser.groovy:148) at org.codehaus.groovy.tools.shell.RigidParser$parse.call(Unknown Source) at org.codehaus.groovy.tools.shell.Parser.parse(Parser.groovy:67) at org.codehaus.groovy.tools.shell.Parser$parse.call(Unknown Source) at org.codehaus.groovy.tools.shell.Groovysh.execute(Groovysh.groovy:130) at org.codehaus.groovy.tools.shell.Groovysh$execute.callCurrent(Unknown Source) at org.codehaus.groovy.tools.shell.Shell.leftShift(Shell.groovy:114) at org.codehaus.groovy.tools.shell.Shell$leftShift$0.call(Unknown Source) at org.codehaus.groovy.tools.shell.commands.LoadCommand$_load_closure1.doCall(LoadCommand.groovy:80) at	0	0	https://issues.jboss.org/browse/SOA-2309
10937	SOA-2308	This could be problematic, because the config file can't templetized because we can't be sure what the user has already changed. They now could have multiple sources that use HSQLDialect. Do we, by default, change all occurrences? Maybe, when the db schema tool is run, and modeshape is determined to be deployed, we add a prompt telling the user that they will need to update the config file with this "DIALECT". Suggestions?	0	0	https://issues.jboss.org/browse/SOA-2308
10938	SOA-2303	The config file needs to be changed from using the default inmemory repository to using hsql. A question here is - should the config define its own hsql (keep its data separate from everyone else) or reference the default repository by jndi name (DefaultDS) ?	0	0	https://issues.jboss.org/browse/SOA-2303
10939	SOA-2282	Regarding ModeShape: changed it to use the soa-users.properties and soa-roles.properties files, therefore added 2 steps to build process to: 1) exclude the modeshape users and and roles files and 2) remove the policy from the modeshape-jboss-beans.xml and added a step to modeshape installation to insert the policy to login-config.xml Teiid: added a step during build process to remove the application-policy from the teiid-jboss-beans.xml and added a setup to teiid installation to add the policy to login-config.xml Didn't change Teiid to use the soa-users.properties because these users represent the JDBC access to EDS.	0	0	https://issues.jboss.org/browse/SOA-2282
10940	SOA-2197	Need to find out: 1 - Has this been tested ? 2 - What configuration is different ?	0	0	https://issues.jboss.org/browse/SOA-2197
10941	SOA-2197	HP Systinet can be deployed just to 4.3 SOA-P, because their web framework does not work on 5.x SOA-P server (see SOA-1919). But it's possible to connect from 5.x SOA-P server to HP Systinet deployed on 4.3 SOA-P (use remote server as remote UDDI registry).	0	0	https://issues.jboss.org/browse/SOA-2197

#	Requirement	Questions	Answered	Successfully Answered	URL
10942	SOA-2183	Len, this seems to already be corrected in the current build. Page 94 reads: 2. SQL: <sql-provider>, <sql-bus>, <sql-listener> and <sql-message-filter>: The <sql-message-filter> can be added to either the <sql-bus> or <sql-listener> elements. Where the <sql-provider> and <ftp-bus> specify the JDBC connection properties, the <sql-message-filter> specifies the message/row selection and processing properties. 3. FTP: <ftp-provider>, <ftp-bus>, <ftp-listener> and <ftp-message-filter>: The <ftp-message-filter> can be added to either the <ftp-bus> or <ftp-listener> elements. Where the <ftp-provider> and <ftp-bus> specify the FTP access properties, the <ftp-message-filter> specifies the message/file selection and processing properties 4. Hibernate: <hibernate-provider>, <hibernate-bus>, <hibernate-listener> : The <hibernate- message-filter> can be added to either the <hibernate-bus> or <hibernate-listener> elements. Where the <hibernate-provider> specifies file system access properties like the location of the hibernate configuration property, the <hibernate-message-filter> specifies what classnames and events should be listened to. 5. File system: <fs-provider>, <fs-bus>, <fs-listener> and <fs-message-filter> The <fs-message- filter> can be added to either the <fs-bus> or <fs-listener> elements. Where the <fs-provider> and <sql-bus> specify the File System access properties, the <fs-message-filter> specifies the message/file selection and processing properties. Is this the section to which Jehudi is referring? If so, it seems to be correct in the current version of the book attached. Cheers.	1	1	https://issues.jboss.org/browse/SOA-2183
10943	SOA-2145	can you make sure the version numbers of the JAR files are up to date ? Then clone this JIRA with the remaining items & set the FixVersion to FUTURE Thanks.	1	1	https://issues.jboss.org/browse/SOA-2145
10944	SOA-2080	Need some more information for the release notes: How does this affect developers/administrators ? Is there anything they might have to change in their deployed applications ?	0	0	https://issues.jboss.org/browse/SOA-2080
10945	SOA-2064	Kevin, what does this mean: At the moment you should be able to achieve this through the help of static methods. I have tried the properties within the jboss-esb.xml the way you suggested by email, and it works, so I dropped that out of my codebase. However, I did not attempt this in my code until I went through JBoss Support, which said it was not possible: See: https://support.redhat.com/jbossnetwork/restricted/caseDetail.html?caseId=832723 As far as properties in the ESB file, can you point me to the point in the documentation where that is noted?	0	0	https://issues.jboss.org/browse/SOA-2064
10946	SOA-2064	As far as this bug goes, the logic being in the constructor also leads to the issue of the whole service failing when the proxied web service is temporarily unavailable on server restart. You suggested: "It is currently a requirement that the wsdl exist when starting the service, although it should be possible to handle this through a barrier of sorts. What is your scenario?" The scenario is that, if SOA-P is bounced, and a webservice is unavailable that it should be proxying, it leads to an exception on deployment of the ESB service, and *no way* to recover the service, even if the remote service comes back online. If this is a requirement of the SoapProxy, it is a design flaw. The SoapProxy should return a SoapFault until the point that the service becomes available.	0	0	https://issues.jboss.org/browse/SOA-2064
10947	SOA-2032	Comments on doc here: http://downtown.englab.bne.redhat.com/drafts/JBoss_Enterprise_SOA_Platform/4.3.CP04/SOA_Release_Notes/ 3.1.4. JBoss Rules SOA-1914 https://jira.jboss.org/jira/browse/JBRULES-2421 If the user was running JBoss Rules 4.0.7 on an IBM virtual machine with JIT enabled, an .RuntimeDroolsExceptions would occur. This was because of a problem with regard to the way in which XXXX instanceof Object expressions were handled. The code in the org.drools.base.ValueType class which handles these expressions has been modified. As a result, JBoss Rules can now be run on IBM virtual machines with JIT enabled, without it causing exceptions. *** Comment: Note that SOA Platform 4.3 CP04 is not certified with the IBM JVM. ---------- https://jira.jboss.org/jira/browse/SOA-1072 The example.base.build.xml used by jBPM examples has now been removed. It has previously caused some problems but is now no longer needed as examples are now being built with Maven rather than Ant. POM files have now been provided to test and run the examples. *** Comment: This bug is still open in CP04 - we need to list it as an unresolved issue, ---------- https://jira.jboss.org/jira/browse/SOA-1463 A java.lang.Exception exception would occur if the user tried to run start, stop, create or destroy deployment operations from within the ESB Statistics level. In order to mitigate this problem, these operations are no longer available at this level. *** Comment: This bug is still open in CP04 - we need to list it as an unresolved issue, ---------- https://jira.jboss.org/jira/browse/SOA-1998 *** Comment: See the JIRA - the description of this issue is changing. ---------- Known Issues in This Release 4.1.1. JBoss Business Process Manager SOA-1948 https://jira.jboss.org/jira/browse/JBPM-2791 If there is a variable mapping between the Enterprise Service Bus and the Business Process Manager and the mapped value happens to be an empty string, then the result will be null. Also, if an object is persisted in an Oracle database with a field set to be an empty sting, the field value will also be null. Currently, there is no work-around for this problem. 4.1.2. JBoss Enterprise Application Platform SOA-1519 https://jira.jboss.org/jira/browse/JBPAPP-2893 There are problems with crash recovery when inflowing transactions are associated with multiple resource managers. As a consequence, a single XAResource can not be returned to drive the recovery for a given Xid. There is currently no work-around solution for this problem. *** Comment: Have we used this phrase in the past: "There is currently no work-around solution for this problem." ---------- SOA-1529 https://jira.jboss.org/jira/browse/JBESB-3197 Currently, there is no service invoker load-balancing policy that prefers local end-point references. This would be of use if users do not want requests to be balanced between different instances in the clutter, as the latter makes debugging harder in the event of a failure. There is no work-around for this lack of functionality. *** This phrase "There is no work-around for this lack of functionality" sounds overly negative - why not simply say that the functionality is not supported by the Platform.	0	0	https://issues.jboss.org/browse/SOA-2032
10948	SOA-2004	Sorry to be a pain - but the names of the new figures: Figure 6.6. Metrics Disabled by Default One Figure 6.7. Metrics Disabled by Default Two Really should be changed - the titles that you used are from the screenshots that I attached to the JIRA, but those were just the filenames. ;-) How about: Figure 6.6. Service Metrics Disabled by Default Figure 6.7. Action Metrics Disabled by Default Thanks!	1	1	https://issues.jboss.org/browse/SOA-2004

#	Requirement	Questions	Answered	Successfully Answered	URL
10949	SOA-2003	Additional comments: ---------- Note in section 4.1.3 In JBoss ESB 4.3.CP03 the name of the property used to enclose NotificationList elements in the <action> target is not validated. Is this^ new in 4.3 CP03? ---------- Section: 4.3.1. ServiceInvoker This paragraph should be removed: From a clients perspective, the Courier interface and its various implementations can be used to interact with services. However, this is still a relatively low-level approach, requiring developer code to contact the registry and deal with failures. Furthermore, since JBossESB has fail-over capabilities for stateless services, this would again have to be managed by the application. See the Advanced chapter for more details on fail-over. ---------- Chapter 6. Example Comment: Better to call this section "An Example" ---------- 7.2.5. Quartz Scheduler Property Configuration I'm not sure if we support this - we removed it from the 5.0 book - will check with Kevin ---------- Duplicated sections: 11.5.3.2.1. Updated HTTP Connection Configuration. This "Configurator" is always applied to HttpClient instances created by the HttpClientFactory. To make the connection configurations in a different way, create a custom configurator that can generate and attach a different HttpConnectionManager instance with its own settings. 11.7.3.1. Updated HTTP Connection Configuration. This "Configurator" is always applied to HttpClient instances created by the HttpClientFactory. To make the connection configurations in a different way, create a custom configurator that can generate and attach a different HttpConnectionManager instance with its own settings.	1	1	https://issues.jboss.org/browse/SOA-2003
10950	SOA-2002	Hello Martin, As per SOA-1998 I have just added information about changes in numbers of retries that Alejandro documented. Is that what the part to which you were referring? If so, it is now in the document (Section 10.4.) Cheers.	1	1	https://issues.jboss.org/browse/SOA-2002
10951	SOA-2002	FYI - can you please look at these question that Len has raised?	1	1	https://issues.jboss.org/browse/SOA-2002
10952	SOA-1954	This will not be considered for SOA 4.3. Keeping 5.x only. Decision by SOA PM team triage meeting.	1	1	https://issues.jboss.org/browse/SOA-1954
10953	SOA-1950	It has been decided for now to keep the manifest file within the product zip. The MANIFEST.MF file will reside in a top level directory called Manifest. [12:37] <ldimaggi_> mschoene, tkirby remember the SOA-P 5.0 manifest files? Was this done as a stopgap measure as we were unable to sign all the jars? [12:38] <mschoene> ldimaggi_, yes, we wanted to have trusted checksum [12:38] <ldimaggi_> mschoene, tkirby what's the state of signing for 5.1? [12:39] <tkirby> tkirby, mschoene broken by the attempt to add a manifest :-) [12:40] <mschoene> tkirby, who is guilty , hen or egg ? [12:41] <jdc> mschoene: Dinsoaurs. [12:41] <mschoene> jdc, :) [12:41] <tkirby> mschoene, signing was working fine, but the manifest appeared in the top level, attempting to move this outside has caused problems so we can [12:42] <tkirby> mschoene, ldimaggi_ a) just sign [12:42] <tkirby> b) sign and have manifest in top level [12:42] <mschoene> tkirby, the signing process creates manifests too , btw [12:42] <tkirby> c) sign and have manifest seperate [12:42] <mschoene> tkirby, staying with b) is great [12:43] <tkirby> mschoene, ok because that is what we had [12:43] <tkirby> griz ping ^^^^^^^^^^ [12:45] <griz> tkirby: ok, I'll put it back to how it was. [12:45] <ldimaggi_> tkirby, griz mschoene great! [12:47] <tkirby> griz is it still in a meta-inf directory? [12:49] <griz> tkirby: I'd changed it to be a freestanding file in it's own zip. I'll take my changes out of the sign-unsigned-jars and then it'll be in the META-INF dir in the product zip. [12:49] <tkirby> griz can we rename that to Manifest? [12:49] <tkirby> griz or is that what you get given? [12:51] <ldimaggi_> tkirby, please put all this ^ into the JIRA - or we will forget it all [12:51] <griz> tkirby: A directory called Manifest in the prod zip containing a file called MANIFEST.MF? Sure. [12:51] <griz> ldimaggi_: will do. [12:51] <tkirby> mschoene, ^^^^ That OK with you? [12:54] <tkirby> griz unless mschoene objects go ahead with that, add the conversation to the Jiras [12:54] <griz> tkirby: will do. [12:54] <mschoene> tkirby, hmm, i am not sure i understand that entirely from the backlog , i'll comment in the jira [12:55] <tkirby> griz we may have a delay to the build so you probably have a day or so to revert it [12:55] <griz> tkirby: shouldn't take long to change and check [12:55] <tkirby> mschoene, the manifes will be call manifest.mf and will be in the top level of the zip under a dir called manifest [12:57] <griz> mschoene: hi. Just to make sure, The MANIFEST.MF will live in a directory called Manifest which will be a top level directory in the product zip. [12:58] <mschoene> griz, that sounds good to me, will that be in next ER? [12:58] <griz> mschoene: Yep. [13:00] <mschoene> griz, ok, then i'll have a look at that file too once available	0	0	https://issues.jboss.org/browse/SOA-1950
10954	SOA-1949	Hey David, can you add this to the release notes? It won't be fixed for GA. The consequence of this issue is that you cannot use the same JON server and agent to manage SOA instances with different versions (4.3 and 5.0). JON will remove the older esb plugin and you won't be able to manage SOA-P 4.3. So either you use the SOA 4.3 plugin to manage SOA 4.3 or you use the SOA 5 plugin to manage SOA 5. A workaround, which allows you to manage both instances at the same time is provided in comments above (rename the plugin name).	1	1	https://issues.jboss.org/browse/SOA-1949
10955	SOA-1913	What version do users require if they wish to use native ?	1	1	https://issues.jboss.org/browse/SOA-1913
10956	SOA-1904	I need some idea of what is happening with this JIRA so I can schedule docs work for 5.0.1 What volume of material? e.g. how many pages, diagrams, etc When will draft material be ready ?	0	0	https://issues.jboss.org/browse/SOA-1904
10957	SOA-1904	Here is the first draft of the document. Darrin, would you please take a look at it? http://community.jboss.org/wiki/jBPM3WebApponSOA-P In addition to the document, it would be great to have a unit-testable sample application that illustrated the concepts. What do you think?	1	1	https://issues.jboss.org/browse/SOA-1904
10958	SOA-1903	Approved for SOA 4.3 CP03, but depends on inclusion into EAP 4.3 CP08 and the ability for SOA 4.3 CP03 to consume EAP 4.3 CP08 in time for release.	0	0	https://issues.jboss.org/browse/SOA-1903
10959	SOA-1884	I think this needs to be an update of the kbase article that was created for SOA-1444 Can this be confirmed ?	0	0	https://issues.jboss.org/browse/SOA-1884
10960	SOA-1876	So what is this optimization? ..... Basically, if the Smooks configuration passes the following tests: 1. Has just an XSL config in it (nothing else). 2. The XSL config is targeted at the whole message i.e. the "#document" or "/" fragment i.e. not a sub-fragment of the message. 3. Is using the default (XML) reader config i.e. is not configured with a reader for CSV, JSON, EDI etc The XSLT will be applied directly, bypassing the Smooks fragment filtering process completely.	0	0	https://issues.jboss.org/browse/SOA-1876
10961	SOA-1769	Len, As per the above, all of the synchronousTimeout info has been added to the document. However, found the statement about documenting this Wiki ambiguous. Do you want it all documented? I ask this because there are significant differences between what is in the Wiki and what is in the upstream docs (for example, the entire "Security" section is missing from the Docs.) Since the Wiki was updated recently (Jan 10) by Tom Fennelley, do we take it as the "source of truth" for the entire HTTP Gateway section? Thanks.	1	1	https://issues.jboss.org/browse/SOA-1769
10962	SOA-1763	Okay I stand corrected, but can this be written to create a "transformer" pool so that each thread running can get their own instance. This would be a big improvement "approximately 20%" or so in my testing performance? I assume in a http_gateway multiple threads are running, is that true, how can that be confirmed? In my testing I saw no errors, so at least with my configuration is it possible this is only single threaded?	0	0	https://issues.jboss.org/browse/SOA-1763
10963	SOA-1763	Sorry, what does the gateway have to do with this?	1	1	https://issues.jboss.org/browse/SOA-1763

#	Requirement	Questions	Answered	Successfully Answered	URL
10964	SOA-1738	Where is the jbpm-console-production.war? where is the jbpm-console-development.war?	0	0	https://issues.jboss.org/browse/SOA-1738
10965	SOA-1730	Section 7.2.5 of the Programmers' Guide describes the ESB's use of a Quartz Scheduler - is this the same JMS/JCA inflow/quartz scheduler?	0	0	https://issues.jboss.org/browse/SOA-1730
10966	SOA-1730	Sorry Jiri, I don't know why this document never updated on the doc-stage site with the content as per my comment on the 3rd of February. I've just rebuilt it now, hopefully it will show up in a few hours. If for some reason it doesn't (don't know why), there is another copy here: http://downtown.englab.bne.redhat.com/drafts/JBoss_Enterprise_SOA_Platform/5.0.0/Administration_Guide/html/message_and_scheduler_service.html Can you please supply what additional content is required.	0	0	https://issues.jboss.org/browse/SOA-1730
10967	SOA-1730	Re your comment above: https://issues.jboss.org/browse/SOA-1730?focusedCommentId=12514180&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12514180 Does that mean that the originally requested content is now wrong ? If so can you please supply updated content for this JIRA ? (I misread that comment originally, feel kind of stupid now)	1	0	https://issues.jboss.org/browse/SOA-1730
10968	SOA-1715	Jiri, Can you give a few xml snippets that illustrate the different configuration possibilities of these parameters? Thanks, Koen	1	0	https://issues.jboss.org/browse/SOA-1715
10969	SOA-1688	Darrin - the release notes are not ready for review yet - will you be adding the new feature list in this doc?	0	0	https://issues.jboss.org/browse/SOA-1688
10970	SOA-1688	Section: 1. Overview Text: Warning - This is a Beta Release Comment: Remove this note. ;-) ---------- Not a technical review comment - but I think it's worth changing the opening of the first chapter from this: ========== JBoss Enterprise SOA Platform is a certified, tested, and supported platform for developing Enterprise Application Integration and Service Oriented Architecture solutions. It integrates several stable and scalable open source frameworks and solutions including Hibernate, Seam, JBoss Transactions, JBoss Clustering, the JBoss Application Server, and JBoss Enterprise Service Bus (ESB) to provide an infrastructure for enterprise SOA applications. These community developed and enterprise certified and supported products have been combined and tested to provide a solid, robust, and scalable platform. Powered by legendary JBoss innovation and backed by Red Hat engineering and quality assurance, JBoss Enterprise SOA Platform is the platform of choice for a new generation of enterprise applications. ========== To this: ========== JBoss Enterprise SOA Platform is a certified platform for developing Enterprise Application Integration and Service Oriented Architecture solutions. The Platform integrates several stable open source frameworks and community developed products including Hibernate, Seam, JBoss Transactions, JBoss Clustering, the JBoss Application Server, and JBoss Enterprise Service Bus (ESB) to provide a solid, robust, and scalable infrastructure for enterprise SOA applications. Powered by JBoss innovation and backed by Red Hat engineering, quality assurance and support, JBoss Enterprise SOA Platform is the platform of choice for a new generation of enterprise applications. ---------- Section 2 - New Features in 5.0 Text: The JBoss ESB now uses Apache jUDDI version 3 as its UDDI registry. This new version implements the UDDIv3 specification and includes a subscription API in addition to a new web console. These new capabilities are detailed in the new jUDDI User Guide. Comment: If we reference the jUDDI guide - then we need at least an initial draft of the guide at GA. Text: JBoss SOA Administration Console (JOPR) Comment: I could not find any direct reference to JOPR in the EAP 5.0 docs - can you check and see how they described the console? Was there a deliberate reason to not reference JOPR? ---------- 2.1. Technology Previews Text: JBoss Rules Fusion Comment: It's worth referencing Fusion as a CEP implementation. ---------- Section: 3. Components of the 5.0.0 Release Text: Apache jUDDI 3.0 Comment: Should be 3.0.1 ---------- Not looking at this section as it is out-of-date: 8. Known Issues in this release	0	0	https://issues.jboss.org/browse/SOA-1688
10971	SOA-1681	I need some idea of what is happening with this JIRA so I can schedule docs work for 5.0.1 What volume of material? e.g. how many pages, diagrams, etc When will draft material be ready ?	0	0	https://issues.jboss.org/browse/SOA-1681
10972	SOA-1660	Are you monitoring that server with JON?	1	1	https://issues.jboss.org/browse/SOA-1660
10973	SOA-1660	Is this consistent with other information used/collected by JON? It's no longer available after JON collects it? Or - is this configurable? Either way - we will need to document this for users.	0	0	https://issues.jboss.org/browse/SOA-1660
10974	SOA-1660	Could this be a difference between full JON (enables historical data) and embedded JON?	0	0	https://issues.jboss.org/browse/SOA-1660
10975	SOA-1660	Hey Tom, since the ESB JON plugin consumes these messages, where does it use this data? Is it exposed somewhere?	0	0	https://issues.jboss.org/browse/SOA-1660
10976	SOA-1660	This is expected behaviour - should expectations be changed, if so what should be expected	0	0	https://issues.jboss.org/browse/SOA-1660
10977	SOA-1660	Clearing the message alerts is expected behavior for the JON plugin. Not sure what the NPE here is - do we know it is definitely related?	0	0	https://issues.jboss.org/browse/SOA-1660
10978	SOA-1660	Question: If the JON Server or the Embedded Console are consuming these alerts then the collected alerts cannot be accessed programmatically by a user via JMX. They will have to access these via the JON API. Where is the best reference to refer users to for this task?	0	0	https://issues.jboss.org/browse/SOA-1660
10979	SOA-1660	Hi Jarek Can you add links to where to download the SOAP server used to reproduce this issue. Also we need the exact reproduction steps, e.g. which jmx console pages were visited to trigger the NPE.	0	0	https://issues.jboss.org/browse/SOA-1660
10980	SOA-1657	Off the top of my head, candidates for removal: Chapter 2 - it's only download/install instructions Chapter 4 - database config directions, this is handled within SOA already Chapter 6 - is minimal, does it apply to SOA anyway? Chapter 8 - is just placeholder content, unless content is supplied from jUDDI project Also are the screenshots correct for what will be shipping in SOA ?	0	0	https://issues.jboss.org/browse/SOA-1657
10981	SOA-1657	My vote: Chapter 2 - it's only download/install instructions Yes, remove it Chapter 4 - database config directions, this is handled within SOA already Yes, remove it Chapter 6 - is minimal, does it apply to SOA anyway? I think that this should stay Chapter 8 - is just placeholder content, unless content is supplied from jUDDI project Kurt did write up something on annotations here: http://apachejuddi.blogspot.com/2009/11/uddi-annotations-how-do-i-self-register.html Haven't seen anything in the project docs yet.	1	1	https://issues.jboss.org/browse/SOA-1657
10982	SOA-1657	Section 1.1 "Introduction" and Section 1.2 "UDDI Registry" are copies of sections 1.1.1.1 "Introduction" through 1.1.1.6 "The Registry and the JBoss Service-Oriented Architecture Platform" in the Services Guide. But... The information is more detailed and presented better in the services guide. ****** Given the amount of overlap and and of jUDDI-project specific material that we are already removing from the SOA-P jUDDI guide, maybe it's time that we drop the standalone SOA-P jUDDI guide and just expand the services guide?	1	1	https://issues.jboss.org/browse/SOA-1657
10983	SOA-1657	Darrin - can you create a new draft of this doc? Some of the changes needed are pretty dramatic - can we review this again before beta?	1	1	https://issues.jboss.org/browse/SOA-1657

#	Requirement	Questions	Answered	Successfully Answered	URL
10984	SOA-1613	Steps to recreate bug: R+D added a unit test case (LogMoveTest.java)[1] which ensures the function. But in addition, I'd like to do a manual integration test in order to confirm that this one-off solves the problem which NTT encountered in SOA-P environment. The manual test procedure is like: ========= 1. Set up 2 SOA-P instances(ESB node and JMS node) and deploy a test ESB service. (I have a zip file which includes all stuff) 2. Start SOA-P instances. 3. Set a breakpoint on BasicAction.phase2Commit(boolean)[line2172:updateState()] using a debugger. 4. Invoke a test ESB client to start a transactional work. 5. After the ESB node stops at the breakpoint, kill(kill -9) the node. 6. Start the ESB node. 7. Priodic recovery (2min by default) kicks in. The following warning is logged because this is an unrecoverable transaction(Resources are already committed, but the tx-log remains) Could not find new XAResource to use for recovering non-serializable XAResource < 131075, 30, 28, 1--3f57f59b:c51c:4a26721f:21b5-3f57f59b:c51c:4a26721f:21b6 > 8. [If the one-off works correctly] After the specified time (12hr by default), AtomicActionExpiryScanner should move the log to "Expired" directory. Then you can see the warning is not being logged any longer. ========== Yes, it's reproducible. Is this acceptable to QE team? It's a bit complex so I think I should test it by myself during the QE phase. [1] https://jira.jboss.org/jira/browse/JBTM-418 Regards, Toshiya	0	0	https://issues.jboss.org/browse/SOA-1613
10985	SOA-1589	Taro, I have actually reworded it slightly in order to make it read better. The new wording is below but can you also please give a little information about who was modifying the format and why? Was it a particular customer? Many thanks. 2.2.1. Format of End Point Reference Contents Has Changed The contents of message.getHeader().getCall().getTo().getAddr().getAddress() are now returned in a different format than that in the 4.3GA and 4.3.CP01 releases. Previously, the format was in the form of: jms://soa01:1200/queue/myRequestQueue but now it is: jms:soa01:1200#queue/myRequestQueue Users should not modify End-Point Reference formats directly because they are specific to implementations of the API. There is no guarantee the format will remain the same for them.	1	1	https://issues.jboss.org/browse/SOA-1589
10986	SOA-1587	There is no such warning. Does not the script support multiple execution now?	0	0	https://issues.jboss.org/browse/SOA-1587
10987	SOA-1575	Smooks defines multiple terms that are unique to the project. Can we a Glossary? 1. Smooks Core: Provides the basic Smooks infrastructure/framework. Smooks Core is responsible for processing a set of Smooks Resource configurations. From these configurations, Smooks Core constructs and manages "Execution Context", "Content Delivery Configuration" and "Filter" (DOM/SAX) components. Smooks Core is the subject of this section. 2. Smooks Cartridges: Provide the real functionality on top of Smooks Core e.g. Javabean/POJO population, Templating support (XSLT, FreeMarker, StringTemplate), Scripting Support (Groovy) etc. These are the resources that will be applied by Smooks Core during filtering (see below). 3. Visitor: Visitor logic is central to how Smooks works. A Visitor is a simple piece of Java logic that can perform a specific task on the message fragment at which it is targeted, such as applying an XSLT stylesheet. You have the choice of supporting your logic through either the SAX or DOM Filters by implementing one or both of the following interfaces: org.milyn.delivery.sax.SAXElementVisitor or org.milyn.delivery.dom.DOMElementVisitor. 4. Java Binding: This is the ability to populate a Java Object Model from the source message. 5. Message Splitting and Routing: This the ability to perform complex splitting and routing operations on the source message. It includes the ability to route to multiple destinations concurrently. Different data formats can also be routed concurrently. These formats include XML, EDI, CSV and Java. 6. Huge Message Processing: This the ability to be able to declaratively "consume" (that is, transform or split-and-route) very large messages without the need for one to write substantial amounts of high-maintenance code. 7. Cartridge - A cartridge is a Java Archive (JAR) that contains reusable Content Handlers. Smooks includes pre-built and ready to use visitor logic so that users can implement solutions with minimal java code. This visitor logic is bundled into categories. These bundles are called cartridges. 8. Content Handler - Can't find a definition in the doc for this. Content Handlers are the cornerstone of the Smooks component model. The following Content Handler types are currently in existance: 1. Stream Readers: Smooks supports filtering of both XML and non-XML data because it allows you to configure a "Stream Reader" for each filter process. If no reader is configured, it defaults to XML. So, the Stream Reader resource is responsible for generating a stream of SAX events from a hierarchical data stream (e.g. XML, CSV, EDI etc.). This stream of SAX events can then be processed by XML Element Visitors (via Smooks). 2. Element Visitors: After Smooks hooks a Stream Parser to the data Stream, it starts receiving a stream of SAX events i.e. "startElement", "endElement" etc. These events are then used by Smooks to select an "ElementVisitor" implementation, which process the event in some way. This is the primary extension point in Smooks, as well as being the mechanism though which Smooks supports a fragment based processing model. See the list of Smooks Cartridges for examples of ElementVisitor implementations already available with Smooks. 3. Element Serializers: DOM based processing supports implementation of DOM Element "Serialization Units", allowing you to implement custom serialization at a fragment level.	1	1	https://issues.jboss.org/browse/SOA-1575

#	Requirement	Questions	Answered	Successfully Answered	URL
10988	SOA-1575	Binding This feature is used to populate a Java Object Model from a data source (such as a CSV, EDI, XML or Java file.) The populated object models can then be used either as transformation results in themselves or, alternatively, as a templating resources from which XML or other character- based results can be generated. This feature also supports Virtual Object Models (maps and lists of typed data), which can be used by both the ETL and templating functionality. Comment: This is the first place in the book where the term "ETL" is used. We should define the term: Extract Transform Load (ETL) here. It's currently defined 2 pages later. ---------- Section 1 - Overview Comment: Invalid links in PDF file: http://www.smooks.org/mediawiki/index.php%3Ftitle=FAQ should be: http://www.smooks.org/mediawiki/index.php?title=FAQ http://www.smooks.org/mediawiki/index.php%3Ftitle=Maven_%2526_Ant should be: http://www.smooks.org/mediawiki/index.php?title=Maven_%26_Ant http://www.smooks.org/mediawiki/index.php%3Ftitle=Smooks_v1.2_Examples should be: http://www.smooks.org/mediawiki/index.php?title=Smooks_v1.2_Examples <<<Open issue: We really do not want to direct Platform customers to project URLs for examples, documentation, etc. In the case of Smooks, avoiding this will be difficult as the Smooks project documentation is in the form of a tree of wiki pages. Can the Smooks project provide the same information in a standalone PDF file? In the case of the example code, if we include Smooks in the Platform, we should also include the examples' code. >>> ---------- Section 2.1. Basic Processing Model Comment: Another invalid link http://www.smooks.org/mediawiki/index.php%3Ftitle=V1.2%3ASmooks_v1.2_Developer_Guide should be: http://www.smooks.org/mediawiki/index.php?title=V1.2:Smooks_v1.2_Developer_Guide <<<Another open issue: It's one thing to send Platform customers to a project URL for wiki pages. But if we are shipping a user guide that references a developer guide, then we should also ship the developer guide doc.>>> ---------- 2.2. Simple Example Comment: I know that the document is adapated from the Smooks' project on-line docs, but I think that it is a mistake to retain this example. The problems with the example are that the code fragement is incomplete and the example does not generate any ouput. What do you think about replacing the example with the example from: http://www.smooks.org/mediawiki/index.php?title=V1.2:xml-to-xml ? And - the name of the interface is: ExecutionContext not EvolutionContext ;-) Text: In this example, no result is produced. There is also no interaction with the execution of the filtering process. This is because the example didn't create an EvolutionContext and supply it to the Smooks.filterSource method call. See http://www.milyn.org/javadoc/v1.2/smooks/org/milyn/container/ExecutionContext.html Comment: This is the first time that the Smooks "ExecutionContext" interface is mentioned in the document. We really should define the need for this interface, and any others that are needed in the example before we show the user the example. ---------- 2.3. Smooks Cartridges Text: Smooks includes pre-built and ready to use visitor logic so that users can implement solutions with minimal java code. This visitor logic is bundled into categories. These bundles are called cartridges. Comment: We should not use the term "bundle" as it may cause user confusion with Java bundles: http://java.sun.com/j2se/1.4.2/docs/api/java/util/ResourceBundle.html How about "This visitor logic is combined into groups called "cartridges." Comment: Incomplete link: http://www.smooks.org/cartri ---------- Section: 2.4. Filtering Process Selection (Should DOM or SAX be Used?) Comment: We should remove the "(Should DOM or SAX be Used?)" question from the section title and add it in an opening sentence for the section. Text: This does not include non-element Visitor resources, like, for example, readers. Comment: What is a "non-element Visitor resource?" Also, in this context, what is a "reader?"	1	1	https://issues.jboss.org/browse/SOA-1575
10989	SOA-1575	Section: Example 7.1. A Simple Configuration Comment: There is a duplicate record in the output. ---------- Section: 7.1.4.1.3. Segment Matching Comment: The reference to Java 1.5 (http://java.sun.com/j2se/1.5.0/docs/api/java/util/regex/package-summary.html) should be replaced with this refenence to Java 1.6 (http://java.sun.com/javase/6/docs/api/java/util/regex/package-summary.html) ---------- Section: 7.1.4.3. Type Support Comment: Reference to this: http://www.smooks.org/mediawiki/index.php?title=V1.2:Smooks_v1.2_User_Guide#EJC_-_Edifact_Java_Compiler URL should be changed to reference section 7.1.4.5. EJC (Edifact Java Compiler) in this document. ---------- Section: 7.1.4.6. EJC Maven Plug-In Comment: Reference to this: http://www.smooks.org/mediawiki/index.php?title=V1.2:Smooks_v1.2_User_Guide#EDI_Mapping_Models URL should be removed - the corresponding information should be added to this document. ---------- Section: 7.1.4.7.1. Using EJC Text: The easiest way in which to commence using EJC is to learn from the example at http:// www.smooks.org/mediawiki/index.php?title=V1.2:ejc. Comment: We should add this example to the guide. The example code is already in the examples .zip file. ---------- Section 9: Rules Comment - We should add a note indicating that in this context, "rules" does not refer to JBoss Rules or Drools. ---------- Section: 9.2.1.1. Useful Regular Expressions Text: A list of useful regular expressions is maintained on the Smooks Wiki at http://www.smooks.org/ mediawiki/index.php?title=Useful_Regular_Expressions. Comment: It's actually a short list. Can we add it to the document. ---------- Section: 10.4. Example10.4. Example A detailed validation example can be see at http://www.smooks.org/mediawiki/index.php? title=V1.2:validation-basic Text: A detailed validation example can be see at http://www.smooks.org/mediawiki/index.php? title=V1.2:validation-basic Comment: We should include the example in the user guide. ---------- Chapter 11: Note on 1st page of the chapter: Text: Be sure to read Chapter 4, Java Binding . Comment: Why? ---------- Chapter 12 - the entire chapter consists of: Message Splitting and Routing Please refer to Section 11.2, "Splitting and Routing". Comment: We should remove this chapter. ---------- Chapter 14 - the entire chapter consists of: Message Enrichment When using the Persistence features of Smooks, the queried data is bound to the bean context (ExecutionContext). You can use the bound query data to enrich your messages e.g. where you are splitting and routing. Refer to Chapter 13, Persistence (Database Reading and Writing) for additional details. Comment: Why not just add this to chapter 13? ---------- Section: 15.2.1. Global Filter Setting Parameters Text: The following global configuration options are available for configuring Smooks filtering. See http://java.sun.com/j2se/1.5.0/docs/api/javax/xml/transform/stream/StreamResult.html and http:// www.milyn.org/javadoc/v1.2/smooks/org/milyn/Smooks.html#filterSource(javax.xml.transform.Source, %20javax.xml.transform.Result...). Comment: Replace the Java 1.5 reference with: http://java.sun.com/javase/6/docs/api/javax/xml/transform/stream/StreamResult.html ---------- Chapter 18: ESB Integration Smooks plugins are available for a number of ESBs: JBoss ESB http://wiki.jboss.org/wiki/Wiki.jsp?page=MessageTransformation Mule http://www.mulesource.org/display/SMOOKS/Home Apache Synapse/WS02 http://esbsite.org/resources.jsp?path=/mediators/upul/Smooks%20Transform%20Mediator Comment: We should remove this section and replace it with a section that covers the JBossESB integration in the SOA Platform. (I'll even help you write it! ;-)	1	1	https://issues.jboss.org/browse/SOA-1575

#	Requirement	Questions	Answered	Successfully Answered	URL
10990	SOA-1575	Assigning back to David - I completed a first pass on the doc - what I have not yet done is to check that all the example code is correct and working. I'd like to look at the other books before coming back to do this. Some global comments on the Smooks User Guide: * References to Smooks project assets (javadocs). For references to pages in Smooks javadocs, we should be able to reference files installed as part of the SOA-P 5.0 distribution, as the javadocs should be included in the distribution. We'll need to add a separate JIRA/task to add the javadocs to the distribution. They are not there in ER2. * References to Smooks project assets (other information). I haven't kept a count, but there are many instances where the user guide directs the user to a smooks wiki page for additional information. It's disruptive to have so many of these, but we may not have a choice, given the large number of wiki pages used by smooks. * Addition of a glossary - What is the technical challenge that blocks the addition of a glossary? Can we simply create a glossary out of the the most important 10-20 terms? * Best practices information for selecting transformation methods - We need to fulfill the SOA-P 5.0 PRD requirement to provide users with guidance in selecting transformation methods. * Example code from projects that we incorporate into platforms - In BRMS 5.0.1, Drools 5.0 added an expanded set of examples. The examples were referenced in the Drools user guide, but the examples.zip file was not shipped as part of the BRMS distribution. There is a similar situation with Smooks 1.2 as the user guide references example programs that are not included in the SOA-P 5.0 distribution. Recommendation: Add the Smooks examples to the SOA-P 5.0 distribution. http://downloads.sourceforge.net/smooks/smooks-examples-1.2.2.zip We will have to verify that the examples can be built and run. * Removing references to project team members - Some of the code examples in the book include the actual name of a smooks developer. This is acceptable for the open source project, but for the Platforms, we don't want to imply that customers should contact developers directly for support. * References to the "user guide" that are also present in the on-line documentation should be removed from the user guide document. (The document IS the user guide as far as customers should be concerned.) * Graphic images are fuzzy/unclear.	1	1	https://issues.jboss.org/browse/SOA-1575
10991	SOA-1575	Done: 1. Content Handler definition added. 2. Definition of acronym "ETL" moved to first occurrence. 3. "Simple Example" replaced with that for XML-to-XML transformations. 4. Definition of ExecutionContext added. 5. Reworded section to avoid use of word "bundle." 6. Question "(Should DOM or SAX be Used?)" removed from header. 7. Broken URL removed and cartridge list added to book. 8. Smooks Resource definition added. 9. Added mention of new XSL Transformation function in SOA-P 5. 10. Removed name of Smooks developer from example code and replaced with made-up names. 11. Clarified that "rules" does not refer to JBoss Rules. 12. References to Online User Guide removed. Added link to section in our own document. 13. Updated reference to Java 1.6 14. Link to section on EJC added. 15. EJC example added to document. 16. Regular expression list added to document. 17. Validation example added. 18. Java Binding chapter reference removed. 19. Splitting and Routing chapter removed. 20. Message Enrichment chapter merged. (May be broken out again in the future if upstream expands this section.) 21. Reference to Java replaced with reference to 1.6 Not done: 1. Glossary cannot be added due to technical limitations/specific business process rule. 2. Images have not been altered, as this a PDF rendering issue only. 3. Addition of a pre-Smooks 1.1 example. (Please point me in the direction of an appropriate one.) 4. Len to rewrite ESB integration chapter, as mentioned above.	1	1	https://issues.jboss.org/browse/SOA-1575
10992	SOA-1575	Hey Len... when you say "... the example uses FreeMarker - which we don't support in SOA-P 5...", what do we mean? Only XSLT is supported? If so, there are issues with this because XSLT is only supported via the Smooks DOM Filter. If this is what we mean, then I don't think it's a good idea.	0	0	https://issues.jboss.org/browse/SOA-1575
10993	SOA-1574	Should now be visible on Documentation-Stage.	0	0	https://issues.jboss.org/browse/SOA-1574
10994	SOA-1558	It is my understanding that we now support import/include for the EBWS feature (in.xsd, out.xsd and error.xsd). Is that true? I will create a change request.	0	0	https://issues.jboss.org/browse/SOA-1558
10995	SOA-1533	"pipeline" of message processors. These process the message before routing the result to the "replyTo" end-point. The action processing that takes place in the pipeline may consist of steps in which the message is transformed by one processor, then some business logic is applied by the next and so on, before the result is routed to eitehr the next step in the pipeline, or to another end-point. Comment: Misspelling: eitehr ---------- Section: 4.1.2. updated Routers Text: Tools like StaticWiretap and StaticRouter can... Comment: StaticWiretap is printed in bold, StaticRouter in italics ---------- Section: 4.3. Updated What is a Service? Text: The JBoss Enterprise Service Bus does not impose restrictions with regard to what should constitute a service. Comment: Section 3.1 stated: '...In the JBoss Enterprise Service Bus, a service is defined as "a list of action classes that process a Message in a sequential manner." How about changing the opening line in 4.3 to: In the JBoss Enterprise Service Bus, a service is defined as "a list of action classes that process a Message in a sequential manner." Within this definition, the JBoss Enterprise Service Bus does not impose restrictions with regard to what should constitute a service. ---------- Section: 6.1.5. Updated. Configuring for a Remote Service Invoker Text: log4j-1.2.14.jar stax-1.2.0.jar scout-1.2.0.aop.jar commons-logging-1.1.jar jboss-aop-jdk50-1.5.6.GA.jar javassist-3.6.0.GA.jar juddi-core-3.0.0.jar juddi-ws-3.0.0.jar (no such file in 5.0 ER5) xercesImpl-2.8.0.jar Comment: In ER5, the files are named: log4j.jar stax-ex.jar scout-1.2.aop.jar commons-logging.jar jboss-aop-jdk50.jar javassist.jar juddi-core-3.0.0.aop.jar xercesImpl.jar ---------- Section 9.1. Overview Text: http://anonsvn.labs.jboss.com/labs/jbossesb/trunk/product/etc/schemas/xml/jbossesb-1.0.1.xsd Comment: 404 not found error. The correct link is: http://anonsvn.jboss.org/repos/labs/labs/jbossesb/trunk/product/etc/schemas/xml/jbossesb-1.2.0.xsd ---------- Section Table 9.6, Table 9.7 Comment: Why are these pages in landscape mode? ---------- 11.1.1. Updated. Transformers & Converters Comment: Formatting of columns is corrupted for XsltAction 11.1.1. SmooksAction Comment: The note refers to smooks 1.1 - we are shipping 1.2 in SOA-P 5.0 ---------- Section: 11.3. Updated. Scripting 9 Text: 1. "script": Path (on classpath) to Groovy script. Properties Comment: There is no footnote #9. Text: ScriptingAction 10 Executes a script using the Bean Scripting Framework (BSF), receiving the message, payloadProxy, action configuration and logger as variable input. Some notes: 11 12 1. JBoss ESB 4.7 contains BSF 2.3.0, which has less language support than BSF 2.4.0 (for 13 14 15 example: no Groovy , and non-functioning Rhino). A future version will contain BSF 2.4.0, 16 17 which will support Groovy and Rhino . 18 2. BSF does not provide an API to pre-compile, cache and reuse scripts. Because of this, each execution of the ScriptingAction will go through the compile step again. Please keep this in mind while evaluating your performance requirements. 19 3. When including BeanShell scripts in your application, it is advised to use a .beanshell 20 extension instead of .bsh, otherwise the JBoss BSHDeployer might pick it up. Comment: The page only inlcudes footnote #10. ---------- Section 11.5. Updated. Routing Subsection: ContentBasedRouter Text: Regex is configured in exactly the same way as XPath. The only difference being that it uses Regex expressions (as opposed to XPath ones.) JBoss Rules: <action process="split" name="ContentBasedRouter" class="org.jboss.soa.esb.actions.ContentBasedRouter"> <property name="cbrAlias" value="Drools"/> <property name="ruleSet" value="MyESBRules-XPath.drl"/> <property name="ruleLanguage" value="XPathLanguage.dsl"/> <property name="ruleReload" value="true"/> <property name="destinations"> <route-to destination-name="express" service-category="ExpressShipping" service-name="ExpressShippingService"/> <route-to destination-name="normal" service-category="NormalShipping" service-	1	1	https://issues.jboss.org/browse/SOA-1533

#	Requirement	Questions	Answered	Successfully Answered	URL
10996	SOA-1533	The edits look fine - except for the footnotes - what's the problem blocking these?	1	1	https://issues.jboss.org/browse/SOA-1533
10997	SOA-1533	See highlights in attachment programmers_2.jpg for explicit reference to Rosetta code. Can we remove this paragraph?	1	1	https://issues.jboss.org/browse/SOA-1533
10998	SOA-1533	will contain the name of the original file from which the message was sourced. org.jboss.soa.esb.gateway.file The file from which the Message was sourced. This will only be present if this gateway is file- based. Comment: It looks like these file descriptions are for the same file. ========== Section: 4.3.1. The "ServiceInvoker" Text: From a client's perspective, the Courier interface and its various implementations can be used to interact with services. Comment: From Kevin: '...The Courier interface is a lower level, transport specific, interface. Not sure why it is being referenced in the docs though as people should really be using the service invoker...' So - can we remove references to the Courier interface? ========== Section: 4.3.4. Updated. The InVM Transport Text (in a note): One may not be able to achieve all of the ACID (Atomicity, Consistency, Isolation, Durability) semantics, particularly when it is used in conjunction with other transactional resources such as databases. This is because of the volatility of the InVM queue. However, the performance benefits of InVM will outweigh this downside in the majority of cases. In situations in which full ACID semantics are required, Red hat recommend that you use one of the other transactional transports, such as the Java Message Service or a database. Comment: This should read "Red Hat recommends" ========== Section: Examples 4.8 and 4.9 Text: The following example illustrates the declaration of a service which, (without being exposed through a web service end-point), "wishes" to validate the request/response messages: Example 4.8. Enabling the GLOBAL InVM Scope for a Service Example 4.9. Enabling GLOBAL InVM Scope for a Service Comment: The exampple titles refer to GLOBAL InVM Scope, not the service declaration. ========== Section: 5.3. Content-Based Routing Text: Sometimes it is necessary for the Enterprise Service Bus to dynamically route messages to their sources. Some scenarios in which this might be the case are when the original destination is no longer available, the service has moved or the application simply wants to have more control over where messages go based on either content, time-of-day or other such factors. The Content-Based Routing mechanism within the JBoss ESB can be used to route messages based on arbitrarily complex rules. These rules can be defined in XPath or JBoss Rules notation. Section: Can we also add the new Regex CBR feature to this section. ========== Section: 7.1. Fail-Over and Load-Balancing Support Text: It is important have redundancy in mind when one is designing mission-critical systems. The JBoss Enterprise Service Bus includes built-in fail-over, load balancing and delayed message re-delivery, all of which will help one build a robust architecture. When one uses a SOA, it is implied that Comment: "When one uses an SOA, it is implied that" ========== Section: 7.1.3. Protocol Clustering Text: Some Java Message Service providers can be clustered. JBossMessaging is one of these, which is why it was chosen as the JMS provider for the Enterprise Service Bus. When one clusters this JMS, one remove a single point of failure from one's architecture (see the next diagram.) Comment: "When one clusters this JMS, one removes a single point" ========== Section: Figure 7.3. Example of Protocol Clustering Using JMS Text: One should read the documentation on "Clustering for JBossMessaging," if one wants to enable this functionality. Both JBoss ESB replication and JMS clustering can be used together, as shown in the following illustratin. Comment: "illustation" ========== Section: 7.1.3. Protocol Clustering Text (in a note) If one is using JMS clustering in this way, one will obviously need to ensure that the configuration is correct. For instance, if one place all of one's ESB services within a JMS cluster, there could be no expected benefit from replication. Comment: "For instance, if one places all" ========== Section: 7.2. Scheduling of Services Text: JBoss ESB 5.0 supports two types of providers: 1. Bus Providers, which supply messages to action processing pipelines via messaging protocols such as JMS and HTTP. This type is	1	1	https://issues.jboss.org/browse/SOA-1533
10999	SOA-1532	Organization of the 1st chapter is awkward - the content at the start of the chapter is at 1 n.n.n.n level - can we make this one level higher? 1.1. What is the Registry? 1.1.1. Introduction The JBoss SOA-P Registry provides applications and businesses with a central point at which information about services can be stored. The Registry is expected to provide both the same level of information and the same breadth of services as does a conventional "marketplace." Ideally, a registry should also facilitate the automatic discovery and execution of electronic commerce by providing a dynamic environment for business transactions. Therefore, a registry is more than a mere "e. business directory." It is a fundamental component of the Service-Oriented Architecture's infrastructure. 1.1.1.1. Why Does One Need It?	0	0	https://issues.jboss.org/browse/SOA-1532
11000	SOA-1532	Section 3.3. Content-Based Routing Using Smooks Comment: We need to add a note that says: There is no native support for Freemarker inside the ESB (no action, no gateway, etc), any use of Freemarker must be in the context of Smooks. Comment: All the graphics in this section are unreadable. ---------- Section: 2.2.2. Rule Set Creation Text: One can create a rule set by using Red Hat Developer Studio. Comment: Should be "JBoss Developer Studio" ---------- Section: 2.2.1. Introduction Text: This the ESB Message, which is going to be inserted into the Rules Engine's working memory. ---------- Section: 3.2.4.1. XPath and Namespaces Text: The namespace-aware statements differ from the others in that it is a requirement thatthe extra expr Comment: "thatthe" ---------- ---------- Section: Figure 3.4. Message Splitting Text: ??? shows how to perform the by-order-item splitting operation Comment: Why the question marks?	0	0	https://issues.jboss.org/browse/SOA-1532
11001	SOA-1532	Section: 3.3. Content-Based Routing Using Smooks Text: This release of the JBoss Service-Oriented Architecture Platform includes Smooks v1.1. This means that the split described above can be undertaken without defining the binding configurations that were demonstrated in Resources One and Two.) For details on how to do this, one needs to refer to the section entitled "FreeMarker Transforms using Node Models" in the Smooks User Guide . That document can be found at http://docs.codehaus.org/display/MILYN/Smooks+User+Guide . Comment: (Thanks to Tom for supplying this): The Node Model Templating functionality (templating without having to define a Java Model) is available in Smooks v1.1+. Sounds like the "split" config talked about in the comment is using a fully defined Java object model, which is fine, and the comment is just trying to tell the user that they also have this other option of using the Node Model functionality, hence doing away with the need to define a Java Object model. I think I'd just remove the reference to Smooks versions in the first comment and say.... "The split described above can be undertaken without defining the binding configurations that were demonstrated in Resources One and Two.) For details on how to do this see the section on <insert link to section on Node Models>" That sound OK? T.	0	0	https://issues.jboss.org/browse/SOA-1532

#	Requirement	Questions	Answered	Successfully Answered	URL
11002	SOA-1531	http://www.jboss.org/community/wiki/ConfiguringMultipleJBossInstancesOnOnemachine ---------- Figure 1.2. Clustering Scenario Two Comment: I think we need to add some explanatory text to this figure. The difference between the 2 diagrams is that he ESB-aware gateways are either clustered, or duplicated. When would we recommend one configuration over the other? I'll try to get some details that we can add to the guide. ---------- Section 1.3. JBossESB Java Message Service Providers Text: Red Hat currently support the use of a number Comment: "supports" ---------- Text: How Can I Configure Them? You can configure JMSListeners and JMSGateways to listen to either "Queues" or "Topics" by setting the following parameters in the Service Configuration file: Comment: We should explicitly refer to the "jboss-esb.xml Service Configuration file." (FYI guys - see: https://jira.jboss.org/jira/browse/SOA-534 - this is a confusing part of the book.) ---------- Section: 1.3.1. JBoss Messaging Text: One should set the parameters for JBoss Messaging to: jndi-URL="localhost" jndi-context-factory="org.jnp.interfaces.NamingContextFactory" connection-factory="ConnectionFactory" destination-type="queue" destination-name="queue/A" Comment: We should replace the reference to "queue/A" with another queue name - maybe "queue/myqueue" - the "A," "B," "C," and "D" queues that we preconfigured in SOA-P 4.2 and 4.3 are not in 5.0. ---------- Section: 1.3.2. Apache ActiveMQ Text: Warning Apache ActiveMQ has not been fully tested & is not a supported JMS implementation. Comment: If it's not supported then we should not document it. Can we remove this section? Text: In your classpath you should have: • activemq-core-4.x • backport-util-concurrent-2.1.jar Both of these JARs can be found in the lib/ext/jms/activemq file. Comment: We don't ship these files as part of SOA-P. ---------- ---------- Section: 1.3.7. Oracle Advanced Queuing (AQ) Text: Oracle AQ has not been fully tested and it is not a supported JMS implementation. Comment: If it's not supported then we should not document it. Can we remove this section? Text: An example jboss-esb.xml configuration file that one may find useful can be found in the samples/quickstarts/helloworld_action/oracle-aq directory. Comment: We don't ship this quickstart with SOA-P. ---------- Section: 1.3.11. Updated Database Configuration Comment: The updated ant output is listed below - the list of supported database as described in the document should match this output: [ldimaggi@ldimaggi schema]$ ant Buildfile: build.xml check.profile: [echo] JBoss SOA Platform Database Configuration Script [echo] ---------- ---- [echo] [echo] This script is used to configure the SOA platform and all its [echo] constituent components against a new RDBMS. Deployment scripts [echo] are currently available for the following databases: [echo] DB2, MSSQL, MySQL, Oracle, PostgreSQL [echo] [echo] ** Warning ** [echo] This script may not work correctly if you have made manual [echo] changes to the database configuration. [echo] ask.profile: [input] Which server configuration are you changing? (production, all, ...): [default] exists.profile: [echo] Using directory /jboss/local/50_ER4/jboss-soa-p.5.0.0/jboss-as/server/default check.clustered: check.backup: do.backup: [echo] Backing up files ... [copy] Copying 1 file to /jboss/local/50_ER4/jboss-soa-p.5.0.0/jboss-as/server/default/deploy [copy] Copying 1 file to /jboss/local/50_ER4/jboss-soa-p.5.0.0/jboss-as/server/default/deploy/jbossesb-registry.sar/META-INF [copy] Copying 1 file to /jboss/local/50_ER4/jboss-soa-p.5.0.0/jboss-as/server/default/deploy/jbossesb-registry.sar [copy] Copying 1 file to /jboss/local/50_ER4/jboss-soa-p.5.0.0/jboss-as/server/default/deploy/jbossesb.esb [copy] Copying 1 file to /jboss/local/50_ER4/jboss-soa-p.5.0.0/jboss-as/server/default/deploy/jbpm.esb [copy] Copying 1 file to /jboss/local/50_ER4/jboss-soa-p.5.0.0/jboss-as/server/default/deploy/jbpm.esb execute: [echo] Databases configured by this script are: [echo] DB2 9.7, MS SQL 2005, MySQL 5.1, Oracle 10g, PostgresSQL 8.3.7 [echo] Please enter the database you wish to configure: [input] [db2-97, mssql2005, mysql51, oracle10g,	1	1	https://issues.jboss.org/browse/SOA-1531
11003	SOA-1531	US/JBoss_Enterprise_SOA_Platform/5.0.0/html/ESB_Administration_Guide/monitoring_and_management.html Figures 6.1, 6.2 and 6.3 1. could be a little bit bigger 2. refer to AS 4 6.1.1. Services Each ESB service is displayed along with the processing time per action, processed count per action, failed count per action and overall message count (per service), as can be witnessed in the following screen-shot: 1. the following screenshot is not aligned with the text above 2. There are no metrics like "processing time per action, processed count per action, failed count per action". The Message Counter keeps track of the number of both successful and failed messages, along with their processing times and the dates. "processing times" should be "processed bytes" Figure 6.3. MBean Counter Maybe "Message Counter" would be better? 6.1.3. Smooks Transformations There is an MBean that keeps track of the count of "processed" Smooks Transformations and the time taken to process each. It also keeps track of the overall count of the Transformation Chain. Could this be explained in more detail? Where to find that MBean? ========== section 6.2 Alerts http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_SOA_Platform/5.0.0/html/ESB_Administration_Guide/alerts.html Chapter is based up on web-console which is no longer shipped AFAICT - but this needs to be clarified. Maybe this or "6.3.1. Message Alerts" would be a good place to document SOA-1656. ========== http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_SOA_Platform/5.0.0/html/ESB_Administration_Guide/JON_for_SOA.html section 6.3 in the eponymously-named quick start, "messagealert." 1. dot should be after " 2. the QS name is plural, messagealerts As one can deduce from this screenshot, the counters can be reset through the JMX console: 1. The scrnshot is too small to be able to deduce it 2. It's not the *counter* that can be *reset* but *alerts* that can be *cleared* Figure 6.4. Clustering Scenario One wrong title Sample code of MyJMSListenerAction.java not properly formatted at the end and following imported classes are never used: javax.management.MalformedObjectNameException org.jboss.mx.util.MBeanProxyExt org.jboss.soa.esb.actions.AbstractActionPipelineProcessor org.jboss.soa.esb.actions.ActionProcessingException I think, here would be a good place to document SOA-1660 Adding your JBoss SOA Platform Server to the JON Inventory missing numbering (6.2) Conveniently, the error message also contains a short-cut link to the "Connection Properties" page. Looks like there is no error message anymore, but the SOA-P is marked as DOWN (see down.png attached) Figure 6.5 is refferring to SOA-P 4.3 instead of 5.0 JBoss SOA-P Enterprise Service Bus Statistics missing numbering (6.3) Once one's JBoss SOA-P server is correctly configured for the JBoss Operations Network tool, one should find the item JBoss ESB Statistics now available underneath the menu entry in Resources on the MONITOR tab. "available underneath the menu entry in Resources" - not sure this is true, attaching screen "underneath.png" for clarification.. I *think* this should read "available above the Resource menu entry" The metrics collected and displayed vary, depending on the Enterprise Service Bus component in question. The available metrics include: this needs to be revisited, I can help with this setting up the environment, just ping me; attaching screens, might be helpful (count1.png - count7.png) Figure 6.6. Metrics Displayed 1. Displayed should not start with a capital D 2. figure could be bigger Managing deployed ESB archives missing numbering (6.4) A new Enterprise Service Bus archive can be deployed by selecting JBoss ESB Deployment from the Create New menu. Then, from the Create New Resource page, specify which archive to deploy, to where it should be sent (which, under normal circumstance, would be one's deploy directory) and, finally, whether it should be sent as a compressed or extracted archive. "whether it should be sent as a compressed or extracted archive" - you can only upload *zipped* archives (files, not directories) - the "Deploy Zipped" option tells	1	1	https://issues.jboss.org/browse/SOA-1531

#	Requirement	Questions	Answered	Successfully Answered	URL
11004	SOA-1531	I have addressed most of the issues raised by yourself and Jaroslaw above. Here are a few that I have not done, with my comments in CAPITALS: 6.1.3. Smooks Transformations There is an MBean that keeps track of the count of "processed" Smooks Transformations and the time taken to process each. It also keeps track of the overall count of the Transformation Chain. Could this be explained in more detail? Where to find that MBean? NEED INFORMATION FROM SUBJECT MATTER EXPERT The metrics collected and displayed vary, depending on the Enterprise Service Bus component in question. The available metrics include: this needs to be revisited, I can help with this setting up the environment, just ping me; attaching screens, might be helpful (count1.png - count7.png) NOT QUITE SURE WHAT THE ISSUE IS THAT JAROSLAW IS TRYING TO RAISE. ARE THE METRICS OUT OF DATE OR INACCURATE? Maybe this or "6.3.1. Message Alerts" would be a good place to document SOA-1656. SOA 1656 IS ASSIGNED TO DARRIN BUT I WILL LIAISE WITH HIM ABOUT THIS.	1	1	https://issues.jboss.org/browse/SOA-1531
11005	SOA-1531	Note sent to Tom Fennelly: Smooks question for you on the ESB/SOA-P Admin Guide. Section 6.1.3 refers to: 6.1.3. Smooks Transformations "...There is an MBean that keeps track of the count of "processed" Smooks Transformations and the time taken to process each. It also keeps track of the overall count of the Transformation Chain..." But, this is the only smooks MBean that I can find - but - it does not keep track of the transformations: http://localhost:8080/jmx-console/HtmlAdaptor?action=inspectMBean&name=jboss.esb%3Adeployment%3Dsmooks.esb Does this MBean still exist?	1	1	https://issues.jboss.org/browse/SOA-1531
11006	SOA-1531	Please remove the references to the smooks MBean: <ldimaggi_> tcunning, will try a smooks sample <tcunning> just tried one <tcunning> it looks to me like that bean is never registered <tcunning> so i think you're okay removing the section i added <ldimaggi_> tcunning, Was it ever registered? I'm guessing no - thanks!	1	1	https://issues.jboss.org/browse/SOA-1531
11007	SOA-1531	Final comments on this version of the guide: Section: 1.4. JBossESB Java Message Service Providers Text: Important In the sections that follow, these assumptions are made: • your JMS provider runs on 'localhost' • the connection factory is called 'ConnectionFactory' • the destination-type is a 'queue' • the destination-name is 'queue/A' Comment: We changed the examples to not all use queue/A, remember? This line should be removed. ---------- 1.4.1. JBoss Messaging Text: Instructions for installing JBoss Messaging can be found on the project website: http://labs.jboss.com/jbossmessaging/docs/userguide-1.4.0.GA/html/installation.html Comment: We should remove the reference to installing JBoss messaging. SOA-P customer will already have it installed. ---------- Section: 1.4.6. Tibco Enterprise Message Service (EMS) Text: One should set the parameters for the Tibco Enterprise Message Service to the following: jndi-URL="tcp://localhost:7222" jndi-context-factory="com.tibco.tibjms.naming.TibjmsInitialContextFactory" connection-factory="QueueConnectionFactory" destination-type="queue" destination-name="<queue-name>" Comment: We should name the queue "myqueue" to be consistent with the JBoss messaging example. ---------- Section: 1.4.9. Updated Database Configuration Text (in the note): As long as there is script availabe for one's database, the Enterprise Service Bus will auto-create the schemas on startup. (By default, the JBoss Enterprise Service Bus is configured to use a JEE data source.) Comment: "availabe" Comment: The ant output references the ER4 build - should be GA. ---------- Section: 1.4.10. Switching Databases Comment: The example refers to a version of Postgres (7.2) that we do not support. ---------- Section: 1.4.11. Using a JSR-170 Message Store Text: The Java Content Repository implementation included is called Apache Jackrabbit. To enable the Java Content Repository's message store, add the following property to the "core" section of the jbossesb-properties.xml file, (which can be found in the root of the jboss-esb.sar or the root of deployers/esb.deployer on AS5): Comment: On a SOA-P server, the file is found in: ./deployers/esb.deployer/jbossesb-properties.xml Also - we should remove the reference to AS5. ---------- Section: 1.4.13. Clustering and Fail-Over Support Text: Because the ServiceInvoker hides much of the fail-over complexity from users, it only works with "native" SOA Platform messages. Furthermore, not all gateways have been modified to use the ServiceInvoker, so incoming messages that are unaware of the SOA Platform may not always be able to take advantage of service fail-over. The Release Notes document contains more details related to this topic. Comment: Is this still described in the rel notes? ----------	1	1	https://issues.jboss.org/browse/SOA-1531
11008	SOA-1530	Chapter 3 should be removed, it was new content in the upstream but mostly related to downloading and setting up the ESB 3.3.2 - might need to be added as a known issue, SOA-1584 3.3.3 - should go into a "migrating applications" section somewhere 3.3.4 - might belong towards the start of the programmer's guide	0	0	https://issues.jboss.org/browse/SOA-1530
11009	SOA-1530	some typos - use of "an" instead of "a" 2.4 - ...do they represent an minimum requirement... 4.3 - ...to Section 5.1, "Launching an Zip installation". 5.1 - Launching an Zip installation 4.1.2 - Copy the development console WAR archive from the tools/resources directory to the *deploy* directory of the server configuration. It should be *deploy/jbpm.esb*. It's the same for both standalone and eap-embedded version. 4.1.2 - Not sure about this one but in the sentence above there's a difference for standalone and eap-embeddet version: for standalone is - to the deploy directory of the *default* server configuration for eap-embedded just - to the deploy directory of the server configuration 4.1.3 - This value (DownloadServerClasses) is configured to false by default in "EAP Embedded" edition. Actually it's false only for "production" config, for all others it's true by default. Maybe it should be mentioned there. 7 - In the very first paragraph there is: "They can be used Insert definition of what the quickstarts are and what they are for." It looks that someone forgot to substitute this for the actual definition :-)	0	0	https://issues.jboss.org/browse/SOA-1530
11010	SOA-1489	until further notice documenting the package files to be: If you are using JON 2.1.2 : rhq-jbossesb-soa-plugin-SOA.4.3.0.GA_CP02.zip If you are using JON 2.3 : jon-plugin-pack-soa-2.3.0.GA.zip checked in	0	0	https://issues.jboss.org/browse/SOA-1489

#	Requirement	Questions	Answered	Successfully Answered	URL
11011	SOA-1478	1 - this procedure is how to configure JBossESB to use the SOA Software Service Manager as it's UDDI registry instead of jUDDI ? ------> That is correct. 2 - from the document: "This text assumes you have a running JBoss Application/ESB Server, and the latest 4.4 CP GA." This corresponds to SOA Platform 4.3 CP01 ? ------> Yes, but this opens up a good question - to date, we have only deployed and tested with SOA Software running on 4.3 SOA-P servers (and being accessed by SOA-P 4.3 and 5.0 servers). Do we want to also support (and test and certify) deploying to 5.0 servers? Will this work? 3 - from the document: "For JBoss ESB 4.4.GA, the Scout jar must be updated to scout-1.0rc2.aop.jar " This would be SOA Platform 4.3 GA ? Where do they obtain the updated Scout jar from ? And which file do they have to replace with it ? what path ? Or do we only want to support this on SOA 4.3 CP01 or later? ------> In SOA 4.3 CP04, the scout file included is 1.0 - In 5.0, the file is version 1.2 - I think both are OK. 4 - Apache Axis is required because SOAP is the only transport we can use for this integration, or the only one we support/test ? ------> I think that it's possible to do something in memory - this is a Tom Cunningham question. 5 - The document describes the progress of setting up "the guest user in Workbench to the guest role". However it doesn't describe why this is required, nor can I find reference to this user in the rest of the document. Why is it needed ? ------> That's a good point - the docs say that this is used for anonymous access - but - don't we have all access go thru the admin account? Also a Tom Cunningham question.	0	0	https://issues.jboss.org/browse/SOA-1478
11012	SOA-1478	------> In SOA 4.3 CP04, the scout file included is 1.0 - In 5.0, the file is version 1.2 - I think both are OK. Do we only want to support this in 4.3.CP04 or later then ?	0	0	https://issues.jboss.org/browse/SOA-1478
11013	SOA-1473	known issues: 1 - re-worded to: "When upgrading from the JBoss Enterprise SOA Platform 4.3.GA release" 2 - removed	0	0	https://issues.jboss.org/browse/SOA-1473
11014	SOA-1473	Thanks for the feedback, I've updated the mentioned items with the exception of the "FIXED" JIRAs which I will deal with separately. Could run your eye over those items & verify they are OK ?	0	0	https://issues.jboss.org/browse/SOA-1473
11015	SOA-1470	It is not obvious where this should be documented. Has this been added to any of the upstream ESB docs ? new "max-total-connections" property, default value 2 Should this be added to the properties for Apache Commons HttpRouter at http://www.redhat.com/docs/en-US/JBoss_SOA_Platform/4.3.CP01/html-single/Programmers_Guide/index.html#section-HttpRouter and added to the properties for SOAPClient at http://www.redhat.com/docs/en-US/JBoss_SOA_Platform/4.3.CP01/html-single/Programmers_Guide/index.html#OOTB-Actions_webservices-soap	0	0	https://issues.jboss.org/browse/SOA-1470
11016	SOA-1457	This was fixed for 4.3.0 CP02 with: http://fisheye.jboss.org/changelog/JBossESB?cs=26923 which changed the ExistsSql statements in: jbossesb.esb/jbossesb-service.xml jbpm.esb/jbpm-service.xml to use: select count(*) instead of: select *	0	0	https://issues.jboss.org/browse/SOA-1457
11017	SOA-1450	"needs at least 8k page size table space to create tables" doesn't parse too well is "requires a tablespace with at least an 8 kilobyte pagesize" valid ? Also I'm assuming kilobyte , not kilobit ?	1	1	https://issues.jboss.org/browse/SOA-1450
11018	SOA-1450	Hi Jim, I'm not familar with DB2 & wanted to fill in some additional details 1 - can you supply the SQL command/s that will be required to set this?	0	0	https://issues.jboss.org/browse/SOA-1450
11019	SOA-1444	I cannot read this doc - see attachments - can you check on this? Thanks!	1	1	https://issues.jboss.org/browse/SOA-1444
11020	SOA-1444	The doc looks good - when will it go public?	1	1	https://issues.jboss.org/browse/SOA-1444
11021	SOA-1440	This was fixed as part of the 4.3.CP05 release. If there is ever another 4.2 series release, we can reopen this JIRA.	1	1	https://issues.jboss.org/browse/SOA-1440
11022	SOA-1419	usage: build.sh --help This script is to facilitate soa integration with the following function: 1) check necessary code from svn 2) build and do some verification 3) do smoking test on products 4) email notification Usage: /home/rzhang/utilities/build.sh -url soa_build_url -email [username@xxxx] basedir Parameter: basedir The soa build base direcotry, if not provided, . will be used Options: -url yoururl Specify the svn url of soa build code -email youremail Specify the email which accept the notifications -smokedir yoursmoketestdir Specify the smoke test code direcotry -update Need svn update operation, usually in the case already checkouted the code -smoketestonly Smoke test only in the case already with latest build -username yoursvnname Input your svn username, otherwise build will prompt to let you input later -password password Input your svn password, otherwise build will prompt to let you input later -help Display this help text and exit --help Display this help text and exit This script needs Internet connections. example, to build and test SOA-4.3.0.GA_CP02_IR1 Run this: build.sh -url https://svn.devel.redhat.com/repos/jboss-soa/build-team-tags/SOA-4.3.0.GA_CP02_IR1/build-tools -email rzhang@redhat.com -username ryanzhang -password xxx A email will be reported at rzhang@redhat.com, like the fowllowing(hopefully): Check out build-tools from https://svn.devel.redhat.com/repos/jboss-soa/build-team-tags/SOA-4.3.0.GA_CP02_IR1/build-tools .....done! ant build .....done! ant release.zip .....done! Check out smoketest code from https://svn.devel.redhat.com/repos/jboss-soa/trunk/soa-ci/tests/smoketest .....done! Check the existence of soa version zip file .....done! Unzip soa version zip .....done! Check smoketest direcotry existence .....done! Do smoking test for soa version jboss-soa-p.4.3.0, config all .....done! Do smoking test for soa version jboss-soa-p.4.3.0, config production .....done! Do smoking test for soa version jboss-soa-p.4.3.0, config default .....done! Remove /home/rzhang/lab/soa-4.3.0-CP02-IR1/build-tools/release/jboss-soa-p.4.3.0 .....done! Check the existence of standalone version zip file .....done! Unzip standalone version zip .....done! Check smoketest direcotry existence .....done! Do smoking test for standalone version jboss-soa-p.4.3.0, config default .....done! Remove /home/rzhang/lab/soa-4.3.0-CP02-IR1/build-tools/release/jboss-soa-p.4.3.0 .....done! Job done! Build on: /home/rzhang/lab/soa-4.3.0-CP02-IR1/smoketest Linux jawa05.englab.brq.redhat.com 2.6.18-92.1.10.el5xen #1 SMP Wed Jul 23 04:11:52 EDT 2008 x86_64 x86_64 x86_64 GNU/Linux This jira is fixed at r3385 on https://svn.devel.redhat.com/repos/jboss-soa/trunk/utilities/build.sh	1	1	https://issues.jboss.org/browse/SOA-1419
11023	SOA-1407	Hello David I believe that Section 11.9 of the SOA 5.0 Programmers' Guide now contains all of the information that you require. Do you have a copy of the WebService_Proxy_v3.odt document, to which Darrin refers in the comment above? Does it contain any information that I need to incorporate into the Programmers' Guide? Many thanks, David Le Sage	1	1	https://issues.jboss.org/browse/SOA-1407
11024	SOA-1407	Do we have any information that matches Burr's request ?	0	0	https://issues.jboss.org/browse/SOA-1407
11025	SOA-1407	What is not resolved apart from the "compare & contrast" content that Burr requested ?	0	0	https://issues.jboss.org/browse/SOA-1407
11026	SOA-1406	How about configuring JBossCache with <attribute name="CacheMode">LOCAL</attribute> for 'default' configuration of the SOA-P (not only for 'all' and 'production')? EhCache may be ok. HashtableCacheProvider would cause problem even in non-cluster environment as its cache entries don't expire.	0	0	https://issues.jboss.org/browse/SOA-1406
11027	SOA-1382	is there any content from the wikipage (version 21) that is not correct or not appropriate for the SOA documentation? or is there anything specific missing that should be included? Also the wikipage has two TODO's at the bottom with the images.	0	0	https://issues.jboss.org/browse/SOA-1382

#	Requirement	Questions	Answered	Successfully Answered	URL
11028	SOA-1382	Is this also content that should be added to the JBoss ESB Programmers Guide ?	0	0	https://issues.jboss.org/browse/SOA-1382
11029	SOA-1382	Who are you waiting for information from?	1	1	https://issues.jboss.org/browse/SOA-1382
11030	SOA-1365	Changing status to resolved - since the doc is ready for review - Darrin - does this have the final bug lists?	0	0	https://issues.jboss.org/browse/SOA-1365
11031	SOA-1365	Final changes verified in attached doc. Will close the JIRA when the final docs .zip is created.	1	1	https://issues.jboss.org/browse/SOA-1365
11032	SOA-1363	How about we use following paragraph to the 'best practices guide': " In the combination of Seam, jBPM and ESB. it is recommended that users are using seam to interact with jBPM in a loosely coupled way, which is through the ESB service gateway, not the jBPM integration annotation directly. In this way, we can see the ESB as the session facade, which provides the entries to all of other service, not only jBPM process, but also other services, like web service, ESB service etc. Now you might feel that seam and jBPM integration is very simple and handy, but it is not a best practice since this is a component-to-component invocation, it is fine that we've just had two components, but if we have more than two components, especially we are having the ESB into play, it is not recommended to do so. In the SOA world, as the ESB term indicates, it is the bus, other components (here like seam, jBPM) should talk to the ESB, not the end-to-end way, which will lead to the spaghetti issue if we have many components, and most of them are talking to component directly. In the future release, we will introduce a set of annotation for seam and ESB integration, it will make users a lot easier to do the seam and ESB integration, just as simple as what we have for the jBPM integration. Before we've had these annotations, it is recommended to do the integration pragmatically that uses the gateway as channel to talk to ESB, and then invoke the target service. (like jBPM process service etc). "	0	0	https://issues.jboss.org/browse/SOA-1363
11033	SOA-1363	I need two things if this is going to make it into the docs: 1 - Is this content agreed upon as appropriate/correct ? 2 - Where should it be included in the docs?	0	0	https://issues.jboss.org/browse/SOA-1363
11034	SOA-1363	I need some idea of what is happening with this JIRA so I can schedule docs work for 5.0.1 What volume of material? e.g. how many pages, diagrams, etc When will draft material be ready ?	0	0	https://issues.jboss.org/browse/SOA-1363
11035	SOA-1357	Hi Len , This issue is initially created for JBPM DB2 support. Could you clarify if we still need this for JBPM ? I found JUDDI also needs the large page size table space to create tables , should we also move the SOA-1450 to 4.3 CP03 ? Thanks Jim	1	0	https://issues.jboss.org/browse/SOA-1357
11036	SOA-1357	Does jBPM require 32k pages or not now? If yes then the requirement should be merged with SOA-1450 and the documentation to be updated properly.	0	0	https://issues.jboss.org/browse/SOA-1357
11037	SOA-1353	Yes and yes! I'll see about getting the page updated... ----- Original Message ----- From: "Burr Sutter" <bsutter@redhat.com> To: "Len DiMaggio" <ldimaggi@redhat.com> Sent: Monday, June 29, 2009 3:05:36 PM GMT -05:00 US/Canada Eastern Subject: clarification We list HP JDK on this page http://www.jboss.com/products/platforms/soa/testedconfigurations/ What OS did we test this against? HP-UX? If so, shouldn't that be added to the list of supported OSs?	1	1	https://issues.jboss.org/browse/SOA-1353
11038	SOA-1351	why the pages 12,14 and 16 are completely empty? is the standard for end of a section? why is not that way between Chapter 1. and Chapter 2 for example?	0	0	https://issues.jboss.org/browse/SOA-1351
11039	SOA-1351	Hey Darrin - I just looked here: https://svn.devel.redhat.com/repos/jboss-soa/branches/4.2.0/build-tools/docs/SOA_Getting_Started_Guide.pdf And the out of date information is still there - can you update it ASAP - this should be the last edit needed for any of the 4.2 CP04 docs. • At this time there are no migration scripts included for PostgreSQL databases. • The ESB Management Console is a tech preview at this time.	0	0	https://issues.jboss.org/browse/SOA-1351
11040	SOA-1350	Seeing failure with WS tests - is OpenJDK more strict on flagging deprecations and on configuring exception handling? http://hudson.qa.jboss.com/hudson/view/SOA-Release/job/temp-soa-ws-OpenJDK16-single/1/console	1	1	https://issues.jboss.org/browse/SOA-1350
11041	SOA-1348	Final changes verified here: https://svn.devel.redhat.com/repos/jboss-soa/branches/4.2.0/build-tools/docs/esb/SOA_ESB_Message_Action_Guide.pdf Will close the JIRA when the final docs .zip is created.	1	1	https://issues.jboss.org/browse/SOA-1348
11042	SOA-1346	Only some minor changes needed: 1.4. SOA Components 1 An SOA is a specific type of distributed system in which the agents are "services" . The key components of a Service Oriented Architecture are the messages that are exchanged, agents that act as service requesters and service providers, and shared transport mechanisms that allow the flow of messages. A description of a service that exists within an SOA is essentially just a description of the message exchange patter between itself and its users. Within an SOA there are thus three critical roles: requester, provider, and broker. ------> pattern misspelled as "pattern" 3.2. Embedded All ESB components (with components we really mean JVMs in this case) can embed the registry and they all can connect to the same database (or different once if that makes sense). ------> ones misspelled as "once" 3.3. RMI using the juddi.war or jbossesb.sar The JBoss SOA includes juddi.war in the ${SOA_ROOT}/server/production/deploy/ jbossesb.sar/ directory. When deployed this brings up the regular web services but also an RMI service. You also need to deploy a datasource which points to your jUDDI database. An example file is supplied for MySQL. ------> We should remove the reference to the MySQL example file and indicate that the schema tool handles configuring the uddi database. 5.2. More Information • For more information see the wiki http://labs.jboss.com/wiki/JudyEvaluation ------> The URL should be changed to http://www.jboss.org/community/wiki/JudyEvaluation	0	0	https://issues.jboss.org/browse/SOA-1346
11043	SOA-1346	Final changes verified here: https://svn.devel.redhat.com/repos/jboss-soa/branches/4.2.0/build-tools/docs/esb/services/SOA_ESB_Registry_Guide.pdf Will close the JIRA when the final docs .zip is created.	1	1	https://issues.jboss.org/browse/SOA-1346
11044	SOA-1344	Point of reference - in the 4.3 CP01 jBPM doc: The section on enterprise archive is not included in the book. The section on document management is labeled as: Important - This is still an experimental feature. If the jBPM Version 3.2.5.SP5 bits are the same - we should keep the docs the same too. Can the 4.2 CP04 jBPM book mirror the features listed in the same guide for 4.3 CP01?	0	0	https://issues.jboss.org/browse/SOA-1344
11045	SOA-1344	Final changes verified here: https://svn.devel.redhat.com/repos/jboss-soa/branches/4.2.0/build-tools/docs/esb/JBPM_Users_Guide.pdf Will close the JIRA when the final docs .zip is created.	1	1	https://issues.jboss.org/browse/SOA-1344
11046	SOA-1328	The issue is not QE ready IMHO. I'd expect a how to document that we should test. We shouldn't be figuring out how to create the required configuration. Kevin, could you please have a look on it?	0	0	https://issues.jboss.org/browse/SOA-1328
11047	SOA-1299	added to release notes as a known issue as per the 4.3.CP01 issue	0	0	https://issues.jboss.org/browse/SOA-1299
11048	SOA-1297	What is the status?	1	0	https://issues.jboss.org/browse/SOA-1297
11049	SOA-1275	We never officially supported JCA/Mail integration in the project because the mail resource adapter was buggy. Has this been fixed? Do we now support the adapter?	0	0	https://issues.jboss.org/browse/SOA-1275

#	Requirement	Questions	Answered	Successfully Answered	URL
11050	SOA-1275	It is my understanding that there is "patch" for the JCA adapter and that it is supported at the product-level. I am not sure if that patch founds its way back into the jboss.org AS stream. With that said, do we agree that the HTTP & Socket areas need better/more complete docs at least?	1	0	https://issues.jboss.org/browse/SOA-1275
11051	SOA-1264	Do we know if this issue occurs in 5.0? We will need to document it in the Release Notes if so. Thanks.	1	1	https://issues.jboss.org/browse/SOA-1264
11052	SOA-1250	Do we have the correct JSF console? Manifest-Version: 1.0 Archiver-Version: Plexus Archiver Created-By: Apache Maven Built-By: aguizar Build-Jdk: 1.6.0_0 Implementation-Title: JBoss jBPM3 - 4JSF Implementation-Version: 3.2.6.SP2 Implementation-Vendor-Id: org.jbpm.jbpm3 Implementation-Vendor: JBoss, a division of Red Hat	0	0	https://issues.jboss.org/browse/SOA-1250
11053	SOA-1246	What do you expect to see apart from JAAS integration??	1	0	https://issues.jboss.org/browse/SOA-1246
11054	SOA-1246	It seems that the purpose of this issue is blurred. What exactly means identity component in SOA-P? Just support for JAAS? Integration between jBPM and PicketLink IDM? I suggest splitting the issue into two or more so we can track what is done and what is not <jpechane> kconner, The subject says SOA-P but description is talking about jBPM and there is a link to unresloved jBPM issue <jpechane> kconner, So my feeling was that the issue was about jBPM using Identity (probably IDM component) <kconner> jpechane: Okay, so was it a mistake to assign it to me? <kconner> jpechane: Identity should just be another JAAS module, not sure what else would need to be done for jBPM <jpechane> kconner, No it was intentional to generate an interest :-) I am really not sure about what the issue is <kconner> jpechane: lol, that makes two of us in that case <jpechane> kconner, Maybe the issue has to be split in two - integration with JAAS which is there and integration with jBPM which is not. IMHO the description is talking about replacing jBPM Idenity componenet with picktelink's IDM <kconner> jpechane: Then I am definitely the wrong assignee :) <tkirby> jpechane, or ask jeff what his intention was when he raised it <jpechane> kconner, I'll comment on the issue in this way and will assign it back to Jeff <kconner> jpechane: Okay, thanks	0	0	https://issues.jboss.org/browse/SOA-1246
11055	SOA-1224	It looks like SOA-1217 has been pushed back to 4.3.CP02 I imagine we want this documented as a known issue. Could you supply some more concrete details please, I'm not familiar with jUDDI. Thanks.	0	0	https://issues.jboss.org/browse/SOA-1224
11056	SOA-1224	hang on, is this just the removal of juddi-service.sar ?	0	0	https://issues.jboss.org/browse/SOA-1224
11057	SOA-1196	I think that the relevant issue here is for 4.3 GA customers that built their DBs manually. They will need to be told about the DB schema changes in CP01 - how will these customers get these changes into their DBs? Will they invoke the DB upgrade scripts manually? Re-opening and assigning to Darrin.	0	0	https://issues.jboss.org/browse/SOA-1196
11058	SOA-1192	The last comment there Mark.... they were JIRAs purely for testing the quickstarts for the last Community GA. What's the relevance to "New Quickstarts"?	1	0	https://issues.jboss.org/browse/SOA-1192
11059	SOA-1190	Also, can we move the stacktrace to a new issue??	0	0	https://issues.jboss.org/browse/SOA-1190
11060	SOA-1155	Closing - it's in the 5.0 Admin guide: How Can I Configure Them? Set the following parameters in the jboss-esb.xml service configuration file in order to configure JMSListeners and JMSGateways to listen to either queues or topics: * jndi-URL * jndi-context-factory * jndi-pkg-prefix * connection-factory * destination-type * destination-name	1	0	https://issues.jboss.org/browse/SOA-1155
11061	SOA-1084	Hi Darrin, Could you please help me add the above Description into the 'Known Issues in this release' chapter of 'Release Notes' book. Thanks Jeff	1	1	https://issues.jboss.org/browse/SOA-1084
11062	SOA-1069	Darrin - I cannot find this reference in the docs - can you confirm that we have this listed in one of the docs? The Getting Started Guide or the Services Guide?	1	1	https://issues.jboss.org/browse/SOA-1069
11063	SOA-1068	Hi Darrin, Could you please help me to add following section into the '1.8 Configuring OpenSSO in SOA-P' of 'JBoss_Enterprise_Platform_Administration_Guide' book. ========== 1.8.2 Install and configure Opensso in SOA-P Due to an opensso.war deployment issue (https://jira.jboss.org/jira/browse/SOA-731), we have provided a workaround to fix it. We are testing against OpenSSOv1 Build 4.5 (http://download.java.net/general/opensso/stable/openssov1-build4.5/) 1. Go to $SOA-P/jboss-as/samples/quickstarts/opensso/opensso-ear, run 'ant'. 2. Copy the opensso-1.0.ear from $opensso-ear/build to $SOA-P/jboss-as/server/$config(such as default)/deploy 3. Start the soa-p from $soa-p/jboss-as/bin/run.sh 4. Open http://localhost:8080/opensso in a browser. 5. Click on "Create Default Configuration". This causes OpenSSO to configure itself. 6. Enter 'adminpass' for the "Default User[amAdmin]" and 'ldappass' for "Default Agent [amldapuser]" 7. Visit http://localhost:8080/opensso again. Log in using the proper credentials. User Name is amAdmin and Password is the password you chose to go with amAdmin. More information about opensso, please visit its website at: http://opensso.dev.java.net ==========	1	1	https://issues.jboss.org/browse/SOA-1068
11064	SOA-1068	Hi Darrin, Could you please help me to include the above message in the 'Administration_Guide' book? Thanks Jeff	1	1	https://issues.jboss.org/browse/SOA-1068
11065	SOA-1068	verifying tomcat installation: the instructions say build 4.5 but it seems to be no longer available on the site. Can we give people more specific instructions on the correct version's location ?	0	0	https://issues.jboss.org/browse/SOA-1068
11066	SOA-1068	Darrin - are the doc changes available for review? Thx.	0	0	https://issues.jboss.org/browse/SOA-1068
11067	SOA-1061	Jiri Pechanec has pointed out a number of faults in this doc. * "3.2 Figure 3.1 and Figure 3.2 - each describes the same but have different structure (missing pieces in the figures)" Diagrams are derived from the ESB Programmers Guide & I believe 3.1 is supposed to be a much generalized image. Guidance please. * "3.2.4 The name for any given named object in a message must be unique or an exception will be thrown. - is it really true? How can this happen?" Verify content in ESB Programmers Guide * "3.2.5 Creating different body types - there is org.jboss.soa.esb.message.body.content.Payload class for this purpose List of different message bodies - there are duplicate items" Verify content in ESB Programmers Guide * "4.1 Routers are in the section name but are not mentioned in the text - also need to explain difference between routers and notifiers Would be also good to mention that it is not necessary to provide the information to listener - the reference to provider for given transport is typically used Update content in ESB Programmers Guide * "5 This chapter should contain also Service orchestration section" Update content in ESB Programmers Guide * "page 73 - Note and page 34 - Note: The latter says that the gateway EPRs are not stored in registry and the former says they are" Verify/update content in ESB Programmers Guide * "Chapter 10 - the is new style functionality related to WS - EBWS (ESB services exposed as Web Services)" Update content in ESB Programmers Guide * "Appendix A - there is nothing in the text that justifies the inclusion of this Appendix - for what the JAXB is used in ESB context?" Verify/update content in ESB Programmers Guide	1	1	https://issues.jboss.org/browse/SOA-1061

#	Requirement	Questions	Answered	Successfully Answered	URL
11068	SOA-1061	>These are the review comments for the Guide for the actual content from the ESB guide, ie its the same in both, and I need someone to verify either that it is correct or provide from info on what it >should be Ok, I'll go through the ProgrammerGuide in the ESB project and take a look at the issues mentioned. Do I need the SOA version of this document to find references to pages mentioned in this Jira? Thanks	1	1	https://issues.jboss.org/browse/SOA-1061
11069	SOA-1061	Hey Darrin, I've looked through the issues here and compared them with the ESB projects documentation. I've added comments to each one below and this should hopefully be enough to create jiras for those issues that require a change to the ESB project docs. * "3.2 Figure 3.1 and Figure 3.2 - each describes the same but have different structure (missing pieces in the figures)" >Diagrams are derived from the ESB Programmers Guide & I believe 3.1 is supposed to be a much generalized image. Guidance please. This what Figure 3.1 looks like in the ESB Projects ProgrammersGuide: <?xml version = "1.0" encoding = "UTF-8"?> <jbossesb xmlns="http://anonsvn.labs.jboss.com/labs/jbossesb/trunk/product/etc/schemas/xml/jbossesb-1.0.1.xsd"> <services> <service category="Retail" name="ShoeStore" description="Acme Shoe Store Service"> <actions> <action name="println" class="org.jboss.soa.esb.actions.SystemPrintln" /> </actions> </service> </services> </jbossesb> The "missing pieces would in this case be the provider configuration and the listener configurations, unless the invm transport is assumed. The two diagrams do not describe the same thing, unless I'm missing something. I don't think that action should be taken here. * "3.2.4 The name for any given named object in a message must be unique or an exception will be thrown. - is it really true? How can this happen?" No, this is not true (the report is correct that is). This is simply a Hashtable and adding a new named object will replace the old one. This is also present in the ESB Programmers Guide on page 28 which needs to be updated * "3.2.5 Creating different body types - there is org.jboss.soa.esb.message.body.content.Payload class for this purpose >List of different message bodies - there are duplicate items" There are no duplicates in the ESB projects programmers Guide. See page 24. * "4.1 Routers are in the section name but are not mentioned in the text - also need to explain difference between routers and notifiers. >Would be also good to mention that it is not necessary to provide the information to listener - the reference to provider for given transport is typically used Routers are not mentioned in the ESB Programmers Guide, not in this section. Update required. * "5 This chapter should contain also Service orchestration section" >Update content in ESB Programmers Guide The "Service orchestration section" is located in the ESB projects document "ServicesGuide.pdf". It could be lifted into the SOA Platform document. * "page 73 - Note and page 34 - Note: The latter says that the gateway EPRs are not stored in registry and the former says they are" Page 41 of ESB ProgrammersGuide: Note: It is possible that multiple EPRs may be present in the Registry with the same Service Name/Category and that some of these EPRs may not be ESB-aware, e.g., Gateways. If the ServiceInvoker receives ESB-unaware EPRs from the Registry then it will ignore them. You may see the warning: "Invalid EPR for service (probably ESB-unaware)". Page 77 of ESB ProgrammersGuide: Note: For each gateway listener defined in a service, an ESB aware listener (or "native") listener must also be defined as gateway listeners do not define bidirectional endpoints, but rather "startpoints" into the ESB. From within the ESB you cannot send a message to a Gateway. Also, note that since a gateway is not an endpoint, it does not have an Endpoint Reference (EPR) persisted in the registry. This will need to be updated. * "Chapter 10 - the is new style functionality related to WS - EBWS (ESB services exposed as Web Services)" This needs to be added to the ESB projecs ProgrammersGuide. This is no currently in the main trunk but will be shortly. * "Appendix A - there is nothing in the text that justifies the inclusion of this Appendix - for what the JAXB is used in ESB context?" The appendix is for "JAXB Introductions" (http://www.jboss.org/community/docs/DOC-10075) and not JAXB isself.	1	1	https://issues.jboss.org/browse/SOA-1061
11070	SOA-1061	Where do you think it should belong? Maybe we can just add a reference to it there?	0	0	https://issues.jboss.org/browse/SOA-1061
11071	SOA-1051	A couple of final comments - it looks like EAP is larger than before so we should update that section. 1.2. Hardware Requirement Recommendations Disk space requirements - for 4.2 CP03 CR7 standalone (inc. docs) = ~195MB (I would say 200 in the guide) embedded (inc docs) = ~497MB (I would say 500 in the guide) 1.7. Installing the SOA Platform Do we ever tell the user to also install the doc's .zip files? I can't see it anywhere else. We should mention it here.	1	1	https://issues.jboss.org/browse/SOA-1051
11072	SOA-1049	2.5.4.2. Insertion When an Object is insert it returns a FactHandle. ------> Comment "inserted" New instance assertions will only return a new FactHandle if a not equal classes have been asserted. ------> Comment "if a not equal classes condition" 2.5.4.3. Retraction Note that it is possible to have rules that depend on the "non existence" of a fact, in which case retracting a fact may cause a rule to activate (see the 'not' and 'exist' keywords). ------> Comment "(see the 'not' - section 5.2.2.5 and 'exist' - section 5.2.2.6 - keywords) 2.5.4.6. Shadow Facts There's generally two ways to guarantee truthfulness. ------> Comment "There are two ways to guarantee truthfulness." 2.5.4.8. Initial Fact To support conditional elements like "not" (which will be covered later on), ------> Comment "(which will be covered later on in chapter 5)" 2.5.6. Stateless Session The execute method iterates the objects inserting each and calling fireAllRules() at the end; session finished. ------> Comment "and calling fireAllRules() at the end when the session is finished." "StatelessSessions also support sequential mode, which is a special optimized mode that uses less memory and executes faster; please see the Sequential section for more details." ------> Comment "please see the Sequential section (section 2.5.10) for more details." ReferenceOriginalGlobalExporter just passes a reference to the original Global Resolver; the latter should be used with care as identifier instances can be changed at any time by the StatelessSession and the GlobalResolver may not be serializable freindly. ------> Comment - just drop the misspelled word "freindly" 2.5.7.2. Agenda Groups Agenda groups are a way to partition rules (activations, actually) on the agenda. At any one time, only one group has "focus" which means that the activations for rules in that group will only take effect - you can also have rules "auto focus" which means the focus for its agenda group is taken when that rules conditions are true. ------> Comment "when those rules' conditions are true" Agenda groups are a handy way to create a "flow" between grouped rules. ------> Comment: The previous paragraph indicated that a flow between rules cannot be relied on. Are we saying that the flow between groups is a way around that limitation? 2.5.10. Sequential Mode 1. Order the Rules by salience and position in the ruleset (just sets a sequence attribute on the rule terminal node). 4 ------> Comment: Delete the "4" 3.1.4.1. Syntax of templates The "$param" place holder is used in templates to indicate where data form the cell will be interpolated ------> Comment "from" For consequences: How snippets are rendered also depends on if there is anything in the row immediately above it. If there is nothing there, the output is simple the interpolated snippets ------> Comment: Change "simple" to "simply" 4.1. Introduction ------> This section refers to "Red Hat Developer Studio." Is this name being used at all? We only refer to it as JBDS. Note you can get the plug in either as a zip to download, or from an update site (refer to the chapter on installation). ------> There is no "installation" chapter in this book. Figures 4.3, 4.4 - these include a very old Rules icon - probably from Rules 3 in 2006 (figure 4.6 shows a date of 2006) - when can we get these updated? If not in 4.2 CP03 then in 4.3 FP01? 5.5.2.1. Pattern The subsections in this^ section are numbered: 5.5.2.1.1. Field Constraints 5.5.2.1.1.1. JavaBeans as facts 5.5.2.1.1.3.1. Operators It's hard to read and reference chapter sections with 7-level-deep section numbers - why not move section 5.5.2.1 "Pattern" to the same level as 5.5 "Rule" - 4-levels is about as deep as most people can follow.	0	0	https://issues.jboss.org/browse/SOA-1049

#	Requirement	Questions	Answered	Successfully Answered	URL
11073	SOA-1049	6.1. Deployment options If you are using the BRMS, you should use the RuleAgent (the detailed documentation for this is in the chapter on the BRMS). ------> comment: There is no chapter on the BRMS in this doc. 6.1.1. Deployment using the RuleAgent The easiest and most automated way to deploy rules is to use the RuleAgent. This is described in detail in the BRMS user guide. ------> comment: Do we have a BRMS user guide? Index BeanShell, 12 ------> Comment: The only reference to BeanShell is a question on page 12. C Collection, 108, 108 ------> Comment: The term is used on pages 1, 17, 42, etc. But not ever really defined in the context of Rules Domain Specific Languages, 128 ------> Comment: It's on pages 79 and 128 DSL, 129 ------> Comment: It's used on pages 128-133 I think we need to remove/replace the entire index. It's very short and includes multiple errors.	0	0	https://issues.jboss.org/browse/SOA-1049
11074	SOA-1049	feedback incorporated with the following exceptions: JBoss Rules Terms - I'll work on this for FP01, see what comes out of the newer docs & port it back in a future CP 2.5.4.2. Insertion New instance assertions will only return a new FactHandle if a not equal classes have been asserted. ------> Comment "if a not equal classes condition" Neither form makes grammatical sense 2.5.7.2. Agenda Groups Agenda groups are a handy way to create a "flow" between grouped rules. ------> Comment: The previous paragraph indicated that a flow between rules cannot be relied on. Are we saying that the flow between groups is a way around that limitation? Need clarfication on intent of content, specifically the meaning of the term 'flow' in this context	0	0	https://issues.jboss.org/browse/SOA-1049
11075	SOA-1049	Nothing major so far - did the doc pick up all the changes from 4.3 CP03? Index Remove the index - it's very incomplete and includes some errors Section 5.5.2. Left Hand Side (when) Conditional Elements below misspelled as bellow Comments from: https://jira.jboss.org/jira/browse/SOA-1049?focusedCommentId=12444597#action_12444597 These are still in the doc 2.5.4.2. Insertion When an Object is insert it returns a FactHandle. ------> Comment "inserted" New instance assertions will only return a new FactHandle if a not equal classes have been asserted. ------> Comment "if a not equal classes condition" 2.5.4.3. Retraction Note that it is possible to have rules that depend on the "non existence" of a fact, in which case retracting a fact may cause a rule to activate (see the 'not' and 'exist' keywords). ------> Comment "(see the 'not' - section 5.2.2.5 and 'exist' - section 5.2.2.6 - keywords) 2.5.4.6. Shadow Facts There's generally two ways to guarantee truthfulness. ------> Comment "There are two ways to guarantee truthfulness." 2.5.4.8. Initial Fact To support conditional elements like "not" (which will be covered later on), ------> Comment "(which will be covered later on in chapter 5)" 2.5.6. Stateless Session The execute method iterates the objects inserting each and calling fireAllRules() at the end; session finished. ------> Comment "and calling fireAllRules() at the end when the session is finished." "StatelessSessions also support sequential mode, which is a special optimized mode that uses less memory and executes faster; please see the Sequential section for more details." ------> Comment "please see the Sequential section (section 2.5.10) for more details." ReferenceOriginalGlobalExporter just passes a reference to the original Global Resolver; the latter should be used with care as identifier instances can be changed at any time by the StatelessSession and the GlobalResolver may not be serializable freindly. ------> Comment - just drop the misspelled word "freindly" 2.5.7.2. Agenda Groups Agenda groups are a way to partition rules (activations, actually) on the agenda. At any one time, only one group has "focus" which means that the activations for rules in that group will only take effect - you can also have rules "auto focus" which means the focus for its agenda group is taken when that rules conditions are true. ------> Comment "when those rules' conditions are true" Agenda groups are a handy way to create a "flow" between grouped rules. ------> Comment: The previous paragraph indicated that a flow between rules cannot be relied on. Are we saying that the flow between groups is a way around that limitation? 2.5.10. Sequential Mode 1. Order the Rules by salience and position in the ruleset (just sets a sequence attribute on the rule terminal node). 4 ----> Comment: Delete the "4" 3.1.4.1. Syntax of templates The "$param" place holder is used in templates to indicate where data form the cell will be interpolated ------> Comment "from" For consequences: How snippets are rendered also depends on if there is anything in the row immediately above it. If there is nothing there, the output is simple the interpolated snippets ------> Comment: Change "simple" to "simply" 4.1. Introduction ------> This section refers to "Red Hat Developer Studio." Is this name being used at all? We only refer to it as JBDS.	0	0	https://issues.jboss.org/browse/SOA-1049

#	Requirement	Questions	Answered	Successfully Answered	URL
11076	SOA-1048	the following words use capital letters where appropriate: Java, jBPM, JBoss, jPDL, XML, Hibernate, Eclipse, Petri net, JNDI, JDBC, XA, LDAP, Ant, Xerces... This occurs on so many places that I suggest mass search&replace. 2) Most of the diagrams in the document have a very low picture resolution / compression issues, making them look ugly. I suggest taking a look at the printout. 3) Things like "graph based," "domain specific" or "process oriented" should actually go by "graph-based," "domain-specific" and "process-oriented." Occurrences of this phenomena are too many to mention. 4) Pretty pretty please, get native English proof-reader - I'm sure I missed quite a lot of grammar-related problems. 5) Please mass replace all the "till"s with "until" and also check if "cause" (based on the context) cannot be replaced for "because." 6) The actual spelling of "particular" is with "a" not "e" (as in "perticular"). Another mass replace, please. 7) I think using "web application" instead of "webapp" would help. Another mass replace. 8) Double colons (and question marks) aren't AFAIK usually preceded by spaces. Many occurrences in the document where they are. 9) Starter's kit is sometimes referred to as "starters-kit" - it'd be cool to unify this. 10) Most of "-ful" words like powerful, handful etc. are actually written with just a single "l" at the end. Occurences are quite a few, perhaps RE could be made to fix that. 11) "Whether" isn't "wether," another candidate for mass replace. 12) "more then" can be mass-replaced with "more than." 13) It's not "ommited" or "ommitted," it's just "omitted." Page by page: (P${pageNumber}, ${chapterNumber}) P1, 1: "... process language is buil*t* on top..." P1, 1.2: First sentence contains "Download" 3 times. P3, 1.5: - "EJB will be writen against" - isn't it written already? Or should it be written by the user? - "We have done great efforts ..." - AFAIK efforts aren't done but taken. P3, 1.6: "... is richer th*a*n the traditional..." (This problem actually occurs in quite some places.) P7, 3: - "The examples can also be fo*u*nd..." - "download jbpm-3.0-[version].zip" - please highlight the file name based on the document conventions, including highlighting the "[version]" in italics. (repeats on P144, 18.1.2.) - Last paragraph is actually an unfinished sentence that appears out of place. P7, 3.1: "The hello world process" - it would be good to highlight "hello world" to indicate that it is some identifier. P9, 3.2: "E.g. a*n* piece" - the "n" shouldn't be there, should it? P20, 4.1.1: - "The IDE is buil*t* around..." - another very common mistake throughout the text. - "When a user saves*,* the compiled class is saved*.*" - May I suggest better use of interpunction? - "Where ten years ago..." - The sentence doesn't make much sense to me by itself. Either remove the "where" or join it with the other sentences in the paragraph somehow. - "... is a common foundation for all these type*s* of domain*-*specific..." - "... to edit *J*ava in the *source* file format*, because* ..." P21, 4.1.2.1: - "An imperative programming language*s* like..." - "... implementation technique is only targets..." - please remove the "is" - "... of the order manager and the*n* wait *until*..." - "This could be a sufficient..." - please remove the " a" P22, 4.2.1: "nodes have leaving*-* and arriving..." - please remove the dash. P25, 4.2.3: "The event method will propagate find..." - please choose only one verb. :-) P26, 4.2.3: "... doubleCheck's task node is adds ..." - the same as immediately above :-) P27, 4.2.5: "... an event is sen*t* to the execution." P29, 4.3.2: - "... mixed up with multi-threaded* *programming." - "... in the process execution relate*s* ..." - please remove the "s". P30, 4.3.3: "... An invoke will start *of* a new sub process." - Please remove the "of". P30, 4.3.4: "Messaging systems are also know*n* as..." P33, 4.5.1.1: "Let's define a business process" should IMHO be quoted. P35, 4.5.2: "... for writing *a* new services..." - please remove the "a" P36, 4.6: "... even the BPM engine's database tables *in* integrated into..." - please remove the "in" P37, 4.7.2: - "This is problematic because it results in monolithic system that is very hard to integrate into a project *SOMETHING MISSING* combines plain OOP software with business processes." - "... with the BPM server *to* become..." - please remove "to" - "As show*n* above, ..." P38, 4.7.3: I suggest either filling the chapter with content or removing entirely. P39, 5: "... project, rather th*a*n installing..." P40, 5.3: "Only used in the script's and	0	0	https://issues.jboss.org/browse/SOA-1048
11077	SOA-1048	you make a copy the default configuration..." - doesn't make much sense. - "The third part are some misc*a*llanious configurations..." P50, 6.3: The whole chapter would look much better in a bullet-point list instead of naming every config file in its own subchapter with one sentence in it. P51, 6.3.7: "The id*'*s are stored..." - please remove the apostrophe. P51, 6.5: "Therefor*e*, by default, the ... *the that* *H*ibernate" - please add "e", remove "the" and make Hibernate use capital "H." P55, 7.1.1: "The persistence API *is integrated with*..." P58, 7.1.2: "As a result, a Hibernate session will *only be opened*"... P59, 7.1.3: - "The transaction attribute..." - please end the sentence with a full stop. Even better, put it in some other paragraph, as it's just one sentence. - The following sentence begins with "the the". P60, 7.2.1: - first paragraph: names of the properties should be highlighted. - seconds paragraph: we have an environment for important notes. P62, 7.3: "and use the session per transaction pattern" - please, either put the "session per transaction" in quotes or join it using dashes. P65, 7.6.3: - "The jbpm.db subproject*,* contains..." - please remove - "... scripts for all database*s*..." P65, 7.7: - "This can eliminate*s* ..." - please remove - Secondly, this enable*s* ..." P66, 7.8: The bullet-point list should either state sentences (first letter capital, last letter followed by a full-stop) or the first two bullet points should end with a comma (",") and the last one with a full-stop. P70, 8.1.3: "... only takes a few minutes in *windows*." - I suggest we either use full name and state whose trademark is that, or we stop using our competitors' products in examples... :-) P72, 8.1.4.1: "*To execution* this script..." - doesn't seem right. P74, 8.1.4.2: - "mysql -u root -p" - the username should be in italic as it depends on an environment. Also, once we do this, most of the next paragraph can be left out. - "Once the script executes, you should have the foll*ow*ing..." P76, 8.1.6: "Well, almost..." - I'd follow that with a line break. P77, 8.1.6: "This section should look like shown in the listing below. There are two changes in this file: the hibernate.connection.datasource property should point to the JbpmDS datasource we created as the first step in this section and the hibernate.dialect property should match the PostgreSQL or MySQL dialect." - properties should be highlighted. P82: "You may have to ALT-TAB to get this application as it may be covered by another window." - I suggest we don't treat jBPM users as idiots who don't know anything about window management. :-) And if we do, Alt+Tab is heavily dependent on user's environment, so I'd left that out. P88, 9.3.4: - The first paragraph would be much better when converted into a bullet-point list with two items. - Also, "Should the decision made by the process..." doesn't make much sense. - "The first paragraph for which the *conditions resolves* to 'true'"... you can't have both words ending with "s" in this case. P89, 9.4: "If more th*a*n..." P93, 9.6.2: "... to distinct between..." - is "distinct" a verb? P94, 9.7: s/serached/searched P97, 9.10: s/implementated/implemented P101, 10.1: "The most basic operations are*:*" P102, 10.4.1: s/independant/independent P103, 10.6: "hibernatable type" - really? :-) (Two occurences.) P104, 10.6: "When you refer *to* your custom..." P105, 10.6 type "LongInstance" should be highlighted. P108, 11.2.2: "... one task instance *are* associated..." P110, fig 11.1: Instances' arrows go into the boxes. P110, 11.3.4: "*T*o specify the pooled actors..." P111, 11.4: "The controller can be used to create*,* populate and ..." P111, 11.5: Please remove the a) and b) from the bullet-point list. P113, 11.6: s/assignements/assignments P119, 12: - "IMPORTANT NOTE": We have an environment for that, don't we? - In the same paragraph, the repository method should be highlighted. - Same comment for the bullet-point list as in "P66, 7.8" - "a word document" - either use the full trademark name of Microsoft Word or don't mention the product at all. :-) P121, 13.1: - "A timer that is specified on a node*,* is not executed" - without the comma, please. - Same comment for the bullet-point list as in "P66, 7.8" P123, 14.2: - "... by looking *at* two examples." - "The first frame*,* show the..." - without the comma, please. P124, 14.2: "Now, we are going to look at the second example, the second example is a variant..." - I smell redundancy... P125, 14.2: "2 separate transactions*,	0	0	https://issues.jboss.org/browse/SOA-1048
11078	SOA-1048	fixes applied up to and including chapter 5. and - "(wow, that a lot of seq's there :-s)." - please, remove entirely. - P167, 20.3: Either extend or remove, please. - P148: "Thanks Gavin! Absolutely awesome! :-)" - I suggest removal 1st 13 items applied as well, except 2 which will be worked on for a future version	0	0	https://issues.jboss.org/browse/SOA-1048

#	Requirement	Questions	Answered	Successfully Answered	URL
11079	SOA-1047	What is JBoss Enterprise SOA Platform Edition 1.4 on page 2? Page 3 - not sure if this is correct: It may be updated dynamically when services first start-up, or statically by an external administrator. Page 6 - Is still Systinet in business? http://en.wikipedia.org/wiki/Mercury_Interactive Page 6 - SystinetRegistryImplentation should be SystinetRegistryImplementation Page 7 - Strange formating of tables and table names Page 7 - There is no directory product/install/jUDDI-registry with juddi sql scripts, same for page 10, files are in ${deploydir}/jbossesb.sar/juddi-sql and warfile is in ${deploydir}/jbossesb.sar Page 14 - the description is definitely not complete, to use AxisTransoprt, one must also put axis.jar to classpath - Tom C. has more info Chapter 4 - JBDS also provides UDDI browser - should not we promote it here? Section 5.1 - is it also Product related? It seems to belong only to project	0	0	https://issues.jboss.org/browse/SOA-1047
11080	SOA-1047	1 - the 1.4 refers to the book edition 2 - need confirmation on its behaviour 3 - I don't think this really matters, although it would be nice to have more current examples if that is the case 4 - FIXED: corrected typo 5- FIXED: the 3 tables want to be a single table with settings on one axis & transport type on the other, but the content is way to wide to fit it on a page :-(I changed the formatting a little to try to make it clearer. 6 - Need more information. "The product/install/jUDDI-registry directory contains db create scripts for you favorite database. "jbossesb.sar/juddi-sql doesn't contain database creation scripts, it has import scripts. jboss-as/server/production/deploy/jbossesb.sar/lib/juddi-2.0rc5.aop.jar seems to have creation scripts. Are users expected to unpack the jar to get at the files? 7 - where does axis.jar come from ? I need more information 8 - FIXED: added a comment about JBDS included a UDDI browser. Also updated some of the language as was done for the corresponding content in the 4.3 Services Guide 9 - need confirmation if this should be removed 10 - spacing in address has been fixed in book template	0	0	https://issues.jboss.org/browse/SOA-1047
11081	SOA-1046	Address on title page 1801 Varsity Drive Raleigh, NC 27606-2072USAPhone: +1 919 754 3700 Phone: 888 733 4281 Fax: +1 919 754 3701 PO Box 13588Research Triangle Park, NC 27709USA - see missing spaces What is JBoss Enterprise SOA Platform Edition 1.3 on page 2? The version number is confusing 1.2. When would you use an ESB? The figures below illustrate some concrete examples where JBossESB would be useful. Although these examples are specific to interactions between participants using non-interoperable JMS implementations, - this is confusing as the images displays file-based participants and/or interoperable JMS participants The descrpitons for images are above the images so the page break should not separate the description from the image 4.1 Chapter is wrong there is also PickUpCourier and TwoWayCourier is derived from both courier and pickuponlycourier. Morever TwoWayCourier interface has different definition 6.1.4 The code example contains <<<<<<< .mine // it's out response! ======= // it's out response! >>>>>>> .r12360 break; } 7.2.5 In first example there should be <jms-listener name="helloWorld2" busidref="quickstartEsbChannel2" instead of <jms-listener name="helloWorld2" busidref="quickstartFtpChannel2" in the cecond there are two listeners defined on the same ftpChannel3 Figure 7.5 - cuts page number in half 7.3.1 and 7.3.2 - the examples are swapped Note on page 73 - Also, note that since a gateway is not an endpoints, it does not have an Endpoint Reference (EPR) persisted in the registry. Note on page 34 - It is possible that multiple EPRs may be present in the Registry with the same Service Name/Category and that some of these EPRs may not be ESB-aware, e.g., Gateways. I think these two statements are in contradiction	0	0	https://issues.jboss.org/browse/SOA-1046
11082	SOA-1046	FIXED Address on title page in page template CLARIFICATION: Edition 1.3 on page 2 refers to the version of the book NOT BEING FIXED IN THIS VERSION: 1.2. When would you use an ESB? images MORE INFORMATION REQUIRED FOR CORRECT CONTENT 4.1 Chapter is wrong there is also PickUpCourier and TwoWayCourier is derived from both courier and pickuponlycourier. Morever TwoWayCourier interface has different definition FIXED 6.1.4 The code example FIXED 7.2.5 example2 FIXED Figure 7.5 - cuts page number in half FIXED 7.3.1 and 7.3.2 - the examples are swapped MORE INFORMATION REQUIRED, IS ONE OF THEM OUT OF DATE OR IS ONE MISSING INFORMATION: Note on page 73 - Also, note that since a gateway is not an endpoints, it does not have an Endpoint Reference (EPR) persisted in the registry. Note on page 34 - It is possible that multiple EPRs may be present in the Registry with the same Service Name/Category and that some of these EPRs may not be ESB-aware, e.g., Gateways. I think these two statements are in contradiction	0	0	https://issues.jboss.org/browse/SOA-1046
11083	SOA-1046	FIXED Address on title page in page template - VERIFIED CLARIFICATION: Edition 1.3 on page 2 refers to the version of the book - OK NOT BEING FIXED IN THIS VERSION: 1.2. When would you use an ESB? images - OK MORE INFORMATION REQUIRED FOR CORRECT CONTENT 4.1 Chapter is wrong there is also PickUpCourier and TwoWayCourier is derived from both courier and pickuponlycourier. Morever TwoWayCourier interface has different definition - Please look into ESB source code, more info probably provided by the dev team FIXED 6.1.4 The code example - VERIFIED FIXED 7.2.5 example2 - VERIFIED FIXED Figure 7.5 - cuts page number in half - VERIFIED FIXED 7.3.1 and 7.3.2 - the examples are swapped - VERIFIED MORE INFORMATION REQUIRED, IS ONE OF THEM OUT OF DATE OR IS ONE MISSING INFORMATION: Note on page 73 - Also, note that since a gateway is not an endpoints, it does not have an Endpoint Reference (EPR) persisted in the registry. Note on page 34 - It is possible that multiple EPRs may be present in the Registry with the same Service Name/Category and that some of these EPRs may not be ESB-aware, e.g., Gateways. I think these two statements are in contradiction - According to my best knowledge this statement is correct "It is possible that multiple EPRs may be present in the Registry with the same Service Name/Category and that some of these EPRs may not be ESB-aware, e.g., Gateways." which means that the other one should be removed. But this has to be consulted with the dev team	0	0	https://issues.jboss.org/browse/SOA-1046
11084	SOA-1045	The new screenshots are even uglier than the previous ones. Could you include higher quality screenshots?	0	0	https://issues.jboss.org/browse/SOA-1045
11085	SOA-1045	What pdf viewer are you using, what version and on what operating system? Please supply a screenshot of images in question at 100% image size	0	0	https://issues.jboss.org/browse/SOA-1045
11086	SOA-1045	- The images look much better now. Thanks. Besides that: - 1.1 Overview: stand-alone ESB server -> standalone ESB server What are the percentage signs supposed to mean? % cd tools/console/management-esb% - Fig. 2.3 MessageCounter: Could you scale it down?	0	0	https://issues.jboss.org/browse/SOA-1045
11087	SOA-1045	FIXED: stand-alone -> standalone FIXED: the % are shell prompts, missed a linefeed in copy-paste, changed them to $ & fixed position FIXED: hmmm, scaling it down made it look worse when viewed at greater than 100%, I made it wider so at 100% it scales down automatically & when viewer at greater it is closer to actual 100%. Does that make any sense? ;-) Also cleaned up the language in the overview a little	0	0	https://issues.jboss.org/browse/SOA-1045

#	Requirement	Questions	Answered	Successfully Answered	URL
11088	SOA-1044	On the title page: ---------- JBoss Enterprise SOA Platform 4.2 SOA ESB Message Action Guide JBoss Enterprise SOA Platform Edition 1.3 ---------- What is the reference to 1.3? Note - the doc is missing its revision page. ========== 1.1.2. LongToDateConverter It looks like this actin was deleted: https://cloud.prod.atl2.jboss.com/jira/browse/JBESB-152;jsessionid=BE9600515828DE2DBD08A8BD0C8F39BE?page=com.atlassian.jira.ext.fisheye%3Afisheye -issuepanel It's not in the ESB source or javadocs - I'll check with Mark and Kevin. ========== 1.1.3. ObjectInvoke We should also state which properties are required/optional - from the javadocs: <Action name="doCustomer" processor="ObjectInvoke"> <property name="class-processor" value="CustomerProcessor" /> <!-- Required. Class processor used to process the message . --> <property name="class-method" value="fooMethodName" /> <!-- Optional. Default to name of action. --> </Action> ========== 1.1.4. ObjectToCSVString We should also state which properties are required/optional - from the javadocs: <Action name="Customer-To-CSV" processor="ObjectToCSVString"> <property name="bean-properties" value="name,address1,address2,phone" /> <property name="fail-on-missing-property" value="true" /> <!-- (Optional) Default of false. --> </Action> ========== 1.1.8. XStreamToObject exclude-package is also optional - from the javadocs: <Action name="doCustomer" processor="XStreamObject"> <property name="class-alias" value="Customer" /> <!-- Optional. Class alias used in call to XStream.alias(String, Class) prior to deserialisation. --> <property name="incoming-type" value="CustomerProcessor" /> <!-- Required. Class for incoming type used to process the message after deserialisation. --> <property name="exclude-package" value="false" /> <!-- Optional. Default "true". Not applicable if a "class-alias" is specified. --> <property name="root-node" value="/root/Customer" /> <!-- Optional. Specify an XPath expression be used to determine the root node used with XStream. Useful when the object to convert is not the root node of the document --> <property name="aliases"> <!-- Optional list of extra aliases to add to XStream --> <alias name="aliasName" class="className" /> <alias name="aliasName" class="className" /> ... </property> <property name="converters"> <!-- Optional list of converter to register with XStream --> <converter class="className" /> <converter class="className" /> ... </Action> What about the optional converters? ========== 1.2.1. jBPM - BpmProcessor JBossESB can make calls into jBPM using the BpmProcessor action. Please also read the jBPIntegrationGuide to learn how to call JBossESB from jBPM. Should be: "also read the jBPMIntegrationGuide (SOA_ESB_JBPM_Integration_Guide.pdf)" StartProcessCOmmand Should be: StartProcessCommand These two API commands: setProcessInstanceVariables, getTokenVariables are not documented in the SOA_ESB_JBPM_Integration_Guide.pdf doc - I think we should remove them from this doc too. ========== 1.3.5. Notifier The code example includes "@" characters. These are used to substitute strings from variables and is used in several of the quickstarts. In this example, however, it may cause some user confusion. I'd just remove the "@" characters. ========== 1.3.5.3. NotifySQLTable The SqlGatewayListener also requires a timestamp field - see https://jira.jboss.org/jira/browse/SOA-1111	1	1	https://issues.jboss.org/browse/SOA-1044
11089	SOA-1044	1 - its the version number of the book 2 - FIXED 3 - need confirmation that LongToDateConverter is removed. 4, 5 - Marked all properties with "required"/"optional"/"optional, defaults to " 6 - the 4.2.cp03 (CR5) javadocs don't show the optional converters. 4,3 does. Can you check this, maybe added since CR5 ? 7 - FIXED, also reworded first paragraph slightly to help translation, please confirm is still appropriate. 8 - removed @ 9 - I'm a little confused by SOA-1111 & related JIRAs. Is the timestamp field supposed to be optional, or is that a bug in this release that will be fixed in CP04 ?	0	0	https://issues.jboss.org/browse/SOA-1044
11090	SOA-1043	I've attached both text & pdf versions of the document. Do we also want this doc added to those on redhat.com/docs ?	0	0	https://issues.jboss.org/browse/SOA-1043
11091	SOA-1032	We already have a tuning section. Plus, which version of SOA-P is this affecting?	0	0	https://issues.jboss.org/browse/SOA-1032
11092	SOA-1008	Hi Thomas! I am a bit confused about the status now. What I found in the codebase was a "Deployment" interface and a DeploymenService in "the API". But how is it included in the ESB? Unfortunately there are no fisheye links to code. And somewhere needs to be a JBoss AS Deployer, right? Is it only targeted for jbpm 4? Then it couldn't go into JBoss ESB 4.3 already, right? Can you quickly explain? Thanks a lot for sheding light on this...	0	0	https://issues.jboss.org/browse/SOA-1008
11093	SOA-1008	afaict, there has not been a promise to the soa platform productization to deliver a jboss deployer for par files. this would have to be discussed at lenght before because it is has a lot of implications. synchronizing pars in a deploy directory with deployments in a jbpm database has a lot of situations that can go bad. writing a jboss deployer that finds and deploys a par file is easy. but how to know at server boot time if a process needs to be redeployed ? how to sync in a cluster? what if the deployment is deleted in jbpm?... etc. all those requirements have to be discussed and agreed before we can go ahead with this.	0	0	https://issues.jboss.org/browse/SOA-1008
11094	SOA-972	Moreover, the registry now contains interceptors as well. So this should be included into the example - see jboss-as/server/produciton/deploy/jbossesb.sar/jbossesb-properties.xml. Anything about caching interceptor SHOULD NOT be there since it doesn't work currently.	1	1	https://issues.jboss.org/browse/SOA-972
11095	SOA-972	3.3. RMI using the juddi.war or jbossesb.sar The JBoss SOA includes juddi.war in the jUDDI-registry directory. When deployed this brings up the regular webservices but also an RMI service. You also need to deploy a datasource which points to your jUDDI database. An example file is supplied for MySQL. This is quite strange, there are two juddi.war files in the SOA-P. First in deploy/jbossesb.sar/ and second in deploy/juddi-service.sar. Only the first contains the configuration shown in the examples later in the text. They are definitely not in the jUDDI-registry directory. So it would be good to correct the directory name.	1	1	https://issues.jboss.org/browse/SOA-972
11096	SOA-972	Why it is not possible to click on the links in the document?	1	1	https://issues.jboss.org/browse/SOA-972

#	Requirement	Questions	Answered	Successfully Answered	URL
11097	SOA-972	outstanding items * not sure what the issue is with the image on page 2 * "Moreover, the registry now contains interceptors as well. So this should be included into the example - see jboss-as/server/produciton/deploy/jbossesb.sar/jbossesb-properties.xml. Anything about caching interceptor SHOULD NOT be there since it doesn't work currently." This needs to be fixed in ESB Services Guide source document * "3.3. RMI using the juddi.war or jbossesb.sar The JBoss SOA includes juddi.war in the jUDDI-registry directory. When deployed this brings up the regular webservices but also an RMI service. You also need to deploy a datasource which points to your jUDDI database. An example file is supplied for MySQL. This is quite strange, there are two juddi.war files in the SOA-P. First in deploy/jbossesb.sar/ and second in deploy/juddi-service.sar. Only the first contains the configuration shown in the examples later in the text. They are definitely not in the jUDDI-registry directory. So it would be good to correct the directory name." This needs to be fixed in ESB Services Guide source document * "5.1. Scout and jUDDI pitfalls • Make sure to put our version of the jaxr-api-1.0.jar, scout-0.7rc2-embedded.jar and the juddi-embedded.jar first. The versions of these libraries that are included with the JBoss Application Server are incompatiable. This should get resolved in future release of JBoss Application Server. I did not find any jaxr*jar or juddi-embedded.jar in SOA-P, so what is that thing reffered as "our version"? I only found scout-1.0rc2.aop.jar or scout.jar (version 1.0rc1). " This needs to be fixed in ESB Services Guide source document *"Page 23, another reference to non-existing jUDDI-registry directory: • Make sure to read the README in the jUDDI-registry directory, for instructions regarding prepopulating your own jUDDI database. " This needs to be fixed in ESB Services Guide source document *"Page 39, "3. xpathGreaterThan" The rule is in fact called "xpathLessThan"." I assume you are referring to the "xpathLowerThan" rule *"Page 45, Figure 10.1 is useless because the texts in the image can't be read. It has poor quality." We are constrained by a maximum image width of 430px, this diagram has to be completely re-thought * "Why it is not possible to click on the links in the document?" Are there any particular links you were having problems with? What PDF viewer, what OS? Versions?	0	0	https://issues.jboss.org/browse/SOA-972
11098	SOA-972	libraries that are included with the JBoss Application Server are incompatiable. This should get resolved in future release of JBoss Application Server. (I don't think that the paragraph is at all applicable to 4.3. It should be deleted.) Section 6.1 (Rules Service) (This section of the guide refers to Rule "facts," but does not define this term in the context of JBoss Rules. Why don't we add some brief definitions such as: Fact (application data) Working Memory (stateful collection of facts) Production Memory (the rules) Inference Engine (JBoss Rules & RETE) Agenda (conflict resolution) Salience (priority) REF: http://www.cooug.org/brsig/event_details/presentations/2006_June_20_DRules.pdf 7.1. Introduction (Rules Services) When a message gets send to the BusinessRulesProcessor a rule set executes over the objects in the message and updates either of those objects or the message. ("sent") 7.3. Rule Service Consumers The consumer of a rule service has little to worry about. There is no need for the consumer to creating rulebases or working memories, inserting facts or firing the rules. ("to be creating") 7.4. Configuration A dsl and dslr (domain specific language) files ("dsl and dslr (domain specific language) files") 9.1. Introduction (Content Based Routing) When a message gets send to the CBR, a certain rule set will evaluate the message content and return a set of Service destinations. ("sent") 9.2. Three Different Routing Action Classes In this case the message is simply filter out. ("filtered") Figure 2 shows an example where the MessageType is used to determine to which destination the Message is going to be send. ("sent") 9.4. XPath Domain Specific Language 2. xpathEquals <element>, <value> : yields true if the element is found and it's value equals the value. 3. xpathGreaterThan <element>, <value> : yields true if the element is found and it's value is greater than the value. 4. xpathLessThan <element>, <value> : yields true if the element is found and it's value is lower then the value. (its, not it's) 9.5. Configuration Now that we have seen all the individual pieces how does it all tie together? It basically all comes down to configuration at this point, which is all done in your jboss-esb.xml. (Now that we have seen all the individual pieces, how does it all tie together? It all comes down to configuration at this point, which is all done in your jboss-esb.xml.) Each ESBMessage is passed on to in this case the ContentBasedRouter action class which is loaded with a certain rule set. (Each ESBMessage is passed on to the ContentBasedRouter action class which is loaded with a certain rule set.) It uses the rule set JbossESBRules.drl, which matches two destinations, name 'xml-destination' and 'serializeddestination'. (It uses the rule set JbossESBRules.drl, which matches two destinations, named 'xml-destination' and 'serializeddestination'.) Table 9.2. CBR Action Configuration Properties Optional property which tells the RuleService to use a stateful session where facts will be remembered between invokations. ("invocations") 9.8. RuleAgent We should not have the same URL duplicated in two separate footnotes on the same page. One reference to the URL in a footnote is sufficient. 10.1. Introduction (Content Based Routing with Smooks) An example of this might be a huge order message with thousands/millions of order items per message. (Is this realistic? Millions of order items in one message?) In this case, we're just binding the data into HahMaps. ("HashMaps") Figure 10.1. Message Splitting - The text in this image is not clear enough to be read. Resource configurations #1 and #2 are there to bind data from the source message into Java Object in the Smooks bean context. In this case, we're just binding the data into HahMaps. The Map being populated in configuration #2 is recreated and repopulated for every order item as the message is being filtered. The populated Java Objects (from resources #1 and #2) are use to populate a FreeMarker template (resource #4), which gets applied on every order item, with the result of the templating operation being output to a FileResourceStream (resource #3). The FileResourceStream (resource #3) also gets applied on every order item, managing the file output for the split messages. (Where are these numbers defined? There are no numbers in the figure.) Chapter 11 - Message Transformation The JBoss ESB supports message	1	1	https://issues.jboss.org/browse/SOA-972
11099	SOA-972	12.5.2. ESBActionHandler under normal processing the signal will be send by the JBossESB callback Service. ("sent") 12.5.3. jBPM to ESB Exception Handling Only if the call was made from the EsbActionHandler it makes sense to report back the exception to jBPM. ("does it make sense") 12.5.4. Scenerio 1: Time-out it maybe that you want to limit how long you want to wait ("may be") 13.1. Orchestrating Web Services JBossESB provides WS-BPEL support via its Web Service components. For details on these components and how to configure and use them, see the Message Action Guide. (We don't have a message action guide. In which of the SOA-P docs is this information?) 13.2. Orchestration Diagram (One of these four footnotes is sufficient.) 1 http://www.active-endpoints.com/ 2 http://www.active-endpoints.com/ 3 http://www.active-endpoints.com/ 4 http://www.active-endpoints.com/ 13.3. Process Deployment and Instantiation In the previous paragraph we create the process definition and we quietly assumed we had an instance of it to explain the process flow. But now that we have created the processdefinition.xml, we can deploy it to jBPM using the IDE, ant or the jbpm-console (as explained in Chapter 1). (That's chapter 12. ;-) Chapter 14 - The Message Store In JBossESB 4.2 the Message Store is also used for storing messages that need to be redelivered in the event of failures. (This is release 4.3.)	1	1	https://issues.jboss.org/browse/SOA-972

#	Requirement	Questions	Answered	Successfully Answered	URL
11100	SOA-971	2.2 Rosetta is built on three core architectural components: but there are FOUR items in the list 3.2 Example 4.3, Figure 3.1 and Figure 3.2 - each describes the same but have different structure (missing pieces in the figures) 3.2.1 header property Action shold be explained in more detail replace MessageIDas with MessageIDs 3.2.4 The name for any given named object in a message must be unique or an exception will be thrown. - is it really true? How this can happen? 3.2.5 Creating different body types - there is org.jboss.soa.esb.message.body.content.Payload class for this purpose List of different message bodies - there are duplicate items 4.1 Routers are in the section name but are not mentioned in the text - also need to explain difference between routers and notifiers Would be also good to mention that it is not necessary to provide the information to listener - the reference to provider for given transport is typically used 5 This chapter should contain also Service orchestration section 7.1.1 Note - should refer to 4.4 version, not 4.2.1 B.4 Arethe versions correct? Should not be 4.4? Versions in general, in text there are multiple references to ESB 4.3 GA. I think it should be 4.4 (maybe CP1) page 73 - Note and page 34 - Note: The latter says that the gateway EPRs are not stored in registry and the former says they are page 76 - enumeration - replace <jms-provder> with <jms-provider>, >services> with <services>, busrefid with busidref Chapter 10 - the is new style functionality related to WS - EBWS (ESB services exposed as Web Services) Appendix A - there is nothing in the text that justifies the inclusion of this Appendix - for what the JAXB is used in ESB context?	0	0	https://issues.jboss.org/browse/SOA-971
11101	SOA-971	* "3.2 Figure 3.1 and Figure 3.2 - each describes the same but have different structure (missing pieces in the figures)" Diagrams are derived from the ESB Programmers Guide & I beleive 3.1 is supposed to be a much generalised image. * "3.2.4 The name for any given named object in a message must be unique or an exception will be thrown. - is it really true? How this can happen?" JIRA to verify content in ESB Programmers Guide * "3.2.5 Creating different body types - there is org.jboss.soa.esb.message.body.content.Payload class for this purpose List of different message bodies - there are duplicate items" JIRA to verify content in ESB Programmers Guide * "4.1 Routers are in the section name but are not mentioned in the text - also need to explain difference between routers and notifiers Would be also good to mention that it is not necessary to provide the information to listener - the reference to provider for given transport is typically used JIRA to update content in ESB Programmers Guide * "5 This chapter should contain also Service orchestration section" JIRA to update content in ESB Programmers Guide * "page 73 - Note and page 34 - Note: The latter says that the gateway EPRs are not stored in registry and the former says they are" JIRA to verify/update content in ESB Programmers Guide * "Chapter 10 - the is new style functionality related to WS - EBWS (ESB services exposed as Web Services)" JIRA to update content in ESB Programmers Guide * "Appendix A - there is nothing in the text that justifies the inclusion of this Appendix - for what the JAXB is used in ESB context?" JIRA to verify/update content in ESB Programmers Guide	0	0	https://issues.jboss.org/browse/SOA-971
11102	SOA-970	of features were developed which special focus on language expressiveness, engine performance and tools availability. The following is a list of the most interesting changes. comment ------> "with special focus" ------ 1.1.1. Language Expressiveness Enhancements Fact attributes auto-vivification for return value restrictions and inline-eval constraints comment ------> "auto-verification" comment ------> Style difference between section 1.1.1 - and 1.1.2 - in 1.1.2, (some) feature names begin with Uppercase letters ------ 1.2.1.1. Working Memory creation Drools 3.0.x had only one working memory type that worked like a stateful working memory. Drools 4.0.x introduces separate APIs for Stateful and Stateless working memories that are called now Rule Sessions. comment ------> Drools 3.0.x had only one working memory type that functioned as stateful working memory. Drools 4.0.x introduces separate APIs for Stateful and Stateless working memories. In Drools 4.0.x, stateful working memories are called now Rule Sessions. ------ 1.2.1.2. Working Memory Actions Drools 4.0.x now supports pluggable dialects and has built-in support for Java and MVEL scripting language. In order to avoid keyword conflicts, the working memory actions were renamed as showed bellow: comment ------>1.2.1.2. Working Memory Actions Drools 4.0.x now supports pluggable dialects and has built-in support for Java and the MVEL scripting language. In order to avoid keyword conflicts, the working memory actions were renamed as shown below: ------ 1.2.2. Rule Language Changes The DRL Rule Language also has some backward incompatible changes as detailed bellow. comment ------>The DRL Rule Language also has some backward incompatiblity changes as detailed below: ------ 1.2.2.1. Working Memory Actions The Working Memory actions in rule consequences were also changed in a similar way to the change made in the API. The following table summarizes the change: comment ------> The following table summarizes the changes: ------ 1.2.2.2. Primitive support and unboxing Drools 3.0.x did not had native support for primitive types and consequently, it auto-boxed all primitives in it's respective wrapper classes. That way, any use of a boxed variable binding required a manual unbox. Drools 4.0.x has full support for primitive types and does not wrap values anymore. So, all previous unwrap method calls must be removed from the DRL. comment ------> Drools 3.0.x did not had native support for primitive types, Consequently, it auto-boxed all primitives in their respective wrapper classes. Any use of a boxed variable binding required a manual unbox. Drools 4.0.x has full support for primitive types and does not wrap values. All previous unwrap method calls must be removed from the DRL. ------ 1.2.3. Drools Update Tool The Drools Update tools is a simple program to help with the upgrade of DRL files from Drools 3.0.x to Drools 4.0.x. Currently its main purpose is to upgrade the memory action calls from 3.0.x to 4.0.x, but future versions will conver additional scenarios. It is important to note that it does not perform a simple search and replace in the rules file, but parses the rules file to ensure it is not doing anything unexpected. For this reason t is a safe tool to use for upgrade large sets of rule files. comment ------> Currently its main purpose is to upgrade the memory action calls from 3.0.x to 4.0.x, but future versions will cover additional scenarios. It is important to note that it does not perform a simple search and replace in the rules file, but parses the rules file to ensure it is not doing anything unexpected. For this reason it is a safe tool to use for upgrade large sets of rule files. ------ 1.2.3. Drools Update Tool comment ------> Having a problem building the update tool - I'll pass this on to Mark Proctor mvn clean install [INFO] Scanning for projects... [INFO] ---------- [ERROR] FATAL ERROR [INFO] ---------- [INFO] Failed to resolve artifact. GroupId: org.drools ArtifactId: drools Version: 4.0.0.SNAPSHOT Reason: Unable to download the artifact from any repository org.drools:drools:pom:4.0.0.SNAPSHOT from the specified remote repositories: central (http://repo1.maven.org/maven2) [INFO] ---------- [INFO] Trace org.apache.maven.reactor.MavenExecutionException: Cannot find parent: org.drools:drools for project: null:drools-	0	0	https://issues.jboss.org/browse/SOA-970

#	Requirement	Questions	Answered	Successfully Answered	URL
11103	SOA-969	native speaker review these documents as I'm sure I didn't catch all the spelling errors, typos and stuff. General comments: - Comments in examples are word-wrapped in a way that usually one word overflows to the next line - this leaves tons of unused space on the page and looks unprofessional. If resolved, document gets shorter and better looking. - Also, the color of the comments (along with the background color) makes them very very hard to read. - When talking about specific methods (such as "take" or "execute") it would be cool if these were somehow "annotated" to show this isn't just another text. A Courier font perhaps? - The previous note also applies to class names. The occurences are too many to mention. - When there are comments in XML examples, their opening and closing chars (<>) are not in the usual comment color. - Make sure that, when JBoss is mentioned, it's JBoss and not jboss. The same goes for jBPM, XML etc. - "Justify" text-alignment looks great for print, I suggest we use it. - 3.5.1.1: "Let's define a business process is a description of the way..." (The 'Let's define a business process' should IMHO be quoted.) - 3.7.3: The section is completely void of any information; removal suggested. - 7.1, last listing - completely wrong syntax highlighting. - 8.9: "For an example on how to do this, check out the Fork and Join node implementation in the jbpm source code :-)." (A link would be good, or just remove the smiley to make it look less like we are challenging our customers." - 10.11: "People interested in contributing should check the jboss jbpm jira issue tracker for more details." (URL would be nice, also it's JBoss and jBPM and JIRA.) - 13.6 and 13.7:and 19.3 and perhaps also someplace else: Looks very unprofessional to me to have these "TODO" sections there, some of them don't even provide any information. I suggest re-phrasing or removal. - 15.1.3: "Setting notify to yes, true or on will cause jBPM to send an email to the actor that will be assigned to this task." (The possible values should be separated from the regular text somehow.) - 16.1: "(wow, that a lot of seq's in there :-s)." (Is that really necessary there?) - Index: contains only one item. Is this intentional? Language-wise problems (delimited by asterisks): - 1.6 The JBoss jBPM identity component: "The model used in the identity component is richer th*e*n the traditional servlet-, ejb- and portlet models." - Tutorial: "The examples can also be *fond* in the jBPM download package in the directory src/java.examples." - 3.1.1: "The IDE is buil*d* around the grammar of a language." - 3.1.1: "*Where* ten years ago, the biggest part of a developer was spend on writing code." - 3.1.2.1: "The Graph Oriented Programming implementation technique *is* only targets process languages that are..." - 3.2.3: "Let's assume that the execute implementation of the doubleCheck's task node *is* adds an entry into ..." - 3.2.4: "An Action is also a command*s* with an execute method." - 3.2.4: "All the actions will be executed with the event *fires*." - 3.2.5: "An execution starts when an event is sen*d* to the execution." - 3.2.5: "By default, all of these propagation*'s* are done as nested method calls." (repeated) - 3.2.5: "So the execution can have trave*l*ed over multiple nodes..." - 3.3.2: "Let's go *though* the steps to extend the Graph Oriented Programming model..." - 3.3.3: "An invoke will start *of* a new sub process." - 3.3.4: "The command executor *take* the message from the queue and invokes..." - 3.4.3: "Graph oriented programming on itself does not*'t* support analytical algorithms..." - 3.6: "When the BPM engine can be completely integrated into a software development project and when even the BPM engine's database tables *in* integrated into..." - 3.7.2: "Then other products added and packaged together with the BPM server *to become* a complete BPM suite." - 6.2.4: "This is done to get people up and running as fast as possible without having to worr*ie* about classpaths." - 6.2.4: "Previous incompatibilities between a jboss version and a p*e*rticular ehcache version..." - 6.4: "Make sure that you bind the datasource to an XA datasource in case you're using more th*e*n 1 resource." - 7.1.4: "To execut*ion* this script with DBVisualizer, you select 'Database->Execute'." - 8.3.4: "The first transition for which the conditions resolve*s* to 'true' will be taken." - 8.4: "If more th*e*n one transition has the same name, the first transition..." - 8.8: "The path of execution of the super	0	0	https://issues.jboss.org/browse/SOA-969
11104	SOA-969	Items still to be done: * Dealing with wrapping of source code samples * tagging of methods, classnames * correct capitalisation of product names (I did the few I noticed) "Justify" text-alignment looks great for print, I suggest we use it. I agree but the document style sheets are out of my control. - 16.1: "(wow, that a lot of seq's in there :-s)." (Is that really necessary there?) Not sure what is being referred to here - Index: contains only one item. Is this intentional? I'm not sure that the story is with the index. It isn't something that is manually included. There is an element somewhere that causes the docbook processor to do 'hey I need an index to store this item'.	0	0	https://issues.jboss.org/browse/SOA-969
11105	SOA-969	>> - 16.1: "(wow, that a lot of seq's in there :-s)." (Is that really necessary there?) > Not sure what is being referred to here There is a sentence "(wow, that a lot of seq's in there :-s)." in section 16.1 which, in my opinion, is totally childish and unnecessary. >> "Justify" text-alignment looks great for print, I suggest we use it. > I agree but the document style sheets are out of my control. Would you, please, suggest to the responsible people to include these changes next time? There's also the syntax highlighting issues that I expect are out of your control for the same reason.	0	0	https://issues.jboss.org/browse/SOA-969
11106	SOA-968	More comments: 1.1 http://wiki.jboss.org/wiki/ConfiguringMultipleJBossInstancesOnOneMachine ----> This has moved to: http://www.jboss.org/community/docs/DOC-9384 1.2. JMS Providers The JBoss Enterprise SOA Platform currently supports JBoss Messaging, IBM Websphere MQ Series (version 5.3 and 6.0) and Tibco EMS as Java Message Service (JMS) providers. (We only tested with MQ 6 - where did the 5.3 reference come from?) Each JMSListener and JMSGateway can be configured to use it's own JMS provider so you can use more than one provider in your deployment. (its) 1.2.5. Tibco Enterprise Message Service (EMS) TibcoEMS version 4.4.1 has tested with JBoss SOA (version 5 has been tested too) 1.4. Switching Databases Why don't we mention the schema tool in this section of the document? /jboss-esb/tools/schema <type-mapping>PostgreSQL 7.2</type-mapping> (We only tested with Postgres 8) 1.7. Clustering and Fail-over support Furthermore, in not all gateways have been modified to use the ServiceInvoker, ("not all gateways") 1.8.1. Installing and configuring OpenSSO in Tomcat (Duplicate text) JAVA_OPTS="$JAVA_OPTS "-Xmx1G" "- Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager" JAVA_OPTS="$JAVA_OPTS "-Xmx1G" "- Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager" JAVA_OPTS="$JAVA_OPTS "-Xmx1G" "- Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager" - JAVA_OPTS="$JAVA_OPTS "-Xmx1G" "- Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager" Example 1.1. Adding max size to JAVA_OPTS JAVA_OPTS="$JAVA_OPTS "-Xmx1G" "- Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager" Example 1.2. Adding max size to JAVA_OPTS (We should direct the reader to the opensso quickstart as it also describes the opensso installation process.) Chapter 5 The Jboss application server documentation ("JBoss") Sections 6.1.2 and 6.1.4 are different descriptions of the same information. 6.3 JON for SOA Existing archives can be deleted by checking them in the list of Child Resources and clicking DELETE. (We should also tell the reader that they have to "uninventory" the archive to remove it from the JON management console.) see the groovy QuickStart at http://wiki.jboss.org/wiki/Wiki.jsp?page=JBossESBQuickStart ----> This has moved to: http://www.jboss.org/community/docs/DOC-10363 http://wiki.jboss.org/wiki/Wiki.jsp?page=JBossESBContentBasedRouting ----> This has moved to: http://wiki.jboss.org/wiki/Wiki.jsp?page=JBossESBContentBasedRouting	0	0	https://issues.jboss.org/browse/SOA-968

#	Requirement	Questions	Answered	Successfully Answered	URL
11107	SOA-968	All these items have been updated with the exception of: 1.4. Switching Databases Why don't we mention the schema tool in this section of the document? /jboss-esb/tools/schema <type-mapping>PostgreSQL 7.2</type-mapping> (We only tested with Postgres 8) The schema tool is not not refered to in the source doc (ESB Admin Guide), that is why it's not here. Should we create a JIRA for someone to create that content ? what should the type-mapping be? <type-mapping>PostgreSQL 8</ type-mapping> or <type-mapping>PostgreSQL 8.0</ type-mapping> or something else?	0	0	https://issues.jboss.org/browse/SOA-968
11108	SOA-968	I've added some content in for the schema tool. Do we still want to include the step-by-step instructions as well.	0	0	https://issues.jboss.org/browse/SOA-968
11109	SOA-953	Will we see any functional impact of this change? For example, in the JON console?	0	0	https://issues.jboss.org/browse/SOA-953
11110	SOA-931	SOA-898 looks like it is just a naming problem for the jar (currently SNAPSHOT). Are you still intending to modify the jar version as part of the build? Will it still be in step with the SOA-P version?	1	1	https://issues.jboss.org/browse/SOA-931
11111	SOA-918	I'm confused. Am I looking at the wrong thing? The only references I can find (in IR5) is this, which says the column is 255 characters, not 2000 , "ACCESS_POINT_URL VARCHAR(255) NULL" jboss-as/server/production/deploy/juddi-service.sar/META-INF/ddl/juddi_create_db.ddl CREATE TABLE BINDING_TEMPLATE (SERVICE_KEY VARCHAR(41) NOT NULL, BINDING_KEY VARCHAR(41) NOT NULL, ACCESS_POINT_TYPE VARCHAR(20) NULL, ACCESS_POINT_URL VARCHAR(255) NULL, HOSTING_REDIRECTOR VARCHAR(255) NULL, LAST_UPDATE TIMESTAMP NOT NULL, PRIMARY KEY (BINDING_KEY), FOREIGN KEY (SERVICE_KEY) REFERENCES BUSINESS_SERVICE (SERVICE_KEY));	0	0	https://issues.jboss.org/browse/SOA-918
11112	SOA-918	What's the status?	0	0	https://issues.jboss.org/browse/SOA-918
11113	SOA-916	I don't want to move this out any further. Can we have this in CP04 please?	0	0	https://issues.jboss.org/browse/SOA-916
11114	SOA-911	Hiya Ian. Could you please clarify what further manifest properties are used by JON and what values are expect? For example, how would you determine that this was a SOA version of run.jar rather than an EAP version? Would the suffix be sufficient? It would be better to nail all these in one go. Thanks	1	1	https://issues.jboss.org/browse/SOA-911
11115	SOA-911	Thanks Ian. It looks like we may only have one outstanding issue in that case, differentiating between our two servers. The configurations we have are - standalone with a default configuration of 'default' - embedded with a default configuration of 'production' From what I can see from your last comment, you only (currently) cater for the embedded version and are not capable of handling the standalone version. Is this correct? If we are then to distinguish between standalone and embedded for the purpose of updating, does this mean that we must use different Implementation-Version strings? Perhaps 4.3.0.GA_SOA and 4.3.0.GA_SOA_STANDALONE? If I have understood your previous comments then this would appear to remove any clashes during the update process but, obviously, would not handle the different default configuration names. Thanks again, Kev	0	0	https://issues.jboss.org/browse/SOA-911
11116	SOA-880	Where are we with this one?	0	0	https://issues.jboss.org/browse/SOA-880
11117	SOA-877	Can we have confirmation that fixes are in place for 4.2CP03.	0	0	https://issues.jboss.org/browse/SOA-877
11118	SOA-872	Can we have confirmation that fixes are in place for 4.2CP03.	0	0	https://issues.jboss.org/browse/SOA-872
11119	SOA-867	Can we have confirmation that fixes are in place for 4.2CP03.	0	0	https://issues.jboss.org/browse/SOA-867
11120	SOA-860	Can we have confirmation that fixes are in place for 4.2CP03.	0	0	https://issues.jboss.org/browse/SOA-860
11121	SOA-855	Can we have confirmation that fixes are in place for 4.2CP03.	0	0	https://issues.jboss.org/browse/SOA-855
11122	SOA-852	Can we have confirmation that fixes are in place for 4.2CP03.	0	0	https://issues.jboss.org/browse/SOA-852
11123	SOA-844	Can we have confirmation that fixes are in place for 4.2CP03.	0	0	https://issues.jboss.org/browse/SOA-844
11124	SOA-842	This is due to the issue that we reported at: https://jira.jboss.org/jira/browse/SOA-731. so we will have opensso deployed in the tomcat in soa-p 4.3 release. does it make sense we add this explanation in the documentation? if it is, I can update it and then ask Darrin to help me updated in the doc book.	0	0	https://issues.jboss.org/browse/SOA-842
11125	SOA-841	Is this in the ESB docs or just the SOA-P docs?	0	0	https://issues.jboss.org/browse/SOA-841
11126	SOA-826	Also what document should this information be in? Service Guide?	0	0	https://issues.jboss.org/browse/SOA-826
11127	SOA-826	I need some idea of what is happening with this JIRA so I can schedule docs work for 5.0.1 What volume of material? e.g. how many pages, diagrams, etc When will draft material be ready ?	1	0	https://issues.jboss.org/browse/SOA-826
11128	SOA-825	Can we just include it "as is"?	0	0	https://issues.jboss.org/browse/SOA-825
11129	SOA-825	What is the status of this? We either include the docs in the next few days or push this out. It can't hold up the release. Worst case: let's include a link to the Smooks docs!	0	0	https://issues.jboss.org/browse/SOA-825
11130	SOA-818	The stacktrace would suggest that the database has not been initialised correctly, specifically that the classification scheme cannot be located. From the scout code public void setClassificationScheme(ClassificationScheme cscheme) throws JAXRException { scheme = cscheme; /* * not 100% clear, but I *think* the JavaDoc indicates that if * our internality dictates that of the scheme. */ 103: ((ClassificationSchemeImpl) scheme).setExternal(isExternal()); } So scheme must be null. public Classification createClassification(ClassificationScheme scheme, InternationalString name, String value) throws JAXRException { Classification cl = (Classification) this.createObject(LifeCycleManager.CLASSIFICATION); 198: cl.setClassificationScheme(scheme); cl.setName(name); cl.setValue(value); ((ClassificationImpl) cl).setExternal(true); return cl; } public Classification createClassification(ClassificationScheme scheme, String name, String value) throws JAXRException { 210: return createClassification(scheme, this.createInternationalString(name), value); } So this value of scheme has come from the JAXRegistryImpl ClassificationScheme cScheme = bqm.findClassificationSchemeByName(findQualifiers, "org.jboss.soa.esb.:category"); Collection<Classification> classifications = new ArrayList<Classification>(); 425: Classification classification = blm.createClassification(cScheme, "category", category); So the scheme is obtained from scout/jUDDI and would appear to be null. Does this classification exist in the jUDDI database? If not, why not?	1	1	https://issues.jboss.org/browse/SOA-818
11131	SOA-818	How is this classification created in the UDDI directory? Is the ESB supposed to create it at startup time?	0	0	https://issues.jboss.org/browse/SOA-818
11132	SOA-818	Do we need to setup categories, etc in SOA Software before starting up the ESB. Similar to how the jUDDI database is initialized?	0	0	https://issues.jboss.org/browse/SOA-818
11133	SOA-818	Tom, is there any way that this can be handled automatically?	1	1	https://issues.jboss.org/browse/SOA-818
11134	SOA-816	I'll add this in the Administration Guide in the Monitoring section is there also a config xml file where this can be set ?	0	0	https://issues.jboss.org/browse/SOA-816
11135	SOA-764	This doesn't sound like it is critical for this release. Correct?	0	0	https://issues.jboss.org/browse/SOA-764

#	Requirement	Questions	Answered	Successfully Answered	URL
11136	SOA-759	"The war file created to automatically expose ESB services through a webservice endpoint does not contain a dependency on the esb archive." This is a specific war file that ships in SOA for this purpose? filename? "the esb archive" in this context is the jbossesb.sar ? This is only an issue if "the esb archive has failed to deploy cleanly". In JBESB-1972 this was due to misconfiguration of, this would also be a factor for otherwise messed up deployments of the ESB into an existing JBossAS deployment?	0	0	https://issues.jboss.org/browse/SOA-759
11137	SOA-731	What solution was implemented for this JIRA? The user still has to install/use Tomcat and deploy OpenSSO to it, yes?	1	1	https://issues.jboss.org/browse/SOA-731
11138	SOA-731	Hi Len, For the readme.txt file update, do you mean add the 'steps 1-6 refer to installing opensso with Tomcat: ' sentence in the file?	0	0	https://issues.jboss.org/browse/SOA-731
11139	SOA-719	just add this note to known issues in release notes? Are there any other specific places where you know it should be noted?	0	0	https://issues.jboss.org/browse/SOA-719
11140	SOA-715	Can you please record the status here? Thanks.	0	0	https://issues.jboss.org/browse/SOA-715
11141	SOA-715	JBPM console works well with JDK 6. I did not find any error when I did the following test : 1. deploy a new processes. 2. mointor the processes status , start/delete a new process instance 3. manually control the task and monitor the task status 4. delete and monitor the all jobs and timers	1	1	https://issues.jboss.org/browse/SOA-715
11142	SOA-714	org.drools.scm.jcr.JcrActionFactoryTest Setting up the repository, registering node types etc. Setting up the repository, registering node types etc. Tests run: 4, Failures: 0, Errors: 3, Skipped: 0, Time elapsed: 0.342 sec <<< FAILURE! Running org.drools.scm.svn.SvnActionFactoryTest org.tmatesoft.svn.core.SVNException: svn: Path 'folder1' not present at org.tmatesoft.svn.core.internal.wc.SVNErrorManager.error(SVNErrorManager.java:49) at org.tmatesoft.svn.core.internal.io.fs.FSCommitEditor.openDir(FSCommitEditor.java:115) at org.drools.scm.svn.SvnActionFactory.openDirectories(SvnActionFactory.java:631) at org.drools.scm.svn.SvnActionFactory$AddDirectory.applyAction(SvnActionFactory.java:333) at org.drools.scm.CompositeScmAction.applyAction(CompositeScmAction.java:33) at org.drools.scm.svn.SvnActionFactory.execute(SvnActionFactory.java:265) at org.drools.scm.svn.SvnActionFactoryTest.testCannotAddDirectoryWithNoParent(SvnActionFactoryTest.java:79) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at junit.framework.TestCase.runTest(TestCase.java:154) at junit.framework.TestCase.runBare(TestCase.java:127) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:118) at junit.framework.TestSuite.runTest(TestSuite.java:208) at junit.framework.TestSuite.run(TestSuite.java:203) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.junit.JUnitTestSet.execute(JUnitTestSet.java:213) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:140) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:127) at org.apache.maven.surefire.Surefire.run(Surefire.java:177) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.booter.SurefireBooter.runSuitesInProcess(SurefireBooter.java:338) at org.apache.maven.surefire.booter.SurefireBooter.main(SurefireBooter.java:997) org.tmatesoft.svn.core.SVNException: svn: File already exists: filesystem '/local/svn/drools-cr2/4.3.0.GA_IR1/drools-repository/target/test-classes/copy_svn_repo_empty/db', transaction '1-1', path '/folder1' at org.tmatesoft.svn.core.internal.wc.SVNErrorManager.error(SVNErrorManager.java:49) at org.tmatesoft.svn.core.internal.io.fs.FSCommitter.makeDir(FSCommitter.java:166) at	0	0	https://issues.jboss.org/browse/SOA-714

#	Requirement	Questions	Answered	Successfully Answered	URL
11143	SOA-714	org.drools.brms.server.security.SecurityServiceImplTest INFO 29-08 16:43:23,870 (SecurityServiceImpl.java:login:45) Logging in user [XXX] Tests run: 2, Failures: 0, Errors: 0, Skipped: 0, Time elapsed: 0.368 sec Running org.drools.brms.server.PopulateDataTest DELETE test repo dir: /local/svn/drools-cr2/4.3.0.GA_IR1/drools-jbrms/repository TEST repo dir deleted. Setting up the repository, registering node types etc. DEBUG 29-08 16:43:27,722 (JackrabbitRepositoryConfigurator.java:registerNodeTypesFromCndFile:127) Attempting to regsiter node type named: {http://www.jboss.org/drools-repository/1.0}categoryNodeType Tests run: 1, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 3.856 sec <<< FAILURE! Running org.drools.brms.server.contenthandler.ContentHandlerTest DEBUG 29-08 16:43:27,745 (ContentManager.java:<init>:29) Loading content properties WARN 29-08 16:43:27,753 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. Tests run: 3, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 0.013 sec <<< FAILURE! Running org.drools.brms.server.util.FileManagerUtilsTest WARN 29-08 16:43:27,760 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:27,763 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:27,765 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:27,787 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:27,856 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:27,858 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. Tests run: 6, Failures: 0, Errors: 6, Skipped: 0, Time elapsed: 0.105 sec <<< FAILURE! Running org.drools.brms.server.files.DeploymentURIHelperTest Tests run: 1, Failures: 0, Errors: 0, Skipped: 0, Time elapsed: 0.004 sec Running org.drools.brms.server.selector.SelectorManagerTest DEBUG 29-08 16:43:27,868 (SelectorManager.java:<init>:27) Loading selectors ERROR 29-08 16:43:27,868 (SelectorManager.java:loadSelectorImplementation:85) java.lang.ClassNotFoundException: XXX DEBUG 29-08 16:43:28,302 (SelectorManager.java:getSelector:57) No selector found by the name of goo DEBUG 29-08 16:43:28,363 (SelectorManager.java:<init>:27) Loading selectors DEBUG 29-08 16:43:28,363 (SelectorManager.java:getSelector:57) No selector found by the name of XX Tests run: 1, Failures: 0, Errors: 0, Skipped: 0, Time elapsed: 0.499 sec Running org.drools.brms.server.util.ClassicDRLImporterTest import goo.wee Tests run: 3, Failures: 0, Errors: 0, Skipped: 0, Time elapsed: 0.005 sec Running org.drools.brms.client.modeldriven.ui.ConstraintValueEditorTest Tests run: 1, Failures: 0, Errors: 0, Skipped: 0, Time elapsed: 0.031 sec Running org.drools.brms.server.rules.BRMSSuggestionCompletionLoaderTest WARN 29-08 16:43:28,405 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:28,412 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:28,415 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:28,417 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:28,431 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:28,437 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. Tests run: 8, Failures: 0, Errors: 6, Skipped: 0, Time elapsed: 0.037 sec <<< FAILURE! Running org.drools.brms.server.ServiceImplementationTest WARN 29-08 16:43:28,445 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:28,448 (RulesRepository.java:getAreaNode:163) The repository appears to have become corrupted. WARN 29-08 16:43:28,451	0	0	https://issues.jboss.org/browse/SOA-714
11144	SOA-690	What tests do you refer to? We were only converting quickstarts - JBQA-1710	0	0	https://issues.jboss.org/browse/SOA-690
11145	SOA-680	Hi Martin, have you tried MetaDataFilter again, does the MetaDataFilter stil have great impact on performance (as Mark said it is not an issue any more...)? and how about other filters? do they have a performance impact that customers should be aware of? Thanks.	0	0	https://issues.jboss.org/browse/SOA-680
11146	SOA-662	Jirka, I have closed the ESB issue as 'Out of Date' given that we have now moved to Drools 5. Do you think this is still an issue with Drools 5?	1	0	https://issues.jboss.org/browse/SOA-662
11147	SOA-652	Tag created: https://svn.jboss.org/repos/labs/labs/jbossrules/soa_tags/4.3.0.GA_IR1/ It includes the following fixes: * [JBRULES-1251] Cannot remove a category if it was used in an archived rule * [JBRULES-1290] BRMS does not allow rule names to contain an apostrophe character (') * [JBRULES-1292] BRMS ignores "initial description" entered when creating new "Business rule using a DSL (text editor)" rule * [JBRULES-1441] ClassTypeResolver failling badly when trying to load a non existing simple name * [JBRULES-1676] Importing of nested classes more than one level deep does not work * [JBRULES-1677] Eclipse examples has incorrect .classpath setup	1	1	https://issues.jboss.org/browse/SOA-652
11148	SOA-647	You mean JBPM-1449 is fixed already? How was it fixed, was the facelets dependency upgraded?	0	0	https://issues.jboss.org/browse/SOA-647
11149	SOA-641	No changes? how can i use MySql database with no changes?	0	0	https://issues.jboss.org/browse/SOA-641
11150	SOA-640	Email sent to soa-engineering-list@redhat.com Hi, trying to get some info together on the SOA platform hardware requirements. I didn't think there would be much point with actual minimum requirements, as people would never actually run that. So I was thinking of a base line of "small to medium single application deployment". CPU: the only references I can find say a 'modern CPU', what should JBoss consider to be the baseline for 'modern'? P4? multi-core? is there any material on other platforms? SPARC, Itanium, PowerPC ? or anyone know of a good comparison reference ? does it matter? Harddrive (based on unzip downloads): standalong ESB: 220meg Full SOA (ie EAP+ESB): 600meg + space for logs Memory: Any recommendations for RAM for standalone ESB & full SOA with "small to medium single app" ? The SOA 4.2 getting started guide says: 768meg (addition to OS & other software requirements) with unspecified load This JIRA SOA-640 recommends 1Gig for EAP as a GUI-less server for "small to medium single app deployment" Darrin.	0	0	https://issues.jboss.org/browse/SOA-640
11151	SOA-640	Proposed inclusion: SOA Platform Hardware Requirement Recommendations ========== These recommendations do not represent an absolute minimum requirement for running the SOA Platform but rather presents a baseline for reasonable performance. The SOA Platform will run on far lesser hardware although your performance will suffer accordingly. Also note that performance is affected by other factors such as the sub-systems of the hardware platform. When using 'non-server class' hardware such as a laptop for development & testing, performance results may be less than those when deployed in your production environment. These guidelines also do not take into account the requirements of the server's operating system or other applications, nor does it include your own applications and services deployed on the SOA Platform. It is the developers' responsiblity to provide performance requirements for their own services & applications to be deployed. CPU === Minimum: Pentium4 class. Recommended: Dual-core or greater. MEMORY ====== 1 Gigabyte of memory is a recommended minimum. Increasing memory is often the most effective means of increasing performance. STORAGE SPACE ============= The full installation of the SOA Platform server requires 600 megabytes of harddrive space. The standalone ESB server only requires 220 megabytes of space. You will also require additional space for log files. 10 Gigabytes is considered adequate for a production environment with log rotation configured.	0	0	https://issues.jboss.org/browse/SOA-640

#	Requirement	Questions	Answered	Successfully Answered	URL
11152	SOA-628	Couple of things need to be verified here. 1) transactions. are the JTA transactions properly integrated with jBPM ? Is there a retry policy for transactions that failed because of optimistic locking failures ? Does it work correctly for rollback scenarios in case of different resources being used by ESB and jBPM and... 2) suite setup. in jpdl-suite, we have tried to create an easy-to-unzip distribution that contains out of the box, jbpm runtime installation, jbpm deployed on jboss with a preloaded jBPM DB containing an example and some example users. and the designer. The preconfigured database, the examples and the documentation needs to be reviewed. Maybe the jBPM build can be improved so that it is easier to build the SOA platform and so that the examples of jBPM 3 actually get tested automatically.	1	0	https://issues.jboss.org/browse/SOA-628
11153	SOA-628	Thomas, where does this currently stand?	0	0	https://issues.jboss.org/browse/SOA-628
11154	SOA-628	For completeness ;) (15:05:06) tdi: I want to have a look at https://jira.jboss.org/jira/browse/SOA-628 (15:05:21) kconner: ok (15:05:31) tdi: whats the esb bracnch I should be using? (15:05:46) kconner: You need to be using a SOA release (15:06:01) kconner: The SOA guys build jBPM and overlay ours (15:06:15) kconner: Which version are you looking at? (15:06:19) kconner: SOA 4.3 GA? (15:06:54) tdi: I'd like to create ESB/jBPM3 integration tests somewhere (15:07:05) kconner: ?? (15:07:59) tdi: shall we talk on the phone for a sec? this might be quicker (15:08:38) kconner: Sorry, I thought the purpose of the task was for someone to check the configuration of jBPM under SOA (15:08:42) kconner: Is that not correct? (15:09:05) kconner: Any 'integration' tests would just be between the jBPM configuration and the SOA platform, not ESB (15:09:27) kconner: Trevor has pasted the download link into #jbosssoa (15:09:44) tdi: There is a jBPM action in the ESB, right? (15:10:13) kconner: Only for sending to an ESB service, nothing more (15:10:20) kconner: We already have tests for that (15:10:27) kconner: This issue is not related to that action (15:11:52) tdi: If I want to answer the question "are JTA transactions properly integrated with jBPM", where do I create a test for that? (15:12:23) kconner: You create any jBPM specific test to see whether the transaction configuration of the job scheduler is correct (15:12:40) kconner: That is the issue (15:12:58) tdi: ok, where do I do that? (15:13:00) kconner: This is about generic jBPM and its configuration, nothing more (15:13:35) kconner: I would guess that the correct place for jBPM tests is within jBPM (15:13:45) kconner: If it is SOA specific then you may want to ask Trev to add it to their build (15:13:48) kconner: (or QE) (15:14:01) kconner: There is no ESB involvement with this (15:14:16) tdi: ok, thanks (15:15:27) kconner: If you do want to look at the ESB code then I can point you in the right direction (15:15:39) kconner: But I think the real issue is the jBPM configuration within SOA (15:15:41) tdi: yes, please (15:17:00) kconner: The current SOA code is http://anonsvn.jboss.org/repos/labs/labs/jbossesb/workspace/platform/JBESB_4_4_SOA_4_3_GA/product/services/jbpm/src/main/java (15:17:08) kconner: I am going to go through it soon and simplify it	1	0	https://issues.jboss.org/browse/SOA-628
11155	SOA-628	Tom, could you please reassign appropriately.	1	0	https://issues.jboss.org/browse/SOA-628
11156	SOA-623	I think the title of this issue needs fixing as it doesn't make sense ;-) (replace jboss-esb with jboss-esb?)	0	0	https://issues.jboss.org/browse/SOA-623
11157	SOA-623	Darrin, where - in which files and/or docs - were these changes made? I'm seeing references to SOA_ROOT, but not ESB_ROOT. (Which actually seems OK - do we have any docs that refer to ESB_ROOT?) GettingStarted Guide: To use the POP3 gateway, replace the mail-ra.rar in ${SOA_ROOT}/server/${CONFIG}/deploy/ with the mail-ra.rar from the directory seam/mail/. Admin Guide: jcr.jndi.path - optional path in JNDI where the repository is found. If this property is not specified, a new repository will be created based on ${SOA_ROOT}/server/${CONFIG}/deploy/ jbossesb.esb/repository.xml , and the repository data will be stored in ${SOA_ROOT}/server/${CONFIG}/data/repository/	0	0	https://issues.jboss.org/browse/SOA-623
11158	SOA-613	Where in the documentation this should be added ?	0	0	https://issues.jboss.org/browse/SOA-613
11159	SOA-613	Still don't know where to put this. Services Guide? Release Notes? Supported ActiveEndpoints BPEL Engine ========== The following versions of the ActiveEndpoints ActiveBPEL Engine are supported on the JBoss Enterprise SOA Platform: * ActiveBPEL 5.0 * ActiveBPEL 4.1 * ActiveBPEL 4.0 * ActiveBPEL 3.1 * ActiveBPEL 3.0 * ActiveBPEL 2.1 * ActiveBPEL 1.2 They can be downloaded from the ActiveEndpoints website at: http://www.activevos.com/community-open-source-engine-download.php The ActiveEndpoints ActiveBPEL Engine open-source license terms can be found here: http://www.activevos.com/community-open-source-license.php	0	0	https://issues.jboss.org/browse/SOA-613
11160	SOA-613	Feedback received: Looking at https://jira.jboss.org/jira/browse/JBESB-1852, it looks as though v5.0.2 is the latest version known to work with ActiveBPEL. As for a full list of versions... I wouldn't have that. Regards, T. Darrin Mison wrote: > Hi, can you confirm what versions of ActiveEndpoints BPEL Engine have > been tested against. > This is for the 4.3GA docs.	0	0	https://issues.jboss.org/browse/SOA-613
11161	SOA-577	Hi Trevor. so what are we going to do regarding this jira? Remove these demos that didt work or add dependencies that required by these failed demos or what?	0	0	https://issues.jboss.org/browse/SOA-577
11162	SOA-570	Please could you have a look on this?	1	1	https://issues.jboss.org/browse/SOA-570
11163	SOA-570	Unfortately I can't use phrases like "vice versa" because they don't translate well. New content: Once a queue is undeployed you cannot deploy a new queue with the same name if one of the queues was a clustered queue and the other queue was not. You must delete the destination from the JBM_POSTOFFICE table before the new queue can be deployed successfully.	0	0	https://issues.jboss.org/browse/SOA-570
11164	SOA-554	Question - are we seeing a related issue with MS-SQL Server DB - where the limit is 8K?	0	0	https://issues.jboss.org/browse/SOA-554
11165	SOA-554	Can you check tha value of @@TEXTSIZE parameter?	0	0	https://issues.jboss.org/browse/SOA-554
11166	SOA-554	Release notes: 7.1. Sybase database cuts the payload according to pagesize by default When a database is created Sybase sets the maximum size of text and image datatypes to the default page size of 2 KB. The longer message payload is truncated without any information or warning. The database parameters have to be updated to change the page size. Further information: http://jira.jboss.com/jira/browse/SOA-554 | http://jira.jboss.com/jira/browse/SOA-553 Is this ok? Or do you want something about MS SQL Server, as well?	0	0	https://issues.jboss.org/browse/SOA-554
11167	SOA-550	The CP02 CR1 build Getting Started Guide includes this text: ---------- 6. Delete the following files and directories from the development/deploy directory: ? cluster-service.xml ? deploy-hasingleton-service.xml ? ejb3-clustered-sfsbcache-service.xml ? httpha-invoker.sar ? jboss-web-cluster.sar ? deploy-hasingleton ? farm ---------- The last two lines should be moved (deploy-hasingleton, farm) - these files/dirs are not in server/development/deploy, but are in server/development	0	0	https://issues.jboss.org/browse/SOA-550

#	Requirement	Questions	Answered	Successfully Answered	URL
11168	SOA-550	Also seeing this failure - need to retry... javax.naming.CommunicationException: Could not obtain connection to any of these urls: 127.0.0.1:1100 and discovery failed with error: javax.naming.CommunicationException: Receive timed out [Root exception is java.net.SocketTimeoutException: Receive timed out] [Root exception is javax.naming.CommunicationException: Failed to connect to server 127.0.0.1:1100 [Root exception is javax.naming.ServiceUnavailableException: Failed to connect to server 127.0.0.1:1100 [Root exception is java.net.ConnectException: Connection refused]]] at org.jnp.interfaces.NamingContext.checkRef(NamingContext.java:1562) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:634) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:627) at javax.naming.InitialContext.lookup(InitialContext.java:351) at org.jboss.util.naming.Util.lookup(Util.java:215) at org.jboss.resource.adapter.jms.inflow.dlq.AbstractDLQHandler.setupDLQDestination(AbstractDLQHandler.java:106) at org.jboss.resource.adapter.jms.inflow.dlq.AbstractDLQHandler.setup(AbstractDLQHandler.java:82) at org.jboss.resource.adapter.jms.inflow.dlq.JBossMQDLQHandler.setup(JBossMQDLQHandler.java:48) at org.jboss.resource.adapter.jms.inflow.JmsActivation.setupDLQ(JmsActivation.java:369) at org.jboss.resource.adapter.jms.inflow.JmsActivation.setup(JmsActivation.java:315) at org.jboss.resource.adapter.jms.inflow.JmsActivation$SetupActivation.run(JmsActivation.java:601) at org.jboss.resource.work.WorkWrapper.execute(WorkWrapper.java:204) at org.jboss.util.threadpool.BasicTaskWrapper.run(BasicTaskWrapper.java:275) at EDU.oswego.cs.dl.util.concurrent.PooledExecutor$Worker.run(PooledExecutor.java:756) at java.lang.Thread.run(Thread.java:595) Caused by: javax.naming.CommunicationException: Failed to connect to server 127.0.0.1:1100 [Root exception is javax.naming.ServiceUnavailableException: Failed to connect to server 127.0.0.1:1100 [Root exception is java.net.ConnectException: Connection refused]] at org.jnp.interfaces.NamingContext.getServer(NamingContext.java:274) at org.jnp.interfaces.NamingContext.checkRef(NamingContext.java:1533) ... 14 more Caused by: javax.naming.ServiceUnavailableException: Failed to connect to server 127.0.0.1:1100 [Root exception is java.net.ConnectException: Connection refused] at org.jnp.interfaces.NamingContext.getServer(NamingContext.java:248) ... 15 more Caused by: java.net.ConnectException: Connection refused at java.net.PlainSocketImpl.socketConnect(Native Method) at java.net.PlainSocketImpl.doConnect(PlainSocketImpl.java:333) at java.net.PlainSocketImpl.connectToAddress(PlainSocketImpl.java:195) at java.net.PlainSocketImpl.connect(PlainSocketImpl.java:182) at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:366) at java.net.Socket.connect(Socket.java:520) at java.net.Socket.connect(Socket.java:470) at java.net.Socket.<init>(Socket.java:367) at java.net.Socket.<init>(Socket.java:267) at org.jnp.interfaces.TimedSocketFactory.createSocket(TimedSocketFactory.java:84) at org.jnp.interfaces.TimedSocketFactory.createSocket(TimedSocketFactory.java:77) at org.jnp.interfaces.NamingContext.getServer(NamingContext.java:244) ... 15 more 22:02:06,112 INFO [HbmBinder] Mapping subclass: org.jbpm.context.log.variableins	0	0	https://issues.jboss.org/browse/SOA-550
11169	SOA-550	Still seeing the above problem - after these steps: vim development/conf/jboss-log4j.xml Logging settings vim development/conf/jboss-service.xml <attribute name="KeepAliveTime">5000</attribute> rm development/deploy/cluster-service.xml rm development/deploy/deploy-hasingleton-service.xml rm development/deploy/ejb3-clustered-sfsbcache-service.xml rm -Rf development/deploy/httpha-invoker.sar/ rm -Rf development/deploy/jboss-web-cluster.sar/ cp -r default/deploy/http-invoker.sar/ development/deploy rm development/deploy/hajndi-jms-ds.xml cp default/deploy/jms-ds.xml development/deploy vim development/deploy/jboss-web.deployer/META-INF/jboss-service.xml Comment out: <!-- depends>jboss.cache:service=TomcatClusteringCache</depends> --> Jarek - what am I missing?	0	0	https://issues.jboss.org/browse/SOA-550
11170	SOA-548	JBAS-5094 was closed by rollback, did the xmlbeans issue in the first commentever get looked at?	1	1	https://issues.jboss.org/browse/SOA-548
11171	SOA-539	Does this issue belong in the FP? Is it new functionality?	0	0	https://issues.jboss.org/browse/SOA-539
11172	SOA-539	We have closed the ESB issues as 'Wont Fix', do we still need to add this?	1	0	https://issues.jboss.org/browse/SOA-539
11173	SOA-521	Other than Docs I'm not sure what you are wanting? jbpm-jpdl at the top level is nothing to do with EAP, it's there because JBDS wanted somewhere they could get at all the jbpm jars. Docs were moved to their current location as a result of a request from josh in https://jira.jboss.org/jira/browse/SOA-361 Can we move wholesale directories around in a CP?	0	0	https://issues.jboss.org/browse/SOA-521
11174	SOA-520	The quickstart has been modifed/updated in the SOA-P 4.3 release (and in ESB 4.5)	1	1	https://issues.jboss.org/browse/SOA-520
11175	SOA-512	Clarification required for documentation: "if the person INITIALLY manually modified the config files, they will be unable to use the schema tool at a later date. " what config files are being refered to here?	0	0	https://issues.jboss.org/browse/SOA-512
11176	SOA-508	Do you have some specific examples that you could share with us? cheers, --josh	0	0	https://issues.jboss.org/browse/SOA-508
11177	SOA-507	Could you please give the incorrect sentence and perhaps some context around it so that I can grep for it. thanks, -josh	1	1	https://issues.jboss.org/browse/SOA-507
11178	SOA-507	seamframework.org is the project site. The jboss.com site has been previously updated so that most of the links go to seamframework.org for the latest information. Pointing to the .com site should be fine as long as is it updated. Who's responsibility is that?	1	1	https://issues.jboss.org/browse/SOA-507
11179	SOA-487	Can you please comment on whether this is still required or obsolete. If it obsolete can you close as such, otherwise can you re-affirm and re-assign to me.	0	0	https://issues.jboss.org/browse/SOA-487
11180	SOA-465	It's close - but not the same. Is this a 'major' priority problem? JMX Console: 4.2.0.GA_CP02_SOA (build: SVNTag=SOA_4.2.0.GA_CP02 date=200806090827) run.jar: Implementation-Version: SOA_4.2.0.GA_CP02 (build: SVNTag=SOA_4.2.0.GA_CP02 date=200806090827)	0	0	https://issues.jboss.org/browse/SOA-465
11181	SOA-462	1.5GB is too large. Will remove all the .svn directories. That should halve the size: mkdir soa-4.2.0 && \ (cd soa-4.2.0 && \ svn co https://svn.corp.jboss.com/repos/soa/trunk/src-tools && \ mv src-tools/* . && \ rm -rf src-tools && \ svn co https://svn.corp.jboss.com/repos/soa/tags/4.2.0GA/build-tools && \ find build-tools -name .svn -exec rm -fr {} \; -prune && \ (cd build-tools && ant materialize.projects) && \ find soa-work -name .svn -exec rm -fr {} \; -prune) && \ find soa-4.2.0 | zip soa-4.2.0.zip -@ && \ md5sum soa-4.2.0.zip > soa-4.2.0.zip.md5	1	1	https://issues.jboss.org/browse/SOA-462
11182	SOA-457	Jervis: I did see a QS named bpm_orchestration4 in jbossesb-server-4.3.GA kit built by myself using trunk. Shall I say this jira can be closed now? Trevor: It looks like it's in the trunk but not in the platform branch. Assign it to me and I'll investigate further.	1	1	https://issues.jboss.org/browse/SOA-457
11183	SOA-454	Can I have some information on this issue for the TP_FP01 release notes please? thanks, --josh	1	0	https://issues.jboss.org/browse/SOA-454
11184	SOA-428	Well GA is out already so I care more about SP1 or whatever next version will be. Weren't we going to have a branch for each project?	0	0	https://issues.jboss.org/browse/SOA-428
11185	SOA-428	Can we close this issue?	0	0	https://issues.jboss.org/browse/SOA-428

#	Requirement	Questions	Answered	Successfully Answered	URL
11186	SOA-428	Alex - are you all set with this? Trevor - isn't this the branch for CP01? https://svn.labs.jboss.com/labs/jbossesb/workspace/platform/JBESB_4_2_1_SOA_4_2/	0	0	https://issues.jboss.org/browse/SOA-428
11187	SOA-422	Are there any other open JIRAs against these issues that I can refer to? Is there a way for the user to distinguish between the spurious WARN and ERROR messages and important ones?	1	1	https://issues.jboss.org/browse/SOA-422
11188	SOA-421	Jiri, could you please update this JIRA with more information, or a pointer to source JIRAs for more information. thanks --josh	1	1	https://issues.jboss.org/browse/SOA-421
11189	SOA-421	What's the roadmap for stateful action replication in a clustered environment? Is there one? Is there a pointer that I can give to more information on developing a replicating stateful action?	1	1	https://issues.jboss.org/browse/SOA-421
11190	SOA-405	Josh, can you add this to the online docs for GA and then pull it into the docs for the CP release?	0	0	https://issues.jboss.org/browse/SOA-405
11191	SOA-377	Josh, we may be creating a workaround for this issue. Can you hold off on this until I get back to you?	0	0	https://issues.jboss.org/browse/SOA-377
11192	SOA-375	jUDDI / Scout compatibility between JBoss Enterprise SOA Platform and JBoss Enterprise Application Platform During testing of the JBoss Enterprise SOA Platform it did not pass a number of JAXR Scout tests from the JBoss Enterprise Application test suite. This indicates that there may be differences between the jUDDI / Scout implementations that affect applications. At this stage there are no known instances of these differences affecting applications. However, if you find an affected application, please inform us. Further Information: http://jira.jboss.com/jira/browse/SOA-354	0	0	https://issues.jboss.org/browse/SOA-375
11193	SOA-359	That version should be from BRANCH_3_2_2_SOA_4_2, right?	0	0	https://issues.jboss.org/browse/SOA-359
11194	SOA-327	Do I have this right: jboss-soa-p.4.2.0/jboss-as/server/all/deploy/jbpm.esb/jbpm-console.war jboss-soa-p.4.2.0/jboss-as/server/production/deploy/jbpm.esb/jbpm-console.war are both the secure war that should be used in production. jboss-soa-p.4.2.0/jbpm-jpdl/deploy/jbpm-console.war is the "insecure" war file to be used for development. To switch from the default secured war a user should copy the currently deployed one from /server/production/deploy/jbpm-console.war to another folder as jbpm-console.war.secure (can they do this in place? i.e: can they simply rename the file like this) and copy in the jbpm-jpdl version. Rinse, lather, and reverse to go from insecure to secure?	1	1	https://issues.jboss.org/browse/SOA-327
11195	SOA-325	My thoughts from an email: SOA-P is based on EAP-4.3, but beside patches available for this EAP version, we will have patches strictly related to the SOA-P - for "subcomponents" like esb, smooks, drools, ... which are not present in EAP. Are we going to double the bits available for download? If no subcomponents are involved for a particular CP then we will have one CP for EAP-4.3 and the same CP available for SOA-P. Or do we want to have all CPs for EAP-4.3 also available for SOA-P and additionally another CP for SOA-P "subcomponents"? I don't think that would make sense. Maybe another approach: Create CPs only for SOA-P which would be also applicable for EAP-4.3. Even if the CP would contain patches for esb, drools, etc, it wouldn't affect the EAP-4.3 because these bits aren't there and would be ignored. The version string for SOA-P is currently the same as in EAP 4.3.	1	1	https://issues.jboss.org/browse/SOA-325
11196	SOA-290	Tom, could you please reassign appropriately.	0	0	https://issues.jboss.org/browse/SOA-290
11197	SOA-236	Trev, can you talk to Mike and one of you look at this as a priority in the new year? When I asked Ryan who did this for EAP he said: "Mr. Crouch spec'd it & implemented the JON logic. I implemented the build tooling. This repsonsibility has now been transfered to release engineering (Chris O'Brien)." So that should give you a handle on contact points.	0	0	https://issues.jboss.org/browse/SOA-236
11198	SOA-235	Why is this a bad thing to bundle the jbpm framework ?	1	0	https://issues.jboss.org/browse/SOA-235
11199	SOA-235	Coming back to this issue. How do you want to solve the problem of a tool supporting multiple runtimes at the same time and offering code completion and 'ctrl-click' functionality to view sources of classes if you don't include these sources? Also, this is really only a matter of what you want to offer your users. If you don't include the sources, than the editor and the create process wizard will still work, only they will not be able to create 'jbpm projects'.	1	1	https://issues.jboss.org/browse/SOA-235
11200	SOA-235	koen - is it actually *all* the sources that is needed or is it like with seamgen in seam where we will extract the "template" part to use for tooling independent of the runtimes... anyway this is a general problem we have that the runtimes are not dev-friendly, but primary prod-friendly. Something we should find a good solution for since it is relevant for all projects.	0	0	https://issues.jboss.org/browse/SOA-235
11201	SOA-235	Is src attachement to classpath the *only* feature this is for ? I thought you had some templates or something you needed from the distribution ?	0	0	https://issues.jboss.org/browse/SOA-235
11202	SOA-235	this would be a big work. i don't think it's worth doing it on the jBPM 3 codebase. we will make sure that this works in jPDL 4. but the solution is not yet clear to me. on the one hand, you want the plugin not to have a dependency on the runtime for creating a new template process eclipse project. on the other hand, we want the jpdl runtime download to contain a command line tool for creating a new template process eclipse project so do we want to end up with 2 versions of the runtime ? one that is contained within the designer and one on the filesystem ? here's my thoughts on this: * if the designer needs things from a runtime, it should get it from a runtime installation on the file system. * runtime should be able to create a new eclipse template project for working with processes (command line tool) * designer should be able to invoke that command line tool from the designer * optionally, the designer could have a new process project that doesn't have a dependency on the runtime. this would be the case for process modelling only. at this point we're not sure yet to what extend this will be realized in the next jPDL 4 designer version.	0	0	https://issues.jboss.org/browse/SOA-235
11203	SOA-235	Who wants command line creation of jbpm specific projects ? I still have not understood what the dependency is all about except the wish to have source attachements in the classpath container - that is not a critical thing and should be handled in a different way IMO (but again that is a future thing) So I ask again: What is the *exact* dependency on the runtime dir in jbpm (besides source attachments) ? e..g for Seam it iis that we tempaltes in the seam-gen dir that we use in our Eclipse Wizards - main functionallity that would not work without a "runtime dir".	1	0	https://issues.jboss.org/browse/SOA-235
11204	SOA-234	Can we close this issue now?	0	0	https://issues.jboss.org/browse/SOA-234
11205	SOA-195	care to elaborate why this is the case? I'm very curious	0	0	https://issues.jboss.org/browse/SOA-195
11206	SOA-182	Yes, but then they have CORBA that is supported in EAP all and production profiles, but not default. So how is this any more confusing?	0	0	https://issues.jboss.org/browse/SOA-182
11207	SOA-180	Is this still an issue? Did SOA-205 fix it?	0	0	https://issues.jboss.org/browse/SOA-180
11208	SOA-151	Wouldn't it be sufficient to run the WS tests against JBAPP?	0	0	https://issues.jboss.org/browse/SOA-151

#	Requirement	Questions	Answered	Successfully Answered	URL
11209	SOA-151	What we really need to do is to run the tests in both a JBPAPP and SOA-P configuration: * Running the tests against a "vanilla" JBPAPP system will give us a baseline for throughput * We can then compare that set of throughput baseline numbers against the same tests running on a SOA-P system - we would hope to not see either functional problems not seen on the JBPAPP system and we would hope to see no significant degredation in throughput for the tests running on a SOA-P system as compared to the JBPAPP system	1	0	https://issues.jboss.org/browse/SOA-151
11210	SOA-106	what is a noauth zip?	0	0	https://issues.jboss.org/browse/SOA-106
11211	SOA-84	SOA-147 is a GA blocker - can be worked around for testing	0	0	https://issues.jboss.org/browse/SOA-84
11212	SOA-84	I fail to see how SOA-75 can still be happening as those calls no longer use the RMIAdapter. Do we have a log from this failure?	0	0	https://issues.jboss.org/browse/SOA-84
11213	SOA-62	Will a script be included to do this, or will this be a sequence of steps that users need to follow manually, and which we will need to provide in the installation guide?	0	0	https://issues.jboss.org/browse/SOA-62
11214	SOA-53	Jarek and Fernando - is this issue still open? Should it still go into the rel notes?	0	0	https://issues.jboss.org/browse/SOA-53
11215	SOA-45	please see http://hudson.qa.jboss.com/hudson/view/SOA/job/drools-4.0.2_SOA-4.2_branch/ Mark, can you verify we run all necessary tests against the SOA-P drools bits? http://hudson.qa.jboss.com/hudson/view/SOA/job/drools-4.0.2_SOA-4.2_branch/64/testReport/	1	1	https://issues.jboss.org/browse/SOA-45
11216	SOA-45	what do you mean by "all remaining tests", can you point to any jiras? Are these new tests? What kind of tests? BTW, in my comment above, by "all necessary tests", I meant all necessary tests included in the native drools testsuite.	1	1	https://issues.jboss.org/browse/SOA-45
11217	SOA-45	Below is an URL to results for running the testsuite against the current SOA-P bits: http://hudson.qa.jboss.com/hudson/job/jboss-soa-platform-rules-testsuite-single/93/testReport/ All errors will go away, when we use finally the correct drools bits in SOA-P, beside two issues which come up only when using JANINO: Could you please have a look at them? http://hudson.qa.jboss.com/hudson/job/jboss-soa-platform-rules-testsuite-single/93/testReport/org.drools.compiler/Java5Test(JANINO)/testJava5Rule/ http://hudson.qa.jboss.com/hudson/job/jboss-soa-platform-rules-testsuite-single/93/testReport/org.drools.integrationtests/MiscTest(JANINO)/testWithInvalidRule/ Thanks	1	1	https://issues.jboss.org/browse/SOA-45
11218	SOA-45	According the latest test the configuration isn't founding janino library, could someone take a look or help me with? http://hudson.qa.jboss.com/hudson/job/jboss-soa-platform-rules-testsuite-single/95/	1	1	https://issues.jboss.org/browse/SOA-45
11219	SOA-45	What is the status of this issue?	1	1	https://issues.jboss.org/browse/SOA-45
11220	SOA-45	Is there a JIRA for the failed tests?	0	0	https://issues.jboss.org/browse/SOA-45
11221	SOA-43	Should I assign this to Fernando to work on?	0	0	https://issues.jboss.org/browse/SOA-43
11222	SOA-43	Fernando - can you add this to the short list of other Drools issues that you are working on? Thanks...	1	0	https://issues.jboss.org/browse/SOA-43
11223	SOA-43	I need to create a unit test to include this, where can I find some example on how write tests against soa-p?	0	0	https://issues.jboss.org/browse/SOA-43
11224	SOA-11	Len, can you email me about this?	0	0	https://issues.jboss.org/browse/SOA-11
11225	SOA-9	Yes, it would be extremely helpful/nice if we would be able to cut down these dependencies somewhat. Would it be possible to isolate what dependencies are actually _needed_ to complete the build?	0	0	https://issues.jboss.org/browse/SOA-9
11226	SOA-4	Received svn commit access today (Aug 24 2007). Some questions: Will this repository be a mirror of https://svn.labs.jboss.com/labs/jbossesb/trunk? (plus the tests that we add) I just tried to access https://svn.corp.jboss.com/repos/soa/ - the only link there is "admin" - when I tried to access it I received a 403 Forbidden - You don't have permission to access /repos/soa/admin/ on this server error.	0	0	https://issues.jboss.org/browse/SOA-4
11227	SOA-2	Can we add some automatic profiling to CI?	1	1	https://issues.jboss.org/browse/SOA-2
11228	JBEWS-455	Permaine, the only information I have at present is this: Enterprise Web Server is now distributed with Red Hat Signed JARs Red Hat signed JARs mean you can be confident that your product distribution is exactly the distribution that was shipped by Red Hat. Is this sufficient? More information about the benefits of signed JARs might be helpful.	0	0	https://issues.jboss.org/browse/JBEWS-455
11229	JBEWS-452	Assigning to you for the exact commands required for each of the above installations. Which would be the most common use cases and in what order should they be included?	0	0	https://issues.jboss.org/browse/JBEWS-452
11230	JBEWS-444	Mladen, could you please take a quick look to see if the attached patch is good for 2.2.17? Thanks!	1	1	https://issues.jboss.org/browse/JBEWS-444
11231	JBEWS-444	httpd22-2.2.17-5.ep5.el4, httpd-2.2.17-5.1.ep5.el5 and httpd-2.2.17-5.1.ep5.el6 built in Brew with the patch. Mladen has also added the patch to the Hudson builds, will be in the build tomorrow. Marking as resolved.	0	0	https://issues.jboss.org/browse/JBEWS-444
11232	JBEWS-439	Documented as a known issue in EWS 1.0.1 release notes.	0	0	https://issues.jboss.org/browse/JBEWS-439
11233	JBEWS-439	This feature is verified to be provided on all platforms in the EWS 1.0.2 release with a known bug on RHEL-4 JBPAPP-6282 . Closing issue.	0	0	https://issues.jboss.org/browse/JBEWS-439
11234	JBEWS-439	I have included this in the New With This Release section of the RNs for 1.0.2 as so: Apache httpd 2.2 contains support for Online Certificate Security Protocl (OCSP). Is this the right place for it? Is there anything else you would like to see written about it? Thanks.	1	0	https://issues.jboss.org/browse/JBEWS-439
11235	JBEWS-429	I've got a program: {code}int main(void) { printf("%s\n",apr_version_string()); }{code} Solaris: {code}# ./apr 1.4.2{code} RHEL5: {code}$ ./apr 1.2.7{code} Why do we use a different apr version on solaris and RHEL? Wouldn't it be less engineering for everybody to support same version?	0	0	https://issues.jboss.org/browse/JBEWS-429
11236	JBEWS-429	If this issue is resolved, does the documentation need to advise that aprutils and apr are required to be installed? ie: https://issues.jboss.org/browse/JBPAPP-3293 Thanks.	0	0	https://issues.jboss.org/browse/JBEWS-429
11237	JBEWS-417	Hi Juraci, and Rajesh, thanks for your comments. @Rajesh, I have removed the references to Validator and Search in the additional modules section of the Annotations Guide. @Juraci, I changed EJB3 to JPA where I could, but I am not familiar enough with the subject to be confident of my changes. Could you please look over the new copy(ies) on stage and let me know A) If I have made a complete mess and B) Where I have missed EJB3--> JPA? There are a number of methodnames etc that I have not changed. I have added a note to the release notes in New With This Release warning customers that JTA is not supported with EWS; would you like to see notes in all of the books as well, or is the RNs sufficient? I have removed the Setup section advising that the compatibility matrix be consulted. Please let me know if I need to replace it. I have made the change from RESOURSE_LOCAL To RESOURCE_LOCAL Thanks again for your feedback.	1	1	https://issues.jboss.org/browse/JBEWS-417
11238	JBEWS-417	Hello Rebecca, I think I'm still seeing the old document: http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Web_Server/1.0/html-single/Hibernate_Entity_Manager_Reference_Guide/index.html Can you please confirm it's the right location?	1	1	https://issues.jboss.org/browse/JBEWS-417

#	Requirement	Questions	Answered	Successfully Answered	URL
11239	JBEWS-416	One of my RHEL friends gave me these lines: [root@localhost sbin]# /usr/sbin/groupadd -g 91 apache [root@localhost sbin]# /usr/sbin/useradd -s /bin/sh -d /home/apache -u 91 -g apache -G apache apache are they correct, should they be in the docs.	0	0	https://issues.jboss.org/browse/JBEWS-416
11240	JBEWS-413	Hi Eva, can you please verify the progress on this, and resolve, please? Thanks.	0	0	https://issues.jboss.org/browse/JBEWS-413
11241	JBEWS-410	Laura, can you please let Jim know of the location of the EWS docs? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBEWS-410
11242	JBEWS-410	Permaine, who would be the best person to get this information from? Currently we have no docs for EWS except the Release Notes. Do we need a separate guide, possibly a common tasks guide of some sort?	0	0	https://issues.jboss.org/browse/JBEWS-410
11243	JBEWS-410	Sorry, I meant the Release Notes, the info is in there. Can you please let us know where the Release Notes is located? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBEWS-410
11244	JBEWS-407	Mladen, can we get rid of the duplicate javadocs in the Solaris/Windows builds? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBEWS-407
11245	JBEWS-407	Documentation is produced as part of build process. For example make install for httpd creates the documentation as well. The final documentation for windows also has DOS line endings which makes it more familiar to the windows users. I would leave that to our customers to decide weather they wish the single download package or multiple packages that would require the careful drop in inside the installation tree to have working and browsable documentation. Since this is certainly not an technical issue, we can discuss to eventually split documentation in later versions. Changing that in the middle of product life cycle is usually not an smart idea, cause the first thing the customer will file as a technical support question would be: Where are the docs?	1	1	https://issues.jboss.org/browse/JBEWS-407
11246	JBEWS-404	You are right, I was starting cmd from Total Commander. But this is real-life scenario - will the client be aware that he can't call service.bat from a cmd launched from a 32 bit app? Even if it's documented, this will be a bit annoying feature. Please consider making this work from 32-bit cmd, at least for EWS 1.0.2.	0	0	https://issues.jboss.org/browse/JBEWS-404
11247	JBEWS-400	Dave, can you please look into this? Thanks!	0	0	https://issues.jboss.org/browse/JBEWS-400
11248	JBEWS-399	Hi Eva, can you please verify if this has been included and resolve? Thanks.	1	1	https://issues.jboss.org/browse/JBEWS-399
11249	JBEWS-399	Hi Rebecca, module descriptions are available at http://httpd.apache.org/docs/2.2/mod/ I have provided the link in the guide. However, I think the list of supported modules should be included in the guide. Could you, please, tell me, who is the right person to talk to? And also: is this sufficient? Thanks	0	0	https://issues.jboss.org/browse/JBEWS-399
11250	JBEWS-397	Rajesh, Can this be closed?	1	0	https://issues.jboss.org/browse/JBEWS-397
11251	JBEWS-396	In the RHEL zips, I've added a dir named extras under jboss-ews-1.0 and moved the log4j.jar and log4j.properties there. Mladen, can you please update the zip for the other platform? Thanks!	0	0	https://issues.jboss.org/browse/JBEWS-396
11252	JBEWS-391	The version of mod_jk in EWS 1.0.1 will be 1.2.28, can you please re-evaluate and let me know? Thanks, Permaine	1	1	https://issues.jboss.org/browse/JBEWS-391
11253	JBEWS-390	Hi Eva, Could you please take a look at this and see if you can work it into the Installation Guide? Thanks! Bec	1	1	https://issues.jboss.org/browse/JBEWS-390
11254	JBEWS-390	Could you, please, provide me the information I need to add to the installation guide? Thank you	1	1	https://issues.jboss.org/browse/JBEWS-390
11255	JBEWS-390	Paths to startup.sh: Solaris: /opt/redhat/ews/share/tomcat5/bin/startup.sh /opt/redhat/ews/share/tomcat6/bin/startup.sh RHEL zip - can be installed anywhere, so, $EWS_HOME/tomcat[5|6]/bin/startup.sh RHEL RPM - I don't know, ask Ivo Studensky Windows: "c:\Program Files\Red Hat\Enterprise Web Server\share\tomcat5\bin\startup.sh" "c:\Program Files\Red Hat\Enterprise Web Server\share\tomcat6\bin\startup.sh" Maybe you could use $EWS_HOME for simplicity...	1	1	https://issues.jboss.org/browse/JBEWS-390
11256	JBEWS-390	I agree that management system should provide instructions on how different components can be used to work with it. Ondrej, can you please ask someone from the JON team to take a look? we don't have experience with JON. Thanks!	0	0	https://issues.jboss.org/browse/JBEWS-390
11257	JBEWS-390	Charles, could you please assign someone to help with this?	1	1	https://issues.jboss.org/browse/JBEWS-390
11258	JBEWS-390	Rewording the paragraph once more. Eva, rather than describing it here, can you give me a link to SVN to this chapter? I'd make a patch and send you.	1	1	https://issues.jboss.org/browse/JBEWS-390
11259	JBEWS-390	Hi Ondrej, what's the progress on this? Did you get the SVN link you were after? Thanks.	0	0	https://issues.jboss.org/browse/JBEWS-390
11260	JBEWS-389	We don't add those in by default, but I can add them in as comment so that it'll be easy for the customers. How does that sound?	0	0	https://issues.jboss.org/browse/JBEWS-389
11261	JBEWS-386	Will these be available for customer to deploy or one would need to download from the apache web site?	0	0	https://issues.jboss.org/browse/JBEWS-386
11262	JBEWS-386	Can we remove the examples (webapps/examples) completely from the distribution? We still have the WEB_INF files lying around. webapps/ROOT/index.jsp should be updated to remove the jsp and servlet examples link from the index page. On tomcat5 the the WEB-INF files are not present. But the index page may still need to be updated.	0	0	https://issues.jboss.org/browse/JBEWS-386
11263	JBEWS-384	Think this one is already fixed by latest build. The latest build contains README file that contains ... Timestamp: YYYYddmmHHMMSS ... Can you check that number?	1	1	https://issues.jboss.org/browse/JBEWS-384
11264	JBEWS-141	Hi all, I'm assigning this to Eva to include in the Installation Guide. I don't think it is a release note issue, as it is being deprecated in the *next* release. Please let me know if you disagree. Eva, could you please amend the package names in the Install Guide, and add a note warning that jkstatus-ant will be deprecated in the next release? Thank you.	1	1	https://issues.jboss.org/browse/JBEWS-141
11265	JBEWS-141	Verified, found "Deprecation of tomcat-jkstatus-ant" in http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Web_Server/1.0/html-single/Installation_Guide/index.html However missing in Release notes. Permaine, is that ok? http://documentation-stage.bne.redhat.com/docs/en-US/JBoss_Enterprise_Web_Server/1.0/html-single/Release_Notes_1.0.2/index.html	0	0	https://issues.jboss.org/browse/JBEWS-141
11266	JBEWS-141	Where do we usually put 'deprecated items'? In the installation guide or release notes?	1	1	https://issues.jboss.org/browse/JBEWS-141
11267	JBEWS-140	Hi Samuel, sorry I couldn't find the description you've quoted from http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Web_Server/1.0/html-single/Release_Notes_1.0.2/index.html Could you please tell me where you find these instructions?	1	1	https://issues.jboss.org/browse/JBEWS-140
11268	JBEWS-140	Hi Misha, yes you are correct, seems the release notes I've found are not the correct one. I've found a standalone installation guide on our official doc site: [Installation|http://docs.redhat.com/docs/en-US/JBoss_Enterprise_Web_Server/1.0/html/Installation_Guide/index.html] Could you please help to check it.	1	1	https://issues.jboss.org/browse/JBEWS-140
11269	JBEWS-140	Weinan, I'd like to close off the issue identified in this JIRA (sudo not included where required for certain commands in the release notes). Can you confirm: a. Are the commands Samuel is linking to required for the EWS 2 release? b. If yes, which ones require sudo?	1	1	https://issues.jboss.org/browse/JBEWS-140
11270	JBEWS-137	John Doyle <jdoyle@redhat.com> made a comment on [bug 899449|https://bugzilla.redhat.com/show_bug.cgi?id=899449] If we have 1.2.36 in 2.0.0, is this actually an EWS 1 request?	1	1	https://issues.jboss.org/browse/JBEWS-137

#	Requirement	Questions	Answered	Successfully Answered	URL
11271	JBEWS-137	Mandar Joshi <majoshi@redhat.com> made a comment on [bug 899449|https://bugzilla.redhat.com/show_bug.cgi?id=899449] @Wei Nan, can you provide CCFR for this bug?	1	1	https://issues.jboss.org/browse/JBEWS-137
11272	JBEWS-137	Mandar Joshi <majoshi@redhat.com> made a comment on [bug 899449|https://bugzilla.redhat.com/show_bug.cgi?id=899449] Added DocText. @Wei Nan Li, can you please review the Doc Text content? Since bug 900793 is already documented, do we need to document this bug? This bug seems to be dependant/duplicate of bug 900793.	0	0	https://issues.jboss.org/browse/JBEWS-137
11273	JBEWS-135	Jean-frederic Clere <jclere@redhat.com> made a comment on [bug 900492|https://bugzilla.redhat.com/show_bug.cgi?id=900492] Where is validationQueryTimeout missing?	0	0	https://issues.jboss.org/browse/JBEWS-135
11274	JBEWS-135	Jean-frederic Clere <jclere@redhat.com> made a comment on [bug 900492|https://bugzilla.redhat.com/show_bug.cgi?id=900492] Where is validationQueryTimeout missing?	0	0	https://issues.jboss.org/browse/JBEWS-135
11275	JBEWS-132	I guess the first step it get a hudson build. Could you have a look to it.	0	0	https://issues.jboss.org/browse/JBEWS-132
11276	JBEWS-131	What are the exact changes needed and is container/catalina/build.xml the only file that needs updating?	1	0	https://issues.jboss.org/browse/JBEWS-131
11277	JBEWS-131	Hi Alex, Can you please specify the exact changes needed, and which files need to be updated? Thanks, Permaine	1	0	https://issues.jboss.org/browse/JBEWS-131
11278	JBEWS-131	Dave, can you please look into this? Thanks!	1	0	https://issues.jboss.org/browse/JBEWS-131
11279	JBEWS-131	Hi, After reviewing the build files, I think the solution is an Ant build script that is owned by the EWS build. As I see it, this Ant script would consolidate the test locations and run each in a sequence. A property can point to the customer's installation. It would assume the EWS directory structure when defining the classpath. The EWS build would contain this build file and inject it into the distribution. This has not been an issue upstream, so I recommend we make this an EWS 'feature' versus a tomcat5 feature. Does that satisfy the requirements? -- dave	0	0	https://issues.jboss.org/browse/JBEWS-131
11280	JBEWS-131	Dave, QE think that's ok as a fix. Do we have enough time to put this into EWS 1.0.2? If so, please move this back to 1.0.2.	0	0	https://issues.jboss.org/browse/JBEWS-131
11281	JBEWS-125	Can you please unify those installation feature requests into one issue? Like for install, uninstall would need a proper installer, and that needs to get approved by the PM.	1	0	https://issues.jboss.org/browse/JBEWS-125
11282	JBEWS-124	Hi Ondrej, is this a docs issue? Can you please assign it accordingly? Thanks!	1	0	https://issues.jboss.org/browse/JBEWS-124
11283	JBEWS-124	Hi David, can you please include these files in the tomcat builds for EWS 2 (as described by Arun) when you have time? Thanks!	1	0	https://issues.jboss.org/browse/JBEWS-124
11284	JBEWS-122	How about Windows Scripting Host? Javascript syntax, can use ActiveX/COM... viruses are written in it so it's surely powerful :) Regarding 2) and 3), I think e.g. EZinstaler supports launching WSH scripts amongst others. And, how about Java-based installer like IzPack?	1	0	https://issues.jboss.org/browse/JBEWS-122
11285	JBEWS-122	Mladen, any decision on this? Or who should I ask? Permaine?	0	0	https://issues.jboss.org/browse/JBEWS-122
11286	JBEWS-120	I fixed that in the jboss-ews build which only takes care of the RHEL builds, these should really be removed in the tomcat zips which the Hudson builds use. Dave, can you please fix that in the tomcat* builds? Thanks!	1	1	https://issues.jboss.org/browse/JBEWS-120
11287	JBEWS-120	Mladen, can you please get rid of the backup files in our Hudson builds? Thanks!	0	0	https://issues.jboss.org/browse/JBEWS-120
11288	JBEWS-117	Could you please attach spec file and policy file for the rpm(s). I guess RHEL5 and RHEL6 are different.	0	0	https://issues.jboss.org/browse/JBEWS-117
11289	JBEWS-117	Weinan, can you check on the status of this? Should it be closed?	0	0	https://issues.jboss.org/browse/JBEWS-117
11290	JBEWS-114	Mladen, can you please update our Windows/Solaris jboss-ews build to get the jars directly from mod_cluster and jboss-logging brew builds so that we can avoid rebuilding tomcats just because of mod_cluster rebuilds? Thanks!	0	0	https://issues.jboss.org/browse/JBEWS-114
11291	JBEWS-114	Not for EWS 1 as there's no future planned release, probably not worth the effort, can we get this in EWS 2? sorry for the confusion.	0	0	https://issues.jboss.org/browse/JBEWS-114
11292	JBEWS-114	For EWS2 sure. Where should those files end up? In share/java ?	0	0	https://issues.jboss.org/browse/JBEWS-114
11293	JBEWS-113	Mladen, could you please ACK this and re-assing to docs team (Jared Morgan)? Thanks.	0	0	https://issues.jboss.org/browse/JBEWS-113
11294	JBEWS-112	Will we add the default config?	1	0	https://issues.jboss.org/browse/JBEWS-112
11295	JBEWS-8	Weinan Li <weli@redhat.com> made a comment on [bug 901126|https://bugzilla.redhat.com/show_bug.cgi?id=901126] Hi Mladen, could you please patch the win/sol postinstall file (for zip bundle) as I did in attachment? Thanks!	1	1	https://issues.jboss.org/browse/JBEWS-8
11296	JBESB-3948	Max, is it possible to just comment the action (which we want disable) as ESB doesn't support this feature? I've been doing this "workaround" in my projects! e.g.: <action class="org.jboss.soa.esb.actions.SystemPrintln" name="Print "> <property name="message" value="print"/> </action> When we mark this action as disabled, the JBDS just do something like this: <!-- Action Disabled blablablabla <action class="org.jboss.soa.esb.actions.SystemPrintln" name="Print "> <property name="message" value="print"/> </action> --> It's very useful when we are developing/debugging.	0	0	https://issues.jboss.org/browse/JBESB-3948

#	Requirement	Questions	Answered	Successfully Answered	URL
11297	JBESB-3906	means caused by BZ902156. The error [1] reproduced by comment 13 of BZ915386 indicates it stems from the invoked web service not SOAPClient. And I confirmed in my environment that if all BZ915386 patches but CXF upgrade is applied the error doesn't happen. I believe the root cause should be found in the following BZs: - [BZ-858926] [CVE-2012-3451] jbossws-cxf, apache-cxf: SOAPAction spoofing on document literal web services - [BZ-896338] [CVE-2012-5633] jbossws-cxf, apache-cxf: Bypass of security constraints on WS endpoints I also doubt that this is a regression of JBoss ESB or SOA-P as no behaviours in JBoss ESB seem to change; what's changed is JBoss WS. So, should we close this BZ and chase it as a JBoss WS regression instead? Thank you! [1] 15:33:39,174 WARNING [PhaseInterceptorChain] Interceptor for {http://webservice_consumer1/helloworld}HelloWorldWSService#{http://webservice_consumer1/helloworld}sayHello has thrown exception, unwinding now org.apache.cxf.interceptor.Fault: The given SOAPAction sayHello does not match an operation. at org.apache.cxf.binding.soap.interceptor.SoapActionInInterceptor$SoapActionInAttemptTwoInterceptor.handleMessage(SoapActionInInterceptor.java:192) at org.apache.cxf.binding.soap.interceptor.SoapActionInInterceptor$SoapActionInAttemptTwoInterceptor.handleMessage(SoapActionInInterceptor.java:164) at org.apache.cxf.phase.PhaseInterceptorChain.doIntercept(PhaseInterceptorChain.java:243) at org.apache.cxf.transport.ChainInitiationObserver.onMessage(ChainInitiationObserver.java:111) at org.apache.cxf.transport.servlet.ServletDestination.invoke(ServletDestination.java:99) at org.apache.cxf.transport.servlet.ServletController.invokeDestination(ServletController.java:431) at org.jboss.wsf.stack.cxf.ServletControllerExt.invoke(ServletControllerExt.java:173) at org.jboss.wsf.stack.cxf.RequestHandlerImpl.handleHttpRequest(RequestHandlerImpl.java:61) at org.jboss.wsf.stack.cxf.CXFServletExt.invoke(CXFServletExt.java:185) at org.apache.cxf.transport.servlet.AbstractHTTPServlet.handleRequest(AbstractHTTPServlet.java:179) at org.apache.cxf.transport.servlet.AbstractHTTPServlet.doPost(AbstractHTTPServlet.java:103) at javax.servlet.http.HttpServlet.service(HttpServlet.java:637) at org.apache.cxf.transport.servlet.AbstractHTTPServlet.service(AbstractHTTPServlet.java:159) at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:290) at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:206) at org.jboss.web.tomcat.filters.ReplyHeaderFilter.doFilter(ReplyHeaderFilter.java:96) at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:235) at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:206) at org.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:235) at org.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:191) at org.jboss.web.tomcat.security.SecurityAssociationValve.invoke(SecurityAssociationValve.java:183) at org.jboss.web.tomcat.security.JaccContextValve.invoke(JaccContextValve.java:95) at	0	0	https://issues.jboss.org/browse/JBESB-3906
11298	JBESB-3830	From Martin Weiler: The reason I raised this question was to avoid ending up with the same situation as with the SignalCommand of the old jBPM3 BpmProcessor, which was removed in ESB 4.5 (SOA-P 5.0) by JBESB-2315. Some citations from the forum: https://community.jboss.org/thread/144559: The SignalCommand is no longer available, [...] because it is not safe to use it without additional information. Any service invocation from jBPM to ESB and back stores the appropriate information within the EPR and is automatically handled. All the ESB service has to do is respond to the correct EPR. https://community.jboss.org/thread/153723: direct use of the signal command within jBPM 3 will fail very quickly once you move away from the simple scenarios, which is why we jump through those hoops. [...] The present integration focusses on interactions between ESB services and jBPM, forcing a 'one time only' execution of the signal. The EPR contains additional information necessary to enforce that guarantee, i.e. one request and one response.	1	1	https://issues.jboss.org/browse/JBESB-3830
11299	JBESB-3611	Unfortunately what you are proposing is likely to break things. Can you create a test case showing what you are trying to fix?	1	1	https://issues.jboss.org/browse/JBESB-3611
11300	JBESB-3353	Creating a more comprehensive contract model in ESB 5 that would support both XSDs and WSDLs would be very useful. If a JIRA is created in the future to track this feature, would you add a reference to this JIRA? Thanks.	1	0	https://issues.jboss.org/browse/JBESB-3353
11301	JBESB-3259	Thanks Kurt, can you attach an example to this issue along with the WSDL being parsed? It looks like it should be fairly straight forward to fix.	1	1	https://issues.jboss.org/browse/JBESB-3259
11302	JBESB-3248	Do we want people to be using jbossesb-1.1.0.xsd or jbossesb-1.2.0.xsd? This section suggests using 1.1.0.	0	0	https://issues.jboss.org/browse/JBESB-3248
11303	JBESB-3244	The Scripting action talks about how 4.7 ships with bsf 2.3.0 but a future version is going to have 2.4.0. Does 4.8 have 2.4.0? I noticed both 2.3.0 and 2.4.0 in the build dir.	0	0	https://issues.jboss.org/browse/JBESB-3244
11304	JBESB-3197	Is this still an issue?	1	0	https://issues.jboss.org/browse/JBESB-3197
11305	JBESB-3139	Daniel, this needs to go in the SOA docs. I'm about to create a SOA jira for this, can you add the details of this change to that JIRA?	0	0	https://issues.jboss.org/browse/JBESB-3139
11306	JBESB-3130	Committed to JBESB_4_7_CP with revision 31130.	1	1	https://issues.jboss.org/browse/JBESB-3130
11307	JBESB-3130	picketlink jars updated from 1.0.0 to 1.0.1 in product/.classpath Eclipse file. Committed to JBESB_4_7_CP with revision 31141.	0	0	https://issues.jboss.org/browse/JBESB-3130
11308	JBESB-3046	1 - Did not know the ESB 4.7 does not support CXF via JBossWS so fine that this is a Feature Request 2 - I am not reading the log files or even attempting to diagnose the real cause, I was just bringing this to your attention. I leave the diagnosis to you :)	1	1	https://issues.jboss.org/browse/JBESB-3046
11309	JBESB-2941	JBESB-2906 is definitely the same and there is a workaround on that issue. Can you test with the workaround? JBESB-2948 is not the same issue however, although it is related to what needs to be done to v3 in order for this to pass :)	1	1	https://issues.jboss.org/browse/JBESB-2941
11310	JBESB-2937	[Show »] Tom Fennelly - 05/Nov/09 06:15 PM Done. Still http gateway issues with redeploy on AS5. What isssues are you seeing with redploy here Tom?	0	0	https://issues.jboss.org/browse/JBESB-2937
11311	JBESB-2902	This would only update the pre-installed server version, nothing more. Is this what you want? It is still possible to install ESB into an app server and use whichever version of jboss messaging is already present.	1	1	https://issues.jboss.org/browse/JBESB-2902
11312	JBESB-2902	How can this be "out of date" when we still have a "jbossesb-server-4.8.zip" download, which still comes with "jboss.messaging.providerVersion=1.4.0.SP3" ???	1	1	https://issues.jboss.org/browse/JBESB-2902
11313	JBESB-2882	A few comments. 1) "When someone sends in a Message that specifies dispose as a Message property, halt() is called on the session, which kills the spawned thread." Is session.dispose() called. If not, the Message property should be named halt (or something similar). Does re-deploying the service also call halt and dispose of the the session? 2) Is there a mechanism to "publish" the stateful session so that other services could process messages through a second entry point?	1	1	https://issues.jboss.org/browse/JBESB-2882
11314	JBESB-2875	It should also support public key, does this not work for you?	1	1	https://issues.jboss.org/browse/JBESB-2875

#	Requirement	Questions	Answered	Successfully Answered	URL
11315	JBESB-2813	Kevin, would you like me to create a change request for SOA-P5 to include Noel's changes? AND/OR do you think this can go out in the 4.3 CP stream as well?	0	0	https://issues.jboss.org/browse/JBESB-2813
11316	JBESB-2805	Hi Ricardo! Thanks for providing this. I still miss some functionality: What if I have an XML file with multiple items in it and I want to split the message into messages, one for every item. So all resulting messages go to the same destination, but with a different payload. Like a for each loop. That would be a very good thing to include as well! Or is there work already done on this? Cheers Bernd	1	1	https://issues.jboss.org/browse/JBESB-2805
11317	JBESB-2760	4.7 is being targeted against the 5.1.0 release and SOA 5/EAP 5. Do you know if this will be an issue for either of these?	0	0	https://issues.jboss.org/browse/JBESB-2760
11318	JBESB-2760	What is the state of this issue? I've checked the trunk version of jbossesb and did not notice any changes indicating that AS6 support is being worked on...	1	1	https://issues.jboss.org/browse/JBESB-2760
11319	JBESB-2711	First off, I did revert some stuff to not mask the initial problem: Revisions 27961 and 27962. Second, it looks to me like the smooks request mapping is being done correctly: a) See attached [smooks-request-report.html] file. On the Result tab you will see the complexObject details. b) In the debugger, I see that the output map in SmooksMapper.java:104 *does* contain the complexObject. What's intersting, then, is why this exception is spit out in the server console when exercising the test: it.javalinux.wise.core.exceptions.WiseRuntimeException: Error during method invocation! at it.javalinux.wise.core.exceptions.WiseRuntimeException.rethrow(WiseRuntimeException.java:44) at it.javalinux.wise.core.client.WSMethod.invoke(WSMethod.java:88) at it.javalinux.wise.core.client.WSMethod.invoke(WSMethod.java:119) at org.jboss.soa.esb.actions.soap.wise.SOAPClient.process(SOAPClient.java:251) at org.jboss.soa.esb.listeners.message.ActionProcessingPipeline.processPipeline(ActionProcessingPipeline.java:634) at org.jboss.soa.esb.listeners.message.ActionProcessingPipeline.processPipeline(ActionProcessingPipeline.java:586) at org.jboss.soa.esb.listeners.message.ActionProcessingPipeline.process(ActionProcessingPipeline.java:420) at org.jboss.soa.esb.listeners.message.MessageAwareListener$TransactionalRunner.run(MessageAwareListener.java:545) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:651) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:676) at java.lang.Thread.run(Thread.java:595) Caused by: java.lang.IllegalArgumentException: wrong number of arguments at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at it.javalinux.wise.core.client.WSMethod.invoke(WSMethod.java:75) ... 9 more So, the question becomes, "why is WSMethod.invoke having a problem invoking PingWS.java's 'public ComplexObject pingComplexObject(@WebParam(name = "complexObject")' method/operation?" If the SmooksMapper is indeed creating the correct result in applyMapping, why can't WISE map that to the "complexObject" @WebParam?	1	1	https://issues.jboss.org/browse/JBESB-2711
11320	JBESB-2709	Check out jboss-5.1.0.GA/server/default/conf/bindingservice.beans/META-INF/bindings-jboss-beans.xml maybe?	1	1	https://issues.jboss.org/browse/JBESB-2709
11321	JBESB-2700	ejbprocessor... Seems to be an issue with EJB deployment... ... 63 more Caused by: javax.naming.NameNotFoundException: SimpleSLSB not bound at org.jnp.server.NamingServer.getBinding(NamingServer.java:771) at org.jnp.server.NamingServer.getBinding(NamingServer.java:779) at org.jnp.server.NamingServer.getObject(NamingServer.java:785) at org.jnp.server.NamingServer.lookup(NamingServer.java:396) at sun.reflect.GeneratedMethodAccessor272.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:294) at sun.rmi.transport.Transport$1.run(Transport.java:153) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:149) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:466) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:707) at java.lang.Thread.run(Thread.java:595) at sun.rmi.transport.StreamRemoteCall.exceptionReceivedFromServer(StreamRemoteCall.java:247) at sun.rmi.transport.StreamRemoteCall.executeCall(StreamRemoteCall.java:223) at sun.rmi.server.UnicastRef.invoke(UnicastRef.java:126) at org.jnp.server.NamingServer_Stub.lookup(Unknown Source) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:726) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:686) at javax.naming.InitialContext.lookup(InitialContext.java:351) at org.jboss.soa.esb.actions.EJBProcessor.getEjb3FromJndi(EJBProcessor.java:303) ... 66 more EJB jar (SimpleSLSB.jar) not getting deployed??	1	1	https://issues.jboss.org/browse/JBESB-2700
11322	JBESB-2699	Is this the same as https://jira.jboss.org/jira/browse/JBESB-2331 perhaps?	1	1	https://issues.jboss.org/browse/JBESB-2699
11323	JBESB-2621	I have checked the files nsswitch.conf but the problem still continues.. I can make telnet to this specific port: Trying 127.0.0.1... Connected to localhost.localdomain (127.0.0.1). Escape character is '^]'. ¬ísrjava.rmi.MarshalledObject|½—ícü>lhashlocBytest[BobjBytesq~xp)°Ûjur[B¬óTàxp&¬íthttp://AGW-CLA1:8083/q~q~uq~Ã¬ísr org.jnp.server.NamingServer_Stubxrjava.rmi.server.RemoteStubé?ÜÉ‹áexrjava.rmi.server.RemoteObjectÓa´' a3xpw7 UnicastRef2 10.64.168.61J,p!óWbù€xConnection closed by foreign host. Also I can ping the localhost: PING localhost.localdomain (127.0.0.1) 56(84) bytes of data. 64 bytes from localhost.localdomain (127.0.0.1): icmp_seq=0 ttl=64 time=0.028 ms 64 bytes from localhost.localdomain (127.0.0.1): icmp_seq=1 ttl=64 time=0.036 ms 64 bytes from localhost.localdomain (127.0.0.1): icmp_seq=2 ttl=64 time=0.035 ms When trying to integrate our solution with other products I get usually java.rmi.RemoteException: java.lang.NullPointerException Can it be because of this problem? Could you please give more specific instructions to solve the issue?	1	0	https://issues.jboss.org/browse/JBESB-2621
11324	JBESB-2581	Can we add more details to this please?	1	1	https://issues.jboss.org/browse/JBESB-2581
11325	JBESB-2524	What was the intention of the ftpLocalDir setting? From what I can see in the code... there's no config option (in the XSD) to specify the ftpLocalDir configTree parameter, but that was probably never the intention. It doesn't seem to be getting set on the FTP EPR (via ListenerUtil.assembleEpr). Has anyone an idea what it was intended for, considering that the FTPGateway has been working away (in some capacity) without it?	1	1	https://issues.jboss.org/browse/JBESB-2524

#	Requirement	Questions	Answered	Successfully Answered	URL
11326	JBESB-2524	Thanks Kev, but that wasn't the question ;) I asked did anybody know what was the intention of capturing the ftpLocalDir setting? Has anyone an idea why is it being set on the ConfigTree, since I can't see it ever being used, which is what I think Gearge means by "RemoteGatewayListener ignores local FTP location"? If it's not being used for anything, then it's just a case of removing that code perhaps i.e. not even any need to change it to use the System property. Maybe I'm blind and it is being used somewhere, but if it is I'd expect it to be a bit more obvious.	1	1	https://issues.jboss.org/browse/JBESB-2524
11327	JBESB-2524	+1... I was also thinking this, but was half afraid to suggest it :) I will remove the code (from RemoteGatewayListener) that suggests an intention to configure this at the EPR level (but doesn't in reality).	0	0	https://issues.jboss.org/browse/JBESB-2524
11328	JBESB-2395	I've been using this myself with Red Hat JBoss application server SOA and Jboss application server 4.2.3 with no troubles. Offhand I couldn't see a test suite to run this against. Is there a set of tests? Happy to pick up any formal testing required and work through getting this patch complying with any additional requirements. If anything else is needed, please let me know and I'll get on to it. Keep up the good work! Take care, Murray	1	1	https://issues.jboss.org/browse/JBESB-2395
11329	JBESB-2362	Jim.... can you elaborate on this a bit please? Can you outline the use case... what's on the client side... what protocol etc etc?	1	1	https://issues.jboss.org/browse/JBESB-2362
11330	JBESB-2362	When use http protocol to send the content type is "binary/octet-stream" or "application/octet-stream" http request to RemotingGatewayListener, the HttpUnmashaller can not process them correctly. Also should we check other content types support ?	1	1	https://issues.jboss.org/browse/JBESB-2362
11331	JBESB-2362	Jim... can we get an example of this please (e.g. a modified quickstart)? There was a somewhat similar issue fixed on the JBR Gateway recently (JBESB-2616 and http://www.jboss.org/index.html?module=bb&op=viewtopic&t=156733). Would be good to have something concrete to go on before we try resolving it.	0	0	https://issues.jboss.org/browse/JBESB-2362
11332	JBESB-2362	Guys... with the new <http-gateway>... do we need to do this now? Can't we just tell users to use the new <http-gatway> for HTTP?	1	1	https://issues.jboss.org/browse/JBESB-2362
11333	JBESB-2335	Can you elaborate on that one please Kev?	1	1	https://issues.jboss.org/browse/JBESB-2335
11334	JBESB-2335	Sorry? What's the "issue" with how it is currently? :)	1	1	https://issues.jboss.org/browse/JBESB-2335
11335	JBESB-2335	What, you mean bringing maven into this as well as ivy when we can fix the ivy issue and just use that?	1	1	https://issues.jboss.org/browse/JBESB-2335
11336	JBESB-2335	Sorry... can I take this issue Kev?	1	1	https://issues.jboss.org/browse/JBESB-2335
11337	JBESB-2288	What is the context of this JIRA please? Support for what additional use case(s) is being added, or is it a fix for s bug? What will work after applying this patch that didn't work before?	1	1	https://issues.jboss.org/browse/JBESB-2288
11338	JBESB-2288	Mark may know more as he moved it into 4.6, but have you checked out what has been done in the attachments?	1	1	https://issues.jboss.org/browse/JBESB-2288
11339	JBESB-2285	If you say so, i was informed by Jboss respresenative in Antwerp. To raise a JIRA for feature requests. But you seem not to share that opinion. This definitly something you should listen to with respect cause customer can always choose another product than JBoss if it not satifiesn their needs. You feedback you got feedback goday. /Daniel	1	1	https://issues.jboss.org/browse/JBESB-2285
11340	JBESB-2284	Okay, let me check with Trev on this as he was the one who asked. Sounds like he got the versions mixed up. (09:38:48) tkirby: kconner, this is for fp01, the version of mvel that is being used in drools is mvel2-2.0-M4-SNAPSHOT.jar, do we know if mvel 2.2 will behave with smooks?	1	1	https://issues.jboss.org/browse/JBESB-2284
11341	JBESB-2273	How will the Stateful Rule session with an ESB be handled in a Cluster? That is, do we impose a sticky session or will there be an ordered JMS queue backing this?	0	0	https://issues.jboss.org/browse/JBESB-2273
11342	JBESB-2239	We cannot clean anything up in the EPR as it is intended to be an opaque structure. As for the token/process instance id that was deliberately removed to prevent anyone other than jBPM from driving the process. If you have a use case for this, can you raise it on the forums? We can discuss it and decide what is appropriate.	1	1	https://issues.jboss.org/browse/JBESB-2239
11343	JBESB-2235	jBPM 3.3.0 was released back in November, are there issues with it?	1	1	https://issues.jboss.org/browse/JBESB-2235
11344	JBESB-2235	This sounds like the 'enterprise' version of jBPM which we do not yet support in the SOA platform. Is there another version?	1	1	https://issues.jboss.org/browse/JBESB-2235
11345	JBESB-2235	I can't find jbpm-jpdl-3.3.0 only version 3.2.3 so I assumed that I should be using the 3.3.0 download. I'm slightly confused as which version I should be using for this. Do you know which one? Thanks	1	1	https://issues.jboss.org/browse/JBESB-2235
11346	JBESB-2235	Can you please document all the changes you have had to make to get this to work? The SOA guys will, almost certainly, have to go through the same process when they build the next SOA platform and will appreciate any help.	1	1	https://issues.jboss.org/browse/JBESB-2235
11347	JBESB-2234	I could not find a GA release (http://www.jboss.org/drools/downloads.html). Can we use 5.0.0.M4.	1	1	https://issues.jboss.org/browse/JBESB-2234
11348	JBESB-2234	Tried with M4 and found no issues. All test pass. Very smooth :) Should I check this in or should we wait for the GA version?	1	1	https://issues.jboss.org/browse/JBESB-2234
11349	JBESB-2234	We cannot force JTA as this is not available outside of the application server/esb server. Can we create two new issues for these please? The second one appears to be an issue with the embedded console and should be looked at by TomC.	1	1	https://issues.jboss.org/browse/JBESB-2234
11350	JBESB-2233	Updated to CP2_CR3, still have two doc changes to make. Changes in revision 24966.	0	0	https://issues.jboss.org/browse/JBESB-2233
11351	JBESB-2224	Kev, Can you expand on what you mean as a proper gateway? There is currently a GroovyGateway in the codebase but I'm sure you know that. Thanks	1	1	https://issues.jboss.org/browse/JBESB-2224
11352	JBESB-2224	Ok thanks, got it. Last time I checked Netty did not have support for UDP but there is a JIRA for that. How about we create a MinaGateway based on what Tom has done here?	1	1	https://issues.jboss.org/browse/JBESB-2224
11353	JBESB-2217	Just checked the SOA issue, it is now against SOA-4.4GA and not the CP release	1	1	https://issues.jboss.org/browse/JBESB-2217
11354	JBESB-2207	I can see no real benefit to adding this within the gateway, and the implementation is likely to cause issues in a transactional environment. Can you please start a discussion on the forums about your requirements so that we can find out what you really want?	1	1	https://issues.jboss.org/browse/JBESB-2207
11355	JBESB-2179	Can we just use "http" instead of "tomcat"? What about the HTTP response codes for a timeout, ESB exceptions, etc?	1	1	https://issues.jboss.org/browse/JBESB-2179
11356	JBESB-2179	I also think "http-provider" or "http-gateway" is good name for it. But I wonder if that will confuse some users because the JBRGatewayListner is also a kind of "http" gateway. There is some overlap between the name "jbr-provider" and "http-provider". Is it a problem? The HTTP response code "408" is for a timeout and "500" for a ESB exception . We can also customize the response code by adding a property in ESB message to fit for particular requirement .	1	1	https://issues.jboss.org/browse/JBESB-2179
11357	JBESB-2168	One of the issues that we currently have with the current trailblazer is the building of it. We have been getting a lot of question about users having issues building and configuring it. But if the trailblazer has been updated it might be worth looking more closely at the building and making sure it works will the different environments (source project, jbossesb-server download etc). What do you think?	0	0	https://issues.jboss.org/browse/JBESB-2168
11358	JBESB-2148	Isn't that it on page 95 of the Programmers Guide? (Image attached.)	0	0	https://issues.jboss.org/browse/JBESB-2148

#	Requirement	Questions	Answered	Successfully Answered	URL
11359	JBESB-2132	Will this request be realized in the 4.5 release ?	0	0	https://issues.jboss.org/browse/JBESB-2132
11360	JBESB-2121	> The util also relies on having a keystore configured but it would be sufficient to have the key(s) automatically generated on startup and use this to encrypt the >session information. If we generate keys upon startup the keys will only be known to that ESB node. If the message switches to a different node then using that solution will not work. That is why we are currently using a keystore so that users can configure a key pair that can be installed on all nodes and they will be abe to decrypt a security context encrypted by other nodes. Is this not what we want?	1	1	https://issues.jboss.org/browse/JBESB-2121
11361	JBESB-2121	If you mean that I'm encrypting the SecurityContext by "encrypting everything" then yes that is correct. And yes, other systems will not be able to decrypt it, at least not after this task is completed. At the moment other systems can decrypt the security context. Is what you are after that the security context should be encrypted per node, and if the message is routed to a different node, it should re-authenticate. Is that correct?	1	1	https://issues.jboss.org/browse/JBESB-2121
11362	JBESB-2121	The other question that spring to mind with the current code is what is the reason for sending credentials around? At present the pipeline will decrypt the context and, if successful, will do no authentication. If this is unsuccessful then it has no access to the credentials to perform an authentication.	1	1	https://issues.jboss.org/browse/JBESB-2121
11363	JBESB-2121	>Even as it stands we have a situation where only a specific set of servers would be able to reauthenticate, those with the correct credentials in their kestore. If the >message were to be sent to another set of ESBs or even to clients without access to those credentials then they have no means of authenticating them. Yes, this is intentional. The requirement here is that the correct credentials exist in the keystores and that the ESBs have these in common. >The other question that spring to mind with the current code is what is the reason for sending credentials around? The reason is that the subject is used when executing the pipeline. >At present the pipeline will decrypt the context and, if successful, will do no authentication. If this is unsuccessful then it has no access to the credentials to perform an authentication. Yes, that is correct. But this has sort of been a moving target and this was not always the case. I'm really not sure anymore what should be passed between the ESB and what requirements we have for interoperability :(So I definitely think that we should talk and see if we can get the requirements set and then take it from there.	1	1	https://issues.jboss.org/browse/JBESB-2121
11364	JBESB-2121	> Yes, this is intentional. The requirement here is that the correct credentials exist in the keystores and that the ESBs have these in common. I believe it was intentional, but was not what I thought you were doing :) > The reason is that the subject is used when executing the pipeline. But why the credentials? The only aspect used for the authorisation are the Principals associated with the subject and not its credentials. > Yes, that is correct. But this has sort of been a moving target and this was not always the case. Well, lets make sure it isn't moving any more and make it explicit. The introduction of the encryption appears to have removed one ability. > I'm really not sure anymore what should be passed between the ESB and what requirements we have for interoperability :(> So I definitely think that we should talk and see if we can get the requirements set and then take it from there. Definitely, lets get together and flesh everything out once and for all.	1	1	https://issues.jboss.org/browse/JBESB-2121
11365	JBESB-2120	Sorry about that one. Missed checking in jboss-xml-binding.jar that I'd placed in the testlib directory. Should I check it in?	1	1	https://issues.jboss.org/browse/JBESB-2120
11366	JBESB-1996	Hello Mark, the code is not perfect at the moment. It works for my needs, but needs some modifications to be perfect and another namespace. We have to talk, how some things have to be done. Do you want to have the code as is?	0	0	https://issues.jboss.org/browse/JBESB-1996
11367	JBESB-1937	I'm ready to do this. I have my own notes I took when doing the performance work wtih Daniel, as well as other sources I can gleam from - like the wiki, forums, blog, jira and the like. I agree that it would be nice for these things to be consolidated. I am leaning to just updating/augmenting the somewhat dated wiki page here: http://community.jboss.org/wiki/JBossESBPerformanceTuning , then either reference it from the admin guide, or (gulp) copy/paste it. Honestly I don't like the sounds of copy/pasting. I don't think we need a separate performance tuning guide (as a lot of the information is stuff you can get from AppServer tuning docs). I think for ESB-specific tuning, we put it in the admin guide, and reference it in the programmer's guide. Kev, what are your thoughts?	0	0	https://issues.jboss.org/browse/JBESB-1937
11368	JBESB-1936	What did you have in mind here Kev? Are you just talking about exposing the wsdl for the remote WS, or are you talking about creating an actual proxy (on the ESB) to the remote WS endpoint?	1	1	https://issues.jboss.org/browse/JBESB-1936
11369	JBESB-1934	Is it still in question or is it already outdated/deprecated requirement?	0	0	https://issues.jboss.org/browse/JBESB-1934
11370	JBESB-1904	Pushing this until we flesh out the requirements more. What sort of facility(s) do we need to provide to the user to configure the dispose/don't-dispose logic based on rules execution?? This could be messy enough to implement!	1	0	https://issues.jboss.org/browse/JBESB-1904
11371	JBESB-1865	Jervis. I added a quick fix to trunk. Can you try it out there please.	0	0	https://issues.jboss.org/browse/JBESB-1865
11372	JBESB-1809	The changes in JBESB-1064 seem to already be in the CP branch. Diff:ing product/tools/console/management shows no diffs. Did I miss something here?	1	1	https://issues.jboss.org/browse/JBESB-1809
11373	JBESB-1809	Did you check the 4.2.1 CP branch or the 4.4 CP branch?	1	1	https://issues.jboss.org/browse/JBESB-1809
11374	JBESB-1809	Ah, I only checked the 4.4 CP branch. Should these changes be merge into JBESB_4_2_1_GA_CP ?	0	0	https://issues.jboss.org/browse/JBESB-1809
11375	JBESB-1754	Is it enough that we only list the packages that are used in jbossesb.sar and jbossesb.esb in install/Deployment.txt? Perhaps every service should have its own Deployment.txt, or we could have sections in the install/Deployment.txt for the different services. One for smooks, soap, slsb, soap... The reason for this would be that a few jar exist in both the jbossesb.sar and in the services. Listing all the jars and their versions might be confusing as users might not know which one is used where. Any thoughts on this?	0	0	https://issues.jboss.org/browse/JBESB-1754
11376	JBESB-1753	I followed the steps in install/readme.txt and tried a few quickstart. Only thing I found was (install/readme.txt) : Deployment into JBoss AS 4.2.2.GA ================================ - - Install JBoss AS 4.2.1.GA into an appropriate location. + - Install JBoss AS 4.2.2.GA into an appropriate location. Is there more to this jira then that?	1	1	https://issues.jboss.org/browse/JBESB-1753
11377	JBESB-1752	Will this be changed in the trunk too? [ldimaggi@ldimaggi product]$ cat -n build.xml | grep FAQ 305 <copy file="docs/FAQ.txt" todir="${org.jboss.esb.installationdirectory}/docs"/> At revision 20138.	0	0	https://issues.jboss.org/browse/JBESB-1752
11378	JBESB-1707	Created the 4.3.CR1 tag in revision 19907. Will change to 4.3.GA once the CR1 has been built and released.	0	0	https://issues.jboss.org/browse/JBESB-1707

#	Requirement	Questions	Answered	Successfully Answered	URL
11379	JBESB-1683	Hmmm... well, looks like the environment (classloading + deployers) are different for the test case vs. when run in .esb mode. ;) ;) ;) I found a couple bugs: 1. helloworld.bsh should be renamed helloworld.beanshell so our gold ol' BeanShellDeployer doesn't pick it up (although it seems hit or miss, otherwise I would have noticed it before). So, rename the script and change the name in jboss-esb.xml. 2. jsr-14.jar can be removed. It's not needed. 3. Here's the nasty one. You won't actually have access to the declared beans (like "message") in your py + js scripts unless classloading is "just so". Basically, jython.jar and js.jar *cannot* be included in the esb archive. They have to be loaded by the same classloader as bsf.jar is. Sooooo, for all of the following scripts to work, one needs to place jython.jar and js.jar in jbossesb-server-4.2.1GA/server/default/lib/. NOT in scripting-quickstart.esb (so remove them from the quickstart lib dir). The problem I see with this is (a) licensing, and (b) it's a manual step for the person running the quickstart. Any ideas? 4) Here are the new scripts I tested that work just dandy when jython.jar and js.jar are next to bsf.jar and bsh.jar in jbossesb-server-4.2.1GA/server/default/lib/. : hellworld.bsh: print("**** Hello World from BeanShell! ****"); print(message) helloworld.gy: println "**** Hello World from Groovy! ****" println message helloworld.js: importClass(java.lang.System); System.out.println("**** Hello World from JavaScript! ****"); System.out.println(message.toString()); helloworld.py: print "**** Hello World from Python! ****" print message.toString()	1	1	https://issues.jboss.org/browse/JBESB-1683
11380	JBESB-1683	Well, it's been horrible that EAP has been stuck on bsh-1.3.0 for so long. Maybe I should put in a JIRA on EAP for them to upgrade to bsh-2.0b4? That version is seriously the best it's gonna get in the foreseeable future. Now, if this is a roadblock, then I would still like to see the ScriptingAction getting added. I can just remove the bsh upgrade piece (which would mean removing it from the unit tests and the quickstart too, I guess, considering I'm pretty sure it will cause classloading collisions). Hmm... that gets me thinking. Is it okay that I added rhino (javascript) and jython jars to the quickstart? I'm thinking probably so considering I'm not adding them to the underlying platform...	0	0	https://issues.jboss.org/browse/JBESB-1683
11381	JBESB-1673	I took a look in the web.xml for the jbpm-console.war see that there are three login-configs there. Two are commented out. I actually looked and thought that BASIC authentication was in use, but it is FORM based authentication which is the default. This has an impact on the upload servlet with regard to authentication. Should the ant task be able to support both authentication methods or just one of them? Thanks, Daniel	1	1	https://issues.jboss.org/browse/JBESB-1673
11382	JBESB-1673	Ignore my previous comment. I've got it working will all three login-config authentication options. I do find the commenting in web.xml a little hard to read though. For example this is what one of the login-config elements looks like: <!-- Login configuration option #2 - use basic auth ==> <login-config> <auth-method>BASIC</auth-method> <realm-name>jBPM Administration Console</realm-name> </login-config> <!== End Login configuration option #2 --> If you look at this quickly you might think that this login-config is in use, when it is in fact commented out. Anyone else think that this is somewhat hard to read, or is it just me? :)	1	1	https://issues.jboss.org/browse/JBESB-1673
11383	JBESB-1641	Why is this different from JBESB-1584?	1	0	https://issues.jboss.org/browse/JBESB-1641
11384	JBESB-1641	Did I say it was different? ;-)	0	0	https://issues.jboss.org/browse/JBESB-1641
11385	JBESB-1574	Was this issue really resolved? I cannot find any real commits to the repository in this issue. Also, FreeMarker transformation still doesn't seem to work as I see the following transformation result in SOAP Header. I use JBoss ESB 4.11 (SOA-P 5.3.0). {code} <soapenv:Header><text><wsse:Security soapenv:mustUnderstand="1" xmlns:wsse="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-secext-1.0.xsd"> <wsse:UsernameToken> <wsse:Username>kermit</wsse:Username> <wsse:Password>thefrog</wsse:Password> </wsse:UsernameToken> </wsse:Security> </text></soapenv:Header> {code}	0	0	https://issues.jboss.org/browse/JBESB-1574
11386	JBESB-1513	Thanks for looking into that! I've some questions though: --------- We don't need ESB bits to be copied to the test profile for SOA-P tests. They actually must not be copied to have valid results as the idea is to test SOA-P, not sources. As well SOA-P uses JBM so JBESB-1524 should not be an issue. So I tried running the tests by changing these properties: chprop src/qa/quickstarts/quickstarts.properties org.jboss.esb.server.home "$WORKSPACE/jbosssoa/jboss-as" chprop src/qa/quickstarts/quickstarts.properties org.jboss.esb.server.config $SERVER_NAME chprop product/build/jbossesb/samples/quickstarts/conf/quickstarts.properties org.jboss.esb.server.home "$WORKSPACE/jbosssoa/jboss-as" chprop product/build/jbossesb/samples/quickstarts/conf/quickstarts.properties org.jboss.esb.server.config $SERVER_NAME chprop product/build/jbossesb/install/deployment.properties org.jboss.esb.server.home "$WORKSPACE/jbosssoa/jboss-as" chprop product/build/jbossesb/install/deployment.properties org.jboss.esb.server.config $SERVER_NAME Where SERVER_NAME is "default" or "all" because for the embedded SOA-P distro we need that profile tested. What am I doing wrong? --------- What about the "product" directory tests? Can they run with test.classpath set to external bits?	1	1	https://issues.jboss.org/browse/JBESB-1513
11387	JBESB-1513	Tom, did you modify http://hudson.qa.jboss.com/hudson/job/jboss-soa-platform-ESB-testsuite to test that? Should I assume it is a working version?	1	0	https://issues.jboss.org/browse/JBESB-1513
11388	JBESB-1513	Not sure what to modify - what targets is that version of hudson running for the jboss-soa-platform-ESB-testsuite? All of the code changes here have gone into project code - I don't think we've touched any platform files.	0	0	https://issues.jboss.org/browse/JBESB-1513
11389	JBESB-1498	Hiya Naveen. This code has just been rewritten (checked into trunk yesterday) so may no longer be an issue. Can you attach a log file showing this error? I would like to look through it and make sure it is no longer applicable. Thanks.	1	1	https://issues.jboss.org/browse/JBESB-1498
11390	JBESB-1498	Hiya Naveen. It looks like you may have specified the HA-JNDI configuration as the default JNDI configuration. The result of this is that any name which cannot be located within the HA-JNDI will cause a recursive invocation (and a StackOverflow). What is your setup? How have you configured the server?	1	1	https://issues.jboss.org/browse/JBESB-1498
11391	JBESB-1498	Is this an issue still Kev?	1	0	https://issues.jboss.org/browse/JBESB-1498
11392	JBESB-1457	But let suppose we have application that communicates with ESB via gateway. How we allow this app to specify replyTo address via JMS header? I suppose we want to decouple apps from ESB implementation.	0	0	https://issues.jboss.org/browse/JBESB-1457
11393	JBESB-1457	when the eclipse ide plugin will support those changes?	1	1	https://issues.jboss.org/browse/JBESB-1457
11394	JBESB-1457	Are you asking about the jboss-esb.xml editor plugin or the jpdl editor plug-in?	0	0	https://issues.jboss.org/browse/JBESB-1457
11395	JBESB-1404	Daniel, could you take a look at this? We may want to update the EPR definition too.	1	1	https://issues.jboss.org/browse/JBESB-1404
11396	JBESB-1404	Thanks Daniel. Can you update the docs if necessary too?	1	1	https://issues.jboss.org/browse/JBESB-1404
11397	JBESB-1404	Could you commit this to the CP branch too? Change the Fix Field if you do.	1	1	https://issues.jboss.org/browse/JBESB-1404
11398	JBESB-1389	Kev, can you close this with an Assignee field please?	1	1	https://issues.jboss.org/browse/JBESB-1389

#	Requirement	Questions	Answered	Successfully Answered	URL
11399	JBESB-1366	stop-server: [echo] Stopping local ESB server [java] Shutdown message has been posted to the server. [java] Server shutdown may take a while - check logfiles for completion [echo] Waiting for ESB server to stop... At this time I only see 1 java process. Presumably the integration build itself. So it is waiting on nothing at that point. Where is it waiting for exactly?	1	1	https://issues.jboss.org/browse/JBESB-1366
11400	JBESB-1366	I added a timestamp and now include the specific log file being searched. Revision is 17083. Can you run the following test? - cd product - ant start-server - ant stop-server	1	1	https://issues.jboss.org/browse/JBESB-1366
11401	JBESB-1366	From what I can see the filename looks correct. Can you see what happens if you use ant 1.6.5? (this is the version you should be using) Also, are you running from within cygwin or through a command prompt?	1	1	https://issues.jboss.org/browse/JBESB-1366
11402	JBESB-1366	I've tried modifying esb-config.xml with this: <waitfor maxwait="@{seconds}" maxwaitunit="second" checkevery="5" checkeveryunit="second" timeoutproperty="shutdown.timeout"> <isfileselected file="@{serverHome}/server/@{serverConfig}/log/server.log"> <contains text="[org.jboss.system.server.Server] Shutdown complete"/> </isfileselected> </waitfor> This works for me with Ant version 1.6.5 and with version 1.7.0. Kurt or Kevin, could you give this a try to verify that this does indeed work before I check anything in? Thanks, Daniel	1	0	https://issues.jboss.org/browse/JBESB-1366
11403	JBESB-1365	Daniel, Can you talk with me if this is based on the current plugin? There is some tidying up that needs to be done with the current source before we develop it further.	1	1	https://issues.jboss.org/browse/JBESB-1365
11404	JBESB-1365	I feel a little stupid asking this, but where can I find the Eclipse plugin ? :) I didn't know that there was one. Or did you mean the config editor?	1	1	https://issues.jboss.org/browse/JBESB-1365
11405	JBESB-1283	Does that mean you fixed it? Can't wait to hear what I did wrong :-)	1	1	https://issues.jboss.org/browse/JBESB-1283
11406	JBESB-1251	With Smooks 1.0 it is possible to use "Virtual models" (maps in maps, with lists, etc...), so is this still an issue?	1	1	https://issues.jboss.org/browse/JBESB-1251
11407	JBESB-1181	This should already be covered by ObjectMapper. Do you have an example of where this doesn't work?	1	1	https://issues.jboss.org/browse/JBESB-1181
11408	JBESB-1177	Which console?	0	0	https://issues.jboss.org/browse/JBESB-1177
11409	JBESB-1121	Tom, could this apply to the new JOPR embedded console for the new message alerts capability? If not then we should "retire" this request.	0	0	https://issues.jboss.org/browse/JBESB-1121
11410	JBESB-1072	Is the preferred fix to move everything to the JBoss logger or the Apache logger?	1	1	https://issues.jboss.org/browse/JBESB-1072
11411	JBESB-1041	org.jboss.soa.esb.services.routing.cbr.ContentBasedRouter is abstract (I'm going to change it to an interface now that some actual code moved into the ServiceInvoker. All CBR need to implement this interface. This allows for plugging in other rules engines then just Drools. org.jboss.soa.esb.actions.ContentBasedRouter is the action that delegates the calls to the class behind the interface. Why is this confusing? --Kurt	1	1	https://issues.jboss.org/browse/JBESB-1041
11412	JBESB-1006	The labs page was updated with new binary/src zips in revision 175439. The labs site is currently down so I have been unable to verify this. Coincidence?	0	0	https://issues.jboss.org/browse/JBESB-1006
11413	JBESB-1000	What's this one about?	1	1	https://issues.jboss.org/browse/JBESB-1000
11414	JBESB-997	Which console? Makes my life easier when I re-assign the task ;-)	0	0	https://issues.jboss.org/browse/JBESB-997
11415	JBESB-939	Release Notes for JBoss ESB Includes versions: 4.2 ** Feature Request * [JBESB-177] Add SQLListener to the Trailblazer demo * [JBESB-182] Web Services support * [JBESB-183] JCA support * [JBESB-198] RMI interface for persistence store service * [JBESB-229] Add "Update" capability on all Admin Console wizards. * [JBESB-239] tools module * [JBESB-374] Getting Started Project Template for Eclipse/Ant * [JBESB-415] JCA container support. * [JBESB-429] Clustering: load balancing * [JBESB-430] Clustering: fail-over * [JBESB-431] Hot redeployment * [JBESB-434] Hibernate listener * [JBESB-462] Oracle AQ JMS Provider * [JBESB-505] Add external webservice invocation action * [JBESB-525] FTPGatewayListener: refactor extract rename functionality. * [JBESB-538] Add method to print entire message for debugging purposes in SystemPrintln action * [JBESB-595] Easy proxy for Spring-POJO * [JBESB-597] Easy proxy for WS * [JBESB-639] Update "stub" so that it throws propagated faults as exceptions * [JBESB-684] Richer message structure/hierarchy * [JBESB-725] Create Scheduled Event Provider/Listener * [JBESB-741] Allow default message type to be overridden at runtime * [JBESB-742] Allow for different Message views of the payload at creation time * [JBESB-761] Add control message type * [JBESB-763] Add section on services within the programmer's guide * [JBESB-806] ServiceInvoker should ignore To field * [JBESB-813] Update the documentation to more clearly describe what we can and can't do in terms of SOA for this release. * [JBESB-816] Clearly define Message in relationship to development approach * [JBESB-822] Change the <simple-schedule> "repeatCount" to be an "execCount" * [JBESB-831] Add an application ability to remove messages from the RDLVR queue * [JBESB-859] Add Message best practices/hints-and-tips to the documentation * [JBESB-877] Streaming transformation? ** Bug * [JBESB-192] Listener implementations need to perform the payload to Message wrapping, not the AbstractListener * [JBESB-268] Make the Gateway Controller's creation + initialisation an atomic opp * [JBESB-270] JMSQueueListener doesn't manage Thread creation for the ActionProcessingPipeline i.e. can be flooded with messages * [JBESB-619] Catch NPE in header re-assembly * [JBESB-632] Util was masking errors during serialization/deserialization * [JBESB-641] ConnectionPools are not closed * [JBESB-642] JBossRulesRouter does not dispose of WorkingMemory * [JBESB-647] JAXRRegistryService doesn't unpack "cause" exception after proxied invocation of Registry * [JBESB-650] Investigate why HTTPS isn't working for the SOAPClient * [JBESB-688] quickstart jbpm_simple1 remove or rename * [JBESB-694] ERROR [ClosedInterceptor] ClosedInterceptor.ClientConsumerDelegate * [JBESB-699] ConnectionManager interface should not be in internal package * [JBESB-705] deadletter quickstart modifications * [JBESB-709] jbpm CommandInterpreter action fails silently within doTheJob * [JBESB-710] remoting error messages when running bpm_orchestration2 as standalone * [JBESB-721] Failure on redeployment of helloworld_hibernate_action * [JBESB-727] SOAPClient Action & .NET Endpoints * [JBESB-728] MessageFilter.process duplicates the code from ContentBasedWiretap.process * [JBESB-734] bpm-based quickstarts don't allow for a signal from jBPM 3.2.1 Admin Console * [JBESB-739] jboss-esb-console is registering over an existing MBean * [JBESB-743] ServiceInvoker ignores ReplyTo * [JBESB-747] Changed body contents is not propagated * [JBESB-749] Invalid events in bpm_orchestration1 and bpm_orchestration2 samples * [JBESB-750] Process definition for bpm_orchestration2 contains refernce to class from bpm_orchestration1 * [JBESB-752] Oracle create SQL should be updated * [JBESB-753] Add a Quickstart for NotifyFiles * [JBESB-755] Upgrade jboss-esb-console to seam 1.2 * [JBESB-756] Contextualised information not being used in event driven gateways * [JBESB-766] JBossWS RequestHandlerImpl "create WSDL" method assumes the invocation context is based on a Servlet context * [JBESB-774] helloworld_hibernate_action - ClassCastException * [JBESB-794] The getting	0	0	https://issues.jboss.org/browse/JBESB-939
11416	JBESB-908	This is in trunk, and will be in the 4.2 ESB GA release.	0	0	https://issues.jboss.org/browse/JBESB-908
11417	JBESB-878	I've added a comment about using a clustered DB backend to the docs. Want to do/add anything more Kurt?	0	0	https://issues.jboss.org/browse/JBESB-878
11418	JBESB-875	Do you think it's worth including them, or a reference?	1	0	https://issues.jboss.org/browse/JBESB-875
11419	JBESB-858	Kurt, can you remove the WS patch from the zip. This should be done through the build process and is covered by a separate JIRA.	0	0	https://issues.jboss.org/browse/JBESB-858
11420	JBESB-831	The DLQ is not used to redelivery (the RDLVR is). So once it is in the DLQ is truly dead. I had thought it'd be a manual process to empty the DLQ, but I'm happy to add a button somewhere. Where would you like the button?	1	1	https://issues.jboss.org/browse/JBESB-831

#	Requirement	Questions	Answered	Successfully Answered	URL
11421	JBESB-823	K: Yeah I agree. So I added a new column called 'classification', which basically enables use to have more types of stores, rather then setting up different db schemas (or even tables). The different types are STORE (this is store as intended), and I added DLQ and RDLVR. So they function independently from one another. Does that answer your question? M: Sort of: I think that addresses the back-end implementation issue. But I was wondering about separate services: they may get configured to be the same implementation, but from a users perspective they could still remain separate. That way if someone wants to have a separate MessageStore because of security implications, we can support that. K:The only problem is that we only have one place to configure the store credentials which is in the jbossesb-properties.xml. A while ago a figured we should move it out and put it with the service configuration, but since there was no real hard reason to make clutter up the service configuration I left it as it was. So it's a config issue only. That said you could run different stores on different nodes and get around it that way. M:OK, so it's possible at the moment, just not ideal. I'll add a JIRA task so we don't forget this conversation and then we can get back to it in a later revision. K:Sounds good. Or you can make the messageStore use the JCR, and the DLQ/RDLVR the DBimpl.	0	0	https://issues.jboss.org/browse/JBESB-823
11422	JBESB-802	Can you explain what you're thinking in a little more detail, using code examples. The ServieInvoker doesn't expose anything physical to the calling code. That was the whole point of it. To hide the EPR + Courier "stuff". The Message object is a different kettle of fish ;-)	1	0	https://issues.jboss.org/browse/JBESB-802
11423	JBESB-802	I might add... I think you're hitting on something that's dear to my heart ;-) ... i.e. the idea that using the physical EPR on the Message (for replyTo etc) is a broken concept (in spite of that WS-A may or may not say). It's anything but loose coupling! If I set an EPR for the replyTo, how can I be sure that that EPR will work for the "other guy". If I set a cat:name, the "other guy" can use that to get all the EPRs for that service and select the one that works best for him. High availability etc.	1	0	https://issues.jboss.org/browse/JBESB-802
11424	JBESB-725	Tom, can you keep this as a separate issue for now. Mark had this down as a 5.0 requirement and I would like to hear his reasoning before development starts.	0	0	https://issues.jboss.org/browse/JBESB-725
11425	JBESB-725	Grrrr... why do JBoss insist on disabling features on software like the forums/Confluence/JIRA... drives me bats!!!	0	0	https://issues.jboss.org/browse/JBESB-725
11426	JBESB-677	Is this available right now? Where do I get it and are there official builds available?	1	1	https://issues.jboss.org/browse/JBESB-677
11427	JBESB-677	OK what do you mean "use"? - In qa? - On our developer boxes? - for CI?	1	1	https://issues.jboss.org/browse/JBESB-677
11428	JBESB-676	Can you give pointers to the places that remain?	1	1	https://issues.jboss.org/browse/JBESB-676
11429	JBESB-675	Since this is a public class, all we can do in 4.2 is deprecate it (which you've done). What about documentation? If Invoker is used in the docs then we should update those too.	1	1	https://issues.jboss.org/browse/JBESB-675
11430	JBESB-675	I checked the docs and ServiceInvoker is discussed in the context of JCA. Can you pull this out and create a separate section on it? Since we talk a lot of EPRs and Couriers, it's necessary to have a separate section on ServiceInvoker.	1	1	https://issues.jboss.org/browse/JBESB-675
11431	JBESB-668	Moving this to 5.0 because of recent JON information.	1	1	https://issues.jboss.org/browse/JBESB-668
11432	JBESB-658	Maybe the JMS expiration value should also be carried with the Message object. This might make sence if one there is no point in delivering the message after the expiration time has expired. In the current impl the JMS Message would be save for ever is it is persistent. Mark, do you have any suggestions as to where this infomation should be stored within the Message object? /Daniel	1	1	https://issues.jboss.org/browse/JBESB-658
11433	JBESB-658	Mark, we discussed this ealier about the URI for RelatesTo and I want to make sure I get this right. This is how I'm creating the URI for the correlationID : URI correlationURI = new URI(JMSEpr.JMS_PROTOCOL, "correlationID" , from.getJMSCorrelationID()); which produces an uri like this: jms:correlationID#YYXX-123456780-GG JMSCorrelationID is set as a fragment on the URI and can be retrieved like this: correlationURI.getFragment(); This correlation id is not used in the ESB code base as the default behaviour is to set outgoing JMSCorrelationID's to the ESB MessageID. But users can override this and could then set the outgoing correlation id to the original correlation id. And this is how the JMSMessageID is being set on the Call object instance: to.getHeader().getCall().setMessageID(new URI(from.getJMSMessageID())); Does this look right to you? Also I'm setting the JMSExpiration as a property on the ESB Message. Thanks, Daniel	1	1	https://issues.jboss.org/browse/JBESB-658
11434	JBESB-655	Marc, Please do provide us with your requirements and thoughts as it relates to this feature. It hasn't been requested all that often (notice: Votes: 0 in Jira at the time this comment is being written). Is your use case C++ remote clients taking to JBoss-hosted server side objects? And now you would like to hook an ESB in between the remote client and the server side object as a mediation point? Perhaps bring this discussion to the Design Forums.	0	0	https://issues.jboss.org/browse/JBESB-655
11435	JBESB-654	Would a JBI Gateway essentially be an implementation of a "Service Engine" that acts as a bridge between us and the NMR? OR would a JBI Gateway act like a BC container?	0	0	https://issues.jboss.org/browse/JBESB-654
11436	JBESB-652	Is it alright if I check this in? Perhaps I should add a note about this somewhere in the docs. Users might have different correlation strategies and having an example of how they can override the JMSRouter might help them. Any ide in which doc this would go into? Thanks, Daniel	1	1	https://issues.jboss.org/browse/JBESB-652
11437	JBESB-651	Shouldn't JMSReplyTo be related to the ReplyTo in the ESB Message header?	1	1	https://issues.jboss.org/browse/JBESB-651
11438	JBESB-651	But EPRs support MessageID and RelatesTo (equivalent of correlation identifier). Can we not use them, rather than putting message related information within the message body?	1	1	https://issues.jboss.org/browse/JBESB-651
11439	JBESB-646	I think we need to fix this, because it will affect performance. Is this something that could be specified in the EPR, for instance?	1	1	https://issues.jboss.org/browse/JBESB-646
11440	JBESB-646	Assuming this has a jboss-esb.xml impact. Would we change the jms-bus? <jms-bus busid="quickstartGwChannel"> <jms-message-filter dest-type="QUEUE" dest-name="queue/quickstart_business_service_Request" persist="true" <!-- or something like that --> /> </jms-bus> OR Could it be an out-of-the-box action that marks the message as needing to be persisted? OR something else entirely different?	1	1	https://issues.jboss.org/browse/JBESB-646
11441	JBESB-646	So should this persistent property be added to JMSEpr in the same way that the destinationType, message selector etc, are? I'll go through and check the pointer where we are sending JMS messages. This one I currently know about are: 1. JMSCourier 2. JMSRouter 3. NotifyQueues 4. NotifyTopics Any others I should know about?	1	1	https://issues.jboss.org/browse/JBESB-646
11442	JBESB-646	Mark made me aware that Persistent, TTL, and Priority are QoS aspects and should not go into the JMSEpr. So, I've been poking around and the found that the JMSCourier is created with an instance of a JMSEpr. So I need some other way of making such QoS aspects available to the the courier. I'm getting the feeling that this might be a bigger issue then the original jira was intended for. A QoS config should be usable will all types of transports. The properties above are easily mapped to JMS but for other transports like ftp something has to be put in place. Any thought on this? I've implemented the three configuration properties above for JMSRouter and NotifyJMS which are independent of the solution for the JMSCourier above. Would it be alright to check this in or should I create a separate JIRA for one of these issues? Thanks /Daniel	1	1	https://issues.jboss.org/browse/JBESB-646

#	Requirement	Questions	Answered	Successfully Answered	URL
11443	JBESB-646	To take a pragmatic approach, can we do this within the config as originally thought?	1	1	https://issues.jboss.org/browse/JBESB-646
11444	JBESB-646	Should I add Priority and TimeToLive aswell of leave them until we find a better place for these values?	0	0	https://issues.jboss.org/browse/JBESB-646
11445	JBESB-626	Is this done? Can the issue be closed or Stop Progress?	0	0	https://issues.jboss.org/browse/JBESB-626
11446	JBESB-623	When can we close this?	1	1	https://issues.jboss.org/browse/JBESB-623
11447	JBESB-623	Issue 1: fixed Issue 3: Almost fixed. I have installed hudson's emma plugin. Does esb testsuite generate emma report by default? If not, please let me know how you currently generate report and where is report file. Issue 2: Our patch has not been made to upstream yet. In worst case, I will use our local build of hudsdon containing this fix. Issue 4: No update. Please let me know on issue 3 and I will proceed accordingly.	1	1	https://issues.jboss.org/browse/JBESB-623
11448	JBESB-623	I haven't seen any Hudson build notifications for the past 3 days. What's up?	1	1	https://issues.jboss.org/browse/JBESB-623
11449	JBESB-623	Still waiting for a status update. Prabhat?	1	1	https://issues.jboss.org/browse/JBESB-623
11450	JBESB-601	Can you link the other JIRA as a dependency for this issue?	1	1	https://issues.jboss.org/browse/JBESB-601
11451	JBESB-597	Is this covered by http://wiki.jboss.org/wiki/Wiki.jsp?page=ExternalWebserviceConsumption and http://jira.jboss.com/jira/browse/JBESB-588?	0	0	https://issues.jboss.org/browse/JBESB-597
11452	JBESB-597	Can we close this now?	1	1	https://issues.jboss.org/browse/JBESB-597
11453	JBESB-595	Can we close this?	1	1	https://issues.jboss.org/browse/JBESB-595
11454	JBESB-581	How would I know a service is stateless?	1	1	https://issues.jboss.org/browse/JBESB-581
11455	JBESB-576	http://wiki.jboss.org/wiki/Wiki.jsp?page=JBossESBHotDeployment In which guide should we dupe this info?	0	0	https://issues.jboss.org/browse/JBESB-576
11456	JBESB-575	So which of the two classes goes?	0	0	https://issues.jboss.org/browse/JBESB-575
11457	JBESB-573	Great, I'll add that fix and the other one and check it in. So, are you working of converting scout to use JAXB and perhaps I should begin working on using JAXB to generate the configuration bindings? Wasn't sure if that was what you were implying :)	1	1	https://issues.jboss.org/browse/JBESB-573
11458	JBESB-573	Sounds good. I'll check in what I have currently done and then create a JIRA for the task. This will be an updated xbean.jar and javadoc fixes in the code base. I tried to add your patch as well but it still complained...could be that they have added more code as I downloaded the current stable release. I'm using version 2.2.0. Just thought that if we are going to use JAXB instead then this might be acceptable to see the warnings and save some time. What do you think?	1	1	https://issues.jboss.org/browse/JBESB-573
11459	JBESB-572	I'm to blame for the unit tests that use an embedded ftp server. Sorry about that. Should I remove/rename the tests do you think? I have a few more in my workspace for http://jira.jboss.com/jira/browse/JBESB-533, but they have not been checked in to the trunk yet. thanks, Daniel	1	1	https://issues.jboss.org/browse/JBESB-572
11460	JBESB-572	Kevin - I just checked out revision 13788 and am still seeing this error with no FTP server configured. Am I missing something? compile: [mkdir] Created dir: /opt/ESB_July25/trunk/product/samples/quickstarts/static_router/build/classes [javac] Compiling 2 source files to /opt/ESB_July25/trunk/product/samples/quickstarts/static_router/build/classes assert-ftp-config: BUILD FAILED /opt/ESB_July25/trunk/product/samples/quickstarts/tests/build.xml:173: The following error occurred while executing this line: /opt/ESB_July25/trunk/product/samples/quickstarts/tests/build.xml:141: The following error occurred while executing this line: /opt/ESB_July25/trunk/product/samples/quickstarts/conf/base-build.xml:435: Please configure ftp properties in quickstart.properties	1	0	https://issues.jboss.org/browse/JBESB-572
11461	JBESB-569	Yeah it'd be nice if you name the project "ESBTest" that it builds an ESBTest.esb archive. I think the classnames and package can be set to some default like: my.action.package.HelloWorldAction as far as the jbossall-client.jar goes we should just find the smallest jar in the AS/client jar that fits our need. I think we only need some JMS API classes. Put our jars in a Library (rules and jbpm do this too) "JBESB Library", I think we can add the JBM messaging jar on the classpath so we're ready to go. Should we make the classes in mypackage.test to be real log4j unittests? So you have somekind of validation you're starting from something that works? build.properties should be defaulted to: org.jboss.esb.server.home=/jbossesb-server-4.2.GA org.jboss.esb.server.config=default update the readme the juddi.properties file can be deleted, since the registry is configured to use an RMI connection.	0	0	https://issues.jboss.org/browse/JBESB-569
11462	JBESB-563	Do you have something specific in mind here Mark?	1	1	https://issues.jboss.org/browse/JBESB-563
11463	JBESB-562	Sorry, I noticed that won't work:) It works with the test class but will not be ok in the jboss-esb.xml file. How about: <property name="aliases"> <alias name="vehicles" class="java.util.ArrayList"/> <alias name="vehicle" class="Vehicle"/> <alias name="serviceType" class="ServiceType"/> <alias name="phonenumbers" class="java.util.ArrayList"/> <alias name="phonenumber" class="PhoneNumber"/> <alias name="comments" class="java.util.ArrayList"/> <alias name="comment" class="Comment"/> </property> I have commited code that works with this configuration (same workspace as mentioned above) and this time tested it in a the jboss-esb.xml. Thanks, /Daniel	1	1	https://issues.jboss.org/browse/JBESB-562
11464	JBESB-560	Derek, can this be closed now? I thought I saw your code added a while back?	1	1	https://issues.jboss.org/browse/JBESB-560
11465	JBESB-560	Thanks. Did this need any associated documentation updates?	0	0	https://issues.jboss.org/browse/JBESB-560
11466	JBESB-555	Hasn't this been done? I thought I'd seen the move happen last week?	1	1	https://issues.jboss.org/browse/JBESB-555
11467	JBESB-546	Is this important for 4.2, or should we push this to 5.0?	1	1	https://issues.jboss.org/browse/JBESB-546
11468	JBESB-546	No longer supported under 4.x. The 5.x work will cover this.	0	0	https://issues.jboss.org/browse/JBESB-546
11469	JBESB-544	Can you make sure these are documented?	0	0	https://issues.jboss.org/browse/JBESB-544

#	Requirement	Questions	Answered	Successfully Answered	URL
11470	JBESB-533	I have added a strategy so that it is possible to specify different strategies for a read-only ftp server (comments removed): public interface RemoteFileSystemStrategy { void init(String configFile) throws Exception; File[] filterFileList(File[] files) throws GatewayException; boolean deleteFile(File file) throws GatewayException; boolean renameFile(File from, File to) throws GatewayException; File getWorkFileName(File file, String suffix); } The default implementation does the same as what the RemoteGatewayListener did before. RemoteGatewayListener just delegates now. This lets us specify different implementations in the esb configuration file. Every time that RemoteGatewayListeners's getFileList method is called it calls the strategies filterFileList method. I have written an implementation (ReadOnlyRemoteFileSystemStrategy.class) that uses JBoss Cache to cache the filenames, and replicate to other members in the current cache cluster. This also has a cache loader configured so that the file list is persisted. The cache can be configured so that the files are only stored in the cache for a certain amount of time (jgroups configuration). In our case the files on the ftp server will be archived every day. Having such a solution will prevent the esb from working with the same file more then once. I think that may be others have similar situations and perhaps this is something that should be checked in? I have not yet changed the ftp-message-filter to have the property read-only. If read-only is true the ReadOnlyRemoteFileSystemStrategy could be used, but this could also be overrridden in the config file if a different impl is to be used. Should I add this property?	1	1	https://issues.jboss.org/browse/JBESB-533
11471	JBESB-533	I'll check this in by the end of the week. I need to add a way to specify the TreeCacheListener to be used. I don't quite know where to put the documentation for this. Where do you think if should put this? Thanks, Daniel	1	1	https://issues.jboss.org/browse/JBESB-533
11472	JBESB-526	Can you prune lib/ext as and when those services get broken out?	1	1	https://issues.jboss.org/browse/JBESB-526
11473	JBESB-523	For a start, can you please let me what you currently do to build and run testsuite? Where is the repository?	1	1	https://issues.jboss.org/browse/JBESB-523
11474	JBESB-523	Kevin, Do you need a running instance of jboss for tets to run? I got following error: clean-jms: [get] Getting: http://localhost:8080/jmx-console/HtmlAdaptor?action=invokeOpByName&name=jboss.mq.destination:service=Queue,name=A&methodName=removeAllMessages [get] To: /qa/services/hudson/hudson_workspace/workspace/jboss-esb-linux/trunk/clean.jms [get] Error opening connection java.io.FileNotFoundException: http://localhost:8080/jmx-console/HtmlAdaptor?action=invokeOpByName&name=jboss.mq.destination:service=Queue,name=A&methodName=removeAllMessages If that's the case, it would have been nice if you could tell the build where jboss is installed and it started jboss before running tests. We have ant task for that as well.	1	1	https://issues.jboss.org/browse/JBESB-523
11475	JBESB-523	Hiya Prabhat. This probably got lost in the previous comment, sorry about that - "Each image runs a local instance of AS, purely for JMS, and requires access to an FTP server. " At present we do not restart the app servers as we do not deploy anything into them, it is purely for JMS. Which ant tasks are you referring to? Are these the ones used in the AS build?	1	1	https://issues.jboss.org/browse/JBESB-523
11476	JBESB-523	Yes, the ant task I am refereeing is from AS build. (server:start and sever:stop). Can you add that in build? I guess since it does not deploy anything, it does not matter which AS version is being used as long as jms queue A,B,C are there. I will first configure it on linux, then I will work on windows.	1	1	https://issues.jboss.org/browse/JBESB-523
11477	JBESB-523	It's set up at http://dev45.qa.atl.jboss.com:8585/hudson/job/jboss-esb-linux and you can see the details of last build at http://dev45.qa.atl.jboss.com:8585/hudson/job/jboss-esb-linux/16/. You need to be behind redhat vpn to access the page. Hudson had some problem exiting the build process when it was configured to start and stop running jboss instance that was configured through its shell script UI. So, I have temporarilty started jboss in the background so that testsuite can run until we start jboss from testsuite itself or I import those tasks in hudson. I have set jbossesb-tests@googlegroups.com,jboss-qa-internal@redhat.com as buiild notificaiton recipients. The last build failed org.jboss.internal.soa.esb.couriers.tests.CourierIntegrationTest because of java.net.UnknownHostException: ftp.ncl.jboss.com: ftp.ncl.jboss.com. What is the correct ftp address? Currently, I start the testsuite with ant -f integration-build.xml -Dorg.jboss.esb.properties=/home/hudson/configs/jboss-esb/product.properties and this is how the product.properties file look like org.jboss.esb.build.deprecation=yes org.jboss.esb.build.optimize=off org.jboss.esb.build.frominstall=no org.jboss.esb.build.debug=yes org.jboss.esb.build.tests=no org.jboss.esb.alltests.build=yes org.jboss.esb.jboss.home=. org.jboss.esb.test.ftp.hostname=ftp.ncl.jboss.com org.jboss.esb.test.ftp.user=lintest org.jboss.esb.test.ftp.pwd=lintest org.jboss.esb.test.ftp.dir=/home/lintest/test Please let me know if some configs have changed.	1	1	https://issues.jboss.org/browse/JBESB-523
11478	JBESB-523	Hiya Prabhat. There are four properties that you need to configure for ftp access. You should change these to match whichever ftp server is accessible in your local environment. org.jboss.esb.test.ftp.hostname - The ftp server hostname org.jboss.esb.test.ftp.user - The user name for login org.jboss.esb.test.ftp.pwd - The password for login org.jboss.esb.test.ftp.dir - The base directory for the tests Kurt has recently asked for the build to be changed to use an embedded ftp server but overriding in this fashion will still be supported. Which email address is hudson using to sent reports to the jbossesb-tests mailing list? We will need to add it to the group. Kev P.S. Did you manage to get a windows integration running?	1	1	https://issues.jboss.org/browse/JBESB-523
11479	JBESB-523	Hi Kevin, I thought ftp.ncl.jboss.com is what you want me to use. I would have to to see which ftp server I can use. Would you need to access this ftp server as well? If so, then I have to work with IT on how long to keep files etc. Hudson uses jboss-qa-internal@redhat.com to send mail to jbossesb-tests@googlegroups.com. Please add this to the safe list. Once I set up linux compeltely, I will work on windows. As requested before, can you add server start and server stop target to the tesstuite to start and stop jboss AS?	1	1	https://issues.jboss.org/browse/JBESB-523
11480	JBESB-523	Hiya Prabhat. ftp.ncl.jboss.com is the local ftp server in the Newcastle office (where we currently run the CI). You would not have access to those machines. I do not believe we would need access to your ftp server as long as we could contact the person administering it. I have asked Mark to add the email address to the test google group. I will have a look at adding similar tasks into our build system (we are not part of AS and do not use buildmagic), but this is likely to be a low priority task for us. Is there an issue with leaving the AS running? Thanks for your help, Kev	1	1	https://issues.jboss.org/browse/JBESB-523
11481	JBESB-523	Hiya Prabhat. Is there any way to configure Hudson so that it will always send an email? One of the benefits of receiving the emails is that we are aware of changes shortly after they happen.	1	1	https://issues.jboss.org/browse/JBESB-523
11482	JBESB-523	Also, is there a reason why you chose 12 hours? At present we are running CI checks/builds every 15 mins on the Newcastle systems.	1	1	https://issues.jboss.org/browse/JBESB-523
11483	JBESB-523	It is possible to configure Hudson to send success emails. Can you please do that for our builds?	0	0	https://issues.jboss.org/browse/JBESB-523

#	Requirement	Questions	Answered	Successfully Answered	URL
11484	JBESB-523	Did you guys get an e-mail? I see that there was a build failure and then a successful build and as per log, it should have sent mail to jbossesb-tests@googlegroups.com. Len, please add yourself to jbossesb-tests@googlegroups.com or if you are a member of jboss-qa-internal then you should have got that as well.	1	1	https://issues.jboss.org/browse/JBESB-523
11485	JBESB-523	Hi Kevin, Agreed. 1. We are almost there with e-mail notificaiton. 2. It can now be done with the latest hudson and is being done. 3. Hudson has emma plugin. Where are emma coverage reports? Reliable: What has been going wrong? This will help me fix the problem.	1	1	https://issues.jboss.org/browse/JBESB-523
11486	JBESB-520	Jason, are you going to be able to work on this?	0	0	https://issues.jboss.org/browse/JBESB-520
11487	JBESB-518	Can you put an example here?	1	0	https://issues.jboss.org/browse/JBESB-518
11488	JBESB-509	Hi Kevin, the only thing I want to add to the deploy target is <ant dir="install" target="deploy"/>. Is that what you meant by 'You should either remove it, in favour of the install build.xml, or have it call the install build.xml to handle the deployment.' ? Thanks	1	1	https://issues.jboss.org/browse/JBESB-509
11489	JBESB-509	That looks much cleaner. I'll change it to that. Do you know whether jbossesb-listeners.jar and jbossesb-services.jar are needed anymore? They are no longer created but I wanted to make sure...	1	1	https://issues.jboss.org/browse/JBESB-509
11490	JBESB-501	Re removing the ability to chain actions... I don't think we should be throwing out the baby with the bath water. Not saying I'm a big lover of action chaining in its current form, but I'd suspect removing the option completely and replacing it with nothing more than "singleaction + service2service" will meet with some disapproval from the likes of Burr (and others). When we got Rosetta first, this was effectively what we had. Plugging "actions" (code that operates on a message) together meant defining separate services and routing the message between them. Otherwise you had to write monolithic actions (un-reuseable). IMO, either option was a total pain in the ass from a usability/reuseability perspective. Why can't we have a happy medium of some sort, where chained actions are still supported but we make "troublesome" actions (for scoping reasons or whatever) into standalone Services as Kurt describes.	1	1	https://issues.jboss.org/browse/JBESB-501
11491	JBESB-500	As part of this, can we remove the services directory from core?	1	1	https://issues.jboss.org/browse/JBESB-500
11492	JBESB-498	How is this different from either saying to people that only ESB aware services should be registered in the ESB registry, or that they shouldn't use EPRs for non-ESB-aware services?	1	1	https://issues.jboss.org/browse/JBESB-498
11493	JBESB-490	What's strange about that ;-)	1	1	https://issues.jboss.org/browse/JBESB-490
11494	JBESB-490	Since you started it (in all but name), do you want to finish it? ;)	1	1	https://issues.jboss.org/browse/JBESB-490
11495	JBESB-490	David, I have been looking at an issue logged against the helloworld quickstart and have made a number of changes in this area. Apologies for this. Are your changes ready to commit? If not, could you do it in stages?	1	1	https://issues.jboss.org/browse/JBESB-490
11496	JBESB-476	Does this mean that customactions,demonstrated by the quickstarts, should use the BeanConfiguredAction as the default?	0	0	https://issues.jboss.org/browse/JBESB-476
11497	JBESB-468	For the MP? What sort of transformation?	0	0	https://issues.jboss.org/browse/JBESB-468
11498	JBESB-465	How important do you think this is for 4.x?	0	0	https://issues.jboss.org/browse/JBESB-465
11499	JBESB-464	Is this still on course for this week?	0	0	https://issues.jboss.org/browse/JBESB-464
11500	JBESB-464	Can you register a new JIRA bug against the issue you're seeing? I assume that everything works fine when the server is not running? If that's the case we'll mention this in the release notes and take a look at it after the release.	1	1	https://issues.jboss.org/browse/JBESB-464
11501	JBESB-453	Kurt, please tell me how the undeliverable service would work. First, a message can't reach it's destination(s). You need to store it. How do you want to store it? The message store currently has a DB implementation which use a table and 3 columns (uri key, message, type) which directly reflects the Message types (XML/SERIALIZABLE) that the ESB has support for. Would you have a new Message type? You need to know within the message store that a given message is an undelivered message. Would you require some other columns to support this? A flag, (delivered vs stil to be delivered, etc) ?? Do you want to remove messages that get re-delivered successfully or flag them, etc? More information is required before we can code away... ;-)	1	1	https://issues.jboss.org/browse/JBESB-453
11502	JBESB-449	Derek, can you check your build.xml. According to the one in project.zip there is supposed to be a WEB-INF/lib and there isn't one. This causes the building of the WAR to fail. Maybe I'm missing something?	1	1	https://issues.jboss.org/browse/JBESB-449
11503	JBESB-447	Does this mean we're going to package up JBM in the standalone JBossESB deployment as it's JMS provider? Note that JBM depends on the 4.x MC architecture.	1	1	https://issues.jboss.org/browse/JBESB-447
11504	JBESB-447	So should we deploy ESB to the same 4.0 MC on which JBM is running, or have our bootstrapper start the JBM MC in parallel?	1	1	https://issues.jboss.org/browse/JBESB-447
11505	JBESB-447	Good question. What are the pros and cons?	1	1	https://issues.jboss.org/browse/JBESB-447
11506	JBESB-447	The pros are that deploying to JBossAS or to the JBM MC is pretty much doing the same thing! So we end up with a simpler unified deployment strategy. In this standalone mode we then support clean shutdowns, and clustering (which needs to be tested). The downside is that ESB will be married to the 4.x MC for it's startup and shutdown infrastructure. We would use no longer use the apache launcher. The ESB maybe persieved heaver. Here is where it gets fuzzy for me. Would it still be light if we'd build our own? From a functionality perspective I think it's a no brainer that we should do it. But keep in mind that what JBM is using is a fully dressed down (naked?) version of JBossAS 4.x.	0	0	https://issues.jboss.org/browse/JBESB-447
11507	JBESB-445	What was the CDATA section issue with the XSD? I remember trying this when we were doing the XSD stuff. I may be wrong, but I thought it actually did support a CDATA section in the properties, but it just didn't make it through to the XMLBeans binding classes.	1	1	https://issues.jboss.org/browse/JBESB-445
11508	JBESB-445	Ah OK. Wouldn't it be useful to inject the properties/ConfigTree into the script (as well as the message). This way you can parameterize the scripts from the config. I wonder how this action will perform? I'd imagine that creating a new instance of it every time will be a bit of a dog. But I suppose this is an issue we have with all our actions - they're non-singleton. I'd suggest that if/when we make the actions (or whatever they'll be called in the new architecture) singleton (by default), we could look at using the GroovyClassloader and require the Groovy scripts to implement our ActionProcessor interface (or whatever we call that). Then we can just treat Groovy scripts like any other java class (plus cache the instance on a per action basis).	0	0	https://issues.jboss.org/browse/JBESB-445
11509	JBESB-443	Do you think this is enough for 4.2? Could we push off any more performance improvements to the 5.0 release?	1	0	https://issues.jboss.org/browse/JBESB-443

#	Requirement	Questions	Answered	Successfully Answered	URL
11510	JBESB-438	Mark, can you elaborate on this please? With the current codebase anything within a single .esb file should already be taken down cleanly. The order used to specify the listeners/gateways is the sequence used by the managed lifecycle controller to initialise/start the lifecycles. When the managed lifecycle stops/destroys the lifecycles then the order is reversed. Each listener/gateway should wait until current work has finished before being destroyed, subject to timeout. Now that there is a move to deploy multiple .esb files we must make sure that the dependencies *between* those .esb files are fully realised. I do not believe the current codebase is in a position to do this and would have been happier if this had been deferred until 5 where we will have proper MC/OSGi integration.	0	0	https://issues.jboss.org/browse/JBESB-438
11511	JBESB-437	What would my Action class need to look like to access the properties that were configured from jbossesb.xml?	0	0	https://issues.jboss.org/browse/JBESB-437
11512	JBESB-437	Can this go in for MR2?	1	1	https://issues.jboss.org/browse/JBESB-437
11513	JBESB-434	Nice work Tom. If you haven't already, can you make sure the docs mention this and cover any configuration/deployment options?	0	0	https://issues.jboss.org/browse/JBESB-434
11514	JBESB-430	Can we do this just for JBoss Messaging on 4.2?	1	1	https://issues.jboss.org/browse/JBESB-430
11515	JBESB-429	Is it documented :-)?	0	0	https://issues.jboss.org/browse/JBESB-429
11516	JBESB-417	I can understand what has to be done to realise this for 5 but am unsure what is required for 4.2. Can you elaborate please?	0	0	https://issues.jboss.org/browse/JBESB-417
11517	JBESB-417	Without fundamental changes to the codebase, is there anything more we can do? The answer may be: no.	0	0	https://issues.jboss.org/browse/JBESB-417
11518	JBESB-417	The current deployer architecture, as it is based on the AS, can be used to provide different profiles for deployments in the same way. Would this suffice for 4.2? If so then we can drop the 4.2 requirement and keep this open for 5. BTW The OSGi development should include supporting repositories for bundles, this could also help in this requirement.	0	0	https://issues.jboss.org/browse/JBESB-417
11519	JBESB-417	Concerning the dev, test, staging, production item... That is more a matter of making sure the steps involved in moving a "project" out of dev and into production are at least tested/well documented. In the future there might be a central repository that is responsible for this type of thing. For instance, if I need to update 2 ESBs with the same services/archives together, can I simply have a shared file system between the 2 servers? Can I ask one ESB server and all its associated services to "stand down" in which case it no longer consume messages and all new messages would be routed to other ESB servers that hosted those same services (matched by category and name)? There may be more questions like this.	0	0	https://issues.jboss.org/browse/JBESB-417
11520	JBESB-417	With the current AS setup it is possible to share these deployments between multiple ESBs providing they have access to the same deployment locations. Taking down one ESB should unregister all the services deployed in that ESB but I am unsure whether failover will work in all circumstances. Do you have specific scenarios that you would like us to check?	0	0	https://issues.jboss.org/browse/JBESB-417
11521	JBESB-415	Can we replace gateways with JCA inflow/outflow?	0	0	https://issues.jboss.org/browse/JBESB-415
11522	JBESB-405	How's that for coincidence then ;-)	1	1	https://issues.jboss.org/browse/JBESB-405
11523	JBESB-400	Is this still relevant?	1	1	https://issues.jboss.org/browse/JBESB-400
11524	JBESB-400	Can we close this issue?	0	0	https://issues.jboss.org/browse/JBESB-400
11525	JBESB-396	Is there anything in particular I can do to reproduce this issue Kev?	1	1	https://issues.jboss.org/browse/JBESB-396
11526	JBESB-393	Is this a GA item?	1	0	https://issues.jboss.org/browse/JBESB-393
11527	JBESB-385	Well there are 2 requirements here actually: 1. A home for the Wink (tool used to create the demos) project files. I would put these into the JBossESB main project in SVN but for the fact that they're very big (~25M each). Is there another location in SVN that we can put them, but not as part of the main ESB Project. 2. A home on the web (jboss.org) where we can link the demos. We could have links to these in a number of places - wiki, labs pages etc. This stuff is "sales" material as well as educational.	1	1	https://issues.jboss.org/browse/JBESB-385
11528	JBESB-376	How is this going?	1	1	https://issues.jboss.org/browse/JBESB-376
11529	JBESB-376	Release is today/tomorrow. Should I move this task until after the release?	1	1	https://issues.jboss.org/browse/JBESB-376
11530	JBESB-370	Should we push this until after the 4.0 GA?	0	0	https://issues.jboss.org/browse/JBESB-370
11531	JBESB-369	Kurt has already done the war deployment, certainly for tomcat. Do we know if there is anything else to do in this?	1	1	https://issues.jboss.org/browse/JBESB-369
11532	JBESB-369	Yes this is testing on tomcat 5.5.20, and seems to only work with the tomcat.zip, not with the tomcat installer as Burr recently found out. So we may want to take a look at that, and maybe deploying it to the other environments. How much work would it be to add an orderly shutdown? There is no "unload-on-shutdown" is there?	1	1	https://issues.jboss.org/browse/JBESB-369
11533	JBESB-366	Does the flash work help here?	1	1	https://issues.jboss.org/browse/JBESB-366
11534	JBESB-365	Object to XML: What are we looking for here - integrating something like XStream/Castor? DOM to XML: ??? Just DOM serialisation? Groovy Scripting: I'll be adding Groovy support for Smooks very soon. Once there, people will be able to dev transformations in Groovy (Vs raw Java) Binary to ASCII Encode/Decode: Fine. Byte Array to Serializable: ??? Mail Body to String: ??? Is this a Mail Gateway we're talking about here? Http Request Body to String: ??? Is this a HTTP Gateway?	0	0	https://issues.jboss.org/browse/JBESB-365
11535	JBESB-363	If you're working on this, can you mark it as active.	1	1	https://issues.jboss.org/browse/JBESB-363
11536	JBESB-353	Why do we need this exactly?	1	0	https://issues.jboss.org/browse/JBESB-353
11537	JBESB-347	Kurt, can you add tests for the Aggregator and CBR actions in the package? Thanks.	0	0	https://issues.jboss.org/browse/JBESB-347
11538	JBESB-347	There are tests for that in the qa package as you need a running server. Can I close this?	1	1	https://issues.jboss.org/browse/JBESB-347
11539	JBESB-327	Is this definitely needed for 4.0 GA?	1	0	https://issues.jboss.org/browse/JBESB-327
11540	JBESB-323	Will relative paths not work on all OSs?	1	1	https://issues.jboss.org/browse/JBESB-323
11541	JBESB-261	Which Email helper class?	1	1	https://issues.jboss.org/browse/JBESB-261
11542	JBESB-238	I think we should drop this in favour of 5.0. What do you think?	0	0	https://issues.jboss.org/browse/JBESB-238
11543	JBESB-233	What would the rule look at in the message to determine if it gets stored or not? Would we use XPath and define a rule that way to determine storing rule(s)? What about non xml messages? Wondering what the strategy should be for this.	1	0	https://issues.jboss.org/browse/JBESB-233
11544	JBESB-208	How does this relate to the jUDDI management tool?	1	1	https://issues.jboss.org/browse/JBESB-208
11545	JBESB-208	What do you mean by "configuring an EPR"?	1	1	https://issues.jboss.org/browse/JBESB-208
11546	JBESB-189	Could you add some more detail here? Cheers.	1	1	https://issues.jboss.org/browse/JBESB-189

#	Requirement	Questions	Answered	Successfully Answered	URL
11547	JBESB-181	And yet another option might be to define a static inner class in the ActionProcessor interface such that ActionProcessor to provide "helper" type methods for accessing message components e.g. ActionProcessor.MessageAccessor.setProcessingObject(message, object); ... Object processingObj = ActionProcessor.MessageAccessor.getProcessingObject(message); ... so the static string is completely hiden and the getting and setting are centrally managed... in fact, we could change where we store it in the future (should we decide to do that) and no ones code would break (in theory ;-))	1	1	https://issues.jboss.org/browse/JBESB-181
11548	JBESB-172	I have sent an email round to groups for comments on the work I have done to resolve this issue. Should I add the details here if we are happy with the groups email?	0	0	https://issues.jboss.org/browse/JBESB-172
11549	JBESB-166	Is this still relevant? With the Quickstarts maybe we no longer need to worry about it?	0	0	https://issues.jboss.org/browse/JBESB-166
11550	JBESB-165	Since you're working on this, can you drive the related discussion on the forum around spawning new listeners?	1	1	https://issues.jboss.org/browse/JBESB-165
11551	JBESB-165	Can you add some unit and QA tests before merging with the trunk?	0	0	https://issues.jboss.org/browse/JBESB-165
11552	JBESB-161	What exactly is being requested here? We merged the old "product/config" and "product/docs/install" resources to the current "product/install" structure. Was there something additional in mind e.g. a replacement for the DomElement configs used on the listeners (MicroContainer, no??)?	1	1	https://issues.jboss.org/browse/JBESB-161
11553	JBESB-161	The other related task is http://jira.jboss.com/jira/browse/JBESB-119 and David D is working on this. Can we close out this task now?	0	0	https://issues.jboss.org/browse/JBESB-161
11554	JBESB-152	How about a conversion sub-package with ByteArrayToString, FileToByteArray, ObjectToCSVString, ObjectToFileWriter ObjectToXStream? A template sub-package with all of the DummyX classes.	0	0	https://issues.jboss.org/browse/JBESB-152
11555	JBESB-141	How is this going?	0	0	https://issues.jboss.org/browse/JBESB-141
11556	JBESB-125	Can you mark this as In Progress ;-)?	1	1	https://issues.jboss.org/browse/JBESB-125
11557	JBESB-119	Dave, can you post a reference to the forum in case others have comments? Thanks.	1	1	https://issues.jboss.org/browse/JBESB-119
11558	JBESB-119	Good day, Mark. Excellent suggestion. However, I believe there needs to be a bit more to a forum post. I have (prior to developing the Alpha schema) taken both the Trailblazer and TestJBossESB samples' config files and transformed them to meet the "new standard". The context of each set of properties requires, in my view, a README to explain, along with these sample config files. I will attach these documents shortly, then create the forum thread. Overkill? Dave	1	1	https://issues.jboss.org/browse/JBESB-119
11559	JBESB-119	Dave, I know you're snowed under. Should I ask Kurt to look at this?	1	1	https://issues.jboss.org/browse/JBESB-119
11560	JBESB-114	As per groups message, should I start to abstract the bus concept as a service in order to facilitate this task? My current line of thought is to provide services/transport with sub service categories for email, ftp, jms, file, jdbc etc.	0	0	https://issues.jboss.org/browse/JBESB-114
11561	JBESB-109	What DB connection are you refering to? The persistence store DB will be using Apache Commons DBCP to handle the connection pooling and validating connections in the pool for stale connections, etc... It used a "validate sql" sending any query you tell it to validate wether a connection from the pool is still valid or not.	1	1	https://issues.jboss.org/browse/JBESB-109
11562	JBESB-105	I have decided to use the JDK 1.5 UUID implementation. For more info: http://java.sun.com/j2se/1.5.0/docs/api/java/util/UUID.html Can we close this one Mark if you have already reviewed the algorithm?	1	1	https://issues.jboss.org/browse/JBESB-105
11563	JBESB-103	Yes. I think we're in agreement. Definitely not for the beta. That's why I set Fix Version to 4.0 (the GA).	1	1	https://issues.jboss.org/browse/JBESB-103
11564	JBESB-90	Can this also include renaming them (as per the discussion during the Newcastle meeting)?	0	0	https://issues.jboss.org/browse/JBESB-90
11565	JBESB-70	The class NagiosStandaloneHeartbeat just contains a constructor. How are we going to integrate? Is there an API to use. Once I implemented a component sending a heartbeat message to a Netsaint monitoring server. I was communicating with it using sockets. Is this the way we want to integrate with Nagios (former Netsaint) too?	1	1	https://issues.jboss.org/browse/JBESB-70
11566	JBESB-50	Not quite sure Mark. That would somehow imply that properties correspond to a package. Do they ? if you leave it at org.jboss.soa.esb.property.XXX they would still belong to ESB but not related to any package in particular In any case, it doesn't make a big difference	1	1	https://issues.jboss.org/browse/JBESB-50
11567	JBESB-46	Is this something that could be dynamically pluggable? Instead of replacing a class, couldn't users simply define their own class(es) and let the "encryption manager" pick them up from a configuration file (for example) and instantiate them at runtime?	0	0	https://issues.jboss.org/browse/JBESB-46
11568	JBESB-39	Tom, about 6 weeks prior to release can you go through the code and let the team know which pieces need more coverage?	0	0	https://issues.jboss.org/browse/JBESB-39
11569	JBESB-38	ESCCore/Tests has been deleted and the classes move to their relevant modules. I will leave this issue as open as I need to revisit the tests which are dependant on a running server - these are functional tests and will come into play once refactoring is nearing to a completion for the first phase and we have set up scripts for the server which also deploy the esb and run the functional tests. Can this issue be downgraded ?	0	0	https://issues.jboss.org/browse/JBESB-38
11570	JBESB-35	Do you have an idea for what this should contain?	1	1	https://issues.jboss.org/browse/JBESB-35
11571	JBESB-22	Where are we with this at the minute: 1. Forum: http://www.jboss.com/index.html?module=bb&op=viewtopic&t=81237, http://www.jboss.org/index.html?module=bb&op=viewtopic&t=82506 2. Wiki: http://wiki.jboss.org/wiki/Wiki.jsp?page=TransformationServices	1	1	https://issues.jboss.org/browse/JBESB-22
11572	JBESB-19	Is there any technical documentation available about implementing this? Are there any guidelines?	0	0	https://issues.jboss.org/browse/JBESB-19
11573	JBESB-19	That's what I initialy understood using the description. But than there was somthing about deployment and I got confused. Ok this is what I'm thinking off : Build a REST servlet which can handle GET, POST and DELETE requests. GET : will do a receive from a certain service and returns the result. For example : GET http://www.mycompany.com/rest/ordersQueue will get the first message on a JMS queue named ordersQueue POST : do a synchronous call to a service endpoint and return the result PUT :do a synchronous call whithout returning a result DELETE : will do a receive from a certain service whithout returning the result Like this every service should be accessible using a HTTP method What do you think about this design? What I would like to know before I can start developing is how I can delegate the request to the service itself.	0	0	https://issues.jboss.org/browse/JBESB-19
11574	JBESB-19	Have these classes or has this functionality been added to the newest version of the JBoss ESB? I can't find these classes anywhere and there are no examples of using synchronous HTTP directly (except for a separate service which uses the ESB directly).	1	0	https://issues.jboss.org/browse/JBESB-19
11575	JBESB-16	The client-side is straightforward. Do we implement a server-side or try to leverage one that already exists?	1	1	https://issues.jboss.org/browse/JBESB-16
11576	JBESB-16	Can you integrate the branch using the old listener classes, which are still there but moved to the 'old' package structure? At least then we'd have everything in the trunk and after the RC you and Esteban can talk about updating to the new model.	1	1	https://issues.jboss.org/browse/JBESB-16

#	Requirement	Questions	Answered	Successfully Answered	URL
11577	JBESB-13	Mark, can you confirm this is done with the code under branches/JBESB-13 ? The basic messaging server (slimmed down JBossAS) and a basic esb kernel are up and running. There is nothing to deploy as there are no entries in the configuration files ie no core services and no general services as there are 2 configuration files I have added, see svn under svn\jbossesb\branches\JBESB-13\ESBCore\etc\server\default\config where there empty configuration files jboss-esb-core.xml & jboss-esb-services.xml If we now add service implementations via the MC, we are not dependant on the MC M3 release - which is what standalone Messaging built on the MC will need.	1	1	https://issues.jboss.org/browse/JBESB-13
11578	JBESB-13	Can you add a couple of stand-alone JMS QA tests, just to confirm the deployment outside of ESB?	0	0	https://issues.jboss.org/browse/JBESB-13
11579	JBESB-11	For this issue, i presume you want me to add a jms and messaging protocol implementation similar to the tcp implementation you have added. Can I continue this on the same branch JBESB-13 ?	1	1	https://issues.jboss.org/browse/JBESB-11
11580	JBESB-11	I am unsure whether this should be closed or not, the initial discussion was about getting rossetta running on jboss messaging, the idea was plug-and-play, or rather it should just drop in and wokr, and it does, as there are no dependencies. If i can assume that is the initial part, i am now assuming that the next step to this issue is to provide a messaging layer (not jms, but jboss messaging) transport. Mark can you clarify ?	1	1	https://issues.jboss.org/browse/JBESB-11
11581	JBSSO-37	When will this issue be solved ? I believe it's the last issue before 1.0GA release comes out.	0	0	https://issues.jboss.org/browse/JBSSO-37
11582	JBSSO-36	Could any provide some more info about this bug? It would be nice to start fixing it.	1	1	https://issues.jboss.org/browse/JBSSO-36
11583	JBSSO-25	With what version of seam do you want to provide a hook? we can do it right now and begin with testing... What kind of hook do you want? please explain...	0	0	https://issues.jboss.org/browse/JBSSO-25
11584	JBFORUMS-268	What this task has to do with JBoss Forums project?	0	0	https://issues.jboss.org/browse/JBFORUMS-268
11585	JBFORUMS-263	Yep, I saw all your comments and I will work on mentioned bugs. Nevertheless, I made an agreement with our graphic designer and he will create buttons for all needed languages. All he needs is a translation of buttons' texts. Especially for images: forum_but_addnewcategory.gif forum_but_addnewforum.gif forum_but_closed.gif forum_but_delete.gif forum_but_edit.gif forum_but_email.gif forum_but_moderate.gif forum_but_newtopic.gif forum_but_postreply.gif forum_but_profile.gif forum_but_quote.gif forum_but_subscribe.gif forum_but_unsubscribe.gif which can be found here: http://anonsvn.labs.jboss.com/labs/jbossforums/branches/forums101P26/forums/src/bin/portal-forums-war/default_graphics/images/ Would you be so kind and attach a text file with their translations?	1	1	https://issues.jboss.org/browse/JBFORUMS-263
11586	JBFORUMS-263	Thanks, although it's a bit hard to say that there is a whole team behind Forums ;-). Nevertheless, I have a problem with the TXT file. I believe you saved it in ASCII instead of UTF-8 coding and French characters are messed up. Could you reattach UTF-8 version?	1	1	https://issues.jboss.org/browse/JBFORUMS-263
11587	JBFORUMS-203	Awesome. Big thanks. Could you please tell me which version of ResourceJSF.properties did you translate, from which version of JBoss Forums? 1.0.0 RC from binary downloads or maybe from svn sources. It is very important for me as I'm adding some things there while doing redesign and new views.	1	1	https://issues.jboss.org/browse/JBFORUMS-203
11588	JBFORUMS-203	Hi Luca, Would you mind checking if it works properly? I cannot find in my web browsers locale switch and I'd like to be sure that your contributorship works as expected. btw. I'll do soon some work on 1.0.1 bundle and I'll need few new lines to be translated. Can I count on you ? Cheers, Rysiek	1	1	https://issues.jboss.org/browse/JBFORUMS-203
11589	JBFORUMS-203	Thanks Luca but you used ASCII encoding and the file has messed up Italian specific characters for me. Additionally, could you replace all Italian characters with their equivalents of \uXXXX notation? This fixes a problem with displaying UTF-8 messages in <h:messages>.	1	1	https://issues.jboss.org/browse/JBFORUMS-203
11590	JBFORUMS-152	I checked this issue with lastPost which you pointed out but it works correctly for me. Could you tell me how to reproduce?	0	0	https://issues.jboss.org/browse/JBFORUMS-152
11591	JBFORUMS-151	Is it possible, while changing the look and feel, to add a mechanism to translate the images ? I'd like to translate the buttons in french, but it's kinda ugly if I simply overwrite the english image files. If you can add such a mechanism, I'll create french buttons for you. I'd need the original files though. Working with compressed gifs is such pain in the a$$...	0	0	https://issues.jboss.org/browse/JBFORUMS-151

#	Requirement	Questions	Answered	Successfully Answered	URL
11592	JBFORUMS-65	[org.jboss.wsf.stack.cxf.client.configuration.JBossWSSpringBusFactory] INITIAL_APP_CONTEXT_CREATION_FAILED_MSG: org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'org.apache.cxf.binding.soap.customEditorConfigurer' defined in URL [vfs:/home/Avishkar/Downloads/jboss-6.0.0.Final/common/lib/cxf-rt-bindings-soap.jar/META-INF/cxf/cxf-extension-soap.fixml]: Initialization of bean failed; nested exception is org.springframework.beans.ConversionNotSupportedException: Failed to convert property value of type 'java.util.ArrayList' to required type 'org.springframework.beans.PropertyEditorRegistrar[]' for property 'propertyEditorRegistrars'; nested exception is java.lang.IllegalStateException: Cannot convert value of type [org.apache.cxf.binding.soap.spring.SoapVersionRegistrar] to required type [org.springframework.beans.PropertyEditorRegistrar] for property 'propertyEditorRegistrars[0]': no matching editors or conversion strategy found at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.doCreateBean(AbstractAutowireCapableBeanFactory.java:527) [:3.0.3.RELEASE] at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.createBean(AbstractAutowireCapableBeanFactory.java:456) [:3.0.3.RELEASE] at org.springframework.beans.factory.support.AbstractBeanFactory$1.getObject(AbstractBeanFactory.java:291) [:3.0.3.RELEASE] at org.springframework.beans.factory.support.DefaultSingletonBeanRegistry.getSingleton(DefaultSingletonBeanRegistry.java:222) [:3.0.3.RELEASE] at org.springframework.beans.factory.support.AbstractBeanFactory.doGetBean(AbstractBeanFactory.java:288) [:3.0.3.RELEASE] at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:190) [:3.0.3.RELEASE] at org.springframework.beans.factory.support.DefaultListableBeanFactory.preInstantiateSingletons(DefaultListableBeanFactory.java:574) [:3.0.3.RELEASE] at org.springframework.context.support.AbstractApplicationContext.finishBeanFactoryInitialization(AbstractApplicationContext.java:895) [:3.0.3.RELEASE] at org.springframework.context.support.AbstractApplicationContext.refresh(AbstractApplicationContext.java:425) [:3.0.3.RELEASE] at org.jboss.wsf.stack.cxf.client.configuration.JBossWSBusApplicationContext.<init>(JBossWSBusApplicationContext.java:63) [:3.4.1.GA] at org.jboss.wsf.stack.cxf.client.configuration.JBossWSSpringBusFactory.createApplicationContext(JBossWSSpringBusFactory.java:105) [:3.4.1.GA] at org.jboss.wsf.stack.cxf.client.configuration.JBossWSSpringBusFactory.createBus(JBossWSSpringBusFactory.java:64) [:3.4.1.GA] at org.apache.cxf.bus.spring.SpringBusFactory.createBus(SpringBusFactory.java:86) [:2.3.1] at org.apache.cxf.bus.spring.SpringBusFactory.createBus(SpringBusFactory.java:64) [:2.3.1] at org.apache.cxf.bus.spring.SpringBusFactory.createBus(SpringBusFactory.java:53) [:2.3.1] at org.jboss.wsf.stack.cxf.client.configuration.JBossWSBusFactory.createBus(JBossWSBusFactory.java:54) [:3.4.1.GA] at org.jboss.wsf.stack.cxf.client.ProviderImpl.createServiceDelegate(ProviderImpl.java:69) [:3.4.1.GA] at javax.xml.ws.Service.<init>(Service.java:57) [:1.0.0.Final] at com.tradeexpress.mobile.order.webservice.Service1.<init>(Service1.java:64) [:] at com.tradeexpress.mobile.service.OrderServiceImpl.attachToMSMQ(OrderServiceImpl.java:1370) [:] at com.tradeexpress.mobile.web.controller.OrderController.placeOrder(OrderController.java:761) [:] at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)	1	0	https://issues.jboss.org/browse/JBFORUMS-65
11593	JBFORUMS-2	I will look at the forums debate. Thanks for pointing me there. But one small question Ales. In what way was I bashing? I was just mislead by the fact that it is not a change to the current forum, but a completely new forum with a new searchengine. Sorry if I gave the impression I was bashing, that was not my intention (maybe the term frustrated was wrongly used here, since I'm not frustrated with the current searchengine, but with the users that do not take the time to read the text on the radionbuttons and they get frustrated by me for shouting Search The Fine Forum ;-))	1	0	https://issues.jboss.org/browse/JBFORUMS-2
11594	ENTESB-7042	[~alexkieling] I logged the two tasks we talked about (ENTESB-7127, ENTESB-7128) for authentication and handling special characters. Is there anything else we need to log separately or do you think we can close this JIRA?	1	0	https://issues.jboss.org/browse/ENTESB-7042
11595	ENTESB-6946	https://github.com/jboss-fuse/fuse-jenkins-pipelines/commit/e07daaf33385f99fdbde8e3c18eaec630d7bf565 re: info for prod - we could do it here - you'd prefer a config file? Any preferred format? Would it be produced only for successful builds? What info should be in it (I'm assuming the tags that are produced by the build, but wasn't sure whether there was anything else that would be necessary)?	0	0	https://issues.jboss.org/browse/ENTESB-6946
11596	ENTESB-6879	Can I mark this as Resolved for Fuse 7?	1	0	https://issues.jboss.org/browse/ENTESB-6879
11597	ENTESB-6878	Hi, Spring Boot uses BOM dependency io.fabric8:fabric8-project-bom-camel-spring-boot. I don't know, if fuse 7 will use also this dependency. If yes, which version i should use ? If not, what should we use to test fuse 7 with Camel on Spring Boot ?	1	1	https://issues.jboss.org/browse/ENTESB-6878
11598	ENTESB-6878	For CXF, it doesn't appear that there is an official BOM (https://issues.apache.org/jira/browse/CXF-6857). It looks like the distribution manifest might be something workable there though : org.apache.cxf:cxf-distribution-manifest - here is a link to the most recent one produced : http://origin-repository.jboss.org/nexus/content/groups/ea/org/apache/cxf/cxf-distribution-manifest/3.1.11.fuse-000030/cxf-distribution-manifest-3.1.11.fuse-000030.pom Karaf : org.apache.karaf:karaf pom seems like the best choice - link to most recent version : http://origin-repository.jboss.org/nexus/content/groups/ea/org/apache/karaf/karaf/4.2.0.fuse-000030/karaf-4.2.0.fuse-000030.pom Hibernate : I think that's being pulled in by Karaf or EAP, depending on container. We're not building it as part of the pipeline. For the karaf kit, the only thing I see shipped is hibernate-validator and it's version corresponds to the version of hibernate-validator in the Karaf pom above. So I think there's a few questions here : - should we change org.jboss.fuse:fuse-parent -> org.jboss.fuse.bom:fuse-parent to match 6.3? - do we want to generate a bom combining some of these?	1	1	https://issues.jboss.org/browse/ENTESB-6878
11599	ENTESB-6878	Hi [~tcunning], we were discussing how users can deploy fuse and here's what we found: * bundle deployed on standalone fuse ** runtime jboss-fuse-karaf, all dependencies provided by karaf as features ** BOM with versions aligned to karaf 4 features * jar/war deployed onto EAP 7 ** runtime EAP7 camel and other stuff provided by subsystem ** BOM with versions aligned to EAP * far jar / spring-boot ** flat classpath built by maven plugin from MRRC ** BOM with spring boot camel starter... ([~llowinge] might elaborate what we need here) * war / tomcat (Enterprise Web Server) ** runtime EWS, all dependencies pulled from MRRC ** BOM do we need to align with EWS ? does it make sense to reuse other bom? * FIS / OpenShift deployment / Sprint Boot ** flat classpath ** BOM there is fabric8 spring BOM afaik * FIS / OpenShift deployment / Karaf ** karaf distro built by karaf maven plugin ** BOM there is fabric8 karaf bom AFAIK * IPaaS / integration project generated from UI ** flat classpath with spring boot ** here it gets complicated with connectors and it's mapping to camel components IMHO the best solution would be to have documentation for each target platform showing * basic project template (consumed by fuse tooling in create project wizard and in Fuse documentation) * which BOM to use * what is provided by target runtime and what is "developers responsibility to package" What do you think about that?	1	1	https://issues.jboss.org/browse/ENTESB-6878

#	Requirement	Questions	Answered	Successfully Answered	URL
11600	ENTESB-6878	[~jludvice] I think you did a good job of documenting it right here! If we want this in the user documentation we should probably file a doc JIRA. Are we missing anything here? I think we've come to consensus for Karaf, EAP, Spring Boot, FIS. The only one I don't really know is the EWS scenario.	1	1	https://issues.jboss.org/browse/ENTESB-6878
11601	ENTESB-6878	Hello [~tcunning] (CC [~tturek] [~ppecka] [~apodhrad]) Thanks for creating FUSEDOC-2080, but what is doc team supposed to put in? I recall some mentions that {{jboss-fuse-parent}} is more for internal use than official BOM. What exactly should it be? Where is link for that file? Can we please agree on concrete *GAV* for each deployment type until we pass it on to docs team?	1	0	https://issues.jboss.org/browse/ENTESB-6878
11602	ENTESB-6878	[~jludvice] * bundle deployed on standalone fuse ** runtime jboss-fuse-karaf, all dependencies provided by karaf as features ** BOM with versions aligned to karaf 4 features : org.apache.karaf:karaf * * jar/war deployed onto EAP 7 ** runtime EAP7 camel and other stuff provided by subsystem ** BOM with versions aligned to EAP : org.jboss.eap:jboss-eap-parent * fat jar / spring-boot ** flat classpath built by maven plugin from MRRC ** BOM with spring boot camel starter : io.fabric8:fabric8-project-bom-camel-spring-boot * FIS / OpenShift deployment / Sprint Boot ** flat classpath ** BOM : io.fabric8:fabric8-project-bom-camel-spring-boot * FIS / OpenShift deployment / Karaf ** karaf distro built by karaf maven plugin ** BOM : io.fabric8: fabric8-project-bom-fuse-karaf * IPaaS / integration project generated from UI ** flat classpath with spring boot ** BOM : io.fabric8:fabric8-project-bom-camel-spring-boot The one idea I have really no idea about is war/tomcat (below), I'm not familiar with it at all. * war / tomcat (Enterprise Web Server) ** runtime EWS, all dependencies pulled from MRRC ** BOM do we need to align with EWS ? does it make sense to reuse other bom?	1	1	https://issues.jboss.org/browse/ENTESB-6878
11603	ENTESB-6878	We were just discussing with [~maschmid] and he says that there were two BOM files for FIS 2.0 because of different camel versions * FIS / karaf used camel 2.17.x / aligned to fuse 6.3 * FIS / spring-boot used camel 2.18.x / because of spring-boot support in camel 2.18 Is there any reason to keep them separate if we're going to have just one version of camel? I did now want to say we need 7 different BOM files. It could be one BOM to rule all runtimes, if it's possible :) Will there be any difference in versions for standalone fuse and for FIS/karaf ?	1	0	https://issues.jboss.org/browse/ENTESB-6878
11604	ENTESB-6849	[~janstey] why don't these quickstart issues show up in fuse 7 backlog? Ignore - set component and now they show up.	0	0	https://issues.jboss.org/browse/ENTESB-6849
11605	ENTESB-6828	Looks like 2012 still has active extended support and 2014 is still within active support lifecycle (SP2 ends maintenance support in 2019). https://support.microsoft.com/en-us/lifecycle/search?alpha=SQL%20Server Agree with adding 2016. So should we adjust to drop 2012, keep 2014, and add 2016?	0	0	https://issues.jboss.org/browse/ENTESB-6828
11606	ENTESB-6828	Could you also list about the driver's version which will be used for the certification test please? I'm sure that user loves to know which sets of DB and driver we test and certified.	0	0	https://issues.jboss.org/browse/ENTESB-6828
11607	ENTESB-6810	Now the Artemis plugin is managed by the AMQ team and is already bundled in AMQ 7.0 with hawtio 1.x line. So who should proceed this task? [~aileenc] [~ataylor]	1	0	https://issues.jboss.org/browse/ENTESB-6810
11608	ENTESB-6794	I don't yet fully understand what's really the issue. Under low load most operations in the ActiveMQ tab execute immediately even on a queue with 10000 messages, so does it occur only under high load such as production? Also, we can find the "Execute" buttons only when invoking an operation in the "Operations" sub menu on a MBean. So is the user invoking MBean operations such as {{browseAsTable()}} directly instead of utilising features provided by the ActiveMQ plugin? Essentially is "the execute button to stay depressed till there is a response from the browse command" what we just need to implement for this enhancement request?	1	1	https://issues.jboss.org/browse/ENTESB-6794
11609	ENTESB-6641	What is the best way to handle these changes - should this all be done in the Jenkins pipeline? [~tcunning] [~chirino] ?	1	1	https://issues.jboss.org/browse/ENTESB-6641
11610	ENTESB-6641	There's no "fuse" startup script in bin directory. Karaf gives me this error - do I need later version of JDK? C:\Fuse\esb-project-7.0.0.fuse-000015>.\bin\karaf.bat Error occurred during initialization of VM java/lang/UnsupportedClassVersionError: java/lang/Exception : Unsupported major. minor version 52.0 C:\Fuse\esb-project-7.0.0.fuse-000015> C:\Fuse\esb-project-7.0.0.fuse-000015>java -version java version "1.8.0_101" Java(TM) SE Runtime Environment (build 1.8.0_101-b13) Java HotSpot(TM) 64-Bit Server VM (build 25.101-b13, mixed mode)	0	0	https://issues.jboss.org/browse/ENTESB-6641
11611	ENTESB-6641	I think the reason is that the startup script is using the {{JAVA_HOME}} environment variable to locate the {{java}} executable. Can you check if it points to a JRE / JDK 8 ?	0	0	https://issues.jboss.org/browse/ENTESB-6641
11612	ENTESB-6637	Is this not covered by https://10.8.63.16/view/wildfly-camel/job/fuse-eap-redhat-7.0 and friends?	1	1	https://issues.jboss.org/browse/ENTESB-6637
11613	ENTESB-6634	Hi Paolo, Thanks for the update.I will take care of formatting going ahead. I agree that keeping the container state as it is in case of wrong decryption value in profile is expected or current behaviour. But customer expecting that his application should fail in this case.In fuse it just throws WARN message and keeps on using old correct decryption value without any change. Customer saying they don't restart their containers or redeploy all applications frequently on any profile changes.So In this case if some user has added bad decryption value by mistake and forgot to restart container or redeploy application then his application will keep on using old value until next restart.Customer expecting his application should fail immediately.It should throw error message instead of WARN. If this is not a bug, can we have it as enhancement?	1	0	https://issues.jboss.org/browse/ENTESB-6634
11614	ENTESB-6634	Hi Paolo, Is it possible to suggest in which release this enhancement can be released, so that same i can update to a customer. Thanks, Swapnil	1	0	https://issues.jboss.org/browse/ENTESB-6634
11615	ENTESB-6603	Hi [~alexkieling] I need your input. I have started toying with the activity and I have just started familiarizing with Patternfly. I have noticed that there multiple distros of Patternfly: https://www.patternfly.org/download/ In particular I notice that there is a {{Angular PatternFly}} . Is that that we are using? I'm asking this mainly because I wanted to use a `DataTable` widget. But I have noticed that it has no automatic integration with Angular data binding. I could use it, but it would be like using jquery, with me directly responsible of all the state transitions. Can you comment on this? Should I follow a different approach?	1	1	https://issues.jboss.org/browse/ENTESB-6603
11616	ENTESB-6603	How are you doing, Paolo? Need any help?	1	1	https://issues.jboss.org/browse/ENTESB-6603
11617	ENTESB-6600	Hi [~stefanutti], What's left with this ticket? I'm trying to resolve ENTESB-6358, and to me most of the works here for the JMX tree seems to be complete already (awesome work!). One minor point I noticed is that the icon for leaf nodes of the tree is also a folder, which was a cog. I think a different icon from the folder nodes is better for leaf nodes. Is it something you are working on?	1	1	https://issues.jboss.org/browse/ENTESB-6600

#	Requirement	Questions	Answered	Successfully Answered	URL
11618	ENTESB-6567	With branch https://github.com/jboss-fuse/wildfly-camel/tree/ENTESB-6567 I can build and start the image successfully {code} $ docker run --rm -ti -e WILDFLY_MANAGEMENT_USER=admin -e WILDFLY_MANAGEMENT_PASSWORD=admin -p 8080:8080 -p 9990:9990 fuse/eap-camel ... 14:40:30,193 INFO [org.wildfly.extension.undertow] (MSC service thread 1-4) WFLYUT0008: Undertow HTTP listener default suspending 14:40:30,196 INFO [org.wildfly.extension.undertow] (MSC service thread 1-1) WFLYUT0008: Undertow HTTPS listener https suspending 14:40:30,195 INFO [org.wildfly.extension.undertow] (MSC service thread 1-4) WFLYUT0007: Undertow HTTP listener default stopped, was bound to 0.0.0.0:8080 14:40:30,200 INFO [org.wildfly.extension.undertow] (MSC service thread 1-1) WFLYUT0007: Undertow HTTPS listener https stopped, was bound to 0.0.0.0:8443 14:40:30,207 INFO [org.wildfly.extension.undertow] (MSC service thread 1-2) WFLYUT0004: Undertow 1.4.11.Final-redhat-1 stopping 14:40:30,308 INFO [org.jboss.as.server.deployment] (MSC service thread 1-5) WFLYSRV0028: Stopped deployment hawtio-wildfly-1.5.0.war (runtime-name: hawtio-wildfly-1.5.0.war) in 198ms 14:40:30,332 INFO [org.jboss.as] (MSC service thread 1-2) WFLYSRV0050: JBoss EAP 7.1.0.Alpha1 (WildFly Core 3.0.0.Beta10-redhat-1) stopped in 211ms {code} but later see {code} $ mvn clean install -Ddocker -am -pl docker,itests/docker ... --- T E S T S --- Running org.wildfly.camel.test.docker.CXFEndpointTest CubeDockerConfiguration: serverUri = tcp://192.168.99.103:2376 tlsVerify = false dockerServerIp = 192.168.99.103 definitionFormat = CUBE clean = false removeVolumes = true dockerContainers = containers: wildfly: alwaysPull: false await: {match: WFLYSRV0025, stdErr: true, stdOut: true, strategy: log, timeout: 15} env: [WILDFLY_MANAGEMENT_USER=admin, WILDFLY_MANAGEMENT_PASSWORD=admin] exposedPorts: [8080/tcp, 9990/tcp] image: fuse/eap-camel killContainer: false manual: false portBindings: [8080/tcp, 9990/tcp] readonlyRootfs: false removeVolumes: true networks: {} Tests run: 1, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 1.972 sec <<< FAILURE! - in org.wildfly.camel.test.docker.CXFEndpointTest org.wildfly.camel.test.docker.CXFEndpointTest Time elapsed: 1.971 sec <<< ERROR! org.arquillian.cube.spi.CubeControlException: Could not create wildfly at org.arquillian.cube.spi.CubeControlException.failedCreate(CubeControlException.java:19) at org.arquillian.cube.docker.impl.model.DockerCube.create(DockerCube.java:123) at org.arquillian.cube.impl.client.CubeLifecycleController.create(CubeLifecycleController.java:15) {code}	0	0	https://issues.jboss.org/browse/ENTESB-6567
11619	ENTESB-6534	Sorry to reopen, but I can't see 'Rest Services' tab any more. Once I fixed so it really showed up, but somehow it doesn't show again (at least for me): https://github.com/hawtio/hawtio-integration/commit/4fcd9a031d16440491c22c46795d9d66e5457495 Can you see the 'Rest Services' tab (or 'REST Services' as you recently changed)?	0	0	https://issues.jboss.org/browse/ENTESB-6534
11620	ENTESB-6533	How to monitor FIS with HOSA is documented here : https://github.com/cunningt/fuse-hawkular-agent/blob/master/minishift-hosa.md Need to add this to a blog post.	0	0	https://issues.jboss.org/browse/ENTESB-6533
11621	ENTESB-6499	What do you think [~grgrzybek], can we fix this / backport to fuse 6.3 ?	1	1	https://issues.jboss.org/browse/ENTESB-6499
11622	ENTESB-6499	Feature handling in Karaf 4 was completely rewritten. All I can do is: *console*: {noformat} JBossFuse:karaf@root> features:install test144 Error executing command: Can not resolve feature test144/0.0.0, resource mvn:org.bouncycastle/bcprov-jdk15/1.44 is not proper OSGi bundle {noformat} *logs*: {noformat} 12:32:35,883 | ERROR | l Console Thread | Console | 38 - org.apache.karaf.shell.console - 2.4.0.redhat-630247 | Exception caught while executing command java.lang.Exception: Can not resolve feature test144/0.0.0, resource mvn:org.bouncycastle/bcprov-jdk15/1.44 is not proper OSGi bundle at org.apache.karaf.features.obr.internal.ObrResolver.resolve(ObrResolver.java:95)[14:org.apache.karaf.features.obr:2.4.0.redhat-630247] at org.apache.karaf.features.internal.FeaturesServiceImpl.resolve(FeaturesServiceImpl.java:813)[11:org.apache.karaf.features.core:2.4.0.redhat-630247] at org.apache.karaf.features.internal.FeaturesServiceImpl.doInstallFeature(FeaturesServiceImpl.java:702)[11:org.apache.karaf.features.core:2.4.0.redhat-630247] at org.apache.karaf.features.internal.FeaturesServiceImpl.doInstallFeatures(FeaturesServiceImpl.java:462)[11:org.apache.karaf.features.core:2.4.0.redhat-630247] at org.apache.karaf.features.internal.FeaturesServiceImpl$1.call(FeaturesServiceImpl.java:432)[11:org.apache.karaf.features.core:2.4.0.redhat-630247] at java.util.concurrent.FutureTask.run(FutureTask.java:266)[:1.8.0_121] at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142)[:1.8.0_121] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617)[:1.8.0_121] at java.lang.Thread.run(Thread.java:745)[:1.8.0_121] {noformat} If this is ok, please create PATCH issue.	1	1	https://issues.jboss.org/browse/ENTESB-6499
11623	ENTESB-6499	I think it's good improvement. It should give user hint where to look for a problem. Thanks [~grgrzybek]. [~sjavurek] [~jsherman] Can you please open PATCH issue ? Goal of this issue is to give developer better hint on where to look for problem when Fuse fails to install feature.	1	1	https://issues.jboss.org/browse/ENTESB-6499
11624	ENTESB-6477	Can we have a few Camel apps running on Karaf? We need to show the OSGi pages to the UX team.	1	1	https://issues.jboss.org/browse/ENTESB-6477
11625	ENTESB-6467	Hi [~grgrzybek], what are the easiest steps to reproduce the issue? I found creating certificates for whole ensemble and creating ssl profile which is not so easy to set up. Thanks.	1	1	https://issues.jboss.org/browse/ENTESB-6467
11626	ENTESB-6460	[~jdyson] Hey this 'curl -XDELETE 'http://<Hostname>:9200/insight-2016-07-06' thing they are doing to solve the space issues, does that also delete data that can be viewed from the UI? What do you think about this case?	1	0	https://issues.jboss.org/browse/ENTESB-6460
11627	ENTESB-6418	hello [~tcunning], is there any special way to verify this issue? If I update soap-mtom quickstart switchyard.xml (Fuse 630269 build on EAP) with: <soap:endpointConfig configFile="META-INF/jaxws-endpoint-config.xml" configName="SwitchYard-Endpoint-Config"/> + add jaxws-endpoint-config.xml to main/resource/META-INF as per ENTESB-6305 SY reproducer, I even cannot build it (mvn clean package), it fails with this exception: testSwitchYardWebService(org.switchyard.quickstarts.soap.mtom.SoapMtomTest) Time elapsed: 0.852 s <<< ERROR! javax.xml.ws.soap.SOAPFaultException: Could not generate the XML stream caused by: org.xml.sax.SAXParseException: cvc-type.3.1.2: Element 'imageResponse' is a simple type, so it must have no element information item [children].. at org.switchyard.quickstarts.soap.mtom.SoapMtomTest.testSwitchYardWebService(SoapMtomTest.java:48) Caused by: org.apache.cxf.binding.soap.SoapFault: Could not generate the XML stream caused by: org.xml.sax.SAXParseException: cvc-type.3.1.2: Element 'imageResponse' is a simple type, so it must have no element information item [children]..	1	1	https://issues.jboss.org/browse/ENTESB-6418
11628	ENTESB-6393	[~alexkieling], [~stefanutti], [~stlewis], Should hawtio-log be yet another separate project? Where will the plugin be best located?	1	1	https://issues.jboss.org/browse/ENTESB-6393

#	Requirement	Questions	Answered	Successfully Answered	URL
11629	ENTESB-6393	Ah, I think Antonin's point is valid and really insightful. FWIW, Fuse 6.x already has two different log plugins for standalone and Fabric8: https://access.redhat.com/documentation/en-us/red_hat_jboss_fuse/6.3/html-single/management_console_user_guide/#FMCUG_Insight So I think it's natural even for hawtio v2 to separate the Log plugin for standalone containers from the plugin for OpenShift / Kubernetes. I thought that this ticket is targeted only for the standalone version, which only takes care of standalone containers like Karaf, EAP, and SpringBoot. For this, now I think _hawtio-log_ is the good starting point. For logging on OpenShift or logging for FIS, I'm really not sure how much capabilities are already implemented in either OpenShift Console or Fabric8 projects, but maybe either of them is the good place to accommodate such a plugin. I noticed Fabric8 already has some hawtio plugins like _hawtio-forge_ in it. Or could _hawtio-kubernetes_ be a good home for such a plugin?	1	1	https://issues.jboss.org/browse/ENTESB-6393
11630	ENTESB-6393	Is this task about the [Logs page|https://access.redhat.com/documentation/en-us/red_hat_jboss_fuse/6.3/html-single/management_console_user_guide/#FMCUG_Container_Logs] or the [Insight plugin|https://access.redhat.com/documentation/en-us/red_hat_jboss_fuse/6.3/html-single/management_console_user_guide/#FMCUG_Insight]? If we are going to have multiple logging solutions, it makes sense having a repository to each one. We can register the Angular modules (hawtio plugins) we need to create the final console. What name should we give to the repository and who want's to work on it?	1	1	https://issues.jboss.org/browse/ENTESB-6393
11631	ENTESB-6364	[~stefanutti] Have you migrated JMX tree on the left hand side to {{patternfly-bootstrap-treeview}}? Or will you?	1	1	https://issues.jboss.org/browse/ENTESB-6364
11632	ENTESB-6358	Maybe we should break this task into smaller tasks. What do you think?	1	1	https://issues.jboss.org/browse/ENTESB-6358
11633	ENTESB-6357	I'm curious about how it's to be implemented in hawtio v2. RBAC is one of the major parts that have been improved in Fuse 6.3 (ENTESB-4055 and ENTESB-5219). But that improvement was implemented mostly at the Java backend. How can we accommodate the same capabilities in the hawtio v2 architecture?	1	0	https://issues.jboss.org/browse/ENTESB-6357
11634	ENTESB-6308	io.fabric8.fabric-agent - 1.2.0.redhat-630193 | Refreshing bundles: 2016-12-01 17:20:29,680 | INFO | 6.0.1-1-thread-1 | Agent | 86 - io.fabric8.fabric-agent - 1.2.0.redhat-630193 | blueprint.xml / 0.0.0 (Bundle will be updated) 2016-12-01 17:20:29,699 | INFO | 6.0.1-1-thread-1 | Agent | 86 - io.fabric8.fabric-agent - 1.2.0.redhat-630193 | Starting bundles: 2016-12-01 17:20:29,700 | INFO | 6.0.1-1-thread-1 | Agent | 86 - io.fabric8.fabric-agent - 1.2.0.redhat-630193 | blueprint.xml / 0.0.0 2016-12-01 17:20:29,701 | INFO | loyer-3-thread-1 | Deployer | 86 - io.fabric8.fabric-agent - 1.2.0.redhat-630193 | Scheduled start for bundle:blueprint.xml with a timeout limit of 60 seconds 2016-12-01 17:20:29,731 | INFO | loyer-3-thread-1 | BlueprintCamelContext | 127 - org.apache.camel.camel-core - 2.17.0.redhat-630193 | Apache Camel 2.17.0.redhat-630193 (CamelContext: camel-2) is starting 2016-12-01 17:20:29,731 | INFO | loyer-3-thread-1 | ManagedManagementStrategy | 127 - org.apache.camel.camel-core - 2.17.0.redhat-630193 | JMX is enabled 2016-12-01 17:20:30,584 | INFO | loyer-3-thread-1 | DefaultRuntimeEndpointRegistry | 127 - org.apache.camel.camel-core - 2.17.0.redhat-630193 | Runtime endpoint registry is in extended mode gathering usage statistics of all incoming and outgoing endpoints (cache limit: 1000) 2016-12-01 17:20:30,594 | INFO | loyer-3-thread-1 | ElasticsearchEndpoint | 128 - org.apache.camel.camel-elasticsearch - 2.17.0.redhat-630193 | Starting local ElasticSearch server 2016-12-01 17:20:30,596 | INFO | loyer-3-thread-1 | node | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] version[2.2.0], pid[16232], build[NA/NA] 2016-12-01 17:20:30,596 | INFO | loyer-3-thread-1 | node | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] initializing ... 2016-12-01 17:20:30,596 | INFO | loyer-3-thread-1 | plugins | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] modules [], plugins [], sites [] 2016-12-01 17:20:30,597 | INFO | loyer-3-thread-1 | env | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] using [1] data paths, mounts [[/home (/dev/mapper/centos-home)]], net usable_space [195.9gb], net total_space [418.4gb], spins? [no], types [xfs] 2016-12-01 17:20:30,597 | INFO | loyer-3-thread-1 | env | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] heap size [683mb], compressed ordinary object pointers [true] 2016-12-01 17:20:30,698 | INFO | loyer-3-thread-1 | node | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] initialized 2016-12-01 17:20:30,698 | INFO | loyer-3-thread-1 | node | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] starting ... 2016-12-01 17:20:30,699 | INFO | loyer-3-thread-1 | transport | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] publish_address {local[2]}, bound_addresses {local[2]} 2016-12-01 17:20:30,700 | INFO | loyer-3-thread-1 | discovery | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] elasticsearch/Bo5ICrC8R0yobD1RvOmilw 2016-12-01 17:20:30,705 | INFO | updateTask][T#1] | service | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] master {new {Captain Marvel}{Bo5ICrC8R0yobD1RvOmilw}{local}{local[2]}{local=true}}, removed {{Captain Marvel}{UESQsLgrS42-IT6OCdDgVw}{local}{local[2]}{local=true},}, reason: local-disco-initial_connect(master) 2016-12-01 17:20:30,706 | INFO | loyer-3-thread-1 | node | 136 - org.apache.servicemix.bundles.elasticsearch - 2.2.0.1 | [Captain Marvel] started 2016-12-01 17:20:30,710 | INFO | loyer-3-thread-1 | BlueprintCamelContext | 127 - org.apache.camel.camel-core - 2.17.0.redhat-630193 | AllowUseOriginalMessage is enabled. If access to the original message is not needed, then its recommended to turn this option off as it may improve performance. 2016-12-01 17:20:30,711 | INFO | loyer-3-thread-1 | BlueprintCamelContext | 127 - org.apache.camel.camel-core - 2.17.0.redhat-630193 | StreamCaching is not in use. If using streams then its recommended to enable stream caching. See more	1	1	https://issues.jboss.org/browse/ENTESB-6308
11635	ENTESB-6270	What needs to be done here? Has this upstream feature already reached a camel-redhat release and is not working in the context of EAP or are we still waiting for this to reach camel-redhat?	1	1	https://issues.jboss.org/browse/ENTESB-6270
11636	ENTESB-6205	Can this issue be closed?	1	0	https://issues.jboss.org/browse/ENTESB-6205
11637	ENTESB-6167	[~nibbio84], [~tirelli] Isn't this JIRA duplicate of [BPMSPL-511]. We already have new component developed by [~mfusco] (see [1]). What will be the difference between components developed by BxMS team and Fuse team? [1] https://github.com/kiegroup/droolsjbpm-integration/tree/master/kie-camel	1	1	https://issues.jboss.org/browse/ENTESB-6167
11638	ENTESB-6141	I've disabled the following jobs and deleted the fabric8 tab: fabric8-1.x-website fabric8-2.0-deploy fabric8-2.0-quickstarts-integration-tests-on-jube fabric8-2.0-quickstarts-integration-tests fabric8-2.0-quickstarts-release fabric8-2.0-release fabric8-2.x-quickstarts-deploy fabric8-2.x-website-deploy fabric8-docker-images-checkin fabric8-docker-images-on-new-tag fabric8-docker-images fabric8-release2 The jobs listed below are still used for Fuse 6.2 and Fuse 6.3 patch builds, and can still be found on the Fuse 6.2 and 6.3 tabs fabric8-6.2.x.redhat-patch-checkin fabric8-6.2.x.redhat-patch-deploy fabric8-6.2.x.redhat-patch-platform fabric8-6.2.x.redhat-patch-pullreq fabric8-6.3.x-redhat-hourly-kicker fabric8-6.3.x.redhat-checkin fabric8-6.3.x.redhat-deploy fabric8-6.3.x.redhat-platform fabric8-6.3.x.redhat-pullrequest	0	0	https://issues.jboss.org/browse/ENTESB-6141
11639	ENTESB-6140	I don't understand why the severity is decreased to minor. if the tests of the project are not running fine everytime, shouldn't it block the release and be the task with the highest priority?	0	0	https://issues.jboss.org/browse/ENTESB-6140
11640	ENTESB-6140	but you ignored tests? https://github.com/jboss-fuse/camel/commit/359fd4e7487a64c02a411eac79c3fb68d90340b0 https://github.com/jboss-fuse/camel/commit/971315d475589279c871e82482f01cb41ed63f14 seems not that the issue is really fixed. Can be ok to disable tests for a moment to move forward on other topics and have a green CI but not for long term. Can you open at least JIRA tasks/issues to fix it separately?	0	0	https://issues.jboss.org/browse/ENTESB-6140

#	Requirement	Questions	Answered	Successfully Answered	URL
11641	ENTESB-6134	You don't actually have to click the "Dispatch Router" link in the nav bar. Going straight to the logging console is enough. [~kearls] Let me give you a detailed reproduction steps, and a video. {noformat} $ wget http://download-node-02.eng.bos.redhat.com/rcm-guest/staging/amq/AMQ-BROKER-7.0.0.ER13-1/A-MQ7-7.0.0.ER13-1-redhat-1-bin.zip $ unzip A-MQ7-7.0.0.ER13-1-redhat-1-bin.zip $ A-MQ7-7.0.0.ER13-1-redhat-1/bin/artemis create amq7er13i0 Creating ActiveMQ Artemis instance at: /home/jdanek/Downloads/amq/amq7er13i0 --user: mandatory: Please provide the default username: admin --password: is mandatory with this configuration: Please provide the default password: admin --role: mandatory: Please provide the default role: admin --allow-anonymous | --require-login: mandatory: Allow anonymous access? (Y/N): y Auto tuning journal ... $ amq7er13i0/bin/artemis run {noformat} Now follow the video. Copy the url from output of last command into a browser, and so on...	1	1	https://issues.jboss.org/browse/ENTESB-6134
11642	ENTESB-6134	Is this still an issue for HawtIO V2 / Fuse 7? [~alexkieling] thought the dependency on flash is gone.	1	1	https://issues.jboss.org/browse/ENTESB-6134
11643	ENTESB-6134	When you say "current Web console" do you mean the one for Fuse 6.x? If this is fixed for Fuse 7 then I think we can mark it as Resolved and if anyone wants the fix backoprted to 6.x then open a PATCH jiira.	1	1	https://issues.jboss.org/browse/ENTESB-6134
11644	ENTESB-6134	[~tadayosi] We deprecated zero-clipboard back in January and I think it's safe to remove it now. We provide a much better alternative based on clipboard.js. As part of the migration from bower to npm, I'm planning to remove zero-clipboard from the hawtio-ui project. Do you see any issues with that?	1	1	https://issues.jboss.org/browse/ENTESB-6134
11645	ENTESB-6125	Is this something we can address in Tooling? It may affect other projects as well, not just quickstarts.	1	1	https://issues.jboss.org/browse/ENTESB-6125
11646	ENTESB-6118	Why the quickstarts are different than for instance https://github.com/jboss-fuse/fuse/blob/master/esb/esb-assembly/pom.xml#L28 or https://github.com/jboss-fuse/fuse/blob/master/shared/shared-commands/pom.xml#L30 ? There are not part of the same release process?	1	0	https://issues.jboss.org/browse/ENTESB-6118
11647	ENTESB-6104	steps... Keycloak says: {code} 16:47:08,667 ERROR [org.keycloak.services] (default task-52) KC-SERVICES0037: Error creating admin: java.lang.RuntimeException: Unable to find client role mappings for client: realm-management at org.keycloak.models.utils.RepresentationToModel.createRoleMappings(RepresentationToModel.java:1498) at org.keycloak.models.utils.RepresentationToModel.createUser(RepresentationToModel.java:1364) at org.keycloak.partialimport.UsersPartialImport.create(UsersPartialImport.java:111) at org.keycloak.partialimport.UsersPartialImport.create(UsersPartialImport.java:38) at org.keycloak.partialimport.AbstractPartialImport.doImport(AbstractPartialImport.java:120) at org.keycloak.partialimport.PartialImportManager.saveResources(PartialImportManager.java:74) at org.keycloak.services.resources.admin.RealmAdminResource.partialImport(RealmAdminResource.java:842) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.jboss.resteasy.core.MethodInjectorImpl.invoke(MethodInjectorImpl.java:139) at org.jboss.resteasy.core.ResourceMethodInvoker.invokeOnTarget(ResourceMethodInvoker.java:295) at org.jboss.resteasy.core.ResourceMethodInvoker.invoke(ResourceMethodInvoker.java:249) at org.jboss.resteasy.core.ResourceLocatorInvoker.invokeOnTargetObject(ResourceLocatorInvoker.java:138) at org.jboss.resteasy.core.ResourceLocatorInvoker.invoke(ResourceLocatorInvoker.java:107) at org.jboss.resteasy.core.ResourceLocatorInvoker.invokeOnTargetObject(ResourceLocatorInvoker.java:133) at org.jboss.resteasy.core.ResourceLocatorInvoker.invoke(ResourceLocatorInvoker.java:101) at org.jboss.resteasy.core.SynchronousDispatcher.invoke(SynchronousDispatcher.java:395) at org.jboss.resteasy.core.SynchronousDispatcher.invoke(SynchronousDispatcher.java:202) at org.jboss.resteasy.plugins.server.servlet.ServletContainerDispatcher.service(ServletContainerDispatcher.java:221) at org.jboss.resteasy.plugins.server.servlet.HttpServletDispatcher.service(HttpServletDispatcher.java:56) at org.jboss.resteasy.plugins.server.servlet.HttpServletDispatcher.service(HttpServletDispatcher.java:51) at javax.servlet.http.HttpServlet.service(HttpServlet.java:790) at io.undertow.servlet.handlers.ServletHandler.handleRequest(ServletHandler.java:85) at io.undertow.servlet.handlers.FilterHandler$FilterChainImpl.doFilter(FilterHandler.java:129) at org.keycloak.services.filters.KeycloakSessionServletFilter.doFilter(KeycloakSessionServletFilter.java:90) at io.undertow.servlet.core.ManagedFilter.doFilter(ManagedFilter.java:60) at io.undertow.servlet.handlers.FilterHandler$FilterChainImpl.doFilter(FilterHandler.java:131) at io.undertow.servlet.handlers.FilterHandler.handleRequest(FilterHandler.java:84) at io.undertow.servlet.handlers.security.ServletSecurityRoleHandler.handleRequest(ServletSecurityRoleHandler.java:62) at	0	0	https://issues.jboss.org/browse/ENTESB-6104
11648	ENTESB-6104	It works, you are right [~sonicaaaa]. Tried on Fuse 6.3 208, Keycloak 2.2.0 and 2.3.0.	0	0	https://issues.jboss.org/browse/ENTESB-6104
11649	ENTESB-6085	[~rhn-support-shiggs] FWIW you can already disable any hawtio tabs by \[user\] (at the top right) -> Preferences -> Plugins -> clicking Enabled icons. Isn't it satisfying enough?	1	1	https://issues.jboss.org/browse/ENTESB-6085
11650	ENTESB-6081	Starting to take a look at this. # downloaded the sample project and unzipped it locally # imported as an existing project # did Maven->Update Project to get things up to snuff # set the target runtime to a local Fuse 6.2.1 server # the project is not resolving the CXF_CLASSPATH_CONTAINER in the project library - where is that coming from?	1	1	https://issues.jboss.org/browse/ENTESB-6081
11651	ENTESB-6081	When you say "ignored" - what does that mean? Ignored in the runtime? The transformers are being created in the tooling, so they're not ignored there by what I'm seeing.	1	1	https://issues.jboss.org/browse/ENTESB-6081
11652	ENTESB-5985	Is this already in our Fuse 7 builds?	0	0	https://issues.jboss.org/browse/ENTESB-5985
11653	ENTESB-5930	Why do you want it as a module?	1	0	https://issues.jboss.org/browse/ENTESB-5930
11654	ENTESB-5922	Ericsson are chasing me on this - can this Jira be re-opened and re-reveiwed as stated in the comment by Susan on 12Nov2015 ? Thanks	0	0	https://issues.jboss.org/browse/ENTESB-5922
11655	ENTESB-5922	[~fbolton]Is this something you think we can document? We use it internally for testing so we should be able to provide guidance for our customers. Fabric8 itests are public and may be a good example: https://github.com/jboss-fuse/fabric8/tree/1.2.0.redhat-6-3-x/itests [1] Will the examples help for now? [1] Pax Exam tests are intermixed in here but there are arquillian examples. Also, please fuse-eng-help if you need further help getting the tests running.	1	1	https://issues.jboss.org/browse/ENTESB-5922
11656	ENTESB-5915	[~thomas.diesler], what are the chances that Fuse on EAP 6.3 will run against EAP 7?	1	1	https://issues.jboss.org/browse/ENTESB-5915

#	Requirement	Questions	Answered	Successfully Answered	URL
11657	ENTESB-5885	Hi, I reopen this issue because i'm not able to have the test working and I'm not sure that the local deployment is working. i have no error during camel start but how can I test that the route is really working? (after I added camel-maven-plugin to allow local deployment) here is my current work on the subject: https://github.com/apupier/fuseide/tree/FUSETOOLS-1981-ActivateLocalLaunchCOnfiguration	0	0	https://issues.jboss.org/browse/ENTESB-5885
11658	ENTESB-5875	Ok. That makes sense. As such, I'm guessing that something along the lines of this CamelSpringTestSupport-based test doesn't make much sense: https://github.com/fusesource/fuseide/blob/master/templates/simple/cbr/simple-fuse-cbr-spring/src/test/java/my/company/camel/cbr/CamelXmlTest.java Are you referring to this example? https://github.com/jboss-fuse/quickstarts/tree/master/cxf/rest Any idea on a Java DSL-based version?	1	1	https://issues.jboss.org/browse/ENTESB-5875
11659	ENTESB-5875	effective model for org.apache.cxf.archetype:cxf-jaxrs-service:bundle:3.0.4.redhat-620133 [WARNING] 'build.plugins.plugin.version' for org.apache.maven.plugins:maven-surefire-plugin is missing. @ line 121, column 21 [WARNING] 'build.plugins.plugin.version' for org.apache.felix:maven-bundle-plugin is missing. @ line 128, column 21 [WARNING] [WARNING] It is highly recommended to fix these problems because they threaten the stability of your build. [WARNING] [WARNING] For this reason, future Maven versions might no longer support building such malformed projects. [WARNING] [INFO] [INFO] -- [INFO] Building Simple CXF JAX-RS service using spring configuration 3.0.4.redhat-620133 [INFO] -- [INFO] [INFO] --- maven-resources-plugin:2.7:resources (default-resources) @ cxf-jaxrs-service --- [WARNING] Using platform encoding (UTF-8 actually) to copy filtered resources, i.e. build is platform dependent! [INFO] Copying 1 resource [INFO] [INFO] --- maven-compiler-plugin:2.5.1-jboss-2:compile (default-compile) @ cxf-jaxrs-service --- [INFO] Nothing to compile - all classes are up to date [INFO] [INFO] --- maven-resources-plugin:2.7:testResources (default-testResources) @ cxf-jaxrs-service --- [WARNING] Using platform encoding (UTF-8 actually) to copy filtered resources, i.e. build is platform dependent! [INFO] skip non existing resourceDirectory C:\git\fuseide\templates\medium\cxf-jaxrs\medium-fuse-cxf-jaxrs-spring\src\test\resources [INFO] [INFO] --- maven-compiler-plugin:2.5.1-jboss-2:testCompile (default-testCompile) @ cxf-jaxrs-service --- [INFO] Nothing to compile - all classes are up to date [INFO] [INFO] --- maven-surefire-plugin:2.19.1:test (default-test) @ cxf-jaxrs-service --- --- T E S T S --- Running com.mycompany.cxf.jaxrs.service.HelloWorldTest Aug 19, 2016 4:21:17 PM org.apache.cxf.phase.PhaseInterceptorChain doDefaultLogging WARNING: Interceptor for {http://localhost:8181/cxf/rest/hello/jsonBean}WebClient has thrown exception, unwinding now org.apache.cxf.interceptor.Fault: Could not send Message. at org.apache.cxf.interceptor.MessageSenderInterceptor$MessageSenderEndingInterceptor.handleMessage(MessageSenderInterceptor.java:64) at org.apache.cxf.phase.PhaseInterceptorChain.doIntercept(PhaseInterceptorChain.java:307) at org.apache.cxf.jaxrs.client.AbstractClient.doRunInterceptorChain(AbstractClient.java:624) at org.apache.cxf.jaxrs.client.WebClient.doChainedInvocation(WebClient.java:1094) at org.apache.cxf.jaxrs.client.WebClient.doInvoke(WebClient.java:893) at org.apache.cxf.jaxrs.client.WebClient.doInvoke(WebClient.java:864) at org.apache.cxf.jaxrs.client.WebClient.invoke(WebClient.java:330) at org.apache.cxf.jaxrs.client.WebClient.post(WebClient.java:339) at com.mycompany.cxf.jaxrs.service.HelloWorldTest.testJsonRoundtrip(HelloWorldTest.java:44) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:50) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:47) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:17) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:325) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:78) at	0	0	https://issues.jboss.org/browse/ENTESB-5875
11660	ENTESB-5874	I have several points for which i'm not able to have the project working: - mvn camel:run for local deployment is not working, it wasn't configured so I added the dependency to camel-maven-plugin but then it complains about asm artefacts - the test is not played during build (skipTests=true) and I'm not able to have it working (ConnectException: Connection refused) And last question is: how can I test that the route is working? what should I call? i tried to call http://localhost:8181/cxf/report/ from web browser but I get connection refused too here is my current work for the subject: here is my current work on the subject: https://github.com/apupier/fuseide/tree/FUSETOOLS-1981-ActivateLocalLaunchCOnfiguration	0	0	https://issues.jboss.org/browse/ENTESB-5874
11661	ENTESB-5847	I've been trying to recreate this error with no success. I have been combing through the code to see where it could be improved. The only obvious point seems to avoid the NPE that happens in [SubscriptionHelper|https://github.com/jboss-fuse/camel/blob/2.15.0.redhat-6-2-x-patch/components/camel-salesforce/camel-salesforce-component/src/main/java/org/apache/camel/component/salesforce/internal/streaming/SubscriptionHelper.java#L251]. But that exception is logged and swallowed in the BayeuxClient, so it's not the root cause of the Camel context shutdown. The primary problem is that BayeuxClient gets an early EOF from Salesforce. It's too late for the component to do anything at that point. BayeuxClient is supposed to be able to recover from a network error and reconnect, with a 1s increment and 30s max backoff, but I looks like it doesn't happen in the fixed 40s subscribe timeout. I suspect the problem might be with the user's JVM version or the network environment, which is causing a low level SSL disconnect. Can the user test Salesforce streaming with the [workbench|https://developer.salesforce.com/docs/atlas.en-us.198.0.salesforce1api.meta/salesforce1api/quick_start_workbench.htm], to see if they can recreate the error. It's not strictly equivalent, since it's browser based and doesn't use Jetty client and BayeuxClient like the component, but it's worth a try to validate the environment. The user could also try with the latest JDK if it's not being used already and see if that makes a difference in the SSL connection handling.	1	1	https://issues.jboss.org/browse/ENTESB-5847

#	Requirement	Questions	Answered	Successfully Answered	URL
11662	ENTESB-5845	When I try to "Run as local camel context", it seems to run fine but I have this stacktrace: {noformat} [Blueprint Extender: 1] DefaultRuntimeEndpointRegistry INFO Runtime endpoint registry is in extended mode gathering usage statistics of all incoming and outgoing endpoints (cache limit: 1000) [KTOP-9NT300B:1099/jmxrmi/camel] DefaultManagementAgent WARN Could not start JMXConnector thread at: service:jmx:rmi:///jndi/rmi://DESKTOP-9NT300B:1099/jmxrmi/camel. JMX Connector not in use. java.io.IOException: Cannot bind to URL [rmi://DESKTOP-9NT300B:1099/jmxrmi/camel]: javax.naming.NameAlreadyBoundException: jmxrmi/camel [Root exception is java.rmi.AlreadyBoundException: jmxrmi/camel] at javax.management.remote.rmi.RMIConnectorServer.newIOException(RMIConnectorServer.java:827) at javax.management.remote.rmi.RMIConnectorServer.start(RMIConnectorServer.java:432) at org.apache.camel.management.DefaultManagementAgent$1.run(DefaultManagementAgent.java:559) at java.lang.Thread.run(Thread.java:745) Caused by: javax.naming.NameAlreadyBoundException: jmxrmi/camel [Root exception is java.rmi.AlreadyBoundException: jmxrmi/camel] at com.sun.jndi.rmi.registry.RegistryContext.bind(RegistryContext.java:143) at com.sun.jndi.toolkit.url.GenericURLContext.bind(GenericURLContext.java:228) at javax.naming.InitialContext.bind(InitialContext.java:425) at javax.management.remote.rmi.RMIConnectorServer.bind(RMIConnectorServer.java:644) at javax.management.remote.rmi.RMIConnectorServer.start(RMIConnectorServer.java:427) ... 2 more Caused by: java.rmi.AlreadyBoundException: jmxrmi/camel at sun.rmi.registry.RegistryImpl.bind(RegistryImpl.java:183) at sun.rmi.registry.RegistryImpl_Skel.dispatch(Unknown Source) at sun.rmi.server.UnicastServerRef.oldDispatch(UnicastServerRef.java:411) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:272) at sun.rmi.transport.Transport$1.run(Transport.java:200) at sun.rmi.transport.Transport$1.run(Transport.java:197) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:196) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:568) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run0(TCPTransport.java:826) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.lambda$run$0(TCPTransport.java:683) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:682) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) at sun.rmi.transport.StreamRemoteCall.exceptionReceivedFromServer(StreamRemoteCall.java:276) at sun.rmi.transport.StreamRemoteCall.executeCall(StreamRemoteCall.java:253) at sun.rmi.server.UnicastRef.invoke(UnicastRef.java:379) at sun.rmi.registry.RegistryImpl_Stub.bind(Unknown Source) at com.sun.jndi.rmi.registry.RegistryContext.bind(RegistryContext.java:141) ... 6 more {noformat} it is as if there is another Camel running but I restarted my machine to be sure that it is not the case. [~lburgazzoli] any idea?	1	1	https://issues.jboss.org/browse/ENTESB-5845
11663	ENTESB-5845	What needs to be modified in order that the template become the best one to take as foundation for a real project? This is the exact goal of the templates.	1	1	https://issues.jboss.org/browse/ENTESB-5845
11664	ENTESB-5845	Luca, can you just make the template project work without the camel:run? As Aurelien said we want to have a nearly perfect template and this template is then only supposed to be deployed to a Fuse Runtime for now, no local launch or debug. Can you update the sources for that?	1	1	https://issues.jboss.org/browse/ENTESB-5845
11665	ENTESB-5845	--- Running com.mycompany.camel.activemq.ActiveMQBlueprintTest [main] CamelBlueprintHelper INFO Using Blueprint XML file: /C:/install/eclipse-neon/eclipse/ws_fusetooling_dev%20with%20space/amq-blue/target/classes/OSGI-INF/blueprint/camel-blueprint.xml [Thread-0] RawBuilder INFO Copy thread finished. [main] Activator INFO Camel activator starting [main] Activator INFO Camel activator started [Blueprint Extender: 1] BlueprintContainerImpl INFO Bundle ActiveMQBlueprintTest/1.0.0 is waiting for namespace handlers [http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.0.0, http://camel.apache.org/schema/blueprint] [Blueprint Extender: 2] BlueprintContainerImpl INFO Bundle camel-activemq-blueprint/1.0.0.SNAPSHOT is waiting for namespace handlers [http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.0.0, http://camel.apache.org/schema/blueprint] [Blueprint Extender: 2] BlueprintContainerImpl INFO Bundle camel-activemq-blueprint/1.0.0.SNAPSHOT is waiting for namespace handlers [http://aries.apache.org/blueprint/xmlns/blueprint-cm/v1.0.0] [Blueprint Extender: 2] BeanRecipe WARN Object to be destroyed is not an instance of UnwrapperedBeanHolder, type: null [Blueprint Extender: 3] BeanRecipe WARN Object to be destroyed is not an instance of UnwrapperedBeanHolder, type: null [Blueprint Extender: 3] BlueprintContainerImpl ERROR Unable to start blueprint container for bundle ActiveMQBlueprintTest/1.0.0 org.osgi.service.blueprint.container.ComponentDefinitionException: Unable to initialize bean .camelBlueprint.factory.amq-example-context at org.apache.aries.blueprint.container.BeanRecipe.runBeanProcInit(BeanRecipe.java:714) at org.apache.aries.blueprint.container.BeanRecipe.internalCreate2(BeanRecipe.java:824) at org.apache.aries.blueprint.container.BeanRecipe.internalCreate(BeanRecipe.java:787) at org.apache.aries.blueprint.di.AbstractRecipe$1.call(AbstractRecipe.java:79) at java.util.concurrent.FutureTask.run(FutureTask.java:266) at org.apache.aries.blueprint.di.AbstractRecipe.create(AbstractRecipe.java:88) at org.apache.aries.blueprint.container.BlueprintRepository.createInstances(BlueprintRepository.java:247) at org.apache.aries.blueprint.container.BlueprintRepository.createInstance(BlueprintRepository.java:230) at org.apache.aries.blueprint.container.BlueprintRepository.create(BlueprintRepository.java:145) at org.apache.aries.blueprint.container.BlueprintContainerImpl.getComponentInstance(BlueprintContainerImpl.java:754) at org.apache.camel.blueprint.handler.CamelNamespaceHandler$CamelDependenciesFinder.process(CamelNamespaceHandler.java:980) at org.apache.aries.blueprint.container.BlueprintContainerImpl.processProcessors(BlueprintContainerImpl.java:529) at org.apache.aries.blueprint.container.BlueprintContainerImpl.doRun(BlueprintContainerImpl.java:361) at org.apache.aries.blueprint.container.BlueprintContainerImpl.run(BlueprintContainerImpl.java:269) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:511) at java.util.concurrent.FutureTask.run(FutureTask.java:266) at org.apache.aries.blueprint.container.ExecutorServiceWrapper.run(ExecutorServiceWrapper.java:106) at org.apache.aries.blueprint.utils.threading.impl.DiscardableRunnable.run(DiscardableRunnable.java:48) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:511) at java.util.concurrent.FutureTask.run(FutureTask.java:266) at	1	1	https://issues.jboss.org/browse/ENTESB-5845
11666	ENTESB-5845	[~ppecka] i was asking because we don't provide the Java dsl based templates. So where did you get it? The failure that you have are also here when you use the fixed readme instructions?	1	1	https://issues.jboss.org/browse/ENTESB-5845

#	Requirement	Questions	Answered	Successfully Answered	URL
11667	ENTESB-5845	What does samples from JBDS9 mean? We are only shipping our wizard which has its own set of templates. If you don't use them you are probably using runtime provided quickstarts the tooling is not responsible for in first line. Though if you find issues with those quickstarts we can raise the issues at ENTESB jira for you.	1	1	https://issues.jboss.org/browse/ENTESB-5845
11668	ENTESB-5845	[~aurelien.pupier] as described by this issue. i found the ones in the git repository, and then i've used JBDS9 directly. and i think the JBDS9 provides the java dsl template for activemq, see https://github.com/fusesource/fuseide/tree/master/editor/plugins/org.fusesource.ide.projecttemplates/template-sources/simple/amq [~lhein] can you be more specific in terms of above mentioned "wizard"? i might be using wrong terminology, althrough in JBDS9 I was using some "wizard" to obtain blueprint and spring example. So we probably need to be clear on purpose of this issue. [~lhein] can you help me here? what exactly is "expected behaviour" one should see to say "fixed"? If the fixed means "customer is provided with working templates" then I suggest correcting readme files.	1	1	https://issues.jboss.org/browse/ENTESB-5845
11669	ENTESB-5806	Hi Stephen, We can look at this in the 7.X timeframe, however, in the near term would they be able to configure fabric? If so, there is a master slave component in there that should handle this. Susan	1	1	https://issues.jboss.org/browse/ENTESB-5806
11670	ENTESB-5791	Hi [~sonicaaaa], Could you please advise me how to test this fix? I have tried to run fuse with *-Djolokia.historyMaxEntries=13*. The default value is 10. If I check it via Hawtio (JMX - jolokia - Config) the value of *History max entires* parameter is 10. Supposed this is a good approach the fix doesn't work or Hawtio reads default value and doesn't take into account the changed value.	1	1	https://issues.jboss.org/browse/ENTESB-5791
11671	ENTESB-5791	[~mastepan] Yes, that should be a correct approach to test it. How have you passed {{-Djolokia.historyMaxEntries=13}} to Fuse? Note you need to use {{EXTRA_JAVA_OPTS}} env variable to pass them to Fuse at startup, something like: {code} $ EXTRA_JAVA_OPTS='-Djolokia.historyMaxEntries=13' bin/fuse {code} or adding the line in {{bin/setenv}}: {code} export EXTRA_JAVA_OPTS="-Djolokia.historyMaxEntries=13" {code}	1	0	https://issues.jboss.org/browse/ENTESB-5791
11672	ENTESB-5765	Why do they need to secure the list of endpoints?	1	1	https://issues.jboss.org/browse/ENTESB-5765
11673	ENTESB-5694	it is published only for SNAPSHOT build. For tagged version, shouldn't we provide it to a specific place? if not, I suppose that we "just" have to modify perfectus build to let him publish also the tagged version	1	1	https://issues.jboss.org/browse/ENTESB-5694
11674	ENTESB-5675	Currently, I cannot build wfc-2.4-redhat ``` [ERROR] Failed to execute goal on project wildfly-camel-modules: Could not resolve dependencies for project org.wildfly.camel:wildfly-camel-modules:pom:2.4.0.redhat-630000-SNAPSHOT: Could not find artifact org.slf4j:slf4j-log4j12:jar:1.7.7.redhat-2 in central (http://repo.maven.apache.org/maven2) -> [Help 1] ``` This artefact is missing in all of these locations ``` Downloading: http://repo.maven.apache.org/maven2/org/slf4j/slf4j-log4j12/1.7.7.redhat-2/slf4j-log4j12-1.7.7.redhat-2.jar Downloading: http://repository.jboss.org/nexus/content/groups/public/org/slf4j/slf4j-log4j12/1.7.7.redhat-2/slf4j-log4j12-1.7.7.redhat-2.jar Downloading: https://maven.repository.redhat.com/ga/org/slf4j/slf4j-log4j12/1.7.7.redhat-2/slf4j-log4j12-1.7.7.redhat-2.jar Downloading: https://maven.repository.redhat.com/earlyaccess/all/org/slf4j/slf4j-log4j12/1.7.7.redhat-2/slf4j-log4j12-1.7.7.redhat-2.jar Downloading: http://repo.springsource.org/libs-release-remote/org/slf4j/slf4j-log4j12/1.7.7.redhat-2/slf4j-log4j12-1.7.7.redhat-2.jar Downloading: http://download.eng.bos.redhat.com/brewroot/repos/jb-eap-6.4-rhel-6-build/latest/maven/org/slf4j/slf4j-log4j12/1.7.7.redhat-2/slf4j-log4j12-1.7.7.redhat-2.jar Downloading: https://repository.jboss.org/nexus/content/groups/ea/org/slf4j/slf4j-log4j12/1.7.7.redhat-2/slf4j-log4j12-1.7.7.redhat-2.jar ``` Jon, could you pls have a look?	1	1	https://issues.jboss.org/browse/ENTESB-5675
11675	ENTESB-5567	Are there any changes needed in Fuse 6.3 or can this issue be closed and a new one created in OSFUSE?	0	0	https://issues.jboss.org/browse/ENTESB-5567
11676	ENTESB-5553	Fuse 6.3 should be updated to use the Red Hat build of winsw. [~pgallagh] Can you add link to the Red Hat build?	1	1	https://issues.jboss.org/browse/ENTESB-5553
11677	ENTESB-5547	Hi, I think this could be implemented by a user app bundle. They just need to use the Thread.html.setDefaultUncaughtExceptionHandler [1] to setup a JVM wide exception handler. Then in the exception handler they can use the Throwable.getStackTrace() and FrameworkUtil.getBundle() [2] to find out which bundle threw the exception and handle the exception accordingly. Note that "which bundle your in" can mean different things. Is it the bundle that threw the exception, or is the bundle that started the thread you were running in? If you want to get fancy and make it pluggable, user bundles could register exception handlers in the OSGi registry. The default handler that was set would just need to find all implementation in the registry and call the first one that was registered in the OSGi registry by the bundle associated with the exception. [1]: https://docs.oracle.com/javase/7/docs/api/java/lang/Thread.html#setDefaultUncaughtExceptionHandler(java.lang.Thread.UncaughtExceptionHandler) [2]: https://osgi.org/javadoc/r4v43/core/org/osgi/framework/FrameworkUtil.html#getBundle(java.lang.Class)	0	0	https://issues.jboss.org/browse/ENTESB-5547
11678	ENTESB-5514	Hi [~rick_wagner], is this what you want? guessing you used # while ? should be used for dynamic query. http://camel.apache.org/sql-component.html {quote} CamelSqlQuery Camel 2.8: Query to execute. This query takes precedence over the query specified in the endpoint URI. Note that query parameters in the header are represented by a ? instead of a # symbol {quote}	1	1	https://issues.jboss.org/browse/ENTESB-5514
11679	ENTESB-5471	Helll [~rkieley], could you please add commit link if it's fixed?	1	1	https://issues.jboss.org/browse/ENTESB-5471
11680	ENTESB-5471	Hi [~kearls], I just checked with 6.3 build 237 and it's still 813 kB for {{java.svg}}. Can you please provide link to commit fix to check where we actually have this fixed?	0	0	https://issues.jboss.org/browse/ENTESB-5471
11681	ENTESB-5452	[~gnodet] Can you comment on whether this is feasible?	1	0	https://issues.jboss.org/browse/ENTESB-5452
11682	ENTESB-5430	Pax web change, [~gnodet] can you comment on this?	1	1	https://issues.jboss.org/browse/ENTESB-5430
11683	ENTESB-5379	Would [Ehcache Replication|http://www.ehcache.org/generated/2.9.0/html/ehc-all/#page/Ehcache_Documentation_Set%2Fehcache_all.1.230.html%23] help here ? [~ffang], [~aileenc] as long term solution a jcache-base implementation in cxf may be a nice to have as switching cache implementation would be much easy (no change to the code base)	0	0	https://issues.jboss.org/browse/ENTESB-5379
11684	ENTESB-5379	[~ffang] fyi, ehcache 3.0 has support for JCache, would you like me to create an JCache OAuthDataProvider ?	0	0	https://issues.jboss.org/browse/ENTESB-5379
11685	ENTESB-5363	Would it make sense to eventually put it under camel-kubernetes upstream and then use it in the future "route policy" API ?	1	1	https://issues.jboss.org/browse/ENTESB-5363
11686	ENTESB-5329	Hi [~oscerd], what's the status of this issue? It's marked as resolved but I didn't see any activity on this issue.	0	0	https://issues.jboss.org/browse/ENTESB-5329
11687	ENTESB-5311	hi Tomohisa igarashi, I've tried it as you recommend (@Named("myCDIBeanRouteBuilder") on RouteBuilder) and it works for @Inject annotation above TestBean b; (630067 fuse build on eap) However it doesn't work for @BeanInject annotation. Since this issue relates to @BeanInject I'm not sure if I can consider this issue closed, could you give me advise?	1	1	https://issues.jboss.org/browse/ENTESB-5311
11688	ENTESB-5290	Should this be on list for fuse 7?	1	0	https://issues.jboss.org/browse/ENTESB-5290

#	Requirement	Questions	Answered	Successfully Answered	URL
11689	ENTESB-5162	[~virchete] - could you please try the following: 1. Run Fuse 6.2.1 on EAP 6.4.x 2. Run 'mvn clean package' in the provided transform-jaxb quickstart 3. cp target/switchyard-transform-jaxb.jar $FUSE_HOME/standalone/deployments/ 4. Start Eclipse/JBDS 5. Import transform-jaxb as maven project 6. In SwitchyardCXFSOAPTest, change ORDER_SERVICE_PORT to 8080 (optional: @Ignore tests testMTOMSimpleAttachment and testMTOMWithImage, as these are related to ENTESB-5161) 7. Run the class as JUnit test Note that the two tests testMTOMexpandSimpleAttachment/testMTOMexpandWithImage pass, but the response is not in the expected format nevertheless: * testMTOMexpandSimpleAttachment/testMTOMexpandWithImage: requests against xopExpand=true endpoint. working, but base64 inline attachment and no mtom response * testMTOMexpand_analyseResponse: sends special requests which forces the server to return an attachment and analyses the httpresponse. take a look at the unittest log.	1	1	https://issues.jboss.org/browse/ENTESB-5162
11690	ENTESB-5150	Verified with fuse 630045 on karaf, by QE. For fuse on EAP I have a question: Some packages were updated, but some of them (httpclient, httpcore, freemarker, jgroups) not. These are in "base" folder only, I'm not sure whether should be updated also. Could you give me advice?	1	1	https://issues.jboss.org/browse/ENTESB-5150
11691	ENTESB-5139	Comments from requesting party. March 30, 2016: {quote} I've seen Claus Ibsen's comments on ENTESB-5139. It is obviously a bit more work to make Camel run with Saxon 9.7. The referenced Camel upstream tickets (CAMEL-8830 and CAMEL-7891) are from 2014 and 2015 and still unresolved. Not sure if possible, but it would be really good if a fix with Saxon 9.7 upgrade would make it into Fuse 6.3. I'm a bit afraid that we have to wait very long until something moves forward if the fix gets delayed to much... Do you think this would be possible? {quote} April 3, 2016: {quote} If you could inform PM that we as a customer are heavily invested in XSLT/Saxon on Fuse so that PM might give this some priority, that would be very helpful. {quote}	0	0	https://issues.jboss.org/browse/ENTESB-5139
11692	ENTESB-5087	If you already have an example, could you please make this a PR? There is some test coverage [here|https://github.com/wildfly-extras/wildfly-camel/blob/master/itests/standalone/basic/src/main/java/org/wildfly/camel/test/quartz/QuartzIntegrationTest.java]	0	0	https://issues.jboss.org/browse/ENTESB-5087
11693	ENTESB-5050	[~gnodet1], What karaf related javadocs you would want packaged in the fuse distro? Could you give me a link to the mvn gavs to the javadoc artifacts?	0	0	https://issues.jboss.org/browse/ENTESB-5050
11694	ENTESB-5028	Is this something we can get into 6.3?	0	0	https://issues.jboss.org/browse/ENTESB-5028
11695	ENTESB-4979	[~rick_wagner], can we close this issue?	0	0	https://issues.jboss.org/browse/ENTESB-4979
11696	ENTESB-4974	[~dbokde] Can you comment on this?	0	0	https://issues.jboss.org/browse/ENTESB-4974
11697	ENTESB-4965	Can we do a prototype for this?	1	1	https://issues.jboss.org/browse/ENTESB-4965
11698	ENTESB-4930	These command bundles are available at start level 50 (standalone / non-fabric): {code} JBossFuse:admin@root> osgi:list -t 0 | grep -i commands | grep "(50)|30])" [14] [Active] [Created] [ ] [50] JBoss Fuse :: ESB :: Commands (6.3.0.redhat-067) [39] [Active] [Created] [ ] [30] Apache Karaf :: Shell :: Various Commands (2.4.0.redhat-630067) [40] [Active] [Created] [ ] [30] Apache Karaf :: Shell :: OSGi Commands (2.4.0.redhat-630067) [41] [Active] [Created] [ ] [30] Apache Karaf :: Shell :: Log Commands (2.4.0.redhat-630067) [42] [Active] [Created] [ ] [30] Apache Karaf :: Shell :: PackageAdmin Commands (2.4.0.redhat-630067) [43] [Active] [Created] [ ] [30] Apache Karaf :: Shell :: Development Commands (2.4.0.redhat-630067) [44] [Active] [Created] [ ] [30] Apache Karaf :: Shell :: ConfigAdmin Commands (2.4.0.redhat-630067) [229] [Active] [ ] [ ] [50] camel-commands-core (2.17.0.redhat-630067) [230] [Active] [Created] [ ] [50] camel-karaf-commands (2.17.0.redhat-630067) [271] [Active] [Created] [ ] [30] Apache Karaf :: Shell :: Web Commands (2.4.0.redhat-630067) {code} These start at level 80 - *unavailable* at "hot standby" {code}JBossFuse:admin@root> osgi:list -t 0 | grep -i commands | grep "(80])" [178] [Installed] [ ] [ ] [80] Apache Karaf :: Shell :: SCR Commands (2.4.0.redhat-630067) [181] [Installed] [ ] [ ] [80] Fabric8 :: Boot Commands (1.2.0.redhat-630067) [184] [Installed] [ ] [ ] [80] Fabric8 :: Karaf Commands (1.2.0.redhat-630067) [190] [Installed] [ ] [ ] [80] Fabric8 :: Patch :: Commands (1.2.0.redhat-630067) [278] [Installed] [ ] [ ] [80] JBoss Fuse :: Support :: Commands (1.2.0.redhat-630067) {code} Is this desired state, or do we want to support patching of "hot standby" container? (assuming we continue to suggest hot standby at start-level 50)	0	0	https://issues.jboss.org/browse/ENTESB-4930
11699	ENTESB-4926	[~oscerd] you're right. Closing issue. Verified in fuse 6.3 build 039	0	0	https://issues.jboss.org/browse/ENTESB-4926
11700	ENTESB-4922	What changed that it's fixed [~oscerd]? Confirm it's working in fuse 6.3 buidl 026.	0	0	https://issues.jboss.org/browse/ENTESB-4922
11701	ENTESB-4890	[~ffang] Can you comment on this?	1	1	https://issues.jboss.org/browse/ENTESB-4890
11702	ENTESB-4890	Can you get this into 6.3?	1	1	https://issues.jboss.org/browse/ENTESB-4890
11703	ENTESB-4874	How much work is this to switch over to pubic APIs? It would be great to get this done for Fuse 6.3 release if possible.	0	0	https://issues.jboss.org/browse/ENTESB-4874
11704	ENTESB-4865	Hi [~lburgazzoli], on solaris 10 I now get this error: {code} ./karaf-service.sh -k /Downloads/jboss-fuse-6.3.0.redhat-026 -u root -g root ./karaf-service.sh: line 67: syntax error near unexpected token `"$OPTARG"' ./karaf-service.sh: line 67: ` e) KARAF_ENV+=("$OPTARG") ;;' {code} it is on solaris 10 9/10 s10x_u9wos_14a X86 Could you please check?	1	1	https://issues.jboss.org/browse/ENTESB-4865
11705	ENTESB-4865	Can you take it ?	1	1	https://issues.jboss.org/browse/ENTESB-4865
11706	ENTESB-4853	Would it be possible to provide some additional information about the scenario the user requires. Is the 'http' endpoint that they wish to configure as 'https' the elasticsearch proxy endpoint provided by RTGov? Or is it the Elasticsearch server's own endpoint? (i.e. are they directly accessing the Elasticsearch server, or going via the endpoint provided by RTGov which includes authorization - http://.../rtgov-ui/elasticsearch) What is their Elasticsearch environment - do they startup separate nodes/servers, or use the embedded node? Would be good to include the details configured in the rtgov-elasticsearch properties in standalone-full.xml to help replicate their environment.	1	0	https://issues.jboss.org/browse/ENTESB-4853
11707	ENTESB-4853	[~rick_wagner] Any update on this? Are we considering doing this enhancement?	0	0	https://issues.jboss.org/browse/ENTESB-4853
11708	ENTESB-4852	Hi Hiram, the feature is missing in Fuse distribution by default: {code} JBossFuse:karaf@root> features:install qpid-jms-client Error executing command: No feature named 'qpid-jms-client' with version '0.0.0' available {code} I can see the feature inside of fabric8 environment, but container does not resolve properly: {code} JBossFuse:karaf@root> profile-edit --feature qpid-jms-client jboss-fuse-full Adding feature:qpid-jms-client to profile:default version:1.0 JBossFuse:karaf@root> container-info ... Provision Status: error - org.osgi.service.resolver.ResolutionException: Unable to resolve root: missing requirement [root] osgi.identity; osgi.identity=qpid-jms-client; type=karaf.feature; filter:="(&(osgi.identity=qpid-jms-client)(type=karaf.feature))" Provision Error: org.osgi.service.resolver.ResolutionException: Unable to resolve root: missing requirement [root] osgi.identity; osgi.identity=qpid-jms-client; type=karaf.feature; filter:="(&(osgi.identity=qpid-jms-client)(type=karaf.feature))" ... {code} Are there any plans to include the feature into JBoss Fuse without fabric8?	0	0	https://issues.jboss.org/browse/ENTESB-4852

#	Requirement	Questions	Answered	Successfully Answered	URL
11709	ENTESB-4806	Is there any more to be done for this or can it be marked as resolved?	1	1	https://issues.jboss.org/browse/ENTESB-4806
11710	ENTESB-4708	* (*r) jclouds-1.9.1 is affected by https://issues.apache.org/jira/browse/JCLOUDS-534, which forces guava at version 16.x	1	1	https://issues.jboss.org/browse/ENTESB-4708
11711	ENTESB-4704	[~virchete], this commit was backported to {{1.1.x}} branch. {{fuse-integration-bom}} does not contains {{kie}} artifacts version in version {{1.1.0.redhat-028}}. Is this somewhere documented? User has to add {{kie-bom}} manually now.	0	0	https://issues.jboss.org/browse/ENTESB-4704
11712	ENTESB-4704	[~virchete], could you please provide *ordered* list of BOMs which user should use to overcome this issue?	0	0	https://issues.jboss.org/browse/ENTESB-4704
11713	ENTESB-4690	[~janstey]cool. Do you think it is patchable?	1	1	https://issues.jboss.org/browse/ENTESB-4690
11714	ENTESB-4689	Docs Comment: Where can we find out more information about this?	1	0	https://issues.jboss.org/browse/ENTESB-4689
11715	ENTESB-4689	Engineering will have details here. To be clear, the context is wrapping camel routes in Switchyard, not removing Camel. What has to be changed, why would someone want to do this (e.g. what benefits), etc..	0	0	https://issues.jboss.org/browse/ENTESB-4689
11716	ENTESB-4689	QE Ambiguity Review Comments: * Outstanding QE Requirement Questions: ** None * Outstanding QE Requirement Implementation Questions: ** What are the mechanics for this migration? ** Will users perform the migration in the SwitchYard editor in a visual/manual manner? ** How will errors be reported to the user - will they be highlighted visually? ** Will the users migrate routes one-by-one? Or can they perform a bulk migration?	1	0	https://issues.jboss.org/browse/ENTESB-4689
11717	ENTESB-4688	QE Ambiguity Review Comments: * Is it our intention to (1) debug the configuration/implementation issues in 6.1 and (2) provide a quickstart to demonstrate it? What is the deliverable for 6.1?	0	0	https://issues.jboss.org/browse/ENTESB-4688
11718	ENTESB-4687	To ensure the javadocs are packaged in a form that will be consistent across FSW, can you provide some guidance/example of how you would prefer it to be done. RTGov uses the maven javadoc plugin: https://github.com/Governance/rtgov/blob/master/pom.xml#L307-L348 Is this what you are looking for?	0	0	https://issues.jboss.org/browse/ENTESB-4687
11719	ENTESB-4687	Reminder - do we want to ship only the public API javadocs? Or everything? (I assume just the public API.)	0	0	https://issues.jboss.org/browse/ENTESB-4687
11720	ENTESB-4685	Reminder - do we want to ship only the public API javadocs? Or everything? (I assume just the public API.)	0	0	https://issues.jboss.org/browse/ENTESB-4685
11721	ENTESB-4685	Is this still relevant for 6.2.1? IMHO, all artifact will be in maven repo with associated javadoc.	1	1	https://issues.jboss.org/browse/ENTESB-4685
11722	ENTESB-4680	This will require a Microsoft Dynamics Server to test against. Do we have this in Red Hat or any one with expertise to set it up? I would prioritise the components that are available as a service rather than those that need an installation.	0	0	https://issues.jboss.org/browse/ENTESB-4680
11723	ENTESB-4679	This will require a server to test against. Do we have this in Red Hat or any expertise to set it up?	1	1	https://issues.jboss.org/browse/ENTESB-4679
11724	ENTESB-4678	This seems like an issue targeting prod and qe. Will ENG have access to pSeries machine to run CI against?	0	0	https://issues.jboss.org/browse/ENTESB-4678
11725	ENTESB-4676	[~mmelko] Hi Matej. Would you be able to comment on this idea?	1	1	https://issues.jboss.org/browse/ENTESB-4676
11726	ENTESB-4676	I have another customer using 6.2 pre-patch no internet, and it doesn't work at all. Should I use this or open another jira?	0	0	https://issues.jboss.org/browse/ENTESB-4676
11727	ENTESB-4674	Chandra, would you double check 6.2 please? It's not likely but we did make a number of changes here. Otherwise, we will revisit for 6.3 or if we do't do 6.3, 7.0	0	0	https://issues.jboss.org/browse/ENTESB-4674
11728	ENTESB-4664	Does [~dbokde]'s component tool kit help with SOAP implementation?	1	1	https://issues.jboss.org/browse/ENTESB-4664
11729	ENTESB-4664	What is missing from REST API and do PayPal plan to complete out the functionality missing in REST API? If so we may be better of sticking with that, implement the functionality it supports today and add new features when PayPal support them in REST API. Reimplementing using Classic API seems like a backwards step.	0	0	https://issues.jboss.org/browse/ENTESB-4664
11730	ENTESB-4664	|*SDK*|OpenSource|https://github.com/braintree/braintree_java| |*License*|MIT|https://github.com/braintree/braintree_java/blob/master/LICENSE| |*OSGi ready*|No|https://github.com/braintree/braintree_java/pull/30| |*Dependencies*|None|altrough, common codec sources have been included in [source tree|https://github.com/braintree/braintree_java/tree/master/src/main/java/com/braintreegateway/org/apache/commons/codec]| |*Support*|yes|support@braintreepayments.com| |*Sandbox*|yes|https://sandbox.braintreegateway.com/| |*Component*|yes|https://github.com/lburgazzoli/camel-braintree| Pro: * single SDK for different servics (paypal, apple pay, etc) * simple API * most of the APIs are common among payment methods Cons: * client/server [interaction|https://developers.braintreepayments.com/start/hello-server/java] * sandbox credential are linked to a person so it makes testing harder (i.e. each contributor need to create and configure a sandbox account, run test in parallel may leads to failures) I've been able to wrap some API via Camel API component framework and so far the only concern I have is about thread safety of the API, i.e. can I share a Gateway or *Request instance among endpoints ?	1	1	https://issues.jboss.org/browse/ENTESB-4664
11731	ENTESB-4661	Freeman would you know if this this something we can do?	1	1	https://issues.jboss.org/browse/ENTESB-4661
11732	ENTESB-4658	There are some other options available in Fuse 6.1.0 like following , But those does not work in Fuse 6.2 (Should that be supported for Fuse 6.2) ? "import-resources.karaf" {code} # Imports resources (ex. xml or properties) into a Fabric Profile # ex. # import-resource "TestProfile" "1.0" "something.xml" "/home/jsensharma/Downloads" # profile-display --display-resources TestProfile import-resource = { _profile = $1 _version = $2 _logical_file_path = $3 _physical_file_path = $4 _file_content = ((cat file://$_physical_file_path) | tac) _sref = $.context getServiceReference "io.fabric8.api.FabricService" _srv = $.context getService $_sref _dstore = $_srv getdatastore $_dstore setFileConfiguration $_version $_profile $_logical_file_path ($_file_content getbytes) } {code} *How to run above script ?* {code} JBossFuse:karaf@root> profile-create --parents default TestProfile JBossFuse:karaf@root> source file:///home/jsensharma/Downloads/import-resources.karaf JBossFuse:karaf@root> import-resource "TestProfile" "1.0" "something.xml" "/home/jsensharma/Downloads" {code} - The above command will upload the "something.xml" file present inside the /home/jsensharma/Downloads" location to a profile names as "TestProfile" (version 1.0)	1	0	https://issues.jboss.org/browse/ENTESB-4658
11733	ENTESB-4656	Should we switch to log4j2 for 7.0?	1	1	https://issues.jboss.org/browse/ENTESB-4656
11734	ENTESB-4634	Have any customers requested this?	1	0	https://issues.jboss.org/browse/ENTESB-4634
11735	ENTESB-4623	Is this still relevant for Fuse 7?	0	0	https://issues.jboss.org/browse/ENTESB-4623
11736	ENTESB-4599	I don't know [~kearls] what dialog are you talking about but I can't see nothing wrong under Camel tab. Except for the page doesn't refresh (update) if I add or remove a new camel context. Hasn't that been the thing to fix? (tried on 6.3 067)	0	0	https://issues.jboss.org/browse/ENTESB-4599

#	Requirement	Questions	Answered	Successfully Answered	URL
11737	ENTESB-4598	Hi [~sonicaaaa], the logs from the cleaning manager are printed out to stdout when using interactive shell: {code} JBossFuse:karaf@root> Removing file: Apr 8, 2016 9:20:08 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/log.206f Removing file: Apr 8, 2016 9:16:38 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/log.1b1a Removing file: Apr 8, 2016 9:09:19 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/log.d36 Removing file: Apr 8, 2016 9:13:27 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/log.1434 Removing file: Apr 8, 2016 9:05:01 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/log.1 Removing file: Apr 8, 2016 9:22:46 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/log.2646 Removing file: Apr 8, 2016 9:20:08 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/snapshot.2644 Removing file: Apr 8, 2016 9:16:38 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/snapshot.206d Removing file: Apr 8, 2016 9:13:27 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/snapshot.1b18 Removing file: Apr 8, 2016 9:09:19 AM /tmp/jboss-fuse-6.3.0.redhat-045/data/zookeeper/0000/version-2/snapshot.1432 {code} I believe they should be logged into the log (and probably debug level only?) Is there anything that could be done about this? If not, then resolve the issue again and I will close it as it is working as expected	1	1	https://issues.jboss.org/browse/ENTESB-4598
11738	ENTESB-4597	Hi [~chirino], the examples/samples for: - *cxf* are not present - *camel* are present Is this intended? There are some examples for camel, but not for cxf. If yes, then I think I can close this issue. Thanks Hiram.	1	1	https://issues.jboss.org/browse/ENTESB-4597
11739	ENTESB-4597	Hi [~vkasala], where did you find the examples for camel?	1	1	https://issues.jboss.org/browse/ENTESB-4597
11740	ENTESB-4597	Hi [~chirino], shouldn't be this issue moved to next Sprint? How can I verify it, if it is not in ER1 or other build? Thanks.	0	0	https://issues.jboss.org/browse/ENTESB-4597
11741	ENTESB-4593	[~tcunning], we should create an issue in Camel project to fix the following method: https://github.com/jboss-fuse/camel/blob/camel-2.16.1.redhat-630015/components/camel-jaxb/src/main/java/org/apache/camel/converter/jaxb/FallbackTypeConverter.java#L149 Bean class without annotation {{@XmlRootElement}} is not handled as JAXB type. I have a simple switchyard example where I use HTTP binding and JAXB transformation. There is missing transformer: {{java:java.io.StringReader => java:<my-custom-type>}}. So I have to use annotation {{@XmlRootElement}}. Otherwise I have to add extra transformer... What do you think?	0	0	https://issues.jboss.org/browse/ENTESB-4593
11742	ENTESB-4592	Prod can change the DDL scripts, so we could increase the size of the clob columns, but to what size? The original ticket suggested an increase to the max supported value, which is 2,147,483,647 bytes. What are the performance/storage implications of this change? How will this affect any other DB2 customers? Furthermore, increasing the clob size in 6.3 will break compatability for DB2 users with existing data. Do we also need to provide an upgrade script for the 6.3 release?	1	1	https://issues.jboss.org/browse/ENTESB-4592
11743	ENTESB-4588	I'm attaching my reproducer with question, [~tcunning] can you please help me figure out what is the real issue here? what the customer really expects to get? empty message? empty soap body? Please see my attachment. 1) mvn clean test and please observe lines 13-16 in src/main/resources/wsdl/PingService.wsdl Caused by: org.switchyard.extensions.wsdl.WSDLReaderException: SWITCHYARD012801: Service operations on a WSDL interface must have exactly one parameter. at org.switchyard.extensions.wsdl.WSDLReader.getParts(WSDLReader.java:388) at org.switchyard.extensions.wsdl.WSDLReader.readWSDL(WSDLReader.java:103) at org.switchyard.extensions.wsdl.WSDLService.fromWSDL(WSDLService.java:100) at org.switchyard.extensions.wsdl.WSDLService.fromWSDL(WSDLService.java:83) at org.switchyard.deploy.internal.Deployment.loadServiceInterface(Deployment.java:418) ... 10 more 2) if we uncomment src/main/resources/wsdl/PingService.wsdl on line 13 - 16, tests are passing	1	1	https://issues.jboss.org/browse/ENTESB-4588
11744	ENTESB-4584	Hello Red Hat engineers, We have a support case associated with this request, the customer is offering to provide an implementation if we suggest the design. Can you please consider, offer some suggestions about high-level design? (i.e. the proper package for the thread pool, the classes that will need to interact with it.) Support will maintain a dialogue with the customer, perhaps their implementation can later be shared back. Thanks, Rick	1	0	https://issues.jboss.org/browse/ENTESB-4584
11745	ENTESB-4582	The new component has been merged in master branch as of https://github.com/apache/camel/commit/0421c24dfcf992f3296ed746469771e3800200e3 so that it'll be available in the upcoming 2.17.0 upstream Camel release.	1	1	https://issues.jboss.org/browse/ENTESB-4582
11746	ENTESB-4558	Is this already in the 7.0 kits built from Jenkins pipeline?	0	0	https://issues.jboss.org/browse/ENTESB-4558
11747	ENTESB-4551	That's definitely possible. How do you envision the notification working if the container detects that a new version is available? In all scenarios, I think it's a situation where we compare the version in use (determined via Maven GAV) to the latest available version (via MRRC query or other catalog).	0	0	https://issues.jboss.org/browse/ENTESB-4551
11748	ENTESB-4534	Hi [~igarashitm], I'm not sure why you need to change leveldb itself. Can't you use it's existing APIs to to implement the filter policy?	1	1	https://issues.jboss.org/browse/ENTESB-4534
11749	ENTESB-4459	[~emil.cervenan], there's ENTESB-2844 where we tried to describe most of http-proxy related configuration. Is this invalid now?	0	0	https://issues.jboss.org/browse/ENTESB-4459
11750	ENTESB-4459	[~emil.cervenan], you describe two issues: * http proxy * maven proxy Assuming that fabric-agent (that uses information from profile-edit) first contacts fabric's main maven proxy (started with fabric-maven-proxy bundle), it should be enough to correctly configure maven resolver inside fabric-maven-proxy. Try this: h3. http proxy either create file {{my-settings.xml}}: {code:xml} ... <proxies> <proxy> <id>my-proxy</id> <active>true</active> <protocol>http</protocol> <username>fuse</username> <password>fuse</password> <host>localhost</host> <port>3128</port> <nonProxyHosts /> </proxy> </proxies> ... {code} and tell fabric-maven-proxy about it: {noformat} profile-edit --pid io.fabric8.maven/io.fabric8.maven.settings=/path/to/maven-settings.xml default {noformat} or do it _inline_: {noformat} profile-edit --pid io.fabric8.maven/io.fabric8.maven.proxies=http:host=foo,port=8080,user=xxx,pass=xxx;https:host=bar,port=9090 default {noformat} h3. maven proxy Append your special maven proxy with this command: {noformat} profile-edit --pid io.fabric8.maven/io.fabric8.maven.repositories='http://fusewin.tpb.lab.eng.brq.redhat.com:8081/nexus/content/repositories/fuse-qe-repo@id=fuse-qa' --append default {noformat} The idea above is to configure maven/http proxy in one place (fabric's internal maven proxy, that should be contacted first by each fabric-agent), instead of using {{profile-edit \--pid io.fabric8.agent/org.ops4j.pax.url.mvn.XXX=yyy}} as a way to configure download manager for fabric-agent. [~sonicaaaa], [~janstey], [~chirino] would you like to add/change something to/in my explanation?	0	0	https://issues.jboss.org/browse/ENTESB-4459

#	Requirement	Questions	Answered	Successfully Answered	URL
11751	ENTESB-4408	Can we disable this for 6.2.1 and direct people to command line?	0	0	https://issues.jboss.org/browse/ENTESB-4408
11752	ENTESB-4340	Was it already shipped in Fuse 6.2 release?	0	0	https://issues.jboss.org/browse/ENTESB-4340
11753	ENTESB-4296	[~grgrzybek] The notes say that the patch management step is only necessary for those on 6.2.0GA, does this mean any version of 6.2.1 should be able to be patched to >6.2.1ER4?	1	1	https://issues.jboss.org/browse/ENTESB-4296
11754	ENTESB-4296	Yes [~jonnychristison], everything looks ok. [~janstey] what can I do more about this issue?	1	1	https://issues.jboss.org/browse/ENTESB-4296
11755	ENTESB-4296	[~jonnychristison] so those "baselines" for 133 should be made out of jboss-fuse-full and jboss-a-mq 6.2.0.redhat-133 versions in brew (was this sentence correct?). Then we have to upload "your" (brew) baselines to Nexus, so we can reference them in jboss-fuse build. When "productization tags" are created, you'll build Fuse/AMQ your way and you'll get "productized" {{patch-management-NEWVERSION-for-fuse-6.2.0.redhat-133.zip}}.	1	1	https://issues.jboss.org/browse/ENTESB-4296
11756	ENTESB-4268	[~jludvice] How do I reproduce the rest-binding dual dependency chain here? I'm trying with Perfectus build 069 and don't see one.	1	1	https://issues.jboss.org/browse/ENTESB-4268
11757	ENTESB-4267	Why is camel-sap feature gone?	1	1	https://issues.jboss.org/browse/ENTESB-4267
11758	ENTESB-4267	Test was ran again with build #69 and feature camel-sap is back. Can we close this issue?	0	0	https://issues.jboss.org/browse/ENTESB-4267
11759	ENTESB-4236	It works. But the workaround with classloader is not mentioned neither in [community documentation|http://camel.apache.org/groovy.html] nor in product documentation. Is it ok? Do you think it is obvious?	0	0	https://issues.jboss.org/browse/ENTESB-4236
11760	ENTESB-4221	It looks like this is working {{etc/org.ops4j.pax.web.cfg}} {code} org.ops4j.pax.web.config.file=etc/jetty.xml # enable http org.osgi.service.http.port=8181 org.ops4j.pax.web.session.cookie.httpOnly=true org.osgi.service.http.enabled=true # enable https org.osgi.service.http.secure.enabled=true org.osgi.service.http.port.secure=8441 # configure keystore org.ops4j.pax.web.ssl.keystore=${karaf.etc}/broker.ks org.ops4j.pax.web.ssl.keystore.type=jks org.ops4j.pax.web.ssl.password=awesomeStorePass org.ops4j.pax.web.ssl.keypassword=awesomeKeyPass # PaxWeb docs https://ops4j1.jira.com/wiki/display/paxweb/SSL+Configuratione {code} It was working also when this part of {{jetty.xml}} was deleted. Can we delete this hardcoded port ? Do we need to define those properties? {code:xml} <New id="httpConfig" class="org.eclipse.jetty.server.HttpConfiguration"> <Set name="secureScheme">https</Set> <Set name="securePort"> <Property name="jetty.secure.port" default="8443" /> </Set> <Set name="outputBufferSize">32768</Set> <Set name="requestHeaderSize">8192</Set> <Set name="responseHeaderSize">8192</Set> <Set name="sendServerVersion">true</Set> <Set name="sendDateHeader">false</Set> <Set name="headerCacheSize">512</Set> </New> {code} {{broker.ks}} created for testing purposes with this command {code} keytool -genkeypair \ -dname "CN=name of your container, O=name of your organization, L=city name, ST=your state or province, C=US" \ -validity 365 \ -alias cert-alias \ -keypass awesomeKeyPass \ -storepass awesomeStorePass \ -keystore broker.ks \ -keyalg RSA {code}	0	0	https://issues.jboss.org/browse/ENTESB-4221
11761	ENTESB-4147	Verified by QA: {code} grep -R "jolokia-version>1" fabric8/pom.xml: <jolokia-version>1.3.2.redhat-1</jolokia-version> hawtio/pom.xml: <jolokia-version>1.3.2.redhat-1</jolokia-version> {code}	0	0	https://issues.jboss.org/browse/ENTESB-4147
11762	ENTESB-4085	Why do you think that this would "allow exploitation of any security vulnerabilities"? The trigger for camel module wiring is currently the existence of the Camel subsystem in the configuration. Even if we followed your advice, it is still possible to do an explicit module load of any module that is in the hierarchy. From that module you can then obtain the classloader and do what you do above just the same.	0	0	https://issues.jboss.org/browse/ENTESB-4085
11763	ENTESB-4058	Is this one resolved then?	0	0	https://issues.jboss.org/browse/ENTESB-4058
11764	ENTESB-4045	Can we make this change for 6.2.1 so it doesn't turn into a Fuse 6.2 / Hibernate last minute issue..	1	1	https://issues.jboss.org/browse/ENTESB-4045
11765	ENTESB-4045	Did we ship 1.0 / 1.1 and 1.2 in previous releases?	1	1	https://issues.jboss.org/browse/ENTESB-4045
11766	ENTESB-4045	Is there still a benefit to aligning to versions in EAP? That way if we need a fix in the core of weld (not OSGi bundle) we have a team in house who can provide that.	1	1	https://issues.jboss.org/browse/ENTESB-4045
11767	ENTESB-4043	I thought this analyser would help us identify artefacts that had already been built from source by another product that we might want to align with, Is that the intention of the analyser?	1	0	https://issues.jboss.org/browse/ENTESB-4043
11768	ENTESB-4043	Hi, what are your expectations from this tool? What data would you like to get from it? This tool is in development state and the new features will be added in the future. This is a first version, which should discover, what the users wants. Let me know what do you want from the tool and some of next versions can contains more useful features for you.	1	1	https://issues.jboss.org/browse/ENTESB-4043
11769	ENTESB-4043	What is the goal of this tool? My expectation was that it would give us a list off all 3rd party dependencies fro Fuse and help identify those that have already been built by other teams so that we can aim to align. Is that something you plan to support in future versions?	1	0	https://issues.jboss.org/browse/ENTESB-4043
11770	ENTESB-4043	[~jbartece] We're having a hard time interpreting the report we got https://10.8.63.16/view/Perfectus/job/ncl-test/5/console I see comments like: Requested GA was not found org.apache.servicemix.specs:org.apache.servicemix.specs.activator:2.5.0 None no list 4 dependencies not built 2 available versions None list: io.fabric8.mq:mq-client:1.2.0.redhat-621048 Can you provide some guidance on what we should be looking for in the report?	1	0	https://issues.jboss.org/browse/ENTESB-4043
11771	ENTESB-4043	[~aileenc] Sorry, we still miss the user's documentation - should come in 2 weeks. You should find some helpful information, when running the script without parameters (it'll print the help). Note: The Dependency Analyser version 0.4 was release 2 weeks ago. Download new scripts, please: https://github.com/project-ncl/dependency-analysis/blob/master/cli/da-cli-0.4.2.zip?raw=true Explanation of the output: Requested GA was not found - bad message (bug for it is reported), It typically means, that the GAV hasn't been internally built (-redhat build wasn't found). org.apache.servicemix.specs:org.apache.servicemix.specs.activator:2.5.0 None no list 4 dependencies not built 2 available versions None == <Best Matched Red Hat Version> - if the artifact was built, then there will be RH version of the artifact no list = = <In black/white list?> 4 dependencies not built == <Number of not built dependencies> 2 available versions == <Available Versions> - internal builds of this artifact (in different versions) None list: io.fabric8.mq:mq-client:1.2.0.redhat-621048 - It means, that this artifact is not in blacklist nor whitelist	1	1	https://issues.jboss.org/browse/ENTESB-4043
11772	ENTESB-4043	I think the report could benefit from some reformatting to make it more readable When is the 0.5 release scheduled?	1	0	https://issues.jboss.org/browse/ENTESB-4043
11773	ENTESB-4043	It should be released in 5 or 6 weeks. The output was reworked once, but we can improve it again. Do you have any proposal of more readable output for you?	1	1	https://issues.jboss.org/browse/ENTESB-4043
11774	ENTESB-4043	[~kearls] There's a new release out. Could you try another run?	0	0	https://issues.jboss.org/browse/ENTESB-4043

#	Requirement	Questions	Answered	Successfully Answered	URL
11775	ENTESB-3924	[~janstey] Can you take care of this Jon?	1	1	https://issues.jboss.org/browse/ENTESB-3924
11776	ENTESB-3924	[~jpav] Do you have the build now, can I mark this as resolved?	1	1	https://issues.jboss.org/browse/ENTESB-3924
11777	ENTESB-3867	Can you log an issue against EAP and add the link here?	0	0	https://issues.jboss.org/browse/ENTESB-3867
11778	ENTESB-3735	Can we fix the version ranges? How likely is it that customers will want to use these together?	0	0	https://issues.jboss.org/browse/ENTESB-3735
11779	ENTESB-3648	[~virchete] bq. when the Switchyard Library Integration Test is executed could you be more specific? what should I do to run those tests?	1	1	https://issues.jboss.org/browse/ENTESB-3648
11780	ENTESB-3648	Is this still an issue? we're moving tests to Arquillian.	0	0	https://issues.jboss.org/browse/ENTESB-3648
11781	ENTESB-3648	[~virchete] is it still related to https://github.com/jboss-integration/fuse-bxms-integ/blob/master/release/karaf/itests/switchyard/src/test/java/org/switchyard/karaf/test/quickstarts/demo/HelpdeskDemoQuickstartTest.java#L26 ? Or is there another case?	1	1	https://issues.jboss.org/browse/ENTESB-3648
11782	ENTESB-3615	[~pdrozd], What is the "password policy"?	1	0	https://issues.jboss.org/browse/ENTESB-3615
11783	ENTESB-3539	[~ffang] What's your view on this? Should we change the default as Gary suggests?	0	0	https://issues.jboss.org/browse/ENTESB-3539
11784	ENTESB-3380	[~hkonsek], does it work for you? (sorry for ask+assign action :)).	1	1	https://issues.jboss.org/browse/ENTESB-3380
11785	ENTESB-3355	Only a subset of Camel components are supported on EAP. Is some one asking for this one to be supported?	0	0	https://issues.jboss.org/browse/ENTESB-3355
11786	ENTESB-3327	Currently the path is {{org/fuse/integration}}, but groupId is {{org.jboss.integration.fuse}}. So, shouldn't the path reflect groupId?	0	0	https://issues.jboss.org/browse/ENTESB-3327
11787	ENTESB-3220	[~gnodet], do you know how to fix this chain?	0	0	https://issues.jboss.org/browse/ENTESB-3220
11788	ENTESB-3185	Will I just close this?	1	0	https://issues.jboss.org/browse/ENTESB-3185
11789	ENTESB-3133	Is there any easy way how to verify that the supported standards are accurate? For example, the overview mentions that WS-Security Core standard in 6.2 is 1.3, but it seems that the latest is 1.1 https://www.oasis-open.org/standards#wssv1.1. Another example is WS-Reliable Messaging (1.3 mentioned, 1.2 latest). There is much more of such examples.	1	0	https://issues.jboss.org/browse/ENTESB-3133
11790	ENTESB-3114	Tomo - can we do the same thing here that we did on the other activemq-noweb-broker quickstarts?	1	1	https://issues.jboss.org/browse/ENTESB-3114
11791	ENTESB-3086	Tom, could you perhaps add your SwitchYard bits?	1	1	https://issues.jboss.org/browse/ENTESB-3086
11792	ENTESB-3086	I think this is already done, can jira be closed?	1	1	https://issues.jboss.org/browse/ENTESB-3086
11793	ENTESB-3044	Hi [~jdyson]. Is this !ganttExchange.png|thumbnail! right appearance of the Gantt chart? You haven't provided much detail what has changed. I'm not sure but the Camel page seems to be without changes and the chart still overlaps text behind it. I don't know how to verify this fix.	0	0	https://issues.jboss.org/browse/ENTESB-3044
11794	ENTESB-3013	Should this be moved over to ENTESB as it affects runtime?	1	1	https://issues.jboss.org/browse/ENTESB-3013
11795	ENTESB-3013	Did u also build the json file needed?	0	0	https://issues.jboss.org/browse/ENTESB-3013
11796	ENTESB-2943	[~virchete] [~janstey] Can you work on getting this in for next week's handover to productisation?	1	1	https://issues.jboss.org/browse/ENTESB-2943
11797	ENTESB-2928	While back, in fuse 6.2 build 66 there was issue with package javax.persistence in features camel-bam and camel-jpa (not affecting current builds) {code} Feature "camel-jpa" (2.14.0.redhat-620071) from camel-2.14.0.redhat-620071 =>←[31mError executing command: Could not start bundle mvn:org.apache.servicemix.bundles/org.apache.servicemix.bundles.spring-orm/3.2.12.RELEASE_1 in feature(s) spring-orm-3.2.12.RELEASE_1: Uses constraint violation. Unable to resolve bundle revision org.apache.servicemix.bundles.spring-orm [712.0] because it is exposed to package 'javax.persistence' from bundle revisions org.apache.geronimo.specs.geronimo-jpa_2.0_spec [321.0] and org.hibernate.javax.persistence.hibernate-jpa-2.1-api [333.0] via two dependency chains. Chain 1: org.apache.servicemix.bundles.spring-orm [712.0] import: (&(osgi.wiring.package=javax.persistence)(version>=1.0.0)(!(version>=3.0.0))) | export: osgi.wiring.package=javax.persistence org.apache.geronimo.specs.geronimo-jpa_2.0_spec [321.0] Chain 2: org.apache.servicemix.bundles.spring-orm [712.0] import: (&(osgi.wiring.package=org.hibernate.annotations)(version>=3.3.1)(!(version>=5.0.0))) | export: osgi.wiring.package=org.hibernate.annotations; uses:=javax.persistence org.hibernate.core [343.0] import: (&(osgi.wiring.package=javax.persistence)(version>=2.1.0)) | export: osgi.wiring.package=javax.persistence org.hibernate.javax.persistence.hibernate-jpa-2.1-api [333.0] {code}	1	1	https://issues.jboss.org/browse/ENTESB-2928
11798	ENTESB-2905	[~ffang], do you know where is the problem?	1	1	https://issues.jboss.org/browse/ENTESB-2905
11799	ENTESB-2880	[~ffang], do you know, how to determine, which bundles/features caused this "two dependency chains" problem? I'm unable to find the right combination that causes it. Here is [logfile | https://issues.jboss.org/secure/attachment/12388407/fuse-6.2-build-92-rhel7-java7.txt] with all installed/uninstalled features. Once I know, how to find out bundles related to dependency chain problem, I can provide better reproducer.	1	1	https://issues.jboss.org/browse/ENTESB-2880

#	Requirement	Questions	Answered	Successfully Answered	URL
11800	ENTESB-2880	I just checked it on fuse 6.2.0.redhat-102 and got only this exception: {noformat} java.lang.Exception: Could not start bundle mvn:org.xerial.snappy/snappy-java/1.1.0.1-redhat.001 in feature(s) camel-avro-2.15.0.redhat-620103, camel-hdfs2-2.15.0.redhat-620103, camel-hbase-2.15.0.redhat-620103: Activator start error in bundle org.xerial.snappy.snappy-java [273]. at org.apache.karaf.features.internal.FeaturesServiceImpl.doInstallFeatures(FeaturesServiceImpl.java:525)[9:org.apache.karaf.features.core:2.4.0.redhat-620103] at org.apache.karaf.features.internal.FeaturesServiceImpl$1.call(FeaturesServiceImpl.java:407)[9:org.apache.karaf.features.core:2.4.0.redhat-620103] at java.util.concurrent.FutureTask.run(FutureTask.java:262)[:1.7.0_76] at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145)[:1.7.0_76] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615)[:1.7.0_76] at java.lang.Thread.run(Thread.java:745)[:1.7.0_76] Caused by: org.osgi.framework.BundleException: Activator start error in bundle org.xerial.snappy.snappy-java [273]. at org.apache.felix.framework.Felix.activateBundle(Felix.java:2204)[org.apache.felix.framework-4.4.1.jar:] at org.apache.felix.framework.Felix.startBundle(Felix.java:2072)[org.apache.felix.framework-4.4.1.jar:] at org.apache.felix.framework.BundleImpl.start(BundleImpl.java:976)[org.apache.felix.framework-4.4.1.jar:] at org.apache.felix.framework.BundleImpl.start(BundleImpl.java:963)[org.apache.felix.framework-4.4.1.jar:] at org.apache.karaf.features.internal.FeaturesServiceImpl.doInstallFeatures(FeaturesServiceImpl.java:521)[9:org.apache.karaf.features.core:2.4.0.redhat-620103] ... 5 more Caused by: java.lang.UnsatisfiedLinkError: no libsnappyjava.so in java.library.path at java.lang.ClassLoader.loadLibrary(ClassLoader.java:1886)[:1.7.0_76] at java.lang.Runtime.loadLibrary0(Runtime.java:849)[:1.7.0_76] at java.lang.System.loadLibrary(System.java:1088)[:1.7.0_76] at org.xerial.snappy.SnappyBundleActivator.start(SnappyBundleActivator.java:52) at org.apache.felix.framework.util.SecureAction.startActivator(SecureAction.java:645) at org.apache.felix.framework.Felix.activateBundle(Felix.java:2154) ... 9 more {noformat} [~jludvice], can you check with this native library available?	1	1	https://issues.jboss.org/browse/ENTESB-2880
11801	ENTESB-2877	this issue also affects camel-spark-rest: {code} feature 482/518 - camel-spark-rest will be tested installation of feature Feature "camel-spark-rest" (2.15.0.redhat-620004) from camel-2.15.0.redhat-620004 failed for reason: ←[31mError executing command: Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.eclipse.jetty.client.util)(version>=8.0.0)(!(version>=10.0.0))) camel-jetty9 package:(&(package=org.eclipse.jetty.client.api)(version>=8.0.0)(!(version>=10.0.0))) camel-jetty9 {code}	0	0	https://issues.jboss.org/browse/ENTESB-2877
11802	ENTESB-2825	No code changes required for this functionality, the behavior comes from https://github.com/jboss-fuse/fuseenterprise/commit/5b211ece28bfe367a8faca2a14ab62d916d002c4 corresponding to ENTESB-1758 What is currently *supported* is: - github https protocol - github pub key authentication protocol (assumes key in {{~/.ssh/ }}folder What is *not supported*: - ssh password authentication. (we need code change for this, see http://www.codeaffine.com/2014/12/09/jgit-authentication/) Note that pub key authentication relies on plain .ssh configuration. So you need to verify from command line that the user that will be running the process is also able to access remote ssh location. It's also required to accept key fingerprint of remote servers. Some example configuration I have tested succesfully: {code} # pubkey github # no user/pass required since authentication information is taken from git client and ssh client configuration profile-edit --pid io.fabric8.datastore/gitExternalUrl=git@github.com:paoloantinori/test fabric {code} {code} # https github. # user and pass are read and used only in case of https protocol ! profile-edit --pid io.fabric8.datastore/gitExternalUrl=https://github.com/paoloantinori/test.git fabric profile-edit --pid io.fabric8.datastore/gitRemoteUser=pantinor@redhat.com fabric profile-edit --pid io.fabric8.datastore/gitRemotePassword=******* fabric {code} {code} # private ssh server, pub key only profile-edit --pid io.fabric8.datastore/gitExternalUrl=pantinor@localhost:/tmp/git_repo fabric 1.1 # this setup required the additional customization of ~/.ssh/config to disable manual interaction to accept server key fingerprint. # Host localhost # Port 22 # VerifyHostKeyDNS no # StrictHostKeyChecking no {code}	1	1	https://issues.jboss.org/browse/ENTESB-2825
11803	ENTESB-2821	[~virchete] : Does this apply any more after ER2 (https://github.com/jboss-fuse/fuse-integration/releases/tag/fuse-integration-1.0.0.redhat-620017)?	0	0	https://issues.jboss.org/browse/ENTESB-2821
11804	ENTESB-2788	The only way is to fork pax-cdi off of {{0.x}} branch. there's pax-cdi 1.0-SNAPSHOT waiting for openwebbeans 1.5 to be released. It's CDI-1.2 only. [~gnodet], what do you think?	1	1	https://issues.jboss.org/browse/ENTESB-2788
11805	ENTESB-2778	Hi Freeman, I noticed that test for camel-rx feature passed on HP UX recently and I'm also able to install it on my laptop *without* editing etc/jre.properties. Was there some change in requirements for this feature since camel 2.15? Tested with jboss fuse 6.2 build 100 there is sun.misc in {code} jboss-fuse-6.2.0.redhat-100/etc/config.properties: org.osgi.framework.system.packages.extra= \ 121 org.apache.xerces.impl.validation; version="2.11.0", \ 122 org.apache.xml.serialize; version="2.11.0", \ 123: sun.misc {code}	0	0	https://issues.jboss.org/browse/ENTESB-2778
11806	ENTESB-2773	Hmm, in upstream sources there is skip tests of camel-leveldb component on windows and AIX platforms defined in [camel-leveldb pom.xml|https://git-wip-us.apache.org/repos/asf?p=camel.git;a=blob;f=components/camel-leveldb/pom.xml;h=75e2b54bd526405ff9733444c17f1d7943e9de8f;hb=HEAD] Feature specification in [features.xml|https://repository.jboss.org/nexus/content/groups/ea/org/apache/camel/karaf/apache-camel/2.14.0.redhat-620073/apache-camel-2.14.0.redhat-620073-features.xml] including JNI driver: {code:xml} <feature name="camel-leveldb" version="2.14.0.redhat-620073" resolver="(obr)" start-level="50"> <feature version="2.14.0.redhat-620073">camel-core</feature> <bundle dependency="true"> wrap:mvn:org.fusesource.leveldbjni/leveldbjni-all/1.8$Bundle-Version=1.8&Export-Package=*;-noimport:=true;version="1.8" </bundle> <bundle dependency="true">mvn:org.fusesource.hawtbuf/hawtbuf/1.11</bundle> <bundle> mvn:org.apache.camel/camel-leveldb/2.14.0.redhat-620073 </bundle> </feature> {code} I checked jar leveldbjni-all-1.8-redhat.002.jar and it contains native library for * linux32 * linux64 * osx * windows32 * windows64 Quoting chat with gnodet: <jludvice_> gnodet: is it possible to define feature in karaf that would use some bundle on RHEL and other bundle on other platforms? <jludvice_> gnodet: https://issues.jboss.org/browse/ENTESB-2773 <gnodet> jludvice_: not with karaf 2.4	1	0	https://issues.jboss.org/browse/ENTESB-2773

#	Requirement	Questions	Answered	Successfully Answered	URL
11807	ENTESB-2773	This issue affects also HP UX 11 and fuse 6.2 build 093 {code} Feature "camel-leveldb" (2.15.0.redhat-620006) from camel-2.15.0.redhat-620006 => ←[31mError executing command: Could not start bundle wrap:mvn:org.fusesource.leveldbjni/leveldbjni-all/1.8$Bundle-Version=1.8&Export-Package=*;-noimport:=true;version="1.8" in feature(s) camel-leveldb-2.15.0.redhat-620006: Unresolved constraint in bundle org.fusesource.leveldbjni.leveldbjni-all [605]: No matching native libraries found. {code}	1	0	https://issues.jboss.org/browse/ENTESB-2773
11808	ENTESB-2763	Just retested with jboss fsue 6.2 build 092 and got following message {code} JBossFuse:karaf@root> features:list | grep camel-jcr [uninstalled] [2.15.0.redhat-620004] camel-jcr camel-2.15.0.redhat-620004 JBossFuse:karaf@root> features:install camel-jcr Error executing command: Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=javax.jcr.security)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons package:(&(package=javax.jcr.observation)(version>=2.0.0)(!(version>=3.0.0))) camel-jcr package:(&(package=org.apache.jackrabbit.value)(version>=2.2.0)(!(version>=3.0.0))) camel-jcr package:(&(package=javax.jcr)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons package:(&(package=javax.jcr.query)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons package:(&(package=javax.jcr)(version>=2.0.0)(!(version>=3.0.0))) camel-jcr package:(&(package=javax.jcr.lock)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons package:(&(package=org.apache.jackrabbit.util)(version>=2.4.0)(!(version>=3.0.0))) camel-jcr package:(&(package=javax.jcr.version)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons package:(&(package=javax.jcr.util)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons package:(&(package=javax.jcr.query.qom)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons package:(&(package=javax.jcr.observation)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons package:(&(package=javax.jcr.nodetype)(version>=2.0.0)(!(version>=3.0.0))) Jackrabbit JCR Commons {code} Which build did you try [~willem.jiang]?	1	1	https://issues.jboss.org/browse/ENTESB-2763
11809	ENTESB-2763	I just ran some test with Karaf and Camel 2.15.0, it's not the Fuse ESB. I will give the last build kit a try tomorrow. Josef, do you know where can I get the kit?	1	1	https://issues.jboss.org/browse/ENTESB-2763
11810	ENTESB-2705	Hi [~emil.cervenan] , can you elaborate more when you say that "Whole io.fabric8.agent.properties file is broken afterwards. " ? I tried your steps, and it fails, because we are passing a not value value for that property. If I invoke : {code} profile-edit --pid 'io.fabric8.agent/org.ops4j.pax.url.mvn.repositories=http://someValue=my' --append default JBossFuse:karaf@root> profile-display default Profile id: default Version : 1.0 Attributes: Containers: Container settings ---------- Repositories : mvn:org.apache.karaf.assemblies.features/standard/${version:karaf}/xml/features mvn:io.fabric8/fabric8-karaf/${version:fabric}/xml/features Features : fabric-agent fabric-git fabric-jaas insight-log jolokia fabric-git-server karaf fabric-web fabric-core Overrides : mvn:org.ops4j.pax.logging/pax-logging-service/1.8.1 mvn:org.apache.servicemix.bundles/org.apache.servicemix.bundles.jsch/0.1.51_1 mvn:javax.mail/mail/1.4.7 Agent Properties : patch.repositories = https://repo.fusesource.com/nexus/content/repositories/releases, https://repo.fusesource.com/nexus/content/groups/ea metadata#org.apache.felix.scr#[1.0,2.0)#Provide-Capability = osgi.extender;osgi.extender=osgi.service.component, osgi.service;effective:=active;objectClass="org.apache.felix.scr.ScrService" optional.ops4j-base-lang = mvn:org.ops4j.base/ops4j-base-lang/1.4.0 resolve.optional.imports = true io.fabric8.agent.download.threads = 4 metadata#org.apache.aries.blueprint.core#[1.0,2.0)#Provide-Capability = osgi.service;effective:=active;objectClass="org.apache.aries.blueprint.services.ParserService" org.ops4j.pax.url.mvn.globalUpdatePolicy = always metadata#org.apache.felix.metatype#[1.0,2.0)#Provide-Capability = osgi.service;effective:=active;objectClass="org.osgi.service.metatype.MetaTypeService" org.ops4j.pax.url.mvn.repositories = file:${runtime.home}/${karaf.default.repository}@snapshots@id=karaf-default, file:${runtime.data}/maven/upload@snapshots@id=fabric-upload, http://repo1.maven.org/maven2@id=central, https://repo.fusesource.com/nexus/content/groups/public@id=fusepublic, https://repository.jboss.org/nexus/content/repositories/public@id=jbosspublic, https://repo.fusesource.com/nexus/content/repositories/releases@id=jbossreleases, https://repo.fusesource.com/nexus/content/groups/ea@id=jbossearlyaccess, http://repository.springsource.com/maven/bundles/release@id=ebrreleases, http://repository.springsource.com/maven/bundles/external@id=ebrexternal, http://someValue=my ... {code} Also, you can ignore the difference between escaped and unescaped values. {{java.util.Properties}} handles them both, and it's directly its method {{http://docs.oracle.com/javase/7/docs/api/java/util/Properties.html#store%28java.io.Writer,%20java.lang.String%29}} that introduces them while saving the properties files. Still being able to load correctly the values that do not present the escape sequence. Regarding the re-ordering issue of entries, we don't currently have a satisfactory solution, unluckily. More advance extensions of Properties class idea exist in the java world, but we have Felix constraints and switching to another implementation might not be easy/possible. For example: http://commons.apache.org/proper/commons-configuration/apidocs/org/apache/commons/configuration2/PropertiesConfiguration.html Can you verify if providing a valid location of a Maven repo solves to problem for you as well? thank you paolo	1	1	https://issues.jboss.org/browse/ENTESB-2705
11811	ENTESB-2702	I just updated the doc source for versions 6.1, 6.2, and 6.2.1, explaining that you need to create a fabric and deploy a cloud profile (for example, {{cloud-aws.ec2}}, {{cloud-openstack}}, and so on).	1	1	https://issues.jboss.org/browse/ENTESB-2702
11812	ENTESB-2688	Ideally we woudl have all the quickstartrs running as smoke tests. How many are missing?	1	1	https://issues.jboss.org/browse/ENTESB-2688
11813	ENTESB-2665	feature pax-cdi-web-openwebbeans has the same issue {code} JBossFuse:karaf@root> features:install pax-cdi-web-openwebbeans Error executing command: Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.eclipse.jetty.server.handler)(version>=9.0.0)(!(version>=10.0.0))) OPS4J Pax CDI Web OpenWebBeans Adapter package:(&(package=org.eclipse.jetty.servlet)(version>=9.0.0)(!(version>=10.0.0))) OPS4J Pax CDI Web OpenWebBeans Adapter {code}	1	1	https://issues.jboss.org/browse/ENTESB-2665
11814	ENTESB-2625	Is this done now or is there work still outstanding?	0	0	https://issues.jboss.org/browse/ENTESB-2625
11815	ENTESB-2625	Hi Willem, thanks for suggestion. Now it works. :) But how often does the clientMode consumer send requests to a server? Does it do only once, or is there some some timer? I had assumption that the clientMode consumer will send next request immediately after ever response from a server. But it's not like this.	1	1	https://issues.jboss.org/browse/ENTESB-2625

#	Requirement	Questions	Answered	Successfully Answered	URL
11816	ENTESB-2625	Hi Willem, do you mean modify netty server to don't close channel after message send ? But connection can accidentally fail (e.g. network problem, server down, ...). What in this case? Reset camel route to inicialize new connection? And will it work for with mina2 too (don't close IoSession after message send?) ? I am trying do do it with mina2+hl7 and it works very strange for me now. I will experiment with it now. Can I catch you on IRC? When and where? Your nickname is njiang, isn't it?	1	1	https://issues.jboss.org/browse/ENTESB-2625
11817	ENTESB-2427	The command was add in https://github.com/jboss-fuse/fuse/commit/59ddc9630b3a4ba700f0fc110f153976242ceaf3 For the {{amq:create-admin-user}} command, wouldn't it make sense to update the activemq.jmx.user/password properties as well as part of this command or do we want to leave it up to user whether or not they will use that?	1	1	https://issues.jboss.org/browse/ENTESB-2427
11818	ENTESB-2320	Why file this as requirement when it looks like a bug and is described like a bug?	1	1	https://issues.jboss.org/browse/ENTESB-2320
11819	ENTESB-2320	Because your engineering colleauges asked me to file it this way. Email chain attached. ========== Thanks Guillaume. Kevin would you log an enhancement request please? Remember to use JBFPL for enhancements. Susan On Oct 15, 2014, at 11:21 AM, Guillaume Nodet <gnodet@redhat.com> wrote: > No, there's currently no way to disable this behavior. > However, this may be improved, as I think the computation of the list brings too many bundles, and there may also be some ways to remove a few dependencies that could help a lot. But this need code changes. > > Guillaume > > ----- Original Message ----- > From: "Kevin Boone" <kboone@redhat.com> > To: fuse-eng-help@redhat.com > Sent: Wednesday, October 15, 2014 12:20:27 PM > Subject: Overzealous bundle restarts in Karaf/fabric > > Hi folks > > I'm sorry -- I've raised this subject before, but I didn't realize at > the time how many customers were affected by the problem. it seems to be > pretty pervasive. > > When a bundle gets redeployed, a bunch of other bundles seem to get > restarted. I'm not sure of the exact circumstances in which this > happens, but from the previous discussion it seemed to be a well-know > phenomenon. For a plain, bundle-by-bundle deployment we can use > "osgi:install -r" to limit the amount of restarting. However, in a > deployment based on features, or in a fabric environment, there doesn't > seem to be an equivalent. > > What this means is that when a large application with many bundles is > redeployed, we end up with Fuse bundles (Spring, Hibernate, Camel, etc) > getting restarted, leading to services being unavailable when the > customer's application starts. Here's a example I saw this morning: > > "Caused by: org.apache.camel.FailedToCreateRouteException: Failed to > create route route1: > Route(route1)[[From[timer:///foo?fixedRate=true&period=36000... because > of Failed to resolve endpoint: > timer:///foo?fixedRate=true&period=3600000 due to: No component found > with scheme: timer" > > I understand from the previous discussion that improvements are > under-way for Fuse 6.2, but we have a fair number of customers with > problems of this sort now, and Fuse 6.2 is some way off. > > Is there anything we can do, in a fabric, profile-based deployment, to > prevent any bundles being restarted at deployment time other than those > in the application itself? > > Best wishes > Kevin >	1	1	https://issues.jboss.org/browse/ENTESB-2320
11820	ENTESB-2313	Hi Susan, I have now drafted some doc for patching standalone ActiveMQ. Would you (or somebody else in GSS) like to review it? You can see a preview here: http://jenkinscat.gsslab.pnq.redhat.com:8080/job/doc-FuseSource-ActiveMQDocs-6-1/lastStableBuild/artifact/publican/admin_guide/tmp/en-US/html-single/index.html#ESBRuntimePatchActiveMQ Cheers, Fintan	1	1	https://issues.jboss.org/browse/ENTESB-2313
11821	ENTESB-2288	What target Linux env are you using?	0	0	https://issues.jboss.org/browse/ENTESB-2288
11822	ENTESB-2288	What OS were you checking this one on and what commands were not available?	1	1	https://issues.jboss.org/browse/ENTESB-2288
11823	ENTESB-2288	For 6.1 and greater we already check for telnet, curl and java before trying to create a container. What error were you getting exactly? I tried this on a box without telnet and I got the following: JBossFuse:karaf@root> container-create-ssh --user janstey --password ********** --host 192.168.2.54 --path /tmp/foo2-fuse foo2 The following containers have failed: foo2: janstey@192.168.2.54:22: received exit status 255 executing --- command --- #!/bin/bash function run { echo "Running: $*" ; $* ; rc=$? ; if ["${rc}" -ne 0]; then echo "Command Failed:Error running installation script: $*" ; exit ${rc} ; fi ; } ... --- output --- Running: mkdir -p /tmp/foo2-fuse Running: cd /tmp/foo2-fuse Running: mkdir -p foo2 Running: cd foo2 Java version is greater than 1.6. Command Failed:Telnet is not installed. --- error --- ------ JBossFuse:karaf@root>	1	1	https://issues.jboss.org/browse/ENTESB-2288
11824	ENTESB-2195	[~ffang] How much work is this? Can we get into 6.2?	1	1	https://issues.jboss.org/browse/ENTESB-2195
11825	ENTESB-2181	[~pklimcza], can you verify this issue using provided links?	0	0	https://issues.jboss.org/browse/ENTESB-2181
11826	ENTESB-2169	It makes no sense for the Fuse quickstarts to have a group id of "eap" or "fabric8". What was it for the Fuse 6.1 release?	1	1	https://issues.jboss.org/browse/ENTESB-2169
11827	ENTESB-2160	Is this in the 2.2.0 Milestone?	0	0	https://issues.jboss.org/browse/ENTESB-2160
11828	ENTESB-2146	{{spring-dm}} feature is (Fuse 6.3.x) referenced from these features: * spring-dm-web * activemq-client * camel-spring * fabric-camel I think we could try removing this... [~davsclaus] I remember you planned to remove this in Camel 2.18, right? [~gtully], [~dbosanac], [~chirino] can this be removed from {{activemq-client}} feature? is spring-dm needed there? I'll try to check if it's needed in {{fabric-camel}} feature.	1	0	https://issues.jboss.org/browse/ENTESB-2146
11829	ENTESB-2146	In Camel 2.17 camel-spring and therefore spring-dm is only installed if you really need this, eg you install camel-spring or camel-cxf. Its only camel-cxf that has a dependency on camel-spring due to namespace parsers and because camel-cxf is a big bloated module. That ought to be split up into camel-cxf-core camel-cxf-blueprint camel-cxf-spring camel-cxf-ws camel-cxf-rs All the other modules that uses spring, do not use spring-dm, but the spring 4.x version. Only a few requires either spring 4.0.x or 4.1.x or 4.2.x, but they have the osgi import range for that set correctly. Most can run with any of the spring 4.x version. With only camel-cxf using camel-spring, the need for separating spring-dm out of camel-spring is less needed. As camel-spring component in itself do not bring much value to the table when you use blueprint etc. Using camel-spring without spring-dm does not make much sense as it is really spring-dm, eg using spring xml files with Camel. Though you can still move spring-dm out of it, and leave camel-spring as kind of abstract component. The only value in there is that it integrates the spring expression language in Camel. Notice its only spring-dm on osgi that is the problem. People not using osgi, have no problem at all, they can use camel-spring and whatnot with spring 4.x without any spring 3.x libraries.	1	1	https://issues.jboss.org/browse/ENTESB-2146
11830	ENTESB-2146	[~davsclaus] thanks for explanation. I wouldn't touch camel-spring - it's well known that in OSGi it's a layer on top of spring-dm, so it's fine. I think the explanation, that spring-dm is inherent part of camel-spring and camel-cxf should be enough to close this issue as "won't fix", [~rogerhui] - do you agree?	1	0	https://issues.jboss.org/browse/ENTESB-2146
11831	ENTESB-2142	Hi [~sonicaaaa], "admin" role was added across {{etc/auth/jmx.acl.*}} config files (ENTESB-3109). Could you include it also in this configuration?	1	1	https://issues.jboss.org/browse/ENTESB-2142
11832	ENTESB-2124	Do I need to write the new READMEs then?	1	1	https://issues.jboss.org/browse/ENTESB-2124
11833	ENTESB-2056	Which property / where should I set the timeout value to change the default one?	1	0	https://issues.jboss.org/browse/ENTESB-2056

#	Requirement	Questions	Answered	Successfully Answered	URL
11834	ENTESB-2033	[~chirino] Can you make this change for next build?	1	1	https://issues.jboss.org/browse/ENTESB-2033
11835	ENTESB-2003	What version of hibernate would we need for this?	1	1	https://issues.jboss.org/browse/ENTESB-2003
11836	ENTESB-1993	Can I close this now? I think it's done.	1	1	https://issues.jboss.org/browse/ENTESB-1993
11837	ENTESB-1992	First off, the final name for this new cart has not yet been decided. It's currently only available as downloadable cart but part of the productization will be making it an RPM cart (ENTESB-1993). What this cart does: It does a 'mvn clean install' build of the OSE app's git repo when the git repo is updated. This similar to how most other OSE carts works. It then exposes it's local maven repository over http so that it's builds can be accessed by other machines. Why Fuse OSE deployments benefit: Deploying user built binaries into a fabric cluster can be challenging since the mvn repo hosted by fabric nodes do not replicate and a node failure could cause a loss of availability of user deployed maven artifacts. Fabric deployments can be configured to add the new 'builder' app as a remote maven repo. Furthermore, OSE users can get a more streamlined continuous delivery system. They just need to check in their sources into a git repo, and the resulting binaries will be available to fabric deployments.	1	1	https://issues.jboss.org/browse/ENTESB-1992
11838	ENTESB-1992	should the user add this cart to all the fuse gears / only the fabric server / completely separate gear?	1	1	https://issues.jboss.org/browse/ENTESB-1992
11839	ENTESB-1992	Hi Hiram, I suppose that the best practice would be to have at least one Fuse/AMQ cart in OSE before creating builder to have system variables available for builder "auto config". Is it correct? This should be clarified in the doc also.	1	1	https://issues.jboss.org/browse/ENTESB-1992
11840	ENTESB-1800	Can you comment on this Jon?	0	0	https://issues.jboss.org/browse/ENTESB-1800
11841	ENTESB-1796	It's also important to note that just because remove the profiles, does not me that we are not shipping the binary bits to support what was in the profile configurations. Should we also attempt to remove the associated bundles from the OSE distro that implement those profiles?	1	1	https://issues.jboss.org/browse/ENTESB-1796
11842	ENTESB-1796	Can this be closed? I think it got done for 6.1 OSE release.	0	0	https://issues.jboss.org/browse/ENTESB-1796
11843	ENTESB-1737	Fintan, From the 6.1 release notes, https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Fuse/6.1/html/Release_Notes/files/FMQReleaseNotesNew.html the technical preview part, there is a: Fabric partition But do not have the Fabric8 Maven plug-in sorry I do not know what is the "Fabric partition" meaning? and what is "Process manager" meaning? I think there should have a brief description about the technical preview ones.	1	0	https://issues.jboss.org/browse/ENTESB-1737
11844	ENTESB-1736	Is this one still valid? Do you have specific details of what should be copied?	1	0	https://issues.jboss.org/browse/ENTESB-1736
11845	ENTESB-1733	I've tried to verify, but getting JMX connection exception. Also the "child" options is still present in Hawtio. Is it possible that your fixes aren't in 390-2 respin? {code} JBossFuse:admin@root> container-create-child root child The following containers have failed: <not available>: java.io.IOException: Failed to retrieve RMIServer stub: javax.naming.ServiceUnavailableException [Root exception is java.rmi.ConnectException: Connection refused to host: root-test.example.com; nested exception is: java.net.ConnectException: Connection refused] {code}	1	1	https://issues.jboss.org/browse/ENTESB-1733
11846	ENTESB-1728	Where can we launch OSE VM? Could we use OS1?	1	1	https://issues.jboss.org/browse/ENTESB-1728
11847	ENTESB-1723	Updating the shared "system" directory for all users (this is the big dir with most of the libs in it) won't help with patching though. It will still help save disk quota though since this can be the same for all users. So when Fuse boots up, the OSGi framework copies the set of runtime bundles into the "data/cache" dir along with other metadata. This dir would likely be different for each user. When patching Fuse, we patch the runtime container (i.e. the "data/cache" dir), NOT the "system" dir.	0	0	https://issues.jboss.org/browse/ENTESB-1723
11848	ENTESB-1709	Did we ship process container in 6.1?	0	0	https://issues.jboss.org/browse/ENTESB-1709
11849	ENTESB-1694	Can you explain a bit more what you mean? I do not understand what it is that you want improved?	0	0	https://issues.jboss.org/browse/ENTESB-1694
11850	ENTESB-1580	Would it be enough to add it to the Console Reference page for the command? This is the page: https://access.redhat.com/documentation/en-US/Red_Hat_JBoss_Fuse/6.1/html/Console_Reference/files/ConsoleFabricCreate.html	0	0	https://issues.jboss.org/browse/ENTESB-1580
11851	ENTESB-1539	Can I mark this as already Resolved?	1	1	https://issues.jboss.org/browse/ENTESB-1539
11852	ENTESB-1491	Is there a way to configure an mbean/path so that any user can access it? I've got to ensure anyone can access hawtio.plugins/* as hawtio has to look for available plugins before login to ensure the right branding/login background etc. is in place. So far that I can tell it seems you have to put in some role, tried * and an empty string with no luck...	0	0	https://issues.jboss.org/browse/ENTESB-1491
11853	ENTESB-1491	Could maybe have an empty string disable role matching for an mbean perhaps? I haven't looked at the JMX guard stuff so not sure how it's making it's permission checks...	0	0	https://issues.jboss.org/browse/ENTESB-1491
11854	ENTESB-1479	What are the steps to follow?	1	1	https://issues.jboss.org/browse/ENTESB-1479
11855	ENTESB-1479	Bugzilla bug 1096325 fixed by Customer Portal Operations team: https://bugzilla.redhat.com/show_bug.cgi?id=1096325 Kyle Benson <kbenson@redhat.com> changed: What |Removed |Added ---------- CC| |amatlack@redhat.com Flags| |needinfo?(amatlack@redhat.c | |om) --- Comment #1 from Kyle Benson <kbenson@redhat.com> --- Thanks Jane, I've updated this in prod, including the version number on the Overview page. The link there now should default to whatever is the most recent release. Allison when able if you can publish that'd be great	1	1	https://issues.jboss.org/browse/ENTESB-1479
11856	ENTESB-1477	What mis configuration. Can you provide an example.	0	0	https://issues.jboss.org/browse/ENTESB-1477
11857	ENTESB-1477	It seems that user can configure Log4j and forward its output to stdout, stderr. [~jpechanec] Would this solution be sufficient?	1	0	https://issues.jboss.org/browse/ENTESB-1477
11858	ENTESB-1467	Would an acceptable fix be to just remove the Note which says JDBC isn't transactional?	1	1	https://issues.jboss.org/browse/ENTESB-1467
11859	ENTESB-1456	{{<headerName>:<headerValue>}} works, but {{body:<value>}} does not work. Shall I report an issue? But I think I can close this issue.	1	1	https://issues.jboss.org/browse/ENTESB-1456
11860	ENTESB-1437	Checked 6.1.x.redhat-patch branch locally and on jenkins (http://10.8.63.16/view/JBoss%20Fuse%206.1%20patch/job/fuseenterprise-6.1.x.redhat-patch-checkin/) both built this module fine. Could you check mq/jboss-a-mq/target/dependency-pom.xml.log for more information on why it is failing on your machine?	1	0	https://issues.jboss.org/browse/ENTESB-1437

#	Requirement	Questions	Answered	Successfully Answered	URL
11861	ENTESB-1436	I noticed that dependencies.txt contains some artifacts of 2.5.6.SEC01 version: {code} grep spring dependencies.txt ... - org.springframework.aop org.springframework:org.springframework.aop:jar:2.5.6.SEC01 - org.springframework.beans org.springframework:org.springframework.beans:jar:2.5.6.SEC01 - org.springframework.context org.springframework:org.springframework.context:jar:2.5.6.SEC01 - org.springframework.core org.springframework:org.springframework.core:jar:2.5.6.SEC01 {code} I didn't find them exported anywhere. So, verified in fuse 6.2 117. Just for the sake of curiosity, how is the dependencies.txt generated?	0	0	https://issues.jboss.org/browse/ENTESB-1436
11862	ENTESB-1429	Is this the relevant page? https://access.redhat.com/site/documentation/en-US/Fuse_ESB_Enterprise/7.1/html/Configuring_and_Running_Fuse_ESB_Enterprise/files/ESBRuntimeRemote.html	1	1	https://issues.jboss.org/browse/ENTESB-1429
11863	ENTESB-1344	it. It does disappear but I keep receiving error messages: {code} Core] Operation getProfiles(java.lang.String, java.util.List) failed due to: io.fabric8.api.FabricException : Version '1.6' does not exist [Core] Stack trace: io.fabric8.api.FabricException: Version '1.6' does not exist at io.fabric8.service.FabricServiceImpl.getVersion(FabricServiceImpl.java:829) at io.fabric8.api.jmx.FabricManager.getProfiles(FabricManager.java:786) at sun.reflect.GeneratedMethodAccessor100.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at sun.reflect.misc.Trampoline.invoke(MethodUtil.java:75) at sun.reflect.GeneratedMethodAccessor12.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at sun.reflect.misc.MethodUtil.invoke(MethodUtil.java:279) at com.sun.jmx.mbeanserver.StandardMBeanIntrospector.invokeM2(StandardMBeanIntrospector.java:112) at com.sun.jmx.mbeanserver.StandardMBeanIntrospector.invokeM2(StandardMBeanIntrospector.java:46) at com.sun.jmx.mbeanserver.MBeanIntrospector.invokeM(MBeanIntrospector.java:237) at com.sun.jmx.mbeanserver.PerInterface.invoke(PerInterface.java:138) at com.sun.jmx.mbeanserver.MBeanSupport.invoke(MBeanSupport.java:252) at com.sun.jmx.interceptor.DefaultMBeanServerInterceptor.invoke(DefaultMBeanServerInterceptor.java:819) at com.sun.jmx.mbeanserver.JmxMBeanServer.invoke(JmxMBeanServer.java:801) at org.jolokia.handler.ExecHandler.doHandleRequest(ExecHandler.java:98) at org.jolokia.handler.ExecHandler.doHandleRequest(ExecHandler.java:40) at org.jolokia.handler.JsonRequestHandler.handleRequest(JsonRequestHandler.java:89) at org.jolokia.backend.MBeanServerExecutorLocal.handleRequest(MBeanServerExecutorLocal.java:109) at org.jolokia.backend.MBeanServerHandler.dispatchRequest(MBeanServerHandler.java:102) at org.jolokia.backend.LocalRequestDispatcher.dispatchRequest(LocalRequestDispatcher.java:98) at org.jolokia.backend.BackendManager.callRequestDispatcher(BackendManager.java:409) at org.jolokia.backend.BackendManager.handleRequest(BackendManager.java:158) at org.jolokia.http.HttpRequestHandler.executeRequest(HttpRequestHandler.java:197) at org.jolokia.http.HttpRequestHandler.handlePostRequest(HttpRequestHandler.java:131) at org.jolokia.http.AgentServlet$2.handleRequest(AgentServlet.java:399) at org.jolokia.http.AgentServlet.handle(AgentServlet.java:287) at org.jolokia.http.AgentServlet.doPost(AgentServlet.java:256) at javax.servlet.http.HttpServlet.service(HttpServlet.java:595) at javax.servlet.http.HttpServlet.service(HttpServlet.java:668) at org.eclipse.jetty.servlet.ServletHolder.handle(ServletHolder.java:684) at org.eclipse.jetty.servlet.ServletHandler$CachedChain.doFilter(ServletHandler.java:1496) at io.hawt.web.AuthenticationFilter$2.run(AuthenticationFilter.java:143) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.Subject.doAs(Subject.java:415) at io.hawt.web.AuthenticationFilter.executeAs(AuthenticationFilter.java:140) at io.hawt.web.AuthenticationFilter.doFilter(AuthenticationFilter.java:108) at org.eclipse.jetty.servlet.ServletHandler$CachedChain.doFilter(ServletHandler.java:1467) at io.hawt.web.CORSFilter.doFilter(CORSFilter.java:42) at org.eclipse.jetty.servlet.ServletHandler$CachedChain.doFilter(ServletHandler.java:1467) at	1	1	https://issues.jboss.org/browse/ENTESB-1344
11864	ENTESB-1344	Fixed up the error messages by switching back to the default version when you delete a version -> http://fusesource.com/forge/git/hawtio.git/?p=hawtio.git;a=commit;h=e5371943ce4ed1395cf4d8301a6328db5205e928 However I do notice that the versions still come back after restarting the container: {noformat} JBossFuse:karaf@root> version-list [version] [default] [# containers] 1.0 true 1 1.1 false 0 1.1.foo false 0 1.2 false 0 1.2.foo false 0 1.3 false 0 1.3.foo false 0 1.4 false 0 1.4.foo false 0 1.5 false 0 {noformat} Guess maybe they're not quite purged out of git?	1	1	https://issues.jboss.org/browse/ENTESB-1344
11865	ENTESB-1343	Did the patch fail to install or was the container just hanging while the patch was being applied?	0	0	https://issues.jboss.org/browse/ENTESB-1343
11866	ENTESB-1343	Do we need to make any changes to the product? Or is this just an update to the script that builds patches?	1	0	https://issues.jboss.org/browse/ENTESB-1343
11867	ENTESB-1334	The Release notes should have a summary of what's new in 6.1 and why you would want to upgrade. Can we add a "highlights" section at the start?	0	0	https://issues.jboss.org/browse/ENTESB-1334
11868	ENTESB-1334	What's new in 6.1 Full support of AMQP 1.0 – provides wire-level compatibility across connections Ability to modify broker configuration without requiring a restart Improved high availability (HA) – embedded message store for shared-nothing HA Certification of the JBoss A-MQ resource adapater with JBoss EAP 6 Support for MQTT 3.1.1 draft spec Vast library of connectors – over 200 out-of-the-box-connector via Apache Camel including new connectors for Salesforce and SAP Tooling to manage integration routes - start, stop, measure and trace Camel routes on-premise or in the Cloud	0	0	https://issues.jboss.org/browse/ENTESB-1334
11869	ENTESB-1290	btw we have disabled the Maven plugin out of the box in JBoss Fuse 6.1, as that plugin would require online internet access. And not everyone like that.	1	1	https://issues.jboss.org/browse/ENTESB-1290
11870	ENTESB-1289	{quote} [4REVIEWER] Need info: What's in the Local tab? is it the same as the Remote tab? if they're identical, there's no point in documenting them separately, and instead I will insert a note that states that they're identical {quote} The local tab only shows up if you've managed to get tools.jar in your classpath. It lists local JVMs running on your machine and lets you inject a jolokia instance and connect to that JVM, so it's pretty handy, however I've not ever been successful getting tools.jar available in the classpath in a Karaf server, so likely we can skip this to be honest. {quote} [4REVIEWER] UI feedback: there are a few entries in the drop-down list but when I select them the list shows a different value from what I selected. Also, none of the buttons next to the list seem to work? {quote} Ah so yes, probably best to make sure you go to "Preferences" and go to the "Reset" tab and clear your local storage. Since stuff changes between builds your local storage in your browser can become stale. So best to clear this when moving between builds. Not sure, maybe it's worth noting somewhere in the document that this could be necessary when migrating to newer versions of FMC. Also the reset button underneath the form (in the "Danger zone" section) just got moved to the preference page (sorry!)	1	1	https://issues.jboss.org/browse/ENTESB-1289

#	Requirement	Questions	Answered	Successfully Answered	URL
11871	ENTESB-1289	{quote} [4REVIEWER] UI feedback: this Create button is very similar to the buttons you use to create containers/profiles in other FMC areas. I was actually expecting to be able to create a container/profile from here, is there some way we can differentiate between the buttons? This is especially confusing because the Wiki page opens to the Profiles folder, so logically the Create button would let me create a profile? {quote} Damn, yeah, that's kind of a "create thing" button, i.e. you can create a property file or a regular file, a Fabric8 profile or even a new Fabric8 version. Oh yes, so yeah "create" should let you create a Fabric8 profile, should be the second item in the list, let me know if it's not there as that's a bug, the Fabric8 entries are dynamically enabled if the fabric has been created.	1	1	https://issues.jboss.org/browse/ENTESB-1289
11872	ENTESB-1289	{quote} . [4REVIEWER] UI feedback: in this area (as well as the other pie charts) the mouse cursor changes to a hand when I hover over the slices, but nothing happens when I click. Intentional? Mistake? What is supposed to happen when I click? {quote} Hmm, yeah, I don't believe anything is supposed to happen when you click the charts. {quote} [4REVIEWER] UI feedback + need info: Firstly, what exactly are loaded classes? The other pie slice text is cut-off and I can't read the full name of whatever the opposite of loaded classes is. Second, the area title has a typo in it. {quote} The text is all pulled straight from JMX for all of these charts (except for the typo, I can fix that :-)). It's LoadedClasses vs. UnloadedClasses, i.e. it tells you how many .class files the JVM has loaded and unloaded, handy for figuring out things like classloader leaks. {quote} [4REVIEWER] UI feedback: The area is missing the title (I assume it's "Threads", and "CPU" should be uppercase everywhere. {quote} Yeah, I'll add that in... Again the text in that area is what we get from the JVM. One comment to make on this section regarding the dashboard, it's not really a static thing, you can customize what's on the dashboard and move stuff around. Generally you look for the "Add to dashboard" link on a given page, for those charts there's icons that show up next to the name of an mbean attribute if there's a chart definition available. It'd be good to describe the "Edit dashboard" page a bit, as you can create more dashboards with their own stuff on 'em. Also when in a fabric, what dashboards are available depend on what profiles are deployed on the container you're looking at, so you can have customized per-profile dashboards. So for example if I have a profile that has a broker configured, I can then add a dashboard to that profile that contains stuff specific to a broker view, i.e. mbeans that show queue sizes and so on.	1	1	https://issues.jboss.org/browse/ENTESB-1289
11873	ENTESB-1278	Is this in a felix release yet?	0	0	https://issues.jboss.org/browse/ENTESB-1278
11874	ENTESB-1248	What about new switchyard quickstarts? Should they contain fabric plugin too?	0	0	https://issues.jboss.org/browse/ENTESB-1248
11875	ENTESB-1235	FYI you can access the apps individually outside of hawtio: kibana3 -> http://localhost:8181/kibana3/index.html eshead -> http://localhost:8181/eshead/ If you go to the kibana3 app it looks like maybe we've missed loading a script would be my first guess, javascript console output: Attr.specified is deprecated. Its value is always true. angular.min.js:41 event.returnValue is deprecated. Please use the standard event.preventDefault() instead. jquery.min.js:4 POST http://minimi:8181/kibana3/es/insight-2014.02.21/camel/_search 500 (Server Error) angular.min.js:106 TypeError: Cannot read property 'hits' of undefined at http://minimi:8181/kibana3/panels/table/module.js:175:59 at i (http://minimi:8181/kibana3/common/lib/angular.min.js:79:437) at http://minimi:8181/kibana3/common/lib/angular.min.js:80:485 at Object.e.$eval (http://minimi:8181/kibana3/common/lib/angular.min.js:92:272) at Object.e.$digest (http://minimi:8181/kibana3/common/lib/angular.min.js:90:142) at Object.e.$apply (http://minimi:8181/kibana3/common/lib/angular.min.js:92:431) at j (http://minimi:8181/kibana3/common/lib/angular.min.js:101:80) at r (http://minimi:8181/kibana3/common/lib/angular.min.js:104:449) at XMLHttpRequest.v.onreadystatechange (http://minimi:8181/kibana3/common/lib/angular.min.js:106:90) angular.min.js:63 POST http://minimi:8181/kibana3/es/insight-2014.02.21/camel/_search 500 (Server Error) angular.min.js:106	1	1	https://issues.jboss.org/browse/ENTESB-1235
11876	ENTESB-1235	Can we document this option somewhere as until now nobody was aware that we have to turn it on to use it ?	0	0	https://issues.jboss.org/browse/ENTESB-1235
11877	ENTESB-1235	Did you create any containers with the Elasticsearch storage profile? Without this there is nowhere to store or query logs. The error message could definitely be a bit nicer though.	0	0	https://issues.jboss.org/browse/ENTESB-1235
11878	ENTESB-1235	I've recreated this now. The profiles you need to assign are insight-elasticsearch.datastore (to create somewhere to store data), insight-logs.elasticsearch (to send logs to elasticsearch) & insight-metrics.elasticsearch (to send metrics to elasticsearch). However, even doing this one the JMX mbeans isn't properly registered if done in the order you've done it. You can solve this by restarting the "Fabric8 :: Insight :: Elastic Search Factory" bundle. Can you try that & see if the issue is fixed? If the elasticsearch profiles are assigned either at the same time or before the other profiles then this works fine. Looks like a problem in refreshing bundles when deploying subsequent profiles.	1	1	https://issues.jboss.org/browse/ENTESB-1235
11879	ENTESB-1223	In latest build (#350) with jclouds 1.7.1, only the jclouds-rackspace-cloudblockstorage-us feature is broken :-) JBossFuse:karaf@root> features:install jclouds-rackspace-cloudblockstorage-us Error executing command: Can not resolve feature: Unsatisfied requirement(s): ---------- package:(&(package=org.jclouds.rackspace.cloudidentity.v2_0.config)(version>=1.7.1)) jclouds Rackspace Next Generation Cloud Block Storage US provider JBossFuse:karaf@root> features:install jclouds-aws-sts JBossFuse:karaf@root>	1	1	https://issues.jboss.org/browse/ENTESB-1223
11880	ENTESB-1212	p17 You can deploy multiple types of containers, such as fabric management containers or application host containers. --> Wonder what you mean by application host containers? Is this a word for containers not-managed by fabric? eg just plain karaf? If so we do not have any way in the web console to create karaf containers. All containers created from the web console are fabric managed containers. After all fabric is the "star piece" of our product.	0	0	https://issues.jboss.org/browse/ENTESB-1212
11881	ENTESB-1212	p27 fabric:join --> explain better that the container_name option is. eg its the name you want this container to have in the fabric. Also give some pointers how to get this zookeeper url? And do we not have other documentation that covers Fuse Fabric better, you can refer the user for more details? eg maybe you dont need this chapter 6 at all in this guide? as this guide is for the web console.	0	0	https://issues.jboss.org/browse/ENTESB-1212
11882	ENTESB-1212	general observation I would like to know more about what and why I would need to use/do that. eg chapter 7 is about creating and deleting profiles. So what is a profile? And what do I use it for? And why do I wont to create a profile? And can I not edit/modify a profile? If so why do you not cover this also?	0	0	https://issues.jboss.org/browse/ENTESB-1212
11883	ENTESB-1212	p10 > You can choose a public, private, or open source cloud environment Is it not a public, private or hybrid cloud environment? And its all open source - eg all we do at Red Hat is open source.	1	1	https://issues.jboss.org/browse/ENTESB-1212
11884	ENTESB-1195	Is this already done in 6.1? Can I close it?	1	1	https://issues.jboss.org/browse/ENTESB-1195

#	Requirement	Questions	Answered	Successfully Answered	URL
11885	ENTESB-1193	After creating child container (fabric:container-create-child root superduper) and connect to it, i can see this output: {code:title=child container} JBossFuse:karaf@root> fabric:container-connect superduper Connected _ ____ _____		_\	____|		|_)	___ ___ ___	|__ _ _ ___ ___ _		_ < / _ \/ __/ __|	__|		/ __|/ _ \	|__|		|_)	(_) __ __ \				|_| __ \ __/ ____/|____/ ___/|___/___/ |_| __,_|___/___| JBoss Fuse (6.2.0.redhat-133) http://www.redhat.com/products/jbossenterprisemiddleware/fuse/ Hit '<tab>' for a list of available commands and '[cmd] --help' for help on a specific command. Open a browser to http://localhost:8181 to access the management console Create a new Fabric via 'fabric:create' or join an existing Fabric via 'fabric:join [someUrls]' Hit '<ctrl-d>' or 'osgi:shutdown' to shutdown JBoss Fuse. JBossFuse:admin@superduper> {code} If i create container with ssh (fabric:container-create-ssh --user sshuser --host 127.0.0.1 --new-user admin --new-user-password admin --password sshpassword superduper2), i can see this output : {code:title=ssh} JBossFuse:karaf@root> fabric:container-connect superduper2 Connected _____ _ _ _____	___|		(_)	_		|_ __ _| |__ _ __ _ ___ \ V /	_/ _`	'_ | '__| |/ __|/ _ \			(_| | |_)				(__| |_|	_| __,_|_.__/|_| |_|_______/ Fabric8 Container (1.2.0.redhat-133) http://fabric8.io/ Type 'help fabric' for a list of fabric related commands. Hit '<tab>' for a list of all available commands and 'help [cmd]' for help on a specific command. Hit '<ctrl-d>' or type 'logout' to logout from this container. Fabric8:admin@superduper2> {code} What should be correct message in child container ? Fuse 6.2.0 build 133	0	0	https://issues.jboss.org/browse/ENTESB-1193
11886	ENTESB-1169	Did that get done after 6.1 release?	0	0	https://issues.jboss.org/browse/ENTESB-1169
11887	ENTESB-1169	Kevin looks like you did a few of these upgrades prior to 6.1 GA? I found these branches: https://github.com/FuseByExample/smx-ws-examples/tree/jboss-fuse-6.1 https://github.com/FuseByExample/camel-dynamic-routing/tree/jboss-fuse-6.1 https://github.com/FuseByExample/rider-auto-osgi/tree/jboss-fuse-6.1 Could you update these and others you were working on to the GA version? Thanks, Jon	1	1	https://issues.jboss.org/browse/ENTESB-1169
11888	ENTESB-1169	Sure. What do I need to do, just merge changes on the jboss-fuse-6.1 branch to master?	1	1	https://issues.jboss.org/browse/ENTESB-1169
11889	ENTESB-1165	p11 Should Fuse Service Works not be named JBoss Fuse Service Works p11 Is there no ActiveMQ related book re-organization?	0	0	https://issues.jboss.org/browse/ENTESB-1165
11890	ENTESB-1165	p14 Is version 6.0 correct? Or should it say 6.1?	0	0	https://issues.jboss.org/browse/ENTESB-1165
11891	ENTESB-1165	p29 Do we need to mention those old production dependencies? eg since Camel 2.7 etc. That is a long time ago. I assume this is a migration guide only with stuff relevant for 6.0 -> 6.1 upgrades.	0	0	https://issues.jboss.org/browse/ENTESB-1165
11892	ENTESB-1108	How much effort is it to run this and at least know how far / close we are to compliance even if we cannot address all the issues for 6.2?	1	1	https://issues.jboss.org/browse/ENTESB-1108
11893	ENTESB-1108	Did you get a report for this yet?	1	1	https://issues.jboss.org/browse/ENTESB-1108
11894	ENTESB-1084	JBossFuse:karaf@root> list |grep 236 [236] [Active] [ ] [ ] [60] Fuse Fabric :: Jolokia (7.2.0.redhat-024) {code} 16:14:24,258 | DEBUG | s,name=root/list | SecureHttpContext | 101 - org.fusesource.fabric.fabric-maven-proxy - 7.2.0.redhat-024 | Login failed javax.security.auth.login.FailedLoginException: User does not have the required role admin at org.fusesource.fabric.maven.impl.SecureHttpContext.doAuthenticate(SecureHttpContext.java:114)[101:org.fusesource.fabric.fabric-maven-proxy:7.2.0.redhat-024] at org.fusesource.fabric.maven.impl.SecureHttpContext.authenticate(SecureHttpContext.java:153)[101:org.fusesource.fabric.fabric-maven-proxy:7.2.0.redhat-024] at org.fusesource.fabric.maven.impl.SecureHttpContext.handleSecurity(SecureHttpContext.java:76)[101:org.fusesource.fabric.fabric-maven-proxy:7.2.0.redhat-024] at org.ops4j.pax.web.service.jetty.internal.HttpServiceServletHandler.doHandle(HttpServiceServletHandler.java:68)[98:org.ops4j.pax.web.pax-web-jetty:1.1.11] at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:137)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.security.SecurityHandler.handle(SecurityHandler.java:534)[82:org.eclipse.jetty.security:7.6.7.v20120910] at org.eclipse.jetty.server.session.SessionHandler.doHandle(SessionHandler.java:227)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.server.handler.ContextHandler.doHandle(ContextHandler.java:1038)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.ops4j.pax.web.service.jetty.internal.HttpServiceContext.doHandle(HttpServiceContext.java:117)[98:org.ops4j.pax.web.pax-web-jetty:1.1.11] at org.eclipse.jetty.servlet.ServletHandler.doScope(ServletHandler.java:374)[80:org.eclipse.jetty.servlet:7.6.7.v20120910] at org.eclipse.jetty.server.session.SessionHandler.doScope(SessionHandler.java:189)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.server.handler.ContextHandler.doScope(ContextHandler.java:972)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:135)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.ops4j.pax.web.service.jetty.internal.JettyServerHandlerCollection.handle(JettyServerHandlerCollection.java:83)[98:org.ops4j.pax.web.pax-web-jetty:1.1.11] at org.eclipse.jetty.server.handler.HandlerWrapper.handle(HandlerWrapper.java:116)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.server.Server.handle(Server.java:359)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.server.AbstractHttpConnection.handleRequest(AbstractHttpConnection.java:483)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.server.AbstractHttpConnection.headerComplete(AbstractHttpConnection.java:920)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.server.AbstractHttpConnection$RequestHandler.headerComplete(AbstractHttpConnection.java:982)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.http.HttpParser.parseNext(HttpParser.java:635)[79:org.eclipse.jetty.http:7.6.7.v20120910] at org.eclipse.jetty.http.HttpParser.parseAvailable(HttpParser.java:235)[79:org.eclipse.jetty.http:7.6.7.v20120910] at org.eclipse.jetty.server.AsyncHttpConnection.handle(AsyncHttpConnection.java:82)[84:org.eclipse.jetty.server:7.6.7.v20120910] at org.eclipse.jetty.io.nio.SelectChannelEndPoint.handle(SelectChannelEndPoint.java:627)[81:org.eclipse.jetty.io:7.6.7.v20120910] at org.eclipse.jetty.io.nio.SelectChannelEndPoint$1.run(SelectChannelEndPoint.java:51)[81:org.eclipse.jetty.io:7.6.7.v20120910] at	1	1	https://issues.jboss.org/browse/ENTESB-1084
11895	ENTESB-1060	Generally it works out of the box after changing the dependencies to: {code:xml} <dependency> <groupId>org.apache.hadoop</groupId> <artifactId>hadoop-common</artifactId> <version>2.3.0</version> </dependency> <dependency> <groupId>org.apache.hadoop</groupId> <artifactId>hadoop-hdfs</artifactId> <version>2.3.0</version> </dependency> {code} but the key is - will this work inside OSGi? Checking it right now.	1	1	https://issues.jboss.org/browse/ENTESB-1060
11896	ENTESB-1060	I'm working on correct smx/bundles for hadoop 2.3.0 version. After that there's only a question - should we provide separate components: * {{camel-hdfs}} * {{camel-hdfs2}} or just force using 2.3.0 version of hadoop client libraries? The difference is only in pom (change from {{hadoop-core:1.2.1}} to {{hadoop-client:2.3.0}}.	1	1	https://issues.jboss.org/browse/ENTESB-1060
11897	ENTESB-1046	Does this apply to standalone containers only, or are these options configurable for fabric containers too?	1	1	https://issues.jboss.org/browse/ENTESB-1046

#	Requirement	Questions	Answered	Successfully Answered	URL
11898	ENTESB-1046	Looking at the etc/org.apache.karaf.shell.cfg file in my jboss-fuse-6.1.0 installation, there are no {{mac =}} or {{cipher =}} entries in it. So are the defaults shown in Guilaume's comment ^^ understood? I would only make *mac* and *chipher* entries if I wanted to change those defaults? As far as I can tell, the other options are *hmac-ripemd160* for mac, and *arcfour128* and *arcfour* for cipher ([https://access.redhat.com/site/solutions/420283]). Is this correct, or are there more?	1	1	https://issues.jboss.org/browse/ENTESB-1046
11899	ENTESB-1031	Can this one be closed? I think everything is in the 370 build.	0	0	https://issues.jboss.org/browse/ENTESB-1031
11900	ENTESB-1031	From reading through the JIRA comments I take it that action item #3 {quote} 3. Patch process allows us to purge unused artifacts. Ie.. activemq-osgi.5.8.0.redhat-60024 gets replaced by activemq-osgi.5.8.0.redhat-60033.. the -60024 artifact is deleted from system/ repo {quote} was not implemented? I am further wondering if point #1 got implemented fully as its not clear from the further case history. And in case it got implemented, how does it work and whose task is it to use the Maven plugin?	0	0	https://issues.jboss.org/browse/ENTESB-1031
11901	ENTESB-1030	What are the valid values for these properties? btw, in the above linked file, there seems to be an error/mismatch between command and property name in line 109 {{cm.setConnectionMaxIdleMinutes(getInt("aries.xa.connectionMadIdleMinutes", 15));}} where the command has *Max*, the property has *Mad*. Is that intentional or a typo?	1	1	https://issues.jboss.org/browse/ENTESB-1030
11902	ENTESB-1030	Torsten, can you provide me more info on what other properties than the pooling ones (I see three of those) I need to document, and some examples of how to use them? Where can I find a detailed description of each?	1	1	https://issues.jboss.org/browse/ENTESB-1030
11903	ENTESB-1008	Can you provide some background for why one would want to configure remote JMX SSL on JBoss Fuse. What value/benefit/advantage does it provide, and is there risk/downside associated with it? Is it a recommended management approach? Is it an alternative to using the Management Console? Trying to determine where this information should be presented.	1	1	https://issues.jboss.org/browse/ENTESB-1008
11904	ENTESB-1008	@Jane, - *What value/benefit/advantage does it provide, and is there risk/downside associated with it?* JMX allows us Remote monitoring and management via MBeans. (local/remote access) Remote monitoring and management requires security, to ensure that unauthorized persons cannot control or monitor your application. So in a secured environment it is a most needed functionality to make sure that the JMX Communication happens over SSL so that the data provided by the MBean Server to the JMX client (like JConsole) should not be read by some hacker. - *Is it a recommended management approach?* - Yes, In most of the secured production environment it is recommended to have the JMX access over the SSL, Sothat a hacker can not decrypt the data communicated b/w the MBean Server and JMX Client. SSL communication will make sure that the data is properly encrypted and decrypted. - *Is it an alternative to using the Management Console?* - Management console is a HTTP based utility however in some scenarios users wants to write their JMX / MBean based clients for notification/management logic. Like if some one wants to get the Activemq statistics using a java based JMX Client (like JConsole) then in that case as well SSL communication will be useful/recommended for a secured environment.	1	1	https://issues.jboss.org/browse/ENTESB-1008
11905	ENTESB-1008	Jay, in the step 1 code example, I assume that the string after -dname "....." is specific to your test environment. Can you give me an example that is more generic to demo the command for customers?	0	0	https://issues.jboss.org/browse/ENTESB-1008
11906	ENTESB-1000	Did Guillaume's comment help? Wondering if we can close this one.	0	0	https://issues.jboss.org/browse/ENTESB-1000
11907	ENTESB-736	For the JDF Quickstarts project, they also created a BOM POM, cfr. https://github.com/jboss-fuse/quickstarts/blob/master/bom/pom.xml Apart from some naming conventions differences for properties, the main difference is that this includes some of our main dependencies (like Camel, CXF, ...) in dependencyManagement, so it can be used in combination with <scope>import</scope> (e.g. https://github.com/jboss-fuse/quickstarts/blob/master/cbr/pom.xml) instead of having to use it as a parent pom. Perhaps we should add those dependencies to this BOM's dependency management section as well, so we can reuse the same POM for the quickstarts?	1	1	https://issues.jboss.org/browse/ENTESB-736
11908	ENTESB-448	Do other people see this on their system?	0	0	https://issues.jboss.org/browse/ENTESB-448
11909	ENTESB-414	Is there a good reason to upgrade? As we're already past code freeze we need to keep changes to a minimum to minimise risk of introducing regressions.	1	0	https://issues.jboss.org/browse/ENTESB-414
11910	ENTESB-374	Is this an issue for versions later than 7.0.x?	0	0	https://issues.jboss.org/browse/ENTESB-374
11911	ENTESB-264	Is this something we could include in the 7.1 release (the upgrade to Blueprint 1.0)?	1	1	https://issues.jboss.org/browse/ENTESB-264
11912	ENTESB-125	Is this true for all of our releases? I'm not sure where we would put the requirement for the 32-bit XP requirement since we do not list XP in our list of supported platforms. See http://fusesource.com/downloads/platforms-supported/	1	1	https://issues.jboss.org/browse/ENTESB-125
11913	ENTESB-61	Stan has a windows7 machine and the windows terminal is improved as you can customize the ansi colors. But on XP you cannot. Only the foreground / popup colors. See screenshot.	0	0	https://issues.jboss.org/browse/ENTESB-61
11914	JBID-206	Not in API but in core. Why is it difficult for folks to understand API means API? ;)	1	1	https://issues.jboss.org/browse/JBID-206
11915	JBID-194	Committed to main trunk with revision: 809. Forum thread about the refactoring: http://www.jboss.org/index.html?module=bb&op=viewtopic&t=161637	0	0	https://issues.jboss.org/browse/JBID-194
11916	JBEE-168	Where is the upstream repository for synching up with the new 1.3 release source changes? I couldn't open the [http://search.maven.org/#artifactdetails%7Cjavax.annotation%7Cjavax.annotation-api%7C1.3%7Cjar] link. Our repo for jboss-annotations-api is [https://github.com/jboss/jboss-annotations-api_spec]. [~wolfc], when we release a new jboss-annotations-api_1.3_spec, should we create a new branch for EAP 7.x, so that we have an upstream for EAP 7.x changes to jboss-annotations-api_spec, that is separate from the new EE 8 changes?	1	0	https://issues.jboss.org/browse/JBEE-168
11917	JBEE-140	Hi Shelly, I've being working with the JBoss JASPI 1.1 API review, specially the javadocs. The following PR holds all javadoc changes: https://github.com/jboss/jboss-jaspi-api_spec/pull/2 I've discussed with Stefan about the TCK signature tests. How can we run them considering all changes I did ? I could not find a way to run them locally. Thanks in advance.	1	1	https://issues.jboss.org/browse/JBEE-140
11918	JBEE-106	I noticed that Geronimo also made their servlet 3.0 spec jar into an OSGi bundle [1]. They export the javax.servlet packages as both version 2.6 and 3.0. Is there some official OSGi list of package versions? I'm just wondering why they used the 2.6. [1]http://repo1.maven.org/maven2/org/apache/geronimo/specs/geronimo-servlet_3.0_spec/1.0/geronimo-servlet_3.0_spec-1.0.pom	0	0	https://issues.jboss.org/browse/JBEE-106
11919	JBEE-97	BTW, just for the sake of curiosity, why is Xalan there ?	0	0	https://issues.jboss.org/browse/JBEE-97
11920	JBEE-97	If we are going to fork Xalan, wouldn't it make sense to also change the group ID to "org.jboss.xalan"? That way we could deploy this version to maven central.	0	0	https://issues.jboss.org/browse/JBEE-97

#	Requirement	Questions	Answered	Successfully Answered	URL
11921	JBEE-92	Hi Marcus, this project has moved to github (https://github.com/jboss/jboss-jacc-api_spec). Can you send a pull request which includes the changes in svn? https://source.jboss.org/changelog/JBossAS6?cs=112258 https://source.jboss.org/changelog/JBossAS6?cs=112264 Thanks!	0	0	https://issues.jboss.org/browse/JBEE-92
11922	JBEE-92	Where is the jboss-jacc-api_1.4_spec-1.0.1.Final tag? The version is published: https://repository.jboss.org/nexus/content/groups/public/org/jboss/spec/javax/security/jacc/jboss-jacc-api_1.4_spec/1.0.1.Final/ and used currently in AS7: <version.org.jboss.spec.javax.security.jacc.jbossjacc-api_1.4_spec>1.0.1.Final</version.org.jboss.spec.javax.security.jacc.jboss-jacc-api_1.4_spec> But I could not locate the tag here: https://github.com/jboss/jboss-jacc-api_spec/tags or https://svn.jboss.org/repos/jbossas/projects/specs/tags/	1	0	https://issues.jboss.org/browse/JBEE-92
11923	JBEE-91	Ping?	1	1	https://issues.jboss.org/browse/JBEE-91
11924	JBEE-91	Marek, we're in the process of migrating the Java EE APIs to git. http://community.jboss.org/wiki/JBossJavaEESpecsGitMigration The specs parent will no longer be used. What's the issue?	0	0	https://issues.jboss.org/browse/JBEE-91
11925	JBEE-91	Marek, is it ok to just add the license file to the trunk of the specs parent? I don't think we want to modify the tag.	1	1	https://issues.jboss.org/browse/JBEE-91
11926	JBEE-73	Shouldn't this be handled by the Surefire Manifest JAR? Which only applies when forking tests in a new VM, I suppose.	0	0	https://issues.jboss.org/browse/JBEE-73
11927	JBEE-73	You could create a separate project called something like jboss-javaee-6.0-all, then make the main artifact for that project an assembly that just unzips all dependencies (with scope set to 'provided' in that dependencySet). In the POM, you'd just put a single dependency on the jboss-javaee-6.0, with a scope of 'provided'. The configuration for the assembly plugin in the POM should be appendAssemblyId == false. The assembly descriptor should set includeBaseDirectory == false. The only caution here is that if a client builds an application using jboss-javaee-6.0-all as a dependency, and _doesn't_ mark that dependency as scope == provided, then those API classes may wind up in the deployable (EAR, WAR, etc.)...duplicating those already in the core classloader for the server. So, if we document that case really well, it should be very simple to create this project. Is this a problem beyond running JUnit tests? If not, is there a simple way to create a test launcher for Arquillian or something? Maybe that's the lazy way out, but I wanted to ask.	1	1	https://issues.jboss.org/browse/JBEE-73
11928	JBEE-67	of course....but where is the source project? Here a better writing in the MANIFEST: Bundle-SymbolicName: JTA API Bundle-Version: 1.0.1.GA Export-Package: javax.transaction;version="1.0.1.GA",javax.transaction.xa;version="1.0.1.GA"	1	1	https://issues.jboss.org/browse/JBEE-67
11929	JBEE-67	Luca, The sources for the jboss-transaction-api project are here: http://anonsvn.jboss.org/repos/jbossas/projects/specs/trunk/jboss-transaction-api_1.1_spec/ Do you have a patch ready?	0	0	https://issues.jboss.org/browse/JBEE-67
11930	JBEE-47	The interceptors 1.1 APIs are here: https://svn.jboss.org/repos/jbossas/projects/specs/trunk/jboss-interceptors-api_1.1_spec/ Latest release: https://repository.jboss.org/nexus/content/groups/public/org/jboss/spec/javax/interceptor/jboss-interceptors-api_1.1_spec/1.0.0.Beta1/ Is that what you need?	1	0	https://issues.jboss.org/browse/JBEE-47
11931	JBEE-33	Anil, can you confirm if there are updates need to these APIs?	1	1	https://issues.jboss.org/browse/JBEE-33
11932	JBEE-23	Why? As spec lead, we are the maintainer and guardians of javax.validation, why would we want to duplicate this work? Is there a way to create a "pom redirect" from your description to <dependency> <groupId>javax.validation</groupId> <artifactId>validation-api</artifactId> <version>1.0.0.GA</version> </dependency> Again, we are the one building the javax.validation:validation-api artifacts	1	0	https://issues.jboss.org/browse/JBEE-23
11933	JBEE-22	The API sources have been added to this project: http://svn.jboss.org/repos/jbossas/projects/specs/trunk/jboss-jaxrs-api_1.1_spec/ These APIs have not been released nor yet added to the javaee6-specs-bom. Would like to do so for the next AS 6 milestone release and add the APIs to jbossas/client and common/lib dirs. WDYT?	0	0	https://issues.jboss.org/browse/JBEE-22
11934	JBEE-21	Not following the purpose of this. We already have this under hibernate. Why do we need another?	1	1	https://issues.jboss.org/browse/JBEE-21
11935	JBEE-19	Carlo, this issue should be resolved with the compliant release of the EJB 3.1 APIs here: https://repository.jboss.org/nexus/content/groups/public/org/jboss/spec/javax/ejb/jboss-ejb-api_3.1_spec/1.0.0.Beta2/ I plan to do a Final release after JavaDoc review. Is there other work this OPEN JIRA represents?	0	0	https://issues.jboss.org/browse/JBEE-19
11936	JBEE-16	connector-api 1.5 has dependency on jboss common-core for SerialVersion class ejb 3.0 and 3.1 have a dependency on jbossws native jacc 1.1 has a dependency on jboss common-core for SerialVersion class jad-api has a dependency on jboss common-core for SerialVersion class jaxr-api has a dependency on javax.activation, can this be replaced?	1	1	https://issues.jboss.org/browse/JBEE-16
11937	JBEE-8	see r82194 comment: add new JCA 1.6 api into javaee/trunk and remove 1.5 api, it based on JCA 1.6 pr-to-pub version and change *InflowContext to *WorkContext	0	0	https://issues.jboss.org/browse/JBEE-8
11938	JBEE-4	Shelly, since i've changed the method signature to: public Object lookup(String name) throws IllegalArgumentException; from public Object lookup(String name); can you please make sure it won't break the tck signature test when we upgrade to 3.0.0.GA ?	1	1	https://issues.jboss.org/browse/JBEE-4
11939	JBJMX-108	Hi how can I send a patch with this feature request? thanks	1	0	https://issues.jboss.org/browse/JBJMX-108
11940	JBJMX-105	This issue is still present in JBoss5.0.0 GA. Does the new MicroContainer have the ability to resolve this?	0	0	https://issues.jboss.org/browse/JBJMX-105
11941	JBJMX-53	SourceForge Username: frimfram . Logged In: YES user_id=950487 How do I go about fixing this bug?	1	1	https://issues.jboss.org/browse/JBJMX-53
11942	JBKERNEL-16	This is related to the "related classes" (sic!) What is required is that DescribeAction adds dependencies for the related classes (not just AOP dependencies as per JBMICROCONT-22) Then the GenericBeanFactoryMetaData automatically adds the "beanClass" as a "related class".	0	0	https://issues.jboss.org/browse/JBKERNEL-16
11943	JBLAB-1003	Mark little responded with: "James, sorry it took so long to get back to you, but we've been busy. The team didn't really like any of the current ideas, but put something together themselves. Can you take a look and maybe suggest something along these lines?"	1	1	https://issues.jboss.org/browse/JBLAB-1003
11944	JBLAB-1003	I (Cheyenne) have created a third round of logos (13 versions), and submitted them to Mark Little on April 4th, 2008. http://www.jboss.org/files/jbosslabs/design/ESB/index.htm - Mark Little responded: "Thanks, I'll pass them along. One question: why is the j in JBoss lowercase in the examples when the BOSS is uppercase? Is this some new corporate logo approach?" - I fixed the upper/lower case issue :) - Mark's Response 4/4/08: "OK, thanks :-)"	1	0	https://issues.jboss.org/browse/JBLAB-1003
11945	JBLAB-1002	What's the status of this? Haven't heard anything on it since Meriah left (which has been a while)...	0	0	https://issues.jboss.org/browse/JBLAB-1002
11946	JBLAB-1001	Tom has added 1 reference from Rebecca's list. He's also checked Salesforce and there are a large number of companies that could be contacted. Who should do this? The requirements for the references page are: It should be editable by anyone, like a wiki. It should be part of the main jBPM project pages on JBoss.org.	0	0	https://issues.jboss.org/browse/JBLAB-1001

#	Requirement	Questions	Answered	Successfully Answered	URL
11947	JBLAB-994	September 17, 2008 Hello, artist-types! Folk keep asking me for a logo for JBoss Cache that they want to use in articles, etc. Could I bug you guys to come up with one? Currently I use the very cruddy attachment, but something with a bit more flair and catchiness would be good. Thanks!	1	1	https://issues.jboss.org/browse/JBLAB-994
11948	JBLAB-994	September 17, 2008 Mark Newton: Hi Manik, Do we want to create a logo for JBoss Cache or are we going to change the name of the community project to something else? Perhaps we should create a logo that can be used either with or without the JBoss Cache name so that we are free to make that kind of decision in parallel? cheers, Mark	1	1	https://issues.jboss.org/browse/JBLAB-994
11949	JBLAB-994	September 17, 2008 James Cobb: Hey Manik, We can definitely create a logo for you. I know we've discussed this in person before and I've had "JBoss Cache Logo" on our whiteboard as a project we need to do. If you're ready... we'll move this project to the front of the list. Do you have time tomorrow to discuss this via skype or IRC? James	1	1	https://issues.jboss.org/browse/JBLAB-994
11950	JBLAB-994	A couple more comments: re: 1, the cross section looks like the letter A. Was this intentional, and if so, what does it represent? :-) Maybe it would look better without the crossing bar? re: 4, it actually does look reasonably good, possibly because if it's simplicity. It's really just the colours that are Terracotta-esque. Perhaps using JBoss colours here may help?	0	0	https://issues.jboss.org/browse/JBLAB-994
11951	JBLAB-991	Hi James, sorry for following up so late. So, if you volunteered to build the initial (empty) pages, that would be a great help for us. We could pick the project up from there. Do you require access to our development environment or will you create it locally? regards, Christian	0	0	https://issues.jboss.org/browse/JBLAB-991
11952	JBLAB-991	Hi Christian, Thanks for picking up this project. Once you get the templates from James, can you provide an estimate on how long it will take to complete - or at least get everything into a sandbox for QA? Thanks, - Jon	0	0	https://issues.jboss.org/browse/JBLAB-991
11953	JBLAB-804	is this in progress?	0	0	https://issues.jboss.org/browse/JBLAB-804
11954	JBLAB-787	Sam, Can we get the list of content and whether or not it has been added? I'd like to resolve this issue quickly. Sincerely, Damon	1	0	https://issues.jboss.org/browse/JBLAB-787
11955	JBLAB-757	Isn't this issue resolved on dev?	0	0	https://issues.jboss.org/browse/JBLAB-757
11956	JBLAB-722	The link takes me to the main labs page. Have a different link?	1	1	https://issues.jboss.org/browse/JBLAB-722
11957	JBLAB-718	"In fact, we had this conversation months ago." Who is we? This is the *first* I'm hearing about this... and TomBen was quite surprised when I mentioned this to him also. Are you saying you had this converted and tested? Ignoring the fact that this was never communicated to the rest of us, how did you get around the problems below: 1) paths incorrect 2) load fails, even after basedir creations ? If you have some insight here, please share so we can get this resolved. I've included info from my testing, as well as where to find the copied repos and test dumps for your review in the next comment.	1	1	https://issues.jboss.org/browse/JBLAB-718
11958	JBLAB-718	Where's the copied repos and test dumps? I don't see any other comments.	1	0	https://issues.jboss.org/browse/JBLAB-718
11959	JBLAB-718	And, we're not suggesting you make any "major changes to the current repositories." We're suggesting you give him his own repository, just like all the rest, using the dump. Which is what this task is all about. How does that affect everyone else?	1	0	https://issues.jboss.org/browse/JBLAB-718
11960	JBLAB-687	The downloads section is actually a portlet that uses the downloads.xml of a project. To access projects downloads you'd go to: http://labs.jboss.org/jbosstm/downloads That said, what I think you are asking for is the general ability to browse directory structures in any given project's freezone directories (i.e., where they store all their content). Maybe by adding a flag in a project like "directories browsable" with a default of true (like in Apache httpd) where we generate an 'index.html' type page if one doesn't exist. As it is now, anyone can create an index.html in any directory, but the way we have it now is that we don't allow browsing of directories for security reasons, not that there's anything in there to worry about, but I chose to err on the safe side. Also, currently, the downloads portlet allows a user to choose to put their downloads in a non-browsable directory or put them in the freezone (or sourceforge for that matter). We are consolidating these directories. I think what we will do is consolidate the downloads and freezone directories, and add a feature that generates html to browse the directory structure of the freezone. Would these changes address your issue?	0	0	https://issues.jboss.org/browse/JBLAB-687
11961	JBLAB-666	nice work. Where is the usage doc?	1	0	https://issues.jboss.org/browse/JBLAB-666
11962	JBLAB-584	Adam, last time I looked at this, there was an issue with putting an image in the tags for our JSF components for the contributor apps. Can you look into this? Or, at least comment? Thanks, Damon	1	1	https://issues.jboss.org/browse/JBLAB-584
11963	JBLAB-579	I would like to sum up needed stats for forum: - total number of posts/topics in forum - number of new posts/topics in last 7/31 days Other ideas: - avg (topic views / number of topics) or avg (topic views) - number of users watching this forum - (number of users watching topic) X number of watched topics are those interesting for you? any other requests?	1	0	https://issues.jboss.org/browse/JBLAB-579
11964	JBLAB-519	Thomas, Just checking in again. Did you replace the ampersand? Does it work now? I'd like to turn on your project page. Application Server has already gone live. Cheers, Damon	1	0	https://issues.jboss.org/browse/JBLAB-519
11965	JBLAB-483	svn asks for a pwd. is this going to be sent to me automatically or do we have to take action to get one ? regards, tom.	1	1	https://issues.jboss.org/browse/JBLAB-483
11966	JBLAB-418	Ah, ok... I was just going from the comments above. Can this servlet be protected by something other than a password that has to be stored plain text (and in an open repo currently, but I realize this is just for testing)? Some sort of IP ACL at least, or some key based auth?	0	0	https://issues.jboss.org/browse/JBLAB-418
11967	JBLAB-403	Can you describe the problem? I might fix it in Kosmos, but don't want to duplicate before it gets clarified.	1	0	https://issues.jboss.org/browse/JBLAB-403
11968	JBLAB-403	Can anybody tell me what it is about ? Nobody knows... What does it mean "by project dnot by directory" ? Every project has it's own directory, so... i don't get it. Tomek	0	0	https://issues.jboss.org/browse/JBLAB-403
11969	JBLAB-400	what abvt this ? how about "Out of date" ?	0	0	https://issues.jboss.org/browse/JBLAB-400
11970	JBLAB-399	What colors are not right, and what are the "right" colors? I can replace them in the main Kosmos code, if it's something generic.	0	0	https://issues.jboss.org/browse/JBLAB-399
11971	JBLAB-399	I'm not sure what this has to do with localization. I'm just talking about adding the label that is on the left... which I assume is a list of JIRA types.. from the JIRA server... and put them *inside* the graph. That's all. There shouldn't be any server configuration--since the labels are *already* on the left of the chart anyway. Am I missing something?	0	0	https://issues.jboss.org/browse/JBLAB-399
11972	JBLAB-396	Oh.. probably should be implemented as jsp taglib tag.	1	1	https://issues.jboss.org/browse/JBLAB-396
11973	JBLAB-355	Adam, I removed a few things in the JSP. Instead of drastically modifying prj-info.jsp or whatever, I'd rather give certain projects the ability to use a freezone for the project info page. But, that will be an admin's decision (i.e., it will be in the 'projects.xml' rather than in the project.xml). What do you think?	0	0	https://issues.jboss.org/browse/JBLAB-355

#	Requirement	Questions	Answered	Successfully Answered	URL
11974	JBLAB-342	These requirements somewhat conflict with the current Kosmos implementation, so here are some questions: Q1: is it required to have a SINGLE portlet per project and implement everything inside that single one? - if N, then we can use one JIRA monitoring portlet and one SVN monitoring portlet from Kosmos per project. It means almost zero modifications, only feature additions -> small workload, high reuseability and easy testing - if Y, then it requires first developing a new "composite" portlet which incorporates both functionality -> higher workload, lower reuseability, harder to maintain for the long run and separate testing for this (My suggestion: definitely N, especially because the portlet titlebar are not visible on the Labs Portal, so it really doesn't make too much difference. Also, for the long run it's better for both parties: free bugfixes and features from both sides [Labs community, Kosmos community].) Q2: is it allowed to run a separate server component like in Kosmos? (It can run in the same JBoss AS instance as a separate web-app!) - if N -> low performance (rescanning the SVN repo is VERY expensive), extra workload to develop this "serverless" behaviour - if Y -> high performance, no extra work, a bit more complicated to deploy (have to do it only once!) (My suggestion: Y, I think the more complicated deployment and management doesn't matter for us, and invisible for users.) If you accept both my suggestions, then I just have to add the new features to Kosmos and reuse it as it is. It also means, that I can deploy the current implementations immediately, so apart from the new features, it can be played with.	1	1	https://issues.jboss.org/browse/JBLAB-342
11975	JBLAB-342	> 1. The same stats/graphs that you have in your JIRA porlet.. for the project... so people can see things like recently resolved tasks/bugs. "recently resolved tasks/bugs" is a new requirement, as it doesn't exist in Kosmos. (Doesn't seem difficult though.) > 2. The number of active committers for that project. How do you define "active"? A committer who committed a modification to SVN in the last X days? This is also new, but probably easy. > 3. A summary log for the last x commits What would that log contain? I can show the following info: date, committer, message, touched files [path and modification type (add/delete/modify)]. This is also new, probably easy (a bit difficult to find out a good UI (a lot of info goes to small screen-space), but I have some ideas). 2) Yes, a committer who committed a modifcation to SVN in the last X days. 3) We just need the date, committer, message, touch files as you say. That's exactly what I'm looking for.	1	1	https://issues.jboss.org/browse/JBLAB-342
11976	JBLAB-342	> Only requirement at the moment is that everything goes into the "NORMAL" view as that is the only view of a portlet that we support on Labs. OK, no problem. 1. Can you give me a list of project which you want to see monitored by the portlet(s)? 2. How do you want to group the projects to the portlet window(s)? a - a single window b - one window per project c - other (plz list the projects by group in this case)	0	0	https://issues.jboss.org/browse/JBLAB-342
11977	JBLAB-255	What do we need to do to get "everything in one place" ?? Get rid of downloads/document/freezone distinction? Make them all a freezone portlet? Can we get this done for the release on the 26th? I need answers! ;)	1	1	https://issues.jboss.org/browse/JBLAB-255
11978	JBLAB-255	yes, technically the documentation portlet and freezone are almost identical with the only difference being the types of links both parse - so if we take all current html pages hosted on labs in the Freezone section and use the freezone to translate them to HTML, we can then override those pages and simply use the 'docs' portlet to parse them - so instead of <a href = "/myDocs/index.html> it will have portal-relative link. And we won't use the ${link: format anymore - so docs becomes Freezone and users can upload/write/link any HTML document they like (in their directory or external) now to make it a bit clearer - it can be done for the release and will be done by then, I just need to know what do we want to support - only regular HTML pages, the way people have always been writing them or combine freezone style with it. that's all	1	1	https://issues.jboss.org/browse/JBLAB-255
11979	JBLAB-245	In the JIRA home page, should not be Kosmos listed in the "Category : JBoss Labs Projects" category?	1	0	https://issues.jboss.org/browse/JBLAB-245
11980	JBLAB-237	Isn't it related to JBLAB-228? Besides, because of having the blog entries in the repo (which would always be central, I suppose), multiple machines shouldn't be a problem.	1	0	https://issues.jboss.org/browse/JBLAB-237
11981	JBLAB-237	... regarding... multiple... machines... hmmm... I have to think about that. You may be right. I'm just wondering about just simple synchronization... can they get out of synch? I suppose there might be some cases where that might happen... but it might not matter as it will only happen/last for a minute. Interesting point.	0	0	https://issues.jboss.org/browse/JBLAB-237
11982	JBLAB-223	Tomek, can you clarify this task a little? I'm not sure what this is referring to, or what you would like to see.	1	1	https://issues.jboss.org/browse/JBLAB-223
11983	JBLAB-167	Can we close this issue? It looks to be done.	0	0	https://issues.jboss.org/browse/JBLAB-167
11984	JBLAB-128	Consider such scenario: we want to generate this report for a given period of time. The following actions happen: 1. developer A closes an issue somewhere between the start and end of the period that the report covers 2. developer B reopens the (same) issue after the end date of the report 3. developer C closes the issue Now, should developer A have this issue counted on his "issues closed" account or not?	0	0	https://issues.jboss.org/browse/JBLAB-128
11985	JBLAB-78	What exactly is missing for projects to safely use the forums ?	1	1	https://issues.jboss.org/browse/JBLAB-78
11986	JBLAB-74	From email from Eric Brown: Ivelin: "BTW, do you anticipate further delay of the Forge hardware deployment? What is the current ETA?" Eric's response: We are testing the base network and firewall setup now. We want to test and validate the failover of the switch infrastructure and firewalls (OpenBSD running CARP and pfsync for stateful failover). We've completed most of this (I should say TomBen has completed most of this) and we hope to wrap up a few remaining details over the weekend. We hope to have this deployed at the hosting facility next week before I go on my trip to China. TomBen will be following up the next week to get the hardware for labs and network in place (the easy part;).	1	1	https://issues.jboss.org/browse/JBLAB-74
11987	JBLAB-28	Can we look at LDAP vs. the existing user db from the website? I'd like to move this off to a directory so we can have a branch for Employees, Community, Partners and Customers. This will allow us to replicate portions of the directory for different systmes, with a single master. Thoughts?	1	0	https://issues.jboss.org/browse/JBLAB-28
11988	JBLAB-13	Marc initiated discussion with Bajz??t R??bert - rbajzat@freemail.hu. Robert is interested. Andy will follow up.	0	0	https://issues.jboss.org/browse/JBLAB-13
11989	LOGMGR-139	This is a rather complicated problem to solve. Since the suffix can be set to nearly any supported by {{java.text.SimpleDateFormat}} it makes it very difficult to parse file names to determine which files would be available for purging. There are some things that could be done to make some assumptions about file names, but having a log manager make assumptions about files on a file system seems wrong. You wouldn't want your log manage to delete files it shouldn't be deleting because it made an assumption about a file name. One option besides a custom handler is to use tools on the OS, like a cron job on Linux, to purge older files. As far as not getting your custom handler to work I'm guessing you mean on WildFly or JBoss EAP. What seems to be the issue there?	1	0	https://issues.jboss.org/browse/LOGMGR-139

#	Requirement	Questions	Answered	Successfully Answered	URL
11990	LOGMGR-139	The trick is given the file name {{server.log.20161606}} how do you get {{20161606}} back into a date format? It's possible you could assume the current {{suffix}} would be the date format, but now there's potential to miss purging files. For example say you get 5 days of log files and shutdown the server every Friday and restart it Monday. If you just take {{current_date - 5 days}} then you'd always miss purging Monday and Tuesdays log files. What I meant more by OS tools was using a cron job to delete files older than x days. For example {code} find ${JBOSS_HOME}/standalone/log -name "server.log.*" -mtime +5 | xargs -I {} rm "{}" {code} Given the configuration that should only rotate when the file reaches 1GB in size and keep 1 back up. So you should only ever have {{server.log}} and {{server.log.1}}. If it's got a date pattern suffix then something else is rotating the file.	0	0	https://issues.jboss.org/browse/LOGMGR-139
11991	LOGMGR-131	It looks like the {{StackTraceElement.toString()}} will already include the module and version. {code:java} public String toString() { String mid = ""; if (moduleName != null) { mid = moduleName; if (moduleVersion != null) mid += "@" + moduleVersion; mid += "/"; } return getClassName() + "." + methodName + "(" + mid + (isNativeMethod() ? "Native Method)" : (fileName != null && lineNumber >= 0 ? fileName + ":" + lineNumber + ")" : (fileName != null ? ""+fileName+")" : "Unknown Source"))); } {code} Do we still want format options for it?	1	1	https://issues.jboss.org/browse/LOGMGR-131
11992	LOGMGR-120	Is the goal here to pass messages through a filter regardless of the log level?	1	1	https://issues.jboss.org/browse/LOGMGR-120
11993	LOGMGR-120	Hi guys, This is the patch I have for this new feature. Can you guys take a look at it? [^force_logging.patch] Thanks! Thiago.	1	0	https://issues.jboss.org/browse/LOGMGR-120
11994	LOGMGR-120	Hi David, Thanks for the quick feedback. Regarding performance: * I added a system property for that so the user can opt-in. * When the system is enabled, the callback interface will be used to dynamically determine whether to force the log or not. It means that a custom implementation of this interface should take care of it; Example: an implementation could use memoization. {quote} ... the core of the logging system itself is the place that this modification should be happening, ... {quote} Hmmmm, I didn't quite get that. Where is the core of the logging system? :)	0	0	https://issues.jboss.org/browse/LOGMGR-120
11995	LOGMGR-120	It sounds like what you're describing is more along the lines of consulting a Filter (instead of doing a level check) to determine if a message is logged during certain operations. Would you say that is an accurate assessment?	0	0	https://issues.jboss.org/browse/LOGMGR-120
11996	LOGMGR-30	Just for curiousity: how is the progress on this issue?	1	1	https://issues.jboss.org/browse/LOGMGR-30
11997	LOGMGR-29	Can you briefly explain why the current behavior is not satisfactory (with use case)? It seemed to me that this method would not likely be used in the container environment.	1	1	https://issues.jboss.org/browse/LOGMGR-29
11998	LOGMGR-29	Interesting. And what is the expected behavior of the returned logger? Are its log messages expected to ultimately appear in the server log?	1	1	https://issues.jboss.org/browse/LOGMGR-29
11999	LOGMGR-29	Is there still a chance this patch could get applied? I did go to the trouble to patch the issue.	1	0	https://issues.jboss.org/browse/LOGMGR-29
12000	JBLOGGING-123	I have downloaded the latest JBoss Logging JAR (3.3.0.Final), but there is no "META-INF/services" folder. Which JAR exposes the interface LoggerProvider as service?	1	0	https://issues.jboss.org/browse/JBLOGGING-123
12001	JBLOGGING-107	No clue .. however I made a small change. basically I changed this line to : this.logger.log(translatedLevel, parameters == null || parameters.length == 0 ? this.messageFactory.newMessage(message) : this.messageFactory.newMessage(String.valueOf(message), parameters), and this works like a charm. I guess the problem is the parameters that are passed. The first parameter which is "null" in fact is the LEVEL... and I guess this is the issue causing the NPE, maybe log4j2 changed something for the final release ? You can see this on the screenshot https://dl.dropboxusercontent.com/u/887821/IDEA.png	0	0	https://issues.jboss.org/browse/JBLOGGING-107
12002	JBLOGGING-99	there are some extra jars which i have added for infinispan...is this issue due to some jar conflict?	0	0	https://issues.jboss.org/browse/JBLOGGING-99
12003	JBLOGGING-99	So ..how can i go forward with it..?should i downgrade infinispan version?will that be stable with jboss 5.1	1	1	https://issues.jboss.org/browse/JBLOGGING-99
12004	JBLOGGING-99	Sorry for the very long delay. Is this still an issue?	0	0	https://issues.jboss.org/browse/JBLOGGING-99
12005	JBLOGGING-95	Wanted to get someone's opinion before I start writing code... Log4j 2 defines an {{AbstractLogger}} with a {{log}} method that takes level, message, arguments, throwable, etc. that should be used for this improvement. (this is nearly identical to how Log4j 1 / {{org.jboss.logging.Log4jLogger}} works). However, this method is abstract, so it can only be called A) from a subclass, or B) from a class in the same package. So, the best/most correct/best performing way to implement this improvement is to add an {{org.apache.logging.log4j.spi.Log4j2Logger}} in the JBoss Logging source. The {{Log4j2LoggerProvider}} could still go in the {{org.jboss.logging}} package. Is this acceptable?	1	1	https://issues.jboss.org/browse/JBLOGGING-95
12006	JBLOGGING-95	Can we please get this merged and backported to 3.1, and release 3.1.4.GA?	1	1	https://issues.jboss.org/browse/JBLOGGING-95
12007	JBLOGGING-95	Huh. I didn't even notice that this had been merged and close. Guess I didn't get the email update or something. Thanks! Since this didn't get backported and included in 3.1.4.GA (sad face), when do we expect 3.2.0 to come out? Log4j 2 is going GA in the next couple of weeks, my book goes to the printers next week, and I'd really like this to all line up. I'd at least like to get 3.2.0.Beta1 rolling in the next week or two. There haven't been many changes at all to this version. :-) Is there anything I can do to help move things along?	1	0	https://issues.jboss.org/browse/JBLOGGING-95
12008	JBLOGGING-88	Can you describe the use case you have in mind?	1	1	https://issues.jboss.org/browse/JBLOGGING-88
12009	JBLOGGING-88	Reviewing pull requests today. I still don't see a clear use case that is solved this way; also, the usage of META-INF/services is not consistent with the ServiceProvider mechanism which uses this directory. Can you clarify a bit?	1	1	https://issues.jboss.org/browse/JBLOGGING-88

#	Requirement	Questions	Answered	Successfully Answered	URL
12010	JBLOGGING-84	I believe this RFE is related to the case I will describe, if not I will open a new RFE. Having come from a JBoss AS background, so it is more related to it. The needs come to have a property to set the default locale of jboss logging, and not depend upon default locale (-Duser.country=BR and -Duser.language=pt). However I see from the current code, it is not possible. org.jboss.logging.Messages.java public static <T> T getBundle(Class<T> type, Locale locale) { String language = locale.getLanguage(); String country = locale.getCountry(); String variant = locale.getVariant(); The simplest case to solve it to create a property at jvm level and in the describing code above, check for that property and use it, using as fallback the current implementation, what do you think ? Why I ask that ? From a application server point of view. 1. Its kind of weird to deal with some translations for technical terms, they just don't sound right. I am not blaming the translation text, only that it doesn't fit in the technical view to analyze logs and messages. 2. Not all messages are localized, so printing only english messages looks more consistent [org.jboss.as.naming] (MSC service thread 1-2) JBAS011802: Iniciando o Serviço de Nomeação [org.jboss.as.connector.subsystems.datasources] (ServerService Thread Pool -- 27) JBAS010403: Implantação do driver compatível-JDBC class org.h2.Driver (versão 1.3) [org.jboss.as.webservices] (ServerService Thread Pool -- 48) JBAS015537: Ativação da Extensão WebServices [org.jboss.as.mail.extension] (MSC service thread 1-4) JBAS015400: Sessão de correio limitado [java:jboss/mail/Default] [org.jboss.as.connector.subsystems.datasources] (ServerService Thread Pool -- 27) JBAS010404: Deploying non-JDBC-compliant driver class org.postgresql.Driver (version 9.1) [org.jboss.ws.common.management.AbstractServerConfig] (MSC service thread 1-3) JBoss Web Services - Stack CXF Server 4.0.4.GA-redhat-1 [org.apache.coyote.http11.Http11Protocol] (MSC service thread 1-4) Starting Coyote HTTP/1.1 on http-localhost/127.0.0.1:8080 [org.jboss.as.connector.subsystems.datasources] (MSC service thread 1-4) JBAS010400: Limite da fonte de dados [java:jboss/datasources/ExampleDS] [org.jboss.as.connector.subsystems.datasources] (MSC service thread 1-4) JBAS010400: Limite da fonte de dados [java:jboss/datasources/KitchensinkQuickstartSSO] [org.jboss.as.server.deployment] (MSC service thread 1-2) JBAS015876: Iniciando a implantação do "jboss-as-cdi-injection.war" [org.jboss.as.remoting] (MSC service thread 1-1) JBAS017100: Escutando no 127.0.0.1:9999 [org.jboss.as.remoting] (MSC service thread 1-3) JBAS017100: Escutando no 127.0.0.1:4447 3. If the user wants to search in the internet for help, it is not going to have success with this term "Sessão de correio limitado" 4. The user who uses an application servers is expected to understand many technical terms in english language, so he already understand relevant english terms to read the english log messages. 5. There are accents in many translated words, when they are displayed in terminal consoles (with different charset encoding), it only show a binary character instead of the accented char.	1	1	https://issues.jboss.org/browse/JBLOGGING-84
12011	JBLOGGING-84	I am looking to contribute this feature, should I fork which version at https://github.com/jboss-logging/jboss-logging ?	1	1	https://issues.jboss.org/browse/JBLOGGING-84
12012	JBLOGGING-78	Can you be specific about what version you are using?	0	0	https://issues.jboss.org/browse/JBLOGGING-78
12013	JBLOGGING-51	no problem....it's my work :) where are the sources of 3.x?	1	1	https://issues.jboss.org/browse/JBLOGGING-51
12014	JBLOGGING-6	I send you .java file for DatasourceAppender class	1	1	https://issues.jboss.org/browse/JBLOGGING-6
12015	JBLOGGING-6	Hi David. where I find jboss logging spi source codes? I would send a patch to transform it in a OSGI bundle	1	1	https://issues.jboss.org/browse/JBLOGGING-6
12016	JBMAIL-194	Is this admin tool has a GUI?	1	1	https://issues.jboss.org/browse/JBMAIL-194
12017	JBMAIL-147	okay so I'm down from > 100 errors to only 63... Will complete it by tommorrow I think.	0	0	https://issues.jboss.org/browse/JBMAIL-147
12018	JBMAIL-146	So I haven't figured out how we're going to autoinstall this in existing copies of JBoss. So for the moment let us just assume: 1. Most but NOT ALL users will install the JBAS that we give them (Probably true for milestone releases). 2. The ones that don't probably don't need the installer to do this and can follow the wiki instructions for now. Issues/difficulties: LDAP login module doesn't "create" schemas and we will NEVER make this config straightforward (but presumably if you are using LDAP you won't find it hard) until/unless we distribute our own. While I really prefer to have an out of the box "it runs" install, I'm not convinced LDAP configuration is something not best left to post installation configuration (preferably via a GUI with a "TEST" button). DB login module doesn't create tables. Easy enough for us to extend/enhance it, but then there are DB portability issues. A Hibernate/EJB3/Whatever login module would be our standard MO for database portability and centralized persistence configuration, unfortunately, I bet this would cause a circular dependency within JBAS. For the moment: I think a "jbms" realm with a users/roles template (see src/GRAPHICAL/templates/....) which replaces the login-config.xml and contains the users from the "define your 5 favorite users" screen in the installer is probably a good idea. Moreover we can switch the default configuration from the StaticUserRepository to the JAASUserRepository. As mentioned I do not wish to confuse JBMS with JEMS at the moment. Question: Do you want to do this? Are you able to do this? Do you want me to do this? Long term: I need to talk to scott about an alternative JAAS configuration system. While I understand the potential security issues that can be avoided by static load configuration, I think that the simplicity of installation/configuration is itself a security feature. However we have enough to bite off and chew this month!	1	1	https://issues.jboss.org/browse/JBMAIL-146
12019	JBMAIL-142	where is this?	0	0	https://issues.jboss.org/browse/JBMAIL-142
12020	JBMAIL-92	I actually want to talk this over with some of the folks who do this for a living. We have all the hooks, its just a matter of plugins. However, not by end of march.	0	0	https://issues.jboss.org/browse/JBMAIL-92
12021	JBMAIL-60	what about cheese ?	0	0	https://issues.jboss.org/browse/JBMAIL-60
12022	JBMAIL-60	We need a common key generation panel/config as this is something we need for jbossas setup of tomcat, jms, remoting ssl. Is this something that can be integrated into the jboss-izpack module for use across installer users?	1	1	https://issues.jboss.org/browse/JBMAIL-60
12023	JBMAIL-42	Hmmm... this looks, er, "interesting"... Have you looked at "Reinventing Email using REST"? http://www.prescod.net/rest/restmail/ Alternatively, one could simply route regular email protocols through HTTP's CONNECT, no? http://www.w3.org/Protocols/rfc2616/rfc2616-sec9.html#sec9.9	1	0	https://issues.jboss.org/browse/JBMAIL-42
12024	JBMAIL-36	Shouldn't a hash function (ie SHA-1) be more appropiate for this task?	1	0	https://issues.jboss.org/browse/JBMAIL-36
12025	JBMAIL-33	I think Mr. Rupp was working on this.. . Where is this? Is it done?	0	0	https://issues.jboss.org/browse/JBMAIL-33
12026	JBMAIL-22	Mike, where is this? Please update.	1	1	https://issues.jboss.org/browse/JBMAIL-22
12027	JBMAIL-9	Do we really need this? JBoss har deployer will do it automatically.	1	0	https://issues.jboss.org/browse/JBMAIL-9
12028	JBMAIL-7	We have stuff for JAAS, what is its state? Has it been tested? This is a feature we should be able to run wiht in M3.	0	0	https://issues.jboss.org/browse/JBMAIL-7
12029	JBMAIL-6	Sending to local user is fixed as well as received: header checking. There is still a problem with header body separation: basically the message looses its header body separation and the body part is a few characters too short. Do I need to add the normally terminating dot to have it being removed by creating the Mail() object? What about dot-escaping?	1	1	https://issues.jboss.org/browse/JBMAIL-6

#	Requirement	Questions	Answered	Successfully Answered	URL
12030	JBMAIL-6	Where is this?	1	1	https://issues.jboss.org/browse/JBMAIL-6
12031	JBMAR-157	I'm wondering if JBMAR should really need to reconstruct the exact same types, or if it would be acceptable to relax the requirement and allow to deserialize into alternative implementation, when a more suited structure is available. Example: an UnmodifiableList which wraps a LinkedList is being marshalled. Wouldn't it be more sensible to unmarshall this is an ArrayList, and wrap with an UnmodifiableList? Generally I think (hope) we pick LinkedList vs ArrayList making a performance decision based on how the list is filled in initially, but once the content is defined (as when it's marked by wrapping into unmodifiable) and a copy is needed, an ArrayList should always be a better choice. When marshalling an unmodifiable list, I'd say it could have a bit in the header to switch modifiable/immutable. Same thoughts could be applied to any concurrent collection which is wrapped into UnmodifiableCollection: we could pick the non-concurrent implementation as it's safe for reads, especially considering that after it being marshalled it's impossible for other threads to mutate the underlying collection.	1	1	https://issues.jboss.org/browse/JBMAR-157
12032	JBMAR-120	I have an issue with some of these changes. # Commits should not reorganize imports; it obscures the actual change. # Are you certain that reporting missing fields as "defaulted" and defaulting them on read matches the JDK behavior and is correct? # Are you certain that simply skipping the exception has predictable and correct semantics? What is the effect on the protocol? Are there cases where the writer could write user data, and the reader (by not reading fields) ends up with corruption, or vice-versa? # Have you tested this in detail at all? The reason I'm cautious is that the marshallers/unmarshallers are at a level of complexity that there are interactions which might be surprising or unpredictable.	0	0	https://issues.jboss.org/browse/JBMAR-120
12033	JBMAR-120	Commits should not reorganize imports; it obscures the actual change. Can fix. Are you certain that reporting missing fields as "defaulted" and defaulting them on read matches the JDK behavior and is correct? Absolutely, from what I remember there is an example of this in the JDK in TimeZone class (could be one of the other time related classes) going from version 6 to 7. JDK 7 did not complain at all when it deserialized a jdk 6 version but river did and after looking I narrowed it down to that issue. Will need to test covariant reads too. Are you certain that simply skipping the exception has predictable and correct semantics? What is the effect on the protocol? Are there cases where the writer could write user data, and the reader (by not reading fields) ends up with corruption, or vice-versa? I tested this before, and I did not have any problems, but it was a simple test using jdk classes going from version 6 to 7. There might be a problem if the protocols/(un)marshallers expect that there are no cases where there can be no fields read or written, if there is any serializable fields that the protocol/(un)marshaller expects certain bytes to be written regardless if they are written or not, if getFields() lazily sets some state that other parts of the unmarshaller needs other than the new/off/fields states that are read in the finish() methods. Have you tested this in detail at all? Enough to pass a simple test case.	0	0	https://issues.jboss.org/browse/JBMAR-120
12034	JBMAR-118	Actually infinispan call the marshallerFactory so: (MarshallerFactory) Thread.currentThread().getContextClassLoader().loadClass("org.jboss.marshalling.river.RiverMarshallerFactory").newInstance(); is it ok for you? What do you think?	1	1	https://issues.jboss.org/browse/JBMAR-118
12035	JBMAR-38	Hi, I would like to fix this bug. can you show me the fixing procedure..	1	0	https://issues.jboss.org/browse/JBMAR-38
12036	JBMESSAGING-1929	This would certainly be due to the 2 timeouts. With the write timeout set (not infinite), maybe a test that used to failed is now passing. I'd be surprised if the other settings affected a test. Is there anything that I can do to help?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1929
12037	JBMESSAGING-1929	Howard, I updated the remoting-bisocket-service.xml file. Could that be pulled into the latest tag? Thank you, Doug	1	0	https://issues.jboss.org/browse/JBMESSAGING-1929
12038	JBMESSAGING-1929	Hi Doug, The tag has been just made so It's too late for it. Could you create a separate Jira? Thanks Howard	0	0	https://issues.jboss.org/browse/JBMESSAGING-1929
12039	JBMESSAGING-1908	Howard, could you commit this file to the JBM source tree so one can easily copy it over the original one when desired?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1908
12040	JBMESSAGING-1890	Can you give org.jboss.jms.server.recovery.MessagingXAResourceRecovery2 a more descriptive name? I'm thinking something like org.jboss.jms.server.recovery.ClusteredMessagingXAResourceRecovery. What do you think?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1890
12041	JBMESSAGING-1890	Why not use a configuration like this: <mbean code="org.jboss.jms.jndi.JMSProviderLoader" name="jboss.messaging:service=JMSProviderLoader,name=JMSProviderRemote"> <attribute name="ProviderName">ClusteredRecoveryJMSProvider</attribute> <attribute name="ProviderAdapterClass">org.jboss.jms.jndi.JNDIProviderAdapter</attribute> <attribute name="FactoryRef">java:/ClusteredXAConnectionFactory</attribute> <attribute name="QueueFactoryRef">java:/ClusteredXAConnectionFactory</attribute> <attribute name="TopicFactoryRef">java:/ClusteredXAConnectionFactory</attribute> <attribute name="Properties"> java.naming.factory.initial=org.jnp.interfaces.NamingContextFactory java.naming.factory.url.pkgs=org.jboss.naming:org.jnp.interfaces java.naming.provider.url=${jboss.bind.address}:1099 </attribute> </mbean> Since the local ClusteredXAConnectionFactory is part of cluster its stub should have all the necessary information to contact all the other nodes in the cluster, right?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1890
12042	JBMESSAGING-1860	Hi Howard, could you please look at this enhancement and implement new behaviour. This enhancement will make our JBM cluster solution more HA. Thank you, Mirek	0	0	https://issues.jboss.org/browse/JBMESSAGING-1860
12043	JBMESSAGING-1857	Hi Roger, Is it a HornetQ issue? It's a different project than JBoss Messaging. Howard	1	0	https://issues.jboss.org/browse/JBMESSAGING-1857
12044	JBMESSAGING-1857	Howard, Yes, this is HornetQ, how do I switch to HornetQ? Should I reopen one?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1857
12045	JBMESSAGING-1841	Gerald, I'm speculating, and don't have time to test it for you, but I think I can imagine a scenario in which the JBM_MSG_REF table would have two rows with the same message ID: clustered topics. How has your modification stood-up to testing? I wouldn't mind trying the performance improvement myself if it has worked consistently for you (and it isn't subject to the hypothetical weakness above). queue the "this should be a forum discussion, not a bug comment" reply from a committer...	0	0	https://issues.jboss.org/browse/JBMESSAGING-1841
12046	JBMESSAGING-1802	[~jsightler], did you see the code comment in Howard's previous comment? It says: {noformat} We DO NOT support Strings written to BytesMessages or StreamMessages or written as keys or values in MapMessages, or as String properties or other String fields having serialized length >= 64K {noformat}	1	1	https://issues.jboss.org/browse/JBMESSAGING-1802

#	Requirement	Questions	Answered	Successfully Answered	URL
12047	JBMESSAGING-1802	Hi Justin, The comment itself essentially describes the bug, IMO. It's described as for "performance reasons", but I don't see how that could apply to only StreamMessage. Why would we support writeObject for an ObjectMessage (containing a String) of >65535, but not for a StreamMessage of the same? Performance concerns and considerations would be the same for both of those cases, IMO. I've tried this with HornetQ, and large messages are accepted with StreamMessage. I believe that competing products do as well (this came up in a recent WebLogic migration). Isn't the fix fairly straightforward anyway (pass the longStrings boolean all the way down)? Or is there more here that I'm not seeing?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1802
12048	JBMESSAGING-1772	JBM is supposed to be in bug fix only mode - i.e. no new features. So there needs to be very good justification for implementing any new features. What is the justification in this case?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1772
12049	JBMESSAGING-1667	We can check the name in DestinationServiceSupport.startService(), if it contains forward slash, throw an exception, that make the deployment fail. if (name.contains("/")) { throw new IllegalStateException("The " + (isQueue() ? "queue" : "topic") + " " + "name should not contain forward slash " + name); }	1	1	https://issues.jboss.org/browse/JBMESSAGING-1667
12050	JBMESSAGING-1667	Nobody called Justin has commented on this issue? Did you enter this comment in the wrong issue? :)	1	0	https://issues.jboss.org/browse/JBMESSAGING-1667
12051	JBMESSAGING-1615	I check the code and this is synchronous - i.e. there is no delay. Have you verifed this?	1	0	https://issues.jboss.org/browse/JBMESSAGING-1615
12052	JBMESSAGING-1591	How is this a JBM bug? Looks like this is a JBR bug in which it doesn't clean up after itself.	1	1	https://issues.jboss.org/browse/JBMESSAGING-1591
12053	JBMESSAGING-1546	Hi, JBM 1.4.3.GA is released specifically for JBOSS AS 5.1. Can you check the version of your JBoss AS? Thanks.	1	1	https://issues.jboss.org/browse/JBMESSAGING-1546
12054	JBMESSAGING-1546	So I think you are using the EAP version, this version has JBM 1.4.0.SP3-CP03 built in. You should be able to see something similar to the following line have appeared in your console: 09:46:16,709 INFO [ServerPeer] JBoss Messaging 1.4.0.SP3-CP03 server [0] started Supposed you are using the correct version as indicated above, the problem seems your network has something blocking your application. Did you check the connection or firewall settings? If the network is OK. can you please upload the piece of code by which you are looking for the connection factory? Also if you have an EAP support account, you can also raise the issue through the support channel, I believe you will get more help on this, and fast. :)	1	1	https://issues.jboss.org/browse/JBMESSAGING-1546
12055	JBMESSAGING-1546	Hi I am new to Jboss, sorry will you please tell me where do i find it and what should be in connection factory. Please let me know the path and file name in Jboss 4.3 1. file path 2. file name 3. what should be the content. if i put that will it resolve my issue? thanks	1	0	https://issues.jboss.org/browse/JBMESSAGING-1546
12056	JBMESSAGING-1546	does it related to the error that i am getting? how do i fix or re-configure this connection-factories-service.xml file to resolve my issue?	1	0	https://issues.jboss.org/browse/JBMESSAGING-1546
12057	JBMESSAGING-1513	Hi Len, Can you brief me what this script for? I had a look at SOA-310 but I don't know much about SOA project. Thanks. Howard	0	0	https://issues.jboss.org/browse/JBMESSAGING-1513
12058	JBMESSAGING-1491	Hi Folks, I've finished working on JBoss Messaging and running the test suite against Ingres community edition 9.2 build 134 on linux 32 bit. I've been using the './build.sh report' version of the tests and have been tuning Ingres so that it works correctly and passes all tests. For those that are interested, the settings that should be used/configured in Ingres are: 1/ DBMS->system_locklevel = row 2/ DBMS->connect_limit = 512 3/ DBMS->table_auto_structure = ON 4/ DBMS->qsf_guideline = huge 5/ Locking->per_tx_limit = 2500 It should be noted that any version on Ingres lower than 9.2 will not work correctly due to the enhancements made in 9.2 specifically around (C/B)LOB's. I have also tried using build 140 and 141 on linux 32 bit and 64bit. That release of Ingres has some problems and should not be used. Ingres is currently in our release cycle for the 9.2 enterprise version(as opposed to the community version). Releases come out for all our different platforms one by one. It is my intention to support JBoss Messaging on Ingres 9.2 enterprise edition going forth. I will continue to test running the JBoss Messaging test suite against Ingres 9.2 enterprise editions that are released if I have hardware and community editions, again if I have hardware. I've made a svn diff file for adding ingres support against the JBoss-Messaging-1.4.2.GA code line for people to use/commit etc. It is attached to this jIRA. I'd like to also add a file documenting suggested Ingres settings for use with JBoss Messaging, and I would think that other db's also have recommended settings, but not sure where to put this. Would it be possible to get this patch applied to the 1.4.3 code so it is included in that release? If not, is there anything I need to do to help out in getting Ingres added to the supported database options?	1	0	https://issues.jboss.org/browse/JBMESSAGING-1491
12059	JBMESSAGING-1491	Murray, great job! I agree we add this DB support in JBM. Tim what's your opinion?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1491
12060	JBMESSAGING-1491	THanks, I've been putting in a lot of hours over this at work(gotta love beening able to do open source as your job at work...) I'm certainly interested in also looking at version 2 and testing ingres against that code line too. Is version 2 code in a good position to look at this, or is it a little early before going there? Cheers Howard. Take care, Murray	1	0	https://issues.jboss.org/browse/JBMESSAGING-1491
12061	JBMESSAGING-1491	Hi Howard, How are things? Just wondering how things are going with getting a QA test setup going for Ingres. The version to use is the enterprise 9.2 version which is a licensed product. I have chased up making some arrangement for the QA testing to use this with an allowable license. Is there anything I can help with? Take care, Murray	1	0	https://issues.jboss.org/browse/JBMESSAGING-1491
12062	JBMESSAGING-1491	Hi Folks, Just wanting to touch base on this issue and try and get things moving along. I've chased up a legitimate Ingres enterprise license for the JBoss Messaging project to use for a copy of Ingres for the continuous test centre. Where do we go from here? Is there anything that I need to do or get done? My aim is to be able to offer Ingres as a supported database platform for JBoss Messaging. Take care, Murray	1	0	https://issues.jboss.org/browse/JBMESSAGING-1491
12063	JBMESSAGING-1416	Hi Howard, The customer has following questions to your comment (03/Feb/09 05:11 AM): Q1. In "The Programming Way" , you show us a way to produce messages in an ordering group with JBossMessageProducer API. It seems that the API is for the producer side. How about the consumer side then? Q2. In "The Configuration Way", is it complete transparent both for producers and consumers without JBoss specific API? Q3. The last line says "In case of Distributed Queue, user should use HASingleton to make sure ordering group works properly". Does this simply mean that this task covers the single server case in the JBoss Messaging 1.4 and not cover the clustered queues, which is targeted for 2.0? Regards, Fusayuki	1	1	https://issues.jboss.org/browse/JBMESSAGING-1416
12064	JBMESSAGING-1416	Hi Howard, Let me confirm the error scenario in JBoss Messaging 1.4. Does your comment (03/Feb/09 05:11 AM) cover the following scenario stated in this JIRA's description ? "In order to guarantee strict ordering, even on rollback or in the presence of multiple consumers or with xa transactions, the queue needs to ensure that no more than one message with the value of ordering group is being delivered at any one time. " Our customer requires that even in the singleton's case. Regards, Fusayuki	1	1	https://issues.jboss.org/browse/JBMESSAGING-1416
12065	JBMESSAGING-1416	Howard, Thanks a lot for the clarification. Now I'd like to write the customer's requirements for Distributed Queue. Should I input them here or in a new JIRA for it? Regards, Fusayuki	1	1	https://issues.jboss.org/browse/JBMESSAGING-1416

#	Requirement	Questions	Answered	Successfully Answered	URL
12066	JBMESSAGING-1416	Hi Howard, I have a requirement from the customer about "selector" limitation of you idea (03/Feb/09 05:11 AM). >* Ordering group shouldn't be used together with message selectors and scheduled delivery. We guess this statement means those just change the messagin ordering. But the customer need to use "selectors" because JBoss ESB uses selectors quite often for request-reply message pattern with queues: http://www.redhat.com/docs/en-US/JBoss_SOA_Platform/4.3.GA/html-single/Programmers_Guide/index.html#sect-SOA_ESB_Programmers_Guide-What_is_a_Service-ServiceInvoker For example, a user sends requests such as req-1, req-2, req-3 then should get the reply-1, reply-2, reply-3 in order, even with selectors for the reply queue. So the customer requires JBM keeps order of messages and doesn't throw exceptions when we even use selectors with ordering group function. Is that OK with you? Regards, Fusayuki	1	1	https://issues.jboss.org/browse/JBMESSAGING-1416
12067	JBMESSAGING-1416	Hi Fusayuki, Thanks for the info. "message selectors" in ordering group has been discussed before here: http://www.jboss.org/index.html?module=bb&op=viewtopic&t=144872 So the decision is not to support 'selector' at all. Can you take a look at the discussion and see if there is any feedback on this? I'll be glad to here any thoughts on this. Howard	1	1	https://issues.jboss.org/browse/JBMESSAGING-1416
12068	JBMESSAGING-1416	Howard, Would you elaborate "not to support selector at all" a little more, please? If JBM ignores the Ordering Group when we create consumers with selectors, it is fine for us. If JBM throws exceptions, then we cannot use the queues with JBoss ESB. Let me explain our background using with message selectors. For request-replay message pattern, we should specify connection factory names for a JBoss ESB service. But we can only specify one connection factory name applying both for request and reply queues. In other words, we cannot simply specify a connection factory name only for the reply queue. So we would like to use the queues with selectors even though it breaks FIFO. Could you just ignore the message ordering when using with selectors, as if Ordering Group is not used? Best regards, Fusayuki	1	1	https://issues.jboss.org/browse/JBMESSAGING-1416
12069	JBMESSAGING-1416	Hi Fusayuki, "Would you elaborate "not to support selector at all" a little more, please? If JBM ignores the Ordering Group when we create consumers with selectors, it is fine for us. If JBM throws exceptions, then we cannot use the queues with JBoss ESB." --- If we turn Ordering Group on, it always tries to keep strick ordering of the messages. Using selectors on ordering group will not result in exceptions, but may result in messages never being delivered. "Let me explain our background using with message selectors. For request-replay message pattern, we should specify connection factory names for a JBoss ESB service. But we can only specify one connection factory name applying both for request and reply queues. In other words, we cannot simply specify a connection factory name only for the reply queue. So we would like to use the queues with selectors even though it breaks FIFO. Could you just ignore the message ordering when using with selectors, as if Ordering Group is not used?" --- I don't understand the use case. If ordering can be ignored, why we turn it on? We just use normal delivery. Can you give a very simple concrete case where ordering and selector are both needed and also must be used together? I don't think single connection factory would be a problem, because we can use JBossMessageProducer's API to control ordering, if you don't want ordering all over the connection factory scope.	0	0	https://issues.jboss.org/browse/JBMESSAGING-1416
12070	JBMESSAGING-1415	Can we have more info on this? I can't see any param numberOfRetries in the pasted config....	1	1	https://issues.jboss.org/browse/JBMESSAGING-1415
12071	JBMESSAGING-1415	Ron - if numberOfRetries is no longer used, what is it replaced by? This was quite an important config param for JBM. Thx	1	1	https://issues.jboss.org/browse/JBMESSAGING-1415
12072	JBMESSAGING-1407	And sorry Howard! Related to your svn account I thought Howard was your last name. Greets from Germany Markus	1	1	https://issues.jboss.org/browse/JBMESSAGING-1407
12073	JBMESSAGING-1359	Have you actually tried googling this? http://www.google.co.uk/search?hl=en&q=mssql+not+reclaiming+space&btnG=Google+Search&meta= It brings up lots of hits. Seems this is not an isolated problem.	1	1	https://issues.jboss.org/browse/JBMESSAGING-1359
12074	JBMESSAGING-1359	I'm assuming we have a MS support contract. Can we just use that? If not, we *should* have a MS support contract, since MS is one of our supported platforms ;)	1	1	https://issues.jboss.org/browse/JBMESSAGING-1359
12075	JBMESSAGING-1359	Have you tried "sp_dboption '<your db>', 'trunc. log on chkpt.', 'true';"?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1359
12076	JBMESSAGING-1359	Alex's suggestion doesn't work. It was already tried before. Nobody was able to tell me whether there's some MS support for RedHat. Len do you have any new information on this?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1359
12077	JBMESSAGING-1359	Changed component to doc - Martin - can you open linked JIRAs in EAP and SOA-P projects - also in doc component - so we can get the information into the relnotes? Thanks!	1	1	https://issues.jboss.org/browse/JBMESSAGING-1359
12078	JBMESSAGING-1356	I can take a stab at it. On which branch do you want it implemented?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1356
12079	JBMESSAGING-1356	How's this going? Have you implemented this using the OrderedExecutorFactory? Should just be a few lines of code...	0	0	https://issues.jboss.org/browse/JBMESSAGING-1356
12080	JBMESSAGING-1356	Do you want me to merge? If yes, where?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1356
12081	JBMESSAGING-1356	Could this be made available in 1.4 branch ? This causes scaling issues.	1	0	https://issues.jboss.org/browse/JBMESSAGING-1356
12082	JBMESSAGING-1296	Make sure we support the equivalent of the JBoss Remoting 2.x functionality from JBM 1.4: I.e. we should be able to specify the port the client binds to. secondaryconnectport in jbr 2.x	1	1	https://issues.jboss.org/browse/JBMESSAGING-1296
12083	JBMESSAGING-1296	Are you sure? We're talking about the *client socket* here - a range.	1	1	https://issues.jboss.org/browse/JBMESSAGING-1296
12084	JBMESSAGING-1293	Question - are we seeing a related issue with MS-SQL Server DB - where the limit is 8K?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1293
12085	JBMESSAGING-1255	Jay- Can you make this change on Branch_Stable too?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1255
12086	JBMESSAGING-1254	Can you make this change on Branch_Stable too?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1254
12087	JBMESSAGING-1247	Brian, I did what you suggested (replacing the ObjectName property channelFactory by a property of type org.jgroups.ChannelFactory). However, I still need to figure out 2 things: * how to I inject this new property in JBAPP postgresql-persistence-service.xml? if I replace <attribute name="ChannelFactoryName">jgroups.mux:name=Multiplexer</attribute> by <depends optional-attribute-name="JChannelFactory">jgroups.mux:name=Multiplexer</depends> I still expect to be injected a JChannelFactoryMBean, not a JChannelFactory. * in our own tests, we recreate "manually" a configuration based on the MBean exposed. In that case too, the only thing I can inject is a JChannelFactoryMBean, not a JChannelFactory. Using a MBean instead of an ObjectName is an improvement but it falls short from what you're suggesting. I think I missed something :)	0	0	https://issues.jboss.org/browse/JBMESSAGING-1247
12088	JBMESSAGING-1238	I don't really understand why it should be set to false. If a server is dead, do we really want to be pinging it ad infinitum?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1238
12089	JBMESSAGING-1217	Is this issue already solved ? It is markes as duplicate but can't find the original one. I've the same problem with JBoss ESB Server 4.5 (AS 4.2.3 + JBM 1.4.0.SP3).	0	0	https://issues.jboss.org/browse/JBMESSAGING-1217
12090	JBMESSAGING-1202	when moving to DLQ, original headers are stored in JBOSS_MESSAGING_ORIG_MESSAGE_ID and JBOSS_MESSAGING_ORIG_DESTINATION headers (plus JBOSS_MESSAGING_ACTUAL_EXPIRY_TIME when the message is moved to the expiry queue) in ServerSessionEndpoint.makeCopyForDLQOrExpiry method	0	0	https://issues.jboss.org/browse/JBMESSAGING-1202

#	Requirement	Questions	Answered	Successfully Answered	URL
12091	JBMESSAGING-1189	So can we close this task and the linked JBAS-5077 ?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1189
12092	JBMESSAGING-1131	The web.xml file has a reference to a MBean as follows: <param-value>jboss.remoting:service=invoker,transport=servlet</param-value> Where do I setup this MBean and what Invoker class should be used?	0	0	https://issues.jboss.org/browse/JBMESSAGING-1131
12093	JBMESSAGING-1097	Can you explain what the motivation for this is? If you don't want messages to survive server shutdown/startup then maybe you should be using non persistent messages, not persistent messages?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1097
12094	JBMESSAGING-1094	Do we really need to do this?	1	1	https://issues.jboss.org/browse/JBMESSAGING-1094
12095	JBMESSAGING-999	We've done a lot of failure testing on this recently and have yet to see a problem with reconnection. Can you be more specific about what you think is not working. Also please open a support ticket with your support contract if you want this investigated further.	1	1	https://issues.jboss.org/browse/JBMESSAGING-999
12096	JBMESSAGING-999	Ok, I am confused here. Markust says it was "fixed" in SP1?? - Can you be more specific? Howard- can you verify if the problem occurs in the Branch_1_4 code? If so, please can you write up a diagnosis. *do not* specifiy a solution at this stage. AFAICT the current code should work. if it doesn't we need to know why not - specifics please.	1	1	https://issues.jboss.org/browse/JBMESSAGING-999
12097	JBMESSAGING-999	Markus- >> Compare 3163 to 3272 I'm still confused - which branch are you looking at? Howard- Can you verify if the problem occurs on Branch_1_4? I'm assuming this was the first thing you did, but I guess not..	1	1	https://issues.jboss.org/browse/JBMESSAGING-999
12098	JBMESSAGING-999	Markus. the only differences I can see in BridgeService between 1.4.0.SP3_CP03 and 1.4.0.SP3_CP04 is: --- BridgeService.java (.../JBossMessaging_1_4_0_SP3_CP03/src/main/org/jboss/jms/server/bridge/BridgeService.java) (revision 5536) +++ BridgeService.java (.../JBossMessaging_1_4_0_SP3_CP04/src/main/org/jboss/jms/server/bridge/BridgeService.java) (revision 5536) @@ -351,6 +351,8 @@ destCff= new JNDIConnectionFactoryFactory(targetProps, targetCFRef); } + bridge.setName(serviceName.getCanonicalName()); + bridge.setSourceConnectionFactoryFactory(sourceCff); bridge.setDestConnectionFactoryFactory(destCff); The svn command I used is: svn diff https://svn.jboss.org/repos/messaging/tags/JBossMessaging_1_4_0_SP3_CP03/src/main/org/jboss/jms/server/bridge/BridgeService.java https://svn.jboss.org/repos/messaging/tags/JBossMessaging_1_4_0_SP3_CP04/src/main/org/jboss/jms/server/bridge/BridgeService.java That change is accounted for by this JIRA: https://jira.jboss.org/jira/browse/JBMESSAGING-1407 I don't see any other changes. I guess you are looking in a different place. Can you be more specific as to where.	1	1	https://issues.jboss.org/browse/JBMESSAGING-999
12099	JBMESSAGING-981	I've made the corresponding properties (ServerPeerID, DefaultQueueJNDIContext, DefaultTopicJNDIContext) writable, with a check that ensures they can only be set before the server peer is started. Will this work for you?	1	1	https://issues.jboss.org/browse/JBMESSAGING-981
12100	JBMESSAGING-981	Great! Exactly what i needed, thank you. Should i close the bug or are you going to do it yourself?	0	0	https://issues.jboss.org/browse/JBMESSAGING-981
12101	JBMESSAGING-964	I uncommented the tests but they are of course failing since the functionality they are testing (pre-configured client IDs) is apparently not implemented or not configured. Should I implement it?	1	1	https://issues.jboss.org/browse/JBMESSAGING-964
12102	JBMESSAGING-957	>but I can no longer get to jconsole How do I do that?	0	0	https://issues.jboss.org/browse/JBMESSAGING-957
12103	JBMESSAGING-957	A logical JMS quueue in JBoss Messaging is actually composed of a set of individual queues, one for each node in the cluster that the logical queue is deployed on. Each queue fragment is owned by that node and will contain a different set of messages than other queue fragments. This is why we need a persistent node id, so the node knows which queue fragments in the database it owns so it can load the queue state on startup and add/remove messages. I'm beginning to think the tomcat approach of specifying the node id in a system property is the best. This allows the actual JBM config to remain the same across each node. Do you see any problems with this approach. In particular I'm thinking cluster installation will become more complex since the user will have to manually specify the unique node id on each run.sh - which they currently don't have to do with AS cluster install. Another possibilty would be for JBM to generate a unique id on startup if it can't find one in some config file. Once it's generated it, it would save it locally. This has the difficulty of guarantee uniqueness of the id across the cluster (the id is currently an int, not a string for performance reasons). Any comments?	1	1	https://issues.jboss.org/browse/JBMESSAGING-957
12104	JBMESSAGING-957	Brian - I've been trying to find the place in the tomcat config (AS 4.0.5) that substitutes the node id, but can't find it. Can you point me to the right file? (And place in the file). Thx	1	1	https://issues.jboss.org/browse/JBMESSAGING-957
12105	JBMESSAGING-957	Brian - why would need a int property *and* a string property? JBM just needs an int property. What are these "other uses"?	1	1	https://issues.jboss.org/browse/JBMESSAGING-957
12106	JBMESSAGING-956	Hi Tim, I assume temporary queue is working for clustered queue. We are trying to implement request/response queues. 1. One clustered request queue. 2. One temporary response queue. Currently we have: 1. jboss-4.2.2.GA / jboss-messaging-1.4.0.SP3 / jboss-remoting-2.2.2.SP4 2. PostOffice is clustered. On the requester: 1. The sender creates a temporary queue: tmpQueue = session.createTemporaryQueue(); 2. The sender creates receiver on this temporary queue: session.createReceiver(tmpQueue); 3. The sender sends messages by setting replyto in the message. 4. The sender periodicly retrieve the message using: msg = receiver.receiveNoWait() On the handler: 1. Receive the message. 2. Get the queue from the message: (Queue)msg.getJMSReplyTo() 3. Send the reply back to queue. The issue is, at step 4 above, the msg that the requester got is always null. But I did not see any exceptions in the log. Do you have any ideas? Thanks.	0	0	https://issues.jboss.org/browse/JBMESSAGING-956
12107	JBMESSAGING-942	Do we need to benchmark ActiveMQ and QPid colocated?	0	0	https://issues.jboss.org/browse/JBMESSAGING-942
12108	JBMESSAGING-893	The logic on the Valve is correct: You can't acquire a write lock until all the readLocks are released, and this will happen when: - All threads leave the valve - All thread try to close the valve while they detect a failure. And this is expected behavior... we can't start failover until all threads using a connection fail. I would reject this issue now... but I will talk to you first.	0	0	https://issues.jboss.org/browse/JBMESSAGING-893
12109	JBMESSAGING-806	Jeff - has Trustin fixed the issues in Netty yet? If so, can you close this task?	1	1	https://issues.jboss.org/browse/JBMESSAGING-806
12110	JBMESSAGING-806	the issue is fixed in netty but we have not integrated the fixed version yet (3.1.0.alpha4). I'll close this issue when we integrate this version and I uncomment the corresponding test. is that ok?	0	0	https://issues.jboss.org/browse/JBMESSAGING-806
12111	JBMESSAGING-804	Tim briefly described what we need. That's correct. This issue planned for 1.2.0.SP2 according to this jira entry. And it's marked as "closed". Is the workaround described by Ovidiu the solution? I'd prefer the Tim's proposal as maintenable solution. Is any work planned for it?	1	1	https://issues.jboss.org/browse/JBMESSAGING-804
12112	JBMESSAGING-653	Do we really need this? We can certainly do it if there are requests for it, but for the time being we could start from the assumption that we provide XARecovery from 4.0.5 on. Also see http://jira.jboss.org/jira/browse/JBMESSAGING-655	0	0	https://issues.jboss.org/browse/JBMESSAGING-653
12113	JBMESSAGING-651	Finally tracked down and now appears to be fixed. Redundant Tx callback registration removed from path JBMESSAGING-407. Should now correctly behave and register callbacks while transactions are replayed on the channel.	0	0	https://issues.jboss.org/browse/JBMESSAGING-651

#	Requirement	Questions	Answered	Successfully Answered	URL
12114	JBMESSAGING-640	I don't like the idea of SVN links... If we want to fix that it would be through repository.jboss.com, and we would still have two repositories anyway. Can't we just reject this issue?	0	0	https://issues.jboss.org/browse/JBMESSAGING-640
12115	JBMESSAGING-624	Tomaz, We don't currently have a QA set up to test messaging against DB2. Is it possible for you to checkout the current Branch_1_0, build, and verify that the changes were applied correctly to work against DB2? Thanks, -- Juha	1	1	https://issues.jboss.org/browse/JBMESSAGING-624
12116	JBMESSAGING-607	Can you add new member's name to the page also? :-)	1	1	https://issues.jboss.org/browse/JBMESSAGING-607
12117	JBMESSAGING-601	Ovidiu Could you please set this up today if possible. I would like to commit changes to get the messaging configuration running tomorrow, so I can start testing my refactored test cases tomorrow as well.	1	1	https://issues.jboss.org/browse/JBMESSAGING-601
12118	JBMESSAGING-601	Hi Ovidiu The artifacts did end up in the repository, however, a couple of problems: 1. you forgot to update the component-info.xml file with the new artifacts which were added in version 1.0.1.GA - I added these this morning. If the artifacts aren't listed in component-info.xml, they don't get processed by the buildthirdparty mechanism (if my understanding is correct) By the way, until I saw in fisheye that you made a change to the version number, I thought I was going crazy - I noticed the incorrect CR5 after I had comitted my change (and missed the CR5 mistake), and after you had corrected it, but before it made it back into the repository - the versions were different and I couldn't understand why! 2. I mentioned in the JIRA issue that these new artifacts need to be added for each version of Messaging which we want to test. I should have been more clear on this. The main problem is that it's not so easy to work with two versions of Messaging in the same build/testsuite, which is the situation we have at the moment. At present, the snapshot version of Messaging is being pulled in to the AS build, for Prabhat's TCK testing. This has the effect of setting the jboss.messaging.classpath path element, among other things, which is used in the tests.classpath path element in the testsuite. If you then want to run the testsuite against a different version of Messaging, say 1.0.1.GA, this becomes difficult, as we need to define a different path element (different from jboss.messaging.classpath) and adjust the classpath accordingly. Can we : (i) assume that we will run TCK tests and integration tests against the same version of Messaging (if not in the long term, then at least in the short term) (ii) add the additional artifacts into the snapshot directory in the repository Richard	1	0	https://issues.jboss.org/browse/JBMESSAGING-601
12119	JBMESSAGING-601	I noticed a build problem when trying to use 1.0.1.GA with the AS build earlier this morning (i.e. I changed the version of Messaging from snapshot to 1.0.1.GA in build-thirdparty.xml). Specifically, in the component-info for Messaging, 1.0.1.GA requires serialization 1.0.2.GA or 1.0.3.GA, but serialization 1.0.1.GA is getting pulled in to the build, and the build is failing. What is the approach to resolving these sorts of tangled dependencies (as changing the version of serialization which gets pulled in may upset some other component)? Trial and error? Or should these types of changes only be made with consent of the owner? Richard	1	0	https://issues.jboss.org/browse/JBMESSAGING-601
12120	JBMESSAGING-578	>Also, an important test that we should add is to make sure that leaving a shared JChannel doesn't leave it in a incosistent state, especially if other components share the channel Isn't that integration tests? We are creating new JChannels using Multiplexor MBean. If we think JGroups could fail in any sense, I belive this should be done under integration tests. We will probably need to write new ones.	0	0	https://issues.jboss.org/browse/JBMESSAGING-578
12121	JBMESSAGING-572	Not finished yet, where are the failover stress tests?	1	1	https://issues.jboss.org/browse/JBMESSAGING-572
12122	JBMESSAGING-556	What was your intention here, over and above identifying which tests were JBossMQ specific?	1	1	https://issues.jboss.org/browse/JBMESSAGING-556
12123	JBMESSAGING-556	I had a look for some existing JBossMQ functional tests in the JBossMQ source directory (jbossmq) and didn't see any. Is it the case that the only tests (whether functional, stress, performance or smoke) for JBossMQ were in the AS testsuite - that is, those I have been looking at? Thus, is your intention to consider the JBossMQ-specific tests which were running in the AS testsuite to be the 'functional' tests for JBossMQ and that these should cotinue to be run in the AS testsuite, although now grouped together as 'functional tests for JBossMQ', as there is no separately developed stand-alone test framework for them to run in, as there is for Messaging. We would then have several sets of JMS provider tests in the AS testsuite: 1. generic JMS tests run against Messaging 2. generic JMS tests run against JBossMQ 3. JBossMQ functional tests run against JBossMQ	0	0	https://issues.jboss.org/browse/JBMESSAGING-556
12124	JBMESSAGING-555	I've started to move testcases over, and things are progressing smoothly. A number of the testcases move over with no change - in particular, those test cases which only assume the existence of a ConnectionFactory object in JNDI and which also make use of createTemporaryTopic()/createTemporaryQueue(). However, for the test cases I have encountered so far which make use of named topics and queues, I have found that these are all deployed dynamically by the test cases themselves - my making a call to JBossTestCase.getDeploySetup() and passing test-destinations-service.xml file. In fact, in Adrian's server configurations for JBossMQ, he only deploys one Queue, named DLQ. This is why all test destinations are deployed dynamically. In addition to this, some of the test cases rely on defined usernames and roles. At present, the Messaging configuration does not match this arrangement, and if we are to keep the test cases more or less in their present state (i.e. keep existing deployments as they are, keep existing security related features as they are), we are going to have to rework the Messaging configuration to essentially match Adrian's JBossMQ configuration, in the following ways: (i) ensure that only one Queue named DLQ is created by the AS at startup (ii) create deployment descriptors for starting the test Topics and Queues (done) (iii) add the users and roles which Adrian uses to messaging-roles.properties and messaging-users.properties In order not to interfere with your use of the existing Messaging configuration with your tests (functional, stress, etc), it seems as though we need to add a target to release-admin.xml which builds a Messaging configuration for the AS testing. Does this sound reasonable?	1	1	https://issues.jboss.org/browse/JBMESSAGING-555
12125	JBMESSAGING-555	Should these changes to the contents of the jboss-messaging.sar be performed in release-admin.xml or in your build? They are changes to the files in the sar, and not any changes relating to the default configuration, and so you have complete control over their contents. Changing these files in the sar would require introducing a different build artifact for AS testing as opposed to your functional testing.	1	1	https://issues.jboss.org/browse/JBMESSAGING-555
12126	JBMESSAGING-555	> In order not to interfere with your use of the existing Messaging configuration with your tests (functional, stress, etc), it seems as though we need to add a target to release-admin.xml which builds a Messaging configuration for the AS testing. Does this sound reasonable? I would rather keep it separated. I see it more of a function of the testsuite build script. The release admin does need not care of security roles in JBoss AS integration test suite. I can create the ant configuration script fragment for you to use, if you wish, it shouldn't be any big deal The changes should be applied post-deployment.	1	1	https://issues.jboss.org/browse/JBMESSAGING-555
12127	JBMESSAGING-554	Where are you currently checking in your changes? The 4.0 branch or the head?	1	1	https://issues.jboss.org/browse/JBMESSAGING-554

#	Requirement	Questions	Answered	Successfully Answered	URL
12128	JBMESSAGING-554	Ideally, I (or anyone else) should be able to do this: - check out the AS trunk - build the AS: ./build/build.sh - run the Messaging installation procedure (this is reduntant, Messaging should come as default provider, but for the time being, this is fine): ant -f release-admin.xml - run ./build/build.sh tests-messaging (or whatever target name you went for) and get a full set of JMS integration tests passing - run ./build/build.sh test-jbossmq and get a full set of JMS integration tests passing How close are you to this?	1	1	https://issues.jboss.org/browse/JBMESSAGING-554
12129	JBMESSAGING-554	At present, the JBossMQ tests are initiated in various targets within the testsuite - for example, some of the basic tests are in the tests-standard-unit target, and others in the tests-standard-stress. There is no target for executing all JBossMQ tests as a group. This can be easily changed so that a separate JBossMQ target is defined and this will initiate the JMS tests for JBossMQ. The Messaging tests have already been set up as you describe, except that you don't need to create the Messaging configuration by hand - the target calls release-admin.xml on your behalf. How close am I to that way of executing the tests? Very close. I would just need to create a new target for JBossMQ. That's 10 minutes work, but i'd rather leave it until I finish getting the refactoring done, if you don't mind.	0	0	https://issues.jboss.org/browse/JBMESSAGING-554
12130	JBMESSAGING-523	is there any documentation on how to use this feature in JBM 2.0 ? I couldnt find any. Besides I only see a 2.0 alpha release but this JIRA states it was implemented in 2.0 beta...	0	0	https://issues.jboss.org/browse/JBMESSAGING-523
12131	JBMESSAGING-519	This task can probably be split into two: 1) The client side management of failover. The maintenance of the client side list of servers, the client side load balancing policy. 2) The server side management of failover. This is how I think failover should work: On the server side, if a node fails, this needs to be detected by the group. The failover node for the node that failed needs to then take over responsibility for any queues that the failed node was handling - it effectively "becomes" that node. Any non persistent messages will be lost. Any persistent messsages will be loaded for those queues from the database. If in memory message persistence is being used then no need to load them from the db. Thing to consider: When the failed node comes back - how to handle the responsibility for the queues back to the original node? This probably isn't possible since new connections may have connected to the failover node and we can't just boot them off, or move them to the other node since they will lose non persistent messages. Note that we DO NOT need to replicate delivery state of sessions to the failover node. This will be too slow. Instead we can get transparent failover by the client "replaying" the session to the failed over server. On the client side, when the client detects server failure, it chooses a new server to try. This must match the failover server. Something to think about: How can we ensure that the client doesn't failover to the new server *before* the new server has failed over it's queues? The client starts a new JMS connection, then recreates automatically all the sessions it had on the pervious failed connection. For each session, the client then tells the server the ids of all the delivered but unacknowldged messages it had at point of failure. The server then recreates its delivery lisst from this information. The sessions can then continue as normal - with perhaps redelivery of any non persistent messages that were in the failed sessions. (persistent messages will be ok)	1	1	https://issues.jboss.org/browse/JBMESSAGING-519
12132	JBMESSAGING-471	Can't we just upload them by hand for now? You can make this a feature request for your Maven build, but I don't think we should do this now. Ruel will be following up with you about Maven early next week.	1	1	https://issues.jboss.org/browse/JBMESSAGING-471
12133	JBMESSAGING-367	Is it possible to indicate somewhere how to resolve the issue when using different versions of Messaging and AS. For example it would be good if somewhere it stated that AS 4.0.4.GA could only work with Messaging 1.0.1.CR1 and above. Or could it be posted how to resolve the issue oter than downloading the latest Messaging.	0	0	https://issues.jboss.org/browse/JBMESSAGING-367
12134	JBMESSAGING-237	Some (old) thoughts I thought I'd copy and paste into here for completeness: I just wanted to summarise this since I know this is a complex discussion that could probably do with being put in context. I think this may be useful, especially for people getting up to speed with this subject. I have extended this somewhat to include a quick summary of the lazy loading queues idea which is a related subject, even thought the task at hand (for now anyway) is to design and provide a default implementation of the Message Store not implementation of lazy loading queue functionality, but I think this provides a background to how the message store fits into the whole picture. In order to cope with very large queues/subs (currently all message refs for queues are stored in memory at once), Adrian came up with the idea of lazy loading queues: With lazy loading queues/subs, we have "special references" in each queue / sub. When one of these special marker references are read from the front of the queue, they cause the next n message references to be loaded into the queue followed by another marker. Also the messages corresponding to those message references are themselves are loaded into the message store. (In some cases they may be in the store already since a queue/sub can maintain a reference to the same message). (Aside - actually I think we could extend this idea to lazy load *before* the message is popped off the front, this could be done in a different thread in order to smooth things out. I.e. prevent a big pause to the client while refs/messages are loaded from the db) The value of n for a particular queue can be configured by the user to suit the operational environment. If messages are sent to a queue and there are already more than n messages in that queue, they go straight in the persistent store (even in the case of a non-persistent message!) and no reference is added to the queue. So this means if the value of n is tuned well for a particular queue, and the messages in that queue have a similar size then we should be able to avoid running out of memory in the message store in most cases. :) However, if the message store *does* start running out of memory , which messages should it evict first? It's likely a default implementation would want to evict those messages that have been in the store the shortest amount of time since they are towards the back of the queues (MRU). (Needs to be pluggable) I guess we should also be able to support eviction of messages in batches (i.e. batch updates in the db for performance). We should also know if a message is already in the db or not without making a db hit at passivation. (E.g. persistent messages will already be there and non-persistent may not) When memory usage recovers how do we load messages back into the store? Again this needs to be pluggable. Perhaps we don't bother and just wait until the message is actually accessed, this is simple and perhaps should be our default implementation. Other implementations could get more complex. We could batch load a subset of the same messages we evicted back into the store in order to utilise free memory well. This means we would need to keep track of the ids of the messages we evicted. When loading them back in we may not be able to load them all back in since free memory may not have recovered to the same level it was at before. In order to know how many to load we would need to store some kind of cumulative message size in the db. (SELECT * from message_refs where cumulative_size < xyz). This all gets complicated and it's my feeling which should probably just stick to a simple default impl for now.	1	1	https://issues.jboss.org/browse/JBMESSAGING-237
12135	JBMESSAGING-187	Ovidiu, Is there anything I can do to help you get your content certified on Labs? I'm here to help. Cheers, Damon	1	1	https://issues.jboss.org/browse/JBMESSAGING-187

#	Requirement	Questions	Answered	Successfully Answered	URL
12136	JBMESSAGING-185	Ovidiu, Did you get a chance to take a look at the files i sent you. rgds Rajdeep "Ovidiu Feodorov (JIRA)" <jira-events@jboss.com> wrote: [http://jira.jboss.com/jira/browse/JBMESSAGING-185?page=all] Ovidiu Feodorov updated JBMESSAGING-185: ---------- Fix Version: 1.0.1 (was: 1.0.0) -- This message is automatically generated by JIRA. - If you think it was sent incorrectly contact one of the administrators: http://jira.jboss.com/jira/secure/Administrators.jspa - For more information on JIRA, see: http://www.atlassian.com/software/jira ---------- New Yahoo! Messenger with Voice. Call regular phones from your PC and save big.	1	1	https://issues.jboss.org/browse/JBMESSAGING-185
12137	JBMESSAGING-153	Will this be done in 2.0? /L Redpill	0	0	https://issues.jboss.org/browse/JBMESSAGING-153
12138	JBMESSAGING-134	Feb 08 2006, London design meeting Spillover: ========== Optimisation: Maintain a reference to the message holder in the message reference to prevent having to do a lookup on every getMessage. We will maintain a high memory limit, and a spillover limit. When memory exceeds the high memory limit, messages are evicted down to the low mem Why is there a channel id as part of the primary key in message reference?	0	0	https://issues.jboss.org/browse/JBMESSAGING-134
12139	JBMESSAGING-134	"get ref from db" markers in the queues when loading refs do we load messages too? we only want to load them if a) They are not already loaded When a new message arrives for a queue, we add the reference in the in memory queue if: 1. the in memory queue size is less than the max in memory queue size and 2. there is sufficient free memory otherwise: we write it to disk keep it simple for 1.0 Case 1: A new message arrives on a queue: 1. The message is referenced, if there is sufficient free memory, the message is store in memory, otherwise it goes immediately to disc (even for non-persistent) 2. If the there is not enough free memory to store the reference, OR, the queue in memory size is greater than maximum in memory size then, a marker reference is written into the queue if there is not one there already, and the reference is written to the database. Case 2: Receiving a message 1. If there is an available ref in the queue the message for that is returned from in memory (it will always be there) 2. If the next ref is a marker ref, then the next x refs are retrieved from the database. x is determined by the queue batch load size, the message corresponding to the message refs are loaded at the same time. However there may not be enough memory to load all the messages or refs, so as messages are loaded from the db (must use a cursor for this) we only create them if there is enough free memory. Case 3: Memory runs low. Memory runs low and the y last messages to be refed are determined to be removed from the in memory store. Since multiple refs can refer to the same message, this isn't necessarily the last refs to to be added. (Important!) We also need to remove and replace with a marker the refs corresponding to the mesages that have been removed. This is done by maintaining in the holder references to the message refs that reference the message. Other method: A new message arrives on queue: 1. The message is referenced, if there is sufficient free memory, the message is store in memory, otherwise it goes immediately to disc (even for non-persistent) 2. If the queue size is greater than the max size, and a marker hasn't already been written then we write a marker and write the ref to disc, otherwise we add the ref to the in memory queue. Receiving a message: 1. If there is an available ref, then try to retrieve the message from memory, if it is not there then retrieve a batch of messages from the db. Problem! How many and which mesages to retrieve? If the server is just started up this is very difficult to calculate. Problem is when the server is started how do we determine which messages to load into memory? possible issue with storing deliveries in separate structure due to passivating message references to disk then on reload, the delivery is cancelled and ends up in the wrong place in the queue remove message reference copy constructor and replace with explicit method on message store that refs given another message reference. Message store takes the decision when memory is low. Each object whose memory size can be measure must implement the interface MemoryMeasurable with operation getBytesSize() The message store or whatever also implements the interface MemoryManager Queue Memory management ---------- - Case 1: Queue Starts spilling over. Memory Manager knows when memory is low so calls message store to ask it to deal with the situation. It would tell message store to reduce memory consumption by 1230000 bytes. MessageStore implements a MRU cache, which is updated on every call to ref to create a new message reference (make sure to remove messagereference copy constructor and replace with an explicit ref(MessageReference other) method on the message store. This means the most recently touched messages in the cache are those that correspond to message references on the ends of queues. The message store is aware of the overall size of the message and any message references that reference it so it knows how many message (and corresponding message references) to evict to bring memory down to the required level. The holder object also needs to keep track of the channel id of any channels that reference the message. When the message is evicted it informs the channel to evict one item from the end of the channel for any channels referencing the message. (This will always be at the end of the channel since the messages are	0	0	https://issues.jboss.org/browse/JBMESSAGING-134
12140	JBMESSAGING-58	Is there a reason why Tim doesn't have cvs access? Also, there is no "remoting" module in jboss-head anymore, so this is obviously a patch against an old version?!?	1	1	https://issues.jboss.org/browse/JBMESSAGING-58
12141	JBMESSAGING-26	JEMS installer currently has a "JMS server" profile which is based on jbossmq. It currently includes jmx microkernel, logging, hsqldb, jndi, jaas, jta, and jbossmq services. So my understanding is that we need to add "JBoss Messaging" profile to the installer. Since user can only select one profile, JBoss Messaging and current jbossmq-based "JMS Server" will be mutually exclusive. Is there a need to have a standalone installer for JBoss Messaging or is it ok to just add a "Standalone JBoss Messaging" profile to the JEMS installer? Which version of JBoss Messaging?	1	0	https://issues.jboss.org/browse/JBMESSAGING-26
12142	JBMESSAGING-26	Needs to be further discussed. Do we really want to support a standalone configuration?	0	0	https://issues.jboss.org/browse/JBMESSAGING-26
12143	JBMESSAGING-25	Initial code drop for pure MC based messaging. Contributed by Arvinder Singh [15/May/06 05:05 PM] Originally attached to JBMESSAGING-26. Comments following attachment: Tim Fox [20/Jun/06 01:47 PM]: I'm not sure what I am looking at here. I don't know much about the MC but it is my understanding that all we need to do is provide POJO versions of our current MBeans. Am I missing something? Arvinder Singh [21/Jun/06 03:39 AM]: Ok it might become clearer with the next set of changes i will upload. Basically I am adding code that belongs in other projects like - NamingService - TransactionService - RemotingService - etc in addition to the MBeans in messaging which I have yet to do. The issue is not providing POJO versions of current messaging MBeans, but other POJO versions from the other sub projects (like above) in order to get this up and running. Let me upload the next set and then we can go over what has been done.	1	1	https://issues.jboss.org/browse/JBMESSAGING-25

#	Requirement	Questions	Answered	Successfully Answered	URL
12144	JBMETA-391	from resource: web-app_2_5.xsd 2015-11-24 04:13:12,406 WARN [org.jboss.metadata.parser.util.XMLResourceResolver] (MSC service thread 1-3) Cannot load publicId from resource: web-app_2_5.xsd 2015-11-24 04:13:13,106 ERROR [org.jboss.metadata.parser.util.XMLSchemaValidator] (MSC service thread 1-3) Cannot get schema for location: http://java.sun.com/xml/ns/javaee/web-app_2_5.xsd: org.xml.sax.SAXParseException; systemId: http://java.sun.com/xml/ns/javaee/javaee_5.xsd; lineNumber: 83; columnNumber: 38; sch-props-correct.2: A schema cannot contain two global components with the same name; this schema contains two occurrences of 'http://java.sun.com/xml/ns/javaee,descriptionGroup'. at org.apache.xerces.util.ErrorHandlerWrapper.createSAXParseException(ErrorHandlerWrapper.java:201) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.util.ErrorHandlerWrapper.error(ErrorHandlerWrapper.java:132) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.XMLErrorReporter.reportError(XMLErrorReporter.java:394) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.xs.traversers.XSDHandler.reportSchemaError(XSDHandler.java:4093) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.xs.traversers.XSDHandler.reportSchemaError(XSDHandler.java:4088) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.xs.traversers.XSDHandler.checkForDuplicateNames(XSDHandler.java:3746) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.xs.traversers.XSDHandler.buildGlobalNameRegistries(XSDHandler.java:1315) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.xs.traversers.XSDHandler.parseSchema(XSDHandler.java:610) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.xs.XMLSchemaLoader.loadSchema(XMLSchemaLoader.java:580) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.xs.XMLSchemaLoader.loadGrammar(XMLSchemaLoader.java:547) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.impl.xs.XMLSchemaLoader.loadGrammar(XMLSchemaLoader.java:513) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at org.apache.xerces.jaxp.validation.XMLSchemaFactory.newSchema(XMLSchemaFactory.java:233) [xercesImpl-2.11.0.SP4.jar:2.11.0.SP4] at javax.xml.validation.SchemaFactory.newSchema(SchemaFactory.java:638) [rt.jar:1.8.0_60] at __redirected.__SchemaFactory.newSchema(__SchemaFactory.java:167) [jboss-modules.jar:1.4.4.Final] at org.jboss.metadata.parser.util.XMLSchemaValidator.getSchemaForLocation(XMLSchemaValidator.java:117) at org.jboss.metadata.parser.util.XMLSchemaValidator.validate(XMLSchemaValidator.java:85) at org.wildfly.extension.undertow.deployment.WebParsingDeploymentProcessor.deploy(WebParsingDeploymentProcessor.java:104) at org.jboss.as.server.deployment.DeploymentUnitPhaseService.start(DeploymentUnitPhaseService.java:147) [wildfly-server-2.0.2.Final.jar:2.0.2.Final] at org.jboss.msc.service.ServiceControllerImpl$StartTask.startService(ServiceControllerImpl.java:1948) [jboss-msc-1.2.6.Final.jar:1.2.6.Final] at org.jboss.msc.service.ServiceControllerImpl$StartTask.run(ServiceControllerImpl.java:1881) [jboss-msc-1.2.6.Final.jar:1.2.6.Final] at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) [rt.jar:1.8.0_60] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) [rt.jar:1.8.0_60] at java.lang.Thread.run(Thread.java:745) [rt.jar:1.8.0_60] 2015-11-24 04:13:13,115 WARN [org.jboss.metadata.parser.util.XMLSchemaValidator] (MSC service thread 1-3) Cannot get schema for location:	0	0	https://issues.jboss.org/browse/JBMETA-391
12145	JBMETA-366	I am not sure it is OK you remove the classname logic and replace it by interface but the javaee_7.xsd seems to accept both, did I miss something?	1	1	https://issues.jboss.org/browse/JBMETA-366
12146	JBMETA-258	John, i see that jboss-metadata-common 2.0.0-alpha-14 has been released. Is this JIRA still relevant?	1	1	https://issues.jboss.org/browse/JBMETA-258
12147	JBMETA-223	What do we need to do to fix in AS trunk?	1	0	https://issues.jboss.org/browse/JBMETA-223
12148	JBMETA-171	Meaningful difference is between the ordering of proxy-factory-config: Branch_4_2 <!ELEMENT proxy-factory-config (client-interceptors? , activation-config? , endpoint-interceptors? , web-class-loader? , JMSProviderAdapterJNDI? , ServerSessionPoolFactoryJNDI? , metadata (the same in both jboss_4_2.dtd and jboss_5_0.dtd) <!ELEMENT proxy-factory-config (client-interceptors?, endpoint-interceptors? , web-class-loader? , activation-config?, JMSProviderAdapterJNDI? , ServerSessionPoolFactoryJNDI? , I don't know for which reason. In Java proxy-factory-config is represented as org.w3c.dom.Element (so, no ordering is defined in Java binding). I am thinking of changing it back to the ordering as in Branch_4_2. Other changes are added comments for the child elements of cluster-config.	0	0	https://issues.jboss.org/browse/JBMETA-171
12149	JBMETA-146	Alexey: In implementing this change I fixed up the jboss.xsd schema and put in some test coverage, but I'm not sure what else you need (more tests, DTD?) to validate all is working as expected. Additionally I have not added any tests for XML-based input. Could you fill in the gaps?	1	1	https://issues.jboss.org/browse/JBMETA-146
12150	JBMETA-60	If I remember correctly the javadoc was generated at the very last "mvn release:perform" step. If it doesn't produce the javadocs now, then we are fine and you can close the issue. Thanks Paul.	0	0	https://issues.jboss.org/browse/JBMETA-60
12151	JBMETA-50	Does this also include AbstractEJBEverythingTest protected MessageDrivenBeanMetaData assertFullMDB(String ejbName, IEnterpriseBeansMetaData<?, ?, ?, ?> enterpriseBeansMetaData, Mode mode) { throw new NotImplementedException(); } What I am currently missing out is ejb local refs defined on message-driven beans... Does this JIRA take care of it or should I create a new one?	1	1	https://issues.jboss.org/browse/JBMETA-50
12152	JBMETA-33	I don't see any dependencies specified in the parent pom. I guess, it should be the version range. Can you do it?	0	0	https://issues.jboss.org/browse/JBMETA-33
12153	JBMDR-60	Did this happen yet?	1	1	https://issues.jboss.org/browse/JBMDR-60
12154	JBMDR-60	I upgraded to 2.2.16.GA. btw: how "big" of a change is your change? e.g. can it be ported to 2.0.x branch w/o breaking any api?	1	1	https://issues.jboss.org/browse/JBMDR-60
12155	JBMDR-60	I avoided breaking any api as recommended on the design forums. It shouldn't hurt if this change gets merged back as part of another change. However, I think there is more work left to do to get full benefit from this ScopeKey change. I don't think the ScopeKey change is that big. There are still places where I put an OptimizedKey (return value from ScopeKey.getOptimizedKey()) as the Map search key but we later perform a Map search with a plain ScopeKey instead of an OptimizedKey. This causes ScopeKey.equals to be invoked with an OptimizedScopeKey object (or maybe the other way around). This is slower than if we used OptimizedScopeKey.equals with an OptimizedScopeKey object. My plan is to find all of the places that cause the mixed OptimizedScopeKey + ScopeKey based searchs to occur during AS boot and attempt fixing them to pass an OptimizedScopeKey. I got blocked on this effort before when I couldn't figure out how to fix this jira on my own. When we start doing parallel deployment, the performance impact of this ScopeKey change will be even more dramatic (since we avoid the need to synchronize on the Map). Do you need a new jira for moving MDR + MC Kernel to AS trunk?	0	0	https://issues.jboss.org/browse/JBMDR-60
12156	JBMICROCONT-329	Where do you see this usage? The only place I was able to find it was in mc\kernel\pom-jdk14.xml, which I think we don't keep it up to date.	0	0	https://issues.jboss.org/browse/JBMICROCONT-329
12157	JBMICROCONT-261	Alexey, Can you please validate the schema at https://svn.jboss.org/repos/jbossas/projects/microcontainer/trunk/aop-mc-int/src/resources/main/schema/aop-beans_1_0.xsd? The included schema lives at https://svn.jboss.org/repos/jbossas/projects/microcontainer/trunk/kernel/src/resources/main/schema/jboss-beans-common_2_0.xsd	1	1	https://issues.jboss.org/browse/JBMICROCONT-261

#	Requirement	Questions	Answered	Successfully Answered	URL
12158	JBMICROCONT-259	Alexey, Can you please validate the schema at https://svn.jboss.org/repos/jbossas/projects/microcontainer/trunk/kernel/src/resources/main/schema/temp-bean-deployer_2_0.xsd? If valid, please assign this issue back to me so I can update the main bean-deployer_2_0.xsd before closing this issue. Thanks	1	1	https://issues.jboss.org/browse/JBMICROCONT-259
12159	JBMICROCONT-178	Hi Scott, how are we looking on this task? The work is done and there is just more testing I believe. http://www.jboss.com/index.html?module=bb&op=viewtopic&t=112436&postdays=0&postorder=asc&start=0	0	0	https://issues.jboss.org/browse/JBMICROCONT-178
12160	JBMICROCONT-107	Hi Adrian Will it mkes sense to keep this one open? And leave this to the deployers only?	1	0	https://issues.jboss.org/browse/JBMICROCONT-107
12161	JBMICROCONT-93	Hi Paul Could you confirm that this issue can be closed? Cheers	0	0	https://issues.jboss.org/browse/JBMICROCONT-93
12162	JBMICROCONT-55	Adrian, Is there anything I can do to help you get your content certfied on Labs? Any instructions missing or unclear? Let me know. I'm here to help. Cheers, Damon	1	1	https://issues.jboss.org/browse/JBMICROCONT-55
12163	JBMICROCONT-41	Try it again now and see what you think?	0	0	https://issues.jboss.org/browse/JBMICROCONT-41
12164	JBMICROCONT-40	can use generics in xsd ? <parameter <?> > Some value </parameter> * policy - method or constructor is overridden * resolve rule - is the best match using the type and then its supertypes <parameter <T>="java.lang.Integer"> Some value< /parameter> * policy - isn't overridden must be clear setting * throw an error when it is ambigous	1	1	https://issues.jboss.org/browse/JBMICROCONT-40
12165	JBMICROCONT-25	Should we proceed even further with cleaning - even more into direction of JoinPointFactory api?	0	0	https://issues.jboss.org/browse/JBMICROCONT-25
12166	JBMICROCONT-5	I'm adding a comment here, because the forums are down.. Lets say we have: a.bean |-b.bean |-c.jar And just 3 deployers, registered in that priority order: (C)lassLoadingDeployer -> (B)eanDeployer -> (J)ARDeployer. After the deployment structure is analyzed (by the bean & jar deployers) we have the corresponding DeploymentContext/subcontexts in place with each one having recorded its own url as a classpath. How 2nd stage deployment processing would work in this case, wrt depth and priority (.jar comes before .bean, according to suffix rules) C(c.jar) C(b.bean) C(a.bean) B(c.jar) B(b.bean) B(a.bean) J(c.jar) J(b.bean) J(a.bean) Or C(c.jar) B(c.jar) J(c.jar) C(b.bean) B(b.bean) J(b.bean) C(a.bean) B(a.bean) J(a.bean) Also, in the 1st case, should the MainDeployer control the traversal, or simply pass the base context to each deployer, and he traverses at will? Finally, what would the ClassLoadingDeployer do? Judging from the 3.x/4.x mechanism it seems to be creating a UCL for the base node, then adding the classpath of all subnodes to this UCL. (but processing here is parent first). Should I just go to the top most deployment, create a plain URLClassLoader, add all classpaths, and point to it from subcontexts?	1	1	https://issues.jboss.org/browse/JBMICROCONT-5
12167	JBMICROCONT-4	I've been experimenting with wrapping apache-vfs as a plug-in to jboss-vdf (virtual deployment framework?) As a starting point I made some packages under: org.jboss.vdf.spi org.jboss.vdf.plugins.vfs and I'm trying to re-write a deployment scanner over this, slowly adding the features required. It still needs a lot of work and testing but the result seems interesting, in that we could have a single scanner potentially scanning different filesystem types over the VFS plug-in by simply supplying different URIs, e.g.: file:/c:/jboss/deploy webdav:http://remote.server/deploy smb: ... (samba filesystem) (you could also have a different VDF plug-in, or plug-ins to the VFS plugin...) Then, I'd probably try to re-write one of the deployers using this, assuming this is the only way to come up with a working VDF api. I'm wondering about the proper location for checking in this. Currently I need to have it in HEAD/system to drive the development using a working server. I just chose package names that will presumably move easily to the HEAD/kernel module at a later time. Is there a better location for this code?	0	0	https://issues.jboss.org/browse/JBMICROCONT-4
12168	JBMICROCONT-4	Can we put discussions in the forums and reserve JIRA for conclusions?	1	1	https://issues.jboss.org/browse/JBMICROCONT-4
12169	JBMCTOOL-2	What / where is JBossCommon Graph?	0	0	https://issues.jboss.org/browse/JBMCTOOL-2
12170	MODULES-256	[~dmlloyd] Does that mean people will still be able to use the old way of declaring modules - using the slot attribute, if they stay at the 1.5 version of XSD namespace? Will that work at runtime?	1	0	https://issues.jboss.org/browse/MODULES-256
12171	MODULES-236	[~dmlloyd] Can we use org.jboss.dmr.ValueExpressionResolver instead of the JDK code? If so, can we add a dependency from jboss-modules to org.jboss.dmr, or do you want to keep it pure, and just copy in the ValueExpressionResolver minus the code we don't need?	1	1	https://issues.jboss.org/browse/MODULES-236
12172	MODULES-236	BTW, if you want to see the implementation I have so far, please check here: https://github.com/jasinner/jboss-modules/tree/MODULES-236 I'm not really sure how to make this classes private? Can you provide some details on that?	1	1	https://issues.jboss.org/browse/MODULES-236
12173	MODULES-236	Thanks for your feedback. I made your suggested modifications in my fork. The 'jboss.home.dir' System Property is being passed to the java executable after the -jar argument, so it's not being set a System Property till later. If we change the standalone.sh as follows, using '${jboss.home.dir} in modules.xml works as expected. {code} diff --git a/bin/standalone.sh b/bin/standalone.sh index 5b9e0f8..6cee397 100755 --- a/bin/standalone.sh +++ b/bin/standalone.sh @@ -301,12 +301,12 @@ while true; do eval \"$JAVA\" -D\"[Standalone]\" $JAVA_OPTS \ \"-Dorg.jboss.boot.log.file="$JBOSS_LOG_DIR"/server.log\" \ \"-Dlogging.configuration=file:"$JBOSS_CONFIG_DIR"/logging.properties\" \ + -Djboss.home.dir=\""$JBOSS_HOME"\" \ + -Djboss.server.base.dir=\""$JBOSS_BASE_DIR"\" \ -jar \""$JBOSS_HOME"/jboss-modules.jar\" \ $MODULE_OPTS \ -mp \""${JBOSS_MODULEPATH}"\" \ org.jboss.as.standalone \ - -Djboss.home.dir=\""$JBOSS_HOME"\" \ - -Djboss.server.base.dir=\""$JBOSS_BASE_DIR"\" \ "$SERVER_OPTS" JBOSS_STATUS=$? else {code} Do you think we should raise this as an enhancement? If so, which project should we raise it under? If you've got a better idea of how to handle this, please let me know.	1	1	https://issues.jboss.org/browse/MODULES-236
12174	MODULES-226	Is this a system property that Maven also uses, or is this a property that we added just for JBoss Modules?	0	0	https://issues.jboss.org/browse/MODULES-226
12175	MODULES-219	What's the use case? Is it just to support these tools (i.e. is there a problem with using them that is being solved)? I wonder what the effect of URL-decoding these class paths might have...	0	0	https://issues.jboss.org/browse/MODULES-219
12176	MODULES-204	In my latest test I got the CNFE from `ModuleClassLoader.findClass#205`, and I already have a subclass of `ModuleClassLoader` but that method is final. Just adding a `protected String classNotFoundExceptionMessage(String className, Module currentModule){}` would be enough to override the message. As for the part where I traverse dependencies in order to find that class file (without having to load it, a simple resource-type lookup would work), well I guess I'd have to ask you how to do it ;)	0	0	https://issues.jboss.org/browse/MODULES-204
12177	MODULES-199	At this point the WildFly / EAP codebase constitutes the 'module repository' for said distro. Thus there is no need for a 'central' repo. What I see is an common layer between several projects (/ products), but without any means to govern that common layer. Why is that common layer not being created?	1	1	https://issues.jboss.org/browse/MODULES-199

#	Requirement	Questions	Answered	Successfully Answered	URL
12178	MODULES-199	[~wolfc] That common layer isn't being created because lack of tooling support in the jboss module ecosystem (or lack of documentation/guidelines if that tooling does exists). But that wouldn't solve the real issues: - Suppose EAP 6 doesn't use jfreechart. Suppose BRMS 6 does and it creates a module for jfreechart. Now suppose EAP 7 does use jfreechart => it will create a duplicate (but different) jfreechart module. What is BRMS 7 to do? Why is EAP reinventing the wheel and wasting manpower? - Suppose EAP 6 doesn't use freemarker. Suppose BRMS 6 does and DataGrid 6 does use it. BRMS 6 and DataGrid 6 are siblings, neither includes the other. So both of them create their own jfreechart module (and they choose a different module identifier). When a customer than wants to combine BRMS 6 and DataGrid 6 in the same module repository, that's a problem. We can theoretically solve this problems by "more communication" between the groups, but in practice there is already too much communication/spam. Instead of optional communication, we need enforced module id uniqueness and support to avoid duplicate work, which boils down to a "module search" function.	1	0	https://issues.jboss.org/browse/MODULES-199
12179	MODULES-198	How does this relate to http://www.smartics.eu/smartics-jboss-modules-maven-plugin/index.html ?	1	0	https://issues.jboss.org/browse/MODULES-198
12180	MODULES-158	Currently clients cannot specify file patterns directly. This could easily be added, what do you think? i.e. null == '*'	1	1	https://issues.jboss.org/browse/MODULES-158
12181	MODULES-153	So here's my idea. Would it be possible if I give jboss-modules in a non-running server 1) a modules directory, and 2) a module string (such as org.jboss.as.web), it would return me a basic model or simply a list of best-guess jars? We would be able to provide overrides, for example "com.sun.jsf-impl, slot 1.2" I could read / parse the standalone.xml myself to get the list of top-level modules we want. I forsee something like this, but this is just an example. {code} String[] modulesINeed = new String[]{"org.jboss.as.web" }; File[] allJars = jbossModulesGetJarRecurse(modulesINeed); OR String[] modulesINeed = new String[]{"org.jboss.as.web" }; String[] overrideId = new String[] {"com.sun.jsf-impl" }; String[] overrideSlot = new String[] { "1.2" }; File[] allJars = jbossModulesGetJarRecurse(modulesINeed, overrideId, overrideSlot); {code} We could also have overrideSlot be some type of osgi version string with ranges, if possible. As a return value, we would expect a list of all dependencies (recursive included) in a best-guess manner. Instead of File[], we could return some object array which would include the File, and information on its restrictions, whether the jar itself is public / private, or if public / private is enhanced to be per-package instead of per jar, some list of restricted packages. In this way, it would be up to tools to get the list of top module names, either by asking the user directly in our UI or parsing the standalone.xml ourselves. jboss-modules job is a lot easier, a lot less busy, AS itself requires no additions to api, and tools still maintains some responsibility at the top level. And best of all, a started server is not required. Does this sound feasible? What work would need to be done? Is it possible I can help? ;) (This message edited for formatting and additional info)	0	0	https://issues.jboss.org/browse/MODULES-153
12182	MODULES-125	One requirement which comes from Jigsaw and elsewhere is that we assume that we are loading modules from a filesystem that we do not have write access to. This would certainly violate that. Also note that the distribution of SOs and JARs is the responsibility of the *packager*; thus run time isn't really the appropriate time to do this anyway. There is no reason why these files can't be installed expanded at installation time that I can see. Is there a use case for per-platform JARs outside of native libraries?	1	1	https://issues.jboss.org/browse/MODULES-125
12183	MODULES-125	Would a fallback to *tmp* filesystem in case of readonly location satisfy that need? (Think that jigsaw recommends that, although they say this yet has to be defined) Run time expansion allows to have the same content with expended one depending on the configuration settings or run time environment. Currently it is possible to have only one instance of the modules/org/jboss/as/web/main/lib/os/arch/resource.name Using dynamic expansion that would be possible by choosing the correct version inside configuration eg. moudules/.../main contains resource-1.jar resource-2.jar resource-n.jar and inside module.xml user can choose the correct version <resource-root name="resource-2.jar"/> I'm not aware how else modules can allow to package multiple versions for the same platform which can be chosen by configuration settings.	1	1	https://issues.jboss.org/browse/MODULES-125
12184	MODULES-125	What is the use case for packaging multiple versions of a library for the same platform which can be chosen by configuration?	1	1	https://issues.jboss.org/browse/MODULES-125
12185	MODULES-125	So it looks like you're aiming for two new requirements for Modules: * Support the ability to distribute native libraries inside of a JAR (presumably for the purposes of simplified installation) * Allow the user to pick between many versions of a native library (or maybe any resource root) with a configuration change at run time Is this correct?	1	1	https://issues.jboss.org/browse/MODULES-125
12186	MODULES-125	The question remains - why do we need this functionality? As far as ease of installation, this is a problem better solved at a different layer. And why do we need the ability to change versions of *only* native libraries at run-time?	1	1	https://issues.jboss.org/browse/MODULES-125
12187	MODULES-117	I'm going to assume here that what we are really talking about is the AS launch and not JBoss Modules in some generic sense. In a generic sense, I understand why wanting jboss-modules Main to support a -P param which it the processes the way the AS does sounds nice. Let users aggregate all their sys props in one file instead of requiring a bunch of -D stuff. But what's the practical effect in the AS case? Here's what we pass on the command line now when we start a standalone server, excluding JVM stuff like -Xmx etc: 1) JDK properties -Djava.net.preferIPv4Stack=true -Djava.awt.headless=true 2) JBoss Modules properties -Djboss.modules.system.pkgs=$JBOSS_MODULES_SYSTEM_PKGS -mp \"${JBOSS_MODULEPATH} -jaxpmodule "javax.xml.jaxp-provider" 3) AS properties -Dorg.jboss.resolver.warning=true -Djboss.server.default.config=standalone.xml // standalone -Djboss.domain.default.config=domain.xml // domain -Djboss.host.default.config=host.xml // domain -Dorg.jboss.server.log.file=$JBOSS_LOG_DIR/server.log -Dlogging.configuration=file:$JBOSS_CONFIG_DIR/logging.properties -Djboss.home.dir=$JBOSS_HOME -Djboss.server.base.dir=\"$JBOSS_BASE_DIR The -jaxpmodule "javax.xml.jaxp-provider" one we want to get rid of anyway. Looking at the others, in general JDK ones are unsafe to put in a .properties file and set via -P handling, because that -P handling might execute after the JDK has already read the property. A number of other ones don't have static values. They have values that are set by the script. So there is no way to just put them in a static .properties file. So, a .properties file could include -Dorg.jboss.resolver.warning=true -Djboss.server.default.config=standalone.xml -Djboss.domain.default.config=domain.xml -Djboss.host.default.config=host.xml Those last 3 are really examples anyway, just illustrating how the user could configure the default config file. I don't see a lot of value in separating those out from the existing scripts into a separate file. Tools would still need to understand how to deal with the other ones.	0	0	https://issues.jboss.org/browse/MODULES-117
12188	MODULES-69	Please see the commit at: https://github.com/jbossas/jboss-modules/commit/aaf4017287598acb This provides a basic mechanism for filtering classes and resources. I'm thinking that as a next step, it would make sense to add optional parameters to dependency construction for import/export filters for classes and resources. Would this satisfy the requirements of the OSGi implementation?	1	1	https://issues.jboss.org/browse/MODULES-69
12189	MODULES-60	Do you have a specific reason for this request?	1	1	https://issues.jboss.org/browse/MODULES-60

#	Requirement	Questions	Answered	Successfully Answered	URL
12190	MODULES-22	Are there remaining issues which prevent this blanket JIRA from being closed? If so, can they be created as subtasks or dependent tasks?	1	1	https://issues.jboss.org/browse/MODULES-22
12191	MODULES-20	Ok, I see if I can switch to the manifest metadata for the OSGi/Modules dependency and service integration (XServices) tests. My first requirement would be to use a notation that clearly demarcates these headers as JBoss proprietary so they cannot get confused with existing (i.e. OSGi) and future standards. So it should be JBoss-Module-Name: instead of Module-Name Learning OSGi and its associated headers is hard enough without this possible confusion.	1	1	https://issues.jboss.org/browse/MODULES-20
12192	MODULES-20	David L: are you thinking about additions to the actual MANIFEST.MF or just a file that follows the same format with a different name? In the case of a different file I would recommend using a Properties file instead of the actual manifest file if you want to move away from XML as properties don't have the silly ~76 char line limitation and other manifest idiosyncrasies.	1	1	https://issues.jboss.org/browse/MODULES-20
12193	MODULES-13	You should be able to do this already using DependencySpec.exportExcludes/exportIncludes (yeah I know they're called export* but it's supposed to control the exports of the dependency and thus your imports). Does this not work?	0	0	https://issues.jboss.org/browse/MODULES-13
12194	MODULES-7	Looks good... I think... but can you give a URL that is based on upstream including just this change? I'm not sure how to do that though :)	1	1	https://issues.jboss.org/browse/MODULES-7
12195	MODULES-3	I created a small maven project that 1. wraps the class files from jboss-modules alongside a very small and simple module containing a main class into a jar (effectively creating a JAR-based repository), 2. registers org.jboss.modules.Main as that jar's main class in MANIFEST.MF, 3. registers my own org.jboss.modules.ClassPathModuleLoader in META-INF/services/org.jboss.modules.ModuleLoader (thus telling jboss-modules to use this as its default ModuleLoader-implementation), and 4. runs a simple integrations test that executes "java -jar my-jar-repo.jar my.module" via java.lang.ProcessBuilder and verifies that the Module identified by "my.module" is indeed loaded and started. Since this test passes I herewith proudly declare that loading modules from a single JAR basically works. There is obviously room for improvement (e.g. logging could be improved, the repository root within the jar could be made configurable via a manifest entry and so forth), yet all in all the present solution seems to be basically usable. Where do we go from here?	0	0	https://issues.jboss.org/browse/MODULES-3
12196	MODULES-2	It will be better to reboot this discussion with a question: What additional features is ModuleClassLoader missing that you need?	0	0	https://issues.jboss.org/browse/MODULES-2
12197	JBNAME-42	Exception { if (useAbsoluteName(env)) name = getAbsoluteName(name); final Object obtained = NamingManager.getObjectInstance(obj, name, this, env); if(obtained instanceof Reference) // <=== problem { final Reference ref = (Reference) obtained; throw MissingObjectFactoryException.create(ref.getFactoryClassName(), name); } return obtained; } In particular, org.jboss.ejb.Container sets org.jboss.naming.ENCThreadLocalKey as the factory for non-local references, and ENCThreadLocalKey.getObjectInstance() returns a new LinkRef(target). But LinkRef is a subclass of Reference, so getObjectInstance() throws a MissingObjectFactoryException even if the reference was found. This problem shows up in org.jboss.test.naming.test.EjbLinkUnitTestCase.testEjbNoLink(), which throws 2010-02-26 15:05:15,797 INFO [org.jboss.test.naming.ejb.TestEjbLinkBean] (WorkerThread#0[127.0.0.1:58696]) failed: javax.naming.NamingException: Could not dereference object [Root exception is java.lang.IllegalArgumentException: missing object factory class name is required] at org.jnp.interfaces.NamingContext.getObjectInstanceWrapFailure(NamingContext.java:1591) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:900) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:911) at org.jnp.interfaces.NamingContext.lookup(NamingContext.java:749) at javax.naming.InitialContext.lookup(InitialContext.java:392) at org.jboss.test.naming.ejb.TestEjbLinkBean.testEjbLinkCaller(TestEjbLinkBean.java:75) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.invocation.Invocation.performCall(Invocation.java:386) at org.jboss.ejb.StatelessSessionContainer$ContainerInterceptor.invoke(StatelessSessionContainer.java:233) at org.jboss.resource.connectionmanager.CachedConnectionInterceptor.invoke(CachedConnectionInterceptor.java:156) at org.jboss.ejb.plugins.StatelessSessionInstanceInterceptor.invoke(StatelessSessionInstanceInterceptor.java:173) at org.jboss.ejb.plugins.CallValidationInterceptor.invoke(CallValidationInterceptor.java:63) at org.jboss.ejb.plugins.AbstractTxInterceptor.invokeNext(AbstractTxInterceptor.java:121) at org.jboss.ejb.plugins.TxInterceptorCMT.runWithTransactions(TxInterceptorCMT.java:350) at org.jboss.ejb.plugins.TxInterceptorCMT.invoke(TxInterceptorCMT.java:181) at org.jboss.ejb.plugins.SecurityInterceptor.process(SecurityInterceptor.java:228) at org.jboss.ejb.plugins.SecurityInterceptor.invoke(SecurityInterceptor.java:211) at org.jboss.ejb.plugins.security.PreSecurityInterceptor.process(PreSecurityInterceptor.java:97) at org.jboss.ejb.plugins.security.PreSecurityInterceptor.invoke(PreSecurityInterceptor.java:81) at org.jboss.ejb.plugins.LogInterceptor.invoke(LogInterceptor.java:205) at org.jboss.ejb.plugins.ProxyFactoryFinderInterceptor.invoke(ProxyFactoryFinderInterceptor.java:138) at org.jboss.ejb.SessionContainer.internalInvoke(SessionContainer.java:650) at org.jboss.ejb.Container.invoke(Container.java:1072) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at	0	0	https://issues.jboss.org/browse/JBNAME-42

#	Requirement	Questions	Answered	Successfully Answered	URL
12198	JBNAME-38	Running the unit tests yields: --- Test set: org.jnp.test.NamingEventsUnitTest --- Tests run: 2, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 0.053 sec <<< FAILURE! testAddRemoveOneLevel(org.jnp.test.NamingEventsUnitTest) Time elapsed: 0.01 sec <<< ERROR! java.lang.UnsupportedOperationException: Naming implementation (org.jnp.server.NamingServerWrapper@30d82d) does not implement NamingEvents at org.jnp.interfaces.NamingContext.addNamingListener(NamingContext.java:1296) at org.jnp.interfaces.NamingContext.addNamingListener(NamingContext.java:1314) at org.jnp.test.NamingEventsUnitTest.testAddRemoveOneLevel(NamingEventsUnitTest.java:88) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at junit.framework.TestCase.runTest(TestCase.java:168) at junit.framework.TestCase.runBare(TestCase.java:134) at junit.framework.TestResult$1.protect(TestResult.java:110) at junit.framework.TestResult.runProtected(TestResult.java:128) at junit.framework.TestResult.run(TestResult.java:113) at junit.framework.TestCase.run(TestCase.java:124) at junit.framework.TestSuite.runTest(TestSuite.java:232) at junit.framework.TestSuite.run(TestSuite.java:227) at org.junit.internal.runners.JUnit38ClassRunner.run(JUnit38ClassRunner.java:81) at org.apache.maven.surefire.junit4.JUnit4TestSet.execute(JUnit4TestSet.java:62) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:140) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:127) at org.apache.maven.surefire.Surefire.run(Surefire.java:177) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.booter.SurefireBooter.runSuitesInProcess(SurefireBooter.java:345) at org.apache.maven.surefire.booter.SurefireBooter.main(SurefireBooter.java:1009)	0	0	https://issues.jboss.org/browse/JBNAME-38
12199	JBNATIVE-7	Can you attach the hs_err_pidxxx.log file? Thanks.	1	1	https://issues.jboss.org/browse/JBNATIVE-7
12200	JBNATIVE-7	Greg - how did you get that "thread hanging stack"? The hs_ file is for some other bug it seems. When my VM hangs on exit - I don't get an hs_ file - it just hangs my windows cmd prompt as I describe in my comment on 22/Mar/07 07:07 PM	1	1	https://issues.jboss.org/browse/JBNATIVE-7
12201	JBNATIVE-7	There is something specific to my app that causes this hang whenever SIGHT is loaded. I just tried a simple java app to try and replicate and I can't. I added a shutdown hook that terminates the lib and it all works. Trying to see how I can replicate this outside of the agent in a small app. One thing I noticed was that when the agent hangs, its log4j log file gets truncated to file size of 0. I don't know why but I'm sure its a symptom of the problem. Whenver the JVM tries to exit and it hangs, the log file seems to go down to a file size of 0.	0	0	https://issues.jboss.org/browse/JBNATIVE-7
12202	JBNADM-3091	[10:07] <jdob> my first thoughts are that its related to what I emailed you yesterday [10:08] <jdob> in that email, where I say I create the resource to commit it, I dont check to see if there was a previously created on with that key that was marked as deleted [10:08] <jdob> so chances are, its creating a new resource with the same key as one that already exists thats marked as deleted [10:09] <jdob> and thats boogering up our discovery [10:09] <mspritz> and in some queries that deleted record gets returned instead. [10:09] <jdob> if the query doesnt use the ID and just things like key, ya, that'd make sense [10:09] <mspritz> yep. [10:09] <jdob> ResourceFactoryManagerBean:209 [10:09] <mspritz> So why are you done yet? [10:09] <mspritz> not [10:09] <jdob> that method needs to check if a resource with that key exists and is marked as deleted [10:10] <mspritz> :D [10:10] <jdob> and just flip the status to committed	1	1	https://issues.jboss.org/browse/JBNADM-3091
12203	JBNADM-3057	Of course, I agree with Jason. The ui rendering would be nice. However, is it really the UI's responsibility? It might be the only solution, but I don't think the UI should have that kind of logic.	1	0	https://issues.jboss.org/browse/JBNADM-3057
12204	JBNADM-2995	New plan: (1:03:38 PM) ccrouc1: ok so we're going to breakout the license generator into its own module, that doesnt end up in the server distro (1:03:48 PM) ccrouc1: and in the src which goes to RHQ change the salt? (1:04:04 PM) ccrouc1: is that the plan? (1:04:22 PM) ips-lunch: sounds good to me (1:04:23 PM) ghinkle: yea	1	1	https://issues.jboss.org/browse/JBNADM-2995
12205	JBNADM-2958	well, we don't have (won't have?) a full listing of our contributors in a copyright.txt file.	0	0	https://issues.jboss.org/browse/JBNADM-2958
12206	JBNADM-2953	Do we want to have a dbupgrade script here as well? Or would we just start with a blank database again?	0	0	https://issues.jboss.org/browse/JBNADM-2953
12207	JBNADM-2953	We just discussed on the call that we are not doing a dbupgrade step for this, but reset the schema version to 2.0 with no dbupgrade. Question remains: when do we start using dbupgrade scripts again?	0	0	https://issues.jboss.org/browse/JBNADM-2953
12208	JBNADM-2577	C Icon is for Configuration. Do you mean the Operations Icon?	0	0	https://issues.jboss.org/browse/JBNADM-2577
12209	JBNADM-2526	rev7689 JBNADM-2529 - can click "VIEW HISTORY" from inventory tab of some resource group to view list of group plugin configuration updates; JBNADM-2527 - can click on row item from JBNADM-2529 to view group update request details, each of those rows are clickable to the individual PluginConfigurationUpdate; JBNADM-2531 - can click "EDIT" from inventory tab of some resource group, which will initiate an async request to update the all resource plugin configurations in the group;	0	0	https://issues.jboss.org/browse/JBNADM-2526
12210	JBNADM-2368	Jess, it sounded from the status call you were looking at this. Can you close it out?	1	1	https://issues.jboss.org/browse/JBNADM-2368
12211	JBNADM-2040	(9:36:05 AM) ghinkle: its already done in the plugin... just need to add the check in the UI configure page (9:37:56 AM) ccrouch: so is this going to be a configurable minimum (9:38:05 AM) ghinkle: no (9:38:09 AM) ghinkle: hard code it to 30 seconds	1	1	https://issues.jboss.org/browse/JBNADM-2040
12212	JBNADM-1196	Isn't that a 1.4 only issue and could be just closed? In 2.0 there is no more status bar.	0	0	https://issues.jboss.org/browse/JBNADM-1196

#	Requirement	Questions	Answered	Successfully Answered	URL
12213	JBNADM-1143	here's what we need: 1) agent sends up inventory report - we should not schedule any metric collections yet since we don't even have the resources imported yet 2) the server gets the inventory report and populates them as NEW in ON_RESOURCE and puts them in the auto-discovery portlet 3) user clicks "import" to commit resources into inventory 4) we send a message to the agent to say, "this set of resources have been imported, you need to go grab their measurement schedules" - this can be sent asynchronously 5) the agent gets the message and then turns around and sends a request to the server asking for the schedules 6) the agent gets the schedules and populates the ResourceContainer with the new schedules 7) the MeasurementManager needs to be told to now collect the data based on the new schedules Questions: a) how do we identify the resources to the agent? can we send down the resource keys only? the parent resource key and the resource key? b) can we combine 4) and 5)? can the server just send down the schedules after the import button is pressed? it can do this asynchronously	1	1	https://issues.jboss.org/browse/JBNADM-1143
12214	JBNADM-1140	Ran out of time in beta1 so moving to beta2. The last comment on this is from April 10th. Is this still an issue?	1	0	https://issues.jboss.org/browse/JBNADM-1140
12215	JBNADM-1138	One more thing... we're going to need customizable error messages for constraint validations... (why does this field need to be between 1024 and 65556?)	1	1	https://issues.jboss.org/browse/JBNADM-1138
12216	JBNADM-1054	Can this be integrated as part of the regular tasks done during a release or development cycle. For instance, at the end of alpha 1, let's make sure we audit all the code that was submitted for alpha 1.	0	0	https://issues.jboss.org/browse/JBNADM-1054
12217	JBNADM-1049	Some first round things this needs to address: Making sure we've secured all the requests (no changing the resource id parameter) Checking for system end-arounds (agents that don't need to register or the like, unnecessary open ports) Check for trust issues with the agent (e.g do we let the agent update inventory for items not local to the agent? Call damaging server apis) File system dangers documented with plugins and artifacts (e.g. can i request "../../../../etc/something") Config, operations, service starting	1	1	https://issues.jboss.org/browse/JBNADM-1049
12218	JBNADM-1049	Facet Interfaces MeasurementFacet ---------- - Need to change "public MeasurementDataNumeric(MeasurementDataPK md, Double value)" to protected - Need to change "public MeasurementDataTrait(MeasurementDataPK, String value)" to protected ArtifactFacet ---------- - ArtifactType instance can be manipulated when passed into artifact discovery - Possible solution: clone before sending in - Possible solution: only send in pieces of artifact type data - ArtifactType instance can be manipulated when passed through ArtifactDetails (applies to revision discovery, create, etc.) - Possible solution: clone when creating ArtifactDetails (ArtifactUtils.artifactToDetails) - ArtifactRevisionDetails only contains artifact key, not the type. This isn't a problem since the type is known by the ArtifactManager when the revision discovery is made. However, the type is not taken into account when associating the revision with the artifact. This means that if two artifacts + on the same resource + of different types, have the same artifact key, we'll have a problem. The model is set up to maintain that relationship, it's just the revision processing that needs to be fixed. - InputStream passed into create/update. Can this potentially be an issue if the plugin sits on the stream and doesn't use it? DeleteResourceFacet ---------- CreateChildResourceFacet ---------- - ResourceType instance can be manipulated when passed into create method - ArtifactType instance can be manipulated when passed into create method - Should resource and/or plugin configuration in CreateResourceReport return copies when gotten from the report? - createResource shouldn't return the CreateResourceReport, it should be populated in place ConfigurationFacet ---------- OperationFacet ----------	1	1	https://issues.jboss.org/browse/JBNADM-1049
12219	JBNADM-1049	Scenario #5 ========== "Install a fresh agent somewhere else and start it fresh (no token). It registers with the server, but this one also wants to have agent name "A". Server accepts the message even though no token exists in message because it is a register request. Server sees agent "A" already exists and returns to the agent the token for A which is "A2-TOKEN". Now let's say this new "A" agent executed "register regenerate" - the server will give it a new token "ANOTHER-TOKEN". * this is bad, as it means any agent can accidentally (or maliciously) steal the token of another agent and can behave as that agent. If this new agent used "register regenerate", then the original agent's token is now invalid because a new token was assigned to it. If that happens, the original "A" can't talk to the server because its token (A2-TOKEN)is now invalid (the new valid token for A is now "ANOTHER-TOKEN). Is the answer to this to NOT allow an agent to "regenerate" a token? Perhaps a token is created once and from thereafter used for that agent - that token is forever linked to that agent?" ---------- This is currently an issue. The test I ran was very similar to the one described, but I'll quickly summarize it. - Start agent A. It is given a token, for example "123". - Start agent B. Agent B is given the same name as agent A. The server recognized the name and sent agent B the token "123". - Entered the command "register" on agent B. Server reflects that it will not regenerate the token. - Entered the command "register regenerate" on agent B. Server indicates it will regenerate the token. Agent B's new token is "456". Problem: Agent A's token is still "123". The server starts to reject commands from agent A. - Entered the command "register" on agent A. It received token "456" from the server. It was back in sync. This does exhibit the problem as described in the original scenario. An agent can disable another agent and act in its place by simply specifying the same agent name.	1	1	https://issues.jboss.org/browse/JBNADM-1049
12220	JBNADM-1049	Scenario #6 ========== This second agent wants to change its name to "B" and asks to re-register again (say, via register prompt command). Server sees token "ANOTHER-TOKEN" and accepts it as valid. Server sees that no agent is named "B" and creates it and returns to the agent a new token "B-TOKEN". * this is bad - because now there is a ghost agent named "A" in the database. In fact it is worse than that because agent A still has all its resource relationships. So all the resources that once belonged to A in the DB still belong to A. Changing the business key of an agent (its name) is bad. Is the answer to this that, as part of the registration request, the agent should send its current token. If the agent already has a token, the server can check this token in the database and see what agent name is associated with it. If it is different than the name of the agent asking for the registration, the server should abort the registration and tell the agent "you are already registered under the name "A" - you cannot change your name to something else" ---------- This test passed. - Start agent A with name "OriginalName". It successfully handshakes with the server. - Stop agent A. - Change the name of agent A in the Java prefs to "NewName". - Start agent A. - Enter command "register". The following output was seen in the agent console: The agent will now attempt to register itself with the server. Regenerating token=[false] Waiting 30 seconds... The server has rejected the agent registration request. Cause: [org.jboss.on.serverclient.core.AgentRegistrationException:The agent asking for registration is already registered with the name [MyAgent], it cannot change its name to [MyNewAgent]] The agent was not able to register itself. Is the server up? The same output was received for a call to "register regenerate". I checked the database as well. There was no entry that referenced "NewName". The contents of the "OriginalName" row remained in tact and unchanged (i.e. its token was not effected).	1	1	https://issues.jboss.org/browse/JBNADM-1049
12221	JBNADM-1049	What about JBNADM-2234 -> JBNADM-2243?	1	1	https://issues.jboss.org/browse/JBNADM-1049

#	Requirement	Questions	Answered	Successfully Answered	URL
12222	JBNMAN-433	What issues are you running into? We have customers that run on 10g without problem. As for version support, we support 8i, 9i, and 10g. (The last of our customers went off 8i just recently)	1	1	https://issues.jboss.org/browse/JBNMAN-433
12223	JBNMAN-410	Sachs can you take care of this? You can use snmp-2.5.19.tar.gz, code base is the same as jboss'	0	0	https://issues.jboss.org/browse/JBNMAN-410
12224	JBNMAN-410	The customer is running Apache 2.0.55 on Solaris 9. Could we please get the compiled SNMP and RT modules for this combination? Thanks Ajit	1	1	https://issues.jboss.org/browse/JBNMAN-410
12225	JBNMAN-405	Jav, can you comment on this?	0	0	https://issues.jboss.org/browse/JBNMAN-405
12226	JBNMAN-405	Looking through some code generated by xdoclet I came across the following in one of my value classes: public void removeAllFeedLightValues() { // DOH. Clear the collection - javier 2/24/03 this.FeedLightValues.clear(); } Do we have documented somewhere the changes Hyperic made to Xdoclet? If not, could they be documented here too.	1	0	https://issues.jboss.org/browse/JBNMAN-405
12227	JBNMAN-405	Thanks Javier. When did you guys take the snapshot of xdoclet's CVS, i.e. most of the contents of \xdoclet\xdoclet and \xdoclet\xjavadoc?	0	0	https://issues.jboss.org/browse/JBNMAN-405
12228	JBNMAN-403	Yes, we want those binaries. Please give us those in the next code drop	1	1	https://issues.jboss.org/browse/JBNMAN-403
12229	JBNMAN-385	The following comment we received from a customer is relevant to how accurate our platform availability measurements are... "Last night we installed maintenance on our windows servers. This involved rebooting the machines. Each reboot takes the server out for a few minutes. After the servers came back up I looked at the availability measurements in the manager. There were a couple of grey dots where the server was down for the reboot and the system uptime measurement dropped to a couple of minutes after the reboot. However, the availability measurement still stated 100% available. It looks like you are not sensing the change in availability properly." Since the metric collection interval for availability of a Windows platform is 1minute we should, in theory, be able to take account of any downtimes longer than that. IIRC the gray dots indicate we couldn't collect availability data for that time period. So it appears we don't immediately update the availability metric just because we don't get a complete set of responses. If that's the case, what's our real level of tolerance when it comes to platform availability: "If the platform is down for longer than X minutes then it will begin to affect the availability %"? What is X in our case? Also presumably when it comes to platform availability, we don't differentiate between a platform being down and the agent on that platform not running?	1	1	https://issues.jboss.org/browse/JBNMAN-385
12230	JBNMAN-385	The unavailability does not get reported immediately. There's is a lag time to account for slow networks, server load, etc. The way the platform unavailability gets reported is as follows: An "unavailable filler" runs every 5 minutes on the server. It looks for availability metrics that have not had received any values at the designated intervals and fills the holes with 0% availability. However, it does so only for "holes" that have occurred longer than 3 intervals ago. So, this means that a platform availability metric that is scheduled for 1 minute intervals may get the unavailable metric reported between 3-5 minutes.	1	1	https://issues.jboss.org/browse/JBNMAN-385
12231	JBNMAN-384	I'm having a tough time reproducing this. In my setup, I'm running everything on the same machine: JON Server JON Agent JBoss 4.0.2 stock When I stop and start the server, the results come back as expected. (Using the 2.2.21 codebase). Does your environment differ in any way? One interesting piece of info would be the agent log with DEBUG logging turned on. If you are seeing abnormal delay, a threaddump may help to. For reference this is what I see in my agent log: 2006-01-25 19:57:20,846 DEBUG [SSLConnectionListener] Handling SSL connection from /127.0.0.1 2006-01-25 19:57:20,967 DEBUG [CommandListener] Dispatching request for 'control:controlPluginCommand' 2006-01-25 19:57:20,968 DEBUG [CommandListener] Method 'net.hyperic.hq.agent.AgentCommand@1d85a26' returned an object result 2006-01-25 19:57:20,981 DEBUG [ActionThread] Adding job 10002 to 2-10051 queue 2006-01-25 19:57:20,981 DEBUG [ActionThread] Running job 10002 2006-01-25 19:57:26,019 INFO [JBossServerControlPlugin] doCommand args=[/home/rmorgan/hyperic/hyperic_hq.jboss/build/agent/background.sh, /home/rmorgan/apps/jboss-4.0.2/bin/run.sh, -c, default] 2006-01-25 19:57:26,030 DEBUG [Execute] Waiting process 2006-01-25 19:57:26,049 DEBUG [Execute] End waiting, stop threads 2006-01-25 19:57:26,049 DEBUG [Execute] Watchdog stopped 2006-01-25 19:57:26,213 DEBUG [Execute] Stream handler stopped 2006-01-25 19:57:26,213 DEBUG [Execute] Done exit=0 /usr/bin/env 2006-01-25 19:57:26,220 DEBUG [Execute] Waiting process 2006-01-25 19:57:26,233 DEBUG [Execute] End waiting, stop threads 2006-01-25 19:57:26,233 DEBUG [Execute] Watchdog stopped 2006-01-25 19:57:28,236 DEBUG [Execute] Stream handler stopped 2006-01-25 19:57:28,236 DEBUG [Execute] Done exit=0 /home/rmorgan/hyperic/hyperic_hq.jboss/build/agent/background.sh /home/rmorgan/apps/jboss-4.0.2/bin/run.sh -c default 2006-01-25 19:57:28,236 INFO [JBossServerControlPlugin] doCommand result = 0 2006-01-25 19:57:38,902 DEBUG [ActionThread] Removed job 10002 from the queue	1	1	https://issues.jboss.org/browse/JBNMAN-384
12232	JBNMAN-379	In one of our EM installations, a monitored Apache instance has the snmpCommunity property set to "public." Is this different from the read-only setting you mention? Also, snmpUser, snmpPassword, and snmpAuthType are defined for SNMP v3. Woud it be possible to take advantage of this to further secure the Apache SNMP agent?	1	1	https://issues.jboss.org/browse/JBNMAN-379
12233	JBNMAN-379	There are discrepancies in the SNMP configuration properties listed on EM's Server Configuration page (Administration->Server Configuration) and its resource configuration page for the Apache server. For example, the set of properties that apply to SNMP v3 are different on each page. Shouldn't these be identical? Also, the Server Configuration edit page tailors the list of properties to the selected SNMP version. Would it make sense to have the Apache server configuration edit page behave that way as well?	1	1	https://issues.jboss.org/browse/JBNMAN-379

#	Requirement	Questions	Answered	Successfully Answered	URL
12234	JBNMAN-375	I've uploaded the server.log and a debug agent.log from the machine running IIS. [BTW one crazy thing with the server.log which we haven't been able to figure out is that he lines 2378 thru 3350 are duplicated from lines 3352 to the end of the file. Except the duplicate lines also have the date and the fully qualified class name ????] From looking at the server.log it appears that we initially had problems with many, but not all, of the IIS virtual hosts, e.g. 15:50:19,128 INFO [RuntimeReportProcessor] Merging Service into inventory: name: aisdev1 IIS 6.x A7900 PayrollExpense 15:50:19,128 INFO [RuntimeReportProcessor] Creating new service: aisdev1 IIS 6.x A7900 PayrollExpense 15:50:19,194 WARN [DefaultMetricsEnablerUtil] Error turning on default metrics, configuration ({protocol=http, iishost=D5805 webDeposit, timeout=10, server.log_track.enable=false, installpath=C:\WINDOWS\system32, port=10001, server.log_track.level=Error, service_name=W3SVC, platform.ip=10.160.1.53, server.log_track.pattern=, platform.name=aisdev1.ais.ad.msu.edu, platform.fqdn=aisdev1.ais.ad.msu.edu, hostname=localhost, path=/, platform.type=Win32}) couldn't be validated: Invalid configuration: Resource reports that it is unavailable 15:50:19,202 WARN [DefaultMetricsEnablerEJBImpl] Caught InvalidConfigException in DefaultMetricEnabler MDB: Invalid configuration: Resource reports that it is unavailable 15:50:19,389 INFO [RuntimeReportProcessor] Completed merging service: aisdev1 IIS 6.x A7900 PayrollExpense into inventory. But then by part way through the log these errors have cleared up. In the error messages the path is always to '/', which must be wrong since the ports and hostnames are always the same too. Does this mean the customer has to manually set the path on every virtual host service?	1	1	https://issues.jboss.org/browse/JBNMAN-375
12235	JBNMAN-375	Re-opening to confirm that the fix will work when the value of the HTTP host header parameter is not reachacle from the machine running the agent. Just to clarify there are currently two properties related to hosts in the IIS vhost configurations properties: "hostname" and "iishost" In my comment I talked about adding a third parameter, may "HTTP Host header". Is that the direction you were going in?	1	1	https://issues.jboss.org/browse/JBNMAN-375
12236	JBNMAN-367	Apparently the customer just resolved this himself! Customers... don't you just love 'em.	0	0	https://issues.jboss.org/browse/JBNMAN-367
12237	JBNMAN-362	I have the agent working w/ IBM 1.4 JRE. Changes so far: - changed agent certificate generation code to use keytool rather using deprecated com.sun classes - changed com.sun.ssl.* usage to javax.ssl.* (agent and lather) - removed lib/jdk1.3-compat libs I just need to do some more testing and cleanup. Also, on our side of the problem, works fine w/ WebSphere 5.1, not sure yet what we'll do about WebSphere 4.0 Since this does not impact you guys, I expect the changes to be in the next code drop, I'll mark this resolved once they are in.	0	0	https://issues.jboss.org/browse/JBNMAN-362
12238	JBNMAN-362	The IBM 1.4.2 jre only works on AIX 5.1, 5.2 and 5.3. It does not work on 4.3.x. This is an issue for us, can you guys checks which version(s) your AIX customers are using?	0	0	https://issues.jboss.org/browse/JBNMAN-362
12239	JBNMAN-361	I refactored all the NestedException stuff... it was almost too easy - I'm thinking something just isn't going to work. I'm going to assign this to Charles Lee - Charles - can you tell me what some of the issues were that caused you to pause wrt this issue? What are some of the things you can see going wrong? It compiles fine for me now, but I haven't run with it yet - i'm scared :) What things would you expect to break?	1	1	https://issues.jboss.org/browse/JBNMAN-361
12240	JBNMAN-361	Charles - can you comment and then reassign back to me? Thanks...	1	1	https://issues.jboss.org/browse/JBNMAN-361
12241	JBNMAN-327	ryan - I ran into this last month, unsure if any of the code drops we got since then fixed the problem. But, also, this issue was more to discuss the testing framework. Can you give us a quick blurb on how you guys use it and how we should use it?	1	1	https://issues.jboss.org/browse/JBNMAN-327
12242	JBNMAN-307	Can you provide a test case where there are more errors not shown? The exceptions are displayed with getStack() method.	1	1	https://issues.jboss.org/browse/JBNMAN-307
12243	JBNMAN-304	Jav, can you fill Maz in on how this stuff works?	1	1	https://issues.jboss.org/browse/JBNMAN-304
12244	JBNMAN-303	Thanks! Have you ever had reason to purge an agent from the database? That was something I think I needed yesterday. What if I did the following: It was determined that my DB was hosed because someone tried to import things from an agent (running on the server box) but there was no network connectivity. 1) Have the customer delete the platform from the inventory (so it deletes all servers/services for that platform too) 2) Purge the agent's row from the eam_agent table 3) Run those DB comments you specified in the earlier comment to clear the agent's AI queue If I do those three things, can it be assumed the customer is running from a "clean slate" with respect to that agent. So, if restart the agent (after getting it some network connectivity :-) it will ask for and get a new token and start anew with being the AI queue from things it discovered. ??	1	1	https://issues.jboss.org/browse/JBNMAN-303
12245	JBNMAN-294	I was trying to be too clever :-). As Rich pointed out, removing the '/b' from the script causes everything to work great. So '/b' means exit the batch file but not the cmd.exe process. Does this issue highlight a problem with long running scripts not getting timed out and marked as failed by JBN EM?	1	1	https://issues.jboss.org/browse/JBNMAN-294
12246	JBNMAN-294	exit /b 0 worked fine for me, i got green lights and the metrics, here my agent log: 2005-09-28 09:56:10,989 DEBUG [MeasurementCommandsServer] Getting real time measurement: Script:system.script:Arg=D%3A\dougm\test\test.bat:Availability 2005-09-28 09:56:11,020 DEBUG [Execute] Waiting process 2005-09-28 09:56:11,083 DEBUG [Execute] End waiting, stop threads 2005-09-28 09:56:11,083 DEBUG [Execute] Watchdog stopped 2005-09-28 09:56:11,083 DEBUG [PumpStreamHandler] Joining it 2005-09-28 09:56:11,083 DEBUG [PumpStreamHandler] Joined 2005-09-28 09:56:11,083 DEBUG [PumpStreamHandler] Joining it 2005-09-28 09:56:11,083 DEBUG [PumpStreamHandler] Joined 2005-09-28 09:56:11,083 DEBUG [Execute] Stream handler stopped 2005-09-28 09:56:11,083 DEBUG [Execute] Done exit=0 D:\dougm\test\test.bat 2005-09-28 09:56:11,083 DEBUG [MeasurementCommandsServer] Result was: 1 i added 'sleep 600' to test.bat and it was killed properly after the 30 second timeout: 2005-09-28 10:15:00,098 DEBUG [Execute] Waiting process 2005-09-28 10:15:23,145 DEBUG [RtScheduleThread] Sleeping 60000 to next batch 2005-09-28 10:15:24,583 DEBUG [RunnableTrackThread] Running loggers 2005-09-28 10:15:30,114 DEBUG [Execute] End waiting, stop threads 2005-09-28 10:15:30,114 DEBUG [Execute] Watchdog stopped 2005-09-28 10:15:30,114 DEBUG [PumpStreamHandler] Joining it 2005-09-28 10:15:30,114 DEBUG [PumpStreamHandler] Joining it 2005-09-28 10:15:31,114 DEBUG [PumpStreamHandler] Joined 2005-09-28 10:15:31,114 DEBUG [Execute] Stream handler stopped 2005-09-28 10:15:31,114 DEBUG [Execute] Done exit=1 D:\dougm\test\test.bat 2005-09-28 10:15:31,114 WARN [ScheduleThread] Collection of metric: 'Script:system.script:Arg=D%3A\dougm\test\test.bat:ExecTime' took: 31016ms I guess there should be an ERROR logged if the script is killed due to timeout, is that what you're looking for?	1	1	https://issues.jboss.org/browse/JBNMAN-294

#	Requirement	Questions	Answered	Successfully Answered	URL
12247	JBNMAN-292	We have observed all the suggested alternatives proposed here and elsewhere prior to reporting this issue, and still are encountering deadlocks with the exact same codebase which did not have these issues in 3.2.3. If you review the entity layer of the application's various subsystems, you'll find that every single getter method is marked read-only. This seems consistent with the dtd's for 4.0.2. We are using a rather old version of xdoclet, and are generating 3.2 descriptors in our ear. I was worried that this might perhaps have changed in 4.x but I can not see any differences in this specific area. Mazz, can you please email Alexey the jboss.xml dtd which is produced as part of the HQ build process? It is found in hq/build/dd/jboss.xml etc. thanks.	1	1	https://issues.jboss.org/browse/JBNMAN-292
12248	JBNMAN-292	Richard Friedman asked: "Alexey, I noticed our EJB's are being generated with xdoclet and jboss 3.0 Does it matter? Should this be upgraded for jboss 4.0?" Alexy responded: "It's not problem for me. The DTD version number affects only allowed elements in the DDs. If you don't use features added in 3.2, it's irrelevant. 4.x and 3.2.x are equal in CMP. You could just update to 3.2."	1	1	https://issues.jboss.org/browse/JBNMAN-292
12249	JBNMAN-292	The dd's are for 3.2 (despite the fact that we're telling XD to use 3.0). This is part of the customization we made to xdoclet in our sources. I figured this would be his response (which is good). Has he had a chance to review the descriptors? read-only methods are *definitely* locking in a way they werent in 3.2.3 Thanks. -j	1	1	https://issues.jboss.org/browse/JBNMAN-292
12250	JBNMAN-292	Javier -- Scott Stark is asking if we can come up with a replicatable test case. Is it possible for you to come up with some test that shows the getId() vs. getPrimaryKey().getId() issue?	0	0	https://issues.jboss.org/browse/JBNMAN-292
12251	JBNMAN-276	http://ant.apache.org/faq.html#unknownelement.taskcontainer ---- Why do my custom task containers see Unknown Elements in Ant 1.6 - they worked in Ant 1.5? The objects added in TaskContainer.addTask(Task task) have changed from Tasks to UnknownElements. There was a number of valid reasons for this change. But the backward compatibility problems were not noticed until after Ant 1.6.0 was released. Your container class will need to be modified to check if the Task is an UnknownElement and call perform on it to convert it to a Task and to execute it. (see apache.tools.ant.taskdefs.Sequential) If you want to do more processing on the task, you need to use the techniques in apache.tools.ant.taskdefs.Antlib#execute() This does make use of one 1.6 method call (UE#getRealObject()), you need to use UE#getTask() instead - this will return null for non tasks (types like fileset id=x). So.. iterate over the tasks, if they are UEs, convert them to tasks, using UE#maybeConfigure and UE#getTask() for (Iterator i = tasks.iterator(); i.hasNext();) { Task t = (Task) i.next(); if (t instanceof UnknownElement) { ((UnknownElement) t).maybeConfigure(); t = ((UnknownElement) t).getTask(); if (t == null) { continue; } } // .... original Custom code } This approach should work for ant1.5 and ant1.6.	1	1	https://issues.jboss.org/browse/JBNMAN-276
12252	JBNMAN-275	allow AgentDListProvider to run standalone as a quick way to dump the values * found in a data file. * * @param args one argument which is passed as the first argument to {@link #init(String)} * * @throws Exception */ public static void main(String args[]) throws Exception { System.out.println("Dumping keys/values found in data file."); System.out.println("#####"); AgentDListProvider listProvider = new AgentDListProvider(); listProvider.init((args.length == 1) ? args[0] : "data:20:50"); Set keys = listProvider.getKeys(); for (Iterator iter = keys.iterator(); iter.hasNext();) { String key = (String) iter.next(); String value = listProvider.getValue(key); System.out.println(key + "=" + value); } System.out.println("#####"); } ---------- When run, my agent's data/keyvals properties are dumped like this: ---------- Dumping keys/values found in data file. ##### runtimeautodiscovery.2-10008-JBoss 4.0.platform.ip=192.168.0.5 runtimeautodiscovery.2-10006-JBN Agent.platform.name=mazzlap runtimeautodiscovery.2-10009-Tomcat 5.5.installpath=C:\mazz\jboss\hq-engine\jboss-4.0.2-bin\server\default\deploy\jbossweb-tomcat55.sar runtimeautodiscovery.2-10009-Tomcat 5.5.platform.name=mazzlap agent.certDN=CAM-AGENT-1126249626177-37537816232016306 1-1205726130352927838 runtimeautodiscovery.2-10008-JBoss 4.0.platform.fqdn=192.168.0.5 covalent.CAMAgentToken=1126249629842-1548575944595594040-1194344452432975171 runtimeautodiscovery.2-10009-Tomcat 5.5.platform.type=Win32 runtimeautodiscovery.2-10006-JBN Agent.platform.type=Win32 agent.hostName=mazzlap runtimeautodiscovery.2-10008-JBoss 4.0.platform.name=mazzlap runtimeautodiscovery.2-10009-Tomcat 5.5.jmxUrl=http://localhost:7080 measurement_srn=AAAAnwAAAAIAACcYAAAAAQAAAAIAACcZAAAAAQAAAAIAACcWAAAAAQAAAAMAACcUAAAAAQAAAAMAACcSAAAAAQAAAAMAACcWAAAAAQAAAAMAACcVAAAAAQAAAAMAACcTAAAAAQAAAAMAACcXAAAAAQAAAAMAACcYAAAAAQAAAAMAACcaAAAAAQAAAAMAACcZAAAAAQAAAAMAACcbAAAAAQAAAAMAACccAAAAAQAAAAMAACcdAAAAAQAAAAMAACceAAAAAQAAAAMAACcfAAAAAQAAAAMAACcgAAAAAQAAAAMAACchAAAAAQAAAAMAACciAAAAAQAAAAMAACcjAAAAAQAAAAMAACckAAAAAQAAAAMAACclAAAAAQAAAAMAACcmAAAAAQAAAAMAACcnAAAAAQAAAAMAACcoAAAAAQAAAAMAACcpAAAAAQAAAAMAACcqAAAAAQAAAAMAACcrAAAAAQAAAAMAACcsAAAAAQAAAAMAACctAAAAAQAAAAMAACc7AAAAAQAAAAMAACc4AAAAAgAAAAMAACczAAAAAgAAAAMAACc8AAAAAQAAAAMAACc2AAAAAgAAAAMAACc9AAAAAQAAAAMAACcwAAAAAgAAAAMAACc+AAAAAQAAAAMAACc3AAAAAgAAAAMAACc/AAAAAQAAAAMAACdAAAAAAQAAAAMAACcvAAAAAgAAAAMAACdBAAAAAQAAAAMAACc6AAAAAgAAAAMAACdCAAAAAQAAAAMAACc5AAAAAgAAAAMAACdDAAAAAQAAAAMAACcxAAAAAgAAAAMAACc1AAAAAgAAAAMAACdEAAAAAQAAAAMAACd0AAAAAgAAAAMAACdGAAAAAQAAAAMAACcyAAAAAgAAAAMAACdHAAAAAQAAAAMAACcuAAAAAgAAAAMAACdIAAAAAQAAAAMAACdJAAAAAQAAAAMAACdKAAAAAQAAAAMAACdLAAAAAQAAAAMAACdMAAAAAQAAAAMAACdNAAAAAQAAAAMAACdOAAAAAQAAAAMAACdPAAAAAQAAAAMAACdQAAAAAQAAAAMAACdRAAAAAQAAAAMAACdSAAAAAQAAAAMAACdTAAAAAQAAAAMAACdUAAAAAQAAAAMAACdVAAAAAQAAAAMAACdWAAAAAQAAAAMAACdXAAAAAQAAAAMAACdYAAAAAQAAAAMAACdZAAAAAQAAAAMAACdcAAAAAQAAAAMAACdaAAAAAQAAAAMAACdbAAAAAQAAAAMAACddAAAAAQAAAAMAACdeAAAAAQAAAAMAACdfAAAAAQAAAAMAACdgAAAAAQAAAAMAACdhAAAAAQAAAAMAACdiAAAAAQAAAAMAACdjAAAAAQAAAAMAACdkAAAAAQAAAAMAACdlAAAAAQAAAAMAACdmAAAAAQAAAAMAACdnAAAAAQAAAAMAACdoAAAAAQAAAAMAACdpAAAAAQAAAAMAACdqAAAAAQAAAAMAACdrAAAAAQAAAAMAACdsAAAAAQAAAAMAACdtAAAAAQAAAAMAACduAAAAAQAAAAMAACdvAAAAAQAAAAMAACdwAAAAAQAAAAMAACdxAAAAAQAAAAMAACdyAAAAAQAAAAMAACdzAAAAAQAAAAMAACd0AAAAAQAAAAMAACd1AAAAAQAAAAMAACd2AAAAAQAAAAMAACd3AAAAAQAAAAMAACd4AAAAAQAAAAMAACd5AAAAAQAAAAMAACd6AAAAAQAAAAMAACd7AAAAAQAAAA	1	1	https://issues.jboss.org/browse/JBNMAN-275
12253	JBNMAN-263	Have to be honest, not sure I know what a TCP trap is, do you have a reference to something I can look at? Do they mean specifying SNMP traps on different ports on a different TCP port? Here is one definition, but it is relative to some product called Trap Console. From the following site: http://www.cscare.com/TrapConsole/download/readme.txt - TcpTrapPort The TCP trap port on which Trap Console listens for SNMP traps. By default SNMP traps are being sent to the SNMP trap port 162. However, especially for testing purposes or for security reasons, it is possible to configure Trap Console to listen for SNMP traps on a different TCP port. Default: TcpTrapPort=162	0	0	https://issues.jboss.org/browse/JBNMAN-263
12254	JBNMAN-247	resourcetype view -platform 10001 shows me the new metric. I guess this is another one in which I'm probably confused about what this command does. What did I create and why doesn't it show up in my platform 10001 (which is a Linux platform)??	1	1	https://issues.jboss.org/browse/JBNMAN-247
12255	JBNMAN-247	The 'metric create' command allows you to create custom metrics that are derived from other metrics that HQ is collecting. Once you created your derived measurement, did you enable it? It may need to be enabled to show up in the metric list. You'll either need to go through the UI and do: MONITOR -> [configure] to enable your metric Or you can try the 'metric configure -platform 10001' to reschedule all the metrics. Give that a quick try and let me know if it works.	1	1	https://issues.jboss.org/browse/JBNMAN-247

#	Requirement	Questions	Answered	Successfully Answered	URL
12256	JBNMAN-246	The 'metric configure' command within the shell does not deal specifically with metrics. It refers to the metric configuration for a particular resource, or entity. When you configure a resource's metric configuration it will apply the configuration and enable all metrics for that resource using the default metric collection intervals. Note that it does not obey the default on/off flag, as this was added recently. When you see the message: [Configuring metric: 2:10005] 2:10005 configured It means that it has configured resource 2:10005 metric configuration. Each resource has it's configuration broken down into 'product' (which is referred to as shared in the UI), 'metric', and 'control'. The shell breaks these down, so to configure a resource you must do: resource configure -server 10001 metric configure -server 10001 control configure -server 10001 In the UI all these options are in a single form and are committed at once. This is the 'Edit Inventory Properties' page. I've committed a fix to the typo you mention, and that will be in 2.1.13.	1	1	https://issues.jboss.org/browse/JBNMAN-246
12257	JBNMAN-230	I think I found it. ResourceManagerEJBImpl.createResource and removeResource. All roads lead to Rome - where these two seem to be Rome. All code that adds/removes resources from inventory eventually make their way here. Javier - can you confirm if that's true?	1	1	https://issues.jboss.org/browse/JBNMAN-230
12258	JBNMAN-230	Hmm... maybe this isn't the best solution. Under what conditions do these appdef events get sent or not sent? First, I put a breakpoint in the first line in DefaultMetricsEnableEJBImpl.onMessage - just to confirm that I am getting events. I confirmed that events are flowing (I can see periodic heartbeat events coming in). I then moved the breakpoint in DefaultMetricsEnablerEJBImpl.onMessage to the line: AppdefEvent ae = (AppdefEvent) obj At this point in the code, its been determined the event coming in is an Appdef event telling us something has been created, deleted, updated or whatever. However, during runtime, that breakpoint is not getting reached! I am not seeing any AppdefEvents flowing. I manually created a new Platform resource (and deleted it) but neither time did I see an AppdefEvent. I do this by going to "Browse Resourses" and do the "Add Platform". I created my own MDB and the same behavior is exhibited (I see heartbeats but no appdef events). Any ideas? Should this be considered a bug or is something amiss in the understanding of the AppdefEvents?	1	1	https://issues.jboss.org/browse/JBNMAN-230
12259	JBNMAN-225	Added support for 5.5 connectors. Also added JK support for 5.0 and 5.5 and removed some unused classes. Will all be in the 2.1.13 drop.	1	1	https://issues.jboss.org/browse/JBNMAN-225
12260	JBNMAN-224	If any user defined externally in LDAP is allowed to register with JBNEM without admin intervention, would it not create a potential security breach? It could be dangerous if the user can assign himself a role, but if that's not permitted, a malicious user could create some confusion about the JBNEM user base. I understand that the feature is provided with ease of administration in mind, but just wanted to learn what the potential risks would be for the customer.	1	1	https://issues.jboss.org/browse/JBNMAN-224
12261	JBNMAN-216	ryan - can you comment on what those yellow lights are trying to tell me? :)	0	0	https://issues.jboss.org/browse/JBNMAN-216
12262	JBNMAN-215	The following doesn't resolve this issue but it does answer the question 'how do I force a new auto-discovery of "services" (a runtime scan)'. The answer I believe is to go to the Inventory tab for the Server which contains the Services which you wish to auto-discovery and hit the 'EDIT...' button at the bottom of the 'Configuration Properties' section. On the new page which is rendered, make sure the Autodiscovery of embedded services box is checked and then hit the OK button, this should discovery any outstanding Services which are contained by this Server.	1	1	https://issues.jboss.org/browse/JBNMAN-215
12263	JBNMAN-214	Another HQ question. See http://help.hyperic.net/installation.html#GETTING-STARTED-INSTALL-GUIDE-PLANNING It talks about a medium infrastructure environment has up to 25 platforms. Anything more is considered large. If an infrastructure environment has 26 or more platforms, is it recommended: a) to use HA mode (i.e. have more than one server) and/or b) use Oracle? What's typically the recommendation for environments with hundred (or hundreds) or management platforms? How do you know if you should use HA, and if so how many HA servers to configure?	1	1	https://issues.jboss.org/browse/JBNMAN-214
12264	JBNMAN-214	Ryan - I just got off the phone with a customer of ours and they asked about this type of functionality - so can I bump this up to get these questions answered? Did you guys test all of this HA capability now that the server is running JBoss 4.0.2?	1	1	https://issues.jboss.org/browse/JBNMAN-214
12265	JBNMAN-213	If I go to the control tab now, I don't see the control actions (as expected) but instead I see this message: "Control Actions are not enabled. To enable control, set Configuration Properties in the Inventory section." This implies I could reenable the controls if I wanted to. I don't think this message is right. I click on the ConfigProps link there and it takes me to the configuration page but I don't think that's right. I think that message should not be there. Ryan - am I missing something?	0	0	https://issues.jboss.org/browse/JBNMAN-213
12266	JBNMAN-211	Could you please attach a patch for this build.xml change. Thanks.	0	0	https://issues.jboss.org/browse/JBNMAN-211
12267	JBNMAN-209	The SNMP module build process uses the prefix "covalent" in the names of the SNMP DSO's. These names are user-visible as they need to be specified in the LoadModule directives the user is required to add in httpd.conf: LoadModule covalent_snmpcommon_module modules/snmpcommon.so LoadModule covalent_snmpagt_module modules/snmpmonagt.so Can we drop the prefix in the Unix build scripts and also from the pre-built Windows modules?	1	1	https://issues.jboss.org/browse/JBNMAN-209
12268	JBNMAN-162	I see this in apache plugin jar's hq-plugin.xml: <scan registry="SOFTWARE\Apache Group\Apache\2.*"> <include name="ServerRoot"/> </scan> I do have a ServerRoot (whose value is my install dir of apache). Note also my Software node is capitalized (not all caps which is what is in this hq-plugin.xml). Its also under HKEY_CURRENT_USER. Because this registry attribute doesn't specify "HKEY_LOCAL_MACHINE" or "HKEY_CURRENT_USER", I don't know what the default is - where does it know where to look under? I guess it defaults to HKEY_LOCAL_MACHINE; if so, would explain why it can't find it using a registry scan.	1	1	https://issues.jboss.org/browse/JBNMAN-162
12269	JBNMAN-162	Understood about the SNMP thing... I wasn't looking to monitor it... just wanted to see if it could be discovered. Does the apache plugin do process scanning or only registry scanning? The fact that it always looked in HKEY_LOCAL_MACHINE is why it can't find it - it is not there. Its in the CURRENT_USER. I wonder if I installed it using the "only for this user" option, as opposed to "all users" (you know how some software can be installed so only the start menu group appears in the installing user's start menu or in all users' start menus? I wonder if this is related to that kind of thing). I can't remember how I installed it - I do know I installed it using the normal Apache installer.	1	1	https://issues.jboss.org/browse/JBNMAN-162
12270	JBNMAN-162	Yeah, it looks like if you don't install with the "All users", you don't get that HKEY_LOCAL_MACHINE reg. key. Can you add the ptql queries for apache plugin on windows? That would solve the problem should a customer install Apache using that "current user" install mode, as opposed to "all users".	1	1	https://issues.jboss.org/browse/JBNMAN-162
12271	JBNMAN-162	Doug - as per your comment in this issue - can you add the ptql queries for apache plugin on windows? If this is simple to do and can be put in the last code drop prior to our R1 release, great. If you can't, then just move this issue to "fix version" of R2.	1	1	https://issues.jboss.org/browse/JBNMAN-162

#	Requirement	Questions	Answered	Successfully Answered	URL
12272	JBNMAN-161	charles, are you on windows or unix? if windows, did you install tomcat using the .exe installer or use the .zip install bundle? the process query for tomcat should match all platforms. can i see the output from the command below w/ tomcat running: java -jar build/agent/pdk/lib/sigar.jar pargs State.Name.eq=java	1	1	https://issues.jboss.org/browse/JBNMAN-161
12273	JBNMAN-161	>charles, are you on windows or unix? Windows XP > if windows, did you install tomcat using the .exe installer or use the .zip install bundle? The version I was trying to autodetect was installed via zip, so no reg entries obviously. [FWIW I tried doing control actions on this instance, and they failed I think because there wasn't an associated Windows Service, so I guess install via .zip really isn't a supported configuration?] >the process query for tomcat should match all platforms. can i see the output from the command below w/ tomcat running: java -jar build/agent/pdk/lib/sigar.jar pargs State.Name.eq=java Well this command definitely finds it, number 2 in the list... C:\usr\projects\current\jboss_network\private_cvs\jboss-network\jbn-admin\hyperic_hq>java -jar build/agent/pdk/lib/sigar.jar pargs State.Name.eq=java pid=460 exe=C:\usr\apps\j2sdk1.4.2_06\bin\java.exe cwd=C:\usr\projects\current\jboss_network\hyperic\runtimes\jboss-3.2-bin\bin 0=>C:\usr\apps\j2sdk1.4.2_06\bin\java<= 1=>-Dprogram.name=run.bat<= 2=>-Xms128m<= 3=>-Xmx512m<= 4=>-classpath<= 5=>;C:\usr\apps\j2sdk1.4.2_06\lib\tools.jar;C:\usr\projects\current\jboss_network\hyperic\runtimes\jboss-3.2-bin\bin\\run.jar<= 6=>org.jboss.Main<= 7=>-c<= 8=>default<= ----------------------- pid=4172 exe=C:\usr\apps\j2sdk1.4.2_06\bin\java.exe cwd=C:\usr\apps\java\tomcat-5.30\bin 0=>C:\usr\apps\j2sdk1.4.2_06\bin\java<= 1=>-Djava.endorsed.dirs=C:\usr\apps\java\tomcat-5.30\common\endorsed<= 2=>-classpath<= 3=>C:\usr\apps\j2sdk1.4.2_06\lib\tools.jar;C:\usr\apps\java\tomcat-5.30\bin\bootstrap.jar<= 4=>-Dcatalina.base=C:\usr\apps\java\tomcat-5.30<= 5=>-Dcatalina.home=C:\usr\apps\java\tomcat-5.30<= 6=>-Djava.io.tmpdir=C:\usr\apps\java\tomcat-5.30\temp<= 7=>org.apache.catalina.startup.Bootstrap<= 8=>start<= ----------------------- pid=3056 exe=C:\usr\apps\j2sdk1.4.2_06\bin\java.exe cwd=C:\usr\projects\current\jboss_network\private_cvs\jboss-network\jbn-admin\hyperic_hq 0=>java<= 1=>-jar<= 2=>build/agent/pdk/lib/sigar.jar<= 3=>pargs<= 4=>State.Name.eq=java<= ----------------------- C:\usr\projects\current\jboss_network\private_cvs\jboss-network\jbn-admin\hyperic_hq>	1	1	https://issues.jboss.org/browse/JBNMAN-161
12274	JBNMAN-161	The output is below. It definitely finds the instance I was interested in (tomcat-5.30). Maybe the issue is that I didn't try to restart the agent to trigger the discovery? Does this query get run when you hit the auto-discovery scan from the web interface? C:\usr\projects\current\jboss_network\private_cvs\jboss-network\jbn-admin\hyperic_hq>java -jar build/agent/pdk/lib/hq-product.jar -m discover -p servlet -Ddiscover-services=fa lse 2 servers detected =-= Server: JBOSS-32F1AADFE Tomcat 5.0 [C:\usr\apps\java\tomcat-5.30] config... product..{jmxUrl=http://localhost:8080} metric...{} control..{startupScript=C:\usr\apps\java\tomcat-5.30\bin\catalina.sh} =-= Server: JBOSS-32F1AADFE Tomcat 5.0 [C:\Program Files\Apache Software Foundation\Tomcat 5.0.30-beta] config... product..null metric...null control..null C:\usr\projects\current\jboss_network\private_cvs\jboss-network\jbn-admin\hyperic_hq>Here is the output	0	0	https://issues.jboss.org/browse/JBNMAN-161
12275	JBNMAN-156	is this part of the code drop? What is the level of effort for this on windows?	0	0	https://issues.jboss.org/browse/JBNMAN-156
12276	JBNMAN-139	Just realized that we have the agent launcher source code. I see this in cam-agent.h: #define AGENT_CLASSPATH \ "-Djava.class.path=" \ "lib\\AgentClient.jar;" \ "lib\\AgentServer.jar;" \ "lib\\lather.jar;" \ "lib\\jdk1.3-compat\\jce1_2_2.jar;" \ "lib\\jdk1.3-compat\\sunjce_provider.jar;" \ "lib\\jdk1.3-compat\\jsse.jar;" \ "lib\\jdk1.3-compat\\xerces.jar;" which is I think what I saw in the classloader when I was debugging. I saw AgentServer and AgentClient jars as well as those jdk1.3 jars. So, the question is - is there anywhere that we can put our new agent jars that would allow them to be picked up without having to recompile the agent launcher C executable? I looks like (from the source) that the answer is pdk\lib. In cam-agent.cpp, there is a method called getClassPath that not only uses those AGENT_CLASSPATH .jars, but it also appends all jars it finds in pdk\lib. Is that how Hyperic expects it to work? If I have a jar that one of my Commands Server jars needs, should I just put it in pdk\lib?? BTW: I see where it sets that agent.mode system property in cam-engine.cpp: static char *jvm_options[] = { ... "-Dagent.mode=thread", //default is process ... };	0	0	https://issues.jboss.org/browse/JBNMAN-139
12277	JBNMAN-138	Ryan - I have added the icon and 2 images in. The license file is going to take a bit. can we proceed without it, we will update it in our cvs when completed. does it get bundled in any of the binaries which just bundle? Is it used in the unix / shell installers?	1	1	https://issues.jboss.org/browse/JBNMAN-138
12278	JBNMAN-136	I'll remove this from our code base - Javier, you mentioned it wasn't used anymore. Can you confirm? Just reassign this issue to me if you know that is true and I'll take care of it. Thanks.	1	1	https://issues.jboss.org/browse/JBNMAN-136
12279	JBNMAN-134	Are these settings that should then be added to the "ADMINISTRATION" tab for the management console?	0	0	https://issues.jboss.org/browse/JBNMAN-134
12280	JBNMAN-128	The mode parameter is required for a DispatchAction derived class to figure out what action method to call. Obviously, we haven't done a "multipart/form-data" in our application yet, so I'm not sure how it all works. Would it be possible to forward to /platform/Inventory.do with the mode parameter in the request URL? i.e. /platform/Inventory.do?mode=view	1	1	https://issues.jboss.org/browse/JBNMAN-128
12281	JBNMAN-128	I'm not sure why it's not picking up the mode parameter, as it's passed in directly in the query (as defined in forward directive), and not as one of the form parameters. Is it possible to forward NewAction "success" to a JSP that can compose the /platform/Inventory.do?mode=view&eid=1:xxxxx and redirect to it?	1	1	https://issues.jboss.org/browse/JBNMAN-128
12282	JBNMAN-128	Could you please expand on the options above... 1) I don't understand how to remove the action chaining? When using Struts validation the "input" attribute in the action mapping must point to either a jsp page or a tile, since we've seen that having it point to an action which checks the request e.g. "/resource/platform/Inventory.do?mode=new", does not work. How do you setup the portal and portlets correctly from a jsp or tile? 2) I was just using the new platform action as an example. I can't see a way to upload file within the HQ portal framework, while also utilizing Struts validation.	1	1	https://issues.jboss.org/browse/JBNMAN-128
12283	JBNMAN-125	Just starting to look at the JBoss plugin enhancements. Running 4.0.2 with the TrailBlazer app, I see 1 DataSourceBinding instance: MBean: org.jboss.resource.adapter.jdbc.remote.WrapperDataSourceService Name: jboss.jca:name=DefaultDS,service=DataSourceBinding 0. Attribute: JndiName = DefaultDS 1. Attribute: UseJavaContext = true 2. Attribute: JMXInvokerName = jboss:service=invoker,type=jrmp 3. Attribute: StateString = Started 4. Attribute: State = 3 5. Attribute: ConnectionManager = jboss.jca:name=DefaultDS,service=LocalTxCM 6. Attribute: BindName = java:DefaultDS 7. Attribute: Name = WrapperDataSourceService Operation: start Operation: jbossInternalLifecycle Operation: destroy Operation: invoke Operation: create Operation: stop Is this the correct MBean? We can easily provide an availability metric and control actions for start/stop for this service.	1	1	https://issues.jboss.org/browse/JBNMAN-125
12284	JBNMAN-120	Actually, we had changed it to "Notify JBN Users". Would you rather have them be "Notify Users"?	1	1	https://issues.jboss.org/browse/JBNMAN-120

#	Requirement	Questions	Answered	Successfully Answered	URL
12285	JBNMAN-119	I'm particularly interested in step3) and how it relates to our intended development of portions of the JBN. Let me provide a brief recap of the architecture we are intending to use for reading and updating the configuration attributes of resources. Then it would be great if you could describe how "custom attributes" fit in with that. Here is the suggested process for viewing configuration attributes of a resource: 1) The user clicks on the Inventory tab for a particular resource e.g. DataSource, JMS Server. 2) The HQ portal will call the APL (JBN invention) passing in the AppDefEntityID. 3) The APL layer will ask the TFS (JBN invention) to go and retrieve the configuration file associated with the resource. 4) The APL will then process the configuration to pull out the relevant attributes for the resource and return them to the HQ portal 5) The HQ portal will then display the attributes in a new section of the Inventory tab for that resource. Updating an attribute would follow a generally similar flow. So how will "custom attributes" play with this approach? Are "custom attributes" complementary or replacements for the above?	1	1	https://issues.jboss.org/browse/JBNMAN-119
12286	JBNMAN-117	Hoe does the metric bar work in terms of understanding the number of plugins to load? Is it dynamic enough, so when things get removed added it will work. -Rich. _____ From: Ryan Morgan (JIRA) [mailto:jira-events@jboss.com] Sent: Tuesday, June 07, 2005 2:14 PM To: Richard Friedman Subject: [JBoss JIRA] Work logged: (JBNMAN-117) NotReadyValve still has Hyperic HQ Issue (View Online <http://jira.jboss.com/jira/browse/JBNMAN-117>) <http://jira.jboss.com/jira/images/bluepixel.gif> Key: JBNMAN-117 <http://jira.jboss.com/jira/browse/JBNMAN-117> Type: Task <http://jira.jboss.com/jira/browse/JBNMAN-117> Task Status: Resolved<http://jira.jboss.com/jira/images/icons/status_resolved.gif> Resolved Priority: Major<http://jira.jboss.com/jira/images/icons/priority_major.gif> Major Resolution: Done Assignee: Ryan Morgan <http://jira.jboss.com/jira/secure/ViewProfile.jspa?name=rpm> Reporter: John Mazzitelli <http://jira.jboss.com/jira/secure/ViewProfile.jspa?name=mazz%40jboss.com> Time Spent: 10 minutes Operations <http://jira.jboss.com/jira/images/bluepixel.gif> <http://jira.jboss.com/jira/images/icons/bullet_creme.gif> View all <http://jira.jboss.com/jira/browse/JBNMAN-117?page=all> <http://jira.jboss.com/jira/images/icons/bullet_creme.gif> View comments <http://jira.jboss.com/jira/browse/JBNMAN-117?page=comments> <http://jira.jboss.com/jira/images/icons/bullet_creme.gif> View change history <http://jira.jboss.com/jira/browse/JBNMAN-117?page=history> NotReadyValve still has Hyperic HQ <http://jira.jboss.com/jira/browse/JBNMAN-117> Updated: 07/Jun/05 02:12 PM Created: 07/Jun/05 03:01 AM The following comment has been added to this issue: Author: Ryan Morgan <http://jira.jboss.com/jira/secure/ViewProfile.jspa?name=rpm> Date: 07/Jun/05 02:12 PM Time Logged: 10 minutes Comment: Switched the startup valve to report JBN rather than HQ. Changed the plugin deployer to report 'Deploying JBN Plugin xxx'. Also fixed the progress bar to be more accurate. The old code assumed we were deploying all of Hyperic's plugins, not just the ones JBoss has access to. Project: JBoss Network - Management <http://jira.jboss.com/jira/browse/JBNMAN> Components: Olives Description Not sure if this was already brought up or if there's a JIRA on this already. If so, sorry for the dup. Anyway, NotReadyValve.java still has "Hyperic HQ" in there. Need to "reskin" that. Do a search-n-replace on that string - its there in a few places. I encountered this on line 219 (even though the comment says we should not get there - in our case, we can get there. JBN is adding another war to the hq.ear in which case this will be reached on error in that war). <http://jira.jboss.com/jira/images/border/spacer.gif> This message was automatically generated by Atlassian JIRA <http://www.atlassian.com/c/JIRA/10140> Enterprise Edition, Version: 3.0.2-#74 - Bug/feature request <http://jira.atlassian.com/default.jsp?clicked=footer> If you think it was sent incorrectly contact one of this server's administrators <http://jira.jboss.com/jira/secure/Administrators.jspa> .	1	0	https://issues.jboss.org/browse/JBNMAN-117
12287	JBNMAN-115	this variable holds the data parsed from hq-plugin.xml the package scope is very much intentional. there are getters and OO-type access, but not to the variable itself. just look at the references to the variable, all of those get* methods can be overridden by plugins. JBossProductPlugin.getClassPath for example uses super.getClassPath and prepends jboss.installpath to each entry. why does your plugin need to directly access this variable?	1	1	https://issues.jboss.org/browse/JBNMAN-115
12288	JBNMAN-115	I haven't run into any specific use case yet that I need this. However, I noticed that some of the subclasses to GenericPlugin (like ControlPlugin in its getActions() method) go directly to this member - not thru any getters. So, i assumed I'll eventually need access to this data member as well, since the other HQ plugin types seem to need it. But because my plugin class it outside of the package of GenericPlugin, I won't be able to. Am I missing something? Where does that data member get accessed outside of its package? I did a search and couldn't find it. I am not sure why the specific plugins (like ControlPlugin and MeasurementPlugin) access that data member directly.	1	1	https://issues.jboss.org/browse/JBNMAN-115
12289	JBNMAN-114	What's the difference between a ProductPlugin and a GenericPlugin? I see ProductPlugin extends GenericPlugin. I also see ControlPlugin extends GenericPlugin (not the subclass ProductPlugin) but JBossProductPlugin extends ProductPlugin. The thing that's confusing me is that ProductPlugin is the place where the plugin types are defined (TYPE_CONTROL, TYPE_MEASURMENT, etc.). But it seems ProductPlugin is the base class for all products' plugins (like JBoss, Apache, etc.). Would have thought those TYPE_ constants would be defined in GenericPlugin and all the plugin types would extend from GenericPlugin (like ControlPlugin, MeasurementPlugin, etc.). Maybe I'm just confused about the hierarchy.	1	1	https://issues.jboss.org/browse/JBNMAN-114
12290	JBNMAN-114	So, does this mean that ProductPlugin really does not do any of the business logic? If I want to create my own plugin type (ADMIN or whatever), I need to extend GenericPlugin and in that subclass (say, AdminPlugin) I do the work. Do I need to write any classes that extend ProductPlugin to support this new Admin plugin type? I'm trying to find out what gets done "magically" for me and what I need to manually create classes for.	1	1	https://issues.jboss.org/browse/JBNMAN-114

#	Requirement	Questions	Answered	Successfully Answered	URL
12291	JBNMAN-114	Another question - this time on the plugin config in hq-plugin.xml. In the jboss plugin's hq-server.xml, I see this: <server name="JBoss" ... <config> <option name="java.naming.provider.url" description="Naming Provider URL" default="jnp://localhost:1099"/> </config> <plugin type="measurement" class="JBossMeasurementPlugin"/> ... 1) So, just to confirm - that <plugin> tag defines what subclass of GenericPlugin this defined server supports (or, if <plugin> is a child of <service>, what the defined service supports). So, in my example, it will be something like this: <plugin type="admin" class="JBossAdminPlugin" /> where JBossAdminPlugin is a subclass of my generic AdminPlugin - JBossAdminPlugin provides specific functionality to run administration tasks for the JBoss product. I could presumably have an IISAdminPlugin and ApacheAdminPlugin too. 2) That <config> element - is this that ConfigResponse?? How does my GenericPlugin subclass (aka AdminPlugin) get access to that <config> information? 3) Assuming I add my own <plugin> elements, how do I "get" the plugin instance from my CommandsServer? Assume I have a PluginManager (AdminPluginManager) that my CommandsServer has access to (I know how to do that). I just don't get how I go from a CommandsServer/PluginManager to an instance of my GenericPlugin subclass associated with a particular resource so the CommandsServer can actually ask the GenericPlugin subclass to "do" something to manipulate that resource.	1	1	https://issues.jboss.org/browse/JBNMAN-114
12292	JBNMAN-114	Ok. Sounds good. What are the APIs I should be looking at for these things. I see there is GenericPlugin.init(PluginManager, ConfigResponse). My subclass (AdminPlugin) will have its init() method called automatically by the HQ infrastructure I assume - which means that my plugin will automatically gets its ConfigResponse (as entered by the user on the config properties GUI page) without me having to do anything. Correct? I am assuming also that the ConfigResponse is read-only from the plugin perspective - the plugin never modifies ConfigResponse, the user enters the data in the inventory GUI and the persisted configresponse data is handed to the plugin via this init() method. (???) As for getting the plugin, I am looking at PluginManager (which is the object the CommandsServer has initially). I see PluginManager.getPlugin(String name). I see parameters named "name" in some of these methods, and that's what is confusing me. Is this "name" parameter a different name that that which is returned by my PluginManager's getName() method (which is something like "control" or "measurement"). I think it is - but I don't know where this "name" parameter comes from or what it is or who defines valid values for name. So, for example, my CommandsServer needs to get the AdminPlugin. I get my AdminPluginManager: AdminPluginManager apm = ... get it ... then I get the plugin: apm.getPlugin("WHAT IS THIS NAME") Not sure what that "name" parameter is and where I get it. Thanks again - this is a very informative thread - glad its being captured in JIRA :)	1	1	https://issues.jboss.org/browse/JBNMAN-114
12293	JBNMAN-114	Question about packing the product specific plugin jars (jboss-plugin.jar, apache-plugin.jar). I see there is a <hqplugin> task that is used to build the xxx-plugin.jar (right?) The way we have our source laid out, we have: com.jboss.jbossnetwork.product which is analogous to net.hyperic.hq.product. that is to say, files in our .product package are getting placed in the hq-product.jar. Now, we will have packages like: com.jboss.jbossnetwork.product.jboss com.jboss.jbossnetwork.product.apache for our product specific plugin classes. Analogous to net.hyperic.hq.product.jboss, etc. Those classes have to go in their respective xxx-plugin.jar files. In build.xml, I see this: <hqplugin name="jboss" /> How do I ensure that our plugin classes go into that jboss-plugin.xml. Same with the other products. Is the fact that our package names are different (com.jboss, not net.hyperic) going to cause problems in say the hp-plugin.xml generation? How do we get our files picked up?	1	1	https://issues.jboss.org/browse/JBNMAN-114
12294	JBNMAN-114	I just found another class (this time on the agent side) that I think needs to be modified when we add a new plugin type. AgentDaemon.getPluginManager(String) has a series of if-then-else statements - for each type it returns a data member containing that plugin manager. Also, startPluginManagers() in that class has the same type of if-then-else statements for each type. I assume we'll need to add another else-statement for each new type of plugin we add. Is that correct?	1	1	https://issues.jboss.org/browse/JBNMAN-114
12295	JBNMAN-113	I'm a little confused. Is this a recent screenshot? If so, the images weren't unpacked correctly. The zip file includes replacements for the top tabs, in-page tabs, and "tab buttons". Please confirm.	1	1	https://issues.jboss.org/browse/JBNMAN-113
12296	JBNMAN-86	as per our conf call yesterday, assigning to HQ folks so we can capture in this JIRA issue any initial information on how events/actions are generated/executed; any classes, methods we should look at as examples or APIs we should use. Is there anyplace in the data model that we may have to modify to add our own events/actions or is the current data model generic enough for us to extend?	0	0	https://issues.jboss.org/browse/JBNMAN-86

#	Requirement	Questions	Answered	Successfully Answered	URL
12297	JBNMAN-80	I want to document what is going on. I just ran some tests and its clear in my head now, I want to document it before I go into Flower-For-Algernon mode and begin to lose it all. When you "import", you are in effect adding/commiting one set of changes to the HYPERIC branch and tagging the branch after the import as Vx_x_x_xxxx_xx_xx. Since you are committing only a single set of changes (and assuming no one else is committing files on the vendor branch - which they should absolutely not be doing) you will never get Conflicts during the import (where a "Conflict" is what you typically think of as a CVS conflict - that being two sets of changes were made to one file and those changes conflict with one another - the import will still probably say there are "conflicts" but they are conflicts as compared with files in the HEAD branch, not conflicts within the HYPERIC branch). Once the import is complete, you have a new version on that branch. Just as if you did a commit on HEAD after making changes to HEAD. It's no different. If we made no changes to the HQ code, and our HEAD mimics that of the HYPERIC branch, the import will say there were no conflicts and we are done. This will never be the case, though. Now, assume I made some changes on the HEAD branch to files that are coming over on the import. This will be the case always - we are adding/removing/modifying files in the HQ code base all the time. When this happens, the import will say "N conflicts created by this import.". N is the number of files that we changed with those coming over in the import. These "conflicts" are NOT necessarily the typical CVS conflicts that require manual hand merging (although they can be). The N conflicts may simply be merged automatically using the "update -j X -j Y" command. This is similar to the situation where two developers made two set of changes and they need to be merged - in that case, the one set of changes is checked into HEAD and the other sets of changes are sitting on your local filesystem (doing an update may automatically merge the checked in changes with your local changes if they truely don't conflict). In this vendor import situation, the two changes are actually checked in - one in HEAD and one in the vendor branch (HYPERIC) - but the situation is the same. We have two files that each have changes in them and they need to be merged. When you do an "update -j -j" on a workspace with HEAD files (or a checkout -j -j), this is actually merging in two sets of changes to your HEAD copy. The first set of changes is that which occurred between the last vendor import and this current one (i.e. what did the vendor change, from V1_0_0 to V2_0_0). The second set of changes is that which occurred between the new vendor file and our copy (i.e. what did we change, from V1_0_0 to HEAD). It could be that these changes do not actually "conflict" in the traditional sense and CVS will be able to automatically merge the two changes. It will store the file in your working space and its up to you to then commit those changes. If, however, the vendor changed a line that we also changed, then a true "C" conflict has occurred and that update -j -j will act like it normally would - it writes out a file with the <<<< and >>>> characters and you have to manually merge the changes into the working space copy. In essence, update -j X -j Y gets one set of changes (between the two branches, X and Y) and merges those with the changes it sees between the latest vendor branch Y and HEAD. Thus moving HEAD from X to Y and maintaining HEAD changes. Don't forget after doing an update -j -j (or a checkout -j -j) and the merging is complete, you need to commit your changes back into HEAD.	0	0	https://issues.jboss.org/browse/JBNMAN-80
12298	JBNMAN-76	As discussed in the conference call on Nov 15, this will likely not get done in the 1.2 Final timeframe. The issue here is the upgradeability of windows installations that use the internal PostgreSQL database. Windows installations that use external PostgreSQL, or Oracle do not suffer from this limitation. The workarounds are to have the customer install on Oracle (which will usually be the case in a large scale production env). For smaller environments that use PostgreSQL, the database can be exported and re-imported into an external db. We also have instructions to upgrading the DB in place using the no-JRE installer. I will attach those instructions, or provide a link if we have one. The eventual resolution of this issue will result in a re-writing of both the text and MSI installers. Ideally, we would like the installer to just un-roll the files. Currently things are put in place using ant tasks, which makes things unnecessarily complicated. The only obstacle is moving the DBSetup process to occur during server startup, rather than in the installer, which is not a trivial task. Since both JBoss and Hyperic are in a code-freeze, it's not prudent to make sweeping changes like this so close to release.	1	1	https://issues.jboss.org/browse/JBNMAN-76
12299	JBNMAN-73	Do you still need info on this? If so, assign it back to me and we'll take care of it.	1	0	https://issues.jboss.org/browse/JBNMAN-73
12300	JBNMAN-70	Can already be done via the 'script' plugin. Currently captures availabiltiy (did it execute), return value, execution time.	1	1	https://issues.jboss.org/browse/JBNMAN-70
12301	JBNMAN-68	I think we can use password stacking and have ProductConfigService add a "webdav roles only" login module. This will free us from modifying JDBCLoginModule but will allow us to define the roles for a user of the WebDAV system. Some text I found floating around: Password stacking allows modules to skip the actual authentication and just provide supplemental roles if an earlier module has already authenticated the user. The modules require the password-stacking option to useFirstPass for this to work. Here is an example: <application-policy name="todo"> <authentication> <login-module code="org.jboss.security.auth.spi.LdapLoginModule" flag="required"> <module-option name="password-stacking">useFirstPass</module-option> <module-option name="java.naming.factory.initial"> com.sun.jndi.ldap.LdapCtxFactory </module-option> <module-option name="java.naming.provider.url"> ldap://localhost/ </module-option> <module-option name="java.naming.security.authentication"> simple </module-option> <module-option name="principalDNPrefix"> uid= </module-option> <module-option name="principalDNSuffix"> ,ou=People,o=jbossnotebook </module-option> </login-module> <login-module code="org.jboss.security.auth.spi.DatabaseServerLoginModule" flag="required"> <module-option name="password-stacking">useFirstPass</module-option> <module-option name="dsJndiName">java:/DefaultDS</module-option> <module-option name="rolesQuery"> select role, 'Roles' from USER_ROLES where login=? </module-option> </login-module> </authentication> </application-policy> Both the LDAP login module and the database login module are marked as required, and they both have the password-stacking option set. Notice that the LDAP configuration omits the roles query information and the database server module omits the principals query. It is better to provide fallback queries even when they aren't used, but it isn't required.	0	0	https://issues.jboss.org/browse/JBNMAN-68
12302	JBNMAN-66	"Update the specified 'view' method to add the portlet ".resource.NNN.manage.list.detail". - where is this specified? There is no "view" method in either ServiceManageAction or ServerManageAction. I see things like: Portal portal = Portal.createPortal("resource.server.Control.PageTitle.Edit", ".resource.service.control.edit"); Do I need to make "resource.server" and/or "resource.service" to "resource.platform"?	1	0	https://issues.jboss.org/browse/JBNMAN-66
12303	JBNMAN-45	Charles Lee logged work on JBNMAN-45: ---------- Time Worked: 3 days Remaining Estimate: 57600 (was: 144000) Time Spent: 86400 Updated license mechanism to be able to specify the feature(s) enabled. Furthermore, metrics are not automatically enabled if the license does not have monitoring feature supported.	1	1	https://issues.jboss.org/browse/JBNMAN-45

#	Requirement	Questions	Answered	Successfully Answered	URL
12304	JBNMAN-22	I saw the following screenshot http://www.microsoft.com/presspass/presskits/windowsvista/images/image003.jpg from Windows Longhorn (Vista). I think it is an excellent example of the sort of start page we are looking for. What do people think? I'm going to try to put something together which is in this vein, with a few links to the more obvious areas of JBN EM. If anyone else has suggestions for links which should go on the start page please add them to this JIRA issue.	1	1	https://issues.jboss.org/browse/JBNMAN-22
12305	JBNMAN-10	rich/doug - is this issue on track for R1? If not, please push out to R2.	0	0	https://issues.jboss.org/browse/JBNMAN-10
12306	JBNMAN-9	What about just being able to do detection and list in inventory? No metrics... Not even availability, But at least it will be an asset? -Rich.	1	1	https://issues.jboss.org/browse/JBNMAN-9
12307	JBNMAN-8	Looks like some great metrics. I don't see any instances of the hibernate MBeans in the EJB3Trail app that Rich sent us. Is there a config switch or something simple I can do to enable hibernate stats with this app?	1	1	https://issues.jboss.org/browse/JBNMAN-8
12308	JBNMAN-7	The "Restructure JBoss plugin using new PDK semantics" portion of this task is complete. In order to implement the enhancement sub tasks, would be great if folks could provide some detail for each on which MBeans I should be looking for to provide metrics and control. I've commented on what I found for the DataSource service, I don't see any MBeans for AOP or Hibernate (other than their deployers), nor do I see any jgroups related MBeans. I'm running 4.0.2 with the latest TrailBlazer app. Thanks.	0	0	https://issues.jboss.org/browse/JBNMAN-7
12309	JBNCERT-13	Is the list of configuration files to capture, configurable?	1	1	https://issues.jboss.org/browse/JBNCERT-13
12310	JBOSGI-762	I am not sure to whom can I assign this. Thomas, is it correct to assign it to you?	1	1	https://issues.jboss.org/browse/JBOSGI-762
12311	JBOSGI-699	Thomas, I see you are awaiting volunteers. Depending on the nature of the beast, I may be able offer some of my (albeit limited) skills. How were you thinking of implementing this requirement? What is currently the main obstacle? My understanding was, and I may be wrong, that the Weld OSGi OSGi/CDI integration was going to be used. I have tested the weld-osgi-2.x branch on JBosgi 2 and Wildfly 8 Alpha 3, but I could not get Weld OSGi and the Weld subsystem to run side by side. I described my attempt in more detail on the Weld forum at [https://community.jboss.org/thread/231047]. What I don't understand though, is what JBoss OSGi can do to resolve this conflict. Is it merely a packaging issue, packaging Weld OSGi as modules and enabling them as a cability? Or does it involve re-implementing some of the Weld-OSGi logic in a more Wildfly-friendly way? Or did you have an entirely different approach in mind?	0	0	https://issues.jboss.org/browse/JBOSGI-699
12312	JBOSGI-699	is this issue a duplicate of this https://issues.jboss.org/browse/JBOSGI-656 ?	1	0	https://issues.jboss.org/browse/JBOSGI-699
12313	JBOSGI-697	My concern is being compatible with future versions of the OSGi enterprise JPA requirements. If AS supplies the implementation of the above services, will that be aligned with what the next release of the specification requires? I think that depends on how the stated sub-goal of OSGi enterprise R5 is reached: {quote} * Integration of established Java EE technology into OSGi {quote} I'm not against proceeding like your suggesting, but it might also be a good time to discuss with the OSGi enterprise expert group what the longer term target is for JPA & OSGi. Maybe handling unmanaged JPA access now and managed JPA access later. However, with the unmanaged JPA access being implemented by AS code, we wouldn't use the OSGi "service creation" code in the persistence provider. I'm not sure if that is an alternative that we must support eventually or not.	0	0	https://issues.jboss.org/browse/JBOSGI-697
12314	JBOSGI-697	Do we need AS7-5476 still?	0	0	https://issues.jboss.org/browse/JBOSGI-697
12315	JBOSGI-688	Why has it been removed?	1	0	https://issues.jboss.org/browse/JBOSGI-688
12316	JBOSGI-684	Can you elaborate? I don't understand the issue.	1	0	https://issues.jboss.org/browse/JBOSGI-684
12317	JBOSGI-584	It's actually not the problem I described above. When I remove the {code}<capability name="javax.servlet.api"/>{code} from the configuration it works as expected.	0	0	https://issues.jboss.org/browse/JBOSGI-584
12318	JBOSGI-572	Why is this a problem? AS7 as well as jbosgi define the versions that are actually used.	0	0	https://issues.jboss.org/browse/JBOSGI-572
12319	JBOSGI-520	Is there an aries issue you can cross reference from here?	0	0	https://issues.jboss.org/browse/JBOSGI-520
12320	JBOSGI-462	Is this true are we abanoning the jbosgi umbrella project? I thought we keep offering a stand alone pure OSGi runtime, but remove AS7 as a supported remote test target.	1	1	https://issues.jboss.org/browse/JBOSGI-462
12321	JBOSGI-417	I guess the update should involve updating the dependency on the felix framework from 3.0.1 to 3.0.4 I noticed that the version of this component seems to follow the felix version (as it's currently 3.0.2-SNAPSHOT, while depending on felix 3.0.1). Should the version of this component stay to 3.0.2-SNAPSHOT or move to 3.0.5-SNAPSHOT?	1	1	https://issues.jboss.org/browse/JBOSGI-417
12322	JBOSGI-404	This code is now on http://github.com/jbosgi/jbosgi-resolver/commits/jbosgi404 could you please review? I can also update jbosgi-felix. Is that codebase supposed to contain the plain and unmodified felix classes?	1	1	https://issues.jboss.org/browse/JBOSGI-404
12323	JBOSGI-394	What is the set of bundles that are required to be deployed in order to be able to run the remote example testsuite infrastructure?	1	1	https://issues.jboss.org/browse/JBOSGI-394
12324	JBOSGI-393	Hmmm - jboss-osgi-webapp.jar has a dependency on org.jboss.osgi.vfs. Should we expose this package as well from the framework?	1	1	https://issues.jboss.org/browse/JBOSGI-393
12325	JBOSGI-393	How far should we go with this? If we patch jboss-as-osgi we should probably also patch jboss.osgi:common/deployment/framework/jmx/resolver/spi and vfs And how about their dependencies? I would probably be best to add the patched jar files to the modules directory structure (without replacing what was there) and then update the module.xml file to point to the new file. In that case you can roll back if needed as the original jar is still there.	1	1	https://issues.jboss.org/browse/JBOSGI-393
12326	JBOSGI-393	Wouldn't it be a better idea for people to simply get a SNAPSHOT build of as7 and run the installer over that. This is already supported in the current installer.	1	1	https://issues.jboss.org/browse/JBOSGI-393
12327	JBOSGI-388	How do I get access to the bundle context in the Arquillian based tests? I have the following code: public class StartLevelArquillianTestCase { @Inject public OSGiContainer container; @Test public void testGeneratedBundle() throws Exception { Archive<?> archive = container.getTestArchive("arq194-bundle"); Bundle bundle = container.installBundle(archive); System.out.println("*** " + bundle.getBundleContext()); It turns out that bundle.getBundleContext() returns null...	1	1	https://issues.jboss.org/browse/JBOSGI-388
12328	JBOSGI-388	How can I turn the Archive into an InputStream in the test? The following doesn't work: Archive<?> testArchive = container.getTestArchive("myBundle"); ZipExporter ze = testArchive.as(ZipExporter.class); InputStream is = ze.exportZip(); (It complains that it doesn't know ZipExporter)	1	1	https://issues.jboss.org/browse/JBOSGI-388
12329	JBOSGI-343	Can we revisit this for 1.0.0.Beta9?	1	1	https://issues.jboss.org/browse/JBOSGI-343
12330	JBOSGI-336	Can we revisit this for 1.0.0.Beta9?	1	1	https://issues.jboss.org/browse/JBOSGI-336
12331	JBOSGI-335	Example tests require docs coverage in the user guide. Could you please add a little.	1	1	https://issues.jboss.org/browse/JBOSGI-335

#	Requirement	Questions	Answered	Successfully Answered	URL
12332	JBOSGI-285	What is the jar suppose to use to fix this? I tried to use versions 0.7.1, 0.7.2 and 0.8.0 but I keep getting the same warning message and JSPs still don't work. They appear as HTML pages.	1	1	https://issues.jboss.org/browse/JBOSGI-285
12333	JBOSGI-279	We need logging update then as well. Does it need to be OSGi-zed?	1	1	https://issues.jboss.org/browse/JBOSGI-279
12334	JBOSGI-268	Thomas, could you comment on http://community.jboss.org/thread/146657#523294 Do you think webservices APIs should be fixed there?	0	0	https://issues.jboss.org/browse/JBOSGI-268
12335	JBOSGI-260	Would it be sufficient for a workaround to wrap EclipseLink in a JBoss-Module to enable JPA in the OSGI-Container (since EclipseLink is already OSGified)?	0	0	https://issues.jboss.org/browse/JBOSGI-260
12336	JBOSGI-260	If HHH-7527 is about adding OSGi support to Hibernate ORM and AS7-5405 is about supporting the OSGi Enterprise JPA specification. Which work is this jira for? Can OSGi Enterprise JPA use EE JPA (container) support or do we need a new subsystem for OSGi JPA? Does OSGi require a dynamic implementation of PersistenceProviderResolver (that could change Hibernate/EclipseLink/OpenJPA/other providers on the fly)? I thought that was required but I don't see it in the OSGi Enterprise Release 5 pdf that I'm looking at. {quote}We are most definitely not going to support multiple competing JPA providers in AS7.{quote} I assume you mean with respect to what is possible with OSGi. On the EE side, we need to support different persistence providers for various reasons (e.g. EE TCK compliance testing). Another reason to support EclipseLink, is for migrations. It is easier to migrate an application to run on JBoss AS, if you don't have to switch persistence providers right away.	0	0	https://issues.jboss.org/browse/JBOSGI-260
12337	JBOSGI-260	{quote} Which work is this jira for? {quote} It documents the availability of OSGi JPA in the JBossOSGi umbrella project. You can ignore it for your Hibernate & AS7 work. For the other questions, see [this comment|https://issues.jboss.org/browse/AS7-5405?focusedCommentId=12714861&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12714861]	1	1	https://issues.jboss.org/browse/JBOSGI-260
12338	JBOSGI-260	Is this fix part of JBoss7.1.1?	1	0	https://issues.jboss.org/browse/JBOSGI-260
12339	JBOSGI-258	Hi Andreas, I looked at the attachment and it seems quite unfinished. A number of tests are empty and also the test bundle that's used by the tests is not created by the buildsystem. Could you provide a more finished patch (or zipfile as you did) that works with the current build (git://github.com/jbosgi/jbosgi-framework.git) please? Thanks, David	1	1	https://issues.jboss.org/browse/JBOSGI-258
12340	JBOSGI-209	Could you please create a dependent issue in the MC module that needs to be looked at in order to have this issue resolved. http://www.jboss.org/index.html?module=bb&op=viewtopic&p=4270169#4270169	0	0	https://issues.jboss.org/browse/JBOSGI-209
12341	JBOSGI-179	Can we revisit this for 1.0.0.Beta9?	1	1	https://issues.jboss.org/browse/JBOSGI-179
12342	JBOSGI-178	Wow, that's impressive. I thought you'd say JBoss TM is not yet compatible with OSGi ? At least that's what I though after reading you blog entry: http://jbossosgi.blogspot.com/2009/06/jboss-osgi-runtime-as-integration.html Basically this feature request will be similar with other JEE services : datasources, JNDI, JTA, JMS Should I create a different feature request for these JEE services as well ?	1	1	https://issues.jboss.org/browse/JBOSGI-178
12343	JBOSGI-143	Could you please create a dependent issue in the MC module that needs to be looked at in order to have this issue resolved. http://www.jboss.org/index.html?module=bb&op=viewtopic&p=4270169#4270169	0	0	https://issues.jboss.org/browse/JBOSGI-143
12344	JBOSGI-81	Does it actually fail the test or are you simply seeing the exception? This message is emitted by Aries when it cannot set certain statics in the VM. The reason that it can't set these statics is because AS7 already sets them so this exception (and another one which is also logged at INFO level) can be ignored. So regardless of the exception, the functionality should still work (if the test passes).	1	1	https://issues.jboss.org/browse/JBOSGI-81
12345	JBOSGI-81	I wonder could it be a timeout or something? I haven't been able to reproduce the failure. I've tried on a Mac with JDK 1.6 and on Fedora 16 with OpenJDK 1.7. Should I try to increase the timeout and see if that helps? The Hudson machine is potentially a lot busier than my machine and may therefore take a bit longer to be finished initializing.	1	1	https://issues.jboss.org/browse/JBOSGI-81
12346	JBOSGI-81	Yes, that exception is expected. It's harmless in our case. Does the test pass?	1	1	https://issues.jboss.org/browse/JBOSGI-81
12347	JBOSGI-75	This needs to hook in with the org.jboss.modules.ModularURLStreamHandlerFactory. I also noticed that there exist a org.jboss.osgi.framework.util.VFSStreamHandlerFactory - is that still used?	1	1	https://issues.jboss.org/browse/JBOSGI-75
12348	JBOSGI-71	com.izforge.izpack.installer.ResourceNotFoundException: Cannot find named Resource: '/res/userInputLang.xml_eng' AND '/res/userInputLang.xml_eng_eng' at com.izforge.izpack.installer.ResourceManager.getLanguageResourceString(Unknown Source) at com.izforge.izpack.installer.ResourceManager.getInputStream(Unknown Source) at com.izforge.izpack.panels.UserInputPanel.init(Unknown Source) at com.izforge.izpack.panels.UserInputPanel.panelActivate(Unknown Source) at com.izforge.izpack.installer.InstallerFrame.switchPanel(Unknown Source) at com.izforge.izpack.installer.InstallerFrame.navigateNext(Unknown Source) at com.izforge.izpack.installer.InstallerFrame.navigateNext(Unknown Source) at com.izforge.izpack.installer.InstallerFrame$NavigationHandler.actionPerformed(Unknown Source)	0	0	https://issues.jboss.org/browse/JBOSGI-71
12349	PCACHE-29	Scott, Sorry, but I don't understand the issue here. Why can't you attach profiler (like JProfiler) to CacheAop? -Ben	1	0	https://issues.jboss.org/browse/PCACHE-29
12350	PCACHE-21	I lost my CVS commit rights, but anyway here's a patch to make the Set.contains() add() and remove() implementations be O(1) (more or less.) If for 2.0, it'd be nice if Fqn in PojoCache could be arbitrary types, then there wouldn't be any reason to "toString" on the Long. To fix the List implementation, I guess I would consider two cases separately. There's the LinkedList and RandomAccess styles of Lists. For LinkedList, you wouldn't have to worry about get(int i) being O(1) so removes in the middle could be fast. Could you change the implementation to detect and match closely what the user used?	1	1	https://issues.jboss.org/browse/PCACHE-21
12351	PCACHE-21	Elias, currently collection.ReplicatedSyncSetTest still fails with one more test. Not sure if it is introduced recently or not. I don't have time to fix it. But would you mind to take a look? :-) Thanks! -Ben	1	0	https://issues.jboss.org/browse/PCACHE-21
12352	PCACHE-11	Also how do we intercept the static field variable in the class scope? E.g., if I do: SomeClass.staticVar = 4; how do I intercept it?	0	0	https://issues.jboss.org/browse/PCACHE-11
12353	PCACHE-4	Why should this be coupled to jdk5 support?	1	1	https://issues.jboss.org/browse/PCACHE-4
12354	PCACHE-4	Tim, we are trying to narrow down the 1.3 roadmap. Would you like to take up this task and see JBossCache as whole on how best to do this? Maybe we can create a helper class to handle this.	1	1	https://issues.jboss.org/browse/PCACHE-4
12355	JBPORTAL-2444	Unfortunately, looking into the issue more closely, it won't be possible to change the DB creation scheme for the 2.7 / EPP 4.3 scope as this would imply a change in the Hibernate mapping. I do agree with you that the current mapping is less than ideal but the solution would require changing window states and modes from aggregrated objects to properly persisted entities (which would indeed remove all duplication and performance issue). However, this is not something that can be done in a backward compatible way and is therefore not possible to be done for the 2.7 / EPP 4.3 scope.	1	1	https://issues.jboss.org/browse/JBPORTAL-2444

#	Requirement	Questions	Answered	Successfully Answered	URL
12356	JBPORTAL-2392	Qestions: Identity admin: It is needed to add XSS testing to search input in identity admin? no error message is popped up and the user (xss string) is not found of course WSRP: producer configuration - typing a string into Registration policy class name gives me an error, but typiing it into Registration property validator class name does not. everything seems to work but the error message is missing. is it ok?	1	1	https://issues.jboss.org/browse/JBPORTAL-2392
12357	JBPORTAL-2353	what do you want know? you need sources or you should wait the next version of binary for portal	0	0	https://issues.jboss.org/browse/JBPORTAL-2353
12358	JBPORTAL-2321	Do you personally have a need for WSRP 2?	1	1	https://issues.jboss.org/browse/JBPORTAL-2321
12359	JBPORTAL-2313	some of the portlets are clear, but some of them... i don't know what they should do: cache test: render call count -> only increases after clicking "action" link in cached portlet charset portlet: testing the range no idea what else to do here encoding test: submiting some text to render phase do nothing but erase the content of "request parameter" box submiting text to action phase do change the content of both text boxes is this correct? javascript test: there is one select box on the screen and submit link... when i select something the page refresh nothing change (just the selected value is unselected). i don't know how to test it... btw. when i select something and after pressing the mouse down move it (do drag or select more values) the form is not submited... mode test: don't know what to do here page event test: proposed test case: click on the "normal" button and check if the portlet is the only one in "normal" mode with the column secure test: no idea window id test: what should "test" link do??	1	1	https://issues.jboss.org/browse/JBPORTAL-2313
12360	JBPORTAL-2313	can you add some description how to test URL test? what about secure test... can i close this jire without having it implemented?	1	1	https://issues.jboss.org/browse/JBPORTAL-2313
12361	JBPORTAL-2312	did we find a way how to test this? i think we're not strong enough to beat captcha :) can i close the issue?	1	0	https://issues.jboss.org/browse/JBPORTAL-2312
12362	JBPORTAL-2309	i've created separate class for priority I taks, and added priority II to existing portal admin testcase. what about end 2 end test cases? should they belong to existing testng testsuite or should they have another "end2end" testsuite... i've created such a testsuite, but my question if i should include them in default one still remains.	1	1	https://issues.jboss.org/browse/JBPORTAL-2309
12363	JBPORTAL-2309	i have implemented most of the phase II, but i still have few questions: - adding a listener how can I verify that a listener was successfully added? my current test adds a listener, then reloads a page and verifies that the listener is currently selected. no deeper logic is tested. - error handling this is a tricky one. i can verify that previously selected values are selected, but i guess it is not enough. any quick ideas where can i find (so i can successfully test it) these cases? for portal: When the page is unavailable When there is an error within the page When the page is not found for page: When the window is unavailable When there is an error within the window When the window is not found	1	1	https://issues.jboss.org/browse/JBPORTAL-2309
12364	JBPORTAL-2308	anything in samples doesn't need to be XSS validated. Where is "search inputs"?	1	1	https://issues.jboss.org/browse/JBPORTAL-2308
12365	JBPORTAL-2301	can this issue be closed?	1	1	https://issues.jboss.org/browse/JBPORTAL-2301
12366	JBPORTAL-2299	Where's the screenshot? :P	1	1	https://issues.jboss.org/browse/JBPORTAL-2299
12367	JBPORTAL-2283	portal hasn't a central system to describe admin roles. It let do to external applications. LDAP has an option 'defaultAdminRole' never used into ldap_identity-config.xml file so it could to be used. But about DB authorization we should to create a new option to know that information. What do you think about a new option 'defaultAdminRole' into identity-config.xml?	0	0	https://issues.jboss.org/browse/JBPORTAL-2283
12368	JBPORTAL-2277	Wesley, this would indeed be convenient if people could change the background colors (CSS) without having to modify the images, I'm not too sure about GIF transparency support for IE (I know it's an issue for PNGs). Could you give it a try ? And thanks Fernando for your latest contributions, i appreciate it !	0	0	https://issues.jboss.org/browse/JBPORTAL-2277
12369	JBPORTAL-2263	What do you propose as a fix for this? There is also the stacking tabs issue when there are more tabs than the screen is wide. So it is really a navigation rewrite when you think about it. Also, The dashboard controls would have to be reworked.	1	0	https://issues.jboss.org/browse/JBPORTAL-2263
12370	JBPORTAL-2223	the job using portal builded on linux was created here: http://hudson.qa.jboss.com/hudson/job/jboss-portal-2.7-selenium-winIE7/ there are still many failing tests which i try to resolve continuously, but some of them (i.e. uploading of files) are not supported by selenium right now. should there be a new testsuite for windows created disabling these tests? is the building of portal on windows machine necessary or we just want to do selenium testing (right now implemented as a dependency of basic selenium hudson job)	1	1	https://issues.jboss.org/browse/JBPORTAL-2223
12371	JBPORTAL-2223	Please don't disable any tests now. Let them fail. This gives us idea how much work needs to be done. Regarding fileupload, is it something we can fix if we were commit patches to Selenium? No building on windows is not necessary. You can use the zip from linux.	1	1	https://issues.jboss.org/browse/JBPORTAL-2223
12372	JBPORTAL-2223	i've renamed the job to http://hudson.qa.jboss.com/hudson/view/JBoss Portal/job/jboss-portal-2.7-selenium-winIE7- i wasn't able to find why i can't delete the workspace, this was the easiest way. unfortunatelly, lots of work must be done until i make the suite stable. anyway, we're not building portal on IE machine. i do the build in this job: http://hudson.qa.jboss.com/hudson/view/JBoss%20Portal/job/jboss-portal-2.7-selenium-tests/ save the bundle on qa machine, and IE7 job just make use of this bundle. is it enough or should i spend time on building of portal under windows?	1	1	https://issues.jboss.org/browse/JBPORTAL-2223
12373	JBPORTAL-2197	Samson, Are you still having problems building it because of the problem you mentioned on IRC? From IRC: (06:18:49 PM) samson_: just wanted to ask you what's the role of the *_DTD.xml files generated by the dtd2docbook ANT task? (06:19:26 PM) samson_: prabhat, i need to put a document root in them as they are breaking the build.	1	1	https://issues.jboss.org/browse/JBPORTAL-2197
12374	JBPORTAL-2197	I did a fresh checkout. When I do make pdf-en-US, I get eror. EEND: xml-en-US Thu Nov 13 11:36:01 CST 2008 START: test-en-US Thu Nov 13 11:36:01 CST 2008 warning: failed to load external entity "tmp/en-US/xml/.xml" make: *** [test-en-US] Error 1 I am able to build it using maven. However, the pdf generated does not have same JBoss Logo as it is for the one you generated using publican and you put at https://svn.jboss.org/repos/portal/docs/enterprise/trunk/Reference_Guide/tmp/en-US/pdf/. Is it intentional?	1	1	https://issues.jboss.org/browse/JBPORTAL-2197
12375	JBPORTAL-2197	Prabhat, should we close this issue now?	1	1	https://issues.jboss.org/browse/JBPORTAL-2197
12376	JBPORTAL-2166	As per the bug found by Mike when he entered a "dot" in portlet insttance name, can you please create a test case for it?	1	1	https://issues.jboss.org/browse/JBPORTAL-2166
12377	JBPORTAL-2143	Sohil, I have created a simple portlet that display last login date of a user. You can see it in action at http://perf11.qa.atl2.redhat.com:8080/portal/portal/default/UsersDetailPage. As you can see, I have not logged in as admin yet so you the value for admin is null but for user is some valid date. As expected this value was populated when "user" logged out. I have not applied the patch yet. The way I get this value in portlet is: user = userModule.findUserByUserName(userName); lastLoginDate = (String) userProfileModule.getProperty(user, User.INFO_USER_LAST_LOGIN_DATE); Does this suffice? I will also show you the results with your patch as well but I wanted you to see if I am in write track. But I want to follow up on JMX way as well because that will be handy in other use cases. When you said invoke UserModule remotely, did you mean twiddle or write a java client to do that? In both situation, I will need your help.	1	1	https://issues.jboss.org/browse/JBPORTAL-2143
12378	JBPORTAL-2118	Can you create a new pool called firefox with machines that have firefox installed?	1	1	https://issues.jboss.org/browse/JBPORTAL-2118

#	Requirement	Questions	Answered	Successfully Answered	URL
12379	JBPORTAL-2104	when i want to add some "google portlets" (widget/google) into the page, i can't see them, there are no choices. there were choices before (when i was developing tests for 2.6.6. GA)... is this correct? any suggestions how to test the google portlets functionality?	1	1	https://issues.jboss.org/browse/JBPORTAL-2104
12380	JBPORTAL-2103	the ids of actions in wsrp -> consumer configuration, are not indicating the number of row they are in. thus, i cannot access them by their id's and i must access them by xpath. the differ in one number, but i cannt determine, what those numbers are. on the other hand, the ids in role managment, in the table where the roles are listed are different again in only one number, but the number is indicating the row they belong to (in the table), so i need to know only the "template" of this id ,and the row number. why is the implementation of computation of ids different in those two tables?	1	1	https://issues.jboss.org/browse/JBPORTAL-2103
12381	JBPORTAL-2103	Can you be more specifiic on id part? What I am looking for is let's say I gave bea's WSDL location, which component's id you are looking for? Just a guess: <table class="registration-prop-table inactive"> <tr> <th class="nameColumn">Name</th> <th class="descColumn">Description</th> <th>Value</th> </tr> You expect tr and th to have ids, right?	1	1	https://issues.jboss.org/browse/JBPORTAL-2103
12382	JBPORTAL-2103	Can you give a try with the latest 2.6 version? By the way, do these tests also run on 2.7? When will they be in the main svn repository?	1	1	https://issues.jboss.org/browse/JBPORTAL-2103
12383	JBPORTAL-2103	Have you checked the generated code on that last update, Prabhat? It seems to me that it could result (depending on how forEach is handled) in duplicated ids in the resulting markup...	1	1	https://issues.jboss.org/browse/JBPORTAL-2103
12384	JBPORTAL-2103	i have difficulties with registering a new BEA consumer after the last svn update of JBOSS 2.6 when i register this consumer name: BEA expire: 100 address: http://wsrp.bea.com:7001/producer/producer?WSDL i get Cause: javax.portlet.PortletException: Could not find any registered converter-class by converterId : faces.convert.RegistrationProperty.Status Message: Could not find any registered converter-class by converterId : faces.convert.RegistrationProperty.Status StackTrace: javax.portlet.PortletException: Could not find any registered converter-class by converterId : faces.convert.RegistrationProperty.Status at org.apache.myfaces.portlet.MyFacesGenericPortlet.handleExceptionFromLifecycle(MyFacesGenericPortlet.java:310) at org.apache.myfaces.portlet.MyFacesGenericPortlet.facesRender(MyFacesGenericPortlet.java:502)	0	0	https://issues.jboss.org/browse/JBPORTAL-2103
12385	JBPORTAL-2102	where can i find resources on how the "search" is working in CMS?	1	1	https://issues.jboss.org/browse/JBPORTAL-2102
12386	JBPORTAL-2101	it seems, that i've found the last dynamic ID. when you "edit" the role (Role Management -> edit role) the input for new role display name has dynamic id	1	1	https://issues.jboss.org/browse/JBPORTAL-2101
12387	JBPORTAL-2086	This is something that could be provided with the out of the box portal in 2.8 (and new theme)	1	1	https://issues.jboss.org/browse/JBPORTAL-2086
12388	JBPORTAL-2063	Has any progress been made on this function?	1	0	https://issues.jboss.org/browse/JBPORTAL-2063
12389	JBPORTAL-2049	Can you mention the junit test that you talked about?	1	1	https://issues.jboss.org/browse/JBPORTAL-2049
12390	JBPORTAL-2036	Viliam, can you not create a super class called JBossPortalSeeleniumTestCase where you have common setup (browser and more) and tear methods? Obviously, the actual test would need to extend this class.	1	1	https://issues.jboss.org/browse/JBPORTAL-2036
12391	JBPORTAL-2035	Wesley, i assign this task to you, it went to Sohil for some obscure reason. Could it be the issue we had with the update ? Dan, Sorry about missing this valuable Jira	0	0	https://issues.jboss.org/browse/JBPORTAL-2035
12392	JBPORTAL-1992	denied at org.jboss.portal.core.cms.ui.admin.CMSAdminPortlet.doView(CMSAdminPortlet.java:178) at org.jboss.portlet.JBossPortlet.doDispatch(JBossPortlet.java:257) at org.jboss.portlet.JBossPortlet.render(JBossPortlet.java:244) at org.jboss.portlet.JBossPortlet.render(JBossPortlet.java:390) at org.jboss.portal.portlet.impl.jsr168.PortletContainerImpl.invokeRender(PortletContainerImpl.java:483) at org.jboss.portal.portlet.impl.jsr168.PortletContainerImpl.dispatch(PortletContainerImpl.java:405) at org.jboss.portal.portlet.container.PortletContainerInvoker$1.invoke(PortletContainerInvoker.java:86) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:131) at org.jboss.portal.core.aspects.portlet.TransactionInterceptor.org$jboss$portal$core$aspects$portlet$TransactionInterceptor$invokeRequired$aop(TransactionInterceptor.java:106) at org.jboss.portal.core.aspects.portlet.TransactionInterceptor$invokeRequired_910396445976640707 2.invokeNext(TransactionInterceptor$invokeRequired_9103964459766407072.java) at org.jboss.aspects.tx.TxPolicy.invokeInCallerTx(TxPolicy.java:126) at org.jboss.aspects.tx.TxInterceptor$Required.invoke(TxInterceptor.java:195) at org.jboss.portal.core.aspects.portlet.TransactionInterceptor$invokeRequired_910396445976640707 2.invokeNext(TransactionInterceptor$invokeRequired_9103964459766407072.java) at org.jboss.aspects.tx.TxPolicy.invokeInCallerTx(TxPolicy.java:126) at org.jboss.aspects.tx.TxInterceptor$Required.invoke(TxInterceptor.java:195) at org.jboss.portal.core.aspects.portlet.TransactionInterceptor$invokeRequired_910396445976640707 2.invokeNext(TransactionInterceptor$invokeRequired_9103964459766407072.java) at org.jboss.portal.core.aspects.portlet.TransactionInterceptor.invokeRequired(TransactionInterceptor.java) at org.jboss.portal.core.aspects.portlet.TransactionInterceptor.invoke(TransactionInterceptor.java:72) at org.jboss.portal.portlet.invocation.PortletInterceptor.invoke(PortletInterceptor.java:38) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:115) at org.jboss.portal.core.aspects.portlet.HeaderInterceptor.invoke(HeaderInterceptor.java:50) at org.jboss.portal.portlet.invocation.PortletInterceptor.invoke(PortletInterceptor.java:38) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:115) at org.jboss.portal.portlet.aspects.portlet.ProducerCacheInterceptor.invoke(ProducerCacheInterceptor.java:58) at org.jboss.portal.portlet.invocation.PortletInterceptor.invoke(PortletInterceptor.java:38) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:115) at org.jboss.portal.core.aspects.portlet.AjaxInterceptor.invoke(AjaxInterceptor.java:51) at org.jboss.portal.portlet.invocation.PortletInterceptor.invoke(PortletInterceptor.java:38) at	1	0	https://issues.jboss.org/browse/JBPORTAL-1992
12393	JBPORTAL-1978	note : portal is delivered with scriptaculous.js v1.7.0 In the RichFaces Faq : \richfaces-ui-3.2.0.GA\docs\faq\en\faq.pdf ... last item "1.73. Why does Java Script conflict?" "RichFaces 3.1.3 is bundled with the same ones as Prototype 1.6.0, Scriptaculous 1.8.0" In the 3.2.0, the scriptaculous.js file says : "script.aculo.us scriptaculous.js v1.8.1, Thu Jan 03 22:07:12 -0500 2008". so upgrade the portal script lib with the scriptaculous lib too...	1	1	https://issues.jboss.org/browse/JBPORTAL-1978

#	Requirement	Questions	Answered	Successfully Answered	URL
12394	JBPORTAL-1972	Hi! If everyone else is happy, I am happy...I made a few tiny changes. I had a few questions: * should "rows containing 'ExceptionPortletinstance?' in the INSTANCE_REF column" be ExceptionPortletInstance (capital I for Instance), like all of the other ones? * does the last section of "4.1.3. Portlet Names", about migrating CMS content, make sense to everyone? * "rows containing NewsPortletInstance2 in the ID column." is this correct, or should it be the "PORTLET_REF" column, like all of the other instances listed? Thanks for your help, Murray.	0	0	https://issues.jboss.org/browse/JBPORTAL-1972
12395	JBPORTAL-1972	Bolek, can i mark resolved or do you still need to look at Murray's questions ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1972
12396	JBPORTAL-1957	This is something we will definitely do for 2.7. Now for 2.6.5 we need to check: 1 - that it doesn't introduce an API change 2 - with QA that it runs successfully If 1 and 2 are ok, we will probably upgrade the jBPM libs for 2.6.5	0	0	https://issues.jboss.org/browse/JBPORTAL-1957
12397	JBPORTAL-1871	Which page names ? You mean portal pages ?	1	1	https://issues.jboss.org/browse/JBPORTAL-1871
12398	JBPORTAL-1858	The portlets are made to be configurable, we don't want to end up with tons of configuration we want to keep the highest flexibility. To achieve that flexibility we use jBPM you can use any workflow (even the one we will never think of). So a) is not possible b) that's what the documentation is for ;) the preferences are usually for end-users anyway, not for the developers. The default settings depends on needs... You can change for *your*portal. OOTB settings are set so that it's easy to start with. Thanks for the feedback.	0	0	https://issues.jboss.org/browse/JBPORTAL-1858
12399	JBPORTAL-1847	Do you think you could contribute that feature ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1847
12400	JBPORTAL-1842	Mmmm, when there is no modal screen is when you directly access a restricted page. So there is no real cancel (we could use back as a form of cancel) Is there a different use case ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1842
12401	JBPORTAL-1842	The page that the modal represents(login.jsp) has a dynamically rendered cancel button. If the page is presented in a modal, it shows the cancel. If a user goes to or is redirected to the login page, then what can they cancel? So, for usability, since this is a page, it did not make sense to have a cancel because the portal made a decision to redirect to that page. Good usability is not trying to write your own back button feature, but to let the browser do its job ;-) Although some browsers still don't do there job very well at times.	0	0	https://issues.jboss.org/browse/JBPORTAL-1842
12402	JBPORTAL-1842	Ok, so here is the question: If a user is redirected to the portal login page (either through a page flow or by clicking a link) what options should they have? We know they need a login form and button. What else? Possibly a register button? In it's simplest form we are clicking a link and going to a page and usually this does not require a cancel or js controlled back button. (i.e - Go to google.com and click on "Sign in", what options do you have? You either login or register. And to go back on track, you must click the browser's back button) and this is the same for most web apps. Maybe I am not getting the mindset of the portal user that you guys are talking about, so please help me out if I am not getting it. One thing we could use is a more intuitive splash screen for the login page, and not just a simple login box. I am open to suggestions.	1	1	https://issues.jboss.org/browse/JBPORTAL-1842
12403	JBPORTAL-1841	Try this: 1. Start up Portal. 2. Access the weather portal page as it is and see Miami weather. 3. Now enter http://localhost:8080/portal/auth/portal/default/Weather/WeatherPortletWindow?windowstate=maximized?mode=edit in browser's address directly and see what you get.	1	1	https://issues.jboss.org/browse/JBPORTAL-1841
12404	JBPORTAL-1799	Where can I find the "basic Richfaces documentation" in the bundled 2.6.3 documentation? I have found some information in the reference guide, but that's only about using the Sun JSF RI and the MyFaces implementation. I know that Richfaces/JBoss Portal has implemented a JSR-301 compliant portlet bridge somewhere (is it in the portal itself, or in the Richfaces API's?). I guess the "basic Richfaces documentation" is about using that bridge when developing JSF Richfaces portlets?	1	1	https://issues.jboss.org/browse/JBPORTAL-1799
12405	JBPORTAL-1799	Too bad you couldn't bundle the RF documentation with the 2.6.3 release. Do you have any estimates when things will be ready for the public?	1	1	https://issues.jboss.org/browse/JBPORTAL-1799
12406	JBPORTAL-1799	Now that this is "fixed" and 2.6.4 is released, where is said documentation? I downloaded the 2.6.4 bundle and binary-only versions and I can't find any documentation. Or does "documentation" simply refer to the code examples here: http://anonsvn.jboss.org/repos/portletswap/portlets/src/framework Am I crazy, or are they not included in any of the JBoss Portal 2.6.4 downloads?	1	1	https://issues.jboss.org/browse/JBPORTAL-1799
12407	JBPORTAL-1797	Do you want me to create another ticket for the missing localized strings ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1797
12408	JBPORTAL-1796	First shot was perfect for the versions ;) Identity relates to the identity services, while your case relates to the UI. But it's a no brainer, thanks for the files ! Also can you please "sign" the contributor agreement : http://labs.jboss.com/contributor/ (You will need to login first then click on "sign an agreement" make sure to select JBoss Portal as a project and if you have any question do not hesitate to ask me) So that i can apply your modifications.	1	1	https://issues.jboss.org/browse/JBPORTAL-1796
12409	JBPORTAL-1794	Can i just apply the latest ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1794
12410	JBPORTAL-1794	Luca, did you apply the patch yourself ? Can i mark this as resolved ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1794
12411	JBPORTAL-1788	What screenshots ? There are only jboss-service.xml files to modify if i remember correctly, and they are shown in the doc, so i don't understand what screenshots you are talking about. If you are talking about using the CMS workflow there are screenshots in the user guide. What is your issue exactly ?	1	1	https://issues.jboss.org/browse/JBPORTAL-1788
12412	JBPORTAL-1777	What would you suggest ? Multiple -object.xml files should demark your transactions. If you don't put this in a transaction you would loose consistency, what if the process breaks in the middle ? Maybe i misunderstood the issue though :)	0	0	https://issues.jboss.org/browse/JBPORTAL-1777
12413	JBPORTAL-1774	What's the status on this?	1	0	https://issues.jboss.org/browse/JBPORTAL-1774
12414	JBPORTAL-1754	Which part exactly is dupplicated? I can't find it	1	1	https://issues.jboss.org/browse/JBPORTAL-1754
12415	JBPORTAL-1747	I think that there should only be one guide that shows the GUI way and the XML way (which is actually expected to progressively disappear) in the same place. It doesn't make much sense to separate that information: why should an admin look at 2 different places to learn everything there is to know about Portal? Chapters could be structured with an overview, then how to proceed in the GUI then a more detailed look (if required). This is kind of the approach I took with the WSRP chapter and I think it works quite well. Moreover, looking at the reference guide, there are already some coverage of the GUI so it's kind of confusing that not everything is in there.	1	0	https://issues.jboss.org/browse/JBPORTAL-1747
12416	JBPORTAL-1730	registerMBean is in charge of registering (and throw an exception in case it's already registered or couldn't register). So the logic shouldn't be there. That said, all the calls to registerMBean are previously called with unregisterMBean which would have the same effect as what you are doing here. So in your case what is happening ? What's calling registerMBean ? What issue do you face, what's the stack trace ?	1	1	https://issues.jboss.org/browse/JBPORTAL-1730

#	Requirement	Questions	Answered	Successfully Answered	URL
12417	JBPORTAL-1730	********** I replace the IdentityServiceController, by another implementation and another JMX description ... but with the same JMXService name. The new service implementation take the place of the jboss original one (replacement). The JBoss original JMX description is put within comment (the original service is off). The replacement service is packaged in another proprietary sar. In the MyIdentityServiceControllerImpl, the IdentityServiceLoader is used exactly the same way as in the original IdentityServiceController code. The extra code in MyIdentityServiceControllerImpl is for specific initialization of the MBeans, not about registering them. When I redeploy the Sar package, this exception is thrown from the IdentityServiceLoader : javax.management.InstanceAlreadyExistsException: portal:service=Module,type=User already registered. at org.jboss.mx.server.registry.BasicMBeanRegistry.add(BasicMBeanRegistry.java:761) at org.jboss.mx.server.registry.BasicMBeanRegistry.registerMBean(BasicMBeanRegistry.java:225) at sun.reflect.GeneratedMethodAccessor1.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) ********** The code I propose will unregister the MBean if needed, before to re-register it, and this solve the problem. ********** Note about the original portal code : In the original code of the portal, it does not seems that the MBean are unregistered before the call of registerMBean. "That said, all the calls to registerMBean are previously called with unregisterMBean " : this does not seems ok... make a test : - run the original JBoss Portal, out of the box - make a touch of the jboss-service.xml files that are in the /server/default/deploy/jboss-portal.sar/portal-identity.sar/META-INF and /server/default/deploy/jboss-portal.sar/META-INF the server start redeploying ... the exception is thrown ... 23:04:21,687 WARN [ServiceController] Problem starting service portal:service=Module,type=IdentityServiceController org.jboss.portal.identity.IdentityException: Cannot initiate identity modules: at org.jboss.portal.core.identity.service.IdentityServiceControllerImpl.startService(IdentityServiceControllerImpl.java:160) at org.jboss.system.ServiceMBeanSupport.jbossInternalStart(ServiceMBeanSupport.java:289) at org.jboss.system.ServiceMBeanSupport.start(ServiceMBeanSupport.java:196) .... Caused by: javax.management.InstanceAlreadyExistsException: portal:service=Module,type=User already registered. at org.jboss.mx.server.registry.BasicMBeanRegistry.add(BasicMBeanRegistry.java:761) at org.jboss.mx.server.registry.BasicMBeanRegistry.registerMBean(BasicMBeanRegistry.java:225) at sun.reflect.GeneratedMethodAccessor1.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) so, I guess that the service is not unregistered before new one is registered ? or did I missed something ? ********** Redeploying the portal sar is not a good idea, as it breaks everything in the portal run. but when the IdentityServiceController is "replaced" by another implementation, that is in another .sar, then it is very usefull (for me) to be able to redeploy this .sar. ********** This problem is the same as the one in the former version. the Identity services have been reprogrammed with MBean instead of JMX services. so it is just about the same solution, that was forgotten to "move" to the new implementation. or I missed something big about all that MBean ? Thanks,	1	1	https://issues.jboss.org/browse/JBPORTAL-1730
12418	JBPORTAL-1713	What't the benefit over setting up JAVA_OPTS ? I guess you do that to get the libraries from repository.jboss.org ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1713
12419	JBPORTAL-1713	Possibly, what do they pass to JAVA_OPTS in the existing? Passing the above did not work for me.	0	0	https://issues.jboss.org/browse/JBPORTAL-1713
12420	JBPORTAL-1713	try ANT_OPTS instead of JAVA_OPS ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1713
12421	JBPORTAL-1712	Can we set that with a higher priority ? At the moment in the 2.6 testsuite, the 2 tests which are failing are related to this issue. For that reason, we are getting used to ignore Hudson error reports which is not good.	1	0	https://issues.jboss.org/browse/JBPORTAL-1712
12422	JBPORTAL-1637	"Mandatory" and "type" are already there :) Nice. Why is it used today ? Does the user portlet check this ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1637
12423	JBPORTAL-1637	If you want to have separate config for UserPortlet I should probably move some of those options to your file because it'll be doubled. But when I think about it.... 2 separate files with similar content - a list of properties but with different parameters... looks like a config nightmare to me :) So to add remove a property you need to put that info into two files. Maybe it would be better to workout common list of meta-data describing user property and extend current property-info.xml. The nature of the API is that its easy to extend and difficult to remove :) The question is: What does it buy us to have separate configuration for UserPortlet for stuff that can be gathered from existing metadata. There will be changes in identity in 2.8 but for this part it should be rather evolution than revolution.	0	0	https://issues.jboss.org/browse/JBPORTAL-1637
12424	JBPORTAL-1631	Note that this solution doesn't aim at only solving the CSS compatibility issues but the different kind of browsers. a) It's actually very easy to have a wrong behavior, try to install IE 5 and IE 6 on a same machine, both will be seen as the latest version you install (i guess it takes it from the windows registry) and the conditonal statement will be fooled. The idea of changing the user-agent is usally to fool the server to avoid some display problem. let's say you get a new for of Netscape called MyGreatBrowser, and the server is trying to send a IE CSS because it doesn't know about this new browser, then the client still has a way to "force" the Netscape stylesheet, it's a feature. The code-like stuff i was talking about should be for all the different themes (actually it should also have a way to define no CSS for some themes on some browsers (maybe you will only have 2 themes for the Opera on the Nintendo Wii, but 10 for traditional browsers) My example about IE8 was to say that if IE8 stop supporting this or do it in another way, you would end up with the same problem as of today, you couldn't define a special CSS for IE8. The additional processing is really cheap, 1 parsing per session is something we can offer. (We can also argue that the other browsers also have to parse this conditional statement to ignore it). Anyway, i am not against supporting the conditional statement, but i didn't see a little change that would make it work (i didn't look very hard). So if we need to do extra work, let's do it the good way, even more since we are re-designing the whole theme/layout support for next major versions. Too bad for the CSS compiler :) JBoss CSS Compiler sounded like a good project name :-D Thanks for your input, don't think i'm just against the idea, i'm trying to get the best from the original idea. I totally agree that not supporting this is a major problem for IE versions. (more than any other browser)	1	1	https://issues.jboss.org/browse/JBPORTAL-1631
12425	JBPORTAL-1618	Do we have obvious use case that it eats a lot of memory the VM ? Writing/reading a file on the filesystem will be also a bottleneck in the scalability side.	0	0	https://issues.jboss.org/browse/JBPORTAL-1618
12426	JBPORTAL-1596	We chose myFaces because we had less troubles with the bridge. What's your motivation for switching ?	1	1	https://issues.jboss.org/browse/JBPORTAL-1596
12427	JBPORTAL-1578	Is this the Content Definition window expanding when you change the content type?	1	1	https://issues.jboss.org/browse/JBPORTAL-1578
12428	JBPORTAL-1578	Can you move your changes from portal_style.css to core-admin/src/bin/portal-admin-war/style.css, please?	0	0	https://issues.jboss.org/browse/JBPORTAL-1578
12429	JBPORTAL-1556	What does it mean ? Is it still relevant ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1556
12430	JBPORTAL-1551	Has someone looked at that ?	1	0	https://issues.jboss.org/browse/JBPORTAL-1551
12431	JBPORTAL-1525	When was released the fixed version ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1525
12432	JBPORTAL-1524	can you explaiin the purpose of this task ? what would you like us to do ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1524

#	Requirement	Questions	Answered	Successfully Answered	URL
12433	JBPORTAL-1524	ok what do you call "send a diff" what tool ?	1	0	https://issues.jboss.org/browse/JBPORTAL-1524
12434	JBPORTAL-1522	What is WebKit ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1522
12435	JBPORTAL-1505	I don't see how the org.jboss.portlet.util.Parameters class is related to CommandFactory class. Can you explain ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1505
12436	JBPORTAL-1505	Sorry, I got confused in the classes. Parameters are not used in the CommandFactory itself... When updating my CommandFactory for 2.6, I got (again) that problem : how to see what params reach to the portlet action or rendering ? (list of keys ?, values ?... no params ?....) So it is still usefull to be able to reach this internal map : to get the list of keys, and view all the params it has in it. JBossRenderRequest has one. I saw a jira about a portlet interceptor, that would show the params that the portlet receive.... Same for org.jboss.portal.common.util.Base64.PortletParametersStateString that contains a private PortletParameters instance. It could be usefull to access this PortletParameters, at least in a inherited class. Without opening widely the API, my idea is : - in custom CommandFactory, put a MySpecialDebugParametersClass inheriting from Parameters (or PortletParametersStateString). - in other classes, use MySpecialDebugParametersClass special debug features to get what's going on etc... So protected would be enough. Public would be the easiest but to much and not clean api. (or with a special method, such as getInternalMap() ?). Thanks,	1	1	https://issues.jboss.org/browse/JBPORTAL-1505
12437	JBPORTAL-1498	I don't really get the point of this issue. What would you like us to change precisely ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1498
12438	JBPORTAL-1486	Patch includes: jsf-portlet/src/java/com/sun/faces/portlet/ViewHandlerImpl.java public String getActionURL(FacesContext context, String viewId) { Object r = context.getExternalContext().getResponse(); if (!(r instanceof RenderResponse)) { logger.info("PS_CSFP0064"); throw new IllegalStateException("Must be a RenderResponse"); } RenderResponse response = (RenderResponse) r; PortletURL actionURL = response.createActionURL(); actionURL.setParameter(VIEW_ID_PARAMETER, viewId); // Patch starts here // A unique value is needed while saving values in the session. // As the render response namespace is unique, it is used. actionURL.setParameter(NAME_SPACE_PARAMETER, response.getNamespace()); //End patch if(logger.isLoggable(Level.FINER)){ logger.log(Level.FINER, "PS_CSFP0065", actionURL.toString()); } return (actionURL.toString()); }	0	0	https://issues.jboss.org/browse/JBPORTAL-1486
12439	JBPORTAL-1486	I changed StringBuilder to StringBuffer so that it can compile with jdk1.4. It's almost there with couple of gotcha. It uses javax/faces/context/ExternalContext#.getRequestCharacterEncoding method which is not in jsf1.1. Same issue with .getRequestCharacterEncoding#encodeAll and javax.faces.application.StateManager#saveView. I am trying the same approach I did with facelet i.e. to proudce just this patched class and update the jar. The project page says 1.2 version depends on jsf1.2. Can we use version 1.1 instead?	0	0	https://issues.jboss.org/browse/JBPORTAL-1486
12440	JBPORTAL-1486	Can this patch be the solution to the error I get mentioned on the forum? http://www.jboss.org/index.html?module=bb&op=viewtopic&p=4060173#4060173	0	0	https://issues.jboss.org/browse/JBPORTAL-1486
12441	JBPORTAL-1466	I don't understand the statement about WAP, can we get more info about that ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1466
12442	JBPORTAL-1466	Bill, yes you have the right information about WAP related to this issue. If you need more information about what WAP does, can we please set up a meeting for that? Thanks!	0	0	https://issues.jboss.org/browse/JBPORTAL-1466
12443	JBPORTAL-1415	I can't really align them without affecting the templates (and i don't want to duplicate stuff). So i just aligned the buttons on the left side. (2 rows)	0	0	https://issues.jboss.org/browse/JBPORTAL-1415
12444	JBPORTAL-1365	In the forthcoming CR2 we have rewritten this function and we delegate it to a JSP page, we also documented it in the doco. Can you please wait CR2 release (it's going to be done beginning next week) and see if it can apply there ? thanks	1	0	https://issues.jboss.org/browse/JBPORTAL-1365
12445	JBPORTAL-1362	Thomas, can you please check this as I don't use Eclipse? Thanks in advance.	0	0	https://issues.jboss.org/browse/JBPORTAL-1362
12446	JBPORTAL-1346	So are you just suggesting just removing the get() method from the parent class and adding it to the child classes (changing the return value for the MultiValuedTestParameterValue to Collection instead of Object): MonoValuedTestParameterValue { public Object get() { return value;}} MultiValuedTestParameterValue { public Collection get() { return value;}} Is this all that all that is required (and of course altering the code that relies on these methods)? or am I missing something else here? I am just a bit confused since a Collection is an Object anyways	1	0	https://issues.jboss.org/browse/JBPORTAL-1346
12447	JBPORTAL-1346	should it only ever get the first element, or should it act like an iterators .next()?	1	0	https://issues.jboss.org/browse/JBPORTAL-1346
12448	JBPORTAL-1325	does it still happen in 2.6 ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1325
12449	JBPORTAL-1325	Antoine, did you check on 2.6 ? I'm not sure i understood your issue, so i cannot try it myself.	1	1	https://issues.jboss.org/browse/JBPORTAL-1325
12450	JBPORTAL-1324	why does values cannot be set on the portal objects rather ? the current implementation use the default object name property to find the default child object.	0	0	https://issues.jboss.org/browse/JBPORTAL-1324
12451	JBPORTAL-1324	The default object name property seems to be (working) only for a page in the portlet. I am quiet sure that it was not working when I did the customization of the factory (jbp2.4.1). (did some test,... ended with this solution). I upgraded without looking at all this for 2.6. Of course it works like that. Sure a better solution is to set it in the portals descriptor file, then the DefaultPortalCommandFactory is set up with the one defined there. For what I remember, the setting of default portal name was kind of hard coded to "default" name. Which is not a bad option... but when the default portal is "your wonderfull portal", marketing wants to have an url with "www.toto.com/.../.../your_wonderfull_portal/..." ;-) CLASSES AND SERVICE INVOLVED ? By the way : what are the classes and services involved in the process of reading the xml descriptor files, and setting the portal object definition ? I may need to add things to this process. example : define a few "windows always in pages", to avoid put these windows definition in all the pages definition. inheriting page definition would be the best... but it is not in the JSR. working at the portal container level would be a more compliant solution, no ? another solution : using some xml include file, and play with these includes (if this works). but anyway interested in tips (documentation ?) about this internal portal process . NAMING SUGGESTION Also : "default" is not a really helping word. when searching in source classes and service descriptors for "where all this is coded", default is a word that is everywhere. using a more specific word would be helpfull. avoiding "portal" also. "defPtal" ? "defaultJBPortal" ? "myPortalName" ? idem for the url with "host/portal/portal/default" all this portal is "search impossible". for my portal, I used : host/pcr/dc/myportalname pcr : portal context root dc : default command for a general jboss release, I suggest : pcr => portalCtx dc => defaultPCommand ending to url : host/portalCtx/defaultPCommand/defaultJBPortal more explicit. ok for newbe that needs to have a nice portal running out of the box helpfull for the one that want to customize helpfull for the guy who write the documentation. loosing the nice : "start with http://localhost:8080/portal" to "start with http://localhost:8080/portalCtx"	1	1	https://issues.jboss.org/browse/JBPORTAL-1324

#	Requirement	Questions	Answered	Successfully Answered	URL
12452	JBPORTAL-1324	I thought i removed all the hardcoded "default" when i closed this task. If not then we need to properly fix this instead of adding a workaround. Did you try "portal.defaultObjectName" on a recent release ? (Say 2.6.4 or later)	0	0	https://issues.jboss.org/browse/JBPORTAL-1324
12453	JBPORTAL-1323	I didn't see any related commit, was it done ?	1	1	https://issues.jboss.org/browse/JBPORTAL-1323
12454	JBPORTAL-1311	ok, thanks for the update. does it really matter for dashboard ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1311
12455	JBPORTAL-1309	I wanted to have an update on this issue for tomorrow's jboss platform call. No rush but when can we have this issue resolved? It would be nice to be able to run testsuite as well. I have a question though, wont we have issue running 2.6 in jboss 4.0.5.GA if we have both Sun JSF and MyFace bundled together?	1	1	https://issues.jboss.org/browse/JBPORTAL-1309
12456	JBPORTAL-1251	Isn't it what I have implemented in the hibernate persistence version ?	1	0	https://issues.jboss.org/browse/JBPORTAL-1251
12457	JBPORTAL-1241	How does the user proceed to login ? is it an explicit login using the user portlet ?	1	0	https://issues.jboss.org/browse/JBPORTAL-1241
12458	JBPORTAL-1241	I think we changed that behavior on purpose. We can add a configuration of the /signout in jboss-portal.sar/META-INF/jboss-service.xml that would allow you to configure it to let you to the current page. Would it be ok ?	1	0	https://issues.jboss.org/browse/JBPORTAL-1241
12459	JBPORTAL-1145	Why not, may I ask? :)	1	0	https://issues.jboss.org/browse/JBPORTAL-1145
12460	JBPORTAL-1141	Can you ask Chris to do work on that on 2.4 ?	0	0	https://issues.jboss.org/browse/JBPORTAL-1141
12461	JBPORTAL-1121	IMHO disabling DnD for shared pages is a good idea. Once users can copy a shared page to a private page, they will be able to use DnD. However, copy is not in 2.6 scope, is it?	1	1	https://issues.jboss.org/browse/JBPORTAL-1121
12462	JBPORTAL-1057	Just wanted to revive an old thread here. Are the JBoss Portal DTDs publicly available? Currently, the URLs to do the DTDs in the documentation do not resolve: http://www.jboss.org/portal/dtd/portlet-instances_2_6.dtd This has a few problems, most important to me being that our IDEs can't resolve and validate our usage of the DTDs. Thanks, Andy	1	0	https://issues.jboss.org/browse/JBPORTAL-1057
12463	JBPORTAL-1004	I'm willing to start working on it. How can we proceed?	0	0	https://issues.jboss.org/browse/JBPORTAL-1004
12464	JBPORTAL-1004	What is the current status of this issue - is anyone working on a solution already? I've been looking at the following use case/potential solution: - a single JBoss server (or cluster of servers) - Apache2 with mod_jk as the front end - multiple `VirtualHost' definitions in Apache2 - each VirtualHost has a JkMount for /portal/* and /portal-core/* - Each virtual site has it's own entry page, template, etc. However, there is no isolation between the sites - a user can construct a URL that let's them see content from any of the sites through any of the VirtualHosts. This is a viable solution if all the sites are related. If the sites are not related (e.g. if one site is for a mining company and another site is for an environment group), then such a solution would be inappropriate.	1	0	https://issues.jboss.org/browse/JBPORTAL-1004
12465	JBPORTAL-895	View the forum thread for more info. Kind regards, dj	0	0	https://issues.jboss.org/browse/JBPORTAL-895
12466	JBPORTAL-895	does it work with the myfaces version that is currently bundled with JBoss 4.0.4.GA ?	0	0	https://issues.jboss.org/browse/JBPORTAL-895
12467	JBPORTAL-895	So that would require to have JBoss Portal 2.4 only works with 4.0.4.SP1 ? What is the release date of 4.0.4.SP1 ? What about 2.2.x which will not upgrade its tomahawk version ?	1	1	https://issues.jboss.org/browse/JBPORTAL-895
12468	JBPORTAL-887	I was planning to add pagination somehow to instances and portlets. Would that improve usability as well ?	0	0	https://issues.jboss.org/browse/JBPORTAL-887
12469	JBPORTAL-882	The Seam portlet is a tricky one, the configuration is not easy. What's wrong with the deployment ?	1	1	https://issues.jboss.org/browse/JBPORTAL-882
12470	JBPORTAL-876	Implementation complete. Will it stand the test of tests? :)	0	0	https://issues.jboss.org/browse/JBPORTAL-876
12471	JBPORTAL-861	What is the encoding of the current file ? Could you use native2ascii to produce a file that would work will all encoding systems.	0	0	https://issues.jboss.org/browse/JBPORTAL-861
12472	JBPORTAL-860	Is it going to be a never closed issue ? which descriptors are targetted by this issue ?	0	0	https://issues.jboss.org/browse/JBPORTAL-860
12473	JBPORTAL-821	This is fixed when using JBoss WS bundled with 4.0.5. Question is what do we do for 4.0.4?	0	0	https://issues.jboss.org/browse/JBPORTAL-821
12474	JBPORTAL-771	We have implemented not a solution of page inheritance, but a solution of "page copy", with possible recursion (meaning copy all sub-pages) and portlet windows. What is the best way to contribute this changes?	1	0	https://issues.jboss.org/browse/JBPORTAL-771

#	Requirement	Questions	Answered	Successfully Answered	URL
12475	JBPORTAL-748	4.0.4.CR2/server/default/deploy/jbossweb-tomcat55.sar/META-INF/jboss-service.xml: <attribute name="Java2ClassLoadingCompliance">true</attribute> <attribute name="UseJBossWebLoader">true</attribute> 02:12:07,093 ERROR [PortalServlet] Invocation exception org.jboss.portal.common.invocation.InvocationException: org.jboss.portal.theme.LayoutException: java.lang.ClassNo tFoundException: org.jboss.portal.theme.impl.render.response.JSPFragmentAssembler at org.jboss.portal.core.command.MarkupCommand.execute(MarkupCommand.java:231) at org.jboss.portal.core.command.ControllerCommand.dispatch(ControllerCommand.java:93) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:142) at org.jboss.portal.core.aspects.controller.EventBroadcasterInterceptor.invoke(EventBroadcasterIntercepto r.java:84) at org.jboss.portal.core.command.CommandInterceptor.invoke(CommandInterceptor.java:37) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:132) at org.jboss.portal.core.aspects.controller.PageNavigationInterceptor.invoke(PageNavigationInterceptor.ja va:77) at org.jboss.portal.core.command.CommandInterceptor.invoke(CommandInterceptor.java:37) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:132) at org.jboss.portal.core.aspects.controller.PolicyEnforcementInterceptor.invoke(PolicyEnforcementIntercep tor.java:80) at org.jboss.portal.core.command.CommandInterceptor.invoke(CommandInterceptor.java:37) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:132) at org.jboss.portal.core.aspects.controller.PortalNodeInterceptor.invoke(PortalNodeInterceptor.java:59) at org.jboss.portal.core.command.CommandInterceptor.invoke(CommandInterceptor.java:37) at org.jboss.portal.common.invocation.Invocation.invokeNext(Invocation.java:132) at org.jboss.portal.common.invocation.Invocation.invoke(Invocation.java:177) at org.jboss.portal.core.command.CommandContext.execute(CommandContext.java:102) at org.jboss.portal.core.command.ExecutionContext.execute(ExecutionContext.java:91) at org.jboss.portal.core.command.ExecutionContext.execute(ExecutionContext.java:78) at org.jboss.portal.core.CoreController.handle(CoreController.java:151) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.mx.interceptor.ReflectedDispatcher.invoke(ReflectedDispatcher.java:155) at org.jboss.mx.server.Invocation.dispatch(Invocation.java:94) at org.jboss.mx.interceptor.AbstractInterceptor.invoke(AbstractInterceptor.java:133) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.interceptor.ModelMBeanOperationInterceptor.invoke(ModelMBeanOperationInterceptor.java:142) at org.jboss.mx.server.Invocation.invoke(Invocation.java:88) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:260) at	1	1	https://issues.jboss.org/browse/JBPORTAL-748
12476	JBPORTAL-738	The tests needs to be in Branch_4_0, is it?	0	0	https://issues.jboss.org/browse/JBPORTAL-738
12477	JBPORTAL-730	is that task useful ?	1	0	https://issues.jboss.org/browse/JBPORTAL-730
12478	JBPORTAL-692	Which version of myfaces are you using? i'm using 1.1.1	0	0	https://issues.jboss.org/browse/JBPORTAL-692
12479	JBPORTAL-686	How these versions of retrotranslator and backport-util-concurrent compare to what is in the src distribution of 4.0.4RC1?	1	1	https://issues.jboss.org/browse/JBPORTAL-686
12480	JBPORTAL-672	implemented and commited in cvs - but test case is commented out in code. As writen in spec: SPEC:17 Testable= true Section=PLT.5.2.4.4: A PortletException signals that an error has occurred during the processing of the request and that the portlet container should take appropriate measures to clean up the request. If a portlet throws an exception in the processAction method, all operations on the ActionResponse must be ignored and the render method must not be invoked within the current client request What is unclear - if one portlet throws PortletException in processAction() should other portlets on page be rendered()? Currently in JBoss Portal when one portlet throws PortletException http code "500" is returned. This behaviour should be configurable by admin	0	0	https://issues.jboss.org/browse/JBPORTAL-672
12481	JBPORTAL-671	The one which left to implementation for this section is: SPEC:7 Testable=true Section=PLT.5.2.2.1 The portlet container may reattempt to instantiate and initialize the portlets at any time after a failure. The exception to this rule is when an UnavailableException indicates a minimum time of unavailability. When this happens the portlet container must wait for the specified time to pass before creating and initializing a new portlet object. What will be the best approach to test it using current testframework abilities?	0	0	https://issues.jboss.org/browse/JBPORTAL-671
12482	JBPORTAL-646	Issues not done: SPEC:95 Testable=true Section=PLT.14.1 The store method must be conducted as an atomic transaction regardless of how many preference attributes have been modified. - It's rather not possible now to test if this was an atomic transaction.... SPEC:98 Testable=true Section=PLT.14.1 The PortletPreferences object must reflect the current values of the persistent store when the portlet container invokes the processAction and render methods of the portlet. - It's hard to understand what is it about. How it can reflect to not current values? SPEC:94 Testable=true Section=PLT.14.1 The store method must persist all the changes made to the PortletPreferences object in the persistent store - this is tested in every other test case of this chapter so it does not make sense to test it separately	0	0	https://issues.jboss.org/browse/JBPORTAL-646
12483	JBPORTAL-636	SPEC:71 Testable=true Section=PLT.12.2.2 Is not testable as after propagating IllegalStateException to Portlet Container '500' code is return so there is no possibility to end test.	0	0	https://issues.jboss.org/browse/JBPORTAL-636
12484	JBPORTAL-636	Issues not done in "PLT.14 Portlet Preferences" chapter: SPEC:95 Testable=true Section=PLT.14.1 The store method must be conducted as an atomic transaction regardless of how many preference attributes have been modified. - It's rather not possible now to test if this was an atomic transaction.... SPEC:98 Testable=true Section=PLT.14.1 The PortletPreferences object must reflect the current values of the persistent store when the portlet container invokes the processAction and render methods of the portlet. - It's hard to understand what is it about. How it can reflect to not current values? SPEC:94 Testable=true Section=PLT.14.1 The store method must persist all the changes made to the PortletPreferences object in the persistent store - this is tested in every other test case of this chapter so it does not make sense to test it separately	0	0	https://issues.jboss.org/browse/JBPORTAL-636
12485	JBPORTAL-630	Can't you use the org.jboss.portal.node.* API packages for this ?	1	0	https://issues.jboss.org/browse/JBPORTAL-630

#	Requirement	Questions	Answered	Successfully Answered	URL
12486	JBPORTAL-620	class and use the current policy (JACC policy) to create the permissions associated to the class that it wants to load 3/ The SecureClassLoader delegates the getPermission to the UCL that overloads it 4/ The policy uses the current jaccc policycontext associated with the current thread to get the permission for the protection domain and finds the lazy permission collection that portal uses 5/ BANG : I made our lazy permission collection elements() throw an unsupported operation exception (so I can track down its usage and I did not implement it) 6/ the exception is swallowed by the classloadingtask : something like : setError(Throwable t) { ... if (trace) { log.error("blah blah" + t); } } it helped me to change the log statement to : log.error("blah blah", t); _____ here is stacktrace I got then : java.lang.UnsupportedOperationException: not implemented at org.jboss.portal.security2.impl.jacc.JACCPortalPermissionCollection.elements(JACCPortalPermissionCollection.java:60) at java.security.PermissionsEnumerator.getNextEnumWithMore(Permissions.java:439) at java.security.PermissionsEnumerator.hasMoreElements(Permissions.java:416) at org.jboss.security.jacc.ContextPolicy.getPermissions(ContextPolicy.java:67) at org.jboss.security.jacc.DelegatingPolicy.getPermissions(DelegatingPolicy.java:237) at org.jboss.mx.loading.UnifiedClassLoader.getPermissions(UnifiedClassLoader.java:240) at java.security.SecureClassLoader.getProtectionDomain(SecureClassLoader.java:155) at java.security.SecureClassLoader.defineClass(SecureClassLoader.java:123) at java.net.URLClassLoader.defineClass(URLClassLoader.java:251) at java.net.URLClassLoader.access$100(URLClassLoader.java:55) at java.net.URLClassLoader$1.run(URLClassLoader.java:194) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:187) at org.jboss.mx.loading.RepositoryClassLoader.findClassLocally(RepositoryClassLoader.java:648) at org.jboss.mx.loading.RepositoryClassLoader.findClass(RepositoryClassLoader.java:628) at java.lang.ClassLoader.loadClass(ClassLoader.java:289) at org.jboss.mx.loading.RepositoryClassLoader.loadClassLocally(RepositoryClassLoader.java:177) at org.jboss.mx.loading.ClassLoadingTask$ThreadTask.run(ClassLoadingTask.java:131) at org.jboss.mx.loading.LoadMgr3.nextTask(LoadMgr3.java:399) at org.jboss.mx.loading.RepositoryClassLoader.loadClassImpl(RepositoryClassLoader.java:496) at org.jboss.mx.loading.RepositoryClassLoader.loadClass(RepositoryClassLoader.java:392) at java.lang.ClassLoader.loadClass(ClassLoader.java:235) at java.lang.ClassLoader.loadClassInternal(ClassLoader.java:302) at java.lang.Class.forName0(Native Method) at java.lang.Class.forName(Class.java:141) at org.apache.jasper.JspCompilationContext.createCompiler(JspCompilationContext.java:235) at org.apache.jasper.JspCompilationContext.createCompiler(JspCompilationContext.java:213) at org.apache.jasper.JspCompilationContext.compile(JspCompilationContext.java:570) at org.apache.jasper.servlet.JspServletWrapper.service(JspServletWrapper.java:293) at org.apache.jasper.servlet.JspServlet.serviceJspFile(JspServlet.java:314) at org.apache.jasper.servlet.JspServlet.service(JspServlet.java:264) at javax.servlet.http.HttpServlet.service(HttpServlet.java:810) at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:252) at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:173) at org.jboss.portal.portlet.impl.PortletRequestDispatcherImpl.execute(PortletRequestDispatcherImpl.java:77) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native	0	0	https://issues.jboss.org/browse/JBPORTAL-620
12487	JBPORTAL-615	What's the problem exactly ? The methods moved to JBossInterceptorStackFactory.java, is that the issue ?	1	1	https://issues.jboss.org/browse/JBPORTAL-615
12488	JBPORTAL-548	Is this related to development or to your agenda ?	0	0	https://issues.jboss.org/browse/JBPORTAL-548
12489	JBPORTAL-536	We extended the object_node object on version 2.4 to at least support a display label on each object node. We have also added other properties, such as "open in external window" and "external url", etc... How can we contribute this changes?	0	0	https://issues.jboss.org/browse/JBPORTAL-536
12490	JBPORTAL-490	Julien, Is there anything I can do to help you get your content on Labs certified? Let me know if you need help with anything. Cheers, Damon	1	1	https://issues.jboss.org/browse/JBPORTAL-490
12491	JBPORTAL-465	Does the fix for 535 resolves this ?	0	0	https://issues.jboss.org/browse/JBPORTAL-465
12492	JBPORTAL-455	I have posted on forum about this : http://jboss.org/index.html?module=bb&op=viewtopic&p=4037894#4037894 question : How to workaround this, now ? What class/service are responsible for reading the descriptor and set the role visible in the portlets ? To overload it, knowing it is not JSR-168 compliant... Thanks,	0	0	https://issues.jboss.org/browse/JBPORTAL-455
12493	JBPORTAL-388	Patches to allow the portal (2.0 branch) components to be imported into Eclipse without errors. Is the setup module still used within portal? If not instead of applying the patch for that module the .classpath and .project files should be removed to prevent it from being imported into Eclipse as there are current build errors in the source.	0	0	https://issues.jboss.org/browse/JBPORTAL-388
12494	JBPORTAL-341	Since creating the original patched there has been some ongoing development that has affected the core and cms projects. The attached archive 'Eclipse Patches 002.zip' contains the required patches. At the same time would it be possible to remove the .classpath and .project files from the 'bridge' module. This module is no longer compiled as part of the portal build and it's contents are out of date so is full of errors when imported into Eclipse.	0	0	https://issues.jboss.org/browse/JBPORTAL-341
12495	JBPORTAL-298	solved 2) as well. The portlet can declare a header-content element in jboss-portlet.xml which in turn will be picked up by a component container plugin. in jboss-portlet.xml add markup like: <portlet-app> <portlet> <portlet-name>LinkList</portlet-name> <header-content> <link rel="stylesheet" type="text/css" href="/portlet-styles/LinkList.css" title="" media="screen" /> </header-content> </portlet> </portlet-app> the layout of the page / portal that will render this portlet needs to contain a tag that picks up the content and injects it: <!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN" "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd"> <%@ taglib uri="/WEB-INF/theme/portal-layout.tld" prefix="p" %> <html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en"> <head> .... <p:headerContent /> </head> ... The tag takes care of resolving the context path etc. so that the link url will be formed appropriately.	0	0	https://issues.jboss.org/browse/JBPORTAL-298
12496	JBPORTAL-286	What do you mean by session replication?	1	1	https://issues.jboss.org/browse/JBPORTAL-286
12497	JBPORTAL-283	Kinda curious, are you guys going to leverage the work on wsrp4j on this?	1	1	https://issues.jboss.org/browse/JBPORTAL-283
12498	JBPORTAL-278	the table renderer was removed (since it has no benefit over the div renderer) improved the divRenderer, and the theme api not sure about the i18n. What are the requirements ? Can we create that as a separate issue: localization and internationalization of layouts and themes ? I definitely need a required feature set to implement this!	0	0	https://issues.jboss.org/browse/JBPORTAL-278

#	Requirement	Questions	Answered	Successfully Answered	URL
12499	JBPORTAL-245	Overall it looks good and the presentation is great. Thanks for the great work! However, I am reading the doc from as a beginner (no knowledge about JBoss portal, only used portals like My Yahoo). So, there are several things I feel need to be clarified in the doc. Here are some specific comments: 1. My understanding is that the portal consists of four parts: the core portal server and database; the core portlets (user, role, permission, CRM); the JBoss packaged portlets (forum); and user installed portlets. I think it will be good if we can make the distinctions clear and explain them in the "overview" and even the "feature list" sections. 2. Is there a distinction between the "admin view" and "user view"? For instance, the portal user should not see the user/role/permission admin interface in normal use. Is that right? It would be nice to show a screen shot of what the portal user see, explain how to get to the admin screen, and show a screen shot of what the portal admin see. 3. Please explain how to install third party portlets. A hello world example will do. 4. I would also like to see discussion of some basic portal features. For instance, can each user customize the portal for his/her preference (i.e., to hide some portlets and re-arrange windows locations). If so, how to do it?	0	0	https://issues.jboss.org/browse/JBPORTAL-245
12500	JBPORTAL-224	define a set or preferences that modify the behaviour of the forums that the user can override to personnalize it. like - how many topic per page - how many post per page - what is the reading direction of the posts : old post to young post or the reverse etc...	1	1	https://issues.jboss.org/browse/JBPORTAL-224
12501	JBPORTAL-213	In adding the Managing Users functionality to the Roles portlet, I decided it needed a redesign. Hence the patch.	0	0	https://issues.jboss.org/browse/JBPORTAL-213
12502	JBPORTAL-184	Looks good. Several notes: 1) I assume exposing the password in the URL is just to illustrate the parameters, but in the implementation the login form will do POST instead of GET. 2) Why is the functionality of creating a ticket token in a filter StoreTicketFilter? Isn't it a one time action following successful authentication?	1	0	https://issues.jboss.org/browse/JBPORTAL-184
12503	JBPORTAL-184	The "Affects version/s" is set to 2.0 Beta, but the "Fix Version/s" is blank. Is it intentional? Only when "Fix version/s" is set a feature will show in the roadmap and release notes.	1	0	https://issues.jboss.org/browse/JBPORTAL-184
12504	JBPORTAL-184	This issue has priority "major". This means that it should not be mentioned in marketing material. Is that the intention? http://www.jboss.org/wiki/Wiki.jsp?page=JBossJIRAIssueTypes	1	0	https://issues.jboss.org/browse/JBPORTAL-184
12505	JBPORTAL-177	Blojsom uses the BSD licence, is it not compatible with LGPL. I wonder.	0	0	https://issues.jboss.org/browse/JBPORTAL-177
12506	JBPORTAL-55	So what i understand from it is that you would like to store an email template in a text file. Then there would be different properties files for internationalizing that template. Some application would read the text file or store its text in memory at startup and apply the correct locale and send out an email. Do you think storing it in a text file is the right thing to do , or would it be a better idea to store the text in the database?	0	0	https://issues.jboss.org/browse/JBPORTAL-55
12507	JBPORTAL-32	Where is the Wiki page or any other document detailing the requirements for this feature? Even inline doc within this issue would suffice.	1	1	https://issues.jboss.org/browse/JBPORTAL-32
12508	JBPORTAL-25	does org.jboss.portal.server.app.spi.UserContext makes sense ?	0	0	https://issues.jboss.org/browse/JBPORTAL-25
12509	JBPORTAL-6	Is there a more detailed description of the functionality of user operations?	1	1	https://issues.jboss.org/browse/JBPORTAL-6
12510	JBPORTAL-2	Isn't this feature big and risky enough to try to fit in Beta?	1	1	https://issues.jboss.org/browse/JBPORTAL-2
12511	JBPORTAL-2	we should break clustering into several parts ? why because the portal state is complex. so that would give us more visibility on that concern, also it would help to identify what is a show stopper.	1	1	https://issues.jboss.org/browse/JBPORTAL-2
12512	JBPROFILER-1	Have this implemented for Memory Profiler already.	0	0	https://issues.jboss.org/browse/JBPROFILER-1
12513	JBTPL-15	brain confusion with JBTPL and portal, could you remove this task please ivelin ?	1	1	https://issues.jboss.org/browse/JBTPL-15
12514	JBTPL-15	What is the confusion?	1	1	https://issues.jboss.org/browse/JBTPL-15
12515	JBREFLECT-39	I'll copy the relevant part of the stack trace to here in case the blog entry disappears Caused by: java.lang.ArrayIndexOutOfBoundsException: 0 at org.jboss.reflect.plugins.MethodInfoImpl.(MethodInfoImpl.java:103) at org.jboss.reflect.plugins.introspection.ReflectMethodInfoImpl.(ReflectMethodInfoImpl.java:78) at org.jboss.reflect.plugins.introspection.IntrospectionTypeInfoFactoryImpl$2.run(IntrospectionTypeInfoFactoryImpl.java:211)	0	0	https://issues.jboss.org/browse/JBREFLECT-39
12516	RELOADED-1	What on earth is "JBoss Reloaded"?	1	1	https://issues.jboss.org/browse/RELOADED-1
12517	JBREM-1283	Is this going to be included in EAP 5.1.1 ? A one-off patch was made available for JBPAPP-6351 . I don't see an equivalent JBPAPP JIRA for EAP 5	1	0	https://issues.jboss.org/browse/JBREM-1283
12518	JBREM-1283	Hi Lami, re: "Is this going to be included in EAP 5.1.1 ?" Yes, it is fixed in Remoting release 2.5.4.SP2, which will ship with EAP 5.1.1. See JBPAPP-5985 "Upgrade to Remoting 2.5.4.SP2". -Ron	1	0	https://issues.jboss.org/browse/JBREM-1283
12519	JBREM-1191	Hey Ben, Setting "socketFactoryClassName" (Remoting.SOCKET_FACTORY_CLASS_NAME) to "org.jboss.security.ssl.JaasSecurityDomainSocketFactory" in the InvokerLocator - i.e., <attribute name="socketFactoryClassName" isParam="true">org.jboss.security.ssl.JaasSecurityDomainSocketFactory</attribute> will cause the Remoting client to create an org.jboss.security.ssl.JaasSecurityDomainSocketFactory. What am I missing?	1	1	https://issues.jboss.org/browse/JBREM-1191
12520	JBREM-1054	Sounds like an interesting idea. Could you show me a conceptual example?	1	0	https://issues.jboss.org/browse/JBREM-1054
12521	JBREM-1054	That will greatly simplify the user code. Is there any reason that we can't do this automatically without annotations?	1	0	https://issues.jboss.org/browse/JBREM-1054
12522	JBREM-1054	Then what happens when a transporter is not used in this case?	1	0	https://issues.jboss.org/browse/JBREM-1054
12523	JBREM-1054	What about @RemoteReference then? I guess it sounds more straightforward but it's up to you. :)	1	0	https://issues.jboss.org/browse/JBREM-1054
12524	JBREM-1054	Are we going to exclude transporter from 3.1.0, or just the annotation?	1	0	https://issues.jboss.org/browse/JBREM-1054
12525	JBREM-975	The following issues have been updated with fix version 2.4.0.Beta1, which reflects when the changes were applied to branch 2.x: [JBREM-653] - allow user to set content-type for http responses [JBREM-732] - When server terminates and has clients, when the server comes back up clients that survived, can't connect. Connection refused when trying to connect the control socket. [JBREM-734] - BisocketClientInvoker constructor should get parameters from InvokerLocator as well as configuration map. [JBREM-741] - Eliminate unnecessary log.warn() in BisocketServerInvoker [JBREM-748] - BisocketClientInvoker should guard agains scheduling on an expired Timer [JBREM-750] - Logger in HTTPClientInvoker should be static. [JBREM-751] - Eliminate unnecessary "Unable to process control connection:" message from BisocketServerInvoker [JBREM-757] - Implement quick Client.removeListener() for polled callbacks. [JBREM-761] - NPE in BisocketServerInvoker$ControlConnectionThread	0	0	https://issues.jboss.org/browse/JBREM-975
12526	JBREM-931	Have created a framework as follows: * server side: A connector is created for each of the transports: bisocket, http, rmi, socket * client side: A launcher periodically creates one of three kinds of clients: 1. MockEJBClient: makes a series of invocations, each of which creates a new org.jboss.remoting.Client 2. MockJBMClient: makes a series of invocations, each of which generates some number of callbacks 3. HeavyComputeClient: makes a single invocation which runs a computationally intensive thread for some period of time	0	0	https://issues.jboss.org/browse/JBREM-931

#	Requirement	Questions	Answered	Successfully Answered	URL
12527	JBREM-920	Another - java.lang.RuntimePermission accessClassInPackage.sun.net.www.protocol.http org.jboss.test.remoting.configuration.HTTPConfigurationTestCase(java_serialization_tomcat).testClientConfiguration java.security.AccessControlException: access denied (java.lang.RuntimePermission accessClassInPackage.sun.net.www.protocol.http) at java.security.AccessControlContext.checkPermission(AccessControlContext.java:269) at java.security.AccessController.checkPermission(AccessController.java:401) at java.lang.SecurityManager.checkPermission(SecurityManager.java:524) at java.lang.SecurityManager.checkPackageAccess(SecurityManager.java:1491) at java.lang.Class.checkMemberAccess(Class.java:1450) at java.lang.Class.getMethod(Class.java:983) at org.jboss.remoting.transport.http.HTTPClientInvoker.getSimulatedTimeout(HTTPClientInvoker.java:582) at org.jboss.remoting.transport.http.HTTPClientInvoker.transport(HTTPClientInvoker.java:138) at org.jboss.remoting.MicroRemoteClientInvoker.invoke(MicroRemoteClientInvoker.java:157) at org.jboss.remoting.Client.invoke(Client.java:1685) at org.jboss.remoting.Client.invoke(Client.java:589) at org.jboss.remoting.Client.invoke(Client.java:577) at org.jboss.test.remoting.configuration.HTTPConfigurationTestCase.testClientConfiguration(HTTPConfigurationTestCase.java:112) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:324) at junit.framework.TestCase.runTest(TestCase.java:154) at junit.framework.TestCase.runBare(TestCase.java:127) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:118) at junit.framework.TestSuite.runTest(TestSuite.java:208) at junit.framework.TestSuite.run(TestSuite.java:203) at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.run(JUnitTestRunner.java:297) at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.launch(JUnitTestRunner.java:672) at org.apache.tools.ant.taskdefs.optional.junit.JUnitTestRunner.main(JUnitTestRunner.java:567) Mar 7, 2008 10:45:24 PM org.apache.coyote.http11.Http11BaseProtocol destroy	1	1	https://issues.jboss.org/browse/JBREM-920
12528	JBREM-920	Many tests are passing despite access failures. I've added a "tripwire" sort of security manager that logs the violations, with interesting results. For example: java.security.AccessControlException: access denied (java.util.PropertyPermission org.jboss.logging.Logger.pluginClass read) at java.security.AccessControlContext.checkPermission(AccessControlContext.java:264) at java.security.AccessController.checkPermission(AccessController.java:427) at java.lang.SecurityManager.checkPermission(SecurityManager.java:532) at org.jboss.test.security.LoggingSecurityManager.checkPermission(LoggingSecurityManager.java:22) at java.lang.SecurityManager.checkPropertyAccess(SecurityManager.java:1285) at org.jboss.test.security.LoggingSecurityManager.checkPropertyAccess(LoggingSecurityManager.java:191) at java.lang.System.getProperty(System.java:662) at org.jboss.logging.Logger.init(Logger.java:359) Do we need an updated jboss logging? Ours is: Implementation-Version: 4.0.4beta (build: CVSTag=Branch_4_0 date=200601031254) How do we decide which version to use? I presume that we're getting the above problem because the Logger isn't doing System.getProperty within a privileged block.	1	1	https://issues.jboss.org/browse/JBREM-920
12529	JBREM-861	Will not complete - 2.2.2.SP3 will be replaced by 2.2.2.SP4.	0	0	https://issues.jboss.org/browse/JBREM-861
12530	JBREM-728	Proposed solution: Add the following line to org.jboss.remoting.transport.http.HTTPClientInvoker.useHttpURLConnection(): metadata.put(HTTPMetadataConstants.RESPONSE_HEADERS, conn.getHeaderFields()); (where the actual value of HTTPMetadataConstants.RESPONSE_HEADERS is "ResponseHeaders". As a result, there will be a single entry in the metadata map passed to org.jboss.remoting.Client.invoke() with the map of all response headers	0	0	https://issues.jboss.org/browse/JBREM-728

#	Requirement	Questions	Answered	Successfully Answered	URL
12531	JBREM-665	application code. I've run a similar, though pure Remoting, example in Windows (jdk 1.4 and jdk 1.5) and linux (jdk 1.5), and in all cases I see the following on the client side: java.io.WriteAbortedException: writing aborted; java.io.NotSerializableException: org.jboss.test.remoting.ExceptionServer$SampleInvocationHandler at java.io.ObjectInputStream.readObject0(Unknown Source) at java.io.ObjectInputStream.defaultReadFields(Unknown Source) at java.io.ObjectInputStream.readSerialData(Unknown Source) at java.io.ObjectInputStream.readOrdinaryObject(Unknown Source) at java.io.ObjectInputStream.readObject0(Unknown Source) at java.io.ObjectInputStream.readObject(Unknown Source) at org.jboss.remoting.serialization.impl.java.JavaSerializationManager.receiveObject(JavaSerializationManager.java:133) at org.jboss.remoting.marshal.serializable.SerializableUnMarshaller.read(SerializableUnMarshaller.java:95) at org.jboss.remoting.transport.socket.MicroSocketClientInvoker.versionedRead(MicroSocketClientInvoker.java:937) at org.jboss.remoting.transport.socket.MicroSocketClientInvoker.transport(MicroSocketClientInvoker.java:579) at org.jboss.remoting.MicroRemoteClientInvoker.invoke(MicroRemoteClientInvoker.java:123) at org.jboss.remoting.Client.invoke(Client.java:1550) at org.jboss.remoting.Client.invoke(Client.java:530) at org.jboss.remoting.Client.invoke(Client.java:518) at org.jboss.test.remoting.ExceptionClient.testException(ExceptionClient.java:106) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(Unknown Source) at java.lang.reflect.Method.invoke(Unknown Source) at junit.framework.TestCase.runTest(TestCase.java:154) at junit.framework.TestCase.runBare(TestCase.java:127) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:118) at junit.framework.TestSuite.runTest(TestSuite.java:208) at junit.framework.TestSuite.run(TestSuite.java:203) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.runTests(RemoteTestRunner.java:478) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.run(RemoteTestRunner.java:344) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.main(RemoteTestRunner.java:196) Caused by: java.io.NotSerializableException: org.jboss.test.remoting.ExceptionServer$SampleInvocationHandler at java.io.ObjectOutputStream.writeObject0(Unknown Source) at java.io.ObjectOutputStream.defaultWriteFields(Unknown Source) at java.io.ObjectOutputStream.writeSerialData(Unknown Source) at java.io.ObjectOutputStream.writeOrdinaryObject(Unknown Source) at java.io.ObjectOutputStream.writeObject0(Unknown Source) at java.io.ObjectOutputStream.writeObject(Unknown Source) at org.jboss.remoting.serialization.impl.java.JavaSerializationManager.sendObject(JavaSerializationManager.java:86) at org.jboss.remoting.marshal.serializable.SerializableMarshaller.write(SerializableMarshaller.java:104) at org.jboss.remoting.transport.socket.ServerThread.versionedWrite(ServerThread.java:780) at org.jboss.remoting.transport.socket.ServerThread.processInvocation(ServerThread.java:585) at	1	1	https://issues.jboss.org/browse/JBREM-665
12532	JBREM-663	Ron, is this available on the 2.x branch? If we release 2.2.0.Alpha5 off that branch, is it included?	1	0	https://issues.jboss.org/browse/JBREM-663
12533	JBREM-657	how we will disconnect a client without making a trip to the server without these methods? Also.. I guess you should open a new JIRA.. since this was already in a release.. (I guess)	0	0	https://issues.jboss.org/browse/JBREM-657
12534	JBREM-657	I mean.. .for us at this point ... this JIRA was already released... Ron reopened it and set a new version... What means... we would loose the feature from a previous version. (It was a solved problem for us)	0	0	https://issues.jboss.org/browse/JBREM-657
12535	JBREM-657	what will be the API then? How will be the API to disconnect a client without a server trip, when we know the server is not available? (Failover)	0	0	https://issues.jboss.org/browse/JBREM-657
12536	JBREM-643	Actually, the issue is still open. Here's an example of a multihome InvokerLocator: socket://multihome/?serializationtype=jboss&homes=127.0.0.1:4444!192.168.7.7 The set of host/port pairs is specified as a "!"-separated list. To create a multihome Connector, just use this new notation and create the Connector as usual: InvokerLocator locator = new InvokerLocator("socket://multihome/?serializationtype=jboss&homes=127.0.0.1:4444!192.168.7.7"); connector = new Connector(locator); The only programmatic examples, at this point, are in org.jboss.test.remoting.multihome.MultihomeTestParent Eventually, I'll add documentation to the Remoting Guide. Note, by the way, that this feature is only in Remoting 2.4.0.	0	0	https://issues.jboss.org/browse/JBREM-643
12537	JBREM-624	So I don't need to change my source file headers from their current "JBoss, Home of Professional Open Source..."? Can I just put the EULA file anywhere within the release binary?	1	1	https://issues.jboss.org/browse/JBREM-624
12538	JBREM-543	The response object type in this case is indeed java.lang.String. I have update the test case to print it out (See ServletInvokerTestClient or SSLServletInvokerTestClient). The output I get from JBossAS is: <html><head><title>Apache Tomcat/5.5.17 - Error report</title><style><!--H1 {font-family:Tahoma,Arial,sans-serif;color:white;background-color:#525D76;font-size:22px;} H2 {font-family:Tahoma,Arial,sans-serif;color:white;background-color:#525D76;font-size:16px;} H3 {font-family:Tahoma,Arial,sans-serif;color:white;background-color:#525D76;font-size:14px;} BODY {font-family:Tahoma,Arial,sans-serif;color:black;background-color:white;} B {font-family:Tahoma,Arial,sans-serif;color:white;background-color:#525D76;} P {font-family:Tahoma,Arial,sans-serif;background:white;color:black;font-size:12px;}A {color : black;}A.name {color : black;}HR {color : #525D76;}--></style> </head><body> HTTP Status 500 - Error occurred processing invocation request. <HR size="1" noshade="noshade"> type Status report message <u>Error occurred processing invocation request. </u> description <u>The server encountered an internal error (Error occurred processing invocation request.) that prevented it from fulfilling this request.</u> <HR size="1" noshade="noshade"> Apache Tomcat/5.5.17 </body></html> I agree that would be best for the http/servlet client invoker to thrown an exception as is done with other transports, but was changed a while back to this format per request from the web services team. I am re-opening this discussion with them (http://www.jboss.com/index.html?module=bb&op=viewtopic&p=3958247#3958247).	1	1	https://issues.jboss.org/browse/JBREM-543

#	Requirement	Questions	Answered	Successfully Answered	URL
12539	JBREM-543	IMHO, this has to be changed back - at least to return an Exception message (if you want, let the exception's message be that string). Ideally, have it configurable - let me say whether I want the exception to be thrown, returned as an Exception or returned as a String. I just posted a comment to that forum link you provided. My problem - in my code I don't know (nor do I want to know) what transport I'm using. If you just return a String, how do I know this was an error that occurred? If I am forced to look at the HTTP response code, that requires that my Client code know that I'm using http or servlet transport. My client doesn't know what transport its using - in fact, my customers can configure it anyway they want - they can use sslsocket, socket, http, servlet, whatever they put in the config file. (BTW: this is one of the great features of remoting - you actually don't even have to know how the messages are getting transported! I just pass an opaque locator URL that I get from a config file and connect my Client to it. I let remoting do everything for me - I just pass things in and get things back from invoke.) Forcing me to examine the HTTP response code just to know if the invocation failed or not breaks that separation.	1	1	https://issues.jboss.org/browse/JBREM-543
12540	JBREM-542	From summary: "This is limited in that requires JMX be available within the web container" I don't see how this fix really isolates you from requiring JMX. You will still need your web container to load in JMX classes won't you? Alot of remoting classes require JMX just to load (e.g. Connector implements the JMX interface MBeanRegistration; InvokerRegistry loads ServerInvoker classes which need JMX). e.g. when this new servlet runs, it makes this call: ServerInvoker[] serverInvokers = InvokerRegistry.getServerInvokers(); But ServerInvoker has alot of dependencies on JMX classes. So this servlet will still require JMX classes to be loaded. Or am I missing something?	1	1	https://issues.jboss.org/browse/JBREM-542
12541	JBREM-542	I have put back in the ability to us 'invokerName' param as well as the 'locatorUrl' (which is basically a copy of the patch provided. Only change is to not use the JBoss jmx implementation specific MBeanServerLocator. So is going to still require the servlet have access to jmx classes, but at least won't have to be the jboss specific implementation now (in the case they are using jdk 1.5 implementation). Since the Connector also requires jmx implementation classes (due to imports), this seems acceptable requirements. Also the point about not having to change config out of the box is a compelling one. web.xml config can now be: <?xml version="1.0" encoding="UTF-8"?> <!DOCTYPE web-app PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN" "http://java.sun.com/dtd/web-app_2_3.dtd"> <!-- The the JBossRemoting server invoker servlet web.xml descriptor $Id: web.xml,v 1.2 2006/07/17 03:44:00 telrod Exp $ --> <web-app> <servlet> <servlet-name>ServerInvokerServlet</servlet-name> <description>The ServerInvokerServlet receives requests via HTTP protocol from within a web container and passes it onto the ServletServerInvoker for processing. </description> <servlet-class>org.jboss.remoting.transport.servlet.web.ServerInvokerServlet</servlet-class> <init-param> <param-name>invokerName</param-name> <param-value>jboss.remoting:service=invoker,transport=servlet</param-value> <description>The servlet server invoker</description> <!-- <param-name>locatorUrl</param-name> <param-value>servlet://localhost:8080/servlet-invoker/ServerInvokerServlet</param-value> <description>The servlet server invoker locator url</description> --> </init-param> <load-on-startup>1</load-on-startup> </servlet> <servlet-mapping> <servlet-name>ServerInvokerServlet</servlet-name> <url-pattern>/ServerInvokerServlet/*</url-pattern> </servlet-mapping> </web-app> where can swap out the 'invokerName' and 'locatorUrl' params.	0	0	https://issues.jboss.org/browse/JBREM-542
12542	JBREM-455	numberOfRetries - specifies the number of times the client invoker will try to get a valid socket connection from the pool to use for invocatoins. Note, if not doing connection checking, which is the default, will not always know if is a valid connection until time of call is made (as in the case where the socket connection has timed out, which will still return true to the Socket.isConnected() method call). Default value is 3 (MicroSocketClientInvoker.MAX_RETRIES). numberOfCallRetries - specifies the number of time the client invoker will try to make an invocation while catching a java.net.SocketException. This is needed because even though the socket connection returned from the pool may appear to be valid, it is possible that it is not (due to timeout), so to protect from this, will retry the call by getting a new socket connection. How socket client invoker works in regards to socket connection checking: When a socket client invoker tries to make an invocation, it will ask its connection pool for the next available socket connection. If the pool does not have a socket connection available for use, a new one will be created. When an invocation has completed, the socket client invoker will see if pool has reached its max capacity (defined by the 'clientMaxPoolSize' property) and if it has not, it will place that socket connection in the pool, else the connection will be closed and disgarded. When the connection pool is asked for a socket connection, it can do one of two checks on the connection to see if it is valid. By default, the pool will just check the socket connection by calling its isConnected() method. This method will just return the state of the socket. It does not however indicate if the connection itself is still valid (as could have timed out, yet still be in connected state). If the 'socket.check_connection' property is turned on, the pool will first make a call to the server sending a single byte to verify that the connection is still good before returning it to the socket client invoker for use. If the connection is not good, meaning the ping to the server failed, the connection will be thrown away and will try again with the next connection in the pool (the number of retries for this is defined by the 'numberOfRetries' property). If the maximum number of retries is reached an no valid connection within the pool could be found, a new socket connection will be created. Therefore, the only way to truely know that the socket connection returned by the pool is valid for use is to check the connection first. Since the socket client invoker can not be absolutely sure that the socket connection from the pool is valid, it will allow for invocation retries. If upon making an invocation on the server a java.net.SocketException is thrown, the socket client invoker will try getting another socket connection and try to make the invocation again. The number of times the socket client invoker will do this is defined by the 'numberOfCallRetries' property. If the socket client invoker reaches its last try before erroring, it will flush the connection pool, so that a new socket connection will be created and will try the invocation one last time with this new connection. In general, not doing the connection check on a socket connection being returned from the pool is more performant as does not require and extra network call before making actual server invocation. In cases where the socket timeout is set to a large number and calls are being made frequently (so that the connections being returned from the pool have not had time to timeout), this can yeild around a 5 percent increase in performance. However, if the socket timeout is set to a low value and calls are not being made frequently, the likelyhood is that the connections within the pool have timeout and are not valid. Therefore, may require several invocation retries and eventually creating a new socket connection before are able to make a sucessful invocation, which is less peformant than the time it takes to check the connection while initially in the connection pool. There are too many factors involved, timeout setting, frequency of calls, size of pool, etc., to be able to make a general recomendation on which configuration would be best for all cases, however the default configuration is targeted for what is considered to provide the best performance under load.	0	0	https://issues.jboss.org/browse/JBREM-455
12543	JBREM-433	This can definitely be done in terms of a one-way invocation. But how would the response aggregation work in a request-reply scenario? Using a connection-oriented protocol, you'd be no better off than if you iterated over a series of connections, sending a request to each, and then waiting until a reply is received for each. It seems like you'd want to utilize UDP broadcast or multicast for this sort of thing. But in this case, you have no reliable way to figure out how many responses to expect.	1	1	https://issues.jboss.org/browse/JBREM-433
12544	JBREM-427	Is this for client, server, or both?	1	1	https://issues.jboss.org/browse/JBREM-427

#	Requirement	Questions	Answered	Successfully Answered	URL
12545	JBREM-366	Invoker Registry reference counting This fundamental change will change the InvokerRegistry to be the true gatekeeper of client invokers. When a Client is created, it will basically be a lame object (meaning it can not perform any real work), until its connect() method is called. The connect() method will call on the InvokerRegistry to get a client invoker to use. The InvokerRegistry will check to see if there has already been a client invoker created for the same invoker locator. If there has been, it will increment the count of Client's that have a reference to that client invoker instance and return it. If not, it will create a new one, set the reference count to one, and return it. When a Client disconnects (via call to disconnect() method), it will call the InvokerRegistry to destory the client invoker. The InvokerRegistry will check the number of Clients that have a reference to the client invoker instance and if is the last one, will disconnect the client invoker instance. Otherwise, it will just return. Client conversational methods (meaning requires being connected before can call it): addConnectionListener() addListener() setMarshaller() setUnMarshaller() If these methods are called before Client is connected, will throw an exception. Note, part of connecting a Client is to create/get the client invoker from the InvokerRegistry. Todos: 1. Set serialVersionUID on Client (only needed if uncomment connect() line within Client.readExternal() method. 2. Test case for multi-threaded, concurrent Client calls to server Issues: 1.Thinking about removing the Client.setInvoker() and Client.getInvoker(). Is there any real need to get this externally? Only problem is that ConnectionValidator uses the getInvoker() directly (and for good reason). 2. Wonder if should automatically connect a deserialized Client within the readExternal() method, which will require user to explicitly call connect on it or try to do it automatically (which would end up throwing serialization exception)? Maybe need a test case for this. 3. Still struggling with having to call connect() before can do invoke() call. Is this really needed? Is it really that big of a pain to do? 4. Should the client invoker ref counter be incremented when local client invoker created? Test case failures: org.jboss.test.remoting.configuration.ClientConfigurationMapTestCase - basically looks like config map passed to Client constructor should be used when creating ClientInvoker instance within InvokerRegistry (which it is), but also looks like this map should impact the end invoker locator value (adding the extra map entries as locator uri parameters). Is this correct? Don't think my changes should impact this. org.jboss.test.remoting.invoker.LocalInvokerTestCase (line160) - this test creates a Client and a Connector with the same locator uri (thus causing a LocalClientInvoker to be created). The Connector is stopped, which will cause the InvokerRegistry to disconnect the LocalClientInvoker. Then a new Connector is created with the same invoker uri and the Client tries to make another call, but fails because its connected state was set to false. So should this be the default behavior or should InvokerRegistry *not* disconnect the LocalClientLocator when the Connector is stopped? org.jboss.test.remoting.lease.LocalLeaseTestCase - invalid test case now because no longer do client leasing for local invokers - http://www.jboss.com/index.html?module=bb&op=viewtopic&p=3945122#3945122 org.jboss.test.remoting.transporter.proxy.TransporterProxyTestCase - failing because connect() is not automatically being called Client in its readExternal() method. Code cleanup: LocalClientInvoker - disconnect() Client - disconnect(), readExternal() MicroRemoteClientInvoker - disconnect()	1	1	https://issues.jboss.org/browse/JBREM-366
12546	JBREM-363	Remaining issues: 1. Need to verify that multiplex is setup to do true multiplex for callbacks 2. Should the CallbackPoller deliver the callbacks on a separate thread? If have a ton of callbacks, could slow down the timer thread (and would possibly even have multiple threads pulling callbacks at the same time). 3. PortUtil.findFreePort() is very slow (and is what is being used by default when creating internal callback server Connector).	0	0	https://issues.jboss.org/browse/JBREM-363
12547	JBREM-316	Assume you are referring to org.jboss.remoting.transport.coyote.ssl.RemotingSSLSupport which reuses code from org.apache.tomcat.util.net.SSLSupport? There is also org.jboss.remoting.transport.coyote.ssl.RemotingSSLImplementation and org.jboss.remoting.transport.coyote.ssl.RemotingServerSocketFactory, which extend tomcat classes, but does not reuse any of that code directly (just via linking). So removing the JBoss header from top of SSLSupport and leaving the rest should be ok?	1	1	https://issues.jboss.org/browse/JBREM-316
12548	JBREM-316	So I totally missed this. Since Remy included them, I thought they were original classes that were needed to comply to the tomcat interface signatures. Didn't realize they were a copy (although looks like were mostly Remy's anyways, so entitled to change license?) . So the following are files that were copied (more or less): org.apache.catalina.connector.ClientAbortException org.apache.catalina.connector.CoyoteInputStream org.apache.coyote.tomcat4.CoyoteOutputStream org.apache.catalina.connector.InputBuffer org.apache.catalina.connector.OutputBuffer I guess Remy did this so didn't have to include the jars they were in on the classpath for just the few classes that were needed? So can we just label the corresponding classes and the RemotingSSLSupport with the original Apache license and leave the rest LGPL? org.jboss.remoting.transport.coyote.CoyoteInvoker, for example, is all JBoss specific. Or is it just easier from management perspective to re-label license for all?	1	1	https://issues.jboss.org/browse/JBREM-316

#	Requirement	Questions	Answered	Successfully Answered	URL
12549	JBREM-311	The following covers which remoting jars and thirdparty jars will be needed depending on what features of remoting will be used. Remoting jars: jboss-remoting.jar (734K) - this binary contains all the remoting classes. This is the only remoting jar that is needed to perform any remoting funtion within JBoss Remoting. Since some may want to better control size of the binary footprint needed to use remoting, the remoting classes have been broken out into multiple remoting binaries based on their function. There are three categories of these binaries; core, detection, and transport. core: jboss-remoting-core.jar (204K) - contains all the core remoting classes needed for remoting to function. If not using jboss-remoting.jar, then jboss-remoting.core.jar will be required. detection: jboss-remoting-detection (40K) - contains all the remoting classes needed to perform automatic discovery of remoting servers. It includes both the jndi and multicast detector classes as well as the network registry classes. transport: jboss-remoting-socket.jar (41K) - contains all the classes needed for the socket and sslsocket transports to function as both a client and a server. jboss-remoting-socket-client.jar (21K) - contains all the classes needed for the socket and sslsocket transports to function as a client only. This means will not be able to perform any push callbacks or sending of streams using this jar. jboss-remoting-http.jar (59K) - contains all the classes needed for the http and https transports to function as a client and a server. jboss-remoting-http-client.jar (18K) - contains all the classes needed for the http, https, servlet, and sslservlet transports to function as a client only. This means will not be able to perform any push callbacks or sending of streams using this jar. jboss-remoting-servlet.jar (10K) - contains all the classes needed for the servlet or sslservlet transports to function as a server only (also requires servlet-invoker.war be deployed within web container as well) jboss-remoting-rmi.jar (20K) - contains all the classes needed for the rmi and sslrmi transports to function as a client and a server. jboss-remoting-multiplex.jar (126K) - contains all the classes needed for the multiplex and sslmultiplex transports to function as a client and a server. Use of this jar also requires jboss-remoting-socket.jar be on classpath as well. Thirdparty jars: This section covers which thirdparty jars are required based on the feature or transport to be used. Remember, whenever see jboss-remoting-XXX.jar mentioned, they can all be replaced with just the jboss-remoting.jar. all remoting servers: - jboss-remoting-core.jar - jboss-common.jar - jboss-jmx.jar - log4j.jar all remoting clients: - jboss-remoting-core.jar - jboss-common.jar - jboss-jmx.jar - log4j.jar - concurrent.jar Note: concurrent.jar needed because of org.jboss.util.id.GUID used to create session id within Client (http://jira.jboss.com/jira/browse/JBREM-549) Remoting requires the use of JMX classes. It does not require the JBoss implementation (jboss-jmx.jar) of JMX in order to function correctly, so can replace jboss-jmx.jar with another JMX implementation library (or exclude it if using jdk 1.5 or higher, which has JMX implementation built in). multicast detection - jboss-remoting-detection.jar - concurrent.jar - dom4j.jar jndi detection - jboss-remoting-detection.jar - concurrent.jar - dom4j.jar - jnpserver.jar (for jndi api classes) The dom4j.jar for use of detection is required because using jboss-jmx.jar. socket server: - jboss-remoting-socket.jar socket client - jboss-remoting-socket-client.jar http server - jboss-remoting-http.jar - tomcat-coyote.jar - tomcat-util.jar - commons-logging-api.jar - tomcat-http.jar Note: need tomcat-apr.jar and tcnative-1.dll/tcnative-1.so on system path if want to use APR based tomcat connector http client - jboss-remoting-http-client.jar servlet server - servlet-invoker.war (deployed in web container) - jboss-remoting-servlet.jar servlet client - jboss-remoting-http-client.jar rmi server and client - jboss-remoting-rmi.jar multiplex server and client - jboss-remoting-socket.jar - jboss-remoting-multiplex.jar if use JBoss Serialization, will need to include jboss-serialization.jar and trove.jar	1	1	https://issues.jboss.org/browse/JBREM-311
12550	JBREM-305	For the user, of course they should be able to use the JAX-WS API - but Remoting does not provide this API. My question is, what does Remoting add here? It sounds like you want Remoting to act not so much as a remote invocation layer, but rather more as a generic transport facade. Is this accurate?	1	1	https://issues.jboss.org/browse/JBREM-305
12551	JBREM-288	Ron, can you add an entry for the multiplex transport and then close this issue? Thanks. -Tom	1	1	https://issues.jboss.org/browse/JBREM-288
12552	JBREM-257	Per my dev mailing list reply: "Remoting will throw the server side exception thrown by the server-side handler as though it was called locally. This is by design and how I believe it should be. Remoting should not force the client caller into handling application exceptions via a different process than calling it locally. This makes it easier for the client caller and is sometimes actually a local call anyways (if running within the same vm). If there is an error that comes from remoting, it will be thrown as a runtime exception and will have the stack trace indicating this came from a remoting error (marshalling error, etc.). I see no good argument for it to be otherwise. " So what is the use case for wanting client stack trace? If is thrown from server, then knowing the call stack on the client side does not seem to be relevant as the issue is on the server side. Is fairly easy to change this behavior to wrap the server side exception in another exception and throw that on the client side, but now need to define what that exception would be. For example, throw RemoteException, like is done with RMI, but now are forcing client user to now know if is a remote vs. local call, which, seems to me to be unnecessary.	1	1	https://issues.jboss.org/browse/JBREM-257
12553	JBREM-257	Why is this needed? Because yes, the exception was initiated on server, but you won't know what piece of code on the client initiated the remote invocation. Also *DO NOT* wrap the exception, why? Well, for EJB for instance, you may be throwing an application exception that you want to be able to catch in the client. Really, setStackTrace is simple.	1	1	https://issues.jboss.org/browse/JBREM-257
12554	JBREM-255	The 'ResponseCode' is present, but is it intenional that there still is key/value pair null=[HTTP/1.1 200 OK] ? 2006-01-16 14:52:09,531 DEBUG [org.jboss.ws.binding.soap.SOAPMessageUnMarshaller] getMimeHeaders from: {X-Powered-By=[Servlet 2.4; JBoss-5.0.0alpha (build: CVSTag=HEAD date=200601101256)/Tomcat-5.5], ResponseCodeMessage=OK, null=[HTTP/1.1 200 OK], Date=[Mon, 16 Jan 2006 13:52:09 GMT], Content-Type=[text/xml], Server=[Apache-Coyote/1.1], HEADER={SOAPAction=", Content-Type=text/xml}, Transfer-Encoding=[chunked], ResponseCode=200}	1	1	https://issues.jboss.org/browse/JBREM-255
12555	JBREM-237	In which environment are you talking (meaning which versin of jboss, which jmx remoting implementation, which jdk, etc.)?	1	0	https://issues.jboss.org/browse/JBREM-237
12556	JBREM-233	Forgot another question. What information do you want about the client? I won't have much available by default, other than client session id and subsystem. You need anything else?	1	1	https://issues.jboss.org/browse/JBREM-233
12557	JBREM-233	I am unsure as to what the API looks like for this. On the server side I register a ConnectionListener, but on what object do I register it? someObject.registerConnectionListener(myServerSideConnectionListener) What is the type of someObject and how do I get a reference to it? Any chance of an example, also more info on the structure of the "client info" would be good.	1	1	https://issues.jboss.org/browse/JBREM-233
12558	JBREM-233	Tom, one thing I don't understand. In our case (JBossMessaging) we call connectionFactory.createConnection() to create a JMS connection. This basically creates a load of JBossMessaging specific objects on the server side, then returns. Then on the client side we create a remoting Client object that represents the physical JMS connection, which is then used for any subsequent traffic for the connection. If no invocations are made on that client object, and then the connection fails, e.g. the client process crashes, then how does the server know the client session id for the client that has failed? If the server doesn't know the client session id then it can't clear up the server side resources. How can we get the client session id plus some arbitrary meta data (in our case the jms connection id) from the client to the server before any invocations have been made? I'm guessing this needs to be done in the Client.connect() method. We need the meta-data so we can map the remoting session id to our connection id, otherwise it is meaningless to us. Am I missing something?	1	1	https://issues.jboss.org/browse/JBREM-233

#	Requirement	Questions	Answered	Successfully Answered	URL
12559	JBREM-233	Once the connect() method is called on the remoting Client class, it will send an initial ping request to the server. Technically, this will happen even if don't call connect() and just start calling invoke(), as will call connect() for you on the first invoke() method. This initial ping request will send the client session id (which is a GUID that is created when the Client is constructed). You are right that there is no good way to correlate this client session id with anything releated to users environment. Since Ron added the code to include metadata Map when creating the client, will see if I can include that information with the ping request, as well as the Client instance that is passed to the ConnectionListener's handleConnectionException() method. Will that work for you?	1	1	https://issues.jboss.org/browse/JBREM-233
12560	JBREM-233	I think there is an edge case here: If the session id/meta-data is passed on the first ping, then what happens if the connection fails after it is established but before the first ping is sent? In that case we still need to clear up server side resources. I'm not sure how we would deal with that apart from doing an application level ping which I'd like to avoid.	1	1	https://issues.jboss.org/browse/JBREM-233
12561	JBREM-233	If I create a Client with some metadata using the new constructor, is that metadata passed on every invocation made using that client? I'm assuming it's only passed on the first invocation, since that's the only time it's needed. (Otherwise we are needlessly pushing data over the wire)	1	0	https://issues.jboss.org/browse/JBREM-233
12562	JBREM-195	Think this was due to a particular implementation problem within the rmi invoker implementation, which should be corrected now. Do you have a specific test case I could use to verify this? Thanks. -Tom	1	1	https://issues.jboss.org/browse/JBREM-195
12563	JBREM-195	Sorry. Deleted previous comment which was: "I wrote a simple test case to try and reproduce (using raw rmi invocations vs jboss remoting rmi). Although is a little better now than was previously, is stil about twice as slow. We have not done any profiling to date to see where can optimize more. Is there any reason that you could not use socket transport? Also same test using raw socket invocations vs jboss remoting sockets and is much better in terms of performance. The test (which just sends and receives simple String), doing 250K calls, produced following on my laptop: rmi remoting: 69821 milliseconds raw rmi: 35661 milliseconds socket remoting: 75118 milliseconds raw sockets: 81277 milliseconds" Was meaning to update the information. These results are actually incorrect. Was meaning to update with correct results. They are hardly comforting. Will be working on optimizations and profiling to get better results. Since this was just a very quick test case I threw together, was wondering if you might have a better "apples" to "apples" test case between plain old rmi and remoting rmi that I could use when profiling? Thanks. -Tom Update results are: rmi remoting: 69821 milliseconds raw rmi: 35661 milliseconds socket remoting (with java serialization): 74567 milliseconds socket remoting (with jboss serialization: 69470 milliseconds raw sockets (1 thread accepting request): 47538 milliseconds raw sockets (10 thread accepting request): 47588 milliseconds	1	1	https://issues.jboss.org/browse/JBREM-195
12564	JBREM-188	Per http://www.ietf.org/rfc/rfc2616.txt?number=2616 9.4 HEAD The HEAD method is identical to GET except that the server MUST NOT return a message-body in the response. The metainformation contained in the HTTP headers in response to a HEAD request SHOULD be identical to the information sent in response to a GET request. This method can be used for obtaining metainformation about the entity implied by the request without transferring the entity-body itself. This method is often used for testing hypertext links for validity, accessibility, and recent modification. The response to a HEAD request MAY be cacheable in the sense that the information contained in the response MAY be used to update a previously cached entity from that resource. If the new field values indicate that the cached entity differs from the current entity (as would be indicated by a change in Content-Length, Content-MD5, ETag or Last-Modified), then the cache MUST treat the cache entry as stale.	1	1	https://issues.jboss.org/browse/JBREM-188
12565	JBREM-185	How was this resolved, where should the keepAliveTimeout be set? Documentation still indicates that keep-alive timeout is not configurable	0	0	https://issues.jboss.org/browse/JBREM-185
12566	JBREM-153	Change committed. Thanks Ron. If you don't mind, could you e-mail me (tom@jboss.org)? Thanks.	1	1	https://issues.jboss.org/browse/JBREM-153
12567	JBREM-152	Made changes to SocketServerInvoker to handle the case that previous invoker has not finished its cleanup. Should be able to click the 'Version Control' link to see the code change. I was unable to create a test case to reproduce this. Would you happen to have one? Would like to check it in as part of the test run.	1	1	https://issues.jboss.org/browse/JBREM-152
12568	JBREM-152	Hey Ron. Think this one is fixed, but wanted to see if you could verify? Don't think I could come up with a test case for it.	1	1	https://issues.jboss.org/browse/JBREM-152
12569	JBREM-93	It just wraps any Throwable thrown when calling callBackClient.invoke(internalInvocation,callback.getRequestPayload()), so can just get original exception by calling getCause() on the HandleCallbackException thrown. Not sure I understand your request. Is it that you want the original exception throw from callBackClient.invoke() to be throw out of handleCallback() method, without wrapping it? Maybe you can explain use case?	0	0	https://issues.jboss.org/browse/JBREM-93
12570	JBREM-93	Yes, I guess I can use getCause() to get the original exception. Is this behaviour documented somewhere, so I can rely on it? Won't it change accross releases? If I didn't look in the code, I had no way of knowing I can do that. Since I wasn't able to find it documentated, I was reluctant to use it. On the other hand, if you have handleCallback() returning an Object, the interface would be much more intuitive. Any chance to include this is any of your upcoming releases?	1	1	https://issues.jboss.org/browse/JBREM-93
12571	JBREM-93	I don't think this is documented anywhere. The real problem lies in that there can be many exceptions that are raised that have nothing to do with the actual callback handler receiving the callback (i.e. network connection problem). Another major issue is that the client may get the callbacks by polling, so even if the client has a problem processing the callback, there is no way to provide this feedback to server (and is up to client if is push or pull, so no way for server side to rely on this behavior). I would prefer to change the HandleCallbackException than return an Object from the handleCallback() method because of the forementioned problem of not knowing if is push or pull (as would always have to return some default value if was pull, like null, since would not be able to get it from the client till sometime later). So if I added something to HandleCallbackException like getUserException() which would return original exception thrown by client (if there was one, otherwise return null), would that work? Also, would it be useful to provide a way of knowing if CallbackHandler is using pull or push model? Maybe a use case would help me understand what you are after a little better.	0	0	https://issues.jboss.org/browse/JBREM-93
12572	JBREM-70	is sun jmx RI and sun jaxp required in classpath?	0	0	https://issues.jboss.org/browse/JBREM-70
12573	JBREM-69	Implementation pretty much done. Only remaining issues are: - is jmxmp required default by spec for the protocol? - what is the default port? Spec says, "0 means a default or anonymous port, depending on the protocol." Assume this mean default port for protocol? For example, 80 for http?	0	0	https://issues.jboss.org/browse/JBREM-69
12574	JBREM-63	Is this just for jboss-head, or is this needed for 4.0 branch as well?	0	0	https://issues.jboss.org/browse/JBREM-63

#	Requirement	Questions	Answered	Successfully Answered	URL
12575	JBREM-51	Have part of the remote classloading implemented for marshall loading from the server. Need to expand this to general remote classloading both ways, but will require implementing security around this.	0	0	https://issues.jboss.org/browse/JBREM-51
12576	JBREM-47	Have part of the remote classloading implemented for marshall loading from the server. Need to expand this to general remote classloading both ways, but will require implementing security around this.	0	0	https://issues.jboss.org/browse/JBREM-47
12577	JBREM-39	Is there an ability to use an external thread pool service like the org.jboss.util.threadpool.ThreadPool, org.jboss.util.threadpool.BasicThreadPool used by jbossmq and naming?	1	1	https://issues.jboss.org/browse/JBREM-39
12578	JBREM-39	Nope. The socket invoker manages threads manually (as is done with the pooled invoker). The http invoker uses the oswego PooledExecutor. The rmi invoker does not have any direct interaction with thread management. I wouldn't mind adding ability to plug in threadpools, only problem is how to make this generic? Would want to be able to accept jboss, oswego, etc., but is no standard interface for these, right?	0	0	https://issues.jboss.org/browse/JBREM-39
12579	JBREM-35	It needs to be discussed who this would work with the WS4EE deployment model for JSE. As a reminder, it uses the <servlet-class> element to the declare the endpoint implemenation bean. What would be the benefit of running server side JBossWS with a dependency on Remoting rather than Tomcat?	0	0	https://issues.jboss.org/browse/JBREM-35
12580	JBREM-27	One additional note: You seem to use the originalURI in the transport() method of HTTPClientInvoker class: String targetURL = getLocator().getOriginalURI(); But i think that getLocator().getURI() is more correct. For example, when using annotations we sometimes use the following: @RemoteBinding(clientBindUrl="http://0.0.0.0:8888/SharedStateManager",jndiBinding="SharedStateManagerhttp") 0.0.0.0 is substituted by the real IP-number of the server during the deployment. Therefore, the originalURI of the locatoris http://0.0.0.0:8888/SharedStateManager, but URI field is correct and set to: http://x.y.z.w:8888/SharedStateManager where x.y.z.w is the IP of the server. With your current code, there is no easy way for a client to get the correct locator, since http://0.0.0.0:8888/SharedStateManager is returned. But may be it is also the same bug somewhere else, which makes use of originalURI, instead of URI. The fix is obvious and we hope to see it soon in the main branch. Best regards, Ivan	1	0	https://issues.jboss.org/browse/JBREM-27
12581	REM3-154	Are you sure it's hanging? This test can take a long time, especially on CPU-constrained hardware. From the stack trace, it looks like it is busy executing the test, not hung.	1	1	https://issues.jboss.org/browse/REM3-154
12582	REM3-154	Is the test using CPU when it times out, or is it idle?	1	1	https://issues.jboss.org/browse/REM3-154
12583	REM3-154	Is it just on Windows or other platforms as well?	1	1	https://issues.jboss.org/browse/REM3-154
12584	REM3-154	Can you give the exact version (or git revision) of Remoting and XNIO being used? I cannot reproduce this locally at all.	1	1	https://issues.jboss.org/browse/REM3-154
12585	REM3-154	Can you include the output of "ifconfig -a" and the contents of your /etc/hosts file? Use a private comment if you wish.	1	1	https://issues.jboss.org/browse/REM3-154
12586	REM3-154	Awesome. Thanks David. Could you link the commit as well?	1	0	https://issues.jboss.org/browse/REM3-154
12587	REM3-89	Here's the patch.. (somewhat dirty due to automatic cleaning in Eclipse)	1	1	https://issues.jboss.org/browse/REM3-89
12588	REM3-56	What MBean properties should be exported? What did R2 expose before?	1	1	https://issues.jboss.org/browse/REM3-56
12589	REM3-56	Here's the draft: http://anonsvn.jboss.org/repos/jbossremoting/remoting3/trunk/jboss-remoting/src/main/java/org/jboss/remoting3/management/ There are two MXBeans: EndpointMXBean and ConnectionMXBean. They are almost identical except that ConnectionMXBean has close/forceClose operations while EndpointMXBean does not and that you can list the registered services with EndpointMXBean while you can't with ConnectionMXBean. They both provide some counters via the Counters object. Once agreed, I will finish documentation. One missing piece though is how we register an MBean for a connection. We could register all EndpointMXBeans automatically but we can't do that for Connections because there will be much more connections and registering that many connections will probably pollute the JMX console with overinformation. What user wants might be to register only the ConnectionMXBeans that meet some criteria..? Long time ago, I wrote a JMX operation that accepts some expression so that only the MBeans that meets the specified condition, but I'm not sure it's a good idea since it's a lot of work.	1	1	https://issues.jboss.org/browse/REM3-56
12590	REM3-56	If the number of connections are expected to be small, it might be fine as you said. As long as we have an option that disables JMX exports, it will be fine. Authentication failures could be represented as rejected clients? (i.e. clientCounters.failure) What information else could be provided by EndpointMXBean / ConnectionMXBean by the way? A user might want to know where exactly the endpoint / connection is pointing at. Will the MBean object name be enough?	1	1	https://issues.jboss.org/browse/REM3-56
12591	REM3-56	Then maybe we could add connectionCounters? +1 for not registering local connections. If we need to provide the information for connections, we could add ConnectionInfo like I did for ServiceInfo. What attribute would you suggest other than the connection ID and its timestamp?	1	1	https://issues.jboss.org/browse/REM3-56
12592	REM3-38	Is this still an issue now that we switched to JBMAR? At least I cannot find CoreSession or any stream detection logic in the stream package.	1	0	https://issues.jboss.org/browse/REM3-38
12593	REM3-29	Is there any progress regarding this issue? Are there other ways/workaround to share client kerberos login credentials amongst all remote sessionbeans?	0	0	https://issues.jboss.org/browse/REM3-29
12594	REM3-29	Are you actually planning to put a new feature in a CR2 release?	0	0	https://issues.jboss.org/browse/REM3-29
12595	REM3-24	What about just returning IoFuture on close()? Perhaps closeAsync() could return IoFuture and close() could be synchronous graceful shutdown. When do we need immediate shutdown btw?	1	0	https://issues.jboss.org/browse/REM3-24
12596	REM3-23	SSLContexts should be allowed to be reused regardless if the connection was dropped, closed cleanly, or there was no prior connection. That is, a user should be able to specify an option when making a connection attempt, just like he/she can do for socket options. On a dropped connection, any in-progress requests should be notified (via IoFuture) so that the client is aware of possibility of duplicate requests on redelivery or dropped requests on non-redelivery. i.e. it should never be retried silently. What other contextual information would we need to retain on disconnection for quick resilence? If it doesn't take much time and it's not CPU intensive, we might not need this feature, except for SSLContext reuse?	1	1	https://issues.jboss.org/browse/REM3-23
12597	JBSER-91	Should we re-package 1.0.2.GA and upload it into sourceforge or should we do it on 1.0.3.GA only?	1	0	https://issues.jboss.org/browse/JBSER-91
12598	CMTOOL-103	Decreasing priority, the server boots if referenced modules are missing, just logging an error: {code} 19:31:21,784 ERROR [org.jboss.as.controller.management-operation] (ServerService Thread Pool -- 49) WFLYCTL0013: Operation ("add") failed - address: ([("subsystem" => "naming"), ("binding" => "java:global/foo/bar/factory")]) - failure description: "WFLYNAM0052: Could not load module org.foo.bar:main" {code}	0	0	https://issues.jboss.org/browse/CMTOOL-103

#	Requirement	Questions	Answered	Successfully Answered	URL
12599	CMTOOL-85	Does it mean that the default configuration of {{elytron}} subsystem will be added during a migration? Just this and nothing more/less?	0	0	https://issues.jboss.org/browse/CMTOOL-85
12600	CMTOOL-85	I cannot see any other CMTOOL jira linked to EAP7-709. Do you plan any other works related to EAP7-709 support?	0	0	https://issues.jboss.org/browse/CMTOOL-85
12601	JBTIS-1082	Hi [~naciao]! [~pleacu] mentioned on Mattermost that we need to remove the SwitchYard-specific doc/help files from the Fuse Tooling docs and migrate them into the SwitchYard world since we're pulling SY out of Fuse as it moves to Dev Studio. Can you or [~tcohen] or [~jmurphey] create a PR against the Fuse Tools docs and let us know what we do to make this transition? Thank you!	1	1	https://issues.jboss.org/browse/JBTIS-1082
12602	JBTIS-1071	[~lhein] [~apodhrad] - Hey - Id like to execute the following command in the following directories to clean up CI cruft - I wanted to pass it by you first: cd /media/TOOLS-ssh/neon/integration/updates/integration-stack/fuse-tooling/ find . -name "org*.jar" -type f -mtime +10 -delete This will clean up all CI org*.jars greater than 10 days old. WDYT?	0	0	https://issues.jboss.org/browse/JBTIS-1071
12603	JBTIS-1057	10800 seconds timeout is set in _tests/pom.xml_ --> _surefire.timeout_. I tried to prolong it to 21600 seconds.	1	1	https://issues.jboss.org/browse/JBTIS-1057
12604	JBTIS-1044	Do you plan to use sha256 for DevstudioIS 10.1.0.GA?	1	1	https://issues.jboss.org/browse/JBTIS-1044
12605	JBTIS-1040	Hey [~apodhrad] - sorry - still confused. I do this: 1. Install Devstudio 10.2.0.GA 2. Add the following vm args to devstudio.ini -Djboss.discovery.directory.url=https://devstudio.redhat.com/10.0/staging/updates/integration-stack/discovery/10.1.0.CR2/devstudio-directory.xml -Djboss.discovery.site.integration-stack.url=https://devstudio.redhat.com/10.0/staging/updates/integration-stack/discovery/10.1.0.CR2/ 3. Start the Devstudio and open Red Hat Central 4. Type 'fuse' into the text field in the middle of the RH Central (nothing is found) 5. Switch to the tab 'Software/Update' 6. Check 'Enable Early Access' 7. Switch back to the tab 'Getting Started' (now you can see the fuse quickstart) new steps 6. Select 'JBoss Fuse Development' 7. Install it 8. When it restarts into the 'Getting Started' tab - the Fuse quickstarts are available. In your scenario, JBoss Fuse Development still isn't installed, I don't see the Fuse quickstart. I still get nothing found when I perform step 7. If I select 'JBoss Fuse Development' and install it - then go back to 'Getting Started', I see the Fuse quickstarts (without EA enabled). This behavior is as I would expect. Are you using a new workspace?	1	1	https://issues.jboss.org/browse/JBTIS-1040
12606	JBTIS-1040	Thanks [~fbricon] - that looks like it works. Now when we switch Eclipse platforms it's possible that Fuse will return to EA - will we have the original issue again?	1	1	https://issues.jboss.org/browse/JBTIS-1040
12607	JBTIS-1032	I reproduced this issue but I am not able to find a cause of the problem. The issue appears only when you run test via command line with bundle eclipse. I found out that there is weird version of drools eclipse plugin (7.0.0.201701021409) although there should be (6.5.0.CR1). Are there any related changes in plugin installation?	1	1	https://issues.jboss.org/browse/JBTIS-1032
12608	JBTIS-1008	[~pleacu] There might be an issue with the Target Platform or the way I consume it, the new features requires to have one bundle in exploded mode (org.eclipse.cdt.core.win32.x86_64) otherwise we hit a JVM crash as I reported in [JBDS-4227] usually I'm using the TP as directory folder but in this case the bundle is jar. I tried to configure my local TP as an update site but I have several errors: !screenshot-1.png|thumbnail! Do you know if it is an issue with the generated TP which should have the jar exploded? or if it is the way that I'm consuming the TP?	1	1	https://issues.jboss.org/browse/JBTIS-1008
12609	JBTIS-1008	[~aurelien.pupier] - Hey Aurelien - I've released the 4.4.1.CR2 JBTIS TP with the missing openshift plugins. I plan on closing this Jira since it's original issue is resolved - can you open a new Jira on the Eclipse issue with the local TP?	1	1	https://issues.jboss.org/browse/JBTIS-1008
12610	JBTIS-1003	Hey [~apodhrad] - sorry - I'm a little fuzzy today :) Where exactly would I add this id? Your POM defined the repositories (or your settings.xml). There already is a composite at http://download.jboss.org/jbosstools/updates/requirements/neon/ The .target file defines a 'repository location'. --p	1	1	https://issues.jboss.org/browse/JBTIS-1003
12611	JBTIS-949	Hey Andrej - by calling it a 'runtime-installer' it seems to imply that it installs just runtimes and not the tooling as well. Can we combine 'standalone' and runtime somehow - wdyt?	1	1	https://issues.jboss.org/browse/JBTIS-949
12612	JBTIS-848	[~jomarko] please provide more details about this issue. Which test is failing? Is there any stacktrace?	1	1	https://issues.jboss.org/browse/JBTIS-848
12613	JBTIS-844	Hi [~aurelien.pupier], could you provide steps how we can verify the issue. please? Btw is this the only one bug fix whoch was added or are there more commits included?	1	1	https://issues.jboss.org/browse/JBTIS-844
12614	JBTIS-819	[~pleacu], where can I see the new icons?	1	1	https://issues.jboss.org/browse/JBTIS-819
12615	JBTIS-799	[~djelinek] I presume this will be a part of your tests. Am I right?	1	0	https://issues.jboss.org/browse/JBTIS-799
12616	JBTIS-792	Do you have any reason for preserving _WARTest.java_ and _DynamicVdbTest.java_? You commented out all test methods, so these classes are useless. I suggest to remove them completely.	1	1	https://issues.jboss.org/browse/JBTIS-792
12617	JBTIS-784	For which version(s) of integration stack are these targetting? Is it only for Mars? What about Neon? There is no automated update for the ide-config.properties file. We update it by hand in github as needed, then push the updated file to download.jboss.org. You have the correct file, but you may need to update it in more than one place since you probably want to update the jboss.discovery.site.integration-stack-sap.url for devstudio 9 and 10, as well as jbosstools 4.3 and 4.4, If you need to source data from TWO different update sites, then you need a composite update site to merge them into one, which can then be referenced by the jboss.discovery.site.integration-stack-sap.url variable But.. http://download.jboss.org/jbosstools/mars/development/updates/integration-stack/extras/ currently only lists one child site, http://download.jboss.org/jbosstools/mars/development/updates/integration-stack/extras/jboss-fuse-sap-tool-suite/8.0.0.Alpha1/all/ And there's a stable release site as well: http://download.jboss.org/jbosstools/mars/stable/updates/integration-stack/extras/ http://download.jboss.org/jbosstools/mars/stable/updates/integration-stack/extras/jboss-fuse-sap-tool-suite/8.0.0.Beta1/all/ Not sure why you guys aren't linking to the newer release. [~pleacu] can maybe explain why the version available in Central for JBT 4.3.x is your 8.Alpha, not 8.Beta. So, best person to talk to here is Paul. I'll move this issue to JBTIS for you.	1	0	https://issues.jboss.org/browse/JBTIS-784
12618	JBTIS-758	There is a problem with container, because number of edit parts in Camel Editor are same for collapsed/expanded situation (eg. collapse/expand 'Choice' component in fuse project with CBR, Spring). How is possible to test collapse/expand feature of Camel Component ? [~apodhrad]	1	1	https://issues.jboss.org/browse/JBTIS-758
12619	JBTIS-758	What about test size of the choice component CamelComponentEditPart.getBounds()?	1	1	https://issues.jboss.org/browse/JBTIS-758

#	Requirement	Questions	Answered	Successfully Answered	URL
12620	JBTIS-755	org.jboss.reddeer.core.handler.MenuHandler.select(MenuHandler.java:113) at org.jboss.reddeer.swt.impl.menu.ShellMenu.select(ShellMenu.java:82) at org.jboss.tools.drools.ui.bot.test.util.TestParent.deleteAllProjectsWorkaround(TestParent.java:189) at org.jboss.tools.drools.ui.bot.test.util.TestParent.cleanUp(TestParent.java:171) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:50) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:47) at org.jboss.reddeer.junit.internal.runner.statement.RunAfters.evaluate(RunAfters.java:76) at org.jboss.reddeer.junit.internal.runner.statement.RunIAfterTestExtensions.evaluate(RunIAfterTestExtensions.java:61) at org.junit.rules.TestWatcher$1.evaluate(TestWatcher.java:55) at org.junit.rules.TestWatcher$1.evaluate(TestWatcher.java:55) at org.junit.rules.RunRules.evaluate(RunRules.java:20) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:325) at org.jboss.reddeer.junit.internal.runner.RequirementsRunner.runChild(RequirementsRunner.java:171) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:57) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:290) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:71) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:288) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:58) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:268) at org.jboss.reddeer.junit.internal.runner.statement.RunBefores.evaluate(RunBefores.java:79) at org.jboss.reddeer.junit.internal.runner.statement.FulfillRequirementsStatement.evaluate(FulfillRequirementsStatement.java:45) at org.jboss.reddeer.junit.internal.runner.statement.RunIBeforeClassExtensions.evaluate(RunIBeforeClassExtensions.java:72) at org.jboss.reddeer.junit.internal.runner.statement.RunAfters.evaluate(RunAfters.java:68) at org.jboss.reddeer.junit.internal.runner.statement.CleanUpRequirementStatement.evaluate(CleanUpRequirementStatement.java:44) at org.jboss.reddeer.junit.internal.runner.statement.RunIAfterClassExtensions.evaluate(RunIAfterClassExtensions.java:59) at org.junit.rules.TestWatcher$1.evaluate(TestWatcher.java:55) at org.junit.rules.RunRules.evaluate(RunRules.java:20) at org.junit.runners.ParentRunner.run(ParentRunner.java:363) at org.jboss.reddeer.junit.internal.runner.RequirementsRunner.run(RequirementsRunner.java:153) at org.junit.runners.Suite.runChild(Suite.java:128) at org.junit.runners.Suite.runChild(Suite.java:27) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:290) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:71) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:288) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:58) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:268) at	1	1	https://issues.jboss.org/browse/JBTIS-755
12621	JBTIS-730	[~bfitzpat] [~bbrodt] - Hey Fitz - Bob has a new Neon BPMN2 Modeler that I've just mirrored and put into the Neon JBTIS. The stack currently won't build because SwitchYard has a non-inclusive range: {code} [ERROR] Software being installed: site-ea raw:4.4.0.'SNAPSHOT'/format(n[.n=0;[.n=0;[-S]]]):4.4.0-SNAPSHOT [ERROR] Missing requirement: org.switchyard.tools.bpmn2.feature.feature.group 2.1.0.Beta3-v20160226-2057-B645 requires 'org.eclipse.bpmn2.modeler.feature.group [1.2.0,1.3.0)' but it could not be found [ERROR] Cannot satisfy dependency: site-ea raw:4.4.0.'SNAPSHOT'/format(n[.n=0;[.n=0;[-S]]]):4.4.0-SNAPSHOT depends on: org.switchyard.tools.bpmn2.feature.feature.group 0.0.0 {code} Bob's bpmn2 is 1.3.0.Final. Could you extend the range [1.2.0,1.3.0] and perhaps build with a Neon JBTIS TP (4.4.0.Alpha1-SNAPSHOT) ? I'll take the range extension if you're not ready for the TP. Thanks!	1	1	https://issues.jboss.org/browse/JBTIS-730
12622	JBTIS-688	Hey [~apodhrad] - since we can't easily go back I guess we go forward. What's the latest issue with Drools? Can we ship with it?	1	1	https://issues.jboss.org/browse/JBTIS-688
12623	JBTIS-674	looks good, but just to understand: where are the runtime features intended to be included if we are deprecating the current feature they get installed with ? And since they are there to use the plugins that are marked as deprecated why are the runtime detectors not deprecated too ?	1	1	https://issues.jboss.org/browse/JBTIS-674
12624	JBTIS-674	Agreed - will deprecate 'JBoss Tools Runtime Detection'	1	1	https://issues.jboss.org/browse/JBTIS-674
12625	JBTIS-674	reopen for clarification sorry, i'm not following - are you saying runtime detection will be deprecated now even for jbpm/drools ? are you planning on removing jbpm/drools ?	1	1	https://issues.jboss.org/browse/JBTIS-674

#	Requirement	Questions	Answered	Successfully Answered	URL
12626	JBTIS-666	Assistent does not contain 'id' value at org.junit.Assert.fail(Assert.java:88) at org.junit.Assert.assertTrue(Assert.java:41) at org.jboss.tools.fuse.ui.bot.test.CamelEditorTest.testCodeCompletion(CamelEditorTest.java:192) {code} DataTransformation {code} org.jboss.reddeer.core.exception.CoreLayerException: No control has focus. Perhaps something has stolen it? Try to regain focus with for example "new DefaultShell()". at org.jboss.reddeer.core.lookup.MenuLookup.getTopMenuMenuItemsFromFocus(MenuLookup.java:226) at org.jboss.reddeer.swt.impl.menu.ContextMenu.<init>(ContextMenu.java:45) at org.jboss.reddeer.swt.impl.menu.ContextMenu.<init>(ContextMenu.java:34) at org.jboss.tools.fuse.reddeer.editor.CamelEditor.deleteCamelComponent(CamelEditor.java:172) at org.jboss.tools.fuse.ui.bot.test.DataTransformationTest.testBasics(DataTransformationTest.java:65) {code} {code} org.jboss.reddeer.common.exception.WaitTimeoutExpiredException: Timeout after: 60 s.: The following jobs are still running Enabling Maven Dependency Management Building workspace at org.jboss.reddeer.common.wait.AbstractWait.timeoutExceeded(AbstractWait.java:173) at org.jboss.reddeer.common.wait.AbstractWait.wait(AbstractWait.java:126) at org.jboss.reddeer.common.wait.AbstractWait.<init>(AbstractWait.java:91) at org.jboss.reddeer.common.wait.AbstractWait.<init>(AbstractWait.java:61) at org.jboss.reddeer.common.wait.AbstractWait.<init>(AbstractWait.java:46) at org.jboss.reddeer.common.wait.WaitWhile.<init>(WaitWhile.java:35) at org.jboss.tools.fuse.ui.bot.test.utils.ProjectFactory.importExistingProject(ProjectFactory.java:162) at org.jboss.tools.fuse.ui.bot.test.DataTransformationTest.testBasics(DataTransformationTest.java:59) {code} DragAndDrop on Windows {code} org.jboss.reddeer.core.exception.CoreLayerException: Exception during sync execution in UI thread at org.jboss.reddeer.core.util.Display.syncExec(Display.java:87) at org.jboss.reddeer.core.util.Display.syncExec(Display.java:53) at org.jboss.reddeer.gef.handler.EditPartHandler.select(EditPartHandler.java:45) at org.jboss.reddeer.gef.impl.editpart.AbstractEditPart.select(AbstractEditPart.java:61) at org.jboss.tools.fuse.reddeer.editor.CamelEditor.selectEditPart(CamelEditor.java:561) at org.jboss.tools.fuse.reddeer.editor.CamelEditor.setProperty(CamelEditor.java:427) at org.jboss.tools.fuse.ui.bot.test.CamelEditorTest.testDragAndDropComponents(CamelEditorTest.java:241) Caused by: java.lang.NullPointerException: null at org.jboss.reddeer.gef.handler.EditPartHandler$1.run(EditPartHandler.java:48) at org.jboss.reddeer.core.util.Display$VoidResultRunnable.run(Display.java:194) at org.jboss.reddeer.core.util.Display$VoidResultRunnable.run(Display.java:1) at org.jboss.reddeer.core.util.Display$ErrorHandlingRunnable.run(Display.java:160) at org.eclipse.swt.widgets.RunnableLock.run(RunnableLock.java:35) at org.eclipse.swt.widgets.Synchronizer.runAsyncMessages(Synchronizer.java:135) at org.eclipse.swt.widgets.Display.runAsyncMessages(Display.java:4155) at org.eclipse.swt.widgets.Display.readAndDispatch(Display.java:3772) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine$4.run(PartRenderingEngine.java:1127) at org.eclipse.core.databinding.observable.Realm.runWithDefault(Realm.java:337) at org.eclipse.e4.ui.internal.workbench.swt.PartRenderingEngine.run(PartRenderingEngine.java:1018) at	1	1	https://issues.jboss.org/browse/JBTIS-666
12627	JBTIS-654	[~pleacu] - is there any chance we can update the summary (step 7) of the installer to include the "additional bits" the user selects? or select all the fuse bits by default on step 6?	1	1	https://issues.jboss.org/browse/JBTIS-654
12628	JBTIS-654	[~pleacu] I'm on fedora and only question I'm asked is to accept the licence. How do I specify to include Fuse?	1	1	https://issues.jboss.org/browse/JBTIS-654
12629	JBTIS-654	[~aileenc] - hmmm - is this what you downloaded? {code} ls -lt jboss-devstudio-is-installer-9.0.0.CR1a.jar -rw-rw-r--. 1 pleacu pleacu 992043163 Apr 14 16:46 jboss-devstudio-is-installer-9.0.0.CR1a.jar {code} downloaded from https://devstudio.redhat.com/9.0/snapshots/updates/integration-stack/jboss-devstudio-is-installer-9.0.0.CR1a.jar Are you running java 1.8? I'm on FC23 - which Fedora are you running? Thkx!	1	1	https://issues.jboss.org/browse/JBTIS-654
12630	JBTIS-654	This is what i see when I run the Installer: no option to select Fuse. What am I missing? ---------- This document can be seen online at http://www.redhat.com/licenses/jboss_eula.html. Archived 2014-07-31 18:00 GMT. press 1 to accept, 2 to reject, 3 to redisplay 1 [Starting to unpack] [Processing package: P2 Director Application (1/3)] [Processing package: Red Hat JBoss Developer Studio 9.1.0.GA (2/3)] [Processing package: Red Hat JBoss Developer Studio 9.1.0.GA Metadata Generation (3/3)] [Unpacking finished] [Console installation done]	1	1	https://issues.jboss.org/browse/JBTIS-654
12631	JBTIS-654	[~aileenc] That's actually good news - you definitely need java 8. Did you try my previous comment on your linux box w/ java 8?	1	1	https://issues.jboss.org/browse/JBTIS-654
12632	JBTIS-654	On the topic of Java 8 vs 7 - based JBoss Tools should actually tell you there are bundles that did not start due to bad java version. Did that not happen ? Also the installer should warn you if using wrong java version - if not that is a bug.	1	1	https://issues.jboss.org/browse/JBTIS-654
12633	JBTIS-654	[~pleacu] Couple of questions regarding runtime configuration in the installer... I agree that not everybody will want a fuse-only tooling installer, but I think from an evaluator point of view it would be awesome to get a one-click-and-go install process that not only installs JBDS + Fuse tooling in a location of the user's choosing, but pre-configures a Fuse runtime as well. If we want to compete in the same space as Mulesoft and Talend, having some sort of all-in-one installer as an option would definitely help get us part of the way. Can we add such options to the installer as we move to the next phase of this effort?	1	1	https://issues.jboss.org/browse/JBTIS-654
12634	JBTIS-654	[~bfitzpat] - that makes sense to me but we would need to proceed carefully. Random thoughts: * We can't get the Fuse Tooling production build (and the JBDSIS build by extension) caught on the Fuse Runtime (we need to preserve the ability to release autonomously). * The stand-alone IS installer today is just shy of 1Gb - do we have an idea what the toal aggregate size would be if we add the Fuse runtime? * If we add the Fuse runtime should we consider a Fuse stand-alone installer (not JBDSIS - just Fuse)? * Should we do both - an aggregate JBDSIS stand-alone installer which includes Fuse Tooling (and the rest of the stack) and a separate Fuse stand-alone isntaller containg just Fuse Tooling and runtime? * Any changes to the JBDS installer mechanism would have to occur in 10.x (Neon) if we find something that needs changes there.	1	1	https://issues.jboss.org/browse/JBTIS-654
12635	JBTIS-654	I'm building this now via Jenkins on-demand: https://devstudio.redhat.com/9.0/staging/updates/integration-stack/9.0.1-SNAPSHOT/devstudio-is-installer-9.0.1-SNAPSHOT.jar Should we close this Jira and create new ones on how we'd like to modify what we now have?	1	1	https://issues.jboss.org/browse/JBTIS-654
12636	JBTIS-654	Why do u want to blow up the size of the installer in that way? Sometimes less is more.	1	1	https://issues.jboss.org/browse/JBTIS-654
12637	JBTIS-654	btw. aren't you also saying it is annoying having to install Fuse after installing devstudio despite it has builtin support to automatically download it when users start the Fuse wizard - does the same reasoning not apply for runtimes ? Some users like it simply to work when installed.	1	1	https://issues.jboss.org/browse/JBTIS-654
12638	JBTIS-654	I guess anything's possible... :) So - it stands to reason that the person who picks the 1.5Gb JBDSIS stand-alone installer with Fuse runtime wants to install Fuse Tooling. Should this installer have *just* Fuse Tooling + Fuse Runtime (or minimally pre-select Fuse Tooling). What about BRMS and Data Virt runtimes? SHould they also have a tooling-specific + runtime stand-alone installer?	1	1	https://issues.jboss.org/browse/JBTIS-654

#	Requirement	Questions	Answered	Successfully Answered	URL
12639	JBTIS-654	That's a good question and a logical extension of this conversation. Can I suggest that we start with the Fuse tooling + runtime installer we've been discussing and use it as a pilot to see if it gets any traction? If it does, we can float this to other product teams as well. Is there any way to track not just downloads but how many time it gets run? We already do some tracking in the tooling... but I would think download tracking of the installer would be one metric we could pretty easily keep track of.	0	0	https://issues.jboss.org/browse/JBTIS-654
12640	JBTIS-654	Just an update - what I've got here [1] is an installer that contains JBDS, JBDSIS (with Fuse Tooling pre-selected) and a Fuse runtime. Running the installer will get you JBDSIS with Fuse Tooling (by just selecting Next). The Fuse runtime is sitting in the jar (used download-maven-plugin/fetch-target--platform to wget it dynamically). What I'm having trouble with is how to extract it out when the installer runs. Basically all I need to do is: {code} jar -xf jbdsis-fuse-installer-9.0.1-SNAPSHOT.jar jbdsis/fuse {code} Which leads me to the other question - should this go into jbdevstudio-install-dir/fuse-runtimes to avoid confusion w/ EAP? Any/all comments welcome... [1] https://devstudio.redhat.com/9.0/snapshots/updates/integration-stack/devstudio-is-fuse-installer-9.0.1-SNAPSHOT.jar	1	1	https://issues.jboss.org/browse/JBTIS-654
12641	JBTIS-654	There are multiple Fuse runtimes - one for karaf, one for eap and separate versins for each supported platform. Which Fuse runtime are you including? Will this be Windows only as that is likely the most common evaluation platform?	1	1	https://issues.jboss.org/browse/JBTIS-654
12642	JBTIS-654	Hey [~dgolovin] [[~akazakov] - I know you guys are extremely busy with the summit release. This is a post-summit enhancement request. I talked with [~maxandersen] about how best to architect this support. He was of the opinion that it would make more sense to parameterize the devstudio installer support to allow it to take a platform-specific runtime as opposed to migrating the p2-director and izpack installer infrastructure over in the integration stack build. Having migrated the two over in a prototype it did amount to a near duplication of the code. I'm currently successfully using Denis' p2 director installer enhancement (AdditionalFeaturesSpec.json) to allow the IS to specify its features into the devstudio installer. We'll release that in the Mars stream. What I'd like is to extend that... Is it possible to get into a future devstudio 10.0.x (neon) release enhanced installer support where a user could specify: 1. A stand-alone installer jar - this would be the devstudio-is (Integration Stack) stand-alone installer. 2. A compressed runtime (zip) that would be expanded into an appropriate runtime folder within the devstudio footprint (using the EAP one or a seperate folder). This would likely go into its own pack ID - "Red Hat JBoss Developer Studio Component Runtime" - something like that? ref: jboss-fuse-full-6.3.0.redhat-075.zip 3. A specific IzPack panel binary - tailored to the platform we're offering. This would be installed into the specified IS installer (1) to create the new extended installer based on what was passed in., The goal would be to provide an abstract way to specify a runtime zip and installer panel associated with the runtime into a pre-existing stand-alone installer. The existing runtime detection would need to be augmented as well. That would allow greater flexibilty in creating full stand-alone installers for Fuse, Teiid Designer, etc and leverage the existing devstudio installer to the greatest extent. I'm open to any other ideas as well. WDYT?	1	1	https://issues.jboss.org/browse/JBTIS-654
12643	JBTIS-654	[~leacu] I really not sure we currently have resources for such enhancements in the devstudio installer code base for the next release. Can you guys try to take care of it yourself? :) [~dgolovin] could help you as a consultant and also help with reviewing your Pull Request.	1	1	https://issues.jboss.org/browse/JBTIS-654
12644	JBTIS-654	As I said when we talked I think reusing all the logic and just allow for in the build to start which runtime to include is better. But what is it that #3 panel is supposed to do that is not just unzipping the content ? Any setup beyond that I would initially have expected to happen inside devstudio. Also, from Aileen it sounds like you have different runtime zips per OS ? is that true ?	1	1	https://issues.jboss.org/browse/JBTIS-654
12645	JBTIS-654	[~akazakov] - reviewed and +1 from Denis. I believe I need your +1 and someone to push my PR. Also - can this get on the 4.4.1.AM3 sprint for Neon? Thkx	1	1	https://issues.jboss.org/browse/JBTIS-654
12646	JBTIS-654	Okay - I can look at the progress bar later - can we get this first one pushed?	0	0	https://issues.jboss.org/browse/JBTIS-654
12647	JBTIS-649	[~apodhrad] - where do you see org.jboss.tools.foundation.security.linux 1.2.1.Beta2 being pulled in? All I can see is {code} [pleacu@pleacu-2-laptop jbosstools-integration-stack]$ grep -R 'org.jboss.tools.foundation.security.linux' * target-platform/src/targets/core-base.target: <unit id="org.jboss.tools.foundation.security.linux.feature.feature.group" version="1.2.1.Final-v20160401-0436-B108"/> {code} !found1.png|thumbnail!	1	1	https://issues.jboss.org/browse/JBTIS-649
12648	JBTIS-648	Hey [~maxandersen] can you give us some guidance here, based on your previous experiences with Eclipse projects & JBDS?	1	1	https://issues.jboss.org/browse/JBTIS-648
12649	JBTIS-602	Andrej, we'll need a lot more here. Can you provide additional details? Or is this a QE-only issue?	1	1	https://issues.jboss.org/browse/JBTIS-602
12650	JBTIS-523	Are we only talking about the install experience via the installer or overall install of Fuse ? The installer is by far the least used afaics so if we only look at this from installer perspective we will not cover the paths of BYOE (bring your own eclipse) for users using eclipse market place. If just caring about installation could be solved by adding a Fuse specific marketplace entry that just contains the Fuse functionallity as done for the IS marketplace entry. This would still install the required JBDS core parts, incl. central. Is this just about having a way to install it easily or also about how Fuse will be configured in eclipse ? Then setting first perspectives are bad/not enough IMO. Maybe Fuse could add entries to JBoss perspective ? and beyond that is making Fuse quickstarts etc. available from JBoss Central so anyone can get to Fuse even without having first installed Fuse and once installed have access to all the Fuse quickstarts.	1	1	https://issues.jboss.org/browse/JBTIS-523
12651	JBTIS-517	Could you unify the background between ga and ea, please?	1	1	https://issues.jboss.org/browse/JBTIS-517
12652	JBTIS-510	And then for 4.2.4/8.0.4, all the content will be 100% released, with nothing still in EA? I wonder... since we're trying to reduce the # of URLs shown in JBT/JBDS, and since we tag EA content w/ "Early Access" in the UI... do we REALLY even need a separate URL for the EA content? Why not simply use ONE URL for everything? [~maxandersen] WDYT? for JBDS 8, could we composite the stuff in /earlyaccess/8.0/ into the /updates/8.0/ site?	1	1	https://issues.jboss.org/browse/JBTIS-510
12653	JBTIS-502	Hey [~bbrodt] - do you have a Mars native version of Drools?	1	1	https://issues.jboss.org/browse/JBTIS-502

#	Requirement	Questions	Answered	Successfully Answered	URL
12654	JBTIS-498	Hey [~nickboldt] - Ref https://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/JBTIS-aggregate-disc/421/console - I've added '--del' parameter to the rsync.sh invocation which is copying the IS discovery plugins directory. The Jenkins job script fragment is: {code} # Publish the plugin JARs . ${WORKSPACE}/rsync.sh -s ${WORKSPACE}/sources/${BUILD_KIND}/discovery/generation/target/plugins \ -t ${TARGET_PLATFORM}/${TARGET_DIR}/updates/integration-stack/discovery/${VERSION}/plugins --del {code} which results in: {code} 16:00:45 + . /qa/hudson_workspace/workspace/JBTIS-aggregate-disc/rsync.sh -s /qa/hudson_workspace/workspace/JBTIS-aggregate-disc/sources/jbosstools/discovery/generation/target/plugins -t mars/staging/updates/integration-stack/discovery/4.3.0.Alpha2/plugins --del ... 16:00:57 [INFO] rsync -arzq --protocol=28 --del /qa/hudson_workspace/workspace/JBTIS-aggregate-disc/sources/jbosstools/discovery/generation/target/plugins/* tools@filemgmt.jboss.org:/downloads_htdocs/tools/mars/staging/updates/integration-stack/discovery/4.3.0.Alpha2/plugins/ {code} rsync will disregard the '--del' option if wildcards are present. It looks like rsync.sh always wildcards here: https://raw.githubusercontent.com/jbosstools/jbosstools-build-ci/jbosstools-4.3.x/publish/rsync.sh#L72 WDYT?	1	1	https://issues.jboss.org/browse/JBTIS-498
12655	JBTIS-498	You're pulling the rsync script directly from github; we now prefer to fetch the zipped versions from nexus so you can see which version you're getting. I've updated your job here: https://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/JBTIS-aggregate-disc/jobConfigHistory/showDiffFiles?timestamp1=2015-09-15_15-14-37&timestamp2=2015-09-17_17-28-27 1. Why do you have 2 jobs for JBTIS-aggregate-disc and JBDS-aggregate-disc ? why not build them at the same time from the same job? 2. Why do you rsync the contents of sources/BUILD_KIND/site-final/target/ [1] AND its child folder sources/BUILD_KIND/site-final/target/repository [2] into the same target folder on dl.jb.org [3]? You're essentially publishing the same repo twice, plus you're publishing a lot of maven metadata that you (probably) will never need. [1] http://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/JBTIS-aggregate-disc/ws/sources/jbosstools/site-final/target/ [2] http://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/JBTIS-aggregate-disc/ws/sources/jbosstools/site-final/target/repository [3] http://download.jboss.org/jbosstools/mars/snapshots/builds/integration-stack/4.3.0.Alpha2/ Do you do anything with category.xml? checkstyle*? local-artifacts? p2agent? p2artifacts? p2content? targetPlatformRepository? or the site*.zip? 3. As to the use of the --del flag... I think the problem is that you're pushing a single folder's contents into a single folder. When I use rsync.sh --del, it's to push a WHOLE folder and replace whatever is already there. So: {code:title=http://jenkins.mw.lab.eng.bos.redhat.com/hudson/view/DevStudio/view/DevStudio_9.0.mars/job/jbosstools-discovery_master/configure-readonly/} . ${WORKSPACE}/sources/publish/rsync.sh --del -s ${WORKSPACE}/sources/jbosstools/org.jboss.tools.central.discovery.earlyaccess/target/discovery-site \ -t mars/snapshots/updates/discovery.earlyaccess/${jbosstools_site_stream}/ {code} becomes {code:title=http://jenkins.mw.lab.eng.bos.redhat.com/hudson/view/DevStudio/view/DevStudio_9.0.mars/job/jbosstools-discovery_master/lastBuild/console} ++ echo '[INFO] rsync -arzq --protocol=28 --del /mnt/hudson_workspace/workspace/jbosstools-discovery_master/sources/jbosstools/org.jboss.tools.central.discovery.earlyaccess/target/discovery-site/* tools@filemgmt.jboss.org:/downloads_htdocs/tools/mars/snapshots/updates/discovery.earlyaccess/master//' {code} and results in this: http://download.jboss.org/jbosstools/mars/snapshots/updates/discovery.earlyaccess/master/ Alternatively, if you use a versioned folder, you'll get an automated composite site and automated cleanup -- no --del flag needed! {code} . ${WORKSPACE}/sources/publish/rsync.sh -s ${WORKSPACE}/sources/jbdevstudio/com.jboss.jbds.central.discovery/target/discovery-site \ -t 9.0/snapshots/builds/jbosstools-discovery.central_${jbosstools_site_stream}/${BUILD_ID}-B${BUILD_NUMBER}/all/repo/ \ -DESTINATION devstudio@filemgmt.jboss.org:/www_htdocs/devstudio {code} results in: https://devstudio.redhat.com/9.0/snapshots/builds/jbosstools-discovery.central_master/ and https://devstudio.redhat.com/9.0/snapshots/builds/jbosstools-discovery.central_master/compositeContent.xml and https://devstudio.redhat.com/9.0/snapshots/builds/jbosstools-discovery.central_master/latest/all/repo/ So... if you can change your method of publishing such that the site you produce in Jenkins is the same as the site you put on dl.b.o or ds.jb.c (without needing a dozen individual rsync/copy/move steps) not only will your local builds == your CI builds, but your release process will be simpler. 4. Why does your CI build sometimes ALSO push bits to /staging/ as well as /snapshots/ ? Surely that promotion step should ONLY ever happen by hand when it's time to do a QE staging release? 5. Why are there two copies of the same directory.xml file?	1	1	https://issues.jboss.org/browse/JBTIS-498
12656	JBTIS-498	Re: #5, I see you now have two identical copies of the directory.xml file: https://devstudio.redhat.com/9.0/staging/updates/integration-stack/discovery/9.0.0.Alpha2/devstudio-integration-stack-directory.xml https://devstudio.redhat.com/9.0/staging/updates/integration-stack/discovery/9.0.0.Alpha2/devstudio-directory.xml Why do you create two of them?	1	1	https://issues.jboss.org/browse/JBTIS-498
12657	JBTIS-492	Sure - a couple of questions... [~nickboldt] Where can I find the upstream directory XML file? (jbosstools-directory.xml) - http://download.jboss.org/jbosstools/mars/staging/updates/central/4.3.0.Beta1/jbosstools-directory.xml ? http://download.jboss.org/jbosstools/discovery/staging/4.3.0.Alpha1/jbosstools-directory.xml ? Is this the correct JBTIS update site? http://download.jboss.org/jbosstools/mars/snapshots/updates/integration-stack/aggregate/4.3.0.Beta1	1	1	https://issues.jboss.org/browse/JBTIS-492

#	Requirement	Questions	Answered	Successfully Answered	URL
12658	JBTIS-492	Central Discovery now builds *SNAPSHOTS* here: http://download.jboss.org/jbosstools/mars/snapshots/builds/jbosstools-discovery.central_master/ (master) http://download.jboss.org/jbosstools/mars/snapshots/builds/jbosstools-discovery.central_4.3.mars/ (currently Beta1a, built from Beta1x branch) then when *STAGING* for QE: http://download.jboss.org/jbosstools/mars/staging/builds/jbosstools-4.3.0.Beta1-build-discovery.central/ http://download.jboss.org/jbosstools/mars/staging/updates/discovery.central/4.3.0.Beta1/ There are also equivalent URLs for the *EARLY ACCESS* site: http://download.jboss.org/jbosstools/mars/snapshots/builds/jbosstools-discovery.earlyaccess_master/ (master) http://download.jboss.org/jbosstools/mars/snapshots/builds/jbosstools-discovery.earlyaccess_4.3.mars/ (currently Beta1a, built from Beta1x branch) http://download.jboss.org/jbosstools/mars/staging/builds/jbosstools-4.3.0.Beta1-build-discovery.earlyaccess/ http://download.jboss.org/jbosstools/mars/staging/updates/discovery.earlyaccess/4.3.0.Beta1/ Once *DEVELOPMENT* milestones are released: http://download.jboss.org/jbosstools/mars/development/updates/discovery.central/ (contains Alpha2) http://download.jboss.org/jbosstools/mars/development/updates/ (contains Alpha2) ---- For JBDS, the URLs are essentially the same as above, but swap http://download.jboss.org/jbosstools/mars/ for https://devstudio.redhat.com/9.0/ and use 9.0 instead of 4.3... https://devstudio.redhat.com/9.0/snapshots/builds/jbosstools-discovery.central_4.3.master/ https://devstudio.redhat.com/9.0/snapshots/builds/jbosstools-discovery.central_4.3.mars/ https://devstudio.redhat.com/9.0/staging/builds/devstudio-9.0.0.Beta1-build-discovery.central/ https://devstudio.redhat.com/9.0/staging/updates/discovery.central/9.0.0.Beta1/ https://devstudio.redhat.com/9.0/snapshots/builds/jbosstools-discovery.earlyaccess_4.3.master/ https://devstudio.redhat.com/9.0/snapshots/builds/jbosstools-discovery.earlyaccess_4.3.mars/ https://devstudio.redhat.com/9.0/staging/builds/devstudio-9.0.0.Beta1-build-discovery.earlyaccess/ https://devstudio.redhat.com/9.0/staging/updates/discovery.earlyaccess/9.0.0.Beta1/ https://devstudio.redhat.com/9.0/development/updates/discovery.central/ (not yet available) https://devstudio.redhat.com/9.0/development/updates/ (not yet available) You can also see this information here: http://download.jboss.org/jbosstools/configuration/ide-config.properties -------- {quote} Is this the correct JBTIS update site? http://download.jboss.org/jbosstools/mars/snapshots/updates/integration-stack/aggregate/4.3.0.Beta1 {quote} Yes, but if you can omit the needless path segment "/aggregate/" that'd be super.	1	1	https://issues.jboss.org/browse/JBTIS-492
12659	JBTIS-492	Sure. [~nickboldt] I was using https://raw.github.com/jbosstools/jbosstools-build-ci/master/publish/publish.sh - now switching to https://raw.githubusercontent.com/jbosstools/jbosstools-build-ci/jbosstools-4.3.x/publish/promote.sh correct?	1	1	https://issues.jboss.org/browse/JBTIS-492
12660	JBTIS-487	I have migrated the four features using the guidelines in https://github.com/jbosstools/jbosstools-devdoc/blob/master/source/how_to_use_license_features.adoc and I am getting this during a build: {code} [ERROR] Failed to execute goal org.eclipse.tycho:tycho-packaging-plugin:0.23.1:package-feature (default-package-feature) on project org.switchyard.tools.feature: E:\Users\bfitzpat\github\tools\eclipse\features\org.switchyard.tools.feature\build.properties: bin.includes value(s) [license.html] do not match any files. -> [Help 1] {code} Does JBTIS expose org.jboss.tools.foundation.license.feature in a way that we need to import it as a required feature in downstream projects?	1	1	https://issues.jboss.org/browse/JBTIS-487
12661	JBTIS-487	When I look at the latest JBT IS 4.3. Alpha2 / JBDS IS 9.0.0.Alpha2 bits I see this: * in JBT: !IS-projects-using-old-license.png! * in JBDS: !IS-projects-using-old-license_2.png! is this because you didn't publish your newer bits into the IS composites/aggregates? Or is the problem in SwitchYard itself? What's the latest version available w/ the fixed features / license terms? Something newer than 2.1.0.v20150821-1353-H598-Alpha1 ?	1	0	https://issues.jboss.org/browse/JBTIS-487
12662	JBTIS-484	Hey [~blafond] - once the Mars build is stabilized - can you take a look at this? It's a mechanical change.	1	1	https://issues.jboss.org/browse/JBTIS-484
12663	JBTIS-483	[~bbrodt] - Could you look at this for your Mars build. The instructions are very straightforward. Thkx!	1	1	https://issues.jboss.org/browse/JBTIS-483
12664	JBTIS-483	Hey [~bbrodt] - it looks like your original PR was modified and pushed - is there a Mars build that has this updated? thxk!	1	1	https://issues.jboss.org/browse/JBTIS-483
12665	JBTIS-483	Hey [~bbrodt] - I don't think I can use that URL - it's a nexus internal URL - can I get one captured as a beta someplace friendlier??? Thks!	1	1	https://issues.jboss.org/browse/JBTIS-483
12666	JBTIS-483	Hey [~mbiarnes], [~KrisVerlaenen] - I'm looking for an Eclipse Mars build of drools 6.4.0 that I can access for the Integration Stack. The current version I'm pulling into the stack is: http://download.jboss.org/jbpm/release/6.3.0.Final/updatesite/ Bob found the above site that appears to be not-for-public-consumption. Is there a http://download.jboss.org/jbpm/... drools site (doesn't have to be .Final) that I could use? Also - could the above license issue be addressed as well?? Thkx!!	0	0	https://issues.jboss.org/browse/JBTIS-483
12667	JBTIS-483	Last chance [~KrisVerlaenen] [[~bbrodt] - do you have a released Mars-based 6.4.0 drools that you'd like to include in JBDSIS 9.0.0.Beta1??? i.e. http://download.jboss.org/jbpm/release/6.4.0.Final/updatesite/ # .Final is not required	1	1	https://issues.jboss.org/browse/JBTIS-483
12668	JBTIS-482	[~nickboldt] - cool script - where can I find find.sh ?	1	1	https://issues.jboss.org/browse/JBTIS-482
12669	JBTIS-471	Hey [~bbrodt] - I think org.jboss.tools.runtime.drools.detector.feature org.drools.eclipse.feature org.guvnor.tools.feature org.eclipse.bpmn2.feature org.eclipse.bpmn2.modeler org.jbpm.eclipse.feature org.kie.eclipse.feature need to change. Can you take a look at the version format of those?	1	1	https://issues.jboss.org/browse/JBTIS-471

#	Requirement	Questions	Answered	Successfully Answered	URL
12670	JBTIS-471	Did more digging and I can't see how your tests depend on JMock. So, here's another commit on the PR [1], which runs them but they fail with a curious error: [1] https://github.com/nickboldt/bpel/commit/52652ad8e6079d76d8a7d7c435945a8b25bea649 {code} [INFO] Expected eclipse log file: /home/nboldt/eclipse/workspace-eclipse/org.eclipse.bpel/tests/org.eclipse.bpel.ui.tests/target/work/data/.metadata/.log [INFO] Command line: [/opt/jdk1.8.0_60/jre/bin/java, -Dosgi.noShutdown=false, -Dosgi.os=linux, -Dosgi.ws=gtk, -Dosgi.arch=x86_64, -Xms64m, -Xmx1024m, -XX:MaxPermSize=128M, -Dorg.eclipse.swtbot.search.timeout=15000, -Dosgi.clean=true, -jar, /home/nboldt/.m2/repository1/p2/osgi/bundle/org.eclipse.equinox.launcher/1.1.1.R36x_v20101122_1400/org.eclipse.equinox.launcher-1.1.1.R36x_v20101122_1400.jar, -data, /home/nboldt/eclipse/workspace-eclipse/org.eclipse.bpel/tests/org.eclipse.bpel.ui.tests/target/work/data, -install, /home/nboldt/eclipse/workspace-eclipse/org.eclipse.bpel/tests/org.eclipse.bpel.ui.tests/target/work, -configuration, /home/nboldt/eclipse/workspace-eclipse/org.eclipse.bpel/tests/org.eclipse.bpel.ui.tests/target/work/configuration, -application, org.eclipse.tycho.surefire.osgibooter.uitest, -testproperties, /home/nboldt/eclipse/workspace-eclipse/org.eclipse.bpel/tests/org.eclipse.bpel.ui.tests/target/surefire.properties, -testApplication, org.eclipse.ui.ide.workbench, -product, org.eclipse.sdk.ide, -nouithread] Java HotSpot(TM) 64-Bit Server VM warning: ignoring option MaxPermSize=128M; support was removed in 8.0 Launching application org.eclipse.ui.ide.workbench... ---------- T E S T S ---------- java.lang.reflect.InvocationTargetException at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray2(ReflectionUtils.java:208) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:156) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:82) at org.eclipse.tycho.surefire.osgibooter.OsgiSurefireBooter.run(OsgiSurefireBooter.java:91) at org.eclipse.tycho.surefire.osgibooter.AbstractUITestApplication.runTests(AbstractUITestApplication.java:44) at org.eclipse.ui.internal.testing.WorkbenchTestable$1.run(WorkbenchTestable.java:71) at java.lang.Thread.run(Thread.java:745) Caused by: java.lang.RuntimeException: Bug in saff's brain: Suite constructor, called as above, should always complete at org.junit.runner.Request.classes(Request.java:72) at org.apache.maven.surefire.junitcore.JUnitCoreWrapper.createRequestAndRun(JUnitCoreWrapper.java:102) at org.apache.maven.surefire.junitcore.JUnitCoreWrapper.executeEager(JUnitCoreWrapper.java:85) at org.apache.maven.surefire.junitcore.JUnitCoreWrapper.execute(JUnitCoreWrapper.java:54) at org.apache.maven.surefire.junitcore.JUnitCoreProvider.invoke(JUnitCoreProvider.java:134) ... 11 more {code} See http://fossies.org/linux/privat/junit-4.12-sources.jar/org/junit/runner/Request.java for the source of that error message.	1	1	https://issues.jboss.org/browse/JBTIS-471
12671	JBTIS-471	Hey again [~bbrodt] - I'm currently picking up a bpel 1.0.5.Final from last August. Are you sure this is 1.0.5? --P	1	1	https://issues.jboss.org/browse/JBTIS-471
12672	JBTIS-471	[~bbrodt] I think the question was: Are these two sites equivalent? * http://download.eclipse.org/bpel/updates/mars/1.0.5 (new with full qualifier) * http://download.eclipse.org/bpel/site/1.0.5 (1.0.5.Final, no qualifier)	1	1	https://issues.jboss.org/browse/JBTIS-471
12673	JBTIS-471	Nick: why can't I call my next 1.0.5 something like 1.0.501 (since you're suggestion a triple digit service release number anyway)? Besides, I'm not planing on making any more changes to BPEL editor - it's basically at EOL anyway (except it still needs to be in Neon) Paul: yes, I haven't changed anything in jbosstools-bpel either.	1	1	https://issues.jboss.org/browse/JBTIS-471
12674	JBTIS-471	I think I've got them all. Can we call this one done?	1	1	https://issues.jboss.org/browse/JBTIS-471
12675	JBTIS-471	[~bbrodt] - Hi Bob - Could you tweak your build versions for consistency? Give a time of 0800 and a build 212: * JBoss Drools Core 6.4.0.201601201107 should be 6.4.0-v201601201107-0800-B212 * BPMN2 Metamodel 1.2.1.Final-v20160119-1829 should be 1.2.1.Final-v20160119-1829-B212 * BPMN2 Modeler - Diagram Editor 1.2.3.CR2-v20160209-2002 should be 1.2.3.CR2-v20160209-2002-B212 * BPMN2 Modeler - jBPM Runtime Extension Feature 1.2.3.CR2-v20160209-2002 should be 1.2.3.CR2-v20160209-2002-B212 [~bfitzpat] - Could you make a small tweak to your versions as well: SwitchYard 2.1.0.v20160211-1448-H632-Beta2 should be 2.1.0.Beta2-v20160211-1448-B632 Same for the other 2 SY features. [~apodhrad] - please comment if the above is correct.	1	1	https://issues.jboss.org/browse/JBTIS-471
12676	JBTIS-447	Hi [~apodhrad], as discussed, I have made the necessary changes to the doc. Can you please have a look? (http://docbuilder.usersys.redhat.com/23023/) Thanks - supriya	1	1	https://issues.jboss.org/browse/JBTIS-447
12677	JBTIS-447	I agree - looks good. One quick thing - we're emphasizing Fuse Tooling - in sections 4 and 5, can we move 'JBoss Fuse Development' to first in the list?	1	1	https://issues.jboss.org/browse/JBTIS-447
12678	JBTIS-438	Hey [~kcbabo] - where did you find jaxb-impl in the TP?? I found a bunch of org.eclipse.jpt.jaxb.feature. for web services but that's not what your want. Just checking... Thkx	1	1	https://issues.jboss.org/browse/JBTIS-438
12679	JBTIS-437	[~maxandersen] [~nickboldt] Seeing what Max and Nick have for an opinion. There appears to be a single required jar (~90kb). Is this a candidate for locus or http://download.jboss.org/jbosstools/updates/requirements/ ??	1	1	https://issues.jboss.org/browse/JBTIS-437
12680	JBTIS-437	Sounds more like 5 jars, if you read Keith's details re: commons-lang(2 or 3) & jackson-* /updates/requirements is where we store p2-mirrored update sites. If these are not p2 artifacts, they'd need to be bundled into Locus for you to use them as osgi dependencies. TPs can only contains osgi plugins/features; AFAIK you can't include maven deps directly w/o first re-wrapping them via Locus (or Orbit). Can you use any of the json/jackson stuff in Eclipse's Orbit? Here's the latest mirror [1] we've pulled, which includes: * com.fasterxml.jackson * org.codehaus.jackson * org.glassfish.jersey.media.jersey-media-json-jackson * com.fasterxml.jackson.jaxrs.jackson-jaxrs-json-provider [1] http://download.jboss.org/jbosstools/updates/requirements/orbit/S20150428165605/plugins/	1	1	https://issues.jboss.org/browse/JBTIS-437
12681	JBTIS-437	Hey [~nickboldt] - could you please :) take a look at this PR: https://github.com/jbosstools/jbosstools-locus/pull/43 I think it's close - need to figure out the dependencies. I've modified the locus TP but can't get an Import-Package to find the two packages I've added: Import-Package: com.fasterxml.jackson.core.jackson-databind, javax.validation.validation-api So - the pom from the import of the package lists a set of dependencies here: https://repo1.maven.org/maven2/org/jsonschema2pojo/jsonschema2pojo-core/0.4.13/jsonschema2pojo-core-0.4.13.pom Most of these dependencies Fuse Data Trans didn't need to specify in their local libs dependencies capture. Should all of these be part of the Import-Package of the jsonschema2pojo manifest? Thkx (in advance...)	1	1	https://issues.jboss.org/browse/JBTIS-437

#	Requirement	Questions	Answered	Successfully Answered	URL
12682	JBTIS-431	Updated JBTIS TP 4.2.3.CR1a repo and nexus. [~phantomjinx] - please try rebuilding teiid designer w/ jbtis tp version 4.2.3.CR1a . Also - Barry gave me a teiid des to use for the first CR1 capture for QE to test: 9.0.2.Beta1 . Should I hold off giving this to QE (waiting for an update with this TP) or should I go ahead with 9.0.2.Beta1? https://github.com/jbosstools/jbosstools-integration-stack/commit/73d3c92caaa66577cc1bacaf6ccf729b2f57d617 http://download.jboss.org/jbosstools/targetplatforms/jbtistarget/4.2.3.CR1a/	1	1	https://issues.jboss.org/browse/JBTIS-431
12683	JBTIS-429	Hey [~bfitzpat] - I've got what I need to spin an initial JBTIS 4.2.2. Do you think you'll have a SwitchYard update soon?	1	1	https://issues.jboss.org/browse/JBTIS-429
12684	JBTIS-429	It's been held up to at least the end of the week. Go ahead without me and let's shoot for a SY update for an 8.0.3?	1	1	https://issues.jboss.org/browse/JBTIS-429
12685	JBTIS-416	Hey [~mhusnain] - I should have generated this in your Jira. Is there any way to get this paragraph into the 8.0.1.GA release notes? If not then just comment as such and I'll close it.	1	1	https://issues.jboss.org/browse/JBTIS-416
12686	JBTIS-416	Thanks Misha - looks good to me. [~apodhrad] - what do you think?	1	1	https://issues.jboss.org/browse/JBTIS-416
12687	JBTIS-412	[~pleacu] where can I find the website with the link to release notes?	0	0	https://issues.jboss.org/browse/JBTIS-412
12688	JBTIS-379	[~nickboldt] - Do you think it's okay to create a stable/ final community p2 update site that has Beta or CR components or should I create two p2 update sites (similar to the jbdevstudio proposal)? What should I name them? http://download.jboss.org/jbosstools/updates/stable/integration-stack/aggregate/4.2.0.Final http://download.jboss.org/jbosstools/updates/development/luna/integration-stack/aggregate/4.2.0.Final-earlyaccess http://download.jboss.org/jbosstools/updates/development/luna/integration-stack/aggregate/4.2.0-earlyaccess	1	1	https://issues.jboss.org/browse/JBTIS-379
12689	JBTIS-355	as discussed during f2f this initially would only happen when they are GA, thus this should *not* (at first) need any new updatesite. If we want to expose development/non-final installs on marketplace for IS *then* that will be required - but then we are talking an install that is will be working/targeting development releases of JBoss Tools / Developer Studio content comes from. This might not be a problem since 8.x should be api backwards compatible anyways but it is just something both core and IS team need to be aware about. Thats why at f2f it was discussed as being done for 7.0 first (we need to cleanup the updatesite first so there is just one updatesite like for devstudio 8) and then do it for devstudio 8 when going GA. btw. I'm not following what this item is about: "In JBDS 8.1, add a new com.jboss.devstudio.core.central.feature that packages up com.jboss.devstudio.core.central plugin; include this in com.jboss.devstudio.core.feature instead of the plugin?" ? what is this solving ? what is it needed for ?	1	1	https://issues.jboss.org/browse/JBTIS-355
12690	JBTIS-355	{quote} btw. I'm not following what this item is about: "In JBDS 8.1, add a new com.jboss.devstudio.core.central.feature that packages up com.jboss.devstudio.core.central plugin; include this in com.jboss.devstudio.core.feature instead of the plugin?" ? what is this solving ? what is it needed for ?{quote} This is there to allow a JBDS IS marketplace install to include the JBDS Central feature. Today, the only way to force-install JBDS IS + Central is to install JBDS IS + JBDS, which for more slimmed-down JBDS IS installs (eg., one for just Fuse tooling) would be excessively disk-heavy. So, by splitting out the JBDS Central branding plugin into its own feature, we can have a Marketplace entry that includes JBDS Central w/o also dragging in all of JBDS. (Marketplace entries cannot include plugins, only features.) I have not yet opened a JIRA for that item because we don't have an actual schedule for JBDS 8.1 yet. If we have agreement on the notion of doing a 8.1 release, then I can open a JBDS JIRA and assign it accordingly to the correct target milestone, with dev_ACK enabled..	1	1	https://issues.jboss.org/browse/JBTIS-355
12691	JBTIS-355	I'm not following how creating a com.jboss.devstudio.core.central.feature helps this situation about having a slimmer central. The splitting up is a jboss tools level issue, not devstudio. For me there are two issues: 1) getting the list of toplevel features to install central as it is today 2) looking at cutting down the footprint of these features IF they are found too excessive. Is the above devstudio.core feature #1 ? or also #2 ? #1 my belief (at least during f2f was that this was doable today) , #2 is something we said we would *possibly* look into. neither of these we need a date for to create the jiras - create the jiras first so they can be triaged would be better.	1	1	https://issues.jboss.org/browse/JBTIS-355
12692	JBTIS-355	why are we now creating these feature groupings ? have we solved the problem of updates of these without requiring respin of them ?	1	1	https://issues.jboss.org/browse/JBTIS-355
12693	JBTIS-355	So - to be clear - the list above differs from JBoss Central connectors by: + adding JBoss Fuse Development - removing SOA 5.x Development Would it make things simpler by having the Eclipse Marketplace entries match JBDSIS JBoss Central connectors exactly? All we would need to do is keep SOA 5.x and add JBoss Fuse. Also - please review: https://docs.google.com/a/redhat.com/presentation/d/1DVVAAghJ7w2meTKG9QLA-VtO1j3jdzzKd1gstacih7U/edit#slide=id.g439643f45_2_8	1	1	https://issues.jboss.org/browse/JBTIS-355
12694	JBTIS-355	With this suggested commit [1] in this PR [2], I've added a JBoss Fuse Development connector to JBDS IS in Central. [1] https://github.com/nickboldt/jbosstools-integration-stack/commit/3948e4478abf00f8fb7d46d18f536bebe862087f [2] https://github.com/jbosstools/jbosstools-integration-stack/pull/265 So the only difference therefore between what I've set up in that PR for *JBDS IS on Eclipse Marketplace* vs. *JBDS IS in JBoss Central* is that Marketplace currently does NOT include a soa-legacy (SOA 5.x) feature. [~alansantos] [~maxandersen] [~pleacu] Could you please decide if you want a 5th "wrapper feature" in the JBDS IS Eclipse Marketplace entry to facilitate the installation of SOA 5.x stuff? My understanding from JBTIS-258 (https://issues.jboss.org/browse/JBTIS-358?focusedCommentId=13027520&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13027520) was you only wanted 4 features listed. Shall I prepare a 3rd PR and put the soa-legacy (SOA 5.x) feature back in?	1	1	https://issues.jboss.org/browse/JBTIS-355
12695	JBTIS-355	JDV is currently going GA. My current plan is to wait until B*MS and FSW goes GA and then enable all JBDSIS tooling in Eclipse Marketplace at the same time. Does that sound okay to you? (esp [~maxandersen] [~nickboldt])	1	1	https://issues.jboss.org/browse/JBTIS-355
12696	JBTIS-355	Why do we want to wait with the marketplace entry ? is it not better to get it live and tested/seen live ?	1	1	https://issues.jboss.org/browse/JBTIS-355
12697	JBTIS-355	I didn't want to stress out the first Luna tooling release since it was minimal (just the DataVirt tools) - the rest as EA. I agree it would be better to get some live testing before Fuse comes along. I'll split the difference and plan on an Eclipse Marketplace presence for JBDSIS 8.0.1 (B*MS). Then we can test all of the JDV and BPM stuff. WDYT?	1	1	https://issues.jboss.org/browse/JBTIS-355
12698	JBTIS-355	[~nickboldt] - these two hidden container features are now available (B*MS and JDV) in JBDSIS 8.0.1.GA: {code} <feature id="com.jboss.devstudio.integration-stack.bpr.feature" path="../../devstudio/site-ga/com.jboss.devstudio.integration-stack.bpr.feature"/> <feature id="com.jboss.devstudio.integration-stack.ds.feature" path="../../devstudio/site-ga/com.jboss.devstudio.integration-stack.ds.feature"/> {code} *at your convenience* - can we see what it will take to get thse in Eclipse Marketplace? Thkx	1	1	https://issues.jboss.org/browse/JBTIS-355

#	Requirement	Questions	Answered	Successfully Answered	URL
12699	JBTIS-355	[~pleacu] we just need some text for this entry and we should be okey. Here is my suggestion based on JBDS from https://marketplace.eclipse.org/content/red-hat-jboss-developer-studio-luna {code} Title: Red Hat JBoss Developer Studio w/Integration Stack <version> Red Hat JBoss Developer Studio w/Integration Stack <version> includes https://marketplace.eclipse.org/content/red-hat-jboss-developer-studio-luna[JBoss Developer Studio] plus: * BRMS Tooling * JBoss Data Virtualization JBoss Developer Studio w/Integration Stack adds support for JBoss and related technologies, such as Teiid and Drools. For a quick overview, see <link to docs/overview?> More videos here: http://docs.jboss.org/tools/movies/ See also: http://www.redhat.com/products/jbossenterprisemiddleware/developer-studio/ This release was built with Eclipse 4.4.2 (Luna SR2), but should also work with other Luna or Eclipse 4.4.x versions. Luna SR2 JEE bundle recommended: http://www.eclipse.org/downloads/packages/eclipse-ide-java-ee-developers... NOTE: Use with latest Java 7 higher than jdk-7u45 or Java 8 versions. NOTE: If you encounter problems, please open a https://jira.jboss.org/jira/browse/JBDS[JIRA]! {code} WDYT ?	1	1	https://issues.jboss.org/browse/JBTIS-355
12700	JBTIS-355	I assume the only two features we need are: {code} com.jboss.devstudio.integration-stack.bpr.feature com.jboss.devstudio.integration-stack.ds.feature {code} since they should auto include devstudio, correct ? btw. what does bpr and ds stand for ?	1	1	https://issues.jboss.org/browse/JBTIS-355
12701	JBTIS-354	Should we close this since now not the priority and we will do it for JBDS 8 based installls instead and don't have to reorg JBDS 7 updatesites to support this then ?	1	0	https://issues.jboss.org/browse/JBTIS-354
12702	JBTIS-348	[~mmusaji] - are you able to access [3] above? Those are the correct rel notes - not sure why the customer can't see them or who translates them...	1	1	https://issues.jboss.org/browse/JBTIS-348
12703	JBTIS-348	[~ganandan-Redhat] and [~lcarlon] - The release notes in [3] above are accessible and correct. Is there another step that's missing here - the customer (in Japan) can't download them and I am unaware of any further process. Thkx	1	1	https://issues.jboss.org/browse/JBTIS-348
12704	JBTIS-348	Thanks, [~mmusaji]. [~nmoore], can you work with Paul to convert and stage the release notes on the CSP?	1	1	https://issues.jboss.org/browse/JBTIS-348
12705	JBTIS-348	Hi Paul, I've spoken with Michelle Murray, and she has provided me with a procedure to create a release notes document which can be downloaded and printed. The procedure involves scraping the release notes text from the relevant Jiras using a script. I have noticed that most (if not all) of the Jiras are closed and do not have any release notes text. This procedure is used by the JDBS documentation team and would be a good idea to follow the same procedure as Michelle suggested that the JDBS IS doco should be hosted on the JDBS documentation portal page. Michelle suggested that we reopen the closed Jiras, update the settings to say that the Jiras affect the Release notes, and add sample release notes text to the Jira. After that the developers would need to add the release notes text, then we can close the Jiras again. I'm happy to do the reopening and setting changes, if I have permission, and Michelle has provided sample release notes text, which I can add to each Jira. Would you be able to persuade the developers to add release notes text to the Jiras? The procedure is outlined in this Mojo doc - https://mojo.redhat.com/docs/DOC-944745 Thanks, Nichola Moore	1	1	https://issues.jboss.org/browse/JBTIS-348
12706	JBTIS-348	Hey [~nmoore] - I think the immediate issue is to simply convert the existing 7.0.3.GA JBDSIS release notes ([3] above) into a CSP version. I don't know what the procedure is for that (does docs create a ticket to update the CSP?). As for the extended procedure mentioned above - it appears great but I'd bet there would be a fair amount of angst re-opening all of those Jiras. I also don't have the authority to get the component owners to accept this procedure.	1	1	https://issues.jboss.org/browse/JBTIS-348
12707	JBTIS-348	Paul Leacu - I've created a version of the Release Notes using our usual style sheet. I know Michelle uses a specific release notes style sheet so hers looks slightly different, but this has all we need (plus a few bits we don't but I can't get rid of those!). What we need now is to update the JBoss Developer Studio splash page so that the customer can see the link to the release notes. Are the JDBS and JDBSIS releases numbered the same? I'm afraid I can't even send you a link so you can have a look. There's no presence on our stage server to link to. I'll try to work it out on Monday. I'm just letting you know that work is in progress and it's all not too far away! Regards, Nichola	1	1	https://issues.jboss.org/browse/JBTIS-348
12708	JBTIS-348	Hi Paul, Finally! Please can you have a look at http://documentation-devel.engineering.redhat.com/site/documentation/en-US/Red_Hat_JBoss_Developer_Studio_Integration_Stack/7.0/html-single/7.0.3_Release_Notes/index.html and let me know if there needs to be any changes. Thanks, Nichola	1	1	https://issues.jboss.org/browse/JBTIS-348
12709	JBTIS-339	{quote}Where "direct-references" is what we do now, with a simple <references> block added to content.jar!/content.xml, and "composite-with-references" is the more complex tree structure we're proposing?{quote} Yes. How much do you like it?	1	1	https://issues.jboss.org/browse/JBTIS-339
12710	JBTIS-339	There is no reference at all. Did you apply that on a site where you have some "associateSites" ?	1	1	https://issues.jboss.org/browse/JBTIS-339
12711	JBTIS-339	So how does an IS component pom change from the old/bad assocSites in order to build an independent update site with references?	1	1	https://issues.jboss.org/browse/JBTIS-339
12712	JBTIS-339	Having compositeArtifacts different is normal, because the "references" repo doesn't contain any artifacts, it contains just references. Are IS component inheriting your pom configuration? If yes, then you only need to change the IS parent pom. Note that Mojo is currently a SNAPSHOT version, so we'll probably need to release it before using it more widely.	1	1	https://issues.jboss.org/browse/JBTIS-339
12713	JBTIS-339	Great! So I'm merging my PR, then repository-utils 0.21.1-SNAPSHOT will contain this feature. When do you plan to make a release of JBT-IS ? We'll make a release of the tycho-plugins just before that.	1	1	https://issues.jboss.org/browse/JBTIS-339
12714	JBTIS-339	Ok. technically, I can do it from now on. However, I prefer to make sure we won't do that for nothing and we won't need to re-release tomorrow. Have you validated, on JBT-IS side that the withReferences vs "plain" approach satisfies your use-cases? On JBT, we'm happy with using the plain URL in discovery and having to composite it with whatever seems necessary, and keeping control on the sites we reference/composite.	1	1	https://issues.jboss.org/browse/JBTIS-339

#	Requirement	Questions	Answered	Successfully Answered	URL
12715	JBTIS-339	Ok - good point. So now JBTIS can create a fully resolved "withReferences" composite update site and in theory the existing JBTIS discovery poms should continue to work. I'll verify that. So JBDSIS (which draws from JBTIS) doesn't carry associateSites - it pulls from a mirror-target-to-repo. Is that not affected by this? We pass into our IS poms upstream discovery and composite urls. https://github.com/jbosstools/jbosstools-integration-stack/blob/master/jbosstools/pom.xml https://github.com/jbosstools/jbosstools-integration-stack/blob/master/devstudio/pom.xml The JBTIS/JBDSIS discovery plugins has embedded 8.0-development/central/integration-stack in them - does that need to go? https://github.com/jbosstools/jbosstools-integration-stack/blob/master/jbosstools/discovery/org.jboss.tools.central.discovery.integration-stack/plugin.xml#L27 https://github.com/jbosstools/jbosstools-integration-stack/blob/master/devstudio/discovery/com.jboss.jbds.central.discovery.integration-stack/plugin.xml#L56 I need to get fuseide functional as an independently installable update site (withReferences) and have it supply a "references" plain site to the IS. IS that just adding the same composite reference strategy? So - I need some help from [~mickael_istria] and [~nickboldt] to just look over the IS poms with me if you can spare the time.	1	1	https://issues.jboss.org/browse/JBTIS-339
12716	JBTIS-339	[~leacu] the "fuseide independent installable update site" is *just* for community testing, right ? the supported version will come from JBDSIS updatesites and planned to be put on marketplace. correct ? btw. if fuseide updates into JBTIS anyway - is there a reason for them to not just tell users to use the JBTIS update site ? Just one quick thing I noticed is that jboss.is.update.url should be changed to match the naming in ide.properties.	0	0	https://issues.jboss.org/browse/JBTIS-339
12717	JBTIS-339	btw. why not just have fabric8 in JBTIS ? excluded in JBDSIS like we had for savara ?	0	0	https://issues.jboss.org/browse/JBTIS-339
12718	JBTIS-339	[~mickael_istria] - I see you've merged the PR for this feature in repository-utils 0.21.1-SNAPSHOT. I need to create the JBTIS 4.2.0.Beta1 release early next week so I'll need the updated tycho-plugins - is it available? I believe the withReferences vs "plain" approach satisfies the primary use case. Thanks!	1	1	https://issues.jboss.org/browse/JBTIS-339
12719	JBTIS-339	Yes 0.21.1-SNAPSHOT with this change is already available in JBoss Snapshot repo (https://repository.jboss.org/nexus/content/repositories/snapshots/org/jboss/tools/tycho-plugins/0.21.1-SNAPSHOT/). You can already use it. I can let you have fun with it and unless you find something wrong, I'll release it as 0.21.1 on Monday morning (european time). Is it good for you?	0	0	https://issues.jboss.org/browse/JBTIS-339
12720	JBTIS-332	Merging these changes broke the build at hudson.eclipse.org. After spending about a day on trying to figure this out, I finally gave up and reverted the changes. The existing build did generate an SDK feature which included source jars for all of the plugins - why was that not sufficient? Anyway, unless someone smarter than I can figure out what the problem is, I'm going to leave things as they are...	0	0	https://issues.jboss.org/browse/JBTIS-332
12721	JBTIS-331	what kind of earlyaccess are we talking here ? on jboss tools I assume this is just being added into jbosstools development updatesite - no special thing here (sure we can mark it as earlyaccess since it wont be there when jbosstools GA's). But for it to be visibiel in JBoss Tools Final and from developer studio the earlyaccess content needs to goto the earlyaccess p2 repository, just changing the certificate ID does not do anything on its own. This paragraph: "It is not necessary to refactor (or copy) your plugin to include .earlyaccess in the name, UNLESS you plan to have one plugin for "early access" and one for "released" content, as we do in JBT/JBDS. (You can also include both types of content in the same plugin.)" sounds *very* wrong to me. Why does IS not need to have separate earlyaccess content like core do ? what makes it okey their updatesites gets automatically added ?	0	0	https://issues.jboss.org/browse/JBTIS-331
12722	JBTIS-331	I'm not talking about the discovery file - that one can be one or two (I prefer two since then it is consistent). What i'm talking about is that the PR's I see and the text above talks as if adding the certificate flag to the connector is enough. We can't add the JBTIS-*development* update site url into our Final or GA discovery catalogs. It has to be on earlyaccess otherwise we are going to mix streams/updatesites/TP's afaics. Users Help > Update will be affected without them or us wanting that. Again, why is JBT-IS not doing it like we do for our earlyaccess ?	0	0	https://issues.jboss.org/browse/JBTIS-331
12723	JBTIS-331	{quote} Also - I assume we are going to set jboss.is.update.url to be http://download.jboss.org/jbosstools/updates/development since thats how it would work ever since https://issues.jboss.org/browse/JBDS-3086 have been made in effect ? (central is now composite into / of the main updatesite) {quote} Do we want to consider JBT-IS as one of the main update sites (just like Core or Central)? If so then you suggestion is probably what we'll need to do. When JBT is still in Beta and JBT-IS too: * Set jboss.is.update.url to http://download.jboss.org/jbosstools/updates/development in ide-config.properties * Make http://download.jboss.org/jbosstools/updates/development also reference JBT-IS development site When JBT is released and JBT-IS is still in Beta * Update JBT-IS discovery jar to add Early-Access * Set jboss.is.update.url to JBT-IS develpment site * Remove JBT-IS development site from http://download.jboss.org/jbosstools/updates/releases/luna When JBT is released and JBT-IS too * Update JBT-IS discovery to remove Early Access * Add JBT-IS release site to http://download.jboss.org/jbosstools/updates/releases/luna * Set jboss.is.update.url to http://download.jboss.org/jbosstools/updates/releases/luna in ide-config.properties	0	0	https://issues.jboss.org/browse/JBTIS-331
12724	JBTIS-331	So latest release of JBT-IS is always on Central site and accessible together with latest release of JBT (at the same releases/luna URL). Milestones of JBT-IS are added to Early-Access catalog (and have their own URLs). Is that it?	0	0	https://issues.jboss.org/browse/JBTIS-331
12725	JBTIS-331	no, ...why is this so hard ;) we have the following sites for our plugins: /nightly /staging /development /stable (edited, was release by error before) Each matching the state we are in various development phases. in past we had integration-stack, central, etc. versions of these handed out as public separate sub-urls The change in JBDS-3086 is that these are now composited so we only need to use the root of these sites (/development, /release) when published to users to consume it. Internally we still have individual sites for compositing/testing, but they aren't supposed to be exposed as stable urls to users. Then the intent is that when we do milestones we make it available in the matching category. i.e. right now we got JBoss Tools CR1 in /development and optimally we would get JBoss Tools IS Beta into /development too and thus made available in jboss central for install. when JBoss Tools core goes Final, IS Beta can't just be listed as being installable thus IF we want it available we have /earlyaccess site to provide this and users will have to go and explicit enable EA to get it. Since you can't just have Final central pointing to /development site since then you will be mixing JBoss Tools 4.2.1.x into the release unknowingly. I can see it is probably a bad thing if when are in Core Alpha/Beta/CR1 that IS won't show up as Earlyaccess when it will be that in Final/GA of tools. Thus we could still use the EA marker, but still at GA time integration stack needs to physically be in /earlyaccess location not in /development.	0	0	https://issues.jboss.org/browse/JBTIS-331

#	Requirement	Questions	Answered	Successfully Answered	URL
12726	JBTIS-331	So we've identified that repository references are making things possibly inconsitent. JBIDE-18541 should provide a tool to avoid direct repository references but still providing ability to link other repos silently. But this is a mid-term solution. For a first iteration, here are a few suggested actions: * Have earlyaccess site compositing http://download.jboss.org/jbosstools/updates/development/luna/integration-stack/aggregate/ * -Have integration stack discovery catalog reference on the JBT URL (from ide-config.properties): jboss.discovery.site.url and jboss.discovery.earlyaccess.site.url (as opposed to what I suggested to Paul with creation of a dedicated property for integration-stack)- We do want to keep the specific is property for testing, and give ability to reference other property, eg: jboss.discovery.site.integration-stack.url|jbosstools|4.2.0=${jboss.discovery.site.url} * Remove references from early-access aggregate site to JBT core development	0	0	https://issues.jboss.org/browse/JBTIS-331
12727	JBTIS-320	Objection noted by Jeff DeLong: On Sep 15, 2014, at 8:58 AM, Alan Santos wrote: Thanks Jeff, does this mean your team is actively using it for partner education Yes or on-boarding? If so can you elaborate on how We have a module in our online FSW training that covers pretty much what is in the video below including a lab as well. I.e. we generate the choreography from a set of scenarios and SY services from the Choreography. We generated service are uploaded into DTGov, and we flush them out in subsequent labs. and are you aware of any partners who are using it in a similar manner? No. Sorry. It is hard to track individual partner usage in all but a few cases with the largest partners where we get engaged with them on supporting their projects. Thanks, Jeff -alan On Sep 15, 2014, at 10:56 AM, Jeffrey DeLong <jdelong@redhat.com> wrote: Alan, - BPMN2 choreography generation from scenario diagrams - SwitchYard service generation from BPMN2 choreography Essentially the features described in http://vimeo.com/53007939 Jeff Jeffrey DeLong Sr. Manager Global Partner Enablement http://www.redhat.com/partners/partner-enablement/ http://youtu.be/EdsvGcfIoJs cell: 801.554.4185 On Sep 15, 2014, at 8:51 AM, Alan Santos wrote: Jeff, can you share how your team and any partners are using Savarra? thanks On Sep 15, 2014, at 10:13 AM, Jeffrey DeLong <jdelong@redhat.com> wrote: I object from removing something we have spent 5 years working on and testing. Jeff Jeffrey DeLong Sr. Manager Global Partner Enablement http://www.redhat.com/partners/partner-enablement/ http://youtu.be/EdsvGcfIoJs cell: 801.554.4185	1	1	https://issues.jboss.org/browse/JBTIS-320
12728	JBTIS-318	Hey [~KrisVerlaenen] - I noticed that jBPM3 is actually generating the source bundles correctly. I think what I need is something like this for each feature pom.xml: <build> <plugins> <plugin> <groupId>org.eclipse.tycho.extras</groupId> <artifactId>tycho-source-feature-plugin</artifactId> </plugin> <plugin> <groupId>org.eclipse.tycho</groupId> <artifactId>tycho-p2-plugin</artifactId> </plugin> </plugins> </build> wdyt?	1	1	https://issues.jboss.org/browse/JBTIS-318
12729	JBTIS-318	Hey Kris - Currently - drools/jBPM6/guvnor are the only IS components not generating source bundles. What's interesting is jBPM3 is generating them correctly. Is this something you could look at fairly soon - it's needed for JBDSIS 8.0.0. Thanks, --paull	1	1	https://issues.jboss.org/browse/JBTIS-318
12730	JBTIS-295	[~maxandersen] - Is this Jira intended for the IS infrastructure itself to add support for integration testing or for the IS components (i.e. SwitchYard) to add packaged bundle testing support?	1	1	https://issues.jboss.org/browse/JBTIS-295
12731	JBTIS-295	[~maxandersen] [~mickael_istria] - Opinions wanted... Looking at some older Jira - I believe all of the IS components now use the IS TP. Both JBTIS and JBDSIS are run through the p2 director install tests here: https://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/jbosstools-install-p2director.install-tests.integration-stack.matrix_master/ Also - SwitchYard uses the JBTIS TP and performs their own unit testing prior to being aggregated. Do you guys see a good use for using Mickael's p2installed surefire runtime test support at the IS level or does it not apply given unit testing at the component level? Thkx!	1	1	https://issues.jboss.org/browse/JBTIS-295
12732	JBTIS-295	Whent SwitchYard (as an example) uses the JBTIS TP to run its unit test against it it will *not* necessarily load all parts of JBDS - it will only load/start the bundles your test and its dependencies declare. With Michael's approach here you can run tests inside a specific version of developer studio (i.e. devstudio 9.1 beta) instead of just the dependency set. Thus I think it would help catch other kind of errors - but that would probably need new tests that are more integration level tests (i.e. does SwitchYard start in JBDS without any errors as opposed to testing SY unit level features works).	0	0	https://issues.jboss.org/browse/JBTIS-295
12733	JBTIS-295	[~apodhrad] - Hey Andrej - I know you've got testing infrastructure that utilizes the tycho testing support where it loads a specific JBDS. I think at one point I tried running those but they were too comprehensive for a smoke test tied to the JBDS build. Is it possible to construct some fast smoke tests for Fuse and maybe Teiid Designer? WDYT?	1	1	https://issues.jboss.org/browse/JBTIS-295
12734	JBTIS-295	Actually, the "fast" suite of smoke tests for Fuse tooling is in progress. Our plan is to run such smoke tests for each PR (just with JBTIS dependencies as a verification job). Moreover, every night we execute some subset of tests for the latest build of Fuse tooling with the latest version of JBDS (corrently only with 9.x). Something similar we plan for the Teiid Designer but not for the other components (there are no changes or only small changes). I'm not sure what's the exact goal of this issue. All work mentioned above will be done on Fuse (resp. TD) side so that developers can immediately see what's wrong.	1	1	https://issues.jboss.org/browse/JBTIS-295
12735	JBTIS-290	Paul & Max, I've added the #2 part to the appropriate ModeShape JCR plugin. Will one of you be doing the #1 part for the ModeShape JCR feature? Or would you like me to do that? Also, when I try to build locally now the *{{usage plugin}}* cannot be found. In the past I've used *{{-Punified.target}}* when building locally. Is the usage plugin found somewhere else now? Thanks	1	1	https://issues.jboss.org/browse/JBTIS-290
12736	JBTIS-290	usage plugin have not moved since it was created. What are you trying to build exactly ?	1	1	https://issues.jboss.org/browse/JBTIS-290
12737	JBTIS-290	You must have put a "settings.xml" file at the root of the ModeShape Tools code because there is not one there. And whatever you have in it is making it all magically work:-) Can you send me yours or at least the parts of it that you think would help. Thanks Paul.	1	1	https://issues.jboss.org/browse/JBTIS-290
12738	JBTIS-290	Hey [~elvisisking] - did this ever get resolved? We need at least one feature resolved for Modeshape for 8.0.x. Have you added: org.jboss.tools.usage;resolution:=optional;x-installation:=greedy into your MANIFEST.MF ? We're currently tracking MODESHAPE("org.jboss.tools.modeshape.rest.feature", "org.jboss.tools.modeshape.rest"), thkx, --paull	1	1	https://issues.jboss.org/browse/JBTIS-290
12739	JBTIS-290	I left comments on the PR but I'll repeat it here in summary: A) the names gets appended into a lenght limited string thus we want as *short* names as we can, thus TEIID is better than TEIIDDESIGNER B) we want to limit what we track to truly unique features - for example we don't no more track ARCHIVES since its always part of Server now. In that context, is Guvnor really relevant to track independently ?	0	0	https://issues.jboss.org/browse/JBTIS-290

#	Requirement	Questions	Answered	Successfully Answered	URL
12740	JBTIS-289	Hey [~bfitzpat] - I'm assembling a kepler maintenance release of the IS now (JBTIS 4.1.6/ JBDSIS 7.0.3) - should we build/promote a kepler version of ESB with your examples fix?	1	1	https://issues.jboss.org/browse/JBTIS-289
12741	JBTIS-289	Hey [~bfitzpat] - the staging area's last build is from July - don't know why it didn't stage. Could you rebuild the 4.1.x version again and let me know when it's ready? I'll promote it form staging and update the stack.	1	1	https://issues.jboss.org/browse/JBTIS-289
12742	JBTIS-278	Hey [~nickboldt] - would the link go here [1] or at the download page for jbds 7.1.1? Also - is there a dummy login so I can see what the CSP JBDS download page looks like? [1] https://access.redhat.com/site/products/red-hat-jboss-developer-studio/	1	1	https://issues.jboss.org/browse/JBTIS-278
12743	JBTIS-278	I would list it with the JBDS 7.1.x content. That's what we did for JBDS 5: https://access.redhat.com/jbossnetwork/restricted/listSoftware.html?downloadType=distributions&product=jbossdeveloperstudio&version=5.0.0 As to a login, you can register using your redhat address. Not sure how? Try this link: https://www.redhat.com/wapps/sso/login.html?redirect=%2Fwapps%2Feval%2Findex.html%3Fevaluation_id%3D1002 Then go here: https://access.redhat.com/jbossnetwork/restricted/listSoftware.html?product=jbossdeveloperstudio&downloadType=distributions	0	0	https://issues.jboss.org/browse/JBTIS-278
12744	JBTIS-253	[~phantomjinx] what is that script that sounds like it is "blindly" setting include/exclude versions on all plugin dependencies ? Sounds like that could lead to some funky limitations in upgrade paths....or that you made something great that we could use elsewhere. Got a link ? Great to hear things seem to work even after our major change in server adapters. Did you try run it too or just build it for now ?	1	1	https://issues.jboss.org/browse/JBTIS-253
12745	JBTIS-245	Last years promises was that we would at least get alpha builds running so we are not yet repeating not having IS tooling at all available for next release of JBDS 8. Could we at least do one that contains the components that will work on Luna out of the box ? We have many things that need IS feedback for the changes we are doing in JBT core/JBDS core like based on bug reports from users and IS QE: * Runtime less server to allow defining a server without having a local runtime * Refactored JBoss server adapter to have better management support (incl. wildfly/eap 6.x new requirements) * Allow installation of earlyaccess components from Central Luna itself is what will provide: * Java 8 support * Support for eclipse on newer Linux versions (GTK3 on fedora 20/Ubuntu) - kepler will and does have issues. * much faster maven import/support ...and of course bugfixes etc.	1	1	https://issues.jboss.org/browse/JBTIS-245
12746	JBTIS-241	Shouldn't this be logged against the TEIID project, not JBTIS?	1	1	https://issues.jboss.org/browse/JBTIS-241
12747	JBTIS-241	Bear in mind that the reason that Server Tools went from 2.4 to 3.0 is big api changes and refactorings. Might also precipitate a similar version change in Teiid. Also, I'm not sure that we should expect the same version of Teiid to work on both Kepler and Luna. Should we?	1	1	https://issues.jboss.org/browse/JBTIS-241
12748	JBTIS-241	you have some points. We do wire into AS api to discover Teiid submodule existance and contributing to the Server view contents. Until we investigate/upgrade to Luna target platform and new AS 3.0 version, we won't know for sure. From a Product standpoint, what will we be supporting for the rest of the year? DV6/DV6.1 == JBDS 7.x ? If so, then compatibility with JBDS 8.0 will be left to JBTIS 8.0 build?	0	0	https://issues.jboss.org/browse/JBTIS-241
12749	JBTIS-241	We also have to set expectations. * For a user wishing to use the combination of JBDS 8.0.alpha2 and Teiid - which tooling should he use - and when will it be available? * Must this user wait for JBDSIS 7.0.1? * What about users wishing to use JBDS 8.0.alpha2 and FSW6 - do we think that they can safely use JBDSIS 7.0.0?	0	0	https://issues.jboss.org/browse/JBTIS-241
12750	JBTIS-235	How about a new (optional) update site which contains doc plugins for contributing components? devstudio.jboss.com/updates/4.0/help ?	0	0	https://issues.jboss.org/browse/JBTIS-235
12751	JBTIS-235	The "eclipse" help generated from docbook using Maven, currently has no "branding" applied. Current "Hibernate" help has none. So when we include "Help" from each project in JBDS, do we build it with "JBDS/Red Hat" branding? or leave it generic? If "generic", then there would be no changes required, moving from Project to Product builds. If you want "branding", then there will be a change in Help "source" that needs to be applied for Product. This will be difficult if we're pulling "bits" from projects to directly create the bits for the product.	0	0	https://issues.jboss.org/browse/JBTIS-235
12752	JBTIS-235	[~pleacu] is this something we could move to JBTIS? Are there resources for such a thing? If not we can just close this as Won't Fix or Rejected, I guess.	1	0	https://issues.jboss.org/browse/JBTIS-235
12753	JBTIS-221	Complete list of affected files: [^anonsvn.jboss.org_refs_in_examples_xml.txt] Workspaces in use include: * [~snjeza] (seam, portlet, resteasy, weld, richfaces, jbossas7) * Denny (esb) * [~elvisisking] (teiid designer) * [~jdoyle] (teiid designer) * Grid (bpel, esb) * [~rcernich] (esb) * [~fbricon] (archetypeRepo?) * [~bfitzpat] (SOA 5.2) * [~dgolovin] (richfaces) So... the people mentioned above (except for Denny and Grid) need to migrate this content to a new location in github. Do we want to host all this stuff in a new repo, or just put it into https://github.com/jbosstools/jbosstools-download.jboss.org/tree/master/jbosstools/examples ? Should the JBTIS stuff be split out into a separate repo? (cc: [~pleacu])	1	0	https://issues.jboss.org/browse/JBTIS-221
12754	JBTIS-221	If we move these files, how many older installations of JBDS pointing to those paths will need to be updated? We're going back 3 or 4 years on some of these examples aren't we? I just worry about all the potential breakage coming down the line. As far as the ESB examples from Denny & Grid, I've moved those to github in the esb repository for 4.1.x and master. They are here now: https://github.com/jbosstools/jbosstools-esb/tree/master/quickstarts	1	1	https://issues.jboss.org/browse/JBTIS-221
12755	JBTIS-221	I'm not suggesting we MOVE old stuff, but rather that we COPY the old stuff into a new place from which ongoing maintenance would occur (like when we migrated from SVN -> Github). So the old stuff, if referenced by JBDS 7 and earlier, would still exist in the old/ugly locations, but from JBDS 8 onward only the CORRECT URLs would be used in the examples *.xml files on download.jboss.org. I also like the idea of co-locating these examples in the projects for which they are relevant, as Brian did for ESB. [~bfitzpat] did you also get the examples from Rob C, and the BPEL examples? And what of the floor show that had been spoken of? In an empty house? In the middle of the night? What diabolical plan had seized Fitz' crazed imagination? [1] [1] http://lmgtfy.com/?q=and%20what%20of%20the%20floorshow?	1	0	https://issues.jboss.org/browse/JBTIS-221
12756	JBTIS-219	Support for secure storage password: <appArgLine>-eclipse.password resources/security/password</appArgLine> where password is a file containing the password	1	1	https://issues.jboss.org/browse/JBTIS-219
12757	JBTIS-197	[~apodhrad]: Just noticed that https://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/jbosstools-integration-stack-tests.aggregate_41 is COMPILING but not RUNNING the tests. Should it be RUNNING them too? (We could just remove the -DskipTests flag in the maven config.)	0	0	https://issues.jboss.org/browse/JBTIS-197
12758	JBTIS-187	Could you set Fix Version? It is in development but what version it is targeted for?	1	1	https://issues.jboss.org/browse/JBTIS-187

#	Requirement	Questions	Answered	Successfully Answered	URL
12759	JBTIS-152	Hey Phantom - We're trying to isolate the component dependencies to the JBTIS TP. For JBDS-IS, we will only be providing plugins from the TP site, so if it's not there, users won't be able to install. Are you using the latest JBTIS TP? https://repository.jboss.org/nexus/content/repositories/releases/org/jboss/tools/integration-stack/target-platform/4.1.2.Final/ https://github.com/jbosstools/jbosstools-integration-stack/blob/master/target-platform/integration-stack-base.target#L111-119 This shows the jackson core/mapper in it.	1	1	https://issues.jboss.org/browse/JBTIS-152
12760	JBTIS-152	This is to allow a user to install Teiid Designer plugins without needing to add the orbit repository as a separate update site in their Eclipse IDE. As it stands, Teiid Designer plugins do not install in a virgin eclipse installation, using the JBT-IS update site, due to lacking access to the jackson plugins. The pull request fixes that -without displaying the contents of orbit to the user- but does insert orbit as an associative update site in the user's available update sites. So is there is another way around this? From Nick, we have had [https://issues.jboss.org/browse/TEIIDDES-1843], which adds the dependencies to the Designer update site. Is that going to be a preferred option, despite my misgivings (see my comment on that JIRA).	0	0	https://issues.jboss.org/browse/JBTIS-152
12761	JBTIS-118	It seems like this is an issue in the target platform (Locus is clean). Can you try including the necessary bundles from Gemini in your target-platform? Locus won't include bundles that are available in other p2 repository.	0	0	https://issues.jboss.org/browse/JBTIS-118
12762	JBTIS-118	[~phantomjinx]: Can you rebuild / reinstall Teiid Designer as you did initially, and report back if the above changes have helped to resolve the missing packages? If not, can you provide steps to reproduce the console messages? I don't see them when building or installing, as noted above.	1	1	https://issues.jboss.org/browse/JBTIS-118
12763	JBTIS-113	Should this issue move to the new JBTIS project? The ability to have more control over dependencies could be resolved by component-specific target platforms (ie each component uses a target-definition or some repositories in its pom to specify dependencies).	0	0	https://issues.jboss.org/browse/JBTIS-113
12764	JBTIS-110	Rob, do you have any ideas on this one?	1	1	https://issues.jboss.org/browse/JBTIS-110
12765	JBTIS-97	Hey Paul - Would you like this for Kepler+ releases (JBTIS 4.1.0 SOA-P 6.0 Beta 1 and beyond)? Is it okay to skip modifying the Juno release? Thanks! --paull	1	1	https://issues.jboss.org/browse/JBTIS-97
12766	JBTIS-93	What version of designer is included in this integration stack build? This change has been implemented in Teiid Designer 8.2.0 (current master). Do you need an alpha release? in what timeframe?	1	1	https://issues.jboss.org/browse/JBTIS-93
12767	JBTIS-93	Hey Barry/Paul - This is needed for the Kepler SOA-P 6.0 Beta 1 which is using the same URL as the Juno Alpha 2: http://download.jboss.org/jbosstools/updates/release/juno/soa-tooling/teiiddesigner/8.1.0.Final/ Would you like me to grab a newer release for the Kepler release? Thanks, --paull	1	1	https://issues.jboss.org/browse/JBTIS-93
12768	JBTIS-93	Hey Barry/Paul - The powers-that-be want a JBDSIS 7.0.0.Beta1 with just SOA 5.3 support in it asap. This replaces the JBDSIS 6.0.0 that was built from Juno. That means I'll need your kepler teiid 8.2.Alpha soon so I can build the production JBDSIS p2 installer for SOA 5.3 support. I'll use the same teiid 8.2.Alpha for the JBTIS 4.1.0 kepler community release (2 for 1!). Please let me know when it's ready (urgency++). Thanks! --paull	1	1	https://issues.jboss.org/browse/JBTIS-93
12769	JBTIS-88	Hey guys, Because emf compare 2.1 and graphiti 0.10 are part of kepler, which is already mirrored, can you just update the target platform to pull these features/bundles from the kepler mirror on jboss.org? Seems like that would be simplest. Also, Paul, once all these changes are complete, can you release a 4.0.1 target platform that I can use for pulling in these dependencies? Thanks, Rob	1	1	https://issues.jboss.org/browse/JBTIS-88
12770	JBTIS-88	Hey [~nickboldt] - No dice on ssh access (to either). I'll work on that - could you mirror these for me so I can create an aggregated TP for JBTIS 4.0.1 (SOA-P 6 Alpha 2)? Thanks!!! --paull http://download.eclipse.org/bpel/site/1.0.3/ http://download.eclipse.org/bpmn2-modeler/updates/kepler/0.2.6/ http://download.eclipse.org/bpmn2-modeler/updates/juno/0.2.6/	1	1	https://issues.jboss.org/browse/JBTIS-88
12771	JBTIS-83	Next up will figure out where to create an examples plug-in and include it in the right places... This is the IRC chat I'm basing some of this on: (6:52:15 AM) maxandersen: pleacu: the idea is not to add *more* connectors, but to make the right examples be part of the *existing* connectors. (6:52:16 AM) pleacu: maxandersen: its commented out anyway - but the idea was to install the exAmples - this will be your legacy-examples connector (6:52:53 AM) pleacu: ahh - ok - could add IUs to the JBoss SOA legacy development connector (6:52:56 AM) maxandersen: pleacu: yeah so the right connector to add these legacy examples is this: https://github.com/jbosstools/jbosstools-integration-stack/blob/master/devstudio/discovery/com.jboss.jbds.central.discovery.integration-stack/plugin.xml#L107 (6:53:08 AM) pleacu: yep (6:53:11 AM) maxandersen: but we'll still need some plugin where this extension gets added. (6:53:59 AM) maxandersen: that extension would depend on central/project examples in core. (6:54:08 AM) maxandersen: I mean that plugin would depend on that. (6:55:59 AM) maxandersen: so if we can find a good place for those we should be good. (6:56:37 AM) maxandersen: hava integration-stack.soa-legacy.examples bundle that provides this ? (6:56:54 AM) maxandersen: or push down into esb / bpel to declare it ? (6:57:16 AM) bfitzpat: do we really want to split 'em up or keep 'em together as we've talked about doing? (6:57:35 AM) pleacu: maxandersen: don't know... (6:58:10 AM) maxandersen: bfitzpat: split up : no need to introduce new plguin at IS stack level, keep together: should be done at IS stack level is how i I see it. (6:58:22 AM) maxandersen: doesn't make much sense esb enables bpel examples and the other way around does it ? (6:58:40 AM) maxandersen: then if you install esb you'll get bpel examples that wont work well and/or vice versa ?	1	1	https://issues.jboss.org/browse/JBTIS-83
12772	JBTIS-83	As these are NEW properties to be added to the ide-config.properties file, then we can safely add the above lines to the file so you can test it. https://github.com/jbosstools/jbosstools-download.jboss.org/pull/387 However, I notice your filepath uses a specific commitID and that the filename has "7.1" in the name, rather than 8.0. Is that intentional? If you used "latest in a branch" rather than "specific SHA hash" you would be able to more easily update this on the fly w/o having to update the IDE config file.	1	1	https://issues.jboss.org/browse/JBTIS-83
12773	JBTIS-83	Ok. Have updated all the paths to the zips in the project examples xml file to NOT be the static revision in the latest PR - https://github.com/jbosstools/jbosstools-esb/pull/56/files	1	1	https://issues.jboss.org/browse/JBTIS-83
12774	JBTIS-83	[~pleacu] - Beautiful! Will we have an updated Kepler one as well once we get BPEL worked out?	1	1	https://issues.jboss.org/browse/JBTIS-83
12775	JBTIS-81	Hey Andrej - The SOA-P 6.0 Alpha for JBTIS 4.0.0 (Juno) is closed to all but show-stopper bugs. Can this wait for the SOA-P 6.0 Beta for JBTIS 4.1.0 (Kepler)? If not then you need to raise this to Anne-Louise, Len DiMaggio and the PMs. --paull	1	1	https://issues.jboss.org/browse/JBTIS-81
12776	JBTIS-81	Hi Paul, does it mean that there will be no final JBTIS update site for Juno? By final I mean a version which can be turned into a product for JBDS.	1	1	https://issues.jboss.org/browse/JBTIS-81
12777	JBTIS-73	Nothing has changed for ESB since 4.0, so do we really need to change anything?	1	1	https://issues.jboss.org/browse/JBTIS-73

#	Requirement	Questions	Answered	Successfully Answered	URL
12778	JBTIS-73	Looking at the jbosstools-integration-stack project, there are no "compositeArtifacts.xml" or "compositeContent.xml" files in the devstudio/site directory. What should I do there?	1	1	https://issues.jboss.org/browse/JBTIS-73
12779	JBTIS-71	Hey Bob - We're picking up the updated BPMN2 modeler feature target dependency from http://download.jboss.org/jbosstools/updates/requirements/bpmn2-modeler/0.2.6.201306192132_S20130619_kepler/ The BPEL component feature we're picking up from: http://download.jboss.org/jbosstools/builds/staging/jbosstools-4.0_stable_branch.component--bpel/all/repo/ for Kepler. Is that correct? If yes then you we can close this JIRA. Thanks	1	1	https://issues.jboss.org/browse/JBTIS-71
12780	JBTIS-70	Hey Paul, Were you going to create an individual JIRA for each component? The next tooling release for SwitchYard will be 1.0. Depending on the project schedule, I'm not sure whether or not the 1.0 release will target Kepler or will remain on Juno. I suspect other projects may already be compatible with Kepler. If that is the case, those projects should not have to provide a specific Kepler build (i.e. their current release may work just fine). Rob	1	1	https://issues.jboss.org/browse/JBTIS-70
12781	JBTIS-66	To answer your other questions, Nick: a) Um...good question. Is "I haven't gotten around to doing that" an acceptable answer? ;) Hint: I could use some help with this b) I'm doing maintenance releases on an as-needed basis; there's no fixed schedule. c) I don't believe there's anything in 0.2.5 that SY needs (Rob Cernich?). This latest build was required for Savara. d) Nope, the BPMN2 metamodel is fixed...it's perfect...no bugs here :)	1	1	https://issues.jboss.org/browse/JBTIS-66
12782	JBTIS-66	[~bbrodt] Anything left to be done here? [~pleacu] I see you don't have a version of BPMN2 Modeler in your Kepler-based target platform site [1] or your aggregate update site [2]. What's left to be done to make that happen? [1] https://github.com/jbosstools/jbosstools-integration-stack/blob/master/target-platform/integration-stack-base.target [2] http://download.jboss.org/jbosstools/updates/integration/kepler/integration-stack/aggregate/4.1.0/	1	1	https://issues.jboss.org/browse/JBTIS-66
12783	JBTIS-64	Closed/rejected PR - this change will wait for JBTIS-66. It will go into the kepler mainline capture for SOA-P 6.0 Beta/ JBDS 7.	1	1	https://issues.jboss.org/browse/JBTIS-64
12784	JBTIS-64	Hey Gary - Would you like for me to take the Savara 2.1.2.Final for both the SOA-P 6.0 Alpha 2 (Juno) and the SOA_P 6.0 Beta 1 (Kepler) from this update site? http://download.jboss.org/savara/releases/updates/2.1.2.Final/ The BPMN2 Modeler is at 0.2.6 now. Thanks! --paull	1	1	https://issues.jboss.org/browse/JBTIS-64
12785	JBTIS-50	OK, looks like we have the latest BPMN2 modeler (0.2.4) Let 'er rip!	1	1	https://issues.jboss.org/browse/JBTIS-50
12786	JBTIS-44	Hey Barry, Could you please reassign to the responsible party and ask Paul to update the component owner? Thanks, Rob	1	1	https://issues.jboss.org/browse/JBTIS-44
12787	JBTIS-32	commit comment says: "JBIDE-11101 move org.jboss.tools.runtime.esb.detector plugin and feature from jbosstools-runtime-soa project to jbosstools-esb project (master branch 1.5.200)", but runtime-soa jbosstools-4.0.x branch still have the same esb runtime detector plugin. The same state is in master branch for both affected components. Should we remove esb runtime detector plugin from jbosstools-runtime-soa?	1	1	https://issues.jboss.org/browse/JBTIS-32
12788	JBTIS-16	Which version of parent pom and target-platform do you use? I think we can integrate that change in Trunk (3.4.0.M1-SNAPSHOT) and branch (3.3.2-SNAPSHOT) target-platforms. Would it be fine for you? You'd just have to change parent version to 3.3.2-SNAPSHOT.	1	0	https://issues.jboss.org/browse/JBTIS-16
12789	JBTIS-16	Doug, This is a dependency specific to SOA plugins. I don't think we can easily add it to the indigo TP. Can you mirror the graphiti stuff and add it to this site you use http://download.jboss.org/jbosstools/updates/indigo/soa-tooling/SR2 ?	0	0	https://issues.jboss.org/browse/JBTIS-16
12790	JBTIS-16	Pulled a Graphiti 0.10 mirror... can that be used instead of 0.9? http://download.jboss.org/jbosstools/updates/requirements/graphiti/0.10.0.v20121219-0836/	0	0	https://issues.jboss.org/browse/JBTIS-16
12791	JBTIS-16	Which project in JBT core does require Graphiti? If none, we should remove Graphiti from JBT & JBDS target and make it part of the soa target extension suggested in JBIDE-11121.	0	0	https://issues.jboss.org/browse/JBTIS-16
12792	JBTIS-13	Paul, I'm project lead for Teiid Designer org.teiid.datatools.connectivity.feature_7.7.2.v20121212-2348-H174-Nightly-SOA.jar org.teiid.designer.feature_7.7.2.v20121212-2348-H174-Nightly-SOA.jar org.teiid.designer.runtime.feature_7.7.2.v20121212-2348-H174-Nightly-SOA.jar Our latest project release is 8.0 and update site is at: http://download.jboss.org/jbosstools/updates/development/juno/soa-tooling/teiddesigner/8.0.0.Final/ Barry From: "Paul Leacu" <pleacu@redhat.com> To: "Maciej Swiderski" <mswiders@redhat.com>, "Kris Verlaenen" <kverlaen@redhat.com>, "Gary Brown" <gbrown@redhat.com>, "William Collins" <bicollin@redhat.com>, "Brian Fitzpatrick" <bfitzpat@redhat.com>, "Lars Heinemann" <lheinema@redhat.com>, "Barry Lafond" <blafond@redhat.com>, "Mark Drilling" <mdrillin@redhat.com>, "Paul Richardson" <parichar@redhat.com>, "Ted Jones" <tejones@redhat.com>, "Dan Florian" <dflorian@redhat.com>, "John Verhaeg" <jverhaeg@redhat.com>, "Bob Brodt" <bbrodt@redhat.com>, "Rob Cernich" <rcernich@redhat.com> Sent: Monday, January 7, 2013 11:33:03 AM Subject: JBoss Integration Tools info Hey - I've been working with Rob C to produce an aggregate/composite update site project for the layered release of the JBoss tools integration stack. The good news is that we've made sufficient progress that we need to know the particulars in more detail. You're receiving this email because I believe that you can provide some portion of the specifics I need. The aggregated IS capture is 162 plugins - could you please review the attached features and plugins file and... 1. Identify which features/plugins you are responsible for (or which most closely 'belong' to you). 2. Specify the location of the most recent stable update site/repository. 3. Identify the full version of the lates release of your plugins. I'll update the target dependencies, composite and site files and see if we can produce an independent layered IS release! Thanks in advance! --paull	1	1	https://issues.jboss.org/browse/JBTIS-13
12793	JBTIS-11	Hi Paul Could you remove the pi4soa feature from site.xml. compositeArtifacts.xml - looks fine, although found it interesting that all other urls were download.jboss.org, except savara which was downloads.jboss.org. Both seem to map onto the same location though, so it is fine, or you might want to change it to be consistent with the others. The target platform looks ok - although worried about the BPMN2 feature still being based on a nightly build. I think we need to coordinate with Bob Brodt to make sure there is a stable release of this project (although Bob is not the lead he may be able to help), and also that the latest version of the BPMN2 modeler is released before the integration stack is released.	1	1	https://issues.jboss.org/browse/JBTIS-11

#	Requirement	Questions	Answered	Successfully Answered	URL
12794	JBTIS-11	Hi Paul The Savara features/plugins all start with org.savara and org.scribble. The stable release is here: http://downloads.jboss.org/savara/releases/updates/2.x (version 2.0.0.Final) This site also bundles the Eclipse BPMN2 and BPEL editors, as well as the WS-CDL tools (features/plugins beginning with org.pi4soa). The BPMN2 and BPEL editors have been available as separate releases for some time, so the ones in this update site can be ignored. However if you could also include the org.pi4soa feature/plugins. The dev release is here: http://downloads.jboss.org/savara/releases/updates/milestones/current (version 2.1.0.M2 currently) In M3 the number of plugins will be reduced and the org.scribble plugins will no longer be distributed as plugins, but embedded in a savara plugin. Regards Gary ----- Original Message ----- > > Hey - > I've been working with Rob C to produce an aggregate/composite > update site project for the layered release of the JBoss > tools integration stack. The good news is that we've made > sufficient progress that we need to know the particulars in > more detail. You're receiving this email because I believe > that you can provide some portion of the specifics I need. > The aggregated IS capture is 162 plugins - could you please > review the attached features and plugins file and... > > 1. Identify which features/plugins you are responsible for (or > which most closely 'belong' to you). > 2. Specify the location of the most recent stable update > site/repository. > 3. Identify the full version of the lates release of your > plugins. > > I'll update the target dependencies, composite and site files and see > if we can produce an independent layered IS release! > > Thanks in advance! > > --paull	1	1	https://issues.jboss.org/browse/JBTIS-11
12795	JBTIS-10	Hi Paul, Here is the info requested. Notice there is one more plugin listed than you had. Also included test features which were not in your list (but probably not needed here). 1. Features/Plugins I'm resonsible for: org.jboss.tools.modeshape.jcr.feature org.jboss.tools.modeshape.jcr.test.feature org.jboss.tools.modeshape.rest.feature org.jboss.tools.modeshape.rest.test.feature org.jboss.tools.modeshape.client org.jboss.tools.modeshape.jcr org.jboss.tools.modeshape.jcr.doc.user org.jboss.tools.modeshape.jcr.ui org.jboss.tools.modeshape.rest org.jboss.tools.modeshape.rest.doc.user org.jboss.tools.modeshape.ui 2. Stable/Development update-site (version 3.0.0): http://download.jboss.org/jbosstools/updates/development/juno/soa-tooling/modeshape/ 3. Staging update-site (version 3.1.0): http://download.jboss.org/jbosstools/builds/staging/modeshape-tools-continuous/all/repo/ Let me know if you need anything else. Thanks, Dan ----- Original Message ----- > > Hey - > I've been working with Rob C to produce an aggregate/composite > update site project for the layered release of the JBoss > tools integration stack. The good news is that we've made > sufficient progress that we need to know the particulars in > more detail. You're receiving this email because I believe > that you can provide some portion of the specifics I need. > The aggregated IS capture is 162 plugins - could you please > review the attached features and plugins file and... > > 1. Identify which features/plugins you are responsible for (or > which most closely 'belong' to you). > 2. Specify the location of the most recent stable update > site/repository. > 3. Identify the full version of the lates release of your > plugins. > > I'll update the target dependencies, composite and site files and see > if we can produce an independent layered IS release! > > Thanks in advance! > > --paull	1	1	https://issues.jboss.org/browse/JBTIS-10
12796	JBTIS-9	Hi Koen - Could you please verify the jBPM configuration in the JBoss Integration Stack and make sure the upstream component target dependencies are correct? Please note the new location of the files to be reviewed: https://github.com/jbosstools/jbosstools-integration-stack/blob/master/aggregate-site/site.xml https://github.com/jbosstools/jbosstools-integration-stack/blob/master/releases/compositeArtifacts.xml https://github.com/jbosstools/jbosstools-integration-stack/blob/master/releases/compositeContent.xml https://github.com/jbosstools/jbosstools-integration-stack/blob/master/target-platform/base.target We need your input ASAP. Thanks!! --paull	1	1	https://issues.jboss.org/browse/JBTIS-9
12797	JBTIS-9	Hi Paul, I have never dealt with the above artifacts. What exactly do you want me to do with them? Cheers, Koen	1	1	https://issues.jboss.org/browse/JBTIS-9
12798	JBTIS-8	Paul, 1. Following features and plugins all all included in one update site *** Features org.guvnor.tools.feature_5.3.5.Final.jar org.drools.eclipse.feature_5.3.5.Final.jar org.jbpm.eclipse.feature_5.3.5.Final.jar org.jbpm.eclipse.task.feature_5.3.5.Final.jar *** Plugins org.guvnor.eclipse.webdav_5.3.5.Final.jar org.guvnor.tools_5.3.5.Final.jar org.drools.eclipse_5.3.5.Final.jar org.jbpm.eclipse_5.3.5.Final.jar org.jbpm.eclipse.task_5.3.5.Final.jar 2. Latest release droolsjbpm 5.5.0.Final https://repository.jboss.org/nexus/content/groups/public-jboss/org/drools/org.drools.updatesite/5.5.0.Final/org.drools.updatesite-5.5.0.Final-assembly.zip 3. Snapshot releases for 6.0 https://repository.jboss.org/nexus/content/groups/public-jboss/org/drools/org.drools.updatesite/6.0.0-SNAPSHOT/ There are some other plugins related to jBPM and Drools as well: The following features / plugins are related to jBPM3. Has the decision been made that they will be included in the aggregated update site? As I think it might be that we will no longer include these for JBDS6. Features org.jboss.tools.jbpm3.feature_4.5.0.Beta2-v20121106-1148-B1927.jar org.jboss.tools.jbpm.common.feature_4.5.0.Beta2-v20121106-1148-B1927.jar org.jboss.tools.jbpm.convert.feature_4.5.0.Beta2-v20121106-1148-B1927.jar org.jbpm.gd.jpdl.feature_4.5.0.Beta2-v20121106-1148-B1927.jar Plugins org.jbpm.gd.jpdl_4.5.0.Beta2-v20121106-1148-B1927.jar org.jboss.tools.jbpm.common_4.5.0.Beta2-v20121106-1148-B1927.jar org.jboss.tools.jbpm.convert_4.5.0.Beta2-v20121106-1148-B1927.jar The following features / plugins are extensions to the first set of droolsjbpm tools, responsible for detecting a drools or jbpm runtime for specific platforms. They were created / maintained by the JBT / JBDS teams however. Features org.jboss.tools.runtime.drools.detector.feature_1.4.0.Beta2-v20121221-1759-B60.jar org.jboss.tools.runtime.jbpm.detector.feature_1.4.0.Beta2-v20121221-1759-B60.jar Plugins org.jboss.tools.runtime.drools.detector_1.4.0.Beta2-v20121221-1759-B60.jar org.jboss.tools.runtime.jbpm.detector_1.4.0.Beta2-v20121221-1759-B60.jar Kris On 01/07/2013 06:33 PM, Paul Leacu wrote: > Hey - > I've been working with Rob C to produce an aggregate/composite update site project for the layered release of the JBoss tools integration stack. The good news is that we've made sufficient progress that we need to know the particulars in more detail. You're receiving this email because I believe that you can provide some portion of the specifics I need. The aggregated IS capture is 162 plugins - could you please review the attached features and plugins file and... > > 1. Identify which features/plugins you are responsible for (or which most closely 'belong' to you). > 2. Specify the location of the most recent stable update site/repository. > 3. Identify the full version of the lates release of your plugins. > > I'll update the target dependencies, composite and site files and see if we can produce an independent layered IS release! > > Thanks in advance! > > --paull	1	1	https://issues.jboss.org/browse/JBTIS-8
12799	JBTIS-2	Is this a done issue?	0	0	https://issues.jboss.org/browse/JBTIS-2
12800	JBTIS-1	Hey Nick, Mickael, We're working on putting a process together for handling things like this for SOA tools. That said, this looks OK with the following exceptions: * The guvnor/jbpm entries should be removed * The bpmn2 modeler need to be from ../site/0.1.0 I have not verified the specific versions of some of the other dependencies (e.g. OCL, bpel). Thanks for putting this together. Rob	1	1	https://issues.jboss.org/browse/JBTIS-1

#	Requirement	Questions	Answered	Successfully Answered	URL
12801	JBTIS-1	[~elvisisking]: how do you know that it's using the suggested soa target platform? That's why I'd like you to try step 1 to 4, so we are sure you are actually testing the new SOA target-platform.	1	1	https://issues.jboss.org/browse/JBTIS-1
12802	JBTIS-1	No problem, Dan. Your use case just shows that I should find a more direct way for you to test. I've worked on something simpler, can you please *ALL* give it a try asap? Just follow *exactly* those instructions: If you're using JBT parent pom # change parent pom version to 4.1.0.JBIDE-11121-SNAPSHOT # run "mvn clean verify -P soa -Dno.jbosstools.site -Dtycho.localArtifacts=ignore" If you're *not* using JBT parent: # Remove/comment from your pom all <repositories> and target-platform-definition-configuration # put in your pom.xml: {code:xml} <plugin> <groupId>org.eclipse.tycho</groupId> <artifactId>target-platform-configuration</artifactId> <version>${tychoVersion}</version> <configuration> <target> <artifact> <groupId>org.jboss.tools.target-platforms.soatools</groupId> <artifactId>multiple</artifactId> <version>4.2.1.a-SNAPSHOT</version> <classifier>multiple</classifier> </artifact> </target> </configuration> </plugin> {code} # Run a "mvn clean verify" Please post your build log there.	1	1	https://issues.jboss.org/browse/JBTIS-1
12803	JBTIS-1	Hey y'all...I know you THINK I'm supposed to know what all of this means, but really I don't have much of a clue. I know there's this thing called a ".target" file which I can create using an eclipse wizard, but beyond that...all of the words here parse like english (well, most of them) but they don't really make much sense to me. What exactly am I supposed to do here? I tried following the steps 1-5 described by Mickael above, but haven't gotten very far. What I see is [ERROR] Plugin org.jboss.tools.tycho-plugins:target-platform-utils:0.0.1-SNAPSHOT or one of its dependencies could not be resolved: Failed to read artifact descriptor for org.jboss.tools.tycho-plugins:target-platform-utils:jar:0.0.1-SNAPSHOT: Could not find artifact org.jboss.tools.tycho-plugins:tar get-platform-utils:pom:0.0.1-SNAPSHOT in tycho-snapshots (https://oss.sonatype.org/content/groups/public/) -> [Help 1] which makes not much sense to me HELP 1!	1	1	https://issues.jboss.org/browse/JBTIS-1
12804	JBTIS-1	Based on what I see here [0], your Juno-based target platform requires: Graphiti 0.9.1 Guava 10.0.1 (from Orbit R20120526062928) OCL 4.0.1 EMF Compare 1.3.2 -> requires EMF 2.8 EMF Query OCL 1.6 -> requires EMF 2.8, OCL 4.0, UML2 4.0 All of which is included in here [1]. As well as: BPMN2 0.7.0 and BPMN2 Modeler 0.1.0 [2] BPEL 1.0.2 [3] and of course... JBT 4.0 JBT Tests 4.0 --- Questions: 1. EMF Compare 1.3 is only compatible w/ Eclipse 3.8, not 4.2 [4]. What depends on this, and can version 2.0 be used instead, which works w/ Eclipse 3.8 AND with 4.2? Or 2. UML2 4.0 is not in your TP, though if I recall correctly (and they haven't removed the requirement) Teiid Designer requires UML2. Should it be added to your TP? I suspect you're resolving that dependency because it's already included in the JBT target platform [5] ? If that's the case then I believe it should be removed from ours and added to yours instead. 3. Graphiti 0.9 is in both your TP and the JBT one [5], yet I don't think anything in JBT requires it. It was most likely added to support something in the SOA-P/BRMS tooling space before you started having your own TP. It too should probably be removed from ours in favour of it being in yours - no need to be in both. Observation: 1. I've added a mirror on download.jboss.org for the only requirement I saw above that's not already there - the older version of the bpmn2 stuff, now mirrored here [2]. As such, you should be able to have all the URLs in your TP point to download.jboss.org URLs instead of the 8 eclipse.org ones, and get better performance when building the target platform site. You can even simplify the TP since of those 8 URLs, 6 can be collapsed down into a single one [1], leaving only BPEL and BPMN2 as separate entries, followed by the JBoss Tools stuff. -- [0] https://jenkins.mw.lab.eng.bos.redhat.com/hudson/job/pleacu-jbosstools-intstack-aggregate/ws/sources/target-platform/integration-tools.target/*view*/ [1] http://download.jboss.org/jbosstools/updates/requirements/juno/201209280900-R/ [2] http://download.jboss.org/jbosstools/updates/requirements/bpmn2-modeler/201211132058/ *NEW* [3] http://download.jboss.org/jbosstools/updates/requirements/bpel/1.0.2.v20121121-1634-CI/ [4] http://wiki.eclipse.org/EMF_Compare#Compatibility [5] http://download.jboss.org/jbosstools/updates/juno/SR1a/REPO/	0	0	https://issues.jboss.org/browse/JBTIS-1
12805	LOCUS-46	not following why you say JBDS has jboss-tools-locus as name and JBT has something different. JBDS naming is for *product artifacts*, thus not really relevant here. and the jbt consistent name would be jbosstools-locus-1.3.0-<version|SNAPSHOT>-updatesite.zip afaik. What other artifacts are you seeing the jbosstools-1.3.0-SNAPSHOT-updatesite-locus.zip format ?	1	1	https://issues.jboss.org/browse/LOCUS-46
12806	LOCUS-45	i'm lost - we did not have locus publish to dl.jboss.org before did we ? why do we need to have these urls configured and linked now ?	1	1	https://issues.jboss.org/browse/LOCUS-45
12807	LOCUS-42	doesn't this require addiitonal dependencies in locus or are they already provided from somewhere else ?	1	1	https://issues.jboss.org/browse/LOCUS-42
12808	LOCUS-42	Hey [~nickboldt], [~mickael_istria] - I know you guys are finishing up on the Beta1 9.0 TP. Could you assign this so it's not lost - it's needed so I can spin a 9.0.x Beta1 JBTIS TP for use with Data Trans in Fuse Tooling. There's only the one jar. Thkx!	1	1	https://issues.jboss.org/browse/LOCUS-42
12809	LOCUS-42	[~pleacu] Which jar(s) do you need us to add from Orbit into the JBT/JBDS target platforms for 4.50.0.Beta1 ? You said there were 4 6 days ago, and today you say there's only one jar. Can you submit a PR to update our TP so I know exactly what you need/want? https://github.com/jbosstools/jbosstools-devdoc/blob/master/building/target_platforms/target_platforms_updates.adoc https://github.com/jbosstools/jbosstools-target-platforms/pulls Or, if this is something that ONLY IS needs, why not just add the jar(s) to your JBTIS/JBDSIS TP? Why does it need to be upstream in JBT/JBDS TP?	1	1	https://issues.jboss.org/browse/LOCUS-42
12810	LOCUS-42	I wasn't thinking of adding it to the JBT TP, I was thinking it belonged in Locus (where I could pick it up from the JBTIS TP). I can find org.apache.commons.lang3_3.1.0.v201403281430.jar and the three jackson jars. It looks at this point like we need the jsonschema2pojo-core jar and two jaxb libraries (ref: https://issues.jboss.org/browse/JBTIS-438) - 3 total Project page/sources: https://jaxb.java.net/ Version: 2.2.7 (no hard requirement on this specific version that I know of) License and owner: CDDL v1.1 and GPL v2 Original p2 repo: unknown JBoss mirror: unknown Include Sources: Yes Affected projects: Fuse IDE Include in JBDS: unknown Type of dependency: unknown non-p2 repository: https://repo1.maven.org/maven2/com/sun/xml/bind/jaxb-xjc/2.2.7/jaxb-xjc-2.2.7.pom I thought the way to add the jars to JBTIS/JBDSIS TP was through Locus - how can I add them directly into my TP??	1	1	https://issues.jboss.org/browse/LOCUS-42
12811	LOCUS-42	Sorry, misunderstood. Thought you meant all the things you needed were already in Orbit (and therefore could be directly added to your TP). So to confirm, you want three new jars in Locus 1.3.0: * jsonschema2pojo-core jar, * jaxb-xjc jar and * jaxb-impl jar I don't see anywhere else in this JIRA or in JBTIS-438 where you've explicitly asked for these three, so I wanted to confirm I've successfully read the tea leaves of this cryptic request. (It'd be cool if you could update your issues' descriptions/titles when they evolve to include more work . :D) Is that correct?	1	1	https://issues.jboss.org/browse/LOCUS-42
12812	LOCUS-42	woow now ....CDDL and GPL v2 ? these are *not* kosher to use. is that truly their license ?!	1	1	https://issues.jboss.org/browse/LOCUS-42
12813	LOCUS-42	actually CDDL is possibly okey, I need to check - we've had issues with it in past. And is jaxb *absolutely* required ? What is it used for ? is it just for annotations or actually needed by jsonschema2pojo ?	1	1	https://issues.jboss.org/browse/LOCUS-42

#	Requirement	Questions	Answered	Successfully Answered	URL
12814	LOCUS-42	Hey [~kcbabo] - hoping that we can revisit this now that Fuse Tooling has shipped. Can you respond to Max's question above concerning JAXB? Also it appears that we're good to go with the jsonschema2pojo-core jar - [~mickael_istria] can we look into getting that into locus? Thkx	1	1	https://issues.jboss.org/browse/LOCUS-42
12815	LOCUS-42	[~mickael_istria] - what's the version branch for the next mars locus release? Is it still 1.3.0.Final? 1.3.1.?? Thkx	1	1	https://issues.jboss.org/browse/LOCUS-42
12816	LOCUS-42	Hey [~kcbabo] - I'd like to use version 0.4.13 of jsonschema2pojo - looks like it came out last week. If it causes trouble we can back off of it. WDYT?	1	1	https://issues.jboss.org/browse/LOCUS-42
12817	LOCUS-42	[~maxandersen] [~kcbabo] [~nickboldt] [~mickael_istria] This is a cleaned up version of the #42 PR. It adds jsonschema2pojo 0.4.13 and modifies the internal locus target platform to add referenced jars from eclipse orbit: com.fasterxml.jackson.core.jackson-annotations com.fasterxml.jackson.core.jackson-core com.fasterxml.jackson.core.jackson-databind javax.validation org.junit https://github.com/jbosstools/jbosstools-locus/pull/47 *This is less than the list of dependencies listed in the source dependencies list:* https://repo1.maven.org/maven2/org/jsonschema2pojo/jsonschema2pojo-core/0.4.13/jsonschema2pojo-core-0.4.13.pom I built this into a local IS TP: https://github.com/jbosstools/jbosstools-integration-stack/pull/386 I also modified the Fuse Data Trans editor plugin to ref the updated IS TP: https://github.com/pleacu/fuseide/tree/LOCUS-42 Everything compiles successfully. If this is agreeable I'd suggest we build a -SNAPSHOT locus and a -SNAPSHOT JBTIS TP and see someone from Fuse Data Trans wants to complete my branch changes. WDYallT?	1	1	https://issues.jboss.org/browse/LOCUS-42
12818	LOCUS-40	Would https://code.google.com/p/powermock-osgi/ fit your needs?	1	1	https://issues.jboss.org/browse/LOCUS-40
12819	LOCUS-40	[~jcantrill] Is this something you would see being installed OOTB in JBoss Tools? Also in JBDS? Or something that we could add to Central as an optional add-on for both JBT/JBDS users?	1	1	https://issues.jboss.org/browse/LOCUS-40
12820	LOCUS-40	The PR only includes one feature - for mockito-junit. What about the other three mocks? {code} org.powermock.easymock-junit.feature.feature.group 1.5.4.1 org.powermock.easymock-testng.feature.feature.group 1.5.4.1 org.powermock.mockito-testng.feature.feature.group 1.5.4.1 {code} Would anyone in JBTIS or JBT benefit from those? WDYT, [~adietish], [~xcoulon], [~phantomjinx] (ref: JBIDE-12971, JBIDE-14315)	1	1	https://issues.jboss.org/browse/LOCUS-40
12821	LOCUS-40	[~nickboldt] I think it depends on what is currently be used in the code base. Do we have any unit tests that utilize 'testng'? Do we have any unit tests that utilize easymock? It is my opinion for uniformity across the code base we should promote one or the other and limit the feature group. The alternative is to add them since it doesn't 'cost' us anything right?	1	1	https://issues.jboss.org/browse/LOCUS-40
12822	LOCUS-40	[~phantomjinx] The use with files is only an example. I am also using the approach where I create tmp-files and kick them off afterwards. There are plenty of cases where you cannot rely on interfaces, you need more power than proxies. Mockito has such means but afaik powermock goes far beyong what mockito has to offer. So far mockito was fully sufficient to me. Up to [~jcantrill] to weigh the benefits/costs. Can you elaborate? [~nickboldt] I dont think that there's any need for the additional testng (we use junit) and easymock (comparable to mockito, no benefit) features.	1	0	https://issues.jboss.org/browse/LOCUS-40
12823	LOCUS-40	btw. why is multiple junit 4.x versions needed? shouldn't we be okey using one specific 4?	0	0	https://issues.jboss.org/browse/LOCUS-40
12824	LOCUS-40	is that mockito-core/javaassist not supposed to be the one from our locus ? it seem to still be the one from powermock?	0	0	https://issues.jboss.org/browse/LOCUS-40
12825	LOCUS-39	[~mickael_istria] I believe it is actually possible to "aggregate" any amount of maven call in one using maven-invoker plugin. Basically it lets you to invoke maven from maven, but from user perspective it looks like single maven call. That would first call EBR builds for required bundles, install them into local repo and then call maven again to build p2repository. I am also curious why we don't use inclusion of actual jar inside bundle and then expose it through OSGi, for examle like that https://github.com/jbosstools/jbosstools-locus/compare/jbosstools:master...dgolovin:locus-dmg-question?expand=1. Are there any class loading related problems?	1	0	https://issues.jboss.org/browse/LOCUS-39
12826	LOCUS-39	{quote}I am also curious why we don't use inclusion of actual jar inside bundle and then expose it through OSGi, for examle like that https://github.com/jbosstools/jbosstools-locus/compare/jbosstools:master...dgolovin:locus-dmg-question?expand=1. Are there any class loading related problems?{quote} This has an extra cost on Equinox performances. It's the main reason why Orbit does this as well. Also, nested jars can confuse PDE which doesn't add them correctly in the classpath and so on. So overall, it's not as good aa repackaging in plain OSGi.	1	0	https://issues.jboss.org/browse/LOCUS-39
12827	LOCUS-39	https://wiki.eclipse.org/Tycho/Packaging_Types#p2-installable-unit Would this feature help?	1	0	https://issues.jboss.org/browse/LOCUS-39
12828	LOCUS-36	This JIRA lacks information. What does this do, why do we need it, and do we want it in the upcoming 1.0.0.Final release, or in master branch for 1.0.1 / 1.1.0?	1	1	https://issues.jboss.org/browse/LOCUS-36
12829	LOCUS-36	It adds Eclipse-SourceReferences to bundles. Cf * http://wiki.eclipse.org/PDE/UI/SourceReferences * http://wiki.eclipse.org/Tycho/FAQ#How_to_generate_Eclipse-SourceReferences_MANIFEST_header.3F Another benefit of this Eclipse-SourceReferences is that it puts in the MANIFEST.MF a reference to the commit/tag that was used to build the bundle. As for the version, I'd be tempted to say 1.0.0.Final, but is it an extra-effort compared to only 1.0.1. As it's a non-blocking improvment, I'd vote for the easiest solution.	0	0	https://issues.jboss.org/browse/LOCUS-36
12830	LOCUS-35	[~fbricon], [~rob.stryker] and [~adietish] you are the users of the dmr library. Can you test that the bundles produced by the #39 PR works with your usecases ? Should be done by removing your embedded dmr or dependency to as.dmr and see if things works.	1	1	https://issues.jboss.org/browse/LOCUS-35
12831	LOCUS-35	So do you want dmr 1.3 or is it ok to merge the PR for 1.2.2 ?	1	1	https://issues.jboss.org/browse/LOCUS-35
12832	LOCUS-28	Here are several possible approaches: * Embed jars in bundle * Explode jars and put the .class files in bundle * Get source and recompile [~maxandersen] Which one do you want?	1	1	https://issues.jboss.org/browse/LOCUS-28
12833	LOCUS-28	As your PR has already been applied, I checked it out, built it, and discovered that there are no class files in the fest jars. So here's a fix for that and a few others cleanup items: https://github.com/jbosstools/jbosstools-locus/pull/2 Unfortunately, it appears I don't have push rights. [~maxandersen], [~mickael_istria] or [~phantomjinx], can you review and apply the PR? [~maxandersen] is there a reason I don't have commit rights? :(	1	1	https://issues.jboss.org/browse/LOCUS-28
12834	LOCUS-25	Cool - not sure why I didn't spot that big blue attachement yesterday :) WDYT about the issue of a possible race condition on this static variable when running concurrent tests ? Is there no way in Mockito to refer to the "current configuration" to set it per mockito 'instance' instead of globally ?	1	1	https://issues.jboss.org/browse/LOCUS-25
12835	LOCUS-25	okey so Mockito keeps it state per-thread ? thats comforting at least.	0	0	https://issues.jboss.org/browse/LOCUS-25
12836	LOCUS-24	anyone know how the eclipse orbit is setup ? do they have a one big source repo with builds for these we could look at ?	1	1	https://issues.jboss.org/browse/LOCUS-24
12837	LOCUS-24	when they update do they do a full rebuild of orbit for this ? i assume not, which would be a reason to have various versions/tags of it.	1	1	https://issues.jboss.org/browse/LOCUS-24

#	Requirement	Questions	Answered	Successfully Answered	URL
12838	LOCUS-24	Nice but why do we need the source for *all* of these ? wouldn't a reference to their maven central GAV be enough for most cases ?	1	1	https://issues.jboss.org/browse/LOCUS-24
12839	LOCUS-24	Hey, if we migrate this to a new project, https://github.com/jbosstools/jbosstools-locus, and fix the namespaces such that s/org.jbosstools/org.jboss.tools/g, do we need all the 1:1 feature wrappers? * We can add plugins to our TP (features are not needed). * We can have our plugins depend on locus plugins, and be installed to an update site, then to a user's Eclipse. * We can have our features include or depend on locus plugins, and be installed to an update site, then to a user's Eclipse. * We only need features IFF there's a use case for a user to be installing Mockito or another Locus component *DIRECTLY*, via p2 Install, or JBoss Central, and you need the *additional metadata that only a feature provides*. So, is there a reason to have org.jboss.tools.locus.*.feature features? As of Tycho 0.16 and Kepler, we can publish all the individual plugins to an update site w/o the need for feature wrappers. Then those plugins can be installed via p2 Install. (Not sure about Central however.) Ref: https://bugs.eclipse.org/bugs/show_bug.cgi?id=378338#c21	1	1	https://issues.jboss.org/browse/LOCUS-24
12840	LOCUS-24	Who should the provider be for these? Sourceforge? Google Code? or JBoss by Red Hat? I would think us, but IANAL.	1	1	https://issues.jboss.org/browse/LOCUS-24
12841	LOCUS-24	+1, this looks good. About license, there is not such license attribute in plugin itself, so I guess p2 installer doesn't show anything for bundles. The remaining question is about qualifiers. Maybe it would make sense to use the jgit-timestamp provider for those since we don't want them to contain some BUILD_STREAM and other JBT specific labels. WDYT?	1	1	https://issues.jboss.org/browse/LOCUS-24
12842	LOCUS-24	on first look this looks very awesome! I want to try execute this and see the actual result. I can't make that today unfortunately. Is the screenshots missing ? the links are broken here and I see no attachements. with respect to license - what does Orbit do ? and with respect to installing - I hope all in locus are stuff that is included by others, not something to install manually.	1	1	https://issues.jboss.org/browse/LOCUS-24
12843	LOCUS-24	the biggest remaining issue is the mockito stuff which breaks the API of that bundle. can we exclude that for M7 without any bad sideffects ? [~phantomjinx] did you ever hear back from mockito guys - seem both your and my recent comment went completely unanswered :(	1	1	https://issues.jboss.org/browse/LOCUS-24
12844	LOCUS-24	I've finally gotten around to publish my README.md PR at https://github.com/nickboldt/jbosstools-locus/pull/1 This is/was very much based on what is found on Orbit docs about how to do these bundles. In doing so i've found some discrepencies compared to Orbit we'll need to react to. 1) we are adding jar's inside a bundle instead of using the actual real jar with updated manifest.mf - and advantage/disadvantage to this ? 2) fest contains multiple jar's incl. javadoc jar - that seems wrong. Why is it not one bundle per jar ? 3) Required execution environment is set to java 6 - is that true for all ?	1	1	https://issues.jboss.org/browse/LOCUS-24
12845	LOCUS-24	I'm looking at locus this morning and noticed when building both http://download.jboss.org/jbosstools/updates/juno and http://download.jboss.org/jbosstools/updates/kepler and both their associated TP sites gets added to the build. Why is that ? Everything in locus is supposed to have *zero* dependencies on these so why are they included ?	1	1	https://issues.jboss.org/browse/LOCUS-24
12846	LOCUS-23	[~mickael_istria] Have tested the PR. * Built it in jenkins under the jbosstools-locus.site_branch-testing project (build #2) - successful * Build published to http://download.jboss.org/jbosstools/updates/nightly/locus/provisional - successful * Re-generated Teiid Designer's target locally against this new provisional build and validated - successful * Executed Designer build and tests - successful On that basic happy for it to be merged. Assuming it is merged, what is the turnaround going to be for seeing this in the integration-stack?	1	1	https://issues.jboss.org/browse/LOCUS-23
12847	LOCUS-23	[~maxandersen] re. JBTIS. No problem. I can raise a JIRA on Paul to take the new version of locus. When would that new version of locus be available to take? re. provisional. I set up this link some time ago so that devs would not have to generate a local TP for Teiid Designer and that the 'new' saxon version would be available to builds in jenkins. It seemed a logical addendum to the publishing script in jenkins to set the TARGET_FOLDER variable, ie. -PARENT_FOLDER locus -TARGET_FOLDER provisional. The objective was to get new versions of Teiid Designer building with the new target-platform (new saxon version) until at such point the JBTIS version can be updated in the former's poms use the new version containing the new locus containing the new version of saxon. Thus, the current Teiid Designer target-platform pom looks like this https://github.com/Teiid-Designer/teiid-designer/blob/master/target-platform/pom.xml - note the commented-out sections and the attaching of the [teiid-designer-extras-4.2.0.Final.target|https://github.com/Teiid-Designer/teiid-designer/blob/master/target-platform/teiid-designer-extras-4.2.0.Final.target] If this objective can be achieved in a better way, can you point it out for me. Apologies if this was the incorrect protocol.	1	0	https://issues.jboss.org/browse/LOCUS-23
12848	LOCUS-21	Can this wait until Kepler M6 is out, or is it more urgent?	1	1	https://issues.jboss.org/browse/LOCUS-21
12849	LOCUS-21	[~mgagyi] Matej, is this something that you could use in your test? FestAssert adds nice syntactic sugar to assertions, my guess is that this is something that you could nicely use when doing tests? If yes, then what about trying this out and report to [~nickboldt] in this jira?	0	0	https://issues.jboss.org/browse/LOCUS-21
12850	LOCUS-21	We don't want to add things that people could use in TP, we really to add things that people want or need to use, so that it's sure they will use them. So [~mgagyi], do you want or need festassert? Will you use it if we introduce it in TP?	1	0	https://issues.jboss.org/browse/LOCUS-21
12851	LOCUS-21	Are you aware of anyone who could use assertJ in an OSGi world? Does it have OSGi bundles? A p2 repo? If yes, then we can probably include it in next (4.50.0.Alpha2-SNAPSHOT) TP. Please make another Jira request dedicated to AssertJ and close this one as out-of-date.	1	0	https://issues.jboss.org/browse/LOCUS-21

#	Requirement	Questions	Answered	Successfully Answered	URL
12852	LOCUS-17	I would think we want a "latest best" URL too, which would initially point to 1.0.0, then later 1.0.1, etc. Convention dictates that everything under /updates/ is then qualified by a TYPE or release: http://download.jboss.org/jbosstools/updates/stable/ http://download.jboss.org/jbosstools/updates/development/ http://download.jboss.org/jbosstools/updates/integration/ http://download.jboss.org/jbosstools/updates/nightly/ And, to avoid having lots of locus-* folders in the root of /updates/, I'd prefer to see locus/<version>/ than locus-<version>/. Thus, I would suggest this: http://download.jboss.org/jbosstools/updates/stable/locus/ --> points to latest stable release -> http://download.jboss.org/jbosstools/updates/stable/kepler/core/locus/ This pattern is the same for JBoss Tools: http://download.jboss.org/jbosstools/updates/stable/ --> points to latest stable release -> http://download.jboss.org/jbosstools/updates/stable/juno/ (currently 4.0.1) In fact, we already (sort of) have this, but only for /integration/ since we haven't had a /development/ or /stable/ release yet: http://download.jboss.org/jbosstools/updates/integration/locus/ (manually updated to point at latest release) -> http://download.jboss.org/jbosstools/updates/integration/kepler/core/locus/ I think we need to decide, like with Orbit (and its R-build releases which are not always backward compatible w/ previous releases), if 1.0.0 is a "Juno+" release, or since it's used for projects which are based on Kepler, whether it should be considered a "Kepler+" release. And then going forward, if we decide a baseline eclipse target based on the stuff that uses it, or what's in Locus & its TP. Personally, I think it makes sense for Locus 1.0.0 to be under /stable/kepler/core/locus/1.0.0.Final/, because: * it's stable * it's used by Kepler-based projects * core is there because it lets me use my existing promote.sh script w/o having to change it (#laziness #reuse) * it's locus (/locus/ is a composite of all the contained releases therein) Then, if we do a Locus release that depends on Luna, it would be /stable/luna/core/locus/1.1.0.Final/, but all subsequent releases depending on Kepler would be 1.0.x. Make sense?	1	1	https://issues.jboss.org/browse/LOCUS-17
12853	LOCUS-15	what does orbit do?	0	0	https://issues.jboss.org/browse/LOCUS-15
12854	LOCUS-13	QE has consistently said they LIKE having BUILD_NUMBER used in qualifiers. Why would we remove them? (qualifier is always changing because of timestamp too)	1	1	https://issues.jboss.org/browse/LOCUS-13
12855	LOCUS-13	Would QE be happy enough with Git commit id? [~ldimaggio] [~mmalina] any opinion on that?	1	1	https://issues.jboss.org/browse/LOCUS-13
12856	LOCUS-13	Yeah, I'm with Nick on this one. I think we're quite happy with the build number - it gives us a distinct ID that we can rely on. I get the point that you cannot reproduce a qualifier twice, but as Nick says, you couldn't have the date either. Having the corresponding git hash somewhere (is there a place for it? manifest? I don't know) wouldn't be bad, but I don't think build qualifier (=part of filename) is the right place. My $0.02 ;)	0	0	https://issues.jboss.org/browse/LOCUS-13
12857	LOCUS-13	I did not see this jira - we discussed early on for Locus that since this site will be used across multiple versions of Eclipse and from different installations it would be important we did not unnecessarily bump the version of the libraries. That said, is this the actual issue [~pleacu] bumped into ?	1	1	https://issues.jboss.org/browse/LOCUS-13
12858	LOCUS-13	I'm not sure how to interpret your answer. Is it ok to remove the SNAPSHOT/Final distinction in build?	0	0	https://issues.jboss.org/browse/LOCUS-13
12859	LOCUS-13	Removing BUILD_ALIAS won't remove the B## suffix. It will only remove the .Final part. Do you ALSO want to remove the BUILD_NUMBER? {quote}The action here would be simply be to configure tycho-packaging-plugin to remove reference to BUILD_NUMBER.{quote} Oh, I see you do. Or did. But don't now? Plugins in http://download.jboss.org/jbosstools/mars/snapshots/builds/jbosstools-locus.site_master/latest/all/repo/plugins/ show names like org.jboss.tools.locus.dmr_1.3.0.Final-v20150430-1349-B77.jar But when we rebuild, even if we remove the .Final bit, we'll get org.jboss.tools.locus.dmr_1.3.0.v20150430-1349-B78.jar And then each build after that, assuming no code change to the plugin (so no timestamp change)... org.jboss.tools.locus.dmr_1.3.0.v20150430-1349-B79.jar org.jboss.tools.locus.dmr_1.3.0.v20150430-1349-B80.jar ...	1	1	https://issues.jboss.org/browse/LOCUS-13
12860	JBTM-2905	This class was added to demonstrate how to support new log record types in the tooling. The bug is that it was included in the default list of types. What is your plan to fix this, are you going to remove it from the default list?	1	1	https://issues.jboss.org/browse/JBTM-2905
12861	JBTM-2905	[~mmusgrov] ok, that could be my wrong in understanding the issue here. Without the header reader I got troubles to get working ExpiredAssumedCompleteScanner. But I could take it from a wrong end. In that case, this change would be superfluous and the fix should go in opposite way - removing from the default list and leaving header reader in the test package. Can you please elaborate a bit on what is the correct way of fixing it? Thanks in advance!	1	1	https://issues.jboss.org/browse/JBTM-2905
12862	JBTM-2835	What does the error look like with your patch? What is the effect of the NPE?	1	0	https://issues.jboss.org/browse/JBTM-2835
12863	JBTM-2822	Hi [~tomjenkinson], during the check ups JBEAP issues I got to this issue and when applying the same test as I have for JBTM-2636 I found out that there is a branch in code where deffered exception is not added - that's when heuristic happens. Is it intentional? This kind of update fixes the issue that I refer to: https://github.com/ochaloup/narayana/commit/cb2ab4252a6c2e1106751d28103daefea42e86f6	1	1	https://issues.jboss.org/browse/JBTM-2822
12864	JBTM-2822	Sorry Ondra - I don't quite follow. Which branch does not have this commit in?	1	1	https://issues.jboss.org/browse/JBTM-2822
12865	JBTM-2733	Hi [~tomjenkinson], I was wondering on your comment https://issues.jboss.org/browse/JBTM-2733?focusedCommentId=13284121#comment-13284121 and now I found some time to go through the XA spec for it. I wasn't able to find reason behind why the recover sequence should be {{xa_recover(START), xa_commit/rb, xa_recover(END)}} and can't be e.g. like {{xa_recover(START|END), xa_commit/rb}}. Would you be so kind and elaborate a bit on this? Or can you point me to spec/article where I got information how this is expected to be used? It's quite possible that I just haven't read/searched properly. I'm just curious. Thanks in advance.	1	1	https://issues.jboss.org/browse/JBTM-2733
12866	JBTM-2686	Where is the link to the PR?	0	0	https://issues.jboss.org/browse/JBTM-2686
12867	JBTM-2660	It looks like the jca spec 1.5 dependency JTA 1.1 {code} +- org.jboss:jboss-transaction-spi:jar:7.3.0.Final:provided [INFO] | \- org.jboss.spec.javax.resource:jboss-connector-api_1.5_spec:jar:1.0.0.Final:provided [INFO] | \- org.jboss.spec.javax.transaction:jboss-transaction-api_1.1_spec:jar:1.0.0.Final:provided {code}	1	1	https://issues.jboss.org/browse/JBTM-2660
12868	JBTM-2634	Can you ask for feedback on the forum now it seems complete.	0	0	https://issues.jboss.org/browse/JBTM-2634
12869	JBTM-2634	It looks like it is in the spring-boot-1.4.0-M2 release now.	0	0	https://issues.jboss.org/browse/JBTM-2634
12870	JBTM-2606	Can we get a link to your branch here please?	1	1	https://issues.jboss.org/browse/JBTM-2606
12871	JBTM-2605	Is this done?	1	1	https://issues.jboss.org/browse/JBTM-2605

#	Requirement	Questions	Answered	Successfully Answered	URL
12872	JBTM-2599	[~lbarreiro] - understood. Will do the release as the PR passed. It will be tomorrow AM (UK) though. I don't think characterising this as a bug (in Narayana) is correct so changing the issue type. Also, please can you add a description? Probably just replicate your comment https://issues.jboss.org/browse/JBTM-2599?focusedCommentId=13145088&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13145088	1	1	https://issues.jboss.org/browse/JBTM-2599
12873	JBTM-2591	Yes, using 5.2.9.Final produces an example that works ... 15:43:35,988 INFO [org.jboss.weld.deployer] (MSC service thread 1-1) WFLYWELD0006: Starting Services for CDI deployment: ROOT.war 15:43:36,055 INFO [org.jboss.weld.Version] (MSC service thread 1-1) WELD-000900: 2.2.16 (SP1) 15:43:36,303 INFO [org.jboss.weld.deployer] (MSC service thread 1-3) WFLYWELD0009: Starting weld service for deployment ROOT.war 15:43:38,196 INFO [stdout] (ServerService Thread Pool -- 58) StartupBean call 15:43:38,217 INFO [stdout] (ServerService Thread Pool -- 58) MyEJB call 15:43:38,217 WARN [com.arjuna.ats.arjuna] (ServerService Thread Pool -- 58) ARJUNA012094: Commit of action id 0:ffffc0a80121:-26a9308d:567c12a6:b invoked while multiple threads active within it. 15:43:38,218 WARN [com.arjuna.ats.arjuna] (ServerService Thread Pool -- 58) ARJUNA012107: CheckedAction::check - atomic action 0:ffffc0a80121:-26a9308d:567c12a6:b commiting with 1 threads active! 15:43:38,602 INFO [javax.enterprise.resource.webcontainer.jsf.config] (ServerService Thread Pool -- 58) Initializing Mojarra 2.2.12-jbossorg-2 20150729-1131 for context '' 15:43:40,112 INFO [org.wildfly.extension.undertow] (ServerService Thread Pool -- 58) WFLYUT0021: Registered web context: / 15:43:40,151 INFO [org.jboss.as.server] (ServerService Thread Pool -- 34) WFLYSRV0010: Deployed "ROOT.war" (runtime-name : "ROOT.war") 15:43:40,264 INFO [org.jboss.as] (Controller Boot Thread) WFLYSRV0060: Http management interface listening on http://127.0.0.1:9990/management 15:43:40,266 INFO [org.jboss.as] (Controller Boot Thread) WFLYSRV0051: Admin console listening on http://127.0.0.1:9990 15:43:40,266 INFO [org.jboss.as] (Controller Boot Thread) WFLYSRV0025: WildFly Full 9.0.2.Final (WildFly Core 1.0.2.Final) started in 14595ms - Started 330 of 510 services (221 services are lazy, passive or on-demand) I won't raise a pull request because maybe there's a reason it shouldn't be 5.2.9.Final. Whoever takes this issue should provide the right fix, i.e., don't wait for me :)	0	0	https://issues.jboss.org/browse/JBTM-2591
12874	JBTM-2564	Because if so we need to do it in a profile at a guess. Unless those arguments work on an IPv6 box. Do we know which quickstarts need those? [~hhovsepy] - can you raise a PR if you have the fix on your box so we can check it works on our CI cluster too please?	1	1	https://issues.jboss.org/browse/JBTM-2564
12875	JBTM-2564	Hi [~hhovsepy] - I thought Mike said you had tested a fix locally? If so, please can you commit it and PR it and I can amend it as necessary? Alternatively just paste in the surefire config you tested as a comment and I will try it here.	1	1	https://issues.jboss.org/browse/JBTM-2564
12876	JBTM-2564	Did you need any fixes to Narayana to make it work on IPv6 because I think [~hhovsepy] is reporting they don't work for him? I think the reason they would work on IPv6+IPv4 is arq will be fine to run with its IPv4 args. I therefore don't think we can mark it as resolved because the QS are running on IPv6 with IPv4 too so that is difference to the summary+description of this issue.	1	1	https://issues.jboss.org/browse/JBTM-2564
12877	JBTM-2564	As I understood, this task was created to add new maven profile into several quickstarts to pass IPv6 parameters (mentioned in earlier comments), right? Is it already done?	1	1	https://issues.jboss.org/browse/JBTM-2564
12878	JBTM-2553	I did some debug with the CrashRecovery2 Run with the following command {code} ./build.sh -f ArjunaJTA/jta/pom.xml clean test -Dtest=CrashRecovery2 -Dorg.jboss.byteman.debug -Dorg.jboss.byteman.verbose -PcodeCoverage {code} check the output of the test and it looks like the rule does not be triggered as it can not find the local binding var "$this" The fail2pc rule is {code} RULE Fail 2PC CLASS com.arjuna.ats.arjuna.coordinator.BasicAction METHOD phase2Commit AT ENTRY BIND NOTHING IF readCounter("phase2commit") == 0 DO debug("Called"), incrementCounter("phase2commit"); com.arjuna.ats.arjuna.coordinator.ActionManager.manager().remove($this.get_uid()); RETURN ENDRULE {code} The failure is {code} start:4 triggerPos:0 end:535 (I add this debug message in agent/src/main/java/org/jboss/byteman/agent/adapter/RuleCheckMethodAdapter.java) RuleCheckMethodAdapter.checkBindings : invalid local variable binding $this checking method phase2Commit(Z)V {code} if we run without the jacoco, the start and end is {code} start:0 triggerPos:0 end:327 {code} [~adinn] how do you think about this ?	1	1	https://issues.jboss.org/browse/JBTM-2553
12879	JBTM-2553	Hi Amos, thanks for reporting this. I don't know why you are seeing that error. It's very odd because $this is the only local variable reference that we can identify when there is no local variable table. Could you try changing the rule to use $0 instead of $this and see if that makes it work.	1	1	https://issues.jboss.org/browse/JBTM-2553
12880	JBTM-2553	Doh! I see Amos is way ahead of me here (nice work, Amos :-) The trace statement Amos inserted shows that the local variable checking code has found an entry in the local var table for $this but one that starts at bytecode offset 4 and ends at offset 535. The injection code is trying to inject the rule at offset 0 i.e. at entry to the code. This is after transformation by Jacoco. With no transformation the local var this is found starting at offset 0. I have not yet looked at the Jacoco code but it seems pretty obvious what is happening. When Jacoco inserts profiling instructions into the method it updates the start and end offsets for entries in the local variable table to take account of the new code. One of those profiling instruction sequences will be injected at the start of the method and Jacoco appears to be shifting the local var start offset for this so that it comes into scope after the profiling code rather than at bytecode 0. No doubt it is also shifting the start position for method parameter entries. This is at best unhelpful; actually, I think it is an error because this and the parameters are /always/ in scope from the start of the method even if code is inserted before the start. I thought that ASM fixed this as a special case. Which version of Jacoco is being used here? The latest Jacoco version uses ASM 5 which shoudl deal with this problem. Have you tried updating to a later Jacoco? If that doesn't work I'll take this up with the Jacoco developers.	1	1	https://issues.jboss.org/browse/JBTM-2553
12881	JBTM-2553	Thanks for the detailed explanation Andrew - it certainly sounds a complex issue. Do you know of any other coverage frameworks that Byteman might interoperate with today? If not we can just disable the byteman tests under coverage testing. If you do get a Jacoco issue reference please can you let us know so we can reference it here? Thanks again!	1	0	https://issues.jboss.org/browse/JBTM-2553
12882	JBTM-2532	Did you put back a commit for the unit tests for now?	1	1	https://issues.jboss.org/browse/JBTM-2532
12883	JBTM-2522	I meant https://github.com/jbosstm/performance/tree/master/comparison, but I'm not sure it's going to work. Is it possible to use JMH to test Arquillian tests? Or applications in WildFly deployed manually?	0	0	https://issues.jboss.org/browse/JBTM-2522

#	Requirement	Questions	Answered	Successfully Answered	URL
12884	JBTM-2495	The failures related to the upgrading to the activemq artimes in WFLY 10. Also in the Step 5, you should run with standalone-blacktie.xml {code} copy docs/examples/configs/standalone-blacktie.xml wildfly-<version>/standalone/configurations bin/standalone.sh -c standalone-blacktie.xml {code} I'm sorry that I can not find in the QS to say use WFLY 9 but I will update the README.md to say it should integrate with the master from WFLY. can you show me where you find the information to run with WFLY 9 ?	0	0	https://issues.jboss.org/browse/JBTM-2495
12885	JBTM-2495	I have tried using Wildfly 10 and it worked. THanks for the help. I used Wildfly 9 as 10 was a beta and was not sure it was stable. Also there was no mention that Blacktie only integrates with master from WFLY.	1	1	https://issues.jboss.org/browse/JBTM-2495
12886	JBTM-2490	Hi Hayk, Would you have chance to work on this at all? You have marked it as a blocker so trying to work out the priority as you didn't mark a fix version. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-2490
12887	JBTM-2481	Hi [~gytis], this has been Coding in Progress for a while. Is it worth marking "Work Stopped" as I don't think you are currently coding on this one?	1	1	https://issues.jboss.org/browse/JBTM-2481
12888	JBTM-2481	Are we not planning to make the release completely automatic and incorporating it into our CI?	1	1	https://issues.jboss.org/browse/JBTM-2481
12889	JBTM-2478	a "Tricks" section to the README.md if the SVG file gets too big. Brendan is working on a Python version of FlameGraph https://github.com/brendangregg/FlameGraph/tree/master/dev which has search support: Search -> com/arjuna -> Purple nodes	0	0	https://issues.jboss.org/browse/JBTM-2478
12890	JBTM-2464	Can this be closed?	0	0	https://issues.jboss.org/browse/JBTM-2464
12891	JBTM-2458	Is this done?	1	1	https://issues.jboss.org/browse/JBTM-2458
12892	JBTM-2458	What was the resolution?	1	0	https://issues.jboss.org/browse/JBTM-2458
12893	JBTM-2440	Do you know why these only seemed to start failing when we upgraded to JDK8?	1	1	https://issues.jboss.org/browse/JBTM-2440
12894	JBTM-2440	Hi Mike, Am I right in thinking that this should be fixed, but we still have JBTM-2454? If so it is worth closing this one. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-2440
12895	JBTM-2440	I asked [~adinn] to provide an opinion about whether it could be the annotation rather than the agent loading functionality that is the root cause of the problem and he suspects that it is most likely a bug in the dynamic agent loading behaviour in the J9 compiler. He also gave a good analysis of the issue which I am including here: {quote} I downloaded IBM's Java 7 SDK and tried to use it to build Byteman. That's enough to thoroughly test the use of the agent. The tests on the Byteman agent maven module are configured to install the agent from the command line using the -javaagent option (I do that rather than use the BMUnitRunner to auto-load because I don't want failures in BMUnit to stop the agent jar being built). These tests all worked ok. So, the IBM JDK does support all the functionality of the Byteman agent when you load it from the command line. The maven tests in the BMUnit maven module rely the BMUnitRunner annotation to auto-install the agent into an already running test JVM and this is what failed. The error output is attached. As you can see from the first exception trace the IBM JVM does implement dynamic loading -- the Attach operation is fielded and J9 tries to install the agent by executing com.ibm.tools.attach.javaSE.Attachment.loadAgentLibrary If you look at the cause of the exception you can see that this eventually calls into org.jboss.byteman.agent.Main.premain which is the Byteman entry point. This tries to register its ClassFileTransformer (InstrumentationImpl.addTransformer) and at this point J9 blows up. It is objecting to the agent's request for the ClassFileTransformer to be installed as a 'retransformable' transformer. That means it wants to be able to handle retransformations of existing classes as well as being given a chance to modify newly loaded classes. This seems a tad odd because when the Byteman agent is loaded on the command line it also asks its ClassFileTransformer to be installed as a retransformable transformer. So, the refusal by J9 must be some sort of policy decision on its part. I suspect that this may be a performance hack -- J9 may be willing to not perform certain optimizations in order to allow for retransformers if you install them from the get go but otherwise may decide to switch on optimizations that mean subsequent use of retransformers won't work. Whatever the reason IBM have done this you are snookered until you can get them to remove this restriction. n.b. it is a legitimate (i.e. JVMTI agent spec-compliant) behaviour but it's not exactly very useful. I will make eqnuiries to see if I can talk to someone in IBM and get this removed or, at least, find a way to relax it but I am not sure I will have any traction. I think there is only one workaround which will fix this if you want to use BMUnit on both OpenJDK/Oracle and IBM. You will need to switch to loading the agent from the command line (whatever JVM you are using). To do this you need to: i) install the agent on the command line of the test JVM by configuring java command line argument -javaagent:/path/to/byteman.jar=listener:true ii) stop BMUnit from loading the agent by adding adding the property setting -Dorg.jboss.byteman.contrib.bmunit.agent.inhibit n.b. you will need to build the path to the byteman jar by referencing the relevant dependency in your maven repo dir. you should be able to do that using by substituting maven property values (repo dir + byteman dependency version) into the path n.b. if you don't do the load yourself then you may want to set some other -javaagent options and/or Byteman system properties. The auto load will do the equivalent of adding boot:/path/to/byteman.jar to the -javaagent line and adding -Dorg.jboss.byteman.allow.config.update as a system property setting. The former allows you to inject into JDK runtime classes. The latter allows BMUnit to pay heed to your any options specified using a BMUnitConfig annotation on your test class. regards, Andrew Dinn {quote}	0	0	https://issues.jboss.org/browse/JBTM-2440
12896	JBTM-2432	Is there anything in Narayana that requires Java8 though? If not then why not still allow that at build time?	0	0	https://issues.jboss.org/browse/JBTM-2432
12897	JBTM-2432	That's WildFly and I do understand why they're moving to 8. However, if someone wants to build a component project using 7 then why wouldn't we allow that within that project until/unless we add features only available in 8?	0	0	https://issues.jboss.org/browse/JBTM-2432
12898	JBTM-2412	Thanks Tom, Which component would that be. Tom	1	0	https://issues.jboss.org/browse/JBTM-2412
12899	JBTM-2398	After the code review I have moved the implementation of the new methods into VolatileStore and TwoPhaseVolatileStore and added a new property to the store configuration bean to control the behaviour. As expected the change has no effect on the benchmarks for the new stores: -i 1 -wi 10 -f 1 -t 10 -r 100 # 10 warm up iterations and then a single iteration lasting 100 seconds using 1 thread c.a.a.j.x.p.VolatileStoreBenchmark.testVolatileStore thrpt 1 637563.988 NaN ops/s c.a.a.j.x.p.TypedVolatileStoreBenchmark.testVolatileStore thrpt 1 185962.192 NaN ops/s c.a.a.j.x.p.TwoPhaseVolatileStoreBenchmark.testVolatileStore thrpt 1 144480.394 NaN ops/s c.a.a.j.x.p.TypedTwoPhaseVolatileStoreBenchmark.testVolatileStore thrpt 1 107737.372 NaN ops/s	0	0	https://issues.jboss.org/browse/JBTM-2398
12900	JBTM-2392	Looks like a duplicate of JBTM-2216 to me. https://github.com/jbosstm/narayana/blame/master/ArjunaJTA/jta/classes/com/arjuna/ats/internal/jta/recovery/arjunacore/XARecoveryModule.java Its fixed in WFLY master version (was fixed in Narayana 5.0.4 which should be in WFLY since 9.Beta1).	0	0	https://issues.jboss.org/browse/JBTM-2392

#	Requirement	Questions	Answered	Successfully Answered	URL
12901	JBTM-2390	I am surprised you think its an anomaly. As an example, when you turn volatile store on our performance is something much more than a 15 fold improvement iirc? Some suggestions: 1. Are you using HQ store (can you use it in this config)? 2. Have you tried enabling the volatile store? In terms of notx vs emptytx I expect that is all down to the creation of the TX endpoints. You could profile the server to get more details there. I imagine your participants are no-ops so in the no tx case you have 2 no-op calls vs the empty tx case where you have 3 no-ops (I will class commit() here as a no-op as I assume it really is very little) and 1 op (the creation of the tx endpoints). Seems reasonable to me. In the emptytx case, can you drop the commit calls (without it falling over) to see the cost of the commit vs the create?	1	1	https://issues.jboss.org/browse/JBTM-2390
12902	JBTM-2379	Is it OK to raise this one now that 2377 is merged thanks?	1	1	https://issues.jboss.org/browse/JBTM-2379
12903	JBTM-2379	Great, please can you let me know when the PR is merged so I can start the 5.1.1 release?	0	0	https://issues.jboss.org/browse/JBTM-2379
12904	JBTM-2366	EDIT: The first link will need to be viewed using: https://drive.draw.io/?#G0B9c6Ahw-picmQklPZXp5M19nalk as the text does not seem to be visible on the Google Docs preview ---------- Gytis has a draft, I have shared it publicly but its still very much in development: https://drive.google.com/file/d/0B9c6Ahw-picmQklPZXp5M19nalk/view?usp=sharing For comparison, this is the existing version: http://narayana.io/images/narayana-architecture-diagram.jpg I think we need to align the cosmetics with other JBoss.org projects. Are isometric boxes still the most graphically clear?	1	1	https://issues.jboss.org/browse/JBTM-2366
12905	JBTM-2358	We should try to get the best JDBC results we can so should try to use a database that is either on the same machine ideally. Does haverstraw host a postgres?	1	1	https://issues.jboss.org/browse/JBTM-2358
12906	JBTM-2345	Would we need a variant of WFLY-1119 (Assign an unique NodeID automatically) to avoid conflicts (ie would different images share the same JDBC store). Also it is possible to point the tooling (ObjStoreBrowser) at a filesystem to manage the logs although it is unsupported (normally it is done via a JMX browser connected to the TM process). Do we need similar functionality for accessing the JDBC store directly?	0	0	https://issues.jboss.org/browse/JBTM-2345
12907	JBTM-2344	Happy to go with either, I wonder if Container could be confused with JEE container though. Maybe Cloud? I had a quick look at other projects and Forge has Docker and LiveOak seemed to use Cloud. I am happy with Docker though even with the potential overloading.	1	1	https://issues.jboss.org/browse/JBTM-2344
12908	JBTM-2343	How is this different to the current class that allows thread-to-transaction association and disassociation?	0	0	https://issues.jboss.org/browse/JBTM-2343
12909	JBTM-2334	Does that address the same use case though? If not then we should re-examine what we're doing.	1	1	https://issues.jboss.org/browse/JBTM-2334
12910	JBTM-2318	Agreed. Maybe create a separate JIRA for the tooling?	0	0	https://issues.jboss.org/browse/JBTM-2318
12911	JBTM-2318	Yeah, I haven't upgraded the logging for the component. The only thing I can think is that maybe the default used to be WARN or something? From looking at the log statements there looks to be three statements, its the way the log messages are formatted that is causing each of the statements to bridge two lines. The logging format looks to be the format generated by JDK logging rather than log4j so I guess this is in an IDE? {code} 2014-12-18 16:30:49,755 [main] INFO (ats.arjuna :88) - ARJUNA012163: Starting service com.arjuna.ats.arjuna.recovery.ActionStatusService on port 51214 2014-12-18 16:30:49,765 [main] INFO (ats.arjuna :327) - ARJUNA012337: TransactionStatusManagerItem host: 127.0.0.1 port: 51214 2014-12-18 16:30:49,765 [main] INFO (ats.arjuna :138) - ARJUNA012170: TransactionStatusManager started on port 51214 and host 127.0.0.1 with service com.arjuna.ats.arjuna.recovery.ActionStatusService {code} I am happy for those statements to move down to debug but the reason they are info is the default option seems to be to choose a random port so it might be the case then that people would all enable debug just to get at that. I can see why with multiple processes on the same machine why electing a port dynamically has some advantages but I know that WFLY at least configure a default of 4713. (4712 for recovery port). Why don't we use those defaults and only then log the statement at INFO if the user changes the port to 0?	0	0	https://issues.jboss.org/browse/JBTM-2318
12912	JBTM-2299	It looks like we don't need it as we has removed the TAO dependency in JBTM-1665 and used the http rest methods to access the TM [~tomjenkinson], Does it make sense ?	1	1	https://issues.jboss.org/browse/JBTM-2299
12913	JBTM-2296	[~tomjenkinson], you had a go with it some time ago. Do you think it's usable enough to be merged into our dockerfiles repo?	1	1	https://issues.jboss.org/browse/JBTM-2296
12914	JBTM-2286	Hi Mark, I clarified the description a bit. Its an addition, not a replacement of the existing standalone set. Its for ease of publishing on the site. Hope you agree it would be beneficial for that. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-2286
12915	JBTM-2279	Wouldn't it have been easier/better/less invasive to have hidden this within a call to something like getCheckedActionFactory()? _checkedAction = _checkedActionFactory.getCheckedAction(get_uid(), type()); + _checkedAction = allowCheckedActionFactoryToBeOverriden ? arjPropertyManager + .getCoordinatorEnvironmentBean().getCheckedActionFactory() + .getCheckedAction(get_uid(), type()) + : defaultCheckedActionFactory.getCheckedAction(get_uid(), + type());	1	1	https://issues.jboss.org/browse/JBTM-2279
12916	JBTM-2270	JTS participant records are not in-lined with the transaction log so they appear in a separate part of the log store. The tooling only shows participants for JTA records which do have their participants inlined. Ondrej, is the Fix version field (6.0.0) ok? I just noticed that Gytis is down as the assignee so I will let him update that field.	0	0	https://issues.jboss.org/browse/JBTM-2270

#	Requirement	Questions	Answered	Successfully Answered	URL
12917	JBTM-2264	Hi Evgeniy, I have raised a pull request for this in our transactional driver code. I have also deployed a snapshot version of Narayana with the change already in there for you to have a go with. I have tested it locally (extensively) and it works for me. I can see the SNAPSHOT over here: https://repository.jboss.org/nexus/content/groups/public/org/jboss/narayana/jta/narayana-jta/5.0.4.Final-SNAPSHOT/. All you should need to do is change your pom to this version. Now for some background. BTM and Atomikos are not strictly spec compliant in my mind as their core transaction processing engine performs simple pointer comparison to determine if two XAResources are equivalent rather than the spec defined approach of using isSameRM. You can see the pointer comparison for yourself over here: https://github.com/bitronix/btm/blob/master/btm/src/main/java/bitronix/tm/internal/XAResourceManager.java#L194 You can see where the spec has pseudo code that demonstrates how the transaction manager is intended to use isSameRM in section "3.4.9 Identifying Resource Manager Instance" over here: https://jcp.org/en/jsr/detail?id=907 If two XAR are not equivalent you are meant to generate a new Xid for the branch and call start with TMNOFLAGS (as BTM/Atomikos will in this case, even though its technically the same RM) and if they are equivalent you use the same Xid and call start with TMJOIN (as we do). It turns out that there are a number of resource managers that return true to isSameRM but cannot handle the TMJOIN flag. It also turns out that this is a known issue with some of the drivers and Ironjacamar (Red Hats JCA implementation in Wildfly) already works around that for a number of drivers by wrapping their XAResources with an is-same-rm-override flag that returns false. The impact of this is that although your test would work in WildFly, there is a slight performance impact as the RM has two branches, plus should you touch the same tables you may be able to deadlock. Anyway, as you were using transactional driver rather than ironjacamar you have hit this issue. I have effectively backported that IJ workaround for a number of drivers which are also not compliant to our transactional driver code. The reason we haven't gone to pointer comparison is: 1. Spec compliance - we don't break that 2. There is an issue with overriding isSameRM which affects app servers performing a process known as transaction inflow. Due to limitations with the transaction inflow API (discussed at length elsewhere, for example https://issues.jboss.org/browse/JBTM-622) a transaction manager is not permitted to generate different Xids for resources. As such if we override the same rm report from the RM, RMs which correctly implement isSameRM and accept TMJOIN would be unfairly penalised as they would get: {code} xaResource1.start(xid1, TMNOFLAGS) xaResource2.start(xid1, TMNOFLAGS) {code} Which is a protocol error. Hope it helps, Tom	1	1	https://issues.jboss.org/browse/JBTM-2264
12918	JBTM-2262	BTW what did these scripts do in the old days? Did they add the jacorb jar explicitly?	1	1	https://issues.jboss.org/browse/JBTM-2262
12919	JBTM-2259	One of the reasons we haven't done this to date is that it's not standards compliant. Why isn't the ordering that JTA allowed in the interposed option sufficient?	0	0	https://issues.jboss.org/browse/JBTM-2259
12920	JBTM-2259	OK, but this is not compliant with the standards and therefore if some other transaction engine is used it will break. The reason ordering of resources and synchronisations isn't supported in the standards is because it becomes very hard to manage global ordering in a heterogenous environment. Are we really happy to add something which basically ties IronJacamar and Narayana together? I'm no so sure. While it's nice to add something like resource ordering, as we do with AbstractRecords, it should be done cautiously, in situations where no other option is possible, or as an optimisation, i.e., so without it the system still works, but may be slower, less efficient etc.	1	0	https://issues.jboss.org/browse/JBTM-2259
12921	JBTM-2245	Hi Tom, Mike, What is the planned fixed version of this task? Is there already a support of shutdown(timeout) in Narayana? I am asking as we have RFE https://issues.jboss.org/browse/EAP7-86 resolved and ready for QA. Thanks, Hayk.	1	0	https://issues.jboss.org/browse/JBTM-2245
12922	JBTM-2231	Hi Jesper, Please can you give me some ideas of what this for? Is it the IJ tracer? I have no problem with adding it, please just provide some context as there could be more I can give you. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-2231
12923	JBTM-2231	Hi Jesper, Is the attached PR what you were thinking? Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-2231
12924	JBTM-2230	As mentioned on the call, we need to be mindful of when we introduce this new interface method as it will require implementations to add the new method when they wish to upgrade. We could add it to a new interface. As there are various options here Mike we should create a design discussion and mention it to various teams. Can you create the discussion and flag it to Jesper (and Andy) to make sure known XARW impls are aware of the requirements (Andy obviously knows but he might want to contribute if Jesper has any observations.	1	0	https://issues.jboss.org/browse/JBTM-2230
12925	JBTM-2224	Can you point me at the forum posting where this was discussed? Then I can take my comments there rather than in JIRA.	1	1	https://issues.jboss.org/browse/JBTM-2224
12926	JBTM-2210	Granted I haven't updated my local copy for a few weeks, but these warnings seemed to come all of a sudden and a lot of them to boot. Was this expected?	0	0	https://issues.jboss.org/browse/JBTM-2210
12927	JBTM-2210	So all of these warnings have been there for 7 months?!	0	0	https://issues.jboss.org/browse/JBTM-2210
12928	JBTM-2210	No, failing on warn is probably not a good idea. I'm surprised that I didn't see it earlier, since I built the source quite a bit only a few weeks ago. I'm also surprised no one else caught it if it's been around that long. Do we regularly check the build logs? Even building in the cloud or an IDE should generate the same warnings (e.g., in Eclipse).	1	1	https://issues.jboss.org/browse/JBTM-2210
12929	JBTM-2209	Hi Jesper, Please can you add a few more details to the Jira on the reasons for the change? Tom	1	1	https://issues.jboss.org/browse/JBTM-2209
12930	JBTM-2209	Thanks, but that's about as much use to me as a chocolate fireguard! ;) Or giving me a scalar number when I need a vector. What was the test used? How many iterations? We could do this ourselves of course, but it'd be so much simpler if you've already done that work :-) And yes, I see the test code you added, but I was looking for some test that worked in situ with the Stack usage within Narayana.	1	1	https://issues.jboss.org/browse/JBTM-2209
12931	JBTM-2209	Great work on the fix and the benchmark, I agree it looks like it is worth merging. Mike, may I ask you to do the honours please as it is your PR now?	0	0	https://issues.jboss.org/browse/JBTM-2209
12932	JBTM-2206	I don't think we should change ArjunaCore/arjuna/classes/com/arjuna/ats/arjuna/coordinator/CheckedAction.java as it is in the public API so unless it has considerable performance win the cost of update apps to use the new API is probably not justified. As to the rest of the changes, can you give us a brief indication of the performance benefit of this particular change?	1	1	https://issues.jboss.org/browse/JBTM-2206
12933	JBTM-2206	Can you include a link to the forum posting here for future reference?	0	0	https://issues.jboss.org/browse/JBTM-2206
12934	JBTM-2203	Hi Jesper, Thanks for creating the separate issues! Can you now share the relative performance benefit of the enhancement please, same for JBTM-2209, JBTM-2207, JBTM-2208 and JBTM-2206? Thanks! Tom	1	1	https://issues.jboss.org/browse/JBTM-2203
12935	JBTM-2203	I don't think we can remove the *Mutex() methods from StateManager as these are part of the public API so people that inherit from that class would no longer have access to them (as you saw in the tests you had to modify). Please can you provide a few more details on what the performance win was with removing the Reentrant lock to help evaluate the benefit vs the drawback of the API change?	1	1	https://issues.jboss.org/browse/JBTM-2203

#	Requirement	Questions	Answered	Successfully Answered	URL
12936	JBTM-2203	Hold on ... doesn't this changed the public API of StateManager and LockManager? Whilst the latter is a superset of the previous API, the former isn't and potentially breaks other classes from users that may inherit from StateManager. Did we deprecate as discussed last July? Did we determine what the cost/benefit is of doing this?	1	0	https://issues.jboss.org/browse/JBTM-2203
12937	JBTM-2202	What does this have to do with REST transactions? If it's local-only, then distribution over REST is not an option.	0	0	https://issues.jboss.org/browse/JBTM-2202
12938	JBTM-2189	For clarity in future if others come to this issue, can you please describe what problem you've seen and hence why you've submitted a patch. "fails to bootstrap properly" isn't sufficient. Thanks.	0	0	https://issues.jboss.org/browse/JBTM-2189
12939	JBTM-2189	After applying the change proposed to ConfigurationInfo, the following exception (which is addressed by the proposed change in AbstractPropertiesFactory) occurs: java.io.FileNotFoundException: \naintegration-1.0-SNAPSHOT.one-jar.jar (The system cannot find the file specified) at java.util.zip.ZipFile.open(Native Method) at java.util.zip.ZipFile.<init>(Unknown Source) at java.util.zip.ZipFile.<init>(Unknown Source) at java.util.jar.JarFile.<init>(Unknown Source) at java.util.jar.JarFile.<init>(Unknown Source) at sun.net.www.protocol.jar.URLJarFile.<init>(Unknown Source) at sun.net.www.protocol.jar.URLJarFile.getJarFile(Unknown Source) at sun.net.www.protocol.jar.JarFileFactory.get(Unknown Source) at sun.net.www.protocol.jar.JarURLConnection.connect(Unknown Source) at sun.net.www.protocol.jar.JarURLConnection.getInputStream(Unknown Sour ce) at java.net.URL.openStream(Unknown Source) at com.arjuna.common.util.propertyservice.AbstractPropertiesFactory.load FromFile(AbstractPropertiesFactory.java:142) at com.arjuna.common.util.propertyservice.AbstractPropertiesFactory.getP ropertiesFromFile(AbstractPropertiesFactory.java:100) ... 33 more	1	1	https://issues.jboss.org/browse/JBTM-2189
12940	JBTM-2189	Hi Chris, Many thanks for your interest in Narayana. Normally for something like this it is best to start with a discussion on our forums over here: https://community.jboss.org/en/jbosstm. That said, please can you add a little more detail into the description as to the nature of the bug. We do already function in Java SE and also complex classloader hierarchies so this may be a corner case which we would be happy to work with you to address. I am not familiar with OneJar, is it a maven plugin that will put all the dependency jar classes into the same Jar as the application jar? The main problem with your report as it stands is that we don't have a well defined reason for _why_ the change needs to be made, if you can express the motivation for the change in more general terms then we can get it merged. As you say, the fix looks like it should work in our currently supported options so it shouldn't be a problem to get it merged once we have more context to backup the need to merge it. Once we understand the problem better, and we proceed with your approach, please can you: 1. Sign the CLA over here: https://cla.jboss.org/ (project is JBoss Transactions) 2. Fork: https://github.com/jbosstm/narayana 3. Raise a pull request on the main repo with your fix in. It should be easier to review inline, plus it will get CI run on it. Many thanks for the contribution! Tom	1	1	https://issues.jboss.org/browse/JBTM-2189
12941	JBTM-2184	What's the procedure for ensuring our artifacts make it into maven central? How are you determining our missing dependencies?	0	0	https://issues.jboss.org/browse/JBTM-2184
12942	JBTM-2184	So my guess is that we need to add the dependency to ArjunaJTA/narayana-jta/pom.xml - what is the easiest way to test that? I ran tattletale over narayana-jta and it reports transitive dependencies on other stuff as well (org.hornetq.core and org.jfree). Does that imply we need to add explicit dependencies on those too?	1	0	https://issues.jboss.org/browse/JBTM-2184
12943	JBTM-2184	Just trying to work out, what were you doing to cause the issue? For example the ArjunaJTA/maven qs doesn't exhibit this behaviour after I fixed it to have the correct runtime dependencies. Some of the quickstarts do import the SPI correctly as they must need to use this behaviour: [tom@dhcppc1 quickstart](JBTM-2182) (15:40:29) $ grepinpom transaction-spi ./jca-and-tomcat/pom.xml ./jta-and-hibernate/pom.xml ./jca-and-hibernate/pom.xml It looks like "provided" scope dependencies of constituent modules have not been included in shaded poms. Some of this goes away if we start to recommend users to use the raw modules rather than the uber jars created by the shade plugin as then the poms are self documenting (like in WildFly). At the moment it is effectively assumed that you would know which jars to include if you use the uber jars as the Jars do not document the requirement. The disadvantage of the none-uber jar is only for JTS as you would be required to provide a configuration file to enable JTS, at the moment it is selected by the jar. Unfortunately I can't get a shaded pom to include provided scope at all, I could make it runtime but then it would still be dependent on the Jars being available which for some situations shouldn't be required. Will have a think...	1	1	https://issues.jboss.org/browse/JBTM-2184
12944	JBTM-2183	Have we seen this behaviour elsewhere in the past? I don't recall it. A change in Oracle? Or just lucky for us so far not to have seen it? As to changing the default, we can't do that until the EAP 7 timeframe even if we wanted to. Plus, we'd better be able to show that this has no impact on performance, reliability etc. in environments where the RM doesn't exhibit this behaviour. Otherwise may be best to err on the side of caution and perhaps catch the error in recovery and swap in the alternate behaviour for the next sweep.	0	0	https://issues.jboss.org/browse/JBTM-2183
12945	JBTM-2180	As a start on this task I have enabled the Oracle® JRockit Flight Recorder for a subset of the performance tests in the perf CI job (http://172.17.131.2/job/narayana-performance/). Flight data is stored in the maven build directory (target/test-classes/*.jfr) and can be read using the jmc tool available in the standard JDK (edition 7u40 and above). What is the best way to evaluate other options given that there are many alternatives (and trying out each one involves engineering work): * JDK Collections (much of our code hasn't been updated to take advantage of newer releases (Jesper has been looking into this); * Goldman Sachs Collections; * Google Collections (Guava); * javolution; * Highly Scalable Java; * ... Tests with flight data currently include: Core: com.hp.mwtests.ts.arjuna.performance.Performance1 com.hp.mwtests.ts.arjuna.performance.Performance2 com.hp.mwtests.ts.arjuna.performance.Performance3 com.hp.mwtests.ts.arjuna.performance.Performance4 JTA: com.arjuna.ats.jta.xa.performance.OnePhase2PCPerformanceDefaultUnitTest com.arjuna.ats.jta.xa.performance.OnePhase2PCPerformanceVolatileUnitTest com.arjuna.ats.jta.xa.performance.OnePhasePerformanceDefaultUnitTest com.arjuna.ats.jta.xa.performance.OnePhasePerformanceVolatileUnitTest JTS com.hp.mwtests.ts.jts.local.synchronizations.Performance com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance1 com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance2 com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance3	1	1	https://issues.jboss.org/browse/JBTM-2180
12946	JBTM-2163	Hi [~zhfeng], Does your current pull request enable the recovery manager? If not we should make sure you have done that as a top priority before closing this issue. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-2163
12947	JBTM-2159	Huh? Assume I can't view the Bugzilla, which may go part way to explaining my confusion ;-) Is this a race condition between recovery and application? If so it should be taken care of already within the transaction code.	1	1	https://issues.jboss.org/browse/JBTM-2159

#	Requirement	Questions	Answered	Successfully Answered	URL
12948	JBTM-2151	Should I remove the fix version number as I don't believe this is work that is currently scheduled?	0	0	https://issues.jboss.org/browse/JBTM-2151
12949	JBTM-2122	[~tomjenkinson] Not fully sure if our discussion at https://issues.jboss.org/browse/JBEAP-610?focusedCommentId=13147405#comment-13147405 is directly connected to this issue but you previous comments seem to me so. What I remember by [~hhovsepy]'s investigation it seemed that JTS works with 2pc for propagated transactions at any time (meaning even with one resource enlisted). Later on that was found that we have (?) some trouble in test and 1pc is used. Your reference of https://c.na7.visual.force.com/apex/Case_View?sbstr=01560781 seems to me not related as it's rather caused by the fact discussing at https://issues.jboss.org/browse/JBEAP-225 and https://issues.jboss.org/browse/WFLY-4608 which is based on issue in container. _plus as action item from f2f I would say that what you mentions in your first comment is not so. if we are not wrong in our settings then jts and jboss remoting run 1pc_	1	0	https://issues.jboss.org/browse/JBTM-2122
12950	JBTM-2119	Hi Lami, Would you be interested in contributing a pull request for this enhancement? If so, please can you sign the CLA over here: https://cla.jboss.org/index.seam (the project is JBoss Transactions) The raise a pull request against either 4.17 or master over here: https://github.com/jbosstm/narayana/pulls Many thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-2119
12951	JBTM-2116	Is this something that needs to be meaningful in a distributed environment? What about a heterogeneous environment?	0	0	https://issues.jboss.org/browse/JBTM-2116
12952	JBTM-2116	I added a forum reference for this which I found. I am extrapolating a little, but is the issue related to calling hashCode in JTS afterCompletion in a subordinate node and are there other scenarios you have found it to be volatile?	0	0	https://issues.jboss.org/browse/JBTM-2116
12953	JBTM-2081	I think that if we are allowing defaults for all values (good idea) then it should be possible for someone to deploy a transactional application without a property file. Maybe give a warning or output in Info?	1	1	https://issues.jboss.org/browse/JBTM-2081
12954	JBTM-2081	Hi Mark, The issue is that jts is not a strict superset of jta. To be able to provide a single binary containing both JTS and JTA, you therefore need a properties file. Otherwise you have to decide the java hardcoded defaults at compile time, which means using code templates in the jta module compile. We could look at only requiring the properties file once you get to JTA or greater level of the framework? What do you think? Tom	1	1	https://issues.jboss.org/browse/JBTM-2081
12955	JBTM-2081	Would the intention also be to no longer ship a default jbossts-properties.xml file in the jars, and to no longer be able to change the default name of the config file in the manifest.mf?	1	1	https://issues.jboss.org/browse/JBTM-2081
12956	JBTM-2081	What values are in the default properties file these days?	1	1	https://issues.jboss.org/browse/JBTM-2081
12957	JBTM-2081	The original Android code should still be in my Android branch on svn. The vert.x code is https://github.com/nmcl/mod-stm Try removing the property file? Note, this is 5.0.0.Final of TS.	1	1	https://issues.jboss.org/browse/JBTM-2081
12958	JBTM-1972	JIRA is not the appropriate place to ask this question. Please can you raise it via the forums at http://www.jboss.org/ejb3	0	0	https://issues.jboss.org/browse/JBTM-1972
12959	JBTM-1955	We should make sure we are using the PRODUCT variable throughout. We should also add variables for the target app server (IMO). This will allow us to use Narayana and WildFly for the community docs and JBossTS and JBoss AS in the product docs. Does this last point make sense?	1	0	https://issues.jboss.org/browse/JBTM-1955
12960	JBTM-1955	I've read through the XTS docs. I've applied the minor changes (mostly typos) here: https://github.com/jbosstm/documentation/commit/238a06aa9e50ff03144c7f542c08f1b36a9e6b93. Gytis, Can you apply this to the 4.17 branch please? I also have some ideas about how to layout the docs for the new feature-work done in Narayana 5. I will discuss this with Gytis when we meet to discuss the XTS docs.	1	0	https://issues.jboss.org/browse/JBTM-1955
12961	JBTM-1943	Ondrej Chaloupka <ochaloup@redhat.com> made a comment on [bug 1009981|https://bugzilla.redhat.com/show_bug.cgi?id=1009981] Hi Tom, I help to Mirek with verification of this issue. The correct name of attribute that is needed to be set is com.arjuna.ats.jta.orphanSafetyInterval (default is 10000) Mirek ran his test with several values of the attribute. When we set the attribute to be 20s (20000) we still get the XAExceptions time to time. When we redefine the time being 40s (40000) the exception disappeared from the log. I think that I can verify this issue just don't you think that it would be better to set the default interval time a bit bigger? For not to get the error as default behavior. But I'm not sure how would it influence the whole recovery process. It's true that the problem occurs on faster machines under the load of receiving a bigger amount of messages. Ondra	0	0	https://issues.jboss.org/browse/JBTM-1943
12962	JBTM-1943	Ondrej Chaloupka <ochaloup@redhat.com> made a comment on [bug 1009981|https://bugzilla.redhat.com/show_bug.cgi?id=1009981] Hi Tom, a one more request. I've added a doc text to this bug but I don't think that my text is well written. Could you please check it and polish it? Thank you Ondra	0	0	https://issues.jboss.org/browse/JBTM-1943
12963	JBTM-1942	Is this issues resolved now? /cc [~paul.robinson]	1	1	https://issues.jboss.org/browse/JBTM-1942
12964	JBTM-1930	The definitive documentation for CLI subsystems are in the XSDs (docs/schema/jboss-as-txn_<version>.xsd). Rather than duplicate this documentation and risk it diverging my preference is to simply point the community docs at this file. Does anyone have any objections with this approach?	1	1	https://issues.jboss.org/browse/JBTM-1930
12965	JBTM-1925	Gytis, I've had a look at this issue and updated the description with what should now work. Tom and I have tried it out. It's going to be quite time-consuming, and I still have a few M5 issues left to do. Can you take a look at this please. Re-assign it back if you don't think you will have time to do it.	0	0	https://issues.jboss.org/browse/JBTM-1925
12966	JBTM-1915	Is this still likely for tomorrow?	1	1	https://issues.jboss.org/browse/JBTM-1915
12967	JBTM-1909	Seems like a few things broke or changed behaviour due to 1875. Can we have a better vetting process in future to either prevent these things creeping in or at least document them such that they don't come as a surprise?	1	0	https://issues.jboss.org/browse/JBTM-1909
12968	JBTM-1908	I notice qa and rts have printSummary set to true. Maybe do that for everything anyway? Whatever we do, we should be consistent though.	1	1	https://issues.jboss.org/browse/JBTM-1908
12969	JBTM-1908	Did we upgrade the surefire version recently?	1	1	https://issues.jboss.org/browse/JBTM-1908
12970	JBTM-1855	[~gytis], [~zhfeng] mentioned that he had asked you to review: https://github.com/zhfeng/narayana/commit/b51faee71db0977898f471e806b5e88162100945 Have you had chance to do this yet?	1	1	https://issues.jboss.org/browse/JBTM-1855
12971	JBTM-1852	I see. Now I understand the issue but not the context nor the solution, thanks. How does the participant recover and who rolls back the participant if there is no transaction.	1	1	https://issues.jboss.org/browse/JBTM-1852
12972	JBTM-1838	Done in 4.17 using EAP: https://github.com/jbosstm/quickstart/commit/5d7a4ea39bea48c90b63de783381ced493ba343e Gytis: can you apply to master using WildFly. I think the commands are the same, but you will need to test them. Also, you will need to rename the occurrences of JBoss EAP 6.1 with WildFly. I think it should be less specific about the actual version downloaded. Just tell the user to go get WildFly from the download page and then unzip it. This will stop it going out of date every time a new version is released.	0	0	https://issues.jboss.org/browse/JBTM-1838

#	Requirement	Questions	Answered	Successfully Answered	URL
12973	JBTM-1821	Conversation with Gytis on a potential solution: pfrobinson: whilst you are at it, can you have a quick check to see if you can add modules to the classloader in that phase too? [12:06pm] gytist: you mean dependencies like org.jboss.xts? [12:10pm] pfrobinson: yeah [12:10pm] pfrobinson: I want to scan for the use of @Transactional and then add the org.jboss.jts depenedency to the deployment [12:10pm] gytist: I think it is, some of the subsystems automatically add dependencies [12:11pm] pfrobinson: yeah, that was my thoughts too [12:11pm] pfrobinson: for example, if persistence.xml is present, you add the jpa deps [12:11pm] gytist: I saw it when I was researching my issue on Monday [12:11pm] gytist: need to find it again [12:12pm] gytist: grep for this class WSDependenciesProcessor [12:12pm] gytist: that's what you need [12:12pm] pfrobinson: great, thanks	0	0	https://issues.jboss.org/browse/JBTM-1821
12974	JBTM-1816	{code} Jul 03, 2013 1:48:11 PM com.arjuna.ats.internal.jts.orbspecific.javaidl.recoverycoordinators.JavaIdlRCServiceInit initORBandOA INFO: ARJUNA022087: RecoveryServer using existing ORB Jul 03, 2013 1:48:12 PM com.sun.corba.se.impl.orb.ORBDataParserImpl getPersistentServerId WARNING: "IOP00910206: (INITIALIZE) Persistent server ID is not set" org.omg.CORBA.INITIALIZE: vmcid: SUN minor code: 206 completed: Maybe at com.sun.corba.se.impl.logging.ORBUtilSystemException.persistentServeridNotSet(ORBUtilSystemException.java:3643) at com.sun.corba.se.impl.logging.ORBUtilSystemException.persistentServeridNotSet(ORBUtilSystemException.java:3657) at com.sun.corba.se.impl.orb.ORBDataParserImpl.getPersistentServerId(ORBDataParserImpl.java:307) at com.sun.corba.se.impl.oa.poa.POAPolicyMediatorBase.<init>(POAPolicyMediatorBase.java:107) at com.sun.corba.se.impl.oa.poa.POAPolicyMediatorImpl_NR_UDS.<init>(POAPolicyMediatorImpl_NR_UDS.java:52) at com.sun.corba.se.impl.oa.poa.POAPolicyMediatorFactory.create(POAPolicyMediatorFactory.java:45) at com.sun.corba.se.impl.oa.poa.POAImpl.initialize(POAImpl.java:383) at com.sun.corba.se.impl.oa.poa.POAImpl.create_POA(POAImpl.java:799) at com.arjuna.ats.internal.jts.orbspecific.javaidl.recoverycoordinators.JavaIdlRCServiceInit.getRCPOA(JavaIdlRCServiceInit.java:110) at com.arjuna.ats.internal.jts.orbspecific.javaidl.recoverycoordinators.JavaIdlRCServiceInit.startRCservice(JavaIdlRCServiceInit.java:177) at com.hp.mwtests.ts.jts.orbspecific.recovery.JavaIdlRecoveryUnitTest.testInit(JavaIdlRecoveryUnitTest.java:74) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:47) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:50) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:222) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:252) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:141) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:112) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at	0	0	https://issues.jboss.org/browse/JBTM-1816
12975	JBTM-1816	How are you running the test. The following syntax used to work but not anymore: {code} ./build.sh -f ArjunaJTS/pom.xml test -Dtest=com.hp.mwtests.ts.jts.orbspecific.recovery.JavaIdlRecoveryUnitTest -Didlj-enabled {code}	0	0	https://issues.jboss.org/browse/JBTM-1816
12976	JBTM-1816	The reason in the difference in the tests is because if JacORB runs first it has put its RC entry in the object store. You can replicate the issue by using the ObjectStore in the attached jacorb-rc-in-object-store.tgz. If you put it as ArjunaJTS/jts/target/test-classes/ObjectStore then even "./build.sh -f ArjunaJTS/jts/pom.xml test -Dtest=JavaIdlRecoveryUnitTest -Didlj-enabled" would fail. iirc the RC manager uses a known UID hence why running JacORB then JDKORB in the same ObjectStore will be a problem. If you notice, the specific test that fails isn't actually the one that is generated in the log you pasted, so success or fail you get those log messages. i.e. those log messages come from tests where failure is expected. That being said, I don't know if there are any tests where success is expected - we should make sure we have some of those as that log message implies that the JDK orb hasn't been configured correctly. i.e. do we have a test that correctly initializes a JDK ORB RC manager? The only tests I see in JavaIdlRecoveryUnitTest are testing failure scenarios.	0	0	https://issues.jboss.org/browse/JBTM-1816
12977	JBTM-1809	Did this get resolved? I saw some commits on the 4.17 branch	1	1	https://issues.jboss.org/browse/JBTM-1809
12978	JBTM-1808	Tom, I was thinking that this would be good for Amos to investigate. Can you add any ideas you may have to the above text and then assign accordingly?	1	1	https://issues.jboss.org/browse/JBTM-1808
12979	JBTM-1808	Maybe this isn't worth the bother? How about we split the job up into 5 1hr pieces and defer this task until we are in a position where we are using too much CI time?	1	1	https://issues.jboss.org/browse/JBTM-1808
12980	JBTM-1789	Hi Paul, Can you list the files that are in the uber jars but not in the individual jars? The AS currently just consumes the uber jars so I want to make sure when we move to using the individual jars in the AS we can anticipate any further modifications required. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-1789
12981	JBTM-1788	Hi Tom, [tom@dhcppc1 quickstart](4.17.4.Final)) $ git name-rev HEAD HEAD tags/4.17.4.Final [tom@dhcppc1 quickstart](4.17.4.Final)) $ find . -name \pom.xml -exec grep -l 4.17.5 {} \; [tom@dhcppc1 quickstart](4.17.4.Final)) $ I can't find any reference to 4.17.5 on that tag (SNAPSHOT or otherwise). Can you try to find which file has that reference in? Maybe somehow your repo is corrupted? Tom	1	0	https://issues.jboss.org/browse/JBTM-1788
12982	JBTM-1772	Is this not something that the build can automate?	1	0	https://issues.jboss.org/browse/JBTM-1772
12983	JBTM-1770	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	1	0	https://issues.jboss.org/browse/JBTM-1770
12984	JBTM-1754	Hi Paul, The dir name and artifact id are effectively fixed at the moment. However, as you say we have more flexibility on the "project name" field. Can you update your table with an example of your preferred project name for those modules you have referenced originally, please? Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-1754
12985	JBTM-1754	Hi Paul, is the "Narayana: " prefix implied and therefore a little redundant (I realise I added the redundancy myself with JBossJTS, so just trying to get your opinion here)? if not the table looks good to me - thanks!	1	1	https://issues.jboss.org/browse/JBTM-1754

#	Requirement	Questions	Answered	Successfully Answered	URL
12986	JBTM-1742	In what way did it break the build? I thought I had built everything when I changed this yesterday.	1	1	https://issues.jboss.org/browse/JBTM-1742
12987	JBTM-1721	Is there a pull request for this on WildFly or does it require M5?	1	1	https://issues.jboss.org/browse/JBTM-1721
12988	JBTM-1707	Hi, Could you please confirm, can we use JBossTS 4.3.GA release for above mentioned problem. Thanks, Sube	1	0	https://issues.jboss.org/browse/JBTM-1707
12989	JBTM-1699	Should we link this to JBTM-1970 as well?	0	0	https://issues.jboss.org/browse/JBTM-1699
12990	JBTM-1695	Hi Weinan, We left as much of the old builds directory structure in place as maven could cope with. Please can you clarify if there is an EAP issue with the code in these folders as we have previously decided not to make cosmetic changes to the structure? Tom	1	1	https://issues.jboss.org/browse/JBTM-1695
12991	JBTM-1694	{quote} we are using a custom fork but only for community testing {quote} That's what our thirdparty-releases repository [1] is for, so ideally we should upload the patched release there. Were the patches also submitted to upstream emma so they can get into a future release? Moving the emma tests to a profile would be good. [1]https://community.jboss.org/wiki/MavenDeployingaThirdpartyRelease	0	0	https://issues.jboss.org/browse/JBTM-1694
12992	JBTM-1689	Which profiles do your team regularly use? I guess we should keep just the most commonly used ones?	0	0	https://issues.jboss.org/browse/JBTM-1689
12993	JBTM-1689	I would be happy to remove: all, core, jta, jts, jts-idlj and stm from /pom.xml [~marklittle] does use "mvn install -Pstm", but perhaps he will be happy to use: "mvn install -pl STM -am" instead? Leave it with me and I will raise the pull request if Mark is happy to stop using the profile. For completeness, this is what would be replaced: all (as used in EAP builds) = mvn install core = mvn install -am -pl ArjunaCore/arjunacore jta = mvn install -am -pl ArjunaJTA/narayana-jta jts = mvn install -am -pl ArjunaJTS/narayana-jts-jacorb jts-idlj = mvn install -am -pl ArjunaJTS/narayana-jts-idlj -Didlj-enabled=true stm = mvn install -am -pl stm	1	1	https://issues.jboss.org/browse/JBTM-1689
12994	JBTM-1680	Please can you add an estimate for this?	1	0	https://issues.jboss.org/browse/JBTM-1680
12995	JBTM-1679	[~tomjenkinson], maybe this would be worth pushing to 5.next?	1	1	https://issues.jboss.org/browse/JBTM-1679
12996	JBTM-1679	How is this different to JBTM-1107? Should we close this as a duplicate?	1	0	https://issues.jboss.org/browse/JBTM-1679
12997	JBTM-1679	I think we should just close it and add it into the details for JBTM-1107 as it seems somewhat intrinsic to the behaviour. If you agree, please can you make the change?	0	0	https://issues.jboss.org/browse/JBTM-1679
12998	JBTM-1672	Please can you clarify the reason to disable the maven-surefire-plugin? If it is still needed, please can you raise a separate Jira for this one?	0	0	https://issues.jboss.org/browse/JBTM-1672
12999	JBTM-1658	Assign this to me, as I think I know what the problem is?	0	0	https://issues.jboss.org/browse/JBTM-1658
13000	JBTM-1645	How about we do this: Split the PR job into a multi config with the same split as Narayana: MAIN XTS QA_JTA QA_JTS_JACORB QA_JTS_JDKORB These will all be scheduled in a random order, which is annoying if you are willing to merge based on a sub-set of tests. It wastes resources, and causes you to wait longer as the tests you are uninterested in may be scheduled first. The solution is to specify in the description which of the above groups you wish to skip. This is similar to what you mention above, but a bit more course-grained and fits in more neatly with the hudson multi-config approach. You can get the PR description with the following command. There's probably a better approach, but this works: curl -s https://api.github.com/repos/jbosstm/narayana/pulls/279 | grep \"body\": I don't think it makes sense to allow MAIN to be skipped, and if you want to skip one lot of QA tests, you probably want to skip them all? So how about we just have these two options: skipXTS skipQA We could add more fine-grained options if it proves useful. I don't see the point in specifying who can do this. If a contributor skips some tests that need to pass, we simply tell them to update the description and re-test.	0	0	https://issues.jboss.org/browse/JBTM-1645
13001	JBTM-1645	How about if the pull just changes a JacORB qa test? Or just an XTS test, that would be rationale to skip MAIN wouldn't it? Agreed on the white listing, we just need to be vigilant :)	0	0	https://issues.jboss.org/browse/JBTM-1645
13002	JBTM-1643	Why focus solely on JTS? Why not support something like this for each component?	1	0	https://issues.jboss.org/browse/JBTM-1643
13003	JBTM-1641	I was unable to duplicate the problem on jboss-as-8.0.0.Alpha1-SNAPSHOT. I was able to duplicate it on EAP. I agree that printing the stacktrace in this instance is excessively verbose. We are in the process of reworking the transactions SPI and as part of that work I will rationalise the debug logging, however we may not make the changes in time for wildfly 8 release. I have added a link from JBTM-1699 (Changes to the transaction SPI) to make sure we don't loose sight of this JIRA. Would you object to me closing this issue "as will not fix".	0	0	https://issues.jboss.org/browse/JBTM-1641
13004	JBTM-1635	Hi Jesper, Do you want me to pretty-print the flag or is the int enough? I don't think you can OR the flags on delist so the int should be enough? Tom	1	1	https://issues.jboss.org/browse/JBTM-1635
13005	JBTM-1602	I just noticed this change in the code and tracked it back to github (no JIRA associated with it, so I couldn't see what the issue was): try { Object syncObject = ((BasicAction.Current() == null) ? getMutex() : BasicAction.Current()); synchronized (this) { synchronized (getMutex()) { synchronized (syncObject) { I'm not sure why this change is in here, but need to check because it doesn't look right. Ignoring the fact that we now have 3 locks being acquired before the actual locksHeld lock, we're also potentially acquiring the same lock twice (syncObject could be getMutex())!	1	1	https://issues.jboss.org/browse/JBTM-1602
13006	JBTM-1598	Hi Weinan, I can do this, but is it not ok for these all to be in the root pom in a dependency management section? I thought that was the preferred route for jboss projects? e.g. grep https://github.com/jbosstm/narayana/blob/master/pom.xml for "common<" Tom	1	1	https://issues.jboss.org/browse/JBTM-1598
13007	JBTM-1584	Should this be done as part of the BLACKTIE-381 / BLACKTIE-202 package?	1	1	https://issues.jboss.org/browse/JBTM-1584
13008	JBTM-1583	How about using rest rather than SOAP?	0	0	https://issues.jboss.org/browse/JBTM-1583
13009	JBTM-1576	So is this replacing sockets with JNI calls - I would have expected sockets (domain sockets on unix) to be faster than JNI. Have you considered java NIO memory mapped files which will be faster than both JNI and sockets.	0	0	https://issues.jboss.org/browse/JBTM-1576
13010	JBTM-1576	Hi Mike, There are two considerations in this: 1. Performance - Amos, please review Mikes comment "Have you considered java NIO memory mapped files which will be faster than both JNI and sockets." 2. Administration - deploying into the AS container for admin seems like it would be easier than managing standalone C processes and an AS? In terms of performance, I am considering the (unlikely, but ramp it down a bit and it might be more likely) case of someone trying to send a 4GB .iso to a service. Would a socket be quicker than JNI call (or your NIO mapped files - out of interest, can these be accessed from C?)? I talked to Andrew Dinn and he thought (maybe my question wasn't phrased well enough to him) that JNI would be faster than socket (we didn't consider NIO). Tom	1	0	https://issues.jboss.org/browse/JBTM-1576
13011	JBTM-1576	Please can you add an estimate for this issue?	1	0	https://issues.jboss.org/browse/JBTM-1576
13012	JBTM-1573	In HybridSessionImpl::send we call apr_socket_timeout_set(stompConnection->socket, 1); and abandon the send if we don't get the message receipt in time. Where in the code are we doing a blocking wait when the transaction timeout is set?	1	1	https://issues.jboss.org/browse/JBTM-1573
13013	JBTM-1568	Please can you add an estimate for this issue?	1	1	https://issues.jboss.org/browse/JBTM-1568

#	Requirement	Questions	Answered	Successfully Answered	URL
13014	JBTM-1568	do you know the jbossas jira for this way to deploy a war as module ? does it need more configuration ? just for merge stempconnect and blacktie-admin-servers ear the estimate is 4h	1	1	https://issues.jboss.org/browse/JBTM-1568
13015	JBTM-1568	I can not access ncl openvpn as my crt has been expired. can you mail me these txt ?	1	1	https://issues.jboss.org/browse/JBTM-1568
13016	JBTM-1568	Would it be possible for you to push the fix directly without the need for a PR (since it currently breaks all of our builds). Also why aren't we using the wildfly version that is built during the main narayana build.	1	1	https://issues.jboss.org/browse/JBTM-1568
13017	JBTM-1565	Hi Crispin, Please can you: 1. Sign our CLA: https://cla.jboss.org/index.seam (choose JBoss Transactions rather than BlackTie) 2. Raise a pull request for your change? This will then run CI over the proposed enhancement and double check it for us I will then merge the PR after it passes :) Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-1565
13018	JBTM-1558	Can you provide more detail to this task, since I thought we already had an SPI?	1	0	https://issues.jboss.org/browse/JBTM-1558
13019	JBTM-1558	There are a few SPI related issues at the moment. Maybe cross-link them?	0	0	https://issues.jboss.org/browse/JBTM-1558
13020	JBTM-1556	What do you mean by a remote thread? We should be using interposition in all distributed cases, which means the Synchronisation will be registered with a subordinate coordinator.	0	0	https://issues.jboss.org/browse/JBTM-1556
13021	JBTM-1556	Regarding the question about what I meant by the "remote thread" case: {quote} <tomjenkinson> smarlow: AS1:EJB1 enlists JPA(XAR1) then calls AS2:EJB2 over JTS <tomjenkinson> smarlow: AS2:EJB2 enlists JPA(XAR2) this is done by a jacorb request processor thread at random <tomjenkinson> smarlow: AS1:EJB1 calls commit on the transaction, JPA(XAR1) afterCompletion is handled in the same thread as before - this is fine - but then the call to AS2 to commit the transaction goes over a different jacorb request processor, now afterCompletion in AS2 is done by a different thread {quote} The above doesn't raise any concurrency concerns since there is no (active) application thread on AS2 at the time that AS2 invokes the Synchronization.afterCompletion. Another case, would be similar to the above but with one difference. AS1:EJB1 enlists JPA(XAR1) then loops on calling AS2:EJB2 over JTS AS2:EJB2 enlists JPA(XAR2) and does processing and returns Looping continues until the transaction times out while current application execution is in AS2:EJB2 or just in AS1:EJB1. I don't think that only checking the reaper thread name covers when the transaction times out when the application thread of execution, is in AS2:EJB2. Question 1: Does the root coordinator (AS1) reaper always handle transaction timeout or can the subordinate coordinator (AS2) reaper also handle transaction timeout? I'm not sure how the subordinate coordinator would handle timeout (maybe by sending an out of band message to its parent coordinator that propagates up to the root coordinator?) Even so, it would be chaotic to have multiple coordinators cancelling the transaction at the same time. Regarding the idea of having globally searchable "being cancelled" (I assume you mean by the reaper) list, sounds better than my current "reaper thread" name check. The globally searchable list would also cover the second (above) case when the application execution is in AS2:EJB2 while (concurrently) an AS2 background orb thread sees the cancel message from the root coordinator.	1	1	https://issues.jboss.org/browse/JBTM-1556
13022	JBTM-1556	When each coordinator handles the timeout, does it also cancel its subordinates (as if the subordinate is one of the resources)? Or is it more like the coordinator times out the subordinate (as a resource) which means blocking until the subordinate has completed the subordinate level (reaper) timing out of its (subordinate) resources? In other words, could the subordinate coordinator Synchronization.afterCompletion() be called from a transport thread ever or only from a reaper thread?	0	0	https://issues.jboss.org/browse/JBTM-1556
13023	JBTM-1526	Please can you itemize the dependencies that you are worried about - thanks	1	1	https://issues.jboss.org/browse/JBTM-1526
13024	JBTM-1526	Hi Weinan, Please can you provide an update on which dependencies you are worried about? We might be cutting a 4.17.4 soon and would like to hear from you, thanks! Tom	1	1	https://issues.jboss.org/browse/JBTM-1526
13025	JBTM-1526	Hi Weinan, Please can you raise a Jira for the second issue (orb.idl) as we don't see that on our machines, it might be environmental so let us know the OS etc and assign it [~mmusgrov] Please can you raise a pull request for the first one? Thanks! Tom	1	1	https://issues.jboss.org/browse/JBTM-1526
13026	JBTM-1503	Just wondered what had necessitated this? Not that it is a problem but we should record in the Jira why this decision was made.	1	1	https://issues.jboss.org/browse/JBTM-1503
13027	JBTM-1485	Does this change represent a possible backward incompatibility?	1	0	https://issues.jboss.org/browse/JBTM-1485

#	Requirement	Questions	Answered	Successfully Answered	URL
13028	JBTM-1482	version of ArjunaJTA/jta/tests/../JTATest.java to throw a runtime in afterCompletion and get this: Sep 19, 2013 2:19:27 PM com.arjuna.ats.internal.jta.resources.arjunacore.SynchronizationImple afterCompletion WARN: ARJUNA016029: SynchronizationImple.afterCompletion - failed for com.hp.mwtests.ts.jta.common.Synchronization@483ee0b7 with exception java.lang.RuntimeException at com.hp.mwtests.ts.jta.common.Synchronization.afterCompletion(Synchronization.java:72) at com.arjuna.ats.internal.jta.resources.arjunacore.SynchronizationImple.afterCompletion(SynchronizationImple.java:96) at com.arjuna.ats.arjuna.coordinator.TwoPhaseCoordinator.afterCompletion(TwoPhaseCoordinator.java:532) at com.arjuna.ats.arjuna.coordinator.TwoPhaseCoordinator.end(TwoPhaseCoordinator.java:101) at com.arjuna.ats.arjuna.AtomicAction.commit(AtomicAction.java:162) at com.arjuna.ats.internal.jta.transaction.arjunacore.TransactionImple.commitAndDisassociate(TransactionImple.java:1170) at com.arjuna.ats.internal.jta.transaction.arjunacore.BaseTransaction.commit(BaseTransaction.java:126) at com.hp.mwtests.ts.jta.basic.JTATest.test(JTATest.java:102) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:47) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:50) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:222) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.eclipse.jdt.internal.junit4.runner.JUnit4TestReference.run(JUnit4TestReference.java:50) at org.eclipse.jdt.internal.junit.runner.TestExecution.run(TestExecution.java:38) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.runTests(RemoteTestRunner.java:467) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.runTests(RemoteTestRunner.java:683) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.run(RemoteTestRunner.java:390) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.main(RemoteTestRunner.java:197) Sep 19, 2013 2:19:33 PM com.arjuna.ats.arjuna.coordinator.TwoPhaseCoordinator afterCompletion WARN: ARJUNA012127: TwoPhaseCoordinator.afterCompletion - returned failure for SynchronizationImple< 0:ffff7f000001:d16e:523af9d9:9, com.hp.mwtests.ts.jta.common.Synchronization@483ee0b7 > Strangely looking at git blame I can't find	1	1	https://issues.jboss.org/browse/JBTM-1482
13029	JBTM-1479	It appears that the commit for JBJCA-1011 made it into IronJacamar 1.0.16.Final. Is the quickstart working now? Where do I find it?	1	1	https://issues.jboss.org/browse/JBTM-1479
13030	JBTM-1479	I've tried this simple test on windows and it also failed with "java.io.IOException: Could not delete PATH_TO_TEMP_DIRECTORY/ironjacamar-jdbc.jar" {code} public class SimpleTest { private static final String JDBC_RAR_FILE_PATH = "src/main/resources/jdbc-xa.rar"; private static Embedded EMBEDDED; @Test public void test() throws Throwable { EMBEDDED = EmbeddedFactory.create(); EMBEDDED.startup(); EMBEDDED.deploy(getURL(JDBC_RAR_FILE_PATH)); System.out.println("Test..."); EMBEDDED.undeploy(getURL(JDBC_RAR_FILE_PATH)); EMBEDDED.shutdown(); } private static URL getURL(final String path) throws MalformedURLException { final File f = new File(path); return f.toURI().toURL(); } } {code}	1	1	https://issues.jboss.org/browse/JBTM-1479
13031	JBTM-1479	For the record, this is the failure: {code} --- Test set: org.jboss.narayana.quickstart.jca.model.CustomerDAOTest --- Tests run: 10, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 2.32 sec <<< FAILURE! org.jboss.narayana.quickstart.jca.model.CustomerDAOTest Time elapsed: 0.063 sec <<< ERROR! java.io.IOException: Could not delete C:\Users\hudson\AppData\Local\Temp\2\iron.jacamar\tmp\jdbc-xa.rar\ironjacamar-jdbc.jar at com.github.fungal.impl.KernelImpl.recursiveDelete(KernelImpl.java:1520) at com.github.fungal.impl.KernelImpl.recursiveDelete(KernelImpl.java:1515) at com.github.fungal.impl.KernelImpl.recursiveDelete(KernelImpl.java:1515) at com.github.fungal.impl.KernelImpl.shutdown(KernelImpl.java:850) at org.jboss.jca.embedded.EmbeddedJCA.shutdown(EmbeddedJCA.java:158) at org.jboss.narayana.quickstart.jca.common.AbstractTest.afterClass(AbstractTest.java:61) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:36) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.apache.maven.surefire.junit4.JUnit4TestSet.execute(JUnit4TestSet.java:35) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:146) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:97) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.apache.maven.surefire.booter.ProviderFactory$ClassLoaderProxy.invoke(ProviderFactory.java:103) at $Proxy0.invoke(Unknown Source) at org.apache.maven.surefire.booter.SurefireStarter.invokeProvider(SurefireStarter.java:145) at org.apache.maven.surefire.booter.SurefireStarter.runSuitesInProcess(SurefireStarter.java:87) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:69) {code}	1	1	https://issues.jboss.org/browse/JBTM-1479

#	Requirement	Questions	Answered	Successfully Answered	URL
13032	JBTM-1479	I think that might be enough information to log a bug report against IronJacamar? Can you link this bug as dependent on the one you raise? Also can you provide the link to the stackoverflow article just in case. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-1479
13033	JBTM-1469	I think we should just take one of the existing quickstarts from the Narayana quickstart repo (multi-hop or multi-service) and simplify it for the following scenario: 1. Client begins TX 2. Client updates a DB 3. Client invokes service 4. Service updates another DB 5. Client commits TX Where client and Service are both EJBs.	1	0	https://issues.jboss.org/browse/JBTM-1469
13034	JBTM-1463	you are right in your concern and that is what I am trying to address but ineloquently I was thinking more about changes to narayana that need a change to jbosstm somehow being able to test both pulls at once. Maybe we could somehow put convention when you raise a pull request for jbosstm that you can put a comment to say something like: "requires: https://github.com/jbosstm/jboss-as/pull/217" We can then somehow use that info in the narayana.sh to pull in that branch rather than 5_BRANCH, what do you think? I know you did some work with using rest to talk to git hub to comment on pulls, could it possibly download the comment for the pull in order to grep it for that info?	1	1	https://issues.jboss.org/browse/JBTM-1463
13035	JBTM-1463	Please can you add an estimate for this?	0	0	https://issues.jboss.org/browse/JBTM-1463
13036	JBTM-1454	Why? Also, will you be testing HTTP versus IIOP initially to get a base?	1	0	https://issues.jboss.org/browse/JBTM-1454
13037	JBTM-1424	Please can you add a component for this?	1	1	https://issues.jboss.org/browse/JBTM-1424
13038	JBTM-1424	Shouldn't this also be fixed for 4.17?	0	0	https://issues.jboss.org/browse/JBTM-1424
13039	JBTM-1423	Please can you add a component for this?	0	0	https://issues.jboss.org/browse/JBTM-1423
13040	JBTM-1412	I think the test didn't have long enough to crash and be detected, it was set for 5 seconds but the fact your build of JTA takes an hour tells me that node is sloooow ;) I have given it a minute now (5 seconds x 20 attempts), hopefully that will be enough. The other failures are cascading off that one but I fixed that now too (was using @AfterClass rather than @After - d'oh) If you want to test this, the following will get you to my code and run the integration unit tests: git remote add tomjenkinson git://github.com/tomjenkinson/jbosstm.git git fetch tomjenkinson git checkout -b JBTM-1412 tomjenkinson/JBTM-1412 ./build.sh -f ArjunaJTS/integration/pom.xml (Assuming you are relatively current with master, if you are quite out of date it will be worth a full build)	1	1	https://issues.jboss.org/browse/JBTM-1412
13041	JBTM-1407	It should be trivial to do that *today* for JDK orb with profiles as you can activate a profile by the presence of a file, e.g. JAVA_HOME/lib/orb.idl. However, JacORB would be in the classpath unless you specifically disabled it with a different profile so its presence would be guaranteed. Now, in the future when AS7-5582 is implemented, the "JDK" orb will be a maven dependency (same as JacORB), i.e. it will be available by default in the classpath. Lets say we made the change today as a convenience until AS7-5582 is available, the question is do we really want to have unpredictable results in the build. Lets say the user attempts this new default, detection based build but does still expect both a JDK orb and JacORB version but we as can't locate orb.idl (say they are on pi) and just build the JacORB version - what would you expect then, log a warning?	1	1	https://issues.jboss.org/browse/JBTM-1407
13042	JBTM-1407	Since I'm still receiving the following using the latest JDK 7 build for OSX ... [ERROR] Failed to execute goal org.apache.maven.plugins:maven-surefire-plugin:2.9:test (surefire-idlj) on project jts: Execution surefire-idlj of goal org.apache.maven.plugins:maven-surefire-plugin:2.9:test failed: java.lang.reflect.InvocationTargetException; nested exception is java.lang.reflect.InvocationTargetException: null: com/arjuna/ArjunaOTS/OTSAbstractRecordPOA: com.arjuna.ArjunaOTS.OTSAbstractRecordPOA -> [Help 1] I'm assuming we haven't made the automatic build the default? Building for both ORBs is something that has only just happened. So users currently really don't expect an IDLJ build. However, as a developer what I'd prefer is for the build system to not cr*p out ;-) if it can only find one ORB. If it can't find both and I asked it to build JTS then that's definitely an error. But if I just specify JTS on the build line I think I'm either meaning "Build JTS with JacORB" (given the last 7 years of history) or "Build JTS and I really don't care which ORB you use". Therefore, I suggest changing the build so that it only fails if it can't find either ORB. If it can't find both ORBs then at the end it should print out a message saying which one failed to build. Plus add a jacorb option to the build to correspond to the idlj option, so someone can specify "build jacorb only". This would mean the following: ./build.sh -P jts --> build with any ORB, I really don't care which ./build.sh -P jts idlj --> build with only the IDLJ ORB ./build.sh -P jts jacorb --> build with only JacORB	0	0	https://issues.jboss.org/browse/JBTM-1407
13043	JBTM-1398	Should org.jboss.as.weld.deployment.WeldAttachments be lifted into its own subsystem that doesn't depend on EJB3?	1	1	https://issues.jboss.org/browse/JBTM-1398
13044	JBTM-1398	I think this is resolved for now. Do you agree Tom?	0	0	https://issues.jboss.org/browse/JBTM-1398
13045	JBTM-1382	Hi Mike, THANKS for raising this, I keep forgetting! Can we put the output in testoutput/ that way we can have a git ignore on just testoutput rather than testoutput TEST*.txt. Tom	1	1	https://issues.jboss.org/browse/JBTM-1382
13046	JBTM-1382	The JIRA asks for output to be sent to qa/testoutput (not qa). For a qa run against both orbs if you look in the testoutput directory you will find two zips (testoutput-jacorb.zip and testoutput-idlj.zip) which contain the testoutput directory and the TEST*txt files corresponding to each run. To see it working look at the latest narayana-dualstack job (I did notice that btny-pulls-narayana CI job does not archive testoutput so I have updated that job config). I think I may have misunderstood your requirement. Did you mean move all the TEST*txt files into testoutput but leave the zip files in the qa directory.	1	1	https://issues.jboss.org/browse/JBTM-1382
13047	JBTM-1369	Can we have an update on this?	1	1	https://issues.jboss.org/browse/JBTM-1369

#	Requirement	Questions	Answered	Successfully Answered	URL
13048	JBTM-1369	We have four JTS performance unit tests - they all run 1000 empty transactions with each test using a different method to begin the txn. I got the following figures from our CI cluster but if you want to try them on your own machine do the following: Build narayana with JDK orb support enabled: {code}./build.sh -Didlj-enabled=true clean install -DskipTests{code} To run just a specific test, such as Performance2 for example: {code}./build.sh -f ArjunaJTS/jts/pom.xml -Didlj-enabled=true -Dtest=com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance2{code} View the results in the surefire report directories: {code}JacORB: ArjunaJTS/jts/target/surefire-reports/com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance1-output.txt JDK ORB: ArjunaJTS/jts/target/idlj-surefire-reports/com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance1-output.txt{code} The results generally show jacorb performance as being significantly better on our CI cluster but, interestingly, on my laptop JDK orb seemed to do better: {code:title=com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance1} 1000 txns (OTSImpleManager.current), no participants (with warm up of 10 txns): idlj Average time for empty transaction = 1.284000039100647 Transactions per second = 778.8161756602676 jacorb Average time for empty transaction = 1.3229999542236328 Transactions per second = 755.8579248680498 {code} {code:title=com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance2} 1000 txns (TSImpleManager.factory().create()), no participants (with warm up of 10 txns): idlj Average time for empty transaction = 3.48799991607666 Transactions per second = 286.6972546045273 jacorb Average time for empty transaction = 2.115999937057495 Transactions per second = 472.5898061181409 {code} {code:title=Test: com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance3} 1000 txns (ArjunaTransactionImple), no participants (with warm up of 10 txns): idlj Average time for empty transaction = 0.05400000140070915 Transactions per second = 18518.5180381656 jacorb Average time for empty transaction = 0.052000001072883606 Transactions per second = 19230.76883399276 {code} {code:title=Test: com.hp.mwtests.ts.jts.local.synchronizations.Performance} 1000 txns (org.omg.CosTransactions.Current) each registering a single synchronization idlj TPS: 191.20458891013382 jacorb TPS: 309.2145949288807 {code}	1	1	https://issues.jboss.org/browse/JBTM-1369
13049	JBTM-1369	Can you paste in the commands you used to get those numbers, please?	1	1	https://issues.jboss.org/browse/JBTM-1369
13050	JBTM-1369	What version of JacORB was used please?	1	1	https://issues.jboss.org/browse/JBTM-1369
13051	JBTM-1369	What hardware architecture is your CI cluster?	1	1	https://issues.jboss.org/browse/JBTM-1369
13052	JBTM-1359	Why has this been linked to the JDBC object store work? I'd like to understand the relative priority of this to other tasks.	1	0	https://issues.jboss.org/browse/JBTM-1359
13053	JBTM-1359	Thanks, but it still doesn't seem that there is clarity on what exactly is required. What is the problem that we're trying to fix? Is it only for JTA? Until we have answers to these and other questions there's no point in implementing anything.	0	0	https://issues.jboss.org/browse/JBTM-1359
13054	JBTM-1357	It's worth pointing out that the script is inherited from AS7's build.sh, do you think you could try to raise this "upstream" so to speak?	1	0	https://issues.jboss.org/browse/JBTM-1357
13055	JBTM-1357	I've been using it for a wile on my machine and I think it's currently too colourful to be adopted across the board. I think we just want to colour: 1. Test failures red 2. overall build success as green 3. overall build failure as red 4. ERROR as red 5. WARN as orange The output should then be mostly default colour with the important stuff highlighted.	0	0	https://issues.jboss.org/browse/JBTM-1357
13056	JBTM-1340	Hi Mark, Please can you upload the test you were using and the name of the tool that outputted the profiling information? We can look at tuning TxStats, I am thinking a listener pattern might help? Before that though, the test and profiler settings would be useful. Thanks, Tom	1	1	https://issues.jboss.org/browse/JBTM-1340
13057	JBTM-1339	When the linked JIRA is ready make sure JTS functions with the JDK orb	1	1	https://issues.jboss.org/browse/JBTM-1339
13058	JBTM-1339	Isn't this out of date now?	0	0	https://issues.jboss.org/browse/JBTM-1339
13059	JBTM-1339	OK, thanks. Does the title of the issue need changing now though? I'm assuming we're not back porting this to AS7 and only focussing on WildFly.	0	0	https://issues.jboss.org/browse/JBTM-1339
13060	JBTM-1338	How do I get access to BaseTransaction in AS7? As JNDI lookup (java:jboss/UserTransaction) returns ServerVMClientUserTransaction.	1	1	https://issues.jboss.org/browse/JBTM-1338
13061	JBTM-1328	Was this a one-off?	1	1	https://issues.jboss.org/browse/JBTM-1328
13062	JBTM-1328	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	1	0	https://issues.jboss.org/browse/JBTM-1328
13063	JBTM-1310	Does this affect all versions or 4.17.1+ in particular?	1	0	https://issues.jboss.org/browse/JBTM-1310
13064	JBTM-1254	Mike, Does that mean that you agree with me that we should resolve this as "won't fix"?	1	0	https://issues.jboss.org/browse/JBTM-1254
13065	JBTM-1251	You don't need to branch the quickstart, narayana, or docs branches to create these jobs. Currently 5.x and 4.17 share a common branch (master) so both sets of jobs build against master. The place in which they differ is that 4.17 builds against 4_BRANCH of our jboss-as fork and 5.x against 5_BRANCH. They also differ in what is tested as TXF is only tested in 5.x. I would suggest we create the jobs now, but the problem is that we are going to double the load on Hudson as any change is going to trigger both the EAP61 and the narayana-500 version of the job. The reason why I'm keen to get some narayana 5.x CI is that TXF is currently not getting any CI. Maybe I could just create some TXF jobs in the mean time until this is resolved?	1	1	https://issues.jboss.org/browse/JBTM-1251
13066	JBTM-1247	What is the use case for this?	0	0	https://issues.jboss.org/browse/JBTM-1247
13067	JBTM-1227	Did this make it into PropertiesFactory?	1	1	https://issues.jboss.org/browse/JBTM-1227
13068	JBTM-1217	Seems to be related to REM3-158. Ivo, I'm on PTO until August 28th. Can you try rebuilding AS7 with XNIO 3.0.6 and then re-testing on the IPv6 node?	1	1	https://issues.jboss.org/browse/JBTM-1217
13069	JBTM-1212	Amos, Has this been done for the docs yet?	0	0	https://issues.jboss.org/browse/JBTM-1212
13070	JBTM-1212	Can we mark this as resolved now?	0	0	https://issues.jboss.org/browse/JBTM-1212

#	Requirement	Questions	Answered	Successfully Answered	URL
13071	JBTM-1199	The only way I have found to check an old version of the AS (but still having our commits from 5_BRANCH) is as follows: 0. cd jboss-as 1. git checkout 5_BRANCH (assumes you haven't checked this out yet, if you already have this such as by having ran scripts/hudson/narayana.sh git checkout master git branch -D 5-BRANCH git checkout 5_BRANCH) 2. git pull --rebase --ff-only upstream master (to get up to date) 3. Go to: https://github.com/jbossas/jboss-as/commits/master and find a commit you want to check against, e.g one from July 3rd when build was still working: https://github.com/jbossas/jboss-as/commit/55fca1401f399591d009c211205f2e6ef94f876b 4. Use this in a git rebase: git rebase -i 55fca1401f399591d009c211205f2e6ef94f876b^ (you need the ^ to include that commit in the rebase) 5. delete all the commits after that commit (dd in vi) but leave our commits in at the bottom of the rebase 6. You will get a rebase issue on the first one, simply open pom.xml and remove the jbossjts part of pom.xml which came from HEAD -- grep in pom.xml for an "=" sign), then: git add pom.xml git rebase --continue 7. You should now be able to build the AS as per that commit but with our commits on top MAVEN_OPTS=-XX:MaxPermSize=256m ./build.sh "$@" clean install -DskipTests 8. cd ../ 9. export JBOSS_HOME=${WORKSPACE}/jboss-as/build/target/`ls -1 ${WORKSPACE}/jboss-as/build/target \| grep jboss-as` 10. ./build.sh -f XTS/localjunit/pom.xml --projects xtstest,crash-recovery-tests -Parq "$@" clean install -Dtest=TestATCrashDuringCommit#MultiParticipantPrepareAndCommitTest After you have verified that that version works, try a later revision of AS, however you will need to delete your local branch first! This will be tedious, so to reduce the number of checks, the first thing to do is identify the day where a commit was added that broke our build, then use a binary tree search style to cut down the options on that day. Suggested commits to try: https://github.com/jbossas/jboss-as/commit/2e2f1a7dfcebebb282987d2d8e2df4407de036ac https://github.com/jbossas/jboss-as/commit/f2d65c311d67c61d953bd17c079f4169c3dbf8aa Tom	1	1	https://issues.jboss.org/browse/JBTM-1199
13072	JBTM-1199	I can confirm that the the build started to break was between the following two jboss-as commits. I know this because: FAILED = https://github.com/jbossas/jboss-as/commit/f2d65c311d67c61d953bd17c079f4169c3dbf8aa PASSED = https://github.com/jbossas/jboss-as/commit/55fca1401f399591d009c211205f2e6ef94f876b Please can you start checking the commits inbetween to see which specific commit broke the build?	1	1	https://issues.jboss.org/browse/JBTM-1199
13073	JBTM-1199	Amos, please can you investigate why https://github.com/jbossas/jboss-as/commit/435c862b9514bfcf0a036bb2b916ac3dc63b4f4a would stop the XTS crash recovery tests from working? I noticed somethings to do with "Reload" perhaps AS7 has some behaviour that reloads the server automatically now?	1	1	https://issues.jboss.org/browse/JBTM-1199
13074	JBTM-1199	I added a comment onto that commit, please enable notifications for updates to that commit, my comment is https://github.com/jbossas/jboss-as/commit/435c862b9514bfcf0a036bb2b916ac3dc63b4f4a#commitcomment-1579773, you can enable updates by going to the bottom of the page and searching for: "Email notifications for these commit notes"	1	1	https://issues.jboss.org/browse/JBTM-1199
13075	JBTM-1192	Can resolve when this: http://172.17.131.2/job/btny-pulls-quickstarts/31/ is seen testing TXF.	0	0	https://issues.jboss.org/browse/JBTM-1192
13076	JBTM-1188	Where did the requirement for a per-tx override come from? Is the existing global config option not sufficient?	1	0	https://issues.jboss.org/browse/JBTM-1188
13077	JBTM-1187	Amos, What made you think "the client handler chain can be added by txframework automatically"? This is a feature we have planned but we don't have it yet. Did you read it in some docs somewhere?	1	1	https://issues.jboss.org/browse/JBTM-1187
13078	JBTM-1182	With recent changes to transactional fileio repo, is this still relevant? Also what about the AS7 RPMs that have already been produced for Fedora: don't they duplicate this?	1	0	https://issues.jboss.org/browse/JBTM-1182
13079	JBTM-1173	How to run the tests: # ipv4: {noformat} mvn clean test {noformat} # ipv6: {noformat} mvn clean test -Dipv6 -Parq-jbossas-managed-7 {noformat} Note that for IPv6 the XTS demo app needs to be compiled with HOSTNAME set to [::1] in jboss.properties file.	1	1	https://issues.jboss.org/browse/JBTM-1173
13080	JBTM-1173	I won't be able to test this before Weds. If it's urgent, could someone else test the pull? It sounds like we have many other IPv6 (JBTM-1156) issues that will be pushed to 5.final, so we could push this too and I'll look at it and JBTM-1156 on Weds when I hope to be back to work.	1	1	https://issues.jboss.org/browse/JBTM-1173
13081	JBTM-1166	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	1	0	https://issues.jboss.org/browse/JBTM-1166
13082	JBTM-1156	Great, are you working with Paul on resolving the XTS issues? +1 on the JDBC tests	1	1	https://issues.jboss.org/browse/JBTM-1156
13083	JBTM-1147	Please can you add an estimate for this issue?	0	0	https://issues.jboss.org/browse/JBTM-1147
13084	JBTM-1116	Good points, lets rework this Jira as an opportunity to simply: 1. Ensure the Javadocs in isolation are still correct, fully explained 2. Ensure the documentation explains the architecture of our configuration mechanism 3. Ensure the documentation builds on point 2 to reference the relevant beans when configuring common deployment scenarios, e.g. which properties do you need to enable hornetq/JDBC object store	1	1	https://issues.jboss.org/browse/JBTM-1116
13085	JBTM-1082	As workaround is possible to define java opt argument -XX:-UseSplitVerifier Put it to file sharedbuild.xml to the tag <junit> like: <jvmarg value="-XX:-UseSplitVerifier" /> I suppose that this problem is caused by the fact that the emma library was compiled with older version of compiler that we are trying to use for compiling TS sources. Some links e.g. here http://stackoverflow.com/questions/7970622/java-7-jvm-verifyerror-in-eclipse Do you think that the workaround is enough for solving this jira or should be tried to recompile emma or dive to some other investigation?	1	1	https://issues.jboss.org/browse/JBTM-1082
13086	JBTM-1073	Worth a link to where this was done?	1	0	https://issues.jboss.org/browse/JBTM-1073
13087	JBTM-1072	I commented on the pull re the password in narayana.sh, not ideal, can it be externalized into the job config? Also, I commented to say what happens if you are running this locally on a checked out pull, can we make the default that the commenting is disabled and job config enables it?	1	1	https://issues.jboss.org/browse/JBTM-1072
13088	JBTM-1072	Tom, btny-pulls-blacktie and btny-pulls-quickstarts are tied to Windows2008. btny-pulls-documentation will run on any machine (including windows). It's not going to be trivial to get this working on windows: {code} function comment_on_pull { if ["$COMMENT_ON_PULL" = ""]; then return; fi PULL_NUMBER=$(echo $GIT_BRANCH | awk -F 'pull' '{ print $2 }' | awk -F '/' '{ print $2 }') if ["$PULL_NUMBER" != ""] then JSON="{ \"body\": \"$1\" }" curl -d "$JSON" -ujbosstm-bot:$BOT_PASSWORD https://api.github.com/repos/$GIT_ACCOUNT/$GIT_REPO/issues/$PULL_NUMBER/comments else echo "Not a pull request, so not commenting" fi } {code} I guess the only way to do this is by having the job run cygwin with a bash script that does the above. I don't fancy finding Windows commands for all of the above! WDYT?	1	1	https://issues.jboss.org/browse/JBTM-1072
13089	JBTM-1063	Hi Amos, Please can you backport this to the branch? Thanks! Tom	1	1	https://issues.jboss.org/browse/JBTM-1063
13090	JBTM-1063	Amos, {quote} and now jboss-as version is 7.1.2.SNAPSHOT and jbossts version is 4.16.2.Final {quote} What do you mean by this?	1	1	https://issues.jboss.org/browse/JBTM-1063

#	Requirement	Questions	Answered	Successfully Answered	URL
13091	JBTM-1063	Hi, I have updated the config on the jobs to use this new script thanks Amos! Paul, please can you give Amos a pointer on the xts commands? Tom	1	1	https://issues.jboss.org/browse/JBTM-1063
13092	JBTM-1063	Paul, I checkout 5_BRANCH on jboss-as and build fail with following: [ERROR] Failed to execute goal on project jboss-as-transactions: Could not resolve dependencies for project org.jboss.as:jboss-as-transactions:jar:7.1.1.Final: The following artifacts could not be resolved: org.jboss.jbossts:jbossjts:jar:5.0.0.M2-SNAPSHOT, org.jboss.jbossts:jbossjts-integration:jar:5.0.0.M2-SNAPSHOT: Failure to find org.jboss.jbossts:jbossjts:jar:5.0.0.M2-SNAPSHOT in http://repository.jboss.org/nexus/content/groups/public/ was cached in the local repository, resolution will not be reattempted until the update interval of jboss-public-repository-group has elapsed or updates are forced It looks like we should change groupId from "org.jboss.jbossts" to "org.jboss.narayana" can you confirm this change? Amos	1	1	https://issues.jboss.org/browse/JBTM-1063
13093	JBTM-1063	Amos, Can you give us an update on when this will be ported to the 4.16 branch?	1	1	https://issues.jboss.org/browse/JBTM-1063
13094	JBTM-1063	Paul, I ran these test commands on my machine and it failed with <testcase classname="com.arjuna.qa.junit.LocalJUnitRunner" name="testCallServlet" time="0.01"> <error message="org/apache/commons/logging/LogFactory" type="java.lang.NoClassDefFoundError">java.lang.NoClassDefFoundError: org/apache/commons/logging/LogFactory at org.apache.commons.httpclient.HttpClient.<clinit>(Unknown Source) at com.arjuna.qa.junit.HttpUtils.accessURL(HttpUtils.java:97) at com.arjuna.qa.junit.LocalJUnitRunner.testCallServlet(LocalJUnitRunner.java:40) Caused by: java.lang.ClassNotFoundException: org.apache.commons.logging.LogFactory at java.net.URLClassLoader$1.run(URLClassLoader.java:202) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:247) </error> </testcase> did I miss some jars ?	0	0	https://issues.jboss.org/browse/JBTM-1063
13095	JBTM-1063	I think that was fixed by https://issues.jboss.org/browse/JBTM-1086 Have you done an 'svn up' on the 4.16 branch?	1	1	https://issues.jboss.org/browse/JBTM-1063
13096	JBTM-1063	how about commands on the txbridge tests?	0	0	https://issues.jboss.org/browse/JBTM-1063
13097	JBTM-1061	I just put this on HORNETQ-274: Hi I just trawled the the Jira and forum postings I ls'd my .m2/repository/org/hornetq find . -name \pom.xml in the AS7 version of hornetq but can't find the dependency to import for this. Do you know that this work was done? If so, do you happen to know what I should replace: <dependency> <groupId>org.hornetq</groupId> <artifactId>hornetq-core</artifactId> </dependency> with? I also can't find it by grepping the internet, e.g. http://mvnrepository.com/artifact/org.hornetq	1	1	https://issues.jboss.org/browse/JBTM-1061
13098	JBTM-1052	Is this still required? We could do with a better description to explain the actual motivation or leave this closed as "incomplete description".	0	0	https://issues.jboss.org/browse/JBTM-1052
13099	JBTM-1051	A better solution would be to have this problem detected in the AS7 testsuite. This would prevent the error from even being accepted into the AS7 codebase. Ivo, can this type of check be made in the allTests junit suite?	1	1	https://issues.jboss.org/browse/JBTM-1051
13100	JBTM-1051	I would add that I believe this issue can be (not formally, but effectively) resolved through education. Who ever makes a change to src/main/resources/standalone/configuration/*xml should be made aware that they should also update src/main/resources/docs/examples/config/*xml This issue should affect: ./messaging/src/test/resources/standalone-with-messaging.xml ./modcluster/src/test/resources/standalonemodcluster.xml ./build/src/main/resources/docs/examples/configs/standalone-hornetq-colocated.xml ./build/src/main/resources/docs/examples/configs/standalone-xts.xml ./build/src/main/resources/docs/examples/configs/standalone-minimalistic.xml ./build/src/main/resources/docs/examples/configs/standalone-jts.xml ./build/src/main/resources/docs/examples/configs/standalone-osgi-only.xml So I see it as broader than XTS. That said, +1 on checking whether the testsuite can detect it, but I don't know if the other configs are tested during testsuite.	0	0	https://issues.jboss.org/browse/JBTM-1051
13101	JBTM-1051	PS, shouldn't this issue be moved to an AS7 Jira?	0	0	https://issues.jboss.org/browse/JBTM-1051
13102	JBTM-1051	I don't even know where standalone-jts.xml is referenced from :) In my quickstarts I show people how to change the config manually. I wonder who added the jts one? Maybe that is an option for XTS? Explain the manual configuration?	0	0	https://issues.jboss.org/browse/JBTM-1051
13103	JBTM-1051	Ivo, Of course, I forgot these tests would do this also. Do your tests fail if the server boots with errors, even though XTS is still functioning correctly? We still have the problem where a minor change is forgotten; one which doesn't break XTS and doesn't fail the server boot. However, these are relatively minor and could probably be checked manually every now and then.	1	1	https://issues.jboss.org/browse/JBTM-1051
13104	JBTM-1051	Can this be moved to an AS7 Jira?	1	1	https://issues.jboss.org/browse/JBTM-1051
13105	JBTM-1031	I've just added 4.17.4 to the fix version. Can you cherry-pick your commit from master and raise a pull request on 4.17?	0	0	https://issues.jboss.org/browse/JBTM-1031
13106	JBTM-1016	I suggest we add a dependency on BYTEMAN-190 for this task. Amos, if you agree, can you add the dependency please? More Actions -> Link -> Depends on	1	1	https://issues.jboss.org/browse/JBTM-1016
13107	JBTM-1016	Please can you add an estimate for this issue?	1	1	https://issues.jboss.org/browse/JBTM-1016
13108	JBTM-1016	Paul, can I delete XTS/localjunit/crash-recovery-tests/check.sh as now we can use maven plugin to check rule ?	1	1	https://issues.jboss.org/browse/JBTM-1016

#	Requirement	Questions	Answered	Successfully Answered	URL
13109	JBTM-999	In case the fpaste expires: {code} package org.jboss.wsf.stack.cxf.interceptor; import javax.management.ObjectName; import javax.xml.ws.WebServiceContext; import javax.xml.ws.handler.MessageContext; import org.apache.cxf.jaxws.context.WrappedMessageContext; import org.jboss.wsf.spi.SPIProviderResolver; import org.jboss.wsf.spi.classloading.ClassLoaderProvider; import org.jboss.wsf.spi.deployment.Endpoint; import org.jboss.wsf.spi.deployment.Reference; import org.jboss.wsf.spi.invocation.EndpointAssociation; import org.jboss.wsf.spi.management.EndpointRegistry; import org.jboss.wsf.spi.management.EndpointRegistryFactory; public class XTSHandler { //.... public Object getEndpointInstance() { Endpoint endpoint = EndpointAssociation.getEndpoint(); Reference ref = endpoint.getInstanceProvider().getInstance(endpoint.getTargetBeanName()); return ref.getValue(); } public Object getEndpointInstanceCXF(MessageContext ctx) { Endpoint endpoint = ((WrappedMessageContext)ctx).getWrappedMessage().getExchange().get(Endpoint.class); Reference ref = endpoint.getInstanceProvider().getInstance(endpoint.getTargetBeanName()); return ref.getValue(); } public Object getEndpointInstanceCXF(WebServiceContext ctx) { return getEndpointInstanceCXF(ctx.getMessageContext()); } public Object getEndpointInstanceUsingRegistry() { ClassLoader cl = ClassLoaderProvider.getDefaultProvider().getServerIntegrationClassLoader(); EndpointRegistryFactory eprFac = SPIProviderResolver.getInstance(cl).getProvider().getSPI(EndpointRegistryFactory.class, cl); EndpointRegistry registry = eprFac.getEndpointRegistry(); Endpoint endpoint = null; for (ObjectName on : registry.getEndpoints()) { Endpoint ep = registry.getEndpoint(on); //somehow identify your endpoint... name? address? if (ep.getAddress().equals("....")) { endpoint = ep; } //... } Reference ref = endpoint.getInstanceProvider().getInstance(endpoint.getTargetBeanName()); return ref.getValue(); } } {code}	0	0	https://issues.jboss.org/browse/JBTM-999
13110	JBTM-968	Quick notes after Andrew explained the XTS structure to me: {quote} XTS project layout WS-AS, Activity service (needed? - yes but need to know why) WS-CF (general purpose context management), useful WS-C Client and coordinator side impl of the oasis standard (current) Activation and registration service WS-T oasis TX protocol specific services for client, coord, ws negotiation (WS-AT and WS-BA) WS-TX JBoss specific APIs which drive WS-C and WS-T services {quote}	1	1	https://issues.jboss.org/browse/JBTM-968
13111	JBTM-956	Hi Amos, is this closed?	0	0	https://issues.jboss.org/browse/JBTM-956
13112	JBTM-948	Ah, ok. So what you are saying is that this will work on Hudson as you have fixed this problem before for other jobs?	1	1	https://issues.jboss.org/browse/JBTM-948
13113	JBTM-948	Paul, what CI environment did you mean in your comment [1]? Yes, I was talking about Hudson as it is the primary environment for QE testing. If I run the tests locally on my laptop I usually don't need headless mode as I want to see how the test is going and thus see its actions performed in Firefox in front of my eyes. This is the reason why the tests look so at the moment. But I can change them to better fit to your requirements. On Hudson if you want to run a selenium test you need to enable 'Run Xvnc during build' and select some of selenium slaves at 'Label Expression'. Then the selenium test should work. [1] https://issues.jboss.org/browse/JBTM-948?focusedCommentId=12661667&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12661667	1	1	https://issues.jboss.org/browse/JBTM-948
13114	JBTM-934	Do we need pseudo support for idlj? [tom@tomsfc15 JBOSSTS_4_15_0_Final]$ find . -name *idl -exec grep -l pseudo {} \; ./install/jacorb/idl/omg/Corba.idl ./install/jacorb/idl/omg/orb.idl ./install/jacorb/idl/omg/CosTSPortability.idl ./install/jacorb/idl/omg/CosNaming.old.idl ./install/jacorb/idl/omg/PortableServer.idl ./ArjunaJTS/build/install/jacorb/idl/omg/Corba.idl ./ArjunaJTS/build/install/jacorb/idl/omg/orb.idl ./ArjunaJTS/build/install/jacorb/idl/omg/CosTSPortability.idl ./ArjunaJTS/build/install/jacorb/idl/omg/CosNaming.old.idl ./ArjunaJTS/build/install/jacorb/idl/omg/PortableServer.idl ./ArjunaJTS/jacorb/idl/omg/Corba.idl ./ArjunaJTS/jacorb/idl/omg/orb.idl ./ArjunaJTS/jacorb/idl/omg/CosTSPortability.idl ./ArjunaJTS/jacorb/idl/omg/CosNaming.old.idl ./ArjunaJTS/jacorb/idl/omg/PortableServer.idl	1	1	https://issues.jboss.org/browse/JBTM-934
13115	JBTM-934	When creating the jobs, you can probably copy: http://172.17.131.2/view/Narayana+BlackTie/job/narayana-java6/ and http://172.17.131.2/view/Narayana+BlackTie/job/narayana-java7/ Maybe rename the old jobs narayana-jacorb-java6 and the new job narayana-idlj-java6 and so on	1	1	https://issues.jboss.org/browse/JBTM-934
13116	JBTM-934	Will it work on OpenJDK?	1	1	https://issues.jboss.org/browse/JBTM-934
13117	JBTM-928	[~tomjenkinson], these test still don't work. Shall I investigate further or keep them ignored?	1	0	https://issues.jboss.org/browse/JBTM-928
13118	JBTM-883	Does this work? Are you ready to create a pull request?	1	0	https://issues.jboss.org/browse/JBTM-883
13119	JBTM-870	Steps: * Agree on suitable set. * Ensure they are testable and tested via Hudson. * Non-standalone ones need a way of grabbing the AS deps * Must follow the AS7 quickstart rules	0	0	https://issues.jboss.org/browse/JBTM-870
13120	JBTM-852	Firstly some background. It is possible to configure different stores for different purposes. Each store is a singleton so each store type is shared. We currently use at least four stores (XTS uses some others): - an action store for transaction and participant logs; - a recovery store (which defaults to the action store); - a state store for use with TxOj and is not used with JTA (or JTS); - a communications store for storing connectivity information such as ORB endpoints or transaction manager status items. Calling arjPropertyManager.getObjectStoreEnvironmentBean() looks up the default store and is equivalent to: BeanPopulator.getNamedInstance(ObjectStoreEnvironmentBean.class, "default") and is used to configure the action store. The other stores can be configured by looking up a specific named ObjectStoreEnvironmentBean. Use "default" or "stateStore" to configure the action and state stores respectively. Use "communicationStore" to configure the communications store. So the actual bug/feature is that looking up the default store environment bean via the properties manager (arjPropertyManager.getObjectStoreEnvironmentBean()) should return a bean that will configure all stores. Howeve this could be risky (for example it could interfere with XTS behaviour) so my preference would be to force the user to be specific about which stores he wants an environment bean for by making the getObjectStoreEnvironmentBean call take the store name as a parameter. But that leaves the difficult issue of how to specify which store to use in the properties file <entry key="ObjectStoreEnvironmentBean.objectStoreType">XYZStore</entry> since this entry is only applied to the default store. We could simply document the behaviour in jbossts-properties.xml and tell users to explicitly set the communications store programatically if required. So the communications store type would have to be set as follows: BeanPopulator.getNamedInstance(ObjectStoreEnvironmentBean.class, "communicationStore").setObjectStoreType(com.arjuna.ats.internal.arjuna.objectstore.VolatileStore.class.getName()); Is this latter approach acceptable to all concerned?	0	0	https://issues.jboss.org/browse/JBTM-852

#	Requirement	Questions	Answered	Successfully Answered	URL
13121	JBTM-846	Further to previous comment ... for instance what do we (or customers) expect from the statistics gathering routines in this new environment? Should it be measuring all transactions created and terminated in the JVM, or on a per transaction manager basis?	0	0	https://issues.jboss.org/browse/JBTM-846
13122	JBTM-846	To address Davids requirement we need to support 1xArjunaCore (for the MSC) and 1xArjunaJTA||ArjunaJTS (for the EJB container). The problem is the recovery manager, are we expecting to launch two of those in the same VM? I presume we have to in order to cleanly separate the MSC from the EJB container (e.g. different object stores)	0	0	https://issues.jboss.org/browse/JBTM-846
13123	JBTM-846	Why would we need to recovery managers? As long as the transactions are unique and things like the nodename are shared by recovery, it should be possible to have one instance.	0	0	https://issues.jboss.org/browse/JBTM-846
13124	JBTM-846	Allowing multiple TMs in the same JVM is a generic statement and although WF needs this, there may be other users who want it for different reasons. I think we should either define this as a generic problem which we can try to solve for other use cases (do we know them all?) or we should refocus this solely on the WF use case and then come back to a more general requirement if/when it comes back up.	0	0	https://issues.jboss.org/browse/JBTM-846
13125	JBTM-845	Isn't this duplicated by JBTM-1451?	1	0	https://issues.jboss.org/browse/JBTM-845
13126	JBTM-829	Hi, Please can you confirm that I am right in thinking you mean that we should remove the version numbers from the remaining dependencies in maven/pom-template.xml where they have a corresponding version number in the AS7 pom tree? Ta	1	1	https://issues.jboss.org/browse/JBTM-829
13127	JBTM-817	I'm in the process of validating the trace output from these tests and I think doing this automatically is not going to be a trivial process. I have some ideas around how the byteman scripts can be re-factored to simplify the validation process and to also improve scalability so that we can introduce more complex, multi-party tests. Therefore I am removing the requirement to validate the trace output from this issue. Validation will come in a later release. Hopefully this will turn this issue into a much more straight forward task. So to confirm, the requirements for this issue are to: # Run the tests as part of the AS7 testsuite (as per Ivo's pattern: https://github.com/istudens/jboss-as/compare/xts_tests#diff-0) # Automate the running of each crash recovery test without any Byteman rules. This provides some test coverage and ensures we have a valid system before the Byteman rules are applied. # Automate each valid combination of Byteman script to TestCase. The Byteman scripts specify which tests are available. # Gather the output from all test runs in separate files for "eyeball validation". Until this validation is automated, I (with the help of Andrew initially) will do this validation. The output files should be named "<bytman script name>:<test class name>". For example: "ATCrashDuringCommit.MultiParticipantPrepareAndCommitTest". Amos: does this make sense? Is this in line with how you are currently implementing the feature?	1	1	https://issues.jboss.org/browse/JBTM-817
13128	JBTM-817	Also, I don't think we need Arquillian to crash the app server as Byteman is doing that for us. This is good as Arquillian is unlikely to have this feature as part of a JBoss AS release soon enough. Amos: Does this make sense, or am I missing something?	1	1	https://issues.jboss.org/browse/JBTM-817
13129	JBTM-813	Email from Jonathan re ArjunaJTS: JTS install guide is same as core except name and: - add first and last bullet point paras for 'preparing your system' section - add 'specifying the ORB to use' and 'the interface repository' sections' - the 'loading the distribution' section is obsolete, just junk it. - add 'installing into non-default location' section - operating system services needs additional note on PUT_JACORB_HOME_HERE - additional 'web services component' section. JTS admin guide is same as core except name and: - intro paragraph adds mention of web services - add 'OTS/J2EE transaction service management' (except it's now 'Java EE' not 'J2EE') - add 'web service transaction management' section - in periodic recovery section, RecoveryEnvironmentBean.recoveryExtensions module list is wrong for JTS - it should be per the xml in ArjunaJTS/jbossts-properties-arjunajts.xml - add 'orb specific configurations' section - add 'initialising JBossTS applications' section side note: the docs dialect is actually 'en_GB' not 'en_US', hence 'initialising' not 'initializing' and such. What is the company policy on this - do we need to standardise it? The docbook layout seems to assume the canonical version is en_US. Likewise A4 vs. US letter for pdf page size. Jonathan.	0	0	https://issues.jboss.org/browse/JBTM-813
13130	JBTM-810	The controlling property is com.arjuna.ats.arjuna.recovery.recoveryListener (in bean RecoveryEnvironmentBean). With a socket based recovery manager the TM/RM can easily detect whether another RM is managing an object store. Changing the default puts the onus on the user to make sure that only one RM is managing it. So my question is: is it the intention to make the user responsible for getting it right?	1	1	https://issues.jboss.org/browse/JBTM-810
13131	JBTM-810	Would the same problem exist if you opened up two RMs on different ports but against the same ObjectStore? If so, +1 on the user getting it right.	1	1	https://issues.jboss.org/browse/JBTM-810
13132	JBTM-809	Is there any documentation that explains what is going on here? Will JBossTS be still usable for JDBC connections? In what would consist the migration? We're using JBossTS in a JTA infrastructure where database queries are done through JDBC. The client code just uses JDBC connections (obtained through the transactional driver), while the global JTA transaction is controlled using Spring JTA transaction manager. The JBossTS backend, through the transactional driver and some of our glue code, is obtaining JTA connections through XA data sources created using the DynamicClass mechanism. We're working outside JEE, just plain Tomcat + JBossTS + Spring. How will the scenario change with JCA?	1	1	https://issues.jboss.org/browse/JBTM-809
13133	JBTM-785	Hi Paul, Is there any update on this one? Is it covered by any QS? Tom	1	0	https://issues.jboss.org/browse/JBTM-785
13134	JBTM-785	Hi [~paul.robinson], Is there any update on this? We can't understand what is needed here. What is a standalone client in this context? Tom	1	0	https://issues.jboss.org/browse/JBTM-785
13135	JBTM-747	I have implemented draft 4 of the specification that resides at http://www.jboss.org/reststar/specifications/transactions.html with some amendments (documented in the docs directory). What's left: - hook it into the build releases target of JBossTS; - port the "Night out" demo from the previous REST-tx prototype; - track updates to the specification	0	0	https://issues.jboss.org/browse/JBTM-747
13136	JBTM-741	[~tomjenkinson], what would you think about postponing this issue until JBTM-2571 is resolved? I think it'd be too dangerous to remove public classes or methods without having code coverage from all available tests. What do you think?	1	0	https://issues.jboss.org/browse/JBTM-741
13137	JBTM-728	Are you sure the remote tests are actually running? Looks to me like they are turned off, hence no coverage. <target name="run.tests"> ... <exclude name="**/remote/**"/> ... </target>	1	1	https://issues.jboss.org/browse/JBTM-728
13138	JBTM-713	Marks description from the duplicate issue: {quote} We're working on WS-BA based compensation annotations. That's great, but can we refactor so that we can support compensations outside of WS-BA? ArjunaCore can support compensations without Web Services. Maybe the WS-BA annotations vanish entirely, or collapse down to a subset of what currently exists to support what is truly unique about WS-BA? {quote}	0	0	https://issues.jboss.org/browse/JBTM-713

#	Requirement	Questions	Answered	Successfully Answered	URL
13139	JBTM-700	I looked at this from the perspective of EAP, i.e., what docs are needed there versus only needed in the community. TS docs outline for EAP distribution: In: 0..1 InstallGuide 1..1 ReleaseNodes (JIRA generated) FailureRecovery Guide Administration Guide (ArjunaCore as basis, with diffs from JTA and JTS) Programmers Guide (ArjunaCore with all TXOJ aspects removed; start with JTA) Separate Programmers guide for JTS and XTS (maybe separate Admin guides then?) Separate all JTS and XTS information into separate documents? JTS Quickstart Guide (should be expanded to cover JTA) Configuration guide (Administration Guide update) Out: Debugging documentation All TXOJ aspects from documentation Maybe: None. Duplicated: Tooling discussions Debugging discussions ObjectStore setup Errors	0	0	https://issues.jboss.org/browse/JBTM-700
13140	JBTM-659	Hi Mike, Trying to work out whether this is something we want to support on 4.16.x branch right now. Do you know if it is fixed in 4.16.x yet? e.g. hoping to put 4.16.1 out on Friday and hopefully that would be the last version for EAP6. Tom	1	1	https://issues.jboss.org/browse/JBTM-659
13141	JBTM-659	When I come to release 4.16.1 it is going to ask me what to do with this bug, is it in a state that it can be closed or does it need to move to 4.16.2 (or ideally 5.0.0.M2)?	0	0	https://issues.jboss.org/browse/JBTM-659
13142	JBTM-659	Is this bug resolved in 4.6.1.CP13?	0	0	https://issues.jboss.org/browse/JBTM-659
13143	JBTM-659	I fixed it in a private branch of JBOSSTS_4_6_1_GA_CP12. I am not aware of any CP13 branch. The details are in https://issues.jboss.org/browse/JBTM-522 including a link to the hudson job that tests the functionality. Has there been a decision on whether the backport will go into a 4.6.1 release? When/if it does then we should ask QA to include it in there test plan.	0	0	https://issues.jboss.org/browse/JBTM-659
13144	JBTM-659	Can we mark this as complete?	1	1	https://issues.jboss.org/browse/JBTM-659
13145	JBTM-608	huh? we already do - the syncs records are ordered and you can register one even during beforeCompletion as long as the new one has a order later in the sequence that the current one.	1	1	https://issues.jboss.org/browse/JBTM-608
13146	JBTM-596	Introduced a set of *EnvironmentBean classes and updated code to use them instead of accessing the property manager directly. The beans have hardcoded defaults that are overridden with values from the property manager (i.e. properties file / system properties) when the bean is instantiated. Thereafter the bean is independent of the property manager. A handful of property manager usage remains, esp. in the orb portability layer which uses dynamically named properties that are difficult to map to fixed beans. Needs some further thought. Some more thought also needed to: capturing running config, since it's not necessarily the same as what's in the properties file. Some sort of bean state dumper, similar to what the module Info classes do? even stronger type safety - e.g. some beans have hostnames vs. InetAddresses, lists of classnames vs lists of classes - consider impact on injection systems and JMX API. dynamic reconfig mechanism - should updating bean properties automatically restart relevant parts of the system e.g. rebind listeners to new port without requiring full restart? but that's all for the next cycle. For now the type safety improvements are enough progress.	0	0	https://issues.jboss.org/browse/JBTM-596
13147	JBTM-595	LogFactoryInterface yields a AbstractLogInterface - how does one get from Logl18n to AbstractLogInterface? I confess this logging infrastructure is pretty baffling to me	1	1	https://issues.jboss.org/browse/JBTM-595
13148	JBTM-595	I stuck with Logl18n initially because it was unclear how LogFactory can receive resource bundle information. I guess the key to solving resource bundles will indeed be some kind of nifty subclassing as we are limited to a single resource bundle instance per category. Would you consider simply moving to java.util.logging wholesale, and removing all the facade stuff completely? And by "you" I guess I mean "me", since it's my dumb idea. The only practical difficulty I see would be that I'd have to add a Level class so that you have DEBUG/TRACE instead of JUL's FINE/FINER/FINEST, but that's not too big a deal. The upside would be removing the dep on commons-logging and deleting a bunch of classes. At first glance I think the resource generation stuff would work OK with JUL.	1	1	https://issues.jboss.org/browse/JBTM-595
13149	JBTM-581	For example ... <!-- Periodic recovery settings. Time values in this section are in seconds. --> <!-- Interval in seconds between initiating the periodic recovery modules. Default is 120 seconds. --> <property name="com.arjuna.ats.arjuna.recovery.periodicRecoveryPeriod" value="120"/> <!-- Interval in seconds between first and second pass of periodic recovery. Default is 10 seconds. --> <property name="com.arjuna.ats.arjuna.recovery.recoveryBackoffPeriod" value="10"/>	1	1	https://issues.jboss.org/browse/JBTM-581
13150	JBTM-575	Hi Richard This does not make much sense to me. JBAS-4238 and JBTM-66 are similar but not equivalent. The JBAS issue is addressed in both releases, but JBossAS 4.2.3 uses TS 4.2.3.SP7, which does not have the JBTM-66 changes. JBossTS AS 5.1.0.GA uses TS 4.6.1.GA , which does. So, you seem to be having problems due to a behaviour change that you requested in the first place?! Ah well, maybe I misunderstood what you wanted. Do you have any preference between the implementation alternatives I outline in http://community.jboss.org/message/515473 assuming either will address your needs?	1	1	https://issues.jboss.org/browse/JBTM-575

#	Requirement	Questions	Answered	Successfully Answered	URL
13151	JBTM-575	though I understand JIRA seems to indicate it does. The fact remains that I (and others, by the look of it) are seeing transaction RollbackExceptions *that do not include their root cause* when using JBoss 5.1.0.GA. For example, I see: --- 08:48:35,468 http-127.0.0.1-8080-6 DEBUG [SQL] delete from Resource where id=? and version=? 08:48:35,484 http-127.0.0.1-8080-6 WARN [JDBCExceptionReporter] SQL Error: 1451, SQLState: 23000 08:48:35,484 http-127.0.0.1-8080-6 ERROR [JDBCExceptionReporter] Cannot delete or update a parent row: a foreign key constraint fails (`avant-ss-production/bug`, CONSTRAINT `FK106546DD208B6` FOREIGN KEY (`assignedTo_id`) REFERENCES `resource` (`id`)) 08:48:35,484 http-127.0.0.1-8080-6 ERROR [AbstractFlushingEventListener] Could not synchronize database state with session org.hibernate.exception.ConstraintViolationException: could not delete: [com.avant.ss.ejb.entity.Staff#402881812627347901 2627c550d60004] at org.hibernate.exception.SQLStateConverter.convert(SQLStateConverter.java:94) at org.hibernate.exception.JDBCExceptionHelper.convert(JDBCExceptionHelper.java:66) at org.hibernate.persister.entity.AbstractEntityPersister.delete(AbstractEntityPersister.java:2569) at org.hibernate.persister.entity.AbstractEntityPersister.delete(AbstractEntityPersister.java:2725) at org.hibernate.action.EntityDeleteAction.execute(EntityDeleteAction.java:97) at org.hibernate.engine.ActionQueue.execute(ActionQueue.java:279) at org.hibernate.engine.ActionQueue.executeActions(ActionQueue.java:263) at org.hibernate.engine.ActionQueue.executeActions(ActionQueue.java:172) at org.hibernate.event.def.AbstractFlushingEventListener.performExecutions(AbstractFlushingEventListener.java:321) at org.hibernate.event.def.DefaultFlushEventListener.onFlush(DefaultFlushEventListener.java:50) at org.hibernate.impl.SessionImpl.flush(SessionImpl.java:1027) at org.hibernate.impl.SessionImpl.managedFlush(SessionImpl.java:365) at org.hibernate.ejb.AbstractEntityManagerImpl$1.beforeCompletion(AbstractEntityManagerImpl.java:504) at com.arjuna.ats.internal.jta.resources.arjunacore.SynchronizationImple.beforeCompletion(SynchronizationImple.java:101) at com.arjuna.ats.arjuna.coordinator.TwoPhaseCoordinator.beforeCompletion(TwoPhaseCoordinator.java:269) at com.arjuna.ats.arjuna.coordinator.TwoPhaseCoordinator.end(TwoPhaseCoordinator.java:89) at com.arjuna.ats.arjuna.AtomicAction.commit(AtomicAction.java:177) at com.arjuna.ats.internal.jta.transaction.arjunacore.TransactionImple.commitAndDisassociate(TransactionImple.java:1423) at com.arjuna.ats.internal.jta.transaction.arjunacore.BaseTransaction.commit(BaseTransaction.java:137) at com.arjuna.ats.jbossatx.BaseTransactionManagerDelegate.commit(BaseTransactionManagerDelegate.java:75) at org.jboss.aspects.tx.TxPolicy.endTransaction(TxPolicy.java:170) at org.jboss.aspects.tx.TxPolicy.invokeInOurTx(TxPolicy.java:87) at org.jboss.aspects.tx.TxInterceptor$Required.invoke(TxInterceptor.java:190) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.aspects.tx.TxPropagationInterceptor.invoke(TxPropagationInterceptor.java:76) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at org.jboss.ejb3.security.RunAsSecurityInterceptorv2.invoke(RunAsSecurityInterceptorv2.java:94) at org.jboss.aop.joinpoint.MethodInvocation.invokeNext(MethodInvocation.java:102) at	1	1	https://issues.jboss.org/browse/JBTM-575
13152	JBTM-575	at com.arjuna.ats.jbossatx.BaseTransactionManagerDelegate.setRollbackOnly(BaseTransactionManagerDelegate.java:123) at org.hibernate.ejb.AbstractEntityManagerImpl.markAsRollback(AbstractEntityManagerImpl.java:421) at org.hibernate.ejb.AbstractEntityManagerImpl.throwPersistenceException(AbstractEntityManagerImpl.java:576) at org.hibernate.ejb.AbstractEntityManagerImpl.throwPersistenceException(AbstractEntityManagerImpl.java:614) at org.hibernate.ejb.AbstractEntityManagerImpl$1.beforeCompletion(AbstractEntityManagerImpl.java:513) What's key here is that, from the transaction manager's perspective, the setRollbackOnly call precedes the exception - hibernate sets the tx to rollback before throwing the exception. So, the setRollbackOnly call IS the root cause of the transaction rolling back. JBTM-66 is not only present but functioning exactly as designed. As mentioned in the forum thread, this is arguably misuse of the api by hibernate. Of the alternatives proposed in the forum thread, chaining the deferredThrowable to the rollbackOnly trace won't work, as there is no deferredThrowable at the point in time that the setRollbackOnly call is made. So the only viable alternative seems to be using the deferredThrowable as the initCause of the rollback in preference to the rollbackOnlyCallerStackTrace in cases where both are available. Will that meet your needs?	1	1	https://issues.jboss.org/browse/JBTM-575
13153	JBTM-575	Johathan, can I ask you another question on a related topic(one that does not involve the client at all). If a business method in EJB session bean throws a RuntimeException in the JBoss log I see only this line: 11:32:12,550 ERROR [TxPolicy] javax.ejb.EJBTransactionRolledbackException No cause, no message, no nothing - the RuntimeException is however successfully relayed to the client, but I remember that in older JBoss version I used to get I full stack trace with the javax.ejb.EJBTransactionRolledbackException. What might be causing this problem?	0	0	https://issues.jboss.org/browse/JBTM-575
13154	JBTM-575	Jonathan, As you asked, I asked someone from the Hibernate team (Steve Ebersole) about casuality. Steve has since responded, as have a couple other people. What is our next step to getting this resolved? Regards, Richard.	1	1	https://issues.jboss.org/browse/JBTM-575
13155	JBTM-575	Was this fixed or not? I can't tell from the JIRA or the forum post. Assuming that it was, I will write this release note: Cause: In ArjunaCore which was included in JBoss Application Server 4.2 and previous, when a transaction rolled back, the <systemitem>RollbackException</systemitem> included the root cause of the exception. In JBoss Application Server 5.1, the exception no longer included this information, even though the information could still be found in a different area of the log file. Consequence: Developers who relied on the information included in the <systemitem>RollbackException</systemitem> in the older versions of JBoss Application Server found that their applications no longer worked as expected. Fix: The debugging information has been put back into the exception, so that it works as it did previously. Result: Developers can once again write code to take advantage of the debugging information found in the <systemitem>RollbackException</systemitem>.	0	0	https://issues.jboss.org/browse/JBTM-575
13156	JBTM-573	This says "Done" but there is no information. Is it done or WONTFIX? If it is done, what is the command-line tool and what is the syntax for using it?	1	1	https://issues.jboss.org/browse/JBTM-573

#	Requirement	Questions	Answered	Successfully Answered	URL
13157	JBTM-555	Maybe running the test with the gc.log and analysing it afterwards would also be a sounded approach ? Basically what is to know of the test is : "Is there memory consumption below a certain threshold ?", this could be answer by a gc.log file... (However, this is rather tied to the Sun JVM implementation). Moreover, analysing the GC log also allow to see how the memory consumtion progress, maybe spotting a memory consumption rising (even if it's stay beneath the appropriate threshold during the test) that wasn't rising in the previous version. That gain this sort of knowledge would requires you to either run test on previous tag release or keep previous GC.log on file, however doing the first one with Git is not really that complicated (in fact, it will be similar to using the 'bissect' command). On other "trick" that I know of is the Ant's for-each statement in the ant-contrib.jar, which will allow to do something like this: <for-each ..> <junit .... > <!-- run one particular test or test group --> <!-- Load memory configuration for an appropriate test from a properties file --> <property name="${testId}-config.."/> <jvmarg value="-Xms${memory.start}"/> <jvmarg value="-Xmx${memory.max}"/> <!-- ... Hope that's help.	1	1	https://issues.jboss.org/browse/JBTM-555
13158	JBTM-535	Another thought about the Listener.finalize() method: it is useless also because if the Listener is running, it can't be finalized, so finalize() will never be called. Viceversa, if finalize() is called it means that the listener has already been stopped, so its finalize() implementation is useless... am I wrong?	1	1	https://issues.jboss.org/browse/JBTM-535
13159	JBTM-535	Could this be targetted for 4.7.0? Resolution should be quite straightforward... Thanks in advance.	1	1	https://issues.jboss.org/browse/JBTM-535
13160	JBTM-522	new Uid ("invalidString", true); should return cleanly with a Uid such that x.valid()== false. Instead it throws NPE. The exception handling probably needs a catch for that in addition to the existing exceptions it knows about.	0	0	https://issues.jboss.org/browse/JBTM-522
13161	JBTM-503	Hi Andrew, any news on this? Could it be considered for 4.7.0?	1	1	https://issues.jboss.org/browse/JBTM-503
13162	JBTM-465	I'm here again regarding this issue. As I wrote in the referenced forum topic, what I'm doing right now is to collect all the created connections (by my JDBCAccess implementation) and shut down them on webapp restart. However, there's a little issue when embedding JBossTS in a stand-alone application. In our case, this is happening when unit-testing our application code. I'm going to explain it better. In out unit testing framework, we do the following: --- for every test case: --- 1. in the setup phase, initialize JBossTS 2. in the test phase, run the tested code, which uses JBossTS 3. in the tear down phase, shutdown JBossTS Step 1. and 3. are almost the same that happen in our web application start and stop phases. In step 3. all the connections returned by our implementation of JDBCAccess are closed. Now, when more than one test case is executed (so, points 1. to 3. are reiterated), during the second, third, etc. iterations unfortunately in step 1. JBossTS does not re-create the Object Store, so at the first attemp to connect to it, connections in com.arjuna.ats.internal.arjuna.objectstore.JDBCImple._theConnection array are all closed (by step 3. of the previous iterations), so any attempt to setup and execute a prepared statement (for instance in com.arjuna.ats.internal.arjuna.objectstore.JDBCImple.write_state(Uid, String, OutputObjectState, int, String)) fails. This is not a big problem, because com.arjuna.ats.internal.arjuna.objectstore.JDBCImple.reconnect(int) is invoked to create a new connection, however if you consider for instance com.arjuna.ats.internal.arjuna.objectstore.jdbc.h2_driver.write_state(Uid, String, OutputObjectState, int, String), because the internal state of the class is not handled in a stack-way (to handle the recursive call in the catch clause), when the outer write_state calls terminates and calls com.arjuna.ats.internal.arjuna.objectstore.JDBCImple.freePool(int), a warning is logged (identified by "com.arjuna.ats.internal.arjuna.objectstore.JDBCImple_18") because the connection has been already freed by the inner call to write_state. As I said, this is not a real issue, but it is annoying in our use case. Anyway, all of this could be avoided if either JDBCImple or the Object Store correctly implemented a shutdown method. In fact, from one iteration to another the Object Store is recycled and not recreated because even if I correctly disable the TxControl in step 3., in step 1. of the following iteration the TxControl is re-enabled, the TransactionStatusManager is correctly recreated and it creates a new ActionStatusService. However, the latter keeps on using its static field com.arjuna.ats.arjuna.recovery.ActionStatusService._objectStore, which is initialized once (at step 1. of the first iteration) and never disposed. Saying it in a simple way: it would be useful to implement a correct lifecycle for the Object Store (and for its connections, of course), bound to the TxControl one.	1	1	https://issues.jboss.org/browse/JBTM-465
13163	JBTM-465	Please can you review this in context of JBTM-1230, I think the changes we made there might have addressed your requirements	1	1	https://issues.jboss.org/browse/JBTM-465
13164	JBTM-465	Hi Tom, please let me understand something. From JBTM-1230 I understand that JDBC connections for the Object Store are now retrieved using JNDI, is this correct? If so, as I already stated already in other circumstances, this is a problem for us, since we're using JBoss TS embedded in an application which is meant to work outside a JEE environment. We would need to setup a fake JNDI repository with Spring to let the object store find JDBC connections. Do you confirm this? Or is there an alternative way to still provide connections without JNDI? If not, I personally consider this, together with the push towards the mandatory use of a JCA connector (in place of the plain JDBC implementation) in JBoss TS 5.x, a step backward. Please note that I've not followed this evolution so tightly, but this is what I understood based on the feedback I received on the bugs I opened myself (and for some of which I also contributed with patches that were never integrated). Thanks in advance.	1	1	https://issues.jboss.org/browse/JBTM-465
13165	JBTM-437	Jonathan, can you take this?	0	0	https://issues.jboss.org/browse/JBTM-437
13166	JBTM-409	This is now working for the 1.1 WS-AT implementation. Need to add supporrt for WS-AT 1.0 next. There is now a new API class UserSubTransaction which subclasses the remote UserTransaction class. It can be used to create an AT UserTransaction when already inside an AT transaction by calling getTransaction. This can be suspended and resumed via the TransactionManager class like any other UserTransaction. Any participants enlisted in the subtransaction will be coordinated using the coordinator of the local host rather than the coordinator of the original TX. See the WSTX tests for details.	0	0	https://issues.jboss.org/browse/JBTM-409
13167	JBTM-400	Done: JBTM-1472. Can you add some more detail to the description? As it stands, my understanding of this issue is that it's a duplicate of JBTM-1472.	1	1	https://issues.jboss.org/browse/JBTM-400
13168	JBTM-394	Added getters and setter for addresses / ports to the jta and jts service mbean interfaces and implementations. Revised assorted jta/jts/core internals to remove knowledge of AS from things that don't need it - the TransactionManagerService puts the values from ServiceBindingManager (or SERVER_BIND_ADDRESS as a fallback) in to the arjuna config objects, so app server integration is transparent to the rest of the JBossTS codebase. Revised XTS AS integration to read same service bindings the AS webserver uses, making it more POJO and less JMX style integration. Whilst checking port bindings, discoved and fixed inconsistencies in various places where 4711 was stated as recovery manager port, rather than the correct 4712. Used 4713 for transaction status manager and 4714 for socketProcessId in the app server. Note that when running standalone these still use random ports - do we want to hardcode these defaults in the standalone version also?	0	0	https://issues.jboss.org/browse/JBTM-394

#	Requirement	Questions	Answered	Successfully Answered	URL
13169	JBTM-390	Most JTS test now run locally. Some issues still exist: no JTA test coverage, no JDBC test coverage, excessive memory usage by ant/junit, excessive buffering of output (may be related to memory usage), occasional unexplained hanging of tests on husdon only (ORB shutdown blocking?). These are relatively low priority. Pushing to next release.	0	0	https://issues.jboss.org/browse/JBTM-390
13170	JBTM-389	Which svn tag are your src changes based on?	1	1	https://issues.jboss.org/browse/JBTM-389
13171	JBTM-387	A prototype package has been added to the TS repository in workspace/adinn/orchestration which supports orchestration of different activities in multi-threaded test cases.The prototype works by injecting triggers for Event Condition Action (ECA) rules at arbitrary points in the code under test. The ECA rules include primitives which enable sequencing, synchronization, delay and abort of the threads executing the trigger class.. The prototype has been used to automate one of the test scenarios for XTS recovery and will be developed to cover other cases. However, the tool is independent of the TS code base and development and use of the tool for recovery testing will continue beyond the 4.4.0 GA release. So this JIRA is being rescheduled for completion after that date .	1	1	https://issues.jboss.org/browse/JBTM-387
13172	JBTM-369	Can you give a use case for this? I'm not seeing any real advantage over just doing e.g. -Dcom.arjuna.ats.arjuna.xa.nodeIdentifier=foo ?	1	1	https://issues.jboss.org/browse/JBTM-369
13173	JBTM-369	Jonathan, How to end-up with a "smooth" integration within the AS? I understand that the TM has its own configuration file when working inside some other runtime, but when runing inside the AS, can't he use the JBoss config files (like jboss-services.xml file) or similar? If not, can we at least implement the same feature set as JBoss AS has, including the property replacer that Roland was suggesting to add? Thanks for your feedback. Cheers, sacha	1	1	https://issues.jboss.org/browse/JBTM-369
13174	JBTM-369	Hi Sacha Please clearly define what you are trying to achieve. So far nobody has given me a convincing, specific real world use case for this feature. As far as I can tell the existing ability to override values from the config file by specifying system properties is equivalent to value substitution, except for cases where you may want to use the same value for multiple parameters. The only such case I can think of is the IP address, for which we have a separate solution anyhow. With the exceptions of the port bindings (which will be done via SBM), recovery (which will be done via the new tx integration API) and multiple 1PC (which is an horrifically bad idea anyhow), users should not need to change values in the tx properties file. I believe this feature will benefit only a small proportion of users. I agree it's better to have consistent behaviour for the properties files, it makes for a better user experience. But that's secondary to having a product that works. Because of these factors, I consider this issue a low priority. It will get done as and when we have time, but not until e.g. crash recovery and JTS integration are finished. That's why it's still not scheduled against any particular release. Perhaps I'm missing something though. If you feel it's worth taking time out from managing the division to comment on the issue, you must feel it's important. Can you say why?	1	1	https://issues.jboss.org/browse/JBTM-369
13175	JBTM-360	When resolved the server manager jar should be included in the JBossTM build (and the server wrapper removed)	0	0	https://issues.jboss.org/browse/JBTM-360
13176	JBTM-355	Can you please explain why the value of this timeout needs to be configured with a larger value and how this value relates to your transaction's 'size'? Also, by 'size' are you referring to the number of participants or the amount of work done by each participant? The transport timeout is only ever employed when dispatch of a WS-C/T protocol message must wait on receipt of a corresponding protocol message. For example, the WS-AT Coordinator engine responds to a prepare request by sending a PREPARE message to the participant and then waiting for a PREPARING message to be returned in response. Processing of the incoming message and dispatching the corresponding response will normally happen very quickly as it only involves a small amount of work on the part of the transaction system. The other variables in the equation are the communication transport delays and, in some cases, a certain amount of processing time added by the participant implementations (in the example above the participant's prepare method will be invoked). The participant may conceivably be delaying the response for more than 30 seconds. If so then this is likely to be due to an error in the participant.	1	1	https://issues.jboss.org/browse/JBTM-355
13177	JBTM-355	I'm still not convinced that configuring this timeout is the correct thing to do. Can you explain i) which database you are using and ii) what your code is doing in the prepare call? I am assuming that this problem arises because entities created in your subordinate TX only get flushed at prepare time due to the delegation involved in bridging. The persistence manager install a Syncrhonization to flush cached SQL updates to the database. This normally runs before the 2PC is entered but with bridging the flush cannot be performed (i.e. the Synchronization cannot be executed) until the delegated prepare arrives which stalls the parent TX. One possibility might be for you to manually flush changes inside your EJB code as you make them. This should leave little or no work to do at prepare time because the DB updates will already be out on disk. Can you try this? Another option might be to do break the create operations into lots of little transactions. Do you have to do this update all in one go? Anyway leaving these questions aside for the moment, if you want to reset the timeout you can easily do so by invoking TransportTimer.setTransportTimeout(long) from your application code.	1	1	https://issues.jboss.org/browse/JBTM-355
13178	JBTM-345	What happens if a module goes away between first and second passes, or during a recovery pass? Do we need blocking unregister calls for modules so that does not happen?	1	1	https://issues.jboss.org/browse/JBTM-345
13179	JBTM-334	Just a clarification: this is not about deprecating what NORMAL does, just changing the name from NORMAL to PERIODIC. Correct?	0	0	https://issues.jboss.org/browse/JBTM-334
13180	JBTM-331	See http://jira.jboss.com/jira/browse/JBTM-20?page=all for the backstory on this. We already downgrade this stuff from INFO to DEBUG when running integrated with the AS. The arjuna level won't be changed, so the options are to change the global rewriting of arjuna messages in AS from INFO->DEBUG to INFO->TRACE which will impact other messages too, or install a filter in log4j. Why is AS even logging at DEBUG by default? It should be INFO, at least for production releases.	1	1	https://issues.jboss.org/browse/JBTM-331
13181	JBTM-319	Yes, we definitely don't want to duplicate code from elsewhere. How about getting the ManagedConnectionFactory instance from JNDI, since that's what creates the WrapperDataSource in the first place?	0	0	https://issues.jboss.org/browse/JBTM-319
13182	JBTM-319	Can you take any queries and discussion to the forums, where it's easier to track? Thanks.	0	0	https://issues.jboss.org/browse/JBTM-319
13183	JBTM-312	Would this work as a general design pattern if you wanted to implement a HA recovery coordinator? 1. Use the JDBC object store. 2. Configure a recovery coordinaor (RC) on each node to recover all (*) nodes. You can only *start* one of these RC at any one time 3. Develop some "watchdog" process runs on every node and monitors the current RC process and if it fails, uses a group protocol to elects one of the other RC's to start 4. Rather than starting the RC on one of the nodes, start the watchdog processes and allow them to enable one of the RC's Would this work for JTS?	1	0	https://issues.jboss.org/browse/JBTM-312

#	Requirement	Questions	Answered	Successfully Answered	URL
13184	JBTM-300	Thank you for backporting. I put 4.2.3.SP5_CP07-brew jars into EAP 4.3.0 CP06 (SOA-P 4.3.0 CP02) and tested AtomicActionExpiryScanner. Then I found that expired tx-logs were copied to "Expired" directory but the original files were not removed. > find data/tx-object-store ... data/tx-object-store/HashedActionStore/defaultStore/StateManager/BasicAction/TwoPhaseCoordinator/AtomicAction/#197#/7f000001_9020_4b025a1e_2f9 data/tx-object-store/HashedActionStore/defaultStore/StateManager/BasicAction/TwoPhaseCoordinator/AtomicAction/Expired/#197#/7f000001_9020_4b025a1e_2f9 ... It's hitting JBTM-418. Could you take care of JBTM-418 for 4.2.3 branches too?	0	0	https://issues.jboss.org/browse/JBTM-300
13185	JBTM-280	Do you mean simply saving them as .odt files? I tried to check one in but got the following error : svn: CHECKOUT of '/!svn/ver/15549/labs/jbosstm/trunk/ArjunaCore/docs/install_guide': 403 Forbidden (https://svn.labs.jboss.com) Am I using the wrong path or don't I have the correct permissions to commit perhaps. Thanks, Daniel	1	1	https://issues.jboss.org/browse/JBTM-280
13186	JBTM-265	Ok, I have managed to modify XTS to run over JBossWS/CXF 3.1.0 but a few caveats apply I had to patch two bugs in CXF so we cannot go ahead with moving to CXF untiil JBossWS assimilates fixes for CXF-2161 and CXF-2167. I had to break XTS working on JBossWS/Native because of several incompatibilities. The most notable difference is that the AddressingProperties class and related classes reside in different packages but there are a couple of others: i) the RelatesTo field in CXF has only one entry whereas in JBoss it is an array of entries ii) when a JaxWS port is created from a W3CEndpoint JBossWS configures its request messaqe properties with a MAP containing any WSA header values derived from the endpoint. this requires the XTS code to merge any extra headers into the existing MAP. CXF stashes the endpoint WSA header values in the port and so XTS has to create a MAP containing the extra headers and hang this off the request message properties. iii) CXF does not support JAX-RPC so the JaxRPC client handlers could not be compiled (the requisite headers are missing from the CXF libraries). This means the old version of the demo has to be dropped on CXF. Ignoring the JAX-RPC issue, it ought to be possible to hide these differences behind a wrapper class which managed the underlying MAP types and handled the task of configuring a JaxWS port. However, this would mean that the build process would have to be modified to build either a JBossWS/Native compatible version or a JBoss/CXF compatible version i.e. the XTS code would for all intents and purposes become dual binary (this is why we really need a fix in JBossWS).	0	0	https://issues.jboss.org/browse/JBTM-265
13187	JBTM-264	Is this resolved as part of the CLI work possibly? e.g. can the AS7 CLI be deployed externally?	1	1	https://issues.jboss.org/browse/JBTM-264
13188	JBTM-264	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	1	0	https://issues.jboss.org/browse/JBTM-264
13189	JBTM-252	Daniel, would you please clarify the scope of the quanitfiers 'some' and 'all' in your issue description? i.e. in which documents have you noticed the format problem? which other documents do you suspect are affected?	1	1	https://issues.jboss.org/browse/JBTM-252
13190	JBTM-144	Mark The DTF QA tests have their own (JBoss internal) repo, due to different licencing for the DTF. Can we close this one?	1	1	https://issues.jboss.org/browse/JBTM-144
13191	JBTM-116	Which lower level? JBossTS understands -1. It means "transactions never time out".	1	1	https://issues.jboss.org/browse/JBTM-116
13192	JBTM-103	Mark I'm not clear on this one. The effective default timeout appears to be 0 regardless of which TransactionManagerImple is used. Am I missing something?	1	1	https://issues.jboss.org/browse/JBTM-103
13193	JBTM-103	In ArjunaJTA BaseTransaction.begin: Integer value = (Integer) _timeouts.get(ThreadUtil.getThreadId()); int v = 0; // if not set then assume 0. What else can we do? if (value != null) { v = value.intValue(); } // TODO set default timeout which doesn't exist in the JTAX BaseTransaction, because that relies on the JTS Current object to manage timeouts, which pulls the default from the JTS property and not the ArjunaCore property (check relevant Environment.java classes to see what I mean). The JTA and JTS implementations should both use the Core definition of default transaction timeout. At the moment, if someone wants to set a default timeout for JTA transactions, that should be used if no explicit timeout is defined, they can only do so via the JTS implementation. Plus, they'd need to make sure the JTS value is the same as the Core value if they want to mix and match transaction implementations (not likely in reality). We should have only a single place to set the default transaction timeout, and this value should be used when no timeout is explicitly defined. A value of 0 or -1 should be taken as "no timeout", i.e., infinite. JTS TransactionFactory does this already: int theTimeout = time_out; if (theTimeout == 0) theTimeout = TransactionFactoryImple._defaultTimeout; and _defaultTimeout could be 0.	0	0	https://issues.jboss.org/browse/JBTM-103
13194	JBTM-98	what is the setup of your test? If I have a single JBoss server talking to oracle then it works fine. When I try the following scenario is when I get the problem (both server 4.0.3 SP1 and JBossts 4.2): Server1 : starts the global distributed transaction, performs local oracle database operations and calls a server2Bean on server2 // this method has the tx "required" attribute in ejb-jar.xml performGlobalTx(){ doSomething(); // do oracle operations server2Bean.doSomething(); // call remote jboss Server2 } Server2: the bean on this server updates the same oracle database: doSomething(){ performDBOperations...; // here is where I get the exceptions reported in the forum <<<<----------here }	1	0	https://issues.jboss.org/browse/JBTM-98
13195	JBTM-95	We have such a thing in place now as far as I can see. Let's close this?	1	1	https://issues.jboss.org/browse/JBTM-95
13196	JBTM-86	Could you please increase the priority of this bug, or at least give it some sort of schedule? Defect analysis becomes a nightmare without any sort of exception chaining.	0	0	https://issues.jboss.org/browse/JBTM-86
13197	JBTM-86	Hi Andrew, hi Jonathan, what is the fix version that includes the whole fix? 4.14? Thanks in advance!	1	0	https://issues.jboss.org/browse/JBTM-86
13198	JBTM-49	Does this still make sense? We should probably ensure our tooling is consistent with AS7 as much as anything.	1	1	https://issues.jboss.org/browse/JBTM-49
13199	JBTM-49	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	1	0	https://issues.jboss.org/browse/JBTM-49
13200	JBTM-20	Can you add something to the documentation (the section around application server integration)? If you want, send me text and I'll put it in place.	1	1	https://issues.jboss.org/browse/JBTM-20
13201	JBTM-5	From Adrian Brock: "Adrian comment: "Will this also be integrated with the JBoss installer and the JBoss Network for "ease" of installation/choice? i.e. we include JBoss Transactions in the installer distrubtion and let them choose between the two implementations."	0	0	https://issues.jboss.org/browse/JBTM-5
13202	JBVFS-194	Aleš, would you mind fixing or removing these tests please? Thank you.	1	0	https://issues.jboss.org/browse/JBVFS-194

#	Requirement	Questions	Answered	Successfully Answered	URL
13203	JBVFS-161	(10:39:27 AM) ALR: baileyje: Up yet? :) (10:39:48 AM) baileyje: Yeah :) (10:40:11 AM) ALR: baileyje: I'd like to get at a TempDir without a ScheduledExecutorService (10:40:21 AM) ALR: The service is just used for cleanup operations in close() (10:40:54 AM) ALR: http://pastebin.com/sBuqcbB2 (10:40:59 AM) baileyje: Yeah. Cleanup with some retry behavior to prevent locking. (10:41:05 AM) ALR: Yep. (10:41:12 AM) ALR: baileyje: Can that be optional? (10:41:32 AM) ALR: And instead go to some deleteOnExit logic if no SES was provided at construction? (10:42:41 AM) ALR: baileyje: I've a case where we shouldn't be flushing VirtualFiles to getPhysicalFile anyway (10:42:56 AM) ALR: So providing a SES is kind unnecessary in the first place. (10:42:57 AM) baileyje: ALR: hmmm. What is the issue with the executor? (10:43:03 AM) ALR: baileyje: I don't have one. (10:43:12 AM) baileyje: ok. You just want to avoid it.. (10:43:13 AM) ALR: baileyje: And in AS, most components are not allowed to make one. (10:43:15 AM) ALR: Yeah. (10:44:15 AM) ALR: baileyje: I'll be happy to make the patch/fix if we can agree on some alternate logic. (10:44:52 AM) baileyje: I will pull dmlloyd in on that. He was the originator of the TempFileProvider. (10:45:17 AM) ALR: baileyje: Cool, thanks. We'll see what he says in a bit. (10:45:44 AM) baileyje: I think if there is a legitimate reason to have a different cleanup strategy, something could be done. (10:46:02 AM) ALR: baileyje: Simplification / ease of ue. (10:46:04 AM) ALR: *use	0	0	https://issues.jboss.org/browse/JBVFS-161
13204	JBVFS-12	Is this fixed or not in Beta9? (JBCTS-747) When commiting code in VFS, we need VFS jira tags, not JBCTS ones, otherwise, it's impossible to track. I want to mark Beta9 as released.	0	0	https://issues.jboss.org/browse/JBVFS-12
13205	JBVFS-8	Hi Scott, Can this one be closed?	1	1	https://issues.jboss.org/browse/JBVFS-8
13206	JBWEB-292	Thanks for the hint Remy, I can see it now in the docs on [cipher-suite|http://docs.jboss.org/jbossweb/7.0.x/config/ssl.html]. Does it mean that the configuration syntax changed between {{jbossweb-7.3.0.Final}} and {{jbossweb-7.4.0.Beta4}} or that it has been so all along and the colon-delimited cypher-suite was simply ignored in {{jbossweb-7.3.0.Final}}? Anyhow, I tried it with: {{cipher-suite="AES128-SHA,ALL,!ADH,!LOW,!MD5,!SSLV2,!NULL"}} and I must say that the error persists: {{JBWEB002081: No cipher match}}.	1	0	https://issues.jboss.org/browse/JBWEB-292
13207	JBWEB-292	Yes, JSSE names did it, this works: {code} <subsystem xmlns="urn:jboss:domain:web:1.5" native="false"> <connector name="https" protocol="HTTP/1.1" socket-binding="https" scheme="https" enabled="true" secure="true"> <ssl name="https" ca-certificate-file="/proper/ca-cert.jks" certificate-key-file="/proper/client-cert-key.jks" certificate-file="/proper/client-cert-key.jks" password="tomcat" verify-client="false" key-alias="javaclient" cipher-suite="TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384,!TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,!TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,!TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256,!TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256,TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,!TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384,TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384,TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA,TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA,TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384,TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384,TLS_ECDH_RSA_WITH_AES_256_CBC_SHA,TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA,TLS_DHE_RSA_WITH_AES_256_CBC_SHA256,TLS_DHE_RSA_WITH_AES_256_CBC_SHA,!TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256,!TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256,!TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA,!TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA,!TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256,!TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256,!TLS_ECDH_RSA_WITH_AES_128_CBC_SHA,!TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA,!TLS_DHE_RSA_WITH_AES_128_CBC_SHA256,!TLS_DHE_RSA_WITH_AES_128_CBC_SHA,TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA,TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA,TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA,TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA,TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA,TLS_RSA_WITH_AES_256_GCM_SHA384,TLS_RSA_WITH_AES_256_CBC_SHA256,TLS_RSA_WITH_AES_256_CBC_SHA,!TLS_RSA_WITH_AES_128_GCM_SHA256,!TLS_RSA_WITH_AES_128_CBC_SHA256,!TLS_RSA_WITH_AES_128_CBC_SHA,TLS_RSA_WITH_3DES_EDE_CBC_SHA" protocol="TLS"/> </connector> <virtual-server name="default-host" enable-welcome-root="true"> <alias name="localhost"/> <alias name="example.com"/> </virtual-server> </subsystem> {code} So, is this just a documentation issue then?	1	0	https://issues.jboss.org/browse/JBWEB-292
13208	JBWEB-276	This is a potential patch which adds a new error message (is using one past the last number correct?) and uses it when a RejectedExceutionException is caught instead of the normal message. It alters the NIO, APR and JIO endpoints	1	1	https://issues.jboss.org/browse/JBWEB-276
13209	JBWEB-276	James Livingston <jlivings@redhat.com> made a comment on [bug 985191|https://bugzilla.redhat.com/show_bug.cgi?id=985191] For the default blocking thread pools, a RejectedExecutionException means a serious problem occurred, yes. If you switch the web container over to using a non-blocking thread pool however, so that it drops requests rather than blocking, it is expected that you would receive these. Currently if you want to make JBoss drop web requests when there are too many, it writes a large number of these errors out into the log, which would have performance implications reducing the usefulness of dropping request under excessive load. If you don't want to hide the message since it is serious for blocking thread pools, is there an alternative which would let people using non-blocking thread pools hide it? It would be possible to move it to a different logging category, which could be turned off, but adding a new logging category for one message seems excessive.	1	1	https://issues.jboss.org/browse/JBWEB-276
13210	JBWEB-276	Paul Gier <pgier@redhat.com> made a comment on [bug 985191|https://bugzilla.redhat.com/show_bug.cgi?id=985191] Can this issue be closed?	0	0	https://issues.jboss.org/browse/JBWEB-276
13211	JBWEB-248	where do you put the valve?	0	0	https://issues.jboss.org/browse/JBWEB-248
13212	JBWEB-241	I have updated http://docs.jboss.org/jbossweb/7.0.x/sysprops.html. Was that what you were looking for?	1	0	https://issues.jboss.org/browse/JBWEB-241
13213	JBWEB-241	I guess that is OK. Thanks Jean-Fredric	0	0	https://issues.jboss.org/browse/JBWEB-241
13214	JBWEB-230	Can you please point me at a source file or a doc page that shows how JSF integrates servlet filters?	0	0	https://issues.jboss.org/browse/JBWEB-230

#	Requirement	Questions	Answered	Successfully Answered	URL
13215	JBWEB-228	OK. I think I see your point. To summarize: your proposal would be to implement a JBoss AS 7 extension (subsystem) for loading the authenticator valve and adding it to the valve chain. Right? I personally think this is much better than configuring it per app in the jboss-web.xml file (because it would not need an additional configuration on application developer/deployer side, besides application developer/deployers don't have to care where to get the classes from). IMO I still think this approach would have 2 drawbacks (when comparing it with the possibility of the registration of the Authenticator in the file Authenticator.properties file): * Developer/Deployer cannot configure the Authenticator by means of web.xml. It's even not possible to see by means of the web.xml file which Authenticator is really in use. As far as I know if a custom Authenticator is configured as a valve, JBoss will always take the valve configuration and will ignore the configuration in the web.xml file. This could be tricky (somehow magic) in terms of support and maintenance * The valve cannot be enabled/disabled by means of a deployment. It will be either always enabled or disabled (if we control this by means of a system property). For instance it is not possible to have 2 different War files deployed in one JBoss where one needs my custom Authenticator and the other one needs a Java EE Standard Authenticator. -> Using the jboss-web.xml approach per deployment this requirement can be solved (coming along with the drawback of a more tricky configuration (and class loading). I don't understand why JBoss cannot just provide the possibility to plugin custom authenticators and the required configuration for them (including the mapping to a auth-method usable in the web.xml file) in the standard subsystem web. From my point of view a custom Authenticator is different from a normal valve in terms of it should support and act the same way as the standard Java EE authenticators. And I don't see a way to achieve that when just treating it like a normal valve (neither when using jboss-web.xml file to configure it, nor using a JBoss Extension subsystem). To detail the requirement - The authenticator should be made available globally (just like the standard Java EE Authenticators are made available by the container) - The authenticator should be configurable / enabled per deployment (the same way like it is for a standard Java EE Authenticator) - And best would be that the configuration mechanism in use is the same than it would be for Java EE standard authenticators (using a custom auth-method in the web.xml file)	0	0	https://issues.jboss.org/browse/JBWEB-228
13216	JBWEB-178	@Remy Maucherat Thank you for your very interesting answer. Can you add in Jboss Web FAQ at http://www.jboss.org/jbossweb/faq.html that : "Jboss Web, the web container in JBoss is no longer compatible with Tomcat. All application vendor that sell web app certified for Tomcat would have to include specific JBoss AS versions in their certification process." It is very important for customers that want that their software editor add new support for Jboss Web. And can you add some words that explain why now Jboss Web, the web container (JSP/Servlet) in JBoss, is no longer compatible with Tomcat. What are the big problems or changes for this compatibility problem ? There is only few information about Comet Advanced IO, Tomcat Versions, NIO but is not very detailed for IT professionnals. Can you update Jboss Web FAQ at http://www.jboss.org/jbossweb/faq.html because then now this line is not truth: "...The latest version of JBoss Web corresponds to the latest version of Tomcat for the patches and security fixes." Thanks for your help.	1	0	https://issues.jboss.org/browse/JBWEB-178
13217	JBWEB-178	Can you add more information in the Jboss Web's FAQ. See https://jira.jboss.org/browse/JBWEB-178?focusedCommentId=12547642&page=com.atlassian.jira.plugin.system.issuetabpanels%3Acomment-tabpanel#action_12547642 Thanks.	1	0	https://issues.jboss.org/browse/JBWEB-178
13218	JBWEB-164	Jean-Frederic, So you are saying that this functionality exists in JBossWeb trunk, is that right? Any idea when the trunk code will make it into a JBoss release?	0	0	https://issues.jboss.org/browse/JBWEB-164
13219	JBWEB-164	The link in the description is broken... My notes only tell about replaceProperties() in the TC7 repo. Could you please put the new link in a comment?	1	0	https://issues.jboss.org/browse/JBWEB-164
13220	JBWEB-142	Since Remy is on vacation, can you please apply this upstream?	1	1	https://issues.jboss.org/browse/JBWEB-142
13221	JBWEB-8	Can someone review the article that is at: http://www.apache.org/~mturk/docs/article/ftwai.html In the article I have focused on determining the best topology and determining the actual configuration to use. Would you like to see something else in the article? Thanks, Mladen	1	1	https://issues.jboss.org/browse/JBWEB-8
13222	JBWEB-8	Finished the article. What's left is publishing	0	0	https://issues.jboss.org/browse/JBWEB-8
13223	JBWEB-5	Released jk-1.2.7-beta-2. This is the last beta release I hope. The next one will be 1.2.8	0	0	https://issues.jboss.org/browse/JBWEB-5
13224	JWS-151	[~fgoldefu], is there any specific configuration is required for monitoring mod_cluster connections? I was under the impression that monitoring apache through mod_bmx would suffice (I could be totally wrong though).	1	1	https://issues.jboss.org/browse/JWS-151
13225	JBWS-4056	Sorry about that, I missed that usage since they are looked up by FQDN class name in String. Anyway I sent a proposal PR https://github.com/jbossws/jbossws-cxf/pull/23, could you have a look?	0	0	https://issues.jboss.org/browse/JBWS-4056
13226	JBWS-4000	[~david.boeren], is it possible to see (privately is fine too) the code for WebPaymentInternal?	0	0	https://issues.jboss.org/browse/JBWS-4000
13227	JBWS-3943	Is this completed / done?	0	0	https://issues.jboss.org/browse/JBWS-3943
13228	JBWS-3916	This is fixed in wildfly9.0.0.CR2. I exclude this test in wildfly 8.x containers.	0	0	https://issues.jboss.org/browse/JBWS-3916
13229	JBWS-3878	I built jbossws using JDK 1.8.0_40. I saw only 2 groups of deprecation WARNINGS. 1.) Logging annotions in org/jboss/wsf/stack/cxf/jaspi/log/Loggers.java I checked in a change for this (Committed revision 19595) 2) deprecated classes org.jboss.wsf.stack.cxf.client.configuration.JBossWSSpringConfigurer org.jboss.wsf.stack.cxf.client.configuration.JBossWSSpringBusFactory which I have not addresses yet. How did you generate the javadoc warnings?	1	1	https://issues.jboss.org/browse/JBWS-3878

#	Requirement	Questions	Answered	Successfully Answered	URL
13230	JBWS-3832	Although it succeeds, I see the followinf exception on screen. {code} Concurrency config is parallel='classes', perCoreThreadCount=true, threadCount=2, useUnlimitedThreads=false org.jboss.test.ws.jaxws.cxf.in_container_client.CustomBusServletTestCase(org.jboss.test.ws.jaxws.cxf.in_container_client.CustomBusServletTestCase$3)java.lang.Exception: "JBAS014803: Duplicate resource [(\"system-property\" => \"org.jboss.ws.cxf.jaxws-client.bus.selector\")]" at org.jboss.as.webservices.deployer.RemoteDeployer.checkResult(RemoteDeployer.java:470) at org.jboss.as.webservices.deployer.RemoteDeployer.applyUpdate(RemoteDeployer.java:457) at org.jboss.as.webservices.deployer.RemoteDeployer.setSystemProperty(RemoteDeployer.java:442) at org.jboss.wsf.test.JBossWSTestHelper.setSystemProperty(JBossWSTestHelper.java:396) at org.jboss.test.ws.jaxws.cxf.in_container_client.CustomBusServletTestCase$3.setUp(CustomBusServletTestCase.java:85) at org.jboss.wsf.test.JBossWSTestSetup$1.protect(JBossWSTestSetup.java:140) at junit.framework.TestResult.runProtected(TestResult.java:128) at org.jboss.wsf.test.JBossWSTestSetup.run(JBossWSTestSetup.java:148) at org.junit.internal.runners.JUnit38ClassRunner.run(JUnit38ClassRunner.java:83) at org.junit.runners.Suite.runChild(Suite.java:128) at org.junit.runners.Suite.runChild(Suite.java:24) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:471) at java.util.concurrent.FutureTask.run(FutureTask.java:262) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at java.lang.Thread.run(Thread.java:745) Running org.jboss.test.ws.jaxws.cxf.in_container_client.ServletTestCase {code}	1	1	https://issues.jboss.org/browse/JBWS-3832
13231	JBWS-3797	Jim, can this be marked as solved?	0	0	https://issues.jboss.org/browse/JBWS-3797
13232	JBWS-3750	I don't think this enhancement is going to function the way you intend. The format for a URL/URI is [scheme:][//[host]:[port]][path][?query][#fragment] If I have the following services declared in a wsdl http://localhost:8080/jaxws-jbws2150/ValidURL http://localhost:8080/jaxws-jbws2150/InvalidURL then there are 2 unique paths, "jaxws-jbws2150/ValidURL" and "jaxws-jbws2150/InvalidURL". This enhancement replaces each path with a single custom path. If I set the custom path to "/my-custom-path/foo/bar" the 2 addresses above become http://localhost:8080/my-custom-path/foo/bar http://localhost:8080/my-custom-path/foo/bar Is this what is intended?	1	1	https://issues.jboss.org/browse/JBWS-3750
13233	JBWS-3738	/** * The user's name. It is used differently by each of the WS-Security functions: * a) It is used as the name in the UsernameToken * b) It is used as the alias name in the keystore to get the user's cert and private key for signature * if {@link SIGNATURE_USERNAME} is not set. * c) It is used as the alias name in the keystore to get the user's public key for encryption if * {@link ENCRYPT_USERNAME} is not set. */ public static final String USERNAME = "ws-security.username"; /** * The user's password when a {@link CALLBACK_HANDLER} is not defined. It is currently only used for * the case of adding a password to a UsernameToken. */ public static final String PASSWORD = "ws-security.password"; /** * The user's name for signature. It is used as the alias name in the keystore to get the user's cert * and private key for signature. - If this is not defined, then {@link USERNAME} is used instead. - If that is also not specified, it uses the the default alias set in the properties file referenced by {@link SIGNATURE_PROPERTIES}. - If that's also not set, and the keystore only contains a single key, that key will be used. */ public static final String SIGNATURE_USERNAME = "ws-security.signature.username"; /** * The user's name for encryption. It is used as the alias name in the keystore to get the user's public * key for encryption. If this is not defined, then {@link USERNAME} is used instead. If * that is also not specified, it uses the the default alias set in the properties file referenced by * {@link ENCRYPT_PROPERTIES}. If that's also not set, and the keystore only contains a single key, * that key will be used. * * For the web service provider, the "useReqSigCert" keyword can be used to accept (encrypt to) any * client whose public key is in the service's truststore (defined in {@link ENCRYPT_PROPERTIES}). */ public static final String ENCRYPT_USERNAME = "ws-security.encryption.username"; provide test for each - public static final String CALLBACK_HANDLER = "ws-security.callback-handler"; - public static final String SAML_CALLBACK_HANDLER = "ws-security.saml-callback-handler"; /** * The Crypto property configuration to use for signature, if {@link SIGNATURE_CRYPTO} is not set instead. * The value of this tag must be either: * a) A Java Properties object that contains the Crypto configuration. * b) The path of the Crypto property file that contains the Crypto configuration. * c) A URL that points to the Crypto property file that contains the Crypto configuration. */ public static final String SIGNATURE_PROPERTIES = "ws-security.signature.properties"; /** * The Crypto property configuration to use for encryption, if {@link ENCRYPT_CRYPTO} is not set instead. * The value of this tag must be either: * a) A Java Properties object that contains the Crypto configuration. * b) The path of the Crypto property file that contains the Crypto configuration. * c) A URL that points to the Crypto property file that contains the Crypto configuration. */ public static final String ENCRYPT_PROPERTIES = "ws-security.encryption.properties"; /** * A Crypto object to be used for signature. If this is not defined then the * {@link SIGNATURE_PROPERTIES} is used instead. */ public static final String SIGNATURE_CRYPTO = "ws-security.signature.crypto"; /** * A Crypto object to be used for encryption. If this is not defined then the * {@link ENCRYPT_PROPERTIES} is used instead. */ public static final String ENCRYPT_CRYPTO = "ws-security.encryption.crypto"; /** * Whether to enable Certificate Revocation List (CRL) checking or not when verifying trust * in a certificate. The default value is "false". */ public static final String ENABLE_REVOCATION = "ws-security.enableRevocation"; /** * The Subject Role Classifier to use. If one of the WSS4J Validators returns a JAAS Subject * from Validation, then the WSS4JInInterceptor will attempt to create a SecurityContext * based on this Subject. If this value is not specified, then it tries to get roles using * the DefaultSecurityContext in cxf-rt-core. Otherwise it uses this value in combination * with the SUBJECT_ROLE_CLASSIFIER_TYPE to get the roles from the Subject. */ public static final String SUBJECT_ROLE_CLASSIFIER = "ws-security.role.classifier"; /**	1	1	https://issues.jboss.org/browse/JBWS-3738
13234	JBWS-3738	What type of documentation is being sought? Where is it desired to reside?	1	1	https://issues.jboss.org/browse/JBWS-3738
13235	JBWS-3724	The option you mention was discarded because in HTTP 1.0 new connections are open for each request/response pair. And after each response the connection would be closed. This lead to some big efficiency problems. I asked to several people and they do not recommend do this. I disabled Chunked encoding in client side, over HTTP 1.1 and i was hoping to do the same in server side, keeping HTTP 1.1 because it means that chunked transfer encoding is an option in HTTP 1.1	0	0	https://issues.jboss.org/browse/JBWS-3724
13236	JBWS-3648	Alessio, would you check this issues? https://community.jboss.org/thread/231657 Invoking @PolicySets-added Service Endpoint Interface is failed. I suggest to make another JIRA to solve this issues. I need your advice deperately	1	1	https://issues.jboss.org/browse/JBWS-3648

#	Requirement	Questions	Answered	Successfully Answered	URL
13237	JBWS-3628	Rebecca, sure, the client has power to change this specific value, as well as using a completely different wsdl contract. The JMS case here is probably also not a perfect fit for the new feature request, as for pure JMS endpoints the wsdl contract is actually not published at all (that's also why in the testcase you mentioned the client is retrieving the wsdl from as a resource in the client jar). This said, it could still be useful to support a server side property expansion mechanism; the user could deploy a contract first endpoint with something like @foo@ in its wsdl, then the jbossws deployer would expand that by reading the value of the "foo" system property when publishing the endpoint wsdl. Do you see what I mean?	1	1	https://issues.jboss.org/browse/JBWS-3628
13238	JBWS-3628	OK, I see. So we can't "fool" CXF into using a different location for the wsdl. Btw, I also remember that CXF has a wsdl cache mechanism that is based on the wsdl url, so playing with the wsdl location url might definitely be risky. The idea of running a filter at the time CXF parses the wsdl into the WSDL4J Definition model is interesting; something I've noticed in the patch is that is applies to the JMS wsdl extensibility element only. So we'd need to clarify in advance which elements actually support property expansion and which do not. Something I've been thinking about is why don't we run the filter directly at WSDL4J level? javax.wsdl.factory.WSDLFactory (which is used by CXF to parse the wsdl) lets us have a different implementation of WSDLFactory (see newInstance() and findFactoryImplName() methods). We could have a custom WSDLFactory impl that runs our filters and eventually delegates to the default WSDLFactory impl (or to what would have been resolved before, but that's an impl detail). The filtering could either be performed through pluggable filters as you did in your last patch for cxf, or (better, I think) be StAX / DOM based on the whole document (depends on which readWSDL method flavor is called). The advantage are no need to have changes in CXF and a solution that applies to any component relying on the WSDLReader.	1	1	https://issues.jboss.org/browse/JBWS-3628
13239	JBWS-3628	CXF would need only a minor modification in order for us to register an custom WSDLFactoryImpl through the bus for CXF's WSDLManager to reference. Our WSDLFactoryImpl class would override the newWSDLReader method and return a custom WSDLReaderImpl class. It would be most desirable to subclass com.ibm.wsdl.xml.WSDLReaderImpl and override select parsing methods, however is this not possible because the methods of interest are "protected". I see two solutions. One, copy the contents of com.ibm.wsdl.xml.WSDLReaderImpl into our WSDLReaderImpl class and make the desired modifications to the protected methods. Two, in our WSDLFactoryImpl override method newDefinition() and return a custom impl of DefinitionImpl and in DefinitionImpl override methods like createService() and return a custom ServiceImpl class (and the like) that does our bidding. This would allow com.ibm.wsdl.xml.WSDLReaderImpl to be used as is. Because of the nature of what we are currently customizing I am more in favor of option one. This provides us full access to the wsdl contents with the smallest number of classes that needs to be maintained. Option two seems like over kill and requires a larger number of custom classes to be maintained by us. What are your thoughts?	1	1	https://issues.jboss.org/browse/JBWS-3628
13240	JBWS-3573	I re-executed Coverity analysis, there are 2 new defects detected: - JBossWS SPI - TRANSIENT_FIELD_NOT_RESTORED: probably design decession - http://station5.brq.redhat.com:8080/sourcebrowser.htm?projectId=10001#mergedDefectId=10219&defectInstanceId=1283&fileInstanceId=2332 - JBossWS Native - DM_DEFAULT_ENCODING: shouldn't be UTF-8 enforced ? - http://station5.brq.redhat.com:8080/sourcebrowser.htm?projectId=10004#mergedDefectId=10220&defectInstanceId=1364&fileInstanceId=3078 JBossWS -- api and spi items updated based on Richard's comment - low priority items and tests marked as Intentional Overall Status: http://station5.brq.redhat.com:8080/projects/index.htm?tableKey=ProjectsTable Please could you check 2 new items? Are defects for JBossWS common and JBossWS CXF section intentional? I would like to add classification to this items to skip them in future reports.	0	0	https://issues.jboss.org/browse/JBWS-3573
13241	JBWS-3549	Michael, sorry I've spotted this just now. Can you please provide more info? It's not clear how to reproduce given what you wrote above. Thanks	0	0	https://issues.jboss.org/browse/JBWS-3549
13242	JBWS-3435	How appclient tests work? ---------- 1) We're providing appclient-ws.xml configuration to configure JBoss remoting to allow for remote access to appclient JNDI env. 2) Our JAX-RPC proxies are de/serializable to be accessible remotely. 3) Our tests are always calling appclient.sh(.bat) scripts internally via java.lang.Process and passing --appclient-config=appclient-ws.xml parameter to use our custom config 4) Every appclient archive must be nested in ear, so we're using foo-appclient.ear#foo-appclient.jar syntax in our JBossWS test helper to specify which ear contains what appclient jar. 5) We're registering AppclientKiller for every appclient.jar that contains no main-class MANIFEST.MF entry. The goal of this class is to detect shutdown mark files and call System.exit when mark file is found.	0	0	https://issues.jboss.org/browse/JBWS-3435
13243	JBWS-3409	Are you able to reproduce this problem against 7.1.0.CR1 ? http://www.jboss.org/jbossas/downloads	1	1	https://issues.jboss.org/browse/JBWS-3409
13244	JBWS-3259	Can you please provide the war that's being deployed when the exception is thrown?	0	0	https://issues.jboss.org/browse/JBWS-3259
13245	JBWS-3255	I experience same using JBossAS-5.1.0,GA and JBossWS-native 3.4.0. I simply drop the @EndpointConfig(configName = "Standard WSSecurity Endpoint") annotation into my Port Impl class and I get the same errors. If I understand Marcus Moyses's comment above correctly this means I need to create my own version of org.jboss.ws.extensions.security.jaxws.WSSecurityHandlerServer since the one provided is tied to a JNDI name that is no longer valid as of 5.1.0 if not earlier. Is that right? Or is another workaround available on 5.1.0?	1	0	https://issues.jboss.org/browse/JBWS-3255
13246	JBWS-3255	Thanks Alessio, but my point is that it DOESN'T work with Native 3.4,0 on top of AS 5.1.0 in spite of your assertions that it SHOULD. Have you actually tried it? For one thing, although the documentation included with the 3.4.0 Test Suite talks about a JBossWS-Security And Attachments Sample supposedly located in the org.jboss.test.ws.jaxws.samples.news package, this package does not in fact exist in the 3.4.0 Test Suite and there is no source code in the 3.4.0 Test Suite examples that contains the @EndpointConfig annotation, although you can find it in the included documentation, which, however, does not seem to pertain to 3.4.0. If 3.4.0 could do this, there should be a sample illustrating how. Maybe there needs to be a 3.4.0a? So I must ask: is it even possible to do WS Security with 3.4.0 over 5.1.0? And if so, how? And, as we know from http://community.jboss.org/wiki/JBossWS-SupportedTargetContainers, 3.4.1 is not compatible with 5.1.0.	1	0	https://issues.jboss.org/browse/JBWS-3255
13247	JBWS-3185	Jim, the CXF jiras mentioned above are fixed in the 2.3.4 version we're using, can you please verify this issue is fixed and mark it as resolved in that case? Thanks!	1	1	https://issues.jboss.org/browse/JBWS-3185

#	Requirement	Questions	Answered	Successfully Answered	URL
13248	JBWS-3115	This is my TODO list related to this issue: Our SPI have dependencies on jboss-logging, jbossxb and jboss-kernel, see: [/home/opalka][/home/opalka/svn/jbossws/spi/trunk]>grep -r "import org.jboss" * | grep -v "\.svn" | grep -v "jboss\.wsf" src/main/java/org/jboss/wsf/spi/deployment/DeploymentAspect.java:import org.jboss.logging.Logger; src/main/java/org/jboss/wsf/spi/metadata/j2ee/serviceref/UnifiedServiceRefMetaData.java:import org.jboss.logging.Logger; src/main/java/org/jboss/wsf/spi/metadata/j2ee/serviceref/HandlerChainsObjectFactory.java:import org.jboss.logging.Logger; src/main/java/org/jboss/wsf/spi/metadata/j2ee/serviceref/HandlerChainsObjectFactory.java:import org.jboss.xb.binding.ObjectModelFactory; src/main/java/org/jboss/wsf/spi/metadata/j2ee/serviceref/HandlerChainsObjectFactory.java:import org.jboss.xb.binding.UnmarshallingContext; src/main/java/org/jboss/wsf/spi/metadata/webservices/PortComponentMetaData.java:import org.jboss.logging.Logger; src/main/java/org/jboss/wsf/spi/metadata/webservices/WebservicesFactory.java:import org.jboss.logging.Logger; src/main/java/org/jboss/wsf/spi/metadata/webservices/WebservicesFactory.java:import org.jboss.xb.binding.ObjectModelFactory; src/main/java/org/jboss/wsf/spi/metadata/webservices/WebservicesFactory.java:import org.jboss.xb.binding.Unmarshaller; src/main/java/org/jboss/wsf/spi/metadata/webservices/WebservicesFactory.java:import org.jboss.xb.binding.UnmarshallerFactory; src/main/java/org/jboss/wsf/spi/metadata/webservices/WebservicesFactory.java:import org.jboss.xb.binding.UnmarshallingContext; src/main/java/org/jboss/wsf/spi/metadata/webservices/WebserviceDescriptionMetaData.java:import org.jboss.logging.Logger; src/main/java/org/jboss/wsf/spi/annotation/AuthMethod.java:import org.jboss.logging.Logger; src/main/java/org/jboss/wsf/spi/management/recording/RecordGroupAssociation.java:import org.jboss.logging.Logger; src/main/java/org/jboss/wsf/spi/serviceref/ServiceRefHandler.java:import org.jboss.xb.binding.UnmarshallingContext; src/main/java/org/jboss/wsf/spi/util/KernelLocator.java:import org.jboss.kernel.Kernel; These dependencies must be removed! To help other J2EE servers developers implement our SPI. Kernel dependency have to be moved to AS IL. Remove dependency on logging + jbossxb = evaluate how! Refactor to AS IL? We should create logging SPI facade to delegate to concrete impl.	1	1	https://issues.jboss.org/browse/JBWS-3115
13249	JBWS-3084	Is it possible to use Dispatch<SOAPMessage> instead of SOAPConnection ? The analysis shows that due to a number of restrictions in SAAJ API (not possible to get an access to the connection headers used by default SOAPConnection instances, even by relying on the custom SOAPConnectionFactories, explicitly added Content-Length header to SOAPMessage.getMimeHeaders() is ignored on the client side), as well as the disconnect between bare SOAPConnections and stack-specific chains, using Dispatch<SOAPMessage> and setting a property on the SOAPMessage which will be recognized by stack-specific HTTP transport modules is probably the only realistic approach.	1	1	https://issues.jboss.org/browse/JBWS-3084
13250	JBWS-3084	Thanks Sergey; can you please comment here with the jira number of the required CXF changes? That might be of use when we will need to port this fix to EAP branches. Thanks.	1	1	https://issues.jboss.org/browse/JBWS-3084
13251	JBWS-3011	If a change is required that will affect all clients or all endpoint why not just edit the "Standard Client" or "Standard Endpoint" definitions directly? Why does a new configuration need to be added then a system property set to specify the new configuration is the new default leaving the existing defaults unused.	1	0	https://issues.jboss.org/browse/JBWS-3011
13252	JBWS-2935	java.lang.RuntimeException: Unable to inject jndi dependency: env/service/wsejbmultipleclientinjectiontest1 into property com.sun.ts.tests.webservices12.ejb.annotations.WSEjbMultipleClientInjectionTest1.Client.service: InputSource must have a ByteStream or CharacterStream at org.jboss.injection.JndiPropertyInjector.lookup(JndiPropertyInjector.java:82) at org.jboss.injection.JndiPropertyInjector.inject(JndiPropertyInjector.java:99) at org.jboss.ejb3.client.ClientContainer.<init>(ClientContainer.java:130) at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:39) at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:27) at java.lang.reflect.Constructor.newInstance(Constructor.java:513) at org.jboss.reflect.plugins.introspection.ReflectionUtils.newInstance(ReflectionUtils.java:149) at org.jboss.reflect.plugins.introspection.ReflectConstructorInfoImpl.newInstance(ReflectConstructorInfoImpl.java:106) at org.jboss.joinpoint.plugins.BasicConstructorJoinPoint.dispatch(BasicConstructorJoinPoint.java:80) at org.jboss.kernel.plugins.dependency.DispatchJoinPoint.run(DispatchJoinPoint.java:47) at java.security.AccessController.doPrivileged(Native Method) at org.jboss.kernel.plugins.dependency.ExecutionWrapper.execute(ExecutionWrapper.java:54) at org.jboss.kernel.plugins.dependency.KernelControllerContextAction.dispatchExecutionWrapper(KernelControllerContextAction.java:109) at org.jboss.kernel.plugins.dependency.KernelControllerContextAction.dispatchJoinPoint(KernelControllerContextAction.java:70) at org.jboss.kernel.plugins.dependency.InstantiateAction.installActionInternal(InstantiateAction.java:66) at org.jboss.kernel.plugins.dependency.InstallsAwareAction.installAction(InstallsAwareAction.java:54) at org.jboss.kernel.plugins.dependency.InstallsAwareAction.installAction(InstallsAwareAction.java:42) at org.jboss.dependency.plugins.action.SimpleControllerContextAction.secureInstallAction(SimpleControllerContextAction.java:67) at org.jboss.dependency.plugins.action.AccessControllerContextAction$1.run(AccessControllerContextAction.java:80) at java.security.AccessController.doPrivileged(Native Method) at org.jboss.dependency.plugins.action.AccessControllerContextAction.install(AccessControllerContextAction.java:103) at org.jboss.dependency.plugins.AbstractControllerContextActions.install(AbstractControllerContextActions.java:51) at org.jboss.dependency.plugins.AbstractControllerContext.install(AbstractControllerContext.java:348) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:1631) at org.jboss.dependency.plugins.AbstractController.incrementState(AbstractController.java:934) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:1082) at org.jboss.dependency.plugins.AbstractController.resolveContexts(AbstractController.java:984) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:774) at org.jboss.dependency.plugins.AbstractController.install(AbstractController.java:540) at	0	0	https://issues.jboss.org/browse/JBWS-2935
13253	JBWS-2818	Removing the declaration "<executable>${java.jdk15.home}/bin/javac</executable> " makes the problem go away. Where is the java.jdk15.home variable set?	1	1	https://issues.jboss.org/browse/JBWS-2818
13254	JBWS-2782	Maybe asking the helpdesk team would be enough to achieve that. Who other could handle this? I don't hink each team configures it on its own.	1	0	https://issues.jboss.org/browse/JBWS-2782

#	Requirement	Questions	Answered	Successfully Answered	URL
13255	JBWS-2752	Alessio, could you back port this issue fix to JBossWS EAP 5 branch, please? JBIDE-4257 depends on it and it's really critical for them to have it there (I discussed this issue with Denny via IRC)	1	0	https://issues.jboss.org/browse/JBWS-2752
13256	JBWS-2751	Hi Stefan, I've just committed changes taken from your sources. The project has already been "mavenized" by me in the mean time and you can find the artifacts on JBoss Maven repositories: http://repository.jboss.org/maven2/jboss/jaxbintros/ http://snapshots.jboss.org/maven2/jboss/jaxbintros/ This said, I'd really appreciate if you could give the latest snapshot (1.0.2-SNAPSHOT) a try, or better update sources and try if everything is working as expected. Btw, is there any other annotation missing that you want to add?	1	1	https://issues.jboss.org/browse/JBWS-2751
13257	JBWS-2741	What have been done: * jbossws-native-jaxws.jar became manifest only file * some useless classes from jbossws-native-jaxws-ext.jar have been removed. * there were added jaxws-api.jar and jsr181-api.jar files to our distributions (APIs)	1	1	https://issues.jboss.org/browse/JBWS-2741
13258	JBWS-2724	Thanks for the patch Hans. BTW, how is the interoperability testing going?	0	0	https://issues.jboss.org/browse/JBWS-2724
13259	JBWS-2723	Thanks for the patch Hans. BTW, how is the interoperability testing going?	1	1	https://issues.jboss.org/browse/JBWS-2723
13260	JBWS-2694	I've always seen that to disable chunking encoding is always related to client configuration. How can it be disabled from server side? What configuration file can be changed in JBoss or JBossWS to do that?	0	0	https://issues.jboss.org/browse/JBWS-2694
13261	JBWS-2685	I am able to run the tests and get the same results with both the Sun 1.6 JDK and OpenJDK: ========== Failed tests: test(org.jboss.test.ws.jaxws.samples.dar.JMSClientTestCase) Tests in error: org.jboss.test.ws.jaxws.samples.dar.ClientTestCase Tests run: 28, Failures: 1, Errors: 1, Skipped: 0 ========== This is after BUILDING with Sun and RUNNING with Sun and OpenJDK. Alessio - can you recheck this? Thx...	1	0	https://issues.jboss.org/browse/JBWS-2685
13262	JBWS-2680	Jim, your PR has been merged; is it possible to have the same fix applied to the 800 an 810 ASIL too?	0	0	https://issues.jboss.org/browse/JBWS-2680
13263	JBWS-2679	Is this issue meant to address compression at the transport level through "Content-Encoding: gzip"?	1	0	https://issues.jboss.org/browse/JBWS-2679
13264	JBWS-2647	Paul, what scripts do you actually refer to? Can you give further details? Currently the jboss-xml-binding.jar is referenced by the testsuite, the jbossws-client manifest and the ws tools scripts. This said, JBossWS Native uses the JBossXB jar provided by the AS. Currently that's jboss-xml-binding.jar. Are you saying there're plans for having jbossxb.jar in AS client dir instead of jboss-xml-binding.jar ?	1	1	https://issues.jboss.org/browse/JBWS-2647
13265	JBWS-2647	OK, just created JBAS-7441 for fixing the AS build first, then we'll fix the WS stuff accordingly. Fine with you?	1	1	https://issues.jboss.org/browse/JBWS-2647
13266	JBWS-2647	If I change the references in the AS build first, won't that break some tests? It seems like this issue should be resolved first, and then when AS updates to jbossws-native 3.3.2 the AS build scripts can be changed.	1	1	https://issues.jboss.org/browse/JBWS-2647
13267	JBWS-2554	Any news on this? Is Remoting 3 going to be included in AS 5.1.0.CR1? Which version should I start looking at exactly? And can I install that version to previous AS releases?	1	1	https://issues.jboss.org/browse/JBWS-2554
13268	JBWS-2554	The community time-line has been reduced quite a bit, so remoting 2 will definitely remain since jbm 1.4 depends on it. This should definitely be updated to another http client impl though, since you have WAY more flexibility (can support things like NTLM auth etc)	1	1	https://issues.jboss.org/browse/JBWS-2554
13269	JBWS-2554	Jason, should I read this as AS 5.1.0.GA will still have Remoting 2 ? Asking as this kind of change does not fit very well into a micro jbws release (3.1.1 we're working on) and we plan a 3.2.0 after that... Regarding feature and flexibility, I agree, perhaps we already have some of them with other stacks, but yes this is a chance for improvements.	1	1	https://issues.jboss.org/browse/JBWS-2554
13270	JBWS-2510	I wouldn't put any effort on this. AS doesn't care about it at all. Why have JBossWS internationalized and other AS parts not? Finally there should be i18n and l10n team inside RH that would take care of that IMHO.	1	0	https://issues.jboss.org/browse/JBWS-2510
13271	JBWS-2424	We're seeing some failures in jaxr TCK5 gui at the moment. There are listed jaxr tests that doesn't exist in commandline version. Note that commandline version of jaxr passes on 100%. I wrote to javaeetck@jboss.org mailing list and I'm waiting for response how is it possible nonexisting tests are displayed in TCK5 GUI version.	1	1	https://issues.jboss.org/browse/JBWS-2424
13272	JBWS-2397	The test have been removed: --- <ropalka> asoldano, regarding https://issues.jboss.org/browse/JBWS-2397 <ropalka> asoldano, this test was using the following native config <ropalka> --- <ropalka> <context-param> <ropalka> <param-name>jbossws-config-name</param-name> <ropalka> <param-value>.NET friendly Endpoint</param-value> <ropalka> </context-param> <ropalka> --- <ropalka> asoldano, and I oversaw that so I thought it can run against CXF too. <ropalka> asoldano, Native contained a lot of hacky fixes to make this .NET friendly endpoints work <ropalka> asoldano, this is not going to work on Apache CXF <ropalka> asoldano, I'm going to delete this test from shared-testsuite <ropalka> asoldano, any complains? Do you want to implement this feature at Apache CXF in near future? <asoldano> ropalka, fine with removing the test. We actually never run it with metro or cxf stack <asoldano> ropalka, the problem would really be another though <asoldano> ropalka, ie. verifying .net compatibility, but I assume that nowadays cxf is tested elsewhere wrt this topic	1	0	https://issues.jboss.org/browse/JBWS-2397
13273	JBWS-2337	Is there any specification mandating single or double quoted attributes for being used in these cases?	1	0	https://issues.jboss.org/browse/JBWS-2337
13274	JBWS-2324	Why postpone this issue constantly? It is a major performance problem?	0	0	https://issues.jboss.org/browse/JBWS-2324

#	Requirement	Questions	Answered	Successfully Answered	URL
13275	JBWS-2279	@HandlerChain on client SEI # [CXF-1261] Handler does not see XOP message properly # [CXF-1510] Cannot handle anonymous bare requests # [CXF-1511] WrappedMessageContext does not implement SOAPMessageContext # [CXF-1513] NPE in JaxWsServiceConfiguration # [CXF-1514] Generated WSDL does not take 'transport-guarantee' in web.xml into account # [CXF-1515] Bad Enumeration value 'extension restriction' # [CXF-1516] Type inheritance with document/literal/wrapped # [CXF-1517] HTTP bindings for Provider # [CXF-1518] Missing fault detail in messages from SLSB Provider # [CXF-1519] Explicitly set the namespace of a WebFault # [CXF-1520] MTOM chunked encoding transfer # [CXF-1521] Exception handling with @WebFault annotation # [CXF-1522] NPE in ServiceUtils # [CXF-1623] Wrong AnnotationHandlerChainBuilder.protocolMatches() method implementation or preconditions # [CXF-1624] XmlSchemaException: Schema name conflict in collection # [JBWS-981] @WebContext defined virtual hosts only being fixed for JBoss 5 # [JBWS-1683] Fix JAXR samples for CXF # [JBWS-1774] Provide tools implementation for CXF # [JBWS-2112] Cannot use CXF with endpoint API # [JBWS-2114] IllegalArgumentException: Wrong target. class # [JBWS-2115] EJB client vehicle not sufficiently isolated # [JBWS-2117] web.xml modified to web.xml.org - Subsequent runs fail # [JBWS-2118] context-root in jboss.xml is ignored # [JBWS-2119] WebServiceRef injection from DD doesn't work # [JBWS-2165] Fix xop tests with maven build # [JBWS-2223] wsconsume for cxf build # [JBWS-2227] Investigate why multiple virtual hosts test fails on CXF FIXME: [CXF-1257] Cannot create proxy for namespaced doc/wrapped FIXME: [CXF-1258] soap:Fault response does not cause SOAPFaultException FIXME: [CXF-1310] Generated WSDL for an WS-RM endpoint does not contain RM policies FIXME: [JBCTS-540] EJB3 proxy does not implement the Home interface FIXME: [JBWS-1655] Add support for endpoint address rewriting FIXME: [JBWS-1665] incorrect wsdl generation FIXME: [JBWS-2203] Implement CXF standard message context properties AS 423: # [CXF-1244] Application Exception not initialized properly # [CXF-1252] Provider PAYLOAD endpoint expects SOAP Envelope # [CXF-1253] CXF does not respect @HandlerChain on client SEI # [CXF-1261] Handler does not see XOP message properly # [CXF-1510] Cannot handle anonymous bare requests # [CXF-1511] WrappedMessageContext does not implement SOAPMessageContext # [CXF-1513] NPE in JaxWsServiceConfiguration # [CXF-1514] Generated WSDL does not take 'transport-guarantee' in web.xml into account # [CXF-1515] Bad Enumeration value 'extension restriction' # [CXF-1516] Type inheritance with document/literal/wrapped # [CXF-1517] HTTP bindings for Provider # [CXF-1518] Missing fault detail in messages from SLSB Provider # [CXF-1519] Explicitly set the namespace of a WebFault # [CXF-1520] MTOM chunked encoding transfer # [CXF-1521] Exception handling with @WebFault annotation # [CXF-1522] NPE in ServiceUtils # [CXF-1623] Wrong AnnotationHandlerChainBuilder.protocolMatches() method implementation or preconditions # [CXF-1624] XmlSchemaException: Schema name conflict in collection # [JBWS-981] @WebContext defined virtual hosts only being fixed for JBoss 5 # [JBWS-1683] Fix JAXR samples for CXF # [JBWS-1774] Provide tools implementation for CXF # [JBWS-2112] Cannot use CXF with endpoint API # [JBWS-2114] IllegalArgumentException: Wrong target. class # [JBWS-2115] EJB client vehicle not sufficiently isolated # [JBWS-2117] web.xml modified to web.xml.org - Subsequent runs fail # [JBWS-2118] context-root in jboss.xml is ignored # [JBWS-2119] WebServiceRef injection from DD doesn't work # [JBWS-2165] Fix xop tests with maven build # [JBWS-2223] wsconsume for cxf build # [JBWS-2227] Investigate why multiple virtual hosts test fails on CXF FIXME: [CXF-1257] Cannot create proxy for namespaced doc/wrapped FIXME: [CXF-1258] soap:Fault response does not cause SOAPFaultException FIXME: [CXF-1310] Generated WSDL for an WS-RM endpoint does not contain RM policies FIXME: [JBCTS-540] EJB3 proxy does not implement the Home interface FIXME: [JBWS-1655] Add support for endpoint address rewriting FIXME: [JBWS-1665] incorrect wsdl generation FIXME: [JBWS-2203] Implement CXF standard message context properties AS 500	1	1	https://issues.jboss.org/browse/JBWS-2279
13276	JBWS-2168	Could you please fix this as well? bpel: [unzip] Expanding: /home/tdiesler/svn/jbossws/stack/native/trunk/modules/testsuite/native-tests/target/bpel-artifacts/bpel.sar into /home/tdiesler/svn/jbossws/stack/native/trunk/modules/testsuite/native-tests/target/bpel-artifacts/jbpm-bpel log4j:ERROR setFile(null,false) call failed. java.io.FileNotFoundException: /test.log (Permission denied) at java.io.FileOutputStream.open(Native Method) at java.io.FileOutputStream.<init>(FileOutputStream.java:179) at java.io.FileOutputStream.<init>(FileOutputStream.java:102) at org.apache.log4j.FileAppender.setFile(FileAppender.java:289) at org.apache.log4j.FileAppender.activateOptions(FileAppender.java:163) at org.apache.log4j.DailyRollingFileAppender.activateOptions(DailyRollingFileAppender.java:215)	0	0	https://issues.jboss.org/browse/JBWS-2168
13277	JBWS-2151	Where is wsrunclient.sh for JBoss AS 7/Wildfly ?	1	0	https://issues.jboss.org/browse/JBWS-2151
13278	JBWS-2106	Hi Andrew, given the current jbossws-cxf leverages apache cxf 2.1.3, I've taken a look at its sources and think it has its own implementation of what you need. In particular I've successfully tried setting the FaultTo in an outgoing message with this: import static org.apache.cxf.ws.addressing.JAXWSAConstants.CLIENT_ADDRESSING_PROPERTIES; import org.apache.cxf.ws.addressing.AddressingProperties; import org.apache.cxf.ws.addressing.AddressingPropertiesImpl; import org.apache.cxf.wsdl.EndpointReferenceUtils; ... Service service = Service.create(wsdlURL, serviceName); proxy = (ServiceIface)service.getPort(ServiceIface.class); Map<String, Object> requestContext = ((BindingProvider)proxy).getRequestContext(); AddressingProperties maps = new AddressingPropertiesImpl(); maps.setFaultTo(EndpointReferenceUtils.getEndpointReference("http://localhost:9999/decoupled_endpoint")); requestContext.put(CLIENT_ADDRESSING_PROPERTIES, maps); proxy.sayHello(); Would this be fine for your usecase when using the cxf stack integration? Can you take a look? Thanks.	1	1	https://issues.jboss.org/browse/JBWS-2106
13279	JBWS-2106	Hi Alessio, Yes, cxf does implement the required AddressingProperties class but in a different package to the one used by Native, as are the other types provided by CXF. I probably could use their implementation to get my code to run on the CXF stack -- it looked to be pretty complete when I used it. The problem is how to use it compatibly with the Native JBossWS implementation. A lot of the methods in my current XTS build are currently parameterised with AddressingProperties instances or employ local values of this type. The same applies for the various dependent types employed in conjunction with AddressingProperties as AddressingBuilder, AttributedURI, Relationship, etc.. I need to be able to manipulate these types in order to construct faultto, replyto etc addresses. I don't really see how to generalize the code so that my code can operate over both stacks. Do you have any ideas how to resolve this?	1	1	https://issues.jboss.org/browse/JBWS-2106

#	Requirement	Questions	Answered	Successfully Answered	URL
13280	JBWS-2106	I have bypassed this problem by creating an abstraction layer in the XTS code which hides all uses of the AddressingProperties class and its close relatives. Alternative implementations for Native and CXF are included (it comprises only 5 classes in all). See: http://anonsvn.jboss.org/repos/labs/labs/jbosstm/trunk/XTS/WS-C/dev/src11-cxf/org/jboss/jbossts/xts/wsaddr/map/ and http://anonsvn.jboss.org/repos/labs/labs/jbosstm/trunk/XTS/WS-C/dev/src11-native/org/jboss/jbossts/xts/wsaddr/map/ A consequence of adopting this abstraction layer is that the XTS code is now dual source/binary. This complicates the build, deploy and test processes. It would be easier if this compatibility library was deployed by JBossWS. This would also provide a starting point for other JBossWS clients who wish to have a portable API for WSA. Perhaps you could look at the code and i) see if it could be extended with any missing capabilities or revised and cleaned up ii) advise on the feasibility of transferring it from TS to WS. I have not yet commented the API of MAP and its related classes so the way it works may not be entirely clear. Amongst the trickier details the following are especially noteworthy. I dealt with the RelatesTo problem by limiting the API to only allow one relationship to be managed. There is some trickery required when attempting to install the outbound AddressingProperties for a port derived from a W3CEndpointReference. In Native the AddressingProperties are pre-installed by the getPort call with To defined and with the ReferenceProperty elements stashed as a property of the AddressingContext. In CXF you get a port which has no AddressingProperties object on the request context. The To/ReferenceProperty values seem to have been stashed away somewhere else so they still get used. So, when you add other WSA properties the routine which fetches the outbound map creates an AddressingProperties object on the fly at first access.	1	1	https://issues.jboss.org/browse/JBWS-2106
13281	JBWS-2106	Hi Alessio, Thanks for this. It looks ok to me -- did you change any of my original API? (apart from adding getReferenceParameters, that is). I didn't spot any diffference.	1	1	https://issues.jboss.org/browse/JBWS-2106
13282	JBWS-2102	Well, that's maybe what the blog entry says but is it actually what the JaxWS committee have decided? Have you asked Thomas Diesler (since he is JBoss' representative) whether this is the official position of the committee? Personally, I think this decision makes provision of the WS Addressing capability in JaxWS utterly pointless. JaxWS might just as well not bother attaching WS Addressing headers to messages if there is no simple API for clients/SEIs to set/detect the header values. If they have to resort to installing SOAP handlers to do this then they might as well either i) implement WS Addressing themselves or ii) invent their own (simpler) custom header protocol. The latter is not an option for me as the WSTX specs I am working to require WS-Addressing headers with quite specific values. So, if this is a final decision then it means I will have to provide my own implementation of the JBossWS AddressingProperties class and associated handlers and ensure it runs on all 3 stacks -- a pointless fork of the code.	1	1	https://issues.jboss.org/browse/JBWS-2102
13283	JBWS-2090	Which mailing lists do you require? Usually for a project we have the following: <projectName>-announce <projectName>-commits <projectName>-dev <projectName>-issues <projectName>-users Mark	1	1	https://issues.jboss.org/browse/JBWS-2090
13284	JBWS-2090	Helpdesk ticket #665595 has been created with the following text: Please could the following mailing lists be created at http://lists.jboss.org metro-users cxf-users	1	1	https://issues.jboss.org/browse/JBWS-2090
13285	JBWS-2090	Mark, could you please make the forum messages go to the above mentioned mailing lists? IOW we need what we currently have with jbossws-users collecting messages coming from http://www.jboss.org/index.html?module=bb&op=viewforum&f=200 Thanks	1	1	https://issues.jboss.org/browse/JBWS-2090
13286	JBWS-2090	Mark, it's a long time we're waiting for this issue to be solved. Can you please give us any update on this or reassign in case you're not the one supposed to work on this? thank you	1	1	https://issues.jboss.org/browse/JBWS-2090
13287	JBWS-2031	I added a test for the addressing client. I am handing over to you because its currently failing. Could you please have a look? Similar tests should be added for the other clients.	1	1	https://issues.jboss.org/browse/JBWS-2031
13288	JBWS-2030	I added a test for the addressing client. I am handing over to you because its currently failing. Could you please have a look? Similar tests should be added for the other clients. tests-run-internal: [junit] Running org.jboss.test.ws.jaxws.samples.dar.AddressingClientTestCase [junit] Sent message ID: urn:uuid:4b9ffcb8-372-4e1e-a01e-8342ce858d55 [junit] Wed Mar 12 14:42:16 CET 2008 Sending request... [junit] Sent message ID: urn:uuid:4113f634-ade-4577-8fb3-b52a172f01f6 [junit] Wed Mar 12 14:42:23 CET 2008 Sending request... [junit] Doing something interesting in the mean time... ;-) [junit] Sent message ID: urn:uuid:7559bab7-b11a-4c2a-9fb8-57517af50c0e [junit] Wed Mar 12 14:42:28 CET 2008 Sending request... [junit] Wed Mar 12 14:42:29 CET 2008 Done. [junit] Tests run: 3, Failures: 0, Errors: 2, Time elapsed: 19.331 sec [junit] Test org.jboss.test.ws.jaxws.samples.dar.AddressingClientTestCase FAILED	1	1	https://issues.jboss.org/browse/JBWS-2030
13289	JBWS-1994	There's the package: src/main/java/org/jboss/wsf/stack/sunri/* Should it be renamed to metro as well?	1	1	https://issues.jboss.org/browse/JBWS-1994
13290	JBWS-1994	"should go" means "should be removed"	0	0	https://issues.jboss.org/browse/JBWS-1994
13291	JBWS-1982	This need to be looked at again. * The knowledge of how to (un)deploy stack X should not leak into framework * If AS has stack X.0.0 deployed and you want to deploy stack Y.0.1, how can framework that comes with Y.0.1 know how to undeploy X.0.0? The changes need to be rolled back until we know how to do this properly.	1	1	https://issues.jboss.org/browse/JBWS-1982
13292	JBWS-1967	Richard, as this is marked as depending on JBAS-5257 which you closed, could you please see what still needs to be done here?	0	0	https://issues.jboss.org/browse/JBWS-1967
13293	JBWS-1928	Apart from * Integrate Metro (wsit+jaxws) into jbossws build * Provide a metro based community release is the required functional;ity all there? i.e. Is the exclude file clean? Are there no unwanted FIXMEs?	0	0	https://issues.jboss.org/browse/JBWS-1928
13294	JBWS-1869	The test is failing for AS50 http://jbws.dyndns.org:8180/hudson/view/Metro%20Tests/job/Metro-Core-AS-5.0.0/35/testReport/org.jboss.test.ws.jaxws.jbws1869/JBWS1869TestCase/testStreamBasedInvocation/ unit.framework.ComparisonFailure: expected:<..."1.0" ...> but was:<...'1.0' encoding='UTF-8'...> at org.jboss.test.ws.jaxws.jbws1869.JBWS1869TestCase.testStreamBasedInvocation(JBWS1869TestCase.java:98) Could you find out why the encoding is added in AS50 and not in AS42?	0	0	https://issues.jboss.org/browse/JBWS-1869

#	Requirement	Questions	Answered	Successfully Answered	URL
13295	JBWS-1842	What exactly needs to be done here? I can find the above mentioned class in https://svn.jboss.org/repos/jbossas/projects/metadata/trunk but it is not clear to me who the producer and consumer of instances for the metadata object is. Is it correct to assume that somehow an instance of ServiceReferenceMetaData will be available through the DeploymentUnit and that [code] public interface ServiceRefHandler { final String BEAN_NAME = "WSServiceRefHandler"; enum Type {JAXRPC, JAXWS}; ServiceRefMetaData newServiceRefMetaData(); Object newChild(ServiceRefElement ref, UnmarshallingContext navigator, String namespaceURI, String localName, Attributes attrs); void setValue(ServiceRefElement ref, UnmarshallingContext navigator, String namespaceURI, String localName, String value); void bindServiceRef(Context encCtx, String encName, UnifiedVirtualFile vfsRoot, ClassLoader loader, ServiceRefMetaData sref) throws NamingException; } [/code] needs to be extended to consume instances of ServiceReferenceMetaData?	1	1	https://issues.jboss.org/browse/JBWS-1842
13296	JBWS-1814	Hi Alessio, This still has a bug wherein sometimes it throws an error saying: Cannot get the certificate for message encryption! Verify the keystore contents, considering the certificate is obtained through the alias specified in the encrypt configuration element or (server side only) through a single key used to sign the incoming This error is generated randomly i.e. sometimes it works but sometime it doesn't. Any help with this would be highly appreciated! Thanks in advance! Regards, Pratik Pai	1	0	https://issues.jboss.org/browse/JBWS-1814
13297	JBWS-1814	Hi Alessio, This problem still persist wherein it throws an error saying: Cannot get the certificate for message encryption! Verify the keystore contents, considering the certificate is obtained through the alias specified in the encrypt configuration element or (server side only) through a single key used to sign the incoming This error is thrown randomly i.e. it goes through sometimes but sometimes it doesnt Any help with this would be highly appreciated! Thanks in advance! Regards, Pratik Pai	1	0	https://issues.jboss.org/browse/JBWS-1814
13298	JBWS-1777	Alessio, I remember you did some progress here. What's the state of this issue?	1	0	https://issues.jboss.org/browse/JBWS-1777
13299	JBWS-1707	Do you guys mind if I tweak the status of some of those resolved bugs, to see what's going on? I'll put everything back how I found it before I'm done.	1	1	https://issues.jboss.org/browse/JBWS-1707
13300	JBWS-1670	I've just done some tests, this could be easily fixed resurrecting the InvocationExceptionHandler stuff I committed last month (just need a little modification to well fit into the latest deployement architecture). However, I suspect there might be something wrong with the InvocationExceptionHandler idea, since it has been deleted... am I missing something? Would you like me to take care of this?	1	1	https://issues.jboss.org/browse/JBWS-1670
13301	JBWS-1655	Hi Thomas, could you please provide more details about this issue, how to reproduce etc, or is there any existing test case? Thanks.	0	0	https://issues.jboss.org/browse/JBWS-1655
13302	JBWS-1612	Tom Fennelly sais: I'm interested in it from a JBossESB perspective i.e. being able to expose JBossWS endpoints through the ESB => making the endpoints consumable over any channel supported by the ESB (HTTP, JMS, FTP etc). I had some success. However, it involved using some JBossWS classes that I'm not sure are part of its public API => meaning they could change in the future, breaking integration with the ESB. The ESB class of most interest is (specifically, the "process" and "getServiceEndpoint" methods): http://anonsvn.labs.jboss.com/labs/jbossesb/trunk/product/core/listeners/src/org/jboss/soa/esb/actions/soap/JBossWSAdapter.java Was just wondering if you could throw an eye over this and let us know what you think. Mark (Little) said that one of your customers has done something like this in the past and ran into issues with it re API changes. If so, what are the possibilities of making the relevant code part of a public API?	1	1	https://issues.jboss.org/browse/JBWS-1612
13303	JBWS-1552	OK, so "rpc/encoded is supported ... as is", but if you use it, it throws a WSException... I'm confused??? Can it be used or not? And if so, how can we use it without getting an exception thrown? Thanks.	0	0	https://issues.jboss.org/browse/JBWS-1552
13304	JBWS-1527	Thomas, Did you verify the ejb21 ws tests as part of your TCK testing?	0	0	https://issues.jboss.org/browse/JBWS-1527
13305	JBWS-1526	How come the interop tests pass with JDK6? All interop tests should run against the public endpoints. http://jbws.dyndns.org:8180/hudson/view/Native%20Tests/	0	0	https://issues.jboss.org/browse/JBWS-1526
13306	JBWS-1525	Jason, could you please have a look if this is critical for 1.2.0? If not please reschedule for 1.2.1	0	0	https://issues.jboss.org/browse/JBWS-1525
13307	JBWS-1525	I fixed JBWS-1529, however the test still fails because the service and client are referring to differing WSDL definitions. The service is using two namespaces, and the service name ends with "Service". The client is using one namespace without the "Service". What should the test be testing. Are you attempting to test the same WSDL that we published with the previous JSR-181 implementation?	0	0	https://issues.jboss.org/browse/JBWS-1525
13308	JBWS-1508	Could you you please incude those pictures that you generated from the schema? This issues is relevant for http://tddell.dyndns.org/mediawiki/index.php/JAX-RPC_User_Guide as well	0	0	https://issues.jboss.org/browse/JBWS-1508
13309	JBWS-1506	Are you sure you want to maintain the BPEL docu in two locations? http://tddell.dyndns.org/mediawiki/index.php/JAX-RPC_User_Guide#WS-BPEL http://tddell.dyndns.org/mediawiki/index.php/JAX-WS_User_Guide#WS-BPEL There should be no duplication of content between th JAX-RPC, JAX-WS user guide	0	0	https://issues.jboss.org/browse/JBWS-1506
13310	JBWS-1496	I ran the test three times in a row, restarted app server, then ran test three more times. I could not reproduce the test failure. prepare: [echo] ---------- [echo] jboss.home = C:/dev/jbossas/branches/Branch_4_2/build/output/jboss-4.2.0.CR1-ejb3 [echo] tomcat.home = ${tomcat.home} [echo] java.home = C:\dev\jdk1.5.0_10\jre [echo] endorsed = C:/dev/jbossas/branches/Branch_4_2/build/output/jboss-4.2.0.CR1-ejb3/lib/endorsed [echo] excludesfile = tests-jboss42-excludes.txt [echo] ---------- one-test: [junit] Running org.jboss.test.ws.jaxrpc.samples.wsbpel.hello.BpelHelloTestCase [junit] Tests run: 2, Failures: 0, Errors: 0, Time elapsed: 5.641 sec I do see the JobExecutorThread error. The problem is the thread starts before the data source inside jbpm-bpel.sar is deployed. Despite the stack trace, the error is harmless as the thread tries again later. This issue will be addressed as part of BPEL-229. What are the steps to reproduce the problem?	0	0	https://issues.jboss.org/browse/JBWS-1496
13311	JBWS-1474	AFAIC, 1.2.0 is production ready for all included pieces. Isn't it?	0	0	https://issues.jboss.org/browse/JBWS-1474
13312	JBWS-1469	Will JBossWS-1.2.0 be made available with JBossAS-4.0.5.SP1 or 4.0.6?	0	0	https://issues.jboss.org/browse/JBWS-1469
13313	JBWS-1449	This feature request is not clear to me. What does the user want? I can understand it two ways: * Either user wants WSDL to be generated in the form WSDL + external file schema imports (schema is not inlined in WSDL but imported instead) * Or user wants Java2Schema command line tool	1	0	https://issues.jboss.org/browse/JBWS-1449
13314	JBWS-1443	Carlo sais: I've made an implementation which passes the test: http://jira.jboss.com/jira/browse/EJBTHREE-757?page=com.cenqua.fisheye.jira:fisheye-tabpanel I turned the tables by pulling in the message context via MessageContextAssociation. Do you see any problems with this? The one thing I don't like is the circular dependencies between ejb3 and ws.	0	0	https://issues.jboss.org/browse/JBWS-1443

#	Requirement	Questions	Answered	Successfully Answered	URL
13315	JBWS-1443	This approach is probably not the right way to go. How about: WebServiceContextCallback callback = new WebServiceContextCallback() { public MessageContext getMessageContext() { ... } public WebServiceContext getWebServiceContext() { ... } }; sessionContext.registerWebServiceContextCallback(callback); The WebServiceContextCallback interface would live in the ejb3 code base. The dependency would be unidirectional, as it is now already.	0	0	https://issues.jboss.org/browse/JBWS-1443
13316	JBWS-1443	Ok, works on jboss50. Are you planning to backport this to jboss40, jboss42?	0	0	https://issues.jboss.org/browse/JBWS-1443
13317	JBWS-1443	what is your target platform? Have a look at http://jbws.dyndns.org:8180/hudson/ for tests that have passed QA	1	1	https://issues.jboss.org/browse/JBWS-1443
13318	JBWS-1442	trunk Branch_4_0 jdk1.4 cannot compile. Does java need additional libs from these required by 1.0 branch?	1	1	https://issues.jboss.org/browse/JBWS-1442
13319	JBWS-1442	Can you provide a summary of problem?	1	1	https://issues.jboss.org/browse/JBWS-1442
13320	JBWS-1442	I followed the instructions in the wiki. When I try compile the testsuite it fails just the same way as the cc build (see build artifacts for details ;) Prabhat, is testsuite compiling by you?	0	0	https://issues.jboss.org/browse/JBWS-1442
13321	JBWS-1442	Thanks for telling me about JBWS-1458 Do you see any problem with jdk15 builds? They seem to me fine.	0	0	https://issues.jboss.org/browse/JBWS-1442
13322	JBWS-1442	Should I remove cruisecontrol builds then? btw I don't see the hudson instance	0	0	https://issues.jboss.org/browse/JBWS-1442
13323	JBWS-1433	Darran, I believe you implemented this. Is this code in trunk?	1	0	https://issues.jboss.org/browse/JBWS-1433
13324	JBWS-1402	What is the relation of this bug to bug http://jira.jboss.org/jira/browse/JBWS-1911? Isn't it fixed already?	1	0	https://issues.jboss.org/browse/JBWS-1402
13325	JBWS-1397	Thomas, Does this affect all TestCases written for JBoss-5.0.x?	0	0	https://issues.jboss.org/browse/JBWS-1397
13326	JBWS-1381	Remy, could you please comment and reasssign?	0	0	https://issues.jboss.org/browse/JBWS-1381
13327	JBWS-1330	http://www.jboss.com/index.html?module=bb&op=viewtopic&p=4041299#4041299 My test clients have code like this Code: InitialContext iniCtx = getInitialContext(); OrganizationHome home = (OrganizationHome)iniCtx.lookup("ejb/RoleSecuredSLSB"); OrganizationRemote bean = null; try { bean = home.create(); fail("Security exception expected"); } catch (Exception e) { // all cool, now try again with valid credentials SecurityAssociation.setPrincipal(new SimplePrincipal(USERNAME)); SecurityAssociation.setCredential(PASSWORD); bean = home.create(); } This no longer works. How is the client now supposed to provide valid principal/credential	1	1	https://issues.jboss.org/browse/JBWS-1330
13328	JBWS-1309	There @WebContext annotation seems to not be fully-covered in the EJB3 deployment descriptor. Please see the discussion on [url=http://www.jboss.com/index.html?module=bb&op=viewtopic&t=140182]this thread[/url] of the JBossWS forum regarding the issue. Could you re-open this issue so that the security-overrides that are missing can be addressed? Evidently the stuff we require	0	0	https://issues.jboss.org/browse/JBWS-1309
13329	JBWS-1308	Can we access these slides externally? Thanks	1	0	https://issues.jboss.org/browse/JBWS-1308
13330	JBWS-1307	is there any way we can get access to this documentation? I can not seem to access jbws.dyndns.org. I can access all the documentation for other projects since they are hosted on labs.jboss or docs.jboss.. Thanks, Gary	1	0	https://issues.jboss.org/browse/JBWS-1307
13331	JBWS-1297	How to obtain the fault bean: 1.a If the exception class has a getFaultInfo method **and** WebFault annotation: (note: how to interpret the "and" above? an exception having a WebFault annotation but not a getFaultInfo method does not comply with the semantics of the annotation. should we bomb?) 1.a.1 invoke getFaultInfo on the exception object; the return instance goes to JAXB 2.b Otherwise 2.b.1 retrieve the bean class from the fault metadata 2.b.2 instantiate the bean class 2.c.2 copy the non-excluded properties from the exception instance to the bean instance	0	0	https://issues.jboss.org/browse/JBWS-1297
13332	JBWS-1297	Marshalling is done. The outgoing SOAP messages look good, altough the tests still fail because unmarshalling is not done yet. I modified JSR181MetaDataBuilder to populate the XML type for the fault bean. Currently, I copy the properties using low-level reflection. Can you point me to the accessor API you are referring to? I'll commit what I have and work on unmarshalling the rest of the evening.	0	0	https://issues.jboss.org/browse/JBWS-1297
13333	JBWS-1296	This thing seems to be really trivial to do. Also this is is apparently used frequently and quite critical as for full blown WS-Security stack. Please consider implementing it in 2.0.3 or something. All Metro, Axis2 and CXF does support this.	1	1	https://issues.jboss.org/browse/JBWS-1296
13334	JBWS-1194	This is the build.properties file I use to run the test script: jboss.jdk14.src.url=http://anonsvn.jboss.org/repos/jbossas/tags/JBoss_4_0_5_GA jdk14.distro.name=jboss-4.0.5.GA #only set this if you need to overwrite the one from jdk14.src.url #provide the absolute dir location #jboss.jdk14.home= jboss.jdk15.src.url=http://anonsvn.jboss.org/repos/jbossas/branches/Branch_4_2 jdk15.distro.name=jboss-4.2.0.CR1-ejb3 #only set this if you need to overwrite the one from jdk15.src.url #provide the absolute dir location #jboss.jdk15.home= tomcat.url=http://archive.apache.org/dist/tomcat/tomcat-5/v5.5.17/bin tomcat.zip.file=apache-tomcat-5.5.17.zip tomcat.home.name=apache-tomcat-5.5.17 #Tomcat home. This is required to run samples test eventhough target server is JBoss. #Will ask WS team to fix this. Only when needs to be overwritten with your tomcat #tomcat.home= #jboss webservices svn repository url jbossws.svn.url=https://svn.jboss.org/repos/jbossws/branches/jbossws-1.0/ #This is needed for logging purpose jbossws.version=1.0.4.GA #jboss-head location jboss-head= #JDK1.5 home java15=/opt/jdk1.5.0_05 #JDK 1.4. Home. This jdk should have BouncyCastle security provider jar. #See the WIKI http://wiki.jboss.org/wiki/Wiki.jsp?page=JBWSFAQBuildAndInstallJBoss40x java14=/home/pjha/j2sdk1.4.2_09 #Ant Home. WS needs some specific jars in Ant. The best way is to #point ant home to jboss-head/tools. You need to copy log4j.jar to #jboss-head/tools/lib though. No need to change this. ant-ws=./ws-tools jdk14.deploy=deploy-jbossws14 jdk15.deploy=deploy-jbossws	1	1	https://issues.jboss.org/browse/JBWS-1194
13335	JBWS-1178	Where did you find the source code for jbossws 1.0.3? It is not available at the download page (only binary, source for newer versions only) and I couldn't find it in SVN either. This is a break of JBoss own licence...	1	1	https://issues.jboss.org/browse/JBWS-1178
13336	JBWS-1172	I cannot find any test for the @SchemaValidation annotation on a webservice, is this already supported ?	1	0	https://issues.jboss.org/browse/JBWS-1172
13337	JBWS-1172	Also, when I add this to my webservice: @SchemaValidation(enabled=true, errorHandler=nl.kb.hrd.dd.karakterisering.impls.ValidationErrorHandler.class, schemaLocation="ergens") and send an invalid request, nothing happens. The request just gets passed, the handler is not called, and I see no errormessage about the invalid schemaLocation. A related question: in the jax-ws RI, it is not necessary to specify the schema location. Why is it needed here ?	1	0	https://issues.jboss.org/browse/JBWS-1172
13338	JBWS-1162	Can we concatenate their text summaries?	1	1	https://issues.jboss.org/browse/JBWS-1162
13339	JBWS-1162	The recent email conversations are an indication that this issue is not resolved properly. As I suggested before, we need a web application that any team member can access and initiate a testrun in a non-conflicting manner. Can this be done?	1	1	https://issues.jboss.org/browse/JBWS-1162
13340	JBWS-1158	Can we discuss this one as well? administrators: http://www.atlassian.com/software/jira	0	0	https://issues.jboss.org/browse/JBWS-1158

#	Requirement	Questions	Answered	Successfully Answered	URL
13341	JBWS-1136	Wil this JIRA allow POJO WAR JSEs to be both authenticated with the configured JBossSX login module and allow them to be authorized with the standard declarative WAR web.xml <security-constraint> tag structure? rjb	0	0	https://issues.jboss.org/browse/JBWS-1136
13342	JBWS-1136	Given I agree with you about what you say on POJO endpoints, we could nevertheless do what the issue title says i.e. allowing a username element in the require list of the wsse configuration the same way we have for the timestamp, for example. This way we could reject requests that do not have a Username Token. What do you think about?	0	0	https://issues.jboss.org/browse/JBWS-1136
13343	JBWS-1119	Hi Jason, was this actually merged to the trunk and eventually removed later? I'm asking this because I can't find this in the trunk and somebody is asking for this feature on the forum. I can work on this and port from tag 1.0.3 to the trunk of course, just want to know if it was not done or was reverted because of some problems discovered later and not mentioned here. Thank you.	0	0	https://issues.jboss.org/browse/JBWS-1119
13344	JBWS-1119	Will this issue be included in the JBoss.native-2.0.4 then? /Michael	1	1	https://issues.jboss.org/browse/JBWS-1119
13345	JBWS-1090	As of https://svn.jboss.org/repos/jbossws/trunk@938 this test is not yet available in trunk, or is it?	0	0	https://issues.jboss.org/browse/JBWS-1090
13346	JBWS-1086	Why should this being resolved "Out of Date"?	0	0	https://issues.jboss.org/browse/JBWS-1086
13347	JBWS-1077	Does this include updating the WS-Security documentation?	0	0	https://issues.jboss.org/browse/JBWS-1077
13348	JBWS-1065	Is there a way to let the build scripts check the Xerces version?	0	0	https://issues.jboss.org/browse/JBWS-1065
13349	JBWS-937	What is the purpose of supporting wscompile options? The old axis based stack jboss-ws4ee is deprecated and will no longer be worked on. All soapui integration efforts should focus on jbossws	1	0	https://issues.jboss.org/browse/JBWS-937
13350	JBWS-925	The referenced document (http://java.sun.com/developer/technicalArticles/xml/fastinfoset/) is over four years old. Since the time it was published several of your open source competitors have integrated Fast Infoset into their stacks. This includes Apache's Axis2 (http://wso2.org/library/2686) which is used by Geronimo (http://geronimo.apache.org/) and Jonas (http://jonas.objectweb.org/). It also includes Sun's reference implementation of JAX-WS (https://jax-ws.dev.java.net/) which is the basis of the Metro Web Service Stack (https://metro.dev.java.net/) used by Glassfish (https://glassfish.dev.java.net/). I don't know anything about commercial vendors, but I assume coverage is good there as well. The reason I bring this up is because I am endeavoring to write the highest performing web service possible with open source software. Right now, we have standardized on JBoss, but the benefits of Fast Infoset look large enough that they might justify changing application servers. Is there any possibility Fast Infoset could be integrated into the 2.0.x codebase rather than the 2.1.x codebase?	0	0	https://issues.jboss.org/browse/JBWS-925
13351	JBWS-925	I see that this ticket was changed to "Provide initial support for FastInfoset." What does "initial support" entail?	0	0	https://issues.jboss.org/browse/JBWS-925
13352	JBWS-922	is there a webservice to remotely list deployed webservices or should we start by parsing the html returned by /jbossws/services?	1	1	https://issues.jboss.org/browse/JBWS-922
13353	JBWS-914	> > How is the download count managed? This is probably not setup to work > with the jbossws download area. Yes, you need to enable counters for your project as defined in this wiki: http://labs.jboss.org/wiki/DownloadCounters	0	0	https://issues.jboss.org/browse/JBWS-914
13354	JBWS-871	According to the schema <complexType name='IntegerArray'> <sequence> <element name='value' nillable='true' maxOccurs='unbounded' minOccurs='0' type='int'/> </sequence> </complexType> A null array can only be expressed like this <Integer_1 xmlns:xsi='http://www.w3.org/2001/XMLSchema-instance'> <ns:value xmlns:ns='http://jbws871.ws.test.jboss.org/jaws' xsi:nil='1'/> </Integer_1> which is no different to an array with a sigle null value. Is it true that the endpoint must be able to deal with that ambiguity?	0	0	https://issues.jboss.org/browse/JBWS-871
13355	JBWS-871	"which is no different to an array with a sigle null value. Is it true that the endpoint must be able to deal with that ambiguity? " Yes, when an array is reprented as a single element definition, in other words no wrapping complex type, then a null value is semantically equivalent to inserting a single null value. Endpoints should use empty arrays instead of null to prevent the obvious problems that come from this. -Jason	1	1	https://issues.jboss.org/browse/JBWS-871
13356	JBWS-868	I got some questions: - where are the user guide docbook sources? - what do I have to do for the existing bpel tests to get packaged in the sample archive?	1	1	https://issues.jboss.org/browse/JBWS-868
13357	JBWS-868	Can you elaborate?	1	1	https://issues.jboss.org/browse/JBWS-868
13358	JBWS-868	BPEL does not explicitly depend on AS, but it requires JMS. Does Messaging run on plain Tomcat? The question is, where do I take bpel.sar from? Do I include it in the resources of the sample (1.4MB!) or get it from the repository? Is there a sample with similar requirements?	1	1	https://issues.jboss.org/browse/JBWS-868
13359	JBWS-868	The testduite still fails Cannot obtain URL for: jaxrpc-samples-wsbpel-hello.war java.lang.IllegalArgumentException: Cannot obtain URL for: jaxrpc-samples-wsbpel-hello.war at org.jboss.test.ws.JBossWSTestHelper.getArchiveURL(JBossWSTestHelper.java:123)	0	0	https://issues.jboss.org/browse/JBWS-868
13360	JBWS-868	Aren't samples intended to be run after (a) the sample suite is built and (b) the sample jars are packaged? It looks like you are trying to run them as a normal unit test. jbpm-bpel.sar is fetched from the remote repository as part of the samples build process. Once you build the samples, please change to the jaxrpc/wsbpel directory and execute the "tests" target there: C:\dev\jbossws\jbossws-1.0\output\jbossws-samples-1.0.4.DEV\jaxrpc\wsbpel>ant tests Buildfile: build.xml init: [echo] jboss.home = C:/dev/jboss/Branch_4_0/build/output/jboss-4.0.5.GA-ejb3 [echo] tomcat.home = ${tomcat.home} [echo] java.home = C:\dev\jdk1.5.0_07\jre [echo] node0 = AGUIZAR compile: [javac] Compiling 1 source file to C:\dev\jbossws\jbossws-1.0\output\jbossws-samples-1.0.4.DEV\output\classes copy-resources: zip-processes: check-jbpmbpel: install-jbpmbpel: [mkdir] Created dir: C:\dev\jboss\Branch_4_0\build\output\jboss-4.0.5.GA-ejb3\server\default\deploy\jbpm-bpel.sar [unjar] Expanding: C:\dev\jbossws\jbossws-1.0\output\jbossws-samples-1.0.4.DEV\lib\jbpm-bpel.sar into C:\dev\jboss\B ranch_4_0\build\output\jboss-4.0.5.GA-ejb3\server\default\deploy\jbpm-bpel.sar generate-sources: [move] Moving 2 files to C:\dev\jbossws\jbossws-1.0\output\jbossws-samples-1.0.4.DEV\output\wstools\java compile-generated: [javac] Compiling 2 source files to C:\dev\jbossws\jbossws-1.0\output\jbossws-samples-1.0.4.DEV\output\classes copy-generated: [copy] Copying 4 files to C:\dev\jbossws\jbossws-1.0\output\jbossws-samples-1.0.4.DEV\output\resources jars: [war] Building war: C:\dev\jbossws\jbossws-1.0\output\jbossws-samples-1.0.4.DEV\output\libs\jaxrpc-samples-wsbpel- hello.war [jar] Building jar: C:\dev\jbossws\jbossws-1.0\output\jbossws-samples-1.0.4.DEV\output\libs\jaxrpc-samples-wsbpel- hello-client.jar store-processes: [storeprocess] target url: http://localhost:8080/jbpm-bpel/deploy?processfile=file%3A%2FC%3A%2Fdev%2Fjbossws%2Fjbossws-1 .0%2Foutput%2Fjbossws-samples-1.0.4.DEV%2Foutput%2Flibs%2Fjaxrpc-samples-wsbpel-hello-process.zip [storeprocess] got response code: 200 tests: init: [echo] jboss.home = C:/dev/jboss/Branch_4_0/build/output/jboss-4.0.5.GA-ejb3 [echo] tomcat.home = ${tomcat.home} [echo] java.home = C:\dev\jdk1.5.0_07\jre [echo] node0 = AGUIZAR test: [junit] Running org.jboss.test.ws.jaxrpc.samples.wsbpel.hello.BpelHelloTestCase [junit] Tests run: 2, Failures: 0, Errors: 0, Time elapsed: 5.015 sec	1	1	https://issues.jboss.org/browse/JBWS-868

#	Requirement	Questions	Answered	Successfully Answered	URL
13361	JBWS-865	Since this fix is slated for JBossWS 1.0.1, will there be a workaround that I can do for the JBoss AS 4.0.4 GA release to get around this issue until JBossWS 1.0.1 is released?	0	0	https://issues.jboss.org/browse/JBWS-865
13362	JBWS-839	Could you please write a wiki about the bouncy castle issue? Can we automate this for our testruns?	1	1	https://issues.jboss.org/browse/JBWS-839
13363	JBWS-803	JBossEntityResolver How to apply the patch (jboss-entity-resolver-V3-patch.zip): 1) Apply the patch-common.txt in your common project (HEAD). (If you are using Eclipse, you can apply the patch: right mouse click to the commons project> team > apply patch) 2) Apply the patch-server.txt in your server project (HEAD). The current implementation of the JBossEntityResolver distinguishes between an in-server mode and a client-mode. In-server-mode: When used inside the application server the jax-ws-catalog.xml is loaded form the ${jboss.server.conf.dir}. This catalog references entities located in ${jboss.server.conf.dir}/catalog/[dtd|schema]. In-client-mode: When used from a remote client, all catalog files located in [META-INF/jax-ws-catalog.xml]and [client-resources/jax-ws-catalog.xml] are considered (class path resources). The jboss-all-client.jar contains a folder (client-resources\[dtd| schema]) where client-side resources are located and referenced form the client-resources/jax-ws-catalog.xml (also included in this jar file) Potential discussion topics: I) Currently the schema/dtd files are located: a) In the server project for the in-server mode -> the build process copy this files to the ${jboss.server.conf.dir} b) In the common project for the client access -> this build process copy this files to the jboss-common.jar and later to the jboss-all-client.jar Someone who knows the build process better tan me should unify this. II) To distinguish between the in-server/ client mode the System.getProperty("jboss.home.dir") is used. If this property is not null, then the resolver is running in the in-server mode. Maybe there is a better way to find out if the JBossEntityResolver is used inside the application server or not.	0	0	https://issues.jboss.org/browse/JBWS-803
13364	JBWS-784	what download link?	1	1	https://issues.jboss.org/browse/JBWS-784
13365	JBWS-733	The online chapter http://labs.jboss.com/portal/jbossws/user-guide/en/html/wstransaction.html Where is the JBossWS source repository? http://wiki.jboss.org/wiki/Wiki.jsp?page=JBWSFAQSourceRepository	1	1	https://issues.jboss.org/browse/JBWS-733
13366	JBWS-722	This webservices.xml does not qualify the wsdl-port <webservices version='1.1' xmlns='http://java.sun.com/xml/ns/j2ee' xmlns:impl='http://org.jboss.ws/samples/jsr109pojo' xmlns:xsi='http://www.w3.org/2001/XMLSchema-instance' xsi:schemaLocation='http://java.sun.com/xml/ns/j2ee http://www.ibm.com/webservices/xsd/j2ee_web_services_1_1.xsd'> <webservice-description> <webservice-description-name>TestService</webservice-description-name> <wsdl-file>WEB-INF/wsdl/TestService.wsdl</wsdl-file> <jaxrpc-mapping-file>WEB-INF/jaxrpc-mapping.xml</jaxrpc-mapping-file> <port-component> <port-component-name>JaxRpcTestServicePort</port-component-name> <wsdl-port>JaxRpcTestServicePort</wsdl-port> <service-endpoint-interface>org.jboss.test.ws.samples.jsr109pojo.JaxRpcTestService</service-endpoint-interface> <service-impl-bean> <servlet-link>TestService</servlet-link> </service-impl-bean> </port-component> </webservice-description> </webservices> org.jboss.ws.WSException: Cannot find port in wsdl: JaxRpcTestServicePort at org.jboss.ws.metadata.JSR109ServerMetaDataBuilder.buildMetaData(JSR109ServerMetaDataBuilder.java:130) at org.jboss.ws.server.WebServiceDeployerJSE.createWebServicesMetaData(WebServiceDeployerJSE.java:187) at org.jboss.ws.server.WebServiceDeployer.create(WebServiceDeployer.java:101) at org.jboss.ws.server.WebServiceDeployerJSE.create(WebServiceDeployerJSE.java:65) at org.jboss.deployment.SubDeployerInterceptorSupport$XMBeanInterceptor.create(SubDeployerInterceptorSupport.java:180) at org.jboss.deployment.SubDeployerInterceptor.invoke(SubDeployerInterceptor.java:91) at org.jboss.mx.server.Invocation.invoke(Invocation.java:90) at org.jboss.mx.server.AbstractMBeanInvoker.invoke(AbstractMBeanInvoker.java:262) at org.jboss.mx.server.MBeanServerImpl.invoke(MBeanServerImpl.java:668) at org.jboss.mx.util.MBeanProxyExt.invoke(MBeanProxyExt.java:190) at $Proxy34.create(Unknown Source) at org.jboss.deployment.MainDeployer.create(MainDeployer.java:953) at org.jboss.deployment.MainDeployer.deploy(MainDeployer.java:807)	1	1	https://issues.jboss.org/browse/JBWS-722
13367	JBWS-688	What's the ETA for the release?	0	0	https://issues.jboss.org/browse/JBWS-688
13368	JBWS-684	Could you pls check that these show up on cruisecontrol?	1	1	https://issues.jboss.org/browse/JBWS-684
13369	JBWS-678	How to apply the patch: The solution is mainly realized in the common project. A new global JBoss oasis catalog file (jax-ws-catalog.xml) is located in the server project. This file is copied during the installation to the conf directory. It contains the references to the entities (located in /docs/dtd, and /docs/schema/). Before, the references were hard coded in the JBossEntityResolver class. The new JUnit tests are located in the testsuite project. 1) Apply the patch-common.txt in your common project (HEAD). (If you are using Eclipse, you can apply the patch: right mouse click to the commons project? team ? apply patch) 2) Apply the patch-server.txt in your server project (HEAD). 3) Apply the patch-testsuite.txt in your testsuite project (HEAD).	1	1	https://issues.jboss.org/browse/JBWS-678
13370	JBWS-678	Reopen The JBossEntityResolver is not only used by server components, but also in remote clients. (i.e. many of the web service tests use application clients to connect to a remote server) From what I can see web service tests fail because entities cannot be loaded on the client side. How is this supposed to work? Also, I suggest that entities are configured on a per server basis and not loaded from a global location like /docs/[dtd|schema]	1	1	https://issues.jboss.org/browse/JBWS-678

#	Requirement	Questions	Answered	Successfully Answered	URL
13371	JBWS-678	!! New patch available> attached to the task: http://jira.jboss.org/jira/browse/JBWS-803 !! How to apply the patch (jboss-entity-resolver-V3-patch.zip): 1) Apply the patch-common.txt in your common project (HEAD). (If you are using Eclipse, you can apply the patch: right mouse click to the commons project> team > apply patch) 2) Apply the patch-server.txt in your server project (HEAD). The current implementation of the JBossEntityResolver distinguishes between an in-server mode and a client-mode. In-server-mode: When used inside the application server the jax-ws-catalog.xml is loaded form the ${jboss.server.conf.dir}. This catalog references entities located in ${jboss.server.conf.dir}/catalog/[dtd|schema]. In-client-mode: When used from a remote client, all catalog files located in [META-INF/jax-ws-catalog.xml] and [client-resources/jax-ws-catalog.xml] are considered (class path resources). The jboss-all-client.jar contains a folder (client-resources\[dtd| schema]) where client-side resources are located and referenced form the client-resources/jax-ws-catalog.xml (also included in this jar file) Potential discussion topics: I) Currently the schema/dtd files are located: a) In the server project for the in-server mode -> the build process copy this files to the ${jboss.server.conf.dir} b) In the common project for the client access -> this build process copy this files to the jboss-common.jar and later to the jboss-all-client.jar Someone who knows the build process better than me, should unify this. II) To distinguish between the in-server/ client mode the System.getProperty("jboss.home.dir") is used. If this property is not null, then the resolver is running in the in-server mode. Maybe there is a better way to find out if the JBossEntityResolver is used inside the application server or not.	0	0	https://issues.jboss.org/browse/JBWS-678
13372	JBWS-669	While you look at this, could you pls explore the possibilities of exposing MBeans as WS endpoints?	0	0	https://issues.jboss.org/browse/JBWS-669
13373	JBWS-669	Have you tried the depends clause in jboss-web?	1	1	https://issues.jboss.org/browse/JBWS-669
13374	JBWS-652	What about JSR-196 (JASPI) which does talk about securing/validating SOAP messages?	0	0	https://issues.jboss.org/browse/JBWS-652
13375	JBWS-546	INFO [HelloJavaBean] hello:Kermit 14:09:39,020 INFO [HelloJavaBean] hello:UserExceptionOne 14:09:39,020 ERROR [ServiceEndpoint] SOAP request exception org.jboss.test.webservice.jbws251.UserExceptionOne: just a test at org.jboss.test.webservice.jbws251.HelloJavaBean.hello(HelloJavaBean.java:35) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.jboss.ws.server.ServiceEndpointInvokerJSE.invokeServiceEndpoint(ServiceEndpointInvokerJSE.java: 87) at org.jboss.ws.server.ServiceEndpointInvoker.invoke(ServiceEndpointInvoker.java:102) at org.jboss.ws.server.ServiceEndpoint.handleRequest(ServiceEndpoint.java:226) at org.jboss.ws.server.ServiceEndpointServlet.doPost(ServiceEndpointServlet.java:110) at javax.servlet.http.HttpServlet.service(HttpServlet.java:717) at javax.servlet.http.HttpServlet.service(HttpServlet.java:810) at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:252) at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:173) at org.jboss.web.tomcat.filters.ReplyHeaderFilter.doFilter(ReplyHeaderFilter.java:96) at org.apache.catalina.core.ApplicationFilterChain.internalDoFilter(ApplicationFilterChain.java:202) at org.apache.catalina.core.ApplicationFilterChain.doFilter(ApplicationFilterChain.java:173) at org.apache.catalina.core.StandardWrapperValve.invoke(StandardWrapperValve.java:213) at org.apache.catalina.core.StandardContextValve.invoke(StandardContextValve.java:178) at org.jboss.web.tomcat.security.CustomPrincipalValve.invoke(CustomPrincipalValve.java:54) at org.jboss.web.tomcat.security.SecurityAssociationValve.invoke(SecurityAssociationValve.java:173) at org.jboss.web.tomcat.security.JaccContextValve.invoke(JaccContextValve.java:74) at org.apache.catalina.core.StandardHostValve.invoke(StandardHostValve.java:126) at org.apache.catalina.valves.ErrorReportValve.invoke(ErrorReportValve.java:105) at org.apache.catalina.core.StandardEngineValve.invoke(StandardEngineValve.java:107) at org.apache.catalina.connector.CoyoteAdapter.service(CoyoteAdapter.java:148) at org.apache.coyote.http11.Http11Processor.process(Http11Processor.java:868) at org.apache.coyote.http11.Http11BaseProtocol$Http11ConnectionHandler.processConnection(Http11BasePro tocol.java:663) at org.apache.tomcat.util.net.PoolTcpEndpoint.processSocket(PoolTcpEndpoint.java:527) at org.apache.tomcat.util.net.MasterSlaveWorkerThread.run(MasterSlaveWorkerThread.java:112) at java.lang.Thread.run(Thread.java:595) and this one in the report Testcase: testUserExceptionOne took 0.31 sec Caused an ERROR Call invocation failed with code [Client] because of: org.jboss.test.webservice.jbws251.UserExceptionOne: just a test; nested exception is: javax.xml.rpc.soap.SOAPFaultException: org.jboss.test.webservice.jbws251.UserExceptionOne: just a test java.rmi.RemoteException: Call invocation failed with code [Client] because of:	0	0	https://issues.jboss.org/browse/JBWS-546
13376	JBWS-522	Tom, is it possible to get at the HTTP return code? The 404 is not availbale from the HTTP headers. 2005-12-13 20:56:19,071 DEBUG [org.jboss.remoting.transport.http.HTTPClientInvoker] connect called for: org.jboss.remoting.transport.http.HTTPClientInvoker@191f801 2005-12-13 20:56:19,071 DEBUG [org.jboss.remoting.transport.http.HTTPClientInvoker] Setting request header with SOAPAction : " 2005-12-13 20:56:19,071 DEBUG [org.jboss.remoting.transport.http.HTTPClientInvoker] Setting request header with Content-Type : text/xml 2005-12-13 20:56:19,081 DEBUG [org.jboss.ws.binding.soap.SOAPMessageUnMarshaller] Got HTTP Header name=[Content-Length] value=[960] 2005-12-13 20:56:19,081 DEBUG [org.jboss.ws.binding.soap.SOAPMessageUnMarshaller] Got HTTP Header name=[Date] value=[Tue, 13 Dec 2005 19:56:19 GMT] 2005-12-13 20:56:19,081 DEBUG [org.jboss.ws.binding.soap.SOAPMessageUnMarshaller] Got HTTP Header name=[Server] value=[Apache-Coyote/1.1] 2005-12-13 20:56:19,081 DEBUG [org.jboss.ws.binding.soap.SOAPMessageUnMarshaller] Got HTTP Header name=[Content-Type] value=[text/html;charset=utf-8] 2005-12-13 20:56:19,081 DEBUG [org.jboss.ws.binding.soap.SOAPMessageUnMarshaller] Got HTTP Header name=[WWW-Authenticate] value=[Basic realm="EJBServiceEndpointServlet Realm"] 2005-12-13 20:56:19,141 DEBUG [org.jboss.ws.soap.MessageFactoryImpl] createMessage: [contentType=text/html; charset=utf-8] 2005-12-13 20:56:19,141 ERROR [org.jboss.ws.binding.soap.SOAPMessageUnMarshaller] Cannot unmarshall SOAPMessage javax.xml.soap.SOAPException: Unsupported content type: text/html; charset=utf-8	0	0	https://issues.jboss.org/browse/JBWS-522
13377	JBWS-522	I see this key=value pair in metadata null=[HTTP/1.1 200 OK] Is the null key value intended? Maybe it should be something more descriptive, like org.jboss.remoting.http-header-line	0	0	https://issues.jboss.org/browse/JBWS-522
13378	JBWS-515	may I ask what version of the spec this is?	0	0	https://issues.jboss.org/browse/JBWS-515
13379	JBWS-454	Jason, can we (a support customer) count on this issue being fixed in 4.0.4?	1	1	https://issues.jboss.org/browse/JBWS-454

#	Requirement	Questions	Answered	Successfully Answered	URL
13380	JBWS-434	Hi Alexey, can you give me an estimate on this issue? I did take a look myself, but it seems to be a fairly complex problem that requires changes at the marshalling and unmarshalling components. The eventing implementation depends on it, and i think about an intermediate solution to keep my subproject moving. Regards, Heiko	1	1	https://issues.jboss.org/browse/JBWS-434
13381	JBWS-434	testNullArrayOfAny() throws Exception { ArrayOfAny inObj = new ArrayOfAny(); assertNull(inObj._any); ArrayOfAny retObj = port.echo(inObj); assertNotNull(retObj); assertNull(retObj._any); } } This test method will fail: Testsuite: org.jboss.test.ws.jbws434.JBWS434TestCase Tests run: 2, Failures: 0, Errors: 1, Time elapsed: 5,098 sec Testcase: testNullArrayOfAny took 3,647 sec Caused an ERROR Call invocation failed: Could not transmit message; nested exception is: javax.xml.soap.SOAPException: Could not transmit message java.rmi.RemoteException: Call invocation failed: Could not transmit message; nested exception is: javax.xml.soap.SOAPException: Could not transmit message at org.jboss.ws.jaxrpc.CallImpl.invokeInternal(CallImpl.java:703) at org.jboss.ws.jaxrpc.CallImpl.invoke(CallImpl.java:408) at org.jboss.ws.jaxrpc.CallProxy.invoke(CallProxy.java:116) at $Proxy3.echo(Unknown Source) at org.jboss.test.ws.jbws434.JBWS434TestCase.testNullArrayOfAny(JBWS434TestCase.java:98) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at junit.extensions.TestDecorator.basicRun(TestDecorator.java:22) at junit.extensions.TestSetup$1.protect(TestSetup.java:19) at junit.extensions.TestSetup.run(TestSetup.java:23) Caused by: javax.xml.soap.SOAPException: Could not transmit message at org.jboss.ws.soap.SOAPConnectionImpl.call(SOAPConnectionImpl.java:162) at org.jboss.ws.jaxrpc.CallImpl.invokeInternal(CallImpl.java:674) ... 23 more Caused by: java.lang.NullPointerException at org.jboss.util.xml.Dom2Sax.process(Dom2Sax.java:55) at org.jboss.util.xml.Dom2Sax.dom2sax(Dom2Sax.java:48) at org.jboss.ws.jaxb.JAXBMarshallerImpl$1.marshal(JAXBMarshallerImpl.java:168) at org.jboss.xb.binding.XercesXsMarshaller.marshalWildcardOccurence(XercesXsMarshaller.java:907) at org.jboss.xb.binding.XercesXsMarshaller.marshalParticle(XercesXsMarshaller.java:789) at org.jboss.xb.binding.XercesXsMarshaller.marshalModelGroupSequence(XercesXsMarshaller.java:1102) at org.jboss.xb.binding.XercesXsMarshaller.marshalModelGroup(XercesXsMarshaller.java:1057) at org.jboss.xb.binding.XercesXsMarshaller.marshalParticle(XercesXsMarshaller.java:729) at org.jboss.xb.binding.XercesXsMarshaller.marshalComplexType(XercesXsMarshaller.java:685) at org.jboss.xb.binding.XercesXsMarshaller.marshalElementType(XercesXsMarshaller.java:414) at org.jboss.xb.binding.XercesXsMarshaller.marshalElement(XercesXsMarshaller.java:380) at org.jboss.xb.binding.XercesXsMarshaller.marshalElementOccurence(XercesXsMarshaller.java:886) at org.jboss.xb.binding.XercesXsMarshaller.marshallInternal(XercesXsMarshaller.java:275) at org.jboss.xb.binding.XercesXsMarshaller.marshal(XercesXsMarshaller.java:222) at org.jboss.ws.jaxb.JAXBMarshallerImpl.marshal(JAXBMarshallerImpl.java:210) at org.jboss.ws.jaxrpc.encoding.JAXBSerializer.serialize(JAXBSerializer.java:100) at org.jboss.ws.soap.SOAPContentElement.getXMLFragment(SOAPContentElement.java:172) at org.jboss.ws.soap.SOAPContentElement.expandToDOM(SOAPContentElement.java:844) at	0	0	https://issues.jboss.org/browse/JBWS-434
13382	JBWS-374	How are we going to address org.jboss.xb.binding.JBossXBRuntimeException: Global type definition is not found: {http://org.jboss.ws/anonymous/types}>root My inital tought was, that the JBossXSModel should support these types of lookup and return the XSModel object that represents the anonymous type associated with the global element {http://org.jboss.ws/anonymous/types}root. If that is what we agree upon, JBWS-374 is not (yet) your issue. Anil, looking at the debug log, I see that the anonymous types are not serialized correctly (they are missing). If this is a serialization problem this issue is minor. If they are indeed missing from the model, it would be serious. Jason, if Alex delegates this back to us and demands JBossXSModel lookup for anonymous type names, could you take care of that? All cases defined by the JSR-109 spec should be supported. See org.jboss.test.ws.anonymous.AnonymousTypesTestCase	0	0	https://issues.jboss.org/browse/JBWS-374
13383	JBWS-374	Alexey, Since the jaxrpc-mapping data is passed to jbossxb (via jbossxb specific mappings), and the JBossXSModel is just a schema model, is it possible for you to add support for the annoymous type expressions that are defined in jsr109? Basically it amounts to using expression based qnames such as typeNS:>>root>inside to specify which anonymous type to map to which field. -Jason	1	1	https://issues.jboss.org/browse/JBWS-374
13384	JBWS-374	I have implemented this in JBossXSModel, and marshalling seems to work properly now. However, unmarshalling fails because JBossXB works off special object heirarchy that is copied from JBossXSModel. Since anonymous types are not true types, and thus not registered under their names the type binding skips them. We need the ability to add type defitions to an existing JBossXB schema binding, perhaps by modifying XsdBinder to expose some of its private methods. If there was a way to add XSTypeDefinitions and XSElementDeclarations, then I could modify JAXBUnmarshaller to scroll through the anonymous types and add them. Alexey, What are your thoughts on this? All anonymous types and elements can be retrieved from JBossXSModel as maps by calling getAnonymousElements() and getAnonymousTypes(). -Jason	1	1	https://issues.jboss.org/browse/JBWS-374
13385	JBWS-374	I've added a couple of methods to the XsdBinder. Is this enough? /** * @deprecated * * This method is added temporary to get anonymous type binding working in JBossWS. * It will be removed when anonymous type binding in JBossWS is implemented properly. * No one else should use this method. * * <p>This method binds a type definition and adds it as a global type to the passed in schema binding. * * @param schema schema binding the type should be added to * @param type type definition to be bound */ public static void bindType(SchemaBinding schema, XSTypeDefinition type) /** * @deprecated * * This method is added temporary to get anonymous type binding working in JBossWS. * It will be removed when anonymous type binding in JBossWS is implemented properly. * No one else should use this method. * * <p>This method binds an element declaration and adds it as a global element to the passed in schema binding. * * @param schema schema binding the type should be added to * @param element element declaration to be bound * @param minOccurs * @param maxOccurs * @param maxOccursUnbounded */ public static void bindElement(SchemaBinding schema, XSElementDeclaration element, int minOccurs, int maxOccurs, boolean maxOccursUnbounded)	0	0	https://issues.jboss.org/browse/JBWS-374
13386	JBWS-355	I'll ask this question again... Why are we writing to "private" Xerces API inside the webservices implementation? This needs to be abstracted away through JBossXB.	0	0	https://issues.jboss.org/browse/JBWS-355

#	Requirement	Questions	Answered	Successfully Answered	URL
13387	JBWS-355	Some of this could be abstracted and moved to JBossXB. I believe the reason for its current location is that currently schema generation is only needed by the webservice module. Since wsdl generation/processing is webservice specific, and since it is based heavily off of schema, the webserivce module needs to use a full schema api. There is no current standard API for this, but there is a W3C submission that xerces implements. http://www.w3.org/Submission/2004/SUBM-xmlschema-api-20040309/ This could be abstracted in JBossXB, but it would require moving schema generation into JBossXB since the schema api is a read-only API and webservices requires the ability to modify the schema (for example if someone manipulates a Call object). Migrating this codebase would be a fairly involved amount of work, and could affect the timelines of both the webservice and JBossXB projects. Alexey and Thomas may have additional reasons that I didn't think of, so I will defer to them. -Jason an the	1	0	https://issues.jboss.org/browse/JBWS-355
13388	JBWS-314	For the above solution, what happens when you have more than one mapping file needed (more than one service...)? Does the above account for that?	0	0	https://issues.jboss.org/browse/JBWS-314
13389	JBWS-299	I am looking at using this option now. We currently have JBoss.NET WS running in production (3.2.6), and some calls return arrays of 50K instances. Axis serialisation just dies for 10 - 15 seconds with these requests (the .NET client on a trader desktop is much faster parsing it). We are looking at migrating the app to JBoss 4 and J2EE 1.4 WS, and nodatabinding will allow us to cache the SOAP response (since the underlying data is static) which should make the call much faster. Are we saying this wscompile is not supported in JBoss 4?	1	0	https://issues.jboss.org/browse/JBWS-299
13390	JBWS-286	Jason says to my earlier comment: "it sort of is, ssl actualy does asym key exchange with negotiated sym keys for tansport" Right, because of the negotiated sym keys for transport, can a claim for non repudiation be made?	0	0	https://issues.jboss.org/browse/JBWS-286
13391	JBWS-271	Does that mean all methods in the SEI that have a void return will be one-way? JSR-181 handles this at the method level @WebService public class PingService { @WebMethod @Oneway public void ping() { } }; Resulting WSDL: <definitions> <message name="ping"/> <portType name="PingService"> <operation name="ping"> <input message="ping"/> </operation> </portType> </definitions>	0	0	https://issues.jboss.org/browse/JBWS-271
13392	JBWS-268	Thomas, this is just one of the vagaries of wscompile. Array handling by wscompile has infact given a lot of headache to the tools project. I have an issue for it at JBWS-227 and spent a lot of time in April trying to match wscompile array handling. This also is related to my comments (item b) on why we cannot make Xerces schema parser strict[JBWS-244]. Schema generated by wscompile is not liked by xerces sometimes. I have had a question on this for few days now, as to how wscompile equates the namespace defined by prefix "ns2" in the first namespace to the targetnamespace in the second. Internally, wscompile does use Xerces for schema parsing. Analyzing the schema generated by wscompile, this is valid schema, I think. In this case, it seems to be a bug with the Xerces schema implementation. Its just that the Xerces schema implementation requires schema elements to be in seperate files. Just as a test, can we try changing the order of the schema elements? In that case the type "HelloObj" will be read into the Xerces parser before it starts handling the second schema with the array type. If it fails, we should file a bug on the Xerces website.	1	1	https://issues.jboss.org/browse/JBWS-268
13393	JBWS-265	tdiesler@satellite /cygdrive/d/projects/jboss-head/testsuite $ ant -Dtest=org.jboss.test.webservice.samples.ServerSideJMSTestCase one-test one-test: [junit] Running org.jboss.test.webservice.samples.ServerSideJMSTestCase [junit] Tests run: 1, Failures: 0, Errors: 1, Time elapsed: 1.272 sec [junit] Test org.jboss.test.webservice.samples.ServerSideJMSTestCase FAILED 2006-02-20 23:53:20,911 WARN [org.jboss.ejb.plugins.jms.JMSContainerInvoker] JMS provider failure detected for OrganizationMDB org.jboss.deployment.DeploymentException: Could not find the queue destination-jndi-name=queue/MessageDrivenEndpointQueue; - nested throwable: (javax.naming.NameNotFoundException: MessageDrivenEndpointQueue not bound) at org.jboss.ejb.plugins.jms.JMSContainerInvoker.innerStartDelivery(JMSContainerInvoker.java:767) at org.jboss.ejb.plugins.jms.JMSContainerInvoker.startService(JMSContainerInvoker.java:839) The temporary queue is not created. The client fails with javax.naming.NameNotFoundException: MessageDrivenEndpointQueue not bound at javax.naming.InitialContext.lookup(InitialContext.java:351) at org.jboss.test.webservice.samples.ServerSideJMSTestCase.testSOAPMessageToEndpointQueue(ServerSideJMSTestCase.java:65)	1	1	https://issues.jboss.org/browse/JBWS-265
13394	JBWS-258	Verification of JBWS-707 needs to be done. Can we schedule it for the next CTS run, Thomas?	0	0	https://issues.jboss.org/browse/JBWS-258
13395	JBWS-219	Your last changes provide access to the EJB3 meta data after the deployment create step but not to the annotations that are present on the bean class. Would it be possible to load the bean class in the create step like this? public void start() throws Exception { initializeClassContainer(); for (int i = 0; i < constructors.length; i++) { if (constructors[i].getParameterTypes().length == 0) { defaultConstructorIndex = i; break; } } initializePool(); createCallbackHandler(); } public void create() throws Exception { clazz = classloader.loadClass(beanClassName); }	1	1	https://issues.jboss.org/browse/JBWS-219
13396	JBWS-219	The annotations that WS is interested in are JSR-181 annotations that will be given at the bean class. My point remains, would it be possible to have the bean class available in the create step?	1	1	https://issues.jboss.org/browse/JBWS-219
13397	JBWS-218	The new build will do it "automagically" and should be extensible to other IDEs? http://cvs.sourceforge.net/viewcvs.py/jboss/jbossbuild/src/main/org/jboss/ant/tasks/component/IDETask.java?rev=1.1.1.1&view=markup	0	0	https://issues.jboss.org/browse/JBWS-218
13398	JBWS-211	Does this list cover all we need in both WS testsuites? If not, what is missing?	0	0	https://issues.jboss.org/browse/JBWS-211
13399	JBWS-158	The parsing is part of the marshalling layer. Alex could you have a look at integrating StAX into JBossXB or comment appropriately?	0	0	https://issues.jboss.org/browse/JBWS-158
13400	JBWS-157	Can this perhaps be resolved by establishing a convention that <soap:address> entries that contain a valid URL with a valid protocol part will not be modified? An additional boolean switch in jboss-service.xml can help not break existing deployments. <attribute name="ModifySOAPAddress">true</attribute>	0	0	https://issues.jboss.org/browse/JBWS-157
13401	JBWS-157	This sounds good to me. Can you also add the <virtual-host></virtual-host> element to the jboss-web.xml element that gets created so that it can be synced to virtual hosts within Tomcat?	0	0	https://issues.jboss.org/browse/JBWS-157
13402	JBWS-154	Note from Thomas: if ("element".equals(localName)) { QName xmlName = navigator.resolveQName(attrs.getValue("name")); QName xmlType = navigator.resolveQName(attrs.getValue("type")); return wsdlTypes.addSchemaElement(xmlName, xmlType); } An element is never the child of WSDLTypes, it is the child of a schema element. How would this code cope with multiple schemas with different namespaces? I don't see a class that abstracts the schema element in WSDLTypes	0	0	https://issues.jboss.org/browse/JBWS-154

#	Requirement	Questions	Answered	Successfully Answered	URL
13403	JBWS-147	I reviewed the tests and seperated the XSDToJavaTestCase into * AnonymousTypesTestCase * ComplexTypesTestCase * NillableTypesTestCase * ReferencesTestCase * SimpleTypesTestCase A) There are a few more directories with xsd files in cvs that I assume to be input for more tests. These are * annotations * schemaimport * schemainclude Are they not covered by tests yet? If this is the case, please provide the tests and the associated test fixtures generated with wscompile. That these additional tests would fail is not a problem. Test failures can be assigned individual JIRA tasks. JBWS-147 is only about generating a comprehensive test suite and *not* about providing the implementation that make it work. B) The xsd2java implementation should be separated from WSDL concerns. The API that it used by the tests seems not logical SchemaBuilder builder = new SchemaBuilder(); XSDSchema schema = new XSDSchema(); schema = builder.parseSchema(xsdFile.toURL(), schema); WSDLToJava wsdl2java = new WSDLToJava(); wsdl2java.generateSEI(schema, dir, packagename, true); Can we clean this up and provide a simple xsd2java API that is well documented in javadoc? This becomes more important as we will potentially fork out some of this work to partners (e.g. Novell) C) The tests generate java code in cvs locations. This should be moved to the output directory. D) Some tests exist in two versions, one set uses an internal xsd object representation the other set uses the xerces xsd parser. What is the situation here? Have we decided to go with xerces for the reading and our own implementation that extends xerces for writing? Thanks for the effort	1	1	https://issues.jboss.org/browse/JBWS-147
13404	JBWS-137	I presume this issue also has dependency on JBWS-644. (http://jira.jboss.com/jira/browse/JBWS-644) Since the JAX-RPC 2.0 implementation is still not there should i work on this issue?	1	0	https://issues.jboss.org/browse/JBWS-137
13405	JBWS-136	i am writing test cases for this conformance. The jaxrpc api classes dont seem to be there for versio 2.0 For example i am not able to find javax.xml.jaxrpc.BindingProvider And javax.aml.jaxrpc.Call interface must extend javax.xml.jaxrpc.BindingProvider... Please let me know if i am missing out something. How do i add jaxrpc2.0 apis to check the conformance.	1	1	https://issues.jboss.org/browse/JBWS-136
13406	JBWS-116	Is this going to be fixed? Obviously the work around is a solution but surely it would be better if it was correctly generated.	1	0	https://issues.jboss.org/browse/JBWS-116
13407	JBWS-104	A month down the line, did you hava a chance to get any progress on this?	0	0	https://issues.jboss.org/browse/JBWS-104
13408	JBWS-75	Thanks for getting started on JBWS-75. The performance and benchmark tests should run in an automated fashion as part of an ant task. This is important because they will run as part of the jboss-labs performance regression tests. Unfortunately, we cannot have any dependencies on third party software that we cannot distribute. What I am initially looking for is a set of ws deployments that cover common ws processing use cases. What immediately jumps to mind are use cases like: 1) trivial soap messages 2) complex user types 3) array types I would assume that we are not the first who think about measuring ws performance. So ideally, there is a set of wsdl available somewhere that form the basis of good performance tests. If this fails, maybe there is a paper available that allows us to derive the needed set of wsdl files. From the wsdl files we derive the deployments and the test cases. Initially we would have a build for jbossws deployments. Later we would add corresponding deployments for JWSDP, Axis, etc to actually do a comparison.	1	1	https://issues.jboss.org/browse/JBWS-75
13409	JBWS-72	I reviewed the matrix version 1.0, here is some feedback. Chapter 4 Generally I don't like the disconnect of command line options and config file. For example this forces me to include the command line options that I use with a particular config.xml as an xml comment to be able to reproduce the same results. jbossws should opt for a configfile with as little command line options as necessary. options 2) why do we need -cp and -classpath 4) what does -define do 5) why are -f:features not supported 10) what is -gen:wsdl needed for 13) what does -import do features 3) donotoverride is useful 7) do we even generate stubs? otherwise what does nodatabinding do? 9) what is nomultirefs 11) is duplicated 16) what is serializeinterfaces 17) strict vs. lenient ? 19) how does useonewayoperations ---------- The matrix could be split up in server and client side. Server side takes SEI as input and generates wsdl, schema, jaxrpc-mapping. This is of higher importance than client side. Client side takes wsdl as input and generates SEI, SI, complex types, holders. One of my biggest questions is: Do we need generated stubs or can we get away with dynamic proxies. Could you please create fine grained jira tasks and associate them with the features you like to support. Java2Wsdl * primitive types, no schema * complex types + schema * rpc/literal * doc/literal * INOUT, OUT paramters * bound and unbound headers * bound and unbound attachments * one way operations * jaxrpc-mapping * what else? wsdl2Java * primitive types, no schema * complex types + schema * schema with anonymous types * service interface * rpc/literal * doc/literal * INOUT, OUT paramters * bound and unbound headers * bound and unbound attachments * one way operations * what else? cheers	1	1	https://issues.jboss.org/browse/JBWS-72
13410	JBWS-72	Thomas::Generally I don't like the disconnect of command line options and config file. For example this forces me to include the command line options that I use with a particular config.xml as an xml comment to be able to reproduce the same results. jbossws should opt for a configfile with as little command line options as necessary. Anil: JbossWSTools will use an xml config file just the way wscompile uses. Thomas: 2) why do we need -cp and -classpath 4) what does -define do 5) why are -f:features not supported 10) what is -gen:wsdl needed for 13) what does -import do Anil: -import reads the WSDL file, generates SEI and a template implementation of the SEI. -gen:wsdl is a jbosswstool option to generate wsdl (that is the only server side artifact that we can generate given a SEI) -define allows you to read a SEI and define a javax.xml.rpc.Service interface for the SEI. -cp and -classpath allow u to locate the SEI compiled classes. Thomas: 3) donotoverride is useful 7) do we even generate stubs? otherwise what does nodatabinding do? 9) what is nomultirefs 11) is duplicated 16) what is serializeinterfaces 17) strict vs. lenient ? 19) how does useonewayoperations Anil: First phase of tools, we will always regenerate the SEI.(Regarding donotoverride) Thomas: One of my biggest questions is: Do we need generated stubs or can we get away with dynamic proxies. Anil: I have been having the same question. Generated stubs are not portable(The easiest way for the users to use JBossWS, though.)	1	1	https://issues.jboss.org/browse/JBWS-72
13411	JBWS-72	Thomas, can you review the JBossWS Tools document to see if I need to add anything to the feature matrix.	1	1	https://issues.jboss.org/browse/JBWS-72

#	Requirement	Questions	Answered	Successfully Answered	URL
13412	JBWS-72	I reviewed the doc, here is some feedback: 1. This issue asks you to backup each feature by a corresponding JIRA task. This still needs to be done. The created tasks should be linked to this issue. 2. The document should reflect the intended usage of the jbossws tools for the first release, not the current state. Otherwise you will have a hard time to keep the documentation in sync. 3. Table 2.1 states command line options for jbossws. AFAIK, we agreed on removing the disconnect between the command line options and the config file. What is the minimal set of cmd line options we need? 4. Use multiple digit serial numbers, so we can refer to them uniquely accross the document 5. Table 2.3 sn-4: exlain the benefit and usage of 'define' (This is repeated feedback from 11/Feb) 6. Table 2.3 sn-8: exlain the benefit and usage of 'import' (This is repeated feedback from 11/Feb) 7. Table 2.4 sn-5: exlain the benefit and usage of 'explicitcontext' 8. Table 2.4 sn-6: exlain the benefit and usage of 'jaxbenumtype' 9. Table 2.4 sn-5: exlain the benefit and usage of 'nodatabinding' 10. Table 2.4 sn-10: exlain the benefit and usage of 'norpcstructures' 11. Table 2.4 sn-11: is duplicated, should this be another feature? 12. Table 2.4 sn-12: change to NOT SUPPORTED 13. Table 2.4 sn-13: exlain the benefit and usage of 'resolveidref' 14. Table 2.4 sn-16: exlain the benefit and usage of 'serializeinterfaces' 15. Table 2.4 sn-18: lets make wrapped the default 16. Table 2.4 sn-19: how do intend to support 'useonewayoperations'? 17. Generally, all supported features should be explained. the description collumn is often not meaningful enough 18. Replace the examples using wscompile with how one uses jbossws. Show the typical use cases. 19. Table 2.6 Do we support CustomType Inheritance, Enumerations 20. Configuring jbossws.sh: There should be a set of tool jars. Remove dependency on JBOSS_HOME 21. Running jbossws.sh: Replace by intended xml configutration 22. 4.1.2 Remove API documentation. This is covered by javadoc. Avoid duplicate/redudant documentation sets Please update the document version an have a date at the front page. I suggest you create JIRA tasks for features that you describe in the document, but are not implemented yet. Then you can reference them from the document, allowing us to continously have up to date information. The relevant comments from the JIRA tasks can then be incorporated in the document and references can be removed with the next revision.	1	1	https://issues.jboss.org/browse/JBWS-72
13413	JBWS-72	Anil wrote: Hi Thomas, here are the minutes of our meeting. Took me sometime to format it. Main goal: Previous exercise was to generate a feature matrix of jbossws against wscompile (JBWS-72). Now it is time to generate a document that just talks of jbossws features and how to use it. (JBWS-321) - Need to get a schema written for the tools configuration. - the attributes in the configuration file should be of mixed case (AND not all lower case). Section 2.1.1.2 jbossws configuration * interface is a bad name. * packagename should be removed. * wsdldest: bad attribute. Should be removed and all directory related attributes should be removed. - assume a sub directory called "jbossws" and generate the artifacts there. * includetypesinwsdl: should be renamed to includeSchemaInWSDL * wsxml element - remove the directory related attribute. - add an attribute for servletlink. Section 2.1.2 Question from Thomas: Is the Service Interface being generated? Answer from Anil: It is being generated but I forgot to mention it in the document. - remove packagename and seidir attribute. Section 2.1.3: WS-I Compliance Remove this section as jbossws will be ws-i compliant. Section 2.1.4: wscompile features section should be completely replaced with jbossws features. Notes on Document Style Web Services:- * jbossws should not doublewrap classes (wscompile does wrapping of wrapped classes). - there should be a option for user to tell jbossws that it is using wrapped classes. - jbossws should validate that the wrapped request structure is valid. * Document/Literal Wrapped is Microsoft technology which is deprecated and no support should be available. For WSDL-> Java * User should be allowed to inform jbossws whether to unwrap the document request/response structures or not. Unwrap should be the default. 2.1.4.10 nodatabinding feature is not required. Section 2.3 JSR-181 Annotations Support * Support should be for both interfaces and implementation classes. If the implementation class has the @interface annotation, then jbossws should read the interface for the methods to expose. * Oneway operations - JSR181 refers to oneway capabilities at the method level. - For unannotated interfaces, there should be a way to specify which methods are oneway. Section 2.3.,4 SoapBinding Annotation - Support custom jbossws annotations to take in the jboss.xml attributes like port-component-uri, context root etc. Section 4: Java To WSDL/XSD Focus on annotated types in the examples. - Remove all the API details. Section 5: Architecture should be removed. First Release should include the following: 1) Documented Schema. 2) User Guide. 3) Annotated endpoint support + oldstyle J2EE 1.4 artifacts.	1	0	https://issues.jboss.org/browse/JBWS-72
13414	JBWS-67	I will start contributing to the WSDL Validators project of Eclipse Web Tools. They are working on WSDL 2.0 validators (in cooperation with the W3C). Why reinvent the wheel?	1	1	https://issues.jboss.org/browse/JBWS-67
13415	JBWS-35	should we promote this issue to resolved status rather than a voting status? We need support for the following : a) Support for SAML Token Profile (Higher Priority). b) Support for Kerberos Token Profile (Low Low Priority). This is important in intranets that hosts windows-centric systems.	0	0	https://issues.jboss.org/browse/JBWS-35
13416	JBWS-31	I'll take this one, though I currently don't have permission to assign it to myself. Could someone give me JIRA rights? Thanks, Jason	1	1	https://issues.jboss.org/browse/JBWS-31
13417	JBWS-28	Is the generated schema going to be stored on disc somewhere? I would wait for the implementation of XSD generation (by Anil?) first.	1	1	https://issues.jboss.org/browse/JBWS-28
13418	JBWS-24	Alex says: As I understand, we are given *.wsdl, webservices.xml and jaxrpc-mapping.xml (optional?). What happens then? We generate marshaller/unmarshaller for jaxrpc parameters? Is there any other source of Java - XML binding info except for jaxrpc-mapping.xml?	0	0	https://issues.jboss.org/browse/JBWS-24
13419	JBWS-24	It seems logical to have the same tool generate Java classes and their mapping to XML elements. Is the mapping file going to be optional? If so, let's first try w/o it. So the client will use the same tool to generate Java classes, serializers/deserializers, mapping file as the one used on the server side? It seems to be a weakness to me, since how can you be sure the client has this same tool? Assuming there is no standard for these tools' output.	0	0	https://issues.jboss.org/browse/JBWS-24
13420	JBWS-24	Is Anil working on Java classes generation? If so I will focus on default Java to XML names mapping support.	0	0	https://issues.jboss.org/browse/JBWS-24
13421	JBWS-24	Ok, then I will work on that. What do we use right now for code generation? wscompile? Some Axis tool?	0	0	https://issues.jboss.org/browse/JBWS-24
13422	JBWS-24	So right now we some Axis tool for generation? In what case do we need to generate Java classes from XSD? When a client deploys a web-service into JBoss, we need to generate only serializers/deserializers, right?	0	0	https://issues.jboss.org/browse/JBWS-24
13423	JBWS-11	This task is completed. wsdl4j does not handle XML Schema ie. <types> in wsdl. A collaboration with JBossXB, which will deal with XSD Types to Java, will handle it.	0	0	https://issues.jboss.org/browse/JBWS-11
13424	JBWS-1	All the CTS tests in jaxrpc, saaj, webservices have one ore more SEI. Generating a valid wsdl comes in two parts: 1. generate wsdl definition 2. generate contained xsd types Checkout if XMLBeans can help with generating XMLSchema from java. Are there other tools? Xerces is suposed to have a schema parser. Julien sais: don't try to write your own xsd parser!	1	1	https://issues.jboss.org/browse/JBWS-1
13425	WFK2-755	one-off patch	1	1	https://issues.jboss.org/browse/WFK2-755
13426	WFK2-755	[~ivassile], according to workflow, this issue should be now in "QA in progress" state. Can you please update it?	1	1	https://issues.jboss.org/browse/WFK2-755

#	Requirement	Questions	Answered	Successfully Answered	URL
13427	WFK2-755	[~ppitonak]: Hold on the testing on this patch. I just received information from the customer, that they are using a patched version of RF 4.3.5 which includes issues RF-13719 and RF-13733. Both issues are not included in RF 4.3.7 and this patch will roll-back the changes. The client prefers a 4.3.5 (WFK 2.5) patch. I'm going to build cumulative patch based on [1]. [1] https://access.redhat.com/jbossnetwork/restricted/softwareDetail.html?softwareId=32893&product=web.framework.kit&version=2.5.0&downloadType=patches	1	1	https://issues.jboss.org/browse/WFK2-755
13428	WFK2-755	RF-13719 is an issue for RichFaces in portal environment. Can you confirm that the customer uses portal and therefore we need to test this patch in portal?	1	1	https://issues.jboss.org/browse/WFK2-755
13429	WFK2-755	Verified, both *RF-13640 : Autocomplete with custom layout: cant select item by mouse* and *RF-13733 : dataTable built-in sorting and filtering* work with this patch. No regressions have been found. Verified on IE 10, IE 11, Firefox 24, Firefox 31, latest Chrome, on EAP 6.2.4.GA. [~pmensik], please, could you check the Portal issue: RF-13719 : rich.fileUpload breaks form action in portal ?	1	0	https://issues.jboss.org/browse/WFK2-755
13430	WFK2-396	Added to WFKDOC-16 jira tracking update of supported configs page for WFK 2.5 by docs team.	0	0	https://issues.jboss.org/browse/WFK2-396
13431	WFK2-396	What is the purpose of this upgrade? an issue?	1	1	https://issues.jboss.org/browse/WFK2-396
13432	WFK2-396	[~manaRH] this update contains @InFrame related bugfixes and support for IE9 - more details here http://arquillian.org/blog/2014/01/31/arquillian-graphene-2-0-1-Final/ Why is JDF issues needed here?	1	0	https://issues.jboss.org/browse/WFK2-396
13433	WFK2-326	[~mmurray] not sure if this task is solved/reviewed. Could you please process it till the end - closing ;-). Thanks in advance!	1	0	https://issues.jboss.org/browse/WFK2-326
13434	WFK2-324	Is it ready to release? please provide staging repository for QE	0	0	https://issues.jboss.org/browse/WFK2-324
13435	WFK2-324	[~jowilson] Based on Pete's comments: https://issues.jboss.org/browse/WFK2-321?focusedCommentId=12925810&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12925810 Shouldn't we change to 1.0.3.GA instead ?	1	1	https://issues.jboss.org/browse/WFK2-324
13436	WFK2-323	Is it ready to release? please provide staging repository for QE	0	0	https://issues.jboss.org/browse/WFK2-323
13437	WFK2-321	Where these Archetypes should be released ? maven.repository.redhat.com or Maven Central ?	0	0	https://issues.jboss.org/browse/WFK2-321
13438	WFK2-321	The version used on HTML5 Archetypes is 2.4.0.Final. Should we change it to 2.4.0.GA or 2.4.0-redhat-x ? Will Spring and Richfaces Archetype also follow the WFK version ?	0	0	https://issues.jboss.org/browse/WFK2-321
13439	WFK2-168	[~burrsutter], I'd like one clarification - should this be considered complete only if TiMo uses a native API? i.e. should your second comment also be fulfilled? The first comment (import TiMo sources into JBDS 7.1) can be done, but with a minor concern - we'd have to load cordova.js without using Require.js. Currently in TiMo, a version and platform specific {{cordova.js}} file is loaded through require.js based on the runtime platform. We'll instead have to create a separate HTML page for cordova apps (like app.html instead of index.html) which * references cordova.js that is managed by the tooling, and * is referenced in the config.xml file as the initial page for the app. I don't see how we can 'fix' this in the tooling to support require.js alongside cordova.js.	1	1	https://issues.jboss.org/browse/WFK2-168
13440	WFK2-153	Hi Sanne, do you plan to finish this issue into wfk 2.5.0, or will it be postponed?	1	1	https://issues.jboss.org/browse/WFK2-153
13441	WFK2-152	Can we change the title of this jira to lose the word "integration"? It always throws me off - it is just Spring + Websockets, not "Spring Integration"	1	1	https://issues.jboss.org/browse/WFK2-152
13442	WFK2-152	[~emil.cervenan] awesome - can you publish an eample/quickstart formatted out to github some place so other parties could look at your setup?	1	1	https://issues.jboss.org/browse/WFK2-152
13443	WFK2-152	I tried your example on JBoss EAP 6.3.0.Alpha1 (AS 7.4.0.Final), but that one doesn't seem to support websockets yet: 16:26:24,654 ERROR [org.jboss.msc.service.fail] (MSC service thread 1-7) MSC000001: Failed to start service jboss.deployment.unit."spring-websockets.war".PARSE: org.jboss.msc.service.StartException in service jboss.deployment.unit."spring-websockets.war".PARSE: JBAS018733: Failed to process phase PARSE of deployment "spring-websockets.war" at org.jboss.as.server.deployment.DeploymentUnitPhaseService.start(DeploymentUnitPhaseService.java:166) [jboss-as-server-7.4.0.Final.jar:7.4.0.Final] at org.jboss.msc.service.ServiceControllerImpl$StartTask.startService(ServiceControllerImpl.java:1980) [jboss-msc-1.1.5.Final.jar:1.1.5.Final] at org.jboss.msc.service.ServiceControllerImpl$StartTask.run(ServiceControllerImpl.java:1913) [jboss-msc-1.1.5.Final.jar:1.1.5.Final] at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) [rt.jar:1.7.0_45] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) [rt.jar:1.7.0_45] at java.lang.Thread.run(Thread.java:744) [rt.jar:1.7.0_45] Caused by: org.jboss.as.server.deployment.DeploymentUnitProcessingException: JBAS018014: Failed to parse XML descriptor "/C:/repos/3rdparty/jboss/jboss-wfk-quickstarts/spring-websockets/content/spring-websockets.war/WEB-INF/jboss-web.xml" at [20,3] at org.jboss.as.web.deployment.JBossWebParsingDeploymentProcessor.deploy(JBossWebParsingDeploymentProcessor.java:79) at org.jboss.as.server.deployment.DeploymentUnitPhaseService.start(DeploymentUnitPhaseService.java:159) [jboss-as-server-7.4.0.Final.jar:7.4.0.Final] ... 5 more Caused by: javax.xml.stream.XMLStreamException: ParseError at [row,col]:[20,3] Message: Unexpected element 'enable-websockets' encountered Where can I find EAP 6.3.0.ER3? Thanks in advance, Geert.	1	1	https://issues.jboss.org/browse/WFK2-152
13444	WFK2-152	[~manaRH], [~trepel]: currently the WFK Spring Guide states "Support for Spring 4 as a tested framework is limited to the features supported by Red Hat JBoss Enterprise Application Platform. Specifically, support is not provided for the new features in Spring 4 related to Java 8, Java EE 7, and WebSockets." Does this needs to be updated in light of this jira?	1	0	https://issues.jboss.org/browse/WFK2-152
13445	WFK2-152	[~trepel], [~manaRH] - I will fix the text in the Spring doc. To make it obvious to users, let include an RN text to state that WebSockets is supported. Can one of you complete the RN text please.	1	0	https://issues.jboss.org/browse/WFK2-152
13446	WFK2-138	How will we handle BOMs in product repository? I assume they still will be delivered in shipped WFK 2.7.0 product repository and not through jboss-developer ways or deployment process - so they will accessible on developer pages as today. Do I understand it correctly?	0	0	https://issues.jboss.org/browse/WFK2-138
13447	WFK2-138	As far as I know, all BOMS ship with their respective product maven repository .zip - they should also land at maven.repository.redhat.com. Is there anything special about the WFK BOMs that require Nexus (other than Staging) which is actually now being architected outside of the Nexus environment?	0	0	https://issues.jboss.org/browse/WFK2-138
13448	WFK2-123	Also, @burrsutter, we don't currently have Topcoat at any tier. Do you think we should add it to tier 4 because of this?	1	1	https://issues.jboss.org/browse/WFK2-123

#	Requirement	Questions	Answered	Successfully Answered	URL
13449	WFK2-121	Apparently I looked at the wrong Backbone-based Kitchensink. :-) What is the URL to the sources you are looking at? [~vineet.reynolds] Do you think there would be value in removing the responsive aspect so that it is mobile all the time? The reason I ask is...concern for learning curve as well as ease of support in the tooling.	1	0	https://issues.jboss.org/browse/WFK2-121
13450	WFK2-115	[~alrubinger]what is the approximate timeframe for this item? Mid-2014? or mid-2015?	1	0	https://issues.jboss.org/browse/WFK2-115
13451	WFK2-113	[~burrsutter], concerning your statement about no "responsive design", what specifically are you asking for here. Some of RWD covers "resizing" based on the viewing area so that it will fit on desktops to mobiles. Is this what you are referring to? Because RWD also refers to the use of JavaScript and JS libs. As you have already asked for jQuery Mobile, I am thinking that you are not talking about building a site without JS. So can you please clarify the 'no "responsive design"' request?	1	1	https://issues.jboss.org/browse/WFK2-113
13452	WFK2-113	[~pmuir] I have started working on this as a "contacts" app as [~burrsutter] requested. Should I continue under a different JIRA or stop work on this?	1	1	https://issues.jboss.org/browse/WFK2-113
13453	WFK2-113	It would be good to have a non-Angular and with-Angular option - so people can compare apples to apples. If we only have time for the with-Angular option that demonstrates create, read, update and delete operations, ideally with the additional fields/properties (e.g. birthdate) then that would work for now.	1	1	https://issues.jboss.org/browse/WFK2-113
13454	WFK2-113	The app I created was tasks based. So if we want to create a non-angular option, we could clone it's UI? To keep them v similar.	1	1	https://issues.jboss.org/browse/WFK2-113
13455	WFK2-113	[~pmuir] & [~burrsutter] do you want to keep this Jira focused on a "contacts" app? Or do you want a modified version of Pete's "task" app that doesn't have AngularJS? If the later, who will build it? I have the time now to do either.	1	1	https://issues.jboss.org/browse/WFK2-113
13456	WFK2-108	Do you know if there are any plans for that Pete?	1	1	https://issues.jboss.org/browse/WFK2-108
13457	WFK2-108	[~petemuir], there seems to be little progress on this jira and it's subtasks. Will it make it for 2.5? Also, how is 'online delivery of BOMs' related to the current tech preview online red hat/jboss maven repo?	1	1	https://issues.jboss.org/browse/WFK2-108
13458	WFK2-108	Hi [~petemuir], will the pages appear in the new design soon enough to make this feature work in WFK 2.6?	1	1	https://issues.jboss.org/browse/WFK2-108
13459	WFK2-108	There is already some WFK 2.5 content on the new Developer site, so it would make sense to update those to WFK 2.6 eventually. Can we update those after WFK 2.6 GA then?	1	1	https://issues.jboss.org/browse/WFK2-108
13460	WFK2-108	[~manaRH], can you clarify what this will look like for 2.7 please? * Quickstarts - from jboss-developer website, github and WFK quickstarts zip on customer portal * TiMo - from jboss-developer website (as for 2.6), github (as for 2.6) and WFK TiMo zip on customer portal * BOMs - from jboss-developer maven repo and WFK maven zip on customer portal Will the RF, HSearch, and Snowdrop demo.zip files continue to be distributed from the customer portal? will these also be available from jboss-developer in some way?	1	1	https://issues.jboss.org/browse/WFK2-108
13461	WFK2-72	[~maschmid] when do you need the release complete by in order for it to make WFK 2.4?	1	1	https://issues.jboss.org/browse/WFK2-72
13462	WFK2-48	Hey Andrew, do you know if that is something that is being/will be worked on?	0	0	https://issues.jboss.org/browse/WFK2-48
13463	WFK2-46	Yes, of course. Do we have due dates for the Use Case documents?	0	0	https://issues.jboss.org/browse/WFK2-46
13464	WFK2-46	Why the reassign, [~rruss]? You'd still like use-case docs from me for this, yes?	1	1	https://issues.jboss.org/browse/WFK2-46
13465	WFK2-45	Same comment as in WFK2-6: what is the work to be done to get this included?	1	1	https://issues.jboss.org/browse/WFK2-45
13466	WFK2-45	These are two separate tasks; for your tracking as well, [~pmuir], do you mind if I split this issue into one each for Persistence and Tx?	1	1	https://issues.jboss.org/browse/WFK2-45
13467	WFK2-45	[~maschmid] when do you need the release complete by in order for it to make WFK 2.4?	1	1	https://issues.jboss.org/browse/WFK2-45
13468	WFK2-25	[~alrubinger]What is the customer impact of this item? At this time, the wildfly/upstream adapter seems to work fine with EAP, at least in the quickstarts/archetypes that leverage Arquillian. Should we continue with reusing the same adapter or is now the time to fork?	1	0	https://issues.jboss.org/browse/WFK2-25
13469	WFK2-25	[~kpiwko]thank you [~alrubinger] what will it take to get this scheduled?	1	0	https://issues.jboss.org/browse/WFK2-25
13470	WFK2-25	Do we have a team in place for maintenance of this component? Also, shipment is via Maven Central or mave.repository.redhat.com? (it should not ship as part of the WFK distribution)	0	0	https://issues.jboss.org/browse/WFK2-25
13471	WFK2-16	I read the doc and wanted to know what you saw that needs to be changed. We are going to changed the directory structure to align with an add-on. And provide a version of the standalone.xml so that no one needs to modify it. What else?	0	0	https://issues.jboss.org/browse/WFK2-16
13472	WFK2-15	Should that certification include also Seam (Seam examples) ?	0	0	https://issues.jboss.org/browse/WFK2-15
13473	WFK2-15	Could Openshift access internal maven repository or extracted zipped maven repository? This probably needs a discussion with openshift team.	0	0	https://issues.jboss.org/browse/WFK2-15
13474	WFK2-15	Can we use OpenShift Enterprise for the testing? I don't see how you intend to get OpenShift online to access something behind the VPN, nor what the OpenShift team could do about this.	1	1	https://issues.jboss.org/browse/WFK2-15
13475	WFK2-15	[~burrsutter] WDYT? Can we claim that this is done if we are missing the necessary instructions in the readmes?	1	1	https://issues.jboss.org/browse/WFK2-15
13476	JBWIKI-106	Is there a reason for which the two packages that you mentioned are not in the repo?	0	0	https://issues.jboss.org/browse/JBWIKI-106
13477	JBWIKI-94	Do i have to add this also to binaries released before acquisition ?	0	0	https://issues.jboss.org/browse/JBWIKI-94
13478	JBWIKI-76	what is standard diff output ?	0	0	https://issues.jboss.org/browse/JBWIKI-76
13479	JBWIKI-17	Tomek, is this finished? I thought that roles aren't included until 1.0 Final. Am I wrong?	1	0	https://issues.jboss.org/browse/JBWIKI-17
13480	JBWIKI-2	What exactly do we need here ? (you wrote sth. about spliting this up)	1	1	https://issues.jboss.org/browse/JBWIKI-2
13481	JBWIKI-2	There should be a link... in every wikipage... (or above.. next to history).. that allows people to subscribe to all changes.. in a wikipage... which gives: 1. The new wikipage 2. The changes made via the 'diff' format. .... as a blog entry in rss/atom/rdf.	1	1	https://issues.jboss.org/browse/JBWIKI-2
13482	JBXB-246	Thanks for doing this. But, could you give an example of what the output should look like now, and which part is the fileName (is it "URI"?). I'm get the following message which I'm having a heck of time interpreting enough to diagnose: "Failed to resolve schema nsURI= location=web-fragment" What is "ns"? why no URI? If it had a URI, what would that be indicating (the filename?)? etc Thanks in advance for any more info that can be provided.	1	0	https://issues.jboss.org/browse/JBXB-246
13483	JBXB-246	In fact...the message I'm seeing is the same as the one noted in the referenced conversation that prompted this change. I'm running with the 6.0 M4 - why am I not getting the updated format? Should I be?	1	0	https://issues.jboss.org/browse/JBXB-246
13484	JBXB-244	Yes, that's exactly what David is asking for. Should we switch AllBinding implementation to UnorderedSequenceBinding?	0	0	https://issues.jboss.org/browse/JBXB-244
13485	JBXB-206	What I meant is that the issue is probably the resolver is not configured properly. It's not enough to just instantiate it. What namespace is resolved with XsdBinder?	0	0	https://issues.jboss.org/browse/JBXB-206

#	Requirement	Questions	Answered	Successfully Answered	URL
13486	JBXB-194	How would that be different from SchemaBinding.setNormalizeSpace() which is false by default and SchemaBinding.setIgnoreWhitespacesInMixedContent() which is true by default?	1	0	https://issues.jboss.org/browse/JBXB-194
13487	JBXB-100	AFAICS, just Source won't be more efficient. Since it gives access only to the systemId of the source. Or am I missing something?	1	0	https://issues.jboss.org/browse/JBXB-100
13488	JBXB-100	So what are you asking for? To add unmarshal(javax.xml.transform.Source)? Then its implementation will be unmarshal(Source s) { return unmarshal(s.getSystemId()); } Or unmarshal(DOMSource ds) and navigate through the Nodes producing SAX-like events?	1	0	https://issues.jboss.org/browse/JBXB-100
13489	JBXB-100	How about: unmarshal(Source source) { if (source instanceof DOMSource) { Node node = ((DOMSource)source).getNode(); ... // process node } else throw new UnsupportedOperationException(); } Or you prefer to use DOMSource class directly in your method signatures?	0	0	https://issues.jboss.org/browse/JBXB-100
13490	JBXB-87	Thanks, Alex. Right in time ;) Still I would need a quick explanation. Could we have a quick skype session this afternoon? /Heiko	1	0	https://issues.jboss.org/browse/JBXB-87
13491	JBXB-65	Sorry, I didn't mentioned it before ... There is called 'content.get(index)' also on other methods in this class. That's why I did't simple used Iterator in my test (I have used ArrayList). The other occurences are more complicated to rewrite to iterator. Are you sure, that it will not slowdown other processes?	1	1	https://issues.jboss.org/browse/JBXB-65
13492	JBXB-65	What's the reason to stick with the linked list? Compared to the ArrayList replacement, the Iterator approach doesn't really have an impact on performance: Without any changes: testCase resultValue dataFile 3 Items 2983.500 data/po.xml 20 Items 1013.800 data/po-20.xml 50 Items 454.800 data/po-50.xml 100 Items 219.200 data/po-100.xml With the Iterator change: testCase resultValue dataFile 3 Items 3410.600 data/po.xml 20 Items 1074.900 data/po-20.xml 50 Items 489.700 data/po-50.xml 100 Items 257.400 data/po-100.xml With the ArrayLIst change: testCase resultValue dataFile 3 Items 4250.000 data/po.xml 20 Items 1415.100 data/po-20.xml 50 Items 650.200 data/po-50.xml 100 Items 342.100 data/po-100.xml	1	1	https://issues.jboss.org/browse/JBXB-65
13493	JBXB-42	Thomas. Can you please discuss things on public mailing lists before scheduling work to others! Otherwise this is a feature request not a task! I'm going to add some of the related "productization" issues. There are others. Which probably makes this task redundant.	1	1	https://issues.jboss.org/browse/JBXB-42
13494	JBXB-40	Alex, Why is this closed without comment?	0	0	https://issues.jboss.org/browse/JBXB-40
13495	JBXB-33	You proposed to add SchemaBindingErrorHandler to not depend on Xerces API? I agree it's a good idea. But then I would also get rid of org.apache.xerces.xs.StringList. Is that ok for you? How do you create an instance of that? Anonymous class?	1	0	https://issues.jboss.org/browse/JBXB-33
13496	JBXB-33	Doesn't the Xerces parser need a StringList to parse a set of Schema files? If there is a way in which you will be ok with a List of schema locations (List<String> or List<URL>), I am fine. The parsing is all yours. All I care about is passing a SchemaBindingResolver and an ErrorHandler to the Util method. :-) We have a RW based Xerces model and I do have a way of constructing a StringList via JBossStringList class. But as you said, lets not tie the util method to StringList.	0	0	https://issues.jboss.org/browse/JBXB-33
13497	JBXB-33	Alexey, can you provide me a resolution for this feature request in the next few days? I will then be able to offload schema parsing from the ws tools layer to the binding layer.	0	0	https://issues.jboss.org/browse/JBXB-33
13498	JBXB-30	How can I do it?	1	0	https://issues.jboss.org/browse/JBXB-30
13499	JBXB-5	Adrian, could you please correct assertions in org.jboss.test.xml.PojoServerUnitTestCase when you have time? You can leave fixing the binding to me. Just let the tests fail. This is not a blocking issues for me right now. I started to work on XSD/runtime binding.	0	0	https://issues.jboss.org/browse/JBXB-5
13500	ORG-3634	[~velias] I understand Libor is sick. Can someone else make this change? It should be quick. [~adelaarreola] can help out and show you exactly what you need to add.	1	1	https://issues.jboss.org/browse/ORG-3634
13501	ORG-3608	I don't think 'mixed content' is the right term here. Essentially, we need to ensure that links to the developer blog use https. These links are coming from the sys_url in the DCP documents. How about you use the same protocol in the sys_url as is used in the feed URL used to configure the river? We then update the URLs for all the blogs that support https.	1	1	https://issues.jboss.org/browse/ORG-3608
13502	ORG-3580	hi Libor. I can't get to that link- it says temporarily unavailable. I see this jira as 'done' status. What's the next step to getting visibility into first response times?	1	0	https://issues.jboss.org/browse/ORG-3580
13503	ORG-3495	Could this be done this week (by August 24), since we are doing the marketing announcement for JDV 6.3 this week.	1	1	https://issues.jboss.org/browse/ORG-3495
13504	ORG-3476	Is there any update of this please? I would like to verify JBIDE-17615 and this blocks me. I tried it again now and the change password page still rejects my current password.	1	1	https://issues.jboss.org/browse/ORG-3476
13505	ORG-3441	Fixed Shall I apply to prod?	0	0	https://issues.jboss.org/browse/ORG-3441
13506	ORG-3314	* "Type" (atricle/solution) information is already available in the DCP, we have two separate types in DCP {{rht_knowledgebase_article}} and {{rht_knowledgebase_solution}} * Last update date is not available on the API we use to pull data into DCP. We have to work with CP team to get it there if possible Also we currently pull data for middleware (JBoss) products only. Should we change this to pull data for all products?	0	0	https://issues.jboss.org/browse/ORG-3314

#	Requirement	Questions	Answered	Successfully Answered	URL
13507	ORG-3314	"Solution", "uri": "https://api.access.stage.redhat.com/rs/solutions/106723", "path": "https://api.access.stage.redhat.com/rs/solutions/106723", "view_uri": "https://access.stage.redhat.com/solutions/106723", "title": "How to use KnowledgeAgent to handle hot deployment for BRMS ?", "allTitle": "How to use KnowledgeAgent to handle hot deployment for BRMS ?", "sortTitle": "How to use KnowledgeAgent to handle hot deployment for BRMS ?", "id": "106723", "solution.id": "106723", "authorSSOName": "rhn-support-alazarot", "lastModifiedBySSOName": "rhn-support-abhumbe", "hasPublishedRevision": "true", "lastModifiedDate": "2014-07-23T12:29:34Z", "createdDate": "2012-04-23T15:06:52Z", "kcsState": "verified", "ModerationState": "published", "accessState": "active", "setLanguage": "en", "abstract": "- How to use KnowledgeAgent to handle hot deployment for BRMS? That topic is explained at [6.1. The Knowledge Agent](https://access.redhat.com/documentation/en-US/JBoss_Enterprise_BRMS_Platform/5/html-single/BRMS_User_Guide/index.html#The_Knowledge_Agent) from [BRMS User Guide Documentation](https://access.redhat.com/documentation/en-US/JBoss_Enterprise_BRMS_Platform/5/html-single/BRMS_User_Guide/index.html). A simple snippet: ~~~ private static KnowledgeBase readKnowledgeBaseCS() throws Exception {final KnowledgeAgent kagent = KnowledgeAgentFactory.newKnowledgeAgent(\"kagent\"); ResourceFactory.getResourceChangeNotifierService().start(); ResourceFactory.getResourceChangeScannerService().start(); kagent.applyChangeSet(ResourceFactory.newClassPathResource(\"changeset.xml\")); KnowledgeBase kbase = kagent.getKnowledgeBase(); return kbase;} ~~~ changeset.xml may refer to PKG url provided by Guvnor but also files in file system / classpath. For example: ~~~ ~~~ ~~~ ~~~ How to use KnowledgeAgent to handle hot deployment for BRMS ?", "detectedLanguage": ["en"], "language": "en", "sbr": ["Business Rule Frameworks"], "tag": ["brms"], "product": ["Red Hat JBoss BRMS"], "publishedTitle": "How to use KnowledgeAgent to handle hot deployment for BRMS ?", "publishedAbstract": "- How to use KnowledgeAgent to handle hot deployment for BRMS?", "caseCount": 7, "kcsRating": 0, "category": ["Developer"], "publication_state": "verified", "requires_subscription": true, "boostProduct": "jboss", "_version_": 1534676601117081600, "timestamp": "2016-05-18T14:27:37.633Z" }, { "documentKind": "Solution", "uri": "https://api.access.stage.redhat.com/rs/solutions/1117493", "path": "https://api.access.stage.redhat.com/rs/solutions/1117493", "view_uri": "https://access.stage.redhat.com/solutions/1117493", "title": "Supported methods in jBPM interface org.jbpm.task.event.TaskEventListener", "allTitle": "Supported methods in jBPM interface org.jbpm.task.event.TaskEventListener", "sortTitle": "Supported methods in jBPM interface org.jbpm.task.event.TaskEventListener", "id": "1117493", "solution.id": "1117493", "authorSSOName": "rhn-support-lywang", "lastModifiedBySSOName": "rhn-support-enewby", "hasPublishedRevision": "true", "lastModifiedDate": "2014-07-03T16:08:43Z", "createdDate": "2014-07-02T06:30:45Z", "kcsState": "unverified", "ModerationState": "published", "accessState": "active", "setLanguage": "en", "abstract": "- I'm using the TaskEventListener to capture human task events, how can I capture the event when a task is assigned to a group or a set of users?- In jBPM community version 5.4 some more methods are supported such as `taskCreated()` and `taskStarted()` as per this [API doc](http://docs.jboss.org/jbpm/v5.4/javadocs/org/jbpm/task/event/TaskEventListener.html), can I have them in JBoss BRMS 5.3? According to [JBoss Enterprise BRMS Component Details](https://access.redhat.com/articles/119953), BRMS 5.3 is based on jBPM 5.2 so the TaskEventListener interface looks like below. Only four methods	0	0	https://issues.jboss.org/browse/ORG-3314
13508	ORG-3239	[~newtonm] [~hwindzio] [~burrsutter] Once we have the list, is there any reason why we couldn't go ahead and get these added this week? Is there a branding concern having RHEL, CDK, etc in the JBoss forums? Who has access to create new product forums?	1	1	https://issues.jboss.org/browse/ORG-3239
13509	ORG-3239	What's the strategy with the Customer Portal discussion forums? They have a developer category and various developer tags. It seems like they should be focused on production vs the developers.redhat.com focused on development discussions. Anyone speak with them yet? I don't know that's a blocker for adding additional products, but figured I'd ask as it may be confusing to developers (and we could also ask GSS to help out with the developers.redhat.com forums instead ;)) Also, are we adding *all* Red Hat products, or only those that are used within the context of development? I haven't seen the list.	1	1	https://issues.jboss.org/browse/ORG-3239
13510	ORG-3239	[~newtonm] [~hwindzio] I propose that we get rid of this space: https://developer.jboss.org/en/products/portal as it's not active and we don't even show that product on RHD. Any reason why we would not want to get rid of this? Is there another related product where we could move these threads?	1	1	https://issues.jboss.org/browse/ORG-3239
13511	ORG-3239	Something to discuss: How do we sunset a product space?	1	1	https://issues.jboss.org/browse/ORG-3239
13512	ORG-3187	[~lkrzyzanek] [~tnolan] - I'd suggest we keep it simple. http://design.jboss.org/redhatdeveloper/website/redhatdeveloper_2_0/pages/images/RHD_jive_main_r1v1.png 1) Change the background to #242424 2) Change the top padding to 5em 3) Add the RHD logo above the content 4) add a dotted pattern to the top right. Libor... is this workable in JIVE?	1	1	https://issues.jboss.org/browse/ORG-3187
13513	ORG-3182	[~lkrzyzanek] As part of this story, can we also add a message at the bottom of the email about the product forums being rebranded? Or would you like that work tracked in a different story?	1	1	https://issues.jboss.org/browse/ORG-3182
13514	ORG-3182	[~mikgue] [~lkrzyzanek] Here's a draft of the suggested changes. Mike - Can you take a look and let me know if you think it needs any other changes?	1	1	https://issues.jboss.org/browse/ORG-3182
13515	ORG-3182	What is "Them Test"? What is "Red Hat JBoss - RHD"? The dual branded forums? If so, will anyone know what RHD means as we use this mostly internally? Please change "Red Hat Developers" to "Red Hat Developer Program" Shouldn't "Following" instead be "Follow"?	1	1	https://issues.jboss.org/browse/ORG-3182
13516	ORG-2898	Review here: https://developer-dev.jboss.org/account/download-thank-you.jspa?filename=aaa.zip What reset password should be used ? The jboss.org one or RHD ?	1	1	https://issues.jboss.org/browse/ORG-2898
13517	ORG-2780	[~lvlcek] do you have a report from the Searchisko load testing you did a while ago? Are the tests reproducible?	1	0	https://issues.jboss.org/browse/ORG-2780
13518	ORG-2780	Yes, this is probably it. Do you have updated figures on what request load we are getting today on the DCP, and what multiplication factor you think we can achieve? Also, how easy is it to reproduce what you did? It would probably be wise to re-run as the data is 1.5yrs old. We've now got a new major version of the DCP and I believe new hardware.	0	0	https://issues.jboss.org/browse/ORG-2780

#	Requirement	Questions	Answered	Successfully Answered	URL
13519	ORG-2763	Hi, I analyzed content of weld webpages (http://weld.cdi-spec.org) and I see two main glitches there. *How to get page content* Weld website is missing some unique css class which allows us to get real content of the page, without common stuff as header, footer, main menu. We can take whole content but it will decrease search relevance for some words present in the common stuff. [~mkouba] should you look at https://github.com/weld/weld.github.io/issues/22 please? Attachment !weld-website-content-div.png|thumbnail! shows which {{div}} should be marked as it seems it is present on all pages of the website. *How to get page description* It is good to have some short description of the page stored in the search index. This description is typically shown on search result page. The description is typically taken from {{<meta name="description" content="short page description">}} tag stored in {{<head>}} section of the page, but it is missing on current weld website. I created https://github.com/weld/weld.github.io/issues/23 for this, [~mkouba] should you look at this please? If not added we should prepare short description as eg. first 150 characters of the page content, but we need glitch 1 to be resolved for this not to take common stuff into the description of each page ;-)	1	1	https://issues.jboss.org/browse/ORG-2763
13520	ORG-2679	[~lvlcek] is it you that's working on this? If so, can you assign it to yourself. Otherwise, if you know who's working on it, can you assign to them?	1	1	https://issues.jboss.org/browse/ORG-2679
13521	ORG-2679	[~unibrew] Who should review this? Can you chase them to get it done?	1	1	https://issues.jboss.org/browse/ORG-2679
13522	ORG-2644	[~newtonm] Can you set the release on this? Also, should this be a business requirement?	1	0	https://issues.jboss.org/browse/ORG-2644
13523	ORG-2644	[~unibrew] is it you that's working on this? If so, can you assign it to yourself. Otherwise, if you know who's working on it, can you assign to them?	0	0	https://issues.jboss.org/browse/ORG-2644
13524	ORG-2632	Burr, isn't https://repository.jboss.org/nexus/#nexus-search good enough for your use case?	1	0	https://issues.jboss.org/browse/ORG-2632
13525	ORG-2544	Hi, both accounts has been deleted. Should you use help@jboss.org to report this kind of problems next time please? We use this jira project for development related tasks, not common short term operational issues. Thanks.	1	1	https://issues.jboss.org/browse/ORG-2544
13526	ORG-2335	Can you explain what you mean by upgrading Zurp in Magnolia? I do not think it was ever used there. [~unibrew] might know more though.	1	1	https://issues.jboss.org/browse/ORG-2335
13527	ORG-2251	The approved product code is "unifiedpush" Can we go ahead and create the product forum? https://developer.jboss.org/<lang>/products/unifiedpush	1	1	https://issues.jboss.org/browse/ORG-2251
13528	ORG-2231	Please review: http://planet-jbossorgdev.rhcloud.com/ * Title link points to planet's blog post page * On compact view title link expands the content * When blog is already expanded then title link works as standard link * Next to Permanent Link is Original Blog link Let me know if it's OK. Thanks.	1	1	https://issues.jboss.org/browse/ORG-2231
13529	ORG-2216	[~pmuir] [~velias] based on analysis above it seems that it's possible to construct safely github and facebook links. Google and Linkedin is problematic. One option how to handle google and linkedin issue could be to pre-fill the field and let user to change it via jive UI if generated URL is wrong. Second option is to skip google and linkedin. WDYT?	1	1	https://issues.jboss.org/browse/ORG-2216
13530	ORG-2214	[~velias] can I also have the same job run against dcpbeta? Will those creds work?	1	1	https://issues.jboss.org/browse/ORG-2214
13531	ORG-2194	Note that you could reconfigure the headers for this rendering path to not need to check for a "modified since" operation, if those resource names increment monotonically at each change? What I mean is that assuming that when you update your avatar picture it changes from "23.png" to "24.png" you can set these to be cached for 10 years at client side. Worth checking how the avatar thing works, but I expect this to be the case as it's a common requirement nowadays to deliver speedy webpages ;-)	0	0	https://issues.jboss.org/browse/ORG-2194
13532	ORG-2180	Thanks for letting us now. Where could we expect that in production?	1	1	https://issues.jboss.org/browse/ORG-2180
13533	ORG-2173	For public social providers like twitter, linkedin, facebook and github it makes sense to set default privacy to everyone. However for google I would keep default value "yourself". WDYT?	1	1	https://issues.jboss.org/browse/ORG-2173
13534	ORG-2173	How is google plus different?	1	1	https://issues.jboss.org/browse/ORG-2173
13535	ORG-2161	Sensitive user data like emails and other ids (github id, jbossorg id, ...) are available at http://dcp.jboss.org/v1/rest/contributor/ but requires authorization (currently DCP admin only). May be we can expose this data to other authorized users once we implement more fine grained roles in Searchisko (cc: [~velias]). What is your specific use case? May be we can find a different solution and you do not have to wait for new version of DCP deployment.	1	0	https://issues.jboss.org/browse/ORG-2161
13536	ORG-2156	Implemented in sso plugin 3.1.3. See https://community-dev.jboss.org/edit-social-accounts!added.jspa [~james.cobb] [~cheyenneweaver] Could you please add Red Hat icon to the social accounts sprite https://static.jboss.org/images/jive/social-accounts-sprite.png ? and then update the https://static.jboss.org/css/jive/jbossdeveloper.css (class .social-account-logo-RedHatInternalProvider). You can force Jive to reload the css/img by increasing number in Jive's sytem property "jbossorg.themelocation" via https://community-dev.jboss.org/admin/system-properties.jsp . Thanks.	1	1	https://issues.jboss.org/browse/ORG-2156
13537	ORG-2156	[~lkrzyzanek] Where are we using this sprite? I'm seeing something different on the SSO page.	1	1	https://issues.jboss.org/browse/ORG-2156
13538	ORG-2141	Do I need access to dcpbeta-searchisko to update this? I know that it is being parse from a JSON, but I am not sure where to put in the information in order to have it added to the Buzz section.	0	0	https://issues.jboss.org/browse/ORG-2141
13539	ORG-2052	Hmm, maybe time to engage linkedin? I can start a thread with them.	1	1	https://issues.jboss.org/browse/ORG-2052
13540	ORG-2050	Hi, will you, please, let me know in what repository should the artifacts from Neo4j be accessible? I can create the proxy, but it would most likely not help you, because you would need to add the repository to your pom.xml.	1	1	https://issues.jboss.org/browse/ORG-2050
13541	ORG-2026	Hi [~gastaldi], it is not problem to add new provider into DCP. We have to clear out exactly data structure you plan to push into DCP to be able to configure provider appropriately. So few questions: * author name - which will be in this field? Full name of person (one or more)? Are you able to push email address, github username or jboss.org username there also, so DCP can correctly normalize against contributor profiles (as normalization based on full name is not so simple). * license - what is this? only textual info about license used? Do you want to be part of fulltext search data for addons or not? * repository url - what is this? Is it URL to some HTML page with addon description and download, or direct link to the some jar file or so? * description - is it plain text or contains some markdown or html formatting? I also expect you will write your own code to push data for addons into DCP. How do you plan to use these data? Some search on forge site?	1	1	https://issues.jboss.org/browse/ORG-2026
13542	ORG-2026	Vlastimil, About the author name, is there an existing model for us to reuse? Like, some sort of key that could be used in here that links to the github/jira/jboss.org username? Thanks	1	1	https://issues.jboss.org/browse/ORG-2026

#	Requirement	Questions	Answered	Successfully Answered	URL
13543	ORG-2026	Here are the attributes comments: Addon Info: - author name: I think we should store the github user name in here? Or perhaps a key? How would you suggest? - description: Markdown and/or HTML formatting should be available - license: Perhaps just the license acronym (EPL, ASL, LGPL, etc) would be interesting to have. Not sure if full text search would be interesting. - repository url: This is a Github URL Documentation Info: - documentation url: A URL pointing to an asciidoc document - content: Asciidoc content for indexing - title - tags	1	1	https://issues.jboss.org/browse/ORG-2026
13544	ORG-2026	I prepared initial configuration of DCP content provider. Expected input format for data is {code} { "sys_title" : "Addon title", "sys_url_view" : "http://github.com/addon/repository", "sys_description" : "Addon description in plain text (or markdown which is nearly plain text)", "sys_content" : "Addon documentation/long description text in Asciidoc format", "sys_license" : "ASL", "sys_url_doc" : "http://some.url/with/documentation", "tags" : ["tag1","tag2"], "author" : "velias", "sys_created" : "2011-04-14T11:32Z", "sys_activity_dates" : ["2011-04-14T11:32Z","2012-11-27T13:55Z","2012-11-27T14:20Z","2012-11-27T14:22Z"] } {code} Where fields contain: * sys_title - Addon name * sys_url_view - URL of Github addon repository * sys_description - Basic addon description in plain text (or markdown which is nearly plain text). We typically restrict this to 400 characters. * sys_content - Addon documentation/long description text in Asciidoc format * sys_license - license acronym (EPL, ASL, LGPL, etc) * sys_url_doc - an URL pointing to an asciidoc document * tags - an array of tags * author - github username of addon author * sys_created - optional - date of Addon creation in ISO format * sys_activity_dates - optional array of dates when some activity around Addon became (eg. created, updated, ...) in ISO format If you agree then I can install provider to DCP testing server and provide you other info about how to push data in and how to use them over search REST API	1	1	https://issues.jboss.org/browse/ORG-2026
13545	ORG-2022	What is the lower security for this ? why is it not the exact same as when I've used my username/password ? (I here assume that the remember me is shown where it will be noticed just as with username/password)	1	0	https://issues.jboss.org/browse/ORG-2022
13546	ORG-2022	Security will be lower if we automatically apply 'remember me' for social login without any explicit user control (without checkbox). On the other hand I never saw website which used Social login with explicit Remember Me control/checkbox by users. We can't use current "Remember me" checkbox as it is too far from social icons, so nobody will know to use it together. We should probably place something like new "Remember me after Social login" checkbox somewhere near/under Social login icons, but I'm really not sure users will understand what it is for as it is not commonly used on web today. [~james.cobb] what do you think about this solution? Or any other idea how to implement this feature?	0	0	https://issues.jboss.org/browse/ORG-2022
13547	ORG-1961	Could you please specify which pages should be modified? I cannot find anything at community.jboss.org. Thanks, Libor	1	1	https://issues.jboss.org/browse/ORG-1961
13548	ORG-1876	Andy, can you please provide more info about curated feeds requirements? Is there anything missing in existing FEED API? http://docs.jbossorg.apiary.io/#feedapi	0	0	https://issues.jboss.org/browse/ORG-1876
13549	ORG-1875	Andy, you mean you want to index iCal [1] file format content? Do you think you can provide a little bit more context about this? Does it mean you want to put an "event" info into DCP or do you want to use it for something else? [1] http://en.wikipedia.org/wiki/ICalendar	0	0	https://issues.jboss.org/browse/ORG-1875
13550	ORG-1818	Could you add me as an administrator?	0	0	https://issues.jboss.org/browse/ORG-1818
13551	ORG-1780	Is there a way I can create forums, or does the .org team need to do this?	1	1	https://issues.jboss.org/browse/ORG-1780
13552	ORG-1745	What would you suggest for expiry time of files stored in project directories? e.g. http://static.jboss.org/example/images/example-banner-1180px.png. For our static.jboss.org/theme/* content, we set 2 days expiry time. Does it seem reasonable for the whole domain ? Obviously the question is regards normal usage of URLs, the FFE is a separate story.	1	1	https://issues.jboss.org/browse/ORG-1745
13553	ORG-1742	I have updated my mirrors section to the following {code} <mirrors> <mirror> <id>jboss-public-repository-group</id> <mirrorOf>*,!jboss-deprecated,!jboss-maven2-brew,!jboss-eap-5</mirrorOf> <name>JBoss.org Public Repository Group</name> <url>https://repository.jboss.org/nexus/content/groups/public/</url> </mirror> </mirrors> {code} However, when I try a mvn clean install I end up with {code} $mvn clean install -pl testsuite/compat/ -DallTests [INFO] Scanning for projects... [INFO] [INFO] ---------- [INFO] Building WildFly Test Suite: Compatibility Tests 8.0.0.Alpha4-SNAPSHOT [INFO] ---------- Downloading: https://repository.jboss.org/nexus/content/groups/public/org/eclipse/persistence/eclipselink/2.4.0/eclipselink-2.4.0.pom [WARNING] The POM for org.eclipse.persistence:eclipselink:jar:2.4.0 is missing, no dependency information available Downloading: https://repository.jboss.org/nexus/content/groups/public/org/eclipse/persistence/eclipselink/2.4.0/eclipselink-2.4.0.jar [INFO] ---------- [INFO] BUILD FAILURE [INFO] ---------- [INFO] Total time: 10.611s [INFO] Finished at: Thu Jul 25 20:56:35 BST 2013 [INFO] Final Memory: 119M/497M [INFO] ---------- [ERROR] Failed to execute goal on project wildfly-testsuite-compat: Could not resolve dependencies for project org.wildfly:wildfly-testsuite-compat:jar:8.0.0.Alpha4-SNAPSHOT: Could not find artifact org.eclipse.persistence:eclipselink:jar:2.4.0 in jboss-public-repository-group (https://repository.jboss.org/nexus/content/groups/public/) -> [Help 1] [ERROR] [ERROR] To see the full stack trace of the errors, re-run Maven with the -e switch. [ERROR] Re-run Maven using the -X switch to enable full debug logging. [ERROR] [ERROR] For more information about the errors and possible solutions, please read the following articles: [ERROR] [Help 1] http://cwiki.apache.org/confluence/display/MAVEN/DependencyResolutionException {code} I see 2.4.2 under https://repository.jboss.org/nexus/content/groups/public/org/eclipse/persistence/eclipselink/ but no 2.4.0, which is also needed. https://github.com/wildfly/wildfly/blob/master/testsuite/compat/pom.xml also seems to need some other versions although I have not looked very closely.	0	0	https://issues.jboss.org/browse/ORG-1742
13554	ORG-1719	Libor, can this be made visible to the user to whom the profile belongs *and* to admins?	1	1	https://issues.jboss.org/browse/ORG-1719
13555	ORG-1710	Ok, thanks for trying. I wish it would have been a trivial change... Did the issue you got show the same stacks as those ones: https://bugs.eclipse.org/bugs/show_bug.cgi?id=411239 ? Maybe you can add a quick note to this bug report to explain the use-case and symptoms you got so it would help plugins author to fix it.	0	0	https://issues.jboss.org/browse/ORG-1710
13556	ORG-1708	Is there an estimated date where this will be ready/live? Thanks, Catherine	1	1	https://issues.jboss.org/browse/ORG-1708
13557	ORG-1708	Hello, Is this FORUM LIVE & KICKING. Cause when the link https://community.jboss.org/en/soap was tried to access it gave a Page Not Found. Kindly do update us the FORUM link for Red Hat JBoss SOA-P Thanks, Star	1	1	https://issues.jboss.org/browse/ORG-1708
13558	ORG-1661	Hi, I agree with you, the user mechanism currently used on our Nexus is not good and I hope I will be able to fix it this year. Can you, please, explain what is Hibernate CI.	1	1	https://issues.jboss.org/browse/ORG-1661

#	Requirement	Questions	Answered	Successfully Answered	URL
13559	ORG-1661	So then is there a chance of this happening? Or are we just going to have to expose the current one used by all projects outside the firewall?	1	1	https://issues.jboss.org/browse/ORG-1661
13560	ORG-1661	Hi, I do not know, if there is a firewall issue with uploading the artifacts to Nexus. I suggest we try it and learn. Please, give me the list of user names and e-mails, that should be added as Nexus users. The names must differ from jboss.org accounts. I suggest we use some postfix, i.e. ebersole_hibernateCI, or something like that. I will create the accounts. What repository will be the target of your deploys? Then you should be able to test it and if there is a firewall issue, we ask Eng-ops to fix it. I do not have access to network configurations, but I will create the user(s) for you.	1	1	https://issues.jboss.org/browse/ORG-1661
13561	ORG-1591	Or maybe to render outgoing links as plain text for unauthenticated users, or for search bots only? This should be better if spam is added due SEO reasons only.	1	1	https://issues.jboss.org/browse/ORG-1591
13562	ORG-1554	I was able to reproduce the problem with the even-odd color problem, however I can not find any page with a problem with the sorting by date. When I click on the date column, it is ordered properly. Can you, please, point me to it?	1	1	https://issues.jboss.org/browse/ORG-1554
13563	ORG-1540	I noticed that we don't have nofollow's in our outgoing links. That's the main motivation for the spammers - they get higher pagerank. Can we please add rel="nofollow"? ===> ORG-1591	0	0	https://issues.jboss.org/browse/ORG-1540
13564	ORG-1466	How's this: https://www.jboss.org/author/confirmation.html?	1	1	https://issues.jboss.org/browse/ORG-1466
13565	ORG-1446	This feature will be present in DCP (Distributed contribution platform = a new content delivery system we develop. For those with access more information can be found here: https://docspace.corp.redhat.com/community/bu/middleware/jbossorg/dev/dcp). Points from discussion with lvlcek and velias: - DCP will return JSON object with the data - DCP will return either whole list of the projects or information about a specific project based on the query. - at this moment we expect project identifier will be the field you can narrow the list of returned projects. Is there any other field (or property) you would like to filter by? I.e. "active" / "archived" projects. - Magnolia will supply all the data to DCP using the logic, that feeds the project matrix paragraph - it may take a while until DCP is fully implemented. In the meantime it would make sense if I publish a Rest API, that will give you the whole list. This API should be considered a temporary solution.	1	1	https://issues.jboss.org/browse/ORG-1446
13566	ORG-1446	Hi David, could you please prepare the mock JSON of jboss.org/projects for us, till the the DCP is ready?	0	0	https://issues.jboss.org/browse/ORG-1446
13567	ORG-1446	The link you have posted isn't valid any more. Is [README.md|https://github.com/jbossorg/dcp-api/blob/master/documentation/rest-api/README.md] what you mean.	0	0	https://issues.jboss.org/browse/ORG-1446
13568	ORG-1411	Hi Darran, as I know content of JIRA "Summary" field is shown in hover popup when you stay with mouse pointer above issue key for a while. See attached screenshot. Isn't this enough?	1	1	https://issues.jboss.org/browse/ORG-1411
13569	ORG-1408	David, " I would go to Google " -- so why would we want the overview list at all then, when we say the user to go to Google anyway? The main issue is that with our project names asdfhasoid and sudzf98wzr and asn348t6e8rgh no user can guess what they try to do. So we need to help him. -> Short description, always visible. Directly below the project name is space. And then as you say, a longer description via hover.	1	1	https://issues.jboss.org/browse/ORG-1408
13570	ORG-1405	Are we allowed to get the mbox file data? Is there any link?	0	0	https://issues.jboss.org/browse/ORG-1405
13571	ORG-1404	Are we allowed to get the mbox file data? Is there any link?	0	0	https://issues.jboss.org/browse/ORG-1404
13572	ORG-1402	Is there any idea where those two mail lists should fall to in terms of project? Does it make sense to cover it under "JBoss AS7" or does it make sense to create a new project filter value for them (i.e. does it make sense to create a new visual project filter category for them in search.jboss.org)?	1	1	https://issues.jboss.org/browse/ORG-1402
13573	ORG-1402	With cdi-dev we have the same situation as with forge-dev. The mail list contains a lot of JIRA event emails (as of writing there are around 1300 of emails like that). I am suggesting not indexing them for now and ignoring all emails that have author like: _some_name_(JIRA) <jira-events@lists.jboss.org> For example: Pete Muir (JIRA) <jira-events@lists.jboss.org> or Dan Allen (JIRA) <jira-events@lists.jboss.org> Such data is now searchable directly in JIRA and it contains high volume of noise. It will be better integrated into search.jboss.org once we decide to include JIRA data. WDYT? Whatever your decision will be it would be better if it is consistent with https://issues.jboss.org/browse/ORG-1400 (as stated above we have the same problem there).	0	0	https://issues.jboss.org/browse/ORG-1402
13574	ORG-1400	Hi, as of writing the forge-dev list contains around 830 emails from <jira-events@lists.jboss.org>. Typically the mbox 'From:' header value is prefixed with the original (author_name+" "+(JIRA)) like: Jason Porter (JIRA) <jira-events@lists.jboss.org> or Lincoln Baxter III (JIRA) <jira-events@lists.jboss.org> I would suggest ignoring such emails for now. I.e. not indexing them. These seem to be result of some mirroring from JIRA events. This data is now searchable directly in JIRA and it contains high volume of noise. These will be better integrated into search.jboss.org once we decide to include JIRA data. WDYT?	1	1	https://issues.jboss.org/browse/ORG-1400
13575	ORG-1398	Where will this discussion happen?	0	0	https://issues.jboss.org/browse/ORG-1398
13576	ORG-1390	What do you mean by all JBoss Community contributors? Do you mean everyone or a subset?	1	1	https://issues.jboss.org/browse/ORG-1390
13577	ORG-1368	Git supports rsync: http://schacon.github.com/git/git-fetch.html http://schacon.github.com/git/git-push.html I'm also curious why ssh isn't an option. If sftp and scp are open, wouldn't that mean ssh is available?	1	0	https://issues.jboss.org/browse/ORG-1368
13578	ORG-1368	Sadly, the rsync implementation inside git seems to use ssh in background and since our chrooted environment on filemgmt.jboss.org is very limited and has ssh turned off completely, we cannot use it. So, after various tries with services we have and protocols available, we've decided to change strategy with this task. We are planning to set up a very simple service which would receive GitHub's webhook JSON messages from your repos - https://help.github.com/articles/post-receive-hooks/. On receive of such message it would simply update the local checkout on filemgmt.jboss.org and therefore update the site. You won't have to login to filemgmt.jboss.org any more, setting up the webhook is just few clicks and only a one time issue and also setting this configuration up should be doable in fairly short manner. Does this sound good to you?	0	0	https://issues.jboss.org/browse/ORG-1368
13579	ORG-1314	Hi David, did you add styles to the page through a paragraph or through the editor by any chance? Those styles are conflicting with the styles I added to the project stylesheet and need to be removed. If not, there are some mystery breadcrumb styles that I need to hunt down :-) Also, I styled the breadcrumb to have a span inside the div. Is it possible to add that back in? Thanks!!	1	1	https://issues.jboss.org/browse/ORG-1314
13580	ORG-1310	I still don't understand. The gunvor-dev mailing appears to be being used to ask questions about Drools Guvnor, not Guvnor Core. What do you currently use for Guvnor Core?	1	1	https://issues.jboss.org/browse/ORG-1310

#	Requirement	Questions	Answered	Successfully Answered	URL
13581	ORG-1310	Until now projects have used their <project>-dev mailing list for issues related to subprojects too. For example the hibernate-dev list contains emails about Hibernate Search in addition to Hibernate Core: http://lists.jboss.org/pipermail/hibernate-dev/2012-January/thread.html Given that there is no traffic on the guvnor-dev list anyway would it be possible to use this also for Guvnor SOA emails to keep things consistent? If you really want a separate list then we can do it but it makes sense to keep things consistent where possible and it also prevents a proliferation of mailing lists as more and more subprojects are created.	1	1	https://issues.jboss.org/browse/ORG-1310
13582	ORG-1302	Could you please make me an administrator of the group? username: dan.j.allen	1	1	https://issues.jboss.org/browse/ORG-1302
13583	ORG-1300	Could you please make me an administrator of the group? username: dan.j.allen	1	1	https://issues.jboss.org/browse/ORG-1300
13584	ORG-1258	Hi Linconl, please propose a blog and I'll approve it. I see in propositions only "OCPsoft" feed. Is it that feed which you mean? Probably you'd like to have separate group "Forg" right? Thanks, Libor	1	1	https://issues.jboss.org/browse/ORG-1258
13585	ORG-1165	Can you, please, specify what exactly do you want to be displayed on that page? I am not sure whether you want to be able to add a paragraph with the twitter accounts of a single project or the list of all the projects in SBS and their twitter user names. If it is the list, we will need to add some properties to SBS groups and users, because there are more groups than projects (they are used for other purposes) and we also might need to specify who is "core developer" and who is not (unless we consider all the developer from SBS to be core developers).	1	1	https://issues.jboss.org/browse/ORG-1165
13586	ORG-1165	OK, so I suggest following solution: I will modify the Project team paragraph. At this moment it is able to display just a photo of the developer, e-mail and it marks a leader of the project. I will add check boxes for other possible contact informations, (such as Twitter, Facebook....). I will also ask Libor to add G+ to SBS, so we can add it there too. That way each project could set what they want to display. Note: If the specific user sets its privacy settings not to display the specific field, it will not be displayed in the paragraph, but it will be displayed for other people from the team, that have the information available to the public. Is this solution suitable for you?	1	1	https://issues.jboss.org/browse/ORG-1165
13587	ORG-1165	And is there a way to aggregate all the twitter ids? Or will I need to search through each person's profile? I would like to follow all JBoss'ers :-)	1	0	https://issues.jboss.org/browse/ORG-1165
13588	ORG-1165	I might have not been entirely clean, so I am adding two screenshots. The first one has the latest project team paragraph (twice: once with the user icon displayed and the second also without it). The second screenshot shows the settings dialogue. What I suggest is to add check boxes for the contact types (such as G+, Twitter etc.). If the group is checked, in the paragraph the links for checked contacts will be displayed too. In case of the displayed user icon (I will refer it as user picture to avoid an ambiquity), I think the best solution will be the clickable icons of the services under the name. In case of the user picture, that is switched off, the best way would be the service icons on the same line as the name of the developer (after the project leader's star, if it is present). Burr, will you, please, confirm, this solution will help you to achieve what you wanted to?	0	0	https://issues.jboss.org/browse/ORG-1165
13589	ORG-1156	Hey David, I've added a silver star for admins and put the gold star on the lead Can we just add text inside the span that either says "lead" or "admin"? Lead Admin (I've attached a screenshot of what I mean) Thanks, cheyenne	1	1	https://issues.jboss.org/browse/ORG-1156
13590	ORG-1127	The decision to only include the latest version of docs in Confluence was because we wanted to focus people's attention on editing the latest version in the quickest possible way whilst providing a bridge to DocBook XML. i.e. it wasn't a requirement to allow easy editing of prior doc versions. There were also concerns about performance given the large number of projects we have. However since we have a clustered setup that works well and maintaining multiple doc versions in Confluence appears to be the de-facto standard it looks like a good solution. If we can resolve the copy space issue and give some better hierarchical structure to the spaces then I don't see why we shouldn't do this. We'd still of course improve the DocBook XML export process as much as possible for releases. Would projects still have a need to produce the HTML and single-HTML formats of docs if we went ahead with this (given that Confluence is already viewed in a web browser), or only the PDF format?	1	1	https://issues.jboss.org/browse/ORG-1127
13591	ORG-1121	Are you interested in having same content in a couple of places or just a paragraph of same type to edit afterwards?	0	0	https://issues.jboss.org/browse/ORG-1121
13592	ORG-1097	One workaround for this issue is to use WebDAV to mount confluence data as files on your computer, and then use standard local tools to find and replace some texts in these files. Related documentation is here: http://confluence.atlassian.com/display/DOC/Configuring+a+WebDAV+client+for+Confluence#ConfiguringaWebDAVclientforConfluence-UsingaWebDAVClienttoWorkwithPages Root URL of Confluence WebDAV access is here: https://docs.jboss.org/author/plugins/servlet/confluence/default/Global/	0	0	https://issues.jboss.org/browse/ORG-1097
13593	ORG-1096	Hmm, that is a very unintuitive place to put it. Is it possible to have it where docbook puts it (foot of content area)?	0	0	https://issues.jboss.org/browse/ORG-1096
13594	ORG-1096	Would it be possible for you to share the code of the macro? This would be really great!	0	0	https://issues.jboss.org/browse/ORG-1096
13595	ORG-1085	Do you possibly have a list of whose blogs should we include in this syndicated filtered copy? I thought that as a source we could simply use a filtered main JBoss Community feed (http://planet.jboss.org/view/feed.seam?name=all&from=0). Basically, everyone should be aggregated in this feed anyway already.	1	1	https://issues.jboss.org/browse/ORG-1085
13596	ORG-1082	Is there already a URL to programmatically download the transformed DocBook XML from Confluence? If so, I could very easily write a simple script to use that until this Maven plugin was available.	1	1	https://issues.jboss.org/browse/ORG-1082
13597	ORG-1075	We need to talk about how you propose to integrate with SBS, i.e. what purpose does the integration serve. Given our recent emails it appears you want to implement a different way to create project pages which isn't a valid idea. All projects should be using Magnolia.	0	0	https://issues.jboss.org/browse/ORG-1075
13598	ORG-1075	Can we please enable the REST API for SBS? There are tons of use cases where we could use this data to help streamline processes in our projects. One idea is to get builds to generate blog entries. Another use case is to be able to get lists of JBoss user profiles to link them to identities in other places. The thing about a REST API is that you just never know what great mashups can occur, both inside the project teams and outside, that lead to innovation and great participation.	0	0	https://issues.jboss.org/browse/ORG-1075
13599	ORG-1071	How to add button to site (not only on search result): http://www.google.com/webmasters/+1/button/index.html#utm_source=webmastercentral&utm_medium=blog&utm_campaign=announcement	0	0	https://issues.jboss.org/browse/ORG-1071

#	Requirement	Questions	Answered	Successfully Answered	URL
13600	ORG-1060	Some questions: # Do I always have to approve content that I wrote? # Is there any way to only display the latest approved content for a given site? (Except for the author of the unapproved content who should probably see their content with the "unapproved" header) Just noting it as "unapproved" doesn't really prevent misinformation from being added. # If we put this in, can we automatically approve everything that was previously written so we don't have to click a zillion approve links?	1	1	https://issues.jboss.org/browse/ORG-1060
13601	ORG-1060	I installed new version to the staging server with some improvements. ad 1 - it's not necessary to approve content written by user with approve permissions ad 3 - I added "Approve all" link to the screen showing all unapproved pages for space ad 2 - what is shown is selected by Confluence internals and I have no chance to change it. I highlighted "unapproved" warning message more. I think this message warns readers that content is not approved so may contain misinformations, and they can to show latest approved content. It's probably best what we can do. We can change wording of "unapproved" warning in some way if you think.	1	1	https://issues.jboss.org/browse/ORG-1060
13602	ORG-1060	Okay, well that's better than nothing. Can we get a special notification for unapproved changes? Or do I just have to watch the space manually?	1	1	https://issues.jboss.org/browse/ORG-1060
13603	ORG-1052	The purpose of the documentation is to describe the final releases so why would you need to include mockups? If you're trying to describe a screen layout then a screenshot would be better. By standardising on using the Gliffy Diagram plugin consistently across the docs it makes it much easier to gain economies of scale and present consistent, high-quality docs to users.	1	1	https://issues.jboss.org/browse/ORG-1052
13604	ORG-1038	Should we be able to embed a metadata reference in a link? When I use a wiki text such as {noformat} [foo {space-metadata-from:bar}|http://...]{noformat} the next time I edit the page the link will be turned into plain text - e.g. {noformat}\[foo {space-metadata-from:bar}|http://...\]{noformat}	0	0	https://issues.jboss.org/browse/ORG-1038
13605	ORG-1005	Hi Laura, should you look this please because related to Authoring Guide? But we are constrained by Confluence and Scroll export plugin features here.	0	0	https://issues.jboss.org/browse/ORG-1005
13606	ORG-1005	I'm not sure we want to add this functionality. I'd say rather than relying on formatting for this, the better way is to make sure the element that you're referring to is clearly distinguished by your language. Would a decent compromise be to add line numbers to the 'code' blocks, or request that the ability to call out specific lines be added to a) Confluence and b) the Scroll plugin? We can use callouts in <programlisting> to identify particular code blocks (shown in the Seam Reference Guide, off the top of my head) but I don't believe Confluence or Scroll can cope with this.	0	0	https://issues.jboss.org/browse/ORG-1005
13607	ORG-980	I expect that we want to have option to add new sponsor to the default list as well, right? Should the user be able to edit each description of sponsor, even those which are available by default ? I guess, they will often prefer just to have logo without the text.	0	0	https://issues.jboss.org/browse/ORG-980
13608	ORG-980	ok, so, what are the defaults? Could some native English person fill them in on http://10.16.93.207:8080/author/rysieksParagraphsTestPage.html ?	0	0	https://issues.jboss.org/browse/ORG-980
13609	ORG-980	Yeah, I was thinking all the text descriptions would be editable. In fact, I just made a lot of the text up and your right that some people will want a logo-only option, so we might want to leave the text off by default(?) So perhaps we could add an option to the dialog that asks if you want to add text, and the class could change depending on whether they do or not: class="sponsors" for no-text class="sponsors-text" for logos+text Does that make sense? Thanks!	1	1	https://issues.jboss.org/browse/ORG-980
13610	ORG-980	Or we could add default text... perhaps we could just paraphrase Drools, Portal, and Microcontainer: JProfiler - EJ-Technologies has provided the ___ project JProfiler licenses for free. NexB - nexB is the preferred Eclipse development partner of the ___ project. IntelliJ - JetBrains has donated an IntelliJ license for use on the open source ___ project. Structure101 - Structure 101 is a valued sponsor of the ___ project. MagicDraw - A Special thanks to Magicdraw for supporting ___ project. (could we auto-fill the name of the project in the blank, or should I try and write it w/o mentioning the project?)	1	1	https://issues.jboss.org/browse/ORG-980
13611	ORG-980	I finished coding, quite a bit of work. Anyway, it's available here http://10.16.93.207:8080/author/sponsorsTest.html for testing. Please check if everything works as you expected. Also for default banners I used links to main pages of the sponsors, that's ok? Or should I use links to the product itself?	0	0	https://issues.jboss.org/browse/ORG-980
13612	ORG-978	If it should always appear on the project page template then, I'll just add it as a part of the template not as a paragraph. The question is: where should it appear?	1	1	https://issues.jboss.org/browse/ORG-978
13613	ORG-965	is the import/parse code possible to run outside of Confluence some how to test the import? I'm getting the not so useful error: Error during DocBook file import: Error attempting to parse XML file (href='extension_reference.xml').	1	1	https://issues.jboss.org/browse/ORG-965
13614	ORG-920	Does this GPL have a version number appended as well, like ASL v2 or LGPL v2.1 ?	1	1	https://issues.jboss.org/browse/ORG-920
13615	ORG-920	Okay, I've created the badge for GPL which is now in staging: http://10.16.93.207:8080/author/cookbook/demo.html? (also see screenshot) Rysiek, we'll need to add it to the paragraph dialog with the option for adding the version number in text like the jsr's. <li class="license-gpl">General Public License version 2.0 (GPL v2.0)2.0 If we're adding the option to change the version number, what should we link to? http://www.gnu.org/licenses/old-licenses/gpl-2.0.html ? Thanks! ~c	1	1	https://issues.jboss.org/browse/ORG-920
13616	ORG-920	Do you think other projects might need a different version number in the near future? Cheers, Cheyenne	1	1	https://issues.jboss.org/browse/ORG-920
13617	ORG-869	The reason AFAIK is because of using P tag for each single line and not using BR tag for break line. More info: http://wiki.moxiecode.com/index.php/TinyMCE:FAQ#TinyMCE_produce_P_elements_on_enter.2Freturn_instead_of_BR_elements.3F	1	1	https://issues.jboss.org/browse/ORG-869
13618	ORG-869	It sounds good to me. I've modified the css in staging. Libor, can we strip the p tags?	0	0	https://issues.jboss.org/browse/ORG-869
13619	ORG-869	I've tweaked the CSS in staging. Shall we strip the p tags?	0	0	https://issues.jboss.org/browse/ORG-869
13620	ORG-869	I have few concerns: 1. RTE editor still display paragraph with 0 margins - it means that it's not WYSIWYG because then paragraphs are rendered with margin and can be confusing for users. 2. Existing documents usually have empty line between paragraphs. Now margin is bigger and rendered document looks differently. For example compare: http://goomba.qa.atl2.redhat.com/wiki/RiftSawInstructionsforUploadingtodocsjbossorganddownloadsjbossorg and http://community.jboss.org/wiki/RiftSawInstructionsforUploadingtodocsjbossorganddownloadsjbossorg 3. Don't know why but doesn't work in Chrome. Firefox is OK	1	1	https://issues.jboss.org/browse/ORG-869

#	Requirement	Questions	Answered	Successfully Answered	URL
13621	ORG-846	We currently have these styles used by various projects: https://www.jboss.org/netty/testimonials Is this something needed for Magnolia or SBS?	0	0	https://issues.jboss.org/browse/ORG-846
13622	ORG-846	How's that formatting supposed to work? The PDF would have to be rewritten on our side with your template? Because I don't think we could require users to use it while writing the success story? It's additional work for them.	1	1	https://issues.jboss.org/browse/ORG-846
13623	ORG-834	From an email thread on the core. I had said: "I actually used WordPress, not Movable Type. So I have both Blogspot and WordPress formatted dumps, but neither can be imported into SBS. As for Movable Type, do you know of any freely available, hosted MT service I could use? The last thing I'm going to do is download and install Movable Type (inc a RDBMS, etc) just to convert a blog. :-)"	1	0	https://issues.jboss.org/browse/ORG-834
13624	ORG-834	If I send you the dump that I have generated, could you try with this?	1	0	https://issues.jboss.org/browse/ORG-834
13625	ORG-830	who should i ask if i want hibernate repo be in the fisheye?	1	1	https://issues.jboss.org/browse/ORG-830
13626	ORG-811	Mark, Is there somewhere I can drill down into downloads.jboss.org to style an example? Not Found for me here: http://downloads.jboss.org/docbook/ Thanks!	1	1	https://issues.jboss.org/browse/ORG-811
13627	ORG-811	Thanks Mark. So assuming that the styles suffice, who would I ping to get this implemented?	0	0	https://issues.jboss.org/browse/ORG-811
13628	ORG-811	Here's a taste of what we do for JBoss Tools: Unstyled Apache default (dir browsing enabled): http://download.jboss.org/jbosstools/builds/staging/ http://download.jboss.org/jbosstools/updates/nightly/ Custom-assembled HTML pages (to help users know what's what): http://download.jboss.org/jbosstools/updates/ http://download.jboss.org/jbosstools/builds/ http://download.jboss.org/jbosstools/updates/helios/ I'm good with the Apache default changing to something prettier, but I'd still like the ability to override (or supplement?) that with our own custom content to explain what is in a directory, and where people should be clicking to find content. Is that possible? Eg., if I drop a README.txt in a folder, you could pick that up and embed it into your templated HTML? Or if there's an index.html it would replace your generated stuff?	1	1	https://issues.jboss.org/browse/ORG-811
13629	ORG-769	I see SteamCannon in the Projects menu on jboss.org but I don't see it in the projects menu on planet.jboss.org. Is that a separate system, and do I need to open a separate issue to get it to show up there?	1	1	https://issues.jboss.org/browse/ORG-769
13630	ORG-745	Added to the css which can be found here: http://10.16.93.207:8080/exampleproj-rightcol.html Rysiek can we add this to the paragraph? Thanks!	1	1	https://issues.jboss.org/browse/ORG-745
13631	ORG-704	Where are we with these? Do have some place we can dump logs inside the Red Hat Infrastructure? 3) If we have it hosted within JBoss / Red Hat, we should be able to pull this all in and search it easily (I would think). There's still value in being able to go back and read logs IMO, if there wasn't a meeting, but some good discussion happened, it's nice to see what actually happened. 4) May be nice, but wouldn't the freenode web interface work for this?	0	0	https://issues.jboss.org/browse/ORG-704
13632	ORG-692	Hi Mark Yep, I can confirm PressGang is sitting on the projects matrix page. Thanks for that. I think when we were first discussing PressGang I seem to remember we also talked about making a Documentation Support category (or something similar) to lump things like PressGang, jbossorg-jdocbook-style, maven-jdocbook-plugin, and the other related items. Just to get them out of the "Other" category. Would that be possible? For Point 2 in this ticket Description, I can't see a listing in the https://svn.jboss.org/repos/ listing for PressGang. I also note it is really out of alphabetical order (as an aside).	0	0	https://issues.jboss.org/browse/ORG-692
13633	ORG-669	Can not someone quickly add the anchor?	1	1	https://issues.jboss.org/browse/ORG-669
13634	ORG-639	Hi Marek, I think it will be better to just move BoxGrinder node from stormgrind/projects/boxgrinder to the root. It should save you some time in moving the content but still some links or header image etc. might break and may need correction. Do you agree? or you prefer a clean project site? Cheers	1	1	https://issues.jboss.org/browse/ORG-639
13635	ORG-639	Is there a chance to copy that instead of move? I would prefer to polish (sic!) the new project before release and then remove unnecessary files from StormGrind. But yah, it'll save a lot of time!	1	1	https://issues.jboss.org/browse/ORG-639
13636	ORG-639	In case of SBS, I think we should do the same to have a clear mapping between Magnolia and SBS. I'd like to block the idea of (creating|renaming StormGrind to) Cloud for now. Btw. what's the point of moving BoxGrinder out of StormGrind if still you find it appropriate to place it inside same subspace if it's named Cloud ?	1	1	https://issues.jboss.org/browse/ORG-639
13637	ORG-620	Can we see what it's going to look like? Can we participate by giving feedback on the design and layout?	1	1	https://issues.jboss.org/browse/ORG-620
13638	ORG-620	First, thank you for sharing the screenshot and answering my questions. I have some feedback regarding the screenshot. * [statement] I agree that the blog posts now appear on the left side when visiting the root URL of planet. We want to present users with the news first, then a list of sources as a second click. This design accomplishes that goal. * [feature request] The posts appear to be truncated to only a short summary. All other planet feeds I know of show the whole post. Can we show the whole post? If we are concerned it's too much scrolling, could we use an expandable block using jQuery as an alternative approach? That should expand w/o reloading the page...it just reveals the hidden content. * [request for clarification] The posts are still missing avatars (profile pictures) for each author. Do you plan on hooking those up with the identities on jboss.org so that the picture is shown? I understand if that doesn't show up in staging, but I would like to know ahead of time if pictures need to be fixed for the most regular authors before we go live. * [feature request] Formatting code has been a problem in the past. Code is a very central part of many of the posts that we write. How will code be handled? Will it show up abnormally, or is there some way we can get a consistent rendering?	1	1	https://issues.jboss.org/browse/ORG-620
13639	ORG-613	I sent a communication to the team, that created the tool. There are two important questions: is it legal, to replicate their functionality? And second question is how to generate the parameters, so we can use it some clever way.	0	0	https://issues.jboss.org/browse/ORG-613
13640	ORG-606	is this applied to the JBoss Website running currently ?	1	1	https://issues.jboss.org/browse/ORG-606
13641	ORG-601	Okay, I added and to projectContent.jsp but migration.jboss.org isn't showing the update. Can you push this change so I can check to see it's working before we schedule a merge? Also, is there another template I need to add these to for the full-width and right column layouts, or will this template take care of all 4 layouts? Thanks, C	0	0	https://issues.jboss.org/browse/ORG-601
13642	ORG-584	I think it is better, if it is not extra paragraph, because it applies to all pages without modifications. However I might following: - hide it on all home pages, so it is not necessary to do this by CSS - add a checkbox into property settings of our community page templates. They would be on by default and if unchecked, the breadcrumb would disappear. What do you think?	1	1	https://issues.jboss.org/browse/ORG-584

#	Requirement	Questions	Answered	Successfully Answered	URL
13643	ORG-577	Regarding to point 2. When you receive another type of feed (e.g. ATOM) then you receive also missing details. Just add "?feedType=ATOM" at the end of URL. Example: http://community.jboss.org/blogs/feeds/posts?feedType=ATOM	0	0	https://issues.jboss.org/browse/ORG-577
13644	ORG-575	So this would cover all SBS pages, not just blogs? cool. How can I have access to this? My google ID is manik.surtani@gmail.com	1	1	https://issues.jboss.org/browse/ORG-575
13645	ORG-575	Sounds perfect. So how do I do this? :)	1	1	https://issues.jboss.org/browse/ORG-575
13646	ORG-575	How would a project such as Infinispan configure the ${gacode} tracker ID variable? Via some admin interface on SBS?	1	1	https://issues.jboss.org/browse/ORG-575
13647	ORG-574	I checked feedburner features and I didn't find any big advantage against Google Analytics (since we have separated stats for each project). Manik, Is there something special that we could use?	0	0	https://issues.jboss.org/browse/ORG-574
13648	ORG-573	Can this not be extended to handle any Google Gadget? See the Google Gadget API: http://code.google.com/apis/gadgets/ so that this will cover not just the re-tweeting gadget but anything else written for Blogger/iGoogle/etc.	1	1	https://issues.jboss.org/browse/ORG-573
13649	ORG-562	Hey David, so what I did was to switch the position of the default tile to be at the very top of the sprite, so that whenever the class that calls a specific project that doesn't have an icon yet (because the project has not been branded yet), it will always default to this very top tile of the sprite which will show the default bullet. This way you don't have to mess with trying to find out if the class already exists or not. We don't have to add or change any html either. You can see a test I set up to show this working here: http://10.16.93.207:8080/author/projects1/matrix.html? I added a class in the css to correctly call the Application Server 7 icon for the corresponding node; .jbossas_icon_16x I also added a version of a hypothetical new project at the bottom there called "SomeNewProject" where you can see the default bullet showing up despite not being called by the class .blank_icon_16x Look good? Cheers, cheyenne	1	1	https://issues.jboss.org/browse/ORG-562
13650	ORG-537	The permanent link to your article (http://community.jboss.org/docs/DOC-15107) gets replaced in the address bar with the Human Friendly URL: http://community.jboss.org/wiki/HowIsArquillianPronounced This capitalizes the first letters of all words which you say is much more readable and fits in the spirit of wiki links. So what exactly is the issue?	0	0	https://issues.jboss.org/browse/ORG-537
13651	ORG-537	Please can you start again and explain what it is you want, I'm very confused :)	0	0	https://issues.jboss.org/browse/ORG-537
13652	ORG-537	Sure thing. I believe that users expect words in a friendly URL to be capitalized. This makes it easier on the eyes to read. If I create a document with the name "A test document", the resulting URL should be: https://community.jboss.org/wiki/ATestDocument This is consistent with how seamframework.org (and many other sites) create friendly URLs. Examples of what looks normal: http://seamframework.org/Community/15GotchasInSeam http://seamframework.org/Community/PostingAboutSeam3 The URL should be the same regardless of what capitalization I use in the document title. For instance, if I create a document with the name "a Test document", the resulting URL should still be: https://community.jboss.org/wiki/ATestDocument This works by capitalizing all words, removing non-alphanum characters (spaces, etc), then appending it to https://community.jboss.org/wiki/ Make sense?	0	0	https://issues.jboss.org/browse/ORG-537
13653	ORG-524	>> Manik, if you seriously haven't been 'able' to use the dropdown then I think we have bigger problems ;) Well, lets put it this way: with 50 or so *unsorted* entries, locating the one I need would require reading through almost each of them till I find the one I want (Infinispan ends up somewhere in the last 2/3rds of the list). It's quicker to search on *all spaces* and then look through the unnecessarily excessive search results for the one I want. Hence I never use the drop down. :) >> I know little issues like this can be frustrating but please try to report them accurately as it'll help us to prioritise things accordingly. In this case the dropdown is usable but not optimized for rapid searching. If you remember the location of your project in the list then you can find it in a few seconds. What's inaccurate about the report? I detail how I get to the menu in question, I explain which drop down in which screen is an issue, and I explain that the issue is that this list is not sorted. And I include a screen shot. Do you want a patch as well? ;) >> As Libor says I do think we can order this alphabetically though so we can improve it. This is *precisely* what I have asked for (see title of this JIRA).	1	1	https://issues.jboss.org/browse/ORG-524
13654	ORG-524	The fact that you say the dropdown is unusable is inaccurate as you can use it. It's just not as easy to use as it could be. You're using the word 'unusable' in a subjective way rather than an objective way which is why we're having this conversation. For example there is no way to differentiate this issue from one where the dropdown 'actually' became unusable because it didn't render. In its current state locating the entry you want simply means remembering where it's ordered in the list of projects. Once you've done this once then it takes about 3 seconds to locate again every time. Can you visually filter the search results to exclude the non-Infinispan related items in less than 3 seconds?	0	0	https://issues.jboss.org/browse/ORG-524
13655	ORG-524	You've woken up the pedant in me. :) >> The fact that you say the dropdown is unusable is inaccurate as you can use it. It's just not as easy to use as it could be. From http://en.wikipedia.org/wiki/Usability : "... usability is the study of the ease with which people can employ a particular tool or other human-made object in order to achieve a particular goal ..." >> In its current state locating the entry you want simply means remembering where it's ordered in the list of projects. Once you've done this once then it takes about 3 seconds to locate again every time. Can you visually filter the search results to exclude the non-Infinispan related items in less than 3 seconds? No. On my browser the viewport only displays a fraction of the complete list. And again my browser doesn't render scrollbars to indicate how much further there is to go to the end of the list (see screenshot), so I end up having to scroll one entry at a time. And keyboard shortcuts (pressing "I") does not scroll to the project in question (for whatever reason), and not even anywhere close. Maybe this is an OS/Browser combination thing (see the "Environment" field that I have filled out).	1	1	https://issues.jboss.org/browse/ORG-524
13656	ORG-524	Hi Manik, in Firefox it works perfectly. See the screenshot. Combobox is visible with scroolbar. Also typing letters works as expected. I tried type "inf" and it jumped to Infinispan. I also tried that behavior on final version of Google Chrome on Fedora12 and it works also perfectly - same like in Firefox. What version of Chrome do you use? Thanks.	1	1	https://issues.jboss.org/browse/ORG-524
13657	ORG-524	Have they been sorted alphabetically? It still seems the same (just checked)...	1	1	https://issues.jboss.org/browse/ORG-524
13658	ORG-517	Hi Libor, http://goomba.qa.atl2.redhat.com/en/hibernate I've updated the css to pull the search bar up into the proj_nav and swap out the search button. I've placed the dogear in jive-userbar, but it's not tall enough to show the entire dogear, so I'm going to need to add two nested divs inside jive-userbar to allow it to sit above the content, like so: (see screenshot also) Is this possible?	1	1	https://issues.jboss.org/browse/ORG-517
13659	ORG-494	It might work without too much complexity if you can just redirect from one page to another. So you would change the title of the page, but create a new page with the old title which redirects to the new title. The problem with the current system is that people using external systems are linking to the URL of the page title because this is the URL displayed. This URL breaks every time the title of the page changes. The unique DOC-XXXX links only work within SBS, but links coming in from external systems are easily broken.	1	1	https://issues.jboss.org/browse/ORG-494

#	Requirement	Questions	Answered	Successfully Answered	URL
13660	ORG-493	Could you share the numbers or show me where I can find them? (I believe you, just curious.) Also, I still believe a fluid layout would be an appropriate solution. It would allow users at any resolution to view the documentation at their preferred scale. A third option, which I did not present in the initial issue, would be to enable a user-setting to choose resolution. But, in reality, a fluid layout solves most of the issues.	1	1	https://issues.jboss.org/browse/ORG-493
13661	ORG-493	Is there a way to drill down and see what those 11.48% of users who's resolution is 1024x768 are doing on the site? I'd be surprised if they are active members of the community.	1	1	https://issues.jboss.org/browse/ORG-493
13662	ORG-485	Can we make the dialog change depending on the permissions of the user? Editable by the sysadmin, uneditable for the users.	1	1	https://issues.jboss.org/browse/ORG-485
13663	ORG-485	Mark: Sure thing, I'll hold off. Who's making the final approval?	1	1	https://issues.jboss.org/browse/ORG-485
13664	ORG-485	Following my call with the TDMs we agree that it doesn't make sense to place this widget on all project pages. Given the wide range of opinions people have on this matter I think we would be better off by leaving the title look&feel as it is so that we can place it anywhere on the page. Is it possible to modify the CSS for the dogear graphic so that it literally' overhangs over any other image placed there, or would that look weird?	1	1	https://issues.jboss.org/browse/ORG-485
13665	ORG-485	Rysiek, what I've done is add the html into a text and image paragraph at the top of the right column. Can we hard code that into the template? http://10.16.93.207:8080/supportedin-test.html?	0	0	https://issues.jboss.org/browse/ORG-485
13666	ORG-469	[~lkrzyzanek] is there any documentation anywhere describing how to enable and/or make use of this feature? Does the solution implemented involve mailman as has been suggested in this comment thread?	1	1	https://issues.jboss.org/browse/ORG-469
13667	ORG-421	Looks like the modifications didn't take. We need to re-position custom.css to overwrite org_common.css Libor, can we make this change in staging? Thanks! -cheyenne	1	1	https://issues.jboss.org/browse/ORG-421
13668	ORG-421	Hi David, I've fixed how CS is handling the sizes and spacing of each h tag which will be visible in our next release (until then see the screenshot I included). If this fix does not end up able to satisfy your needs, would you please move the request for a bulleting system, or FAQ template into another jira? Thanks! -cheyenne	1	1	https://issues.jboss.org/browse/ORG-421
13669	ORG-407	Can we assign this issue to some new release? It was reopened so it was not handled within the 4.1.0 release and probably will be in some future one. I just don't think we should have unresolved tasks in Changelog.	0	0	https://issues.jboss.org/browse/ORG-407
13670	ORG-400	Similar issue from Vicky Kak: I have started getting notifications from various sub groups with subject line starting with [DataSource*, [Installation*, [Security*, [Embedded JBoss Development], [Beginner's Corner], [Performance Tuning] etc.. I never used to get mail notification from these forums before, is there a way to disable email notification from these forums ? With older nukes forums I was subscribed to AS,JCA,TX,JMX,AOP sub groups and things were working fine.	0	0	https://issues.jboss.org/browse/ORG-400
13671	ORG-320	The main issue here is this. What the heck is that symbol all about? It's so cryptic. If it gave any indication of what happens when you click it, then this wouldn't be such an issue. Something like curly brackets: {...}	0	0	https://issues.jboss.org/browse/ORG-320
13672	ORG-297	Take your meds if "JBoss Users" and "JBoss Portal Users" share two first words ;-) Anyway, I know what you mean so I guess we will think about this longer. Still, what other condition can you do in GMail, there are very few options? What could we add to improve this?	0	0	https://issues.jboss.org/browse/ORG-297
13673	ORG-297	ok ok... I was about to write something else. But with this naming schemes it is not so unlikely for the whole prefix be a first part of other one ;) You made me experiment a bit with filters once again. So it looks like emails from lists.jboss.org are now recognized properly to have mailing list id in header and it is easy create proper filter with one click. I'm pretty sure gmail only worked in this way with google groups so I don't know if this is an improvement on gmail or on lists.jboss.org side. So it is nice for things like commit/jira notifications that go thru lists.jboss.org atm. Anyway... I looked a bit into email headers and the magic seems to be: List-Id: JBoss JIRA notifications <jboss-jira.lists.jboss.org> and then "Has the words" filter like "list: "<jboss-jira.lists.jboss.org>"" works perfectly... On the other possibilities - maybe encode more info into sender field? like "John Kowalski <xxx-yyy@communities.jboss.org>"? Other workaround is to setup mailing list channel for each forum :)	0	0	https://issues.jboss.org/browse/ORG-297
13674	ORG-297	Can't a header be added like for example what we have from the mailing lists? e.g. List-Id: Developer Forums discussion archive <jboss-dev-forums.lists.jboss.org>	0	0	https://issues.jboss.org/browse/ORG-297
13675	ORG-297	Sorry if I'm trying to be to smart as my SBS knowledge is minimal... but could this be helpful?: https://community.jivesoftware.com/message/263499#263499	1	1	https://issues.jboss.org/browse/ORG-297
13676	ORG-284	What's our current view on the ohloh link in project nav?	1	1	https://issues.jboss.org/browse/ORG-284
13677	ORG-260	Libor, we also have a parser (based on an antlr grammar) which you could take a look at in Seam - http://anonsvn.jboss.org/repos/seam/branches/community/Seam_2_2/seam-text.g. It will convert from a wiki-esque syntax to html. We also have a reasonable validator for it which is able to identify which character caused the error. Might be useful?	0	0	https://issues.jboss.org/browse/ORG-260
13678	ORG-260	Hi Ondrej, could you please provide more details? Which page did you edit? Did you try Undo button in your browser (Main menu -> Edit -> Undo) ? Thanks.	0	0	https://issues.jboss.org/browse/ORG-260
13679	ORG-247	How about an optional IRC nickname field as well?	1	1	https://issues.jboss.org/browse/ORG-247
13680	ORG-234	Since we have our own plugin for SBS&JIRA integration feel free to create new Enhancement ticket and describe what kind of functionality would you like to have. Thank you. Closing this ticket.	0	0	https://issues.jboss.org/browse/ORG-234
13681	ORG-227	Perfect (almost)! I have enabled this on one of my FAQ pages: http://www.jboss.org/community/wiki/InfinispanTechnicalFAQs But I don't want my Q's to be H1 (too big), so I use H5. But this means the questions are heavily indented in the TOC. Is there any way to tweak this to remove the indenting? Perhaps what we'd need is a separate style for TOC? Also, How can I have the TOC displayed as a numbered list rather than a bulleted one? Thanks Manik	1	1	https://issues.jboss.org/browse/ORG-227

#	Requirement	Questions	Answered	Successfully Answered	URL
13682	ORG-208	After some considerations, there are three possible solutions coming out: 1. One central repository In this case, admin will set the URL of the downloads repository centrally in the module configuration. Common users will use a control with behavior like current "Internal link", i.e. after clicking the button a new window with the repository tree will appear, where user will select the file to download. PROS: easier to develop, centralized = easier to maintain and manage possible problems CONS: only one repository for all projects 2. Separate repository for each paragraph instance There will be two tabs in the paragraph's dialog, in the first tab user will define the URL of the repository, in second tab will add items. The "Add item" control will work the same way as in #1, only the repository tree will be from the repository defined by user. PROS: more flexible (for "independent" projects like Hibernate) CONS: much more complicated development, possible production problems with user-defined repositories 3. Mix of #1 and #2 In the first tab of the paragraph's dialog, user will chose whether to use the repository predefined by admin in module's config, or to use different repository (in which case user has to define its URL). Second tab works the same way as in #2. PROS: flexible enough for "independent" projects, but offers "default" for standard projects CONS: development even more complicated than in #2 (but only a bit), possible production problems with user-defined repositories Mark (& Rysiek?), could you help me to choose the right option? Whether to invest more energy to get flexible solution, or we can last with one central repository? Eventually, we can start with #1 but have #3 as the final goal...? What's the priority & time frame?	1	1	https://issues.jboss.org/browse/ORG-208
13683	ORG-208	Please can we go for the simplest solution for now, i.e. defining the repository's URL in the module configuration. I don't want us to attempt a risky solution given the workload we're currently under.	1	1	https://issues.jboss.org/browse/ORG-208
13684	ORG-204	Putting the date in front of the month is very European. With the month first, you see a natural grouping of when the posts were created. May 05 2009 May 14 2009 Jun 08 2009 Jun 21 2009 I looked at many of the feed sources for our blogs and all of them that I saw had the day (Monday, Tuesday, etc.) One question we should ask, is the time important? Is it that critical that the user knows that a post went out in the morning or afternoon? I could see if we want to prove that we've posted something prior to a competitor to stay in front of say... system outages (zing!). What about this format "Tue May 05 2009 07:21 EDT by Pete Muir" or this one... "Tue May 05 2009 by Pete Muir"	0	0	https://issues.jboss.org/browse/ORG-204
13685	ORG-203	I think this feature request is handled already, or should we change anything still in this subject?	1	1	https://issues.jboss.org/browse/ORG-203
13686	ORG-118	Please can we rename the paragraph type and the 'Project Navigation' and 'Useful Links' styles to: Paragraph Type: General Link List Style 1: Lines Style 2: Block Conceptually the vertical navigation paragraph type you've created is a 'General Link List' with different styles: Lines - the list items are separated by lines Block - the list items are contained within a block We've got to be careful that we don't confuse people with terminology. Project Navigation is achieved using the horizontal bar at the top of a project page. When you add a new page to Magnolia using AdminCentral then we automatically add this to the Overview dropdown menu. This allows users to navigate around the project pages. This behaviour is the same for all projects using the default JBoss.org template so users only have to learn this once. If the project teams want to put other links on the page to other websites or maybe a couple of additional links to their project pages then they can use this 'General Link List'. It's not intended to replace the links automatically generated in the 'Overview' dropdown menu however. thanks, Mark	1	1	https://issues.jboss.org/browse/ORG-118
13687	ORG-59	Blocks and Issues not yet resolved: -figure out why "Tasks" is not allowed by admin -'new discussion' dropdown should be hidden? -make sure 'customize' links are positioned correctly -why are announcements coming back like zombies the next morning after "expire"? -change meta title "Jboss.org community: Home" -reload favicon with transparent background -re-org. template.ftl to move footer below wrapper and give it a z-index of 2 to solve the follover issues -make search bar text disappear on click	0	0	https://issues.jboss.org/browse/ORG-59
13688	ORG-59	x-unhide tooltip issue -reset your password page has a weird issue at the top, then the following dialogue as well -re-skin PRINT preview! THIRD RE-SKIN AFTER MIGRATION: x-add header name x-add favicon x-add "Search wiki..." x-add space under table header -Avatars are getting hidden in table structure for liquid layout in MULTIPLE places. So add css for max-width and overflow: hidden -Test announcements issue coming back like a zombie (still happening) -Move "personalize" and "customize" back and wrap a tab around it (because when you close the dialogue abt how to personalize- it stays closed) -"Guest" seems to be author of everything - is this a migrating issue? x-re-organize main page from admin- rename "spaces" "projects" and move widgets around x-what happend to avatars? x-in "projects" (jboss wiki) there's an empty tab bos showing - hide this. -Should I also hide the "Overview" tab here (as it's the only tab) or will there be other tabs once people comment on it, etc? x-fix missing underline in tables styles -make different table styles for the project pages, and people search pages??? -add specific poles for POC ease of use -"Recent History" and "Projects" headers have too much room below them because of tabs turned off - lessen this. -"Create Task" email still doesn't log you in and sends more than one reminder email x-Hide the overflow on the history dropdown x-move jboss hr right a few pixels x-redirect "show recent history" link in history dropdown to documents page not discussions x-hide all tabs in history x-un-hide all dropdowns and remove from templates x-make sure all "discussions" and "blogs" links are gone from right column navigation too (Your Stuff) x-take out whole "Your Blogs" box -maybe also take out "Recent Activity" box on the homepage -"recent content" box on personalize your homepage has overflow issues - hiding the authors and avatars again! -TURN OFF of both "projects" and "tasks" - widgets, links, right-col boxes -add more room between border-left and people icon? -table in "People search" is cutting off at the right AND is generally fucked up -wrap a tab around "personalize" and it should say "Your view" next to it -Authors on the main page "Recent Content" say "Guest" but should have the name of the real authors!!!!! x-Green favicon comes back on "Spaces" pages x-jboss.org logo is gone on a few pages	0	0	https://issues.jboss.org/browse/ORG-59
13689	JBPM-6137	[~cristiano.nicolai], It occurred to me the other day that Process Definitions perspective won't have new search mechanism at it's disposal. If we remove the old search there, what alternative will users have to do any kind of filtering of the process definitions list?	1	0	https://issues.jboss.org/browse/JBPM-6137
13690	JBPM-6137	[~bsremac] sure, we should add the quick search UI also to the process definition list. [~swiderski.maciej], for the sake of UI component consistency, where all of the runtime views are based on the dataset API but the definition list, would be possible that we use a data set query instead of the service call to retrieve the list of process definitions? Do you see any issues with that?	1	1	https://issues.jboss.org/browse/JBPM-6137
13691	JBPM-6006	[~roger600] Hi Roger, maybe couple thing to think of: -a sequence flow can have an expression condition property set which the user might not want to lose (as they can sometimes be very long and would have to retype). When removing node if incoming/outgoing connections have expression set maybe a confirmation that they are ok with losing this is ok before removing would be good? - in case of when the incoming connection is the "default" (also property on sequence flow) sequence flow coming out of a gateway, if removing it maybe we can remind user to set a different sequence flow as the new default? - a node can also have undirected associations (for example to text annotation) so when deleting the node those should probably also be deleted. - case where the removed node is the only node inside a container (lane, ad-hoc subprocess etc). In this case does it make sense to delete the container as well?	1	1	https://issues.jboss.org/browse/JBPM-6006

#	Requirement	Questions	Answered	Successfully Answered	URL
13692	JBPM-6006	Hey [~tsurdilovic], Thanks for the feedback! Lemme share my comments: * Good point about default connection - but where is this attribute located in the SequenceFlow bean? Is it the ConditionExpression (when empty String)? * Yeah good point too, associations are not supported and will require a different Edge Content type, so we can address this once this type is present, thanks! * Well not sure, I had been looking at some process modelers and no one does this... so for now I would let this feature for the future, if customers requires so, but you have more background than me, so will let you Tiho and [~KrisVerlaenen] decide, np for me! :) Thanks!	1	1	https://issues.jboss.org/browse/JBPM-6006
13693	JBPM-6005	[~uxdlc] Hi Liz, could you please provide some designs how the type selection could look like?	1	0	https://issues.jboss.org/browse/JBPM-6005
13694	JBPM-5967	Hi [~roger600], thank you for update! As I can see this PR have some issues and didn't merge yet. Can you check are reported issues are also reproducible in Stunner?	1	1	https://issues.jboss.org/browse/JBPM-5967
13695	JBPM-5925	[~pere.fernandez] ok sure i will make pr. do you know if opening empty form now is ok ?	1	1	https://issues.jboss.org/browse/JBPM-5925
13696	JBPM-5925	[~jsoltes] Could you provide a screenshot or better video of what you are experiencing because I cannot reproduce. You say the Start form is not generated (do you mean the process form?) and then that it is generated (correctly/incorrectly?) but not visible in what library?	1	1	https://issues.jboss.org/browse/JBPM-5925
13697	JBPM-5903	hey [~srambach], [~uxdlc]!! Can you please take a look at the latest comments [on this PR|https://github.com/kiegroup/kie-wb-common/pull/843]? So we did the palette is floating on top of canvas, in the top left position, but we realize that trying to open some existing processes (from other editors too), we cannot guarantee that some of the shapes could be hidden by the palette.... (see screenshot on the PR - the green start event cannot be selected for example....). Any ideas how to fix this? Adding some paddings, margins? Thanks in advance!!! CC [~federico.locci]	1	1	https://issues.jboss.org/browse/JBPM-5903
13698	JBPM-5902	Great, thanks for the feedback [~tsurdilovic]! One question anyway, right now Stunner does not need to go to server to perform validations, validations are done on the client while modelling... so probably this validations you mention, that are server side, could be only applied once clicking on save or validate button? This way there will be some structural validations that are applied while modelling , an other more complex structure and data model validations that apply once saving/validating (server side). WDYT?	0	0	https://issues.jboss.org/browse/JBPM-5902
13699	JBPM-5898	I believe the point of this query is to not ignore variable details but actually to include them in the response. so it seems there is no way to actually provide such support, is that valid conclusion?	1	0	https://issues.jboss.org/browse/JBPM-5898
13700	JBPM-5886	Are we sure we don't have this anywhere else in the codebase? The PR does not seem to touch the platform BOM, so other projects can still be using this.	1	1	https://issues.jboss.org/browse/JBPM-5886
13701	JBPM-5881	and where is the stack trace of the failure of console?	1	1	https://issues.jboss.org/browse/JBPM-5881
13702	JBPM-5880	Hi [~swiderski.maciej], can you set AsyncMode in an XML directly? Seems like Business Central won't let you do that. How do I add the AsuncMode property [here|http://pastebin.test.redhat.com/502878]? (Otherwise the default configuration)	1	1	https://issues.jboss.org/browse/JBPM-5880
13703	JBPM-5808	Hey [~cchianel] [~kgaevski] Thanks for your feedback guys! Well had been looking again and each modeller does different.. looking at genmymodel, it places the shape based on the top-left corner of the mouse location, but looking at the existing jbpm designer it does place it in the center. In fact no matter for me which one to use, so just double checking with UX guys about this, let's them point the right solution, is it good for you? Please [~mstevens1], [~uxdlc], what do you think about this? Once dragging an element from the palette, where it should be placed - the top-left of the created shape at the drag location, or the center of the shape at the drag location? Thanks in advance!!	1	0	https://issues.jboss.org/browse/JBPM-5808
13704	JBPM-5808	Hello [~kgaevski] [~cchianel] [~roger600] [~uxdlc] I agree with both of the key points above. I think the shape should be placed top left. But I also agree that, because the cursor is centered on the "shadow" object, that the user might think it will be placed in the center. Is it possible to move the cursor to be positioned top left, _off_ the shadow object? You might also consider changing the cursor to a pointer arrow instead of the four sided arrows. The latter makes it feel like you are pinpointing a specific (center) position. And finally, it looks like lane is the only object that doesn't have an icon / shadow object. But I assume that will change. Thank you, Matthew	1	0	https://issues.jboss.org/browse/JBPM-5808
13705	JBPM-5808	Hey [~kgaevski] [~cchianel] [~mstevens1] I'm agree with [~mstevens1], I think actually it's better using top-left corner. Mostly because the current approach is to show the icon, while dragging from the palette, rather than the shape, so the user, while not "dropping", doesn't know about the resulting size for the shape. So using the center approach, all cardinal directions as from the drop's target location, can be potentially used by the new shape, meanwhile using top-left approach, at least the resulting shape will require only space in the E/SE/S directions. I attached a video for the center approach, using the lane, where this happens, and IMO is not expected. Good point too [~mstevens1] about the cursor! Please [~cchianel] can you try to change the icon too, sounds easy? BTW [~mstevens1] about the lane icon, yeah that video is a bit old, actually it's using the lane icon you provided guys. Attached a palette screenshot where you can see it - while dragging from the palette, same icon appears over the canvas. Thanks for the feedback and good work guys!!	1	0	https://issues.jboss.org/browse/JBPM-5808
13706	JBPM-5808	Great [~cchianel], Please let's try to fix this "ficker" that appears on the mouse pointer yeah! AFAIK neither the workbench editors neither lienzo would be changing the cursor here, so probably it's stunner who, while dragging the icon from the palette, tries to change the icon at same time as you... Forgot to say this, please look for usages for [this|https://github.com/kiegroup/kie-wb-common/blob/master/kie-wb-common-stunner/kie-wb-common-stunner-core/kie-wb-common-stunner-commons/kie-wb-common-stunner-client-common/src/main/java/org/kie/workbench/common/stunner/core/client/canvas/AbstractCanvas.java#L96], and well you can probably use this too for changing the cursor while dragging, not sure tbh, anyway this can be the conflicting point, is Stunner calling it as while you drag too?	0	0	https://issues.jboss.org/browse/JBPM-5808
13707	JBPM-5671	Ah, I see. Thanks for the explanation [~swiderski.maciej]. Couple follow up comments / questions: * How does this partitioning functionality relate to permissions (e.g., I'm in Finance and can't see or create HR processes). If we are using permissions to restrict who can start / see the process definitions in different departments, wouldn't the other departments just be invisible to those without access * Who do we see as the current target user of this page? Ops or End Users? ** From what I understand, ops users would primarily be interested in deploying and managing servers and instances. If we have an Ops / Admin user can see across multiple departments, it makes their job harder to have to keep going to that dropdown to browse all definitions. If an Ops user can't see across to another department for security reasons, then access could simply be restricted to them. ** Similarly, if I'm an end user who wants to start a process, wouldn't I just want to see all the process definitions I personally can start?	1	1	https://issues.jboss.org/browse/JBPM-5671

#	Requirement	Questions	Answered	Successfully Answered	URL
13708	JBPM-5671	{quote}How does this partitioning functionality relate to permissions (e.g., I'm in Finance and can't see or create HR processes). If we are using permissions to restrict who can start / see the process definitions in different departments, wouldn't the other departments just be invisible to those without access{quote} this is handled by the kie server itself so if you have access to certain deployments then you will see them otherwise they are hidden. But that is quite limited and it's based on required roles in deployment descriptor of the project {quote}Who do we see as the current target user of this page? Ops or End Users?{quote} in my opinion business central UI for processes and tasks is certainly Ops/Admin focus as it's very generic so real users will be lost as it lacks the domain context. {quote}From what I understand, ops users would primarily be interested in deploying and managing servers and instances. If we have an Ops / Admin user can see across multiple departments, it makes their job harder to have to keep going to that dropdown to browse all definitions. If an Ops user can't see across to another department for security reasons, then access could simply be restricted to them.{quote} how I see it, less is better. If I have to go over 5 pages to find the definition I am looking for but I know it's in Finance department then it's way easier to look at only finance processes rather than all. Remember that they are still run on different servers so showing all of them at once hides the potential issues as it might give you a false impression that everything is fine because those that do have issues won't be in the initial list and trying to find them there manually will be time consuming. {quote}Similarly, if I'm an end user who wants to start a process, wouldn't I just want to see all the process definitions I personally can start?{quote} do you see end users (business users) to not be focused on their work area? If I need to start new payment process I know what I'd like to start so there is no need to show all process definitions as that just makes it fuzzy by forcing user to look over multiple pages to find what they need to do. Or sort by some stuff or filter etc. Anyway, that UI is not really for end users (business users) as it has no business context in my opinion. So overall, I see server template selection as a way for initial filtering/scoping to what you really want to work with. Some time ago we talked about having this as part of preferences so user directly can be brought to what they need.	1	0	https://issues.jboss.org/browse/JBPM-5671
13709	JBPM-5624	could you elaborate a bit on what do you expect? here you can find example of using jBPM with spring boot: https://github.com/mswiderski/jbpm-examples/tree/master/spring-boot-jbpm	1	0	https://issues.jboss.org/browse/JBPM-5624
13710	JBPM-5624	Many BPM engines have ready maven dependencies. For example Activiti and Camunda BPM have: <dependency> <groupId>org.activiti</groupId> <artifactId>spring-boot-starter-basic</artifactId> <version>${activiti.version}</version> </dependency> In my opinion such features increase developer efficiency, reduce boiler code and make work much simpler. You provided a good example in github. Could you please add maven-dependency so that people could just include it in pom.xml file and work with jbpm. So my point is, less code and less configuration far better.	0	0	https://issues.jboss.org/browse/JBPM-5624
13711	JBPM-5623	[~bsremac] this has been changed now, using the quick search we return all process instances related to that definition in the search tab. Are you ok for this issue to be closed?	1	0	https://issues.jboss.org/browse/JBPM-5623
13712	JBPM-5623	[~cristiano.nicolai] Yes, this is now history taking quick search feature into account. Closing it....	0	0	https://issues.jboss.org/browse/JBPM-5623
13713	JBPM-5607	hasNextFireTime method on the following line: long now = kruntime.getSessionClock().getCurrentTime(); Full Method: public Date hasNextFireTime() { Date date = orig.hasNextFireTime(); if (date == null) { return null; } long then = date.getTime(); long now = kruntime.getSessionClock().getCurrentTime(); // overdue timer if (then < now) { return new Date((now + OVERDUE_DELAY)); } else { return orig.hasNextFireTime(); } } 09:00:37,454 INFO [org.kie.server.services.impl.KieServerImpl] (EJB default - 1) Container 1 (for release id com.rhc.jbpm:com.rhc.jbpm.timer.demo:1.0.1) successfully started 09:09:45,342 INFO [stdout] (EJB default - 7) Timer is working triggered from host docker 09:09:45,345 WARN [org.drools.core.time.impl.DefaultTimerJobInstance] (EJB default - 7) Unable to execute timer job!: java.lang.NullPointerException at org.jbpm.process.instance.timer.TimerManager$OverdueTrigger.hasNextFireTime(TimerManager.java:465) at org.drools.core.time.impl.DefaultTimerJobInstance.call(DefaultTimerJobInstance.java:75) at org.drools.persistence.jpa.JpaTimerJobInstance.internalCall(JpaTimerJobInstance.java:59) at org.drools.persistence.jpa.JDKCallableJobCommand.execute(JDKCallableJobCommand.java:35) at org.drools.persistence.jpa.JDKCallableJobCommand.execute(JDKCallableJobCommand.java:21) at org.drools.core.command.impl.DefaultCommandService.execute(DefaultCommandService.java:36) at org.drools.core.command.impl.AbstractInterceptor.executeNext(AbstractInterceptor.java:41) at org.drools.persistence.SingleSessionCommandService$TransactionInterceptor.execute(SingleSessionCommandService.java:604) at org.drools.core.command.impl.AbstractInterceptor.executeNext(AbstractInterceptor.java:41) at org.drools.persistence.jpa.OptimisticLockRetryInterceptor.execute(OptimisticLockRetryInterceptor.java:82) at org.drools.core.command.impl.AbstractInterceptor.executeNext(AbstractInterceptor.java:41) at org.drools.persistence.jta.TransactionLockInterceptor.execute(TransactionLockInterceptor.java:73) at org.drools.persistence.SingleSessionCommandService.execute(SingleSessionCommandService.java:397) at org.jbpm.process.core.timer.impl.GlobalTimerService$DisposableCommandService.execute(GlobalTimerService.java:306) at org.jbpm.persistence.timer.GlobalJpaTimerJobInstance.call(GlobalJpaTimerJobInstance.java:77) at org.jbpm.persistence.timer.GlobalJpaTimerJobInstance.call(GlobalJpaTimerJobInstance.java:47) at org.jbpm.services.ejb.timer.EJBTimerScheduler.executeTimerJob(EJBTimerScheduler.java:78) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.jboss.as.ee.component.ManagedReferenceMethodInterceptor.processInvocation(ManagedReferenceMethodInterceptor.java:52) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:340) at	1	1	https://issues.jboss.org/browse/JBPM-5607

#	Requirement	Questions	Answered	Successfully Answered	URL
13714	JBPM-5607	Maciej pointed out I had run some tests with the wrong datasource configuration, so I think my previous results should be disregarded as invalid. My testing now validates failover is working and no duplication is occurring; however, I do see the following error on one node each time the timer fires: 11:53:39,006 ERROR [org.jboss.as.ejb3] (EJB default - 1) WFLYEJB0164: Exception running timer task for timer [id=c766bba7-005b-42eb-9cda-cb5cf3069568 timedObjectId=kie-server-redhat.kie-server-redhat.EJBTimerScheduler auto-timer?:false persistent?:true timerService=org.jboss.as.ejb3.timerservice.TimerServiceImpl@459df624 initialExpiration=2017-01-26 11:53:39.0 intervalDuration(in milli sec)=0 nextExpiration=2017-01-26 11:53:39.0 timerState=ACTIVE info=EjbTimerJob [timerJobInstance=GlobalJpaTimerJobInstance [timerServiceId=1-timerServiceId, getJobHandle()=EjbGlobalJobHandle [uuid=1-1-1]]]] on EJB kie-server-redhat.kie-server-redhat.EJBTimerScheduler: java.lang.RuntimeException: org.postgresql.util.PSQLException: ERROR: could not serialize access due to concurrent update at org.jboss.as.ejb3.timerservice.persistence.database.DatabaseTimerPersistence.shouldRun(DatabaseTimerPersistence.java:421) at org.jboss.as.ejb3.timerservice.TimerServiceImpl.shouldRun(TimerServiceImpl.java:1110) at org.jboss.as.ejb3.timerservice.TimerTask.run(TimerTask.java:124) at org.jboss.as.ejb3.timerservice.TimerServiceImpl$Task$1.run(TimerServiceImpl.java:1214) at org.wildfly.extension.requestcontroller.RequestController$QueuedTask$1.run(RequestController.java:497) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1142) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:617) at java.lang.Thread.run(Thread.java:745) at org.jboss.threads.JBossThread.run(JBossThread.java:320) Caused by: org.postgresql.util.PSQLException: ERROR: could not serialize access due to concurrent update at org.postgresql.core.v3.QueryExecutorImpl.receiveErrorResponse(QueryExecutorImpl.java:2455) at org.postgresql.core.v3.QueryExecutorImpl.processResults(QueryExecutorImpl.java:2155) at org.postgresql.core.v3.QueryExecutorImpl.execute(QueryExecutorImpl.java:288) at org.postgresql.jdbc.PgStatement.executeInternal(PgStatement.java:430) at org.postgresql.jdbc.PgStatement.execute(PgStatement.java:356) at org.postgresql.jdbc.PgPreparedStatement.executeWithFlags(PgPreparedStatement.java:168) at org.postgresql.jdbc.PgPreparedStatement.executeUpdate(PgPreparedStatement.java:135) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.postgresql.ds.PGPooledConnection$StatementHandler.invoke(PGPooledConnection.java:423) at com.sun.proxy.$Proxy10.executeUpdate(Unknown Source) at org.jboss.jca.adapters.jdbc.WrappedPreparedStatement.executeUpdate(WrappedPreparedStatement.java:537) at org.jboss.as.ejb3.timerservice.persistence.database.DatabaseTimerPersistence.shouldRun(DatabaseTimerPersistence.java:410) ... 8 more	1	1	https://issues.jboss.org/browse/JBPM-5607
13715	JBPM-5601	[~cristiano.nicolai] [~swiderski.maciej] the kie-server role was introduced to secure access to the kie-server remote REST api (as this is in general more powerful than the UI so we want to have more restricted access that just existing roles user, developer, etc.). With the separation of kie-server it seems all users now require this role, which will break this more strict enforcement. Assuming we can't secure kie-server remote api anymore separately (there's no way to distinguish a jbpm-console request from a remote one), I think just removing the kie-server role and giving users with roles like admin,analyst,etc. access to kie-server would be required. Wdyt?	1	1	https://issues.jboss.org/browse/JBPM-5601
13716	JBPM-5597	just one comment, why does it have to have "blank" in the name? isn't that enough to have kjar-archetype as I'd say you'll anyway assume to be empty... and then when we have kjars with some content maybe they should have additional part in the name.	1	0	https://issues.jboss.org/browse/JBPM-5597
13717	JBPM-5572	Added comments from [~KrisVerlaenen]: Connections - can't seem to select a connection - selecting an endpoint for reconnection seems difficult often (delay, need hold) - no bendpoints Toolbar - if nothing is selected and you click delete it first asks for confirmation and then shows an error, shouldn't action be disabled if no selection?	0	0	https://issues.jboss.org/browse/JBPM-5572
13718	JBPM-5546	Hey [~cchianel] [~kgaevski] Thanks for looking into this. Lemme share my comments: * I'm agree on just selecting the shape on right click on it, and avoid starting a drag operation too, it results an unexpected behavior * This means that in Stunner we need to accept the right button as well on the SelectionControl. This way the shape will be selected on right click. * But on the other hand, trying to avoid starting a drag operation once right clicking on the canvas, this happens as well on pure lienzo applications. Try it for example [here|https://www.lienzo-core.com/lienzo-ks/#WIRES_DOCKING]. In fact there already exist a bug for Lienzo about this, [the issue 145|https://github.com/ahome-it/lienzo-core/issues/145]. IMO this is a bug on lienzo side, as I think nobody expects starting a drag with the right button, in fact I tested other different editors and right button never starts a drag. So in order to address this bug, the [issue on lienzo|https://github.com/ahome-it/lienzo-core/issues/145] must be analyzed and addressed first, once this part is clear, we can decide the actions to take in Stunner. WDYT guys? Thanks!	1	1	https://issues.jboss.org/browse/JBPM-5546
13719	JBPM-5467	[~jsniezek], I have two questions: Should the main palette have some tooltip or text? What is the expected behaviour when the user presses the "more" link?	1	1	https://issues.jboss.org/browse/JBPM-5467
13720	JBPM-5467	Hi [~pere.fernandez], sure we can help! Not sure exactly what the need is though. Could you explain?	1	1	https://issues.jboss.org/browse/JBPM-5467
13721	JBPM-5467	Hey [~jsniezek]!! Sorry for bothering again with the palette, but as from [~mark.proctor]'s feedback, there are some customers that have some "complains"/suggestions about the palette design. Probably we should talk all together about ti on next Monday's meeting, but I think the main reason is about "it uses too much white-spaces and icons seem too big". Re-opening the ticket for keeping track of it. Thanks Joe! CC [~pere.fernandez]	1	1	https://issues.jboss.org/browse/JBPM-5467
13722	JBPM-5250	Can you add [~hkonsek] to this issue?	1	1	https://issues.jboss.org/browse/JBPM-5250
13723	JBPM-5250	This has been fixed on 6.5.x branch and will be available in the next release. [~jreimann-1] do you need this backported also into older branches? Also, would you be able to test with the latest 6.5.0-SNAPSHOTs to verify your use case if working?	1	1	https://issues.jboss.org/browse/JBPM-5250
13724	JBPM-5239	not sure I understand this request - org.kie.api.executor.Command is purely dedicated to the jbpm executor that executes these commands as background tasks and thus the method is expected to return results from the execute method as it will be invoked by background thread anyway. So what would be the benefit to spin another thread for running the method execute which was invoked by already separate thread? This command (org.kie.api.executor.Command) is not meant tbe be used directly by application code - like calling its execute without jbpm executor.	1	1	https://issues.jboss.org/browse/JBPM-5239

#	Requirement	Questions	Answered	Successfully Answered	URL
13725	JBPM-5223	first of all [~jludvice] there is no need to apologize quite the opposite, thanks for the time you put into describing and filing the jira! now let's go over the design of work items and work item handlers - the most important concept here is that work item is design time component, so this is the activity in the process that should happen. It defines inputs (what needs to be given so it can perform operation) and expected output (what shall be returned). On the other hand, work item handler is a runtime component that is responsible for handling the action defined by work item. So it's a matter of separation of concerns. If we would merge these two we loose a huge advantage that is - possibility to change handlers without changing the process. So the master component here is work item and not the handler. With that said, - we cannot build work items (and their parameters) based on discovered handler on the time when the process is created (authored) as it is quite possible that on runtime environment that handler will not exists - so it's a false assurance that it will work - especially going towards v7 where workbench won't include runtime at all - idea to improve the handler with annotation to produce valid and readable documentation- +1 to that it's really great to make it self-explanatory to promote reuse of handlers - improve work item definition to allow to specify required and optional parameters, validation schema so we can check if given value is valid or not does this make sense to you guys?	0	0	https://issues.jboss.org/browse/JBPM-5223
13726	JBPM-5220	I believe we should add one more endpoint (or enhance existing one) which would accept the ReleaseId and would return only the containers with such kjar deployed. Deciding if an kjar is deployed can be simply done on the client -- client would request the containers for specified ReleaseId and in case the returned list is empty it means the kjar is not deployed (this needs to somehow take into account the container state as well). I currently see the following options: 1) New endpoing with ReleaseId as part of the URL Example: {{.../containers/releaseId/org.example:artifact:1.0.0.Final}} - easy to use - consistent with the current API (e.g. using URI parts) 2) ReleaseId as URL parameter for existing endpoint Example: {{.../containers?releaseId=org.example:artifact:1.0.0.Final}} - easy to use - not really consistent with current API as URL parameters are not being used 3) ReleaseId encoded in the HTTP body; using the existing endpoint Example: {{.../containers}}, with body {code} <requested-release-id> <groupId>org.example</groupId> <artifactId>artifact</artifactId> <version>1.0.0.Final</version> </requested-release-id> {code} - harder to use (one needs to mess with the HTTP request, instead of just using the URL) - not really obvious (it least that's how it looks to me) [~swiderski.maciej], [~ksuta] WDYT? Which option would you prefer? Or do you perhaps see some better option?	1	1	https://issues.jboss.org/browse/JBPM-5220
13727	JBPM-5220	I have a question. Can I use for filtering only part of the name (like: org.ex*) or I must know whole name (org.example)?	1	1	https://issues.jboss.org/browse/JBPM-5220
13728	JBPM-5189	very good feature request! would you mind extending that and submitting PR?	1	1	https://issues.jboss.org/browse/JBPM-5189
13729	JBPM-5173	I think the easiest and least intrusive solution is the third one. We just need to remove unused *activate* parameter from constructors and add *Vetoed* annotation on those classes. [~swiderski.maciej], what do you think? If you agree, I can make those changes.	1	0	https://issues.jboss.org/browse/JBPM-5173
13730	JBPM-5163	Hi [~davidredhat], tried to reach you on freenode but I think we missed each other. Is the documentation subjira planned to be resolved soon, or is it put on the backburner for now?	0	0	https://issues.jboss.org/browse/JBPM-5163
13731	JBPM-5107	Hello [~abhisheksit2004] Can you explain better what do you want to achieve? Maybe if I understand it we can find a workaround. Cheers!	1	1	https://issues.jboss.org/browse/JBPM-5107
13732	JBPM-5071	Will the impacts of this particular feature be documented anywhere, [~swiderski.maciej]? I'm specifically concerned about the asset management, which the users will sure be looking for and we should make the switch as easy for them as possible.	1	1	https://issues.jboss.org/browse/JBPM-5071
13733	JBPM-5068	Hi there, will it be a complete REST API in API in 7.0.0 Final?	0	0	https://issues.jboss.org/browse/JBPM-5068
13734	JBPM-5036	[~selrahal] see this issue was approved only for 7.0 - the request to backport it to 6.4.x was rejected. See RHBPMS-4778 Is this ok ?	0	0	https://issues.jboss.org/browse/JBPM-5036
13735	JBPM-4899	that usually means it was not discovered. Could you please describe the steps you have taken and maybe upload the jar that contains the implementation of the listener so we could examine it? Would be actually good to move this discussion to user forum instead of jira: http://jbpm.org/community/forum.html	1	0	https://issues.jboss.org/browse/JBPM-4899
13736	JBPM-4881	[~davidredhat], is there any support from dashbuilder to parse query strings like the one mentioned in here? I'm guessing something that would read the string and build a DataSetFilter/ColumnFilter(s).	1	0	https://issues.jboss.org/browse/JBPM-4881
13737	JBPM-4857	[~psiroky], how is the PR coming along?	1	0	https://issues.jboss.org/browse/JBPM-4857
13738	JBPM-4851	Hi [~davidredhat], I saw the PR and it looks like this jira can be closed. Are there any other change which will be implemented? If not, could you please close it? It is part of BPMSPL-257 and I would like to know if I can check it off in my list. Thanks.	0	0	https://issues.jboss.org/browse/JBPM-4851
13739	JBPM-4849	Just double-checking, should they be excluded from current and any new filter in the process definition list correct?	1	1	https://issues.jboss.org/browse/JBPM-4849
13740	JBPM-4847	[~jsniezek], as discussed, can you please amend the details (screenshot) which I can't access.	1	1	https://issues.jboss.org/browse/JBPM-4847
13741	JBPM-4847	Agreed, I'd say that storing XML files is a bit of an extreme use of process variables, and I'm not sure how common something like this is. As for solutions, I like the idea of a popover, but I think there's more design risk with that option (how will it look at different content sizes, in practice?). The edit popup with the text-area would be most effective for this case, particularly for large amounts of content. And perhaps we could _also_ do a tooltip (with a character limit, so it doesn't take up the whole screen), for the "normal" usage of process variables, where process variables are too long for the table column, but are not entire files. What do you think [~cristiano.nicolai]? Joe	1	1	https://issues.jboss.org/browse/JBPM-4847
13742	JBPM-4847	Great, looks good. So, there will be both the "edit" button to bring up the text box (Selection_038.png), and clicking on the row will bring up the pop-over(Tooltip_036-7). Correct?	0	0	https://issues.jboss.org/browse/JBPM-4847
13743	JBPM-4847	[~cristiano.nicolai]: Hello Cristiano, if I mouse over and click on a process variable value, I cannot see the tooltip. Tested on 6.4.0-SNAPSHOT. Can you see the tooltip?	1	1	https://issues.jboss.org/browse/JBPM-4847
13744	JBPM-4838	I'm trying to reproduce it. No success though. There is no limit in the number of columns that can be created in a tabular report. Can you please add further details about how to reproduce it please? As a mention, notice that tables don't allow to modify their columns once the "Group by" setting is on. In order to configure a table, the group by setting must be disabled. Once the configuration of the table heading is done, the group by setting can be enabled.	1	0	https://issues.jboss.org/browse/JBPM-4838
13745	JBPM-4832	Maciej, I think we should aim to support the 3 marshalling formats without exceptions to avoid surprises. What is the problem of supporting XStream if we support JAXB as well?	0	0	https://issues.jboss.org/browse/JBPM-4832
13746	JBPM-4832	after reviewing pull requests I think that building this as new KIE Server extension can bring more value where some of them are: * users can enable or disable extension if not needed * encapsulate UI (or sort of UI) facing capabilities * users can provide new provider types that will be discovered on startup * we can then add more features to this kind of UI extension like return process diagram (png) with active and completed nodes what do you think?	1	1	https://issues.jboss.org/browse/JBPM-4832

#	Requirement	Questions	Answered	Successfully Answered	URL
13747	JBPM-4812	Hey Maciej, Sorry for slow response here - been travelling. I don't see anything on the REST API here. Do you want that addition in this issue or a new one created?	1	1	https://issues.jboss.org/browse/JBPM-4812
13748	JBPM-4812	Just read through the PRs and I think I've got my head wrapped around the impl. Overall looks good, but I have two concerns: # Current API as described in [droolsjbpm-integration/pull/216|https://github.com/droolsjbpm/droolsjbpm-integration/pull/216] does not appear to support multiple task variables in the query. This seems to be an API limitation, not a limitation of the indexing mechanism and supporting JPA data structure but I could be wrong. # Current API returns a [TaskSummary|https://github.com/droolsjbpm/droolsjbpm-integration/blob/master/kie-server-parent/kie-server-api/src/main/java/org/kie/server/api/model/instance/TaskSummary.java] which does not have a data structure to return task variables. [~swiderski.maciej] Can you confirm that these are correct of the current impl and are they correct for the impl we are targeting for jbpm 6.4/bpms 6.3? If so, we probably need to work together see what can be done to enhance the feature in our timetable. Customers will need to be able to query with multiple variables. And the result will need to be able to display some subset of task variables in the resultsSet. We may need to explore a variable expansion approach to limit the data exchange.	1	1	https://issues.jboss.org/browse/JBPM-4812
13749	JBPM-4812	Maciej, All reasonable points. The problem is that this implementation falls short of the use case that our key customers want to solve, which is the ability to create a custom task dashboard that displays a small number of task variables along with each individual task. They then want to filter down the list displayed in this dashboard by using searching on one or more of the task variables. So while it may be possible to write a customer indexer to support multi-variable queries, I still have no way to display these variables in the dashboard. Is there a way to explore a variable expansion to the API which returns the indexed string if a certain param is passed to the API?	1	0	https://issues.jboss.org/browse/JBPM-4812
13750	JBPM-4779	pull request created to allow override default KieStoreServices class and further define command service class. Would that work for you?	1	1	https://issues.jboss.org/browse/JBPM-4779
13751	JBPM-4723	Marco, do we want the WADL to be published on an URL of the REST API as WSDL is? For example, Jersey generates and publishes WADL by default https://jersey.java.net/documentation/latest/wadl.html. But of course, we have to support more implementations than just Jersey.	1	1	https://issues.jboss.org/browse/JBPM-4723
13752	JBPM-4723	[~rsynek] I'd rather not publish the WADL: it doesn't accurately reflect the operations that we provide and may be confusing to users. For example, in the documentation, I list several task operation REST URL's: {noformat} [POST] /task/{taskId}/activate [POST] /task/{taskId}/claim {noformat} et cetera.. (see the [REST Task documentation|http://docs.jboss.org/jbpm/v6.2/userguide/jBPMRemoteAPI.html#d0e15775]). But in the WADL, the URL's will show up differently, since in the code, I don't have 10+ different methods for each task operation, but [one method that accepts a couple different possibilities|https://github.com/droolsjbpm/droolsjbpm-integration/blob/master/kie-remote/kie-remote-services/src/main/java/org/kie/remote/services/rest/TaskResourceImpl.java#L125-L177]. Does that make sense?	1	0	https://issues.jboss.org/browse/JBPM-4723
13753	JBPM-4723	[~rsynek] Huh, I think we've lost sight of what we're talking about now. Or otherwise, I was mistaken about the discussion in the first place! Are we talking about how to check the backwards compatibility of the REST API or are we talking about whether or not to publish the WADL of the REST service as part of kie-wb or BPMS? Please let me know! If it's in regards to publishing the WADL, then we should at least move that to another Jira. When you say "publish", I assume you mean that it should be made publically available via http, as opposed to simply being stored as part of the {{kie-remote-services}} jar. Regardless of whether or not the WADL is _published_, we can always still use it for testing backwards compatibility. From what I understand from Ivo Bek (see the mojo document I referenced), the WADL can *not* fully accurately describe our REST API. However, again: Are we discussing generating the WADL for backwards compatibility tests *OR* are we discussing publishing the WADL? If we are discussing publishing the WADL, then that discussion is orthogonal to generating the WADL for backwards compatibility tests?	1	0	https://issues.jboss.org/browse/JBPM-4723
13754	JBPM-4705	Thanks Is it possible to use LATEST? eg. org.jbpm.test:single-project:LATEST "When using process that signals with external scope (process id: external-project.throw-external-signal) it will start a process (process id: single-project.start-with-signal) as it has signal start event and will trigger process that waits in intermediate catch event (regardless of what strategy is used) assuming the SignalDeploymentId was set to org.jbpm.test:single-project:1.0.0-SNAPSHOT on start of the process"	1	1	https://issues.jboss.org/browse/JBPM-4705
13755	JBPM-4672	Should this be closed?	0	0	https://issues.jboss.org/browse/JBPM-4672
13756	JBPM-4663	We came across the need to use these concepts. Beneficiary - is the entity the case is raised for. Can be a person, a team, a department, an organization. Usually the beneficiary is the party who benefits from the work done on the case and pays for it. Beneficiary can be internal or external to organization. Owner - is the entity responsible to deliver the work done by the case. Owner is the main contact for the case. Participants - are the entities working together on the case to achieve the desired outcome Contacts - when a message is sent between pools there is a need to specify contact information. One entity can have several contacts; for example a person can be contacted to several email addresses or an organization can have multiple representatives related to a case. Also contacts can have attributes such as main contact or delivery contact. At the moment the jBPM implementation for messages is just technical and not usable in practical terms. A message requires the sender to know the target process ID; this is an impossible requirement for an external participant (other pool) as that participant does not have visibility to technical details. External participants can send and receive messages based on contact information only and not on some internal technical detail such as process id.	0	0	https://issues.jboss.org/browse/JBPM-4663
13757	JBPM-4661	Is there any validation to be done on this new field? Does it need to be serialized as CDATA or just xml attribute value?	1	1	https://issues.jboss.org/browse/JBPM-4661
13758	JBPM-4648	This is a very important feature. What happens if a process is not perfect? After the process is started user realizes that additional tasks need to be performed. Adding ad-hoc tasks to a running process solves the problem. Processes can be reviewed and improved based on the history of ad-hoc tasks that were started for that type of process. Also ad-hoc tasks could be added against the case itself. For example we do not have a hire process but we have a "hire" type case. The user can start using the case based on ad-hoc tasks only. After a while an analyst can look to past "hire" cases and find out what kind of tasks were required. Then the analyst can investigate the sequence of steps and convert the case to a BPMN model. The key is that cases should have types: sales, warranty, diagnostics, ... ; some of these cases could be turned into processes if they are stable and repeatable.	1	0	https://issues.jboss.org/browse/JBPM-4648
13759	JBPM-4648	Hi Mauricio I can bring on board the business person who drove our requirements. We have a system in production and are very happy with it. I can talk to him tomorrow morning. Our system is simple but handles all flexibility issues we need. Also there are other concepts that are very useful for us such as "can start date/time" and "must start before date/time".	1	0	https://issues.jboss.org/browse/JBPM-4648
13760	JBPM-4648	What is your timezone? We are in GMT +10. Regards	1	0	https://issues.jboss.org/browse/JBPM-4648

#	Requirement	Questions	Answered	Successfully Answered	URL
13761	JBPM-4648	Ready for Google Hangouts in 10 minutes? My email address is dan.cimpoesu [at] clearbridge.com.au and alex.cimpoesu [at] clearbridge.com.au If you want us to invite you what email address should we use?	1	0	https://issues.jboss.org/browse/JBPM-4648
13762	JBPM-4629	what would be the way to do this using the deployment Service? I tried: @ApplicationScoped @MyKjar public class MyDeploymentServiceCDIImpl extends DeploymentServiceCDIImpl { protected RuntimeEnvironmentBuilder boostrapRuntimeEnvironmentBuilder(KModuleDeploymentUnit deploymentUnit, DeployedUnit deployedUnit, KieContainer kieContainer, MergeMode mode) { DeploymentDescriptor descriptor = deploymentUnit.getDeploymentDescriptor(); if(descriptor == null){ descriptor = new DeploymentDescriptorImpl(); deploymentUnit.setDeploymentDescriptor(descriptor); } descriptor.getConfiguration().add(new NamedObjectModel("drools.processInstanceManagerFactory", MyJPAProcessInstanceManagerFactory.class.getName())); deploymentUnit.setDeploymentDescriptor(descriptor); return super.boostrapRuntimeEnvironmentBuilder(deploymentUnit, deployedUnit, kieContainer, mode); } } But this can't work because the configuration cannot be changed this way.	1	1	https://issues.jboss.org/browse/JBPM-4629
13763	JBPM-4609	could you provide reproducer for this? what surprises me is that you say it uses CMT but then code fails with No transaction exceptions which indicates there is no transaction - meaning CMT somehow did not work for that case? btw, I see you run it with CR4 have you tried with 6.2.0.Final?	0	0	https://issues.jboss.org/browse/JBPM-4609
13764	JBPM-4553	Where can I find the Jbpm7.0.0.Final ? source code any idea I don't see it in the git hub any branch yet. I am not sure that this is right way to ask this question sorry about that if that so .	1	0	https://issues.jboss.org/browse/JBPM-4553
13765	JBPM-4501	not sure I got the point behind this request... could you please elaborate or provide an example of what you try to achieve?What sort of an ID are you after? Note that there is already a process id that is an ID of given process definition. Placing any ID in the bpmn2 itself will make it shared across all process instances created out of that definition.	1	1	https://issues.jboss.org/browse/JBPM-4501
13766	JBPM-4501	Andrea, you should be able to achieve that with relying on UniqueId taken from meta data as long as the process definition is not updated. So when running it by the engine the UniqueId should remain unchanged regardless of how many process instances have been started and if the server was restarted or not. What do you mean by change it - do you change the model with process designer or?	1	1	https://issues.jboss.org/browse/JBPM-4501
13767	JBPM-4496	provided possibility to re-queue failed jobs via UI as currently executor itself do not expose REST api. Changing job data most likely will be addressed post 6.2 community. Duncan, do you mind dropping another jira for it?	0	0	https://issues.jboss.org/browse/JBPM-4496
13768	JBPM-4484	So, I started to attempt an implementation and it didn't take too long to run into a problem. UserGroupCallback defines a method "getGroupsForUser" but does not have a "getUsersForGroup" method. Would this be something that could be added? Is there a reason it wasn't there in the first place? It seems like a natural compliment to getGroupsForUser	1	1	https://issues.jboss.org/browse/JBPM-4484
13769	JBPM-4484	Ahhh...that's why I asked if there was a reason. Would an optional new interface that could be provided work? I realize this is getting messier. As a potential workaround, I was able to add a listener, and afterTaskAdded do the check on my own, and do a claim if there is only one. Right now I'm calling back into the taskService, which causes another event to be fired. I suppose I could just set the actuallyOwner the way the AddTaskCommand is doing and be done with it.	1	1	https://issues.jboss.org/browse/JBPM-4484
13770	JBPM-4482	[~swiderski.maciej] You're the jbpm-audit lead :) so I've left it assigned to you because I'd like your feedback on it. Of course, I *love* working with reflection (have you seen any of my tests? ;D), so this looks like an interesting feature to me. On the other hand, because of the possible performance costs (reflection..) and other pitfalls (it's a feature that deserves careful consideration because it may contain bugs that will be very hard to support), I want to make sure to get your and [~KrisVerlaenen]'s feedback about this.	1	1	https://issues.jboss.org/browse/JBPM-4482
13771	JBPM-4482	I would say that for this purpose it would be better to use variable persistence strategies (unless we see any drawbacks of using that) instead of extending variable instance log capabilities. With persistence strategies we can easily persist variables into an external storage where the easiest will be to use data base as backend but persist the variables as JPA entities. In fact we already have out of the box support for it. Some details how variable persistence strategies can be used are described here: http://mswiderski.blogspot.com/2014/02/jbpm-6-store-your-process-variables.html Marco, what do you think? Would that be sufficient or you see it will lack some of the requested features? Access to it will be directly via process instance variables, and not via history. History will still have string representation of the variable.	1	0	https://issues.jboss.org/browse/JBPM-4482
13772	JBPM-4482	[~swiderski.maciej] I agree that variable persistence strategies cover a lot of what users already need, but I don't think it's very user friendly. - currently, variable persistence does not work out-of-the-box and it's not simple for beginner developers to set it up. - It can't be used automatically with BPMS/kie-wb/console: there's no way for users to easily set it up with these deployables Does that make sense? One of the main reasons for this jira is part of the query work I've done recently. This jira is also actually for a better internal framework with regards to process instance variables. For example, at the moment, it's possible to retrieve process instance variable information in 3 different ways: 1. query the user-defined process instance variable strategy tables, if they are in use 2. query the (jbpm-audit) variable instance logs. 3. make sure to hold a reference to the process instance itself, and retrieve the variables from that instance. There's no interface that allows us (internal developers) or users to generically query variable information regardless of whether configuration/strategy 1, 2 or 3 is being used. Querying 3 is in some cases impossible. What I'd like to acheive, at the very least, is a step in that direction: adding optional (internal) persistence for process instance variables is a good first step, I think. What do you think?	1	1	https://issues.jboss.org/browse/JBPM-4482

#	Requirement	Questions	Answered	Successfully Answered	URL
13773	JBPM-4482	{quote} currently, variable persistence does not work out-of-the-box and it's not simple for beginner developers to set it up. {quote} could you explain what does not work out of the box? {quote} It can't be used automatically with BPMS/kie-wb/console: there's no way for users to easily set it up with these deployables {quote} you can do it very easy with deployment descriptors, see here:https://github.com/droolsjbpm/jbpm-playground/blob/master/translations/src/main/resources/META-INF/kie-deployment-descriptor.xml#L11 {quote} What I'd like to acheive, at the very least, is a step in that direction: adding optional (internal) persistence for process instance variables is a good first step, I think. {quote} I see persistence variable strategies to be actually that, a way to have optional and flexible persistence options for process variables. Although I agree that this in general lacks of searching capabilities. And by that I believe VariableInstanceLog is the simplest way (yet quite powerful) to get some degree of search ability by process variables. At some point we where thinking to include indexing engine like lucene to help with searching options but lack of time stops us a bit ;) I might not know all the details you already thought about so maybe I don't see the right things here so would be great to get more details on what you planned to have here. So either drop some notes here or we can talk about it too. anyway, effective searching through process variables is more then needed question is how to approach it. From my point of view we have two options: - as you mentioned to implement another way to persist variables - implement additional layer dedicated for searching I am in favor of the second one mainly because it is designed for given purpose. As generic persistence option is more complex and might overlap persistence strategies. P.S. nice to have this design discussion back in place :) thanks!!!	1	0	https://issues.jboss.org/browse/JBPM-4482
13774	JBPM-4478	Maciej, First, thank you for implementing this feature request. I hope it is of benefit to not only me but many other users of the application. Could you just clarify how to use the jar with 6.2 if we don't want to use auto-discovery? By patch do you mean personally modifying the source and recompiling jbpm-runtime-manager? Or is there a way to programatically do it within the current APIs. Thanks Again,	1	0	https://issues.jboss.org/browse/JBPM-4478
13775	JBPM-4478	I'm not too clear what the release schedule is (is it published somewhere?), but is there any chance this could be squeezed into 6.2? It sounds like it wouldn't impact a lot of the other code. Thought I would ask.	1	0	https://issues.jboss.org/browse/JBPM-4478
13776	JBPM-4477	[~gangli], we currently support the ability to activate/deactivate a process deployment. I believe that matches the requirements you mentioned in here. Once a process deployment is deactivate, you can't create new instances. Would you mind verifying and let us know if we could close this issue, thanks.	0	0	https://issues.jboss.org/browse/JBPM-4477
13777	JBPM-4456	Can you please specify the use case that you are trying to create when you get this exception? Im quite sure that there are alternative methods that are working. If a method is not currently supported it is probably due the fact that we provided a better alternative for it.	0	0	https://issues.jboss.org/browse/JBPM-4456
13778	JBPM-4443	Toshiya, there are already bunch of tests for timers including Quartz based: https://github.com/droolsjbpm/jbpm/tree/master/jbpm-test/src/test/java/org/jbpm/test/timer Is there anything in particular that you would like to have included?	1	1	https://issues.jboss.org/browse/JBPM-4443
13779	JBPM-4443	Thanks Maciej, With the master, I still see NullPointerException with PerRequestRuntimeManager/PerProcessInstanceRuntimeManager during the test. Increasing "org.jbpm.timer.delay" and Thread.sleep() in getManager() doesn't help. Don't you see the Exception? {noformat} java.lang.NullPointerException at org.jbpm.persistence.timer.GlobalJpaTimerJobInstance.call(GlobalJpaTimerJobInstance.java:67) at org.jbpm.persistence.timer.GlobalJpaTimerJobInstance.call(GlobalJpaTimerJobInstance.java:46) at org.jbpm.process.core.timer.impl.QuartzSchedulerService$QuartzJob.execute(QuartzSchedulerService.java:281) at org.quartz.core.JobRunShell.run(JobRunShell.java:216) at org.quartz.simpl.SimpleThreadPool$WorkerThread.run(SimpleThreadPool.java:549) {noformat}	1	1	https://issues.jboss.org/browse/JBPM-4443
13780	JBPM-4443	Toshiya, I cannot see any NPE while running the test cases for any of the runtime managers - trying the test case delivered by you GlobalQuartzDBTimerServiceTest.testContinueTimer can you check if you have that commit included? https://github.com/droolsjbpm/jbpm/commit/d03c4330a230871c8b223398f67d85069ee0956f	1	1	https://issues.jboss.org/browse/JBPM-4443
13781	JBPM-4443	Hi Maciej, Yes, I confirmed that the commit is included. On my end, it's constantly reproducible when I leave this line. {code:java} public class GlobalQuartzDBTimerServiceTest extends GlobalTimerServiceBaseTest { ... public static Collection<Object[]> persistence() { Object[][] data = new Object[][] { { 1 }, { 2 }, { 3 } }; {code} In other words, if I test only 3 (=PerProcessInstanceRuntimeManager), it's not reproducible. I comment out other tests in GlobalQuartzDBTimerServiceTest/GlobalTimerServiceBaseTest for convenience but it doesn't seem to affect. With some debugging, I noticed that the NullPointerException occurs when Timer is triggered after manager.close(). It means that the scheduler doesn't correctly shutdown on manager.close(). I found that setUp()/tearDown() makes QuartzSchedulerService.timerServiceCounter confused. When I comment-out the following lines, the issue has gone. {code:java} @Before public void setUp() { cleanupSingletonSessionId(); System.setProperty("org.quartz.properties", "quartz-db.properties"); testCreateQuartzSchema(); // globalScheduler = new QuartzSchedulerService(); // ((QuartzSchedulerService)globalScheduler).forceShutdown(); } @After public void tearDown() { try { // globalScheduler.shutdown(); } catch (Exception e) { } System.clearProperty("org.quartz.properties"); } {code} Hmm, should we move testContinueTimer() (= intends to test the default RuntimeManager behaviour) to a different test class?	1	1	https://issues.jboss.org/browse/JBPM-4443
13782	JBPM-4392	Marco, could you please elaborate little bit about this? What we want to achieve with it (I like the idea just need more info to confirm I know what to do)	1	1	https://issues.jboss.org/browse/JBPM-4392
13783	JBPM-4380	Hi Kris, We're investing on this commit and see what we can do based on this commit. After a quick view, we guess that this is the back-end api support in jbpm core of the description for process instance, right? So we can add front-end display function based on it, is it?	1	1	https://issues.jboss.org/browse/JBPM-4380
13784	JBPM-4379	Seems that we did not see this "Administrator" property for human task now in BPMS 6.0.1, so which version does it support? BTW does this variable support a group? You know that sometime the admin should be a group which is supplied from usergroupcallback.	1	1	https://issues.jboss.org/browse/JBPM-4379
13785	JBPM-4379	HI, Kris, I've tried the BusinessAdministratorId property in Human Task node, and found that the user id of this property will be persisted in PeopleAssignments_BAs, does this means the user has been set to Business Administrator of this task? BTW, even if this is what has been done in core engine, actually these is no support in front-end for this feature yet because i did not find any changes for task operations for this task BA. So i think this is what we can do next step. But before that, i still have confusions in several task operations in the user guide and i did not find where to do these operations in UI now: fail, forward, nominate, remove, skip, stop, suspend, resume. Especially for remove, i can not understand if a task has been removed by a admin, what's the status of the process instance which this task belongs to?	0	0	https://issues.jboss.org/browse/JBPM-4379

#	Requirement	Questions	Answered	Successfully Answered	URL
13786	JBPM-4379	Hi,Kris I have same issue like Marcell. this code (((org.jbpm.process.instance.impl.ProcessInstanceImpl) kcontext.getProcessInstance()).getMetaData().get("OwnerId"))) dose not work. I debugging this function, i found this reason is the CDI runtime(DeploymentServiceCDIImpl) is not aware of that listener(IdentityProviderAwareProcessListener), If you want to get more details , you can read email (Re: next step for UI works on JBPM-4379 RegisterableItemsFactory getRegisterableItemsFactory). So i want to modify it. But i found InjectableRegisterableItemsFactory (in DeploymentServiceCDIImpl) very differnet to KModuleRegisterableItemsFactory (in KModuleDeploymentService). Can u give some suggests about it?	1	1	https://issues.jboss.org/browse/JBPM-4379
13787	JBPM-4343	This error is thrown because the service task should reference an interface (that will be called), but the process parser cannot find any interfaces defined. Could you attach the process XML you created? It might be that the order of the elements is not as expected, are there any interface elements in the XML, and is it located before the process tag?	0	0	https://issues.jboss.org/browse/JBPM-4343
13788	JBPM-4341	Thx Maciej. I have discovered something new "SHARED_EM" ;-) Is this property "STANDALONE_LOCAL_SPRING_SHARED_EM" like the other strategies documented somewhere ? KIE_SESSION, STANDALONE_JTA, STANDALONE_JTA_SPRING_SHARED_EM, STANDALONE_LOCAL_SPRING_SHARED_EM, STANDALONE_LOCAL; The official doc is very poor concerning JPA & Tr (http://docs.jboss.org/jbpm/v6.0.1/userguide/jBPMPersistence.html#d0e4914).	0	0	https://issues.jboss.org/browse/JBPM-4341
13789	JBPM-4337	Hi, that's already happening inside the human task service, do you mean to add that feature in the User Interface? Can you please elaborate more?	1	1	https://issues.jboss.org/browse/JBPM-4337
13790	JBPM-4337	Hi, I mean I can see that there is a global administrator role ("admin" or "administrators"? What are the differences between these two?): - When I'm connected to the Business Central GUI as an admin (a user that I created with the "admin" role with the add-user script) or an admininistrator (during jbpm6.0.1+EAP6.1.1 installation, I had to create an admin which has all the jBPM roles (admin,analyst,...) in the User Interface), I'm not allowed to do "revoke" or "complete" on tasks that a simple user has claimed and started. Indeed via the User Interface I necessarily need to be the owner of the task to do these operations - when you say "that's already happening inside the human task service", do you mean that even if the User Interface doesn't show me the possibility to do "revoke" on tasks, I should be able to use the API "revoke" in a webapp? I would like to create a local administrator (supervisor of several users): - I can see that there is a task.setBusinessAdministrator API (human tasks have PeopleAssignments with PotentielOwners, BusinessAdministrators and ActualOwners), but I cannot find the notion of Business Administrator in the User Interface (neither in the process designer nor in menus/buttons/task displays). - What roles/rights do businessAdministrator(s) have on human tasks? I would like to know if I can "revoke" tasks that potential owners have claimed even if I'm not a PotentialOwner of the task, but only a BusinessAdministrator of the task. Thank you	1	0	https://issues.jboss.org/browse/JBPM-4337
13791	JBPM-4337	Thank you. Actually my request is not only about the UI. I've noticed that the feature is not in the UI, but I'm interested in using this feature by calling the APIs (revoke, complete,...) in a webapp. I can see that the notion of BusinessAdministrator exists; OK so it's not usable via the UI yet, but my hope Is that the APIs are operational however? Is the parameter BusinessAdministrator implemented in the human task (as are PotentialOwner and ActualOwner)? Can I set a BusinessAdministrator on a task by calling the API task.setBusinessAdministrator and then call getTaskAssignedAsBusinessAdministrator to get the task ids list? What are the possible actions/rights on these tasks as a BusinessAdministrator (only viewing? I hope he can act on it)? You say "We have this feature in our plans which is to build a task list for administrators that can see all the tasks not only the ones that they own", will the administrator only be able to "see" the tasks list? I would also be very grateful if you could answer my first question about the different notions of admin; it could help me to understand how to do what I need. Thank you! Regards	0	0	https://issues.jboss.org/browse/JBPM-4337
13792	JBPM-4320	Hi I tried to use metadata with workbench following the example provided on github. When I import the .bpmn2 file in workbench the metadata section is stripped out. Is this a bug in workbench or there is a way to save metadata in .bpmn2 file?	0	0	https://issues.jboss.org/browse/JBPM-4320
13793	JBPM-4316	It looks like perhaps a custom implementation of IdentityProvider along with ServicesAwardAuditEventBuilder might work. Is that correct? It's hard to see from the barely documented classes.	1	1	https://issues.jboss.org/browse/JBPM-4316
13794	JBPM-4307	Yes. You are right. I have just downloaded the current master branch (6.1.0-betaX) and the business task is working properly. Thanks for your help.	0	0	https://issues.jboss.org/browse/JBPM-4307
13795	JBPM-4249	I believe that documenting that is better choice as not always CMT is used and by that it should not introduce any additional overhead. So I would still opt to have it disabled by default but well documented. wdyt?	0	0	https://issues.jboss.org/browse/JBPM-4249
13796	JBPM-4233	Is this also available in the Java REST client?	1	0	https://issues.jboss.org/browse/JBPM-4233
13797	JBPM-4217	Shouldn't the implementation classes have the @Alternative annotation?	0	0	https://issues.jboss.org/browse/JBPM-4217
13798	JBPM-4178	[~cdsosi] Could you explain the use case a little bit? Why would you need a method that returns all available process definitions? The reason I ask is that this is obviously contrary to the whole philosophy of limiting process definitions (and ksessions and runtimes, etc.) to deployments.	1	1	https://issues.jboss.org/browse/JBPM-4178
13799	JBPM-4178	Of course I can explain. To my mind the philosophy is to make jBPM6 much more simple to use and to avoid complexity in handling sessions and so on. I think we agree on that point ;) I want to make a generic tooling so that I can retrieve process definitions first and then possibly interact with some process instances especially as far as human task are concerned. I may be wrong but I do not want if possible to know in advance what is available or not in a repository and that's not the case with proposed REST API because you always need to know at least a valid deployment id to use Remote REST AP. Moreover the AuditLog service retrieve all kind of information even if a specific deployment id is specified (that's what 6.0.0.Final user documentation tells) so I think that this parameter (deploymentId) appears to be a little useless to my opinion of course. Querying database through AuditLog service should be more generic, don't you think so? You can also refer to a post on forum where I had some exchange with Maciej about that, this may give you further information => https://community.jboss.org/thread/234895 HTH	1	1	https://issues.jboss.org/browse/JBPM-4178
13800	JBPM-4178	[~cdsosi], thanks, this helps a little bit. I agree with your point about querying the database through the AuditLogService -- that that should be more generic. Unfortunately, the <deploymentId> parameter is still needed becuase of.. how the AuditLogService is set up. Once 6.0 is done, I'll work on making the <deploymentId> parameter not needed.. All right, I'm starting to understand why you would want this, but there's one more challenge I want to bring up: you can have 3 deployments, "A", "B" and "C", and all 3 deployments contain the exact same process definition. Hmm.. it's now obvious to me why this operation would be helpful. I won't be able to deliver it with 6.0.0.Final, but I'll make sure it get's in the next release.	1	1	https://issues.jboss.org/browse/JBPM-4178

#	Requirement	Questions	Answered	Successfully Answered	URL
13801	JBPM-4178	Thanks Marco for your comment / explanations. Just some other questions to clarify : Now with jBPM 6.0.0 is it possible to have 3 differents business processes with same id but in different depoyment unit? Does the deployment unit represent a kind of namespace? >>there's one more challenge I want to bring up: you can have 3 deployments, "A", "B" and "C", and all 3 deployments contain the exact same process definition Exact same process definition means ? 1. Same process id only (with different content) 2. Same process id and actually same bpmn2 xml file Franck	1	1	https://issues.jboss.org/browse/JBPM-4178
13802	JBPM-4178	Hi Marco, That will be great. We are also looking for this feature to be available. Since we had this feature available in 5.4 it is very necessary feature to have the backward compatility. Do you have any estimate timeframe for 6.1.x release?	1	1	https://issues.jboss.org/browse/JBPM-4178
13803	JBPM-4178	Marco, just to add to the conversation as it might be quite useful - what would you say to expose same sort of information we have in UI, mainly following: - process definition id - process name - process version - deployment id - list of variables defined for given process definition including their types - optionally user and group involved in the process all these information we already have for UI in jbpm-kie-services so remote API might be simply taking that data, wdyt?	1	1	https://issues.jboss.org/browse/JBPM-4178
13804	JBPM-4178	Hi Marco, Has this feature been resolved? Can you please tell what is the rest API? Thanks.	1	1	https://issues.jboss.org/browse/JBPM-4178
13805	JBPM-4178	Hi Marco, There must be some confusion - the URL you specified provides the details of a specific process definition from a specific deployment (you have to know the deployment id and process definition id in order to construct the URL). My understanding is that this enhancement should instead provide all process definitions regardless of the deployment id - so I would expect the REST URL not to depend on a <deploymentId> or a <process-def-id>. In this way for example, if I want to start a process instance using REST API from my web app I can display for the user a list of process definitions to select from (just like in the jBPM console), instead of expecting him to manually enter some id's - the end user shouldn't care about those. Can you please clarify this? Thank you.	1	1	https://issues.jboss.org/browse/JBPM-4178
13806	JBPM-4178	Hi guys, Thanks for the questions! Seriously, critical attentiont to features always makes them better. The other new operation that I didn't mention is: {{http://<host>:<port>/<instance>/rest/runtime/<deploymentId>/process/}} (also {{GET}}, of course) If you want to retrieve every single process definition, then you can do the following: 1. call {{.../rest/runtime/<deploymentId>/process}} for the list of process def ids 2. for each process def id, retrieve the process def (see operation above). I had assumed this would be enough. Some users will have 100's of process definitions per deployment, so that retrieving a total list (wihout paging) is not a feature that's practical to implement. However, I can definitely add a feature that retrieves a *paged* list of the process definitions (default page size of 10 or so). It seems like both of you would want that: is that correct?	1	1	https://issues.jboss.org/browse/JBPM-4178
13807	JBPM-4178	[~kgeis], If I understand correctly, you're deploying the process from a different component than your front end. So this is what's happening: 1. Component A deploys process to new deployment id. 2. BPMS/Console deploys process, etc.. 3. Component B needs to find new process definition, but doesn't have info from component A Is this correct?	1	1	https://issues.jboss.org/browse/JBPM-4178
13808	JBPM-4178	[~kgeis], So: 1. You (person) author and deploy process in the console 2. BPMS/Console deploys process.. 3. Front-end doesn't have the deployment id (which only you know), so can't find process Is that correct? (If not: I guess what I'm trying to establish is why one part of your application doesn't know what the deployment id is. Sorry that I'm not understanding this.. :/ Could you explain which parts of your application there are and why the deployment id information isn't available to certain parts of your application?) Thanks!	1	1	https://issues.jboss.org/browse/JBPM-4178
13809	JBPM-4178	[~dragos.snk], Could you attach the stack trace (NPE) to this jira? That is, save the stack trace to a file (text) and then attach the file to this jira? That would be very helpful! Thanks.	1	1	https://issues.jboss.org/browse/JBPM-4178
13810	JBPM-4166	Duncan, could you take a look at the prototype that I believe should cover your case with bpmn2 standard way - transformations. Here you can find pull requests for it: https://github.com/droolsjbpm/droolsjbpm-knowledge/pull/55 https://github.com/droolsjbpm/jbpm/pull/184 kie-api where single DataTransformer interface is defined so it can be referenced in jbpm and if needed used to provide custom implementation of transformers. jbpm where we have default implementation of transformers and registry where such transformers are placed so runtime engine can take advantage of them. Currently implementation covers task nodes (work item based) so whenever transformation is used the variable will be transformed on the fly without actually changing process variable value that is its source. Out of the box, we provide two (at the moment) implementations of DataTransformer: - MVEL based - Java Script based - this is implementation that uses javax.script support from Java SE 6 which is JSR 223. this gives us extremely useful option to add more languages based on users request by simply dropping given scripting engine on class path and it will be automatically found and registered. By default Java 6 and above provides JavaScript engine based on Rhino impl. There are quite few engines already supporting JSR 223 like groovy, freemarker, etc (more here https://java.net/projects/scripting/sources/svn/show/trunk/engines) with that users can take advantage of the preferred scripting language as part of the bpmn2 transformation. Moreover new/custom ones can be provided by anyone, like the one you created based on JAXB, then registered in the engine and off you go... What do you think?	1	1	https://issues.jboss.org/browse/JBPM-4166
13811	JBPM-4164	When you say dumper for process instances, I assume you're referring to dumping an instance of the (RuleFlow)Process class as BPMN2 XML? I do believe that is valid use case however (and it's for example used like this for dynamic process instances). What's the reason for believing this might not be a good idea? Similarly, what would be the reasons for forcing to hide this? What is the risk of this being (ab)used? Note that it's currently only part of the internal implementation classes. That means that people can use it at their own risk, but there's no guarantee we won't change the API etc. It is however used internally to turn a Process object into an BPMN2 XML representation (to be able to restore it later). If people want to use it for similar use cases, I don't see a problem with that. Note however that it doesn't address other features like for example preserving the XML structure (like comments etc.) or simulation data etc.	0	0	https://issues.jboss.org/browse/JBPM-4164
13812	JBPM-4144	Is there a plan to work on this? If so, in which release can we expect this?	1	0	https://issues.jboss.org/browse/JBPM-4144
13813	JBPM-4135	How do I add process model view in my project? I want same debug option in my project so can you guide me?	0	0	https://issues.jboss.org/browse/JBPM-4135
13814	JBPM-4125	Would it be possible to get a small test component that illustrates exact usage of jbpm with spring in your case? Just to avoid any misunderstanding especially with resource local configuration.	1	1	https://issues.jboss.org/browse/JBPM-4125
13815	JBPM-4110	How to reproduce this behaviour: 1. Add the kie-wb-6.0.0.CR2-tomcat-7.war to a Tomcat 7 installation 2. Read the README.txt file and do some configurations 3. Start the application and look if it starts clearly (it should) 4. Build the https://gist.github.com/mbiarnes/6379664 branch and fetch the droolsjbpm/jbpm-dashboard/jbpm-dashboard-distributions/target/jbpm-dashbuilder-tomcat-7.war 5. Deploy it to tomcat/webapp 6. Start the server Now you should see in the logs this error.	1	1	https://issues.jboss.org/browse/JBPM-4110

#	Requirement	Questions	Answered	Successfully Answered	URL
13816	JBPM-4071	Hi Thomas, This sounds as a good community contribution, but I have doubts about storing things in the database. Can the Projection/Prediction be calculated on the fly? This kind of feature can be really interesting if you pick up the simulation data, that can give you business estimates about how much time in average a user needs to wait for the prediction to happen. You also need to take into account the fact that assignments can be done with variable data, which means that sometimes you don't know who the potential owner is.	1	0	https://issues.jboss.org/browse/JBPM-4071
13817	JBPM-4071	Hi Mauricio, Yes actually in my current implementation I also try to resolve the variables that are assigned to Human Task by looking at the Process Variable Map. About resolving things on the fly, it surely can be done but the problem is the user UI is a grid that list all the process instances which he/she may be involved in the future. So if I resolve things on the fly I thought it will take so long considering there should be a lot of active process instances. What do you think? I can share the model code if needed.	1	0	https://issues.jboss.org/browse/JBPM-4071
13818	JBPM-4071	Can you share a couple of screenshots of the UIs that will be affected? I though that the user will be able to see the Future tasks in their task list. I didn't understood that the process instance list will be affected by this feature. I guess that you have both. By the way are you working with the jbpm console ng? and jBPM 6? If you have a fork of the jbpm repository with the changes I can take a look. Cheers	0	0	https://issues.jboss.org/browse/JBPM-4071
13819	JBPM-4046	Marco, I am not quite sure the incoming and outgoing are optional. According to xsd they are marked with minOccurs = 0 but that does not have to mean it's optional. As there might be activities that do not have either incoming or outgoing sequence flow at all, for example start and end events respectively. In my opinion, for custom parsers it is even better to have these elements there as it makes the parsing easier from structure point of view. In most of the cases bpmn2 xml is built like this: - item definitions - properties - activities (nodes) - sequence flows - BPMN DI - relationship (out simulation) so when parsing nodes (assuming stream/event based approach is used) it is quite difficult to refer to sequence flows at that time as sequence flows were not yet parsed thus having its incoming and outgoing node directly makes it much easier. What is the issue with having them in bpmn2 xml? Sorry if I did not get it directly.	0	0	https://issues.jboss.org/browse/JBPM-4046
13820	JBPM-4029	Could you upload the bpmn2? Or even better if you could provide test case. In general nested OR gateway should work fine We cannot rely on outgoing sequence flows on diverging gateway because there can be another split in between or timer that will generate more instances. I reopen this issue again and try the example attached	1	1	https://issues.jboss.org/browse/JBPM-4029
13821	JBPM-4029	indeed the fix added for loops caused the nested gateways to be wrongly evaluated. Have done bit of changes to be able to keep track over looks (node executions) and now both scenarios are handled properly, at least based on the test I added. Marcos, could you please give it a try and let me know how does it look like from your point of view? Once comment about the loop example you have uploaded to the linked jira - it will not be properly handled as there is a mix between XOR and AND gateways - in other words it might be considered as invalid processing as the AND gateway will not wait for loop iterations as it's not aware of it because XOR gateway in front of it. To make it proper switch places of XOR and AND gateways to be more like this -see attached diagram. Will keep this jira open until you comment on it. Thanks	1	1	https://issues.jboss.org/browse/JBPM-4029
13822	JBPM-4029	Marcos, could you give a try once again? Just pushed some clean ups and improvements for the mechanism that decides when the OR converging gateway should be triggered. Added the loop inside OR gateway example to test cases.	1	1	https://issues.jboss.org/browse/JBPM-4029
13823	JBPM-4029	What you mean by this case? Could you upload zip/tar files instead of rar. If you could upload process image too so I know what case you have in mind.	1	1	https://issues.jboss.org/browse/JBPM-4029
13824	JBPM-4008	we should probably group "Front, Back, Forward, Backward | Alignment, Sizing" into a single dropdown menu rather than having them all next to each other (sample image attached). wdyt?	1	1	https://issues.jboss.org/browse/JBPM-4008
13825	JBPM-3989	I assume you created your ruleflow using the Drools&jBPM Eclipse plugin, version 5.5.0.Final? Would it be possible to attach the full XML file for reference?	0	0	https://issues.jboss.org/browse/JBPM-3989
13826	JBPM-3989	this xml seems to be valid, I was able to import it into jbpm web designer without any issue. Though it might have been fixed in version 6. Do you mind checking and confirming (or not) if this is still the case?	0	0	https://issues.jboss.org/browse/JBPM-3989
13827	JBPM-3987	jbpm aligns common libraries with versions used by Jboss Application Server to ensures the best compatibility. Current version we use is 1.4.5 which in general comes out of the box. Although there is nothing against to upgrade that version when embedding jbpm in the system. Have you tried with 1.4.5? Maybe if you could provide more details about the issue we could look into it as if that is actually an issue when using gmail that would be rather serious. Personally have tried several times with gmail account and never ran into any problems.	1	1	https://issues.jboss.org/browse/JBPM-3987

#	Requirement	Questions	Answered	Successfully Answered	URL
13828	JBPM-3979	room. (Ping timeout: 245 seconds) [07:48am] maciejs: salaboy: what you mean suspend? process instances or... [07:48am] salaboy: maciejs: yes a process instance to the state SUSPENDED [07:49am] salaboy: maciejs: we did a review session yesterday and most of the guys were freaking about aborting a process instance.. they prefer to suspend it instead [07:49am] salaboy: maciejs: I remember to see in the Enum the SUSPENDED state [07:49am] maciejs: salaboy: to be honest I don't think we have that supported (or at least I am not aware of it) [07:50am] salaboy: maciejs: cool.. I'm just checking because it ring some bells here.. but I don't remember where that code is.. so I will look at it.. don't worry [07:50am] salaboy: maciejs: public static final int STATE_SUSPENDED = 4; [07:50am] salaboy: maciejs: I'm not than crazy yet hehe [07:50am] maciejs: salaboy: yes, I know about the state, but don't recall anything that sets it [07:51am] salaboy: maciejs: I'm going back to work now.. i need to meet prakash in 1h [07:51am] maciejs: salaboy: have fun [07:51am] salaboy: maciejs: ttyl.. enjoy the rest of the day [07:51am] maciejs: salaboy: same to you [07:53am] mriet: salaboy, re suspended state: isn't that what happens when the process hits a human task or similar? [07:53am] mriet: salaboy, I don't think you can "set" that state via the api. [07:53am] salaboy: mriet: I'm checking that [07:54am] mriet: salaboy, I don't think it's a state for "external use", so to speak.. maybe krisv knows? [07:54am] salaboy: mriet: I remember setting that state in Drools Flow.. in the early days [07:54am] mriet: salaboy, [07:54am] salaboy: mriet: yeah.. the code will tell us the truth [07:56am] salaboy: mriet: doing a grep to all the jbpm code it seems that the SUSPENDED state is not even used in the jbpm code [07:59am] chm007 left the chat room. (Quit: Computer has gone to sleep.) [08:01am] omolinab is now known as omolinab|trn. [08:01am] salaboy: mriet: maciejs looking at the abortProcess method it looks like are just missing the method for suspending a process which just set the ABORTED state.. in the ProcessRuntimeImpl [08:01am] salaboy: mriet: maciejs i will write a mail to krisv to see what his opinion on the matter [08:03am] mriet: salaboy, don't forget persistence complications.. (suspended <=> process instance cache?) [08:03am] jrenaat left the chat room. (Ping timeout: 252 seconds) [08:03am] salaboy: mriet: it should be the same as abort.. is there something special happening when we abort a persistent process? [08:04am] salaboy: mriet: a cache? why is htat? [08:04am] salaboy: mriet: it's too early here for me to think about that kind of stuff hehe [08:04am] mriet: salaboy, JPAProcessInstanceManager [08:04am] salaboy: mriet: let me check that [08:05am] maciejs: mriet: salaboy suspend should be to sort of temporarily exclude process instance from execution [08:05am] maciejs: mriet: salaboy meaning when you call like a signal so it should not be triggered [08:05am] mriet: salaboy, sorry.. time for a coffree break? [08:06am] salaboy: mriet: hmm it's just a filter IMO.. but I will check it [08:06am] salaboy: mriet: enjoy your coffee [08:06am] maciejs: mriet: salaboy but I see potential issues with human tasks - complete human task when process instance is suspended we could end up with inconsistent state [08:07am] salaboy: mriet: maciejs yeah.. I was thinking about that.. it shoudn;t be so hard.. I mean.. you can just suspend all the tasks .. we should have the same policies with aborted processes [08:07am] salaboy: maciejs: mriet I can imagine the same kind of problems with suspend and abort [08:08am] • mriet is going to grab some coffee.. brb. [08:08am] maciejs: salaboy: yeah, but that would mean to extend workitemhandler with suspend and resume capabilities [08:08am] salaboy: maciejs: so I guess that we should use our process event listener to abort all the tasks from an aborted process [08:09am] maciejs: salaboy: for abort it's not an issue we already have a way to abort work item via handler [08:09am] salaboy: maciejs: not really.. I mean if you suspend a process all the tasks are suspended.. if you bring it back to life you will just change the state of the task not of the work item handler.. [08:09am] maciejs: salaboy: and how are you going to communicate between process instance and tasks? [08:10am] salaboy: maciejs: i see your point.. but I also see the business side of it.. [08:10am] maciejs: salaboy: I completely agree with the need for suspend and resume	0	0	https://issues.jboss.org/browse/JBPM-3979
13829	JBPM-3968	Cristiano, I'm not sure how moving the examples module to it's own (maven?) profile would help. Could you explain the proposed solution a little bit more?	1	1	https://issues.jboss.org/browse/JBPM-3968
13830	JBPM-3947	Will this be implemented? According to a related jira this is not an issue but I disagree and think it's a fundamental problem with sessions being disposed and persistent timers	1	1	https://issues.jboss.org/browse/JBPM-3947
13831	JBPM-3946	does this issue refer to providing a custom editor on the Properties tab when selecting a custom WorkItemHandler? If so I could really use something like this to massively improve the tooling experience in the "Continuous Deployment" solution for OpenShfit / Fabric8 using jBPM. Basically we have a BuildTrigger custom WorkItemHandler which has 2 input strings; when a user clicks on the BuildTrigger node in the designer we really want to show our own simple web form by default which shows those 2 strings (along with our own smart defaulting / auto-completing widgets if possible). Right now the Properties view shows mostly stuff thats not relevant to the user (see screen shot https://issues.jboss.org/secure/attachment/12388661/buildtrigger-page1.png); then if the user spots the very first line called "Assignments" then clicks the Down icon a popup appears that lets them edit the values (with no validation/custom widgets) - see screenshot 2 https://issues.jboss.org/secure/attachment/12388662/buildtrigger-page2.png. While this works; its not totally awesome or massively intuitive ;). Any improvements in this area would be awesome.	1	1	https://issues.jboss.org/browse/JBPM-3946
13832	JBPM-3946	I have added a change in master that does the following: * If in your WID you define a workitem for example (note the customEditor property set !) ["name" : "MyTest", "parameters" : ["FirstName" : new StringDataType(), "LastName" : new StringDataType(), "Address" : new StringDataType(), "Country" : new StringDataType()], "customEditor" : "something", "displayName" : "MyTest", "icon" : "defaultemailicon.gif"] Designer will treat each parameter as an individual property entry. The Assignments and Data Input / Data Output properties will be hidden. In addition these properties will be displayed in the "Assignments" section of the property panel (new section that is created just for these) Here you can see what that looks like: http://i.imgur.com/ykAenUh.png So all you have to do to get this is set the customEditor property to anything really (except a blank string)..any value will do. The underlying BPMN2 generated remains the same with the proper ioSpecification section created with correct data inputs / outputs and assignments. Currently its an all/nothing deal (you cannot pick/choose which parameters in the wid you want to externalize into their own properties..they are all done) just to get you moving forward. In the future we will update this so users can pick /choose as well as we will look at the parameter type and associate it with the correct individual properties editors (like boolean, enumeration, script, etc etc). Please give this a try and let me know if it works for you.	1	1	https://issues.jboss.org/browse/JBPM-3946
13833	JBPM-3905	Would it be possible to test this on a more recent version? Spring support was only added after the 5.2.0.Final release, so it is possible that the issue is already addressed in a future release. Would it be possible to for example try with the 5.2.5.Final or the 5.4.0.Final release?	1	0	https://issues.jboss.org/browse/JBPM-3905

#	Requirement	Questions	Answered	Successfully Answered	URL
13834	JBPM-3905	Thanks Kris, We tried to test the same with a latest version: 5.4.0.Final release, we were not able to get the same error anymore. However, when the app server is started it throws below error log: Hibernate: select task0_.id as col_0_0_, deadline1_.id as col_1_0_, deadline1_.deadline_date as col_2_0_ from Task task0_, Deadline deadline1_ where (task0_.archived=0)and((deadline1_.id in(select startdeadl2_.id from Deadline startdeadl2_ where task0_.id=startdeadl2_.Deadlines_StartDeadLine_Id))or(deadline1_.id in(select enddeadlin3_.id from Deadline enddeadlin3_ where task0_.id=enddeadlin3_.Deadlines_EndDeadLine_Id)))and(deadline1_.escalated=0) order by deadline1_.deadline_date Hibernate: select SESSIONINFO_ID_SEQ.nextval from dual Hibernate: insert into SessionInfo (lastModificationDate, rulesByteArray, startDate, OPTLOCK, id) values (?, ?, ?, ?, ?, ?) 83388 23/01 11:10:33,112[pool-2-thread-1] ERROR transaction.support.TransactionSynchronizationUtils.invokeAfterCompletion - TransactionSynchronization.afterCompletion threw exception java.lang.AbstractMethodError: org.drools.container.spring.beans.persistence.DroolsSpringJpaManager.clearPersistenceContext()V at org.drools.persistence.SingleSessionCommandService$SynchronizationImpl.afterCompletion(SingleSessionCommandService.java:448) Even after getting the above error, Everything works fine and we are able to complete and end process for jbpm human task. But, After manually executing delete script for jbpm workflow tables, when we start process and than tries to complete human task, we get below error: Exception: Row was updated or deleted by another transaction ... If I restart my app. server and start process and than tries to complete human task, it works fine. Could you please provide us with any hint to avoid restarting the app server, everytime we delete or update jbpm tables.	1	0	https://issues.jboss.org/browse/JBPM-3905
13835	JBPM-3905	Noor, did you already get rid of this exception? (java.lang.AbstractMethodError: org.drools.container.spring.beans.persistence.DroolsSpringJpaManager.clearPersistenceContext()V) If yes, how?	0	0	https://issues.jboss.org/browse/JBPM-3905
13836	JBPM-3902	Is there a patch to backport this fix to 5.4.0 Community?	1	0	https://issues.jboss.org/browse/JBPM-3902
13837	JBPM-3889	Attached an example, which demonstrates that commit per node transition with BRMS 5.3.0. The customer doesn't feel nice with it because it requires development on their side and more development if they consider recovery which is described in (b). FYI : "jBPM5 Developer Guide" Chapter 11 suggests persisting BusinessEntity object to hold ksessionId/processInstanceId/workItemId relation in WorkItemHandler for async WorkItem but it has the same concern... Isn't it too much effort on application side? Can jBPM provide some help out-of-box?	1	0	https://issues.jboss.org/browse/JBPM-3889
13838	JBPM-3889	Toshiya Kobayashi <tkobayas@redhat.com> made a comment on [bug 958386|https://bugzilla.redhat.com/show_bug.cgi?id=958386] +++ This bug was initially created as a clone of Bug #889091 +++ Description of problem: a) Being able to commit a transaction per node transition regardless of synchronous or asynchronous. Hopefully by configuration basis (not additional development effort by users) b) Recovery support (continuing process instance execution) after JVM reboot (including JVM crash case). The key point is less development effort by users --- Additional comment from JBoss JIRA Server on 2013-01-06 21:48:37 EST --- Toshiya Kobayashi <tkobayas@redhat.com> made a comment on jira JBPM-3889 Attached an example, which demonstrates that commit per node transition with BRMS 5.3.0. The customer doesn't feel nice with it because it requires development on their side and more development if they consider recovery which is described in (b). FYI : "jBPM5 Developer Guide" Chapter 11 suggests persisting BusinessEntity object to hold ksessionId/processInstanceId/workItemId relation in WorkItemHandler for async WorkItem but it has the same concern... Isn't it too much effort on application side? Can jBPM provide some help out-of-box?	1	1	https://issues.jboss.org/browse/JBPM-3889
13839	JBPM-3869	Marco Rietveld <mrietvel@redhat.com> made a comment on [bug 880874|https://bugzilla.redhat.com/show_bug.cgi?id=880874] Hisao, it doesn't look like we'll be adding this functionality to BRMS 5.x, but it will be present in BPMS 6.0+. This issue applies to BRMS. What would you like to do with this issue? I'd like to set it to CLOSED / NEXTRELEASE.	1	1	https://issues.jboss.org/browse/JBPM-3869
13840	JBPM-3869	Hi, Marco. Thank you for your work. I'm fine with that this functionality will be present in BPMS 6.0+. Could you set to CLOSED/NEXTRELEASE please? Thx	1	0	https://issues.jboss.org/browse/JBPM-3869
13841	JBPM-3820	QQ, it seems like there might be a mistake in your configuration -- did you forget to add a " at the end of the <beans element? Also, could you attach your file to the issue?	1	1	https://issues.jboss.org/browse/JBPM-3820

#	Requirement	Questions	Answered	Successfully Answered	URL
13842	JBPM-3814	obvious to you, but I want to document this for others as well, which is why I'm adding more documentation than necessary) To start with, the problem can only happen if the user takes control of the transaction before a the synchronized SSCS.execute() method is called. If no one else takes control of the transaction, then we don't have a problem, as far as I know. Multiple processes started on the same ksession won't cause this, I think: when a process is started (or otherwise run), it's run from inside the synchronized execute() method. Any subsequent process starts must go through the wait and go through the same synchronized method. So unless the ksession is not the owner of the transaction, running multiple processes on the same ksession can never cause a problem. And this is where you're right: running a ksession in an environment where the ksession _is *not* the transaction owner_ will cause problems. In fact, the race condition that you describe will then be present! So, what we can do is use a reentrant lock to fix this problem -- like you suggested. This, like any code that tries to solve a race condition, makes the code complexer. Does it introduce any _new_ race conditions that weren't previously present? Let's take the simplest situation: the transaction owner is _not_ the ksession but something else, and a ksession method is called that, for example, starts a process and thus goes through the SSCS.execute() method. What happens? Well, the the lock is gotten, the process is started, and since the SSCS is _not_ the owner of the transaction, the transaction is not committed. Then the SSCS.execute() method is completed, and the ksession.startProcess() method (or whichever method) returns control to the user code. If we call another ksession operation in the same thread, then the reentrant lock will allow us to lock again (it's reentrant ;)) and we're fine. Eventually, the transaction will commit, the lock will be released, and everything is good! Okay, let's make the situation a little more complex: two threads, and the transaction owner is again *not* the ksession but something else. When the first thread calls, it's the same as the first situation: everything's good! Then the second thread calls ksession method that involves the SSCS.execute block. At this point, # the transaction has not yet been committed and # the first thread owns the lock. When the second thread (and any thread after that) call an operation, then it will block until the transaction is committed. Unfortunately, having the ksession behave this way will break existing implementations for users who expect the following: # I can share the same ksession object with multiple threads # I can take control of the transaction. Users who do this will see their performance decrease drastically, I think. And they will also be working with code that might break on them because of the race condition that you've described. The heart of the problem is the following: - the code assumes that the ksession has control of the transaction So when the code (SSCS) does _not_ have contol of the transaction, then weird and bad things happen. If we try to fix the code, we will be ["going against the grain"|http://dictionary.reference.com/browse/go+against+the+grain] and, as I describe above, possibly introduce other problems. I like your solution -- it's a nice piece of code -- but one of the responsibilities we have is to protect the API. To some degree, the code is like a fortress: users are crazy and there will always be users who are trying to climb in through the chimney instead of using the front door, "because the architecture is based around chimneys". We then have a choice: we can change things around so that the fire isn't burning in the fireplace, or say "NO, you may not enter through the chimney". What we then have to decide is: is it more important the the house stays warm, or was the chimney always actually more of an upstairs door? In this case, I really think that the house has to stay warm. There's no other way to keep it warm in the winter, as far as I know. I'm getting a little carried away with this metaphor.. ;) In other words, I think we have to make clear that having other components take control of the transaction is not something that we support and that if you *do* allow something else to take control of the transaction, then you're on your own. (At this point, we get into a discussion about the architecture of the human-task component, but that's another discussion -- one	1	1	https://issues.jboss.org/browse/JBPM-3814
13843	JBPM-3814	Hi Kris, can you give me some some sample code on how I can get a handle on the underlying command service and how to configure the lock? Thanks.	0	0	https://issues.jboss.org/browse/JBPM-3814
13844	JBPM-3814	Marco, what do you think about keep this configurable for 6.0 release? I mean, using reentrant locks could be configured by jBPM API in some way even affecting the performance at least for this 6.0 release. This is very important because many developers use Spring or EJBS with a couple of methods implicitly using transaction around, which calls ksession/jBPM API in some moment - so keeping it thread safety for these scenarios even with performance implication is better than let the whole context in in a non reliable state. Making this decision configurable, it will not impact in designs where the user transaction is not shared with the jBPM/session API, letting the session creating and managing the transaction/entity manager, without unnecessary locks.	1	0	https://issues.jboss.org/browse/JBPM-3814
13845	JBPM-3803	Does jBPM work with OpenJPA?	0	0	https://issues.jboss.org/browse/JBPM-3803
13846	JBPM-3799	What was in a *whole* lot of tests was the following code: {code:java} protected void tearDown() throws Exception { super.tearDown(); client.disconnect(); server.stop(); } {code} The problem is that super.tearDown() refers to BaseTest.tearDown(), which closes the entityManager and entityManagerFactory instances. Except that TaskServerHandler.messageReceived() calls taskSession.dispose() -- which closes these instances as well. This situation would not really be a problem in a production environment (it's internal code, after all), but it can mean that errors involving persistence and tx's can be thrown during tests (especially with hibernate 4). And those errors can be either alarming or confusing, when working on a complicated bug. Therefore, this bug, and the changes to reorder the above method so that it ends up being this: {code:java} protected void tearDown() throws Exception { client.disconnect(); server.stop(); super.tearDown(); } {code} Fortunately, there was a lot of duplicate code I could refactor into inherited methods that did this.	0	0	https://issues.jboss.org/browse/JBPM-3799
13847	JBPM-3774	It would be really useful to extend this to a) add other Out of box service implementations (e.g. Email? b) make this extensible so user could plug in their own implementations and reference them in the various design tools with a text value, e.g., Implementation = "foo" where foo is their custom workitem handler. This later capability would be instead of creating custom tasks with custom workitem handlers. This would allow for somewhat more portable BPMN2 models.	0	0	https://issues.jboss.org/browse/JBPM-3774

#	Requirement	Questions	Answered	Successfully Answered	URL
13848	JBPM-3769	hardcoded urls (admittedly this is not as bad as some of the other components): > grep -r localhost . ./org.drools.guvnor.Guvnor/hosted.html:var $hosted = "localhost:9997"; ./WEB-INF/classes/preferences.properties:designer.url=http://localhost:8080 ./WEB-INF/classes/preferences.properties:guvnor.url=http://localhost:8080/jboss-brms Is it unrealistic to centralise these for the configuration of BRMS / jBPM as a whole instead of diverse locations I have to modify to get it all working? I have sent a separate list of the same for the designer.war to Tiho with these: > grep -r localhost . ./defaults/customeditors.json: "Actors" : "http://localhost:8080/designer/customeditors/sampleactorseditor.html" ./defaults/formwidgets/backboneformsinclude.fw:<script src="http://localhost:8080/designer/lib/backbone-forms-min.js" charset="utf-8"></script> ./defaults/formwidgets/backboneformsinclude.fw:<script src="http://localhost:8080/designer/lib/backbone-forms-list-min.js" charset="utf-8"></script> ./defaults/formwidgets/backboneformsinclude.fw:<link href="http://localhost:8080/designer/lib/backbone-forms-default.css" rel="stylesheet" /> ./js/Core/main.js: * url:"http://localhost:8080/oryx/stencilsets/bpmn1.1/bpmn1.1.json" ./js/Plugins/bpel2bpmn.js: eRDF = '<?xml version="1.0" encoding="utf-8"?><html xmlns="http://www.w3.org/1999/xhtml" xmlns:b3mn="http://b3mn.org/2007/b3mn" xmlns:ext="http://b3mn.org/2007/ext" xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#" xmlns:atom="http://b3mn.org/2007/atom+xhtml"><head profile="http://purl.org/NET/erdf/profile"><link rel="schema.dc" href="http://purl.org/dc/elements/1.1/" /><link rel="schema.dcTerms" href="http://purl.org/dc/terms/ " /><link rel="schema.b3mn" href="http://b3mn.org" /><link rel="schema.oryx" href="http://oryx-editor.org/" /><link rel="schema.raziel" href="http://raziel.org/" /><base href="http://localhost:8080/backend/poem/new" /></head><body>' }); }indow.show();ndowue,Bpmn2Bpel.ApacheOdeUrlInputTitle,w').close();el_window').close();sport.responseText);Failed); ./js/Plugins/bpmn2pn.js: * getRdfUrl("http://localhost:8080/backend/poem/model/7/self") ./js/Plugins/bpmn2pn.js: * => "http://localhost:8080/backend/poem/model/7/rdf" ./js/Plugins/bpmn2pn.js: * getRdfUrl("http://localhost:8080/backend/poem/model/7/rdf" => "http://localhost:8080/backend/poem/model/7/rdf") ./xslt/BPMN2XForms.xslt: <xsl:variable name="test1" select="'http://localhost:8180/backend/poem/model/237/self#oryx_606560F1-BE0B-4130-81CC-F77202E2F27A'"/> ./xslt/BPMN2XForms.xslt: <xsl:variable name="test2" select="'http://localhost:8180/backend/poem/model/237/self#oryx_E6E357E0-6407-4221-9C4F-E4CE9F5FE42F'"/> ./xslt/BPMN2XForms.xslt: <xsl:variable name="test3" select="'http://localhost:8180/backend/poem/model/237/self#oryx_1A426703-F199-4123-BE6B-452A2900ADDF'"/> ./xslt/BPMN2XForms.xslt: <xsl:variable name="test4" select="'http://localhost:8180/backend/poem/model/237/self#oryx_249FD0F7-EEFF-4A08-ACB5-D3F23F1E0A42'"/> ./xslt/BPMN2XForms.xslt: <xsl:variable name="test5" select="'http://localhost:8180/backend/poem/model/237/self#oryx_24055331-44F5-483D-A612-6ABB2E609096'"/> ./xslt/BPMN2XHTML-extended.xslt:<!-- http://localhost:8180/backend/poem/model/205/self#oryx_F15AEAF4-3AB2-4107-91A9-CDB909AE175C --> ./xslt/BPMN2XHTML-extended.xslt:<g id="node1" class="node"><title>http://localhost:8180/backend/poem/model/205/self#oryx_F15AEAF4-3AB2-4107-91A9-CDB909AE175C</title> ./xslt/BPMN2XHTML-extended.xslt:<!-- http://localhost:8180/backend/poem/model/205/self#oryx_C6DF5A68-6719-47E5-BD8D-FA2DBD0C229F --> ./xslt/BPMN2XHTML-extended.xslt:<g id="node2"	1	0	https://issues.jboss.org/browse/JBPM-3769
13849	JBPM-3769	Thanks Eric, from your list there should be some Designer Jiras that come from it: 1) Remove unneeded and unused xslt files from Designer (that clears most of your list there) 2) Add ability to look up the actual host name in form widgets instead of using localhost by default 3) Add ability to use the actual host name in custom editors configuration rather than defaulting to localhost As far as this Jira goes (initial request): a) Designer does not require a login, so not sure where single signon would fit in? b) Add Jira asking for a common system variable to be used..I think that would be the easiest to do at this point. wdyt? Thanks.	0	0	https://issues.jboss.org/browse/JBPM-3769
13850	JBPM-3764	Hmm, 'difficult' you say? I'd say impossible. How should I to identify from a message like {code} Errors occured loading flows: [7,10]: [ERR 102] Line 7:10 mismatched input 'data' in rule "RuleFlow-Split-info.magnolia.workflow.SimpleWorkflow-10-9-DROOLS_DEFAULT" [0,0]: Parser returned a null Package ProcessLoadError: unable to parse xml : Exception class org.drools.RuntimeDroolsException : invalid package name ProcessLoadError: unable to parse xml : Exception class org.drools.RuntimeDroolsException : invalid package name 2013-04-11 16:23:35,237 ERROR org.drools.builder.impl.KnowledgeBuilderImpl : [7,10]: [ERR 102] Line 7:10 mismatched input 'data' in rule "RuleFlow-Split-info.magnolia.workflow.SimpleWorkflow-10-9-DROOLS_DEFAULT" 2013-04-11 16:23:35,237 ERROR org.drools.builder.impl.KnowledgeBuilderImpl : [0,0]: Parser returned a null Package 2013-04-11 16:23:35,237 ERROR org.drools.builder.impl.KnowledgeBuilderImpl : ProcessLoadError: unable to parse xml : Exception class org.drools.RuntimeDroolsException : invalid package name 2013-04-11 16:23:35,237 ERROR org.drools.builder.impl.KnowledgeBuilderImpl : ProcessLoadError: unable to parse xml : Exception class org.drools.RuntimeDroolsException : invalid package name {code} that I've forgot to switch one condition from "rule / mvel" to "code / java"?	0	0	https://issues.jboss.org/browse/JBPM-3764
13851	JBPM-3763	Jiri Svitak <jsvitak@redhat.com> made a comment on [bug 852044|https://bugzilla.redhat.com/show_bug.cgi?id=852044] Cannot view human task list on EAP 6.0 with JPA 2.0 and PostgreSQL due to: org.postgresql.util.PSQLException: Large Objects may not be used in auto-commit mode. How to reproduce: 1. Set up your environment using the steps in this document: https://docspace.corp.redhat.com/docs/DOC-111367 Ignore the parts about setting up Hibernate 3, instead of that use attached persistence.xml files with JPA 2.0. 2. Run the attached test case. It starts a process instance and connects to the human task server and creates a task there. 3. Open http://localhost:8080/business-central and view your tasks (click on refresh button) of your user (in my case it's jsvitak, but you can easily reconfigure it to krisv for example; in these files: testproc.bpmn2 file, brms-users.properties, brms-roles.properties and DefaultUsers.mvel). 4. You can see an exception in your server, see attached server.log for details. Business central shows no human tasks, but the task from the testproc should be there. With configured Hibernate 3 you can see the test task and complete it too. My environment: EAP 6.0 (with default Hibernate 4) BRMS 5.3.0 GA patch01 (reconfigure to JPA 2.0 jbpm-human-task.war and business-central-server.war) PostgreSQL 9 Added also relevant JIRA issue based on this discussion: https://community.jboss.org/thread/203384	0	0	https://issues.jboss.org/browse/JBPM-3763
13852	JBPM-3758	Kris, not sure how console could help with this, especially with creating new version of a process based on an existing one?? One thing that could be useful is to change the parsing mechanism for process - as soon as there is a version attribute specified then make a sort of composite id of the process automatically: process.id-version so it will be by default possible to host different versions of the process in same knowledge base. For starting such processes a new command could be implemented... Wdyt? Or I completely misunderstood this issue?	0	0	https://issues.jboss.org/browse/JBPM-3758
13853	JBPM-3758	Guys, this is a really important feature that must be supported on the knowledge base primarily. The great problem is that every object in the engine that wants to refer a process definition uses only its id (a String). What happens if I have multiple instances of a process with different versions? They point to the same processId but this is wrong.	1	1	https://issues.jboss.org/browse/JBPM-3758

#	Requirement	Questions	Answered	Successfully Answered	URL
13854	JBPM-3758	Kanchan Desai <kadesai@redhat.com> made a comment on [bug 917949|https://bugzilla.redhat.com/show_bug.cgi?id=917949] Hi Marek, Doc Text has been added. Can you please review the CCFR.	0	0	https://issues.jboss.org/browse/JBPM-3758
13855	JBPM-3741	Implementation for this is in Designer master now. Can you please test it and let me know if there is anything else that need to be done? Thanks.	1	1	https://issues.jboss.org/browse/JBPM-3741
13856	JBPM-3724	This is very nice +100. One thing that is often requested by community is to be able to execute commands against the session via the console REST api - would this also be a feature worth considering here as well?	0	0	https://issues.jboss.org/browse/JBPM-3724
13857	JBPM-3724	One thing that can be sort of reused is to add factories (I think we already have an jira issue for that) for it as it relies on different components (no session and no knowledge base there). And the "template" is web.xml in that case. Is that what you meant Marco? If not could you please elaborate bit more?	1	1	https://issues.jboss.org/browse/JBPM-3724
13858	JBPM-3724	short update on JNDI binding - it can be done in generic way by using JEE components to initialize it: - servlet - servlet context listener as AS 7.1.x will consider this as part of the deployment process of an application and will do all required preparations so there is no need for any JBoss AS specific code at all. Important to note here is that works for AS 7.1.x and does not work for AS 7.0.2 which we use by default. So here it is what we could do: 1. implement servlet context listener that will boostrap console (knowledge base and session based on given configuration) - console code 2. since boostrap is done in JEE managed object it will allow to bind session into JNDI - console code 3. move to JBoss AS 7.1.1 as default server to be able to utilize this feature - installer This mechanism is generic and supported by all app servers and servlet containers and provides us with additional advantage - any errors in console will be visible during server startup and not after first logon attempt. Wdyt?	1	1	https://issues.jboss.org/browse/JBPM-3724
13859	JBPM-3724	Maciej, yes that's what I meant. You wrote: "One thing that can be sort of reused is to add factories (I think we already have an jira issue for that) for it as it relies on different components (no session and no knowledge base there). And the "template" is web.xml in that case. Is that what you meant Marco? If not could you please elaborate bit more?" Good point in that not a lot of what will be used for console can be reused, because HT does different things -- but being able to reuse the general framework (and also have "A Way to Configure Components") is good. Thanks!	1	1	https://issues.jboss.org/browse/JBPM-3724
13860	JBPM-3723	Is there any progress on this issue ?	0	0	https://issues.jboss.org/browse/JBPM-3723
13861	JBPM-3723	Hello. I am still encountering this problem in jBPM 5.4.0 when persistence is enabled. Does your comment from 10/02/12 indicate otherwise? I see that the jira is still open so I wasn't sure... If I add a DefaultProcessEventListener to the ksession, override the afterProcessCompleted() method, and call kession.dispose() inside of that method using jBPM 5.4.0, I still receive: ERROR [org.drools.persistence.SingleSessionCommandService] (Thread-95) Could not commit session: org.jbpm.workflow.instance.WorkflowRuntimeException: ... Caused by: java.lang.IllegalStateException: Illegal method call. This session was previously disposed. at org.drools.reteoo.DisposedReteooWorkingMemory.getEnvironment(DisposedReteooWorkingMemory.java:286) [drools-core-5.5.0.Final.jar:5.5.0.Final] at org.drools.impl.StatefulKnowledgeSessionImpl.getEnvironment(StatefulKnowledgeSessionImpl.java:363) [drools-core-5.5.0.Final.jar:5.5.0.Final] at org.jbpm.persistence.processinstance.JPASignalManager.getProcessInstancesForEvent(JPASignalManager.java:39) [jbpm-persistence-jpa-5.4.0.Final.jar:5.4.0.Final] at org.jbpm.persistence.processinstance.JPASignalManager.signalEvent(JPASignalManager.java:23) [jbpm-persistence-jpa-5.4.0.Final.jar:5.4.0.Final] at org.jbpm.workflow.instance.impl.WorkflowProcessInstanceImpl.setState(WorkflowProcessInstanceImpl.java:264) [jbpm-flow-5.4.0.Final.jar:5.4.0.Final] at org.jbpm.workflow.instance.impl.WorkflowProcessInstanceImpl.setState(WorkflowProcessInstanceImpl.java:269) [jbpm-flow-5.4.0.Final.jar:5.4.0.Final] at org.jbpm.workflow.instance.impl.WorkflowProcessInstanceImpl.nodeInstanceCompleted(WorkflowProcessInstanceImpl.java:435) [jbpm-flow-5.4.0.Final.jar:5.4.0.Final] at org.jbpm.workflow.instance.node.EndNodeInstance.internalTrigger(EndNodeInstance.java:71) [jbpm-flow-5.4.0.Final.jar:5.4.0.Final] at org.jbpm.workflow.instance.impl.NodeInstanceImpl.trigger(NodeInstanceImpl.java:126) [jbpm-flow-5.4.0.Final.jar:5.4.0.Final] ... 30 more 14:45:17,665 WARN [com.arjuna.ats.jta] (Thread-95) ARJUNA016029: SynchronizationImple.afterCompletion - failed for org.drools.persistence.jta.JtaTransactionSynchronizationAdapter@5f7ebab1 with exception: java.lang.IllegalStateException: Illegal method call. This session was previously disposed. at org.drools.reteoo.DisposedReteooWorkingMemory.getProcessRuntime(DisposedReteooWorkingMemory.java:270) [drools-core-5.5.0.Final.jar:5.5.0.Final] at org.drools.impl.StatefulKnowledgeSessionImpl.getProcessRuntime(StatefulKnowledgeSessionImpl.java:874) [drools-core-5.5.0.Final.jar:5.5.0.Final] Can you please confirm whether this was fixed or not in version 5.4.0? Thanks	0	0	https://issues.jboss.org/browse/JBPM-3723
13862	JBPM-3707	Kristofer, Yes, it is true that jBPM doesn't support the full BPMN2 specification yet. I don't think any (open-source) vendor does at this point (as the specification defines a huge amount of constructs, where not all of them are as important or commonly used probably, see http://www.bpminstitute.org/resources/articles/bpms-watch-how-much-bpmn-do-you-need). The scope of the BPMN2 specification is also much broader than just executable business processes. So instead of trying to reach 100% conformance to the BPMN2 specification, we have selected (what we believe is) the most important subset and are constantly trying to extend our support for other constructs. Note that the BPMN2 specification has been defined by a lot of large players in the market, many of the participants can be considered BPM gurus. JBoss / Red Hat participated but can hardly be "blamed" for the fact that the BPMN2 specification did not try to implement each and every possible workflow pattern. As workflow patterns are probably similar to BPMN2 constructs when it comes to this: not all of them are as important or commonly used. And for some of the patterns, the BPMN2 specification might offer a suitable alternative. Could you explain why you believe it should be one of our important goals to implement all of these patterns? Regarding support from jBPM, note that the latest version of jBPM that was evaluated was jBPM 3.1.4 (Jan 2007!). So all the improvements we did over the last 6 years probably haven't been incorporated in that evaluation yet. Note that there have been tons of initiatives related to workflow languages in the past 10 years (WfMC, XPDL, WF patterns, BPEL, etc.), we simply can't support everything out there (but feel free to help out if you believe something is missing). Also note that we currently do support the local task data in jBPM5, as every task has a so-called workItem associated with it, which has parameters that are specific to that instance of the task. Kris	1	0	https://issues.jboss.org/browse/JBPM-3707

#	Requirement	Questions	Answered	Successfully Answered	URL
13863	JBPM-3707	Kristofer, Yes, I am aware of this is a possibility. Unfortunately, we haven't had the time to do something like this ourselves. We welcome anyone in the community who might be willing to help out here. Would you mind if I ask why you need this information? While the workflow patterns were a common approach to evaluate / compare the expressiveness of different workflow engine a while ago, I believe the breakthrough of the BPMN 2.0 specification has significantly changed this: most people now seem more interested in using a standard for describing their process definitions and therefore would like to see information about which BPMN 2.0 constructs are supported rather than which workflow patterns could be expressed using these constructs. Kris	1	0	https://issues.jboss.org/browse/JBPM-3707
13864	JBPM-3703	I will take a look on this issue.. it shouldn't be so hard to configure. As a user, do you want to know about the I18N feature always?	0	0	https://issues.jboss.org/browse/JBPM-3703
13865	JBPM-3700	eagerly anticipate this example. Can you send me a draft version for testing. many thanks.	0	0	https://issues.jboss.org/browse/JBPM-3700
13866	JBPM-3698	Who had try to use jBPM 5.3, JBoss7.0.2 on this issue?	0	0	https://issues.jboss.org/browse/JBPM-3698
13867	JBPM-3698	Tried with both 5.3.Final and master and both work fine. The report template file is located inside jbpm-gwt-console-server.war/WEB-INF/reports. Could you please check if you have that file there?	1	1	https://issues.jboss.org/browse/JBPM-3698
13868	JBPM-3698	kind of weird that it did not help. could you provide bit more information about your environment and how do you install jbpm and then upgrade this report core library to be able to try to reproduce the issue.	0	0	https://issues.jboss.org/browse/JBPM-3698
13869	JBPM-3698	just to double check, did you remove the previous version of report code from jbpm-gwt-console-server.war/WEB-INF/lib? since attached one has different version could be that you have both on classpath...	0	0	https://issues.jboss.org/browse/JBPM-3698
13870	JBPM-3678	Julian, which SessionFactory class do you mean? Could you perhaps provide the fully qualified name (package name + class name) of the class you're referring to? Thanks!	1	1	https://issues.jboss.org/browse/JBPM-3678
13871	JBPM-3671	My bad with previous misinformation about this case, in fact you can do this with installer environment, here is an example of how to get hold of actual owner of a task in a freemarker template: <#if task.taskData??> Actual owner is: ${task.taskData.actualOwner.id} </#if> So, I think most of the cases can be already supported as this provides access to both task and taskData information. With that, I leave that to you to decide if the issue should be in any way enhanced or can be closed.	1	0	https://issues.jboss.org/browse/JBPM-3671
13872	JBPM-3652	See https://bugzilla.redhat.com/show_bug.cgi?id=818936 Things to think about: - What should happen with the task? Should it be stored in the db, but without a user? Should the task be rejected? - Should an exception be thrown, and if so, where in the human-task logic? Is this a persistence problem or a human-task logic problem (with regards to where things are checked)?	1	1	https://issues.jboss.org/browse/JBPM-3652
13873	JBPM-3635	You should file this issue under Guvnor and not jBPM. Designer uses the Guvnor REST api for retrieving/storing process assets and currently this REST API is not capable of retrieving alternative representations of an asset - it only returns a single one which in case of business processes is the bpmn2. What would solve this issue is Guvnor REST api allowing for example: /rest/packages/defaultPackage/assets/myProcessAssetName/source?type=svg /rest/packages/defaultPackage/assets/myProcessAssetName/binary?type=png /rest/packages/defaultPackage/assets/myProcessAssetName/source?type=bpmn2 ... We can currently store the png of the process on the on-save operation, but what good does that do unless you can actually retrieve it? Thanks.	1	1	https://issues.jboss.org/browse/JBPM-3635
13874	JBPM-3635	You don't deploy a business process into Guvnor. Designer is a BPMN2 asset editor OF Guvnor. Guvnor initiates the saving of the process, not Designer. But either way this is kind of irrelevant because the requirement is the same - when storing the business process, the associated process image should be stored as well. The question is - should it stop there? Should we also re-generate all the process forms on on-save operation? Should we also generate the proces PDF, and store the SVG, and the process JSON? If you look at the big picture, what I described in the previous post is the solution (multiple representations of 1 asset in Guvnor) and not a partial solution that you are requesting.	1	1	https://issues.jboss.org/browse/JBPM-3635
13875	JBPM-3635	I'm still open to suggestions. Don't take my previous post as the conversation being "done" by any means :) Regenration of process image is high priority as it is used in monitoring. We can implement a solution for this independently of Guvnor now if that is requested here. WDYT?	1	1	https://issues.jboss.org/browse/JBPM-3635
13876	JBPM-3635	Tihomir, from your POV, seen in the Guvnor scenario this is completely true. From what I see, there are still a lot of people that, like me, are migrating from JBPM 3.x and prefer to use the old process design & deployment flow, with a standalone designer (not Guvnor dependent): 1) Model process in the designer (allowing also load/save from external files) 2) Deploy the definition (process + image) directly in the application, using a servlet or whatever. However it seems that JBoss abandoned this approach. What do you think?	1	1	https://issues.jboss.org/browse/JBPM-3635
13877	JBPM-3635	Do you have some advices/examples to integrate easily Guvnor (and designer) in my app (JSF based)?	1	1	https://issues.jboss.org/browse/JBPM-3635
13878	JBPM-3635	The version number in the bpmn2 definition (version attribute in the process properties) is not updated when you save the process and is not the same as the version number of the process asset in the Guvnor JCR repository (which does get incremented every time you save). Users who are modeling their processes might save many times but do not want their version in the process definition updated (should stay the same). Is this what we are talking about?	1	1	https://issues.jboss.org/browse/JBPM-3635
13879	JBPM-3635	What you say is correct Tihomir. I mean the image asset version number. It remains always 1, even if I update the process (e.g. version 2 of the BPMN). Should the image follow the same process versioning?	1	0	https://issues.jboss.org/browse/JBPM-3635
13880	JBPM-3635	It seems that SVG file is saved after the refresh of the designer and for this reason the bounding box it is not the minimum bounding box. So the SVG file contains big empty space. Is it possible to save an SVG with the minimum bounding box?	1	0	https://issues.jboss.org/browse/JBPM-3635
13881	JBPM-3635	As far I see there no 2.4.1 version to include this feature. Is the latest designer version compatible with Guvnor 5.5.0 Final? Is there a compatibility matrix between Drools and the designer?	0	0	https://issues.jboss.org/browse/JBPM-3635
13882	JBPM-3629	There's also the question of backwards compatibility for Oracle users: - If a Oracle user is already using jBPM 5. - and starts using a version in which the annotations have been updated to specify sequences (but with hbm2ddl=validate). Does the code/database continue to use HIBERNATE_SEQUENCE when updating entries? This needs to be confirmed or denied.	0	0	https://issues.jboss.org/browse/JBPM-3629
13883	JBPM-3606	As discussed briefly on IRC, maybe we could think about a single way of configuring Human Task server/service? In addition to persistence/transaction handling there are few more: 1. UserGroupCallback 2. UserInfo (for escalation and notification) 3. Escalation handler 4. loading users and groups 5. defining transport and its attributes For human task server deployed as web app we do that via init parameters on web.xml but it could be refactored to make use of the single mechanism for configuration. Wdyt?	1	1	https://issues.jboss.org/browse/JBPM-3606

#	Requirement	Questions	Answered	Successfully Answered	URL
13884	JBPM-3600	jbpm4.4 seems can not set business.day.expressed.in.hours, and jbpm3 can set day.expressed.in.hours to any hour in jbpm.business.calendar.properties as below: business.day.expressed.in.hours= 8 business.week.expressed.in.hours= 40 business.month.expressed.in.business.days= 21 business.year.expressed.in.business.days= 220 can you make jbpm 5.x can set that as jbpm 3?	0	0	https://issues.jboss.org/browse/JBPM-3600
13885	JBPM-3600	how to delete this empty comments?	0	0	https://issues.jboss.org/browse/JBPM-3600
13886	JBPM-3600	Maciej, can you explain better this step? "It is configurable by putting BusinessCalandar implementation into environment under jbpm.business.calendar key." Thank you!	0	0	https://issues.jboss.org/browse/JBPM-3600
13887	JBPM-3600	Hi Maciej, How can i specify time interval for work day ? e.g.: http://docs.jboss.com/jbpm/v3/userguide/businesscalendar.html weekday.monday= 9:00-12:00 & 12:30-17:00 ... This is an example of my businesscalendar properties file: business.hours.per.week = 6 business.hours.per.day = 36 business.start.hour = 8 business.end.hour = 15 business.weekend.days = 7 business.holiday.date.format = yyyy-DD-mm business.holidays = 2013-01-01,2013-05-01,2013-06-02,2013-06-13:2013-06-16,2013-08-15,2013-12-25 Could we have an example of a new "jbpm.business.calendar.properties" file Thanks a lot	1	0	https://issues.jboss.org/browse/JBPM-3600
13888	JBPM-3575	Tihomir Surdilovic <tsurdilo@redhat.com> made a comment on [bug 806428|https://bugzilla.redhat.com/show_bug.cgi?id=806428] Reading this thread I don't see prior discussions of this particular issue so not sure what the repeated "as i said" really mean or try to accomplish? Can you please link me the URL where I can download the ER build so can test locally and fix? Thanks.	0	0	https://issues.jboss.org/browse/JBPM-3575
13889	JBPM-3575	Tihomir Surdilovic <tsurdilo@redhat.com> made a comment on [bug 806428|https://bugzilla.redhat.com/show_bug.cgi?id=806428] Could I please get a response on the last comment? Thanks.	0	0	https://issues.jboss.org/browse/JBPM-3575
13890	JBPM-3575	Jiri Svitak <jsvitak@redhat.com> made a comment on [bug 806428|https://bugzilla.redhat.com/show_bug.cgi?id=806428] I have cleared my browser cache, but it did not help. Even if I forget on the attached business process and use another new process the problem persists. I cannot add any notification to human task node notifications. I have noticed that there are small red triangles for columns: Expires At, From, To Users, To Groups, Reply To, Subject What these small red triangles mean? Also I was not able to add body to notification in the bigger editor for message body. Returning back to the attached business process created by Maciej. Have you asked Maciej about the commas in the "Expires At" column? I can confirm that changing the existing two notifications is possible. Cannot add third notification there though. So it's not possible to add new notifications in new processes. Existing saved notifications seems to be possible to edit.	0	0	https://issues.jboss.org/browse/JBPM-3575
13891	JBPM-3575	Tihomir Surdilovic <tsurdilo@redhat.com> made a comment on [bug 806428|https://bugzilla.redhat.com/show_bug.cgi?id=806428] The red triangle means that that column cannot be blank or that the value has changed. Can you please give me steps to reproduce this issue? Neither Maciej nor I are able to reproduce on ER4. Currently you cannot leave to users , to groups, subject and body as blank. if you do designer will ignore your entry (we should probably show an alert message here for users rather than just ignore - please open a BZ for that).	0	0	https://issues.jboss.org/browse/JBPM-3575
13892	JBPM-3575	Jiri Svitak <jsvitak@redhat.com> made a comment on [bug 806428|https://bugzilla.redhat.com/show_bug.cgi?id=806428] Hello, I will write again the steps to reproduce the problem: 1.) Download brms-p-5.3.1.ER4-deployable-ee6.zip from http://jawa05.englab.brq.redhat.com/candidate/BRMS-5.3.1-ER4/. MD5 hash and name of the file: 225bf1b40a1190acfa3267f4c7c2c10b brms-p-5.3.1.ER4-deployable-ee6.zip 2.) Download EAP 6 from here: http://download.eng.brq.redhat.com/released/JBEAP-6/6.0.0/zip/ 3.) I have deployed the brms 5.3.1 ER4 ee6 into EAP6, you can use this guide for that (default human task users and groups, default console users and groups, persistence using MySQL 5.5): https://docspace.corp.redhat.com/docs/DOC-111367 4.) Start the server in standalone mode: jboss-eap-6.0/bin/standalone.sh 5.) Create .dodeploy files for all jBPM .war directories in jboss-eap-6.0/standalone/deployments. 6.) Open Firefox 12.0, click on Tools->Clear Recent History, check Cache to be deletede, time range Everything. 7.) Open new tab and type the address: http://localhost:8080/jboss-brms/ 8.) Log in as user krisv, password krisv. 9.) Click on Knowledge Bases on the left panel. 10.) Click on Create New->New Package. 11.) Fill in the name of the package com.sample and click Create package. 12.) Click on Create New->New BPMN2 Process. 13.) Fill in the name HumanTask and Create in Package com.sample. 14.) HumanTask process is opened in Designer. Click on the left panel and place on the canvas ordinary StartEvent. Add a human task. Add an end event. 15.) Click on the human task. Fill in Name: Human Task, TaskType: User, Actors: krisv, Task Name: HumanTask. 16.) Click on PNG on the bottom panel to generate process image. Generate task forms for human tasks. Click on Validate button, all should validate without problems. 17.) Click on File->Save changes->Check in. 18.) Click again on HumanTask, open right panel and click on Notifications. 19.) Fill in type: Not Started, Expises At: 10s, From: krisv, To Users: mary, To Groups: PM, Subject: subject. 20.) Try to fill Body, the editor is opened. Fill in "body" and click on OK. String "body" is not saved. 21.) Try to save the rest of the notification with filled fields Type, From, To Users, To Groups, Subject. 22.) The notification did not save. NEW FINDINGS: However later I have noticed another editor for the Body. I think that this is the core problem. Two editors for body of the notification have been opened. After I closed the second body editor, which was hidden behind other dialogs, I was able to add a notification on second attempt without problems. Why there are two Body editors opened? It prevents from saving the body of the notification and then it prevents from saving the notification itself. After the first unsuccessful attempt I was able to add new notifications without problems. I see two following problem now: user friendliness - user doesn't know which fields are mandatory, no warnings or validations are available - two editors for Body notification can be opened and user can easily miss this fact and thinks that it's impossible to add a notification POSSIBLE SOLUTION: I think we can agree that there is no time for implementing the validation or program help for the mandatory parameters of notifications. Regarding the second problem - if you have two editors for "Body" opened, please try to fix it - that only one correct editor is opened, which is able to save the body of the notification. If you cannot see this problem, then don't fix it and we'll leave this bug.	0	0	https://issues.jboss.org/browse/JBPM-3575
13893	JBPM-3575	jgargula@redhat.com made a comment on [bug 867819|https://bugzilla.redhat.com/show_bug.cgi?id=867819] After another playing with notifications I've found that there is a problematic field "Reply To" on Notifications tab. I'm not sure how it exactly works but this field lost its value after closing the dialog by pressing OK. How to reproduce: 1. Fill all fields and click "Add" 2. Check filled values by clicking on just saved notification (everything seems to be all right) 3. Click OK and dialog disappears. 4. Double click on the user task and show notifications tab again. 5. After clicking on the previously saved notification there is missing value of "Reply To" field. I noticed that if you include @ sign in Reply To then the situation is different and it influences several fields in this form. They lost theirs values but these values can be seen in some other field.	0	0	https://issues.jboss.org/browse/JBPM-3575

#	Requirement	Questions	Answered	Successfully Answered	URL
13894	JBPM-3543	Rajesh Rajasekaran <rrajasek@redhat.com> made a comment on [bug 870424|https://bugzilla.redhat.com/show_bug.cgi?id=870424] This issue was agreed as a blocker, and has pm_ack. If we can commit to fixing this in the product branch, I will set the release flag for 5.3.1. Kris, can you comment?	1	1	https://issues.jboss.org/browse/JBPM-3543
13895	JBPM-3528	This is odd... The domain of values for MultiInstanceLoopCharacteristics.inputDataItem is constructed from the list of ioSpecification.dataInputs but as far as I can tell, neither the SubProcess nor Process have an ioSpecification. Where should these values come from?	1	1	https://issues.jboss.org/browse/JBPM-3528
13896	JBPM-3512	Meciej, I see you are the new owner of this issue. Is there an ETA for getting this issue resolved? Thanks! :)	1	1	https://issues.jboss.org/browse/JBPM-3512
13897	JBPM-3481	Not sure what you mean, these methods are available in the TaskClient class and are implemented. While a UI might decide to only offer a subset of the available methods, removing these methods from the implementation while they could be used, for example programmatically or by using another task UI that does support these operations. Could you elaborate on what you would like to see removed where and why?	1	1	https://issues.jboss.org/browse/JBPM-3481
13898	JBPM-3481	These methods should be fully implemented, and there should be junit test for most of these as well. What part of the code makes you think those operations are not supported?	1	1	https://issues.jboss.org/browse/JBPM-3481
13899	JBPM-3481	You mean the response handlers in TaskClientHandler? That is because there is one generic response handler for all task life cycle methods like claim, assign, etc.: TaskOperationResponseHandler. Or are you referring to something else?	1	1	https://issues.jboss.org/browse/JBPM-3481
13900	JBPM-3481	I see that, but also only see code in code in TaskServiceSession, for example, to handle Claim, Complete, Fail, Skip, Remove and Register. What am I missing here?	1	1	https://issues.jboss.org/browse/JBPM-3481
13901	JBPM-3445	Looked first at how drools-persistence-jpa handles a DB2 (97) database. The error that's thrown in about half of those tests is "SQL Error: -302, SQLState: 22001". On the following forum post, it's explained that this is because a value that's larger than the field in the db is being persisted: - http://www.dbforums.com/db2/1652962-getting-error-db2-sql-error-sqlcode-302-sqlstate-22001-sqlerrmc-null.html#post6445333 In this case, it's the lob in SessionInfo: the value being saved is larger than the (TINYBLOB) field. {noformat} WARN hibernate.util.JDBCExceptionReporter.logExceptions - SQL Error: -302, SQLState: 22001 ERROR hibernate.util.JDBCExceptionReporter.logExceptions - DB2 SQL Error: SQLCODE=-302, SQLSTATE=22001, SQLERRMC=null, DRIVER=3.61.75 ERROR event.def.AbstractFlushingEventListener.performExecutions - Could not synchronize database state with session org.hibernate.exception.DataException: could not update: [org.drools.persistence.info.SessionInfo#201] at org.hibernate.exception.SQLStateConverter.convert(SQLStateConverter.java:100) at org.hibernate.exception.JDBCExceptionHelper.convert(JDBCExceptionHelper.java:66) at org.hibernate.persister.entity.AbstractEntityPersister.update(AbstractEntityPersister.java:2443) ... at org.hibernate.ejb.AbstractEntityManagerImpl$1.beforeCompletion(AbstractEntityManagerImpl.java:504) at bitronix.tm.BitronixTransaction.fireBeforeCompletionEvent(BitronixTransaction.java:397) at bitronix.tm.BitronixTransaction.commit(BitronixTransaction.java:147) at bitronix.tm.BitronixTransactionManager.commit(BitronixTransactionManager.java:95) at org.drools.persistence.session.JpaPersistentStatefulSessionTest.testUserTransactions(JpaPersistentStatefulSessionTest.java:213) ... Caused by: com.ibm.db2.jcc.am.SqlDataException: DB2 SQL Error: SQLCODE=-302, SQLSTATE=22001, SQLERRMC=null, DRIVER=3.61.75 at com.ibm.db2.jcc.am.ed.a(ed.java:668) ... at org.hibernate.persister.entity.AbstractEntityPersister.update(AbstractEntityPersister.java:2425) ... {noformat}	0	0	https://issues.jboss.org/browse/JBPM-3445
13902	JBPM-3445	We can override the JPA anno setup by having an XML mapping file that overrides the appropriate fields (and adds a @Column(length=...)). There are 2 things to do then: # making sure this happens for testing? # saving the mapping file and documenting it that users know where it is and how to use it {noformat} <mriet> The problem is that, if you just have a @Lob anno on the field of an entity (and no other anno's), it looks like DB2 then automatically translates that to a "TINYBLOB" field in the DB. <mriet> A TINYBLOB field is only 255 long though (bytes?) <mriet> so, when we have a serialized session, or workItem info or proc info, etc. that's longer, BOOM: SQL -302/22201. <mriet> One way of fixing this is adding @Column(length=<long enough>) to the same @Lob field. <mriet> but what's long enough? <mriet> max (64k/65k depending on def of k)? But are we then wasting space? <mriet> and how do other db's relate to this? <krisv> mriet, if it works for other databases as well, that's fine, if not, can you override the annotations for example in an xml config file for db2? <mriet> oh yeah -- mapping-file overrides in the xml for DB2. {noformat}	0	0	https://issues.jboss.org/browse/JBPM-3445
13903	JBPM-3443	Has this been resolved? It seems like it..	0	0	https://issues.jboss.org/browse/JBPM-3443
13904	JBPM-3428	Thanks Kris for the suggestion. It works as long as I keep the action node as an script task. When I switched back to use to the Action interface approach, I got the following exception, We would like to use the Action interface approach so that developer can debug into the action node in the flow. Otherwise, he/she has to rely on println, which is not ideal. So, is it possible for us to using the XOR splitnode + Action node interface + internal API to register the flow (aka without xml tranformation)? Exception in thread "main" java.lang.ClassCastException: org.jbpm.workflow.core.DroolsAction cannot be cast to org.jbpm.workflow.core.impl.DroolsConsequenceAction at org.jbpm.process.builder.ActionNodeBuilder.build(ActionNodeBuilder.java:21) at org.jbpm.compiler.ProcessBuilderImpl.processNodes(ProcessBuilderImpl.java:215) at org.jbpm.compiler.ProcessBuilderImpl.buildNodes(ProcessBuilderImpl.java:201) at org.jbpm.compiler.ProcessBuilderImpl.buildProcess(ProcessBuilderImpl.java:165) at com.sample.SplitnodeTest.registerProcess(SplitnodeTest.java:34) at com.sample.SplitnodeTest.main(SplitnodeTest.java:147)	1	0	https://issues.jboss.org/browse/JBPM-3428

#	Requirement	Questions	Answered	Successfully Answered	URL
13905	JBPM-3428	are the *NodeBuilder primarily needed for script tasks? Or, is this ActionNode specific? <?xml version="1.0" encoding="UTF-8"?> <definitions id="Definition" targetNamespace="http://www.jboss.org/drools" typeLanguage="http://www.java.com/javaTypes" expressionLanguage="http://www.mvel.org/2.0" xmlns="http://www.omg.org/spec/BPMN/20100524/MODEL" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://www.omg.org/spec/BPMN/20100524/MODEL BPMN20.xsd" xmlns:g="http://www.jboss.org/drools/flow/gpd" xmlns:tns="http://www.jboss.org/drools"> <itemDefinition id="_isTrueItem" structureRef="java.lang.Boolean" /> <process processType="Private" isExecutable="true" id="splitnode.test.process" name="splitnode.test.process" tns:packageName="splitnode.test.process" > <!-- process variables --> <property id="isTrue" itemSubjectRef="_isTrueItem"/> <!-- nodes --> <scriptTask id="_305" name="True path" scriptFormat="http://www.java.com/java" > <script>System.out.println("We are on TRUE branch...");</script> </scriptTask> <scriptTask id="_1000" name="isTrueActionNode" /> <startEvent id="_1" name="Starting..." /> <scriptTask id="_2000" name="False path" scriptFormat="http://www.java.com/java" > <script>System.out.println("We are on FALSE branch...");</script> </scriptTask> <exclusiveGateway id="_2001" gatewayDirection="Converging" /> <endEvent id="_40" name="End" /> <exclusiveGateway id="_1010" name="splitOnIsTrue" gatewayDirection="Diverging" /> <!-- connections --> <sequenceFlow id="_1010-_305" sourceRef="_1010" targetRef="_305" name="is True?" > <conditionExpression xsi:type="tFormalExpression" language="http://www.java.com/java" >return isTrue;</conditionExpression> </sequenceFlow> <sequenceFlow id="_1-_1000" sourceRef="_1" targetRef="_1000" /> <sequenceFlow id="_1010-_2000" sourceRef="_1010" targetRef="_2000" name="Not true" > <conditionExpression xsi:type="tFormalExpression" language="http://www.java.com/java" >return !isTrue;</conditionExpression> </sequenceFlow> <sequenceFlow id="_305-_2001" sourceRef="_305" targetRef="_2001" /> <sequenceFlow id="_2000-_2001" sourceRef="_2000" targetRef="_2001" /> <sequenceFlow id="_2001-_40" sourceRef="_2001" targetRef="_40" /> <sequenceFlow id="_1000-_1010" sourceRef="_1000" targetRef="_1010" /> </process> </definitions> Test XOR split node... variable 'isTrue' is set to TRUE Exception in thread "main" java.lang.RuntimeException: unable to execute Action at org.jbpm.workflow.instance.node.ActionNodeInstance.internalTrigger(ActionNodeInstance.java:49) at org.jbpm.workflow.instance.impl.NodeInstanceImpl.trigger(NodeInstanceImpl.java:122) at org.jbpm.workflow.instance.impl.NodeInstanceImpl.triggerConnection(NodeInstanceImpl.java:185) at org.jbpm.workflow.instance.node.SplitInstance.internalTrigger(SplitInstance.java:100) at org.jbpm.workflow.instance.impl.NodeInstanceImpl.trigger(NodeInstanceImpl.java:122) at org.jbpm.workflow.instance.impl.NodeInstanceImpl.triggerConnection(NodeInstanceImpl.java:185) at org.jbpm.workflow.instance.impl.NodeInstanceImpl.triggerCompleted(NodeInstanceImpl.java:150) at org.jbpm.workflow.instance.node.ActionNodeInstance.triggerCompleted(ActionNodeInstance.java:55) at org.jbpm.workflow.instance.node.ActionNodeInstance.internalTrigger(ActionNodeInstance.java:51) at org.jbpm.workflow.instance.impl.NodeInstanceImpl.trigger(NodeInstanceImpl.java:122) at	0	0	https://issues.jboss.org/browse/JBPM-3428
13906	JBPM-3378	Marco, could you comment on what the current issues are if you try to compile with JDK5? Are they code issues, build issues, or something else?	1	0	https://issues.jboss.org/browse/JBPM-3378
13907	JBPM-3366	example mriet: do you have a jira? {noformat}{noformat} tsurdilo: sorry i was wrong - https://issues.jboss.org/browse/JBPM-3366 tsurdilo: there is also ability to execute a process from designer tsurdilo: thats a bigger change tsurdilo: idk how to do that yet..but woudl be cool :) {noformat}{noformat} tsurdilo: well, it would be great to be able test stuff tsurdilo: and just wanted to provide some means for the user to test their proces they are creating mriet: makes sense! tsurdilo: so..either we go with creating a java junit test (cut/paste into eclipse and it runs) tsurdilo: or..we go the big route and integrate with jbpm-console from designer (via rest api) to actually run the process tsurdilo: ... or have both :) tsurdilo: ..or something completely different that's better {noformat}{noformat} tsurdilo: about https://issues.jboss.org/browse/JBPM-3366 tsurdilo: one thing i was thinking is that tsurdilo: i would create a maven project for the user tsurdilo: that they can just run mvn test and thats it tsurdilo: because if it's just a .java file source for a junit test tsurdilo: user would have to set up eclipse, jbpm plugin, etc tsurdilo: i want to be able to execute the test outside of eclipse cause eclipse is a competitor to designer :) :) mriet: right -- why not skip all that maven/etc. stuff and just have the user execute the test _from within designer_? tsurdilo: even better ! tsurdilo: designer already has depends on jbpm and drools jars tsurdilo: not all drools jars tho tsurdilo: so maybe we need more ? mriet: maybe, first we have to figure out exactly how we want to do that: tsurdilo: yap ill leave that up to you :) mriet: sure -- but you're welcome to "think with me", as we say in dutch. tsurdilo: nice will do tsurdilo: idk how to do assertions tsurdilo: like what's expected and what's not mriet: in any case, I'll post my thoughts to the jira the coming days -- feel free to reply or whatever. {noformat}{noformat} mriet: bpm test framework! tsurdilo: its ok to start really small and then expand if it gets any tracktion tsurdilo: you have access to the entire bpmn2 of the process at any time mriet: hehe.. Lol. yeah -- iterative development is definitely easier than big bang. ;D tsurdilo: so you can inspect the process, find out what process vars are defined, user task inputs/outputs etc mriet: right, so thinking out loud, we can divide this into the following parts: mriet: 1. running the process from the designer mriet: 2. testing the process mriet: and 1. is then. mriet: exporting the process (jar? fork a thread? who knows?) mriet: a. ^^ mriet: b. starting the process mriet: c. integration with designer so that the process can be "seen" by user? mriet: d. returning end state of proces (same as c?) mriet: does that make sense? tsurdilo: yap, from what you said i really like the "all-in-designer" testing process tsurdilo: it would not be that hard to create a .java file and execute it in designer tsurdilo: the problem is with workitems and human tasks ..how to do those {noformat}{noformat} mriet: re workItems and human tasks, there are 3 options: mriet: 1. write the java code in designer... :(mriet: 2. have an environment or jars that can be imported mriet: 3. mocking -- in conjunction with the "bpmn test framework"" tsurdilo: ok. maybe just even so simple that we just tell users "paste the expected output (if any) of workitem" mriet: Exactly!!! mriet: the simpler for the users, the better.. tsurdilo: for human task the same really tsurdilo: timers are another thing.. tsurdilo: so its not simple thing tsurdilo: so how to take something not simple and make it simple ? :) mriet: those are the really hard problems. :) tsurdilo: another thing we could do tsurdilo: is instead of trying to run inside designer tsurdilo: have the designer interact via the console rest api mriet: re timers: I think it's probably best to get to that once we get 1. running the process and 2. testing the process done.. 3. will be timers and that kind of stuff. tsurdilo: and just run the process test in console? tsurdilo: idk what would be best :) mriet: yeah, that makes sense. tsurdilo: but that introduces also issues like how do you do junit asserts ? {noformat}{noformat} mriet: when I say "inside" the designer, the only thing I mean (at this point) is simply being able to run the process from the designer -- from the users point of view. tsurdilo: yap that would be already a big thing mriet: so, the whole junit framework is reflection based (btw me == big	1	0	https://issues.jboss.org/browse/JBPM-3366

#	Requirement	Questions	Answered	Successfully Answered	URL
13908	JBPM-3366	+JUnit tests+: Indeed, this is _not_ about JUnit tests or (as far as I'm concerned) any sort of result that only a developer would use. +jBPM console+: integrating with the console is definitely a good idea. But because this has to do being able to test processes as you design them (and you can't design processes in the console), part of this has to happen in the designer, I think. +Simulation/replay+: perfect! Where is this being done? This almost 50% of the work for this idea. I would very much like to start looking at how to integrate what's being done. A lot of this idea has to do more with being able to create a process and then _immediately_ run it, simulate running it, and/or test it against some sort of system of constraints. This Jira is primarily for being able to start/simulate the process out of the designer -- and maybe that starts up a console within the designer, who knows? Lastly, depending on how the simulation/replay stuff works, it would be interesting to be able to write "tests" for a bpmn process. For example, the user can fill in an "assert" to check if the process does this, or that, when a or b happen.	1	0	https://issues.jboss.org/browse/JBPM-3366
13909	JBPM-3340	Thanks ur reply. This is my task "web-based designer based on Oryx that can export jPDL". Thanks great. So, as per you said jBPM 4.x releases is possible. Can you tell me where can get the source code(Oryx repository detail). Also how to download jBPM 4.x	0	0	https://issues.jboss.org/browse/JBPM-3340
13910	JBPM-3308	The build (mssql2008R2) is broken. - It also kind of looks like the MSSQL 2008 R2 database is missing functionality: {code} javax.transaction.xa.XAException: com.microsoft.sqlserver.jdbc.SQLServerException: Failed to create the XA control connection. Error: "Could not find stored procedure 'master..xp_sqljdbc_xa_init_ex'." {code} This is due to the MSSQL database does not have XA transaction functionality installed. Working on this..	0	0	https://issues.jboss.org/browse/JBPM-3308
13911	JBPM-3308	There seem to be some minor problems with creating tables in MS SQL: {noformat} 17/02 11:33:53,296[main] ERROR tool.hbm2ddl.SchemaUpdate.execute - Unsuccessful: create table Notification_email_header (Notification_id numeric(19,0) not null, emailHeaders_id numeric(19,0) not null, mapkey varchar(255) null, primary key (Notification_id, mapkey), unique (emailHeaders_id)) 17/02 11:33:53,296[main] ERROR tool.hbm2ddl.SchemaUpdate.execute - Cannot define PRIMARY KEY constraint on nullable column in table 'Notification_email_header'. 17/02 11:33:53,296[main] ERROR tool.hbm2ddl.SchemaUpdate.execute - Unsuccessful: alter table Notification_email_header add constraint FKF30FE3441F7B912A foreign key (emailHeaders_id) references email_header 17/02 11:33:53,296[main] ERROR tool.hbm2ddl.SchemaUpdate.execute - Cannot find the object "Notification_email_header" because it does not exist or you do not have permissions. 17/02 11:33:53,296[main] ERROR tool.hbm2ddl.SchemaUpdate.execute - Unsuccessful: alter table Notification_email_header add constraint FKF30FE34430BE501C foreign key (Notification_id) references Notification 17/02 11:33:53,296[main] ERROR tool.hbm2ddl.SchemaUpdate.execute - Cannot find the object "Notification_email_header" because it does not exist or you do not have permissions. {noformat}	0	0	https://issues.jboss.org/browse/JBPM-3308
13912	JBPM-3249	Ok... So I'm looking at this in two phases... The Service Task object in the BPMN2 spec really doesn't have much to it, but the jbpm5 version does, so this will definitely have to be split out into the jbpm5 runtime area: <itemDefinition id="_2_InMessageType" structureRef="java.lang.String" /> <message id="_2_InMessage" itemRef="_2_InMessageType" /> <interface id="_2_ServiceInterface" name="org.jbpm.examples.junit.HelloService"> <operation id="_2_ServiceOperation" name="hello"> <inMessageRef>_2_InMessage</inMessageRef> </operation> </interface> <process processType="Private" isExecutable="true" id="ServiceProcess" name="Service Process" > <property id="s" itemSubjectRef="_sItem"/> <serviceTask id="_2" name="Hello" operationRef="_2_ServiceOperation" implementation="Other" > <ioSpecification> <dataInput id="_2_param" name="Parameter" /> <dataOutput id="_2_result" name="Result" /> <inputSet> <dataInputRefs>_2_param</dataInputRefs> </inputSet> <outputSet> <dataOutputRefs>_2_result</dataOutputRefs> </outputSet> </ioSpecification> <dataInputAssociation> <sourceRef>s</sourceRef> <targetRef>_2_param</targetRef> </dataInputAssociation> <dataOutputAssociation> <sourceRef>_2_result</sourceRef> <targetRef>s</targetRef> </dataOutputAssociation> </serviceTask> ... </process> So we need the ability to: 1) Define a property (is it something that should be defined in the Palette? Or something simply maintained in the context of the bpmn file?) 2) Define an item definition, message, interface, and operation so it can be referenced by the custom service task. These seem to be outside the process element in the file, so would need to be added somehow... 3) Figure out why the implementation of the inputs & outputs is so convoluted... a) You define an individual parameter (incoming) and an individual result (outgoing) b) Then you define the set of input parameters (the inputSet) and a set of output parameters (the outputSet) c) Then you associate the property "s" with the inputs and again with the outputs. Is there a reason it's not handled in a single pair of elements instead of in three separate chunks? And this is just from one example (BPMN2 - ServiceProcess.bpmn2 from the jbpm5 examples). What other options/values are available here or is it just a Java service interface that's available at the moment?	1	1	https://issues.jboss.org/browse/JBPM-3249
13913	JBPM-3248	Do we have any idea what properties need to be made available to the user for this one? An example? A spec?	1	1	https://issues.jboss.org/browse/JBPM-3248
13914	JBPM-3248	So this screen shot is of the Constraint editor dialog for a diverging gateway in the Drools-BPMN editor (BPMN2 Process Editor vs. the new BPMN2 Visual Editor). This entire dialog is Drools runtime specific. So here lies another place where we need the ability to hook in runtime-specific components via some sort of extension mechanism to allow for other runtimes to do the same thing. Can we assume that all BPMN2 runtimes will be able to specify constraints on a gateway? The sequenceFlow/conditionExpression seems to map to a connection/expression or connection/formal expression in the BPMN2 Model Editor... I'm guessing that's what we need to maintain.	1	1	https://issues.jboss.org/browse/JBPM-3248
13915	JBPM-3223	Thank you, Tihomir, for doing this! Are you also doing this for tests? A general log policy with regards to tests would be nice. Something like: - info for important info - debug for more detailed - trace for exact system output The output from the BPMN tests starts to drive me nuts if I end up running mvn clean install too many times.. argh.	0	0	https://issues.jboss.org/browse/JBPM-3223
13916	JBPM-3199	Can you copy paste the query or source code line ? I got some JPA-QL experience and might be able to whip up fast JPA-QL instead.	0	0	https://issues.jboss.org/browse/JBPM-3199
13917	JBPM-3181	What do I need to to do pass this patch into the service repo?	0	0	https://issues.jboss.org/browse/JBPM-3181
13918	JBPM-3172	Hey Gary - as a BPMN2 n00b, can you put together a sample or more concrete reproductive steps so I can better understand this one?	1	1	https://issues.jboss.org/browse/JBPM-3172
13919	JBPM-3172	I believe this has been resolved. Gary, can you verify? If this needs more work, please comment in bugzilla. Thanks!	1	1	https://issues.jboss.org/browse/JBPM-3172
13920	JBPM-3170	Is the solution to this problem available to look at? We have a use case which requires this feature and I'm reluctant to build something custom if there is something that has already been implemented.	1	0	https://issues.jboss.org/browse/JBPM-3170
13921	JBPM-3170	Esteban, When a timer, fires, it sends a signal that is then picked up by the process instance that is waiting for that timer. This signal is sent through the command service, which synchronizes and makes sure multiple threads are not executing simultaneously. Do you have a test case that shows this issue? If so, could you open a new JIRA for this and attach the issue?	1	0	https://issues.jboss.org/browse/JBPM-3170

#	Requirement	Questions	Answered	Successfully Answered	URL
13922	JBPM-3170	In our case the session is disposed. We have up coming use cases where we need to utilise timers. "In more complex set-ups, where multiple sessions are used and sessions could be disposed in between, a centralised session could be used to manage the timers when the other sessions are disposed." --> Is there a timeframe for when this might be available? or is there any early access to the code? We need to implement timers and I'd rather look a consistent solution so when this becomes available we can remove our copy.	1	0	https://issues.jboss.org/browse/JBPM-3170
13923	JBPM-3170	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 786402|https://bugzilla.redhat.com/show_bug.cgi?id=786402] Kris Verlaenen <kris.verlaenen@gmail.com> made a comment on jira JBPM-3170 Esteban, When a timer, fires, it sends a signal that is then picked up by the process instance that is waiting for that timer. This signal is sent through the command service, which synchronizes and makes sure multiple threads are not executing simultaneously. Do you have a test case that shows this issue? If so, could you open a new JIRA for this and attach the issue?	0	0	https://issues.jboss.org/browse/JBPM-3170
13924	JBPM-3170	JBoss JIRA Server <jira-update@redhat.com> made a comment on [bug 786402|https://bugzilla.redhat.com/show_bug.cgi?id=786402] Amin Mohammed-Coleman <aminmc@gmail.com> made a comment on jira JBPM-3170 In our case the session is disposed. We have up coming use cases where we need to utilise timers. "In more complex set-ups, where multiple sessions are used and sessions could be disposed in between, a centralised session could be used to manage the timers when the other sessions are disposed." --> Is there a timeframe for when this might be available? or is there any early access to the code? We need to implement timers and I'd rather look a consistent solution so when this becomes available we can remove our copy.	0	0	https://issues.jboss.org/browse/JBPM-3170
13925	JBPM-3170	Kris Verlaenen <kverlaen@redhat.com> made a comment on [bug 786402|https://bugzilla.redhat.com/show_bug.cgi?id=786402] Lukas, can I change the status of this issue to CLOSED - NOT_A_BUG?	0	0	https://issues.jboss.org/browse/JBPM-3170
13926	JBPM-3170	Lukáš Petrovický <lpetrovi@redhat.com> made a comment on [bug 786402|https://bugzilla.redhat.com/show_bug.cgi?id=786402] (In reply to comment #10) > Lukas, can I change the status of this issue to CLOSED - NOT_A_BUG? Sure. Even better, I'll do it for you. (By the way - I think that this used to be an issue in the early days of the release. It must have been fixed somewhere along the road.)	0	0	https://issues.jboss.org/browse/JBPM-3170
13927	JBPM-3138	It should be possible to define simple expressions like that inside #{}. For example, the following test uses #{Person.name} where Person is a process variable. https://github.com/krisv/jbpm/blob/master/jbpm-flow-builder/src/test/java/org/jbpm/integrationtests/ProcessWorkItemTest.java Did you define myData as a variable and is the method param defined? what type is myData in this case? Your expression should follow MVEL syntax.	1	0	https://issues.jboss.org/browse/JBPM-3138
13928	JBPM-3093	I think this is fixed. Gary, can you verify & close?	1	1	https://issues.jboss.org/browse/JBPM-3093
13929	JBPM-3064	I've tried to apply the patch but I end up with several compile time errors after doing so. It is possible that more recent changes have made this patch invalid. Would it be possible to update from git again and regenerate the patch?	0	0	https://issues.jboss.org/browse/JBPM-3064
13930	JBPM-3064	never mind found it. I can see that it's been a while since last comment here but I think it was not applied. Can I find mentioned pull request somewhere?	0	0	https://issues.jboss.org/browse/JBPM-3064
13931	JBPM-3063	Could you clarify how you are editing the action script? Because the eclipse editor should support autocompletion for actions.	0	0	https://issues.jboss.org/browse/JBPM-3063
13932	JBPM-3055	Jiangsiyu, Which database are you using? (Oracle? Sybase? DB2?) Please also upgrade to the latest version of jbpm (see the jboss nexus/repository). Thanks, Marco	1	0	https://issues.jboss.org/browse/JBPM-3055
13933	JBPM-3048	So, what's been done is the following: - ProcessInstanceInfo no longer contains _any_ Hibernate *or* JPA annotations. - When using jBPM with Hibernate 3.x, we make sure to add a line that imports a hibernate mapping file in the persistence.xml. This takes care of correctly mapping the ProcessInstanceInfo to a database entity. - There is now a "META-INF/JBPMorm-JPA2.xml" file available with the jbpm-persistence-jpa.jar. This file is of course an entity-mapping file for use with JPA2 and includes a JPA2 mapping for the ProcessInstanceInfo class. - You can add the following line to your persistence.xml: <mapping-file>META-INF/JBPMorm-JPA2.xml</mapping-file> The reason we don't include JPA 2 annotations in the ProcessInstanceInfo.java file/class is that jBPM is still obligated to be JPA _*1*_ compliant. The element-collection mapping is only available in JPA2..	0	0	https://issues.jboss.org/browse/JBPM-3048
13934	JBPM-3044	@Kris, should I close this? I'm not responsible for the designer itself. I think that the engine it's fine for escalations.	0	0	https://issues.jboss.org/browse/JBPM-3044
13935	JBPM-3044	What editor is this pertaining to? The Eclipse BPMN2 Editor or the BPMN2 Web Designer?	0	0	https://issues.jboss.org/browse/JBPM-3044
13936	JBPM-2971	Actually I had to modify the JbpmSchema tool to deal with Sybase. I have ASE 15.0.3 installed in my workstation. When given a script without any "go" command batch terminator, the Sybase isql tool parses the statements but never actually executes them. See http://www.sybase.com/detail?id=938&contentOnly=true#batch I figure that script execution via JDBC interprets statements differently. What method invocation produces the exception you posted?	1	1	https://issues.jboss.org/browse/JBPM-2971
13937	JBPM-2920	Have you tried to use condition tag for transitions? <transition name="transition" to="next step"> <condition expr="#{1==1}"/> </transition>	1	1	https://issues.jboss.org/browse/JBPM-2920
13938	JBPM-2920	Just consider about this. Should we let joinActivity auto-detect when process instance should continue. Just like InclusiveGatewayActivity in BPMN2. Now, JoinActivity used incoming transitions number to decide when process instance should continue. if the actual number of children executions is not equals as incoming transitions, we should set the multiplicity attribute to let JoinActiivty know when should continue. But if we want to support N of M paths in forkActivity, should we ask people to set multiplicity for each time?	1	1	https://issues.jboss.org/browse/JBPM-2920
13939	JBPM-2904	In general it looks good to me. However I have few comments: - please remove system.out.println statements from method - why you use equals method to compare booleans?! isn't it enough to just use them in if statement? - perhaps it is worth mentioning it in docs as well BTW, nice to see that you found a way to set this within configuration file	1	1	https://issues.jboss.org/browse/JBPM-2904
13940	JBPM-2896	Hi, Just one question, the activity name is available in HistoryTaskInstanceImpl so is it possible to just add it in the HistoryTask interface ? That is all my need ! Today I have done a derivative HistoryTask and it directly use HistoryTaskInstanceImpl to populate activity name and it works without problem... Is there something wrong or a negative impact in that approach ? Regards	0	0	https://issues.jboss.org/browse/JBPM-2896
13941	JBPM-2893	Do we have any new date for this release?	0	0	https://issues.jboss.org/browse/JBPM-2893
13942	JBPM-2873	complete implementation of the query operand functionality. supports following query operands on most of the parameters for all query types: - Equals - Like - In - Not in Do you think it is good enough to be incorporated with trunk?	1	1	https://issues.jboss.org/browse/JBPM-2873

#	Requirement	Questions	Answered	Successfully Answered	URL
13943	JBPM-2855	User id needs to be provided prior to adding comment: processEngine.setAuthenticatedUserId("johndoe"); taskService.addTaskComment(taskId, "hi, how are you guys?") processEngine.setAuthenticatedUserId("joesmoe"); taskService.addReplyComment(helloCommentId, "i'm doing fine, thanks"); So I think this issue can be closed.	1	0	https://issues.jboss.org/browse/JBPM-2855
13944	JBPM-2845	jBPM 4.4 and I want to try if jBPM 4.5 solves my problem. org.jbpm.api.JbpmException: couldn't acquire block of ids at org.jbpm.pvm.internal.id.DatabaseDbidGenerator.getNextId(DatabaseDbidGenerator.java:65) at org.jbpm.pvm.internal.lob.Lob.<init>(Lob.java:81) at org.jbpm.pvm.internal.type.variable.BlobVariable.setObject(BlobVariable.java:54) at org.jbpm.pvm.internal.type.Variable.setValue(Variable.java:95) at org.jbpm.pvm.internal.type.variable.BlobVariable.setValue(BlobVariable.java:66) at org.jbpm.pvm.internal.tx.DeserializedObject.flush(DeserializedObject.java:63) at org.jbpm.pvm.internal.tx.AbstractTransaction.flushDeserializedObjects(AbstractTransaction.java:54) at org.jbpm.pvm.internal.tx.SpringCommandCallback.doInTransaction(SpringCommandCallback.java:47) at org.springframework.transaction.support.TransactionTemplate.execute(TransactionTemplate.java:130) at org.jbpm.pvm.internal.tx.SpringTransactionInterceptor.execute(SpringTransactionInterceptor.java:49) at org.jbpm.pvm.internal.svc.EnvironmentInterceptor.executeInExistingEnvironment(EnvironmentInterceptor.java:47) at org.jbpm.pvm.internal.svc.EnvironmentInterceptor.execute(EnvironmentInterceptor.java:42) at org.jbpm.pvm.internal.svc.RetryInterceptor.execute(RetryInterceptor.java:56) at org.jbpm.pvm.internal.svc.ExecutionServiceImpl.getVariables(ExecutionServiceImpl.java:140) at com.agbar.aquacisxf.component.workflow.helpers.ProcessHelper$18.execute(ProcessHelper.java:609) at com.agbar.aquacisxf.component.workflow.helpers.ProcessHelper$18.execute(ProcessHelper.java:605) at org.jbpm.pvm.internal.svc.DefaultCommandService.execute(DefaultCommandService.java:42) at org.jbpm.pvm.internal.tx.SpringCommandCallback.doInTransaction(SpringCommandCallback.java:45) at org.springframework.transaction.support.TransactionTemplate.execute(TransactionTemplate.java:130) at org.jbpm.pvm.internal.tx.SpringTransactionInterceptor.execute(SpringTransactionInterceptor.java:49) at org.jbpm.pvm.internal.svc.EnvironmentInterceptor.executeInNewEnvironment(EnvironmentInterceptor.java:53) at org.jbpm.pvm.internal.svc.EnvironmentInterceptor.execute(EnvironmentInterceptor.java:40) at org.jbpm.pvm.internal.svc.RetryInterceptor.execute(RetryInterceptor.java:56) at org.jbpm.pvm.internal.processengine.ProcessEngineImpl.execute(ProcessEngineImpl.java:300) at com.agbar.aquacisxf.component.workflow.helpers.ProcessHelper.getVariables(ProcessHelper.java:605) at com.agbar.aquacisxf.component.workflow.helpers.ProcessHelper.getVariables(ProcessHelper.java:591) at com.agbar.aquacisxf.component.workflow.impl.WorkflowComponentImpl.getVariables(WorkflowComponentImpl.java:392) at com.agbar.ovs.business.core.impl.WorkflowBCImpl.getVariables(WorkflowBCImpl.java:272) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at	0	0	https://issues.jboss.org/browse/JBPM-2845
13945	JBPM-2830	Sebastian, Is this a 'regression' compared other jBPM 4 releases? Meaning another field was used before? The name should *not* be used for a descrition. That is where the description element is for	0	0	https://issues.jboss.org/browse/JBPM-2830
13946	JBPM-2830	Now, in jBPM-4.4, BPMN use id for the name of process definition. If there is no id, it will throw a deployment exception. And it use name for the key of process definition. Shall we use name for the description of process definition? and just let the key of process definition always empty?	0	0	https://issues.jboss.org/browse/JBPM-2830
13947	JBPM-2830	Oh, now BPMN used the document element of process as the description of process definition. So shall we revert the behaviour to jBPM-4.3?	0	0	https://issues.jboss.org/browse/JBPM-2830
13948	JBPM-2809	Sounds like a good idea, not sure how though what the best way is to get a transaction into the MailActivity... Could it be simply injected?	0	0	https://issues.jboss.org/browse/JBPM-2809
13949	JBPM-2776	Looks great to me. Can you commit please?	1	1	https://issues.jboss.org/browse/JBPM-2776
13950	JBPM-2771	Reading through the section async executions in the documentation I was just wondering how to start a process without blocking the client. It came to my mind that this should be possible via the API without the need to specifiy this in the process defnition explicitly. So additional to fixing this does it make sense to provide a startProcessInstance() -method enabling users to start process definitions in a non-blocking way (= not in the client thread)?	1	1	https://issues.jboss.org/browse/JBPM-2771
13951	JBPM-2764	You could findActivity() by activityName from ProcessDefinition, then you could get the activityName, activityType, activityDescription. Is it enough?	0	0	https://issues.jboss.org/browse/JBPM-2764
13952	JBPM-2759	Attaching a patch to support business time subtractions. Ronald, could you please double-check this? Thanks.	1	1	https://issues.jboss.org/browse/JBPM-2759
13953	JBPM-2759	I will, but why did you change 12 februari to 19? 12 is my birthday ;-)	0	0	https://issues.jboss.org/browse/JBPM-2759
13954	JBPM-2759	Thanks for giving it a try and creating further tests. I ran out of time to create those, but might be able to take a look at this again in the next days. Could you send them to me, please?	0	0	https://issues.jboss.org/browse/JBPM-2759
13955	JBPM-2741	What's mean of this issue? How could we handle this?	0	0	https://issues.jboss.org/browse/JBPM-2741
13956	JBPM-2661	there seems to be a bug with the behaviour in 4.3 When I put three outcomes, two with a expression and one without a expression whereby one expression can be evaluated to true, the engine will fork the evaluated and the empty expression transition. Can somebody please test that?	1	0	https://issues.jboss.org/browse/JBPM-2661
13957	JBPM-2661	Can you point to the location in the spec on which base it being a bug.. afaik, being able to take more than one transition is valid BPMN2 behaviour... All transitions that evaluate to true (including ones without a 'guard') should be taken. This differs from jPDL!!!	0	0	https://issues.jboss.org/browse/JBPM-2661
13958	JBPM-2648	I understand there are many priorities in the draft, but could someone tell me if that's at least a possible Bug I'm not asking for an urgent solution, but at least answer simply because it seems quite serious.	0	0	https://issues.jboss.org/browse/JBPM-2648
13959	JBPM-2648	Thank you for your work, Maceij. And could you marked this issue as resovled?	0	0	https://issues.jboss.org/browse/JBPM-2648
13960	JBPM-2644	To put it another way, not being able query information about a task when it starts is a major drawback. Are we missing something here? I think this issue should have been considered in the wider context of how jBPM is likely to be used before being rejected.	1	0	https://issues.jboss.org/browse/JBPM-2644
13961	JBPM-2644	Hi Huisheng, Thanks for your reply. Having a TASK_STARTED event which fires once the task is available would solve a lot of our issues I'm fairly sure that if this was done, you could then remove the START event entirely since I can't think of a scenario where you would need to distinguish between START and STARTED. Would it be possible to re-open this Jira?	1	0	https://issues.jboss.org/browse/JBPM-2644

#	Requirement	Questions	Answered	Successfully Answered	URL
13962	JBPM-2633	This is marked as "Out of Date", but has anything been done on this? I have the same problem: upgrading jBPM 3.2.x from one version to another is a pain. This is because it implies: # changing the JARs (easy) # seek for information about the need to update the DB schema # seek for information about changed APIs (either model APIs or infrastructure APIs, since we're embedding jBPM in a non-JEE application) # adapt our extension of JobExecutor and related classes (since its design is not extension-friendly) # adapt our own patches for unsolved bugs (well, at least this is unavoidable) So, is there some source of information (apart from JIRA) I can look at to make this process easier?	0	0	https://issues.jboss.org/browse/JBPM-2633
13963	JBPM-2576	Is the HistorySessionChainBinding part of some configuration file parsing?	1	1	https://issues.jboss.org/browse/JBPM-2576
13964	JBPM-2570	Sebastian, Do you have a real full unittest for this? I 'fixed' something in https://jira.jboss.org/jira/browse/JBPM-2749 which might be related.	0	0	https://issues.jboss.org/browse/JBPM-2570
13965	JBPM-2533	I have started a discussion about similar issue in the forum: http://community.jboss.org/message/528104#528104 I posted as well some information about what was implemented by me to be able to satisfy business requests. Are we talking about the same things? If so, are there any plans for implementation of this, especially if it could be part of upcoming release 4.4.	0	0	https://issues.jboss.org/browse/JBPM-2533
13966	JBPM-2524	Joram, are you sure this is a priority for 4.2 ? Can you explain the use case. Bear in mind that the install scripts should produce the configuration that you want. So can you elaborate with the use case when you need this ? In order to see if it is a developer feature or a userguide feature.	1	1	https://issues.jboss.org/browse/JBPM-2524
13967	JBPM-2509	Since the upgrade scripts will be checked into version control, I do not see why a separate DbUpgrade tool is required, separate from the standard Ant SQL task [1]. Is it just to leverage the known locations of the upgrade scripts in the distribution? If so, could DbUpgrade be simply a macro, or extended from SQLTask? Otherwise, is the intention to enable programmatic invocation as well? [1] http://ant.apache.org/manual/CoreTasks/sql.html	1	1	https://issues.jboss.org/browse/JBPM-2509
13968	JBPM-2508	I'm using jbpm 4.2 with tomcat 6.0.20 and hsqldb(standalone) mode. When I launch jbpm console using http://localhost/:8080/jbpm-console after entering userid/password peter/password , it shows creating workspace and then breifly shows Tasks, etc and then crashes IE. What is the alternative to this? How can I view/manage the process definitions created etc.... This is a high priority issue. The workaround of logout then login does NOT work.	0	0	https://issues.jboss.org/browse/JBPM-2508
13969	JBPM-2508	Would it work on Jboss 5.0.0 ?	0	0	https://issues.jboss.org/browse/JBPM-2508
13970	JBPM-2508	This crashing business doesn't seem to be a problem for Firefox. It crashes IE for me also, though I have not checked out Safari. Does this problem have anything to do with the timeout issues associated with remoting callbacks? There was a similar issue in SEAM a while back that apparently got resolved... https://jira.jboss.org/jira/browse/JBSEAM-3459	0	0	https://issues.jboss.org/browse/JBPM-2508
13971	JBPM-2506	When declaring variables already in the process definition it is possible to specify whether or not they should be taken into account for the history. So the same configuration should be possible via the API for variables which are created during run-time. btw: Why is the history disabled by default? What are the reasons against enabled by default?	0	0	https://issues.jboss.org/browse/JBPM-2506
13972	JBPM-2506	Hi Maciej, Glad to see that you are working on this issue. Do you have any plan for this? Do we need create a new post in the forum and talk about it?	0	0	https://issues.jboss.org/browse/JBPM-2506
13973	JBPM-2506	you can leave JBPM-2840, I have already solved that. But thanks for pointing me to that issue so I will mark it as well when I am done. I am considering custom type variables as well, as a matter of fact I will make some tests today to see how difficult will it be.	0	0	https://issues.jboss.org/browse/JBPM-2506
13974	JBPM-2506	Part of the work is done now: - declare variable in jpdl - mark declared variable as history enabled - create variable via API that is history enabled - retrieve variables from history via API - support for both simple text and EL on declared variables in jpdl If it comes to non string variables that is bit more complex. Mainly due to the structure of the data base where value column is only of type string (varchar). I did some tests with JavaBeans as variable and simple solution that I implemented is that such variables will be still persisted as string based on implementation of toString() method. In my opinion that is quite good solution at the moment that enables us to simply query for history process instances based on variable's values. Of course this is not a final solution but at least something to let users hove some added value from variable history. Difficult to say what was initial idea for history of variables ... any news are more than welcome. In addition, such approuch could be used as well to allow queries based on variables for execution variables as well. I mean that if some once would be interested to query based on variable value (s)he could implement toString with information that should be searchable. Soon will upload a first version of a patch for review. What do you think based on the description so far?	0	0	https://issues.jboss.org/browse/JBPM-2506
13975	JBPM-2505	We are currently evaluating open source BPMS. The search capability is really crucial for us. We have a single workflow that gets startet hundred times a day. So we have thousands of instances. How should a supervisor find critical processes? I really love jBPM so far, but this could be a show stopper :-/	0	0	https://issues.jboss.org/browse/JBPM-2505
13976	JBPM-2505	Do you have any updates on this issue? We are deploying v4.3 and at this point will probably have to create our own tables to enable searches. We like JBPM a lot and are going forward either way, but his feature would be incredibly valuable to us.	0	0	https://issues.jboss.org/browse/JBPM-2505
13977	JBPM-2477	It might be related to the fact that a newer driver is used here compared to the one jBPM is tested with. Can someone check? In the meantime I'll ask somebody to take a look at the 'patch'	0	0	https://issues.jboss.org/browse/JBPM-2477
13978	JBPM-2465	can you add a code snippet of how you like to use this or another description of what you exactly mean ? and why activity.getDefaultOutgoingTransition is not convenient/clean enough.	0	0	https://issues.jboss.org/browse/JBPM-2465
13979	JBPM-2432	the description is insufficient. can you please elaborate?	1	1	https://issues.jboss.org/browse/JBPM-2432
13980	JBPM-2426	I think this should only be tackled if we can build all the features to execute following demo scenario: - process variable is a open office (or word) document - review task - notification that includes a process variable as attachment - assignee can then make updates to document and reply - a simple mail listener will capture the replied document, update the process variable and complete the task. i think the attachment feature will only be valuable if we also build the reply functionality. reply functionality should be done by an email-deamon that is configured very much like the job-executor. it would have 1 thread that periodically scans the jbpm inbox using JavaMail only. what do you think? would you consider this feasible in this release? another extension idea is the introduction of a VersionedFile object that can be stored and retrieved as a process variable.	1	1	https://issues.jboss.org/browse/JBPM-2426
13981	JBPM-2417	What is 'specifiying authentication' in this regard? I have no clue For me it is like 'basic authentication', 'certificate authentication' or something like that.	1	0	https://issues.jboss.org/browse/JBPM-2417
13982	JBPM-2416	How is this issue related to JBPM-2417? I don't see the connection. Could anyone clarify? Thanks.	1	1	https://issues.jboss.org/browse/JBPM-2416

#	Requirement	Questions	Answered	Successfully Answered	URL
13983	JBPM-2415	I am not sure if I understand the description correctly? Does this mean there is already an implementation to declare process variables but it needs to be extended so that you can declare within a group activity definition? Thanks for any explainations.	1	0	https://issues.jboss.org/browse/JBPM-2415
13984	JBPM-2414	what kind of for-each did you have in mind? There are differences in creating 10 tasks for the same user with different variables like task: learn to play instrument 10 times, each with e.g. another instrument or 10 identical tasks for different users or 10 different subprocesses, which ones based on e.g. a processvariable etc... So a little more specific would be nice. The most requested ones are the last two. But often in combination with an n-out-of-m join and cancelling the remaining tasks.	1	1	https://issues.jboss.org/browse/JBPM-2414
13985	JBPM-2414	What I have implemented for now: <for-each name="fork" vars="#{myCollection}" var="item"> <transition to="myTask" name="to_task"/> </for-each> <task name="myTask"> <transition to="join" name="to_join"/> </task> <join name="join"> <transition to="myTask" name="to_end"/> </join> What this does is create new executions for each element in a collection (set, list, collection, array) and put a variable with the name 'item' that in a local scope for the execution. In e.g. the task you can do whatever you want with that variable.... Assignments, do things in event-listeners whatever.... Limitations: - it can only have one outgoing transition - A corresponding join with one incomming transition is needed - Complex nesting things between this for-each and join can and most likely will result in strange behaviour (although I think this is very acceptable) Additions: - The join now accept an EL for multiplicity. A Long/Integer, String (that will be parsed runtime) and a collection (size will be determined runtime). Things like 70% of the created executions has to be passed before it will continue can work now... cool isn't it? One thing I maybe like to change is not call it 'foreach/for-each' but put it in the fork... This makes it visible in the designer, *now*. Comments? After that, I will add the patches here.	1	1	https://issues.jboss.org/browse/JBPM-2414
13986	JBPM-2414	Is this already in the repository? Are there any disadvantages when using <fork foreach> instead of <foreach>?	1	1	https://issues.jboss.org/browse/JBPM-2414
13987	JBPM-2414	No it is not in repo yet. I'll attach the 'patches´ tonight. Why it is not in repo is because I did not (well, not quite at least) receive feedback from the core devs. Why can be read on the post by Alejandro Guizar http://www.jboss.org/feeds/view/jbossjbpm. The disadvantage of using either that is it is not easy to model it in the BPMN way (with a || sign on e.g.a repeating task) using the 'group' thing would make this easier. Otherwise there is no real disadvantage afaics	1	1	https://issues.jboss.org/browse/JBPM-2414
13988	JBPM-2414	Is there any chance to get hold of Ronald's work? That would be very helpful with moving forward with this issue. Anyone is capable to contact Ronald? Thanks Maciej	1	1	https://issues.jboss.org/browse/JBPM-2414
13989	JBPM-2414	While waiting for Ronald's work I tried to implement it based on the discussion given here and some other ideas. In general, it is pretty much how Ronald described this, with one difference, it can have more than one transition in fork node. In such case all transitions will be multiplied by number of items in for-each collection. Please have a look at the patch that includes both code changes and few test cases to illustrate (and of course verify) working fork - for each. Let me know what do you think about it. P.S. If you need more details about how does it work, just let me know and I will try to describe it.	1	1	https://issues.jboss.org/browse/JBPM-2414
13990	JBPM-2414	This is a great feature! Could I suggest a completion expression is also added to the join? This would allow more control when releasing the join, similar to the dynamic partial join pattern, by default the multiplicity behaviour would apply. Tim	1	1	https://issues.jboss.org/browse/JBPM-2414
13991	JBPM-2414	Tim, if I understood your comment correctly that should already be included. multiplicity attribute of join indicates when execution should move on (number of incoming transitions) and it accepts expression. Is that what you ment?	1	1	https://issues.jboss.org/browse/JBPM-2414
13992	JBPM-2414	While waiting for this change to be officially released, what is the preferred way to implement Wfp16MiWithoutAPrioriRuntimeKnowledge in jbpm 4.2/4.3? All references I can find point to createTaskInstance() which exists in v3, but is no longer available in v4. Thanks!	0	0	https://issues.jboss.org/browse/JBPM-2414
13993	JBPM-2413	After reading the group activity description in the developer documentation I do not see the connection either. http://docs.jboss.com/jbpm/v4/devguide/html_single/#groupconcurrency A group activity is a kind of sub process for me which has a static definition just like a normal process. Of course there can be concurrent executions within a group activity if the group activity is modelled like this. So you have 2, 3 or more concurrent executions but just as many as you modelled during design time. In contrast the foreach activity could have several concurrent executions. How many exactly would be decided kind of dynamically - maybe even at runtime, right? Or you could define and set a property for the group activity at design time telling the number of concurrent executions which should be created. is the later one in the manner you think of, Tom?	0	0	https://issues.jboss.org/browse/JBPM-2413
13994	JBPM-2412	I would to, but what about ditching the target-ejb-home and method thing completely and using just one <java action="#{myEjbBean.myMethod} .... The target-expr is only usefull if you want to specify what bean to call in a processvariable. Does that happen often? I've never come across it. My example with <java action="..." could also be used to calls webbeans/seam pojo's/spring beans etc... Wouldn't that be much nicer?	1	1	https://issues.jboss.org/browse/JBPM-2412
13995	JBPM-2404	was this wiki page ever created ? I am finding tons of posts related to deploying jbpm-console to WLS 10.3 but none of them are really conclusive whether this is possible or not. Is WLS 10.3 a supported platform for jbpm 4.x at all ?	0	0	https://issues.jboss.org/browse/JBPM-2404
13996	JBPM-2403	Mark, Do you have an example process in which the fork activity fails? The fork tests in our test suite aren't failing.	0	0	https://issues.jboss.org/browse/JBPM-2403
13997	JBPM-2399	This issue related to the readContext and writeContext. How could we handle it? If there is no plan for read/write context, should we just marked it as Out of Date?	0	0	https://issues.jboss.org/browse/JBPM-2399
13998	JBPM-2398	work started on the test. test needs further elaboration. also test the task comment replies. how should users display the full comment discussion?	0	0	https://issues.jboss.org/browse/JBPM-2398
13999	JBPM-2398	patch that introduces some enhancements in history handling: - super process instance relationship - process initiator - task name and due date available on history task - query for tasks by process instance id - query for process instance by super process instance id (to find all subprocesses) - query by initiator of process instance history details are added for following operations: - change assignee of a task (previous and current assignee are stored) - change priority of a task (previous and current values are stored) - change due date of a task (previous and current values are stored) - change of variable values that has history enabled (previous and current values are stored) To make it doable following entities/tables were extended: Execution: - ExecutionImpl/JBPM4_EXECUTION new column added initiator/INITIATOR_ History: - HistoryProcessInstanceImpl/JBPM4_HIST_PROCINST new columns added superProcessInstanceId/SUPERPROCID_ and initiator/INITIATOR_ - HistoryTaskImpl/JBPM4_HIST_TASK new columns added name/NAME_ and duedate/DUEDATE_ Comments? Do you think it is good enough to be committed?	1	1	https://issues.jboss.org/browse/JBPM-2398

#	Requirement	Questions	Answered	Successfully Answered	URL
14000	JBPM-2378	Koen, are there fixes that we can still apply in the next 2 days ? For which you think it is pretty clear ? The rest, we need to address in 4.1. Can you make a separate jira issue for XSD related matters in 4.1 and add good descriptions of each item that you consider not logical ?	1	1	https://issues.jboss.org/browse/JBPM-2378
14001	JBPM-2368	[14:23] <kconner> tomb: If this is behaving the same way as previously then it would suggest that the thread is still running after being undeployed but then the question would be... is this problem only appearring when jboss shuts down ? or when the jbpm service is shut down. in that case we should synchronize on ending of the job executor. and in that case the issue would not be a blocker, i think.	1	1	https://issues.jboss.org/browse/JBPM-2368
14002	JBPM-2359	heiko, i have an oracle installation locally. is there some way i could help you with this issue by trying something out ?	1	1	https://issues.jboss.org/browse/JBPM-2359
14003	JBPM-2359	is there something i can do to help? can i just try to update the queries in the report directly? could it be just a semicolon at the end ?	1	1	https://issues.jboss.org/browse/JBPM-2359
14004	JBPM-2354	Heiko, I've moved the SQLStmtTest to a normal package name: org.jbpm.test.reporting And I added it to the exclude list. Could you have a look at this ?	0	0	https://issues.jboss.org/browse/JBPM-2354
14005	JBPM-2341	Is this really a change that people need? The implementations I have done all had requirements that the assigned user for swimlanes stayed the same during the process. So if this is implemented I'd make the default *remember* instead of not remember.	1	1	https://issues.jboss.org/browse/JBPM-2341
14006	JBPM-2339	heiko, can you check if my latest commit fixed the problem ?	1	1	https://issues.jboss.org/browse/JBPM-2339
14007	JBPM-2329	What is actually the reason for the schema change and why is this postponed again? Grateful for any light shed on this.	0	0	https://issues.jboss.org/browse/JBPM-2329
14008	JBPM-2320	can you explain your use case a bit better ? first explain what you want to do and next, which features and methods you need for it. thanks	1	1	https://issues.jboss.org/browse/JBPM-2320
14009	JBPM-2293	the sources are attached in the distro profile. this way we avoid the creation of all these jars during normal development. reopen if you don't think this is a good solution, reopen and discuss why	1	1	https://issues.jboss.org/browse/JBPM-2293
14010	JBPM-2291	Which is the actual fix version? 3.3.2? 3.2.9? 3.2.6.SP1 is affected by this problem.	0	0	https://issues.jboss.org/browse/JBPM-2291
14011	JBPM-2281	tried, but wasn't able to reproduce. my environment: OS: windows XP C:\Software\jbpm-4.0-SNAPSHOT\jboss>java -version java version "1.5.0_11" Java(TM) 2 Runtime Environment, Standard Edition (build 1.5.0_11-b03) Java HotSpot(TM) Client VM (build 1.5.0_11-b03, mixed mode, sharing) C:\Software\jbpm-4.0-SNAPSHOT\jboss>ant -version Apache Ant version 1.7.0 compiled on December 13 2006 Maybe it's the ant version. Can you verify that, Burr ?	1	1	https://issues.jboss.org/browse/JBPM-2281
14012	JBPM-2251	The mail session configuration is in jbpm-pvm/src/main/resources/jbpm.default.cfg.xml. It was placed there because all other sessions are (pvm-db-session, message-session, timer-session, etc.). What we could do to avoid having explicit properties there is externalize to a properties file, as in the hibernate configuration case. <mail-session> <mail-server> <session-properties resource="jbpm.mail.properties" /> </mail-server> </mail-session> The mail template is in jbpm-examples/src/test/resources/jbpm.cfg.xml. If needed we can also take this one out through the import mechanism. <import resource="jbpm.mail.templates" /> What are your thoughts? When you reply, please assign this issue to me. JIRA will send me a notification.	1	1	https://issues.jboss.org/browse/JBPM-2251
14013	JBPM-2250	Heiko, could you rename those tabs in the console as well to Personal Task List and Group Task List ?	0	0	https://issues.jboss.org/browse/JBPM-2250
14014	JBPM-2242	To implement suspend and resume on deployments and cascade it to executions, jobs and tasks, I started with synchronizing following properties: * Deployment.state * Execution.state * Job.state * Task.state (some of those had an isSuspended boolean) So i switched to all state (String) properties. Two operations are added to the repository service: suspendDeployment and resumeDeployment. This cascades to all the related executions, jobs and tasks.	0	0	https://issues.jboss.org/browse/JBPM-2242
14015	JBPM-2214	koen, see the user guide for the actual schema and meaning on sub-process parameter-in and parameter-out can you resolve this issue when the necessary designer updates have been done ?	0	0	https://issues.jboss.org/browse/JBPM-2214
14016	JBPM-2203	> Yes. Are you using Unified EL now? Or another EL flavour... i've plugged in jsr 223 with scripting.dev.java.net. in the configuration, you can configure a named list of expression languages and also indicate which is the default expression language. in the jpdl language, for each expr attribute, there is an optional lang attribute. > Possibly - I'm on PTO quite a lot in May, but I think we can sort > something out. How about 18th onwards? ok. my expectation is that you drive the seam integration. i'll be ready to supply any jbpm update that you'll need and re-upload a snapshot. does that work for you ? regards, tom. Pete Muir wrote: > > On 14 Apr 2009, at 14:50, Tom Baeyens wrote: > >> the initial work needs to be done by us, i think. >> >> implementation wise there are a couple of directions, but there is still a bit of thinking involved. >> >> mostly, the current integration is based on integrating the naming resolutions while resolving UEL expressions. apart from that there is only the special SEAM annotations, which will boil down to jbpm service invocations. >> >> * setup of the seam contexts. find the proper places where we need to intercept to set up the seam contexts. this could potentially be very easy. we have a set of services now (==session facades). we also have a mechanism to add interceptors to those. a seam interceptor could ensure that seam is properly initialized around jBPM operations. > > Great, sounds easy enough :-) > >> * expanding the jbpm variable resolver so that it also delegates to the seam contexts. > > Yes. Are you using Unified EL now? Or another EL flavour... > >> * tx around the jbpm timers and messages. we should be able to leverage our container integration. since we'll be deployable as a service, we'll build our own tx interceptor for managing the JTA tx. this is leveraged when using seam. > > Kev has been doing some work on this - but I agree, we should just use JTA here - Seam will join or start a TX as required. > >> >> >> * pageflow. build the page activity. create a pageflow parser. >> >> this is all not cristal clear though. requires getting our hands dirty to see if this plans are complete. i think we should plan a week where we together set out the initial structure. then we can involve productization people to finish the integration. > > Ok. > > ... >> >> regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-2203
14017	JBPM-2203	Is it possible to make the extensability of the variable resolver configurable like it was in jbpm 3? We need this for a project where data is stored in our own datastore and not as variables in the tasks. (see also http://www.jboss.org/index.html?module=bb&op=viewtopic&t=160042) Wkr, Sofie	1	0	https://issues.jboss.org/browse/JBPM-2203
14018	JBPM-2200	Classloading was fixed on 4.2 trunk. @Tom: can you update the doc (chapter 4)	0	0	https://issues.jboss.org/browse/JBPM-2200
14019	JBPM-2193	For the images, how will they get in the console ? through a servlet ? is there a servlet that will process a request in the line of http://localhost/gwt-console/processimage?procdef=InsuranceClaim-8374 then that servlet will go to the repository service. currently there is a resource method to look up a resource by deploymentDbid InputStream getResourceAsStream(long deploymentDbid, String resourceName); but that would not be sufficient. we probably need a method like InputStream getResourceAsStream(String processDefinitionId, String resourceName); Is that also how you see it ? A servlet that depends on a new -to-be-created method getResourceAsStream(String processDefinitionId, String resourceName) ?	0	0	https://issues.jboss.org/browse/JBPM-2193

#	Requirement	Questions	Answered	Successfully Answered	URL
14020	JBPM-2193	Heiko, One more question: is a relative link good between the process definition file and the image file ? Given the processDefinitionId, I can look up the resource name of the process file. Then I can replace .jpdl.xml with .png for instance to get the image file. I think having the link between process file and image file based on the same file name, but with different extensions is good. It's also easy for the designer. Alternatively, we would have to maintain a reference from the process definition to its image file... Doesn't seem like a good idea to me. WDYT ?	0	0	https://issues.jboss.org/browse/JBPM-2193
14021	JBPM-2172	first thing we need to do is establish the concrete first goal. i think this should be the following: * create a jboss installation where - jboss idm is deployed as a service archive and identity session factory is published in jndi - jbpm is deployed as a service archive and is configured so that it uses the jboss idm to get the jbpm installation, start by installing jbpm into the as using the installer (build it with mvn -Pdistro clean install from source) then run through the installer. then update the jbpm configuration files to use the jboss idm component does that make sense ?	1	1	https://issues.jboss.org/browse/JBPM-2172
14022	JBPM-2172	later this week, i'll think about how we can create a profile in our build so that the test suites use jboss idm instead of our build-in identity component. in order to do that, i'll need to integrate idm schema creation into our build. how will that work ?	1	1	https://issues.jboss.org/browse/JBPM-2172
14023	JBPM-2172	Right now, it creates the idm schema and populate it in the runtime. I've already built the schema in the build time, I need to use the jboss mechanism to populate the schema into database in the following days. how did jbpm handle this?? Thanks Jeff	1	1	https://issues.jboss.org/browse/JBPM-2172
14024	JBPM-2172	btw, what property do we have to set in which build script of jbpm ? i think it should be adding -Didentity.component=jboss.idm to line ant -f qa/build.xml $ANT_PROPERTIES integration.testrun.setup in file qa/hudson-jbpm4-jboss.sh can you verify that that script works when specifying that property?	1	1	https://issues.jboss.org/browse/JBPM-2172
14025	JBPM-2172	[15:27] <tomb> jeffyu: you around ? [15:27] <jeffyu> tomb, Hi tom, Yes. [15:27] <tomb> jeffyu: hi [15:28] <jeffyu> tomb, was still looking into the Transaction issue... [15:28] <jeffyu> tomb, I find an minor issue that is in the JTATransaction [15:28] <tomb> jeffyu: in the jbpm trunk, i see that the jboss idm version still points to org.jboss.identity.idm:idm-api:jar:1.0.0-SNAPSHOT [15:28] <tomb> jeffyu: is that intentional ? [15:29] <jeffyu> tomb, yes.. because there were two test cases failed, I didn't make an alpha release for the update. [15:29] <jeffyu> was thinking relied on the new published SNAPSHOT. [15:29] <tomb> jeffyu: you mean you couldn't make an idm alpha release ? [15:29] <jeffyu> after resolved those two test case, I will make another alpha release for the jBPM usage. [15:30] <jeffyu> tomb, I can, but I was thinking I might do some little update due to those two failed test case. related to the Transaction. [15:30] <jeffyu> tomb, you want me to do another alpha release now? [15:30] <jeffyu> I saw you want to make the CR release, right? [15:31] <jeffyu> tomb, if you want to, I think we can do another idm release now. [15:31] <bdawidow> jeffyu: Anil will need one anyway for assembly module fix [15:31] <bdawidow> jeffyu: just give me a moment to fix [15:31] <tomb> jeffyu: my cr release is already underway [15:32] <tomb> jeffyu: so you don't have to do it for that release [15:32] <jeffyu> tomb, ok, yes, the SANPSHOT is intentional.:) [15:32] <jeffyu> tomb, ok. [15:32] <tomb> jeffyu: but it's always good to have it on a specific release [15:32] <tomb> jeffyu: i mean it is always better to ... [15:33] <tomb> jeffyu: when do you think you could do that on trunk ? [15:33] <jeffyu> tomb, agreed... was thinking the Transaction issue might be easier to fix... and didn't realized the CR release is going to happen in these days... sorry. [15:33] <tomb> jeffyu: no prob. if it mattered, then i would have brought it up [15:33] <jeffyu> tomb, you mean make another idm release on idm trunk. [15:34] <tomb> jeffyu: instead of using idm snapshot in jbpm trunk, i would like to switch to using idm releases asap [15:34] <jeffyu> bdawidow, ok, may I know whats the problem for current assembly module?:) [15:34] <tomb> jeffyu: if you're about to do a release then we can wait for that, but i want to plan that [15:34] <bdawidow> jeffyu: looking at it now. Anil says it doesn't contain all stuff atm [15:35] <jeffyu> tomb, I can do another release tomorrow. (after the bdawidow committed the fix) [15:35] <jeffyu> tomb, and then I will update the idm version on the jbpm trunk if you don't mind. [15:35] <tomb> jeffyu: that would be great [15:37] <jeffyu> tomb, ok, will give you an update tomorrow.. btw, JTATransaction.class uses the 'java:comp/UserTransaction' as JNDI name, I thought it should be 'UserTransaction' in the JBoss.	1	1	https://issues.jboss.org/browse/JBPM-2172
14026	JBPM-2170	Does it deploy to the db directly or via the console (http upload)?	1	1	https://issues.jboss.org/browse/JBPM-2170
14027	JBPM-2166	can anyone with schema knowledge paste the right piece of schema to be added to all jpdl elements so that the language becomes extensible ? i'm looking for the proper usage of the any and anyAttributes. the goal is that elements and attributes can be added in any namespace (also the jpdl namespace)	1	1	https://issues.jboss.org/browse/JBPM-2166
14028	JBPM-2166	Are you already working on this? If I got it right, the schema should be extended in that way, that you can add additional elements like for example customized task-elements and additional attributes to existing elements while validation through the schema still works? Might be helpful: http://www.w3schools.com/schema/schema_complex_anyattribute.asp	0	0	https://issues.jboss.org/browse/JBPM-2166
14029	JBPM-2166	What if there were a different approach? I feel like the available node set is adequate, with the existing custom node. The "killer" feature that's missing is the ability to enhance or replace the existing node behavior. So, what do you think about optionally allowing an activity behavior chain? Any activity could specify a list of ActivityBehavior implementations that would be executed on activity start. Additionally, we could specify 2 additional flags: include-default-behavior and default-behavior-location. The attribute include-default-behavior would specify to include the default (a boolean attribute), pre-implemented behavior in the node execution and would default to true if not specified. The attribute default-behavior-location is a enumerated attribute of "first" or "last" and would specify where in the chain the default behavior would be executed. It would be defaulted to "first" if not specified. ActcivityBehavior would be enhanced to return a boolean from the execute(ActivityExecution) method and would indicate if processing in the chain should continue. All of the default implementations would return "true", indicating the execution chain should continue. A return value of "false" would indicate that no further processing should continue. This would allow custom behavior to do things like conditionally create tasks in the task node, or notify task candidates after the task has been created. It could open up a bunch of possibilities... I'll upload an example xml file of what I've described.	1	1	https://issues.jboss.org/browse/JBPM-2166
14030	JBPM-2154	is this a big impact ?	0	0	https://issues.jboss.org/browse/JBPM-2154

#	Requirement	Questions	Answered	Successfully Answered	URL
14031	JBPM-2147	---------- Tests run: 1, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 2.015 sec <<< FAILURE! testTaskSwimlane(org.jbpm.examples.task.swimlane.TaskSwimlaneTest) Time elapsed: 1.953 sec <<< FAILURE! junit.framework.AssertionFailedError: Expected a single task in johndoe's task list expected:<1> but was:<0> at junit.framework.Assert.fail(Assert.java:47) at junit.framework.Assert.failNotEquals(Assert.java:282) at junit.framework.Assert.assertEquals(Assert.java:64) at junit.framework.Assert.assertEquals(Assert.java:201) at org.jbpm.examples.task.swimlane.TaskSwimlaneTest.testTaskSwimlane(TaskSwimlaneTest.java:68) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at junit.framework.TestCase.runTest(TestCase.java:154) at org.jbpm.test.BaseJbpmTestCase.runTest(BaseJbpmTestCase.java:81) at junit.framework.TestCase.runBare(TestCase.java:127) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:118) at junit.framework.TestSuite.runTest(TestSuite.java:208) at junit.framework.TestSuite.run(TestSuite.java:203) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.apache.maven.surefire.junit.JUnitTestSet.execute(JUnitTestSet.java:213) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:140) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:127) at org.apache.maven.surefire.Surefire.run(Surefire.java:177) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.apache.maven.surefire.booter.SurefireBooter.runSuitesInProcess(SurefireBooter.java:338) at org.apache.maven.surefire.booter.SurefireBooter.main(SurefireBooter.java:997) ---------- Test set: org.jbpm.examples.concurrency.graphbased.ConcurrencyGraphBasedTest ---------- Tests run: 1, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 1.266 sec <<< FAILURE! testConcurrencyGraphBased(org.jbpm.examples.concurrency.graphbased.ConcurrencyGraphBasedTest) Time elapsed: 1.25 sec <<< FAILURE! junit.framework.AssertionFailedError: execution ConcurrencyGraphBased.4 should not exist at junit.framework.Assert.fail(Assert.java:47) at junit.framework.Assert.assertTrue(Assert.java:20) at junit.framework.Assert.assertNull(Assert.java:233) at org.jbpm.examples.concurrency.graphbased.ConcurrencyGraphBasedTest.testConcurrencyGraphBased(ConcurrencyGraphBasedTest.java:98) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at	1	1	https://issues.jboss.org/browse/JBPM-2147
14032	JBPM-2136	note that the attached state diagram does not reveal the full lifecycle of a task. only the state related to assignment.	1	1	https://issues.jboss.org/browse/JBPM-2136
14033	JBPM-2124	Why this property was renamed? dbid sounds a bit weired, id would be the much more natural choice, or not?	1	0	https://issues.jboss.org/browse/JBPM-2124
14034	JBPM-2110	is this still relevant ? if yes, can you explain this a bit.	0	0	https://issues.jboss.org/browse/JBPM-2110
14035	JBPM-2065	Burr Sutter wrote: > I was speaking with client XXX today, they have a very large system based on > jBPM3 and they don't plan to ever migrate it. > > They want to see jBPM3 and jBPM4 both running in a production system > (e.g. EAP/SOA-P), same JVM, same DB. > > What do you think?	0	0	https://issues.jboss.org/browse/JBPM-2065
14036	JBPM-2043	Martin, how do I run the single test in 3.2.2 against mysql? In 3.2.5 this would be > mvn -Ddatabase=mysql -Dtest=AsyncCallTest test	0	0	https://issues.jboss.org/browse/JBPM-2043
14037	JBPM-2043	A build of that branch fails for me with BUILD FAILED /home/tdiesler/svn/jbpm/jbpm3/branches/jpdl-3.2.2-SOA-4.2/build/build.xml:35: The following error occurred while executing this line: /home/tdiesler/svn/jbpm/jbpm3/branches/jpdl-3.2.2-SOA-4.2/build/base.project.build.xml:619: The following error occurred while executing this line: /home/tdiesler/svn/jbpm/jbpm3/branches/jpdl-3.2.2-SOA-4.2/build/base.project.build.xml:624: The following error occurred while executing this line: /home/tdiesler/svn/jbpm/jbpm3/branches/jpdl-3.2.2-SOA-4.2/jpdl/userguide/build.xml:17: Cannot find /home/tdiesler/jbpm/repository/docbook-support/support.xml imported from /home/tdiesler/svn/jbpm/jbpm3/branches/jpdl-3.2.2-SOA-4.2/jpdl/userguide/build.xml Could you please fix the build or how the dependency on docbook-support is handled?	0	0	https://issues.jboss.org/browse/JBPM-2043
14038	JBPM-2043	Your numbers are for HSQLDB, right? Can you please redo this for MySQL?	0	0	https://issues.jboss.org/browse/JBPM-2043
14039	JBPM-2042	Unschedule from 3.2.6. When done, please set the fix version appropriately	0	0	https://issues.jboss.org/browse/JBPM-2042
14040	JBPM-2038	Tom, what do you mean by "this should be a developer feature"? Please also see my post in the forum: https://community.jboss.org/thread/150549?tstart=0 The "Configuration.setHibernateSessionFactory(SessionFactory)" doesn't seem to work... What's the best way to deal with externaly managed transactions and hibernate session factories?	0	0	https://issues.jboss.org/browse/JBPM-2038
14041	JBPM-2006	I think this is really a cool feature going to boost jBPM. Is there support from the Drools guys?	1	1	https://issues.jboss.org/browse/JBPM-2006
14042	JBPM-2006	I have a question regarding the integration approach. I don't have much experience in relation to rule engines but there is JSR 94: Java Rule Engine API. AFAIK as of version 2.0 Drools is a fully JSR94-compliant rule engine. Maybe we should take this as an integration approach and thus not aime for Drools integration or support but for support for JSR94-compliant java rule engines? What do you think about this?	1	1	https://issues.jboss.org/browse/JBPM-2006
14043	JBPM-2006	Which approach did Jeff Delong follow in jBPM 3.? Did he use serialized rules or a repository?	0	0	https://issues.jboss.org/browse/JBPM-2006
14044	JBPM-1996	Heiko, is it ok for you if I take this issue ?	0	0	https://issues.jboss.org/browse/JBPM-1996
14045	JBPM-1995	I added this section in the userguide: <section id="databasemigration"> <title>Database migration</title> <para>TODO: document the instructions how users should apply the automatic migration</para> </section> Alejandro, can you complete that docs as part of this jira issue ?	1	1	https://issues.jboss.org/browse/JBPM-1995
14046	JBPM-1995	Does this really belong in the user guide? I don't think we want to support upgrading between arbitrary versions. Furthermore, because it requires checking out the source code and working with Maven, it is hardly an end-user feature. I'll write the instructions in the user guide for the time being, but I think it really belongs to the dev guide.	0	0	https://issues.jboss.org/browse/JBPM-1995
14047	JBPM-1974	"this way an activity will refer to the model element, as well as to the combination of the model element and the behaviour." Do you really think that an overlay in concepts on a specific term makes things easier to understand? It's actually the other way around. Furthermore Activity is a BPMN term. If we use a mixed approach in jpdl4 then we need to avoid overlap in terms between BPMN and jbpm.	1	1	https://issues.jboss.org/browse/JBPM-1974
14048	JBPM-1927	do i get this correctly that you're planning to bind the file upload deployment in the console to the jboss deployer ? if yes, then i don't think this is a good idea. programmatic deployment is also part of the api. and that will work in any container. also tomcat.	1	1	https://issues.jboss.org/browse/JBPM-1927

#	Requirement	Questions	Answered	Successfully Answered	URL
14049	JBPM-1924	Thomas, The participant production-ready forms can be created with JBoss tools.... they have a palet for many components, jBPM should not provide that themselves. The basic form, generated by jBPM can be adapted. I've done it. The 'visual task form designer'? What is that? I have no clue... jBPM does not provide anything visual for this at all... so to say it is relatively weak is to much credit ;-). All changes can be deployed in the par file... so.... The JSF console had indeed the problem that it showed all tasks. I fixed that once, but maybe something got lost in the process. I'll have a look again. The JBoss logo is customizable... it's an image... replace background.gif in the images directory in the webapp... We can also change it that it is customizable via css or whatever... Or do I miss something here? And critical??? hmmmm... strange report	1	1	https://issues.jboss.org/browse/JBPM-1924
14050	JBPM-1924	Hello Ronald, There is a "visual form designer" in JBoss Tools that almost works well with jBPM's generated task forms, ideally the generation logic for task forms would be much more clever, using specialized form input types (e.g. drop-down list box for an enumeration), it shouldn't have to be "adapted" for simple applications. Concerning the logo - I think it is bundled into the .WAR file and I (the business analyst) can't easily crack open, fix and repackage a .war file. I have found that the current JSF-based console has been the single most complained about item in jBPM. Many folks who approach a BPM problem and are looking for a solution (e.g. jBPM) are not too interested in expending lots of hours of highly talented (and costly) Java programming time to get a prototype created, tested, re-worked (rinse, wash, repeat) and then finally to production. This means the tooling has to improve so that it makes better decisions about its code generation, allow for more selections in the property sheets, better usability overall. Currently the JBoss Tools task form generation and visual designer do not work in concert, nor is the Admin console ready to be put in front of an end-user so he/she can work on his/her tasks (e.g. review & approve expense reports). Another way to say all of this is...focus more on the business analyst, give him/her better tools so they can do more for themselves and bypass the Java propeller heads.	1	1	https://issues.jboss.org/browse/JBPM-1924
14051	JBPM-1922	When does this 'BusinessExceptionFlow' have to start? I mean, what is the definition of the 'error' that can occur? And don't come up with 'everythingh' because I'll answer 42 then ;-) And what is 'without' java programming? Meaning that there may not be any actionhandler written (once!!) in java and configured (once!!) in each process? These are the most important parts to know. The rest is 'trivial' I think.... You can start any process from an actionhandler. Provide the full path to the node to it and a reference to the process id for the 'callback' or to retrieve additional data from it. First thing this businessExceptionFlow shoud do is to suspend the calling process. Then it can do anything it wants. At the end of this process you can end the calling process or resume it in the node it was or go to another node if that was passed in or whatever... So if you give an answer to the first two questions, I'll provide the rest and not bill them (or should I). Personally I think this is a great question for the userforum or JBoss consultancy and does not belong in the issues list (besides for keeping track of the customers question).	0	0	https://issues.jboss.org/browse/JBPM-1922
14052	JBPM-1919	I think everything should be fixed here. Can you please verify and close the Bug Thomas?	1	1	https://issues.jboss.org/browse/JBPM-1919
14053	JBPM-1919	Can you please point me to the exact incompatibility in the BatchSignalCommand? I don't find it at the moment...	1	1	https://issues.jboss.org/browse/JBPM-1919
14054	JBPM-1919	Hi Thomas, I double checked the BatchSignalCommand, I don't find any incompatibilities. Or does it make a difference if the method has moved up to the superclass? If that matters what would help is to introduce a method in the BatchSignalCommand which just delegates to super? This would be easy by using Eclipse but quite verbose in the code :-/ Every line marked with a "#" I found in the new class as well... 3.3.0 public class org.jbpm.command.BatchSignalCommand extends java.lang.Object implements org.jbpm.command.Command{ # public org.jbpm.command.BatchSignalCommand(); # public java.lang.Object execute(org.jbpm.JbpmContext) throws java.lang.Exception; # public java.lang.String getProcessName(); # public void setProcessName(java.lang.String); # public long getProcessVersion(); # public void setProcessVersion(long); # public java.lang.String getStateName(); # public void setStateName(java.lang.String); # public long[] getTokenIds(); # public void setTokenIds(long[]); # public java.lang.String getTransitionName(); # public void setTransitionName(java.lang.String); # public java.util.Date getInStateAtLeastSince(); # public void setInStateAtLeastSince(java.util.Date); } 3.3.1 public class org.jbpm.command.BatchSignalCommand extends org.jbpm.command.AbstractTokenBaseCommand{ public org.jbpm.command.BatchSignalCommand(); public java.lang.Object execute(org.jbpm.graph.exe.Token); public java.lang.String getTransitionName(); public void setTransitionName(java.lang.String); public java.util.Date getInStateAtLeastSince(); public void setInStateAtLeastSince(java.util.Date); public java.lang.String getAdditionalToStringInformation(); public org.jbpm.command.BatchSignalCommand transitionName(java.lang.String); public org.jbpm.command.BatchSignalCommand inStateAtLeastSince(java.util.Date); from AbstractTokenBaseCommand protected org.jbpm.JbpmContext getJbpmContext(); public java.lang.Object execute(org.jbpm.JbpmContext) throws java.lang.Exception; public abstract java.lang.Object execute(org.jbpm.graph.exe.Token); public void setTokenIds(long[]); public void setTokenId(long); public java.lang.String getAdditionalToStringInformation(); public java.lang.String getProcessName(); public void setProcessName(java.lang.String); public int getProcessVersion(); public void setProcessVersion(int); public java.lang.String getStateName(); public void setStateName(java.lang.String); public long[] getTokenIds(); public long getTokenId(); public java.lang.String toString(); public org.jbpm.command.AbstractTokenBaseCommand tokenIds(long[]); public org.jbpm.command.AbstractTokenBaseCommand tokenId(long); public org.jbpm.command.AbstractTokenBaseCommand processName(java.lang.String); public org.jbpm.command.AbstractTokenBaseCommand processVersion(int); public org.jbpm.command.AbstractTokenBaseCommand stateName(java.lang.String); }	1	1	https://issues.jboss.org/browse/JBPM-1919
14055	JBPM-1909	I committed the change to add the setActor to the AuthenticationService and added the factory. So it is at least consistent now, but I think is still room for improvement, but maybe this isn't done in jbpm 3 any more? Shall we close this issue?	1	1	https://issues.jboss.org/browse/JBPM-1909
14056	JBPM-1905	Some methods were named "setProcessId" or the like. Since this in ambiguous (is it ProcessDefinitionId or ProcessInstanceId) I set these methods to deprecated and created new ones with the right names.	1	1	https://issues.jboss.org/browse/JBPM-1905

#	Requirement	Questions	Answered	Successfully Answered	URL
14057	JBPM-1887	Okay, I could reproduce the behavior. If the ExceptionHandler itself throws an exception it can catch, it will hang in an infinitive loop. In my eyes this is a bug and I would resolve it for next jbpm version (even if I think it doesn't affect a lot of people). Unfortunately the CVS-history wasn't important in the new SVN and my "old" CVS access doesn't work any more (don't know why), so I cannot easily check what has changed and why it was not there in jbpm 3.2.GA. Anyway, I would say, an exception caused by an action executed by exception handler should NOT be handled by the jbpm exception handler. This could be implemented by checking if the exception is set on the ExecutionContext already: GraphElement 280: // If an exception is already set we are in an action of an exception handler // in this case don't give the exception to the handlers but throw it to the client if (executionContext.getException()==null) { // if an exception handler is available raiseException(exception, executionContext); } else { // if there is no parent we need to throw an action exception to the client if (exception instanceof JbpmException) { throw (JbpmException) exception; } else { throw new DelegationException(exception, executionContext); } } What do you guys think? Worked for my local test case. I could add the test case and commit this change if it doesn't break anything of the test suite... Thoughts?	1	1	https://issues.jboss.org/browse/JBPM-1887
14058	JBPM-1883	Comments made earlier today suggest that a deployer was never implemented, is this true? If this task did not cover the deployer then what did it cover?	0	0	https://issues.jboss.org/browse/JBPM-1883
14059	JBPM-1835	why is this closed? the fact that embedding the jsf forms does not work is just one option. I think it should be left open, the description a little changed and something new maybe develloped (embedded xforms?)	0	0	https://issues.jboss.org/browse/JBPM-1835
14060	JBPM-1822	thomas, can you have a look at this ? i can help you find the old release notes if that turns out a problem.	1	1	https://issues.jboss.org/browse/JBPM-1822
14061	JBPM-1821	Thomas, could you update the installer so that it also produces a db directory in the installation, containing the schema creation DDL scripts ?	0	0	https://issues.jboss.org/browse/JBPM-1821
14062	JBPM-1782	could this be due to the examples not having a retry on an optimistic locking exception ?	1	1	https://issues.jboss.org/browse/JBPM-1782
14063	JBPM-1767	Tom, could you please take care of this docs issue?	1	0	https://issues.jboss.org/browse/JBPM-1767
14064	JBPM-1764	Heiko, could you please work with Alejandro to resolve this?	0	0	https://issues.jboss.org/browse/JBPM-1764
14065	JBPM-1761	Can you elaborate on this? I don't understand what you mean.	0	0	https://issues.jboss.org/browse/JBPM-1761
14066	JBPM-1752	Tom Baeyens: > Let's fix this in jPDL 4.0 Thomas Diesler: > A deferred issue cannot have a fix version Then why not leave it just open and put a fix version of 4.0 on it? That way it does not get lost. Isn't that what versions are for? Now this issue is not visible anywhere anymore	1	0	https://issues.jboss.org/browse/JBPM-1752
14067	JBPM-1740	work going on here? Otherwise I take it....	0	0	https://issues.jboss.org/browse/JBPM-1740
14068	JBPM-1735	I couldn't reproduce it in eclipse, but only with maven. That is weird. That seems to indicate that some parts of the environment are not the same. Have you got an idea how that could have happened ? Maybe explaining the maven profiles and how we use it to test it on different DB's is a good topic for the team meeting...	0	0	https://issues.jboss.org/browse/JBPM-1735
14069	JBPM-1728	Thomas, Do you have specific examples of what you mean by stale content? Just 2.0 stuff or also the 3.1 stuff? Now jboss.org has a new wiki, is it still the intention to move to a dyndns.org wiki? Personally I'd think not, but you guys might think otherwise	0	0	https://issues.jboss.org/browse/JBPM-1728
14070	JBPM-1724	What will we do with SEAM support in general ? It's going to be hard to avoid. If we do support SEAM, how will the test suite look like. Should the jPDL/pageflow tests be in jBPM or should they be in SEAM ? Thinking of it, probably pageflow should move to seam as that is hard for us to test cause it (the page node) has a direct dependency on SEAM iirc. Apart from that there is the timer transaction problem: JBPM-1166	1	1	https://issues.jboss.org/browse/JBPM-1724
14071	JBPM-1724	Alejandro, you can re-introduce the seam pageflow test, but then you have to make sure that it succeeds. I believe that a full pageflow test should be in SEAM, not in jBPM. If we can test the specific aspect that caused JBPM-1707, then we could consider writing a test for only that specific aspect. If that aspect is only reproducable in SEAM, then the test should be in the SEAM codebase, not ours. WDYT?	1	1	https://issues.jboss.org/browse/JBPM-1724
14072	JBPM-1720	Thomas, can you have a look at this ? I have fixed the configuration problem. Now the underlying issue pops up. java.lang.NoSuchFieldError: name at org.slf4j.impl.Log4jLoggerAdapter.<init>(Log4jLoggerAdapter.java:75) at org.slf4j.impl.Log4jLoggerFactory.getLogger(Log4jLoggerFactory.java:75) at org.slf4j.LoggerFactory.getLogger(LoggerFactory.java:92) at org.apache.jackrabbit.core.config.RepositoryConfig.<clinit>(RepositoryConfig.java:64) at org.jbpm.jcr.jackrabbit.JackrabbitJcrServiceFactory.getRepository(JackrabbitJcrServiceFactory.java:48) at org.jbpm.jcr.impl.AbstractJcrServiceFactory.openService(AbstractJcrServiceFactory.java:46) at org.jbpm.svc.Services.getService(Services.java:156) Which seems to be an slf4j classloading problem according to http://www.velocityreviews.com/forums/t635509-hibernate-tutorial-problem.html	0	0	https://issues.jboss.org/browse/JBPM-1720
14073	JBPM-1710	Can you elaborate on your externalization proposal? You mean, set the pageflow schema resource in jbpm.cfg.xml or something similar?	0	0	https://issues.jboss.org/browse/JBPM-1710
14074	JBPM-1471	Same here. I have running subprocess, and when the parent process is ended (using end-complete-process="true"), subprocess is ended, token returns to parent process, and jbpm tries to execute it normally (causes usually unexpected behavior, e.g. something wrong with hibernate proxies). Don't see it being any special case or anything, just subprocess in a wait state, and in meantime parent process is ended. Is there any reason to execute token, if it has ended (end not null)? cvl	1	1	https://issues.jboss.org/browse/JBPM-1471
14075	JBPM-1471	not sure, how throwing an exception would help here. In my view, the process _usually_ ends, when the subprocesses are _not_ ended yet. Would it be possible for token to simply _not_ follow the execution, if it has ended? At least when it is signaled at returning from subprocess. I see throwing an exception would be valid, if token is explicitly signaled (e.g. through external trigger), but what about simply ending process, and having not ended subprocesses? Regards.	0	0	https://issues.jboss.org/browse/JBPM-1471
14076	JBPM-1460	[12:01] <tomb> porcherg: thanks. what do you think about this: https://jira.jboss.org/jira/browse/JBPM-1460 [12:01] <tomb> ? [12:01] <tomb> can i remove class ObjectReference and use Descriptors instead ? [12:02] <tomb> os do you see a downside of that [12:02] <tomb> or do you see a downside of that [12:04] <porcherg> tomb: i think it can be removed [12:04] <tomb> ok [12:04] <porcherg> maybe there will be a db performance impact [12:04] <tomb> i'll do it in the next couple of days [12:04] <tomb> should be minimal [12:04] <tomb> or even for the better [12:05] <porcherg> because the descriptor table will have a lot of entries [12:05] <tomb> with an index on a long column that is not a problem [12:05] <tomb> index is only maintained during deployment [12:05] <porcherg> but this will but easier to understand [12:05] <tomb> don't think so [12:06] <tomb> as most ObjectReferences will have a descriptor memberfield anyway [12:06] <porcherg> but this will be easier to understand [12:06] <tomb> so it removes an extra indirection [12:06] <porcherg> i think having only one Descriptor is easier [12:06] <tomb> ok	1	1	https://issues.jboss.org/browse/JBPM-1460
14077	JBPM-1451	Hi Trevor, what are the Tier-1 DBs?	0	0	https://issues.jboss.org/browse/JBPM-1451
14078	JBPM-1451	Hi Len, could you please resolve the dependent issue	1	1	https://issues.jboss.org/browse/JBPM-1451

#	Requirement	Questions	Answered	Successfully Answered	URL
14079	JBPM-1450	Thomas, Resolving this issue simply requires a change in the jBPM configuration. I understand SOA-P uses its own jBPM and Hibernate configuration files. We should change those files directly. Julian, can you confirm whether SOA-P uses a separate jBPM configuration Thomas, Resolving this issue simply requires a change in the jBPM configuration. I understand SOA-P uses its own jBPM and Hibernate configuration files. We should change those files directly. Julian, can you confirm whether SOA-P uses a separate jBPM configuration file? If so, please let me know how can I change it. Alternatively I can post the required changes in a comment to the Jira issue. -Alejandro El lun, 24-11-2008 a las 09:08 +0100, Thomas Diesler escribió: > > Hi Alejandro, > > > > this issue https://jira.jboss.org/jira/browse/JBPM-1450 is still open, > > but CP03 has been tagged already. What do we need to do about it? > > > > cheers > > -thomas > >	1	1	https://issues.jboss.org/browse/JBPM-1450
14080	JBPM-1446	How can you get a TimeZone from a Locale? The country alone is not enough to determine the exact time zone, as some countries span multiple time zones. I was thinking of parsing the Date HTTP header. I reviewed the Seam source code and did not (yet) find anything besides TimeZone.getDefault(). Any suggestions welcome...	1	1	https://issues.jboss.org/browse/JBPM-1446
14081	JBPM-1423	standard java ========== 1) environment block controls the jdbc transaction through hibernate. hibernate being one resource in the standard transaction. 2) user controls the jdbc transaction with hibernate and environment block has got to participate in that same tx. Here, the tx will have a bigger scope then the environment block. So special care should be taken to verify all usages of transaction-end-listeners (like synchronisations) 3) user controls the transaction with spring and environment block has got to participate in that same tx. Same comment: tx scope is bigger then environment block scope. can we find a way so that the API tests are all executed in each of these 3 deployment configurations ? enterprise java ========== In enterprise Java, we can always assume that the transaction is controlled by JTA and that the transaction has a bigger scope then the environment block. The most common case will be CMT controls the transaction and pvm is addressed through the ExecutionService local SLSB osgi ==== testing only one of the tx configurations in osgi will be sufficient. To create a bundle out of PVM, the	1	1	https://issues.jboss.org/browse/JBPM-1423
14082	JBPM-1395	Tom, could you please reassign appropriately.	1	0	https://issues.jboss.org/browse/JBPM-1395
14083	JBPM-1375	But will jBPM suffer of the same problem? Mauro.	1	0	https://issues.jboss.org/browse/JBPM-1375
14084	JBPM-1375	Sorry, I meant: will jBPM4 suffer of this problem regarding the determination of the node type? Mauro.	1	0	https://issues.jboss.org/browse/JBPM-1375
14085	JBPM-1373	i don't think a special instance needs to be set up. JBPM.3 and JBPM.3_pgsql82 are already available. isn't this just a matter of enabling the publication checkbox in the configuration page of the hudson jobs ?	0	0	https://issues.jboss.org/browse/JBPM-1373
14086	JBPM-1372	what is the purpose of the installer ? i see 3 potential targets: 1) deploy a jbpm installation into a version of jboss. what are all the things that the installer should do? * add a jBPM queue to the JMS message service ? * add a jPDL DB * populate jPDL DB with example processes * add identity DB * populate identity DB with example users * configure a security domain * put the deployment artifact (the war?) in the deploy folder this is what we now have in an ant target somewhere. given that this is part of the simply unzippable suite distribution, i don't really see the urgent need for this. 2) install the jpdl-suite in a file system, which includes the jpdl lib, sources, docs, a preconfigured jbpm server and the eclipse plugin currently this is just unzipping the distribution. so i don't see the need for an installer here 3) installing the gpd into an existing eclipse installation. i think the eclipse update mechanism would be better for that. currently we have a feature distribution package that can just be unzipped in an eclipse. Koen is now working together with Kris on a joint new designer. That seems to be working great. In the meantime, the problem with the milestone versions of the designer dependencies has disappeared. So in this new designer, the standard eclipse versioning mechanism can and will be used.	0	0	https://issues.jboss.org/browse/JBPM-1372
14087	JBPM-1355	jBPM does not install duplicate jars {code} [tdiesler@tdvaio jboss-4.2.2.GA]$ find server/default/deploy/jbpm/ -name "*.jar" server/default/deploy/jbpm/jbpm-enterprise.jar server/default/deploy/jbpm/jbpm-service.sar/jboss-bpm-api.jar server/default/deploy/jbpm/jbpm-service.sar/jbpm-jpdl-core.jar server/default/deploy/jbpm/jbpm-service.sar/jbpm-jpdl-identity.jar server/default/deploy/jbpm/jbpm-service.sar/jbpm-jpdl-integration.jar server/default/deploy/jbpm/jbpm-console.war/WEB-INF/lib/jsf-facelets-1.1.14.jar server/default/deploy/jbpm/jbpm-console.war/WEB-INF/lib/commons-io-1.4.jar server/default/deploy/jbpm/jbpm-console.war/WEB-INF/lib/jbpm-jbpm4jsf-3.3.0-SNAPSHOT.jar server/default/deploy/jbpm/jbpm-console.war/WEB-INF/lib/gravel-1.0.0.GA.jar server/default/deploy/jbpm/jbpm-console.war/WEB-INF/lib/commons-fileupload-1.2.1.jar server/default/deploy/jbpm/jbpm-console.war/WEB-INF/lib/jsf-facelets-1.1.11.jar {code} The jar ordering in jboss is defined, so a potential duplicate would be hidden anyway. In case we have duplicates and have an incompatibility reported we deal with it on a case by case basis. Automating a scan as part of the Hudson QA is also possible but a superfluous effort at the moment ls -1 lib lib/endorsed server/default/lib | grep jar | sort	1	1	https://issues.jboss.org/browse/JBPM-1355
14088	JBPM-1352	> Yes. Are you using Unified EL now? Or another EL flavour... i've plugged in jsr 223 with scripting.dev.java.net. in the configuration, you can configure a named list of expression languages and also indicate which is the default expression language. in the jpdl language, for each expr attribute, there is an optional lang attribute. > Possibly - I'm on PTO quite a lot in May, but I think we can sort > something out. How about 18th onwards? ok. my expectation is that you drive the seam integration. i'll be ready to supply any jbpm update that you'll need and re-upload a snapshot. does that work for you ? regards, tom. Pete Muir wrote: > > On 14 Apr 2009, at 14:50, Tom Baeyens wrote: > >> the initial work needs to be done by us, i think. >> >> implementation wise there are a couple of directions, but there is still a bit of thinking involved. >> >> mostly, the current integration is based on integrating the naming resolutions while resolving UEL expressions. apart from that there is only the special SEAM annotations, which will boil down to jbpm service invocations. >> >> * setup of the seam contexts. find the proper places where we need to intercept to set up the seam contexts. this could potentially be very easy. we have a set of services now (==session facades). we also have a mechanism to add interceptors to those. a seam interceptor could ensure that seam is properly initialized around jBPM operations. > > Great, sounds easy enough :-) > >> * expanding the jbpm variable resolver so that it also delegates to the seam contexts. > > Yes. Are you using Unified EL now? Or another EL flavour... > >> * tx around the jbpm timers and messages. we should be able to leverage our container integration. since we'll be deployable as a service, we'll build our own tx interceptor for managing the JTA tx. this is leveraged when using seam. > > Kev has been doing some work on this - but I agree, we should just use JTA here - Seam will join or start a TX as required. > >> >> * pageflow. build the page activity. create a pageflow parser. >> >> this is all not cristal clear though. requires getting our hands dirty to see if this plans are complete. i think we should plan a week where we together set out the initial structure. then we can involve productization people to finish the integration. > > Ok. > >> ... >> >> regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-1352

#	Requirement	Questions	Answered	Successfully Answered	URL
14089	JBPM-1319	* Downloads * Documentation * Jira * Forums * jBPM Users * jBPM Developers * see if the other forums can be deleted or deprecated * would be great if we add usage instructions like: when should people use forums vs mailing lists. * Mailing list * jbpm-dev (is that one public or internal only) * dev-jbpm * cvs-commits (sf.net) * other sf.net mailing lists * Sources * jbpm3 : jpdl3 & identity * jbpm4 : pvm, jpdl, task, identity * Wiki's	0	0	https://issues.jboss.org/browse/JBPM-1319
14090	JBPM-1319	what is the difference between jbpm-dev and dev-jbpm (I cannot seem to find the latter) Would be nice though if a small text was added where the forums are for: jBPM developers is for people developing jBPM itself, not developing their own apps WITH jBPM.	1	1	https://issues.jboss.org/browse/JBPM-1319
14091	JBPM-1301	Koen, can you take this one ? A patch file should be extracted with eclipse that shows contains the diffs between jbpm-3.2.3 and the current head. That patch file should be assigned to this jira issue. That shouldn't be much. But it might be convenient for Alejandro to apply his fixes again. Then all the files need to be deleted. Then a readme needs to be added to the root jbpm.3.	1	1	https://issues.jboss.org/browse/JBPM-1301
14092	JBPM-1299	what is meant with this ?	1	1	https://issues.jboss.org/browse/JBPM-1299
14093	JBPM-1296	this is a component used in the console. so same comment as the console. only in this case, testing it in isolation will even be harder as JSF components are even harder to test then web pages. and no one knows what those JSF components are supposed to do.	0	0	https://issues.jboss.org/browse/JBPM-1296
14094	JBPM-1292	Hey Thomas. I am a bit confused by now. Where is the current code base hosted? Not in CVS any more, right? Can you give me the right direction, then I try to have a look into it...	1	1	https://issues.jboss.org/browse/JBPM-1292
14095	JBPM-1291	it is a community extension to our core jBPM product. we should define a way how we deal with such extensions. something like incubation maybe ?	0	0	https://issues.jboss.org/browse/JBPM-1291
14096	JBPM-1285	Guillaume, can you describe the proposal you have to fix it ?	1	1	https://issues.jboss.org/browse/JBPM-1285
14097	JBPM-1280	an important part of this task is the systematic approach. meaning that i don't want tests to be created ad hoc cause then we'll never know what we test and what not. that doesn't mean that all possible combinations have to be tested. but i want that we at least thought about which combinations we want to see tested. by having an overview that has a systematic approach, we can get a good appreciation of the coverage of the test suite for the job executor	1	1	https://issues.jboss.org/browse/JBPM-1280
14098	JBPM-1271	Are we sure we want to keep using jbpm4jsf as it is? Big dependency on Gravel, which is not the easiest to learn. The jBPM 4 console will be seam based, maybe with some custom JSF components. If jBPM 3.2.4 should be mavenised, than this should as well, but It will be only short lived I think (hope)	0	0	https://issues.jboss.org/browse/JBPM-1271
14099	JBPM-1269	Could you please help me to resolve this?	0	0	https://issues.jboss.org/browse/JBPM-1269
14100	JBPM-1227	Guillaume, there are still some tests commented in the HibernateJobDbSessionTest. Can you evaluate if those are still relevant or if they should be removed ? Just execute as you think it should be done. Once that is done, this issue can be reloved. All methods are now tested except for the base DbSession methods void flush(); void forceVersionUpdate(Object entity); void lockPessimistically(Object entity); I created JBPM-1684 for that.	1	1	https://issues.jboss.org/browse/JBPM-1227
14101	JBPM-1223	This proposal was already taken in PVM. Not sure it is worth backporting, tough. There aren't many votes for it, though the community may not be aware of it. Which brings another question: how can we communicate with the community about issues for which we would like input/votes?	0	0	https://issues.jboss.org/browse/JBPM-1223
14102	JBPM-1212	is this something more than variable declaration?!	1	1	https://issues.jboss.org/browse/JBPM-1212
14103	JBPM-1212	I think the static process variables is something like processDefinition id, key, name or version. And everything could be shared between all of process instances. E.g, we could declare partnerLinks in processDefinition on BPEL environment, then all of process instances could refer them. Furthermore I think the differences between static process varaible and normal process variable is the static process variable is readonly. Any ideas?	1	1	https://issues.jboss.org/browse/JBPM-1212
14104	JBPM-1210	Hi Tom, this weekend we were working trying to solve Jira#PVM-44 (spring configuration support). At the moment we are trying to learn wich part of your configuration framework has to be done with Spring. We first did an SpringEnvironmentFactory extending EnvironmentFactory. About this we have this first impressions: * There is some coupling with all the parser lifecycle, meaning that in Spring there is no such concept (at the high level, of course). As we understood the intention is that we must make our own Parser that receives the config file location, parse it until we get the SpringEnvironmentFactory. This is not posible without going to deep into the spring architecture. In spring you just create an XmlApplicationContext instance from the file and you ask it for defined beans. The solution was to make an static method SpringEnvironmentFactory#parse(InputStream) -just to follow your conventions, but using this concept in spring is not really usual- that accomplish this task but ignoring all the other way. * We defined a Context in the Environment containing a hibernate session. At first time we tried to make a SimpleContext just implementing the Context interface and keeping it as pojo as possible, to be able to fill it with Spring. We made it and succesfully register the HibernateSession on it. But the problem began when we tried to get it by class (type). The hibernate session get registered as SessionFactoryImpl and in the code it is invoked as get(SessionFactory.Class). Digging into the code we realize that all this registrations had to be done by a layer of Descriptors through the WireContext. Now the problem is that all this layer is not really so much pojo-oriented in the sense that there are no getters and setters for every property. * Maybe the whole package org.jbpm.pvm.internal.wire.* should be avoided in the Spring configuration. But when we try to get this we realize that there are some coupling between the usage of services and the way that they get 'injected' and even this approach will translate in a lot of duplications in the code. Have you any suggestion about this topics? Beside how we resolve this, is necesary to add in your ivy repository the Spring dependencies. At the moment we have <dependency org="springframework" name="spring-context" rev="2.5.4" conf="compile->default" /> <dependency org="springframework" name="spring-core" rev="2.5.4" conf="compile->default" /> <dependency org="springframework" name="spring-beans" rev="2.5.4" conf="compile->default" /> <dependency org="springframework" name="spring-orm" rev="2.5.4" conf="compile->default" /> <dependency org="springframework" name="spring-jdbc" rev="2.5.4" conf="compile->default" /> <dependency org="springframework" name="spring-tx" rev="2.5.4" conf="compile->default" /> Thanks. Mauricio Salatino - Diego López León	1	1	https://issues.jboss.org/browse/JBPM-1210

#	Requirement	Questions	Answered	Successfully Answered	URL
14105	JBPM-1209	TODO: creating tests that use an expression for each situation where it is possible. E.g. * an expression that refers to a method that implements behaviour of a node * an event listener expression. on every possible event. * transition conditions. create an activity that takes the first transition for which the condition evaluates to true * setting timers dynamically based on processvariables * choosing a specific subprocess runtime instead of designtime * maybe event listeners should get a condition as well ? separately, the name resolution should be checked in all of these situations: * can the expression 'see' process variables * can the expression see the process definition properties * can the expression see the named environment and environment-factory objects * can the expression easily navigate a process domain object (stored as process variable) other ideas ?	1	1	https://issues.jboss.org/browse/JBPM-1209
14106	JBPM-1205	I would prefer if they could be renamed to XxxTestArchived. That way, mvn will not pick them up and I will still have them available for archive purposes when looking for "how did i do that before. I know i have done that in the past somewhere..."	0	0	https://issues.jboss.org/browse/JBPM-1205
14107	JBPM-1194	First I think the subject is incorrect. EL expressionevaluation *is* in mail templates, just not in the attributes of the variable value element. Correct? Besides that, if you want *specific* subjects, you can set those in the template and override the default one. The fact that the default currently only is static is not a big deal imo. The advantage only comes in when using complex EL in many places. Thirdly, I'd make a totally different patch and change the mailvariableresolver to parse the value of the value attribute as an EL and not do it in the getters as you do now	1	0	https://issues.jboss.org/browse/JBPM-1194
14108	JBPM-1179	Did you commit only to HEAD ? Or also to the SOA branch ? Who's going to do the test on WebLogic ? Who are we waiting for ?	1	1	https://issues.jboss.org/browse/JBPM-1179
14109	JBPM-1179	Only to head. Should all work on the enterprise module (including this issue) also be ported to the soa branch? Martin is testing on WebLogic.	0	0	https://issues.jboss.org/browse/JBPM-1179
14110	JBPM-1172	hmm... I see that the template is only resolved run-time (not deploy time), so it is not possible to detect a non existent template and to a fall back. Would it be better to change this to do a deploy time check? I personally think so.	1	1	https://issues.jboss.org/browse/JBPM-1172
14111	JBPM-1171	Hi, I just switched to 4.4 recently and found out that the patch is not there and in the trunk either. What happened?	0	0	https://issues.jboss.org/browse/JBPM-1171
14112	JBPM-1148	Hi. What is the current status on this? Anybody did some work here already? Maybe I could have a look at it but want to avoid duplicate work. And Fix Version 3.3 is too late I think, so target release should be 3.3.1? The proposed solutions is a good direction but still has some flaws (see forum entry), but a working solutions in terms of classloading is vital, especially if you think about jBPM ESB integration, scoping of deployments and so on...	1	1	https://issues.jboss.org/browse/JBPM-1148
14113	JBPM-1148	The problem with the proposed fix was that it leads to an endless loop, since the classloader is already needed to load the jbpm.cfg.xml in the JbpmConfiguration (InputStream jbpmCfgXmlStream = ClassLoaderUtil.getStream(resource);), but this wants to read which classloader to use. Exception StackTrace: at org.jbpm.util.ClassLoaderUtil.getClassLoader(ClassLoaderUtil.java:50) at org.jbpm.util.ClassLoaderUtil.getStream(ClassLoaderUtil.java:83) at org.jbpm.JbpmConfiguration.getInstance(JbpmConfiguration.java:279) at org.jbpm.JbpmConfiguration.getInstance(JbpmConfiguration.java:257) at org.jbpm.JbpmConfiguration$Configs.getObjectFactory(JbpmConfiguration.java:418) at org.jbpm.JbpmConfiguration$Configs.hasObject(JbpmConfiguration.java:426) at org.jbpm.util.ClassLoaderUtil.getClassLoader(ClassLoaderUtil.java:50) So my solution is to introduce a special method "getJbpmConfigurationStream" which is responsible especially for loading the jbpm.cfg.xml. It looks first in the ClassLoaderUtil.class.getClassLoader (because this is the "old behavior") and if there is nothing found in the Thread.currentThread().getContextClassLoader(). This method is used by the JbpmConfiguration and the XmlUtil to load jbpm.cfg.xml. And in the ObjectFactoryImpl I changed to lazy load the classloader, otherwise it is queried too early (before the jbpm.cfg.xml is processed)... Additionally the described patch is applied, but currently I am not sure if we still need the customClassLoader configuration? Shouldn't all use cases be already solved by using the context class loader and overwrite the ProcessClassLoader? Third thing I did was to introduce a ProcessClassLoaderFactory which can be configured via the jbpm.cfg.xml. So you can easily exchange the ProcessClassLoader to whatever you need in your project. The three new configuration parameters are: <jbpm-configuration> ... <string name="jbpm.classLoader" value="jbpm" /> <string name="jbpm.customClassLoader.className" value="com...MyClassLoaderUtil" /> <bean name="jbpm.processClassLoader" class="org.jbpm.instantiation.DefaultProcessClassLoaderFactory" singelton="true" /> /> All these values have defaults (which doesn't change the current behavior), so they are not required. All tests are still running, I added a short new test, which also works... I will commit it, so everybody interested can review it... The whole conifguration / load during startup / Config part is a bit tricky, didn't understand each and everything there, would be nice if you, Tom, can have a look at it? If it keeps like that I can also make a short note in the documentation about these classloading configuration possibilities...	1	1	https://issues.jboss.org/browse/JBPM-1148
14114	JBPM-1148	Bernd, what is the status of this?	1	1	https://issues.jboss.org/browse/JBPM-1148
14115	JBPM-1144	Here is the description of the duedate-timer attribute: "the duration (optionally expressed in business hours) that specifies the time period between the creation of the timer and the execution of the timer. See Section 15.1, "Duration" for the syntax." So I thought in the workflowdefinition I'm only able to input a constant value for the duedate. (e.g. "7 days"). But suppose, I want to send a report on the first day of the next month. How can I formulate this in the workflow? I found no solution. So I customized the Class "CreateTimerAction". And with the code above, I'm able to set the timer to the first day of the next month. Thanks in advance, Thorsten	1	1	https://issues.jboss.org/browse/JBPM-1144
14116	JBPM-1137	Alejandro Just wanted to confirm a few things. Running test.jpdl runs these 905 tests. Is that all? [junit] Running org.jbpm.AllTests [junit] Tests run: 905, Failures: 0, Errors: 0, Time elapsed: 88.214 sec Also we test patches by building from the release branch and then replace with the attached patch jar, we do not use the patch branch (jpdl_3_2_2_JBPM-1137) directly. The test target kept rebuilding the jpdl jar replacing the patched one that i copied, hence i had to disable the build part in the build files. For patch testing, we do not cover multiple databases (this is covered as part of SOA platform testing), unless the patch directly affects database. Can you tell me if there is any way or testcase to test this leap year date calculations?	0	0	https://issues.jboss.org/browse/JBPM-1137
14117	JBPM-1127	" Are you saying that the tests will provide us with zero additional (DB-specific) test coverage? " exactly	0	0	https://issues.jboss.org/browse/JBPM-1127

#	Requirement	Questions	Answered	Successfully Answered	URL
14118	JBPM-1117	This is what I get from this discussion: * You can use the ant deploy target, providing it with a configuration file that points to a jBPM DB. * The jBPM Console also contains a config file pointing to the jBPM DB. * So if you update one of those configuration files without the other then the deployed process is not seen in the console. Would the following be a solution ? * Make sure that a jBPM installation always has a deployment servlet * That way, the ant task and the designer can both deploy to a URL * So there is only 1 configuration of the database. And the ant task and the designer only need to be configured with a URL We'll take this into account in building jPDL 4. I don't think we can still prioritize this for jBPM 3.	1	0	https://issues.jboss.org/browse/JBPM-1117
14119	JBPM-1116	do you have a unittest (or a patch for an existing unittest) that demonstrates the problem?	1	0	https://issues.jboss.org/browse/JBPM-1116
14120	JBPM-1089	Is this issue resolved?	1	1	https://issues.jboss.org/browse/JBPM-1089
14121	JBPM-1089	I've the test suite running on hudson. So many thanks for fixing enterprise module! There are some jpdl failures because of wrong version included into SOA-P RC1 so please ignore them. But there are some enterprise module failures as well. Could you look at http://hudson.qa.jboss.com/hudson/job/jboss-soa-platform-JBPM.3-testsuite/59/testReport/ where the two org.jbpm.msg test failures can be seen. For the related server logs, please look under "Build Artifacts" on http://hudson.qa.jboss.com/hudson/job/jboss-soa-platform-JBPM.3-testsuite/59/ Let me know if you need any additional information.	0	0	https://issues.jboss.org/browse/JBPM-1089
14122	JBPM-1079	why is this major?	1	1	https://issues.jboss.org/browse/JBPM-1079
14123	JBPM-1066	Hi. Does this mean this issue has been fixed and the requested functionality exists ? Or that the issue will not be resolved ? Thanks, Ross	1	0	https://issues.jboss.org/browse/JBPM-1066
14124	JBPM-1060	afaik, I already added this feature reqeust in CVS, can you verify?	0	0	https://issues.jboss.org/browse/JBPM-1060
14125	JBPM-1040	like the colspan and rowspan in html tables?	0	0	https://issues.jboss.org/browse/JBPM-1040
14126	JBPM-1016	Is there a new issue made than to track this? Just closing and mentioning it will be in 4.0 is plain wrong.	0	0	https://issues.jboss.org/browse/JBPM-1016
14127	JBPM-998	simplifying, updating and improving the main jbpm.org website is on its way; it's one of my highest priorities atm. apart from that, i just double checked and there is little i can do to remove the SF CVS repo. i disabled cvs and forums. Hopefully that already helps.	1	1	https://issues.jboss.org/browse/JBPM-998
14128	JBPM-993	great... could you also make a small unit test for this? and maybe one with a CustomExpressionAssignmentHandler? Then I'm prepared to check this in.	0	0	https://issues.jboss.org/browse/JBPM-993
14129	JBPM-993	Actually according to http://jira.jboss.com/jira/browse/JBPM-429 this issue has been already fixed in 3.2.2. But I don't see any changes neither in 3.2.2 nor in FishEye. It would be nice to rely on jira status and jBPM documentation. Please write me what kind of unit test are you expecting?	0	0	https://issues.jboss.org/browse/JBPM-993
14130	JBPM-986	Can you comment on this. I looked at the schema and cannot see what you mean	1	0	https://issues.jboss.org/browse/JBPM-986
14131	JBPM-977	How about making the EL pluggable, so that you can retain back compatibility, but Seam 1.3 can use jboss-el?	1	1	https://issues.jboss.org/browse/JBPM-977
14132	JBPM-977	Thanks Tom! Can I get a new build?	0	0	https://issues.jboss.org/browse/JBPM-977
14133	JBPM-955	so the same way as you can map variables to sub processes, you want to map swimlanes to sub processes. would that cover your feature request ?	1	1	https://issues.jboss.org/browse/JBPM-955
14134	JBPM-955	> so the same way as you can map variables to sub processes, you want to > map swimlanes to sub processes. would that cover your feature request ? I guess you're referring to for example: <process-state name="initial interview"> <sub-process name="interview" /> <variable name="a" access="read,write" mapped-name="aa" /> ... </process-state> ...which could then be enhanced to allow for <process-state name="initial interview"> <sub-process name="interview" /> ... <swimlane name="candidate" /> ... </process-state> Is that what you propose? If so: - maybe some others would like to set an explicit actor, not the swimlane (so that would need an <assignment> mapping as well) - if a sub-process is invoked from within another sub-process then one may not know the hardcoded swimlane name ("candidate" in the example above, or "client" and "employer" in my first post) So I guess it would cover my current needs, but I'm not sure if the above would yield the most flexible solution. And it can probably just as easy be implemented using some custom component: <task name="survey"> <assignment actor-id="#{myCustomTaskAssignerComponent.surveyActor.id}"/> </task> Maybe it would be more flexible to either have jBPM search for an assignment (swimlane or actor, whichever is found first) anywhere up in the calling tree, or to make the current swimlane somehow available in expressions. As for the latter, one could then simply use: <task swimlane='#{jbpmSwimlane}' /> If this is doable, then I guess others might like the current actor to be available in expressions as well. Does this answer your question? Thanks, Arjan.	1	0	https://issues.jboss.org/browse/JBPM-955
14135	JBPM-955	" <process-state name="initial interview"> <sub-process name="interview" /> ... <swimlane name="candidate" /> ... </process-state> Is that what you propose? " yes. but also with a mapped-name that allows for translation of the swimlane instance to be stored under a different name in the sub process instance. i still don't get quite what you propose. with looking upward. in general, i'm very sceptical about references maintained between process instances. we'll probably end up in more problems then we actually solve. the proposal that i sketched makes sure that information is copied in from and potentially extracted from the sub process. so the two process instances are independent.	1	1	https://issues.jboss.org/browse/JBPM-955

#	Requirement	Questions	Answered	Successfully Answered	URL
14136	JBPM-955	> also with a mapped-name that allows for translation of the swimlane > instance to be stored under a different name in the sub process instance. A mapping is indeed good to have, as a sub-process might of course define multiple tasks, each of which could have its own swimlane (or assignment) as well. Sorry for not getting it after your first comment, and excuse me for giving yet another example just to be sure. So, I assume you're proposing: <process-definition name="hire"> <swimlane name="candidate" /> <swimlane name="colleague" /> <swimlane name="manager" /> ... <process-state name="initial interview"> <sub-process name="interview" /> <!-- First interview: assessment by future colleague --> <swimlane name="colleague" mapped-name="interviewer" /> <swimlane name="candidate" mapped-name="interviewee" /> ... </process-state> ... <process-state name="second interview"> <sub-process name="interview" /> <!-- Second interview: assessment by future manager --> <swimlane name="manager" mapped-name="interviewer" /> <swimlane name="candidate" mapped-name="interviewee" /> ... </process-state> ... <process-definition> ...using <process-definition name="interview"> <!-- NOTE: no need for expression here, such as "#{interviewer}" -- right? --> <swimlane name="interviewer" /> <swimlane name="interviewee" /> ... <task swimlane="interviewer"> ... </task> <task swimlane="interviewee"> ... </task> ... </process-definition> Right? That looks good! It would also allow for sub-processes within sub-processes. Maybe a future GPD could even ensure the swimlanes are set by the calling process-state. The above would surely cover my needs and is much better than my proposal, which I couldn't even make clear so you should forget about ;-) If you're indeed proposing the above for <swimlane>, then what about doing the same for <assignment> -- would that make sense at all, and would it be doable? I guess this would only make (some) sense for pooled-actors, but would look much better using swimlanes: <process-state name="initial interview"> <sub-process name="interview-using-actor" /> <assignment pooled-actors="colleagues-pool" mapped-name="interviewer-pool" /> <assignment actor-id="candidate-actor" mapped-name="interviewee-actor" /> <assignment expression="previous --> group(hierarchy) --> member(boss)" mapped-name="supervisor" /> ... </process-state> ...and <process-definition name="interview-using-actor"> ... <task> <!-- NOTE: unlike the swimlane, using "#{..}" here --> <assignment pooled-actors="#{interviewer-pool}" /> ... </task> <task> <!-- NOTE: unlike the swimlane, using "#{..}" here --> <assignment actor-id="#{interviewee-actor}" /> ... </task> ... </process-definition> There might be better ways to do this (such as some custom Java classes to set the variable value), and probably most is covered using swimlanes like you propose? Thanks for reading this ;-) Arjan.	1	0	https://issues.jboss.org/browse/JBPM-955
14137	JBPM-952	currently, we don't plan to have configurable out of the box support for authorization. i wan't to provide a more generic solution to this. you could build it yourself by making use of the identity component, though. but then you need to tweak the web console. you could define an authroization-group for each process. all members of that group can then start the process. you might create a link between the authorization group and the process by a name convention on the group. e.g. the group has groupType process-authorization and the group is called [process-name]-authorization-group or something like that. in one of the later versions we might work out some thing similar out of the box, but i want to take the time to analyse this carefully	0	0	https://issues.jboss.org/browse/JBPM-952
14138	JBPM-891	david, i remember that you worked on this. can you give an update and a pointer to the solution if there is one already. i think i forgot to mention that this should go somewhere into the jpdl/dist/src/main/resources/runtime/release.notes.html	1	1	https://issues.jboss.org/browse/JBPM-891
14139	JBPM-890	ronald, could you please have the patience to wait until AFTER my afternoon nap before you close this task ? otherwise i might loose track and do a nap twice, which would be a waste of that effort.	0	0	https://issues.jboss.org/browse/JBPM-890
14140	JBPM-889	who the hell starts using our jira for typing lessons ?!? :-)	0	0	https://issues.jboss.org/browse/JBPM-889
14141	JBPM-850	this seems already committed to cvs. Why is it still open? Does it need a test?	1	1	https://issues.jboss.org/browse/JBPM-850
14142	JBPM-848	Fady, i don't get your remark. As far as I remember, it should be possible for people to specify the property jbpm.mail.from.address in their jbpm.cfg.xml file. Can you try to reproduce and compare the problem report of Sebastian with the docs and the code ? In case our code is the problem, how do you think we should fix it ?	0	0	https://issues.jboss.org/browse/JBPM-848
14143	JBPM-845	subContextInstance.setTransientVariables(superContextInstance.getTransientVariables()); That line is in the process state since 3.1.2. Does that resolve your requirements as well ? In your proposed solution, there is a lot of overhead in the DB schema and the XML schema. So i prefer the approach that is currently taken: Just transfer *all* the transient variables *allways*. What do you think ?	1	1	https://issues.jboss.org/browse/JBPM-845
14144	JBPM-845	Have a look at my first and 4th comment: there I am pointing out, that this solution is too simple. I'll try to explain it by example: Given: A process, that creates a "Contact" in the database (A), taking a transient variable called "contact". A process, that creates a "Domain" in the database (B), taking a transient variable called "domain" and 4 transient variables called "contactAdminc", "contactTechC", "contactZoneC","contactHolderC". Now, every Domain needs to create the contacts. Simple approach: process B uses 4 times process A in a processState to create Contacts (like a function call). The problem: With the actual implementation this is not possible, since a transient variable must have the same name in super- and subprocesses, there is no variable scoping! Do you see the limits of the actual implementation? For us this is a no-go limit... Regarding the overhead: sure ... it uses the same datastructures like persistent variables and that's a lot. But is this really a problem? Concerning to the necessary changes in the DB schema and the XML schema: new Feature --> new XML stuff, new DB stuff - that's the way it goes :-) I think, the only real problem is that there is no backward compatibility. May only help writing good release notes... a making this an optional feature (how?)	1	0	https://issues.jboss.org/browse/JBPM-845
14145	JBPM-838	you mention a number of good improvements. but please consider the following for each of the individual points you make: * Currently it's a feature of jBPM to be runnable on a 1.4.2 JVM. will this still be the case after your proposed change ? * How will the library dependencies of jBPM be impacted ? There are many different deployment environments: from embedded into a user application to standalone on an app server. * What is the actual benefit for the jBPM users ? please, for any of the points that you think must still be addressed in jBPM 3.x, create a new, separate issue and explain previously mentioned aspects of the issue. thanks.	1	1	https://issues.jboss.org/browse/JBPM-838
14146	JBPM-835	Jeff, there are 2 issues about library dependencies for examples. As the number of examples will grow, so will the libraries needed for the examples. If we include them all in the plain download, the size will become bigger, only for the examples. The advantage of course is that it can run out of the box. What about consistently fetching the libs for an example into the lib directory in the build script of the example itself ?	1	0	https://issues.jboss.org/browse/JBPM-835
14147	JBPM-820	I thought, along with this feature, you 'fixed' passing a variablemap when starting a process as well? Correct? Then http://jira.jboss.com/jira/browse/JBPM-412 can be closed as well.	0	0	https://issues.jboss.org/browse/JBPM-820

#	Requirement	Questions	Answered	Successfully Answered	URL
14148	JBPM-810	isn't this a very clumsy way of starting to build a scripting language in jbpm ? maybe we should revert to a real scripting language then ?	1	0	https://issues.jboss.org/browse/JBPM-810
14149	JBPM-808	hmmm... someone care to explain this? Why not use a decision node for this and not go into the tasknode at all...? Now the designer has to be extended to show this condition as well.... I'm not saying that I'm against this, but do not see the advantage immediately, unless it is for ad-hoc tasks that are not part of a process.	1	0	https://issues.jboss.org/browse/JBPM-808
14150	JBPM-807	Tom, do you consider this proposal as setting variables *inside* the expression? To me, it looks simpler: evaluate the expression and set the result as the value of the named variable.	0	0	https://issues.jboss.org/browse/JBPM-807
14151	JBPM-807	pete, can you elaborate on this feature request a bit ? does it mean that variables should be settable through a value binding expression ? or does this just translate to a setVariable on the business process context ?	0	0	https://issues.jboss.org/browse/JBPM-807
14152	JBPM-799	Cool that lots of the webconsole issues have been done, compliments. Wouldn't it be better to assign a 'fix version/s' to these so they will show up on the releasenotes?	0	0	https://issues.jboss.org/browse/JBPM-799
14153	JBPM-770	Ok, I forgot to update the Hibernate mapping for processstate, I've done so as such: <property name="signalProcess" column="SIGNALPROCESS_" /> I think it's working now - is there anything else I need to do?	1	1	https://issues.jboss.org/browse/JBPM-770
14154	JBPM-760	Jeff i didn't find this issue when resolving the problem yesterday. I just did a good effort of trying to add many indexes on each many-to-one where i thought it was possibly being used in a query. I think i have many more indexes defined then you have specified. If you have a chance, could you have a look and see in the scripts if the once i created make sense ?	0	0	https://issues.jboss.org/browse/JBPM-760
14155	JBPM-743	I do not agree, what is the use of swimlanes then at all? Same thing can be achieved with direct assignments, making swimlanes an unneeded special feature.	1	1	https://issues.jboss.org/browse/JBPM-743
14156	JBPM-743	I do not agree either, and agree with Ronald's assertion that in that case swimlanes themselves are a special feature that isn't required. To my mind if a swimlane represents an actor/role in a process, there's a good set of use cases for that actor's physical human changing during the running of a long running process. In that case isn't it still the same swimlane being used? As I read a process (or for that matter generate a nice html picture/report) from a given process isn't a swimlane actually a far more accurate representation of what I'm attempting to accomplish? Furthermore there's no need to force a jPDL author to include an extra attribute in their jPDL - evaluate="always" would cause the behavior I'm seeking, and if the attribute didn't exist the current behavior would be used. Existing processes still parse and work as before, but a valueable feature is added, the lack of which makes for a messier and potential less accurate representation of a process.	0	0	https://issues.jboss.org/browse/JBPM-743
14157	JBPM-721	Can i contribute with this? i want to test and create new webservices functionallity.. could you recommend me something that is missing right now?	0	0	https://issues.jboss.org/browse/JBPM-721
14158	JBPM-720	Hi..Diego Leon and I wanna contribute with this (because we have the time).. I was testing jbpm in 3.2 in JBoss 4.2.2.GA and it works fine.. I was thinking to implement this WS Interfaces as a separate War Maven project using JBoss WS 2.0, with this we are able to use JAX-WS and make compatibility with Metro and CXF.. In the main branch of the repository (jbpm.3) we only find a few classes about this under jpdl/ws.. this is all about this? Ronald can you contact me in private? or we discuss this here? Just in case my mail is salaboy at gmail dot com Thanks	1	0	https://issues.jboss.org/browse/JBPM-720
14159	JBPM-718	This issue still active? Can i help you with something.. I am very interested in this.. but i do not know the state of this issue..	0	0	https://issues.jboss.org/browse/JBPM-718
14160	JBPM-713	Ryan, can you elaborate on this? You mean scheduling them and having them fired all at once? That would mean the scheduler has to be multi threaded... Maybe a good moment to try to switch to quartz for the scheduling	1	1	https://issues.jboss.org/browse/JBPM-713
14161	JBPM-708	can you expand on "Each time the server is restarted, a new version is deployed." ? it's the user's code that should call the, jbpmContext.deployProcessDefinition(). so you should look for the fix where this is called when the server is restarted. does that help you further ?	0	0	https://issues.jboss.org/browse/JBPM-708
14162	JBPM-705	Ryan, Can you comment on what functionality you expect in this area? In the webapp there is not a lot of functionality at the moment, so only duplicating that (without the visible display of the processdiagram) gives you little info. IMO we should expose a JMX interface for some things, but the webapp itself should be the only display functionality we provide. OTOH, there is something like http://www.camunda.com/toolkit_for_jbpm/toolkit_for_jbpm.html which looks nice	1	1	https://issues.jboss.org/browse/JBPM-705
14163	JBPM-705	Triage:::: Tom, can this be closed? Or should the subject change to: " expose management info via jmx"	0	0	https://issues.jboss.org/browse/JBPM-705
14164	JBPM-704	is there anything missing in these docs ? http://docs.jboss.org/jbpm/v3/userguide/deployment.html#thirdpartylibraries	1	0	https://issues.jboss.org/browse/JBPM-704
14165	JBPM-699	Prabhat, I thought I had similar issues, but atm I do not seem to have them using cvs-head. Can you point out what version of jbpm you are trying to use.	0	0	https://issues.jboss.org/browse/JBPM-699
14166	JBPM-699	jbpm 3.1 branch?	0	0	https://issues.jboss.org/browse/JBPM-699
14167	JBPM-675	In TaskInstance.java, i replaced the resume() method with: public void resume() { if(end == null) { // Only re-open the task if the task was not previously completed isOpen = true; } isSuspended = false; } The isSuspended=false; could be moved up inside the if..block, but since the suspend() method is still marking ended tasks as suspended, I figured might as well unsuspend them. Question goes back to this original bug. Do we want to error on suspending a closed taskinstance? Or just send a warning to the log and do nothing when a closed task is suspended?	0	0	https://issues.jboss.org/browse/JBPM-675
14168	JBPM-675	works on 3.2, should this still be fixed in 3.1? I assign it to 3.1.5, and see what happens	1	0	https://issues.jboss.org/browse/JBPM-675
14169	JBPM-672	What is an action controller? Did I mis something?	1	1	https://issues.jboss.org/browse/JBPM-672
14170	JBPM-672	Do I need repost? -- ---------- RODRIGO CARVALHO REZENDE ----------	1	0	https://issues.jboss.org/browse/JBPM-672
14171	JBPM-661	The solution for this issue was developed, but it was never merged into the jbpm.3 codebase due to its dependency on jdk 1.5. How to deal with that?	1	0	https://issues.jboss.org/browse/JBPM-661
14172	JBPM-653	Would this mean that if the jBPM engine is called from a servlet (either jsf or webservice) the web context classloader would be used? That would be great, then versioned classes in the process can have access to files in a war file... nice..	0	0	https://issues.jboss.org/browse/JBPM-653
14173	JBPM-633	on what is this statement based: "tokens are locked in memory; not at a database level" ? afaik, we always use (optimistic) db locking. we plan to switch to quartz in the future for execution of asynchronous stuff. but no timing has been worked out yet. also we don't plan to overcome the limitations that come with maximum 1 command-executor-thread. JMS implementations can be used in case you want to remove those limitations.	1	1	https://issues.jboss.org/browse/JBPM-633

#	Requirement	Questions	Answered	Successfully Answered	URL
14174	JBPM-633	Would it be possible to attach a process definition example and action handler java file that demonstrates this behavior? That would really complement the nice description above.	0	0	https://issues.jboss.org/browse/JBPM-633
14175	JBPM-633	Sorry I've taken a while to get back to you. The statement about tokens being locked in memory refers to the call to token.lock() within the executeAction method of org.jbpm.graph.def.GraphElement. JBPM Javadoc states that token.lock(): "locks a process instance for further execution. A locked token cannot continue execution. This is a non-persistent operation. This is used to prevent tokens being propagated during the execution of actions." However my ActionHandler's execute method starts a new thread on which to perform some long-running process without holding up the single CommandExecutorThread. This new thread also needs to propogate the token and so opens a new JbpmContext and hence a new hibernate session from which to get the token and propogate it; the token will never appear to be locked to this new session (as it is a non-persistent operation). There is however a chance that my new thread will try to do exactly what the token.lock() call is trying to prevent; i.e. the token being propogated during the execution of an action. It's a race condition; it depends on whether the execute method of the action completes before my new thread creates a new JbpmContext and propogates the token. There is also the chance that my new thread gets a copy of the token from a new session before the action's execute method has completed and saved its copy, making the token in the new thread out of date. I don't have a problem with there being a single CommandExecutorThread; what I do need to do is to be able to perform long-running actions that propogate the token without holding up the single CommandExecutorThread from executing other actions. This is achievable by starting a new thread from within the execute method of the ActionHandler and creating a new JbpmContext on the new thread and progresssing the token from within the new context. The problem is that this will sometimes lead to the optimistic locking issues referred to where one thread is left holding an out-of-date copy of the token; hence the solutions that I have suggested. Is there another way of achieving this behaviour where the CommandExecutorThread can execute long-running actions asynchronously so that it can immediately start to process the action specified in next message? I can attach a process definition and action handler that show this behaviour but the behaviour is, as you say, to be expected due to optimistic locking - what I am suggesting is a feature that allows me to wriyte asynchronous actions handlers without needing to worry about optimistic locking - I would have expected this to be a common requirement given the single CommandExecutorThread. I'm afraid that we haven't currently implemented the work-arounds that I suggested as we have leaned towards using our own bespoke command executor thread and message implementation with a view to maybe replacing these with a jBPM solution once this becomes available. [I noticed futher optimistic locking issues of the same nature between the SchedulerThread and the CommandExecutorThread and, more importantly, between the CommandExecutorThread and the client thread - I have reported these in JIRA issue JBPM-626. I think that these latter two issues can be catagorized as defects rather than feature requests.] Thanks for your help so far - I hope that this response has clarified my previous explanation. Mark Shotton	1	0	https://issues.jboss.org/browse/JBPM-633
14176	JBPM-633	I have created a process flow application which can be used by multiple users concurrently. As these users may decide to save the work in various states and it may take a long time like a week to reach the end point - I am trying to use asynchronous continuations. As there is only a single CommandExecutorThread on which all async ActionHandler execute methods are called for different processes what should we do to ensure all the threads are safe? Also, will this be an issue to have so many threads in wait state for so long? I appreciate any feedback Thanks Prasad	1	0	https://issues.jboss.org/browse/JBPM-633
14177	JBPM-617	should the processdefinition.xml by default be in the DB? and should the above config filter it out if needed? Looking at the implementation of FileArchiveParser and FileFilter, this is not implemented yet. Reopening....	1	1	https://issues.jboss.org/browse/JBPM-617
14178	JBPM-616	I don't get this issue, can anybody elaborate?	1	1	https://issues.jboss.org/browse/JBPM-616
14179	JBPM-615	What is the status for this issue? I use already my port of JBPM-libs as Maven2-driven modules. Just the problem is to synchronize it with generic JBPM-updates. I found already in the latest update (3.1.4) some POMs... ...but it is only basic build without tests, reports, e.t.c... ...The first reason to port it to maven - my server platform. I use Geronimo. Currently my "JBPM-Brach" is compatible with geronimo-1.1. Most actual Geronimo utilize fully Maven infrastructure. It will be nice to get more deep integration of JBPM on Geronimo. What are the plans to solve this issue? It is really not a big step to get "MAVEN_INSIDE" :) Current structure of projects alredy good designed. Its needs just minor "refactoring"... So, I will appreciate any Reply.	0	0	https://issues.jboss.org/browse/JBPM-615
14180	JBPM-605	Behaviour? Should selecting a task automagically assign it to 'me' or not. Does there need to be an additional 'assign to me' element in the list? If so what is the behaviour of the ui? go to act upon the task then? or refresh and see the personal tasklist where the just assigned task can be on a completely different page (in case of paging). What does the behaviour of 'cancel' mean in a task that is acted upon from the group task list? Does it go to his/her personal tasklist? or does there need to be an additional button besides save and cancel that says assign to me? and what about the save button? automagically to personal tasklist? with maybe a 'reminder' of this? or back to the group tasklist? Please comment.....	1	0	https://issues.jboss.org/browse/JBPM-605
14181	JBPM-599	only with input from the process archive? What do you want to put in there then? there are things like: - dates only in the past, in the future, more than 18 years ago, between ? - regular epressions? How far do you want to go? And if you do. why not put a variable type in there to. Would make a great enhancement in the automagically generated forms.	0	0	https://issues.jboss.org/browse/JBPM-599
14182	JBPM-599	ok, but I suspect you want to configure this in the processdefinition and in that case not each time a processvariable is used in a task. So you are probably thinking of introducing something like there was in 2.0 (or was it 1.0) to have a global definition of variables and a reference to them from the individual tasks? Or do you just want to do it in the forms and not in the PD at all? OR a kind of a roadmap to first do it in the forms and later in the pd?	1	1	https://issues.jboss.org/browse/JBPM-599
14183	JBPM-583	promoted to highest feature request priority. it will be delivered on the 3.2, but also on the 3.1 branch. the actual versions of delivery are indicative.	0	0	https://issues.jboss.org/browse/JBPM-583
14184	JBPM-582	Quote from Gavin: "So, we need to make pageflow composition work. As I understand it, you are resolving composed process definitions on the database. I have all process definitions in a Map. Can you make it possible for me to plug in a ProcessDefinitionResolver somewhere?"	1	1	https://issues.jboss.org/browse/JBPM-582

#	Requirement	Questions	Answered	Successfully Answered	URL
14185	JBPM-582	Yes, you're right. What about public interface SubProcessResolver { ProcessDefinition findSubProcess(org.dom4j.Element subProcessElement, ProcessDefinition superProcessDefinition); } That way, the plugged in implementations could also take the content of the sub-process element into acount. During implementations, mind self referencing sub processes. In case of the SEAM pageflow, the subprocesses have to be deployed before superprocesses. This implies that bidirectional references are impossible. In the case of jPDL, we plan to support this by means of offering binding capabilities in the jBPM webapp admin console. regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-582
14186	JBPM-581	Is there any way you might be able to suggest an alternative if this isn't something that can be done? I am trying to create a presentation layer for JBPM (where transitions are represented by buttons and states are represented by different screens). Obviously this is for task notes. Do you have any suggestion as to how I could record data that is meaningless to the workflow definition but meaningful to an engine presenting the workflow? Any help would be greatly appreciated. Thanks. Joe	1	0	https://issues.jboss.org/browse/JBPM-581
14187	JBPM-580	Is this implemented in the process designer? If not, when will it be?	0	0	https://issues.jboss.org/browse/JBPM-580
14188	JBPM-578	leave them in, but keep them deprecated? or 'undeprecate' them? Or remove them in 3.3 or 4.0? Still, shouldn't we remove the usage of these in jbpm itself...	1	1	https://issues.jboss.org/browse/JBPM-578
14189	JBPM-578	deprecating them means to express the intend to remove them in the not so distant future. I would therefore not 'won't fix' them, but reopen the issue and do a 'fix version of 3.3 or 4)	0	0	https://issues.jboss.org/browse/JBPM-578
14190	JBPM-569	koen, could you have a look at this ?	1	0	https://issues.jboss.org/browse/JBPM-569
14191	JBPM-569	Could this be the problem of the trailing blanks?	0	0	https://issues.jboss.org/browse/JBPM-569
14192	JBPM-541	Tom, can you go ahead and upload the jar which you wish to work with into the repository? I've made the necessary changes locally, but would like to go ahead include and your jar before into the build script before i commit the changes.	1	1	https://issues.jboss.org/browse/JBPM-541
14193	JBPM-541	Also, don't check out the entire repository. It's huge. Just check out a single folder (e.g. jbossws) which you can use as a template.	1	1	https://issues.jboss.org/browse/JBPM-541
14194	JBPM-541	alex, what do you need from the jboss jbpm integration ? afaik, the BPEL extension does not use the jbpm service archive, right ? what integration between jbpm and jboss must work in order for you to get started with the bpel-jboss integration ? regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-541
14195	JBPM-541	The BPEL extension does not use the jBPM service archive directly; instead, it reuses the MBean with a different service name, has different service mappings and some other libraries specific to the BPEL extension. Regarding the integration, all I need is to agree with you on where to place the libraries, dependencies and configuration files for jBPM and BPEL without creating conflicts. Before, we were planning to add two service archives to the "all" configuration of the appserver. Now we can either share a single server configuration and deal with the conflicts, or simply have two separate custom server configurations. The latter option seems easier to me. What do you guys think?	0	0	https://issues.jboss.org/browse/JBPM-541
14196	JBPM-541	Ruel, can you give us a pointer to an example of a script that builds a server configuration in the jboss test environment ? We have such a script, but don't know how this should be integrated into the jboss test stuff. Alex, i think that we should create 1 jbpm server configuration, which is based on the all configuration + jBPM jPDL + jBPM BPEL. then we should be able to run the whole jboss test suite to make sure that our jbpm deployments don't break anything. then we should start building our own test suite. regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-541
14197	JBPM-537	can you provide a test case that shows the problem ? preferrably with the process definition inside of the test case. regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-537
14198	JBPM-509	Koen, can i assign this task to you ? can you take the lead in this ? regards, tom.	1	0	https://issues.jboss.org/browse/JBPM-509
14199	JBPM-492	Oops, wrong comment. Issues related, since paging becomes more difficult if sorting other than the primary key (if it is ordered at all) comes into play. I personally have no idea how to solve this. - by querying the database and have it sort the result set - keeping the list in memory and have custom sorters while keeping the result set limited to e.g. 100 entries - what about if the user wants to sort on a different column? Re-query the db? I think so - .... If anyone has any experience with this, please let us know.	1	1	https://issues.jboss.org/browse/JBPM-492
14200	JBPM-490	The BPEL extension doesn't conform to the LGPL terms. Should I add the headers anyway?	1	1	https://issues.jboss.org/browse/JBPM-490
14201	JBPM-479	is this still an issue in 3.2?	1	1	https://issues.jboss.org/browse/JBPM-479
14202	JBPM-459	A solution for this would be to use the key <taskname:transitionname> in the formfile and use it the same as with presenting the form for the task itself. Shall I look into this?	1	1	https://issues.jboss.org/browse/JBPM-459
14203	JBPM-459	not sure if we are talking about the same thing here. i meant that after a task is completed, the process moves to the next state. the feedback i'm talking about is new assignments of tasks, new wait states and process completion. i don't see yet how a taskname:transitionname would help. could you elaborate ?	1	1	https://issues.jboss.org/browse/JBPM-459
14204	JBPM-454	Is this likely to be added as a feature in the near future? Cheers. Adrian	1	1	https://issues.jboss.org/browse/JBPM-454
14205	JBPM-454	Is there any documentation for this feature? Is it supported in the jPDL designer for Eclipse? Thanks.	1	0	https://issues.jboss.org/browse/JBPM-454
14206	JBPM-448	Hi, Just checking to see if there is anything I can do to help you certifiy your project information pages on Labs. Is there something that I can help you with or something that needs more explaining or instructions? Let me know. Cheers, Damon	1	0	https://issues.jboss.org/browse/JBPM-448
14207	JBPM-441	the namespace has been set to urn:jbpm.org:jpdl-3.1 so you can use the following ways to refer to the jpdl schema: <process-definition xmlns="urn:jbpm.org:jpdl-3.1" name='...' > in case you have you map the urn:jbpm.org:jpdl-3.1 namespace to the http://jbpm.org/jpdl-3.1.xsd URL in your local catalog or <process-definition xmlns="urn:jbpm.org:jpdl-3.1" xmlns:xsi ="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="urn:jbpm.org:jpdl-3.1 http://jbpm.org/jpdl-3.1.xsd" name='...' > in case you want to include the reference to the actual URL that matches the namespace. regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-441
14208	JBPM-436	Oops, wrong comment. Issues related, since paging becomes more difficult if sorting other than the primary key (if it is ordered at all) comes into play. I personally have no idea how to solve this. - by querying the database and have it sort the result set - keeping the list in memory and have custom sorters while keeping the result set limited to e.g. 100 entries - what about if the user wants to sort on a different column? Re-query the db? I think so - .... If anyone has any experience with this, please let us know.	1	0	https://issues.jboss.org/browse/JBPM-436
14209	JBPM-426	Would be great to have this for 3.3. Can anyone estimate the amount of work?	1	1	https://issues.jboss.org/browse/JBPM-426
14210	JBPM-423	Tiles I agree. (not completely but enough), but is this dependend on Tomahawk? Aren't those nice to have jsf components? date entry etc... Assign this one to me. I can remove it in a lost minutes.	1	1	https://issues.jboss.org/browse/JBPM-423
14211	JBPM-412	It is in the ProcessInstance class, but should this be added to JbpmContext as well?	1	0	https://issues.jboss.org/browse/JBPM-412

#	Requirement	Questions	Answered	Successfully Answered	URL
14212	JBPM-403	gavin, you mentioned that you could provide me the code for the following 3 functions: 1) determine wether an object is an ejb3 entity or not. 2) extract the id and the classname from an ejb3 entity 3) obtain an ejb3 entity from the classname and the id for this issue, i am waiting on this input. what library dependencies does the provided code have ? regards, tom.	1	0	https://issues.jboss.org/browse/JBPM-403
14213	JBPM-403	Tom, do you need any further help on this one?	0	0	https://issues.jboss.org/browse/JBPM-403
14214	JBPM-403	what are the library and JDK version dependencies ? i don't know yet what the impact will be on the current jbpm project structure (which needs to be revised, btw) i didn''t have time now. 3.1 release is really under pressure. it can be scheduled for one of the next releases. how soon it can be done depends on how easy it is to set up tests in the current jbpm project structure. if, in the mean time you want to add this to the seam code base, i can help you with customizing the configuration. regards, tom.	1	0	https://issues.jboss.org/browse/JBPM-403
14215	JBPM-399	i added the last line to org.jbpm.graph.node.ProcessState: [code] // feed the readable variableInstances if ((variableAccesses != null) && (!variableAccesses.isEmpty())) { ContextInstance superContextInstance = executionContext.getContextInstance(); ContextInstance subContextInstance = subProcessInstance.getContextInstance(); subContextInstance.setTransientVariables(superContextInstance.getTransientVariables());[/code] can you make the same update to the beta3 sources and check if it works. if you could write a simple test case that i could add to our test suite, that would be great. regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-399
14216	JBPM-297	How can I get this to work? I have a similar issue with 3.1.1. I notice that SessionFactory.openSession is called without regards to any 'current' session. Hibernate3 provides a standard method on SessionFactory, getCurrentSession(), which can work in a variety of ways: - transaction-aware with Spring (by LocalSessionFactoryBean) - open session in view filter with Spring - CurrentSessionContext (custom plug-in if you are managing thread-local session yourself) It seems like it should be pretty simple for jBPM to take advantage of this.	0	0	https://issues.jboss.org/browse/JBPM-297
14217	JBPM-275	Can't the taskcontroller be configured to accept all kinds of configuration? Including variablenames etc... The thing you have to do is implement processing of variables yourself, but that is why it is a custom taskcontroller. Ronald	1	0	https://issues.jboss.org/browse/JBPM-275
14218	JBPM-258	hmm... isn't that like reinventing the wheel? the portal has something like that. Can't we reuse that in one way or another for real crud? For a simple version use a user.xml, group.xml or whatever file to configure it.	1	1	https://issues.jboss.org/browse/JBPM-258
14219	JBPM-255	Was not assigned to the webapp, but to nothing. This is i18n, internationalization, just translating text, labels etc for webapp specific issues. The ability to define e.g. multiple labels for a variable in the processdefinition is not solved by this. That should be another issue. l10n (localization) is not taken care of, but I think we should. e.g. dates should be presented in the way they are in that specifica region/language. Who defines the priority for this? I put it on major since we should take it into account from the beginning.	0	0	https://issues.jboss.org/browse/JBPM-255
14220	JBPM-255	i'm happy to see that the current i18n is based on plain JSF without any dependencies on struts or any other technology, right ? but i would like to know the target we want to achieve with i18n. as you mention, we don't have support for dates or enumerated types translations. my concern is that making all this translatable will make the webapp over complex. And the question i have is : is it worth to introduce a bit of complexity for the plain translation as we currently have, knowing that we won't be able to support full translation of all items... i'm in doubt here. regards, tom.	1	0	https://issues.jboss.org/browse/JBPM-255
14221	JBPM-255	Why don't we create samples of internationalization, let's say adding french and some eastern language? An earlier issue tackled the handling of eastern encoding, we might combine the two issues together and get it sorted out	1	0	https://issues.jboss.org/browse/JBPM-255
14222	JBPM-228	Is this still open for discussion? Would make a nice addition.... the Java type should imo be a 'render-hint' and not a real type. The real type can be gotten from the instance of the variable.	1	1	https://issues.jboss.org/browse/JBPM-228
14223	JBPM-228	2 minor issues with the proposed proposal :) 1) it is one specific gray value on the scale between black and white. let me try to explain. there's guis that do a lot of generation of the form, based on a few input parameters. that is typically not so flexible. on the other hand, there's jsp's: a lot of information to be specified but more flexible for form layout and parameter types. the proposed solution is one specific version somewhere in the middle. 2) dependency on tomahawk. i intend to remove it (see other jira issue for reasons. search for tiles & tomahawk) let me propose a counterproposal: we could give each task a form element. the content of the form element is free text or elements (as with the any stuff). then in the runtime environment, a form renderer can be specified. a form renderer takes the task's form element and the task instance as input. then it has to render the task form. maybe the task renderer could be a servlet or something like that... your proposal could then be one implementation of such a task form renderer. what do you think. ? regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-228
14224	JBPM-228	lets start simple and not configure it anywhere but let it depend on the type in the DB. If someone wants a specific variable to have a certain type, set it in the process-start event? btw, how can I access this 'any' type? is it stored somewhere?	1	1	https://issues.jboss.org/browse/JBPM-228
14225	JBPM-228	Is this feature supported in 3.2.2? If yes, where could I find docs about it? Thank you.	1	0	https://issues.jboss.org/browse/JBPM-228
14226	JBPM-228	what feature? The any content? The issues is closed, so yes	0	0	https://issues.jboss.org/browse/JBPM-228
14227	JBPM-220	Do you realy want this? I do, but I can imagine that others will not be so happy since it leads to exclusion of jbpm running on a plain old servlet engine or the abilty to use it standalone. People should get a warning when a process using async is used from a system that does not support jms.	0	0	https://issues.jboss.org/browse/JBPM-220
14228	JBPM-185	Is this separate xml file a new one or can we reuse forms.xml? I think the config files should not get to fragmented	1	1	https://issues.jboss.org/browse/JBPM-185
14229	JBPM-185	can you explain how you want to implement this ? regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-185
14230	JBPM-183	can you update the schema ?	1	1	https://issues.jboss.org/browse/JBPM-183
14231	JBPM-165	slipping this to 3.3? Or will it realy still go into 3.2? I would let it slip...	0	0	https://issues.jboss.org/browse/JBPM-165
14232	JBPM-165	Bernd, could you commit that ? I want to have a look and see if we can base an experimental implementation on that.	1	1	https://issues.jboss.org/browse/JBPM-165
14233	JBPM-165	how do you handle the bidirectional reference between the old and the new process instance version ?	1	1	https://issues.jboss.org/browse/JBPM-165
14234	JBPM-165	Regarding the current version of ChangeProcessInstanceVersionCommand Does it make sense to use newDef.getTaskMgmtDefinition().getTask(oldTask.getName()); instead of // find new task Query q = jbpmContext.getSession().getNamedQuery("TaskMgmtSession.findTaskForNode"); q.setString("taskName", oldTask.getName()); q.setLong("taskNodeId", newNode.getId()); This will fix a problem when tasks are defined in the process level instead of a task node level. In addition, does it make sense to set the task management definition in the migrated process to the new definition? Similar to the below: pi.getTaskMgmtInstance().setTaskMgmtDefinition(newDef.getTaskMgmtDefinition());	0	0	https://issues.jboss.org/browse/JBPM-165
14235	JBPM-158	Partly solved by JBPM-816.... The top level of this issue is core, one subtask is core and one subtask is GPD... How do we resolve THIS issue?	1	1	https://issues.jboss.org/browse/JBPM-158

#	Requirement	Questions	Answered	Successfully Answered	URL
14236	JBPM-130	can you check again. or submit the body of a test case that shows the problem ? regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-130
14237	JBPM-125	What should be the approach for multiple tasks for different users? Should they all operate on the same variables? Should they get their own local ones? Should we detect if they post and someone else allready has? We should think about these issues a little and at least describe the behaviour.	0	0	https://issues.jboss.org/browse/JBPM-125
14238	JBPM-113	Magnolia seems to be using jsr170 to, and are going to include a workflow engine in the next release. I think it is openwfe (wrong choice ;-)) JSR-170 should (imo) be used in the CMS of the portal as well. I already made a note for that in the portal forum. Couldn't we make it a joint effort with the portal guys? Security has to be solved in an integrated way. The CMS editor of Magnolia is great btw... I've suggested it to my employer early this year as a valid replacement of something called smartsite. Unfortunately they (management) seem to be moving in the M$ direction. :-(Ronald	1	1	https://issues.jboss.org/browse/JBPM-113
14239	JBPM-108	email support replaced a generic notification framework. future versions of the identity component (merged by that time with portal's identity component?) could store notification type preferences. depending on that, we could send out notifications over IM, email, SMS, ...	0	0	https://issues.jboss.org/browse/JBPM-108
14240	JBPM-71	realy? can't we make it more explicit? use an attribute on the variable. This is old fashioned imo	1	0	https://issues.jboss.org/browse/JBPM-71
14241	JBPM-46	If the license permits it, I think we should look at Quartz again. It has clustering features. With a few classes that should be doable. How far is jbpm 3 on that part. Is the basic scheduler node there and only the backend not implemented?	1	1	https://issues.jboss.org/browse/JBPM-46
14242	JBPM-46	the fact that it supports clustering doesn't bring us new features, i think. we only use one transactional resource (jdbc connection) so that is easy clusterable. but we should make sure that every job is only executed once. i was thinking of using the jboss singleton for that. i think the main beef of quartz is the fact that you can express complicated functions of time a la cron (only time). this is something we don't need because our jobs are always related to process events rather then time. another aspect is that we must be able to control the classloader because the classes that are scheduled could be versioned. although this was one of the reasons why i thought quartz would be a problem, now i think this would be easy solvable. how much KB (MB?) does it require in required runtime jars ? but i'm always open for this discussion. we should not be doing what already has been done. why do you think we should use quartz ? regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-46
14243	JBPM-46	I notice this one is marked Done but there are no Subversion commits logged. Could you point me at what has been done here as I'm just about to attempt a Quartz implementation.	1	0	https://issues.jboss.org/browse/JBPM-46
14244	JBPM-46	The SchedulerServlet was available in 3.1.4 but no longer seems to be in 3.2.GA, am I missing something? The basis for my Quartz implementation will be to create an implementation of SchedulerService and SchedulerServiceFactory, can you confirm that is a sensible approach? Thanks, Matt.	1	0	https://issues.jboss.org/browse/JBPM-46
14245	JBPM-30	Did all work except for the ANT tasks.	0	0	https://issues.jboss.org/browse/JBPM-30
14246	JBPM-9	I guess this means that all workflow patterns (www.workflowpatterns.com) will be implemented? When I looked at the source I saw already some constructs (milestone, interleaveStart). Are these towards the workflowpatterns? Maybe, when all is done, a report could be created showing how each pattern is implemented in jBpm. I'm happy to assist in this :-)	0	0	https://issues.jboss.org/browse/JBPM-9
14247	JBPM-9	I am trying to find out information about workflow patterns supported by jBPM but i cant get success. Can u please give me the answers of following questions. 1. Is jBPM supports all the workflow patterns (www.workflowpatterns.com) ? 2. if no then which are patterns not supported by jBPM. -- Vijay	1	0	https://issues.jboss.org/browse/JBPM-9
14248	JBPM-2	Is this still a BLOCKING issue? If not, could it please be marked Closed? Those red issues really scary admins trying to use jBPM in production.	0	0	https://issues.jboss.org/browse/JBPM-2
14249	JBPM-2	for jBPM 3 this means that it is only possible to have an extra page of graphics inside the current console, directly based on the logs. is the current plan any different from that approach ?	0	0	https://issues.jboss.org/browse/JBPM-2
14250	JBPM-1	can you explain what you mean with adding components ? what are components and where do i find them ? regards, tom.	1	1	https://issues.jboss.org/browse/JBPM-1
14251	BPEL-313	Aleksander, the jBPM code has been migrated to Subversion, including the BPEL runtime. jBPM BPEL 1.1.1 was released from SVN. Would you rather incorporate your work into our SVN repo if you had write access to it?	1	1	https://issues.jboss.org/browse/BPEL-313
14252	BPEL-311	A proof of concept patch. This patch enables caching of VariableDefinitions. It has cut a single process execution's JVM time by 20% for us, as indicated by JProfiler. Note that I've used a ConcurrentHashMap, which requires Java 1.5. This is because I'm not certain whether findVariable can be called concurrently on the same CompositeActivity. Feel free to change it to ordinary HashMap if that won't cause problems. However, the patch needs verifying for correctness. Is it OK to cache VariableDefinitions, or can they get stale over time? What will happen on hot redeployment of BPEL process definition? Will CompositeActivity instances be dropped together with their VariableDefinition caches?	0	0	https://issues.jboss.org/browse/BPEL-311
14253	BPEL-305	Looks like Charles Charles Souillard from Bull S.A.S. has already implemented this for jBPM-BPEL 2: http://fisheye.jboss.org/browse/JBPM/lightning/bpel2/org/jbpm/bpel/xml/query/xpath/jaxen/DOMDoXslTransformFunction.java?r=1.1 Which unfortunately has never been released. Can this be easily imported to jBPM-BPEL 1 branch?	0	0	https://issues.jboss.org/browse/BPEL-305
14254	BPEL-268	Hi, wondering on that status of this. Evaluating jBoss jBPM and (ofcourse) human tasks came up in the discussion. Where can I find this example? Thanks!	1	0	https://issues.jboss.org/browse/BPEL-268
14255	BPEL-244	Has there been any progress on this during last 2.5 years?	0	0	https://issues.jboss.org/browse/BPEL-244
14256	BPEL-193	Burr Sutter wrote: "Would you mind adding an "all" task that runs all the steps necessary for Hello World to build, gen, pack, deploy, etc? It is easier to demo where I change the .bpel file and simply run "All" to go through the whole process."	0	0	https://issues.jboss.org/browse/BPEL-193
14257	BPEL-84	Ok, how can I run these tests?	1	1	https://issues.jboss.org/browse/BPEL-84
14258	BPEL-84	The procedure you describe is not relevant to this issue, it uses the axis based implementation in jboss-4.0.x What I am asking for is: [quote] Please provide initial integration tests at jboss-head/webservice/test/java/org/jboss/test/ws/integration/bpel There should be an test-bpel ant target that runs as part of the tests target. [/quote]	1	1	https://issues.jboss.org/browse/BPEL-84
14259	BPEL-59	Hi, i'd like to make some editions to this request. in order to monitor the process flow i need some how to identify entities inside BPEL xml file and that identifier might be entity's name. but inside BPEL names are not unique. each of the JBPM events come per node definition ID. Is it possible to return BPEL xml file with little addition like real system id of the entity inside JBPM. it would be very helpful, because this way i can know exactly which entities have been already executed.	1	1	https://issues.jboss.org/browse/BPEL-59
14260	GPD-395	weird, I just did clean build and then apply this patch and it works as expected. I run it on MacOSX could this be the issue?	0	0	https://issues.jboss.org/browse/GPD-395
14261	GPD-306	how are users suppose to do this now ? can that serve as a workaround ?	1	0	https://issues.jboss.org/browse/GPD-306

#	Requirement	Questions	Answered	Successfully Answered	URL
14262	GPD-299	Jiri, is there still stuff to do for this?	0	0	https://issues.jboss.org/browse/GPD-299
14263	GPD-293	indeed post ga is better. i leave this issue in 4.0.x, but i think that if you address some of these that you should create individual issues for them	0	0	https://issues.jboss.org/browse/GPD-293
14264	GPD-282	My current ideas are to have four tabs in the tabbed properties view: - General: to contain name (and later description attributes) - Details: to contain all relevant details (e.g. expression for an exclusive gateway); ideally this should be based on metadata. - Graphics: the graphical attributes if any - Advanced: the tabular properties view for all properties not supported by the above tabs Problems to solve: - How to deal with collections of config data? - How do the metadata look to compose the Details tab?	0	0	https://issues.jboss.org/browse/GPD-282
14265	GPD-239	Please tell us a little more about this release * where are the binaries? * where is the GPD site hosted? * etc...	1	1	https://issues.jboss.org/browse/GPD-239
14266	GPD-226	what exactly does this mean ? you already have a build, right ?	0	0	https://issues.jboss.org/browse/GPD-226
14267	GPD-206	Yep, showing in between « » would be great. Would asking for color coded swimlanes be asking too much? :-)	0	0	https://issues.jboss.org/browse/GPD-206
14268	GPD-203	The patch is applied on both HEAD as 3.1.x. Thanks for the contribution Bernd-Christoph. Can anyone from the ESB guys verify that it works as expected. I will attach a build of the plugin here but I'll first have a look at GPD-202 and GPD-204	0	0	https://issues.jboss.org/browse/GPD-203
14269	GPD-165	hmmm I was sure I also added a comment... here we go - old version of GPD, very old version - GPD has it's own project in jira, called GPD: http://jira.jboss.com/jira/browse/GPD - very little info - different requests in one issue Now more constructive answers 1: totally up to you. use the console, create/use ant tasks, create you own app... pick one...if you want different deployment methods between test and production, then do so.. if not, keep them the same. 2: Runtime at the 'process' level? e.g. when a certain task is encountered? then file a jira issue for the GPD. This is *not* just a documentation issue but a very complex thing to build. You can do this in unit tests very easily btw...	0	0	https://issues.jboss.org/browse/GPD-165
14270	GPD-111	Matthew, Did you sign the contributor's agreement? It would be better to do this to avoid later trouble. Also, i can leave your name as the author of the contributed classes then ;-) Cheers, Koen	1	1	https://issues.jboss.org/browse/GPD-111
14271	GPD-111	Hi Matthew, I have worked on this today and I think the feature should be working now. I took the liberty of changing your approach somewhat. - I sticked to the changes in your extension point definition - I changed the interface name to XmlElementMapper, this interface only has the accept(Node) method - The registry will register all the possible configuration elements for an xml element in a set directly sorted according to the specified priority - The xml adapter factory will iterate through the sorted set and once it finds a mapper that accepts the dom node create the adapter found in the configuration element With this approach, I was able to simplify the api (the interface has only one method), to simplify and optimize the implementation of the adapter creation (sorting does only happen once) and to get rid of the default implementation of the mapper. The only thing I do not know is if this approach works for you, so could you please test it for me? The changes are committed to HEAD Regards, Koen	1	1	https://issues.jboss.org/browse/GPD-111
14272	GPD-111	I don't understand. What are the exact issues you have? Why do you have to give different priorities to the nodes?	1	1	https://issues.jboss.org/browse/GPD-111
14273	GPD-111	I made a mistake implementing the priority comparator. I think it should work now, can you verify?	1	1	https://issues.jboss.org/browse/GPD-111
14274	GPD-111	Is it by intention, that the default priority is MAXINT, if you do not set a value for "mapperClassPriority"? Because of that, it is not possible to activate a mapping on an existing xml types, since the existing semantic element is always first in the SortedSet. So I have to set mapperClassPriority for this existing xml mapping. Or am I missing something?	1	1	https://issues.jboss.org/browse/GPD-111
14275	GPD-111	The value MAX_VALUE was originally chosen because this would be the lowest priority (huh?). The (second) mistake was that I assumed wrongly the wrong order within the sorted set. I think it would indeed be best to do as you suggest and assume MIN_VALUE for the default priority and let contributors specify higher priorities if they want their mapper to take over. If I am not mistaken, the ordering should be from high to low now with the latest commit. Can anyone verify?	1	1	https://issues.jboss.org/browse/GPD-111
14276	GPD-111	Matthew, can you have a check if this works for you and close the issue if it does?	1	1	https://issues.jboss.org/browse/GPD-111
14277	GPD-111	ok this appears to be working now - heres a patch from eclipse from the eclipse "Create Patch" - is that the best way to provide these? if not i can attach the entire file... ### Eclipse Workspace Patch 1.0 #P org.jbpm.gd.jpdl Index: src/org/jbpm/gd/common/xml/XmlAdapterFactory.java === RCS file: /cvsroot/jbpm/jbpm.3/designer/jpdl/org.jbpm.gd.jpdl/src/org/jbpm/gd/common/xml/XmlAdapterFactory.java,v retrieving revision 1.8 diff -u -r1.8 XmlAdapterFactory.java --- src/org/jbpm/gd/common/xml/XmlAdapterFactory.java 7 Aug 2007 13:02:45 -0000 1.8 +++ src/org/jbpm/gd/common/xml/XmlAdapterFactory.java 8 Aug 2007 20:16:40 -0000 @@ -99,6 +99,17 @@ Logger.logError("Neither xmlElement nor nameProvider are present on mapping configuration element for " + jpdlElement.getElementId(), e); } } + } else { + try { + XmlAdapter adapter = (XmlAdapter) configurationElement.createExecutableExtension("adapterClass"); + INodeNotifier element = (INodeNotifier)document.createElement(elementName); + element.addAdapter(adapter); + adapter.setNode((Node)element); + adapter.setFactory(this); + return adapter; + } catch (CoreException e) { + Logger.logError("Unable to create XML Adapter for " + jpdlElement.getElementId(), e); + } } return elementName == null ? null : createAdapterFromElementName(elementName); }	1	1	https://issues.jboss.org/browse/GPD-111
14278	GPD-111	Matthias, I have simply *moved* the setFactory() before adapt(). Does that work as well?	1	1	https://issues.jboss.org/browse/GPD-111
14279	GPD-77	It may work in this way, but it doesn't feel right as it blurs the lifecycle of the semantic element. Let me think about this and try to come up with a working solution. Btw. I infer from your previous comment that you guys are doing xpdl on top of jbpm? Can you elaborate on this? Are you planning to contribute something?	1	1	https://issues.jboss.org/browse/GPD-77
14280	GPD-77	I've been experimenting with (on and off) adding GD support for "xpdl", which would become a domain specific "extension" of jpdl. Because of its proprietary nature it wouldn't be very useful to the general public. I plan to keep contributing patches and post feature requests that would make it possible/easier to use GD for editing "xpdl", and this would hopfully also make it easier for other developers to add support for their own dsls. Does this sound like a good approach? Thanks, John	1	1	https://issues.jboss.org/browse/GPD-77
14281	GPD-77	xpdl is a public standard so it is definitely not proprietary. xpdl support is on the roadmap of our new pvm approach. we are collaborating with the guys from the bonita engine and with the eclipse jwt project to get this going. or is this a name clash?	1	1	https://issues.jboss.org/browse/GPD-77
14282	GPD-77	ROTFL Maybe you need to go into the 7-character acromym space to still find something useful in the workflow/BPM space ;-) As for the problem at hand. Isn't it possible to do use the XmlAdapter.initialze(SemanticElement) method to override the (default) values that were set in the SemanticElement.initialize() method by the values that are specified in the xml? This method was created for this purpose in any case.	1	1	https://issues.jboss.org/browse/GPD-77
14283	GPD-77	John, What is wrong with doing this in the constructor of the SemanticElement? Regards, Koen	1	1	https://issues.jboss.org/browse/GPD-77
14284	GPD-65	This issue has been here for almost 2 years now, is it ever going to be resolved?	1	1	https://issues.jboss.org/browse/GPD-65

#	Requirement	Questions	Answered	Successfully Answered	URL
14285	GPD-54	I am not clear about what is really needed here. Is it some kind of an annotation or text box that appears as a descriptive element on the graphical canvas?	1	1	https://issues.jboss.org/browse/GPD-54
14286	GPD-33	Koen, I tried to generate a form for a task with a few variables mapped in the task controller. In the description, you say "we currently start from the data entered in the task controller". Is that behavior available in 3.1.3, in the works, or planned for the future?	1	0	https://issues.jboss.org/browse/GPD-33
14287	GPD-33	Was there a problem with this task form generation? It should normally work fine...	0	0	https://issues.jboss.org/browse/GPD-33
14288	GPD-14	I'm looking at creating an implementation of this feature that pulls the graph structure into another package that does layouts and then generating move commands. Is that a viable approach or is there a specific set of algorithms you'd need to use? I would then simply include a button or menu item to apply the layout(s). I may have some time over the next month or so to put it together.	0	0	https://issues.jboss.org/browse/GPD-14
14289	GPD-3	The most important thing is to decide if we want to have - specific rendering per type (date, string, (longer) text, collections (dropdown, list, radio)) (for each of these there is a jsf component) - the type specifiied in the processdefinition.xml or in the gpd.xml (the processdefinition.xml has my preference since read/write/required is also there) - we want to specify other limitations (regexp on input etc...) for me, specifying this per variable in the pd in a way like xforms uses is nice. xforms has the following attributes: required This property, when set to true() or 1, signals that the attached instance data must not be blank, and will prevent submission until that condition is met. relevant This property controls whether the attached instance data is relevant or not. Non-relevant items are usually not rendered (although this can be fine-tuned with stylesheets) and even not submitted. readonly This property controls whether the attached instance data is read-only, in which case it can't be user-modified. calculate This property bears a full XPath expression, which can include mathematical operators and function calls, which is used to automatically provide a value to the attached instance data nodes. Multiple calculations can interact, and like a spreadsheet, they will recalculate in the proper order for dependencies to work out. constraint This property bears a full XPath expression which must evaluate to true() in order for the attached instance data to be considered valid. Therefore, all constraints must be met before submission can succeed. type This property associates a datatype (taken from XML Schema) with the attached instance data nodes. Datatypes can be useful in refining the view into data, for example using a calendar control to enter a date. p3ptype This property associates a datatype (taken from P3P) with the attached instance data nodes. P3P datatypes are useful for privacy monitoring, namely smart user agents that keep track of what information they give out. It can also serve as a flag for the client to pre-fill certain allowed kinds of information. Which is imo all we need. In xpath these can all take an xpath to another item (e.g. this input is required if another input has a certain value), In jBPM we could use EL for this instead of xpath since we do not have a complex document structure but an 'array' of variables. This way it is even possible to do things like 'required if user is in role' Xforms besides that has some predefined controls: input Entry of free-form value edit box, voice prompt textarea Entry of large amounts of free-form text email body, weblog entry secret Entry of sensitive information password prompt select1 Choice of one-and-only-one item from a list radio buttons, drop-list select Choice of one or more items from a list checkbox group, listbox range Selecting a value from a range slider, volume control upload Selecting a data source file picker, digital camera interface trigger Activating a defined process button, hyperlink submit Activating submission of the form submit button output Display-only of form data inline text I'm not saying we should do it identically to xforms but at least look at what they did and why (I myself would e.g. not have a string based 'read, write, required' but have separate attributes (read="true" write="false" required="${var1 >10}) e.g... would be much nicer. Building a freemarker template what gets this feeded per task and generates .xhtml files with jsf controls in it would not be to difficult for people who want to have it customized. Generating a .xhtml form runtime (jbpm supports this indirectly) based on such a template for having automagically generated forms would be possible then as well. Most work will go into the GPD for having the ui for filling all the attributes described above... Shoot, this has become such a long post that I think we should put it in a wiki page as a design document. What do you think?	1	1	https://issues.jboss.org/browse/GPD-3
14290	GPD-3	Most important imo is to still have the automagically generated pages in 3.2 with the addition of automagically correct input elements. So deciding to store the type somewhere is imo the most important thing the gpd should add. The way forms can be custom build by the user (something different than customizing generated forms) can be described in a 'how do I design custom pages'	1	1	https://issues.jboss.org/browse/GPD-3
14291	GPD-1	Translated palette entries. Not much left to finish it. Feel free to add your own localization, simply find all the files messages.properties and copy them with a new name messages_??.properties, where ?? is two letter signature of your language. All the strings values should be in UTF-8 encoding, for example russian word "???" will look like: "\u0412\u0438\u0434". If you need a utility to do such a translation tell me about it and I'll post it here.	0	0	https://issues.jboss.org/browse/GPD-1
14292	JBRULES-3714	Toshiya Kobayashi <tkobayas@redhat.com> made a comment on [bug 894222|https://bugzilla.redhat.com/show_bug.cgi?id=894222] Hi, This BZ has been approved for the next roll up patch for BRMS 5.3.1 (BZ873846). Could you cherry-pick the commit https://github.com/droolsjbpm/drools/commit/a9a93876f2ead9468fd50eba715083c9a7e8a52a to 5.3.x branch? Thanks! Toshiya	0	0	https://issues.jboss.org/browse/JBRULES-3714
14293	JBRULES-3699	Edson Tirelli <etirelli@redhat.com> made a comment on [bug 958387|https://bugzilla.redhat.com/show_bug.cgi?id=958387] Mario, can you please take a look at this one? Is this still happening in 6?	0	0	https://issues.jboss.org/browse/JBRULES-3699
14294	JBRULES-3693	Hi guys, Although I think this has been resolved, can you please verify and close the ticket if it is resolved. Thanks George.	1	0	https://issues.jboss.org/browse/JBRULES-3693
14295	JBRULES-3679	What are the ramifications on Guvnor (that's a WAR) that will need to build KProject's?	1	1	https://issues.jboss.org/browse/JBRULES-3679
14296	JBRULES-3635	This question is important for us, can be processed as quickly as possible?	0	0	https://issues.jboss.org/browse/JBRULES-3635
14297	JBRULES-3619	This can be done quite cleanly as an integration exercise via a memory file system and a custom classloader. For this reason I have closed the pull request, and have not merged it. I've pasted the comments I left in the pull request below: --- not sure I like this, it seems very kludgy. The knowledge agent is for monitoring external resources, but this is applying them. It's possible to make a MemoryFileSystem and actually use that, without changing KnowledgeBuilder at all. I recently wrote a MemoryFileSystem here: https://github.com/droolsjbpm/drools/tree/master/drools-compiler/src/test/java/org/drools/kproject/memory And here you can see the vfs in action, programmatically, to allow me to build resources in memory without touching disk. See testCompileAndCDI, where I build up the contents of the MemoryFileSystem. https://github.com/droolsjbpm/drools/blob/master/drools-compiler/src/test/java/org/drools/kproject/KProjectTest.java What makes it work is the classloader, scroll down to line 455: MemoryFileSystemClassLoader As you can see the builder factory is using just standard knowledge builder public apis, but they are actually being resolved from in memory. https://github.com/droolsjbpm/drools/blob/master/drools-compiler/src/main/java/org/drools/kproject/KBaseBuilder.java You can actually mix both in memory and disk based file systems, using the in memory to "overlay" the disk.	0	0	https://issues.jboss.org/browse/JBRULES-3619

#	Requirement	Questions	Answered	Successfully Answered	URL
14298	JBRULES-3520	Hi, where can I find the CamelServerApp.class file? tx a lot Maurizio	1	0	https://issues.jboss.org/browse/JBRULES-3520
14299	JBRULES-3470	Can you let me know when the new update sites are available, eg.: http://download.jboss.org/drools/release/5.4.0.CR1/org.drools.updatesite http://download.jboss.org/drools/release/5.3.2.Final/org.drools.updatesite	1	1	https://issues.jboss.org/browse/JBRULES-3470
14300	JBRULES-3470	I see that 5.4.0.Final has a site. http://download.jboss.org/drools/release/5.4.0.Final/org.drools.updatesite/ Can it be linked from https://www.jboss.org/drools/downloads, so the community can find it?	1	1	https://issues.jboss.org/browse/JBRULES-3470
14301	JBRULES-3470	Done, thanks for reporting. http://www.jboss.org/drools/downloads Below is the text, can you verify that the text is accurate (especially the part about JBoss Tools)? Eclipse update site Click here to go to the Drools and jBPM update site 5.4.0.Final. The Drools plugin for Eclipse is also part of JBoss Tools. Click here to go to the JBoss Tools Eclipse update site. Alternatively, you can download the Drools and jBPM tools zip, unzip it and use the directory there as a local updatesite.	1	1	https://issues.jboss.org/browse/JBRULES-3470
14302	JBRULES-3470	Drools 5 and jBPM 5 (as well as jBPM 3) are not strictly part of JBoss Tools anymore, so there's little point linking to http://www.jboss.org/tools/download.html. Once the SOA Tools ship as part of the 3.3.0.Final release, they will provide a link to themselves from within JBoss Central (that is, if you've installed JBoss Tools you'll be able to discover Drools/jBPM3/jBPM5 right in Eclipse). As to the "you can download the Drools and jBPM tools zip", it's not clear which zip you mean. Perhaps place a link in the text to the zip you mean, or reference it using the *exact same text* as in the table above?	1	1	https://issues.jboss.org/browse/JBRULES-3470
14303	JBRULES-3470	Would it be possible to symlink your content in http://download.jboss.org/drools/release/5.4.0.Final/org.drools.updatesite/ from http://download.jboss.org/jbosstools/updates/stable/indigo/soa-tooling/droolsjbpm/5.4.0.Final/ ? You can achieve this with some shell script at the end of your publishing job in Jenkins: {code} ECLIPSERELEASE=indigo PROJECTNAME=droolsjbpm SOURCESERVER=tools@filemgmt.jboss.org:/downloads_htdocs SOURCEPATH=/drools/release DESTINATION=tools@filemgmt.jboss.org:/downloads_htdocs/tools/updates/stable/${ECLIPSERELEASE}/soa-tooling PROJECTVERSION=$(echo "ls *.Final" | sftp ${SOURCESERVER}${SOURCEPATH} 2>&1 | grep "\.\+Final" | grep -v sftp | sort | tail -1); PROJECTVERSION=${PROJECTVERSION%%/*}; echo ${PROJECTVERSION} # if folder already exists, sftp will throw an error but should still be able to proceed echo "mkdir ${PROJECTNAME}" | sftp ${DESTINATION} # symlinking appears to only works with relative paths echo "ln ../../../../../..${SOURCEPATH}/${PROJECTVERSION}/org.drools.updatesite ${PROJECTVERSION}" | sftp ${DESTINATION}/${PROJECTNAME}/ echo "Click here: http://download.jboss.org/jbosstools/updates/stable/${ECLIPSERELEASE}/soa-tooling/${PROJECTNAME}/${PROJECTVERSION}" {code} Alternatively, instead of symlink, could also fetch from one site & post to the other (uses twice as much disk space, however) {code} rsync -aPz --rsh=ssh --protocol=28 ${SOURCESERVER}${SOURCEPATH}/${PROJECTVERSION}/org.drools.updatesite/* /tmp/org.drools.updatesite-${PROJECTVERSION} rsync -aPz --rsh=ssh --protocol=28 --delete /tmp/org.drools.updatesite-${PROJECTVERSION}/* ${DESTINATION}/${PROJECTNAME}/${PROJECTVERSION}/ # cleanup rm -fr /tmp/org.drools.updatesite-${PROJECTVERSION} echo "Click here: http://download.jboss.org/jbosstools/updates/stable/${ECLIPSERELEASE}/soa-tooling/${PROJECTNAME}/${PROJECTVERSION}" {code}	1	1	https://issues.jboss.org/browse/JBRULES-3470
14304	JBRULES-3470	Caching problems? Can you link to a servicedesk ticket w/ details?	1	1	https://issues.jboss.org/browse/JBRULES-3470
14305	JBRULES-3465	Hi, Can i please have an update on whether username or password is supported at all in KnowledgeAgentConfiguration? In my project it is discouraged to check-in any username or password on subversion, so we are passing it from the command line. Thing is, I am getting HTTP 401 when I move out the code to set username/password from change-set.xml onto the KnowledgeAgentConfiguration this way: KnowledgeAgentConfiguration conf = KnowledgeAgentFactory.newKnowledgeAgentConfiguration(); conf.setProperty ("enableBasicAuthentication", "true"); conf.setProperty ("username", username); conf.setProperty ("password", password); KnowledgeAgent kagent = KnowledgeAgentFactory.newKnowledgeAgent("kagent", conf); kagent.applyChangeSet(ResourceFactory.newClassPathResource("rules/change-set.xml")); StatelessKnowledgeSession ksession = kagent.newStatelessKnowledgeSession(); Please advise how I can set the username/password to connect to Guvnor programmatically. Thank you.	0	0	https://issues.jboss.org/browse/JBRULES-3465
14306	JBRULES-3402	Davide, does your fix also cover the use-case in https://issues.jboss.org/browse/GUVNOR-1980 whereby expressions can be used in CEP operators?	0	0	https://issues.jboss.org/browse/JBRULES-3402
14307	JBRULES-3389	Mark, Which kind of unit do you want ? To test OSGI bundles, we should deploy it in an OSGI runtime. Until now, to make tests like that we are using for Apache Karaf, Camel, ... projects Pax-Exam (http://team.ops4j.org/wiki/display/paxexam/Pax+Exam). What do you think ? Regards, Charles	1	0	https://issues.jboss.org/browse/JBRULES-3389
14308	JBRULES-3318	@Jim Do you mean Guvnor data or drools-persistence-jpa data? Could you give an example of such a table name, so we can verify to witch part you're talking about. Guvnor doesn't use any table names directly, because it uses JCR. If you use the default JackRabbit implementation, it writes to the file system, not a database. If you reconfigure that or use Modeshape instead of JackRabbit, you can write to a database (or even infinispan). Modeshape uses hibernate-jpa, so it might support changing the table names. If it's about drools-persistence-jpa, we have to move this to JBRULES. Since that module uses JPA, in theory you can change the table name structure by creating and editing META-INF/orm.xml in that drools-persistence-jp jar (not tested yet). If there is anything that breaks that (for example native SQL queries which aren't there as far as I know), please create an issue in JBRULES. @Rikkola probably all tables, like they can in JPA by configuring an orm.xml in it.	1	1	https://issues.jboss.org/browse/JBRULES-3318
14309	JBRULES-3318	From JBRULES-3305 it's clear that they mean drools-persistence-jpa, so I 've moved this issue to JBRULES. @Jim can you verify if the META-INF/orm.xml way is viable or not?	1	1	https://issues.jboss.org/browse/JBRULES-3318

#	Requirement	Questions	Answered	Successfully Answered	URL
14310	JBRULES-3305	Moved from Guvnor's jira: Guvnor itself doesn't have any hibernate specific annotations. It does depend on hibernate-validator (but that's unrelated to hibernate persistence and/or JPA) and even then it only uses the bean validation JSR annotations. Drools does depend on hibernate-core, annotations and entitymanager, but it doesn't use any hibernate specific annotations. JBPM does use hibernate specific annotations: {code} Targets String 'import org.hibernate' Found usages (5 usages) Production (5 usages) jbpm-bam (4 usages) org.jbpm.process.audit (4 usages) HibernateUtil.java (2 usages) (19: 1) import org.hibernate.SessionFactory; (20: 1) import org.hibernate.cfg.Configuration; ProcessInstanceDbLog.java (1 usage) (21: 1) import org.hibernate.Session; WorkingMemoryDbLogger.java (1 usage) (28: 1) import org.hibernate.Session; jbpm-persistence-jpa (1 usage) org.jbpm.persistence.processinstance (1 usage) ProcessInstanceInfo.java (1 usage) (35: 1) import org.hibernate.annotations.CollectionOfElements; {code} What does the customer want? - Is removing that one @CollectionOfElements in jBPM enough? - Or do all the org.hibernate imports in jBPM have to be removed? - Or do all the hibernate-* dependencies be set to scope runtime (which might not be feasible)?	1	1	https://issues.jboss.org/browse/JBRULES-3305
14311	JBRULES-3282	Geoffrey De Smet <gdesmet@redhat.com> made a comment on [bug 732085|https://bugzilla.redhat.com/show_bug.cgi?id=732085] Is this is a duplicate of these? https://bugzilla.redhat.com/show_bug.cgi?id=741219 https://bugzilla.redhat.com/show_bug.cgi?id=741287 https://bugzilla.redhat.com/show_bug.cgi?id=751319 https://bugzilla.redhat.com/show_bug.cgi?id=751326	0	0	https://issues.jboss.org/browse/JBRULES-3282
14312	JBRULES-3238	Oct 06 12:37:55 <mriet> ge0ffrey, which repo is droolsjbpm-introduction-docs in? Oct 06 12:38:10 <ge0ffrey> mriet: droolsjbpm-knowledge	0	0	https://issues.jboss.org/browse/JBRULES-3238
14313	JBRULES-3192	I am attempting to run Drools and Camel in Karaf (all using the Fuse X.02-20 release). There is an incompatibility from JAXB because Drools OSGi declares dependency on com.sun.tools.xjc. Is this dependency really necessary? Can it be removed? I found a comment related to it in JBRULES-2351.	0	0	https://issues.jboss.org/browse/JBRULES-3192
14314	JBRULES-3157	After some tests it seems to work well and integrated into 5.4.0.Final version. Do you want that I modify the state of this ticket? Jérémie.	1	1	https://issues.jboss.org/browse/JBRULES-3157
14315	JBRULES-3154	Comment from OP (Vincent Legendre): Should add a warning in the doc then. May be that the fix is shorter to do... The Date problem is common with XLS (POI has exactly the same problem), but there, the displayed value is not the one rendered in the DRL, which is a problem (in my opinion).	0	0	https://issues.jboss.org/browse/JBRULES-3154
14316	JBRULES-3152	How do you manage the KA lifecycle? The user is expected to control the monitoring threads: ResourceFactory.getResourceChangeNotifierService().start()/stop(); ResourceFactory.getResourceChangeScannerService().start()/stop(); and to dispose the KA when no longer needed: kAgent.dispose(); ---------- At the moment, the KA name is not unique and they are not tracked. Could you provide a few more details on your use case so I can see if it is necessary/reasonable to implement an agent registry?	1	0	https://issues.jboss.org/browse/JBRULES-3152
14317	JBRULES-3063	kylin <ksoong@redhat.com> made a comment on [bug 724753|https://bugzilla.redhat.com/show_bug.cgi?id=724753] As [1] hints the issue is fixed in drools 5.3.0.Final, but I test via BRMS 5.3.0 the issue still existed [1] https://issues.jboss.org/browse/JBRULES-3063	0	0	https://issues.jboss.org/browse/JBRULES-3063
14318	JBRULES-3063	Edson Tirelli <etirelli@redhat.com> made a comment on [bug 724753|https://bugzilla.redhat.com/show_bug.cgi?id=724753] Can you please clarify how you tested? By default, BRMS does allow rule replacement... it will only raise errors when the engine is configured to raise errors for rule duplication.	0	0	https://issues.jboss.org/browse/JBRULES-3063
14319	JBRULES-3056	Sure, attach a patch or submit a pull request and we'll add it. Can this be done without any additional dependency jars?	0	0	https://issues.jboss.org/browse/JBRULES-3056
14320	JBRULES-2884	We need more info. Is this in the build, in eclipse, in the API? Can you show some code how AS, WS and ESB do this? Note that for the build, we build and deploy the source jars (which includes javadocs) to the public jboss repository.	1	0	https://issues.jboss.org/browse/JBRULES-2884
14321	JBRULES-2833	This issue is not resolved. 1. The file pattern to include/exclude .package files is incorrect: should be "*.package" instead of {code}"***..package"{code} 2. DirectoryScanner.scan() is not called after the includes/excludes are overridden. 3. Most importantly: Overriding the includes which are already defined in the ant task seems like bad design. Seems more logical to me to add a Comparator on the file list which prioritises files with extension .package 4. unittest? ;-)	1	0	https://issues.jboss.org/browse/JBRULES-2833
14322	JBRULES-2829	Why do you call the solution a hack? the parser generated code is huge now as it includes basically the full java grammar in there + all the DRL constructions. The DFAs are pretty big as well. The only way to deal with that that I found so far is increasing the stack size...	0	0	https://issues.jboss.org/browse/JBRULES-2829
14323	JBRULES-2816	Do we want to include all optional jars? Hmm maybe not the felix stuff etc if that is optional? Do we want to mix them in the release zip with non-optional jars? Yea, doesn't matter since the module classpaths are mixed anyway so it's a mess anyway (=> it's not decent dep management) The README.txt and README_dependencies.txt file are out of date.	0	0	https://issues.jboss.org/browse/JBRULES-2816
14324	JBRULES-2792	Wolfgang, I have added some metadata accessors recently, could you please check if the current APIs are sufficient or something is still missing? Thanks	0	0	https://issues.jboss.org/browse/JBRULES-2792
14325	JBRULES-2759	Could you create a junit test for this problem?	1	0	https://issues.jboss.org/browse/JBRULES-2759
14326	JBRULES-2759	Yes, I see, that this test is ok, but if I start process from remote application on gwt-console-server, I get always -1 value in processInstanceId field of Task table. I've posted Task.patch, there isn't in latest snapshot and I think so, but my solution works fine with remote application and I've created every patches, which needed. Have you got any idea about this problem and solution?	0	0	https://issues.jboss.org/browse/JBRULES-2759
14327	JBRULES-2694	would you keep the entire SQL and plain java comparison in it? need to take a closer look... probably only for presentation... not sure it is worth maintaining the code for it ge0ffrey: I am in favor of keeping it, but not necessary for every rule	1	1	https://issues.jboss.org/browse/JBRULES-2694
14328	JBRULES-2685	Geofrey: Can you check if this perf regression has been fixed.	1	1	https://issues.jboss.org/browse/JBRULES-2685

#	Requirement	Questions	Answered	Successfully Answered	URL
14329	JBRULES-2640	Is this a {{group-by}} intended to be in conjunction with {{accumulate}} ? IFF so, normally I try to normalize it as per the below example; possibly doesn't cover all possible use-cases, but I thought maybe could be helpful in the time being. {code} declare GroupByString groupId : String groups : String[] end rule "init Metro hash of my interest" no-loop when then GroupByString grp = new GroupByString(); grp.setGroupId("metroHash"); grp.setGroups(new String[]{"#M1", "#M2"}); insert(grp); end rule "Multiple tweets about a metro line hash" no-loop when $group : GroupByString(groupId == "metroHash") $metroHash : String() from $group.groups accumulate ($tweet : Status(text contains $metroHash, inReplyToUserId == -1) over window:time(5m) ; $list : collectList($tweet), $count : count($tweet); $count >= 2) not Alert(condition == "Multiple tweets about a metro line hash "+$metroHash) over window:time(5m) then Status[] theTweets = new Status[$list.size()]; $list.toArray(theTweets); // TODO check: I've tried more concise one-liner syntax but fail KB compilation. Alert<Status[]> alert = new Alert<Status[]>(sessionCurrentTime(kcontext), "Multiple tweets about a metro line hash "+$metroHash, AlertType.INFO, theTweets); insert(alert); LOG.info("Multiple tweets for " + $metroHash + " count " + $count + " list: " + $list); end {code} taken from here https://github.com/tarilabs/reex2014-rules/blob/master/src/main/resources/tweet.drl	0	0	https://issues.jboss.org/browse/JBRULES-2640
14330	JBRULES-2604	Can you check against our most recent spring stuff, do we have something that does this for you now? If so please close this jira, or come back with additional requests against what we have now.	0	0	https://issues.jboss.org/browse/JBRULES-2604
14331	JBRULES-2471	Was this left out intentionally when first programmed and was expected to work as is or was this a bug that was found out later? Does this seem like a lot of work to persist the event listeners? Is it reasonable to expect this feature in the 5.1 release?	1	1	https://issues.jboss.org/browse/JBRULES-2471
14332	JBRULES-2459	The WS-HT spec doesn't really define reassign, so I think this also should be called delegation (just different behavior depending on the state of the task you are trying to delegate). Could you elaborate why you think delegation is not working (as there are some junit tests showing that it should be)? Could you give an example of your starting task you are trying to delegate (what state is it in, what are the potential owners or assignee etc.), how you delegate and what you would expect as the result (but apparently not getting)?	0	0	https://issues.jboss.org/browse/JBRULES-2459
14333	JBRULES-2455	Given a decision table column: VALUE1, VALUE2, VALUE3 VALUE4, VALUE5 Currently, only a token index specific tag is available, giving snippet: property == "$1" && property == "$2" && property == "$3" which results in the following post compilation DRL: property == "VALUE1" && property == "VALUE2" && property == "VALUE3" property == "VALUE4" && property == "VALUE5"&& property == "" This gives unintended results of matching an empty string and puts an upper bound on the number of tokens.	1	1	https://issues.jboss.org/browse/JBRULES-2455
14334	JBRULES-2448	The fix is an improvement, but wouldn't it be better to avoid the hibernate-specific @CollectionOfElements annotation? Also the way things are I don't think I can override the table name 'EventTypes' because it is assumed here, in this fix.	0	0	https://issues.jboss.org/browse/JBRULES-2448
14335	JBRULES-2409	I like the second idea where you expect the developer to deliberately escape the quotes in the cells. Although I think it would be more natural for a developer to write: $myObject.setField("A 12\" rule") in a cell. A business user will never write this, he will write 'A 12" rule'. Ideally, the escaping "rule" should be intelligent enough to see the cell contains actual code, so no escaping is necessary (since the cell is written and maintained by a developer). But I guess there will always be scenarios slipping through the holes of the safety net. Could we provide some kind of indication in a cell (or column) to tell Drools the cell contains actual code, so its safe to expect the developer properly escaped everything?	1	1	https://issues.jboss.org/browse/JBRULES-2409
14336	JBRULES-2409	I really like the parser to be intelligent. But I agree that it is very hard (much work) to make it work properly. Even in case of Michael's idea what should we do in the rare case a business user writes \"? He would then get " in the drl without expecting it. I know chances are small but still. As ar as I'm aware there is curently no mechanism in the decision table to add meta data to a specific cell (Or is there?). The way I see it is that in most cases you want to enable/disable quote escaping per sheet. Developers and business users typically do not edit the same rule sheet. The busines user should not be bothered with stuff that he does not understand or may not edit. Given that the sequential parameter is also a sheet based parameter, I therefor think it is best to use that mechanism. Kenneth would you be happy with a global (e.g. sheet) parameter?	1	1	https://issues.jboss.org/browse/JBRULES-2409
14337	JBRULES-2409	Here is the pull request: https://github.com/Jospfh/drools/pull/1 Do I need to do anything else?	1	1	https://issues.jboss.org/browse/JBRULES-2409
14338	JBRULES-2409	Hi Jos, thanks for the fix :) We are however having problems merging your pull request into droolsjbpm :(The pull request is not appearing in droolsjbpm. Can you please re-try making the pull request ensuring you have droolsjbpm\drools@master as the target repo?	0	0	https://issues.jboss.org/browse/JBRULES-2409
14339	JBRULES-2351	The mylin Manifest you provide has no exports, is that correct?	0	0	https://issues.jboss.org/browse/JBRULES-2351
14340	JBRULES-2342	Do we need to close this one for 5.1?	0	0	https://issues.jboss.org/browse/JBRULES-2342
14341	JBRULES-2301	Any news on this? Should the sysouts be removed alltogether as was done for jbpm?	0	0	https://issues.jboss.org/browse/JBRULES-2301
14342	JBRULES-2300	Would it be possible to test this with the latest snapshot release? Or could you find a way to reproduce this issue somehow so we can take a look what's going on at runtime?	0	0	https://issues.jboss.org/browse/JBRULES-2300
14343	JBRULES-2264	It's confusing: what is the current official update site? I know of these ways: - 1) According to http://www.jboss.org/drools/downloads the update site is http://downloads.jboss.com/drools/updatesite3.5/ (note that none of the milestones have been uploaded there) (note that its not split up per ganymede, galileo, helios like JBoss Tools) - 2) But also on http://www.jboss.org/drools/downloads you can download the eclipse zip to have a local update site directory (which is available for the milestones too) - 3) JBoss tools includes drools, so there are a couple of update site on http://www.jboss.org/tools/download/ such as http://download.jboss.org/jbosstools/updates/stable/ganymede/ - 4) In development, to build from source, it's a matter of following https://github.com/droolsjbpm/droolsjbpm-tools/blob/master/drools-eclipse/org.drools.updatesite/pom.xml#L17 Do we want to keep and maintain 1)? If so, updating it should be part of the release checklist. If not (which I am in favor of), we just link the download page to JBoss Tools download page. Either way, we should keep 2) (but clearly explain the difference).	1	0	https://issues.jboss.org/browse/JBRULES-2264
14344	JBRULES-2264	We're no longer maintaining our own update site. You can either use - (2) to download the eclipse zip and use it as the local update site - (3) to download the latest version of the Drools plugin from the JBossTools update site I tested http://download.jboss.org/jbosstools/updates/stable/ on my Helios installation and it seems to work fine. Is it a requirement to use Ganymede (as this is already quite old)? Could you upgrade to a more recent version? I haven't tested on Ganymede.	0	0	https://issues.jboss.org/browse/JBRULES-2264
14345	JBRULES-2243	Hello, I can see this issue was moved from milestone 1 to milestone 2. Do you think you will work on it for milestone 2 ? Do you have a planned release date for milestone 2 ? Thanks Benoit	1	1	https://issues.jboss.org/browse/JBRULES-2243
14346	JBRULES-2147	Is this still bad in 5.2?	0	0	https://issues.jboss.org/browse/JBRULES-2147

#	Requirement	Questions	Answered	Successfully Answered	URL
14347	JBRULES-2124	Jurgen, this feature is available since December 6 2008, so I think it is already present in 5.0.1.FINAL. What you have to do is: kagent.applyChangeSet(ResourceFactory.newUrlResource(yourUrl);	0	0	https://issues.jboss.org/browse/JBRULES-2124
14348	JBRULES-2040	Is this still a problem in Drools 5.2, can you re-test?	0	0	https://issues.jboss.org/browse/JBRULES-2040
14349	JBRULES-2020	The screenshot looks good? What's the problem? I am working on the styles and about to release them. Either we fix it now or much later :)	0	0	https://issues.jboss.org/browse/JBRULES-2020
14350	JBRULES-1908	Hi Darrin, These changes are very welcome, but because that file has been changed 11 times since than and the Insertion content is changed on about every line: https://github.com/droolsjbpm/drools/tree/master/drools-docs/drools-expert-docs/src/main/docbook/en-US/Obsolete-Chapter-Rule_Engine/Section-The_Drools_Rule_Engine.xml So these changes are stale :(Would it be possible to create a pull request with the not-stale changes (or even apply them directly as you're the documentation lead)? Poke me if you decide the pull request way: we don't want this to go stale again.	1	1	https://issues.jboss.org/browse/JBRULES-1908
14351	JBRULES-1900	Hi Mauricio, I find what you posted really useful for what i am doing at the moment, but i am trying to get the files you uploaded and keeps saying they dont exist anymore, could you please upload them again so i can have a look at them please? Thanks Chrissie	1	1	https://issues.jboss.org/browse/JBRULES-1900
14352	JBRULES-1826	Did you tried to escape the {}? The new DSL parser is much more robust and strict. It interprets anything inside non escaped {} as a variable name. Try the following: [consequence][]log an error "{errorCode}"= String[] errorMessage = \{$name\}; errorlist.add(new UWError("{errorCode}",$driverKey, errorMessage)); []s Edson	1	1	https://issues.jboss.org/browse/JBRULES-1826
14353	JBRULES-1721	Any chance to make the latest 5.0.0.SNAPSHOT build, with the patch applied, available through the jboss snapshots maven2 repository? Should i create a new issue for this?	1	1	https://issues.jboss.org/browse/JBRULES-1721
14354	JBRULES-1701	I'm not too keen on increasing the parser/build complexity for what is sugar. eval(unique(A, B,C, D)) that seems adequatate to me instead of : unique(A, B,C, D) Which we would expand to: eval(unique(A, B,C, D)) Btw you know there is the remove identity equals option that will disallow any match where a fact is matched more than once? See RuleBaseConfiguration Mark	1	0	https://issues.jboss.org/browse/JBRULES-1701
14355	JBRULES-1615	I see in the documentation for 5.1 the expander keyword is still being used. What does the expander keyword deprecation mean? I would also like to be able to break up my DSL into multiple files but because of this editor limitation I am unable to. I'm using version 5.0.1.FINAL.	0	0	https://issues.jboss.org/browse/JBRULES-1615
14356	JBRULES-1590	is the only difference that the first catch sets the clone object to null again? Don't think I can see any other differences?	0	0	https://issues.jboss.org/browse/JBRULES-1590
14357	JBRULES-1504	This would be a nice one to do (would make support a lot easier !).	0	0	https://issues.jboss.org/browse/JBRULES-1504
14358	JBRULES-1471	Currently: - accumulate - getResults - accumulate - getResults - other restrictions Mark thinks about putting the result from a from node on the agenda, instead of processing it directly. That might result in worse performance, for those that use an accumulate as an early restriction of which the other restrictions are expensive to evaluate: - accumulate -> schedule getResults - accumulate -> schedule getResults - other restrictions - getResults I've been wondering if it's possible to do a best of both worlds? Schedule it on the agenda, but before the other restrictions are evaluated: - accumulate -> schedule getResults - accumulate -> schedule getResults - getResults - other restrictions Does that look possible to implement?	0	0	https://issues.jboss.org/browse/JBRULES-1471
14359	JBRULES-1431	you can add packages and add/remove rules and packages already in Drools. What is it that we are missing?	0	0	https://issues.jboss.org/browse/JBRULES-1431
14360	JBRULES-1431	What is missing is doing it from a rule, as in a Meta rule. How can a rule add and remove rules? currently there is no key word or built in function that a meta rule could use to create rule. The best way to see how that is done is to try it in CLIPS and JESS.	1	0	https://issues.jboss.org/browse/JBRULES-1431
14361	JBRULES-1431	You can use the MVEL dialect to avoid class generation. I'm not sure what your problem is with dynamic rules. It's possible in drools to do something like this. What more do you require? function void addRule(String pkgName, String rule, RuleBase ruleBase) { PackageBuilder builder = new PackageBuilder(); builder.addPackgeFromDrl(new StringReader(rule)); ruleBase.addPackage (builder.addPackage()); } rule X when then String rule = "package org.drools.test\n global java.util.List list\n rule xxx\n when\n i:Integer()\nthen\n list.add(i);\nend"; addRule("org.domain", rule, drools.getWorkingMemory().getRuleBase()); end	1	1	https://issues.jboss.org/browse/JBRULES-1431
14362	JBRULES-1430	How would one do that from a rule, if "workingmemory" is not a reserved word in DRL language? http://downloads.jboss.com/drools/docs/4.0.4.17825.GA/html_single/index.html#d0e2641 Would it be better for the engine to automatically pass that and have it part of the Function interface, so that "workingmemory" doesn't have to be a reserved word and change the parser? It's common in machine learning applications to have a Meta rule that creates new rules. To do this in clips and JESS is trivial. I suppose one could statically pass the workingmemory object to the functions, but that feels a bit hack-ish.	1	1	https://issues.jboss.org/browse/JBRULES-1430
14363	JBRULES-1430	If you can provide a patch we'll consider it. However this requires text mangling, which is hard to do and as users can just pass the workingmemory themselves it's low on our priority. By this I mean in the consequence how do you detect a function call and how do you mangle the text to insert the workingmemory parameter at the end. This isn't something that can be done simply with some regular expression.	0	0	https://issues.jboss.org/browse/JBRULES-1430
14364	JBRULES-1307	Isaac. Could you please improve the DT documentation just a little for 4.0.x and trunk.	1	1	https://issues.jboss.org/browse/JBRULES-1307
14365	JBRULES-1273	can you attach it to our JIRA for us to look at?	1	0	https://issues.jboss.org/browse/JBRULES-1273
14366	JBRULES-1273	Rainer, Fernando has moved on from JBoss, so not sure what the status of his code review us. Will you be able to get this finished for Drools 5.0, meaning in the next 2 weeks? With unit tests and documentation?	0	0	https://issues.jboss.org/browse/JBRULES-1273

#	Requirement	Questions	Answered	Successfully Answered	URL
14367	JBRULES-1273	Hi Geoffrey, where does this requirement for incremental build support originate from? - is it a requirement for all Maven 3 plugins? or - is it a requirement for inclusion of the plugin into the Drools mainline? If it is a requirement to be accepted for mainline inclusion: does the drools compiler API support incremental compilation? Can you point me to the relevant interfaces and/or methods? As far as I know, the KnowledgeBuilder interface only has add() methods to add resources and does not support removal of packages or other artifacts from the knowledge base which originally were inserted as result of compilation of a given resource file. I'm not sure about current features of the knowledge agent, but last time I looked, the agent was also only capable of *adding* new files to a knowledge base. In case of rule refactoring (e. g. moving rule R from package A to package B, and an existing compiled knowledge base where rule R is already compiled into package A), how can the compiler act sensibly during "incremental compilation"? As we cannot tell from which file the rule R in package A originated, and there is no methot to remove this from the knowledge base either, I see no use of *adding* another rule R to package B (as the build would be totally inconsistent with what is stored in the files on disk. This is different to java compilation, where each java file is compiled into a number of classes in a reproducible manner. Given a modified java file, one can tell exactly which classes have to be overwritten. Maybe this is just my personal taste, but since incremental compilation introduces a lot of chances for inconsistencies, I'd suggest this can be postponed to a later version of a compiler plugin, in favour of bringing value to potential users earlier with a basic feature set. What do you mean by "drl compiler plugin MUST support incremental builds"?	1	0	https://issues.jboss.org/browse/JBRULES-1273
14368	JBRULES-1273	We can comprise to leave that "increment build" requirement out of the first version but with the info in the URL's it should be doable. Any non-beta version of a compiling plugin such as this one, should support - incremental builds - bulk compilation (do not compile 1 file at a time): so it can scale to 100's of DRL files (and more). Of course, both might need changes in the drools-api (haven't checked yet).	0	0	https://issues.jboss.org/browse/JBRULES-1273
14369	JBRULES-1260	Can we please re-test this again as another user has confirmed that this also happens to him.	0	0	https://issues.jboss.org/browse/JBRULES-1260
14370	JBRULES-1260	Are you using Eclipse3.2?	1	1	https://issues.jboss.org/browse/JBRULES-1260
14371	JBRULES-1260	Jesse: Would you please share your solution (if you remember).	0	0	https://issues.jboss.org/browse/JBRULES-1260
14372	JBRULES-1260	I don't think this is fixed. You still have to add that library-dependency yourself for a drools project (at least for me, fresh Eclipse 3.5 with JBoss Tools), which is a bit tricky, at least for newbies. Isn't there any easier way available?	0	0	https://issues.jboss.org/browse/JBRULES-1260
14373	JBRULES-1195	Could you attach a zip of your stuff.	1	1	https://issues.jboss.org/browse/JBRULES-1195
14374	JBRULES-1195	haven't had a chance to look into it - but will. How is this different from current DT support? Configurable CH is a good idea - although the GWT is hard coded as reflection/lazy loading is now allowed (but there is probably a way around it by loading the config on the server and squirting it to the client to cache as a map).	0	0	https://issues.jboss.org/browse/JBRULES-1195
14375	JBRULES-1185	Kris, can you please take a look and fix the Audit View for the new events?	1	1	https://issues.jboss.org/browse/JBRULES-1185
14376	JBRULES-1176	Got this idea from bug reporting: Should be able to see "my rules" - rules that you are the creator of, or perhaps recently contributed to. Also (permission willing) be able to specify other user names, and be able to do the same with them.	0	0	https://issues.jboss.org/browse/JBRULES-1176
14377	JBRULES-1155	forget decision tables for the moment and show me how this would be done in standard DRL. How can the consequence bind to some data matched on the LHS? how does it know which pattern or which field?	0	0	https://issues.jboss.org/browse/JBRULES-1155
14378	JBRULES-1155	So say you have a "variable" which is really a service, you could pass in "answerService". Then in the actions of the rules, you can do: answerService.setAnswer(true); etc... [Show »] Michael Neale [15/Sep/07 01:09 AM] If you use Variables - you can set a type and a name, and then pass this into the working memory (stateful session) - the variables is then accessible on the RHS of a rule. So say you have a "variable" which is really a service, you could pass in "answerService". Then in the actions of the rules, you can do: answerService.setAnswer(true); etc... Have you tested this? B/c this is how I got started down this road in the first place. I tried to do what you're suggesting and it didn't work unless the variable was defined on the LHS.	0	0	https://issues.jboss.org/browse/JBRULES-1155
14379	JBRULES-1119	Is the ant task for this something you have written? as the DSL file can be read in from anywhere (just not sure what can be done to do this - if you are loading rules from the API you can store things however you want). Or is this the drools-ant module?	0	0	https://issues.jboss.org/browse/JBRULES-1119
14380	JBRULES-1040	As far as I can tell, it's not working: rule "constraintsBroken" when $multipleQueensHorizontal : Long() from accumulate($q1 : Queen($id : id, $y : y); Queen(id > $id, y == $y); , count($q1)); $multipleQueensAscendingDiagonal : Long() from accumulate($q2 : Queen($id : id, $ascendingD : ascendingD); Queen(id > $id, ascendingD == $ascendingD); , count($q2)); $multipleQueensDescendingDiagonal : Long() from accumulate($q3 : Queen($id : id, $descendingD : descendingD); Queen(id > $id, descendingD == $descendingD); , count($q3)); then scoreCalculator.setScore(- $multipleQueensHorizontal - $multipleQueensAscendingDiagonal - $multipleQueensDescendingDiagonal)); end this gives problems on the ";" at each end. so I tried this: rule "constraintsBroken" when $multipleQueensHorizontal : Long() from accumulate($q1 : Queen($id : id, $y : y) Queen(id > $id, y == $y) , count($q1)); $multipleQueensAscendingDiagonal : Long() from accumulate($q2 : Queen($id : id, $ascendingD : ascendingD) Queen(id > $id, ascendingD == $ascendingD) , count($q2)); $multipleQueensDescendingDiagonal : Long() from accumulate($q3 : Queen($id : id, $descendingD : descendingD) Queen(id > $id, descendingD == $descendingD) , count($q3)); then scoreCalculator.setScore(- $multipleQueensHorizontal - $multipleQueensAscendingDiagonal - $multipleQueensDescendingDiagonal)); end which yields: aused by: org.drools.rule.InvalidRulePackage: [69,6]: unknown:69:6 mismatched token: [@178,2093:2093=',',<11>,69:6]; expecting type RIGHT_PAREN[69,14]: unknown:69:14 Unexpected token '$q1' at org.drools.rule.Package.checkValidity(Package.java:412)	1	1	https://issues.jboss.org/browse/JBRULES-1040
14381	JBRULES-1023	This is a nice to have, but for 4.2 at the earliest.	1	1	https://issues.jboss.org/browse/JBRULES-1023
14382	JBRULES-959	// Method used to generate a single random fact: private BOBOMessage getRandomMessage() { Random generator = new Random(); BOBOMessage msg = new BOBOMessage(); String randomMsg = (generator.nextBoolean()?"BOBO ":" ") + (generator.nextBoolean()?"BIBI ":" ") + (generator.nextBoolean()?"BABA ":" ") + (generator.nextBoolean()?"CACA ":" ") + (generator.nextBoolean()?"POOP ":" ") + (generator.nextBoolean()?"SUPER ":" ") + " this is some text to make the message longer, don't you like it?"; msg.setMessage(randomMsg); logger.debug("Random message generated: " + randomMsg); return msg; }	1	1	https://issues.jboss.org/browse/JBRULES-959

#	Requirement	Questions	Answered	Successfully Answered	URL
14383	JBRULES-959	// Method called to retrieve from the cache if available or recompile and put in the cache otherwise public RuleBase getGroupRuleBase(RuleGroup rulegroup) { String ruleIdString = rulegroup.getId().toString(); logger.debug("RuleService - Asked to provide rulebase for group with id=" + ruleIdString); RuleBase ruleBase = null; boolean nullCache = false; if (rulesCache == null || this.cacheDisabled) nullCache = true; else { nullCache = false; Object retrieved = null; try { retrieved = rulesCache.get("/mitto/rules/group/", rulegroup.getId()); } catch (CacheException e) { logger.warn("Could not find rule group with id=" + ruleIdString + " in the cache", e); retrieved = null; } ruleBase = (retrieved == null)?null:(RuleBase) retrieved; } if (EntityRule.isGroupInvalidated(rulegroup) || ruleBase == null) { logger.debug("RuleService - Generate rulebase for rule group with id=" + ruleIdString); Reader dsl = new StringReader(rulegroup.getPackage().getDSL()); Reader drl = new StringReader(EntityRule.getDRLByGroup(rulegroup)); logger.debug("DRL of the group with id=" + ruleIdString + ": " + EntityRule.getDRLByGroup(rulegroup)); /*try { logger.debug("Expanded DRL of the rule with id="+ ruleIdString + ":" + getExpandedDRL(EntityRule.getDRLByGroup(rulegroup), dsl)); } catch (DroolsParserException e) { logger.error("Drools parsing error", e); } dsl = new StringReader(rulegroup.getPackage().getDSL());*/ PackageBuilderConfiguration pkgBuilderCfg = new PackageBuilderConfiguration(); //pkgBuilderCfg.setCompiler(PackageBuilderConfiguration.JANINO); pkgBuilderCfg.setCompiler(PackageBuilderConfiguration.ECLIPSE); PackageBuilder builder = new PackageBuilder(pkgBuilderCfg); try { builder.addPackageFromDrl(drl, dsl); } catch (DroolsParserException e) { logger.warn("Drools could not parse the DRL for rule group with id=" + ruleIdString, e); } catch (IOException e) { logger.warn("IO error when Drools could not parse the DRL for rule group with id=" + ruleIdString, e); } if (builder.hasErrors()) logger.error("Drools errors happened when trying to compile the rule group with id=" + ruleIdString + ":" + builder.getErrors().toString()); Package pkg = builder.getPackage(); ruleBase = RuleBaseFactory.newRuleBase(); try { ruleBase.addPackage(pkg); } catch (Exception e) { logger.error("Drools error happened when trying to add the package to the Rulebase for rule group with id=" + ruleIdString, e); } if (!nullCache) { try { rulesCache.put("/mitto/rules/group/", rulegroup.getId(), ruleBase); } catch (CacheException ex) { logger.error("Cache error when tryin to put rule group with id= " + ruleIdString + " into the cache", ex); } } EntityRule.setGroupInvalidated(rulegroup, false); } else logger.debug("RuleService - Found rulebase in cache for rule group with id=" + ruleIdString); return ruleBase; }	1	1	https://issues.jboss.org/browse/JBRULES-959
14384	JBRULES-894	is this finished?	0	0	https://issues.jboss.org/browse/JBRULES-894
14385	JBRULES-851	I'm sure I'm not the first to think of this, but just to throw it out there: Would it be possible to contain this functionality within an Applet, and communicate the data to/from javascript using exposed methods on the applet? I promise my suggesting use of an Applet isn't an April Fool's joke :)	1	0	https://issues.jboss.org/browse/JBRULES-851
14386	JBRULES-834	how do we specify svn ignore files?	1	1	https://issues.jboss.org/browse/JBRULES-834
14387	JBRULES-834	ok i'm adding stuff now. We tend to want .project, .setting and .classpath commited. what is local?	1	1	https://issues.jboss.org/browse/JBRULES-834
14388	JBRULES-833	hmmm, I just remembered. if the .project and .classpath files are committed, why are they in svn:ignore in drools-core? It's a trick issupose?	1	1	https://issues.jboss.org/browse/JBRULES-833
14389	JBRULES-833	Do we use maven's Tycho to build the drools eclipse plugin? If not, we probably should: http://www.sonatype.com/people/2008/11/building-eclipse-plugins-with-maven-tycho/ (there might be a more up to date page then that somewhere). from that page: "Can we build eclipse plugins with maven? yes we can!"	0	0	https://issues.jboss.org/browse/JBRULES-833
14390	JBRULES-791	Is there any update on this issue?	1	0	https://issues.jboss.org/browse/JBRULES-791
14391	JBRULES-706	Fer - how is this one going? this is pretty critical for CR release, otherwise we will really annoy people.	0	0	https://issues.jboss.org/browse/JBRULES-706
14392	JBRULES-693	when a rule is in a package, we can validate it in the repo. This can happen on save or checkin. The question - shoudl this happen on save - if so, how should the feedback of errors be provided? should it be form the rule view, or the package view, should one rule at a time be compiled - or a package as a whole?	1	1	https://issues.jboss.org/browse/JBRULES-693
14393	JBRULES-693	Michael, I'm not sure about the meaning of "saving" a rule vs "checking in" a rule, so I will give my opinion using different words... :) My understanding is that a user must be allowed to do draft versions of their rules and only when he tries to promote them is that the validation should be enforced. So, from tooling perspective, I think validation may occur automatically (like when saving a rule in eclipse) or there may be a button that triggers it (like in antlrworks), but that validation should not block any action. When the user tries to promote a rule, then the validation should be enforced and in case problems are found, the rule is not promoted. Is that in line with what you are thinking?	1	0	https://issues.jboss.org/browse/JBRULES-693
14394	JBRULES-671	(Comments from Burr Sutter, some may no longer be relevant as this was from an old mockup). 1) My Fact list could get really long therefore the drop-down list to choose a fact in the add condition pop-up dialog might be cumbersome. 2) It is not clear how I add facts to the overall system. Facts might also have security aspects as well. 3) I'm not sure if "package" makes sense to an average end-user. I would have thought the screen above was under "rules". 4) I'm not sure how to save my edits 5) It would be nice to see the generated .drl and have free-form editing. While I realize this is an end-user tool it seems to me that the person wielding this tool will still need to be very rules savvy to deal with "facts", "constraints", "condition type", "condition templates", etc. There is a lot of jargon being used here, some of which I don't know. Like what is a template condition/action? 6) In the pop-up dialogs (from the green plus sign), the close button should be more clearly labeled and the red X should be in the top right corner. I would offer the user both ways to escape the modal dialog box. 7) It is unclear to me how you pull a series of If-Then "packages" into a single .drl for deployment. 8) Will there be a screen that shows me the "version" associated with a deployed/deployable unit? 9) How do I know which version is deployed into production? 10) If you click around the menu too many times it stops working altogether (in FireFox 1.0.7 on Windows). The work around is to restart Firefox. 11) What will the security model be like? Rule reviewers, rule editors, rule packagers (the deployable unit), rule publishers (pushing the unit to production)? 12) Will we have the concept of development, qa, staging and then production to allow for testing? There really needs to be a solution to basic rule editing and running basic tests. Perhaps have some test facts loaded into the repository. 13) If testing is allowed how does someone run the automated tests?	0	0	https://issues.jboss.org/browse/JBRULES-671
14395	JBRULES-664	Has this been checked in yet? I didn't see any changes some in from SVN	0	0	https://issues.jboss.org/browse/JBRULES-664
14396	JBRULES-626	Fer - do you think you will have time for this in the next week or 2? if not, we can punt this to 4.1 to make way for more critical issues (ie we can live without it, but the tables we have are very ordinary, visually and otherwise, so we will need to revisit them).	1	1	https://issues.jboss.org/browse/JBRULES-626
14397	JBRULES-576	can we have individual pom.xml's? having some trouble applying the patch file.	0	0	https://issues.jboss.org/browse/JBRULES-576

#	Requirement	Questions	Answered	Successfully Answered	URL
14398	JBRULES-546	What town pays 100000 to teachers? (disabled ones at that). Is this affirmative action on steriods?	1	1	https://issues.jboss.org/browse/JBRULES-546
14399	JBRULES-537	Evaluators can be used to compare objects/primitives of different types. If we cast how can we comapare Integer to a BigInteger?	1	1	https://issues.jboss.org/browse/JBRULES-537
14400	JBRULES-450	is this still deprecated? the test is in trunk and passing?	0	0	https://issues.jboss.org/browse/JBRULES-450
14401	JBRULES-443	Juergen, Do you still really need this? our build system isn't setup to handle multiple jars, and its quite a hassle to do. I think I prefer people to change the maven setting and build themselves, as building is very easy now. you ok with that?	1	0	https://issues.jboss.org/browse/JBRULES-443
14402	JBRULES-442	HotSpot(TM) 64-Bit Server VM (build 1.5.0_09-b03, mixed mode) Linux version 2.6.17-1.2142_FC4smp (bhcompile@hs20-bc1-6.build.redhat.com) (gcc version 4.0.2 20051125 (Red Hat 4.0.2-8)) #1 SMP Tue Jul 11 22:59:20 EDT 2006 Steps: - download 3.0.5-src + extract - mvn install Exception: (after adding try-catch-printstacktrace-rethrow in all 3 tests of org.drools.reteoo.ReteooBuilderTest) [surefire] Running org.drools.integrationtests.waltz.LeapsWaltzTest [surefire] Tests run: 1, Failures: 0, Errors: 0, Time elapsed: 3.782 sec com.thoughtworks.xstream.converters.ConversionException: Cannot externalize class java.lang.Class ---- Debugging information ---- required-type : org.drools.reteoo.ReteooRuleBase cause-message : org.drools.reteoo.BetaNode; local class incompatible: stream classdesc serialVersionUID = 4254294774027280987, local class serialVersionUID = -2179114169553759563 class : org.drools.reteoo.ReteooRuleBase message : Cannot externalize class java.lang.Class line number : 578 path : /org.drools.reteoo.ReteooRuleBase cause-exception : java.io.InvalidClassException ---------- at com.thoughtworks.xstream.converters.reflection.ExternalizableConverter.unmarshal(ExternalizableConverter.java:113) at com.thoughtworks.xstream.core.TreeUnmarshaller.convertAnother(TreeUnmarshaller.java:38) at com.thoughtworks.xstream.core.ReferenceByXPathUnmarshaller.convertAnother(ReferenceByXPathUnmarshaller.java:39) at com.thoughtworks.xstream.core.TreeUnmarshaller.start(TreeUnmarshaller.java:99) at com.thoughtworks.xstream.core.ReferenceByXPathMarshallingStrategy.unmarshal(ReferenceByXPathMarshallingStrategy.java:12) at com.thoughtworks.xstream.XStream.unmarshal(XStream.java:552) at com.thoughtworks.xstream.XStream.unmarshal(XStream.java:540) at com.thoughtworks.xstream.XStream.fromXML(XStream.java:497) at org.drools.reteoo.ReteooBuilderTest.checkRuleBase(ReteooBuilderTest.java:86) at org.drools.reteoo.ReteooBuilderTest.testThreeColumnsWithConstraints(ReteooBuilderTest.java:30) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at junit.framework.TestCase.runTest(TestCase.java:154) at junit.framework.TestCase.runBare(TestCase.java:127) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:118) at junit.framework.TestSuite.runTest(TestSuite.java:208) at junit.framework.TestSuite.run(TestSuite.java:203) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.apache.maven.surefire.battery.JUnitBattery.executeJUnit(JUnitBattery.java:242) at org.apache.maven.surefire.battery.JUnitBattery.execute(JUnitBattery.java:216) at org.apache.maven.surefire.Surefire.executeBattery(Surefire.java:215) at org.apache.maven.surefire.Surefire.run(Surefire.java:163) at org.apache.maven.surefire.Surefire.run(Surefire.java:87) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at	0	0	https://issues.jboss.org/browse/JBRULES-442
14403	JBRULES-380	whats the status of this? is ti just asking for examples, or do we now have examples to add?	0	0	https://issues.jboss.org/browse/JBRULES-380
14404	JBRULES-363	Have you made sure that the build system still works for JDK1.4 build systems? As we build the relese distribution on JDR1.4	1	1	https://issues.jboss.org/browse/JBRULES-363
14405	JBRULES-360	Is this a huge project? Does anyone else care about this? Are there any major pitfalls or concerns that would make this more difficult than it seems?	0	0	https://issues.jboss.org/browse/JBRULES-360
14406	JBRULES-296	just built the new plugin, content assist is starting to get really good. Would it be possibly to make sure only valid operators and values are suggested, for the current location? So it should only list valid variables for comparators, which includes valid types and also it must be defined before the current location.	1	1	https://issues.jboss.org/browse/JBRULES-296
14407	JBRULES-266	If assertion behaviour is WM_BEHAVIOR_EQUALS (either for normal or logically asserted facts) there are potential problems in the case that a modify of a fact happens, that makes it equal to some other fact. Do they "merge" into one fact or continue to exists equal but separate? Currently I believe they continue to exist as equals, but that yields problems, as some checks I assume can't distinguish between the two facts or they appear as one. Testcase test_ModifyToEquality.drl attached. I had to modify org.drools.Cheese to CheeseEqual and implement hashCode() and equals(). Classfile can be found at: http://jira.jboss.com/jira/browse/JBRULES-233 This is in general an interesting problem and more a matter of definition of behaviour of modify I guess. If modify is interpreted as rectract+assert then the facts should be merged. But that would yield problems as e.g. the 2 fact-handles pointing to 2 separate objects would then either point to the same or one fact-handle would become invalid. In case of merging and logically assertion: Justifications would have to be updated, for the 2 merging objects + objects logically depending on these two. JESS 7.0b7 does also have problems with that. Attached modifyToEquality.clp, yields errors for cases described.	1	1	https://issues.jboss.org/browse/JBRULES-266
14408	JBRULES-266	logical assertions are always equality based. Stated facts can now use identity and equality - using your configuration settings. I have added another configuration setting to dictate how logical over-ride works. Does it discard the original fact, or turn it into two stated fact handles.	0	0	https://issues.jboss.org/browse/JBRULES-266
14409	JBRULES-240	are there any unit tests to go with these?	1	1	https://issues.jboss.org/browse/JBRULES-240
14410	JBRULES-240	ok I've added the two classes. I've made them impl an interface and changed the method to use them. However we still really need unit tests for these. Can you clean up what you have and improve it, so its suitable?	0	0	https://issues.jboss.org/browse/JBRULES-240

#	Requirement	Questions	Answered	Successfully Answered	URL
14411	JBRULES-233	I assume a logically asserted object should also be retracted (or a justification removed), if the rule that caused the assertion is removed (so A => B means that B is not only justified by A, but by the existence of rule A => B). Logically asserted objects are not retracted if a rule is removed. There seems to be some ClassCastException in the combination of modify object + logically assertion + remove rule. Comment out the currently failing assertions in the supplied testcase: ------- // list = workingMemory.getObjects(c1.getType().getClass()); // assertEquals("remove of rule should retract objects logically asserted based on the rule", 2, list.size()); // assertTrue("remove of rule should retract objects logically asserted based on the rule", list.contains(c1.getType())); // assertTrue("remove of rule should retract objects logically asserted based on the rule", list.contains(c2.getType())); // assertFalse("remove of rule should retract objects logically asserted based on the rule", list.contains(c3.getType())); ----- This will yield for "workingMemory.modifyObject(h, c2);" in the Testcase: java.lang.NoClassDefFoundError: org/drools/test/Rule_rule1_0 at org.drools.test.Rule_rule1_0PredicateOInvoker.evaluate(Rule_rule1_0PredicateOInvoker.java:13) at org.drools.rule.PredicateConstraint.isAllowed(PredicateConstraint.java:84) at org.drools.common.BetaNodeBinder.isAllowed(BetaNodeBinder.java:57) at org.drools.reteoo.LeftInputAdapterNode.modifyObject(LeftInputAdapterNode.java:205) at org.drools.reteoo.ObjectSource.propagateModifyObject(ObjectSource.java:138) at org.drools.reteoo.ObjectTypeNode.modifyObject(ObjectTypeNode.java:162) at org.drools.reteoo.Rete.modifyObject(Rete.java:164) at org.drools.reteoo.RuleBaseImpl.modifyObject(RuleBaseImpl.java:265) at org.drools.reteoo.WorkingMemoryImpl.modifyObject(WorkingMemoryImpl.java:725) at org.drools.reteoo.WorkingMemoryImpl.modifyObject(WorkingMemoryImpl.java:685) at org.drools.integrationtests.IntegrationCases.testLogicalAssertionsDynamicRule(IntegrationCases.java:1966) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:324) at junit.framework.TestCase.runTest(TestCase.java:154) at junit.framework.TestCase.runBare(TestCase.java:127) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:118) at junit.framework.TestSuite.runTest(TestSuite.java:208) at junit.framework.TestSuite.run(TestSuite.java:203) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.runTests(RemoteTestRunner.java:436) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.run(RemoteTestRunner.java:311) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.main(RemoteTestRunner.java:192)	1	1	https://issues.jboss.org/browse/JBRULES-233
14412	JBRULES-231	"not contains" was not implemented, since we have "excludes". Do you want me to add "not contains" as a synonym for "excludes"?	0	0	https://issues.jboss.org/browse/JBRULES-231
14413	JBRULES-210	Unless we create our own editor, I think we'll have to put the drl files in a directory structure that represents the drl's package. Should this be validated (error if rule file is not in correct directory)? Should check whether this has any impact on rule loading (specifying the drl should also contain the directory of course). Create Rule wizard should be altered to create folders representing package too (if necessary). Example project should be changed to that drl files are in correct package.	0	0	https://issues.jboss.org/browse/JBRULES-210
14414	JBRULES-210	ideally this should have been done before RC :(Can we do it quickly?	0	0	https://issues.jboss.org/browse/JBRULES-210
14415	JBRULES-149	another error in 2.1 its Production Memory, as its asserted did you mean it's rather than its? its Production Memory, as it's asserted	0	0	https://issues.jboss.org/browse/JBRULES-149
14416	JBRULES-144	I'm going to use concurrency classes from Dawid Kurzyniec's java 1.4 backport of Doug Lea's java.util.concurrent. http://dcl.mathcs.emory.edu/util/backport-util-concurrent/ Are there other aspects we want to make asynchronous? I'm thinking of making a ConcurrentWorkingMemory that's just a proxy for another WorkingMemory implementation. (And will itself implement WorkingMemory) To start with only the assert() methods would be made asynchronous. Everything else would be a synchronous passthrough call. Sound good?	0	0	https://issues.jboss.org/browse/JBRULES-144
14417	JBRULES-144	I've attached an example of how I think the asynch assert should be implemented. It's an asychronous wrapper of any WorkingMemory that guarantees the lifecycle methods will be executed in a seperate thread, in the order they were invoked. I used the backported concurrency classes to do this: http://dcl.mathcs.emory.edu/util/backport-util-concurrent Is this a worthy approach? Method return values and exception handling are facilitated by the Future returned by each method. If a FutureTask throws an exception while executing, it will result in an ExecutionException being thrown when you call get() to retrieve the result. Methods that normally return void now return null from get().	0	0	https://issues.jboss.org/browse/JBRULES-144
14418	JBRULES-144	I don't have a problem with "breaking" the working memory interface in the case of concurrent. but I don't want the original one to change. The working memory interface implies semantics that are synchronous. I don't think that it is a concern that you can't just "swap" out and "plug in" the async version: moving to an async model is not trivial for a user. So any interface should make it obvious how to use the concurrent features. Looks good to me though ! what sort of environments will this be used in? As on the server, we are little limited as to what we can do with threads. But JMS is far to heavy to be useful for queuing asserts.	0	0	https://issues.jboss.org/browse/JBRULES-144
14419	JBRULES-144	Some general thoughts on this. From past experience with running rule engines in reactive mode for real-time OMS. It isn't necessary to return null in fireAllRules. Take for example JESS, fire() in JESS just returns an int count of the rules fired. When JESS runs in reactive mode, there's no gaurantee an assert will result in rule activation. Therefore running in async mode, the rule engine should provide a way to return results. In my slot-based rule engine, I was considering 3 potential techniques. 1. use listeners to pass rule fire events. the specific technique I was thinking of is to use the asserted object as the source for the event. Each thread using the shared engine then filters the events to see which ones were caused by facts they asserted. 2. provide a generic getRulesFired() method that returns a List. this approach is very simple, but not pratical for a couple of reasons. I. with many threads it quickly blows chunks because each threads has to keep inspecting the list II. how does an application know when to look for results? III. the list could get huge very quickly. especially for pre-trade compliance 3. use a messaging system like JMS to handle inbound objects In a typical event driven system, the data is sent using JMS and the response is published to all recipients. In fact I wrote a JMS client for JESS to do exactly that. Back to the point I was going to make. There are several ways to get the concurrent stuff to fit within the existing WorkingMemory interface. One way to do that is to actually start an instance of the engine in a thread, then the wrapper can simply call asserts using executor. hope that provides some food for thought. peter	1	1	https://issues.jboss.org/browse/JBRULES-144
14420	JBRULES-135	Are these shoots as you wish ?	1	1	https://issues.jboss.org/browse/JBRULES-135

#	Requirement	Questions	Answered	Successfully Answered	URL
14421	JBRULES-118	Is there another way? this is really not nice. What would be nice woudl be if the parser could punt to the semantic compiler in place... just an idea.... but I guess that is against the whole AST idea.	1	1	https://issues.jboss.org/browse/JBRULES-118
14422	JBRULES-71	Is there any way to specify "not matches"? There's no (?x)-style switch you can put inside the java regex. I need something like rule when Email(from matches "me@.*", to !matches "you@.*") then doSomething(); end I can imagine some ways to do what I need with activation groups, and it's easy with an eval, but I don't see a good way to do what I have above.	1	1	https://issues.jboss.org/browse/JBRULES-71
14423	JBRULES-61	yeah the reading part is the harder (and more important) one for now. Once that is reliable, and with a solid XML schema, exporting should be quite easy. It only a rare case that people will actually export, but its a nice item to have, as it means less apparent "lock on" to the DRL (at least according to PHBs). Why are we designing for PHBs? ;)	1	1	https://issues.jboss.org/browse/JBRULES-61
14424	JBRULES-48	whats the status on this, can we close?	1	1	https://issues.jboss.org/browse/JBRULES-48
14425	JBRULES-46	Will this feature be available in 5.0? Thanks.	1	1	https://issues.jboss.org/browse/JBRULES-46
14426	JBRULES-46	Is there any update on backward chaining support in Drools?	0	0	https://issues.jboss.org/browse/JBRULES-46
14427	JBRULES-46	So, I take it these updates mean that backward chaining is being/has been developed? I would love to know more if there was some documentation somewhere. Is backward chaining semi-usable in trunk?	0	0	https://issues.jboss.org/browse/JBRULES-46
14428	JBRULES-39	Yes have something basic working with 2.x. The question is: should it generate a static main example rulebase invoker, or a TestCase? Further to this, for this to properly work, it will need to have a drools "project" concept where a new project is created with the classpath deps already setup (otherwise the example will not compile). It may also be possible to modify the projects classpath to include appropriate deps, like when you chose to create a new JUnit testcase in eclipse.	0	0	https://issues.jboss.org/browse/JBRULES-39
14429	JBRULES-11	has this been done, can we close it off?	0	0	https://issues.jboss.org/browse/JBRULES-11
14430	JBRULES-5	Do we need to set something up? Why not just use the codehaus one and point a domain there? I'm sure bob will set you up with admin priveliges etc. My reasoning for this is there are already quite a few people on irc.codehaus.org so it will make it easy for these people to "pop in" and see other channels. Or the alternative is to request bob to allow irc.codehaus.org to live at jboss, which I think he would also be happy with. Just trying to think of a way to leverage communities.	1	1	https://issues.jboss.org/browse/JBRULES-5
14431	JBRULES-5	Shall I mark this as a won't fix? Even the idea of pointing irc.jboss.org at codehaus?	1	1	https://issues.jboss.org/browse/JBRULES-5
14432	JBRULES-3	Mark, How is this going? Anything I can do to help? I'd like to get this task completely resolved and JBoss Rules turned on in Labs. Cheers, Damon	1	1	https://issues.jboss.org/browse/JBRULES-3
14433	JBRULES-2	Will we continue to use gmane as our main archive site? Or do you plan to host the gmane software internally. Also speak to Bob about mailing lists as I believe he said he has some good configurations/advice or something along those lines :)	0	0	https://issues.jboss.org/browse/JBRULES-2
14434	JBINSTALL-346	PATCH NAME: JBINSTALL-346 PRODUCT NAME: JEMS Installer VERSION: 1.2.0.GA SHORT DESCRIPTION: Fix for Automated installer, generates appropriate InstallConfigRecord.xml LONG DESCRIPTION: The problem before this fix was the InstallConfigRecord.xml generated by Installer (after user completes the GUI Installation) wasn't correct and hence Unattended mode was failing MANUAL INSTALL INSTRUCTIONS: NA COMPATIBILITY: NA DEPENDENCIES: NA SUPERSEDES: NA SUPERSEDED BY: NA CREATOR: amit bhayani DATE: 22 Aug 2007	1	1	https://issues.jboss.org/browse/JBINSTALL-346
14435	JBINSTALL-319	Below said is how I will design the Tree Application Server |_ everything that it has currently |_ jmx-core |_ docs |_ client |_ RMI |_ JCA |_ JTA |_ JMX |_ Clustering |_ minimal EJB EJB3 JMS Portal (In case of Portal profile selected) The idea is that once user selects top level node everything should work out of the box. For example in current scenario Application Server and JCA are peers but one cannot work without other being selected. Let us keep it very simple. Also I see that once I de-select top level node and select it back doesn't cause all the child nodes to be selected automatically. Isn't it suppose to behave like that? The child nodes becomes disabled and I cannot select them anymore.	1	1	https://issues.jboss.org/browse/JBINSTALL-319
14436	JBINSTALL-319	"Also I see that once I de-select top level node and select it back doesn't cause all the child nodes to be selected automatically. Isn't it suppose to behave like that? The child nodes becomes disabled and I cannot select them anymore. " Yes, it behaves like that because it uses the PacksModel behaviour (the same a in PacksPanel). But i agree that selecting a root node should select all child nodes with all their dependencies would be the best behaviour, i was already experimenting on that (there are some corner cases like excludeGroups and preselected="no" items)	0	0	https://issues.jboss.org/browse/JBINSTALL-319
14437	JBINSTALL-297	Just committed the latest http://repository.jboss.com/jems-installer-web-repository/test-7-6-2007/jems-web-installer-1.2.1.RC2.jar Is this really supposed to go in 1.2.1 GA?	0	0	https://issues.jboss.org/browse/JBINSTALL-297
14438	JBINSTALL-296	Should we give an option to user if he/she wants .zip of installed components? May be in same screen as that of User Short Cut Screen or we do it by default anyway? Giving an option will require to modify the Short Panel.	1	1	https://issues.jboss.org/browse/JBINSTALL-296
14439	JBINSTALL-212	The script should picked up the above 3 strings if the ID mapping works properly, then PO files should include such strings. Could you please kindly confirm if eng.xml got that strings within?	1	1	https://issues.jboss.org/browse/JBINSTALL-212
14440	JBINSTALL-207	great! do you always create desktop shortcuts or only when user wants to?	0	0	https://issues.jboss.org/browse/JBINSTALL-207
14441	JBINSTALL-150	Won't fix because it's it sync with the ejb3 server coming from the 4.0.5.GA build.	0	0	https://issues.jboss.org/browse/JBINSTALL-150
14442	JBINSTALL-148	Is the fix for JBAS-3839 already in JBoss_4_0_5_GA_EJBTHREE-798 branch?	1	1	https://issues.jboss.org/browse/JBINSTALL-148
14443	JBINSTALL-140	Alex, there are about 10 other languages enabled for selection in the JEMS installer and we are not planning to translate into these languages soon. Should I disable them for now? (note that the izpack team has done some translation for the basic interface in these langs, but nothing jboss-specific)	1	1	https://issues.jboss.org/browse/JBINSTALL-140
14444	JBINSTALL-134	what is the scope of this task? InstallationGroupPanel? UserInoutPanel's?	1	1	https://issues.jboss.org/browse/JBINSTALL-134
14445	JBINSTALL-133	1. In InstallGroupPanel - the column names are not localizable and are stuck in english. 2. In the same panel - the installation group names are not localizable. 3. In the packs selection panel the "Requires:" and "Excludes" words are stuck in english, but this will be fixed in the next patch of izpack.jboss. 4. There are localized strings from the original izpack translations that can't be externally modified (*but they are already translated*). Do we care about these? 5. The localized installer has different dimensions when a non-english langpack is selected (possibly to fit the new strings?) 6. the "Created with Izpack - http://www.izforge.com" string appears in the non-english version (translated in each of the languages - it's built-in, but externally modifiable). 7. Command line localization is not supported at all by izpack. Do we need it?	1	1	https://issues.jboss.org/browse/JBINSTALL-133

#	Requirement	Questions	Answered	Successfully Answered	URL
14446	JBINSTALL-133	on the first 7 issues you found: 1. InstallGroupPanel - separate issue, fix and sync up with izpack 2. Same as 1 3. Consider branching izpack and doing all I18N fixes in the branch 4. We should not need to care about this for now. I18N team will be testing these translations (provided to izpack by native speakers supposedly), and if they are not happy, we will just have to accept their translations and build a custom I18N version of the izpack. 5. Do you mean that window size is different from what's defined in install.xml? 6. How can you modify it? 7. Already commented on this - submit a feature request JIRA. EULA and Image localization: looks like this is supported by izpack. Documentation mentions that it's possible to have a separate resource for each localization suffixed with language code.	1	1	https://issues.jboss.org/browse/JBINSTALL-133
14447	JBINSTALL-131	Hello, That was a very good analysis and you have identified an issue that will happen down the track. I would like to add a couple of points with regards to RHEL5. [snip] I tried the fallback folder solution on Fedora 6 with JVM 1.5 and 1.4. In order the japanese text to render correctly the following lines have to be executed: # cd /usr/java/jre1.5/lib/fonts # mkdir fallback # cd fallback # ln -s /usr/share/fonts/japanese/TrueType/sazanami-gothic.ttf # ln -s /usr/share/fonts/japanese/TrueType/sazanami-mincho.ttf [snip] The above issue has been fixed in the fonts-japanese package with the symlink work around, not in the fallback dir though as RHEL5 dont have jre installed as default. An ideal solution should probably be something more dynamic like getting the font info using fc-match (fontconfig). Bug 217821: Cannot use japanese fonts on Java application with JRE <http://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=217821> [snip] for chinese: # cd /usr/java/jre1.5/lib/fonts # mkdir fallback # cd fallback # ln -s /usr/share/fonts/japanese/TrueType/ukai.ttf # ln -s /usr/share/fonts/japanese/TrueType/uming.ttf [snip] The situation is a bit different in RHEL5 because we are using ZYSong (commercial font) since previous version of RHEL. Hence, workaround is not required as JAVA 5 still reference to RHEL4 font configuration. [snip] Each java UI component has it's own pre-defined default font - it is one of the logical fonts and which one exactly depends on the GUI manager in Java (this is why in linux you can see some of the GUI componentes rendering correctly east-asian characters and other do not). [snip] This is going to be an issue later when we extend the localisation effort to include South Asia (indic locales) as in RHEL4, there are only 5 supported indic locale. In RHEL5, there are 12 supported indic locale. We can discuss on this in detail later.	0	0	https://issues.jboss.org/browse/JBINSTALL-131
14448	JBINSTALL-114	Resolved this. Please test taking source code from trunk jems-installer If you select Standard, Security will not be implemented. Standard is equivalent to .zip distro. Is this ok or by default we want to implement security?	1	1	https://issues.jboss.org/browse/JBINSTALL-114
14449	JBINSTALL-93	Amit, I have already done it. Should I commit it to the repo to check it? (xmlresolver is not included)	1	1	https://issues.jboss.org/browse/JBINSTALL-93
14450	JBINSTALL-72	I've updated the ejb3 profile to fix a couple missing jars and verified that the seam demos (namely the booking app and the dvd store) run against it. The non-seam dvdstore also runs fine. I'm still forced to use -Djava.net.preferIPv4Stack=true to run the jboss instance, though. Would it make sense to make this a default?	1	1	https://issues.jboss.org/browse/JBINSTALL-72
14451	JBINSTALL-46	I will make the Installer to put install-config.xml in the $INSTALL_PATH folder (the destination folder where the AS is being installed). The users can be simply instructed to send this file back. Is this ok?	1	1	https://issues.jboss.org/browse/JBINSTALL-46
14452	JBINSTALL-46	we need to make sure that it's going to be easy to fully reproduce customer's install by using install-config.xml. Inlcuding pack selections and all user input. Is this going to be possible to run installer and pass a reference to install-config.xml?	1	1	https://issues.jboss.org/browse/JBINSTALL-46
14453	JBINSTALL-46	Alex, when I generate the auto script using the FinnishPanel and when I run the install with this script everything seems to work fine. Do you remember what exactly is broken? And could you clarify what you mean by "silently uses it", I am not sure I undetstand?	1	1	https://issues.jboss.org/browse/JBINSTALL-46
14454	JBINSTALL-44	Should we include all default profile packs into ejb3/ejb3-clustered profiles?	1	1	https://issues.jboss.org/browse/JBINSTALL-44
14455	JBINSTALL-39	can we implement our own version of HelloPanel (hopefully by inheriting from it) in order to avoid having to modify izpack?	1	1	https://issues.jboss.org/browse/JBINSTALL-39
14456	JBINSTALL-39	We can do it. I think its good idea to not touch izpack so that we can upgrade to latest version without many modicfications. However I see that there are already many custom changes in our local izpack-installer. For example HelloPanel and FinishPanel is already modified to make use of GridBagLayout which is not there in actuall izpack project. The concern is do we implement all the changes in our own jboss specific classes and keep it external to izpack? I really like this idea as I dont want to change the original code.	0	0	https://issues.jboss.org/browse/JBINSTALL-39
14457	JGRP-2172	[~rvansa] If you look at my previous comment, are you suggesting the behavior I described is *not* suitable? IMO it is; the causality between P1 and P2 is only given if P1 is blocking and P2 is only sent once P1 returns. In this case, we'll never deliver P2 before P1. Otherwise, there is no causality and then it doesn't matter in which order P1 and P2 are delivered. Do you have an example which shows unwanted behavior?	1	1	https://issues.jboss.org/browse/JGRP-2172
14458	JGRP-2155	Hi Yves, the JGroups JAR is 2MB in size, which isn't really big. On top of that people can remove classes they don't need from the JAR. What's the goal of stripping down the JAR file size even further? Will JGroups run faster as a result? If you've conducted any tests that show a speedup, I'm willing to take a look. Otherwise, since there's quite some work involved (not sure I like the annotations based approach), I'm not inclined to do this.	1	1	https://issues.jboss.org/browse/JGRP-2155
14459	JGRP-2155	Hi, even if it is 2.5mb, i face the problem because Android limits the APK size. Removing class if not so easy. It is a try & fail tedious process because of the two main configuration files, the class instantiation by reflexion and all the dependencies. It is not at all a performance concern as i beleave that this point is governed by the protocol exchange. Another point is related to the security: I have not look this point deep, but i realized that knowing the connect name is enought to communicate with a JGroups instance. As all protocols are present, it may be easy for anyone to send commands. So i think it is important to limits the class to the minimum needed by the project configuration. As you still widely use the Annotations i think this may be the easiest way but you can do it as you will. Anyway, i think that you should move the demo and test packages to comply with maven recommendations. While i am here, can i ask you to use the sl4j framework and leave the logging choice to the user? This will also save the logging package. Thanks a lot for your work. It saved me a lot of sweat. Yves	1	0	https://issues.jboss.org/browse/JGRP-2155
14460	JGRP-2129	What did you do to reproduce this?	1	1	https://issues.jboss.org/browse/JGRP-2129
14461	JGRP-2109	I'm not an expert on S3, although I wrote the original S3_PING (copying code from the Amazon example). I know there are other S3 discovery implementations (you listed Zalando's impl). Have you asked them if they were planning on adding support for this?	1	1	https://issues.jboss.org/browse/JGRP-2109
14462	JGRP-2109	I recently rewrote nativ-s3-ping [1], can you take a look? I'd like to deprecate S3_PING as it started as a copy of AWS sample code, and doesn't use the official SDK. I suggest use AWS_PING or NATVE_S3_PING as replacement. [1] https://github.com/jgroups-extras/native-s3-ping	1	1	https://issues.jboss.org/browse/JGRP-2109
14463	JGRP-2109	Yes: https://issues.jboss.org/browse/JGRP-2170. Are you interested? :-)	0	0	https://issues.jboss.org/browse/JGRP-2109

#	Requirement	Questions	Answered	Successfully Answered	URL
14464	JGRP-2099	If you want to use a single thread pool, I recommend looking into {{ForkJoinPool}} - that's the one I'm thinking of switching Infinispan to. One of the things I like about FJP is that when you have to block, you can use {{ManagedBlocker}} to tell the pool it should allocate a new thread. FJP is also much better than using a classical {{ThreadPoolExecutor}} with a {{SynchronousQueue}} when it comes to context switches: {{SQ.add()}} has to block until a worker thread is free, so for each task either the submitter wakes up a worker, or the submitter parks, then a worker finishes a task and wakes up the submitter. When you use a FJP pool, and the submitter thread is from the same FJP, it will just add the task to its own queue, and the task may be processed without involving any other threads. I've read a long time ago about about an extreme version of this, a variation of {{WithinThreadExecutor}} that queues the task in a thread-local and always runs the tasks on the same thread. (Well, if you run on your own {{Thread}} subclass, you don't technically need a thread-local.) Sounds very interesting, but I think the FJP approach is more practical :)	1	1	https://issues.jboss.org/browse/JGRP-2099
14465	JGRP-2097	Just out of curiosity, when do you plan to remove the Channel class altogether, e.g. 4.1, 4.2, 5.0?	1	0	https://issues.jboss.org/browse/JGRP-2097
14466	JGRP-2063	[~rhusar] I want to release 3.6.10 today, is this a quick change?	1	1	https://issues.jboss.org/browse/JGRP-2063
14467	JGRP-2063	When will this get more thought? :-) I'm pushing it into 4.1, so we have more time.	0	0	https://issues.jboss.org/browse/JGRP-2063
14468	JGRP-2044	How do you see boxing/unboxing in JFR? And what was the percentage you were seeing? I've looked into this before in JGRP-1982 and decided against it. I've changed this in my private branch to use Trove's {{TLongObjectHashMap}}, but this * introduces a dependency which I don't want and * {{TLongObjectHashMap}} is not concurrent (as {{ConcurrentHashMap}}, and therefore every access needs to be synchronized, slowing things down	0	0	https://issues.jboss.org/browse/JGRP-2044
14469	JGRP-2021	Is the fix for this going to be included in EAP 7?	1	1	https://issues.jboss.org/browse/JGRP-2021
14470	JGRP-2021	OK, I'm done with the backport to the 3.6 branch. What are the next steps? For a review of the code on 3.6, use branch [1]. The document to start with is ./doc/design/MaliciousAttacks.txt. The 2 unit tests are {{SYM_ENCRYPT_Test}} and {{ASYM_ENCRYPT_Test}}. The 2 sample configs are sym-encrypt.xml and asym-encrypt.xml (in ./conf). [1] https://github.com/infinispan/private-jgroups/tree/JGRP-2055	1	0	https://issues.jboss.org/browse/JGRP-2021
14471	JGRP-2010	When sending a message with an AnycastAddress, no copy is needed, either	0	0	https://issues.jboss.org/browse/JGRP-2010
14472	JGRP-1989	TCPConnection.send() only makes a copy if you pass a direct ByteBuffer to it, otherwise it doesn't. Are you referring to code further up the stack?	1	1	https://issues.jboss.org/browse/JGRP-1989
14473	JGRP-1989	Actually, buffer pools can be used more widely: * Sending of message batches and single messages ** Also for {{TCP_NIO2}} once JGRP-1991 is in place * Reception of message batches ** The message batches are de-serialized from the buffer, so the buffer can be reused immediately after deserialization (in {{TP.handleMessageBatch}}) * Reception of single messages ** Can we assume that - after the {{receive()}} callback returned - we can reuse the buffer? ** Asynchronous Invocation API? ** User cannot store the buffer somewhere, but needs to copy if it wants to hang on to the buffer ** Possibly make this configurable	1	1	https://issues.jboss.org/browse/JGRP-1989
14474	JGRP-1985	Hmm, by default we have the following headers * Transport: {{UDP}}, {{TCP}} or {{TCP_NIO2}} * Reliable retransmission and ordering protocol: either {{NAKACK2}} or {{UNICAST3}} * RequestCorrelator for RPCs That's 3 headers most of the time. Which other headers did you see? {{FRAG2}}, {{FORK}}? Just out of curiosity, I'll change the default to 4 anyway.	1	1	https://issues.jboss.org/browse/JGRP-1985
14475	JGRP-1980	[~sannegrinovero] This is interesting: {{decrementIfEnoughCredits()}} is called when *sending* a message, and it is executed on a {{Credit}} object in the {{sent}} table. Method {{decrementAndGet()}} is called when a message is *received* on is executed on a {{Credit}} object in the {{received}} table, so the 2 {{Credit}} objects should be different, thus no contention! Which test did you run to come up with this? I'd like to reproduce this...	1	1	https://issues.jboss.org/browse/JGRP-1980
14476	JGRP-1976	I didn't want to make this the default because the current default tries to keep the existing coordinator, wheres the 'winning' policy just picks the first largest subview. What's wrong with setting {{GMS.membership_change_policy}} in your config ?	1	1	https://issues.jboss.org/browse/JGRP-1976
14477	JGRP-1966	Not sure I understand what the issue is. Also, I've never used either JDK logging or sl4j. Can you submit a pull request to fix what you perceive to be the issue? Meanwhile, I'm closing this one, feel free to reopen once you have a PR.	0	0	https://issues.jboss.org/browse/JGRP-1966
14478	JGRP-1962	Actually, I need this file to run my unit tests. Why can't you simply use {{-Dlog4j.configuration=/your/home/log4j2.xml}} ?	0	0	https://issues.jboss.org/browse/JGRP-1962
14479	JGRP-1945	Which version are you referring to ? JGroups switched from LGPL to ASL a few years ago...	1	1	https://issues.jboss.org/browse/JGRP-1945
14480	JGRP-1933	Can you submit a PR to fix this ? Yes, a patch is possible, too. I assume this would be a change in pom.xml ? I'm not too familiar with OSGi.	1	1	https://issues.jboss.org/browse/JGRP-1933
14481	JGRP-1933	Thank you, what's a PR? Also, yes, it is just in the pom file, i.e: {code} <plugin> <groupId>org.apache.felix</groupId> ... <configuration> <instructions> <Bundle-SymbolicName>camel-blueprint</Bundle-SymbolicName> <Import-Package>*</Import-Package> <DynamicImport-Package>*</DynamicImport-Package> </instructions> </configuration> ... </plugin> {code} Also, customer originally downloaded the jar from: http://mvnrepository.com/artifact/org.jgroups/jgroups/3.5.1.Final so it is a community version as opposed to an RH distribution	1	1	https://issues.jboss.org/browse/JGRP-1933
14482	JGRP-1927	Let's say A leaves the relay cluster first, then B tries to send a message {{m}} to the backup site, and then A leaves the local cluster. When B then becomes coordinator of the local cluster and site master, will RELAY2 on B pick up {{m}} and send it to the backup site? If B resends {{m}}, everything should be fine. But if B doesn't resend {{m}}, we'll eventually get a {{TimeoutException}} in Infinispan and we won't retry, so that wouldn't work for us.	1	1	https://issues.jboss.org/browse/JGRP-1927
14483	JGRP-1927	Do I still need to look into this, or can this issue be closed ?	1	1	https://issues.jboss.org/browse/JGRP-1927
14484	JGRP-1909	Do we have an Azure account that we can use to test a new discovery protocol ?	1	1	https://issues.jboss.org/browse/JGRP-1909
14485	JGRP-1901	Do you plan to ack the PREPARE_VIEW without leaving the old view?	1	1	https://issues.jboss.org/browse/JGRP-1901
14486	JGRP-1898	Good. Do you want me to release a version (which one, 3.4.7?), or is a tag good enough ? If you're consuming a Brew-created JAR, then a tag should do the job.	1	1	https://issues.jboss.org/browse/JGRP-1898
14487	JGRP-1898	Sorry for the incorrect backports (fixed now) ! So I'll release 3.4.7. For 3.6.1 I wanted to wait a bit... do you really need it, or is a tag good enough ?	1	1	https://issues.jboss.org/browse/JGRP-1898

#	Requirement	Questions	Answered	Successfully Answered	URL
14488	JGRP-1883	OK, now I understand. But what's the point of selling SASL for authentication and encryption while in fact the latter isn't even done via SASL mechanisms, but by ENCRYPT code facaded by the SASL protocol ? Why can't we say that SASL (the JGroups protocol) only supports authentication as SASL (the mechanism) isn't good enough to provide encryption ? I don't see anything wrong in saying "JGroups supports SASL authentication". The above sounds like a marketing rather than technical argument to me... *If* we decide to go ahead and do this, then we'd have to refactor the components that ENCRYPT is made of for them to be reused in SASL. Let's avoid copy&paste here. Tristan [~NadirX] ?	1	0	https://issues.jboss.org/browse/JGRP-1883
14489	JGRP-1883	No, don't close it yet ! I was just assessing your argument and if marketing/prod mgmt thinks that we need to support SASL all the way, then fine. Is this the case ? It's just that I wouldn't be able to do this, so if you or Tristan have time, no problem... Your argument that the SASL tools get enhanced outside of JGroups and then we benefit from that work by upgrading the SASL lib is certainly valid. I think you and Tristan need to discuss this (and perhaps Divya)...	1	0	https://issues.jboss.org/browse/JGRP-1883
14490	JGRP-1875	Hmm, there's a problem when a new connection is established: * A (conn-id=1) sends A1(first)-A10 to B * B receives A:3 first. Since this isn't tagged as {{first}}, B sends a SYNC to A * A hasn't seen that SYNC yet, so it changes its {{conn-id}} to 3 * B now receives A:1(first) and creates the connection with {{conn-id}} == 2 * B also receives the other messages and adds them to the connections recv-window * Now B receives the {{SYNC-OK}} message with {{conn-id}} == 3 ** B now removes the recv-window for A and creates a new one with {{conn-id}} == 3 This race between the first message and SYNC happens because we cannot guarantee that the first message from A is always received _first_ by B. I'm starting to dislike this whole SYNC business: we're fixing an edge case that happens 0.001% of the time, but introduce a complexity that's not warranted ! So what are the alternatives ? h5. Do nothing (status quo) / first message creates the recv-window * With {{conn-close-timeout}} set to a reasonably large value, JGRP-1873 can almost never happen * The first message creates the recv-window ** If we get some messages before the first, they're dropped on the ground, but as soon as the first message is received (or retransmitted), the recv-window is created. Not too bad, as the first message is usually a JOIN or JOIN-RSP, and only then does message traffic start h5. Use SYNC altogether (replace first message above) * This means that the first N messages would always get dropped until a SYNC and subsequent SYNC-OK have been received ** This would be completely new logic, replacing a proven algorithm h5. Provide connection establishment a la TCP (SYN / SYN-ACK - ACK) * Complex state model (including timeouts and dreaded states such as TIME-WAIT) * Redundant if run over TCP as transport * Blocking semantics: threads need to block until connection has been established	1	1	https://issues.jboss.org/browse/JGRP-1875
14491	JGRP-1875	Or should we make {{conn-id}} a long and use {{nanoTime()}} to set it ? This would guarantee that conn-ids can compare conn-ids and discard messages sent with previous conn-ids. Also, we wouldn't have to add another long, but turn the short conn-id into a long...	1	1	https://issues.jboss.org/browse/JGRP-1875
14492	JGRP-1860	I don't understand why a custom Marshaller impl for handling requests and responses wouldn't fix this. Can you elaborate ? BTW: there is no {{objectFromByteBuffer()}} method in 3.2.13, only {{objectToBuffer()}}. Are we looking at different versions (I'm looking at tag JGroups-3.2.13.Final) ?	1	1	https://issues.jboss.org/browse/JGRP-1860
14493	JGRP-1860	So you want an implementer of a custom marshaller to be able to (re)use the {{Util.objectToFromBuffer()}} methods, but still be able to pass in a class loader ? This in order to prevent having to implement something equivalent themselves ? Makes sense, can you confirm this ?	1	1	https://issues.jboss.org/browse/JGRP-1860
14494	JGRP-1860	> BTW: there is no objectFromByteBuffer() method in 3.2.13, only objectToBuffer(). Are we looking at different versions > (I'm looking at tag JGroups-3.2.13.Final) ? Sorry about that, I've been looking at it in different versions. The method was objectFromByteBuffer in earlier versions, but was renamed to objectFromBuffer. (and the same for objectTo*Buffer). > So you want an implementer of a custom marshaller to be able to (re)use the {{Util.objectToFromBuffer()}} methods, but still be > able to pass in a class loader ? This in order to prevent having to implement something equivalent themselves ? That's correct. Since the classloader is currently hard-coded, the entire marshalling code has to be reproduced in a custom marshaller. With an additional request to be able to set a ClassLoader on RpcDispatcher and have it just handle it, to also skip the step of creating a custom marshaller. (In a non-flat classloader environment, setting the classloader would be a common need for users of RpcDispatcher, so it should be as simple as possible).	1	1	https://issues.jboss.org/browse/JGRP-1860
14495	JGRP-1860	Are you sure that setting a classloader in RpcDispatcher to be used to deserialize requests is a good idea ? AFAICS, ObjectInputStream.resolveClass(), when passed in a classloader, does not attempt to walk a CL chain when a user-defined class is not found. So I'm not sure that setting a single classloader to be used for all requests in {{RpcDispatcher.handle()}} is a good idea... I think adding the classloader ar parameter to the objectFromXXX() methods is good enough. BTW: setting a classloader in RpcDispatcher would also require setting one in RequestCorrelator, too... Thoughts ?	0	0	https://issues.jboss.org/browse/JGRP-1860
14496	JGRP-1856	A,B,C,D are all added to {{suspected_mbrs}} and the coordinator will remove them regardless of their timestamps. So I'm not sure I understand what you mean. Can you decribe this in more detail ?	1	1	https://issues.jboss.org/browse/JGRP-1856
14497	JGRP-1856	For example, we have 6 members A, B, C, D, E and F. The FD_ALL is configured with interval=20 and timeout=60. A, B, C, D are dead at same time. On the FD_ALL timeout scan, the time from since last heartbeat looks like: A (61), B (59), C (56), D(45), E(12), F(15). A is timed out, we suspect it and view change. During view change, the FD_ALL initializes timeout on remaining nodes: B (0), C (0), D(0), E(0), F(0). B, C and D are live another 60 sec. Does it make sense?	1	1	https://issues.jboss.org/browse/JGRP-1856
14498	JGRP-1851	The 'block' option may need to be defined per call, so perhaps a {{NON_BLOCKING_RSVP}} flag needs to be added ? But then we may also want to define timeout, resend interval etc, which leads to a permutation of flags... Should a header be (ab)used to define 'options' to an RSVP call ? Hmm, not nice...	0	0	https://issues.jboss.org/browse/JGRP-1851
14499	JGRP-1844	Why can't you create multiple fork channels ? Each fork channel is independent from the others, while they have the same view and address. The reason for sharing the views is that I made a big mistake a couple of years ago by introducing "service views" which complicated things and were never quite correct compared to JGroups views (google for jgroups "service view"). Can you clarify your use case above with an example ? I'm pretty dead set on removing the shared transport, unless there's a convincing use case / requirement against it...	1	1	https://issues.jboss.org/browse/JGRP-1844

#	Requirement	Questions	Answered	Successfully Answered	URL
14500	JGRP-1844	The situation I am in is that two applications share some Infinispan caches, but also have some of their own independent caches. The short explanation is that I have to create three separate cache managers (one shared, one for each application), each of which requires its own JGroups channel. Infinispan assumes that every node in the view is participating in the cache manager, which isn't true for the app-specific cache managers when using a fork channel. That is arguably more of an Infinispan problem, and I'm going to file some issues in the Infinispan JIRA for those, but I don't know if or when those issues will ever be resolved. Shared transport is a useful workaround in the meantime. I also have my own fork channel that is used by only one of the two applications, and I need to be able to determine a coordinator within that subset of members. JGroups currently doesn't do anything to help with that, so instead of forking the channel that is shared by both applications, I'm forking the app-specific channel. It would be nice if JGroups did something for applications using fork channels, so the application can easily determine which cluster members are participating in the fork. Alternatively, is it possible to put the FORK protocol lower in the stack, below the GMS protocol? Then forks would have their own separate views. I haven't tried that but I doubt it would work based on what I've read.	1	1	https://issues.jboss.org/browse/JGRP-1844
14501	JGRP-1832	The current test focuses mainly on identifying mcast issues which is a great diagnostics start , and we use the largestate test to identify issues such as dropped packets etc.. But the largestate test uses a much bigger stack, is more complex to configure and run, and it uses unicast. 1) It would be great if the mcast sender had a feature , that will fire a stream of sequentially numbered datagrams(for example 1,2,3,...100) to the mcast address , and the receiver will identify and log any gaps in sequence. This will also allow us to run multiple receivers, one on each node, with one sender and complete testing of multiple nodes in one shot , rather then running the test multiple times on "node couples" 2) Can the test detect inefficient mtu fragmentation and defragmantation along a route? 3) In order to track packet loss issues , we sometimes use jgroups at debug or trace log level. But, enabling jgroups debug/trace logging in a production environment , is tricky and can slow down the server, especially if debug level is enabled on multiple servers. It would be great if we can run a receiver instance(lets call it listener) that will listen to the mcast cluster traffic and log relevant data(but obviously will not interfere) This allows us to trace mcast in a separate process or even a different machine without any impact to the prod server. The important aspect here is the ability to correlate the receiver log lines to the log lines on servers for which we did enable jgroups debug. 4) The information that the listener is recording could be further improved , and provide great help , I have quiet a few suggestions on that front but i don't want to make this comment too long one example would be that the listener log can be further precessed by an analysis tool to detect split-brain , or cluster collisions. let me know what you think? more suggestions to come ;)	1	1	https://issues.jboss.org/browse/JGRP-1832
14502	JGRP-1830	Also handle the case where a crash leads to a gap, e.g. A gets 3, B gets 4 and C gets 5, but B crashes before sending message #4. Everyone receives 3(2) -> 5(4), but 4 is missing. SOLUTION: * The coord needs to keep track of seqnos dished out to members (e.g. 4 to B) and send a message to skip the missing ones on a crash of B. * This state needs to be kept until everyone has received B#4. How do we know that? * What happens with this state when the coord itself crashes? ** The reconciliation phase described above needs to be able to reconstruct this state	0	0	https://issues.jboss.org/browse/JGRP-1830
14503	JGRP-1830	Another aspect: uniform delivery, e.g. * C sends m5 (global seqno=5) * A and B receive and deliver m5 * C crashes * D and E haven't received m5 (e.g. because the packet was dropped by the network) Is there a way to provide _uniformity_ for m5?	0	0	https://issues.jboss.org/browse/JGRP-1830
14504	JGRP-1830	What's the initial seqno to be delivered when a new member starts up? The highest seqno currently in use needs to be somehow transferred to a new member... Note that {{SEQUENCER2}} is broken exactly because of this.	0	0	https://issues.jboss.org/browse/JGRP-1830
14505	JGRP-1828	Hmm, how do we do the actual parsing ? XMLConfigurator does the parsing *for the entire XML config*, currently. We'd have to provide hooks for a protocol do do part of this parsing on its own. E.g. after all attributes of a given protocol have been parsed, we look for children nodes. If there are any, we delegate parsing of this subtree to the protocol. The FORK protocol already parses from an input stream, and we could simply add a second parse() method which accepts the DOM subtree (node). The default {{Protocol.parse(Element)}} method would be a no-op.	0	0	https://issues.jboss.org/browse/JGRP-1828
14506	JGRP-1825	What exactly do you want to achieve? Can't you implement the locking/isolation semantics with cache-local locks and using JGroups only as message transport? I'm not entirely happy with the current {{LockService}}, as it doesn't handle partitions well. The {{j.u.c.l.Lock}} interface isn't amenable to takeing locks away from members, so I'm thinking about designing a new partitionable lock service. Not at the top of my todo list though... Since I'm not happy with the current lock impl, I don't want to add something atm.	1	0	https://issues.jboss.org/browse/JGRP-1825
14507	JGRP-1825	Do you have any details on the design ? How and what for you'd use r/w locks ?	1	0	https://issues.jboss.org/browse/JGRP-1825
14508	JGRP-1790	On 02/04/2014 10:50 AM, Bela Ban wrote: > Hi Richard, > > On 04/02/14 15:23, Richard Achmatowicz wrote: >> Hi Bela >> >> I have been thinking about how we might exclude those tests which are not relevant to EAP testing from the JGroups testsuite. >> >> The testsuite has a lot of structure which is not easily customised: >> - runtest gets passed a lot of parameters which determines the suites to run and the runtime parameters to use >> - the report processing depends on runtest output being in a certain format and in certain directories >> - not easy to add in lists of new tests which should be excluded for each runtest target >> >> Because you don't want to add in EAP-related stuff into the JGroups community release, the only clean approach I can see is to do this in the Jenkins job which runs the testsuite: >> - checkout and run the entire JGroups 3.2.13 testsuite >> - after the testsuite has completed and before the reports are generated, delete all test report directories relating to non-EAP tests >> - generate the EAP-only reports and let Jenkins display them > > > An alternative (#1) could be to compile everything, but before running the individual tests (functional, udp, tcp etc), a task would remove all files matching files on the exclusion list from the ./classes directory. > > The advantage over removing xml files is that we don't need to run all of the tests. > >> The "delete non-EAP test results" step can be packaged up as an ant file which deletes the specified directories and can be used on all hosts which run the job. >> >> WDYT? Otherwise, we are going to have to customise the JGroups testsuite in some way, after the branch is checked out and before it is executed. > > > Hmm, not really nice. Also, I don't really like my alternative #1. Perhaps adding an "unsupported" element into the groups annotation attr is not so bad after all, e.g. > > @Test(groups={Global.FUNCTIONAL,Global.UNSUPPORTED}) public void testFoo() {} > > I'm thinking we could create an eap-build.xml which uses <import> to import the regular build.xml and overrides the test targets. Or perhaps this could be done inside of the existing build.xml, in a more elegant way. > > OK, while writing this email, I realized that adding an "unsupported" element to the groups tag is not so bad, so - unless you object - this is the way we're going to do this. > > Do we agree on "unsupported", or would you like it to be named differently ? >	1	1	https://issues.jboss.org/browse/JGRP-1790

#	Requirement	Questions	Answered	Successfully Answered	URL
14509	JGRP-1790	For this to work: - you'll need to add the group "unsupported" to an excludeGroups parameter on every instance of runtest (or add it in by default to the macro), so that each instance of runtest excludes tests marked as unsupported - i'll need to go in an edit every test case file to be excluded to add the "unsupported" annotation parameter. Is that the way you envison things working? It seems fine conceptually, but in practical terms, having to go in an edit a test case to exclude it seems a bit much. It would have been nicer to be able to just specify a list of test case names to be excluded using a system property, somehow ...But testNG doesn't really make life easy. If you would like to proceed that way, i'm fine with it. I'll just have to modify my test report modifying script to modify test cases instead.	1	1	https://issues.jboss.org/browse/JGRP-1790
14510	JGRP-1790	Thanks, Bela. I had written an ant task to update the test cases with the new group, but this is better! I think we also need to mark these tests as eap-excluded: /junit-functional/org/jgroups/tests/FCTest /junit/org/jgroups/tests/LargeStateTransferTest /junit/org/jgroups/tests/StateTransferTest /junit/org/jgroups/tests/StateTransferTest2 /byteman/org/jgroups/byteman/ForwardToCoordFailoverTest Are you going to push these changes to Branch_JGroups_3_2 as well?	1	1	https://issues.jboss.org/browse/JGRP-1790
14511	JGRP-1781	For now, I'll write the length field even for single messages, to have the same code for parsing of single messages and message batches. With NIO.2 in 4.0, this will change anyway	0	0	https://issues.jboss.org/browse/JGRP-1781
14512	JGRP-1773	[~pruivo] was also interested in this, and asked why we can't be clever in UDP's org.jgroups.protocols.UDP.PacketReceiver.run() Wouldn't it be nice to allocate only the receive buffer size of the minimal size possible? I'd say it should not use a size larger than the IPv4/6 required size at least, but maybe we can even look into the header to understand how long the buffer for the copy needs to be?	1	1	https://issues.jboss.org/browse/JGRP-1773
14513	JGRP-1773	hmmm... I don't know why but when I saw the code last week, it looked like this array (https://github.com/belaban/JGroups/blob/master/src/org/jgroups/protocols/UDP.java#L673) was created every time you read a datagram packet. I was wrong :)	0	0	https://issues.jboss.org/browse/JGRP-1773
14514	JGRP-1765	Hi Bela, Could you apply the enhancement without deprecating and ignoring this prop ? TP.looback=false in our configuration.	1	1	https://issues.jboss.org/browse/JGRP-1765
14515	JGRP-1765	And why would that be a problem ?	1	1	https://issues.jboss.org/browse/JGRP-1765
14516	JGRP-1762	Comments: Andre Dietlisheim: {quote} From what I have in mind when packaging libraries to bundles (long time ago, sorry) 1.) will cause havoc to most classpath partitioning frameworks (osgi, jboss modules etc). The code snippet in JGRP-1762 is very likely to fail if bundled. The current thread, your util class and the class you're trying to load may reside in different modules/bundles and visibility rules apply. It is thus is very likely that current thread classloader wont be able to load your class. You should start with 2), the classloader that is accessible via the class you try to load (clazz.getClassLoader() in your snippet). There are most likely further hints that I have forgotten from those times that others here may come up. Jason Green wrote a nice blog that is describing best practices for module-aware classloading: https://community.jboss.org/wiki/ModuleCompatibleClassloadingGuide {quote} Jason Greene: {quote} What I can say is the TCCL lookups only make sense if A) you have a separate API module that discovers a pluggable implementation that the user provides and is discovered lazily during a call from the user's application (e.g. JAXP DOM and STAX) And B) You have a reliable convention that says the TCCL will point to something reliable (like a Java EE container which sets it to the deployment) {quote} James Livingston: {quote} The correct answer is usually "it depends on what is loading the class and why". Is it some shared framework code (JGroups here) which is loading an EE application class passed to it by name? Is it loading a class from a name passed over the network to deserialise an object that's about to follow? Something else? The problem with things like your Util.loadClass() is that you don't know how it's being used so anything you do will be wrong in some situations. Often the best thing to do is to have the API take the classloader to use as a parameter, and push the decision upwards to the caller who either knows the search order or can push the decision upwards to their caller. {quote}	0	0	https://issues.jboss.org/browse/JGRP-1762
14517	JGRP-1750	Would you want to set an {{AddressGenerator}} both for (1) nodes in the local cluster and the bridge cluster, or (2) only for nodes of the bridge cluster ? The latter case is simple: we could allow to set an {{AddressGenerator}} in Relayer, and use it to generate an address which would then we wrapped with a SiteUUID. The former case is more generic: in essence, this would allow for multiple {{AddressGenerator}}s to build an address. So is your use case (1) or (2) ?	0	0	https://issues.jboss.org/browse/JGRP-1750
14518	JGRP-1750	[~ammous] Not sure I understand the need for an AddressGenerator to be set in a bridge. Meanwhile, I created an {{ExtendedAddress}} and {{SiteUUID}} and {{SiteMaster}} are both subclasses of it, so you can add any data to them. I modified UnicastTestRpc slightly and used a config containing RELAY2, and data added to a SiteUUID is actually seen at the receiving end in a different site. Do you have a small test case that shows your problem ? (This is on branch JGRP-1750)	0	0	https://issues.jboss.org/browse/JGRP-1750
14519	JGRP-1743	This is problematic: what is huge and when should this check be done? * Is it the number of messages in the NAKACK table ? * Is it the entire size occupied by the NAKACK table ? In the latter case, I'd have to iterate through all messages and get their length. This is already provided via {{NAKACK2.getSizeOfAllMessages()}}. * When should this be done ? Periodically ? Note that the above method is exposed as a JMX operation, so it could be queried from a script, e.g. {{probe.sh op=NAKACK2.getSizeOfAllMessages}}	1	1	https://issues.jboss.org/browse/JGRP-1743

#	Requirement	Questions	Answered	Successfully Answered	URL
14520	JGRP-1742	h5. Update h6. Serialization of state requests * State transfer requests should be handled *sequentially* at the state provider ** This ensures that a BARRIER that is closed by a new state transfer is not accidentally openend by a completed state transfer h6. State provider * Close BARRIER (see below) only for fetching of digest * When blocked, punch a whole (PUNCH_HOLE) for *unicasts from the state requester* and pass messages tagged with INTERNAL | OOB ** PUNCH_HOLE is needed e.g. for UNICAST3 ACKs or retransmissions * When digest has been fetched, open BARRIER again ** (Leave hole for state requester open, as a different state request might block BARRIER again, before state request has completed) * When state transfer to the state requester has completed, remove hole for state requester (CLOSE_HOLE) h6. State requester * Close BARRIER to set digest *and apply state* * Punch hole for *state provider* ** This is not critical as the state requester can block for some time, before becoming operational ({{JChannel.connect()}} returns) * When the state transfer has completed, remove the hole for the state provider again and open BARRIER h6. Closing of BARRIER / flushing of pending requests * CLOSE_BARRIER can return a boolean (or throw an exception?) if it wasn't able to flush all pending threads ** If false, the state transfer fails and an exception is sent to the state requester * Flushing of requests might fail: ** Requests might be executed as newly spawned threads. This means we'll return immediately, but we're not guaranteed that the state has been modified *** Possible solution: use {{Receiver.block()/unblock()}}: when a {{block()}} is received, the application has to make sure that all pending requests that modify application state have completed h6. Misc * Should we discard traffic when blocked or queue it ? ** Discarding means we'll retransmit, but we might have to kick off a stability round to handle the last-message-lost problem ** Queueing requires keeping track of how many bytes have been added to the queue and ensuring the boundedness of it. We might also overwhelm the thread pools by passing a lot of messages and message batches to them on opening BARRIER	0	0	https://issues.jboss.org/browse/JGRP-1742
14521	JGRP-1742	There are all sorts of problems when a BARRIER is closed and discards messages: * JGRP-1751: a view received during state transfer is not installed * The coordinator doesn't deliver its own view if BARRIER is closed. Only when stability kicks in, will the coord receive its view. This is bad as new views won't get created correctly * When the BARRIER is closed, flushing of pending threads might take a long time So let's think about what the purpose of BARRIER is: # Provide a digest which corresponds to messages *applied to the state* # It doesn *not* need to be active during state transfer (BARRIER can be opened after the digest has been fetched) Re 1: if a message P:25 was added to the digest, but not yet applied to the application, it should be excluded from the digest so we can retransmit it and apply it to the state later. So perhaps a better idea is to *remove BARRIER altogether* and instead use an _application digest_. This would be the digest of the highest message seqnos *applied* to the state. There would be an aditional GET_APP_DIGEST event (implemented by NAKACK and NAKACK2) which returns a digest of seqnos applied to the state. Contrary to the normal digest, this digest is updated when a message has been *delivered* to the application ({{up(Event)}} or {{up(MessageBatch)}} returned) rather than when the message is *received*. Example: * We have digest {{A:10,B:2,C:16}} * Now a message batch {{B:3-10}} is received and a single message {{C:17}} is also received * The normal digest in NAKACK(2) is now {{A:10,B:10,C:17}} * However, the delivery of the message batch hasn't completed yet. The {{up(message)}} method for {{C:17}} has completed though * The _application digest_ returned to the state requester is now {{A:10,B2,C:17}}: ** B's highest seqno is 2 because the message batch {{B:3-10}} hasn't returned yet, so we don't know which messages in it have been delivered to the application (and potentially caused state changes). ** {{C:17}} is part of the application digest as we know that it returned and therefore any potential state change was performed h5. Observations * It is not important to know whether a message really changed the state or not. However, it is important to know whether the message was delivered to the application or not. If not sure (e.g. messages {{B3-7}} of the above message batch might have gotten delivered to the application, but not messages {{B8-10}}), we return the highest seqno that we know was delivered to the application. In the above case, messages {{B:3-10}} would get retransmitted and applied to the state (even if 3-7 were already applied) * As a consequence of the above, *messages need to be idempotent*, or else the state would not be correct. In other words, the state should not be different regardless of whether messages {{B:3-7}} get delivered once or multiple times. * If the above is not possible, then {{FLUSH}} should be used to prevent any sending during a state transfer	0	0	https://issues.jboss.org/browse/JGRP-1742
14522	JGRP-1742	OK, so the following things might solve this puzzle: h5. Coordinator blocked during fetching of digest * Since BARRIER is closed during a state transfer, the coordinator will not only drop messages from other members (except from members for which holes were punched into BARRIER), but also *from itself* * Currently, the only message that causes problems when dropped is a VIEW change multicast * SOLUTION: when multicasting a view V, *have the coord install V locally before multicasting it*. When receiving V, it will get dropped as it is already installed h5. BARRIER skips threads in BLOCKED or WAITING state * Don't skip these, as a blocked thread might simply block on a lock, before changing state (e.g. Infinispan) * If message P:10 was blocked, and we skipped it when fetching the digest, we'd include P:10 in the digest, but not in the state. This would mean that the state requester will never get P:10, neither as part of the state, nor as a retransmission from P h5. Flushing of threads in BARRIER should time out * We cannot wait forever for the threads to time out * The timeout passed to {{getState(timeout)}} should be used to bound the max duration for flushing the threads. If we run into a timeout, either at the state requester or the provider, state transfer (the {{getState()}}) call will fail * A timeout of 0 means wait forever * Closing the channel should terminate the flush h5. Things not tackled * While flushing of threads might succeed in BARRIER, if the application has its own thread pool (e.g. using the _Asynchronous Invocation API_) to handle requests, then flushing will return quickly * However, this is not a guarantee that all incoming threads have completed their changes to the application state * A possible solution might be to call the {{block()}} and {{unblock()}} callbacks in the application. The former would have to wait until all current threads are done modifying the application state. The latter would be called when the digest has been fetched and the application pool can continue making modifications. ** Not very nice, but state transfer should not be used for very large states (taking a long time) anyway ** This will not be addressed by this JIRA. Perhaps it will be tackled in a later release.	0	0	https://issues.jboss.org/browse/JGRP-1742
14523	JGRP-1711	As the implementation of incoming messages delay would be more complicated, I think that implementing outcoming delay could be sufficient. Anyway, I can't distinguish where the message got delayed on network. What was the reason to provide both delays in DELAY?	1	1	https://issues.jboss.org/browse/JGRP-1711
14524	JGRP-1706	Perhaps publish to jcenter rather than nexus? http://ant.apache.org/ivy/history/trunk/resolver/bintray.html	0	0	https://issues.jboss.org/browse/JGRP-1706
14525	JGRP-1672	Hi Bela, couple of comments: 1. Very exciting that you are doing this! 2. The knowledge of 'affinity' in address will be very important to us, we consider the current ISPN 5.x KeyAffinityService API slightly delinquent in this regard (see https://issues.jboss.org/browse/ISPN-3112) 3. By using Java memory mapped files to Linux kernel /dev/shm data do you think you might be able to realize a (C++ like) "call by reference" capability from the Java view? In other words, will you be able to do this in a way that no /dev/shm data ever needs to be copied from the kernel address space into the Java VM's user address space? Thx, Ben	1	0	https://issues.jboss.org/browse/JGRP-1672

#	Requirement	Questions	Answered	Successfully Answered	URL
14526	JGRP-1669	Is it possible to make the similar fix in 2.x release (e.g. make a 2.12.3 with that fix)?	1	0	https://issues.jboss.org/browse/JGRP-1669
14527	JGRP-1669	Hi Aleksandr, sorry, but 2.12.x is not officially maintained anymore. Can't you make that change in your own private repo and build a 2.12.x release yourself ?	0	0	https://issues.jboss.org/browse/JGRP-1669
14528	JGRP-1649	Added max_site_masters (default=1). If > 1, the first N members become site masters. When routing to a remote site, a random member with rank 1..N is picked.	0	0	https://issues.jboss.org/browse/JGRP-1649
14529	JGRP-1625	Thanks for creating this test. The rational for the test is to be able to spot regressions in RPC performance between JGroups release versions. Significant changes to protocol layers can be identified by changed protocol names (MERGE -> MERGE2, for example), so in trying to isolate the effect of changes to JGroups, we can revert back to the original protocol stack layers. But not so for changes to RPC components. Given the dependence of the AS clustering code on SYNC and ASYNC RPCs, comparative tests on the RPC level will increase confidence in version changes.	0	0	https://issues.jboss.org/browse/JGRP-1625
14530	JGRP-1622	Manik mentioned there was a command line tool but it's no longer in the distro... Can it be added back?	1	1	https://issues.jboss.org/browse/JGRP-1622
14531	JGRP-1622	It already exists: Chat. Can you resolve this JIRA ?	1	0	https://issues.jboss.org/browse/JGRP-1622
14532	JGRP-1614	Which protocol ? NAKACK, NAKACK2, UNICAST, UNICAST2, UNICAST3 ? A message needs to be buffered by the sender until it has been received by the peer (unicast) or by all cluster nodes (multicast). You should discuss this on the mailing list before creating a JIRA. Also provide more information if you create a JIRA.	0	0	https://issues.jboss.org/browse/JGRP-1614
14533	JGRP-1613	OK, based on a conversation with Sanne at Red Hat Summit, we decided to narrow the scope. One of the main reasons not to implement a cactus stack like architecture like in the picture above was that - as grafts could be dynamically created (as a matter of fact, Sanne would need a graft to be created *per deployed application*) and deleted - not all nodes would have all grafts. This means that GMS would have to maintain something like the dreaded service views (google for Multiplexer service views, e.g. [1]). This means that if we have view \{A,B,C,D\}, but only nodes B and C have a certain graft created, then the *service view* for that graft would be \{B,C\}. Another reason was that a number of protocols (discovery, failure detection, merge) would have to be rewritten to become singletons. It turned out that it's unlikely this could be genericized, e.g. FD_ALL maintains mappings between logical addresses (which are per channel) and physical addresses. It is therefore tied to a channel and would have to be rewritten to only deal with physical addresses. I'll keep this task for a later day... So the new, limited, functionality should be: * FORK can only be the top protocol in a stack (or towards the top of the stack) * An app gets a call createChannel() which takes a string that has to be unique. We create a fork-channel, which is a subclass of JChannel and has a ref to the main channel * The createChannel() method initially carries a list of instantiated protocols. Later we might also accept XML snippets * The close() or disconnect() method on the fork-channel does *not* close or disconnect the main channel, only the fork-channel * In other words, a fork-channel is a very light weight channel, and we might create hundreds of them without a big penalty * That string is used to mux/demux messages to the app * A header is added to each message including that max-id so we know how to mux/demux * There is a *main channel* and the dynamically created channels refer to it, e.g. their lifetime is less than or equal to the main channel * When the main channel is closed, sending of messages on the fork-channel throws an exception * The view and address of the fork-channels is the same as that of the main channel [1] https://community.jboss.org/wiki/MigrationFromMultiplexerToSharedTransport	1	1	https://issues.jboss.org/browse/JGRP-1613
14534	JGRP-1613	Hmm... Are your components going to be created on *all* cluster nodes ? COUNTER and CENTRAL_LOCK require a coordinator to maintain cluster wide unique counters / locks. However, if we have view \{A,B,C,D\} where A's the coordinator, but the components is only created on B and D, then COUNTER / CENTRAL_LOCK won't work as they send their requests to A, which is the same as to /dev/null as the component isn't running on A...	1	1	https://issues.jboss.org/browse/JGRP-1613
14535	JGRP-1596	Can you elaborate ? The properties map contains property names and associated values ? It would be best to show an example of how such a constructor would be used.	1	1	https://issues.jboss.org/browse/JGRP-1596
14536	JGRP-1585	Hi David, here are my requirements: # Add vararg log methods to Log [1], e.g. trace(String format, Object... params); # Replace log4j, possible candidates so far: ## log4j2 ## jboss-logging # I currently support JUL and log4j as backends, and these need to be supported. If more backends can be supported, even better # Minimize the number of JARs to be used # No pre-processing phase (mvn or ant) # Ability to change the log level at runtime. This is an issue for me with log4j2, which doesn't support this anymore. It's not enough to be able to periodically reparse the config file, as I want to do this across a cluster, via probe.sh op=UDP.setLevel["trace"] ## Derived from above: extensibility. I might want to extend a logging impl to provide this feature. Or users might ant to provide their own logging frameworks. If you tell me that I need only 1 JAR (plus the backend, but that's up to the deployer), than that sounds appealing. I looked at jboss-logging on GitHub, but am unsure which project I should use there. Do you have documentation on jboss-logging ? I googled, but I think I hit old documentation... [1] https://github.com/belaban/JGroups/tree/master/src/org/jgroups/logging [2] https://github.com/jboss-logging	1	1	https://issues.jboss.org/browse/JGRP-1585
14537	JGRP-1585	jboss-logging: * meta-logging framework, does the same as what JGroups org.jgroups.logging already does * Easy to use * Provides logging adapters for JBoss, log4j, JUL and slf4j * Does *not* provide adapter for log4j2 or logback I think I'm going to go with log4j2 for now, with the following changes: * Hack support for setting the level at runtime * Use only the CORE JAR and copy the relevant part of the API JAR (creation of loggers) into JGroups's Log4j2LogImpl	0	0	https://issues.jboss.org/browse/JGRP-1585
14538	JGRP-1576	Hi Mattias. Do you perhaps have a test case that reproduces this problem?	0	0	https://issues.jboss.org/browse/JGRP-1576
14539	JGRP-1576	Any progress on this ? I could simply do what I did in master, namely remove all ThreadGroups (not recommended anyway). WDYT ?	1	1	https://issues.jboss.org/browse/JGRP-1576
14540	JGRP-1564	When sending messages batches, the transport header (TpHeader) is now removed and the cluster name is instead sent as part of the message batch header. This greatly reduces the size of the marshalled message batch ! So if we have a batch of 60 messages and they are sent in a cluster called "my-cluster", then we *add* the length of the cluster name (+12) to a batch, but remove every TpHeader from all 60 messages, that's a saving of (2+2+12) * 60 = 840 -12 = 828 bytes for the message batch ! This effectively fixes https://issues.jboss.org/browse/JGRP-872 !	0	0	https://issues.jboss.org/browse/JGRP-1564
14541	JGRP-1540	This could enable us to use message bundling for *all* messsages, even those tagged with OOB and possibly even DONT_BUNDLE msgs. Since a message marked as DONT_BUNDLE never has to wait for the max bundling timeout to kick in, but instead only for a batch to fill up (should be quick unless the batch size is big !) or no messages remaining in the queue, even sync RPCs might be fast. Need to bench this... e.g. how quickly do 60K fill up when a batch of 1K messages is sent (a few microsecs?)	0	0	https://issues.jboss.org/browse/JGRP-1540

#	Requirement	Questions	Answered	Successfully Answered	URL
14542	JGRP-1539	Sometimes, out-of-order messages are caused by OOB messages: as OOBs are not bundled, a message 20 might be sent as OOB while the bundler has messages 10-19, which will be sent up to bundler_delay ms later. Perhaps look into treating gaps created by OOB message reception differently ? We do know whether a message is OOB or not.	0	0	https://issues.jboss.org/browse/JGRP-1539
14543	JGRP-1535	How do you reproduce this ? In 3.3 (master), I check if a NIC is UP and has one or more addresses associated. How did you create a NIC that doesn't have an adress associated ?	1	1	https://issues.jboss.org/browse/JGRP-1535
14544	JGRP-1535	OK, has this been resolved (config error) ? If not: - Can you show me how to reproduce this on Linux ? - Or can you provide me with a VMWare image that I could download to try this out ? I don't have access to Windows.	1	1	https://issues.jboss.org/browse/JGRP-1535
14545	JGRP-1533	Hi Bela Ban, can you show you email to me? or gtalk? * What's the actual problem Two JGoup Server can't communicate each other when I use UDP mcast_addr=239.190.1.1;mcast_port=32986 * Stack trace how to get the Stack trace? * Logs where the log loaction * How to reproduce etc Server run many days, we found it can't communicate each other	0	0	https://issues.jboss.org/browse/JGRP-1533
14546	JGRP-1528	How would this fit with sync replication calls? I guess they might still block for the batching duration, which I I guess would be much smaller than (max_forward_attempts * forward_sleep milliseconds), i.e. would be in x * 100-millisecond timeout. Also API wise, do you plan to return a future or ISPN-2423 is still relevant even when this will be in place?	1	1	https://issues.jboss.org/browse/JGRP-1528
14547	JGRP-1523	Can you elaborate a bit more ? The suspected_mbrs field is updated also when there is a view change, as all non-members are removed. Other members need to keep suspected members in their suspected_mbrs list, as they may become coordinator themselves. The 'wasteful' doesn't really bother me, as there won't be many suspected members in a list at any given time, and they're cleared after the new view excluding them is installed. I do care about the 'wrong' part of your statement. Can you elaborate ?	1	1	https://issues.jboss.org/browse/JGRP-1523
14548	JGRP-1523	I still don't understand. What do you mean by [old_suspect, new_suspect, me, ...] ? Can you write down a step-by-step description of the scenario that leads to the incorrect state ? You can also ping me on IRC (#jgroups @ irc.freenode.net) if you want to discuss this interactively, or we can have a phone call. (I do speak German if that's of any help) I'd like to release JGroups 3.2.0.Final by the end of October (actually: 1 week before the end of Oct), so if you want this to be fixed and included, as need to act asap. Cheers,	1	1	https://issues.jboss.org/browse/JGRP-1523
14549	JGRP-1523	[1:26pm] bela: ok, simple solution: [1:26pm] bela: if msg_counts_as_heartbeat is true and we do receive a message (or a heartbeat) from a member in suspected_mbrs, we remove it from suspected_mbrs [1:28pm] jboehm: why not recheck the timestamps in suspect()? [1:28pm] jboehm: should be faster [1:29pm] jboehm: suspect is only called if a new member is suspected, not every message, like update() in this case [1:30pm] bela: So how would you remove A from suspected_mbrs then ? [1:32pm] jboehm: in FD_ALL.suspect() instead of eligible_mbrs.removeAll(suspected_mbrs); iterate over suspected_mebrs, get the time stamp, check if it is OK and remove only if not [1:33pm] jboehm: or something like that [1:34pm] bela: aha. And remove members from suspected_mbrs whose timestamps are below the timeout. Hmm, would work I guess	1	1	https://issues.jboss.org/browse/JGRP-1523
14550	JGRP-1495	That further thinking convinced me that the _real_ question is: why is C broadcasting a message, when it knows perfectly well that it is not coordinator? It should have forwarded to B (who is coordinator). The answer to that is that the variable 'flushing' may not be set soon enough. I've updated my pull request to fix that instead.	0	0	https://issues.jboss.org/browse/JGRP-1495
14551	JGRP-1493	Shouldn't this situation correct itself, when D drops the dead members in its view after some time ? At least, if you use FD_ALL, D will exclude B' and C' after some time, and become coordinator of {D,A}, and then the merge should succeed. Re pull request #59: I don't like the solution of spawning a separate thread to send the merge request. I'd rather add a new message flag to bypass fetching the physical address, or make the timeout configurable...	1	1	https://issues.jboss.org/browse/JGRP-1493
14552	JGRP-1493	I have FD and FD_SOCK, but not FD_ALL. D doesn't exclude B' and C', and the situation does not resolve itself. I don't think that physical_addr_max_fetch_attempts is sufficient here: - I agree that this limits the time it takes to look up one member - But there may be more than one dead member. So you might repeatedly fail to find physical addresses when sending MERGE_REQs, so that you still almost surely hit the merge killer. Perhaps the simplest answer would be to use FD_ALL. Is there any reason I wouldn't want to include that in my stack (either in addition to or in place of the other FD protocols)?	1	0	https://issues.jboss.org/browse/JGRP-1493
14553	JGRP-1493	First off, I think D should form a cluster of D and A even with FD (it will be faster with FD_ALL): - We have {B', D, A, C'} - With FD, D pings A and A pings C' - After not receiving acks from C' for a number of times, A excludes C' from its ping set and starts sending out SUSPECT(C') messages - A now starts pinging B' - After some time, A excludes B' from its ping set and starts pinging D. It now starts sending SUSPECT(C', B') messages - D gets the SUSPECT(B') message, takes over as coordinator, creates view {D,A} and installs it So it'll take approximately 2 times (timeout * max_tries) milliseconds to establish the new view. With FD_ALL, this is faster: after timeout (+ (possibly) timeout_check_interval), B' and C' will get excluded. So using FD_ALL instead of FD should only be faster, but both protocols should end up with view {D,A}. Once this happens, a merge should succeed. If the view is not established, I'd rather investigate why this is the case, and not work around this ! WDYT ? I'll see if I can come up with a unit test for the root cause...	1	1	https://issues.jboss.org/browse/JGRP-1493
14554	JGRP-1493	Per my last-but-one - will FD_ALL always help?	0	0	https://issues.jboss.org/browse/JGRP-1493
14555	JGRP-1475	I don't understand what you mean, Jason: I do have unique IDs for logs messages, see [1]. Did you intend to say that I have to have IDs from the ID pool of [2] ? If that's the issue, no problem, I can easily change that. [1] https://github.com/belaban/JGroups/blob/master/conf/jg-messages.properties [2] https://community.jboss.org/wiki/LoggingIds	1	0	https://issues.jboss.org/browse/JGRP-1475
14556	JGRP-1448	I see that the fix version is set to v3.x, should I interpret that as there will not be a 2.12.4?	1	1	https://issues.jboss.org/browse/JGRP-1448
14557	JGRP-1448	Sure I can do that but I was a bit unsure how to go about the correct way. I got the code from a Git repo but the SCM tag in the pom file points to a CVS so was not sure how to create the patch in the "correct" way. Is there any documentation that specifies how JGroups want's the patch file to be created? Sorry for using the comment field for this discussion, didn't know how else to respond.	1	1	https://issues.jboss.org/browse/JGRP-1448
14558	JGRP-1405	Misc comments: - Should we serialize individual messages as they're being added to the bundler ? So serialization would be done on the sending thread, not by the message bundler - Should the message bundler bundle only *multicast* (and not unicast) messages ? This way, we could use a queue instead of a hashmap, as there is only 1 destination (null = send to all). We could then also optimize various headers away, e.g. TpHeader needs to be present only once ! As a matter of fact, we could bundle only the payloads (byte buffers) and send them in one big message with only *one set of headers* !	0	0	https://issues.jboss.org/browse/JGRP-1405

#	Requirement	Questions	Answered	Successfully Answered	URL
14559	JGRP-1396	What was done is the following: - Created Table as a copy of RetransmitTable, but with a lock, and LOW, HD and HR - In NAKACK2: replaced NakreceiverWindow with Table<Message> - In NAKACK2: created only 1 retransmitter task, which periodically scans all tables and triggers retransmission	0	0	https://issues.jboss.org/browse/JGRP-1396
14560	JGRP-1389	Yes, we'll use a separate protocol (RSVP): this *is* separate functionality, not directly tied to NAKACK or UNICAST2. This way, we can use it over UNICAST or even a stack without NAKACK or UNICAST2 ! Design: - RSVP checks if the RSVP flag is set in a message - If so, it creates a unique number (possibly a UUID), and adds it plus all recipients to a hashmap (maybe via AckCollector ?) - Then it adds a header with the UUID and pass the message down - A timer task is then started which periodically sends a LAST_SEQNO event down, telling the NAKACK or UNICAST2 protocols to send their highest seqno (sent or received) to a specific recipient (unicast) or all recipients. - The recipient(s) check the seqno sent by P against their own entry for P and ask P for retransmission if it's missing - Then it blocks on AckCollector until a timeout occurs or all ack have been received - When a timeout occurs, we can throw an exception or simply return (configurable) - When RSVP receives a message with flag=RSVP, it sends an ack back to the sender. This means, we'll only ack messages that have been *delivered*, not when they're *received* - We could first pass the message up, and *then* ack it, or ack it and then send the message. This should be configurable	0	0	https://issues.jboss.org/browse/JGRP-1389
14561	JGRP-1379	We can reduce the number of merge responses using the algorithm below. On a merge request from P: #1 If we're the coordinator --> send a response OR #2 If P is the coordinator of our sub-cluster --> send a response Else suppress the response E.g. if we have sub-clusters {A,B,C,D} and {D,E,F,G}: - A multicasts a discovery request - B, C, and D see that the request's sender is A (the coord), so they reply - D replies, too, because it is the coord of its own subcluster - However, E, F and G don't reply as A is neither their coord, nor are they coords ! This is important if we have large clusters and many subclusters	0	0	https://issues.jboss.org/browse/JGRP-1379
14562	JGRP-1360	Is TP.getBindAddress() and TP.setBindAddress() the only occurrence ? Ichecked other protocols, but didn't find this inconsistency. Fixed.	1	1	https://issues.jboss.org/browse/JGRP-1360
14563	JGRP-1355	Hi William, are you going to be able to resolve this issue in 3.1 (say by Feb 2012) ? If not, I'll push it into 3.2... Cheers,	0	0	https://issues.jboss.org/browse/JGRP-1355
14564	JGRP-1343	We are also using the JBoss Cache 2.1.0 with which the compatible versions of JGroups is 2.6.2. Could you suggest anything else that we could try?	1	1	https://issues.jboss.org/browse/JGRP-1343
14565	JGRP-1317	Sending the digest and associating it (at the receiving end) with a view proved to be problematic in some cases: * If we sent a digest with view-id V1, but the sender already installed V2, the digest would not be associated with a view * In some cases, a view was meaningless, e.g. ** On a merge, each subgroup coordinator (e.g. X in \{X,Y,Z\}) needs to get the digest from all of its subgroup members ** Each member returns its own digest, so Y returns Y=[50:51]. However, Y's view would be \{X,Y,Z\}, so this doesn't work, as Y would return a digest whose seqnos for X and Z would be 0. * On a state transfer, if the membership is \{A,B,C\}, we'd return the state and digest \{A,B,C\}. However, ** If D joined meanwhile, the receiver would have a different view-id and would not be able to associate the received digest with a view. ** Even if the receiver was able to get the view, if B left, the digest for \{A,B,C\} would be associated with view \{A,C\}, which would lead to incorrect seqnos Therefore, digest needs to be changed to include either a view (*reference digest*) or a membership list (*full digest*). The full digest is a bit bigger in memory and when marshalled over the wire, but most digests are unmarshalled, and then some entries in an xmit-table are added/changed or removed, and then the digest is discarded, so this shouldn't be an issue. The only place where we keep a digest around is in the coordinator in STABLE. This shouldn't increase memory consumption. So when would we use full digests and when ref digests ? # JOIN: the existing members would not receive any digest with the new view. The joiner would receive the new view and the *full digest*. # STABLE: members send the *ref digest*; as it is associated with a view-id, the coordinator discards STABLE messages which have a different view-id. # STATE transfer: a state provider always returns the state and the *full digest* # MERGE: the digests for the subgroup are always returned as a *full digest* (as discussed above). However, this only includes *1 member* anyway, so no harm done. The MERGE-REQs always return the *full digests* as well. However, the merge view installation includes the *ref digest* only, as the MergeView and the ref digest are shipped together, and the latter can have a reference to the former. So if we assume that STABLE messages are the most frequent (ref digest), followed by view installation to existing members (ref digest), followed by JOINs (full digest), followed by STATE transfer (full digest) and merge (full and ref digest), then we can see that the cost incurred is not so bad.	0	0	https://issues.jboss.org/browse/JGRP-1317
14566	JGRP-1302	ISPN-658 was pushed into a later release (and may get dropped altogether), so we're pushing this issue into 3.0 as well. If ISPN-658 becomes moot, this issue will be moot, too.	0	0	https://issues.jboss.org/browse/JGRP-1302
14567	JGRP-1290	Hmm, I don't have access to Window 7... does this apply only to Windows 7, or to Windows (XP, 2008) in general ?	1	1	https://issues.jboss.org/browse/JGRP-1290
14568	JGRP-1268	[comment by Jayesh Seshadri] What would be the address visible to other entities (i.e. where would they send their packets to)? Note that currently if you use the external address 'override' for that, jgroups resolves the external_addr property using InetAddress.getByName on the node that the external_addr was assigned to, prior to advertising that address to the peers. A simple fix would be to skip that resolution on the node that the external_addr is assigned and simply advertise the supplied external_addr string as-is to the other nodes, requiring that to be resolvable from the other nodes. So if I set external_addr = 'foobar.corp.com', peers would get the peer address as "foobar.corp.com", and attempt to send packets to "foobar.corp.com", instead of the single IP returned by InetAddress.getByName().	0	0	https://issues.jboss.org/browse/JGRP-1268
14569	JGRP-1268	Hmm, not so sure this makes sense... Even if a user can set the bind address to be 0.0.0.0, at the end of the day they still have to pick an address (e.g. external_addr) that's going to be advertized as the node's address. So we're just pushing the issue of picking an address from bind_addr to external_addr. If we have to pick an address, might as well pick it for bind_addr (as is currently the case). Comments ?	0	0	https://issues.jboss.org/browse/JGRP-1268
14570	JGRP-1262	How about a bind port ? Would you want to pick the bind address only, or also a bind port ?	0	0	https://issues.jboss.org/browse/JGRP-1262
14571	JGRP-1231	I'm not sure what a "SPOF" is. Something like what Mircea made for the JDBC stores in Infinispan would be great: it auto-creates the tables and provides configuration options to define the SQL types. Actually, wouldn't it be great if people could plugin their own implementation in SPI style? So additional dependencies wouldn't be an issue, and people might easily add any exotic alternative strategy.	1	1	https://issues.jboss.org/browse/JGRP-1231
14572	JGRP-1231	to create a complete test I might need to add some dependencies (for testing only), like an in-memory database as H2 or HSQLDB. Would you be ok with that, or would you prefer me to skip merging this kind of tests.	1	1	https://issues.jboss.org/browse/JGRP-1231
14573	JGRP-1231	I will definitely do the testing via Infinispan to "feel" a real scenario, but wouldn't you want an automated test to prevent future regressions? This would be a testing-only dependency, but I'll do as you wish.	0	0	https://issues.jboss.org/browse/JGRP-1231
14574	JGRP-1219	Sounds good, two comments: - Is there a chance that in one cluster we'll have different JBoss versions? Upgrade? If yes, we need to specify always our own prefix. - Maybe prefix with "cluster" at the beginning: "cluster-jboss-6-m4", "cluster-jboss-7-ga" would be nicer :)	1	0	https://issues.jboss.org/browse/JGRP-1219

#	Requirement	Questions	Answered	Successfully Answered	URL
14575	JGRP-1207	{Email Dennis] The fix that works for this customer is bind()'ing the multicast socket on creation to the correct IP, instead of relying on setInterface. Do you think that will cause issues on any other OS's? That's easy enough on 2.4.x. But on 2.6.x and above, I think the separate send/receive sockets would need to be restored, so receive can bind to the multicast IP and send can bind to the bind_addr. Do you agree?	1	0	https://issues.jboss.org/browse/JGRP-1207
14576	JGRP-1207	IMO, IPv6 should use the IP address passed to it in MulticastSocket.setInterface() as *source* address. If it doesn't do that, then I suggest we talk to our IP gurus to see why not ! MulticastSocket.setInterface() sets the interface, but as I understand it, the kernel can pick any address on that interface. I cannot use the MulticastSocket() constructor taking a SocketAddress and pass in the bind address, as we pass in the multicast address to prevent promiscuous traffic ! e.g. new MulticastSocket(new InetSocketAddress("224.5.6.7", 7500)); Check Util.createMulticastSocket() for details. BTW: there's a bunch of JIRAs on this. I suggest you write a simple program which creates a MulticastSocket as described above, and calls setInterface() with an IPv6 address. Then we create a couple of aliases on the same interface as the IP address and observe that the src address is different every time. Someone from Red Hat (preferably Neil Horman) should take a look...	1	0	https://issues.jboss.org/browse/JGRP-1207
14577	JGRP-1207	what do you want me to do then ?	1	0	https://issues.jboss.org/browse/JGRP-1207
14578	JGRP-1202	I guess injecting a SocketFactory instance into a stack is the way to go. However, we have to think about the following case. Say we have a shared transport ("udp") that's used by standard-session-cache and sfsb-cache. All sockets created in sfsb-cache would be created by a SocketFactory instantiated with the name "sfsb-cache". All sockets created in standard-session-cache would be created by a SocketFactory instantiated with the name "standard-session-cache". So, for example, we could have the following config: standard-session-cache: FD_SOCK.bind_port: 7300 sfsb-cache: FD_SOCK.bind_port: 7400 However, since we have a shared transport, what would the naming scheme be for it ? It's used by both standard-session-cache and sfsb-cache, so we cannot give it the name of either one, or should we ? Or should we just call it "udp" ? One solution would be give the SocketFactory in a shared transport the name of its singleton_name attr, e.g. in the above example "udp". When setting the SocketFactory, e.g. in JChannel.setSocketFactory(), we'd create an event SET_SOCKET_FACTORY (with an instance of a SocketFactory) and send it down the stack. When we have a singleton (shared) transport, the TP$ProtocolAdapter would catch the event and install the socket factory, then discard the event (not send it down to the transport). The transport's SocketFactory would have to be set separately. When there's no singleton transport, the SET_SOCKET_FACTORY event would be caught by the transport (TP) and the factory would be installed. A getSocketFactory() method (inherited from Protocol) would return the socket factory of TP$ProtocolAdapter if a singleton transport is used and the call originated from above TP, or the one from TP if the call originated in TP. If no singleton transport is used, the socket factory of TP would always be returned.	0	0	https://issues.jboss.org/browse/JGRP-1202
14579	JGRP-1189	Question is should we hard-code the prefix (e.g. FF0e::) or make it configurable	0	0	https://issues.jboss.org/browse/JGRP-1189
14580	JGRP-1156	I cannot reproduce this, can the customer reproduce this ? I tried bringing the interface down (ifconfig eth1 down), stopping the networok (service network stop), to no avail...	0	0	https://issues.jboss.org/browse/JGRP-1156
14581	JGRP-1154	[Email dated Sept 10 2010 to jg-dev] FYI, I've added 2 new flow control protocols: MFC (Multicast Flow Control) and UFC (Unicast Flow Control). They're an implementation of [1], which is about splitting the FC functionality into 2 separate protocols. The advantages are that this simplifies the code, the code becomes more maintainable, that we can enable unicast or multicast flow control individually, that I could remove some coarse grained locking, and that I created helper classes, e.g. to manage sender or receiver credits, which can be unit tested individually, apart from the use in the protocol. The downside is that we now have MFC and UFC instead of just FC. My goal for performance was that MFC and UFC should be about the same as FC, but not slower. This morning, I ran a quick perf test [2] on a 4 node cluster, with 4 nodes (all senders) sending 1 million messages (of varying size) each. The numbers are MiBs/sec/node. The numbers for 4 nodes are: Size FC MFC/UFC 1K 135 135 2K 145 145 4K 145 145 8K 130 131 The numbers for 8 nodes are (2 processes per physical box): Size FC MFC/UFC 1K 72 73 2K 83 85 4K 87 93 8K 107 113 As can be seen, performance is about the same, or even slightly better with MFC/UFC when running 2 perf JVMs / physical box. However, the second test on 8 nodes is to be taken with a grain of salt, as the 2 JBMs/box are competing for CPU and the network, and therefore we don't saturate the network (bandwidth)... To use the new folw control protocols, replace <FC max_credits="2M" min_threshold="0.6" /> with <UFC max_credits="2M" min_threshold="0.6" /> <MFC max_credits="2M" min_threshold="0.6" /> I will change the configs shipped with JGroups after some more testing and updating the documentation. [1] https://jira.jboss.org/browse/JGRP-1154 [2] http://community.jboss.org/docs/DOC-11594 -- Bela Ban Lead JGroups / Clustering Team JBoss	1	0	https://issues.jboss.org/browse/JGRP-1154
14582	JGRP-1150	We also don't need to marshal 'range' when the request type is MSG ! (Only for XMIT_RSP and XMIT_REQ)	0	0	https://issues.jboss.org/browse/JGRP-1150
14583	JGRP-1121	What's the benefit over the existing solution here ?	1	1	https://issues.jboss.org/browse/JGRP-1121
14584	JGRP-1108	What's the status of this, Brian ? Is it still relevant ? If not, could we remove ChannelFactory from JGroups ?	1	1	https://issues.jboss.org/browse/JGRP-1108
14585	JGRP-1106	Just out of curiosity, why would this be a jgroups change? The last thing we need is for every project to have a complicated logging framework of its own.	0	0	https://issues.jboss.org/browse/JGRP-1106
14586	JGRP-1106	Yes, this makes sense, thx for the contribution Shay ! I applied this change, but changed the system prop to jgroups.logging.log_factory_class. Also, I'm not too sure about whether Class.forName() is good. We usually use Util.loadClass() to make sure to pick the right classloader. This is in CVS, can you take a look ? I'll close this issue, feel free to re-open if you run into an issue.	0	0	https://issues.jboss.org/browse/JGRP-1106
14587	JGRP-1105	OK, here's what we can do: We create new SeqnoRanges every N milliseconds, e.g. all seqnos added in T=0-30ms go into range 1, 31-60ms go into range 2 and so on. When a seqno is received and the current time - last_time is > MAX_TIME, we create a new range and add the seqno to the range. Ditto when the current range's size is exceeded. Algorithm: add(seqno): - if current range has all received and size not exceeded: - add to current range (reuse that range) - start the retransmission task (with a new TIMEOUTS copy) else if not all received and time - last_time > MAX_TIME or size exceeded: - create new SeqnoRange (current range) - set last_time = time - If retransmission task is not running: start retransmission task - add seqno to current range Task run: - retransmit all missing messages - If all received: cancel the task TODO: - when do we remove a range and retransmission task ?	0	0	https://issues.jboss.org/browse/JGRP-1105
14588	JGRP-1104	Vladimir: If oob message is added to NRW it is immediately delivered by its own oob thread propelling that message up, right? So, what is the scenario when #1 occurs ? Bela: I forgot, but this *can* occur. That's why I wrote that code. I think the use case was the OOB message was dropped and then retransmitted (as reg message).	1	1	https://issues.jboss.org/browse/JGRP-1104

#	Requirement	Questions	Answered	Successfully Answered	URL
14589	JGRP-1095	The issue is that MPING wants to send a request back to the requester via the transport (TCP), but a socket to a lonk local address with a scope-id fails. Without the scope-id, it works. TODO: remove the scope-id when creating a socket. Dito for TUNNEL, TCPGOSSIP	0	0	https://issues.jboss.org/browse/JGRP-1095
14590	JGRP-1090	Analyze which protocols above NAKACK send multicast OOB messages (STABLE for sure, maybe others ?)	0	0	https://issues.jboss.org/browse/JGRP-1090
14591	JGRP-1090	For now, the bundlers were changed to store messages per *cluster and destination address*, e.g. messages with dest=null for "cluster1" are stored separately from messages to "cluster2". This way, messages sent on a shared transport to different clusters won't get handled by the same thread at the receiver(s). Also, multicast messages won't interfere with unicast messages. The old solution stored messages under only the destination as key, and didn't take the cluster name into account. So the only messages that will be delivered by the same thread are messages from the same sender to the same destination and cluster. To be investigated: should we neverthless bundle OOB messages as well ? This would entail having a separate hashmap in the bundlers, and handle every OOB message of a message bundle with a separate thread at the receiver(s).	0	0	https://issues.jboss.org/browse/JGRP-1090
14592	JGRP-1084	Is this really needed ? IIRC, Infinispan's CommandAwareDispatcher passes a list of target addresses, but removes the local address before that...	1	1	https://issues.jboss.org/browse/JGRP-1084
14593	JGRP-1078	Bela, I tried this locally and the element in question is present and I don't get -Djgroups.xml.validate=true that you refer to. Using JDK6 on Snow Leopard. What JDK are you using?	1	1	https://issues.jboss.org/browse/JGRP-1078
14594	JGRP-1078	Interesting, my local schema is valid too! What JDK/OS are you using Sei?	0	0	https://issues.jboss.org/browse/JGRP-1078
14595	JGRP-1078	I don't use maven, I use ./build.sh create-schema, which command do you use ?	0	0	https://issues.jboss.org/browse/JGRP-1078
14596	JGRP-1068	treated during @Property processing. I found this example in MPING: public void init() throws Exception { super.init(); String str=Util.getProperty(new String[]{Global.MPING_MCAST_ADDR}, null, "mcast_addr", false, null); if(str != null) mcast_addr=InetAddress.getByName(str); str=Util.getProperty(new String[]{Global.MPING_MCAST_PORT}, null, "mcast_port", false, null); if(str != null) mcast_port=Integer.parseInt(str); str=Util.getProperty(new String[]{Global.MPING_IP_TTL}, null, "ip_ttl", false, null); if(str != null) ip_ttl=Integer.parseInt(str); if(mcast_addr == null) { mcast_addr=InetAddress.getByName("230.5.6.7"); } } This started me thinking about how we process properties in JGroups, especially now with @Property. Does the following concur with your understanding of how properties should be processed? 1. Usera can specify defaults for properties in one of two ways. The usual way is to provide a default for the value in a stack configuration file. Over and above that, they may provide a default on the command line using a named system property. The system property should override the default in the stack configuration file through system property substitution within the stack config file. (NOTE: we could also systematically check system properties during property conversion, with no need to use variable substitution in config files - but forget about this when reading point 2.) 2. Once the stack config file has been processed, each @Property in a protocol gets a value assigned through PropertyConverter processing. At this stage, the users intentions for the variable should have been determined. No further system variable substitution for @Property values should be required. 3. After property processing, @Property variables may have valid values(e.g. user assigned), invalid values(e.g. user assigned IPv6 address, but no scope) or no value (null). We don't really have any systematic way of checking for the case of no (default) value at present. One way to do this is via sanity checks, after property processing and before init() is called on each protocol. This is the intention for the InetAddress related stuff - to call inetAddressSanityCheck() to assign a consistent IP version and check IPv6 scopes before init() starts. The good thing about this is that we have access to all @Property values in the stack and can assign defaults accordingly. This is needed for determining the correct IP version to use. On the other hand, there are a lot of variables in a lot of protocols which may require programmer assigned defaults, and assigning them in this centralized location through table lookup doesn't scale well. It would be more convenient to assign a programmer assigned default value through the @Property mechanism itself. But because @Proeprty processing uses converters to assign values to variables, the assignment of programmer assigned default values would need to be incorporated into the ProeprtyConverter mechanism. For example, in the case of jgroups.udp.mcast_addr, we used to specify a default string value "224.x.y.z" which would get converted in init() to an InetAddress if mcast_addr == null, just like in the code above. But now this happens too late. Is it worth adding an attribute to @Property to specify a programmer supplied default value for the property which would be passed through the converter if no user specified value was present? At the moment, during property processing, we check if a user assigned value is present: if yes, we call the converter; if no, we don't call the converter. This means that some fields, which take complex data structures which need special conversion from strings, like initial_hosts, have null for a default and have to be picked up after property processing has completed. We could incorporate programmer assigned defaults as follows: (i) introduce a new field in @Property called default (ii) during property processing, we check if a user assigned value is present: if yes, we call the converter; if no, we call the converter using the default value, if there is one; if there is no default value, we do nothing. In this way, all properties have either a user assigned or programmer assigned default value. The default values are referenced by Global.jgroups.protocol.variable name which picks up the default from a table during processing. With this in place, we eliminate checking for default values in the sanify	0	0	https://issues.jboss.org/browse/JGRP-1068
14597	JGRP-1050	So what about OSGI support in jgroups? I"m seeing in the binaries jgroups 2.8 and morover and there are not bundle informations inside the MANIFEST.MF file. Testing the sources (2.12.0.Alpha1) I see a maven-bundle-plugin but it is disabled because it need the row <packaging>bundle</packaging> in the pom.xml and actually it is configured as jar.	0	0	https://issues.jboss.org/browse/JGRP-1050

#	Requirement	Questions	Answered	Successfully Answered	URL
14598	JGRP-1043	Another related contention point I see when profiling is for AckReceiverWindow.msgs. Calls the lead to acquiring that lock come from UNICAST.handleDataReceived(). Looking at that I see a couple possible improvements: 1) if(win.smallerThanNextToRemove(seqno)) sendAck(msg.getSrc(), seqno); The win.smallerThanNextToRemove(seqno) call has to acquire the AckReceiverWindow.msgs lock. My interpretation of this is that it's purpose is to send an ack to a message retransmission, where the original message has already been received. Why would this be necessary? A) ACK of original message didn't arrive in time to cancel retranmission B) ACK of original message got lost. If B) occurred, the retransmission will go on forever, so to prevent that we have this code to ack the retransmission. Now, why would win.smallerThanNextToRemove(seqno) return true? a) It's a retransmission, i.e. the case described above. In that case "added" from the earlier "boolean added=win.add(seqno, msg);" will be false. So, checking for !added would handle this. b) Another thread is removing messages from the window and between when the thread adding the message called win.add(seqno, msg) and later called if(win.smallerThanNextToRemove(seqno)), the "removing thread" already removed it. In that case the "removing thread" will ack the message and the adding thread also acking it is redundant. Bottom line, I think the above can be changed to the following, eliminating contention on AckReceiverWindow.msgs: if(!added) { // Must be a retransmission of a msg we already acked. Ack it to ensure retransmission task // is cancelled, in case our ack of original message got dropped sendAck(msg.getSrc(), seqno); } 2) if(!added && !win.hasMessagesToRemove()) { // no ack if we didn't add the msg (e.g. duplicate) return; } I get the feeling this code is leftover cruft, as the "// no ack if we didn't add the msg (e.g. duplicate)" comment doesn't fit the surrounding logic that determines whether there's an ack. The effect of this block is just to cause the thread to return and not drop into the section where it checks if it should become the thread that removes messages. That check is now quite cheap, so why not just remove this? This one is no big deal in the normal case, since added will usually be true so the win.hasMessagesToRemove() call won't be made. But in the overload case where there are starting to be a lot of retransmissions, the win.hasMessagesToRemove() call will start being made, adding to the contention on the lock.	1	1	https://issues.jboss.org/browse/JGRP-1043
14599	JGRP-1043	Further to the issue of contention for AckReceiverWindow.msgs. For acking OOB messages, the relevant logic in handleDataReceived is: if(msg.isFlagSet(Message.OOB)) { if(added) up_prot.up(new Event(Event.MSG, msg)); Message oob_msg=win.removeOOBMessage(); if(oob_msg != null) sendAckForMessage(oob_msg); if(!(win.hasMessagesToRemove() && undelivered_msgs.get() > 0)) return; } ..... logic to possibly become thread that passes messages up, if yes, then... while(true) { Message m=win.remove(processing); if(m == null) { released_processing=true; return; } // discard OOB msg as it has already been delivered (http://jira.jboss.com/jira/browse/JGRP-377) if(m.isFlagSet(Message.OOB)) { continue; } num_regular_msgs_removed++; sendAckForMessage(m); up_prot.up(new Event(Event.MSG, m)); } As we discussed, I believe there is a race where win.removeOOBMessage() can return null, resulting in the OOB message not getting acked, since the logic in the later while loop won't ack an OOB. Wait! Just realized it's not just the race we discussed, where OOB msg 1 gets added and then another thread adds msg 2 and removes 1 before the first thread calls win.removeOOBMessage(). This can happen in the simple case where AckReceiverWindow.next_to_remove is only up to say seqno 5 and an OOB msg 10 comes in. In that case win.removeOOBMessage() will return null and the msg will not get ACKed. (As we discussed, AckSenderWindow.ack() will help clean up after this, by cleaning 10 once it gets an ack for 11). Anyway, back to my original point. :) The win.removeOOBMessage() call is adding contention on AckReceiverWindow.msgs. Just acking OOB messages in the while loop seems cleaner. Is the goal of the win.removeOOBMessage() and if(!(win.hasMessagesToRemove() && undelivered_msgs.get() > 0)) return; to try and provide an escape route for OOB threads so they don't become the thread responsible for passing message up? OK, later comments about JGRP-781 indicate that's the case. In any case the test can be if(!(undelivered_msgs.get() > 0 && win.hasMessagesToRemove())) Do the cheap test before the expensive one. Perhaps the whole thing becomes: final AtomicBoolean processing=win.getProcessing(); if(msg.isFlagSet(Message.OOB)) { if(added) up_prot.up(new Event(Event.MSG, msg)); // JGRP-781. If no thread is processing the window and passing // messages up, we may need to take on that duty even // though we are an OOB thread if (!processing.get()) { // First try and clear our own message from the window Message oob_msg=win.removeOOBMessage(); if(oob_msg != null) sendAckForMessage(oob_msg); // Check and see if there is processing work to do; if not // return and let whatever thread adds work become the // processing thread if(!(undelivered_msgs.get() > 0 && win.hasMessagesToRemove())) return; } else { return; } } I think the undelivered_msgs tracking would be better placed inside AckReceiverWindow. Right now the win.add() and undelivered_msgs.incrementAndGet() calls are not atomic, nor are the remove() and decrement calls. So you can get races plus it makes handleDataReceived more complicated.	1	1	https://issues.jboss.org/browse/JGRP-1043
14600	JGRP-1043	Re: !added: DOH! Good example of why I post comments rather than change code! is this useful? if(!added && win.smallerThanNextToRemove(seqno)) We limit the win.smallerThanNextToRemove to cases where the message isn't the next_to_remove. Hmm, probably not a huge benefit, unless even w/ thread pools in TP it's still the norm that messages get added to the window in seqno order.	1	1	https://issues.jboss.org/browse/JGRP-1043
14601	JGRP-1043	The DelayQueue (based on PriorityQueue) used by ScheduledThreadPoolExecutor seems to be inefficient regarding adding and removing elements. These methods apparently need rebalancing of the tree. ConcurrentSkipListSet promises log(n) cost for add() and remove(), maybe we should take a look at replacing ScheduledThreadPoolExecutor in TimeScheduler with this ! Disadvantage: ConcurrentSkipListSet requires JDK 6 as base line...	0	0	https://issues.jboss.org/browse/JGRP-1043
14602	JGRP-1040	How do you separate tests when they're not sequential ? @Test(sequential=false) ?	1	0	https://issues.jboss.org/browse/JGRP-1040
14603	JGRP-1040	Good question... Doesn't your test listener which collects output from each test handle that? I seem to remember that you dealt with that issue using local thread storage. If not, i'll have to go back and check.	0	0	https://issues.jboss.org/browse/JGRP-1040

#	Requirement	Questions	Answered	Successfully Answered	URL
14604	JGRP-1038	to everybody else took 16 seconds (see below) ! Did you have members send unicasts from everyone to everyone else on purpose, or should those messages have been multicast ? Also, your stack is non-standard: * You don't set multicast and unicast receiver buffer sizes in UDP, this makes a big diff when receiving large messages (as we have to do fewer retransmissions due to dropped messages) * For UDP, you may also want to configure thread pools * Your FC2 is in the wrong place (for FC2_STACK): it is just above NAKACK, but below UNICAST so your unicasts are not flow controlled at all !!! * STABLE should have max_bytes, a gossip time of 5000 is too long for sending a lot of messages * I tested this on 2.8, which version did you use ? I created a JIRA for FC2 to take a look at in 2.9, but unless performance of the existing FC is really bad, I don't think FC2 is necessary. ---------- GMS: address=linux-39675, cluster=x, physical address=192.168.1.5:33472 ---------- . ---------- GMS: address=linux-12313, cluster=x, physical address=192.168.1.5:34644 ---------- . ---------- GMS: address=linux-56774, cluster=x, physical address=192.168.1.5:55607 ---------- . ---------- GMS: address=linux-15964, cluster=x, physical address=192.168.1.5:38462 ---------- . ---------- GMS: address=linux-49489, cluster=x, physical address=192.168.1.5:56906 ---------- . ---------- GMS: address=linux-43361, cluster=x, physical address=192.168.1.5:52679 ---------- . ---------- GMS: address=linux-55342, cluster=x, physical address=192.168.1.5:53446 ---------- . ---------- GMS: address=linux-16696, cluster=x, physical address=192.168.1.5:43925 ---------- . ---------- GMS: address=linux-31728, cluster=x, physical address=192.168.1.5:38416 ---------- . ---------- GMS: address=linux-19539, cluster=x, physical address=192.168.1.5:53866 ---------- . -- sending message #0 of 5242880 bytes to linux-39675 -- received message of 5242880 bytes on linux-39675 Received 1 messages, sent 149 messages and 0 retransmits, Received 375 packets, sent 366 -- sending message #0 of 5242880 bytes to linux-12313 -- received message of 5242880 bytes on linux-12313 Received 2 messages, sent 237 messages and 0 retransmits, Received 506 packets, sent 496 -- sending message #0 of 5242880 bytes to linux-56774 -- received message of 5242880 bytes on linux-56774 Received 3 messages, sent 325 messages and 0 retransmits, Received 633 packets, sent 624 -- sending message #0 of 5242880 bytes to linux-15964 -- received message of 5242880 bytes on linux-15964 Received 4 messages, sent 413 messages and 0 retransmits, Received 765 packets, sent 754 -- sending message #0 of 5242880 bytes to linux-49489 -- sending message #0 of 5242880 bytes to linux-43361 -- received message of 5242880 bytes on linux-49489 Received 5 messages, sent 503 messages and 0 retransmits, Received 899 packets, sent 887 -- received message of 5242880 bytes on linux-43361 Received 6 messages, sent 589 messages and 0 retransmits, Received 1028 packets, sent 1015 -- sending message #0 of 5242880 bytes to linux-55342 -- sending message #0 of 5242880 bytes to linux-16696 -- received message of 5242880 bytes on linux-55342 Received 7 messages, sent 680 messages and 0 retransmits, Received 1165 packets, sent 1151 -- received message of 5242880 bytes on linux-16696 Received 8 messages, sent 765 messages and 0 retransmits, Received 1293 packets, sent 1278 -- sending message #0 of 5242880 bytes to linux-31728 -- received message of 5242880 bytes on linux-31728 Received 9 messages, sent 853 messages and 0 retransmits, Received 1424 packets, sent 1408 -- sending message #0 of 5242880 bytes	1	1	https://issues.jboss.org/browse/JGRP-1038
14605	JGRP-1038	Looks like FC2 handles only *unicasts*, not multicasts. Is this what FC2's supposed to do ? If so, it could be an alternative to UFC... A good perf test would be UnicastTestRpcDist.	0	0	https://issues.jboss.org/browse/JGRP-1038
14606	JGRP-1030	Vladimir, can you look into this ? thx !	1	1	https://issues.jboss.org/browse/JGRP-1030
14607	JGRP-1025	[~belaban] Can we close this one? See given long enough timeline all bugs go away? :-)	1	0	https://issues.jboss.org/browse/JGRP-1025
14608	JGRP-1011	Test results: I list the test which failed repeatedly, followed by the tests which failed, but passed on retry, organized by IP version. ipv4: repeatable failure: * CloseTest(tcp).testMultipleConnectsAndDisconnects (also fails on 2.6.10.GA - no regression) * FlushCloseOpenTest(tcp).testAndLoop (also fails on 2.6.10.GA - no regression) passed on retry: * none ipv4.local: repeatable failure: * CloseTest(tcp).testMultipleConnectsAndDisconnects (also fails on 2.6.10.GA - no regression) * FlushCloseOpenTest(tcp).testAndLoop (also fails on 2.6.10.GA - no regression) passed on retry: * SharedTransportTest(udp).testSendingOfMessageAfterChannelClose * ConnectStressTest (timeout) ipv6: repeatable failure: * CloseTest(tcp).testMultipleConnectsAndDisconnects (not testsed on 2.6.10.GA) * FlushCloseOpenTest(tcp).testAndLoop (not tested on 2.6.10.GA) passed on retry: * MergeTest(udp).testMerging4Members ipv6.local: repeatable failure: * CloseTest(tcp).testMultipleConnectsAndDisconnects (not testsed on 2.6.10.GA) * FlushCloseOpenTest(tcp).testAndLoop (not tested on 2.6.10.GA) * MergeTest(udp).testMerging4Members/testMerge4Members (this test is really broken for IPv6 as it uses IPv4 default mcast addresses and needs to be recrafted) passed on retry: * ConnectStressTest (timeout)	0	0	https://issues.jboss.org/browse/JGRP-1011
14609	JGRP-1009	Merging and join/leaves should be separate and can never overlap. When a merge is in progress: - Join and leave requests should get discarded (or queued ?) When a join or leave is processed: - MERGE events should be discarded - MERGE-REQ requests received from other partitions should be *rejected*	0	0	https://issues.jboss.org/browse/JGRP-1009
14610	JGRP-1007	Is the issue anyway related to JBAS-6807 at https://jira.jboss.org/browse/JBAS-6807. Could this JIRA task provide some inputs?	1	1	https://issues.jboss.org/browse/JGRP-1007
14611	JGRP-988	I have no objection, but I'm questioning in general whether we should trash JChannelFactory. It is IMO pretty useless, new apps inject channels anyway, and that mechanism is external. Can't we move this entire class into JBoss AS ? I'll definitely trash it in 3.0...	1	0	https://issues.jboss.org/browse/JGRP-988
14612	JGRP-983	Hi Galder, a couple of questions: - is there the usual 3-4 day deadline on QAing this? I'm ugly busy at the moment. - is there a test case in here to check for crosstalk? Don't seem to remember if there is one i n the testsuite. If not, do you have a description of a manual procedure for verifying that cross talk does not occur?	0	0	https://issues.jboss.org/browse/JGRP-983
14613	JGRP-983	Re: - is there the usual 3-4 day deadline on QAing this? I'm ugly busy at the moment. Indeed, it's 5 business days AFAIK Re: - is there a test case in here to check for crosstalk? Don't seem to remember if there is one i n the testsuite. If not, do you have a description of a manual procedure for verifying that cross talk does not occur? I don't have a test case as such but the following could be tested on a HP-UX machine (customer did this): - Start two EAP 4.3.0.GA_CP04 server instances passing a different -g and -u to each and see if the clusters start up independently. If there's cross talk, they'd see each others messages.	0	0	https://issues.jboss.org/browse/JGRP-983
14614	JGRP-982	Fixed in 2.4.x, 2.6.x, 2.7.x and trunk. Bela, what are the fix versions for 2.6.x, 2.7.x and trunk?	1	1	https://issues.jboss.org/browse/JGRP-982
14615	JGRP-974	Can you post the link to the fix from Brian? Thanks.	1	0	https://issues.jboss.org/browse/JGRP-974
14616	JGRP-960	Comment from Bob: Is there any race conditions we need to look out for when working with S3? Or do you need to write S3-based lock-manager? With XA and a pony? :) Or a separate thought, treat S3 more directory-like, and instead of updating a file, each node puts (and removes, optionally) its own file within the dir? s3://path/to/my/cluster/node1.conf s3://path/to/my/cluster/node2.conf s3://path/to/my/cluster/node3.conf s3://path/to/my/cluster/node4.conf s3://path/to/my/cluster/node5.conf Just thinking if 100 nodes are booting up at the same time, each trying to write to a shared cluster.conf without trampling each other?	0	0	https://issues.jboss.org/browse/JGRP-960

#	Requirement	Questions	Answered	Successfully Answered	URL
14617	JGRP-951	Bela, One question. The updates I made to allow the 2.8 testsuite to use IPv6 addressses didn't address one issue in build.xml. You call runtest as follows, for udp.xml: <runtest testname.ext="udp" threadcount="10" parallel="tests" initial.mcast.addr="235.10.10.10" initial.mcast.port="25100" initial.tcp.port="25100" testng.xmlfile="testng-udp.xml" /> That is, you are hardcoding IPv4 values for initial.mcast.addr, and you are using different address ranges across the different targets (udp, tcp, udp-flush, ...)for some reason. This makes switching the testsuite from IPv4 to IPv6 rather difficult. I would like to replace these hardcoded values by the system property ${jgroups.udp.mcast_addr} so that the value gets picked up from build.properties. Same for the mcast_port. Is this OK? There shouldn't be any need to use different ranges of mcast addresses when calling runtest as these are run in sequence anyway.	0	0	https://issues.jboss.org/browse/JGRP-951
14618	JGRP-908	What's the diff to SMACK ? No acks ?	1	1	https://issues.jboss.org/browse/JGRP-908
14619	JGRP-908	I don't see why you use seqnos: the code doesn't use them ! Every message contains a seqno, but the receivers don't order them, nor do they prevent duplicate seqnos...	1	1	https://issues.jboss.org/browse/JGRP-908
14620	JGRP-892	Dennis, do you have a pull request ?	0	0	https://issues.jboss.org/browse/JGRP-892
14621	JGRP-886	Bela, as I know NIO2 is already available in JDK 7 builds for several months (or years?). Oracle delayed the release of JDK 7 due to some internal reasons, but it seems to be good enough, I tried it with our production system for several days. Are you going to hard-code NIO2 usage inside JGroups or add some abstraction layer (aka TransportProvider abstraction) so we could use such network frameworks as mina, netty, etc (what if something goes wrong with the new NIO2)?	1	1	https://issues.jboss.org/browse/JGRP-886
14622	JGRP-856	Hi, I'll be happy to make our 64bit machine available for any kind of testing thats deemed appropriate here. Would it help, for example, to just run one or more of the unit tests between it and another box?	1	1	https://issues.jboss.org/browse/JGRP-856
14623	JGRP-856	Excuse my ignorance, but what's a 64 bit Windows OS ? XP, Vista ? Which one do you use ? I don't know how I could test this, unless I purchased a license or used a VMWare image...	1	1	https://issues.jboss.org/browse/JGRP-856
14624	JGRP-856	I assume this image requires a user id and password, can you send them to me via email (bela at jboss dot org) ? Possibly also the admin user, so I can create new users and/or change the pwd myself. Happy Holidays to you too !	1	1	https://issues.jboss.org/browse/JGRP-856
14625	JGRP-822	How do we remove scopes ? If we have web sessions as scope, for example, we'll keep creating new sessions, therefore scopes in SCOPE. How do we purge them, e.g. when a session expires ? We should definitely associate scopes with the member with which they were created. When a member P leaves the cluster, all scopes associated with P should be removed in SCOPE. However, that doesn't solve the issue of creating and expiring sessions (= scopes) within P, when P doesn't leave... Maybe there should be a expire(Scope... scopes) method in SCOPE. This triggers a multicast to all SCOPE instances, and causes purging. The API to use this could be in Message, e.g. removeScope(Scope id), although this feel not right. Maybe a utility method in Util, which - given a channel - grabs the SCOPE protocol and calls expire() method ?	0	0	https://issues.jboss.org/browse/JGRP-822
14626	JGRP-813	Set disable_canolicalization to true by default: remove canonicalization ! This uses a hashmap lookup every time an IpAddress is created, inefficient ! We will canonicalize IpAddresses with LogicalAddresses to a much more efficient format (shorts) later anyway...	0	0	https://issues.jboss.org/browse/JGRP-813
14627	JGRP-806	The current Message (1.84) is not using the ConcurrentHashMap for the headers anymore. The CCHM is used for the canonical adrresses which is a class static instance. The message headers are currently handled using the default java HashMap class. It might be a good solution if there are many headers attached to a common message object. As far as I could see (debug) from our first test application there was just one header attached to messages recevied from the defalt channel (multicast-localhost). How many headers are added by the protocol stacks to sent/received messages? The header name/value pair could be just simply stored in an Object array and iterated for every header put and every header get call. The array would be extended when there are more headers than preallocated size (parameter) required. The second issue is with the synchronization object currently used for the headers access. The ReentrantReadWriteLock is pretty expensive class to be created for every message sent or received by the protocol stacks. Is there really a situation that the same message object is modified (putHeader) and processed (getHeader) from multiple threads? The message is either sent (down the stack) or received (up the stack) by a single thread execution. I think that an object monitor should be considered for the headers synchronization . # # EXAMPLE # private final Object headersMonitor = new Object(); private Object[] headers = new Object [DEF_HEADERS << 1]; ... public Header getHeader(String key) { synchronized(this.headersMonitor) { for(int i = 0; i < header.length; i+=2) { if(this.headers[i] == null) { return null; } else if(this.headers[i].equals(key)) { return (Header) this.header[i + 1]; } } } return null; } ...	1	1	https://issues.jboss.org/browse/JGRP-806
14628	JGRP-806	I don't think this would be very fast, especially as you need to iterate over the array to find the right header. The putHeader() method would also have to see whether a key of the same name already exists, and also iterate to do that. Do you have any measurements which show the opposite ?	1	1	https://issues.jboss.org/browse/JGRP-806
14629	JGRP-806	Ok, I've added MessageStressTest to JGroups. It measures the time to create messages, add and get headers. I implemented your suggested changes and saw a marginal improvement from 605'000 msgs/sec to 620'000 msgs/sec. While this is not really great, it is actually EXCELLENT because I'll create much less garbage. The only thing is that for putHeader(), I need to iterate through the array to see whether the key is already present, and not just find the first free space. Can you attach your impl to this case, so I can compare mine to yours ? I haven't checked in my impl yet...	1	1	https://issues.jboss.org/browse/JGRP-806
14630	JGRP-806	Excellent ! I changed the code locally today (no commits): #1 Use of IDs (ints) rather than strings as keys: performance was about the same, but space overhead was bigger #2 Use of intern() with all keys. Performance dropped to 50% ! So it seems that the Headers class is quite efficient ! Would you be able to contribute your changes to MessageStressTest to JGroups ?	1	0	https://issues.jboss.org/browse/JGRP-806
14631	JGRP-777	This works on Linux, but does NOT work on MacOS. I'll see whether I can add a flag checking for the presence of MacOSx. Or, probably better: I check for the presence of Linux and use the MulticastSocket() constrcutor which takes a SocketAddress, and port otherwise. Cross talking does not happen on Windows and MacOSX	1	1	https://issues.jboss.org/browse/JGRP-777
14632	JGRP-757	Seems that removing Event.INFO is easier path. Do we want to remove set/getInfo API from the JChannel and along with it info map that belongs to a JChannel? Should we store previous INFO values into additional data?	0	0	https://issues.jboss.org/browse/JGRP-757

#	Requirement	Questions	Answered	Successfully Answered	URL
14633	JGRP-755	Thinking a bit more about this, I realized that a stopping channel removes all of its submitted tasks from the timer, so if we have 3 channels which hold refs to a singleton timer, and they are closed, the timer has no references to any threads anymore, so it should get GC'ed. Oh, shit, not true: the TimeScheduler is a subclass of ScheduledThreadPoolExecutor, which maintains a pool of threads. Unless they are daemon threads, this will prevent the timer from being GC'ed. Do you agree on this ? I can make the timer's threads daemon threads, which would allow timer to be GC'ed.	0	0	https://issues.jboss.org/browse/JGRP-755
14634	JGRP-739	re-opened to apply patch. Vladimir, can you look into this ?	1	1	https://issues.jboss.org/browse/JGRP-739
14635	JGRP-735	[Michael Newcomb] Upon further thinking on this issue, it seems that when discovering existing 'members', PingRsp objects that come back with PingRsp.is_server == true need to be treated differently than those PingRsp objects with PingRsp.is_server == false... Right now, they are treated the same and *both* count towards num_initial_members which can cause problems in concurrent startup situation because the concurrently starting members might receive PingRsp objects from other concurrently starting members (PingRsp.is_server == false), hit the num_initial_members threshold, and trigger formation of a new group. So, I think there needs to be a shortcut in Discovery.Responses when a PingRsp is added, it should be checked to see if it is a real member (PingRsp.is_server == true). If it is it could signal the awaiting promise? This would effectively eliminate all wait times associated with joining an existing group because as soon as you receive a response from a real member, you would join his coordinator. I think you should join the first coordinator you see, because it is the most likely case and if you have more than one coordinator, a merge is most likely in place and even with the voting process, there is no guarantee that you selected the winning coordinator anyway (may have missed the 'lowest addresses' coordinator or missed a PingRsp that would have made one coordinator have more votes than another, etc...). If you are not a coord (PingRsp.isCoord() == false), you should behave as normal and just count towards the 'num_initial_members' threshold or the full timeout as normal.	1	1	https://issues.jboss.org/browse/JGRP-735
14636	JGRP-734	With the changes, ENCRYPT14KeyStoreTest fails with 2 errors, this wasn't the case with the old BouncyCastle JAR. Can you consult with the Red Hat expert to see what's going on ?	1	1	https://issues.jboss.org/browse/JGRP-734
14637	JGRP-733	I agree with Vladimir. The question is how do we expose individual channels (and their stacks) in AS 5? I think you already do this as part of the startup of an MBean don't you Brian ? I'm closing this issue	1	0	https://issues.jboss.org/browse/JGRP-733
14638	JGRP-726	Are we postponing this one until JDK 7 and internal GR overhaul that we talked about?	1	1	https://issues.jboss.org/browse/JGRP-726
14639	JGRP-719	Vladimir, can you do this ?	0	0	https://issues.jboss.org/browse/JGRP-719
14640	JGRP-719	Can you upload patches separately? I cannot apply this patch locally! I am using Eclipse as well.	1	1	https://issues.jboss.org/browse/JGRP-719
14641	JGRP-718	Message is read from stream passed to a Message object. I assume discard_messages_larger_than flag is a property of TP? How can we use discard_messages_larger_than flag from within a Message object? Maybe we can make this flag a JVM property?	1	0	https://issues.jboss.org/browse/JGRP-718
14642	JGRP-718	Or we can make a discard_messages_larger_than a static field of TP. Automatic assignment of discard_messages_larger_than will work using @Property mechanism. Are you ok with this kludge?	1	0	https://issues.jboss.org/browse/JGRP-718
14643	JGRP-705	Possibly provide implementations (subclasses) of java.net.Socket and ServerSocket ? This would allow existing code to run unchanged, but provide the flexibility of JGroups, e.g. add a compression protocol.	0	0	https://issues.jboss.org/browse/JGRP-705
14644	JGRP-705	Issue: what do we do with connections to non-existing members ? Failure detection doesn't run because it only starts running if we have 2 or more members, and since we don't have views for unicast channels, we will never start the failure detector ! We might have to modify FD slightly (or write a new unicast failure detector, FD and FD_SOCK), which emits a notification if one of the destinations ceases to respond to heartbeats. When this is the case, we need to remove the sender window for that destination, e.g. in UNICAST, or else we will retransmit messages to the non existing destination forever !	0	0	https://issues.jboss.org/browse/JGRP-705
14645	JGRP-701	The only downside is scalability. How is FD_ALL going to scale to very large clusters?	1	1	https://issues.jboss.org/browse/JGRP-701
14646	JGRP-663	Multiplexer does not treat OOB message between MuxChannels in any special way. It is added to fifo queue just as any other message for a certain MuxChannel. Should we bypass the fifo queue and deliver OOB directly to MuxChannel?	0	0	https://issues.jboss.org/browse/JGRP-663
14647	JGRP-653	- Receiver gets an abstraction in the form of an NIO selectable channel. Handback object needs to contain sender's and receiver's address. We don't need headers, as the app can write/read additional data directly to the channel - On the sender side, we have 1 output stream for the multicast address, and N for the unicast destinations. There's an abstraction (StreamableJChannel ?) which allows us to fetch the right output stream based on the destination - Provide an implementation of blocking method calls based on StreamingJChannel - Provide a new impl of RpcDispatcher/MessageDispatcher, based on an input and output stream - How do we send OOB messages (flags in general) ?	1	1	https://issues.jboss.org/browse/JGRP-653
14648	JGRP-653	Too many open questions, e.g. - this solves an edge case in JBossCache, so the priority is not high - Do we need changes in RpcDispatcher as well ?	0	0	https://issues.jboss.org/browse/JGRP-653
14649	JGRP-644	Are you referring to renaming of a thread that pushes message up the stack on a particular shared transport? For example, if a thread pushes message up the stack on a singleton "udp" then "udp" should be appended to a thread name?	0	0	https://issues.jboss.org/browse/JGRP-644
14650	JGRP-640	Question is do we support 'cactus stacks' ? How do we ensure that e.g. a a SUSPECT event is not demultiplexed to the same protocol N times (where N == number of multiplexed protocols) ?	0	0	https://issues.jboss.org/browse/JGRP-640
14651	JGRP-637	fixed in CVS head and the 2.6 branch (2.6.1)	0	0	https://issues.jboss.org/browse/JGRP-637
14652	JGRP-608	Michael, I was aware of the these problems but as we were in the process of simplifying flush (JGRP-598) I never got around to verify how concurrent flush works with a revised FLUSH version. Would you please verify FLUSH.java version 1.75 that I checked in this morning and see if you still experience these problems. Vladimir	1	1	https://issues.jboss.org/browse/JGRP-608
14653	JGRP-608	Ok. Lets reproduce it if possible.I'll abstract a unit test that attempts to recreate this scenario. Look at FlushTest that we have in a testsuite. It checks for event sequence. What we need is a version of ConcurrentStartupTest and ConcurrentStateTransferTest that checks for event sequence. I'll see to add it today. Have you had a look at the code changes?	1	1	https://issues.jboss.org/browse/JGRP-608

#	Requirement	Questions	Answered	Successfully Answered	URL
14654	JGRP-608	197 private boolean startFlush(Event evt, int numberOfAttempts, boolean isRetry) { 198 boolean successfulFlush = false; 199 if(!flushInProgress.get() || isRetry){ 200 flush_promise.reset(); While reviewing the new code, I came across the above in FLUSH.java. Does isRetry really outweigh the fact that a flush is in progress? If a flush is in progress, is there a need to retry starting a flush? Also, further down, it is possible for a flush to complete sometime after the initial timeout waiting for the flush to complete and the time the flush is attempted again. In fact, a flush attempt that timed out, could even complete a flush has just been reset (line 200). I think that flush messages need to be tracked by member/viewid/flush# or something. So, each flush message can be checked against a flush 'id'. Also, I think that having a flush 'object' tracking all the participants, acknowledged members, and even the rejected members would also aid in the debugging process. Anyway, I'm taking the latest code into our lab and will report back soon.	1	1	https://issues.jboss.org/browse/JGRP-608
14655	JGRP-608	JGRP-618 seems to have resolved this. Well, I no longer see consecutive 'unblock' calls with 5 simultaneous joining/state-requesting group members. As a side note, do you have an estimate of when connect-with-state-transfer will be done? I could really use it for my application. Thanks again, Michael	1	1	https://issues.jboss.org/browse/JGRP-608
14656	JGRP-570	Bela, what do you think? It would be very nice to have ThreadFactory and create threads from factory. Did you talk about exposing this API so users can hook their own factory? Maybe 2.6 afterall?	0	0	https://issues.jboss.org/browse/JGRP-570
14657	JGRP-567	I don't think we need this; substates are rarely used anyway... can we close this issue ?	1	0	https://issues.jboss.org/browse/JGRP-567
14658	JGRP-528	Just a couple of questions on this issue, and hope someone could advice. (I'm sorry, if this is not the correct place to ask questions - I'm new to using this website account) Basically my software is currently using version 2.4.1 SP1. 1) Using version 2.4.1 SP1, I've encountered this issue (spurious retransmission requests) as well. Is this normal? 2) When this issue happened, the CPU shoot up to 100%. Will this issue (spurious retransmission requests) caused CPU to shoot up to 100% ? 3) Is this issue solved in version 2.6.0? 4) If I upgrade to 2.4.1 SP4, will this issue happen, and caused CPU to shoot up to 100% as well? Thanks in advance.	1	0	https://issues.jboss.org/browse/JGRP-528
14659	JGRP-518	#1 What do you want to do if none of the responses are acceptable ? Timeout ? #2 How should 'discarded' (not accepted) responses look like in the RspList ? As not-received ? I would probably have to create a new mode GET_WITH_FILTER, which requires a valid filter. Hmm, I could also *not* do this and leave it up to the user to make sure that the filter always returns: for example, if you have mode GET_ALL, and discard only 1 valid response, and don't have a timeot, you'd block forever. Idea: if I marked every response I received which PASSED the filter as RECEIVED and every response which FAILED the filter as RECEIVED_NOT_ACCEPTED, then we'd know exactly which responses passed and which ones failed. Hmm, not good either: GET_FIRST would return as soon as 1 passed or failed response is received; it it is failed then we'd return with a response which didn't pass the filter. Okay, so the simple solution is then to just drop all responses which don't pass the filter, and it is the programmer's responsibility to make sure the call doesn't hang forever ! Example: - We expect responses from A,B,C - Mode is GET_ALL and timeout is 0 (wait forever) - Get get a response from A, discard it - Get responses from B and C ==> We hang forever	1	1	https://issues.jboss.org/browse/JGRP-518
14660	JGRP-518	Timeouts are ugly. I prefer letting the filter contain logic around returning. So how about if the filter looks like: public interface ResponseFilter { boolean isAcceptable(Object response); boolean needMoreResponses(); } Does that work for you? So every time a response comes in and you use the filter to test whether the response is acceptable, you also test whether you need to wait for more responses or you can just break and return the responses received so far.	1	1	https://issues.jboss.org/browse/JGRP-518
14661	JGRP-515	I've also seen this issue on Redhat Enterprise Linux 4 (not just Fedora). I don't think specifying the class D address you want to listen to in the MulticastSocket constructor would pick the wrong NIC though I haven't tested this because I've only run this on a single homed machine. What about another option in the UDP protocol stack to specify this?	1	1	https://issues.jboss.org/browse/JGRP-515
14662	JGRP-515	Forwarding to Bela for his comment as what Neil Horman said _may_ relate to a discussion he's been having with a user on the JGroups list. Unless Bela comes back and says Neil's comment has changed his opinion on how this stuff should work, I think we should pursue the issue. Ryan Campbell wrote: > Brian, We opened this issue in the RHEL Bugzilla: > > http://wiki.jboss.org/wiki/Wiki.jsp?page=PromiscuousTraffic <http://wiki.jboss.org/wiki/Wiki.jsp?page=PromiscuousTraffic>https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=231899 > > As you can see, we even attached a program to reproduce the issue. They have closed the issue as NOTABUG. > > Aleksandar suggested that we write the program as the suggested and perhaps reopen the bug: The only thing left to check is if binding to the proper socket will result in the desired behavior. I mean 2 processes can bind to the same address(group)/port as well not receive mcast messages to other groups. > > Should we pursue the issue? I think we have a fair chance of getting it fixed if we can demonstrate that they violate the RFC. > > Thanks, > Ryan > > > ---------- > > Subject: > [Fwd: [Bug 231899] multicast incorrect handling] > From: > Aleksandar Kostadinov > Date: > Tue, 22 May 2007 16:25:13 +0300 > To: > Ryan Campbell > > To: > Ryan Campbell > > > seems reasonable... > http://tldp.org/HOWTO/Multicast-HOWTO-2.html#ss2.4 > > The only thing left to check is if binding to the proper socket will > result in the desired behavior. I mean 2 processes can bind to the same > address(group)/port as well not receive mcast messages to other groups. > > > ---------- > ------- Additional Comments From xxxx 2007-05-22 08:55 EST ------- > According to the email as I read it, and from what I understand of the RFC, This > is working exactly as designed. the mcreceive program binds to INADDR_ANY > before adding multicast membership. That means the socket in each mcreceive > process should receive all multicast traffic from all groups that any process is > added to, which is exactly what you are seeing. If you want only the first > process to receive multicast frames, then you should write a separate program to > run in step 2 that binds to 225.12.12.14:1444, rather than INADDR_ANY. > > the Jboss kbase entry should be fixed to reflect that. > -- Brian Stansberry Lead, AS Clustering JBoss, a division of Red Hat	1	1	https://issues.jboss.org/browse/JGRP-515
14663	JGRP-512	What do you mean by this Brian ? I don't think JGroups uses the 'jboss' name space anyway.	1	1	https://issues.jboss.org/browse/JGRP-512
14664	JGRP-489	Vladimir, can you take a look ?	1	1	https://issues.jboss.org/browse/JGRP-489
14665	JGRP-489	Resolved. I was able to reproduce this issue frequently while working on http://jira.jboss.com/jira/browse/JGRP-158. The core of the problem was that SocketThread at GossipRouter assigned to handled input stream from each RouterStub (client) would read incorrect values from stream. One of the values was byte array size that needs to be allocated to deserialize message from RouterStub. If the value of the size was enormous such a byte array could not be allocated resulting in OOME. So how did this happen to begin with? I noticed that RouterStub allowed multiple thread to enter block of code that was writing a messsage to GossipRouter. When I refactored this block of code to be atomic I have never encountered manifestation this problem again.	0	0	https://issues.jboss.org/browse/JGRP-489

#	Requirement	Questions	Answered	Successfully Answered	URL
14666	JGRP-474	Done. When running both 2.4.1 SP2 and 5.0 beta-2, we see On 2.4: [WARN] [IncomingPacketHandler (channel=DrawGroupDemo)] TP.handleIncomingPacket(): packet from 192.168.5.2:4993 has different version (4416) from ours (2.4.1). Packet is discarded On 2.5: [WARN] [Incoming Thread,DrawGroupDemo,192.168.5.2:4992] TP$IncomingPacket.run(): packet from 192.168.5.2:4989 has different version (2.4.1) from ours (2.5.0). Packet is discarded	0	0	https://issues.jboss.org/browse/JGRP-474
14667	JGRP-430	Is GR list at RouterStub static or dynamic? Do we update the list with push or pull?	1	1	https://issues.jboss.org/browse/JGRP-430
14668	JGRP-375	Can we put this one in 2.6.2 as well? It is a very simple feature to implement. Instead of sending one GET_MBRS_REQ at a time we submit each in a separate thread to the thread pool. public void sendGetMembersRequest(String cluster_name) { for(Iterator<Address> it = initial_hosts.iterator();it.hasNext();){ final Address addr = it.next(); final Message msg = new Message(addr, null, null); msg.setFlag(Message.OOB); msg.putHeader(name, new PingHeader(PingHeader.GET_MBRS_REQ, cluster_name)); if(log.isTraceEnabled()) log.trace("[FIND_INITIAL_MBRS] sending PING request to " + msg.getDest()); timer.submit(new Runnable() { public void run() { try{ down_prot.down(new Event(Event.MSG, msg)); }catch(Exception ex){ if(log.isErrorEnabled()) log.error("failed sending discovery request to " + addr, ex); } } }); } }	0	0	https://issues.jboss.org/browse/JGRP-375
14669	JGRP-340	[email by ayyagarikiran at gmail.com] I have added BasicConnectionTable for routing. I have some doubts in that regard (please suggest me after reading this at your convenient time, will continue my dbugging till then) 1. I think the AddressEntry class is no longer needed due to the use of BasicConnectionTable 2. The Gossip requests CONNECT and DISCONNECT are also not needed 3. For every incoming request the BasicConnectionTable's readPeerAddress() method is getting called and is expecting a cookie and the address information so I am able to send requests only after writing that information from the GossipClient.(please see the writeGossipData() method in the attached GossipClient file) Is this way of writing cookie and Address info is correct (Connection.sendLocalAddress() didn't help as it is called only as part of getting a connection) 4. am able to send and register (after following the 3rd step) the GossipRouter.REGISTER request only if I add a check the network bytes available condition in BasicConnectionTable at line no. 588 if(in.available() > 0) { // read int Is this change acceptable? Not having that check is resulting in the below exception 10626 [TRACE] GossipRouter$SocketThread.receive(): - REGISTER(mygroup, 127.0.0.1:7000) 10626 [TRACE] BasicConnectionTable$Connection.run(): - exception is java.io.EOFException 10626 [TRACE] BasicConnectionTable$Connection.run(): - ConnectionTable.Connection.Receiver terminated 5. The ConnectionTable is expecting a non null address for every incoming request which may not be possible with the Gossip GET request.	0	0	https://issues.jboss.org/browse/JGRP-340
14670	JGRP-324	unicast address or a multicast address. Both types of packets have the same payload format (i.e. when we strip off the UDP header, we are left with a payload with the same format). It is this payload which Wireshark will have to dissect. JGroups uses message bundling to put several Messages into one UDP packet. Bundling can either be turned on or turned off, so we have to be able to handle both formats. UDP payload (non-bundled) ---------- version (2 bytes) flags (1 byte) Message (variable length) UDP payload (bundled) ---------- version (2 bytes) flags (1 byte) length (4 bytes) Address (variable length) Message (first bundled message - variable length) Message (second bundled message variable length) ... Message ('length' bundled message- variable length) NOTES: 1. In order to dissect these packets, we need to be able to dissect Address and Message. Address is read using Util.readAddress(). Message is read using Message.readFrom() 2. version format: e.g. JGroups 2.6.2 * X = 0-31 for major versions * Y = 0-31 for minor versions * Z = 0-63 for micro versions See the org.jgroups.Version class for algorithm on how to decode the version. 3. flags format: LIST (0x1) MULITCAST (0x10) OOB (0x100) 4. We can tell whether bundling is turned on by the value the expression (flags & LIST): * if (flags & LIST == false), then after reading version and flags, we expect only one message in the packet * if (flags & LIST == true), then after reading version and flags, we read length and expect 'length' messages in the packet In order to dissect the UDP payload, we need to be able to dissect serialized Address and Message objects. ---------- --- Address wireformat ---------- Address is a JGroups class which represents an 'address'. Depending on the transport we are using, this could be an <IP address, port> combination, or a JMS topic. The in-memory value of an Address is written/serialized using Util.writeAddress(). The wireformat is: isnull (variable length) isIpAddress (variable length) contents (variable length) The values isnull ("is the address equal to null?") and isIpAddress("is the address an IpAddress instance") are encoded Boolean values. Contents is variable and involves shipping classes across the wire. See the method org.jgroups.util.Util.writeAddress() for the specifc details. ---------- Message wireformat ---------- Message is a JGroups class and represents data shipped around the system. Messages contain the following data elements: dest_addr - destination address of the Message src_addr - source address of the Message flags - flags data - payload of the message list of headers - headers which Protocols use to store information Messages are written to the wire using the method org.jgroups.Message.writeTo(). The wireformat is: leading (1 byte) flags (1 byte) dest_addr (variable length) src_addr (variable length) buflen (4 bytes) buf ('buflen' bytes) numhdrs (2 bytes) (Key, Header) (first Header - variable length) (Key, Header) (second Header -variable length) ... (Key, Header) ('numhdrs' Header - variable length) NOTES: 1. leading is a set of flags indicating which elements have been written to the wireformat and which have not. DEST_SET (0x0000001) SRC_SET (0x0000010) BUF_SET (0x0000100) IPADDR_DEST (0x0010000) IPADDR_SRC (0x0100000) SRC_HOST_NULL (0x1000000) 2. flags indicate properties of the Message, whether it is out of bound, etc. OOB (0x001) LOW_PRIO (0x010) HIGH_PRIO (0x100) 3. dest_addr and src_addr are written using either IpAddress.writeTo() or Util.writeAddress(), depending on whether the Address is an IpAddress instance or not. I suspect that almost all will be IpAddress instances. See the method org.jgroups.Message.writeTo() for the specifc details on how Message is serialized. ---------- (Key, Header) wireformat ---------- Header is a JGroups class, and each protocol will have its own versions (e.g. TPHeader, FRAGHeader, etc.) which are subclasses of Header. Each Message has zero or more Headers associated with it. Each Header entry consists of a key and a value (e.g. ("UDP", TPHeader), ("FRAG", FRAGHeader), etc.) The wireformat is: key (variable length) magic (2 bytes) class (variable length, only if magic == -1) size (2 bytes) contents (variable length) NOTES: 1. key is the UTF8 encoding of the key String (e.g. "FRAG") 2. magic is the magic number	1	1	https://issues.jboss.org/browse/JGRP-324

#	Requirement	Questions	Answered	Successfully Answered	URL
14671	JGRP-324	Hello Yamato-san I spent a little more time looking at JGroups wireformats. I sent you earlier the wireformat for JGroups messages as sent by UDP/TCP. I wanted to add to that a concrete protocol as well, and I have done so for the lowest level protocol sitting on top of UDP/TCP - the PING protocol. I have included the wireformat for that message at the end of this email. Just as a reminder, we have the following structure for things sent across the wire in JGroups (up to and including PING): * Ethernet frame ** UDP/TCP packet *** JGroups TP packet, containing a header and one or more bundled Messages **** Message containing a PING header 1. I believe we now have enough information to start putting together a dissector which will handle completely PING messages. A PING message is going to be Message structure which contains a single PING header. We should make it our aim to get this working as a next step. The key problems I see in getting this working are: (i) dealing with bundled messages - Messages can contain one or more bundled Messages and as a first step in receiving a TP packet, need to be able to unbundle it if necessary. (ii) writing small C routines to dissect Address instances - Address instances are embedded both in TP packets, serialized Messages, as well as serialized PingHeaders. Solving (i) will involve reading up on how Wireshark handles bundled messages and getting some basic framework working to display the individual packets within a bundle in some reasonable way on the Wireshark GUI. Solving (ii) will simply be a matter of writing a Wireshark routine for displaying an arbitrary JGroups Address - mimicking the code in Util.readAddress(). Shouldn't be too difficult. 2. If you need a svn repository to store your Wireshark code, I can set up a repo for you on our internal QA repository. What do you think about 1 & 2? Richard PING wire format (PingHeader) ========== byte type: (1=PingHeader.GET_MBRS_REQ, 2=PingHeader.GET_MBRS_RSP) PingRsp response: can take value null, or a PingRsp instance PingRsp class ========== Address local ; Address coord ; boolean is_server ; These elements are serialised using PingRsp.writeTo() and PingRsp.readFrom() methods, which use DataStream methods to write bthe small pieces which make up the serialzed form. For example, Addresses are written using a combination of Util.writeStreamable() and Address.writeTo().	1	1	https://issues.jboss.org/browse/JGRP-324
14672	JGRP-324	I've created a new CVS project for the JGroups Wireshark plugin on SF under the name wireshark-plugin and uploaded a *draft* version of the plugin to SF. You can visit the project at http://javagroups.cvs.sourceforge.net/javagroups/wireshark-plugin. The project contains the sources for the plugin, some sample Wireshark capture files created using the JGroups Draw application, and instructions on how to get the plugin working with Wireshark. This plugin has been only minimally tested over UDP, and will be subject to continued development over the next couple of months. Up until the time we feel it is stable and release a "version", please keep in mind it is in a state of development and use is "as is". A few points: 1. The plugin should be able to display complete JGroups bundles (i.e. the TP header, all Messages in a bundle, as well as all Headers within a Message. At present, only TP and PING headers are fully processed; for the others, we just display the header key (e.g. STABLE) and the header size. 2. There is a user preferences section available from Edit->Preferences->Protocols->JGroups. The main options here you will need are: (i) telling Wireshark which port to associate with JGroups messages (ii) turning on debugging fields - which show pretty much all fields in the JGroups Message, as opposed to those which users would really be interested in. The next immediate development steps will be trying to get the UI presentation right and then testing over more JGroups configurations. Later development will deal with handling Wireshark conversations.	1	1	https://issues.jboss.org/browse/JGRP-324
14673	JGRP-324	Richard, it would be good if we can tie the type of dissector to the version of JGroups. This should be simple IMO, as the version ID (the short) is the first thing a JGroups message contains. I have a feeling the wire format's going to change once more in 2.8, with logical addresses, so it would be good if we can already now tie the dissector to the version, so when/if the wire format changes, we can still parse messages from previous versions. Do you think this is doable ?	1	0	https://issues.jboss.org/browse/JGRP-324
14674	JGRP-324	One comment, one question: 1. Comment: At present, the JGroups wireshark dissector (and every other dissector i've looked at) is written in this way: while (still more data in buffer) { - read some expected data item from the payload buffer and deserialize if necessary (using knowledge of the format of the data) - display the data on the user interface - advance the offset into the buffer by the length of the data read } I structured the code for the above loop into a set of functions, so it really looks like this: while(still more data in buffer) { - read and display TP header - get number of messages in bundle - for each (message in bundle) { - read and display message header - read and display message payload - get number of headers attached to message for each (header in message) { - read and display header } } } So, I have a lot of functions to carry out this work (2500 lines of C), from reading Digests and ViewIDs to deserializing class info based on magic numbers. And each function which involves reading from the payload buffer includes assumptions about the format of the payload buffer and its contents. So, knowledge of the wireformat is at present buried and scattered all over the code. One approach to handle different versions would be to scatter version checks throughout the single copy of the dissector (as in the old C code which runs on different platforms). Another approach would be to have a dissector for each version, but there is so much common code.....and every time we added a new chunk of code to the dissector, we would have to add it for each version. So, its unclear to me how to proceed with handling different versions at present. Another thought is to have wireformat specifications for each version, specifying the order in which data appears, the type, the display names, and so on, and base the dissector to work off that wireformat. But i'm not sure if that would be any better than including version information within the dissector itself (as in the old C style approach). 2. Question How many times has the wireformat changed, for which versions, and for which reasons? (i.e. what part of the format has changed)? Maybe if we look at the changes that have taken place (and might take place), a solution will present itself?	0	0	https://issues.jboss.org/browse/JGRP-324
14675	JGRP-278	Should the SocketFactory SPI be a singleton ? If yes, what are the implications regarding classloading ? We don't want to create 1 factory per classloader !	0	0	https://issues.jboss.org/browse/JGRP-278
14676	JGRP-236	How are we going to handle multiple concurrent joins(view bundling) where some joining members require transfer and others not?	0	0	https://issues.jboss.org/browse/JGRP-236

#	Requirement	Questions	Answered	Successfully Answered	URL
14677	JGRP-196	Email to Scott: Definitions: * Major: 1, minor: 1.1, micro: 1.1.1 * API compatibility: breaking is allowed between major versions * Binary compatibility: breaking is allowed between minor versions * Configuration compatibility: breaking is allowed between minor versions * Changes between micro versions: only implementation changes So, while incompatible API changes are not allowed between minor releases, there are allowed between major releases. That's why we have queued API changes to JBossCache and make them only in release 2.0. However, adding to an API, which doesn't break code using the API, is okay. We discussed design of a framework which allows for pluggable marshallers. Each message (or state transfer in JBossCache) is preceeded by a version number (a short in JGroups). That version number is key into a hashmap, whose values are marshallers. A marshaller knows how to unmarshall a message with a given version number to the one used by the current node. So, by convention, it is always the receiver, never the sender who has to be able to unmarshall messages with a version number different from its own. The marshalling tables are populated ar startup time, based on a directory (or shipped within the JAR and defined in META-INF, TBD). Each JAR has one or multiple marshallers and contains information in the META-INF about the version numbers they handle. This has not been thought through 100%, but I think it will work, and provide both backward and forward binary compatibility. Forward compatibility is an interesting one, as we could simply make new JARs available when new binary-incompatible versions are released, which can be downloaded (e.g. via JBoss Network), and simply be placed into the plugin directory. So, to upgrade a cluster to a new version, you'd copy the new JAR into the plugin dir, and restart the cluster in a staggered fashion. Once this is done, we'd simply restart each node in turn, with the new version. NOW, this is quite some work... I know, some folks are anal about binary compatibility around here, but IMO we're spending 90% of the time solving the 1% problem. For example, when customers switch from JDK 1.4.2 to 5, or from 5 to 6, both cases result in binary incompatibility, which our mechanism won't fix. So, customers will have to restart their clusters altogether. The other thing is that, even assuming the JDK migration is compatible, the JDKs themselves have changed serialVersionUIDs over and over between minor releases, so this won't help either. I intend to tackle this issue in JGroups 2.4: JGroups: http://jira.jboss.com/jira/browse/JGRP-196. Manik has a similar JIRA task, and we'll try to come up with a common solution. Thoughts ?	0	0	https://issues.jboss.org/browse/JGRP-196
14678	JGRP-188	which system properties can you not read from an applet ?	1	1	https://issues.jboss.org/browse/JGRP-188
14679	JGRP-188	Julien, what do you suggest we change ?	1	0	https://issues.jboss.org/browse/JGRP-188
14680	JGRP-181	Question is should we remove queues altogether ? That would mean we have only 1 set of semantics, but we'd have to go through each protocol to see whether it is reentrant	0	0	https://issues.jboss.org/browse/JGRP-181
14681	JGRP-181	- Remove up and down queues - The Channel.send() method runs on the caller's thread - Messages/events have (1) a reliability flag (http://jira.jboss.com/jira/browse/JGRP-180) and (2) a FIFO flag - When the reliability is set, when passing an event up/down, and an exception occurs, we retry (after N ms pause) to deliver the event, until no exception is encountered (this prevents intra-stack loss) - We have a thread pool at the transport level - The receiver thread passes the message to the threadpool if the FIFO flag is *not* set. This means the message can essentially be treated as an out-of-band message, and it can be passed up the stack on its own thread from the thread pool, asynchronous to other messages. When the FIFO flag is set, we pass the message up in the right order viz other messages, this can be done by enqueueing it and having a single thread handle the queue, or by using the receiver's thread (the one that called DatagramSocket.receive()). - The thread pool would be util.concurrent's Executor, and is externalized (with a default), so devs can plug in their own thread pool impl	0	0	https://issues.jboss.org/browse/JGRP-181
14682	JGRP-178	Hi Takayoshi, so you are saying that you did see the NAKACK problem with 2.2.9/2.2.9.1 but *not* with 2.3 beta ? Was that NAKACK problem reproducible ? Cheers,	0	0	https://issues.jboss.org/browse/JGRP-178
14683	JGRP-178	Looks like the jgroups NAKACK error has been fixed in 2.3beta. But we are still seeing the runPrepare phase failed during load from JBossCache that leads to timeout exception. So can it be that JGroups ate the TImeoutEception previously that leads to this error?	0	0	https://issues.jboss.org/browse/JGRP-178
14684	JGRP-178	Is this issue still present or can I close this case ?	1	1	https://issues.jboss.org/browse/JGRP-178
14685	JGRP-130	Is there an ETA on this functionality/2.6?	0	0	https://issues.jboss.org/browse/JGRP-130
14686	JGRP-130	Implementing GMS.reject_join_from_existing_member pretty much fixed this issue. I don't see the need for logical addresses any more. I added a comment to http://jira.jboss.com/jira/browse/JGRP-129. What do people think, is there still a compelling reason we need logical addresses ?	1	1	https://issues.jboss.org/browse/JGRP-130
14687	JGRP-130	This is resolved...is there a release with the new code in it? Do I need to compile from svn?	0	0	https://issues.jboss.org/browse/JGRP-130
14688	JGRP-130	Bela, seems the problem can be resolved without logic addresses, I looked into detailed logs and can see that almost everything works perfectly: the member who sent JOIN_REQ receives JOIN_RESP but ignores it, I do not know why. Maybe my tcp configuration is incorrect or I missed something?	1	1	https://issues.jboss.org/browse/JGRP-130
14689	JGRP-129	Has this code changed since JGroups 2.2.8? We have recently upgraded from JGroups 2.2.8 to 2.2.9.1 and we have a member in a cluster which utilizes a Virtual IP. With JGroups 2.2.8, the member with the virtual IP had no problems joining the group, however after we upgraded to 2.2.9.1, we noticed that that particular server was sending join requests to the coordinator, the coordinator received the requests, sent the join response to the virtual IP, but the participant never received the join response. As a result, the participant continues to send join requests and is never able to join the group. After removing the virtual IP, the node was able to connect to the group successfully. We would like to use the virtual IP for our system for failover, where a passive server can claim the same IP if the primary server goes down for some reason. Is there a hot-fix or workaround that we can use in our system to allow our configuration to work? Thanks. This is the config that we are using: <config> <UDP loopback="false" mcast_recv_buf_size="16384" mcast_send_buf_size="8192" mcast_addr="228.1.2.3" mcast_port="45566" ip_ttl="32" /> <PING down_thread="false" up_thread="false" timeout="5000" num_initial_members="2" /> <FD down_thread="false" up_thread="false" timeout="5000" max_tries="3" shun="true" /> <VERIFY_SUSPECT down_thread="false" up_thread="false" timeout="5000" num_msgs="5" /> <pbcast.NAKACK down_thread="false" up_thread="false" gc_lag="20" retransmit_timeout="300,600,1200,2400,4800,7000,9000" max_xmit_size="4096" discard_delivered_msgs="true" /> <UNICAST down_thread="false" up_thread="false" /> <pbcast.STABLE down_thread="false" up_thread="false" /> <pbcast.GMS down_thread="false" up_thread="false" join_timeout="5000" join_retry_timeout="2000" shun="true" print_local_addr="true" /> <FLOW_CONTROL down_thread="false" up_thread="false" /> <FRAG down_thread="false" up_thread="false" frag_size="4096" /> <COMPRESS down_thread="false" up_thread="false" min_size="50000" /> <VIEW_ENFORCER down_thread="false" up_thread="false" /> </config>	1	1	https://issues.jboss.org/browse/JGRP-129

#	Requirement	Questions	Answered	Successfully Answered	URL
14690	JGRP-129	I'm not sure I want to introduce logical addresses. This is regarding JGRP-129 (logical addresses) and JGRP-130 (problems with reincarnation). JGRP-130 is about a member that leaves and rejoins *under the same IP address and port (using bind/start_port)* before failure detection kicks it out. I implemented a fix for this, which won't allow a member into the group that's already in the group. In the above case, with A and B, B would have to wait until the cluster is {A} before it can rejoin. JGRP-129 is about providing a logical address (possibly host + system time in ms (plus 1 for same system times) for a member, which can be mapped to a physical host/port combo. It would add 4 bytes/address, and I don't see the benefits of this. The original motivation was that when a member gets shunned, it rejoins under a new identity. Sacha didn't like this and therefore added 'logical addresses' in JBoss, which creates a logical address for a member. So when that member is shunned and rejoins under a *new* JGroups address, it would still keep the *old* JBoss logical address. I tend to leave it this way, as I don't want to add something which - adds clutter (Host:26634425614 looks less readable than Host:4536) - adds 4 bytes / IpAddress - has some additional lookup in the transport to find the real physical address to which the logical address maps - needs some ARP-like traffic for establishing local tables which map logical to physical addresses What do people think ?	0	0	https://issues.jboss.org/browse/JGRP-129
14691	JGRP-118	Can the state_id be a byte[]? This would allow the app to use any type of object, for instance a serialized Fqn (which cannot always be accurately represented as a string).	1	1	https://issues.jboss.org/browse/JGRP-118
14692	JGRP-118	Can't you use a string rep of the Fqn ? I'll think about it...	0	0	https://issues.jboss.org/browse/JGRP-118
14693	JGRP-111	- RpcDispatcherSpeedTest -mode method now runs 230 methods/sec (was: 122/sec) - Added MessageDispatcherSpeedTest, SynchronousMessageSpeedTest - RequestCorrelator.Header now supports Streamable - GroupRequest: changed synchronization and streamlined code - GroupRequest: added mode "id" - Added GroupRequestTest unit test	0	0	https://issues.jboss.org/browse/JGRP-111
14694	JGRP-104	Okay, I'm adding this to JGroups, so that if jboss.partition.udpGroup is set, it will override the property mcast_addr in the XML file. While at it, would you want to override the mcast_port as well (jboss.partition.udpPort) ? This will be in JGroups 2.2.9	1	1	https://issues.jboss.org/browse/JGRP-104
14695	JGRP-104	"Okay, I'm adding this to JGroups, so that if jboss.partition.udpGroup is set, it will override the property mcast_addr in the XML file. While at it, would you want to override the mcast_port as well (jboss.partition.udpPort) ? " Does this still allow somebody to explicitly set the mcast_addr? e.g. somebody might want all the channels on one jboss.partition.udpGroup except one which is part of a greater federated cluster. I was thinking that the system property could be used to allow users to not specify the value instead using the value specified at JVM boot time through the system property. It would also be good (but probably requires a separate feature request) to use something like a gossip server to automatically lease unique udp addresses like a DHCP server does for ip addresses, such that developers on a network can isolate themselves during testing.	0	0	https://issues.jboss.org/browse/JGRP-104
14696	JGRP-100	Further points to look at: GMS: - Is the single coordinator a bottle neck ? When a coord hangs, it won't be able to admit new nodes into the cluster, until it is suspected and excluded - Use view_bundling Failure detection: - Replace FD with FD_ALL TP, timers: - Increase thread pools NAKACK: - Set use_mcast_xmit="true" if UDP is used. UDP is recommended anyway for large clusters - Use stats for xmit_timeouts ? PBCAST: - Revisit ? SMACK: - Use instead of NAKACK ? Doesn't require GMS - Improve: don't ack every single message ! FC: - Increase max_credits FRAG2: - Use instead of FRAG - Get latest version (fix for https://jira.jboss.org/jira/browse/JGRP-800) Streaming API (?): - For large values ? Discovery: - Use only servers, not clients (https://jira.jboss.org/jira/browse/JGRP-735) - Workaround: set num_responses to a large value, so we do get some server responses	0	0	https://issues.jboss.org/browse/JGRP-100
14697	JGRP-84	// Author: Bela Ban // $Id: JmxSupport.txt,v 1.2 2005/06/14 14:00:59 belaban Exp $ JMX support for JGroups ========== In 2.2.9, JMX support was added to JGroups. There were several design goals: - No dependencies on JMX if JMX support is not required. It should be possible to ship a JGroups version without any JMX support, and without jmxri.jar (for JDK 1.4.x). Therefore channels and protocols are not turned into MBeans, e.g. FD doesn't implement FDMBean. Instead, a parallel MBean and MBeanImpl hierarchy was created, where the MBeanImpl delegates to the real protocol. This hierarchy is located in package org.jgroups.jmx. For a protocol org.jgroups.protocols.P, there is a corresponding org.jgroups.jmx.protocols.PMBean MBean and an implementation org.jgroups.jmx.protocols.P - No negative effect on performance, minimal effect on memory. JMX support can be disabled in the channel, and all protocols (stats="false") - Support for running in a pre-JMX JDK (1.4, for example), e.g. under JBoss. There is a jgroups-service target in the ANT build which generates a jgroups-service.sar that can be deployed into JBoss. The channel defines in the META-INF/jboss-service.xml file of the SAR defines a JGroups channel MBean, which will automatically create MBeans for all protocols and register them with the MBeanServer - Support for running in JDK 5 standalone. JDK 5 has JMX support built in. The jconsole tool can be used to inspect the MBean hierarchy. To programmatically create MBean support for JDK 5, the following code can be used: JChannel channel; channel.connect(groupname); ArrayList servers=MBeanServerFactory.findMBeanServer(null); if(servers == null || servers.size() == 0) { throw new Exception("No MBeanServers found;" + "\nJmxTest needs to be run with an MBeanServer present, or inside JDK 5"); } MBeanServer server=(MBeanServer)servers.get(0); JmxConfigurator.registerChannel(channel, server, "JGroups:channel=" + channel.getChannelName() , true); JmxConfigurator has a number of helper methods which register and de-register the channel and/or te protocols attached to a given channel. Provided information ---------- Some information provided includes (all info can be reset), besides configuration: - UDP: messages sent/received, bytes sent/received (per member) - Discovery: manual triggering of discovery - NAKACK: retransmissions per member, last N retransmitted messages, lowest and highest sequence numbers per member - GMS: membership, history of views - FD/FD_SOCK: last N suspected members - STATE_TRANSFER: number of state requests, total bytes transferred - STABLE: manual triggering of STABLE protocol Example ------- The Draw program was modified to support JMX. The following steps describe how to configure and run a JDK 5 VM and look at the resulting Draw application via jconsole. - Run the Draw program (the JDK needs to be 5 !): java -Dcom.sun.management.jmxremote org.jgroups.demos.Draw -props c:\fc-fast.xml -jmx (the -jmx flag enables registering the Channel and protocols with the MBeanServer) - Start jconsole and connect to the local Draw process - Under 'MBeans', choose the JGroups channel and browse the channel and the protocols - [optional] Add the STATS protocol to the top of the stack; it display information about messages sent and received	1	1	https://issues.jboss.org/browse/JGRP-84

#	Requirement	Questions	Answered	Successfully Answered	URL
14698	JGRP-82	Here's a copy of FLUSH.txt FLUSH protocol for default stack ================================ Author: Bela Ban Version: $Id: FLUSH.txt,v 1.2 2006/01/27 13:00:00 belaban Exp $ Overview -------- Flushing means that all members of a group flush their pending messages. The process of flushing acquiesces the cluster so that state transfer or a join can be done. It is also called the stop-the-world model as nobody will be able to send messages while a flush is in process. When is it needed ? ---------- (1) On state transfer. When a member requests state transfer, the coordinator tells everyone to stop sending messages and waits for everyone's ack. Then it asks the application for its state and ships it back to the requester. After the requester has received and set the state successfully, the coordinator tells everyone to resume sending messages. (2) On view changes (e.g.a join). Before installing a new view V2, flushing would ensure that all messages *sent* in the current view V1 are indeed *delivered* in V1, rather than in V2 (in all non-faulty members). This is essentially Virtual Synchrony. Design ------ As another protocol FLUSH which handles a FLUSH event. FLUSH sits just below the channel, e.g. above STATE_TRANSFER and FC. STATE_TRANSFER and GMS can be configured with a boolean use_flush attribute which determines whether flushing is used on state transfer and/or view changes. If so, the protocol just sends a SUSPEND event up the stack, where it is handled by the FLUSH protcol. The SUSPEND_OK ack sent back by the FLUSH protocol let's the caller know that the flush has completed. When done (e.g. view was installed or state transferred), the protocol sends up a RESUME event, which will allow everyone in the cluster to resume sending. FLUSH protocol: - On SUSPEND event: multicast START_FLUSH to everyone (including self) - On reception of START_FLUSH: - set suspended flag to true, all subsequent down() methods will block on mutex - *multicast* a FLUSH_OK message - On reception of FLUSH_OK at non-coordinator: discard - On reception of FLUSH_OK at coordinator: - Record FLUSH_OK in vector (1 per member) - Once FLUSH_OK messages have been received from all members --> send down SUSPEND_OK event (in case of crash notification: mark that member's entry as received, so we don't hang forever) - On RESUME event: multicast STOP_FLUSH message - On reception of STOP_FLUSH message: unblock down() - On view change: - Mark crashed member as 'received' (see above) - On suspect(P) - If P is the coordinator and we're next in line - Run the FLUSH protocol from scratch Notes ---- *Multicasting* (as compared to unicasting) the FLUSH_OK message is required so we can be sure that the FLUSH_OK message sent by P is the last message received from P before P stopped sending messages. Since we're waiting on everyone's FLUSH_OK, we can be assured that everyone's previous messages were received (thanks to reliable mcast).	1	1	https://issues.jboss.org/browse/JGRP-82
14699	JGRP-74	Test is StateTransferTest	0	0	https://issues.jboss.org/browse/JGRP-74
14700	JGRP-65	From Eric: For the transmit queue: $ /sbin/ifconfig [eth0|eth1] The txqueue size is currently 1000. To get the kernel queue size for incoming packets: $ sysctl -a | grep net.core.netdev_max_backlog For socket buffer info: $ sysctl -a | grep net.core net.core.divert_version = 0.46 net.core.hot_list_length = 128 net.core.optmem_max = 10240 net.core.message_burst = 50 net.core.message_cost = 5 net.core.mod_cong = 290 net.core.lo_cong = 100 net.core.no_cong = 20 net.core.no_cong_thresh = 20 net.core.netdev_max_backlog = 300 net.core.dev_weight = 64 ** these are the default and max sizes for socket buffer memory net.core.rmem_default = 65535 net.core.wmem_default = 65535 net.core.rmem_max = 131071 net.core.wmem_max = 131071 Here are some other links: Squeeze Your Gigabit NIC for Top Performance http://www.enterpriseitplanet.com/networking/features/article.php/3497796 How to achieve Gigabit speeds with Linux http://datatag.web.cern.ch/datatag/howto/tcp.html Enabling High Performance Data Transfers http://www.psc.edu/networking/projects/tcptune/ >> #2 What's the MTU (do you use jumbo frames, or the regular 1500b) ? In general, we have not tuned this environment. This needs to be addressed as part of the QA Lab upgrades we're targeting for next month. I'll be creating some Jira tasks and working with Ryan's team on this, so we should discuss the network needs along with other system relates topics (oracle tuning, etc.). Regular 1500 for now. We need to revisit this for some upcoming QA lab updates were doing. If we need a dedicated GIGE switch for testing, we need to get this in the budget (I think it's worth the 1.5k myself). If we can't get a dedicated switch, we'll see if we can enable the jumbo frames on the private side subnet (but this requires ALL connected machines be gige w/ jumbo frame support - not an issue w/ the current hw, but something to think about).	1	1	https://issues.jboss.org/browse/JGRP-65
14701	JGRP-28	[Lez Haslewoord] We don't want to use the bind.address property, as it sets that value for the entire VM and all components that reference it (i.e. clustering for DefaultPartition, Tomcat-Cluster, etc). The client requirement is to run JBoss in clustering mode at all times, but only have a node join the group when desired. When testing, the developers just want to bind to the localhost to ensure the mcast packets aren't sent to other members on the network, thereby preventing an undesireable group membership. The only change needed to go to production would be to change the bind address to the default (eth0) ip/interface. Given the following UDP stack config on 2 separate machines (or 2 jboss instances on the same machine): <UDP mcast_addr="228.1.2.3" mcast_port="45566" bind_addr="127.0.0.1" ip_ttl="32" ip_mcast="true" mcast_send_buf_size="800000" mcast_recv_buf_size="150000" ucast_send_buf_size="800000" ucast_recv_buf_size="150000" loopback="true"/> we verified the following w/ ethereal on a 3rd machine in the same network (this machine isn't running JBoss - we're just using it as a network sniffer): 1. During startup of node 1 and node 2, no JGroups UDP packets originating from either host are being sent on the network. (so far, so good) 2. Long after both nodes have started up, no JGroups UDP packets originating from either host are being sent on the network. (still good) 3. If either node shuts down their JBoss instance, 1 and only 1 JGroups UDP packet originating from that node is being sent to the network during shutdown. (not so good). This means that packet was not swallowed by the loopback interface, and intead it was sent via the default interface (eth0), even though we've explicitly configured the UDP stack to _not_ do that. I'm not sure why that 1 packet would be sent (maybe its related to broadcasting a member leave event?). The concern in this case is that we don't want other machines on the network to receive that packet, implying that other nodes do exist. If it _is_ a leave event, the other members would probably ignore it since the node that sent it was never part of their group to begin with. That may be the case, but it would be nice to know why this packet ever goes to the network on the default interface, instead of being swallowed by the loopback interface as it should have been.	0	0	https://issues.jboss.org/browse/JGRP-28

#	Requirement	Questions	Answered	Successfully Answered	URL
14702	JOPR-320	(8:49:29 AM) ccrouch: i'm not sure these new names are going to be any clearer to a basic user (8:49:34 AM) ccrouch: what do you think? (8:51:12 AM) ccrouch: are the following better?: (8:51:12 AM) ccrouch: Enterprise Apps (8:51:12 AM) ccrouch: EJB2 Apps (8:51:12 AM) ccrouch: EJB3 Apps (8:51:12 AM) ccrouch: Embedded Web Apps (8:51:12 AM) ccrouch: Resource Archives (8:51:12 AM) ccrouch: Standalone Web Apps (8:51:21 AM) ccrouch: ozizka: comments? (8:56:00 AM) ozizka: ccrouch: Hi, I don't know how native english speakers call those app types, (8:56:27 AM) ozizka: but afaik, everyone I know just uses the abbrevs. (8:57:13 AM) ozizka: If it were up to me, I would use something like EAR (Enterprise Archive) and similar (8:58:06 AM) ccrouch: ozizka: yeah i think we need something more than three letter acronyms here (8:59:21 AM) ips: ccrouch - there's also the res description which contains a more verbose version (8:59:42 AM) ips: i like keeping the res type names concise for the sake of the leftnav (9:00:04 AM) ccrouch: but aren't we just working around the real problem (9:00:13 AM) ccrouch: that you cant drag the left nav to make it wider? (9:00:17 AM) ips: don't you think any jboss admin/oper will know what war, ear means? (9:01:24 AM) ips: yes and no, even once we support that, it's nice to not have to make the leftnav overly wide, since it steals real estate from the content (9:05:58 AM) ccrouch: try, but i think we can do better than what we have now, maybe smaller font or something along the hybrid approach ozizka mentioned (9:05:58 AM) ccrouch: anyways i'm reopening and moving to 1.3 so we can think about it some more	1	1	https://issues.jboss.org/browse/JOPR-320
14703	JOPR-320	(9:06:48 AM) ips: ccrouch - ok. of course, we need to finalize the names before jon 2.3, since the names are also the keys for restypes (9:07:15 AM) ccrouch: yep (9:07:31 AM) ccrouch: but don't we have a displayname and a real name? (9:07:33 AM) ips: (why did we do that? we should have both name and displayName like we have for every other metadata entity) (9:07:42 AM) ccrouch: ha (9:07:43 AM) ips: :) (9:08:31 AM) ccrouch: +1, thats bad. We should see about adding a displayName, which if not present just defaults to the real name (9:09:03 AM) ccrouch: one other thing (9:09:03 AM) ccrouch: EJB2 is not strictly accurate since EJB1.x beans would go in there too? (9:10:30 AM) ccrouch: also for consistency we should think about embedded ejb jars i guess, or are they included already	1	1	https://issues.jboss.org/browse/JOPR-320
14704	JOPR-296	<ips> jweiss - the reason it was closed is bc if we don't fix it for 2.3, we probably won't ever fix it <jweiss> ips: you mean you *can't* fix it which should have upped the urgency to fix it now <ips> bc it would require changing resource types, which we don't even support ie - once we get the as5 plugin out in the wild, we wouldn't be able to cleanly upgrade users' existing DS and CF resources <jweiss> ips: i hear you. i think that's a great reason to fix it now <ips> so if you really feel strongly about it, i'd try to convince greg or charles to fix it for 2.3 <jweiss> ips: ok. ccrouch ^ let me know what you think <ips> i personally am in favor of keeping it as is, bc i think 3 restypes for the different types of DS's makes more sense since they have disparate properties <jweiss> ips: so you think the as4 way was wrong and it's too late to change that one <joseph> ips, jweiss: using the enhancements to the resourceType/plugin drop-down, can you filter it such that you can see all datasources on a page? <ips> yeah, too late to change the as4 one because we don't support renaming or merging/splitting restypes joseph: the AS5 plugin has a Datasources subcat so you can see all datasources using the leftnav <joseph> i mean from res browser <ips> we could potentially incorporate subcat headings into other pulldown menus, but i think right now they're only used by the leftnav <jweiss> hm. ok ips you convinced me. since the ideal fix is not possible, and making it consistent would be the "wrong" fix, i guess there's no other option but to leave it as is	0	0	https://issues.jboss.org/browse/JOPR-296
14705	JOPR-262	Is this in Jopr 2.3.1? Rhq-agent is not able to detect my JBoss 5 instance with a secured JMX console. Do I have to give rhq-agent the JBoss path or specify some other configuration?	0	0	https://issues.jboss.org/browse/JOPR-262
14706	JOPR-258	Do we have a JON logo to replace them with or should there just not be any logo?	1	1	https://issues.jboss.org/browse/JOPR-258
14707	JOPR-258	> Charles Crouch: From greg: Still need to replace the web icon I assume you mean the favicon. What do we want to use - the new JBoss logo?	1	1	https://issues.jboss.org/browse/JOPR-258
14708	JOPR-230	Lukas, does the SBM actually expose a property which gives us the list of available binding sets? If it does, we could present that as a readonly list to the user, so they could choose the correct one	1	1	https://issues.jboss.org/browse/JOPR-230
14709	JOPR-230	The users can see the list of the defined binding sets as children of the Service Binding Manager in the resource tree. I originally planned for exactly this feature (i.e. the selection being a combobox) but, for the reasons Charles mentioned, I failed at it. We could certainly expose the readonly list directly on the edit page but I didn't do that because I found it a bit confusing. Additionally, as I mentioned above, the user is able to get at the list by expanding the Service Binding Manager resource in the nav tree. I agree with Ondrej that this isn't the most bulletproof approach but at least there is a check present that will prevent the user from entering a non-existant binding set name. I.e. an attempt to save the configuration with an invalid binding set name will result in an exception. An attempt to delete a binding set that is currently active throws an exception as well. For now, I will add a note in the property description describing the dangers and effects of changing the active binding set name. What other options do we have to make this more intuitive and foolproof?	0	0	https://issues.jboss.org/browse/JOPR-230
14710	JOPR-210	rev751 only added support for 2 metrics, create and remove counts, and an operation fro retrieving method invocation statistics for session beans. We need to make sure we implement as much of the metrics mentioned in JBAS-6625 as possible. This is a table of what was requested and what Profile Service offers: SLSB Create Calls = CreateCount Remove Calls = RemoveCount Method-Ready Beans = ? Pool Size = ? Max Pool Size = ? SFSB Create Calls = CreateCount Remove Calls = RemoveCount Method-Ready Beans = ? Passive Beans = ? Entity Beans Create Calls = CreateCount Remove Calls = RemoveCount Method-Ready Beans = ? Pooled Beans = CacheSize? Cache Size = CacheSize? Passivated Count = ? Current Size = ? Max Size = ? MDB Create Calls = CreateCount Remove Calls = RemoveCount Method-Ready Beans = ? ? = MessageCount The DetypedInvocationStatistics managed component provided by Profile Service offers these properties: - concurrentCalls - maxConcurrentCalls But what do they mean exactly?	1	1	https://issues.jboss.org/browse/JOPR-210
14711	JOPR-178	Where would I find these call-time metrics? I have deployed a sample EJB3 jar but I see nothing that indicated what, if any, call-time metrics exist or should exist.	0	0	https://issues.jboss.org/browse/JOPR-178

#	Requirement	Questions	Answered	Successfully Answered	URL
14712	JOPR-177	Okay, I've seen that JON 2.2 already contains a feature like this (Type = DiskSource). Unfortunately, I am not able to fill it with EARs or WARs. What I did: 1. Create new Content Source 2. Root Directory = [somewhere on my local disk] 3. Filename Filter: .*\.ear 4. Package Type Name: ear 5. Architecture: noarch 6. Resource Type Name: JBossAS Server 7. Resource Type Plugin Name: JBossAS I put a file called TestSeam-1.0.0.ear on the path specified by [2]. After pressing on Synchronize, I get the following log: 2009-05-06 01:05:42,146 INFO [org.rhq.enterprise.server.plugin.content.ContentSourcePluginContainer] Scheduled content source sync job [b99c6466--dead438--121130681fc:syncContentSource] to fire now at [Wed May 06 01:05:42 CEST 2009] for [ContentSource: id=[500050], name=[My Apps]] 2009-05-06 01:05:42,193 INFO [org.rhq.enterprise.server.plugin.content.ContentSourceAdapterManager] synchronizeContentSource: [My Apps]: loaded existing list of size=[0] (6)ms 2009-05-06 01:05:42,282 INFO [org.rhq.enterprise.server.plugin.content.ContentSourceAdapterManager] synchronizeContentSource: [My Apps]: got sync report from adapter=[PackageSyncReport: new=[1], updated=[0], deleted=[0]] (85)ms 2009-05-06 01:05:42,300 WARN [org.rhq.enterprise.server.content.ContentSourceManagerBean] Content source adapter found a package of an unknown package type [ear|ResourceType[id=500568, category=Server, name=JBossAS Server, plugin=JBossAS]] Skipping it. 2009-05-06 01:05:42,314 INFO [org.rhq.enterprise.server.plugin.content.ContentSourceAdapterManager] synchronizeContentSource: [My Apps]: report has been merged (28)ms 2009-05-06 01:05:42,324 INFO [org.rhq.enterprise.server.scheduler.jobs.ContentSourceSyncJob] Content source [My Apps] is fully synchronized now. It is marked for lazy loading - bits will not be downloaded now. So what am I doing wrong here?	0	0	https://issues.jboss.org/browse/JOPR-177
14713	JOPR-171	I built and deployed the app attached to this jira, and can view the xml feed, but I cannot download "patches/jboss-5.0.0.GA_CP01.zip", it tells me it's not found if I try to manually download it -- that is, despite the fact that it's quite evidently there on the filesystem. Is there something that needs to be configured on my end to make this test app work? Update: I hacked up the XML to host the patch on a different server. This appears to get me past the issue above. However, attempting to deploy to any EAP5 server returns a Failure with an NPE that reads akin to: java.lang.NullPointerException at org.rhq.plugins.jbossas5.ApplicationServerContentFacetDelegate.deployPackages(ApplicationServerContentFacetDelegate.java:112) at org.rhq.plugins.jbossas5.ApplicationServerComponent.deployPackages(ApplicationServerComponent.java:279) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:616) at org.rhq.core.pc.inventory.ResourceContainer$ComponentInvocationThread.call(ResourceContainer.java:525) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:334) at java.util.concurrent.FutureTask.run(FutureTask.java:166) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1110) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:603) at java.lang.Thread.run(Thread.java:636)	1	1	https://issues.jboss.org/browse/JOPR-171
14714	JOPR-169	Asaf, please look into whether we can do restart/stop via the Profile Service, or whether we will have to use JMX. Also, please describe your impl plan - do you plan on providing the option to use either a script or a JMX (or PS) operation for restart and stop, as we have in the as4 plugin? Start will, of course, have to be done via script :)	0	0	https://issues.jboss.org/browse/JOPR-169
14715	JOPR-166	What exactly do you have in mind? Just the plain JDBC check if the JDBC-url + credentials are valid?	1	0	https://issues.jboss.org/browse/JOPR-166
14716	JOPR-162	Shelly can you attach the .rar you are using for this test	0	0	https://issues.jboss.org/browse/JOPR-162
14717	JOPR-159	The amount of flexibility in the ServiceBindingManager is precisely what makes managing this stuff difficult. You can use the service or not and you can use it different ways. In the end it can be difficult for Jopr to figure out where the data came from. Was it xslt'd? etc. To be honest, it's probably easier for us to go around to each different service file and gather up the ports and display it on one page then to work backwards through these bindings. Or do you think with AS5 we have to just assume their using the new default setup of the SBM?	0	0	https://issues.jboss.org/browse/JOPR-159
14718	JOPR-137	Farah, Do you want to try tackling this one? It should consist of just adding an additional resource-config property to the as5-plugin's descriptor.	1	1	https://issues.jboss.org/browse/JOPR-137
14719	JOPR-137	Ian, I tried adding the following resource-config property locally to the as5-plugin's descriptor: <resource-configuration> <c:simple-property name="partitionName" displayName="Partition Name" description="The name of the cluster partition this app server instance belongs to" type="string"/> </resource-configuration> However, when I deployed the console and tried to click on the "Configuration" tab for "JBoss AS 5 (default)", the following error occurred: 14:49:56,910 ERROR [PluginContainerResourceManager] Error retrieving configuration for resource: Resource[id=-3, type=JBossAS Server, key=/home/fjuma/jbossas5/jbpapp_5_0/JBPAPP_5_0/build/output/jboss-5.0.0.Beta/server/default, name=JBoss AS 5 (default), parent=tool.yyz.redhat.com, version=5.0.0.Beta] org.rhq.core.clientapi.agent.PluginContainerException: Cannot load Resource configuration for [-3] at org.rhq.core.pc.configuration.ConfigurationManager.loadResourceConfiguration(ConfigurationManager.java:174) at org.jboss.on.embedded.manager.pc.PluginContainerResourceManager.getResourceConfiguration(PluginContainerResourceManager.java:235) at org.jboss.on.embedded.ui.configuration.resource.ResourceConfigurationUIBean.resourceConfiguration(ResourceConfigurationUIBean.java:144) I took a look at BaseServerConfig.java (https://svn.jboss.org/repos/jbossas/projects/bootstrap/tags/jboss-bootstrap-1.0.1/src/main/java/org/jboss/bootstrap/BaseServerConfig.java) and it looks like the "partitionName" property is not marked as ViewUse.CONFIGURATION. Could this be causing the error I'm seeing? (I deployed the console to the EAP branch.)	1	1	https://issues.jboss.org/browse/JOPR-137
14720	JOPR-137	Farah, In my comment I said the partition name should be a resource config prop, but it should actually be a trait. I'm sorry about that! As for the exception you're seeing, is there more to the stack trace? It seems like the root cause is missing.	1	1	https://issues.jboss.org/browse/JOPR-137
14721	JOPR-137	Has all AS4 traits plus "Server home dir", "Server Name". rev880	0	0	https://issues.jboss.org/browse/JOPR-137
14722	JOPR-134	Jopr currently doesn't support raw config editing, which is what it would take to allow editing of run.bat/run.sh. However, we could fairly easily add Configuration support for adding/updating/deleting environment variables from run.conf/run.conf.bat. I think 95% of what people would update in the start scripts is env variables (e.g. JAVA_OPTS) anyway. Would you find that useful?	1	0	https://issues.jboss.org/browse/JOPR-134

#	Requirement	Questions	Answered	Successfully Answered	URL
14723	JOPR-108	It looks like there may be more of an issue here than just a missing type in the enum. In the JBAS 5.x source, I see: https://anonsvn.jboss.org/repos/jbossas/branches/Branch_5_x/connector/src/main/org/jboss/resource/metadata/mcf/TxConnectionFactoryDeploymentMetaData.java which defines the @ManagementComponent(type="ConnectionFactory",subtype="XA") componentType. and: https://anonsvn.jboss.org/repos/jbossas/branches/Branch_5_x/connector/src/main/org/jboss/resource/metadata/mcf/NoTxConnectionFactoryDeploymentMetaData.java which defines the @ManagementComponent(type="ConnectionFactory",subtype="NoTx") componentType. But I do -not- see anything which defines the @ManagementComponent(type="ConnectionFactory",subtype="LocalTx") componentType. I see: https://anonsvn.jboss.org/repos/jbossas/branches/Branch_5_x/connector/src/main/org/jboss/resource/metadata/mcf/ManagedConnectionFactoryDeploymentMetaData.java but that appears to just be a base class for the specific CF types. Scott, can you comment?	1	1	https://issues.jboss.org/browse/JOPR-108
14724	JOPR-78	The new RTfilter does its best to determine the contextRoot - get it from the container if servlet spec 2.5 else get it from jboss-web.xml if present else (still to be implemented) get it from application.xml else fall back to parts of the webapp name (which will change if the user calles the app foo-1.war, foo-2.war etc).	0	0	https://issues.jboss.org/browse/JOPR-78
14725	JOPR-78	I'm not sure its the responsibility of the RTFilter to be doing this. Shouldnt this be done as part of the webapp discovery?	0	0	https://issues.jboss.org/browse/JOPR-78
14726	JOPR-73	So how did this get implemented?	1	0	https://issues.jboss.org/browse/JOPR-73
14727	JOPR-73	(5:12:12 AM) fdrabek: about JOPR-73 Add support for changing arguments passed to start/stop operations I thought that it is allready possible to change the arguments in scripts (5:13:00 AM) fdrabek: Was that request about it or did you mean something else? (5:18:52 AM) ccrouch: ok let me describe what i was thinking about when i raised this (5:21:43 AM) ccrouch: <but first off I should say this is probably a lower priority fix that we shouldnt be looking at right now :-)> (5:22:47 AM) fdrabek: oki. I will contact you when I will finish the rest of my tasks .... (5:23:47 AM) ccrouch: so if the issue with the start and stop operations is that they are a different sort of script based operation compared to the operations on a JBAS Script resource type (5:24:42 AM) ccrouch: for start and stop all you can do is to change in the config properties for the JBAS instance where the scripts reside, i.e. their path (5:25:00 AM) ccrouch: you can't change the arguments that we actually pass to those scripts (5:28:04 AM) fdrabek: you can change the Command Line Arguments and enviroment variable ... (5:30:44 AM) ccrouch: where? (5:37:09 AM) ccrouch: i can see that you alter the arguments to a regular script, but thats not how the stop/start JBoss ops work (5:48:40 AM) ccrouch: for instance the command line arguments passed to the start script are hardcoded inside of JBossASServerOperationsDelegate.start() (5:53:17 AM) fdrabek: <operation name="execute" displayName="Execute Script" (5:53:17 AM) fdrabek: description="execute the script; NOTE: environment variables to be (5:53:17 AM) fdrabek: passed to the script can be configured via this Script (5:53:17 AM) fdrabek: service's connection properties (under its Inventory tab)"> (5:53:17 AM) fdrabek: <parameters> (5:53:18 AM) fdrabek: <c:simple-property name="commandLineArguments" type="longString" required="false" (5:53:18 AM) fdrabek: description="the command-line arguments (if any) to pass to the script; each (5:53:19 AM) fdrabek: command-line argument must be on a new line; (5:53:19 AM) fdrabek: the arguments can contain properties (5:53:20 AM) fdrabek: with the syntax %propertyName%; the script plugin will interpolate these with (5:53:20 AM) fdrabek: the current values of the corresponding properties from the (5:53:21 AM) fdrabek: script's parent JBossAS server's connection properties"/> (5:53:21 AM) fdrabek: </parameters> (5:53:22 AM) fdrabek: <results> (5:53:22 AM) fdrabek: <c:simple-property name="exitCode" type="integer" description="the script's exit code (0 indicates success, any other value indicates failure)"/> (5:53:22 AM) fdrabek: <c:simple-property name="output" type="longString" description="the script's output (stdout and stderr interleaved)"/> (5:53:23 AM) fdrabek: </results> (5:53:24 AM) fdrabek: </operation> (5:53:38 AM) fdrabek: the parameters is taken from c:simple-property name="commandLineArguments" (5:54:12 AM) ccrouch: right, and all of that is completed ignored for the start/stop operations :-) (5:54:18 AM) ccrouch: completely (5:54:26 AM) ccrouch: (5:48:40 AM) ccrouch: for instance the command line arguments passed to the start script are hardcoded inside of JBossASServerOperationsDelegate.start() (5:58:00 AM) fdrabek: bingo (5:58:29 AM) fdrabek: thanks for your time charles :) (5:59:21 AM) ccrouch: so what JOPR-73 was about was externalising those start/stop parameters so that they could be changed based on the scripts being called, e.g. some scripts may want to use something other than -c or -b to pass in the config and bind address params	1	0	https://issues.jboss.org/browse/JOPR-73
14728	JOPR-39	Are you really sure you are testing with the latest version of the JBAS plugin? after rev 384 ?	1	1	https://issues.jboss.org/browse/JOPR-39
14729	JSFUNIT-296	This issue ties into the ideas expressed here, ARQ-578. Basically moving JSFUnit 'out' of the container, and instead have 'assertion' objects piggy back on 'any' http request that get executed incontainer on e.g. AfterRender	1	0	https://issues.jboss.org/browse/JSFUNIT-296
14730	JSFUNIT-273	Looks nice. Will non-CDI deployments be able to use method injection for JSFSession, JSFClientSession, and JSFServerSession? public class MyTest { // no instance vars @Test @InitialPage("/index.faces") public myTest(JSFSession jsfSession, JSFClientSession client, JSFServerSession server) {} }	1	1	https://issues.jboss.org/browse/JSFUNIT-273
14731	JSFUNIT-244	I'm updating the fix version for this one to 1.4.0.Final. It's probably doable. Is it something you still need? Would you like to contribute a patch to RichFacesClient?	0	0	https://issues.jboss.org/browse/JSFUNIT-244
14732	JSFUNIT-225	Hi Sanket, JSFUnit/HtmlUnit should behave the same way as a browser. So if your application redirects, then JSFUnit/HtmlUnit should do the same. Exactly how is the redirect initiated? Do you have an example where this is not working? Stan	1	1	https://issues.jboss.org/browse/JSFUNIT-225
14733	JSFUNIT-203	I done some profiling of my own this week. It looks like request/response handling is not the bottleneck at all. It always comes down to HtmlUnit and Rhino. I profiled a non-AJAX application. In that case, the bottleneck is in scanning the HTML page. So the problem is in DOM processing. A secondary bottleneck is (surprisingly), initializing the javascript engine. If you aren't using javascript then telling HtmlUnit that you aren't using javascript should speed things up. But how realistic is that? Even one of the standard JSF components, <h:commandLink>, uses javascript. When I profiled an AJAX application, I found that javascript processing was the bottleneck. This means that to speed things up, you have to speed up Rhino, which HtmlUnit is using for javascript processing. So it looks to me like anything we do to speed up JSFUnit will involve helping the HtmlUnit team to come up with optimizations of their code. I've attached my JProfiler snapshot files.	1	1	https://issues.jboss.org/browse/JSFUNIT-203
14734	JSFUNIT-194	Fixed. The latest snapshot of Mojarra 2.0 has an error page that is almost exactly like the Facelets page. So I can easily check for either error page.	0	0	https://issues.jboss.org/browse/JSFUNIT-194
14735	JSFUNIT-140	Hi Maly, Would you be willing to help out with this and other static analysis issues?	0	0	https://issues.jboss.org/browse/JSFUNIT-140

#	Requirement	Questions	Answered	Successfully Answered	URL
14736	JSFUNIT-96	Thanks Brian! Normally, I try to collect the whole component tree so that JSFUnit tests can assert server-side state on unrendered components. But as you discovered, that's rather impractical with a large UIData. I just committed a fix that does this: int rowsToDisplay = component.getRows(); if (rowsToDisplay == 0) rowsToDisplay = component.getRowCount(); for (int i=0; i < rowsToDisplay; i++) if setRows() was never called on the component the getRows() will return zero, which means "display everything". So I have to use both methods. Can you give that a try and let me know if it solves your problem?	1	1	https://issues.jboss.org/browse/JSFUNIT-96
14737	JSFUNIT-47	Have you tried the third option? That one may work. Please let me know if it does. I suspect that we have the same issue as with Seam conversation scope. In Seam, since the scope instance is referenced with a ThreadLocal, we need to move that reference to the test thread at the end of the HTTP request. We do this currently with FacesContext so that FacesContext.getCurrentInstance() will work. We will look into this for both Seam and Trinidad for Beta 2.	1	1	https://issues.jboss.org/browse/JSFUNIT-47
14738	JSFUNIT-25	Shouldn't this issue be moved to StaticAnalysis?	0	0	https://issues.jboss.org/browse/JSFUNIT-25
14739	JSFUNIT-2	Is there any plans to continue looking into this effort? or are the current ant tasks sufficient? Also, what is the current status with renaming it to "jsfunity"	1	0	https://issues.jboss.org/browse/JSFUNIT-2
14740	KEYCLOAK-5069	Shouldn't we also add test for this? [~sebastienblanc] What do you think?	1	1	https://issues.jboss.org/browse/KEYCLOAK-5069
14741	KEYCLOAK-5018	KeycloakConfigResolver let's you do whatever you want. keycloak.json already supports sys props and env variables. So what's the problem?	0	0	https://issues.jboss.org/browse/KEYCLOAK-5018
14742	KEYCLOAK-5004	This is not a bug, nor is it security sensitive.	0	0	https://issues.jboss.org/browse/KEYCLOAK-5004
14743	KEYCLOAK-4989	I'm assuming this is when passwords are updated through Keycloak. Is it not the responsibility of the LDAP server to hash the passwords when the password is being updated?	1	0	https://issues.jboss.org/browse/KEYCLOAK-4989
14744	KEYCLOAK-4989	Can we change this to a feature request and update the summary to be more specific? I don't think it needs to be security sensitive either.	1	0	https://issues.jboss.org/browse/KEYCLOAK-4989
14745	KEYCLOAK-4986	[~vmuzikar] I don't think this is a bug right? Is this by any change breaking something? Unless this is blocking the community version for some reason, I don't see the point of having two versions of package.json	1	1	https://issues.jboss.org/browse/KEYCLOAK-4986
14746	KEYCLOAK-4982	[~stianst] I suddenly understood why you need all the modules, it's for your set-version.sh script, correct ? So what about creating a profile `bump-version` that contains all the modules and that you activate when you run the set-version script ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4982
14747	KEYCLOAK-4971	Hey [~stianst], I've played a bit with export/import of client and I've found out I'm not quite sure about some stuff. When I export a client, lets say account to account.json. Then remove account client from list of clients. Now account management is not available, http://localhost:8080/auth/realms/master/account shows empty page and 'Admin' -> 'Manage account' navigation returns me back to admin console. And then I create a new client using import from the file. But now http://localhost:8080/auth/realms/master/account shows "We're sorry ... No access" and in log there is {noformat} WARN [org.jboss.resteasy.resteasy_jaxrs.i18n] (default task-80) RESTEASY002142: Multiple resource methods match request "GET /realms/master/account". Selecting one. Matching methods: [public javax.ws.rs.core.Response org.keycloak.services.resources.AccountService.accountPageJson(), public javax.ws.rs.core.Response org.keycloak.services.resources.AccountService.accountPage()] {noformat} Q1: Should be really possible to delete 'account' client and also 'security-admin-console'? (when you delete security-admin-console admin console shows just Not found) Q2: In case it should be possible to delete 'account' . Shouldn't account management work when you re-create account client from the file? Q3: In account.json there are also IDs. like "id": "ddb3b034-0fff-4ac5-9aee-0cc445ef5949" and I've tried following scenario: 1. export account to account.json 2. create new client using import from account.json -> error message "Client account already exists" -> OK 3. change the name of the client in input field clientId to account1 -> "Client account1 already exists" -> not OK, client account1 doesn't exists, the problem is that protocolMapper of the client already exists. So change the message? Can you please look at it? Thanks	1	1	https://issues.jboss.org/browse/KEYCLOAK-4971
14748	KEYCLOAK-4946	Would it make sense to do something similar also for auth-utils?	1	0	https://issues.jboss.org/browse/KEYCLOAK-4946
14749	KEYCLOAK-4915	[~stianst] [~pcraveiro] At what point does the UI review need to be done? Do I need to review what is in 7.1 or should I wait until the 7.2 feature is completed?	1	0	https://issues.jboss.org/browse/KEYCLOAK-4915
14750	KEYCLOAK-4914	[~bill.burke] Can you close this issue once you've reviewed all authorization services APIs and are happy with them?	0	0	https://issues.jboss.org/browse/KEYCLOAK-4914
14751	KEYCLOAK-4895	I've got one question though. [~ssilvert], in the PR you are mentioning that you'll discuss the possibilities of downloading the JS dependencies using Maven and not Node.js which then creates the non-relevant node_modules dir structure etc. In the end, are there any plans for this? Or at least to get rid of the node_modules folder? Do we really need to use Node.js (on top of Maven) for it?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4895
14752	KEYCLOAK-4881	Out of curiosity what's the use-case for this?	0	0	https://issues.jboss.org/browse/KEYCLOAK-4881
14753	KEYCLOAK-4874	The fix should be included in 3.1.0.Final. [~bettina_h], are you using this version ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4874
14754	KEYCLOAK-4827	Hey [~stianst] I've tried to create such a test [1] what you suggested. But I cannot reproduce the issue the the test what I wrote. Can you take a look? Maybe I am missing something. I'm testing it on 2.5.5.Final. The test adds several clients, all of them have the same redirect uri pointing to the existing app. Then it performs login request and then accesses different clients in multiple threads and checks if the user is logged in. Thank you. [1] https://github.com/vramik/keycloak/commit/d1bb1e65312e4a09748083880207f739a5fc9f50	1	1	https://issues.jboss.org/browse/KEYCLOAK-4827
14755	KEYCLOAK-4727	Hey [~bill.burke], what do you plan to do with https://github.com/keycloak/keycloak/blob/master/testsuite/integration-arquillian/tests/base/src/test/java/org/keycloak/testsuite/admin/FineGrainAdminLocalTest.java#L58?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4727
14756	KEYCLOAK-4727	[~vramik] What do you mean? These tests are for some new features that I'm working on.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4727
14757	KEYCLOAK-4666	Shouldn't only the social provider that was used to authenticate be logged-out?	0	0	https://issues.jboss.org/browse/KEYCLOAK-4666
14758	KEYCLOAK-4644	Wouldn't this be covered by KEYCLOAK-4184? Basically a user profile should only be updated through identity broker if: * The user was initially created as login from that specific provider * The user has not been updated through the account management console or admin console	1	1	https://issues.jboss.org/browse/KEYCLOAK-4644
14759	KEYCLOAK-4627	Hi [~hmlnarik]. This sounds really good. One question though: apart from being able to configure these new action tokens separately what we are also looking for is a way to persist them. In our flow these tokens will be very long lived (as in weeks) and we want very much for them to survive Keycloak restarts/redeploys because we do those quite often (because we have integrated the build of Keycloak in our automated delivery pipeline). Are these tokens persisted in the Keycloak database or if not can this be done (by us) without too much effort?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4627

#	Requirement	Questions	Answered	Successfully Answered	URL
14760	KEYCLOAK-4627	Hi [~hmlnarik], if they survive Keycloak restart, that is good enough for us. :-) Just for my understanding: how can they survive restarts if they are not persisted? This seems to indicate that they are not very unique and not very secure then? I would expect that in order to create (and validate) them you use something (pseudo)random and something with the timestamp?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4627
14761	KEYCLOAK-4626	Can we have the same ID for authentication session and user session. There should only ever be one authentication session per user session. For example if user logs in to app1, but abonds the login on the OTP page, then tries to login to app2 later the user should be prompted with the first step in the authentication flow and eventually be logged-in to app2.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4626
14762	KEYCLOAK-4580	I guess we have a use case which fits this feature request, but I'm not sure. Our scenario is: We have a central OIDC Provider (not Keycloak-based) which creates access and identity tokens. We want to use the identity token with the [JWT Bearer Flow|http://example.com] to retrieve Keycloak tokens from Keycloak. The Keycloak tokens are than used to access Keycloak secured services. There is also a [Token Exchange Spec |https://tools.ietf.org/html/draft-ietf-oauth-token-exchange-08] (currently draft), which in my opinion offers much more possibilities and should also be able to satisfy the above use case. This feature request seems to be more tailored for the Token Exchange Spec. Am I right with my assumptions? P.S. I guess this feature request would also be the correct way to implement this issue https://issues.jboss.org/browse/KEYCLOAK-910.	0	0	https://issues.jboss.org/browse/KEYCLOAK-4580
14763	KEYCLOAK-4560	Is there an official workaround? Deployments with downtime are essentially a deal breaker when using KC for end user authentication and you have thousands of active users 24/7... :(	0	0	https://issues.jboss.org/browse/KEYCLOAK-4560
14764	KEYCLOAK-4517	I'm currently working on this issue. Unfortunately I'm randomly getting this error for some tests: {code:java} com.gargoylesoftware.htmlunit.FailingHttpStatusCodeException: Too much redirect for http://localhost:8180/offline-client {code} Do you have any idea how to solve this? WIP can be found here: https://github.com/daklassen/keycloak/commit/18814b13fc1237703045cddd954dccf2cf7800d1	1	0	https://issues.jboss.org/browse/KEYCLOAK-4517
14765	KEYCLOAK-4517	In those cases, I usually tries to simulate issue manually and debug the cause. What is helpful is, that I temporarily add this snippet in the beginning of test method: {code} @Test public void yourTest() { Thread.sleep(1000000000); // Here is rest of the test, which you can simulate manually through your browser } {code} and manually simulate the actions done in test. Another suggestion might be to try run the test with different browser to isolate if it's htmlUnit specific issue. For example you can run test with this property: {code} -Dbrowser=phantomjs {code} and test will use PhantomJSDriver instead of default htmlUnit.	0	0	https://issues.jboss.org/browse/KEYCLOAK-4517
14766	KEYCLOAK-4512	Is something like this (http://lists.jboss.org/pipermail/keycloak-dev/2017-February/008891.html) meant here? I made a custom authenticator which does the following: Users are logging in via SAML IdP initiated SSO request, and if their realm has configured other IdPs, we are showing a page where they can log in to the other IdPs so in the end the Keycloak account is linked to the two external IdPs. We've managed to do this with custom first/post broker login authentication flows, with a custom authenticator which shows a .ftl with the buttons you would normally see on the login screen if you have external IdPs configured. Basicly the authenticator does this; {code} public void authenticate(final AuthenticationFlowContext context) { final List<IdentityProviderModel> unlinkedIdps = getUnlinkedIdps(context); if (!unlinkedIdps.isEmpty()) { LOG.debug("unlinkedIdps not empty, show IDP links page"); final LoginFormsProvider loginFormsProvider = context.form(); loginFormsProvider.setAttribute("unlinkedIdps", unlinkedIdps); loginFormsProvider.setInfo("unlinkedIdpsMessage", getLinkedIdpsAsString(context)); context.challenge(loginFormsProvider.createForm("show-idp-links.ftl")); } else { LOG.debug("unlinkedIdps empty, no need to show IDP links page"); context.success(); } } ... // logic for determining unlinkedIdps {code} the ftl: {code} ... <#list social.providers as p> <#list unlinkedIdps as unlinkedIdp> <#if unlinkedIdp.alias == p.alias> <span class="text">${p.displayName} </#list> ... {code}	0	0	https://issues.jboss.org/browse/KEYCLOAK-4512
14767	KEYCLOAK-4501	edit: I've found https://github.com/keycloak/keycloak/commit/b4f625e1ce572bb18cefb7535b8b26d99b7caaeb so it seems the tests are there. What remains is to solve KEYCLOAK-4534	0	0	https://issues.jboss.org/browse/KEYCLOAK-4501
14768	KEYCLOAK-4444	[~vmuzikar] looks like this feature request is a duplicate with https://issues.jboss.org/browse/KEYCLOAK-4445. Could you please take at the PR attached (https://github.com/keycloak/keycloak/pull/4017) and see if that's enough for you?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4444
14769	KEYCLOAK-4437	There's a fairly good chance KEYCLOAK-4338 has fixed this issue. Do you know how to reproduce this issue?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4437
14770	KEYCLOAK-4437	Is there any updates of this issue? I had the same problem here. So, I'd like to know if this issue still occurs in the newer version of keycloak.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4437
14771	KEYCLOAK-4437	[~stianst], thanks for reply! Sure! I tested in 3.0.0 version, so I'll update my app. Thanks a lot!	1	1	https://issues.jboss.org/browse/KEYCLOAK-4437
14772	KEYCLOAK-4435	We actually do have tests in Node.js https://github.com/keycloak/keycloak-nodejs-connect/tree/master/test. But I'm 110% sure that they must be improved, maybe we can get what you have here https://github.com/keycloak/keycloak/tree/master/testsuite/integration-arquillian/tests/other/nodejs_adapter and migrate to Node.js? Also the usage of Node.js will be pretty much as a backend. I'm not against Webdriver, but do you see a scenario where it should be used?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4435
14773	KEYCLOAK-4429	I don't disagree that often it's useful to have the account created locally. There are sceario's though where you want to rely solely on the information from an external IdP. Our use case is we have users (e.g. userA) who want to access our system using an external IdP. They login using the external IdP and are created an account in our keycloak db. Our keycloak relam also has a read-only federation with LDAP for our corporate users. At some point later on, userA is provisioned an account in our LDAP (corpA). Now they have two accounts in our keycloak realm. Ideally, userA was never created and instead they can now easily link corpA with their external IdP login.	1	0	https://issues.jboss.org/browse/KEYCLOAK-4429
14774	KEYCLOAK-4429	[~stianst] Yes. Linking the two accounts, internal and ldap federated, would probably work. How would / could this be done? My initial feeling is that linking the two might actually complicate the situation rather than making it simpler.	0	0	https://issues.jboss.org/browse/KEYCLOAK-4429
14775	KEYCLOAK-4412	By the way why would you need multiple directories even with full server? Just point it to "themes/src/main/resources/theme" all the sources are there.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4412
14776	KEYCLOAK-4412	Actually, let's take a step back. Is what you're trying to do is to run the console tests with quick changes to the themes? Or are you trying to start a full server and allow quickly changing the themes?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4412
14777	KEYCLOAK-4412	First, I haven't been able to get these tests running with Undertow. Any tips? There is no documentation for the console tests. But I thought we were trying to get away from using Undertow anyway? Second, how would I deal with tests that use custom themes? This requires the ability to change themes in two directories without rebuilding everything.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4412

#	Requirement	Questions	Answered	Successfully Answered	URL
14778	KEYCLOAK-4412	We are getting away from the old testsuite, but the new testsuite supports different containers (embedded Undertow, full server dist and server overlay on wf/eap). Undertow is the default container so just running the console tests from the IDE should do the trick. If that doesn't work there's a bug in there. Not sure about the custom themes. That's a very specific set of tests though, so is that really a big deal? In either case that should be sorted out in embedded server as that's what should be used for quick testing during development. It should be done specifically for the embedded KeycloakServer, not by adding a undocumented/untested feature in the server dist. Could be done with a custom theme loader or something like that.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4412
14779	KEYCLOAK-4379	I think this is an interesting topic. When will this be available or at least if we get time frame would be nice.	1	0	https://issues.jboss.org/browse/KEYCLOAK-4379
14780	KEYCLOAK-4376	I wasn't sure if I should open an issue or just leave a comment here since this seems very related. We recently have run into this issue with building a custom account theme and using multi-valued attributes. Since the AccountBean only returns the first element in the multi valued array its is impossible to display these in the template: [AccountBean.java|https://github.com/keycloak/keycloak/blob/2.5.1.Final/services/src/main/java/org/keycloak/forms/account/freemarker/model/AccountBean.java#L49] We have figured out how to post these multi values attributes through the form but cannot display them to the user after they leave the page and return. I see a TODO in the code but as a short term work around could the AccountBean expose these properties through a delimited string {code}attribute1=aaaa##bbbb##cccc{code} similiar to the admin interface? Or do you have a better idea, maybe the solution to this issue will fix the problem. Can the attributes that get injected into freemarker template just be an array?	1	0	https://issues.jboss.org/browse/KEYCLOAK-4376
14781	KEYCLOAK-4365	Are you using OpenAM SP (client libraries) with Keycloak as the IdP? If so can OpenAM SP be configured to either not use PasswordProtectedTransport or to accept unspecified? It will help us decide on priorities for this issue.	1	0	https://issues.jboss.org/browse/KEYCLOAK-4365
14782	KEYCLOAK-4333	Can you explain why it's a problem to change the password through LDAP when Kerberos is used?	0	0	https://issues.jboss.org/browse/KEYCLOAK-4333
14783	KEYCLOAK-4333	Hello, In my case, what I want to do is just to allow to change a principal Kerberos password through KeyCloak. This means that the user is taken from the LDAP but the authentication is done by Kerberos. So the LDAP is only used to verify that the principal exists. Keycloak would have a open feature to verify if a principal exists and change the principal kerberos password. I would not mind to start to work on it on my free time, but I never add an add-on nor know how to add a new feature in Keycloak. Could you send me some info privately about that? Cheers	0	0	https://issues.jboss.org/browse/KEYCLOAK-4333
14784	KEYCLOAK-4333	[~mposolda] can you comment on this issue please? Perhaps give some hints on how it could be implemented if it makes sense to add?	1	0	https://issues.jboss.org/browse/KEYCLOAK-4333
14785	KEYCLOAK-4271	[~mposolda] can you pls review Migration test for offline tokens for manual migration via applying sql script? There is link to log where you can see test results (for db2 in this case): https://keycloak-jenkins.rhev-ci-vms.eng.rdu2.redhat.com/view/vramik/job/keycloak-jpa-migration-manual/DB=db2-105,jdk=mw_oraclejdk8,label_exp=rhel,migrated_version=test-70-migration/123/testReport/ Auto migration works fine.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4271
14786	KEYCLOAK-4246	I do not want to rewrite/refactor the LDAP provider, nor do I ever want to take on this task. :) We only refactor things if we want to expose the LDAPMapperProvider as an SPI, IMO.	1	0	https://issues.jboss.org/browse/KEYCLOAK-4246
14787	KEYCLOAK-4204	I have only added messages for English for the GUI, where can I get help to get this translated?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4204
14788	KEYCLOAK-4204	We don't maintain the non-English translations and expect someone from the community to translate it at some point. It'll fallback to English in either case.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4204
14789	KEYCLOAK-4204	[~stianst], I don't want to be a bother but is there a possibility someone can take a new look at the PR I made three months ago? I have fixed a few merge conflicts but the PR haven't been merged still - is there a reason why you guys haven't merged it? It would be nice if someone could shed some light on this so that we are not waiting in vain, because this is really a critical feature for our use case. :)	1	1	https://issues.jboss.org/browse/KEYCLOAK-4204
14790	KEYCLOAK-4136	Based on the existing code in the ProtectedResource i'm currently using this code to call the update endpoint: {code:java} try { String token = this.authzClient.obtainAccessToken().getToken(); URI uri = new URI("http://127.0.0.1:8080/auth/realms/photoz/authz/protection/resource_set/" + resource.getId()); final RequestBuilder requestBuilder = RequestBuilder.put(uri); final HttpMethod httpMethod = new HttpMethod(this.configuration, requestBuilder); Object response = httpMethod .authorizationBearer(token) .json(JsonSerialization.writeValueAsBytes(resource)) .response().json(Object.class).execute(); LOG.info("Update Resource Response of object {}", response); } catch (IOException | URISyntaxException e) { throw new RuntimeException("Could not update the given resource", e); } {code} The request contains the content-type: application/json as you stated. But still no succes. As requested I checked the PUT request in the admin console, but the requests differ: Admin console: {code} http://127.0.0.1:8080/auth/admin/realms/photoz/clients/afdefd3f-d121-425a-80a6-77267a75d72b/authz/resource-server/resource/7fa1bf65-8f41-414d-9194-2a7313b954ed {code} AuthzClient: {code} http://127.0.0.1:8080/auth/realms/photoz/authz/protection/resource_set/7fa1bf65-8f41-414d-9194-2a7313b954ed {code} Are they using the same endpoint underwater? I attempted to alter my URI to the one used in the admin console, but that results in a Forbidden result. So that did not do the trick..	1	1	https://issues.jboss.org/browse/KEYCLOAK-4136
14791	KEYCLOAK-4109	Sounds great, is there a list of all feature names somewhere? Currently one needs to look at the source: https://github.com/keycloak/keycloak/pull/3676/files#diff-a8e35fe551f4631ce385ec4a00a14d6aR38	1	0	https://issues.jboss.org/browse/KEYCLOAK-4109
14792	KEYCLOAK-4098	The tests for protocol mappers for 2.3.0 looks fine. I think that for 2.5.0 we can maybe have a test for https://github.com/keycloak/keycloak/pull/3630 (Support for duplicate emails). Assuming we will merge this PR. After migration from the 2.4.0, the switch "loginWithEmailAllowed" should be enabled and "duplicateEmailsAllowed" should be disabled.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4098
14793	KEYCLOAK-4098	@Vlastislav Ramik everything should be added now. Can you double check and close once you're happy?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4098
14794	KEYCLOAK-4098	Are offline tokens not exported as part of the export?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4098
14795	KEYCLOAK-4098	[~mposolda] Currently it is not possible to "Login and request offline token" when 1.9.8 starts (before stop the 1.9.8). Start 1.9.8, stop 1.9.8 and migration is done before the tests are run.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4098
14796	KEYCLOAK-4087	Yeah, ok. So just to clarify, you have something like: uid: john uidname: john uidnumber: 123 And on group you have: memberUID: 123 Is it correct? If yes, I will likely go with what you propose (Have both attributes configurable somehow). Fact is, that we can't just change it like in your PR as for example in the case of some other deployments, I am aware, there is something like: user uid: john entryUUID: 123456 group memberUID: john	1	1	https://issues.jboss.org/browse/KEYCLOAK-4087

#	Requirement	Questions	Answered	Successfully Answered	URL
14797	KEYCLOAK-4052	IMO this is not something specific to LDAP. Instead of having flag on LDAP storage provider, shouldn't we rather enforce Keycloak password policies BEFORE the UserStorage is invoked? The call to UserStorage.updatePassword will happen just if keycloak policies matches?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4052
14798	KEYCLOAK-4052	+1 Seems like this would be best handled outside each individual provider. Q should it be optional?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4052
14799	KEYCLOAK-4018	[~jlpettersson], can you provide some tests for this new policy ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4018
14800	KEYCLOAK-4018	[~jlpettersson], I think I found an issue. The attribute "kc.client.id" maps to the azp (issueFor in AccessToken) and Keycloak puts there the client_id (e.g.: my-client) and not the client identifier. We need to change the {{ClientPolicyProvider}} to obtain a client based on its client_id and then do the check. Do you want me to change that or do you want to change by your self ? Thanks.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4018
14801	KEYCLOAK-4018	I am not sure, what from {{Evaluation evaluation}} I should use for lookup. My previous code worked with {{evaluation.getContext().getIdentity().getId()}}. But how am I sure that the token is from keycloak (aud). Do you think that I should use *azp* anyway, and lookup that client in some way? I am not sure what or how I should lookup here now :) :(	1	1	https://issues.jboss.org/browse/KEYCLOAK-4018
14802	KEYCLOAK-4018	Right. So we are sure that {{kc.client.id}} / azp is unique? Or can anyone create a client with the same name? And we are sure that the accessToken is from keycloak, or can it be from any other ID-server, just with the same azp-name (that is mapped to kc.client.id? If so, it looks good now.	1	1	https://issues.jboss.org/browse/KEYCLOAK-4018
14803	KEYCLOAK-4009	I'm not sure we want to support AD LDS. [~mposolda] what do you think about this PR?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4009
14804	KEYCLOAK-4009	[~stianst] I propose to merge the simplified version of this PR, which will include just the mapper. [~sldab] I propose that in your PR you include just the code of your new MSADLDSUserAccountControlStorageMapper, but NOT any other changes. Especially NOT include: - adding new vendor to the list of supported vendors - adding new LDAPConfig.isActiveDirectoryLds() method (and all related changes) If we include just the new mapper, then people, who want MSAD LDS will just: - Select the existing vendor "Active Directory" and save their federation provider - Then they will need to go to "Mappers" tab and manually remove the existing "MSAD mapper" and instead add your "MSAD LDS Mapper" . This should work fine if I understand correctly, right? The reason is, that we rather prefer smaller list of OOTB supported vendors, but the flexibility to re-configure things manually if people want to use their preferred vendor. WDYT?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4009
14805	KEYCLOAK-4009	Hi [~mposolda], I see your point in reducing vendor-specific logic and I am fine with your proposal. There's just one thing to consider: If we only include the mapper and one would select "Active Directory" as the vendor, the password update would brake. It would try to modify the {{unicodePwd}} attribute instead of {{userPassword}}. If one would select "Other" as vendor, then UUID handling would brake or one would need to resign from having an UUID. Would it be OK for you if we assume, that the {{objectGUID}} attribute is MS-specific and replace the {{isActiveDirectory()}} check with {{"objectGUID".equalsIgnoreCase(config.getUuidLDAPAttributeName())}} in {{LDAPOperationsManager.decodeEntryUUID(..)}} and {{LDAPOperationsManager.getFilterById}}?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4009
14806	KEYCLOAK-4009	Updated PR. Added {{isObjectGUID()}} method to {{LDAPConfig}}. I noticed that the {{processFailedPasswordUpdateException(..)}} method in the existing {{MSADUserAccountControlStorageMapper}} isn't used anywhere. Ist this still work in progress?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4009
14807	KEYCLOAK-4009	I've merged your PR. Thanks. Could you please summary here what are needed steps in configuration to have Keycloak working with MS LDS?	1	1	https://issues.jboss.org/browse/KEYCLOAK-4009
14808	KEYCLOAK-3965	[~bill.burke] is this something which I could help? If yes, could you please elaborate more about "no-import"?	1	0	https://issues.jboss.org/browse/KEYCLOAK-3965
14809	KEYCLOAK-3877	You've lost me.. WF/EAP console is served on management port (9990) right? So I'd expect dynamic keycloak.json to be served there as well (http://localhost:9990/console/keycloak.json). Then the JS in WF/EAP console configures Keycloak JS with this link to make it use the correct auth-server-url, realm and client-id. Correct? I can't see any of that being the responsibility of our subsystem. Then there's a separate issue of making config available for WARs that are secured with a secure-deployment so that they can configure Keycloak JS in a similar way. That's a completely separate issue IMO and the two are not linked at all.	1	0	https://issues.jboss.org/browse/KEYCLOAK-3877
14810	KEYCLOAK-3747	[~zschwarz] the PR is clearly only about the testsuite, so why did you remove the component I added?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3747
14811	KEYCLOAK-3731	Asked about this one via mailing list as well, but [~chris.brandhorst] pointed me in the right direction here. We have a significant need for this one. Is it currently being worked on? If not, I might pick it up as we have a pressing need for it.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3731
14812	KEYCLOAK-3731	Thanks for the work! We retried the configuration with two 2.5.0-Final instances, but sadly we get the same result (StaleCodeMessage from instance 'B'). From the PR it seems that the endpoint I mentioned should still be used, is this correct? Thanks for any pointers.	0	0	https://issues.jboss.org/browse/KEYCLOAK-3731
14813	KEYCLOAK-3712	I just hit this limitation. My plan was to dynamically register a client and its resources. But I dont see any way to register resources without first manually enabling "Authorization" in the admin ui. Is there a workaround? Are there any plans to fix these limitations in the near future?	1	0	https://issues.jboss.org/browse/KEYCLOAK-3712
14814	KEYCLOAK-3678	[~sonicaaaa] I've merged your PR. Thanks. Would you have a chance to update also our documentation for fuse adapter? It's here https://github.com/keycloak/securing_apps_guide	1	1	https://issues.jboss.org/browse/KEYCLOAK-3678
14815	KEYCLOAK-3672	Do you know if this is still an issue with 3.x?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3672
14816	KEYCLOAK-3661	[~praseodym] If you specify "Search scope" as "subtree" in the configuration of LDAP federation provider, it should find also groups under sub-OU . For example if you configure Group LDAP mapper with the "LDAP Groups DN" like "ou=groups,dc=example,dc=com", it will find also group like "cn=group1,ou=subou,ou=groups,dc=example,dc=com" . Is it working for you?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3661
14817	KEYCLOAK-3643	[~sonicaaaa]] I've added another commit with the bit updated README for hawtio https://github.com/keycloak/keycloak/pull/3284 . Especially it's now pointing that Hawtio integration will be available from Fuse 6.3.1 but not with older fuse versions (unless people uninstall the provided hawtio and replace it with newer community version). Is it correct? Can we resolve this JIRA if you are ok with my change in the README?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3643
14818	KEYCLOAK-3631	Hi [~stianst], in the meantime, could you think of any workaround for the issue we are having? I guess we can increase the Session SSO timeout to a couple of days. From a security and performance point of view of course not ideal but we could probably live with it for a little while. That still leaves the issue that if Keycloak is restarted for whatever reason we lose all the client sessions and hence the 'user action tokens' making them invalid. Ideally these tokens would be stored in the Keycloak database so that they survive a restart.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3631
14819	KEYCLOAK-3631	Hi [~hmlnarik]. This sounds really good. One question though: apart from being able to configure these new action tokens separately what we are also looking for is a way to persist them. In our flow these tokens will be very long lived (as in weeks) and we want very much for them to survive Keycloak restarts/redeploys because we do those quite often (because we have integrated the build of Keycloak in our automated delivery pipeline). Are these tokens persisted in the Keycloak database or if not can this be done (by us) without too much effort?	0	0	https://issues.jboss.org/browse/KEYCLOAK-3631
14820	KEYCLOAK-3613	Do we want to move this to 3.x? This isn't just an SPI issue and requires changes to LDAP, SSSD, admin console and account service.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3613

#	Requirement	Questions	Answered	Successfully Answered	URL
14821	KEYCLOAK-3605	Do you have an example for such a custom header? I'm currently working on KEYCLOAK-4163	0	0	https://issues.jboss.org/browse/KEYCLOAK-3605
14822	KEYCLOAK-3592	[~jcain] I'll find someone that can review this, but a few things from my point of view (I've not had a chance to look at it detail though): * Can you squash to a single commit? * Can you use the new profile mechanism (https://keycloak.gitbooks.io/documentation/content/server_installation/topics/profiles.html). By default it should be default, but then enabled through the profile options. * Can you send docs updates to https://github.com/keycloak/keycloak-documentation (we've merged all doc repositories into one) * We need to somehow test this. Ideally automated testing with Docker. Not sure how we'd do that. Also, it seems kinda stupid that Docker won't just allow OpenID Connect providers. Seems like a feature they will eventually have to add. Did you check if there is any progress on that?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14823	KEYCLOAK-3592	We'll need functional tests rather than unit tests for this one. Can we configure the local Docker and just invoke that from a test? If setup needs to be done prior to running the test that's fine as long as it's covered in HOW-TO-RUN.md in the testsuite and also possible to automate in Jenkins.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14824	KEYCLOAK-3592	The problem I've run into isn't invoking a `docker login` from the client - Docker actually has a functional Java adapter for that. The problem is interacting with a docker registry. I don't know of a way to wire that up in a test harness without introducing dependencies on some binaries on the host machine. How much leeway does your current Jenkins/Travis instance give you for system configuration?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14825	KEYCLOAK-3592	[~jcain] Travis wouldn't run this as we're already overloading what we can do there. It would rather have to be part of separate CI that runs everything. I believe we can setup what's needed. Can you explain what that would actually be, then I can ask our QE guys if it's doable?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14826	KEYCLOAK-3592	[~stianst] sure! So the big showstopper for putting this on any kind of CI environment for me has been the requirement for a docker registry to test against. Per above ^ I've tried to make some quickstarts pretty simple out of the box (although I've clearly failed to make them as portable as they should be), but there's likely going to need to be at least some docker binaries/tools on the testing box. Past that, everything is pretty straightforward. Docker's Java client works well for automated testing, and y'all have a really solid setup for spinning up Keycloak instances and realm config. So just some docker + tools (compose perhaps) binaries. What do you have in mind? How would I go about providing you the tests you need in such a scenario?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14827	KEYCLOAK-3592	How do we start the server with the correct profile to be able to select the docker protocol ? bin/standalone.sh|bat -Dkeycloak.profile.feature.<feature name>=disabled	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14828	KEYCLOAK-3592	I can't get it to work... The docker cli is making the request right : {quote}GET /auth/realms/docker/protocol/docker-v2/auth?account=###&client_id=docker&offline_token=true&service=docker-registry HTTP/1.1 User-Agent: docker/17.03.1-ce go/go1.7.5 git-commit/c6d412e kernel/4.4.0-71-generic os/linux arch/amd64 UpstreamClient(Docker-Client/17.03.1-ce \(linux\)) Authorization: Basic ############### Accept-Encoding: gzip Connection: close{quote} but Keycloak answer with the HTML login page... which lead to "Error response from daemon: Get http://registry/v2/: unable to decode token response: invalid character '<' looking for beginning of value" I performed the following actions : * add "docker" realm * add "docker-registry" client * save (after setting * as valid redirect url) I noticed that "BASIC authentication is configured for you realm. Since docker auth requires HTTP Basic auth, this should be the only authenticator configured for the realm hosting the docker registry client." but didn't get it. Can u help so it work with manual setup on clean realm ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14829	KEYCLOAK-3592	Just one question regarding the profile. So, if I understand it correctly, the default settings is keycloak.profile.feature.docker=disabled. So every time you want to use the docker features you need to explicitly enable it. At least, currently it seems to work like that for me. Shouldn't it be enabled by default?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14830	KEYCLOAK-3592	Also, I've found some issues I'd like to discuss. * When creating a new Client, docker-v2 is the default Client Protocol. I believe the default should remain openid-connect as it is obviously much more used. * All built-in Clients in a Realm (account, admin-cli, broker, ...) look like docker-v2 clients in the Admin Console. But of course they are openid-connect. This doesn't concern additionally created Clients. * A subject to discussion: I don't like the idea to have to break the browser flow just to make docker work. I can imagine the scenario when users are stored in Keycloak (directly, without using User Federation) and they'd need to be able to use Keycloak both for docker authentication and a web service authentication at the same time. Additionally, it's a bit complicated to set up docker in Keycloak with all the flows juggling. Would it be a problem to categorize the docker flow as an extra flow, not as browser flow, so it wouldn't be necessary to manipulate with flows at all (docker flow would work out of the box)? EDIT: One more insignificant detail. Is there any specific reason for README.md (which is included in the downloaded YAML ZIP from Keycloak) to have content like {{java.util.zip.ZipFile$ZipFileInflaterInputStream@770c8fa0}}. Shouldn't it be instead a human readable markdown, some instructions?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14831	KEYCLOAK-3592	{quote}* When creating a new Client, docker-v2 is the default Client Protocol. I believe the default should remain openid-connect as it is obviously much more used.{quote} See discussion [here|https://github.com/keycloak/keycloak/pull/3565#discussion_r107416768] {quote}* All built-in Clients in a Realm (account, admin-cli, broker, ...) look like docker-v2 clients in the Admin Console. But of course they are openid-connect. This doesn't concern additionally created Clients.{quote} See discussion [here|https://github.com/keycloak/keycloak/pull/3565#issuecomment-290832697] {quote}A subject to discussion: I don't like the idea to have to break the browser flow just to make docker work ...{quote} That's a really good point. Ideally enabling the docker feature would provide an already configured new flow called docker, so no additional configuration would be required. But then - how does the server know when a specific type of auth flow should kick in (is it browser flow or is it docker flow?). Currently a set of available flow types is hardcoded in realm model so there's no SPI to extend. If 'docker' flow type is added then an if-then should probably be added to wherever the decision logic that triggers browser flow is implemented. And whether to activate 'docker' flow would have to be guided by something present in the request (target uri probably?).	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14832	KEYCLOAK-3592	[~jcain] I see your point. What I meant was really just the thing that after setting up the flows for docker you cannot use the same realm anymore to standard browser authentication as before, e.g. to authenticate the user via browser to access other non-docker resources. Or am I mistaken? Perhaps I'm missing something. Anyway, if you think that this scenario is not really relevant, let's proceed with the tests. Just two more previous questions we haven't discussed yet: * Shouldn't be the docker profile enabled by default? * Content of README.md included in the YAML ZIP	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592

#	Requirement	Questions	Answered	Successfully Answered	URL
14833	KEYCLOAK-3592	As I'm thinking about it, I still feel we need to clear up one thing. If I understand it correctly, to use docker with a realm we need to completely replace the browser flow in such realm. Am I correct? There is no other way of using docker with a realm? So, this leads to breaking up the standard browser login flow, user's can't login to this realm using browser anymore, so that this realm can't be used for much other things than docker only. Do we really want this - the need of separate realm just to use docker? I believe it makes a perfect sense to use one realm for docker and e.g. some other web service. But if we replace the browser flow, we can't use the realm for login using browser. Perhaps, I'm missing something. [~stianst] What do you think? One more thing. We should also mention in the docs that the docker profile is disabled by default in Keycloak (unlike the other profiles).	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14834	KEYCLOAK-3592	[~vmuzikar] I still don't see a problem with this, as the docker client will never use any other flow. There will be a functional separation between what Docker uses and what standard web browsers use within a ream that could be argued as being better expressed in terms of different realms. With that said, I'm up for whatever finally gets this PR accepted. [~stianst] -Docker requires HTTP Basic auth - iirc we're not going to be able to use something like an ootb direct grant flow.- * See comment on PR, I think you're right in that direct grant flow is a better fit here. Would switching to that flow be sufficient?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14835	KEYCLOAK-3592	[~jcain] With the current approach (replacing the browser flow), are users able to login to the same realm using browser login (when authenticating to some other web resource) and docker at the same? Or is it like that after setting up the realm for docker, users won't be able to login using browser? From the docs PR: bq. it is recommended that you configure the docker registry client in a realm other than 'master', since the HTTP Basic auth flow will not present forms. I've done some testing and after replacing the browser flow for Docker Authenticator I can't access following URL {{localhost:8080/auth/realms/{docker-realm}/account}}. I'm getting an error: "Unexpected error when handling authentication request to identity provider."	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14836	KEYCLOAK-3592	instructions (removing the docker binary and installing docker-ce again) didn't do anything. Even installed a fresh Fedora 25 virtual machine and installed fresh docker-ce - still the same results, same exception "Retry limit reached while trying to pull image" Only thing that works for this exception is Red Had docker binary. However, I'm getting another exception using this binary: Failed to get D-Bus connection. {code} 11:45:13,232 INFO [TestEventsLogger] [DockerClientTest] shouldPerformDockerAuthAgainstRegistry() STARTED 11:45:14,661 INFO [KeycloakDronePostSetup] Configuring driver settings. implicitWait=750, pageLoadTimeout=10000 11:45:15,759 DEBUG [AbstractKeycloakTest] updating admin password 11:45:15,760 INFO [URLUtils] starting navigation 11:45:15,763 DEBUG [KeycloakDronePostSetup$HtmlUnitInterceptor] Ignored JS as we are on about:blank page now 11:45:15,764 INFO [URLUtils] current URL: about:blank 11:45:15,764 INFO [URLUtils] navigating to http://localhost:8180/auth 11:45:16,226 INFO [URLUtils] new current URL: http://localhost:8180/auth/ 11:45:16,226 INFO [URLUtils] navigation complete 11:45:16,228 DEBUG [AbstractKeycloakTest] loading test realms 11:45:16,611 INFO [AbstractKeycloakTest] importing test realms 11:45:16,611 DEBUG [AbstractKeycloakTest] importing realm: docker-test-realm 11:45:19,529 INFO [org.testcontainers.dockerclient.DockerClientProviderStrategy] Found docker client settings from environment 11:45:19,565 INFO [org.testcontainers.dockerclient.DockerClientProviderStrategy] Looking for Docker environment. Tried Environment variables, system properties and defaults. Resolved: dockerHost=unix:///var/run/docker.sock apiVersion='{UNKNOWN_VERSION}' registryUrl='https://index.docker.io/v1/' registryUsername='vmuzikar' registryPassword='null' registryEmail='null' dockerConfig='DefaultDockerClientConfig[dockerHost=unix:///var/run/docker.sock,registryUsername=vmuzikar,registryPassword=<null>,registryEmail=<null>,registryUrl=https://index.docker.io/v1/,dockerConfig=/home/vmuzikar/.docker,sslConfig=<null>,apiVersion={UNKNOWN_VERSION}]' 11:45:19,567 INFO [org.testcontainers.DockerClientFactory] Docker host IP address is localhost 11:45:19,881 INFO [org.testcontainers.DockerClientFactory] Connected to docker: Server Version: 1.12.6 API Version: 1.24 Operating System: Fedora 25 (Workstation Edition) Total Memory: 7861 MB 11:45:23,588 INFO [org.testcontainers.DockerClientFactory] Disk utilization in Docker environment is 0% (10179 MB available) 11:45:23,610 INFO [[docker.io/registry:2]] Creating container for image: docker.io/registry:2 11:45:24,508 INFO [[docker.io/registry:2]] Starting container with ID: f0747eb4459e232d3efce11eb3e0a432fc85c8a5bc90ce59ed73785309492c1c 11:45:25,108 INFO [[docker.io/registry:2]] Container docker.io/registry:2 is starting: f0747eb4459e232d3efce11eb3e0a432fc85c8a5bc90ce59ed73785309492c1c 11:45:25,209 INFO [[docker.io/registry:2]] Container docker.io/registry:2 started 11:45:25,398 INFO [[testcontainers/cydmykqpamcrvxu6]] Creating container for image: testcontainers/cydmykqpamcrvxu6 11:45:26,042 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-11T09:45:25Z" level=warning msg="No HTTP secret provided - generated random secret. This may cause problems with uploads if multiple registries are behind a load-balancer. To provide a shared secret, fill in http.secret in the configuration file or set the REGISTRY_HTTP_SECRET environment variable." go.version=go1.7.3 instance.id=c865d45c-aae3-4267-bdca-070bfffccbd3 version=v2.6.1 11:45:26,042 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-11T09:45:25Z" level=info msg="redis not configured" go.version=go1.7.3 instance.id=c865d45c-aae3-4267-bdca-070bfffccbd3 version=v2.6.1 11:45:26,042 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-11T09:45:25Z" level=info msg="Starting upload purge in 42m0s" go.version=go1.7.3 instance.id=c865d45c-aae3-4267-bdca-070bfffccbd3 version=v2.6.1 11:45:26,043 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-11T09:45:26Z" level=info msg="using inmemory blob descriptor cache"	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14837	KEYCLOAK-3592	[~vmuzikar] I think this is the pertienent line for the above ^ test: {quote}11:45:28,060 INFO [DockerClientTest] Cannot connect to the Docker daemon. Is the docker daemon running on this host?{quote} It would appear that your docker daemon isn't started. Can you please: {quote}systemctl start docker{quote} And try the tests again?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14838	KEYCLOAK-3592	[~vmuzikar] - so help me understand here. Are you saying that it is now working with the Red Hat docker fork? I'm trying to help as best I can here, but the above ^ error you gave me is due to the docker daemon not being started. With that resolved, all I could do is guess as to why it wouldn't work - without errors, environment, etc. I'm open to suggestions, but so far this has appeared like the best option.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592

#	Requirement	Questions	Answered	Successfully Answered	URL
14839	KEYCLOAK-3592	08:41:06,741 INFO [TestEventsLogger] [DockerClientTest] shouldPerformDockerAuthAgainstRegistry() STARTED 08:41:07,978 INFO [KeycloakDronePostSetup] Configuring driver settings. implicitWait=750, pageLoadTimeout=10000 08:41:08,852 DEBUG [AbstractKeycloakTest] updating admin password 08:41:08,853 INFO [URLUtils] starting navigation 08:41:08,856 DEBUG [KeycloakDronePostSetup$HtmlUnitInterceptor] Ignored JS as we are on about:blank page now 08:41:08,857 INFO [URLUtils] current URL: about:blank 08:41:08,857 INFO [URLUtils] navigating to http://localhost:8180/auth 08:41:09,324 INFO [URLUtils] new current URL: http://localhost:8180/auth/ 08:41:09,325 INFO [URLUtils] navigation complete 08:41:09,326 DEBUG [AbstractKeycloakTest] loading test realms 08:41:09,787 INFO [AbstractKeycloakTest] importing test realms 08:41:09,787 DEBUG [AbstractKeycloakTest] importing realm: docker-test-realm 08:41:12,332 INFO [org.testcontainers.dockerclient.DockerClientProviderStrategy] Found docker client settings from environment 08:41:12,353 INFO [org.testcontainers.dockerclient.DockerClientProviderStrategy] Looking for Docker environment. Tried Environment variables, system properties and defaults. Resolved: dockerHost=unix:///var/run/docker.sock apiVersion='{UNKNOWN_VERSION}' registryUrl='https://index.docker.io/v1/' registryUsername='vmuzikar' registryPassword='null' registryEmail='null' dockerConfig='DefaultDockerClientConfig[dockerHost=unix:///var/run/docker.sock,registryUsername=vmuzikar,registryPassword=<null>,registryEmail=<null>,registryUrl=https://index.docker.io/v1/,dockerConfig=/home/vmuzikar/.docker,sslConfig=<null>,apiVersion={UNKNOWN_VERSION}]' 08:41:12,356 INFO [org.testcontainers.DockerClientFactory] Docker host IP address is localhost 08:41:12,623 INFO [org.testcontainers.DockerClientFactory] Connected to docker: Server Version: 17.03.1-ce API Version: 1.27 Operating System: Fedora 25 (Workstation Edition) Total Memory: 7861 MB 08:41:17,613 INFO [org.testcontainers.DockerClientFactory] Disk utilization in Docker environment is unknown (unknown available) 08:41:17,630 INFO [[registry:2]] Pulling docker image: registry:2. Please be patient; this may take some time but only needs to be done once. 08:41:26,772 INFO [[registry:2]] Creating container for image: registry:2 08:41:26,870 INFO [[registry:2]] Starting container with ID: 9e3ae6cfaa5ba54be2692d3c96c04129fd57cff1066ed27b851dd92a4006b648 08:41:27,144 INFO [[registry:2]] Container registry:2 is starting: 9e3ae6cfaa5ba54be2692d3c96c04129fd57cff1066ed27b851dd92a4006b648 08:41:27,229 INFO [[registry:2]] Container registry:2 started 08:41:27,235 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-16T06:41:27Z" level=warning msg="No HTTP secret provided - generated random secret. This may cause problems with uploads if multiple registries are behind a load-balancer. To provide a shared secret, fill in http.secret in the configuration file or set the REGISTRY_HTTP_SECRET environment variable." go.version=go1.7.3 instance.id=24b24843-569b-4aa0-8bce-860d621e8f8b version=v2.6.1 08:41:27,236 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-16T06:41:27Z" level=info msg="redis not configured" go.version=go1.7.3 instance.id=24b24843-569b-4aa0-8bce-860d621e8f8b version=v2.6.1 08:41:27,236 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-16T06:41:27Z" level=info msg="Starting upload purge in 11m0s" go.version=go1.7.3 instance.id=24b24843-569b-4aa0-8bce-860d621e8f8b version=v2.6.1 08:41:27,236 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-16T06:41:27Z" level=info msg="using inmemory blob descriptor cache" go.version=go1.7.3 instance.id=24b24843-569b-4aa0-8bce-860d621e8f8b version=v2.6.1 08:41:27,237 INFO [Slf4jLogConsumer] STDOUT: time="2017-05-16T06:41:27Z" level=info msg="listening on [::]:5000, tls" go.version=go1.7.3 instance.id=24b24843-569b-4aa0-8bce-860d621e8f8b version=v2.6.1 08:41:27,237 INFO [Slf4jLogConsumer] STDERR: 2017/05/16 06:41:27 http: TLS handshake error from 172.17.0.1:51178: EOF 08:43:28,650	1	1	https://issues.jboss.org/browse/KEYCLOAK-3592
14840	KEYCLOAK-3576	Is the conclusion now to introduce a new role {{manage-user-federation}} and let the administrators configure it as they like?	0	0	https://issues.jboss.org/browse/KEYCLOAK-3576
14841	KEYCLOAK-3573	Hi [~stianst] and [~dlofthouse], Sent the PR with Elytron SAML and OIDC Adapters. We are using the latest release version of Elytron and related dependencies. The new adapter can be tested by passing a {{elytron.wildfly-version}} system property when running tests using {{app-server-wildfly}}. Something like: {code} mvn clean install -Pauth-server-wildfly -Papp-server-wildfly-elytron -Delytron.wildfly.version=11.0.0.Alpha1-SNAPSHOT -Papp-server-wildfly -Dtest=org.keycloak.testsuite.adapter.WildflyOIDCAdapterTest {code} [~stianst], would like your feedback about distribution of these adapters. My PR also changes adapter disttribution package to include these adapters. Is that OK ? Thanks.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3573
14842	KEYCLOAK-3573	Why is the Elytron adapter being added to the WildFly adapter? See here: https://github.com/keycloak/keycloak/pull/3978/files#diff-c461e2913d8fc274023f25ede83bfdbaR38 That makes no sense to me. Shouldn't this adapter be distributed as a separate WildFly 11 adapter than be co-located with the WildFly 10 / EAP 7.0 adapters? Seems like it could cause issues having both the old Undertow and the new Elytron adapter in the same distribution. Why is there a folder called "elytron" in the root of Keycloak? That makes no sense. With regards to testing we need to test with the latest builds of WF11 and EAP 7.1, not with Elytron installed to older WF versions. The latter is not a scenario we should document or support IMO. Also, did you add any updates to documentation around this?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3573
14843	KEYCLOAK-3573	[~stianst], can you please review changes to documentation ? Are you OK with them ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3573
14844	KEYCLOAK-3519	[~puntogil] thanks, out of curiosity why did you care enough to send this PR? ;)	1	1	https://issues.jboss.org/browse/KEYCLOAK-3519
14845	KEYCLOAK-3505	Yes, the UserAttributeMapper has subclass AbstractClaimMapper, which uses just AccessToken and IDToken. I would go for the approach1 and rather add the missing functionality to AbstractClaimMapper.getClaimValue , so it's able to read stuff from accessToken, idToken or userInfo. Would you mind working on it and send PR? I hope you will be able to reuse some code from AbstractJsonUserAttributeMapper as it has some static helper methods to retrieve the value of claim from JsonNode.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3505
14846	KEYCLOAK-3498	Can you elaborate on why it doesn't work for you?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3498
14847	KEYCLOAK-3498	We're a payments processing company providing internet acquiring services for our clients. Our clients are merchants, owners of internet shops. They integrate our products in their internet shops in order to accept credit card payments online. We're providing powerful capabilities of payment form customization. Our merchants implement html code of credit card payment forms and just include our JS library. This JS lib a "Tokenizer" grabs credit card data from the html document, sends it to our API endpoint, and receive a credit card token. This token afterwards is being sent to merchant's backend which contacts our backend and performs actual credit card charge. In such a way we have two common service types using Keycloak. The first one is usual JS web-application "Merchant Personal office" which uses Keycloak infrastructure for authorisation and authentication/SSO using JWTs. This app uses wide range of scopes described in JWT for accessing different application elements. This app uses our API backend with obtained JWT and renders a view. The second one is that JS credit card data Tokenizer. It also uses our API backend but with very limited scope, actually only one API endpoint: "tokenize this credit card data". So our clients, Merchants, have to be able to obtain Keycloak JWTs in order to have access to their Personal office web-app and in the same time to obtain another JWT from this web-app which they will use for authorization of JS tokenizer. We think that such a functionality is very similar of Github's Personal access tokens (https://github.com/settings/tokens). Could you please clarify whether such business logic could be implemented by existing Keycloak functionality and if not would you accept our pull request implementing it? Thanks!	1	1	https://issues.jboss.org/browse/KEYCLOAK-3498
14848	KEYCLOAK-3485	Great. Love the responsiveness :) When is the next release planned for?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3485
14849	KEYCLOAK-3468	Just out of curiosity, why is it very important to you?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3468

#	Requirement	Questions	Answered	Successfully Answered	URL
14850	KEYCLOAK-3468	A few things that needs to be considered when upgrading: * Test HTTP/2 support * Test and document automatic generation of self-signed certs. This involves documenting steps for clients as well. * Test clustering as Infinispan is upgraded * Review load-balancing profile. Do we drop ours and use the load-balancing profile provided by WF instead?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3468
14851	KEYCLOAK-3468	Even though the core KC server isn't being upgraded to WF 10.1 soon, can I still add the KC OpenID WF adapter to a WF 10.1 install? There is a bug fix I need for my app that is in WF 10.1 so I need to upgrade our app server.	1	0	https://issues.jboss.org/browse/KEYCLOAK-3468
14852	KEYCLOAK-3416	Is this not optional?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3416
14853	KEYCLOAK-3407	What version of JBoss are you actually talking about? We support EAP 6, JBoss AS 7 and WildFly, that's it.	1	0	https://issues.jboss.org/browse/KEYCLOAK-3407
14854	KEYCLOAK-3370	[~stianst] One example of this is the stylistic differences between something like [access.redhat.com|https://access.redhat.com] and the [Openshift Console|https://api.preview.openshift.com/]. Openshift has explicitly asked for the ability to display an Openshift-branded login/registration page to users. A single theme doesn't really seem to suite both of those applications simultaneously. We would certainly include some SSO-specific elements (banner, title text, etc.) to preserve coherence across application boundaries. However, for this particular use case, OpenShift users are more commonly going to interact with a single application than with multiple apps under the SSO umbrella. In light of the usage patterns, does it not make sense to have the option to provide a more consistent experience for those users? And those two are a relatively close example. When you start comparing them with something like [http://www.jboss.org/], it's easy to see why users might want some sort of context on the page that denotes an understanding of what app they're coming from.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3370
14855	KEYCLOAK-3370	+1 This is very important for OpenShift Online's usage. Is there any idea when this could be completed and we could adopt?	0	0	https://issues.jboss.org/browse/KEYCLOAK-3370
14856	KEYCLOAK-3370	Do you actually have that requirement? Just curious ;)	1	1	https://issues.jboss.org/browse/KEYCLOAK-3370
14857	KEYCLOAK-3359	Can you not just change the email after the service account is enabled?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3359
14858	KEYCLOAK-3359	I'm not quite following what you want. Are you simply wanting to have an email attribute somewhere on the client/service account directly? Or do you want to set it at a realm level or something?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3359
14859	KEYCLOAK-3353	I don't think a separate issue was opened to solve this problem. Please correct me if I'm wrong. Shall I just open an issue with "Support for session iframe for confidential clients" as a subject and copy the description or reference this issue? Thanks.	1	0	https://issues.jboss.org/browse/KEYCLOAK-3353
14860	KEYCLOAK-3337	Could you provide an example, please ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3337
14861	KEYCLOAK-3317	Has any progress been made on this subject? We use XACML policies for authorizations in our webapps and webservices. It would be great if we could integrate similar policies in Keycloak.	1	0	https://issues.jboss.org/browse/KEYCLOAK-3317
14862	KEYCLOAK-3316	Hi, About the Implicit Flow : if 'scope=openid' is missing the request is pure OAuth2 so the 'nonce' parameter should not be 'Required' anymore since it does not exist in the OAuth2 specs. The same for 'redirect_uri' which is not mandatory in OAuth2 if already registered at client registration step. What do you think ? -Nicolas	1	1	https://issues.jboss.org/browse/KEYCLOAK-3316
14863	KEYCLOAK-3310	I'm honestly not following what you are asking for here. Are you talking about adding custom attributes to the user? If so that's already supported through mappers both for identity brokering and user federation.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3310
14864	KEYCLOAK-3310	To my understanding the mapping features are applicable under the precondition that a social media link is still established for a user. My request refers to the intention to become able to create this link to an identity provider as part of the process when a user is created in Keycloak by user federation mechnism. Perhaps it is a more reasonable way to override the isValid-method by adding an admin-REST-call to establish this link instead of extending/refactoring the user federation provider SPI for that purpose?...	1	1	https://issues.jboss.org/browse/KEYCLOAK-3310
14865	KEYCLOAK-3310	Oki, I finally get what you want. How about using a required action for it? You can have a default custom required action that retrieves the social media link and associates it with the user on the first login?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3310
14866	KEYCLOAK-3309	What container is this actually referring to? Our Keycloak images are here https://github.com/jboss-dockerfiles/keycloak, but there are no Keycloak+FreeIPA images	1	0	https://issues.jboss.org/browse/KEYCLOAK-3309
14867	KEYCLOAK-3309	[~mposolda] what's the purpose of the keycloak-freeipa-docker image? Is it just something you've used to help test FreeIPA?	1	0	https://issues.jboss.org/browse/KEYCLOAK-3309
14868	KEYCLOAK-3309	What's the reasoning behind requiring people to build it from source themselves?	1	0	https://issues.jboss.org/browse/KEYCLOAK-3309
14869	KEYCLOAK-3279	[~charlee_ch], do you have anything else to add to this issue ? May I assume that you get it working ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3279
14870	KEYCLOAK-3279	Thanks ! Can we close this issue ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3279
14871	KEYCLOAK-3237	Re-opening this as scope=openid has not been added to the JavaScript adapter. Did you add it to all the other adapters?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3237
14872	KEYCLOAK-3209	It could be optional with a warning telling that there will be javascript injected. The http header content-security-policy should also automatically reflect that so that there will still be no inline javascript. You say you are reluctant to add javascript, is that a platform feature or a personal trophée (not trying to be offensive)? For the security part, couldn't there be throttling ? When an user fails registration because an username or an email it means he has to retype 3 field : username/email, password, and confpassword. What if the realm has a lot of user, and the new user tries 7 times ? My guess is that most of them are gonna leave before even attaining this number. That's what I meant by extremely bad, which might be an overstatement. What's more important to a webmaster : having no javascript on its register page or not losing potential subscribers ? Most of the popular website have this feature, so there should be a secure way of doing it.	1	0	https://issues.jboss.org/browse/KEYCLOAK-3209
14873	KEYCLOAK-3209	Hello, I'm just checking by to see if you changed your mind about this. In other words, may I expect this feature in the future ?	1	0	https://issues.jboss.org/browse/KEYCLOAK-3209
14874	KEYCLOAK-3182	The blog post for the 2.3.0.CR1 release (http://blog.keycloak.org/2016/10/keycloak-230cr1-released.html) mentioned that Keycloak now passes the OpenID Connect certification tests. Would it be possible to attach the certification results to this issue? :).	1	1	https://issues.jboss.org/browse/KEYCLOAK-3182

#	Requirement	Questions	Answered	Successfully Answered	URL
14875	KEYCLOAK-3175	I've run the OIDC conformance test as per instructions on http://openid.net/certification/testing/ . What I did was: - Setup of openshift instance with Keycloak 2.0.0.CR1 on https://keycloak-mposolda.rhcloud.com/auth . Created confidential client "openidd" with redirect_uri https://op.certification.openid.net:60720/authz_cb . This points to the testing client deployed by OIDC conformance testsuite. Used user "admin" with some claims filled (email, first name, last name) and realm "master" - Then configured the testing instance like this (second line are my answers): Does the OP have a .well-known/openid-configuration endpoint? Yes Do the provider support dynamic client registration? No (just for easier start) redirect_uris: https://op.certification.openid.net:60720/authz_cb client_id: openidd client_secret: 98d90dd1-9d2e-43ad-a46b-1daeec3f5133 Which subject type do you want to use by default? Public Which response type should be used by default? Code Select supported features: JWT signed with algorithm other than "none" Test specific request parameters: nothing filled So basically I've tested "Basic" profile for startup. The results from testing application: 24 PASSED tests 2 UNKNOWN test results (both defacto passed, but it may be some manual actions involved) 5 FAILED tests 9 WARNING tests (but 5 of them are related to scope parameter, so have the same cause) I will add details to the JIRA for basic profile certification. But in shortcut: - It's easy to setup and run testing application and run/rerun conformance tests anytime we want - Test works in a way that you need to manually click and run individual tests. Some tests require some manual actions (eg. delete cookies) - The keycloak must be available online, hence openshift was the easiest possibility for me to run this though. It allows to restart the server and eventually upload files (jars) through SSH. So we can easily upload snapshot with JAR to openshift and restart the instance when we have some fix. Then we can rerun particular test scenario. - It doesn't seem to be easy (or even possible) to run automated test against the conformance application inside our testsuite. However what we can easily do is "copy" test scenarios into our testsuite. IMO we should have automated tests for all the scenarios required by conformance testsuite (even for those, which we already pass).	1	1	https://issues.jboss.org/browse/KEYCLOAK-3175
14876	KEYCLOAK-3175	Can you document somewhere how to run the conformance testing? Adding it to here https://github.com/keycloak/keycloak/tree/master/misc maybe?	1	0	https://issues.jboss.org/browse/KEYCLOAK-3175
14877	KEYCLOAK-3135	*1) Usecase* Wonder about support classic social use-cases : - I want to share my photoalbum with user X . - I want to share my photoalbum with group of users Y (Google+ circles) . Note that group Y is not "global" group, but is "scoped" to user, who is owner of the group and has permission to add/remove other member to this group as he want . In other words, it's like Google+ circle where I can add/remove any users to some of my circles. Ideally I want to manage authz directly from my application (Click on "Share photoalbum" from inside the photoalbum application). *2) How to implement. ACL?* For supporting this, it might be needed to have some support for ACL policy. The ACL table may look like this: Resource_id | scope (or list of scopes) | user_id (or circle_id) - The ACL records must be possible to manage remotely - It might be special "scopes" for sharing. For example i click "Share the photoalbum with user X" . Then user X has possibility to view the photoalbum, but he has also possibility to share it with other different users, who can share it again etc. - User must be able to add other users to his circles. This should be again possible to do remotely It seems this is quite complex and will likely require much more though to be supported properly. *3) Support OOTB?* Not sure if this use-case is something to support OOTB or is rather something "advanced", which people need to implement by themselves if they want it? Maybe we should wait for feedback from community if this is common usecase, which people want?	0	0	https://issues.jboss.org/browse/KEYCLOAK-3135
14878	KEYCLOAK-3135	I am also interested in an API for Policy Management. Mostly for User-Based Policy and a new Client-Based Policy, to be used for UMA. A UMA Resource need a Permission as in the Admin GUI today, but it would be good to be able to add Policies directly for the resource? Or should all policies be added via a Permission? As in the Resource Registration RFC https://docs.kantarainitiative.org/uma/rec-oauth-resource-reg-v1_0_1.html#rfc.section.2.2.1 A "user_access_policy_uri" field with the URL for policy management should be returned in the response when registering a resource.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3135
14879	KEYCLOAK-3135	[~jlpettersson], I think we can support Client-based policies OOTB. Are you interested in contributing ? Regarding the Policy API, today you can CRUD policies and permissions using Keycloak Java Administration Client. All endpoints used by the administration console are exposed and accessible from this client. Indeed we are using this endpoints in a few tests we run for authorization services. However, we have a lot of improvements to do in this area in order to provide a more clean and simple API for remote policy management.	1	0	https://issues.jboss.org/browse/KEYCLOAK-3135
14880	KEYCLOAK-3135	We have changed the AuthZ REST APIs in order to provide a better usability when managing policies remotely. For now, policy management can be performed via Keycloak Java Admin Client. Can we close this issue ?	0	0	https://issues.jboss.org/browse/KEYCLOAK-3135
14881	KEYCLOAK-3124	You're right about the first point. Not sure about the second point. Couldn't we just as well add adapter tests with WildFly to the default build as a Undertow adapter test.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3124
14882	KEYCLOAK-3103	What about logs ? Could you provide some ? It would also help if you attach the encrypted response, so I can look its structure.	1	0	https://issues.jboss.org/browse/KEYCLOAK-3103
14883	KEYCLOAK-3090	Are you sure? Quick look at the code suggests user should be available. It's not used in the default template, but should be available. For example ${user.username} should work in that template.	1	1	https://issues.jboss.org/browse/KEYCLOAK-3090
14884	KEYCLOAK-3065	This doc is no longer updated, can you check new Gitbook version?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3065
14885	KEYCLOAK-3065	[~stianst]My bad, I am seeing it works when you have the format like in your second format. Sorry for false alarm.	0	0	https://issues.jboss.org/browse/KEYCLOAK-3065
14886	KEYCLOAK-3065	Sent PRs to keycloak master and 1.9.x to remove provider from keycloak-server.json: https://github.com/keycloak/keycloak/pull/2901 https://github.com/keycloak/keycloak/pull/2902 [~stianst] I am thinking about add the note to gitbook too to emphasize that people need to remove "provider" to have enabled switch working. As I suppose that in keycloak-server.json for product, the line with provider will be there as we don't have more builds, so just switch to "enabled": false won't work. WDYT?	0	0	https://issues.jboss.org/browse/KEYCLOAK-3065
14887	KEYCLOAK-3060	Do we want to move this to 3.x? This isn't just an SPI issue and requires changes to LDAP, SSSD, admin console and account service.	0	0	https://issues.jboss.org/browse/KEYCLOAK-3060
14888	KEYCLOAK-3056	Would it be possible with some new configuration options, to verify signature against a configured certificate (or the one bundled in the assertion, if there is an option to verify the common name). I have to analyze this some more, but without this, how can people trust the assertion? IP-filtering on the saml broker endpoint?	1	1	https://issues.jboss.org/browse/KEYCLOAK-3056
14889	KEYCLOAK-3056	[~stianst]: Has there been any work on this? As there is already signature verification on the response, I might attempt a pull request for the same in the assertion...	1	1	https://issues.jboss.org/browse/KEYCLOAK-3056
14890	KEYCLOAK-3048	Isn't SNI supposed to just work with newer Java and HttpClient library?	0	0	https://issues.jboss.org/browse/KEYCLOAK-3048

#	Requirement	Questions	Answered	Successfully Answered	URL
14891	KEYCLOAK-3012	What about adding two buttons to the password policy page "Require all users to update password" and "Require users with insufficient passwords to update password", then also add a new conditional update password policy required action?	0	0	https://issues.jboss.org/browse/KEYCLOAK-3012
14892	KEYCLOAK-2994	Just did a fresh PR concerning this issue. Found quite a lot of typos in the german translation.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2994
14893	KEYCLOAK-2984	But if you had local users and later added an LDAP directory as an option, you *could* have duplicates. How are those handled? Wouldn't that be the same means of resolution? I can see the concern about duplicates, but seems like there should be some way to handle it. AFAIK, no other IDM tool would treat what could even be a transient connection error to one source as a Fatal error blocking management of *all other sources* of user identities. That creates tons of operational issues and makes the entire system only as robust as the least reliable user store.	1	0	https://issues.jboss.org/browse/KEYCLOAK-2984
14894	KEYCLOAK-2984	It's been a while since I played in the Keycloak UI, but my setup was using the Keycloak DB for one set of users and LDAP in read-only mode for another. Since repositories like LDAP can be read-only, there already needs to be a consideration in Keycloak that transient errors in LDAP should *not* block adding local users when in read-only mode for sure. As far as I can recall, you are aware of where you are creating users, are you not? i.e. if you are intending to create local users in the Keycloak DB, you could do that. It wouldn't just arbitrarily decide where to create them, right? If that's the case, maybe that's the root of the problem. My specific case is *not* about fallback. It's about maintaining maximum availability of features in the face of one provider not being reachable. i.e. creating local users explicitly not in LDAP.	1	0	https://issues.jboss.org/browse/KEYCLOAK-2984
14895	KEYCLOAK-2984	I second that if users are supposed to be created locally then failure of a user storage provider should not prevent creation of users. If users are going to be created in a user storage provider that is not working that obviously user creation should fail. [~mposolda] how do we actually deal with where the user is created?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2984
14896	KEYCLOAK-2984	[~bill.burke] Yes, we can do the callback like this. However this issue is not about failed "addUser", but failed user lookup before the addUser is called. In UsersResource.createUser we have this: {code} if (session.users().getUserByUsername(rep.getUsername(), realm) != null) { return ErrorResponse.exists("User exists with same username"); } {code} Assume the LDAP is down. What happens is: - Admin wants to create new user "john" - session.users().getUserByUsername is called. User "john" doesn't exists in local Keycloak DB - So UserStorage is queried for "getUserByUsername" . Since LDAP is unavailable, this will throw an exception. Hence new user won't be created because LDAP is down and the lookup of user throwed the exception. Note that this issue will happen even if LDAP doesn't have registration enabled. As the issue is in userLookup, not in the user creation itself.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2984
14897	KEYCLOAK-2979	please help us what is the Please use user mailing list for support?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2979
14898	KEYCLOAK-2966	Comments from Marek: {quote} Few things: - It will be good to have some OOTB support for multivalued attributes. You will be able to define if attribute is multivalued and then in registration/account pages, users will see something like we have in admin console for "redirect uris" or "web origins" in client detail page. - Besides validation, it may be useful to add some "actions" when attribute is changed? For example if user changes email, there will be the optional action, which will switch "emailVerified" to false and put the "VerifyEmail" required action on him. When he changes mobile number, it will send him SMS and he will need to confirm it somehow (perhaps again through required action), etc. - It will be probably useful to allow admin to skip validation (and actions) for certain attributes. Maybe validators could have an option like "Skip admin" or something like that? Or should we always skip the validations for admin? {quote}	0	0	https://issues.jboss.org/browse/KEYCLOAK-2966
14899	KEYCLOAK-2966	Is there a way to help advance this feature request? Some work in progress somewhere? I did not see anything on the mailing lists, but I may have searched badly.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2966
14900	KEYCLOAK-2966	I haven't had time to get very far with this so I don't have much to share other than a few things I've shared on the mailing list the last few months. You should be able to find it using the search on our website. You could start describing what you need and why you're interested. Also, how much are you willing to get involved?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2966
14901	KEYCLOAK-2962	I'm not sure how so, can you explain it? Ajax requests are sent to the very same URL as normal GET and POST requests. This can't be influenced, this is how the JSF framework works. We are not talking about custom REST services (just to make that clear). So I can't make up how to apply the multi-tenancy example to this problem because they need different paths as far as I understand the multi-tenancy solution.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2962

#	Requirement	Questions	Answered	Successfully Answered	URL
14902	KEYCLOAK-2951	find our configuration... one Realm with two separate User Federations (SUBDOMAIN1 and SUBDOMAIN2). {noformat} SUBDOMAIN1 LDAP Connection URL ldap://subdomain1.mydomain.com:389 Users DN ou=users,dc=subdomain1,dc=mydomain,dc=com Kerberos Kerberos Realm SUBDOMAIN1.MYDOMAIN.COM Server Principal HTTP/myappli.mydomain.com@MYDOMAIN.COM KeyTab /keytabs/myappli.mydomain.com.keytab SUBDOMAIN2 LDAP Connection URL ldap://subdomain2.mydomain.com:389 Users DN ou=users,dc=subdomain2,dc=mydomain,dc=com Kerberos Kerberos Realm SUBDOMAIN2.MYDOMAIN.COM Server Principal HTTP/myappli.mydomain.com@MYDOMAIN.COM KeyTab /keytabs/myappli.mydomain.com.keytab {noformat} The current authenticated user on the Windows OS is {noformat} MYUSERID@SUBDOMAIN2.MYDOMAIN.COM {noformat} The following messages show that both identity store registries are created... SUBDOMAIN1 first, then SUBDOMAIN2. {noformat} 2016-11-25 09:36:49,589 INFO [org.keycloak.federation.ldap.LDAPIdentityStoreRegistry] (default task-3) Creating new LDAP based partition manager for the Federation provider: SUBDOMAIN1, LDAP Configuration: {serverPrincipal=HTTP/myappli.mydomain.com@MYDOMAIN.COM, pagination=true, connectionPooling=true, usersDn=ou=users,dc=subdomain1,dc=mydomain,dc=com, userAccountControlsAfterPasswordUpdate=true, useKerberosForPasswordAuthentication=true, bindDn=MYKDCUSERID@mydomain.com, usernameLDAPAttribute=sAMAccountName, vendor=ad, uuidLDAPAttribute=objectGUID, allowKerberosAuthentication=true, connectionUrl=ldap://subdomain1.mydomain.com:389, syncRegistrations=false, authType=simple, debug=true, searchScope=2, keyTab=/keytabs/myappli.mydomain.com.keytab, useTruststoreSpi=ldapsOnly, kerberosRealm=SUBDOMAIN1.MYDOMAIN.COM, userObjectClasses=person, organizationalPerson, user, rdnLDAPAttribute=cn, editMode=UNSYNCED, batchSizeForSync=1000} 2016-11-25 09:36:49,605 DEBUG [org.keycloak.federation.ldap.idm.store.ldap.LDAPOperationManager] (default task-3) Creating LdapContext using properties: [{java.naming.security.principal=MYKDCUSERID@mydomain.com, java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory, com.sun.jndi.ldap.connect.pool=true, java.naming.ldap.attributes.binary=objectGUID, java.naming.provider.url=ldap://subdomain1.mydomain.com:389, java.naming.security.credentials=[C@328d1274, java.naming.security.authentication=simple}] 2016-11-25 09:36:49,609 INFO [org.keycloak.federation.ldap.LDAPIdentityStoreRegistry] (default task-3) Creating new LDAP based partition manager for the Federation provider: SUBDOMAIN2, LDAP Configuration: {serverPrincipal=HTTP/myappli.mydomain.com@MYDOMAIN.COM, pagination=true, connectionPooling=true, usersDn=ou=users,dc=subdomain2,dc=mydomain,dc=com, userAccountControlsAfterPasswordUpdate=true, useKerberosForPasswordAuthentication=true, bindDn=MYKDCUSERID@mydomain.com, usernameLDAPAttribute=sAMAccountName, vendor=ad, uuidLDAPAttribute=objectGUID, allowKerberosAuthentication=true, connectionUrl=ldap://subdomain2.mydomain.com:389, syncRegistrations=false, authType=simple, debug=true, searchScope=2, keyTab=/keytabs/myappli.mydomain.com.keytab, useTruststoreSpi=ldapsOnly, kerberosRealm=SUBDOMAIN2.MYDOMAIN.COM, userObjectClasses=person, organizationalPerson, user, rdnLDAPAttribute=cn, editMode=UNSYNCED, batchSizeForSync=1000} 2016-11-25 09:36:49,609 DEBUG [org.keycloak.federation.ldap.idm.store.ldap.LDAPOperationManager] (default task-3) Creating LdapContext using properties: [{java.naming.security.principal=MYKDCUSERID@mydomain.com, java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory, com.sun.jndi.ldap.connect.pool=true, java.naming.ldap.attributes.binary=objectGUID, java.naming.provider.url=ldap://subdomain2.mydomain.com:389, java.naming.security.credentials=[C@74951089,	1	0	https://issues.jboss.org/browse/KEYCLOAK-2951
14903	KEYCLOAK-2950	[~mposolda] for [KEYCLOAK-3829|https://issues.jboss.org/browse/KEYCLOAK-3829] I was considering something like: {code:java} public class BrokeredIdentityContext { @Override public boolean equals(Object o) { if (o == null) return false; UserModel userModel = (UserModel) o; if (id != null ? !id.equals(userModel.getId()) : userModel.getId() != null) return false; if (username != null ? !username.equals(userModel.getUsername()) : userModel.getUsername() != null) return false; if (email != null ? !email.equals(userModel.getEmail()) : userModel.getEmail() != null) return false; if (firstName != null ? !firstName.equals(userModel.getFirstName()) : userModel.getFirstName() != null) return false; return lastName != null ? !lastName.equals(userModel.getLastName()) : userModel.getLastName() != null; } } {code} Or if we think override equals is a terrible idea, I just add an utility method. I think this can solve the problem here, too. Don't you think? For example on IdentityBrokerService or IdentityProviderMapper we call: {code:java} if !(brokeredIdentityContext.equals(federatedUser)) { //update } {code} Does it make sense?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2950
14904	KEYCLOAK-2950	[~stianst] Few points: 1) The built-in mappers still doesn't handle this 100% correctly and they write to the keycloak DB even if the value is the same. So pointless cache invalidations are still there. For example looking at org.keycloak.broker.oidc.mappers.UserAttributeMapper#updateBrokeredUser I can still see (and just tested) that for example: -- Mapper for basic attributes (firstName, lastName, email) always update the user account and invalidate the user from the cache even if the old value is same https://github.com/keycloak/keycloak/blob/2.5.0.Final/services/src/main/java/org/keycloak/broker/oidc/mappers/UserAttributeMapper.java#L149 -- For non-basic attributes, there is check for CollectionUtil.collectionEquals, which is good. However in case that attribute is empty, there is always call to "user.removeAttribute" and pointless cache invalidation. We probably have similar stuff in other mapper impls too. 2) In last comments https://issues.jboss.org/browse/KEYCLOAK-2943 we discussed the possibility, that broker login will update user account just in case that particular broker is "owner" of the Keycloak user account. Which means that Keycloak account was registered through login of this broker. So in case that existing Keycloak account was linked to the broker "xy", the additional login of this user through broker "xy" won't update the account and won't call the "IdentityProviderMapper.updateBrokeredUser" . This will also fix the KEYCLOAK-2943 . WDYT? Anyway, this stuff should be re-visited IMO as we still have a space to improve here.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2950
14905	KEYCLOAK-2950	I've reopened KEYCLOAK-3829 which should cover not updating the user unless some values have indeed changed. Also added a separate issue KEYCLOAK-4184 to cover not updating the profile unless it's the IdP that created the account. Does that cover it and we can close this issue?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2950
14906	KEYCLOAK-2939	Will this feature mean that it will be possible to call the following from a browser? {code} POST /auth/realms/<realm>/protocol/openid-connect/logout ContentType: x-www-form-urlencoded Data: - client_id : <client> - refresh_token : <refresh-token-from-login-token> Headers: - Authorization Bearer <IdToken> {code} Currently it fail because the OPTIONS request does not contain the CORS-Headers which were set in the keycloak-configuration.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2939
14907	KEYCLOAK-2939	For browser based logout you should redirect to logout endpoint, this will redirect back to your application afterwards. Otherwise cookies on the Keycloak side are not invalidated. Are you using keycloak.js? There's a logout function there that does the job for you.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2939
14908	KEYCLOAK-2939	I know that there is the possibility with a redirect, but I wanted to do it without the redirect and opening another window. But if you say that there is no other option because of the cookie invalidation in keycloak itself, I will have to think about something else. What is the difference to the POST-Request the keycloak-spring-security-adapter does?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2939

#	Requirement	Questions	Answered	Successfully Answered	URL
14909	KEYCLOAK-2939	Why would you open login in another window? It redirects back to the application afterwards.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2939
14910	KEYCLOAK-2939	Well, I'm asking why don't you want a redirect? That's the only way to do a proper logout and properly invalidate all cookies on the Keycloak side.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2939
14911	KEYCLOAK-2939	Okay, I thought it would be easier with angular 2. But if you recommend that I will do it that way :) But back to my question, will it be possible after this ticket is fixed, or do I get the OpenID specification wrong?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2939
14912	KEYCLOAK-2926	What do you mean about unused required action?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2926
14913	KEYCLOAK-2909	Should I also remove tests under account.custom package?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2909
14914	KEYCLOAK-2858	I'm talking about this: username = *UNICODECHARNOCRLF Which of course does not imply case sensitivity is a must, only the charset, but they also write: "Unless otherwise noted, all the protocol parameter names and values are case sensitive." What do you think?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2858
14915	KEYCLOAK-2858	All usernames in Keycloak are lowercased when saved in the db. Same goes for username in tokens. We're not going to change this as it would break backwards compatiblity. There are also good reasons why we're doing it like this. For example to prevent a situation where you have "userA" and "usera" in the system which refer to different users. Do you actually have a real problem with this or is it just because?!	1	1	https://issues.jboss.org/browse/KEYCLOAK-2858
14916	KEYCLOAK-2857	Isn't there a password configuration type that you can use? There's one for authenticators	0	0	https://issues.jboss.org/browse/KEYCLOAK-2857
14917	KEYCLOAK-2830	Well, they are automatically added by the keycloak wildfly adapter in our configuration. And they are also added in the verify step. In class RedirectUtils line 126, the baseUri.getPort() gives 443 in our configuration, which matches the URI we get from the wildfly adapter, but doesn't match the one from the javascript adapter. We are using an apache HTTPS reverse proxy in front of keycloak, maybe that is the reason. The ProxyPeerAddressHandler from wildfly always sets the specific port, that was send as X_FORWARDED_PORT. I assume if the wildfly does the SSL termination itself and runs on port 443, it will just leave the port empty and then everything works. Would it be ok, if the RedirectUtils would allow both URIs with and without ports for well known port numbers? So when the scheme is HTTP and the baseUri.getPort() is 80, I would allow the URIs http://example.com/ and http://example.com:80/ and for HTTPS and baseUri.getPort() 443 allow https://example.com/ and https://example.com:443/? This should handle both cases. Any ideas?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2830
14918	KEYCLOAK-2808	ok, how can we make SSO possible between different applications of whose not all can "speak" openid-connect? We´ve tried many many ways and relying on the Identity Token seems to be the easiest way.	1	0	https://issues.jboss.org/browse/KEYCLOAK-2808
14919	KEYCLOAK-2797	I gave this a try but the problem is that the counter based OTP works different than the timer based approach and since I need to support "old hardware OTP generators" I need to use the timer based OTP auth ... but as said need to ensure that the token can only be used once. How about adding an option like "allow token reuse: true/false" with default "true" to the "Time Based" OTP type in the OTP Policy page? Then the server needs to store the most recent successfully accepted token together with the servertime on the credentials record. With this information a token that was already used can be rejected.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2797
14920	KEYCLOAK-2784	Can you please confirm if this is still an issue with 1.9.2?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2784
14921	KEYCLOAK-2784	Stian, Is it possible to add this fix to version 2.0.0.CR1? This problem is stopping us from going to production with many of our products.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2784
14922	KEYCLOAK-2777	I'm not following, can you elaborate?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2777
14923	KEYCLOAK-2751	Which REST API is this?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2751
14924	KEYCLOAK-2750	I appreciate that it's far from ideal behavior. There are quite a few "usability" issues with the rest endpoints and we're discussing whether we're going to try to resolve them in the current API, or introduce a new v2 of the API. This work will not be done until 2.x though.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2750
14925	KEYCLOAK-2749	This is going to have to wait for 2.x. Is that ok?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2749
14926	KEYCLOAK-2749	+ 1.9.2 is pretty much done and I don't want more changes in it, except bug fixes, we're past feature freeze and I'm certainly not going to accept features/enhancements for 1.9.3. Is there a way you can work around this? For example to use the redirect uri in your authenticator instead?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2749
14927	KEYCLOAK-2736	??A user can import a JSON realm by mistake via Manage/Import since that's the only visual thing you see when you enter the management console. When you do that, the real is not added properly.?? I've got no idea what you are saying here ??To add a realm, you have in fact to click on "Master" on top left and then click on "Add Realm" and then you can import a JSON file.?? Do you have a better suggestion? It seems pretty clear to me	0	0	https://issues.jboss.org/browse/KEYCLOAK-2736
14928	KEYCLOAK-2721	Hello. I know this issue is marked as closed, however I still get the same error with version 2.5.1 (both keycloak client and core). In my configuration, the client is instantiated in a ejb.jar while I access it from a .war file (all in an .ear deployment). Is there anything I can do, or do I have to wait of a future release? Thank you in advance.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2721
14929	KEYCLOAK-2709	[~stianst] [~bill.burke] Can I provide PR for this? This is really UX fail, try to login Red Hat Customer Portal where you are looking few seconds to the the blank page. I looked into code and it is probably not so easy to use realm Login theme there, so my proposal is: # Change page {{<title>}} to "Authentication redirect" # Add visible paragraph into body with text "We are redirecting you due to authentication, wait a moment please." Should be shown above {{<noscript>}} section. This is really small change in {{org.keycloak.saml.BaseSAML2BindingBuilder.buildHtml(String, String, boolean)}} but I believe it is a big improvement for end users.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2709
14930	KEYCLOAK-2709	If there's anything beyond a blank page it shouldn't be hardcoded. When it's used by identity brokering it should be fine to use login theme for that, but are these not also used on the adapter side?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2709
14931	KEYCLOAK-2684	Did you try the Jetty 9.2 adapter?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2684
14932	KEYCLOAK-2684	Hi guys, can you please also check if this adapter works with jetty running stateless (NO_SESSIONS) and keycloak adapter running the TokenStore with cookie ? This has been an issue with the adapted for 9.2.x already.. See: https://issues.jboss.org/browse/KEYCLOAK-2514 Thanks Gregor	1	0	https://issues.jboss.org/browse/KEYCLOAK-2684
14933	KEYCLOAK-2677	This should be high priority and something we should work on. I'll try as much as possible to make sure this happens soon, but it's going to be 3.x at the earliest. Due to a lot of other things happening at the moment I can't make any promises though. Can you elaborate on progressive user profiling? I'm not sure I fully understand what you are after.	1	0	https://issues.jboss.org/browse/KEYCLOAK-2677
14934	KEYCLOAK-2677	[~velias] I get what you're after now and will make sure once we do user profile SPI it will cover (or allow customizing) to support what you are after. One last question do you have an idea on how you'd specific when the additional fields would be required? Could perhaps be done with scopes? Could have user attributes that are always required, required if client asks for a specific scope or optional. Whenever a client tries to authenticate and requests a scope where the user hasn't provided the required attributes a user profile form would pop up.	1	0	https://issues.jboss.org/browse/KEYCLOAK-2677

#	Requirement	Questions	Answered	Successfully Answered	URL
14935	KEYCLOAK-2654	The option should be available to all social providers except Twitter. Is it only available for the Google provider?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2654
14936	KEYCLOAK-2605	sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.test.impl.TestContextHandler.createClassContext(TestContextHandler.java:92) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.test.impl.TestContextHandler.createSuiteContext(TestContextHandler.java:73) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:145) at org.jboss.arquillian.core.impl.ManagerImpl.fire(ManagerImpl.java:116) at org.jboss.arquillian.test.impl.EventTestRunnerAdaptor.before(EventTestRunnerAdaptor.java:108) at org.jboss.arquillian.junit.Arquillian$4.evaluate(Arquillian.java:241) at org.jboss.arquillian.junit.Arquillian.multiExecute(Arquillian.java:426) at org.jboss.arquillian.junit.Arquillian.access$200(Arquillian.java:54) at org.jboss.arquillian.junit.Arquillian$5.evaluate(Arquillian.java:259) at org.jboss.arquillian.junit.Arquillian$7$1.invoke(Arquillian.java:319) at org.jboss.arquillian.container.test.impl.execution.ClientBeforeAfterLifecycleEventExecuter.execute(ClientBeforeAfterLifecycleEventExecuter.java:99) at org.jboss.arquillian.container.test.impl.execution.ClientBeforeAfterLifecycleEventExecuter.on(ClientBeforeAfterLifecycleEventExecuter.java:72) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.invokeObservers(EventContextImpl.java:99) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:81) at org.jboss.arquillian.container.test.impl.client.ContainerEventController.createContext(ContainerEventController.java:142) at org.jboss.arquillian.container.test.impl.client.ContainerEventController.createBeforeContext(ContainerEventController.java:124) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:497) at org.jboss.arquillian.core.impl.ObserverImpl.invoke(ObserverImpl.java:94) at org.jboss.arquillian.core.impl.EventContextImpl.proceed(EventContextImpl.java:88) at org.jboss.arquillian.test.impl.TestContextHandler.createTestContext(TestContextHandler.java:130) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at	0	0	https://issues.jboss.org/browse/KEYCLOAK-2605
14937	KEYCLOAK-2603	I am seeing on my MySQL that all mentioned indexes are already available. Indeed we have foreign key for USER_ID column in both USER_ROLE_MAPPING and CREDENTIAL table. Similarly we have unique constraint for USERNAME in USER_ENTITY. So MySQL is adding BTREE indexes by default for those cases. Is it any more info, which I am missing why those 3 indexes are necessary? Especially: - What RDBMS was used for performance test? - What is the tested scenario? This is what I am seeing on my MySQL: {code} mysql> show indexes in USER_ROLE_MAPPING; +----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+	Table	Non_unique	Key_name	Seq_in_index	Column_name	Collation	Cardinality	Sub_part	Packed	Null	Index_type	Comment	Index_comment	+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+	USER_ROLE_MAPPING	0	PRIMARY	1	ROLE_ID	A	8174	NULL	NULL		BTREE				USER_ROLE_MAPPING	0	PRIMARY	2	USER_ID	A	8174	NULL	NULL		BTREE				USER_ROLE_MAPPING	1	FK_C4FQV34P1MBYLLOXANG7B1Q3L	1	USER_ID	A	8174	NULL	NULL		BTREE			+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+ 3 rows in set (0.00 sec) mysql> show indexes in CREDENTIAL; +----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+	Table	Non_unique	Key_name	Seq_in_index	Column_name	Collation	Cardinality	Sub_part	Packed	Null	Index_type	Comment	Index_comment	+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+	CREDENTIAL	0	PRIMARY	1	ID	A	1887	NULL	NULL		BTREE				CREDENTIAL	1	FK_PFYR0GLASQYL0DEI3KL69R6V0	1	USER_ID	A	1887	NULL	NULL	YES	BTREE			+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+ 2 rows in set (0.00 sec) mysql> show indexes in USER_ENTITY; +----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+	Table	Non_unique	Key_name	Seq_in_index	Column_name	Collation	Cardinality	Sub_part	Packed	Null	Index_type	Comment	Index_comment	+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+	USER_ENTITY	0	PRIMARY	1	ID	A	1937	NULL	NULL		BTREE				USER_ENTITY	0	UK_DYKN684SL8UP1CRFEI6ECKHD7	1	REALM_ID	A	1937	NULL	NULL	YES	BTREE				USER_ENTITY	0	UK_DYKN684SL8UP1CRFEI6ECKHD7	2	EMAIL_CONSTRAINT	A	1937	NULL	NULL	YES	BTREE				USER_ENTITY	0	UK_RU8TT6T700S9V50BU18WS5HA6	1	REALM_ID	A	1937	NULL	NULL	YES	BTREE				USER_ENTITY	0	UK_RU8TT6T700S9V50BU18WS5HA6	2	USERNAME	A	1937	NULL	NULL	YES	BTREE			+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+ 5 rows in set (0.00 sec) {code}	1	1	https://issues.jboss.org/browse/KEYCLOAK-2603
14938	KEYCLOAK-2590	Do I understand this correctly, when running a test we want to be able to set Time.offset, and so we need a new REST endpoint to allow this, or maybe one exists already, and just has to be exposed through admin-client because it's missing, but I suppose not since you say to make sure that manage realm permission is checked? And for checking you mean to add at the beginning of the method the following line - and that permission is not current realm specific but cuts across all realms - Time.offset is static field: {code} auth.requireManage(); {code}	1	1	https://issues.jboss.org/browse/KEYCLOAK-2590
14939	KEYCLOAK-2537	Any progress? Roadmap?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2537

#	Requirement	Questions	Answered	Successfully Answered	URL
14940	KEYCLOAK-2489	Yes I can give it a spin. Where should the scripting infrastructure tests be placed? in keycloak-services? I guess the tests for the ScriptBasedAuthenticator should be implemented in the keycloak-integration test project? Given the new documentation infrastructure where should I put the docs for this?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2489
14941	KEYCLOAK-2480	Not quite sure what you are requesting here. The default setting will remain whatever WildFly provides as they have more experience with setting this up then we do. Further it's possible to configure it as you want. So what are you actually requesting?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2480
14942	KEYCLOAK-2480	The default settings (standalone-ha.xml) needs modifying for production in any case. You need to add a proper database, configure SSL, etc.. So doesn't really make that much difference that you need to apply different config for JGroups. What we could do is provide examples/documentation on how to setup JGroups for different environments (Docker, OpenShift, Amazon, etc..). We could maybe talk to WildFly guys to document that so we can reference it in our documentation. With regards to applying the config to a new version, have you had a look at jboss-cli? It's a great way to apply changes to standalone-ha.xml and can be executed offline. For example take a look at our JBoss CLI script[1] that can be used to install Keycloak onto an existing WildFly installation. It can be executed with: {code} bin/jboss-cli.sh --file=bin/keycloak-install-ha.cli {code} [1] https://github.com/keycloak/keycloak/blob/master/distribution/server-overlay/cli/keycloak-install-ha.cli	1	0	https://issues.jboss.org/browse/KEYCLOAK-2480
14943	KEYCLOAK-2474	Ok, great! Argh, I made my changes on the 1.8 version and the whole Keycloak project structure changed between 1.8 and 1.9! :-s Plus every file now includes a licence header, so the git patch is also useless. I guess I'll have to re-apply all changes manually ... oh well, keeps you busy :-) Is the current project structure and file contents stable now after those 'big impact' changes? Otherwise I'll wait a little more.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2474
14944	KEYCLOAK-2474	Sorry about that. We had to do it at some point. Doesn't git rebase work? It should figure out files have moved and show you diffs.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2474
14945	KEYCLOAK-2453	[~stianst] what exactly do you mean? Just document how to enable it?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2453
14946	KEYCLOAK-2438	[~stianst] correct me if I'm wrong, but don't you think that adding display name would be an overlap with "name" from [IdentityProviderFactory|https://github.com/keycloak/keycloak/blob/c7a8742a368bd8d76301145b08bb1e4af1b010e9/services/src/main/java/org/keycloak/social/facebook/FacebookIdentityProviderFactory.java#L31-L34] ? I was wondering if we could reuse the "name" attribute.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2438
14947	KEYCLOAK-2438	[~stianst] fair enough. Does it make sense to you? https://github.com/abstractj/keycloak/commit/5ac2c7bd7a24f1db9d44b0e289cc48091c25cac9 If you are ok with this change (just a raw idea), I will move forward on this.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2438
14948	KEYCLOAK-2422	Stian - this is a critical piece before you can make KC a product as all the Organizations typically use commercial certificates. KC generated certificates/keys will not be acceptable. Can you address this ASAP? Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-2422
14949	KEYCLOAK-2406	[~stianst] [~dadirkes] I'm new to keycloak but very much interested in using this project. is there a way to plug the persistence framework in Keycloak? I don't want to bother the awesome developers of this project with my own agenda, Red hat has its own. Although I'm more experienced using Cassandra, I'm just trying to learn Mongo for the sake of using keycloak with a nosql persistence framework. After reading about mongo replica sets and configuration and sharding servers, it just feels to me that keycloak would be so much more powerful if used in conjunction with Cassandra	1	0	https://issues.jboss.org/browse/KEYCLOAK-2406
14950	KEYCLOAK-2406	Hi Stian, Has a Jira case number been defined for the changes to UserProvider/UserFederationProvider you mentioned, so that we can know when that has been completed ? If not yet defined, would you be able to add a comment here when one has been defined ? Thanks	1	1	https://issues.jboss.org/browse/KEYCLOAK-2406
14951	KEYCLOAK-2392	would it be possible to log both username and and UUI together?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2392
14952	KEYCLOAK-2382	Given that many people have requested this enhancement, would it make sense to increase the size by 2 or 3 as a starting point to accommodate some of the use cases? This likely has negligible impact on performance, but will accommodate many more customers.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2382
14953	KEYCLOAK-2378	+1 The current version of the npm repo is 2.0.0. Please update it to the latest one - 2.1.0 final. Also, why do you call it a bower project and feature a bower installation on the npm page? I find it misleading.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2378
14954	KEYCLOAK-2359	Does this make sense for all emails or do you have some dedicated ones in mind?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2359
14955	KEYCLOAK-2359	Hi [~tdarimont] I don't have one at hand that works but I would suspect it would be something like this (for example in the error.ftl I would add the following) {code:java} <#if client?? && client.baseUrl?has_content> ${msg("backToApplication")} <#elseif (request.getParameter("client_id")?has_content> <#assign clientBaseUrl = ${msg("clientBaseUrl_"+request.getParameter("client_id"))}> <p>${msg("backToApplication")} </#if> //Then in my messages.properties file, I will have a key like this clientBaseUrl_DummyClientID=redirect url //*** NOTE the above is not tested.*** {code} But as you can see, this is not good and elegant. So was wondering if its better to send the client base url as well in the parameter list so we can then redirect it to that url without much fuss. I think we still need the client ID as well as we can then accordingly use it to load different style sheets that will style the messages as per the client style sheet. But if we can get the base url as well, I think that would be great	0	0	https://issues.jboss.org/browse/KEYCLOAK-2359
14956	KEYCLOAK-2351	OAuth2 allows that [1] but OpenID Connect doesn't for implicit nor hybrid flow to use response_type=token alone without "id_token" or "code" [2] [3] . So we're not going to support that as we want to be OpenID Connect compliant. Isn't it possible to configure swagger UI to use "response_type=id_token token" ? [1] https://tools.ietf.org/html/rfc6749#section-4.2.1 [2] http://openid.net/specs/openid-connect-core-1_0.html#ImplicitAuthRequest [3] http://openid.net/specs/openid-connect-core-1_0.html#HybridAuthRequest	1	1	https://issues.jboss.org/browse/KEYCLOAK-2351
14957	KEYCLOAK-2343	We currently need that functionality - is this currently in progress or may I submit a PR for this? Cheers, Thomas	1	0	https://issues.jboss.org/browse/KEYCLOAK-2343
14958	KEYCLOAK-2343	Can we just have an additional boolean param on searchForUser() on whether to make query exact or not? Same with a new @QueryParam("exact")?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2343
14959	KEYCLOAK-2343	Hi, I gave this a quick spin and sent a PR with a PoC implementation for (JPA, MongoDB, and some of the other providers) according to the suggested change. At the moment I'm using an ugly checkbox next to the user-search field to mark the search as exact (when checked) and unexact (when not checked, default). Would this be suitable (after some UI polishing and some additional tests?). Cheers, Thomas	1	0	https://issues.jboss.org/browse/KEYCLOAK-2343
14960	KEYCLOAK-2343	How about performing an exact search when the search term is wrapped in single / double quotes?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2343
14961	KEYCLOAK-2339	Shouldn't simply putting the impersonation info in the impersonated {{UserSessionModel}} be enough? Then mappers can be used to expose this information as claims if desired.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2339
14962	KEYCLOAK-2337	Did you completely remove the Keycloak theme? You need to leave keycloak/common as that contains resources used by the base theme as well (Angular for example).	1	1	https://issues.jboss.org/browse/KEYCLOAK-2337

#	Requirement	Questions	Answered	Successfully Answered	URL
14963	KEYCLOAK-2337	Did you update keycloak-themes.json? For keycloak it should only include common	1	1	https://issues.jboss.org/browse/KEYCLOAK-2337
14964	KEYCLOAK-2302	Would it be ok to include the subsystem by default if this PR is merged? I'd like it to be as simple for the administrators to enable this functionality as possible. I'll probably need some help at least on reviewing the Wildfly configuration as it's not our standard operating environment.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2302
14965	KEYCLOAK-2271	[~mposolda] how critical is this and how much effort would it be?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2271
14966	KEYCLOAK-2266	Is there a documentation which requests are valid for the keylcoak token introspection endpoint? With RFC7662 there are several possibilities to introspect a token.	1	0	https://issues.jboss.org/browse/KEYCLOAK-2266
14967	KEYCLOAK-2210	[~stianst] do you think it would be possible to push it to 1.8.CR2 release? Is it big change? It would be great to have it there because it's very misleading from UX perspective.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2210
14968	KEYCLOAK-2210	[~lkrzyzanek] we don't have time to do more features and improvements in 1.8.x now. If there are things that you must have let us know and we can see if we can make an exception. It has to be must haves rather than nice to haves at this point though.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2210
14969	KEYCLOAK-2210	What about: The link you clicked is not valid. I'm trying to change the impression from "hey your server has some trouble" to "OK it's my fault." There is no need to put much details what happened. Just have something else than "general error message".	1	1	https://issues.jboss.org/browse/KEYCLOAK-2210
14970	KEYCLOAK-2210	What is wrong with this message? "The link you clicked is a old stale link and is no longer valid. Maybe you have already verified your email?"	1	1	https://issues.jboss.org/browse/KEYCLOAK-2210
14971	KEYCLOAK-2170	What we have now is: - Tests for various response_types in package "org.keycloak.testsuite.oidc.flows" . This was added when working on OIDC conformance testsuite, which is passing now with implicit flow - Adapter tests done by QA team - AbstractJSConsoleExampleAdapterTest and it's subclasses	0	0	https://issues.jboss.org/browse/KEYCLOAK-2170
14972	KEYCLOAK-2141	Does the openid connect spec actually state it's recommended to not make email unique? Or does it just say it can change over time.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2141
14973	KEYCLOAK-2108	Could you send a mail about this to keycloak-dev please?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2108
14974	KEYCLOAK-2108	Sure. I registered to the dev-mailing list this morning and try to work my way trough the threads ... is there also a google searchable group for this? (I'm currently using this "hack": http://johnpoelstra.com/importing-mailing-list-archives-to-thunderbird/ to read the list via thunderbird - but there must be a better way to do this...) I'll post the question there... Btw. was the keycloak issue board witht the 3 lanes deleted? It is really hard to keep track of what happens now... When I try to access the board: https://issues.jboss.org/secure/RapidBoard.jspa?rapidView=1202 I see the message: "The requested board cannot be viewed because it either does not exist or you do not have permission to view it." Would be great of this would be opened / recreated again.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2108
14975	KEYCLOAK-2089	Is this related to the case where a <secure-deployment> without "secret" would imply that the KC adapter subsystem would try to auto-register the deployment with the server?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2089
14976	KEYCLOAK-2089	I tried to find which issue would be related to the case I mentioned, but I wasn't able to locate it. Is there already a JIRA for that?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2089
14977	KEYCLOAK-2084	[~stianst] for the 3 scenarios you described. Is this what you expect? 1. Admin username/password {code:bash} ./client-registration.sh -u john #Prompt for the password Password: ****** {code} 2. Service Account client id+secret, or JWK keys {code:bash} ./client-registration.sh --client-id 42424242 #Prompt for the secret Secret: ****** #JWK keys ./client-registration.sh --key client.jks {code} 3. Initial access token and registration access token {code:bash} ./client-registration.sh --initial-access-token #Prompt Token: ./client-registration.sh --registration-access-token #Prompt Token: #JWK keys ./client-registration.sh --key client.jks {code} Does it make some sense?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2084
14978	KEYCLOAK-2076	Few thoughts from me around reauthentication support on Keycloak server: h4. Authentication flow for Reauthentication Common Keycloak authentication is performed in "Browser" flow. Question is if we should upgrade this fow to support reauthentication, or define new one. The flows will be typically very similar, providing same authentication options, but also differing in behaviour: h5. Cookie * must be performed as first step. If user is not authenticated from cookie then authentication flow must be performed in next steps, else we can perform steps in Reauthentication mode. h5. Kerberos * new kerberos handshake should be performed to reauthenticate user. What to do if user differs from the original one we are reauthenticating, show error or continue in the process with new user? h5. Forms * we need new form template(s) different from one used during user login - another title - no username field, maybe username shown as readonly? - password field - in case of OTP enabled for current user we should also have field for one time password right on the form. But I'm not sure if current authentication flow mechanism allows to have same form for two underlying authenticators. h5. Rest of flow * Should we perform "Required actions" defined for given user after successful reauthentication? I think yes as it is the only point where we can force them during long lived user session. * After reauthentication flow is finished then authentication timestamp must be stored in user session, and then control returned back to SSO/client protocol to return user info back to client app. h4. Should we allow reauthentication over Identity providers? * I think yes, as user should originally authenticate over some of it. Probably not to automatically forward to the original identity broker, but show list of identity broker buttons on "Reauthenticate" form to allow user use them. * We should maybe allow realm scoped configuration of which identity broker may be used for reathentication and which not. * We also have to decide what to do if user returned from identity provider differs from the original one we are reauthenticating. Show error message or continue in login process with new user? * If identiy provider protocol supports reauthentication, should we use it (so force reauthentication in the provider also)? * how to handle "Authenticate by Default" marked identity providers in case of reauthentication? Should we care of them?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2076
14979	KEYCLOAK-2070	This is a *very* important feature for our project. Is there a workaround that we could implement waiting this feature to be available ? We are using OpenID connect with the javascript adapter.	0	0	https://issues.jboss.org/browse/KEYCLOAK-2070
14980	KEYCLOAK-2070	Is this a duplicate of KEYCLOAK-2775?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2070
14981	KEYCLOAK-2067	Is it always the case that people can sort this at DNS level? Don't some organizations work in a way that admin responsible for Keycloak is not responsible (and doesn't have access) to manage DNS server and viceversa? People can sort this at DNS level already, but "auth-server-url-for-backend-requests" might be an alternative for those where DNS is not an easy option.	1	0	https://issues.jboss.org/browse/KEYCLOAK-2067

#	Requirement	Questions	Answered	Successfully Answered	URL
14982	KEYCLOAK-2054	Upvote from me, I recently ran into the same problem where a HTTP proxy is between the keycloak server and the Java backend that uses keycloak client to authenticate a web user. The proxy configuration needs to be set on the {{org.apache.http.client.HttpClient}} which is build by the {{org.keycloak.adapters.HttpClientBuilder}}. An example configuration can be found on the [HTTPClient website|https://hc.apache.org/httpcomponents-client-ga/httpclient/examples/org/apache/http/examples/client/ClientProxyAuthentication.java]. Looking at the source code it looks like we could add an additional configuration section to the {{org.keycloak.representations.adapters.config.AdapterConfig}} process this config in {{org.keycloak.adapters.KeycloakDeploymentBuilder}} and add the proxy configuration to the HTTPClient in {{org.keycloak.adapters.HttpClientBuilder.build()}}. Does that sound like the way to do this? Looking at {{org.keycloak.adapters.HttpClientBuilder}} it appears it could do with some work as it uses Apache HTTP Client deprecated APIs.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2054
14983	KEYCLOAK-2054	[~nbertram] that sounds about right. Do you fancy sending a PR for this? We don't have any resources to work on it ourselves atm.	1	1	https://issues.jboss.org/browse/KEYCLOAK-2054
14984	KEYCLOAK-2054	I gave this a spin, would you mind having a look? https://github.com/keycloak/keycloak/pull/2381 Could someone add this as a github PR to this issue? It seems that I'm not allowed to do this...	0	0	https://issues.jboss.org/browse/KEYCLOAK-2054
14985	KEYCLOAK-2040	[~tdarimont] looks very interesting. Can you send a mail about it to keycloak-dev mailing list?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2040
14986	KEYCLOAK-2035	Hello [~stianst], Is the support on the UI totally necessary for it to be accepted?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2035
14987	KEYCLOAK-2019	This is what I am looking for as well. Can someone help us with some tips on how to do this?	1	1	https://issues.jboss.org/browse/KEYCLOAK-2019
14988	KEYCLOAK-2019	I honestly don't understand what this is about. Can you please give an example when you would propagate an error message from the underlying user federation provider directly to the users?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2019
14989	KEYCLOAK-2008	Can you please add some more details to this issue? How did you configure it, what was the error, what is the app backend, etc..	1	0	https://issues.jboss.org/browse/KEYCLOAK-2008
14990	KEYCLOAK-2008	On investigating, Found out that the Username of the external user is stored by his email address in Keycloak-"Users" while the username is configured as "LDAP entry dn" for internal users in keycloak. The authentication fails even after passing the DN of the user in the assertion claim. I wish to get the DN of this external user in the username field of Users. Do we need to configure a mapper for the same?	0	0	https://issues.jboss.org/browse/KEYCLOAK-2008
14991	KEYCLOAK-2003	[~mposolda] could you please elaborate more about what you have in mind? Would you like to prevent an attacker from attempting to brute force the client secret against KC?	1	0	https://issues.jboss.org/browse/KEYCLOAK-2003
14992	KEYCLOAK-1998	[~stianst] seems like this workflow is already working today. What I tested: 1. Deployed the latest version of KC 2. Went to my realm configuration, enabled "User Registration", included "totp" as required credential. 3. Tried to register a new user After the registration screen, the OTP is mandatory. Let me know if I'm misundestanding what you want.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1998
14993	KEYCLOAK-1974	[~stianst] ok, can you then set some fix version to not forgot on this, please?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1974
14994	KEYCLOAK-1949	I've got an implementation proposal: https://github.com/gerbermichi/keycloak/commit/0ef36f0ac446fcf70272f2aed05320c3a5083635 I would like to discuss some points of it: - What do you think of the "create-new-session" option? - What do you think of the execution order? - Where is the right place for the tests?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1949
14995	KEYCLOAK-1925	Can you specify more precisely what this task entails? Including what work is done, and what is still to be done here ...	0	0	https://issues.jboss.org/browse/KEYCLOAK-1925
14996	KEYCLOAK-1922	Hi Stian, you are right. These links are orginally used by the admin console and no client is available. But in the rest api doc (http://www.keycloak.org/docs/rest-api/index.html#_send_a_update_account_email_to_the_user) you can set a client id. If you do so, the client is available in the template, but <# skipLink??> evaluates to false. I use this call to trigger a set initial password action, after the user is created by a 3rd party application. In this case I want the user after setting his initial password, to be able to click on a back to to application link. Here is my workaound: {code} ${message.summary} <#if client??> <#if client.baseUrl??> ${msg("backToApplication")} </#if> </#if> {code} I am using Keycloak 1.9.2.	0	0	https://issues.jboss.org/browse/KEYCLOAK-1922
14997	KEYCLOAK-1912	Stian - I am presuming that you are referring to dynamic registration. Is that correct? If not, can you look into the below spec? https://openid.net/specs/openid-connect-registration-1_0.html Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-1912
14998	KEYCLOAK-1905	It would be better if it wasn't an admin endpoint, but instead something that was associated with users themselves and such protected by the user token, not an "admins" token. However, not sure this is a generic enough feature for us to add it. We are planning to add pluggable REST APIs though, so you could add it on yourself. Would that be an ok approach?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1905
14999	KEYCLOAK-1905	Right, a user based endpoint, as you described, would be more suitable - I just took the admin URI pattern as an example from the REST API doc :) In my point of view, it is a very generic feature - one could also think about requiring the user to enter a TOTP token again when he wants to change some import profile / account information (billing, bank account etc.) or after a longer time of inactivity. Could you give me some pointers where / how such a user based REST endpoint could be added? Having said that - if the feature doesn't make it in the product - pluggable rest apis could be one way to workaround that - do you have the JIRA issue for that at hand?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1905
15000	KEYCLOAK-1904	Can you create a bug for that issue?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1904
15001	KEYCLOAK-1903	hi, I already translate all messages I have to correct some sentences and create a PR what do you think about it ? best regards,	1	1	https://issues.jboss.org/browse/KEYCLOAK-1903
15002	KEYCLOAK-1903	Great, would you also be willing to maintain the translation in the future? If we don't have someone that can maintain it it'll quickly become outdated	0	0	https://issues.jboss.org/browse/KEYCLOAK-1903
15003	KEYCLOAK-1902	Could this be included in 2.0.0.Final?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1902
15004	KEYCLOAK-1902	I would very appreciate to see this feature as soon as possible. We have actually this requirement and would have to implement a custom deployment on Keycloak server in order to be able to search the users by custom attributes and roles. Is there a plan, when this feature will be supported natively?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1902
15005	KEYCLOAK-1901	Not sure about this one. I'm not to convinced about us using RestEasy client in the first place and I don't want to expose it. Can we do it by passing in the HttpClient instead?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1901
15006	KEYCLOAK-1893	For a hardware token how would you configure the device in Keycloak? I assume the totp secret is hard coded in the device itself and would have to be copied to Keycloak?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1893

#	Requirement	Questions	Answered	Successfully Answered	URL
15007	KEYCLOAK-1893	For hardware tokens I assume only an admin should be able to register the token with the user as the user themselves wouldn't have the knowledge to configure the token seed? With regards to 4, are we not spec compliant already? We have https://tools.ietf.org/html/rfc4226 and https://tools.ietf.org/html/rfc6238 already AFAIK.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1893
15008	KEYCLOAK-1893	It seems like you're willing to contribute what's required to get hardware tokens working properly in Keycloak? If so we should discuss the best approach through the mailing list. I think we have two options, either introduce a new OTP authenticator or make the existing authenticator configurable through the admin console. Can you start a general thread on the Keycloak dev mailing list with what you're trying to achieve and how? Also, include that you are willing to contribute it ;)	1	1	https://issues.jboss.org/browse/KEYCLOAK-1893
15009	KEYCLOAK-1885	Do you want to just add those that you need?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1885
15010	KEYCLOAK-1872	Will this feature be added in the near future? -- Thanks, Suresh	1	0	https://issues.jboss.org/browse/KEYCLOAK-1872
15011	KEYCLOAK-1851	Should this be rejected? The Realm settings forces verify email for all users. Required Actions is just for one user.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1851
15012	KEYCLOAK-1784	What do you think about configuring the Tomcat Valve in the global context.xml and then having conf/keycloak.xml instead of having to include both in the war? keycloak.xml would have to have secure-deployment elements like standalone.xml does	1	1	https://issues.jboss.org/browse/KEYCLOAK-1784
15013	KEYCLOAK-1744	How could Keycloak know if the password is expired when it is controlled by LDAP?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1744
15014	KEYCLOAK-1737	What custom events are you talking about? Your application shouldn't use Keycloak as an event store. Instead you should have a separate event store and make Keycloak and your application use that.	0	0	https://issues.jboss.org/browse/KEYCLOAK-1737
15015	KEYCLOAK-1735	I'm not following, what has to be re-developed?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1735
15016	KEYCLOAK-1735	I can't see the issue KEYCLOAK-528 in Jira(the duplicate). Can we see it? Thanks, Lucian	1	0	https://issues.jboss.org/browse/KEYCLOAK-1735
15017	KEYCLOAK-1733	Hi, this would be useful for us in scenarios where we can not modify the headers. For example downloads or file uploads. We have most of our app now working using keycloak (using javascript-client and wildfly backend) but its these edge-cases that prevent us from full usage right now.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1733
15018	KEYCLOAK-1724	Does it work when auth-method is set to KEYCLOAK? If so we need to update instructions on http://keycloak.github.io/docs/userguide/html/ch08.html#tomcat-adapter, feel free to create a PR for that.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1724
15019	KEYCLOAK-1720	Excuse me for questioning this, since I am probably missing some bits and pieces of the big picture. Why merging roles is not an option, if a user sets its app to behave like that? In case of role collision, the role would be in that list in any case, what's the issue with that? What I can't really understand is why at least that option can't be controlled by the developer so we can customize keycloack to every possible scenario.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1720
15020	KEYCLOAK-1720	We can probably improve on what we have. I'm not sure what you are proposing is correct though, because there's no way for the client in that case to distinguish if a role is associated with a client or a realm. Can you send a mail about this to the mailing list so we can get input from others as well?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1720
15021	KEYCLOAK-1720	I understand where you are going now. That makes sense to me. Knowing where those roles come from sounds like a must to have. So perhaps simply merging is not a good solution. The thing is, not having some roles at all is also an issue as far as I understand. If my client is in a realm. Both have roles and I need them both in my app, What am I supposed to do? We could also workaround the issue by making realm roles composite, and including some substitutive client-level role. But that approach shows itself unmanageable as more and more clients are added. Where could I find the mailing list? Is there any convention / protocol to share this discussion, or should I just send a regular mail with the information of this ticket.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1720
15022	KEYCLOAK-1719	I assume you're talking about the ability to pass a token by reference so it can be passed around without exposing the details of the token itself? In that case I assume only certain clients should be permitted to invoke the validation api to retrieve the actual jwt? Is there a spec for this or does the commercial vendors provide their own proprietary mechanisms?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1719
15023	KEYCLOAK-1719	If any client can invoke the validation api it seems to negate the whole point of pass by reference. In that case anyone can just invoke the validation api to obtain the details from the jwt. I don't quite follow your example on how it should work. Are you simply saying that the client that invokes the verification api needs to have the scope to obtain all details in the token? So if the token contains a few realm roles, a few client roles and the users profile the client needs to have the corresponding scope configured in the KC admin console? Further only confidential clients should be permitted to invoke this endpoints, but I don't think that's a problem as the public client type is for client-side applications, not services. This would probably require some token cache built in to the adapter otherwise this could end up with generating a lot of traffic to the Keycloak server.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1719
15024	KEYCLOAK-1718	Can you elaborate a bit on your use case? Do you have an Spring MVC based REST backend and a JS frontend that uses said backend and you want to transparently authenticate via Keycloak? Did you see the spring security XML configuration example under: http://keycloak.github.io/docs/userguide/html/ch08.html#d4e1223 ?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1718
15025	KEYCLOAK-1691	note, you guys *do* ship a -server-feature-pack.zip, which is cool, but lacks some modules like keycloak-undertow-adapter, etc.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1691
15026	KEYCLOAK-1690	Can you elaborate on what you need? We support authenticating users with Kerberos through SPNEGO/GSSAPI. This is used as a bridge between web SSO and desktop SSO. We don't support applications authenticating with Keycloak using Kerberos though.	0	0	https://issues.jboss.org/browse/KEYCLOAK-1690
15027	KEYCLOAK-1596	If you've been using Keycloak's adapter with Cordova or AeoroGear, we'd welcome help. My thinking was that we'd add a mobile section to the documentation covering Cordova, briefly covering AeroGear and links to AeroGear's oauth libraries. Not sure what you mean about general information, can you elaborate?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1596
15028	KEYCLOAK-1578	Any news on this topic ? Would you welcome a pull request ?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1578
15029	KEYCLOAK-1542	Would this not just re-implement what's already available in the WF/EAP console?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1542
15030	KEYCLOAK-1542	There are plans to allow login to WF/EAP console with KC. When that's available in WF we can configure it like that OOTB in KC. You'd still need URL though and it's on a separate port as well. TBH I'm a bit undecided on this one. On one side it would be nice to have everything available from one console, but on the other side duplicating functionality isn't ideal either. Let's bring this one to keycloak-dev and see what others think. Do you want to send a mail about it?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1542

#	Requirement	Questions	Answered	Successfully Answered	URL
15031	KEYCLOAK-1530	(see exception below). Not sure if the wrong schema is used (like public - but as far as I can see the code is ok) or if it's a quoting problem (realm vs "realm" - as described here: http://stackoverflow.com/questions/25692362/liquibase-with-posgresql-relation-issue). I would also like to suggest to rename the property "defaultSchema" (which is mapped to hbm2ddl.auto) to "schemaMode" or "schemaDDLMode" (or whatever reflects best what is done here) and the property "schema" to "defaultSchema" or "dbSchemaName". I just think using the name "defaultSchema" for the hbm2ddl.auto property is confusing. 13:34:46,077 ERROR org.keycloak.connections.jpa.updater.liquibase.LiquibaseJpaUpdaterProvider Change Set META-INF/jpa-changelog-1.1.0.Beta1.xml::1.1.0.Beta1::sthorger@redhat.com failed. Error: liquibase.exception.CustomChangeException: Failed to add realm code secret: liquibase.exception.UnexpectedLiquibaseException: liquibase.exception.CustomChangeException: Failed to add realm code secret at liquibase.change.custom.CustomChangeWrapper.generateStatements(CustomChangeWrapper.java:185) at liquibase.database.AbstractJdbcDatabase.executeStatements(AbstractJdbcDatabase.java:1208) at liquibase.changelog.ChangeSet.execute(ChangeSet.java:550) at liquibase.changelog.visitor.UpdateVisitor.visit(UpdateVisitor.java:43) at liquibase.changelog.ChangeLogIterator.run(ChangeLogIterator.java:73) at liquibase.Liquibase.update(Liquibase.java:200) at liquibase.Liquibase.update(Liquibase.java:181) at org.keycloak.connections.jpa.updater.liquibase.LiquibaseJpaUpdaterProvider.update(LiquibaseJpaUpdaterProvider.java:84) at org.keycloak.connections.jpa.DefaultJpaConnectionProviderFactory.lazyInit(DefaultJpaConnectionProviderFactory.java:144) at org.keycloak.connections.jpa.DefaultJpaConnectionProviderFactory.create(DefaultJpaConnectionProviderFactory.java:35) at org.keycloak.connections.jpa.DefaultJpaConnectionProviderFactory.create(DefaultJpaConnectionProviderFactory.java:25) at org.keycloak.services.DefaultKeycloakSession.getProvider(DefaultKeycloakSession.java:99) at org.keycloak.models.jpa.JpaRealmProviderFactory.create(JpaRealmProviderFactory.java:34) at org.keycloak.models.jpa.JpaRealmProviderFactory.create(JpaRealmProviderFactory.java:16) at org.keycloak.services.DefaultKeycloakSession.getProvider(DefaultKeycloakSession.java:99) at org.keycloak.models.cache.DefaultCacheRealmProvider.getDelegate(DefaultCacheRealmProvider.java:70) at org.keycloak.models.cache.DefaultCacheRealmProvider.getRealm(DefaultCacheRealmProvider.java:163) at org.keycloak.services.managers.ApplianceBootstrap.bootstrap(ApplianceBootstrap.java:40) at org.keycloak.services.managers.ApplianceBootstrap.bootstrap(ApplianceBootstrap.java:31) at org.keycloak.services.resources.KeycloakApplication.setupDefaultRealm(KeycloakApplication.java:157) at org.keycloak.services.resources.KeycloakApplication.(KeycloakApplication.java:87) at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:62) at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:45) at	0	0	https://issues.jboss.org/browse/KEYCLOAK-1530
15032	KEYCLOAK-1515	Just to clarify: would this feature enable a user that was federated via Activate Directory / LDAP to automatically log-in via their AD credentials? To me it sounds like it, since the user@corp email that came with the federated user attributes would implicitly imply to use of the configured AD identity provider (for @corp), or is there some additonal mapping for federation provider -> identity provider needed?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1515
15033	KEYCLOAK-1455	Do we not eventually need to get rid of it? Or rename it to isSecondFactorAuthEnabled?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1455
15034	KEYCLOAK-1445	it seems that there is an implementation for that in the keycloak github repository already. Unfortunatly one has to build it manually since it isn't yet consumable from maven central, I added a link to the implementation. What's the status of this? Is this tentative or already useable?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1445
15035	KEYCLOAK-1445	[~tdarimont] If it works fine we can include it in 1.7. I don't know the status of it though. Did you get a chance to try it out?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1445
15036	KEYCLOAK-1434	A few open questions here: * Should it be an icon, a label or both * Does [Edit] make sense? It's also to view.	0	0	https://issues.jboss.org/browse/KEYCLOAK-1434
15037	KEYCLOAK-1432	The blank-slate link is broken - did you mean: empty-state https://www.patternfly.org/patterns/empty-state/ ?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1432
15038	KEYCLOAK-1430	Is this issue meant exclusively only for form buttons and that too for the form header buttons ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1430
15039	KEYCLOAK-1427	[~agalante] Can you please just add a delete glyph icon and a vertical separator in realm selector menu as shown in screenshot 'keycloak-delete-client.png'. I can deal with the latter part of adding a github like modal to confirm deletion of realm. I tried adding delete glyph icon with the separator using custom css and raw values. But, IMHO, using raw values might be a poor practice. Moreover, that separator is a little pain to deal with. The best I could reach upto using raw and custom css is in screenshot 'realm_selector_menu.png'	0	0	https://issues.jboss.org/browse/KEYCLOAK-1427
15040	KEYCLOAK-1423	Query param can be modified by user so can be used as hint only. What about configuration option for Keycloak "Client"?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1423
15041	KEYCLOAK-1423	Why would it be a problem that query param can be modified by user? Configuration option for Client would probably be good as well (or as only option).	1	1	https://issues.jboss.org/browse/KEYCLOAK-1423
15042	KEYCLOAK-1402	Is this still a problem? We're just evaluating Keycloak 1.4.0.Release with Postgresql 9.4 and haven't seen any issues so far - thought we haven't done any migrations yet.	0	0	https://issues.jboss.org/browse/KEYCLOAK-1402
15043	KEYCLOAK-1384	Do we support federated logout for SP's using SAML libraries other than Keycloak ?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1384
15044	KEYCLOAK-1383	Currently our largest deployment (*.redhat.com) uses a separate application for user registration. We have added hooks into our PicketLink deployment to allow for apps like this to programtically authenticate a user after registration. Will KC support scenarios where the user registration can occur outside of KC itself?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1383
15045	KEYCLOAK-1383	should I create a new "feature request" or we can open this one ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1383
15046	KEYCLOAK-1382	Are there any plans to support delgating authn to the container (similiar to JBoss/Picketlink)? For example, we are using org.jboss.security.auth.spi.LdapExtLoginModule for our internal IdP deployment - this module lets us use a PAM passthrough to our OTP server for the majority of users, and at the same time define accounts with static passwords in LDAP for certain services, etc. Being able to use our existing jboss login-modules would make migrating from PL to KC easier for our external IdP as well.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1382
15047	KEYCLOAK-1377	Is there also the ability to do back-end lookups (based on authenticated identity), from LDAP or database for example?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1377
15048	KEYCLOAK-1376	Implemented this, but I'm not sure what you mean by "Ability to pass RelayState as additional parameter". Is this a configurable option? Or do you want to be able to pass it as a URL Query parameter? Or both?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1376
15049	KEYCLOAK-1326	Does this require firing/logging an event including the number of counts in case a user authenticates via google authenticator or FreeOTP ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1326
15050	KEYCLOAK-1267	[~stianst] I see this has been removed from 1.3.0 Beta. Does that it means that its not scheduled for any release coming soon? Do you guys have any plans of a way to approach this solution? If so I would like to provide help fixing this if u can give me some guidelines.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1267

#	Requirement	Questions	Answered	Successfully Answered	URL
15051	KEYCLOAK-1267	That would be really useful! As I see it one would need to do the following things: * Introduce new two new persistent properties to {{RealmModel}} ** rememberMeSsoSessionMaxLifespan ** rememberMeSsoSessionIdleTimeout (those need to be propagated to the Jpa, Infinispan, MongoDB model) ** RealmAdapters have to be adjusted ** when creating a realm initialize the rememberMeSso settings from the corresponding SsoSessionXXX settings. ** Add the two new columns to the liquibase jpa-changelog-2.2.x as a new changeset * use the new rememberMeSso timeouts in case of the clientSession has the rememberMe note set ** As far as I see it this would be in {{org.keycloak.services.managers.AuthenticationManager}}: *** {{AuthenticationManager#createIdentityToken}} *** {{AuthenticationManager#createLoginCookie}} *** {{AuthenticationManager#isSessionValid}} * expose the new rememberMeSsoSession settings in the admin-console ** see {{theme/base/admin/resources/partials/realm-tokens.html}} ** add labels, descriptions to the admin resource bundle {{theme/base/admin/messages/admin-messages_en.properties}} * Testing ** Verify that everything still works ** Add acceptance test(s) for user logins with remember me set... and check whether the correct cookies are generated Did I miss something?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1267
15052	KEYCLOAK-1250	for a sidebar which supports expand-on-hover (when slimmed), would it also be sensible to give it some transparency? Something like: !http://i.imgur.com/pX4QtxS.gif!	1	0	https://issues.jboss.org/browse/KEYCLOAK-1250
15053	KEYCLOAK-1204	[~prabhalar] This issue is rather vague. Is this required for identity brokering when external OpenID Connect IdPs are signing tokens with ECDSA? For HMAC I'm not sure how that would work as that would require having a shared secret between Keycloak and the external IdP.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1204
15054	KEYCLOAK-1201	Hi [~lohithac] I am also struggling with this issue where the "aud" claim is the name of the keycloak client and not an array of accessible resources. The workaround with the custom protocol mapper would be awesome, but I can't seem to configure it. Inside my realm I go to "clients > client > Mappers" and then click on "Create". As Mapper type I select "Hardcoded claim" but then I don't have the option to choose "Multivalued" or to choose an array as my Claim JSON type Can you give me some insight on how to proceed from this point?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1201
15055	KEYCLOAK-1195	Do you guys want to do a PR for this?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1195
15056	KEYCLOAK-1167	I'm looking for that right now - are the relevant openid connect urls listed here: https://github.com/keycloak/keycloak/blob/1c38bb7158fdeb9c2a094dcf8dfd15c3f41aa7f3/core/src/main/java/org/keycloak/constants/ServiceUrlConstants.java#L9 ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1167
15057	KEYCLOAK-1155	Can Kerberos support also be added to the authentication of DB Connections?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1155
15058	KEYCLOAK-1129	Stian - Can this also be implemented in 1.4 version along with Client Credentials grant? If not, I will be happy to get Client Credentials Grant first. Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-1129
15059	KEYCLOAK-1129	Stian - I agree with your explanation about the basic vs implicit flow but we will not be able to consider Keycloak for our enterprise needs unless support is provided for implicit flow. So can you make sure that this feature is implemented in 1.5 version without fail? Thanks once again, Raghu.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1129
15060	KEYCLOAK-1129	Hi Stian - Can we get this feature in 1.6, if not 1.5 ? Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-1129
15061	KEYCLOAK-1129	Hi Marek - While you are at this basic flow, I believe, with a little deviation even the Hybrid flow can be addressed which will make this a more complete implementation of OIDC. What do you think? I will be happy to create a JIRA for hybrid flow. Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-1129
15062	KEYCLOAK-1118	Doesn't make sense. Couldn't you create your own theme to implement that functionality?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1118
15063	KEYCLOAK-1109	Could you please let me know when I could expect having this in master? I'm asking because of implementing KEYCLOAK-1042. Thanks.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1109
15064	KEYCLOAK-1075	What we're currently doing for LDAP is update backend always when any setter is called. See WritableLDAPUserModelDelegate . It's not very great to call backend per each setter call and I have JIRA to address this for LDAP - KEYCLOAK-601 . The improvements to this could be creating custom type of KeycloakTransaction and then call something like: {code} keycloakSession.getTransaction().enlistAfterCompletion(new MyCustomTransaction()); {code} when MyCustomTransaction.commit method will invoke your REST backend. Once you enlist transaction by the call above, the commit callback will be called by Keycloak at the end of HTTP request, so update of REST backend is postponed and called just once per request. We are already using very similar approach for mongo or infinispan models. I will try to take a look if it's possible to use more "generic" task/transaction approach so that federation providers don't need to deal with writing similar code again. Until it's available, you can either call backend per each setter or write custom transaction by yourself. Does it work?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1075
15065	KEYCLOAK-1075	Can I expect this to be added anytime soon? I'm facing the same issue here, having a backend service which has several required fields and not being able to register users using the addUser method on UserFederationManager class. is there any other workaround apart from the transaction one? thanks	0	0	https://issues.jboss.org/browse/KEYCLOAK-1075
15066	KEYCLOAK-1074	We thought about javascript hack also, but it is not reliable, and mainly it doesn't work for Social registrations (eg Google, GitHub etc). I think realm option should be best solution as it can affect both frontend (hide username field from registration form) and backend (validations on register form submit, and mainly value passed to the "registerUser" method for both normal and social registrations). I think this feature is generic enough (brings simplified user registrations without additional username) to be added. As it looks not too complicated I should try to code it myself and contribute back. WDYT?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1074
15067	KEYCLOAK-1074	[~stianst] This doesn't really fit into custom claim work. I'm focused on mapping user data into token claims. This is a bit different as you want one input item to populate two user attributes. Would we really need to have something that maps registration form data into UserModel attributes anyways? I need to be in registration code anyways for custom attributes. I can do this work if you want...BUT....if you are itching to contribute, it is always nice getting somebody new as a contributor. I've scheduled this to be fixed for 1.2.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1074
15068	KEYCLOAK-1073	Input validation is critical, especially for a security application. Can you add some more explanation of how this is okay to be closed out? There is no explanation and the other item was closed as a duplicate. Is there comprehensive input validation now? Is it customizable?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1073
15069	KEYCLOAK-1069	Are you suggesting to use refresh tokens with a lifespan of 180 days (or so) to get a remember-me functionality? Until now I thought about refresh token usage like a session not like a remember-me cookie. The refresh token uses the session idle timeout, not the preferred session max lifespan (see KEYCLOAK-1068 for details). The remember-me cookie of the login page uses the session max lifespan as well.	0	0	https://issues.jboss.org/browse/KEYCLOAK-1069
15070	KEYCLOAK-1069	Could this be clarified a bit? We are also looking into this option. What are the implications of increasing the refresh token expiration times to a few months? If that's okay then it's an easy solution. If not, there should be a way to achieve "remember me"-like functionality with direct grant.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1069
15071	KEYCLOAK-1065	Hi Marek - Can you make it multiple OUs instead of just two (We have users in multiple OUs)? Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-1065

#	Requirement	Questions	Answered	Successfully Answered	URL
15072	KEYCLOAK-1064	Actually your NPE happens on this line: https://github.com/picketlink/picketlink/blob/v2.7.0.CR3/modules/idm/impl/src/main/java/org/picketlink/idm/ldap/internal/LDAPIdentityStore.java#l753 . Wonder if it could be just case-sensitivity issue like "createTimestamp" vs "createTimeStamp" ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1064
15073	KEYCLOAK-1047	Are these apps not adding support for OpenID Connect and SAML?	1	1	https://issues.jboss.org/browse/KEYCLOAK-1047
15074	KEYCLOAK-1045	[~bill.burke] look at the http://www.jboss.org/ - right side of header part contains link to login/register, if you are logged in then it is changed to user avatar image with link to user profile (by javascript which performs AJAX REST call to SSO server to get info about current user and change header is user is logged in). Now imagine that we use Keycloak for the SSO of this site, how can we implement this? Note that the site is static HTML server by plain Apache web server, no any dynamic server side is available.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1045
15075	KEYCLOAK-1045	We don't support doing this in an iframe as it opens for exploits (clickjacking, csrf). Why can't you redirect main browser, you could even do it before displaying anything and the user wouldn't really know (other than longer request time on first request)? I would assume on jboss.org you could have long-lived user sessions, which means it'd only redirect once every 30 days or something for a user, and only if they are already logged-in. There's only the one redirect, which if you store the refresh token in local storage would only be for the first request. You need to swap code for token, but that's a async request (which looks like that's how it's working atm). User profile is included in the token, so you don't need to do a separate request for that.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1045
15076	KEYCLOAK-1045	Are you interested in session management at all for anything hosted at jboss.org? i.e. Single logout, tracked app visits in a sign-on-session, admin initiated logouts, etc. Do you just use CAS to populate some cookie that you extract data from every request?	1	0	https://issues.jboss.org/browse/KEYCLOAK-1045
15077	KEYCLOAK-1045	I guess what I'm asking, do you need *any* of our token architecture? Maybe you are just interested in 1) Is the user logged in. and 2) who the user is.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1045
15078	KEYCLOAK-1045	The OIDC iframe currently passes realm-name and userid. Is that enough information Vlastimil? Maybe we need to expand what information the iframe has and can transmit (like we do any application). @Stian, I guess for performance reasons we need to support implicit flow.	1	0	https://issues.jboss.org/browse/KEYCLOAK-1045
15079	KEYCLOAK-1045	Added a POC to http://stianst.github.io/jbossorg/ backed by Keycloak on OpenShift. How it works: User not-logged in to Keycloak: # Load session iframe from Keycloak - this is a static resource and is cached in the browser (default max-age is 2592000, but is configurable) # As there's is no session cookie no more requests are made User logged-in to Keycloak, but not to static web pages: # Load session iframe from Keycloak (from browser cache) # As there is a session cookie the browser is redirected to Keycloak login pages to obtain code # Ajax request to Keycloak is used to exchange code for token - we can probably add support for the implicit flow to prevent this # Refresh token and access token is stored in sessionStorage if available or cookie if not - this makes it important to give reduced scope and claims to the application in the Keycloak admin console as storing refresh token in cookie/sessionStorage may be a small vulnerability User logged-in to static web pages: # Refresh token and access token is loaded from sessionStorage/cookie and keycloak.js is initialized with it # Load session iframe from Keycloak (from browser cache) # If session is valid (session cookie exists and id is same as in refresh token) the user profile is loaded from the access token. If the access token is expired the refresh token is used to retrieve a new access token from Keycloak, which returns updated user profile # If session is not valid, the flows from above are used Have a look at it and let me know if it fits your needs. I had to do some minor fixes to keycloak.js to support this, which is not in master yet, but added directly to keycloak.js on http://stianst.github.io/jbossorg/	1	1	https://issues.jboss.org/browse/KEYCLOAK-1045
15080	KEYCLOAK-1045	Hi [~stianst], I had no chance to look it yet as I have had a bit more important work on our UserFederationProvider implementation. Hope I'll be able to look it tomorrow. One question around "User logged-in to Keycloak, but not to static web pages:" case. Is user profile obtained as part of ajax request done in step 3? Or is another ajax request necessary to get user profile info? If only one request is used there then it will be probably OK, at least for first release. We will see what Pete Muir and his website guys will say to this solution ;-) Adding "implicit flow" which will prevent step 3 completely should be good still as it improves response times for users. Thanks	1	1	https://issues.jboss.org/browse/KEYCLOAK-1045
15081	KEYCLOAK-1044	Nice. BTW we should think about how often this SPI will be called, who and how can cache UserInfo. Point is that some informations in backend system can be updated often and we need this information to be immediately propagated to the clients. Which leads me to solution when Keycloak shouldn't cache data provided from this SPI at all (so SPI implementation will be responsible for caching), or it should be configurable how long Keycloak can cache data obtained from this endpoint (and maybe provide some REST endpoint where backend system can invalidate this cache for user?).	1	1	https://issues.jboss.org/browse/KEYCLOAK-1044
15082	KEYCLOAK-1044	Is that change still valid? I noticed that _UserDetailProvider_ is missing from master.	0	0	https://issues.jboss.org/browse/KEYCLOAK-1044
15083	KEYCLOAK-1042	I don't like changing details to <String, Object> it's intentionally strings to make it simpler to deal with. We shouldn't expose servlet or jax-rs classes. Further, we also need to make sure an event listener isn't able to change the request. How about adding "HttpRequest Event.getRequest" where HttpRequest has: * String getHeader * Map<String, String> getHeaders * String getQueryParam * Map<String, String> getQueryParameters * Cookie getCookie * Map<String, Cookie> getCookies We'd also create our own Cookie representation	0	0	https://issues.jboss.org/browse/KEYCLOAK-1042
15084	KEYCLOAK-1042	Trying to understand what you want here. What exactly are you going to do in the EventListener? Are you going to plug in your own and do something special? Or are you just going to parse this cookie information from a log file?	0	0	https://issues.jboss.org/browse/KEYCLOAK-1042
15085	KEYCLOAK-1024	Can you make that configurable? Let users of Keycloak choose it based on the risk level	0	0	https://issues.jboss.org/browse/KEYCLOAK-1024
15086	KEYCLOAK-1015	I propose the following solution: * Add 'accessCodeLifespanLogin' to Realm with default 300 seconds * Add 'Back to application' link on error page if application has baseUrl set ** Admin console can't set baseUrl atm, but we'll add that Would that fix your problems? [~gerbermichi] I'm moving this to 1.2.0.Beta1, it'll be released in a week or two and we don't have plans to do a 1.1.1.Final release.	1	1	https://issues.jboss.org/browse/KEYCLOAK-1015
15087	KEYCLOAK-996	Like this feature. Out of curiosity, can't we have the application choose the identity provider while registering with Keycloak?	1	0	https://issues.jboss.org/browse/KEYCLOAK-996
15088	KEYCLOAK-996	You mean, specify a default identity provider for an application ? So when an user tries to access this app he is automatically redirected to the default identity provider ?	0	0	https://issues.jboss.org/browse/KEYCLOAK-996
15089	KEYCLOAK-977	tested on EAP 6.3, 6.4-alpha and also WF 8.2: works fine	0	0	https://issues.jboss.org/browse/KEYCLOAK-977
15090	KEYCLOAK-971	Is there a good reason for this to be included in 1.9.3? We need to prioritize testing in 1.9.3 so this should be postponed to 2.x unless absolutely necessary in 1.9.	1	1	https://issues.jboss.org/browse/KEYCLOAK-971
15091	KEYCLOAK-953	What is the reason for the "Won't Fix"	0	0	https://issues.jboss.org/browse/KEYCLOAK-953
15092	KEYCLOAK-953	It was closed due to lack of demand and the initial use-case from the reporter has been resolved in another way. Is this a feature that you want?	0	0	https://issues.jboss.org/browse/KEYCLOAK-953

#	Requirement	Questions	Answered	Successfully Answered	URL
15093	KEYCLOAK-953	This feature would be very nice. Is someone already working on it?	0	0	https://issues.jboss.org/browse/KEYCLOAK-953
15094	KEYCLOAK-943	Has this been merged in yet?	1	1	https://issues.jboss.org/browse/KEYCLOAK-943
15095	KEYCLOAK-943	[~mtrue] No. It's still a work in progress. You can see what I have now at https://github.com/stianst/keycloak/tree/ACCOUNT-REST. I've got tests for the user profile and session endpoints. Also, CORS tests. I still need to add more (log, password update, etc.) and also add permissions tests.	1	1	https://issues.jboss.org/browse/KEYCLOAK-943
15096	KEYCLOAK-943	How can I get REST API for all operations available in the account management console. This includes view/update profile, reset password, view/update social links	0	0	https://issues.jboss.org/browse/KEYCLOAK-943
15097	KEYCLOAK-941	Bill/Stian - Can you please bump up the prioritization of this Jira and get it addressed in 1.3 Beta? This is turning out to be a major use case for a number of our clients and we require this functionality ASAP. If not, can you give me some pointers on how this can be implemented? I will give it a try once I find some time. Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-941
15098	KEYCLOAK-941	OK. Thanks Stian. In that case, along with certificates, can you also support Kerberos Keytabs? We create Kerberos System accounts for client applications to interact with each other - typically a Client application uses its Kerberos account and Keytab to authenticate against MIT Kerberos and obtains a service ticket applicable for the downstream application and presents it to it as a way of authenticating. Instead of service ticket, it can obtain an access token that can be used.	1	1	https://issues.jboss.org/browse/KEYCLOAK-941
15099	KEYCLOAK-941	Thank for taking this up in the 1.4 version. We also have a requirement of exchanging access tokens for a different scope of equal or lesser value. Our use case is as below: We have a portal like app that hosts links to a number of other service applications (service A, service B etc). This portal app has a number of components (web server, app server 1, app server 2) that reside on different hosts separated by firewalls. When a user requests content, the first component of the portal app gets an access token for a global scope (Service A, Service B Service C) which is then passed on to the other components (app server 1, app server 2). There will be validation of the token at each point and finally app server 2 (this is a sort of policy/authorization engine) will decide where to route the user request and then it will exchange the access token (scope : Service A Service B Service C) with the Authentication server (keycloak) for another access token with the scope downgraded to either "Service A" or "Service B" from "Service A Service B Service C" Can this also be included? Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-941
15100	KEYCLOAK-932	Has there been any talk of reconsidering this decision? LGPL dependencies (brought in by WildFly) are a real problem for us.	1	0	https://issues.jboss.org/browse/KEYCLOAK-932
15101	KEYCLOAK-912	Where should code for this be placed? https://github.com/keycloak/keycloak/tree/master/wildfly might be suitable. Any plans for hwo the widlfly CLI extension should look (which commands, which parameters)? Should the CLI interact with the keycloak-server subsystem directly or just use the REST endpoints via admin-client?	1	1	https://issues.jboss.org/browse/KEYCLOAK-912
15102	KEYCLOAK-904	Can this be prioritized? This would be useful for my application.	0	0	https://issues.jboss.org/browse/KEYCLOAK-904
15103	KEYCLOAK-904	Would it be possible to use offline tokens in deep links for example in emails? I mean the following use case: client application sends a notification email to user; user clicks a deep link with an "offline" token and is automatically authenticated and authorized with a reduced set of roles	1	1	https://issues.jboss.org/browse/KEYCLOAK-904
15104	KEYCLOAK-883	Hi Pedro - I was just trying out the Identity brokering using KC as SAML service provider against ADFS (using latest source code as of now) . Tried to upload the ADFS's metadata and it failed with the error message "File cannot be uploaded. Please verify the file". That metadata file is being used by multiple clients. Tried even other meta data files including KC as well (another instance that was setup as IDP). Even that failed. Can you please look into it and fix the issue? Logs 22:09:11,652 ERROR [io.undertow.request] (default task-15) UT005023: Exception handling request to /auth/admin/realms/SP/identity-provider/: java.lang.RuntimeException: request path: /auth/admin/realms/SP/identity-provider/ at org.keycloak.services.filters.KeycloakSessionServletFilter.doFilter(KeycloakSessionServletFilter.java:54) [keycloak-services-1.2.0.Beta1-SNAPSHOT.jar:1.2.0.Beta1-SNAPSHOT] at io.undertow.servlet.core.ManagedFilter.doFilter(ManagedFilter.java:60) [undertow-servlet-1.1.0.Final.jar:1.1.0.Final] at io.undertow.servlet.handlers.FilterHandler$FilterChainImpl.doFilter(FilterHandler.java:132) [undertow-servlet-1.1.0.Final.jar:1.1.0.Final] at io.undertow.servlet.handlers.FilterHandler.handleRequest(FilterHandler.java:85) [undertow-servlet-1.1.0.Final.jar:1.1.0.Final] at io.undertow.servlet.handlers.security.ServletSecurityRoleHandler.handleRequest(ServletSecurityRoleHandler.java:61) [undertow-servlet-1.1.0.Final.jar:1.1.0.Final] at io.undertow.servlet.handlers.ServletDispatchingHandler.handleRequest(ServletDispatchingHandler.java:36) [undertow-servlet-1.1.0.Final.jar:1.1.0.Final] at org.wildfly.extension.undertow.security.SecurityContextAssociationHandler.handleRequest(SecurityContextAssociationHandler.java:78) at io.undertow.server.handlers.PredicateHandler.handleRequest(PredicateHandler.java:43) [undertow-core-1.1.0.Final.jar:1.1.0.Final] at io.undertow.servlet.handlers.security.SSLInformationAssociationHandler.handleRequest(SSLInformationAssociationHandler.java:131) [undertow-servlet-1.1.0.Final.jar:1.1.0.Final] at io.undertow.servlet.handlers.security.ServletAuthenticationCallHandler.handleRequest(ServletAuthenticationCallHandler.java:56) [undertow-servlet-1.1.0.Final.jar:1.1.0.Final]	1	1	https://issues.jboss.org/browse/KEYCLOAK-883
15105	KEYCLOAK-883	One question - How do I get the metadata for the KC SAML service provider? Our IDP (ADFS) requires either a URL or metadata that should be uploaded. If there is no URL, is there any easy way of generating it? Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-883
15106	KEYCLOAK-883	Can you try from latest upstream ? Just pushed some changes.	1	1	https://issues.jboss.org/browse/KEYCLOAK-883
15107	KEYCLOAK-883	Thanks Pedro. Now I am able to upload the IDP meta data and create a provider. But when I click the link of IDP provider data to view the uploaded configuration, I get an error "Page Not Found". Can you fix that as well? BTW I can't do any more testing unless I know how to get the KC SAML metadata. Any pointers?	1	1	https://issues.jboss.org/browse/KEYCLOAK-883
15108	KEYCLOAK-883	Can you attach the metadata you are using ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-883
15109	KEYCLOAK-883	That is a big use case for us. Can I put in an enhancement request to process the Roles and populate that data using the new functionality that Bill is putting together under Custom profiles? If not, can you point me to some direction on how I can process the Roles? If needed, I will write my own processor and I should be able to plug it in? Regarding metadata, an end point will be very useful as there should be some interfaces through which the external systems can interact with KC in its capacity as Identity Broker. Generating one should be simple enough, as long as we know a few attributes like assertion consumer URL, Issuer name etc	1	1	https://issues.jboss.org/browse/KEYCLOAK-883

#	Requirement	Questions	Answered	Successfully Answered	URL
15110	KEYCLOAK-832	Was thinking of just generating a UUID and storing it as a user attribute of {clientId}.saml.persistent.name.id. Then we just re-use it if it has already been generated. Another question for you is, should Keycloak enforce a certain policy per client? Like for google client only allow transient, but for an internal enterprise app support others?	1	1	https://issues.jboss.org/browse/KEYCLOAK-832
15111	KEYCLOAK-828	Hi Raghu, Thanks for the feedback! I believe you can try it out if you build latest master and then clear browser cache. Clearing cache should hopefully fix "Error, Not Found" . Regarding (1), I don't know what exactly can be done for this? Just expect that token sent in "Authorization: Negotiate <token>" is not SPNEGO token but Kerberos token? I believe that can be done easily by using Oid of Kerberos mechanism instead of Spnego mechanism. I can add configuration option for allowing use different mechanism. Will it suits your needs? Regarding (2), what exactly is the usecase? Currently I have delegated credential from SPNEGO Authentication, which can then be theoretically used for further calls. See https://github.com/keycloak/keycloak/blob/master/broker/kerberos/src/main/java/org/keycloak/broker/kerberos/impl/SPNEGOAuthenticator.java#L87 . But if you have some use-case what exactly you want to support, it would help. Thanks! Marek	1	1	https://issues.jboss.org/browse/KEYCLOAK-828
15112	KEYCLOAK-828	Thanks Marek. That's similar to what I do at home to build and it works. My issue is that not all the libraries are downloadable at my Org. I managed to build the jar files (keycloak*.jar) and I also have the latest auth-server.war. Is there any way I can overlay all the keycloak*.jar and auth-server.war on the top of the existing 1.1 Final and would that work with some changes in the settings? If so, what are the steps?	1	1	https://issues.jboss.org/browse/KEYCLOAK-828
15113	KEYCLOAK-828	OK. finally got the latest code (as of now) working in my organization. Tried out the Kerberos functionality. Saw a login page that has a button for Kerberos. For SPNEGO, I didn't expect to see any login page as KC should do the negotiation and once it finds the proper SPNEGO token, it should allow the user to get in. But on clicking the button on the login page, got the below error: "Registration of new users is not allowed. Please ask admin to register you and login to account management to link the account" The logs just show the below. I think KC authenticated me as I got an error message earlier when I used the wrong principal. 22:30:49,162 INFO [stdout] (default task-3) Debug is true storeKey true useTicketCache false useKeyTab true doNotPrompt true ticketCache is null isInitiator false KeyTab is xxxxxxxx refreshKrb5Config is false principal is xxxxxxxxxxxx tryFirstPass is false useFirstPass is false storePass is false clearPass is false 22:30:49,167 INFO [stdout] (default task-3) principal is xxxxxxxxxxx 22:30:49,167 INFO [stdout] (default task-3) Will use keytab 22:30:49,175 INFO [stdout] (default task-3) Commit Succeeded 22:30:49,175 INFO [stdout] (default task-3) 22:30:49,203 INFO [stdout] (default task-3) [Krb5LoginModule]: Entering logout 22:30:49,204 INFO [stdout] (default task-3) [Krb5LoginModule]: logged out Subject Registration of new users is not allowed. Please ask admin to register you and login to account management to link the account. Tried setting up my user info in the realm but even that didn't help. How should I proceed? Any pointers? Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-828
15114	KEYCLOAK-828	The support for SPNEGO with MS IE "Integrated Windows Authentication" is going to be great. Looking forward to trying this off master as soon as it's ready. Does this support for Kerberos mean that one could integrate with a MS Active Directory LDAP server now without the user syncing to local Keycloak store? I like the current the LDAP sync support as Federation option, but could see certain enterprise users insisting on a check against the MS AD server on every authentication by Keycloak. Perhaps I misunderstand the sync however.	1	1	https://issues.jboss.org/browse/KEYCLOAK-828
15115	KEYCLOAK-828	[~prabhalar] I will try to test with MIT Kerberos today or tomorrow. From what I googled, the "Defective token detected" may be caused if you have some DNS issue. Maybe it will help to use virtual host like "www.myhost.org" (you can manually add it to /etc/hosts) and kerberos domain MYHOST.ORG. But I am just guessing.... Will have more info for you once I try it on MIT Kerberos [~guydavis] I am going to work on linking Kerberos with LDAP authenticated user. The users from MSAD domain will be able to login either via desktop (Kerberos ticket) or Keycloak login screen. In both cases it will be same user. For login via Keycloak login screen, it will use the credentials from MSAD server	1	1	https://issues.jboss.org/browse/KEYCLOAK-828
15116	KEYCLOAK-828	We do have createTimestamp attribute in our standard LDAP and probably some other attribute (can you let me know what attributes you are using for a search?) caused that issue. Thanks for suggestion on the temporary alternative. I think that should be doable, now that I have your implementation of LDAP+Kerberos as reference. I agree that perhaps you have to refactor the SPI into Federation (user info) and authentication (Kerberos, SecureID and others). That would be more intuitive.	1	1	https://issues.jboss.org/browse/KEYCLOAK-828
15117	KEYCLOAK-815	Is this a problem for the adapters in general or just admin-client?	1	1	https://issues.jboss.org/browse/KEYCLOAK-815
15118	KEYCLOAK-815	Hi, What are the results of the discussion on the mailling list? What is the suggested alternative to using the admin API? Accessing the admin API by REST calls?	1	1	https://issues.jboss.org/browse/KEYCLOAK-815
15119	KEYCLOAK-785	Is this implemented? It seems it is not possible to select a different key for signing JWT, than to encrypt the SAML authentication request. Both chooses the key with highest priority in the realm.	0	0	https://issues.jboss.org/browse/KEYCLOAK-785
15120	KEYCLOAK-782	I am interested in fixing this bug. I am a candidate for a GSOC stipendist potentially working on Certificate Management project. Could you assign this issue to me?	1	1	https://issues.jboss.org/browse/KEYCLOAK-782
15121	KEYCLOAK-725	In the example you have added yesterday you are using RealmResource update method ! RealmRepresentation rep = realm.toRepresentation(); rep.setSsoSessionIdleTimeout(123); rep.setSsoSessionMaxLifespan(12); realm.update(rep); Does this exist in 1.0.1.Final ? If not where can i get a nightly build. Again thanks for ur help.	1	1	https://issues.jboss.org/browse/KEYCLOAK-725
15122	KEYCLOAK-725	Good ! What will be the next release data if i may ask ?	1	1	https://issues.jboss.org/browse/KEYCLOAK-725
15123	KEYCLOAK-692	hmm... can't it be cache issue? For me it works as expected - ie when on http://localhost:8081/auth/admin/master/console/#/realms/master/applications/master-realm I am not seeing "Scope" or "Claims" tabs. Similarly when I switch application type to "Bearer-only" the tabs are hidden immediatelly as expected. But btv. found another minor issue with oauth clients during testing https://issues.jboss.org/browse/KEYCLOAK-693	0	0	https://issues.jboss.org/browse/KEYCLOAK-692
15124	KEYCLOAK-684	Hi Stian - Can you please comment on the timelines for this functionality? Thanks, Raghu	1	0	https://issues.jboss.org/browse/KEYCLOAK-684
15125	KEYCLOAK-654	Would it be enough to have a add additional credential form in the credentials tab of a user? This would allow admins to register a HW token secrets.	1	0	https://issues.jboss.org/browse/KEYCLOAK-654
15126	KEYCLOAK-599	Can you make a GUI available to choose different attributes for different client applications? This will allow us to choose all the attributes from LDAP and customize them based on client needs.	1	1	https://issues.jboss.org/browse/KEYCLOAK-599
15127	KEYCLOAK-598	Where should the user profile image be stored and are there limits on image dimensions / filesize / type? Instead of storing images, one could provide a configurable / dynamically parameterizable profile image url that could be referenced / resolved in the profile page (with appropriate origin / content checks of course).	0	0	https://issues.jboss.org/browse/KEYCLOAK-598

#	Requirement	Questions	Answered	Successfully Answered	URL
15128	KEYCLOAK-598	For that would you even need support in Keycloak? For example if app knows profile image is available at: http://mycdn/profiles/{user id} Why add anything?	1	1	https://issues.jboss.org/browse/KEYCLOAK-598
15129	KEYCLOAK-587	Rename admin console from Audit to event?	0	0	https://issues.jboss.org/browse/KEYCLOAK-587
15130	KEYCLOAK-586	Rename everything to event from audit? Admin console, etc.?	0	0	https://issues.jboss.org/browse/KEYCLOAK-586
15131	KEYCLOAK-565	Stian - I am presuming that the user can use the otp code from any of the devices and Keycloak would loop through all the registered devices and validate the code. Is that correct? Thanks, Raghu	1	0	https://issues.jboss.org/browse/KEYCLOAK-565
15132	KEYCLOAK-531	This would help the users of the UnifiedPush Server as well. Would it be possible to be done for the RC-1 ?	0	0	https://issues.jboss.org/browse/KEYCLOAK-531
15133	KEYCLOAK-522	stan is this done? Do I need to schedule this?	1	1	https://issues.jboss.org/browse/KEYCLOAK-522
15134	KEYCLOAK-450	Is there something in OpenID Connect for this? Or should be add something Keycloak specific?	0	0	https://issues.jboss.org/browse/KEYCLOAK-450
15135	KEYCLOAK-401	Can you add even Kerberos Keytab as a form of validating Service accounts? Using the keytab, the client would get a Kerberos token directly from the KDC and send in the authorization header which can be validated by the keycloak server to issue access/refresh tokens. Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-401
15136	KEYCLOAK-392	* Do you need Child type/ID, client ID, and User ID? Wouldn't all of those be covered by the Resource Type/ID? * Operation should really be Operation Type and be a static enum (READ, WRITE, UPDATE, DELETE). * Operation should be able to map to things like "CHANGED PASSWORD", "RESET PASSWORD", "LOGGED OUT ALL", "PUSHED NOT BEFORE", "INVALIDATED SESSION"	0	0	https://issues.jboss.org/browse/KEYCLOAK-392
15137	KEYCLOAK-392	I am not sure about tracking just single container and child. IMO it could be better to log the info about all the "parents" under the path to the actual resource (basically all the path variables). For example if someone invokes the endpoint for update protocol mapper (http://docs.jboss.org/keycloak/docs/1.2.0.Beta1/rest-api/admin/realms/%7Brealm%7D/applications-by-id/%7Bapp-name%7D/protocol-mappers/models/%7Bid%7D/index.html#PUT). I would log all of the: - realm - client - protocol mapper Also instead of having 3 things (Realm ID, Client ID, User ID), we can instead have single "auth" object, which will encapsulate all the info about authenticated user. In addition it could also encapsulate accessToken (or at least available roles of realmAdminApp). Basically it can wrap all the info like current AdminAuth object wrapps. This could be useful if people have issues with authorization, they can just send the log with their event to keycloak-user list and we can easily track which role (or scope) is missing. Another thing, do we really need OperationType? Isn't it sufficient (and less confusing) to just log the HTTP method of the request (GET, POST, PUT, DELETE)?	0	0	https://issues.jboss.org/browse/KEYCLOAK-392
15138	KEYCLOAK-349	Should we also include roles in token as scope rather than roles?	1	1	https://issues.jboss.org/browse/KEYCLOAK-349
15139	KEYCLOAK-349	I amazed that this isn't required for certification, given that currently Keycloak leaks personal information without permission (i.e. email address shouldn't be given out without the appropriate scope being authorized by the user). Is the fix for this going to be in the guts of Keycloak or is it something that could be implemented via a plugin? Is there anything I can do to help? Thanks.	1	1	https://issues.jboss.org/browse/KEYCLOAK-349
15140	KEYCLOAK-349	[~yaytay] If you switch "consent required" for your client and also for your "email" protocolMapper within this client, the consent screen will be shown to the user. Will it address your usecase?	1	1	https://issues.jboss.org/browse/KEYCLOAK-349
15141	KEYCLOAK-349	Hi, I am newby at keycloak, but I want to use the scope as above, can I do that using the admin GUI to set the mapper? I cannot get it to work can you help what to enter to get the SCOPE_PARAM thanks	0	0	https://issues.jboss.org/browse/KEYCLOAK-349
15142	KEYCLOAK-311	Bill/Stian - Regarding the population of custom attributes into a profile, can you make one of the Providers as the source of those attributes? Just to give you an idea of what we do (hopefully this is what even you are planning to do), we authenticate users against multiple stores like Kerberos, SecureID etc and use the subject (or principal) to retrieve all the attributes from other repositories like LDAP and DB which are then made available to our clients in the form of SAML or id_token attributes. The client has the capability to choose the available attributes and change the names to fit their needs during registration. Is that something that you can make available as a part of Keycloak. If not, can you provide some hooks to achieve that functionality? Thanks, Raghu	1	1	https://issues.jboss.org/browse/KEYCLOAK-311
15143	KEYCLOAK-301	I have analyzed this a bit, and ui_locales is part of OpenID Connect 1.0 specification. If we wanted to pass this flag to KeyCloak login page, we would need to adopt keycloak.js and other adapters. While it seems to be pretty easy to adopt keycloak.js to pass a new parameter (ui_locales) and put this to: OAuth2Constants, for me the main hurdle is how to do this with jetty and tomcat adapters. In those cases situation is much harder and I can only think of reading accept-language header flag as I don't know how to pass ui_locale query parameters. Is there any plan to implement a generic servlet based filter (adapter) for example that would work on slightly higher level than Tomcat's valves and give more control to developers what values they should pass, etc?	1	1	https://issues.jboss.org/browse/KEYCLOAK-301
15144	KEYCLOAK-301	[~matiwinnetou] Sorry, I haven't seen that you want to do that task. I've already started with the Realm config. Have a look at my code: [https://github.com/gerbermichi/keycloak/tree/i18n] Do you want to finish that task alone or together?	1	1	https://issues.jboss.org/browse/KEYCLOAK-301
15145	KEYCLOAK-301	[~gerbermichi] it's awesome that you've started working on this! Can you send an email to keycloak-dev on what you're doing, including some details, so we can make sure it aligns with what we need? Also, do you have an idea roughly how long it'll take you? I'm also hanging out on #keycloak on FreeNode if you fancy a chat.	1	1	https://issues.jboss.org/browse/KEYCLOAK-301
15146	KEYCLOAK-301	We don't want to require JS for these screens so the language selector needs to be converted to a css drop-down menu with links instead of a select * Is there somewhere such a css drop-down? Or should I create one? The 'ui_locale' should be 'ui_locales' and should support a list where we should iterate through the list until we find one that is available (see http://openid.net/specs/openid-connect-core-1_0.html#rfc.section.3.1.2) * OK sorry Found a few places where you do replace("'","'") why is this? * I use the MessageFormat to format the messages: {noformat} Within a String, a pair of single quotes can be used to quote any arbitrary characters except single quotes. For example, pattern string "'{0}'" represents string "{0}", not a FormatElement. A single quote itself must be represented by doubled single quotes '' throughout a String. For example, pattern string "'{''}'" is interpreted as a sequence of '{ (start of quoting and a left curly brace), '' (a single quote), and }' (a right curly brace and end of quoting), not '{' and '}' (quoted left and right curly braces): representing string "{'}", not "{}". {noformat} So, I introduced this workaround. Otherwise don't would be displayed like dont. But, I think I should remove this workaround and document that a single quote in the message.properties needs to be two single quotes. What do you think?	1	1	https://issues.jboss.org/browse/KEYCLOAK-301
15147	KEYCLOAK-301	[~stianst] I think I'm finished with this task. Do you have any open issues?	1	1	https://issues.jboss.org/browse/KEYCLOAK-301
15148	KEYCLOAK-292	can this be closed?	1	1	https://issues.jboss.org/browse/KEYCLOAK-292

#	Requirement	Questions	Answered	Successfully Answered	URL
15149	KEYCLOAK-228	Can this be configurable as well. I am thinking when we use KC in non saas environments, we might want to change it without having to wait and upgrade to the newer KC version.	1	1	https://issues.jboss.org/browse/KEYCLOAK-228
15150	KEYCLOAK-228	Currently the instructions for configuring a social provider is "hard-wired" into the admin console. In the future this should come from the JAR that implements the provider, so you could update the instructions by updating the single jar for the provider. Would that work, or do you have something else in mind?	0	0	https://issues.jboss.org/browse/KEYCLOAK-228
15151	KEYCLOAK-142	Why are the jpa entities using Strings for ids?!?! I am not able to deploy to any real database either. What is the intended database? I have tried MS SQL and POSTGRL. Both complain about string ids and deployment fails. This is a serious issue and the developers don't seem to be taking it seriously. They refuse to even reply to questions. Has there been some internal meeting where it was decided to treat their community with disdain? I work for a large government organization where I am trying to do a proof of concept and convince the powers that be to use JBOSS software. I am realizing it may not be such a good choice.	1	0	https://issues.jboss.org/browse/KEYCLOAK-142
15152	KEYCLOAK-117	What are the restrictions here? * implementation language * testing environment (AS7/EAP/WF) * test location * how stable is the ui itself? I could write some Arquillian + Graphene2 based tests (Graphene is Selenium on steroids), however that would be Java based tests. PhantomJS is support out of the box. I've noticed there are already Java based tests in testsuite. Arquillian also have support for QUnit, if that helps. Let me know if you'd be interested in such tests. If you want to have javascript based implementation for test to keep them next to admin-ui, npm + grunt + mocha + webdriver could be the combination you're looking for.	1	0	https://issues.jboss.org/browse/KEYCLOAK-117
15153	KEYCLOAK-108	This message could have the same style of the one in the Google Authenticator Setup page: http://ejsclient-cardosogabriel.rhcloud.com/saas-totp-setup.html . Agree?	1	1	https://issues.jboss.org/browse/KEYCLOAK-108
15154	KEYCLOAK-68	problem is in Keycloak or PicketLink. Basically it's caused by ApplicationEntity. When trying to load an ApplicationEntity from PicketLink, PicketLink tries to assign a Realm to ApplicationEntity.resourceUser. Full stack trace is: {code} org.picketlink.idm.IdentityManagementException: PLIDM000405: Could not load partition for type [class org.keycloak.models.picketlink.mappings.ApplicationData] and identifier [2-1383073500442]. at org.picketlink.idm.internal.DefaultPartitionManager.getPartition(DefaultPartitionManager.java:369) at org.picketlink.idm.internal.DefaultPartitionManager.createIdentityManager(DefaultPartitionManager.java:291) at org.keycloak.models.picketlink.ApplicationAdapter.getIdm(ApplicationAdapter.java:41) at org.keycloak.models.picketlink.ApplicationAdapter.addRole(ApplicationAdapter.java:156) at org.keycloak.models.picketlink.ApplicationAdapter.addRole(ApplicationAdapter.java:27) at org.keycloak.models.picketlink.RealmAdapter.addApplication(RealmAdapter.java:611) at org.keycloak.test.ApplicationModelTest.before(ApplicationModelTest.java:39) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:24) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:27) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:271) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:70) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:50) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:238) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:63) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:236) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:53) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:229) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:26) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:27) at org.junit.runners.ParentRunner.run(ParentRunner.java:309) at org.junit.runner.JUnitCore.run(JUnitCore.java:160) at com.intellij.junit4.JUnit4IdeaTestRunner.startRunnerWithArgs(JUnit4IdeaTestRunner.java:77) at com.intellij.rt.execution.junit.JUnitStarter.prepareStreamsAndStart(JUnitStarter.java:195) at com.intellij.rt.execution.junit.JUnitStarter.main(JUnitStarter.java:63) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at com.intellij.rt.execution.application.AppMain.main(AppMain.java:120) Caused by: org.picketlink.idm.IdentityManagementException: Could not create [6da6e9bb-6f33-470e-aed0-38a1cd315802 from entity [org.picketlink.idm.jpa.model.sample.simple.PartitionTypeEntity@7abdb100]. at org.picketlink.idm.jpa.internal.mappers.EntityMapper.createType(EntityMapper.java:202) at	1	0	https://issues.jboss.org/browse/KEYCLOAK-68
15155	KEYCLOAK-66	Can we close this Marek?	1	1	https://issues.jboss.org/browse/KEYCLOAK-66
15156	KEYCLOAK-65	When creating a new user through admin-console, if I add a user with same email as an existing user, the UI error message is "Error! An unexpected server error has occurred.". There's nothing in the UI to hint the user it's related to the duplicate email address. 14:06:04,588 ERROR [org.hibernate.engine.jdbc.spi.SqlExceptionHelper] (default task-12) Unique index or primary key violation: "UK_QYR7VP9OE4MRWM2JFMT5C5K7Q_INDEX_9 ON PUBLIC.USERENTITY(REALM, EMAIL) VALUES (/* key:3 */ null, 'dm@metcalf.net', null, null, null, null, null, null, null, null, 'keycloak-admin')"; SQL statement: update UserEntity set authenticationLink_id=?, email=?, emailVerified=?, enabled=?, firstName=?, lastName=?, loginName=?, notBefore=?, realm=?, totp=? where id=? [23505-173] [...]	0	0	https://issues.jboss.org/browse/KEYCLOAK-65
15157	KEYCLOAK-59	Is this issue solved? Where is the corresponding duplicate issue?	1	0	https://issues.jboss.org/browse/KEYCLOAK-59
15158	KEYCLOAK-24	I assume that this "logged in user" is your persona Linda, is that correct? Please let's refer to our users using their names, this is one of the uses of the personas. https://docs.google.com/document/d/1VHcllu37pCoiWiypzVkIMtq5pQloTX-3UB2ONyA_ZLM/edit	0	0	https://issues.jboss.org/browse/KEYCLOAK-24
15159	KEYCLOAK-17	Isn't this ticket related to the https://issues.jboss.org/browse/KEYCLOAK-24? Another doubt: Instead of giving access to "Google Account", we are giving access to "Keycloak". Is that correct? So basically create a similar page just changing the names? I guess I need to understand better how it works...	0	0	https://issues.jboss.org/browse/KEYCLOAK-17
15160	KEYCLOAK-16	Do we need to design a page for this? Aren't not those modal boxes provided by the services (aka Facebook, Google)? Attached is a screen asking for Facebook permissions. I was redirected for the facebook page (in background) in this case...	1	1	https://issues.jboss.org/browse/KEYCLOAK-16
15161	KOSMOS-25	If i wanted it slow, when could it be available you think :) ? Cheers, Tomek	1	0	https://issues.jboss.org/browse/KOSMOS-25

#	Requirement	Questions	Answered	Successfully Answered	URL
15162	KOSMOS-1	> Could you create a fisheye portlet which shows the latest changelog for a given CVS directory? Using page-scraping technique, it'd be very easy. I'm not exactly sure how would it play together with a CVS monitoring portlet? Are they replacement of each other? (Not really...) > Another possibility is to create a portlet which displays an arbitrary Fisheye QL query. A quick look at Fisheye API (http://fisheye.cenqua.com/api/) suggests that it's possible with their current API. I like this idea better, because it's a kinda generalization of the previous one, but seems a bit dangerous as the API is not solified yet.	0	0	https://issues.jboss.org/browse/KOSMOS-1
15163	LIVEOAK-745	Can you please followup on this.	1	0	https://issues.jboss.org/browse/LIVEOAK-745
15164	LIVEOAK-732	[~mwringe], here is a proposal for the Android chat. What do you think?	1	1	https://issues.jboss.org/browse/LIVEOAK-732
15165	LIVEOAK-732	awesome, its looking much better than before :) I was thinking of using the same application to support both authenticated and non-authenticated chats. Essentially instead of seeing the 'enter username' you would be directed to a separate login page (most likely a webview so there is nothing you would need to design for that). If we do want to support both authenticated (with separate login screen) and unauthenticated chats, would it make sense to have the enter username on a separate screen instead of as a dialog (http://developer.android.com/design/building-blocks/dialogs.html)? Trying to think about the consistency between the two apps. For the third screen screen, is there any way to determine what the users current username is? Also, if you don't want the 'send' button we can make it more like the web based one and send the message when the user clicks on the 'enter' key (or add a 'send' key) in the on screen keyboard. I have no idea which one has better UX on a touch screen and it doesn't really matter to me either way. Just letting you know that is a possibility. We could also use some app launcher icons, and icons to use for the notifications. There is some cool stuff regarding android notifications available at http://developer.android.com/design/patterns/notifications.html but we probably just need a few simple icons for it.	1	1	https://issues.jboss.org/browse/LIVEOAK-732
15166	LIVEOAK-732	[~mwringe], can you send me the specs for the icon image so I can generate the proper sizes? Also, for the notification, I believe we can use the same logo, but maybe without the shadows, right? Do you know the sizes for that too? Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-732
15167	LIVEOAK-732	I generated the 5 sizes for each icon, so you have the launcher, notification and small notification icon. You can get the files on: https://www.dropbox.com/sh/pxi0mcav3z79b4v/AACRatouaFLBTryy18XtSUMSa?dl=0 I believe the logo will look too small for the small notification icon. Can you send me a screenshot when you have it implemented? Thanks!	1	1	https://issues.jboss.org/browse/LIVEOAK-732
15168	LIVEOAK-732	That's great [~mwringe]. Can you send me some screenshots so I can take a look at the result? Thanks!	0	0	https://issues.jboss.org/browse/LIVEOAK-732
15169	LIVEOAK-723	[~kenfinni], I tried to run it but got authentication errors. How should I proceed to have it running?	1	1	https://issues.jboss.org/browse/LIVEOAK-723
15170	LIVEOAK-723	[~kenfinni], I updated the css to fix the "logout" link issue. I also updated the chat.js to include a "delete" class for the button, which I had forgotten. Can you check if that works now? https://github.com/liveoak-io/liveoak-example-chat-html-secured/pull/3 Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-723
15171	LIVEOAK-705	[~mwringe], can you point me the pages where this is happening? In the meanwhile, should the user be able to do other actions in the console or should we block the screen while the call finishes?	1	1	https://issues.jboss.org/browse/LIVEOAK-705
15172	LIVEOAK-705	I went to the storage configuration page and changed the Datastore from Default to Custom, changing the host and port to 1. I got the feedback that it was successfully saved instantly, but there was a sign on the list item indicating "Unable to reach storage server. Please check the configuration". So I believe I couldn't replicate your scenario. Or was it fixed to save instantly and inform the user that the config is wrong? Considering this to be an issue to different pages, should we display a loading feedback while liveoak is connecting to the database sever? Should this loading prevent the user to click anywhere else on the console, or is this a process that runs in background, so we just need a feedback that stays on the screen but allows the user to keep navigating?	1	1	https://issues.jboss.org/browse/LIVEOAK-705
15173	LIVEOAK-705	If you went to the configuration page and it allowed you to save without waiting for the server to respond back, then that is a bug :/ {quote} Considering this to be an issue to different pages, should we display a loading feedback while liveoak is connecting to the database sever? {quote} The console does not connect to the database server directly, it connects to LiveOak (which then connects to the database server). So it needs to wait for LiveOak to respond back, and since the console connects to LiveOak on pretty much every change, its not restricted to the storage page. {quote} Should this loading prevent the user to click anywhere else on the console, or is this a process that runs in background, so we just need a feedback that stays on the screen but allows the user to keep navigating? {quote} My initial opinion on this is that it should probably stop the user from clicking anything in the console. But maybe it might be appropriate to have some tasks run in the background.	1	1	https://issues.jboss.org/browse/LIVEOAK-705
15174	LIVEOAK-705	Adding [~alexandrem], who is also taking care of the UI. What are you thoughts about how we should care of the feedback for rest calls?	1	1	https://issues.jboss.org/browse/LIVEOAK-705
15175	LIVEOAK-693	Would this be related to LIVEOAK-406 at all?	1	1	https://issues.jboss.org/browse/LIVEOAK-693
15176	LIVEOAK-693	As the generator works, is this done?	1	1	https://issues.jboss.org/browse/LIVEOAK-693
15177	LIVEOAK-674	Right now, example applications have the "example":true field, so we can use that to determine if the example app was installed (the name is the same + example === true). but what to do once we know it? you can't install the example app, because the app with same name already exists. Is there a way for me to install this example app with some custom name?	1	1	https://issues.jboss.org/browse/LIVEOAK-674
15178	LIVEOAK-674	Agree. Can you ping me once the changes on the server side were done? Or just point me, what to fix there. I will then quickly fix the console part. Edit: I see now, that I have to edit the example field in application.json file per each example git.	0	0	https://issues.jboss.org/browse/LIVEOAK-674
15179	LIVEOAK-667	Several questions. # Why have a distinction between "Reports" and "Graphs"? Instead, couldn't the console present graphs and, if a report is valuable, simply offer a link to view a report of the same information? Immediately displaying the graphs gives one less step the user has to do before viewing useful information. # Can the time period be chosen by the user? If so, does the user select a date range, or do they select the window (e.g., last minute, last hour, last day, last week, last year)? # I think overlaying all information on a single graph is best, especially if the user can use a filter to specify which info they want to see. This also would reduce the steps required before actually seeing useful information. It's possible to present multiple graphs for different windows (e.g., stats for each of the last 60 minutes, or stats for each of the last 24 hours, etc.) but again all filtering would apply to all graphs. As a side note, how are all of these data going to be collected, and how are the stats (for the various time windows) going to be calculated and rolled up?	0	0	https://issues.jboss.org/browse/LIVEOAK-667
15180	LIVEOAK-667	Are we renaming "Dashboard" to be "Analytics", or is that just for the prototype?	1	1	https://issues.jboss.org/browse/LIVEOAK-667
15181	LIVEOAK-667	The scenario described a session called "Analytics". In my understanding we wouldn't have the dashboard anymore. Would we?	1	1	https://issues.jboss.org/browse/LIVEOAK-667
15182	LIVEOAK-666	Yes, I was thinking about giving the user the README.md file to be downloaded, just that. I didn't think about a viewer, do you think we need to care about this?	1	1	https://issues.jboss.org/browse/LIVEOAK-666
15183	LIVEOAK-654	[~mwringe], is this issue to support creating a storage resource by selecting a shared datasource instead of entering information directly?	1	1	https://issues.jboss.org/browse/LIVEOAK-654

#	Requirement	Questions	Answered	Successfully Answered	URL
15184	LIVEOAK-654	Can you take a look at how we could modify the "Storage" page, for new and edit, so that we can support the current state of capturing "host", "port" and "db" while offering the chance to select a "datastore" that defines the "host" and "port". Essentially it's a choice of specifying "host"/"port" or "datastore", where "datastore" is one of possibly many that can be selected. Not sure whether "datastore", "datasource", or an entirely different name is best to represent what is being captured.	0	0	https://issues.jboss.org/browse/LIVEOAK-654
15185	LIVEOAK-637	This is definitely fixed for a row, but the issue appears to be related to a column. [~gcardoso], was the issue meant to refer to row or you did mean column?	1	1	https://issues.jboss.org/browse/LIVEOAK-637
15186	LIVEOAK-599	We already have a generic registry for Resources implementation. https://github.com/liveoak-io/liveoak/blob/master/modules/common/src/main/java/io/liveoak/common/DefaultMountPointResource.java allows you to define a Resource that holds a map of resources that can be registered and unregistered from the resource. Do we need something beyond that?	1	1	https://issues.jboss.org/browse/LIVEOAK-599
15187	LIVEOAK-588	[~gcardoso], can you come up with a UX design for the best place(s) to link to Keycloak and Aerogear UPS admin consoles? Thanks Ken	1	1	https://issues.jboss.org/browse/LIVEOAK-588
15188	LIVEOAK-588	[~farrukhny], are you wanting the entirety of the other consoles to be "embedded" within the LiveOak console under those buttons, or simply opening them in another tab?	1	1	https://issues.jboss.org/browse/LIVEOAK-588
15189	LIVEOAK-588	Before jumping straight to a solution, I'd like to understand the requirement better. Does the user can currently access LiveOak / Keycloak / UPS with the same login? Is there a need to browse between the three applications? What would the user do in these three applications? [~mcarrano] was also concerned about the experience of changing between consoles, so I'm adding him to the thread.	0	0	https://issues.jboss.org/browse/LIVEOAK-588
15190	LIVEOAK-588	It doesn't seem that there is a strong need for such links. How about implementing them when there is a strong need for such flow?	0	0	https://issues.jboss.org/browse/LIVEOAK-588
15191	LIVEOAK-587	Your jira covers quite a few different areas, so I'll try and cover each of them. There are not any current plans to include all functionality of the Keycloak and Aerogear Unified Push Server consoles into LiveOak's console. There would simply be too much content for that to be seamless and workable. What we may do is provide links out to the other consoles to perform any management that isn't offered within LiveOak. Could you create a separate JIRA for each item you've mentioned (we already have one for Analytics), as that will make it much easier for us to separate whatever features you would like to see added? Thanks Ken	1	1	https://issues.jboss.org/browse/LIVEOAK-587
15192	LIVEOAK-584	Do you mean getting all the subscriptions that a Stomp Client has subscribed to?	0	0	https://issues.jboss.org/browse/LIVEOAK-584
15193	LIVEOAK-583	What can be done is: - Creating new UserProfile object which will have methods like getFirstName, getLastName, getEmail etc - New method 'getUserProfile' can be added to SecurityContext - Actually securityProfile is filled from calling keycloak endpoint and it's done here: https://github.com/liveoak-io/liveoak/blob/master/modules/keycloak/src/main/java/io/liveoak/keycloak/TokenResource.java#L51 . This is the place where userProfile info can be added as well from the keycloak accessToken. - This TokenResource is called from AuthInterceptor with local client call here - https://github.com/liveoak-io/liveoak/blob/master/modules/keycloak/src/main/java/io/liveoak/keycloak/interceptor/AuthInterceptor.java#L64 . So this is the place where retrieved profile data from ResourceState can be added to SecurityContext	1	1	https://issues.jboss.org/browse/LIVEOAK-583
15194	LIVEOAK-573	Is the proposal then to have ResourceSink not implement AutoCloseable?	1	1	https://issues.jboss.org/browse/LIVEOAK-573
15195	LIVEOAK-573	Why would it be the sole responsibility of the container to know when ResourceSink.close() can be appropriately called? Surely inside Resource.readMembers() a Resource would know when it's completed the work on the sink and can close it? Unless the Resource is swallowing exceptions, they would be thrown out before a ResourceSink.close() call anyway. I agree that removing AutoCloseable makes sense, but why the close() method too? How do we close a ResourceSink without it?	1	1	https://issues.jboss.org/browse/LIVEOAK-573
15196	LIVEOAK-534	Here's a scenario for being able to create/edit scripts. Could you check with [~mwringe] that I didn't mess up what should be entered for some of the values? Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-534
15197	LIVEOAK-534	Hi, I've created some mockups for the new script scenario. Questions: - If the name is optional, does the system attribute an automatic name in case of not specified? - A dropdown from 0 - 10 should be fine, I believe. Or do we expect to have more than 10 scripts? - Does the user need some guidance to fill in the resource path? - Is the script required? - Is it possible to import a file (script) or only write a script? - How is this Web IDE to write/edit the script? Is there a plugin for that?	0	0	https://issues.jboss.org/browse/LIVEOAK-534
15198	LIVEOAK-534	Another thing that I noticed is that it's not possible to specify the triggers in the new/edit script page. Is that correct [~mwringe] Also, should the user delete a script?	1	1	https://issues.jboss.org/browse/LIVEOAK-534
15199	LIVEOAK-528	Are you referring to the icon in the list of example applications, or having an icon in the application list?	1	1	https://issues.jboss.org/browse/LIVEOAK-528
15200	LIVEOAK-512	Ken, [~alexandrem], what do you think about this proposal? https://github.com/vrockai/liveoak/commit/a72d31ecf39d742d82363cdc283aa4ebe220a769 The code is commented so it's easier to get how it works. If you want to try it, build this branch, create first application and then check the Rest 1 and Rest 2 pages (available from the left menu). If you open another browser window and create new apps in there, everything should be propagated to the first window on Rest pages + in the application list in the left menu.	1	1	https://issues.jboss.org/browse/LIVEOAK-512
15201	LIVEOAK-512	I've tried to come up with better solution. I've just commited it to my LIVE branch. How it works? It stores the loaded data in a "cache" (which is actually a field in the $resource get method, you're trying to call). Once you use the LiveLoader it loads the data and registers the callback functions. After loading data it will create the "live" object inside, so you can reference to actual live data handled by callbacks. What this version does is actually "getting better use of subs", because if no sub event was fired, then we can assume that data haven't changed, so we will just return lastly loaded data. This version again has two pages Rest 1 and Rest 2 accessible from the left menu, but it presents both features: "cache" and "live data". If you create new application in another browser, the data in the "application list" in left menu are automatically updated. If you click between the Rest 1 and Rest 2 pages, no new requests to the server are fired (unless some sub event was fired in the meantime).	1	1	https://issues.jboss.org/browse/LIVEOAK-512
15202	LIVEOAK-512	Tried out your latest branch. I was expecting the application list in Rest 1 and Rest 2 to automatically update with new applications when they're created in another browser. Is that not setup to listen for subscriptions right now? Switching between the Rest screens I saw a network call to /admin/applications, is that to be expected?	1	1	https://issues.jboss.org/browse/LIVEOAK-512
15203	LIVEOAK-508	Maybe we could increase the services page with the languages supported / integrated solutions? Does that make sense? If we decide to have three boxes as secondary message, what could we emphasize, besides services? I'll check with Alex about the features list. Thanks!	1	1	https://issues.jboss.org/browse/LIVEOAK-508

#	Requirement	Questions	Answered	Successfully Answered	URL
15204	LIVEOAK-508	Attached a visual design proposal for the Homepage. Please download the file and open it in a browser to have an idea of how it would look like. [~alexandrem] is helping me with the Features session, selecting the items and working on the texts. What do you think?	1	1	https://issues.jboss.org/browse/LIVEOAK-508
15205	LIVEOAK-508	I like this a lot, looks really good. One thought I had was around making the logo more prominent. I'm not sure if it's possible, or easy, but maybe the large green area can be changed to be a "slice" through an enlarged version of the logo, so that it shows the trunk and/or branches super large. Does that make sense?	1	1	https://issues.jboss.org/browse/LIVEOAK-508
15206	LIVEOAK-508	Maybe use a "Fork Me" angled link near the top of the page?	1	1	https://issues.jboss.org/browse/LIVEOAK-508
15207	LIVEOAK-508	[~kenfinni], regarding the homepage, I tried to integrate the logo in the banner, but I couldn't find a good solution for this. [~vrockai], I included the "View on GitHub" button. For the release notes, we could make the text below the buttons as a link to it, and also have a link for it in the downloads page. Regarding the videos, It would be nice if we have frequent updates. Is that the case?	1	1	https://issues.jboss.org/browse/LIVEOAK-508
15208	LIVEOAK-508	[~gcardoso], would it be possible to incorporate something from http://design.jboss.org/liveoak/, such as the project banner, or a slice of the desktop wallpaper, or the Twitter or Google backgrounds? Might be possible to use the many overlaid logos from one of these instead of the green banner? We have reasonably frequent video updates, but in time it should become more frequent	0	0	https://issues.jboss.org/browse/LIVEOAK-508
15209	LIVEOAK-506	[~kenfinni], I added the link to open the application in both pages. Also, I updated the application details to cover also clients and logic, which where missing. What do you think?	1	1	https://issues.jboss.org/browse/LIVEOAK-506
15210	LIVEOAK-484	[~mstruk], did you implement this for Alpha02?	1	1	https://issues.jboss.org/browse/LIVEOAK-484
15211	LIVEOAK-480	I am proposing to use the same box style that we have for lists, like applications list, storage list. The difference here would be the "Show script", which expands the box. What do you think?	0	0	https://issues.jboss.org/browse/LIVEOAK-480
15212	LIVEOAK-480	Looks good! Does the description of the script disappear, or move somewhere, when the scripts is shown? Might be worth keeping the description above the script if its shown	0	0	https://issues.jboss.org/browse/LIVEOAK-480
15213	LIVEOAK-480	Looking good. Is 'FirstScript' and 'AnotherScript' the names of the IDs of the script? We should probably have both (since multiple scripts could potentially have the same name and description). I also don't see the priority or triggers/provides content. For this version, is everything just read only or editable? For a more complex version later on, we should see if we can get a nice javascript editor to use for the scripting part instead of just having it read only.	0	0	https://issues.jboss.org/browse/LIVEOAK-480
15214	LIVEOAK-480	Yes, those are the names. I can put the IDs too. How would an ID look like? How can I represent the priority? What do you mean by 'priority: integer specifying the priority'? I'll include something like Triggers: preupdate Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-480
15215	LIVEOAK-480	You mean, a dropdown to change priority? Or is it still read-only?	1	1	https://issues.jboss.org/browse/LIVEOAK-480
15216	LIVEOAK-477	What's your idea of the actual use-case? If we just add the arrows to the current solution, we will loose expected functionality (left/right to move cursor within the cell we are editing). So I assume the best way would be to do the same thing as in i.e. google spreadsheets. 1st click will just select the input and we would need another click (or start typing) to actually edit it. After 1st click, you would be able to use arrows to navigate, but after second click (= going to edit mode), you'll have the same functionality as now. If this is required we won't be able to quickly copy-paste values from table, since it would require 2 clicks for each cell. [~gcardoso], wdyt about this? Mimicking the google spreadsheet functionality would be a complex task to do, so we need to have clear requirements before the work starts.	1	0	https://issues.jboss.org/browse/LIVEOAK-477
15217	LIVEOAK-477	[~vrockai], I agree with what you said, I just didn't understand why "we won't be able to quickly copy-paste values from table, since it would require 2 clicks for each cell". It's not possible to paste a content when you have the cell selected (one click)? Having a solution like google spreadsheets would be great, but I think it is more a matter of priorities and cost/benefit here. Does the effort worth the benefit in the context that we are now? I'd let this to [~kenfinni]	1	0	https://issues.jboss.org/browse/LIVEOAK-477
15218	LIVEOAK-470	I'm not sure what you're suggesting here, can you explain? if you access /admin then the console is loaded.	1	1	https://issues.jboss.org/browse/LIVEOAK-470
15219	LIVEOAK-454	Do the examples that use push need to be updated in any way with the latest AeroGear UPS? Does http://liveoak.io/docs/install/ups/ need to be updated?	1	1	https://issues.jboss.org/browse/LIVEOAK-454
15220	LIVEOAK-454	I don't think we need to worry about the co-location scenario, as it would likely only be used in a local dev situation and not production. As long as we document that it's not possible to run UPS on the same WildFly instance as LiveOak, that should be fine. Will you be updating the instructions based on how you got it working?	1	1	https://issues.jboss.org/browse/LIVEOAK-454
15221	LIVEOAK-449	Since we are providing the possibility of creating a collection when there are no collections created, it seems to make sense to provide this same functionality when all the collections has been secured. So in that case, if the user clicks on Secure a collection (screen 04), we present the same modal. What do you think [~alexandrem]?	1	1	https://issues.jboss.org/browse/LIVEOAK-449
15222	LIVEOAK-449	This improves the overall user experience a lot, thanks! I have a few questions: * What happens if the application has more than one storage resource ? It should show a dropdown for picking the storage resource where the collection is to be created. * What happens if the application has no storage resources ? Should we allow to create one ? starts to be more complex.. * I think the ability to create a collection can be always present. It probably confuses the user that sometimes a functionality is available (when there are no collections or when all collections are secured) other times it's not (when only some collections are secured) , wdyt ? Plus, it's not likely that the user will only come to this page to secure a collection for the first time, so having all collections secured won't necessarily mean he'll want to create a new one to secure when he visits it.	1	1	https://issues.jboss.org/browse/LIVEOAK-449
15223	LIVEOAK-449	- Good point. In case of having more than one storage resource, he should choose into which storage he wants to create the collection. A dropdown works. - I believe to not have any storage resources is an edge case. I don't think we should invest in such complex logic for an edge case. For this case, we could proceed as we do today, tell them to create a storage resource first. WDYT? - I agree with your third point, but this will probably add some step in the flow when the user chooses to secure a collection (he will need to specify if he wants to secure an existing collection or create a new one). But I believe it won't hurt much :)	0	0	https://issues.jboss.org/browse/LIVEOAK-449
15224	LIVEOAK-444	[~mwringe], as we can specify security for a mongo database, does that cover this use case?	1	1	https://issues.jboss.org/browse/LIVEOAK-444
15225	LIVEOAK-414	What do you think [~kenfinni] and [~alexandrem]?	1	1	https://issues.jboss.org/browse/LIVEOAK-414
15226	LIVEOAK-414	[~gcardoso], looks really good. One thing I'm not sure on is "Type of source". Would "Import method" with "Local file system" and "Git" work? With respect to the message about an invalid application folder when importing from file system, do you expect the UI to transfer to the importing screen and then jump back to the previous screen when it fails for that reason? If so, that's fine, just wanted to clarify.	1	1	https://issues.jboss.org/browse/LIVEOAK-414

#	Requirement	Questions	Answered	Successfully Answered	URL
15227	LIVEOAK-414	Thanks Alex! Ok, let's keep them together. How about the badge? Will we allow application update or not? Sorry, I didn't follow you Ken.	0	0	https://issues.jboss.org/browse/LIVEOAK-414
15228	LIVEOAK-414	Makes sense. It would be a free input text field, right ? Or should we think about fetching branches from the repo and show in dropdown ?	1	1	https://issues.jboss.org/browse/LIVEOAK-414
15229	LIVEOAK-414	What's the POST look like? I hadn't seen an error like that in my testing	1	1	https://issues.jboss.org/browse/LIVEOAK-414
15230	LIVEOAK-409	I recall you mentioning that this might have already been done, is that correct?	0	0	https://issues.jboss.org/browse/LIVEOAK-409
15231	LIVEOAK-408	[~mstruk], what kind of configuration is possible for GridFS? Should the user expect objects within the "drive" to be visible in a collection of GridFS storage?	1	1	https://issues.jboss.org/browse/LIVEOAK-408
15232	LIVEOAK-408	Just had a quick look at the config JSON returned from gridfs vs mongo, and they seem to be identical. Where would a per user quota for gridfs be set? Is there a config value that can be used, which currently doesn't appear when retrieving the config because there is no value?	1	1	https://issues.jboss.org/browse/LIVEOAK-408
15233	LIVEOAK-406	Is something you wanted to put off until after your PTO?	0	0	https://issues.jboss.org/browse/LIVEOAK-406
15234	LIVEOAK-406	The example is here: https://github.com/vrockai/liveoak-example-ionic Would be nice if someone tested it before it's put into official example space. [~alexandrem], could you please spend some time on this? Big thanks!	1	1	https://issues.jboss.org/browse/LIVEOAK-406
15235	LIVEOAK-396	[~mstruk], thanks for looking into this, and providing a detailed analysis of the outcome. Is this even worth leaving as an open JIRA, or should we just close it as won't fix?	1	1	https://issues.jboss.org/browse/LIVEOAK-396
15236	LIVEOAK-389	[~mwringe], I think this JIRA came out of the F2F. Is there anything in particular needs to be done with the Mongo resource? Or is it more a case of just fixing things that might crop up?	1	0	https://issues.jboss.org/browse/LIVEOAK-389
15237	LIVEOAK-389	[~mwringe], is there anything required for this?	1	0	https://issues.jboss.org/browse/LIVEOAK-389
15238	LIVEOAK-388	Do we still want the examples to be deployable as separate cartridges? Or just as part of the main LiveOak one?	1	1	https://issues.jboss.org/browse/LIVEOAK-388
15239	LIVEOAK-388	Is there a benefit/advantage to having them separate? If there isn't a significant one, or several, then I'm ok with them just installing into the main cartridge.	1	1	https://issues.jboss.org/browse/LIVEOAK-388
15240	LIVEOAK-386	Ken, can you list all the resources you want to be secured?	1	1	https://issues.jboss.org/browse/LIVEOAK-386
15241	LIVEOAK-386	Would them be Storage, Push and Clients? Do the same options make sense to these other resources? We currently have: - Can access the collection (role) - Can create an entry (role) - Entry creator can (R,U,D) - Can read all entries (role) - Can update all entries (role) - Can delete all entries (role)	1	1	https://issues.jboss.org/browse/LIVEOAK-386
15242	LIVEOAK-386	In terms of design, we just need to update some labels and feedback message then. Will Alex or Villiam work on that to make it work? If it is only a screen change, I can do the changes myself. Otherwise, they can change them easily.	1	1	https://issues.jboss.org/browse/LIVEOAK-386
15243	LIVEOAK-386	This page calls for some questions. Right now it's only possible to secure a collection. When no collection is available, we're redirected to "Create new collection" page. When we add Push and Clients to the security list, what should be the default page? "Nothing to secure" or redirecting somewhere? Right now, in the dropdown with collection list we got an option "Create new collection...". I guess this will go away once we extend the page for Push and Clients. WDYT [~alexandrem]?	1	1	https://issues.jboss.org/browse/LIVEOAK-386
15244	LIVEOAK-386	I'm thinking that it might be confusing to mix up all things in the same UI. Thus, a decision step could be interesting. Something like: What would you like to secure? And give the options: Collections or Push or Clients. I can work on a prototype for that. What do you think?	1	1	https://issues.jboss.org/browse/LIVEOAK-386
15245	LIVEOAK-386	[~alexandrem], I worked on a prototype to include the new security resources. Since we have tabs, we can adjust the labels according to the resource. Ex.: - Storage - Can access the storage The interaction would be the same as we have now when the user is inside a tab, so he can still select the kind of resource he is in. What do you think?	1	1	https://issues.jboss.org/browse/LIVEOAK-386
15246	LIVEOAK-386	[~kenfinni] can you please validate these are the right rules to work with: {panel:title=Storage (root)} {noformat} { uriPattern: "/[app-name]/storage", requestTypes: ["CREATE", "READ", "UPDATE", "DELETE"], allowedRoles: [...] } {noformat} * Can access the storage (role) * Can create a storage (role) * Storage creator can (R,U,D) * Can read all storages (role) * Can update all storages (role) * Can delete all storages (role) {panel} ---- {panel:title=Storage (Collections)} {noformat} { uriPattern: "/[app-name]/storage/[col-name]", requestTypes: ["CREATE", "READ", "UPDATE", "DELETE"], allowedRoles: [...] } {noformat} * Can access the collection (role) * Can create an entry (role) * Entry creator can (R,U,D) * Can read all entries (role) * Can update all entries (role) * Can delete all entries (role) {panel} ---- {panel:title=Push} {noformat} { uriPattern: "/[app-name]/push", requestTypes: ["CREATE", "READ", "UPDATE", "DELETE"], allowedRoles: [...] } {noformat} * Can access push configuration (role) * Can create push configuration (role) * Push configuration creator can (R,U,D) {panel} ---- {panel:title=Clients} {noformat} { uriPattern: "/[app-name]/clients/", requestTypes: ["CREATE", "READ", "UPDATE", "DELETE"], allowedRoles: [...] } {noformat} * Can access the clients (role) * Can create a client (role) * Client creator can (R,U,D) * Can read all clients (role) * Can update all clients (role) * Can delete all clients (role) {panel} ---- {panel:title=Business Logic} {noformat} { uriPattern: "/[app-name]/logic/", requestTypes: ["CREATE", "READ", "UPDATE", "DELETE"], allowedRoles: [...] } {noformat} * Can access the scripts (role) * Can create a script (role) * Script creator can (R,U,D) * Can read all scripts (role) * Can update all scripts (role) * Can delete all scripts (role) {panel}	1	1	https://issues.jboss.org/browse/LIVEOAK-386
15247	LIVEOAK-386	With Storage (root), should most of the bullet points be about create/edit/read/delete a collection and not a storage? Do we need a way to define security at the App level? Just a thought I had. [~mwringe], can you take a look at the storage and logic access rules to see if they make sense?	1	1	https://issues.jboss.org/browse/LIVEOAK-386
15248	LIVEOAK-377	What exactly is this issue about? It seems like an overly generic issue, maybe we should close it and haven people open individual issues if a PUT operation is not working as expected?	1	0	https://issues.jboss.org/browse/LIVEOAK-377
15249	LIVEOAK-365	So, is the idea here to be able to define application.json templates on the server, so what when you create an application through the endpoint you can select which template you want to use? And also, alternatively be able to fully specify the application.json if you need to customize things. I guess kindof like how you can specify a full mongo server configuration for you application, or just select a system created datastore.	1	1	https://issues.jboss.org/browse/LIVEOAK-365
15250	LIVEOAK-361	[~gcardoso], could you provide some design guidance for [~vrockai] about how this should look? Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-361
15251	LIVEOAK-361	Is that because it's a lot of work and there are still lots of console issues?	1	1	https://issues.jboss.org/browse/LIVEOAK-361
15252	LIVEOAK-361	Now that I see a comparison, I think the sketch Alex did more accurately reflects what I've typically seen in UIs. It also got me thinking that for multiple columns, we need some way to specify AND/OR, so that we know whether all the columns are being AND'd or OR'd together. Right now I don't think we need to get complex by supporting both in a single search. [~gcardoso], was your design based on something that PatternFly has been working on, or just your ideas?	1	1	https://issues.jboss.org/browse/LIVEOAK-361

#	Requirement	Questions	Answered	Successfully Answered	URL
15253	LIVEOAK-350	What do you consider the advantages to be, as opposed to using awestruct as now?	1	1	https://issues.jboss.org/browse/LIVEOAK-350
15254	LIVEOAK-350	Can you elaborate on how LiveOak hosting a static content site would showcase it? I'm not following where you view the benefits to this approach.	1	1	https://issues.jboss.org/browse/LIVEOAK-350
15255	LIVEOAK-345	Are you suggesting that the console is deployed to myliveoak.com and not deployed at myliveoak.com/admin, or simply that we perform a redirect to the console when the root is accessed?	1	1	https://issues.jboss.org/browse/LIVEOAK-345
15256	LIVEOAK-340	Could you look at a good place in the Admin Console to provide a download of LiveOak's client.js file? Possibly from the dashboard somewhere?	1	1	https://issues.jboss.org/browse/LIVEOAK-340
15257	LIVEOAK-340	Is this liveoak.js client file related to all applications running on LiveOak or is this related to a specific app? What is exactly a client file? Do we need to say something (short text) or only presenting a button/link would be enough?	0	0	https://issues.jboss.org/browse/LIVEOAK-340
15258	LIVEOAK-284	Matt, I think this is resolved, can you confirm?	0	0	https://issues.jboss.org/browse/LIVEOAK-284
15259	LIVEOAK-278	[~gcardoso], am I correct in thinking that wireframe 03 and 04 are no longer required as we don't have the settings page any longer? Also, would it be beneficial to provide the ability to upload, etc from the dashboard, maybe in the right hand pane?	1	1	https://issues.jboss.org/browse/LIVEOAK-278
15260	LIVEOAK-278	It doesn't make sense anymore. How important is this functionality? I'm asking that to understand if only a link in the applications list is enough. We could have it in the dashboard too.	1	1	https://issues.jboss.org/browse/LIVEOAK-278
15261	LIVEOAK-278	Do you want me to update this prototype according com the new navigation, [~kenfinni]?	1	1	https://issues.jboss.org/browse/LIVEOAK-278
15262	LIVEOAK-277	Ken, before I propose wireframes for this, I'd like to have a better understanding of it. What exactly is a Single Page App? Is it something that will appear in the console or in the application for the final users? Also, do you think the dashboard is a good place to the upload button? Does Josh need to have an overview of the Single Page Apps that he has uploaded?	1	1	https://issues.jboss.org/browse/LIVEOAK-277
15263	LIVEOAK-277	Yes, Alex helped me on that :) What do you think about the wireframes?	0	0	https://issues.jboss.org/browse/LIVEOAK-277
15264	LIVEOAK-271	Was this done, or does it shift to Alpha2?	0	0	https://issues.jboss.org/browse/LIVEOAK-271
15265	LIVEOAK-266	I don't see this with the current code. Can you update to the latest and try again?	1	1	https://issues.jboss.org/browse/LIVEOAK-266
15266	LIVEOAK-262	Interesting. How would you see a new application being created? Currently there's a button for it on the application list page, but if we were to not have such a page then we'd need another way to create an application.	1	1	https://issues.jboss.org/browse/LIVEOAK-262
15267	LIVEOAK-262	Thanks for the feeback Marko. I agree, logging in at the first time and seeing a dashboard is not the experience I want for the console. I already had discussions with Ken about that, we need a special flow for the first time user, informing that one default app has been created, what he can do and so on, instead of throwing him in a dashboard screen. There is a Jira for that. Are you saying that when there is no applications, you see the page “applications list” empty? If so, that should not happen, since one application is created by default in the first login, isn’t it? What do you mean by available apps are not listed anywhere and available at a glance? The user is redirected to the apps list when there is more than an app right after logging in, and he can change application in the dropdown if he is inside one application. I would like to check if more people have problems to recognize the application link and the dropdown. I can try to make the application link more like a link instead of a label. Same thing with the dropdown button. One thing we must consider is that once the user recognises the application link as a link and the dropdown as a way to switch between applications, he won’t have more problems. That said, we should optimize the UI for future usages. Of course, it is desirable that the user understands the link and dropdown at first, but not that bad if he explores the UI a bit to get the picture, after all, this is something new to him. The problem I see if your proposal is that it brings more pollution to the design trying to optimize something that will be quickly understood (the dropdown interaction). Also, the accordion solution works better when there is different content to show below each option, I don't think it is a good pattern for this navigation. That said, I would stick with the current design. One good strategy is to apply a short survey with Alpha one users to get their insights. The more feedback, the better. Finally, regarding the Add Application button, that was discussed with some people from Patternfly, and we agreed that, since "create a new application" is not an action the user will do all the time, it is better to have it inside applications. This follows the concept of Progressive Disclosure.	0	0	https://issues.jboss.org/browse/LIVEOAK-262
15268	LIVEOAK-259	Oops, this is for the console, not for the resource implementation. Sorry, re-opening.	0	0	https://issues.jboss.org/browse/LIVEOAK-259
15269	LIVEOAK-243	This is a proposal with a visual language that best matches the console. I didn't have many elements to play with, so hopefully we can go with this for now. As I said, after we have a logo I can revisit it. What do you think [~stianst]?	1	1	https://issues.jboss.org/browse/LIVEOAK-243
15270	LIVEOAK-232	I know for url path of creating a storage, it was discussed on the mailing list (http://mail-archive.liveoak.io/New-Storage-screen-design-td724.html) that we'd add an info icon with a hint for that field. Created LIVEOAK-233 to cover this particular usage. [~gcardoso], what are your thoughts on adding hints to all input fields in the console?	1	1	https://issues.jboss.org/browse/LIVEOAK-232
15271	LIVEOAK-232	Ken, don't you think that hints would be good for the UPS config?	1	0	https://issues.jboss.org/browse/LIVEOAK-232
15272	LIVEOAK-214	[~gcardoso], is this complete, or are you still tweaking the HTML for it?	0	0	https://issues.jboss.org/browse/LIVEOAK-214
15273	LIVEOAK-207	Instead of having the users as a tab in security, we could provide access to them with a link "Users" in the left sidebar. What do you think [~stianst]? Also, it is quite confusing that editing the user inside an application affects other applications. Shouldn't we let the admin make only the changes related to that application in this page, and take him to a different place to edit the fields related to all applications? Do we plan on supporting more than a realm for liveoak soon?	1	1	https://issues.jboss.org/browse/LIVEOAK-207
15274	LIVEOAK-206	Ken, the idea is to display a link to security directly from the list of applications, so Josh can go directly to security, push or storage if he wants to. What would a "security policy" be? By "a more details configuration", which fields are you referring to? When Josh chooses "ACL", was he presented to a dropdown with some security options? In the "ACL" configuration, are all fields input text?	1	1	https://issues.jboss.org/browse/LIVEOAK-206
15275	LIVEOAK-206	I had a hangout to Marko and some interesting questions arrived regarding the terminology: - What constitutes a security policy? Is it what you call "ACL policies for resources" plus some "ACL entries for the application", or each "ACL policies for resources" / "ACL entries for the application" is a "security policy" which Josh sees in the "Policies" view? - I also didn't understand very well where the creation of "ACL policies for resources" and "ACL entries for the application" fit the flow to "add a new security policy". Are them add later or part of this creation flow? It is probably a good idea to have a hangout with [~stianst] tomorrow, since he wrote this specification. Thanks	0	0	https://issues.jboss.org/browse/LIVEOAK-206
15276	LIVEOAK-199	Looks good, and I like the idea of identifying which is an example app. Should it be a label to identify it, or have a separate (maybe collapsible) section for example applications deployed to LiveOak?	1	1	https://issues.jboss.org/browse/LIVEOAK-199

#	Requirement	Questions	Answered	Successfully Answered	URL
15277	LIVEOAK-199	Ok. So can I jump to the implementation of these screens? Ken, I'd need the list of example applications that we have, screenshots and some text about them. Can you help me with that?	0	0	https://issues.jboss.org/browse/LIVEOAK-199
15278	LIVEOAK-193	Are you thinking of storing whatever UPS configuration an application might need in liveoak only, or a copy of what UPS holds?	1	0	https://issues.jboss.org/browse/LIVEOAK-193
15279	LIVEOAK-193	My concern with this approach is that we're duplicating data between two systems. Is there a way for LiveOak to be notified of changes to the UPS configuration that occur directly in UPS? Granted we're talking about a short term solution here, but want to make sure that this doesn't become more short term than intended.	1	0	https://issues.jboss.org/browse/LIVEOAK-193
15280	LIVEOAK-186	Gabriel, I'd forgotten about keycloak realms as part of the new application process, which I added to the scenario this afternoon. Could you update the wireframe(s) to account for that change? Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-186
15281	LIVEOAK-186	Hi Ken. I finished the wireframes for both scenarios: new blank and basic app. Could you take a look at them?	1	1	https://issues.jboss.org/browse/LIVEOAK-186
15282	LIVEOAK-186	What do you mean by additional info?	1	1	https://issues.jboss.org/browse/LIVEOAK-186
15283	LIVEOAK-186	ah ok, do we need that for the wireframes or just the html impl?	1	1	https://issues.jboss.org/browse/LIVEOAK-186
15284	LIVEOAK-186	Gabriel, can we change "Blank" to "Do-It-Yourself"? Thomas suggested it and I think it makes it clearer the difference with "Basic"	1	1	https://issues.jboss.org/browse/LIVEOAK-186
15285	LIVEOAK-185	[~kenfinni]I like the blank option also. I guess the "devil is in the details" here, if the downstream tasks of configuring users, data, push, 3rd party services, etc are relatively easy, then I can see how blank would be preferable, even to the newbie. For me, the bottom line is - how easy is either path? basic vs blank - it could be that the end-user experience is nearly the same number of clicks, requiring approximately the same level of mental dexterity.	1	1	https://issues.jboss.org/browse/LIVEOAK-185
15286	LIVEOAK-185	Guys, do you think that presenting the dashboard after the creation of a blank application provides a good user experience? What will this user see if nothing is configured and no data will be displayed? Perhaps it would be more useful to tell him what he could do next. Same thing for the basic application. Won't the dashboard be filled only with zeros in the stats area? It would be better to tell him what to set up next.	1	1	https://issues.jboss.org/browse/LIVEOAK-185
15287	LIVEOAK-185	So my question would be, what do they need after creating these two kinds of applications? We could have links to what they need to do in this "what to do next..." page. Still better than a dashboard with zeros, and we wouldn't be putting additional roadblocks in their way.	1	1	https://issues.jboss.org/browse/LIVEOAK-185
15288	LIVEOAK-185	In the Basic App scenario, do you think Josh could not have all that info (Path, URL, ID, Secret) in the moment of the creation of a new app? Will we make all these fields required or can he move on without filling in some of them? Another question, after creating a basic app, how would the process of connecting an app be?	1	1	https://issues.jboss.org/browse/LIVEOAK-185
15289	LIVEOAK-185	Ken, for the KeyCloak realm field, is it a select box with some realms available to be selected (and application as default) plus the possibility to create a realm? Also, in the blank scenario, when you say that "the application is added to the default 'application' realm of KeyCloak", should we inform that to the user somehow?	1	1	https://issues.jboss.org/browse/LIVEOAK-185
15290	LIVEOAK-152	[~mwringe], is this covered by the system instances?	1	1	https://issues.jboss.org/browse/LIVEOAK-152
15291	LIVEOAK-148	How about this screen for the initial set up?	1	1	https://issues.jboss.org/browse/LIVEOAK-148
15292	LIVEOAK-147	Ken, if I understood correctly, to set up push notifications Josh should have already configured an application in Aerogear, is that correct? Are you predicting a scenario where Josh still hasn't configured an app in Aerogear? What are the "configuration" options shown upon saving the first piece of information?	1	1	https://issues.jboss.org/browse/LIVEOAK-147
15293	LIVEOAK-147	Do you think we need to go to read only after saving it? Isn't there a scenario where Josh needs to update this info, or maybe he just want to test with a different Aerogear config before putting the final config in the form? I'm asking that because we usually don't change to read only in Keycloak and in the Storage config. page...	1	1	https://issues.jboss.org/browse/LIVEOAK-147
15294	LIVEOAK-141	What's the flow for that? Click on the "Export Collection Data" button and present the OS dialog for the user to choose where to save the file?	1	1	https://issues.jboss.org/browse/LIVEOAK-141
15295	LIVEOAK-141	How about the option "Export data to JSON" above "Delete All Data", in the table dropdown?	1	1	https://issues.jboss.org/browse/LIVEOAK-141
15296	LIVEOAK-141	I think that works. Do you think we should provide the option to specify a filename to be exported to, or just export as "{collection name}-export.json" and let them rename it locally as they see fit?	1	1	https://issues.jboss.org/browse/LIVEOAK-141
15297	LIVEOAK-139	Is this an action related to a specific collection or import data creates a new collection?	1	1	https://issues.jboss.org/browse/LIVEOAK-139
15298	LIVEOAK-139	If Josh chooses to import it to the current collection, what would happen if the current data, if existing? It would be overwritten or would we add new lines in the end of the table?	1	1	https://issues.jboss.org/browse/LIVEOAK-139
15299	LIVEOAK-125	Will users use this same theme as the 'default' for login to their apps?	0	0	https://issues.jboss.org/browse/LIVEOAK-125
15300	LIVEOAK-125	Couldn't have formulated that any worse.. Let's try again: Are we creating one LiveOak theme for both the console, and also to be the default theme used by applications?	0	0	https://issues.jboss.org/browse/LIVEOAK-125
15301	LIVEOAK-125	Would it not make sense to just use the same theme for apps as well by default? That's how KC works, we use the same theme for admin console as well as the default for all realms. MADs should obviously be able to create their own, but we need some default theme available	0	0	https://issues.jboss.org/browse/LIVEOAK-125
15302	LIVEOAK-122	[~kenfinni] and [~mstruk], this is a very simple draft for the API Viewer. What do you think about a hangout to discuss the wireframes? I have some doubts and I'm sure you'll have lots of suggestions too.	1	0	https://issues.jboss.org/browse/LIVEOAK-122
15303	LIVEOAK-122	Do we have plans to work on this feature, [~kenfinni]?	1	0	https://issues.jboss.org/browse/LIVEOAK-122
15304	LIVEOAK-118	Marek, is this still an issue?	1	0	https://issues.jboss.org/browse/LIVEOAK-118
15305	LIVEOAK-112	Matt, I think this is actually done, can you confirm?	1	1	https://issues.jboss.org/browse/LIVEOAK-112
15306	LIVEOAK-104	Ken, how about this page? I included the "Edit Table" button, so the data would come as view-only initially. Is that the common use case, or should we display the table with editable fields right away?	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15307	LIVEOAK-104	Regarding your comments: - Ok, I removed the ordering functionality - It is a better practice to reinforce the collection where the user is with a title above the table rather than showing only the link in an active state. This leaves the context clearer for the user, and it won't take much vertical space. - I've updated the search to a filter. What do you think? - I've added the editing options. Some questions: - Is it a use case to remove a row? I've added this button - Whenever Josh needs to edit the table, will he make lots of changes or mostly edit a few strings? This proposal works better for the case of not having many strings to edit. I decided to present the items as view only initially because it is easier to read the data as text than in a bunch of inputs. Wait for the feedback. Gabriel	1	1	https://issues.jboss.org/browse/LIVEOAK-104

#	Requirement	Questions	Answered	Successfully Answered	URL
15308	LIVEOAK-104	How would the filtering look if multiple conditions were added? Does the header area grow in height to account for the larger space needed, or does something else happen? Yes, we definitely need the ability to remove a row. For adding a row, we don't need the concept of adding above or below a current row, we just add it to the end. So I think adding a row would work better as a button in header, as opposed to row selection. Most of the time I would not expect a large amount of data to be edited at one time. To make a field editable, do we need the click if saying you want to edit it, can we just have it that when Josh clicks in the field it becomes editable? What do you see the "Columns" drop down having values of other than "All"? It would certainly be nice to offer ability to show/hide columns, but don't know best way to portray that. We also need to be able to add/remove an entire column from the collection, and thus remove any data in that column for each row.	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15309	LIVEOAK-104	Yes, the header area grows to accommodate more conditions. I've put the actions "Add Row" and "Add Column" at the header. They are complementary. I updated the way to edit a field. When the mouse is over it, a background color appears and the cursor changes to the "editing text" cursor. This should be enough and less annoying than the tooltip. Regarding the "Columns" drop down, I thought on providing all the columns names multiple selection items, so Josh can search in one or more specific columns. I had understood that you asked me for this, but we can remove it if not necessary. To show / hide columns, we could display a dropdown icon when the mouse is over the column. By clicking on it, we could display options, like "hide column", "delete column" and "clear column", which you've just asked for. Will we allow sorting by column (A - Z and Z - A)? What do you think?	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15310	LIVEOAK-104	This is looking good. Regarding "Columns", that sounds fine. I guess I hadn't thought of it that way, so threw me in the wireframe as to what the options would be, but a multi select list seems reasonable. Would hiding a column then adjust the "Columns" multi select to be all columns, except the one that was just hidden? And therefore showing a column again means selecting it in "Columns"? What would "clear column" do? Remove the data from every row in that column? Makes sense to provide sorting A-Z and Z-A on a column. But I think it needs to be single sort only, and ordering a separate column resets everything and just orders that column. Couple of other things I've thought of. We probably need to provide the ability to delete the Collection from the resource and also to delete all data within the collection (while retaining the column structure). Do we need to provide a way to delete multiple rows? Or is it sufficient to be able to delete a single row or all rows?	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15311	LIVEOAK-104	As we're intending to display data as JSON, can we change all the plain text to be surrounded with double quotes? So 1 becomes "1", John becomes "John", etc. Also, I think we need to further show how JSON objects are handled. To that end, can we change email address to be "Address", and we just have one entry for someone that is something like: { "line1" : "54 My Street", "city" : "MyCity", "zip" : "34567", "country" : "USA" } This can be compressed to remove new lines and spaces, I just wanted to expand it here so you could easily see what it was. Not sure if this is possible to show on a wireframe, but we need the ability to shrink and expand the width of columns, and when editing a large JSON object, like the address, if it all isn't visually displayed when read-only we need to expand the column to see it all for editing.	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15312	LIVEOAK-104	I've updated the wireframe with the Columns dropdown. I believe the filtering should happen in all columns, not only the visible ones. Not searching the hiding columns might cause error, cause the user have forgotten that some columns are hidden. For the best UX, we should inform that some results belongs on a hiding column, and ask him if he wants to make those columns visible. But if this is too complex, we can make it simpler, like you proposed. Yes, "clear column" removes the data from every row in that column. Would you like to propose a better label? What do you mean by "single sort only" and that this action "resets everything"? On top of deleting a collection, shouldn't we provide an option to add a new collection? When the user deletes all data inside the collection, should we remove all rows? Probably leave only one row, without data and clear would be good, no? Regarding deleting rows, we need to consider how often the user will need to delete a column. If it is a frequent action and he needs to delete lots of rows at a time, a multiple selection would be better. However, consistency wise, it is better to keep the inline delete, since it is only possible to delete a single row at a time as well. I've updated the data to JSON. To shrink and expand the with of columns, we need to display a different cursor when the mouse is over the edges of the columns. In case of editing a cell with too much information, we could open the input or text area in a bigger box floating above the content.	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15313	LIVEOAK-104	"clear column" - my thought was whether its a common use case or not. What I meant by "single sort only" is that you can only sort on a single column at a time, as opposed to multiple sorts, ie. sort by ColA, then ColB, etc. So a result of that means selecting a different column to sort "resets everything" that was done to sort before. Yes, I'd forgotten about add collection, we definitely need that. Delete collection should probably remove all rows and not leave a blank one, as that may cause confusion that a row still exists with blank values (which is valid with JSON). For when the delete collection icon is visible, what do you consider to be part of the "area of the collection"? Just want to make sure its visible more than it isn't to make it clear how to delete the whole thing. I'm not entirely sure whether the more common use case is deleting multiple rows or a single one. We can leave it as is for now and see if we get it raised as an issue to address. I like the idea of a larger floating box for editing large amounts of JSON data. What do you mean by 'add new column / row comes with inputs'?	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15314	LIVEOAK-104	Some other services uses "Clear column" to refer to the action of deleting data from every row in that column. Ok, now I understand what you mean by "single sort only". That's the way I was thinking it would work. I've updated the wireframe with the "Add Collection" button. What happen after clicking on it? What's the difference between "Delete collection" and "Delete all data from the collection"? Do we need both features? I consider the "area of the collection" a imaginary rectangle around the title of the collection + table. By "add new column comes with inputs", I meant that after clicking on this button, a new column will be created and all rows will come with an input, ready to be filled in.	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15315	LIVEOAK-104	For the Add Collection flow, on top of the name, the user needs to specify the number of columns (and rows) at a certain time. For a best interaction, we could provide a "sample table" with 3 columns and 3 rows, for example (or another number that you find better). This is less annoying than having to go through another step that asks the numbers of rows and columns. About the "Delete all data from collection" button, I suppose that this is an unusual action, is it? If so, we should keep it more hidden than "Add row" and "Add column". I was thinking about the search/filter. We are providing many UI elements for this action in the beginning. I believe we could offer a simple search box that searches the string in all columns, and provide a link "advanced filter" that opens a box with the columns selector, conditions, possibility to add conditions etc. What do you think?	1	1	https://issues.jboss.org/browse/LIVEOAK-104
15316	LIVEOAK-103	Hi Ken. I believe this scenario was partially covered by the "Add new data storage scenario", wasn't it? https://issues.jboss.org/browse/LIVEOAK-85 Could you talk more about this concept of "collections" that are new to me? What is a collection? How many collections will be available per storage? Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-103

#	Requirement	Questions	Answered	Successfully Answered	URL
15317	LIVEOAK-103	Some questions: When you talk about "selecting the storage", you mean that Josh could have many storage resources inside the app and HAS TO select only one from the list of storage? My design only provides the list without the possibility of choosing one... From the "Storage List", it seems possible to access the "Collections" page and also the "Edit Storage" page. Which page will be accessed more often? We need to define better how to redirect him to these pages from the Storage list. In the collections page, if Josh needs to select a collection, we probably should display the collections initially collapsed, since they can be long, right? In case of a collection have too many rows, will we provide continuous scroll or pagination? Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-103
15318	LIVEOAK-103	Yes, viewing data within a storage resource requires choosing one from possibly many. I would expect the ability to view the data within the collections of a storage resource will be used much more frequently than viewing/editing the configuration of a storage resource. Not sure the best way to represent it, but my thoughts were that we'd have a sidebar section, or the like, where the list of collections is present. Then selecting one of those then displays the data grid for that collection to the right of the collection list. I probably didn't make that super clear in the scenario, sorry. The idea is to be able switch which collection you're looking at easily, without trying to see all the data for all collections at once. I might lean towards continuous scroll, but am not stuck on either. Would depend on how we show options to remove rows/columns, add rows/columns, etc. Does that clear it up? Also, I've just realized I probably didn't make it clear that I think View and Edit Data are essentially the same page, just different purpose.	0	0	https://issues.jboss.org/browse/LIVEOAK-103
15319	LIVEOAK-102	Is this still needed? Or is the fact we provide a lot of the functionality across various examples ok?	1	0	https://issues.jboss.org/browse/LIVEOAK-102
15320	LIVEOAK-92	Ok, I won't put the progress bars in the boxes on my visual design proposal. Regarding the main box, I didn't want to have four equal boxes just because something asymmetrical is more interesting. 5 boxes would work fine as well, with the two that you mentioned at the top and three below. Do you think we have info for a 5th box? I haven't thought about the detailed screen for each box. This can be done in the future, we can stick only with the dashboard for now. I am not sure if I understood how the sidebar would be with many resources. Could you play a bit with Balsamic to illustrate this case? This will make easier for me to decide on how to handle this case. Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-92
15321	LIVEOAK-92	Should we use the term "View all Storage"? I think we agreed that Storage sounded better than Data Resource on another screen. Do you mean that we'd only show the 3 last used storage resources under the Storage section, or something else?	1	1	https://issues.jboss.org/browse/LIVEOAK-92
15322	LIVEOAK-86	Gabriel, Here's some more feedback on the data screens: - For the right hand section of the new storage page, I don't think a list of apps using a given resource is relevant anymore, as a single resource will only ever be used by a single application. What might be useful is to have the database name that has been configured for use with that storage resource. - Not sure that "Data Resource" is an appropriate name for the collapsible section where mongo, etc is selected. Maybe "Storage Provider"? - "Doesn't have it installed?" should probably be "Don't have it installed?" - Within the "Settings" collapsible section, I'm wondering if "URL Path" should come before "Database Name", as then all the configuration that specifies a particular database (name, host, port) is together as its all related to what Mongo database you want to use, whereas "URL Path" is how you want it exposed in LiveOak.	0	0	https://issues.jboss.org/browse/LIVEOAK-86
15323	LIVEOAK-86	If we just replace the "Apps using it" section with "DB Name". So the MyFirstStorage resource is connected to MongoDB with DB Name "MyDB". As we have a resource which connects to a particular Mongo instance and database, but each resource that connects to Mongo could be connecting to different databases within the same instance (host/port). Does that make sense? My only uncertainty is whether showing DB Name here would be too much information. Thoughts [~bdaw]?	1	1	https://issues.jboss.org/browse/LIVEOAK-86
15324	LIVEOAK-86	I'm not getting it, I probably need to understand better the relation between a storage resource, the storage provider and the database name. So one app can have multiple storage resources. But isn't true that one storage resource is only associated to one storage provider and on database name? If so, what exactly do you want to show in the sidebar?	0	0	https://issues.jboss.org/browse/LIVEOAK-86
15325	LIVEOAK-86	If we put this information in the sidebar, aren't we replicating the information presented in the storage list page (previous page)? Let's take a step back and try to think about the usage scenario. When Josh is creating or editing a new resource, how important is to see if there are other storage resources and detailed information about them, considering he already came from a previous screen containing this data (storage list page)? Is that really helpful to him? If not, what would be more helpful?	1	1	https://issues.jboss.org/browse/LIVEOAK-86
15326	LIVEOAK-85	I illustrated the case when he has only one application and is redirected to the application Dashboard directly. Once in the dashboard, he clicks on "Data Settings" (right sidebar) and is lead to a list of the data resources from that application. This page also provides a lateral navigation for the other areas related to that application. By clicking on "Add Resource", Josh accesses a new page to add a new resource. A frequente discussion on Interaction Design lists is whether we should redirect the user to the (resources) list after saving or the keep him on the resource page. If we take him to the resources list, it would be good to highlight the modified resource from the list for a few seconds, on top of the success feedback, so the user knows where the item is on the list. This might be more complicated in case of many resources or pagination. How many resources do you think an user will have on average?	1	1	https://issues.jboss.org/browse/LIVEOAK-85
15327	LIVEOAK-85	The list of data resources and creating a new data resource look good. My only thought, or concern, is that on the data resource list and add pages the breadcrumb for "Applications -> MyFirstApp -> Data" is nearer the center of the page. On the dashboard it's on the very left. Wondering if that will be confusing to change where the breadcrumb lives? Should the breadcrumb on list and add pages be below the divider under the app selection dropdown, as it is on the dashboard?	1	1	https://issues.jboss.org/browse/LIVEOAK-85
15328	LIVEOAK-84	Perhaps the possibility of adding a new source on the dashboard is too specific. I was thinking about providing a link like "Data Settings", and then he can do that in the data settings page. What do you think?	1	1	https://issues.jboss.org/browse/LIVEOAK-84
15329	LIVEOAK-81	Gabriel, Could you actually add the balsamiq mockup to the jira, not just a png image? That way any updates can be handled in the jira too. Thanks Ken	1	1	https://issues.jboss.org/browse/LIVEOAK-81
15330	LIVEOAK-81	I've updated the files. Is that correct now? Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-81
15331	LIVEOAK-81	Gabriel, these look good. Couple of questions: - For screen 2, would you see the items listed under the app, such as "Mongo DB" and "Secured", being links to the details for that particular item? Kind of a shortcut to bypass app dashboard? - For both means of editing the application name, would it be better to do away with the idea of a popup and simply have the cursor change to a text cursor when hovering over the name, and clicking converts it from a read only field to an editable one? With the enter key capturing the change?	1	1	https://issues.jboss.org/browse/LIVEOAK-81
15332	LIVEOAK-81	Ken, I created a new wireframe, "6 - Application Settings". What do you think about it?	1	1	https://issues.jboss.org/browse/LIVEOAK-81

#	Requirement	Questions	Answered	Successfully Answered	URL
15333	LIVEOAK-81	Just had another look at Keycloak admin and think I understand better now. So the application role list is just a pure list of current roles for the application, with ability to add/remove as needed. Default roles is then as the wireframe, where the left hand of available roles is based on the application role list, and you select which roles are the default. Client scope left hand of available roles is also based on the application role list? Did I get that right [~stianst]?	1	1	https://issues.jboss.org/browse/LIVEOAK-81
15334	LIVEOAK-81	[~stianst], when you say that "application roles is a list of roles where you can add/remove roles to it" you are talking about that table in Keycloak that leads to another page to create a role, is it? Considering that Josh can add multiple clients, it wouldn't be a bad idea to split this page in two: General Settings and Clients. What do you think Ken?	1	1	https://issues.jboss.org/browse/LIVEOAK-81
15335	LIVEOAK-81	Do you mean making the configuration page, into multiple pages? Not sure how that would look, so maybe put together a quick wireframe to show it?	1	1	https://issues.jboss.org/browse/LIVEOAK-81
15336	LIVEOAK-80	Any idea how are we going to name the app by default? NewApp, NewApp1, NewApp2? MyApp? NameMe? :)	0	0	https://issues.jboss.org/browse/LIVEOAK-80
15337	LIVEOAK-77	Where is the scenario description?	1	1	https://issues.jboss.org/browse/LIVEOAK-77
15338	LIVEOAK-77	Gabriel, could you attach the balsamiq mockups instead? Should be able to export the XML from where you created them and then import them into the balsamiq extension for jira.	1	1	https://issues.jboss.org/browse/LIVEOAK-77
15339	LIVEOAK-77	Few commens. What I like: - All needed info there. Stats, app keys, app status, and links to data and security config. What may be missing - Link or banner of some tutorial. Could be placed around the stats. Like "How to start building your app". This could take to a separate section or modal listing tutorials with different topics On the first sight there is a bit of a lot of info... especially with those 2 additional popups. Just an impression. Still I like the fact that app keys and status of base services is exposed right away and doesn't require looking around.	0	0	https://issues.jboss.org/browse/LIVEOAK-77
15340	LIVEOAK-77	I added the link for the tutorials. How important is to inform the user that "we have created an app for him"? I guess this is part of your "magic" and we should reinforce it. To minimize the informational load, I propose hide the content below a dark background and only display the initial information (2 First Application Dashboard - Welcome). After clicking on "OK", the tooltips are hidden and the user goes to the page (last image). What do you guys think?	1	1	https://issues.jboss.org/browse/LIVEOAK-77
15341	LIVEOAK-76	I liked the way you wrote the scenario (concise, straight to the point) and the topics at the bottom. Some comments: We had discussed the idea of NOT need to log in as admin to be asked to change the password right after. Did that change or it was just a mistake in the scenario description? (https://gatein.mybalsamiq.com/projects/liveoak/First%20login) Does Josh have to change the application name at some time after this "default" app has been created? If so, I believe that a initial screen welcoming him and asking him to name and create an application is more intuitive because he knows that he is creating an app and expects to see some feedback and information about the app right after. In the scenario you described, it is unexpected and maybe confusing (what is this? I just logged in, I already have an application?) (https://gatein.mybalsamiq.com/projects/liveoak/Create%20first%20application) Do you have a clear vision of what pieces of information goes inside the application dashboard after the app is created? I would need to have a list of data and functional elements to be able to prototype it. Thanks	1	1	https://issues.jboss.org/browse/LIVEOAK-76
15342	LIVEOAK-76	Hi Ken, Today I presented the dashboard for this scenario to other JBoss scenario and we realised that maybe this approach is not the best one for someone who just starting on LiveOak. Right after logging in, to present statistics with big numbers showing a zero has no value for the user. For this scenario, instead of showing numbers, we probably should guide the user, telling him what he can or has to do. In this sense, I believe the previous initial screen that I prototyped work better: https://gatein.mybalsamiq.com/projects/liveoak/Initial%20overview In my understanding, after logging in at first, the user should be informed that an application with data, security and push has been created and see what to do next. That's the part we are missing. I believe application details are important too, but you can confirm. What are your thoughts? Gabriel	1	1	https://issues.jboss.org/browse/LIVEOAK-76
15343	LIVEOAK-76	It would be magic if we auto create everything you mentioned. Won't Josh really need to set up anything at all? The best thing to do is tell him what he has to do next, so if he has to connect his mobile app to LiveOak, I agree that the tutorial would be our best option.	1	1	https://issues.jboss.org/browse/LIVEOAK-76
15344	LIVEOAK-76	Ok! And do you guys know who will write this "how to connect your mobile app to LiveOak" tutorial? Do you think we should update the wireframe with some sample text or wait for the real content to do the design?	1	1	https://issues.jboss.org/browse/LIVEOAK-76
15345	LIVEOAK-56	[~mposolda] is this still required? Or is this not needed because its covered by a rewrite of uri policies?	1	0	https://issues.jboss.org/browse/LIVEOAK-56
15346	LIVEOAK-51	I believe that after merging this PR https://github.com/liveoak-io/liveoak/pull/67 the issue can be resolved. What's the PR merging policy? should be merge our own PRs?	1	1	https://issues.jboss.org/browse/LIVEOAK-51
15347	LIVEOAK-22	Shouldn't this wait until we have LIVEOAK-31 in place?	0	0	https://issues.jboss.org/browse/LIVEOAK-22
15348	LIVEOAK-17	What exactly are you trying to do as part of this issue?	0	0	https://issues.jboss.org/browse/LIVEOAK-17
15349	LIVEOAK-10	Is this done per the change from [~bob.mcwhirter] that changed DirectConnector?	0	0	https://issues.jboss.org/browse/LIVEOAK-10
15350	LIVEOAK-1	Currently we place somewhat fast-and-loose with updating resources. The current update(state) method on current Resource implementations is ostensibly a partial-update-allowed sorta implementation. Should we tighten up that contract/docs to indicate that update(state) is a full-replacement call? And then provide for an update(diff) method that allows partial updates? I don't think simply applying the rule of "update the fields mentioned in the provided state" is sufficient for PATCH, since ultimately, you then don't have a way to unset things, but only change their value to other non-null values. ie, if a resource is ostensibly: {code}{ name: 'bob', favorite-food: 'tacos' }{code} And we call the current update() method like: {code}update({ name: 'bob' }){code} it would basically be a no-op if we allow just partial state objects. It'd set my name to 'bob' (as it was to start) and would ignore the un-mentioned favorite-food. If I want to remove any favorite-food, I'd need to do a full update and exclude it from the inbound state. If we had {code}update({ remove: [favorite-food] }){code} as a diff-centric update, then we could remove it without altering saved state beyond the removal of the property. To make life easier for people, we could also provide a default implementation for a diff-centric update, in psuedo-code {code}default void update(ResourceDiff diff) { ResourceState before = doMagicToObtainTheBeforeDiffState(); ResourceState after = applyDiff(before, diff); update(after); }{code} This would allow simple Resources to only care about update(ResourceState state) and assume it's always a full update, while the container would still provide the ability to apply diffs from outside. A smarter or more optimized Resource could implement their own update(ResourceDiff diff) to operate more efficiently than a fetch/apply/update the default implementation.	1	0	https://issues.jboss.org/browse/LIVEOAK-1
15351	LOGTOOL-124	Do we want to use a {{@Producer}} annotation or just make assumptions if the parameter is a {{java.util.function.Function}}?	1	1	https://issues.jboss.org/browse/LOGTOOL-124

#	Requirement	Questions	Answered	Successfully Answered	URL
15352	LOGTOOL-124	Just to confirm. For a {{Function}} the signature must resolve to {{java.util.function.Function<String, Throwable>}} where {{String}} is the message. For a {{BiFunction}} the signature must resolve to {{java.util.function.BiFunction<String, Throwable, Throwable}} where the {{String}} is the message and the first {{Throwable}} is the cause. Does that sound correct? Example of the current output for a function: {code:java|title=Interface} @Message("The operation failed due to %s") <T extends Throwable> T operationFailed(@Producer Function<String, T> function, String op); {code} {code:java|title=Generated} private static final String operationFailed = "The operation failed due to %s"; protected String operationFailed$str() { return operationFailed; } @Override public final <T extends Throwable> T operationFailed(final java.util.function.Function<String, T> function, final String op) { final T result = function.apply(String.format(getLoggingLocale(), operationFailed$str(), op)); final StackTraceElement[] st = result.getStackTrace(); result.setStackTrace(Arrays.copyOfRange(st, 1, st.length)); return result; } {code}	1	1	https://issues.jboss.org/browse/LOGTOOL-124
15353	LOGTOOL-101	Failing that (cloning), an alternative is to allow overloading: {code} @Message(id = 1234, "Binding to %s failed: %s") IOException bindFailed(SocketAddress bindAddress, @CopyStackFrom @Pos(2) IOException cause) BindException bindFailed(SocketAddress bindAddress, @CopyStackFrom @Pos(2) BindException cause) {code} Which I think is allowed today (right?).	1	1	https://issues.jboss.org/browse/LOGTOOL-101
15354	LOGTOOL-101	Do you want the stack trace just appended to the created exceptions stack trace? So something like: {code:java} final StackTraceElement[] resultStackTrace = result.getStackTrace(); final StackTraceElement[] causeStackTrace = cause.getStackTrace(); final int newLen = resultStackTrace.length + causeStackTrace.length; final StackTraceElement[] newStackTrace = Arrays.copyOf(resultStackTrace, newLen); System.arraycopy(causeStackTrace, 0, newStackTrace, resultStackTrace.length, causeStackTrace.length); result.setStackTrace(newStackTrace); {code} WRT cloning I'm not sure I get what you mean. Note I'll also fix the {{@Pos}} documentation as it's rather poor :)	1	0	https://issues.jboss.org/browse/LOGTOOL-101
15355	LOGTOOL-101	If we add a parameter to the {{@Producer}} this might not be needed. Maybe something like {{@Producer(copyStackTrace=false}}, then the {{Function}} could just attempt to clone the exception. {code:java} @Message("Binding to %s failed") <T extends IOException> T bindFailed(@Producer(copyStackTrace=false BiFunction<String, IOException, T> function, @Cause IOException cause, SocketAddress address); throw ValidMessages.MESSAGES.bindFailed((msg, cause) -> { final IOException copy; if (cause instanceof BindException) { copy = new BindException(msg + ": " + cause.getMessage()); } else { copy = new IOException(msg + ": " + cause.getMessage()); } copy.setStackTrace(cause.getStackTrace()); return copy; }, new IOException("Failed IO exception"), address); {code} Note the {{copyStackTrace}} here refers to the reset of the stack trace removing the interface entry from the stack trace. A better name may be needed. E.g the default implementation looks something like: {code:java} @Override public final <T extends IOException> T bindFailed(final java.util.function.BiFunction<String, IOException, T> function, final IOException cause, final SocketAddress address) { final T result = function.apply(String.format(getLoggingLocale(), bindFailed$str(), address), cause); final StackTraceElement[] st = result.getStackTrace(); result.setStackTrace(Arrays.copyOfRange(st, 1, st.length)); return result; } {code} So we'd just need property to indicate the resulting stack trace should not be set. Maybe {{setStackTrace=false}}?	0	0	https://issues.jboss.org/browse/LOGTOOL-101
15356	LOGTOOL-84	Yes, it seems related. In fact, {{@FormatWith}} appears to be what I'm after, just learned about it. I'll give that a try. Though what I had in mind here would go a bit farther, as it basically would specify default {{@FormatWith}} 's based on the parameter type, without the need to specify the formatter type for each single parameter. Btw. the pull request for LOGTOOL-55 is merged, still the issue itself is open. Shouldn't it be marked as resolved?	1	0	https://issues.jboss.org/browse/LOGTOOL-84
15357	LOGTOOL-79	What about just a {{@Signature}} with a {{causeIndex}} and {{messsageIndex}} attribute?	1	1	https://issues.jboss.org/browse/LOGTOOL-79
15358	LOGTOOL-44	Since the JBoss Logging library is compiled to Java 6 I'm not sure this would work anyway. Can you give me the version of JBoss Logging you're using?	0	0	https://issues.jboss.org/browse/LOGTOOL-44
15359	LOGTOOL-44	> Since the JBoss Logging library is compiled to Java 6 Are you using any Java 6 specific features. Sure Java 5 is out of date, but if any framework should be conservative regarding JVM versions it should be logging.	0	0	https://issues.jboss.org/browse/LOGTOOL-44
15360	LOGTOOL-44	> Since the JBoss Logging library is compiled to Java 6 I'm not sure this would work anyway. Right, there is a problem. See https://hibernate.onjira.com/browse/HV-569 So even though Hibernate Validator itself would still run in a Java 5 environment, by switching to JBoss logging we basically force users to use Java 6. Couldn't you just compile JBoss Logging with target 1.5?	0	0	https://issues.jboss.org/browse/LOGTOOL-44
15361	LOGTOOL-44	{quote} Let me make it perfectly clear that 1.5 compatibility is a non-requirement at the moment. {quote} If you are not using any Java 6 specific features it is just a question of adding 'target = 1.5' to your compiler flags. If it is just about _Arrays#copyOfRange_, just add method to the code base. I think the advantage is outweighing the fact that you have "duplicated" code. Even though Java 5 is out of date, I think there are still a lot of deployments out there which run on this JVM. I just see it as an advantage if they have the chance to upgrade parts of their system (as they could w/ Validator if not for the logging framework) {quote} If you want jboss-logging to actually execute on 1.5, you're going to have to use a retroweaver of some sort to do it. {quote} How would that look like? I am not sure this is possible.	0	0	https://issues.jboss.org/browse/LOGTOOL-44
15362	LOGTOOL-20	FYI, this is the output with my patch: those are clearly invalid messages from Infinispan, not in your validation. Or feel free to use them as tests. {quote}[INFO] --- maven-compiler-plugin:2.3.1:compile (default-compile) @ infinispan-core --- [INFO] Compiling 613 source files to /home/sanne/workspaces/infinispan/infinispan/core/target/classes Format specifier 's' Parameter count does not match for format 'Could not register object with name: %s (%s)'. Required: 2 Provided: 1 Format specifier 's' Parameter count does not match for format 'Found WebSphere TransactionManager factory class [%s], but couldn't invoke its static 'getTransactionManager' method'. Required: 1 Provided: 0 [INFO] ---------- [INFO] BUILD SUCCESS {quote}	1	0	https://issues.jboss.org/browse/LOGTOOL-20
15363	MPJDOCBOOK-86	Yes, the whole renaming process won't be easy. Do you think we should go for it anyway? I mean I reported the error as maven was complaining about it, but I don't know if it is worth the trouble.	1	0	https://issues.jboss.org/browse/MPJDOCBOOK-86
15364	MPJDOCBOOK-32	thanks. the website is back online. could you update the mailing list pages or is there a forum available?	0	0	https://issues.jboss.org/browse/MPJDOCBOOK-32
15365	MPJDOCBOOK-19	I was actually not sure how to accomplish this. Do you have any ideas?	1	1	https://issues.jboss.org/browse/MPJDOCBOOK-19
15366	MPJDOCBOOK-6	hmmm sounds interesting. can't wait to see how it comes along for you ;) I'd prefer to not do this as a report plugin as I actually already tried that route. But the site/report stuff is perhaps the worst documented of al the badly documented parts of maven. What did you have in mind when you say "or maybe perhaps just include the raw artifacts in the site"?	1	0	https://issues.jboss.org/browse/MPJDOCBOOK-6
15367	MIGR-235	@pmuir what's the plan for this missing content?	1	1	https://issues.jboss.org/browse/MIGR-235
15368	MIGR-235	Do we have anywhere the previous content backed up in git or not?	1	1	https://issues.jboss.org/browse/MIGR-235

#	Requirement	Questions	Answered	Successfully Answered	URL
15369	MIGR-235	[~paul.robinson] and [~manaRH]: Has this been JIRA been resolved?	1	0	https://issues.jboss.org/browse/MIGR-235
15370	MIGR-233	Why is this here? Are you using MIGR as a personal JIRA list?	1	0	https://issues.jboss.org/browse/MIGR-233
15371	MIGR-233	Marc: >> Maybe we will need a parameter to say if we need the EAP 6 or EAP 7 BOM. (I do not think that there will be any other BOM) I have created the indexer, and I can make it create several different Lucene indexes. So that the user would have to choose which one to use, or Windup would have to query each one. What we can do: a) Have one combined index with the information in which BOM the APIs are available. b) Pick the right index based on the target platform. On the other hand, I don't want to make it too complicated - the feature is basically still "what artifact contains this package"; the BOM is currently only used when creating the index. >> For the war/ear/jar structure, it is also fully sufficient in a first step just to list the analyzed top-level applications. I mean that creating maven project for embedded-self-written (or shared) libraries is not necessary. >> It is perfectly fine to release it without all exhausite provided dependencies. I would recommend to combine it with the implementation of the technological-overview report. For both you will need the same information. >> Adding an option/parameter for the company pom-GAV is fully sufficient. There is no need for detecting it automatically I think. If no parameter has been set, something like "com.mycompany.mavenized:jee-example-app-parent:1.0-SNAPSHOT" is fully sufficient.	0	0	https://issues.jboss.org/browse/MIGR-233
15372	MIGR-233	What do do with cases like these? wimssecurity-client.jar Name: com.mformation.* (from MANIFEST)	0	0	https://issues.jboss.org/browse/MIGR-233
15373	MIGR-214	{code} .addRule() .when(XmlFile.matchesXpath("persistence[not(@version)]")) .perform(Classification .as("JPA Configuration") .withEffort(0) .and(XSLTTransformation.using("transformations/xslt/jboss4-persistence-to-jboss5.xsl") .withDescription("JPA Descriptor - JBoss 5 (Windup-Generated)") .withExtension("-jboss5.xml"))) {code}	0	0	https://issues.jboss.org/browse/MIGR-214
15374	MIGR-211	Actually, if you could clarify what the actual use case is, that might help. What does the user want? Do they want generation of data-sources? Do they like a file backup? (Which is weird since we wouldn't modify files?)	0	0	https://issues.jboss.org/browse/MIGR-211
15375	MIGR-198	Lincoln: "Need to check these two things: Is there newer addon version online? Is it compatible?" https://github.com/forge/furnace#usage	1	1	https://issues.jboss.org/browse/MIGR-198
15376	MIGR-195	Is this covered by Forge?	1	1	https://issues.jboss.org/browse/MIGR-195
15377	MIGR-188	Is this done?	1	1	https://issues.jboss.org/browse/MIGR-188
15378	MIGR-177	Is there anything else we want to do here or is this issue resolved?	1	1	https://issues.jboss.org/browse/MIGR-177
15379	MIGR-177	I have only one question on this. We suggest folks go to the migration guide in the customer portal. What if they aren't customers - which is a strong possibility? The other doc links are to sites anyone can get to? Can we provide a link that is not on the Customer Portal?	0	0	https://issues.jboss.org/browse/MIGR-177
15380	MIGR-175	The IBM link noted above is obsolete, however, a search turned up this page: http://www-01.ibm.com/support/docview.wss?uid=swg24034637 [~rsearls]: Does that look right?	1	1	https://issues.jboss.org/browse/MIGR-175
15381	MIGR-175	[~clichybi]: This link in your document gives an error: http://www-01.ibm.com/support/docview.wss?uid=swg24019368 Should it be replaced with this link? http://www-01.ibm.com/support/docview.wss?uid=swg24034637	1	1	https://issues.jboss.org/browse/MIGR-175
15382	MIGR-175	[~clichybi] In your document, you say to "download jboss-as-tomcat-6.0.0.Final.jar from JBoss AS6". I used the following steps but I must be looking in the wrong release: 1. Browse to http://www.jboss.org/jbossas/downloads 2. Find JBoss AS 6.0.0.Final in the list and click Download. You see the message: Community projects represent the latest development releases and are not supported. Click 'Continue Download' 3. Find jboss-as-distribution-6.0.0.Final.zip in the list and click on it. It downloads the follwoing file: http://sourceforge.net/projects/jboss/files/JBoss/JBoss-6.0.0.Final/jboss-as-distribution-6.0.0.Final.zip/download I don't see jboss-as-tomcat-6.0.0.Final.jar in that distribution. Can you tell me where a customer needs to go to get that file?	1	1	https://issues.jboss.org/browse/MIGR-175
15383	MIGR-174	How to connect to JBoss EAP 6 using JConsole: Standalone server: https://access.redhat.com/site/solutions/149973 Managed domain: https://access.redhat.com/site/solutions/443033 I'm not clear why you need to set the JAVA_HOME within the shell script and configuration files though as we don't do that with the helloworld-mbean quickstart. I need to investigate	0	0	https://issues.jboss.org/browse/MIGR-174
15384	MIGR-166	Would you proivde a link to the Eclipse Plugin archive URL so I can proivded it on the web site? I also need a link to the video for this. If you have specific text you want for this new feature please provide it here and I will use it.	1	0	https://issues.jboss.org/browse/MIGR-166
15385	MODCLUSTER-528	The conclusion doesn't make any sense to me - which systems? Which tomcats? Why would we bring Tomcat 7 dependencies to Tomcat 8 game? When I compile with Tomcat 7 profile, I should need Tomcat 7 dependencies, when I compile with Tomcat 8 profile, I should need only Tomcat 8 dependencies without patching *.java files and without patching pom files. The aforementioned is not possible due to a bug in the Java source code and pom structure. According to the pom, it looks like we offer Tomcat 7 and Tomcat 8 profiles, while in fact we don't (Tomcat 8 profile still depends on Tomcat 7 libs). h3. Suggestion * drop Tomcat 6 from 1.3.x branch and from master, it is not needed there * drop JBossWeb/AS7 from 1.3.x branch and from master, it is not needed there * move Catalina stuff around so as *-PTC7* needs Tomcat 7 dependencies, *-PTC8* needs Tomcat 8 dependencies, *-PTC85* needs Tomcat 8.5 dependencies etc.	0	0	https://issues.jboss.org/browse/MODCLUSTER-528
15386	MODCLUSTER-528	> The conclusion doesn't make any sense to me - which systems? Which tomcats? For instance Fedora 24 build system (i.e. ecosystem). > Why would we bring Tomcat 7 dependencies to Tomcat 8 game? See below. > I should need only Tomcat 8 dependencies without patching *.java files That's a terrible approach. Building for RPM or whatever other system needs to follow what maven is doing -- that's the build that is used and tested. introducing new code in e.g. RPM building is a quiality risk. > drop Tomcat 6 from 1.3.x branch and from master, it is not needed there > drop JBossWeb/AS7 from 1.3.x branch and from master, it is not needed there Long time dropped from master. > move Catalina stuff around so as -PTC7 needs Tomcat 7 dependencies, -PTC8 needs Tomcat 8 dependencies, -PTC85 needs Tomcat 8.5 dependencies etc. This is setup in pre-master, never used and never worked. This is dropped in master (commented out). Doing this means unnecessary code duplication. The base module could then only contain code that can be compiled with *every* Tomcat version we support. The approach that we have been shipping for years is to compile base against single version and override code that needs to be changed for newer tomcat version.	0	0	https://issues.jboss.org/browse/MODCLUSTER-528

#	Requirement	Questions	Answered	Successfully Answered	URL
15387	MODCLUSTER-518	org.jboss.logging.processor.generator.model.ClassModel.generateModel(ClassModel.java:122) at org.jboss.logging.processor.generator.model.ImplementationClassModel.generateModel(ImplementationClassModel.java:81) at org.jboss.logging.processor.generator.model.MessageLoggerImplementor.generateModel(MessageLoggerImplementor.java:92) at org.jboss.logging.processor.generator.model.ClassModel.create(ClassModel.java:104) at org.jboss.logging.processor.apt.ImplementationClassGenerator.processTypeElement(ImplementationClassGenerator.java:63) at org.jboss.logging.processor.apt.LoggingToolsProcessor.process(LoggingToolsProcessor.java:145) at com.sun.tools.javac.processing.JavacProcessingEnvironment.callProcessor(jdk.compiler@9-ea/JavacProcessingEnvironment.java:961) at com.sun.tools.javac.processing.JavacProcessingEnvironment.discoverAndRunProcs(jdk.compiler@9-ea/JavacProcessingEnvironment.java:877) at com.sun.tools.javac.processing.JavacProcessingEnvironment.access$2100(jdk.compiler@9-ea/JavacProcessingEnvironment.java:107) at com.sun.tools.javac.processing.JavacProcessingEnvironment$Round.run(jdk.compiler@9-ea/JavacProcessingEnvironment.java:1184) at com.sun.tools.javac.processing.JavacProcessingEnvironment.doProcessing(jdk.compiler@9-ea/JavacProcessingEnvironment.java:1292) at com.sun.tools.javac.main.JavaCompiler.processAnnotations(jdk.compiler@9-ea/JavaCompiler.java:1229) at com.sun.tools.javac.main.JavaCompiler.compile(jdk.compiler@9-ea/JavaCompiler.java:903) at com.sun.tools.javac.api.JavacTaskImpl$1.call(jdk.compiler@9-ea/JavacTaskImpl.java:104) at com.sun.tools.javac.api.JavacTaskImpl$1.call(jdk.compiler@9-ea/JavacTaskImpl.java:100) at com.sun.tools.javac.api.JavacTaskImpl.handleExceptions(jdk.compiler@9-ea/JavacTaskImpl.java:135) at com.sun.tools.javac.api.JavacTaskImpl.doCall(jdk.compiler@9-ea/JavacTaskImpl.java:100) at com.sun.tools.javac.api.JavacTaskImpl.call(jdk.compiler@9-ea/JavacTaskImpl.java:94) at org.codehaus.plexus.compiler.javac.JavaxToolsCompiler.compileInProcess(JavaxToolsCompiler.java:125) at org.codehaus.plexus.compiler.javac.JavacCompiler.performCompile(JavacCompiler.java:171) at org.apache.maven.plugin.compiler.AbstractCompilerMojo.execute(AbstractCompilerMojo.java:886) at org.apache.maven.plugin.compiler.CompilerMojo.execute(CompilerMojo.java:129) at org.apache.maven.plugin.DefaultBuildPluginManager.executeMojo(DefaultBuildPluginManager.java:134) at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:208) at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:153) at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:145) at org.apache.maven.lifecycle.internal.LifecycleModuleBuilder.buildProject(LifecycleModuleBuilder.java:116) at org.apache.maven.lifecycle.internal.LifecycleModuleBuilder.buildProject(LifecycleModuleBuilder.java:80) at org.apache.maven.lifecycle.internal.builder.singlethreaded.SingleThreadedBuilder.build(SingleThreadedBuilder.java:51) at	0	0	https://issues.jboss.org/browse/MODCLUSTER-518
15388	MODCLUSTER-499	Where can the tarballs be downloaded from? If its from an official apache hosting, that sounds best to me atm.	1	0	https://issues.jboss.org/browse/MODCLUSTER-499
15389	MODCLUSTER-492	What exactly is here to do? * {{maxAttempts}} defaults to 1, so the native balancer *will* try another worker * {{stickySession*}} attributes are configurable on worker side...	1	0	https://issues.jboss.org/browse/MODCLUSTER-492
15390	MODCLUSTER-449	Wouldn't it be better addressed in the native part? The problem is not the load number itself, it's the {{elected}} counter. I suggest treating "newcomers" gently in {{internal_find_best_byrequests}}. IMHO one could guess that info from a large discrepancy between the elected counters of other nodes and this new one. I'm gonna try to verify I'm not just being full of ideas...	0	0	https://issues.jboss.org/browse/MODCLUSTER-449
15391	MODCLUSTER-449	I've developed a new load metric that has a "warm up" period where the capacity starts at 0 and then grows linearly to the true capacity over a configurable number of seconds (default 120). Would this be something you'd be interested in me submitted a pull request for?	1	0	https://issues.jboss.org/browse/MODCLUSTER-449
15392	MODCLUSTER-448	At least on the Apache mod_cluster side, could we have it refuse to register 0.0.0.0 and dump out a warning/error?	0	0	https://issues.jboss.org/browse/MODCLUSTER-448
15393	MODCLUSTER-438	can I use both ajp and WebSocket on one apache for loadbalance? e.g. when use http client access to apache,it will redirect to a http webserver. when use WebSocket client access to the same apache,it will redirect to a websocket webserver.	1	0	https://issues.jboss.org/browse/MODCLUSTER-438
15394	MODCLUSTER-438	[~mbabacek] What's the status of the ws tunnel integration? Do we know, is this supported in Undertow-based implementation?	1	0	https://issues.jboss.org/browse/MODCLUSTER-438
15395	MODCLUSTER-438	[#Jean-Frederic Clere] I am troubling about how to support http and websockt at the same time rather than how to support websocket. To support WebSocket for mod_cluster, i have to set option "EnableWsTunnel" in configuration of apache, and use http instead of ajp. However, it makes all the http to ws. The result is websockt supported but http not. Is there any way to for apache to support the two protocol ,http and ws, at the same time?	0	0	https://issues.jboss.org/browse/MODCLUSTER-438
15396	MODCLUSTER-438	[~gfleury] Thank you. The current version of httpd/mod_cluster supports websockets via proxy_wstunnel. Does nginx/mod_cluster anything differently? e.g. what happens when worker dies? Does the balancer try to re-negotiate handshake with a different worker or it's up to the client to do so? We arrived at the conclusion that there is no meaning in having websocket "ha", so when worker dies, the communication is broken and client has to initiate it again - balancer would give the client another, healthily worker then.	1	0	https://issues.jboss.org/browse/MODCLUSTER-438
15397	MODCLUSTER-438	Is this issue able to include the feature for the configuration Undertow as a Load Balancer Using mod_cluster [1] ? Then the changes will be in mod_cluster, not native. Or is it should be created as a separate issue ? [1] https://access.redhat.com/documentation/en-us/red_hat_jboss_enterprise_application_platform/7.0/html/configuration_guide/configuring_high_availability#configure_undertow_load_balancer_mod_cluster	1	0	https://issues.jboss.org/browse/MODCLUSTER-438
15398	MODCLUSTER-434	Our need is that, in a sticky-session context, when server load exceeds some specified threshold, no more requests are sent to this server until server load drops below another threshold. So that sessions will be moved to other servers as new requests from this sessions arrive. It is not necessary to move all sessions instantly, they could be moved gradually as requests arrive, until server is ready for new requests. The former threshold (upperThreshold) must be grater than the latter (lowerThreshold) so that the server is not taken out and in so frequent. I add an alternative: 4. Add this two thresholds attributes to a new AbstractLoadMetric subclass. In DynamicLoadBalanceFactorProvider.getLoadBalanceFactor(): {code} ... if (previousLoadValueOfThisMetric > 0) { if (currentLoadValue <= metric.getUpperThreshold()) { //do actual stuff with currentLoadValue; } else { // stop evaluating metrics and return -1 return -1; } } else { if (currentLoadValue <= metric.getLowerThreshold()) { //do actual stuff with currentLoadValue; } else { // stop evaluating metrics and return -1 return -1; } } ... {code} I don´t know which are the "consequences" of returning -1 (ERROR STATE) to the balancer. Perhaps we could create a new -3 OVERLOADED STATE. WDYT? If you want, we (with [~godiedelrio]) could try to develop a proposal and upload it for evaluation. Thanks	0	0	https://issues.jboss.org/browse/MODCLUSTER-434

#	Requirement	Questions	Answered	Successfully Answered	URL
15399	MODCLUSTER-420	I think we should align with EAP7 and EWS version that is going to be release then. Do we know which httpd is it going to be based on?	1	1	https://issues.jboss.org/browse/MODCLUSTER-420
15400	MODCLUSTER-390	Jean-Frederic, could you elaborate on how this actually could be used with mod_cluster so as I can test it? AFAIK, it does not work at the moment, one can't specify any such BalancerMember directives and set any specific error codes.	1	1	https://issues.jboss.org/browse/MODCLUSTER-390
15401	MODCLUSTER-390	2.2.26 from EAP 6.3.0 distro * Two EAP 6.3 servers, one called _jboss-eap-6.3_ and the other being _jboss-eap-6.3-2_ * mod_cluster modules build from the latest 1.2.x upstream This is my mod_cluster.conf: {code} MemManagerFile "/dev/shm/mod_cluster-eapx/jboss-ews-2.1/httpd/cache/mod_cluster" ProxyPass / balancer://balancerxxx stickysession=JSESSIONID|jsessionid nofailover=on failonstatus=203,204 ProxyPassReverse / balancer://balancerxxx ProxyPreserveHost on ServerName 192.168.122.204:2181 <IfModule manager_module> Listen 192.168.122.204:8847 LogLevel debug <VirtualHost 192.168.122.204:8847> ServerName 192.168.122.204:8847 <Directory /> Order deny,allow Deny from all Allow from all </Directory> KeepAliveTimeout 60 MaxKeepAliveRequests 0 ServerAdvertise on AdvertiseFrequency 5 ManagerBalancerName balancerxxx AdvertiseGroup 224.0.5.79:65009 EnableMCPMReceive CreateBalancers 1 <Location /mcm> SetHandler mod_cluster-manager Order deny,allow Deny from all Allow from all </Location> </VirtualHost> </IfModule> {code} EAP config... {code} <subsystem xmlns="urn:jboss:domain:modcluster:1.2"> <mod-cluster-config advertise-socket="modcluster" sticky-session="true" sticky-session-remove="false" sticky-session-force="false" connector="ajp"> <dynamic-load-provider> <load-metric type="busyness"/> </dynamic-load-provider> </mod-cluster-config> </subsystem> {code} h3. Test This dead simple [servlet|https://github.com/Karm/clusterbench/blob/simplified-and-pure/clusterbench-common/src/main/java/org/jboss/test/clusterbench/common/jvmroute/CommonHTTPCodeServlet.java#L40] sets an arbitrary response return code according to the given parameter. The output is the following: Warming up, {{http_code=200}} is not one of the defined error codes, so everything should be normal. The session is created on the first request, the subsequent ones are the same, no failover, cool... {code} mbabacek@localhost:~$ curl http://rhel6x86-64:8847/clusterbench/httpcode/?http_code=200 -c cookie.txt -b cookie.txt -i HTTP/1.1 200 OK Date: Thu, 11 Sep 2014 11:02:22 GMT Set-Cookie: JSESSIONID=gMPIq7FfCeiPwzM2iZLyJiR4.jboss-eap-6.3; Path=/clusterbench Content-Type: text/plain;charset=UTF-8 Content-Length: 122 Connection: close Done. HTTP Code was: 200 JVM route: jboss-eap-6.3 Session ID: gMPIq7FfCeiPwzM2iZLyJiR4.jboss-eap-6.3 Session isNew: true mbabacek@localhost:~$ curl http://rhel6x86-64:8847/clusterbench/httpcode/?http_code=200 -c cookie.txt -b cookie.txt -i HTTP/1.1 200 OK Date: Thu, 11 Sep 2014 11:02:24 GMT Content-Type: text/plain;charset=UTF-8 Content-Length: 123 Connection: close Done. HTTP Code was: 200 JVM route: jboss-eap-6.3 Session ID: gMPIq7FfCeiPwzM2iZLyJiR4.jboss-eap-6.3 Session isNew: false mbabacek@localhost:~$ curl http://rhel6x86-64:8847/clusterbench/httpcode/?http_code=200 -c cookie.txt -b cookie.txt -i HTTP/1.1 200 OK Date: Thu, 11 Sep 2014 11:02:27 GMT Content-Type: text/plain;charset=UTF-8 Content-Length: 123 Connection: close Done. HTTP Code was: 200 JVM route: jboss-eap-6.3 Session ID: gMPIq7FfCeiPwzM2iZLyJiR4.jboss-eap-6.3 Session isNew: false {code} Now, let's try {{http_code=203}}, one of the predefined error codes. What I would expect is an immediate failover from _jboss-eap-6.3_ to _jboss-eap-6.3-2_, although this is what happens: {code} mbabacek@localhost:~$ curl http://rhel6x86-64:8847/clusterbench/httpcode/?http_code=203 -c cookie.txt -b cookie.txt -i HTTP/1.1 203 Non-Authoritative Information Date: Thu, 11 Sep 2014 11:02:39 GMT Content-Type: text/plain;charset=UTF-8 Content-Length: 123 Connection: close Done. HTTP Code was: 203 JVM route: jboss-eap-6.3 Session ID: gMPIq7FfCeiPwzM2iZLyJiR4.jboss-eap-6.3 Session isNew: false mbabacek@localhost:~$ curl http://rhel6x86-64:8847/clusterbench/httpcode/?http_code=203 -c cookie.txt -b cookie.txt -i HTTP/1.1 203 Non-Authoritative Information Date: Thu, 11 Sep 2014 11:03:03 GMT Content-Type: text/plain;charset=UTF-8 Content-Length:	1	1	https://issues.jboss.org/browse/MODCLUSTER-390
15402	MODCLUSTER-390	Hmm, apart from the message, could you comment on my question: {quote} @[~jfclere] Hmm, the forced failover doesn't happen at once, it takes several requests...is that the correct behaviour? {quote} I understand that the _forcerecovery_ ain't supported, but there should be an immediate failover if the listed status is hit, IMHO.	1	1	https://issues.jboss.org/browse/MODCLUSTER-390
15403	MODCLUSTER-390	Very well, I've written an automated test that reflects the _comment #4_ and it works: h3. mod_cluster before the feature was implemented Some initial requests, HTTP 200... {code} Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=200, response code = 200 HTTP Code was: 200 JVM route: jboss-eap-6.2 Session ID: 4VKyGpxfi2MvHCOKC4A-l-5i.jboss-eap-6.2 Session isNew: false Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=200, response code = 200 HTTP Code was: 200 JVM route: jboss-eap-6.2 Session ID: 4VKyGpxfi2MvHCOKC4A-l-5i.jboss-eap-6.2 Session isNew: false {code} Now, let's use HTTP 201, the *failonstatus=201* setting... {code} Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=201, response code = 201 HTTP Code was: 201 JVM route: jboss-eap-6.2 Session ID: 4VKyGpxfi2MvHCOKC4A-l-5i.jboss-eap-6.2 Session isNew: false {code} The response comes from the same node, that's fine, but the subsequent requests come from that node as well, i.e. the feature is not implemented: {code} Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=200, response code = 200 HTTP Code was: 200 JVM route: jboss-eap-6.2 Session ID: 4VKyGpxfi2MvHCOKC4A-l-5i.jboss-eap-6.2 Session isNew: false Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=200, response code = 200 HTTP Code was: 200 JVM route: jboss-eap-6.2 Session ID: 4VKyGpxfi2MvHCOKC4A-l-5i.jboss-eap-6.2 Session isNew: false {code} h3. current mod_cluster 1.2.9, feature implemented Some initial requests, HTTP 200... {code} Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=200, response code = 200 HTTP Code was: 200 JVM route: jboss-eap-6.3 Session ID: cETRQkQtnJk9R0I4eaGaoril.jboss-eap-6.3 Session isNew: false Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=200, response code = 200 HTTP Code was: 200 JVM route: jboss-eap-6.3 Session ID: cETRQkQtnJk9R0I4eaGaoril.jboss-eap-6.3 Session isNew: false {code} Now, let's use HTTP 201, the *failonstatus=201* setting... {code} Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=201, response code = 201 HTTP Code was: 201 JVM route: jboss-eap-6.3 Session ID: cETRQkQtnJk9R0I4eaGaoril.jboss-eap-6.3 Session isNew: false {code} The response comes from the same node, that's fine, but the subsequent requests come from a different node, i.e. the feature is implemented: {code} Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=200, response code = 200 HTTP Code was: 200 JVM route: jboss-eap-6.3-2 Session ID: cETRQkQtnJk9R0I4eaGaoril.jboss-eap-6.3-2 Session isNew: false Trying to retrieve content from http://192.168.122.204:2181/clusterbench/httpcode/?http_code=200, response code = 200 HTTP Code was: 200 JVM route: jboss-eap-6.3-2 Session ID: cETRQkQtnJk9R0I4eaGaoril.jboss-eap-6.3-2 Session isNew: false {code} h3. Regarding Status As to your other question, the *jboss-eap-6.3* worker was in *NOTOK* for a period of time just until the next *STATUS* message was received. After that, both *jboss-eap-6.3* and *jboss-eap-6.3-2* have been in *OK* status. So, what shall we do about it? I think it's probably all right. 1. I have a certain set of HTTP codes that I want to trigger a failover 2. it works, failover occurs (on the *second* request) 3. the worker is put back in business, because there is no reason to believe it's broken, i.e. the fact that a web app returned a specific code does not suggest that the whole box should be kept in *NOTOK*... WDYT?	1	1	https://issues.jboss.org/browse/MODCLUSTER-390
15404	MODCLUSTER-343	Doesn't it because of MODCLUSTER-430?	0	0	https://issues.jboss.org/browse/MODCLUSTER-343

#	Requirement	Questions	Answered	Successfully Answered	URL
15405	MODCLUSTER-342	MemManagerFile /var/cache/httpd does it exist? Is it writable by the httpd user? Any SELinux/or UAC security like on the box?	1	1	https://issues.jboss.org/browse/MODCLUSTER-342
15406	MODCLUSTER-342	MemManagerFile /var/cache/httpd does it exist? Yes. it is. Is it writable by the httpd user? Yes, the user is root. Any SELinux/or UAC security like on the box? I don't know, but I think that anything.	1	1	https://issues.jboss.org/browse/MODCLUSTER-342
15407	MODCLUSTER-342	not yet we do have a box to investigate the problem. Could you use strace (or something similar) to track the system call that have failed? Additionally the output of ipcs (or something similar) could help.	1	1	https://issues.jboss.org/browse/MODCLUSTER-342
15408	MODCLUSTER-341	Weird the hosts are removed in: https://github.com/modcluster/mod_cluster/blob/master/native/mod_manager/mod_manager.c#L1440 how do you find that they aren't removed?	1	1	https://issues.jboss.org/browse/MODCLUSTER-341
15409	MODCLUSTER-334	smumford@redhat.com made a comment on [bug 958991|https://bugzilla.redhat.com/show_bug.cgi?id=958991] The JIRA suggests this is fixed, is that the case? And for release notes, does the user have to add that code? And if so, where does it go?	1	1	https://issues.jboss.org/browse/MODCLUSTER-334
15410	MODCLUSTER-329	Is there anything new on this feature?	1	0	https://issues.jboss.org/browse/MODCLUSTER-329
15411	MODCLUSTER-315	So is this essentially resolved? Which AS/EAP version is included?	0	0	https://issues.jboss.org/browse/MODCLUSTER-315
15412	MODCLUSTER-315	Okay, now I am confused. Is this in EAP 6.1 and upstream in AS8?	0	0	https://issues.jboss.org/browse/MODCLUSTER-315
15413	MODCLUSTER-306	I had a customer using this backported fix through the JBPAPP-9571 patch. They use this along with Siteminder and see STATUS MCMP requests were randomly failing/500ing without any real reason. Come to find out that any failed STATUS is missing the Host header and so siteminder fails it: {noformat} 10.170.104.52 - - [06/Mar/2013:12:44:52 -0800] "STATUS / HTTP/1.1" 500 535 "-" "ClusterListener/1.0" 0 375 10.170.104.52 - - [06/Mar/2013:12:44:52 -0800] "STATUS / HTTP/1.1" 500 535 "-" "ClusterListener/1.0" 0 343 10.170.104.52 - - [06/Mar/2013:12:44:56 -0800] "STATUS / HTTP/1.1" 500 535 "-" "ClusterListener/1.0" 0 468 10.170.104.51 - - [06/Mar/2013:12:45:01 -0800] "INFO / HTTP/1.1" 200 5088 "-" "ClusterListener/1.0" 0 848 ina-internal-ch-prd-web1.uboc.com:8466 10.170.104.51 - - [06/Mar/2013:12:45:01 -0800] "INFO / HTTP/1.1" 200 5088 "-" "ClusterListener/1.0" 0 757 ina-internal-ch-prd-web1.uboc.com:8466 10.170.104.51 - - [06/Mar/2013:12:45:01 -0800] "CONFIG / HTTP/1.1" 200 - "-" "ClusterListener/1.0" 0 574 ina-internal-ch-prd-web1.uboc.com:8466 [CSmHttpPlugin.cpp:278][ERROR] Unable to resolve server host name. Exiting with HTTP 500 server error '10-0004'. {noformat} Here's DefaultMCMPHandler.sendRequest from the 1.0.10.GA_CP branch currently: {code:title=DefaultMCMPHandler.java|borderStyle=solid} String head = builder.toString(); int length = encoder.getLength(); char[] body = encoder.getBuffer(); // Require exclusive access to proxy socket synchronized (proxy) { try { String line = null; StringBuilder proxyheadBuilder = new StringBuilder(head); proxyheadBuilder.append(proxy.getAddress().getHostName() + ":" + proxy.getPort()); String proxyhead = proxyheadBuilder.toString(); try { line = sendRequest(proxy, proxyhead, body, length); } catch (IOException e) { // Ignore first write failure } if (line == null) { // Retry failed read/write with fresh connection proxy.closeConnection(); line = sendRequest(proxy, head, body, length); } {code} So it looks like they occasionally use that second sendRequest, which wasn't updated to use the proxyHead string that was built with the info for the Host header. I can see we have it proper in master and 1.2 branches though: {code:title=DefaultMCMPHandler.java|borderStyle=solid} try { line = sendRequest(proxy, proxyhead, body); } catch (IOException e) { // Ignore first write failure } if (line == null) { // Retry failed read/write with fresh connection proxy.closeConnection(); line = sendRequest(proxy, proxyhead, body); } {code} Can we correct this in the 1_0_10_GA_CP branch?	1	1	https://issues.jboss.org/browse/MODCLUSTER-306
15414	MODCLUSTER-303	Hi Vladimir, do we have any documentation how to set this up ?	0	0	https://issues.jboss.org/browse/MODCLUSTER-303
15415	MODCLUSTER-303	I am going to try to target this for WildFly 8. Hm, but seems like there is an issue in the native part - MODCLUSTER-282 ?	0	0	https://issues.jboss.org/browse/MODCLUSTER-303
15416	MODCLUSTER-296	what back-end are you using?	1	0	https://issues.jboss.org/browse/MODCLUSTER-296
15417	MODCLUSTER-282	I would like to eventually finish implementing this (i.e. keep this) but we haven't reached consensus on the topic yet.	0	0	https://issues.jboss.org/browse/MODCLUSTER-282
15418	MODCLUSTER-235	What is this issue duplicated by?	0	0	https://issues.jboss.org/browse/MODCLUSTER-235
15419	MODCLUSTER-211	i remember you had the problem while we were preparing the docs in Neuchatel. Do you remember the solution?	1	1	https://issues.jboss.org/browse/MODCLUSTER-211
15420	MODCLUSTER-211	Some engineering work is required per the note above, who is the core/main engineer on this project?	1	1	https://issues.jboss.org/browse/MODCLUSTER-211
15421	MODCLUSTER-211	Where does that parameter get put?	1	1	https://issues.jboss.org/browse/MODCLUSTER-211
15422	MODCLUSTER-211	Why does the port need to be changed?	0	0	https://issues.jboss.org/browse/MODCLUSTER-211
15423	MODCLUSTER-211	When SE Linux is enabled it gets in the way of random ports getting opened by application, and will block that access it SE Linux is turned on. By default several ports are "open" and port 10001 is one of those ports. Let me know if you have other questions?	0	0	https://issues.jboss.org/browse/MODCLUSTER-211
15424	MODCLUSTER-211	Hi guys, can someone explain this to me in simple terms for the release notes, please? Thanks.	0	0	https://issues.jboss.org/browse/MODCLUSTER-211
15425	MODCLUSTER-211	I was just a reporter of this issue, I would think a test case in a test plan would need to be created to verify this is either corrected or not. Do you agree?	0	0	https://issues.jboss.org/browse/MODCLUSTER-211
15426	MODCLUSTER-211	Thanks, that's what I was looking for. Can you also confirm that this is also not a problem for mod_cluster master (or 1.3.0.Final) on httpd 2.4 on Fedora 20/RHEL7?	0	0	https://issues.jboss.org/browse/MODCLUSTER-211
15427	MODCLUSTER-210	I have added the pdf to links in web site, is that was to you wanted?	0	0	https://issues.jboss.org/browse/MODCLUSTER-210
15428	MODCLUSTER-178	this ticket could be simple, however can become tricky if we want to expand this functionality. use case one is where we stick to a node (which is supported as we already know) second use case stick to a domain (which is what this ticket describes) third use case (which would be tricky i think) is to support combination of both !cluster-small.png! first we stick to a node, then in case of a failure of that node we failover to another node within same domain, and only in case of a second node failure, we will send request to Domain2. do you think we would ever want to implement third use case? i was looking into implementing this one.	1	1	https://issues.jboss.org/browse/MODCLUSTER-178
15429	MODCLUSTER-178	Rado, is there a special configuration that i should do to achieve case 3? i've tested it with sticky session true and force to false, and i am getting 503 if node goes down, here is my info: balancer: [1] Name: mycluster Sticky: 1 [JSESSIONID]/[jsessionid] remove: 0 force: 1 Timeout: 0 maxAttempts: 1 node: [2:2],Balancer: mycluster,JVMRoute: MyClusterE1,LBGroup: [MyDomain2],Host: 127.0.0.1,Port: 4453,Type: http,flushpackets: 0,flushwait: 1,ping: 10,smax: 0,ttl: 60,timeout: 0 host: 1 [myClusterAlias2] vhost: 1 node: 2 host: 2 [mycluterAlias1] vhost: 1 node: 2 context: 1 [/cluster] vhost: 1 node: 2 status: 1 my session cookie: id: 6a1f20eda1ec119e7fa168f766f2309bb5f42692.MyClusterD route: MyClusterD as you see node is MyClusterD and route is MyClusterD which died, however there is MyClusterE that serves /cluster context, however mod_cluster doesn't elect that node	1	1	https://issues.jboss.org/browse/MODCLUSTER-178

#	Requirement	Questions	Answered	Successfully Answered	URL
15430	MODCLUSTER-168	Thank you, Jean, do you know which release versions this switch will be available in? I mean 1.1.?	0	0	https://issues.jboss.org/browse/MODCLUSTER-168
15431	MODCLUSTER-165	It might, and this is the use case. I have setup httpd, I think I am good, but for whatever reason I never see my JBoss instances show up in the mod_cluster-manager. Wouldn't it be nice from the JBoss machine see if I can get across the network to that port. I am always amazed at the security lock downs that are at customer sites, and the headaches our massive use of ports creates for us in these environments.	0	0	https://issues.jboss.org/browse/MODCLUSTER-165
15432	MODCLUSTER-165	That is different then I was thinking, as I always have a browser or telnet available to test out #.#.#.#:6666 from my remote systems. What does this latch on to, and how does it work?	0	0	https://issues.jboss.org/browse/MODCLUSTER-165
15433	MODCLUSTER-160	@[~rhusar] It has been quite some time... Did you like the result configs? It seems to be largely outdated at the moment though.	1	0	https://issues.jboss.org/browse/MODCLUSTER-160
15434	MODCLUSTER-144	Isn't domain what you are looking for?	1	1	https://issues.jboss.org/browse/MODCLUSTER-144
15435	MODCLUSTER-125	So the process to permanently remove the node would be: 1) to set "auto-enable contexts" to false in mod-cluster-jboss-beans.xml 2) invoke the MBean to disable the node in the http proxy. 3) restart the node do your validation 4) set "auto-enable contexts" to true in mod-cluster-jboss-beans.xml 5) invoke the mbean to enable the node in the http proxy. Is that how you see things playing out? Any concern with triggering a deploy of mod-cluster by modifying the config on the file system? Of course "auto-enable contexts" could always be false if that is the desired behavior. However this is only the desired behavior when something has gone wrong so I wouldn't anticipate many folks setting "auto-enable contexts" to false as the default.	1	1	https://issues.jboss.org/browse/MODCLUSTER-125
15436	MODCLUSTER-90	Can you add the version to the /mod_cluster_manager handler as well ?	0	0	https://issues.jboss.org/browse/MODCLUSTER-90
15437	MODCLUSTER-84	Do you want to have MCMP commands corresponding to that? So that the httpd part of mod_cluster supports for a new set of ENABLE-DOM, DISABLE-DOM and STOP-DOM?	0	0	https://issues.jboss.org/browse/MODCLUSTER-84
15438	MODCLUSTER-75	In fact the balancers and the workers are not created correctly in mod_cluster. 1 - watchdog in server_rec principal 2 - via STATUS through isup in the VirtualHost that receives STATUS (6666 in the example). 3 - during the request in proxy_cluster_trans() in the VirtualHost that creates the request. What should we support: - create the balancers (and workers) in all VirtualHosts. - don't create the balancers, create the workers of the balancers defined by a ProxyPass and co, like: ProxyPass ... balancer://mycluster1	1	1	https://issues.jboss.org/browse/MODCLUSTER-75
15439	MODCLUSTER-54	We could maybe open this issue to community contributors? Since in 2.4 it is shipped by default IIUC. Do we know what are the main differences or is that a whole lot different API?	0	0	https://issues.jboss.org/browse/MODCLUSTER-54
15440	MODCLUSTER-25	I have added in buildbin.mod_cluster.sh: mvn -P dist package and copied the produced tar file in the output directory so that it get publish. Is that enough?	1	1	https://issues.jboss.org/browse/MODCLUSTER-25
15441	MODCLUSTER-1	It has been some time guys. What shall we do about this one?	0	0	https://issues.jboss.org/browse/MODCLUSTER-1
15442	MODE-2701	Same as for MODE-2698, [~open4storage], given that you wrote the initial implementation of the S3 backend, what do you think of this proposal?	0	0	https://issues.jboss.org/browse/MODE-2701
15443	MODE-2700	[~dalbani]: index providers don't have the ability to retrieve the stored text information really. Rather, they only store reverse indexes from the analyzed content of a text property towards the existing nodes which have that property. In other words, the entire architecture of indexes is based on being able to retrieve {{NodeKeys}} based on certain criteria. If you wanted to implemented this feature like it is in JR, you'd have to at least: * enhance the JCR-SQL2 supported grammar with an additional selector [see (this package|https://github.com/ModeShape/modeshape/tree/master/modeshape-jcr/src/main/java/org/modeshape/jcr/query/model)] and [parser support|https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/query/parse/JcrSql2QueryParser.java]) * add an additional custom column for each result set (see https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/query/JcrSqlQueryResult.java) * add the ability to extract the excerpt (most likely somewhere in the ScanningQueryEngine) I would recommend first debugging a JCR-SQL2 test case from https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/test/java/org/modeshape/jcr/JcrQueryManagerTest.java until you understand how the query parsing, optimization etc work and after that try to implement this feature.	0	0	https://issues.jboss.org/browse/MODE-2700
15444	MODE-2699	[~hchiorean] I've hacked a (very) dirty workaround in {{ScanningQueryEngine.java}}: {noformat} protected RowFilter createRowFilter(final Constraint constraint, final QueryContext context, Columns columns, QuerySources sources) { ... if (constraint instanceof FullTextSearch) { ... IndexDefinition ftsIndexDefinition = context.getIndexDefinitions().getIndexDefinitions().values().stream().filter(indexDefinition -> indexDefinition.appliesToProperty(propertyName)).findFirst().orElse(null); if (ftsIndexDefinition != null) { IndexProvider indexProvider = indexManager.getProvider(ftsIndexDefinition.getProviderName()); // if (indexProvider.supportsFullTextSearch()) ... return null; } ... {noformat} But given that you qualified the enhancement as "complex", I suppose that I completely missing the point, right? By the way, how come a FTS-based query get handled in {{ScanningQueryEngine.java}} (thus with _scanning_) — when does {{IndexQueryEngine}} comes into action then?	1	0	https://issues.jboss.org/browse/MODE-2699
15445	MODE-2698	Hey [~open4storage], given that you wrote the initial implementation of the S3 backend, what do you think of this proposal?	1	0	https://issues.jboss.org/browse/MODE-2698
15446	MODE-2698	[~dalbani] I actually didn't write the initial implementation. Just updated what was there to included username/passwords for 3rd party S3 compatible storage. That said. The lifecycle IMO is a good choice, it needs a way to manage expiration changes/updates to objects. Users will also need to be aware that a later restore won't bring back objects that expired (if these aren't backed up in some other way) in the event of a restore.	0	0	https://issues.jboss.org/browse/MODE-2698
15447	MODE-2696	How to test: 1/ Launch modeshape with *modeshape-cmis.war* 2/ Connect to a cmis client (like *CMIS workbench*) with the url *http://localhost:8080/modeshape-cmis/services11/cmis?wsdl* 3/ Go to *Types* next *Create type* next *Load type definition from xml* 4/ Select the *my-document.xml* file present on this Jira. => The exception is thrown *org.apache.chemistry.opencmis.commons.impl.jaxb.CmisException: Not supported!* I have test *my-document.xml* file with *OpenCMIS InMemory Repository*, the file is valid	1	0	https://issues.jboss.org/browse/MODE-2696
15448	MODE-2674	I have set {{DEFAULT_MIN_IDLE_POOL_SIZE = 5}}. Would 0 be a better default value for this configuration? See: https://github.com/ModeShape/modeshape/pull/1638/files#diff-7a0174f60ddf8693be6c0df41c7e249bR54	1	1	https://issues.jboss.org/browse/MODE-2674

#	Requirement	Questions	Answered	Successfully Answered	URL
15449	MODE-2634	Thanks for trying. There is nothing that would make the cache internally deadlock except for recursive computations, which is disallowed by ConcurrentHashMap's JavaDoc. Since only get and put style operations are used, I think the issue is a broken test assumption. Can you please try using Caffeine.executor(Runnable::run)? In CLHM the eviction and maintenance work is amortized on calling threads. In JDK8 we have ForkJoinPool.commonPool() that we can delegate to by default, rather than using calling threads or requiring our own. This means the eviction is asynchronous, so a race condition in the test is a likely culprit. The blocking on LazyCachedNode and racy inspection of the cache (e.g. count) might result in a flawed wait condition. By setting the executor the test code is simpler and might help identify the problem. The other possibility is that the test is relying on LRU behavior, whereas Caffeine takes into account frequency. This seems less likely, but is a significant difference.	1	1	https://issues.jboss.org/browse/MODE-2634
15450	MODE-2630	{quote} First, this is not a feature request. Second, in this particular form this looks like a question, not a JIRA issue. {quote} Fair point. I initially wanted to mark this as a bug, but since I do not know the expected behavior, I ended up classifying this as a feature request. {quote} Since there is no "default" number of DB calls involved when creating a node, everything is *highly context dependent*. {quote} I did not ask for the default number of DB calls and I hope I did not imply that anywhere, since I understand that such a metric varies based on the number of nodes. I only pointed to a pattern that I observed, where performance of the library degrades drastically as the number of existing nodes increases. In terms of being dependent on the context, I provided you with the context I have, i.e. a sample application that is linked in the description of this JIRA. That application has a trivial use case, which is based purely on ModeShape JCR implementation, without any additional context. {quote} if you suspect any incorrect behavior, you should profile and investigate your local application {quote} Using a profiler, I observed a pattern that stood out to me and I brought it up here. I do not know if it is correct or not, this is why I wanted to know your thoughts. {quote} ModeShape has been designed to efficiently handle a single node having a large number (100K-200k) of child nodes out-of-the-box. However, if an application stores a lot more nodes than that under a single parent, performance in terms of latency and memory consumption starts to decrease proportionally to the number of stored children. {quote} The quote above is taken from [here|https://docs.jboss.org/author/display/MODE50/Large+numbers+of+child+nodes]. Could you please explain what is meant by efficiency in handling a large number of nodes, preferably, by providing some numbers or pointing to the documentation that outlines the results of the abilities testing, which I assume was performed to come up with those conclusions? Note, that this question was also brought up in [here|https://issues.jboss.org/browse/MODE-2629?focusedCommentId=13293888&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13293888]. [~hchiorean], any help is greatly appreciated. My apologies if I misunderstood any of your points. By the way, thanks a lot for fast and helpful responses in here and on the forum.	1	1	https://issues.jboss.org/browse/MODE-2630
15451	MODE-2630	{quote} I only pointed to a pattern that I observed, where performance of the library degrades drastically as the number of existing nodes increases. In terms of being dependent on the context, I provided you with the context I have, i.e. a sample application that is linked in the description of this JIRA. That application has a trivial use case, which is based purely on ModeShape JCR implementation, without any additional context. {quote} the issue here is that you're describing a low level DB query (based on the internal ModeShape implementation) without mentioning anything about how you're using the JCR API - which is the only client interaction point. You mention you're adding versionable nodes, but without an exact code sample this is extremely vague considering that: * your nodes are versionable - are you using any versioning operations in between ? (if yes, then this is hugely expensive performance wise since every checkin/checkout creates a number of nodes under {{jcr:system}} in additional to the regular nodes) * are you using SNS or not - this can have a significant performance impact because not using SNS means additional reads and validations are required whenever a new node is added to a parent * are you using a short-lived or long lived sessions for each of the persist operations * the issue description keeps mentioning the word "query" - does that mean you're using JCR-SQL2 queries ? If yes, what are they ? ...etc I hope this explains better why without a running example/test case this type of issue is impossible to handle and is extremely vague. {quote} Using a profiler, I observed a pattern that stood out to me and I brought it up here. I do not know if it is correct or not, this is why I wanted to know your thoughts. {quote} hopefully I explained above why without something that is simple enough for me to run and look at locally it's impossible to say if this represents a problem or not. There are simply too many variables. {quote} The quote above is taken from here. Could you please explain what is meant by efficiency in handling a large number of nodes, preferably, by providing some numbers or pointing to the documentation that outlines the results of the abilities testing, which I assume was performed to come up with those conclusions? Note, that this question was also brought up in here. {quote} that statement refers to fact that we've tried to implement the node structure's internal representation so that it works (in a reasonable amount of time) in that particular case. However, if you look at that documentation page, the first thing is a large (hopefully) warning which explains why you should not use deep hierarchies. Being an API, you can use JCR essentially in an infinite number of ways (not to mention that you can also define your node hierarchy in an infinite number of ways with respects to the structure and datatypes). So there are no published numbers since there are an infinite number of uses cases really. What we did profile and test locally was done using the test cases from [this test suite and the JDBC subclass|https://github.com/ModeShape/modeshape/blob/master/modeshape-performance-tests/src/test/java/org/modeshape/test/performance/InMemoryPerformanceTest.java], which only tests reading/writing with or without SNS nodes based on {{nt:unstructured}}. The database that we did run these under were H2, MySql and Postgres primarily.	1	1	https://issues.jboss.org/browse/MODE-2630
15452	MODE-2613	[~oleg.kulikov]: can you please take a look at this ? When working on this, you can create several PRs especially if you have any questions or there are issues around the build. For example we can integrate first just the Maven module (without any tests) but with everything (i.e. Docker) configured and running and then integrate the tests.	1	1	https://issues.jboss.org/browse/MODE-2613

#	Requirement	Questions	Answered	Successfully Answered	URL
15453	MODE-2581	I'll try to express more precisely what I mean. 1. When I need indexing I add the following lines into WildFly config file: {noformat} <!-- Index provider --> <index-providers> <index-provider name="lucene" classname="lucene" module="org.modeshape.index-provider.lucene" lockFactoryClass="org.apache.lucene.store.NoLockFactory" directoryClass="org.apache.lucene.store.RAMDirectory" codec="Lucene53"/> </index-providers> <!-- A pair of indexes --> <indexes> <index name="ns:foo" provider-name="lucene" kind="text" synchronous="true" node-type="nt:base" columns="ns:foo(STRING)"/> <index name="ns:bar" provider-name="lucene" kind="text" synchronous="true" node-type="nt:base" columns="ns:bar(STRING)"/> </indexes> {noformat} Then Wildfly is started and I see: {noformat} 21:43:26,278 DEBUG [org.modeshape.jcr.RepositoryQueryManager] (modeshape-reindexing-7-thread-1) Performing full reindexing for repository 'test' and workspace 'default' ... 21:43:26,689 DEBUG [org.modeshape.jcr.index.lucene.LuceneIndexProvider] (modeshape-reindexing-7-thread-1) Enabled index 'ns:foo' from provider 'lucene' in workspace 'default' after reindexing has completed 21:43:26,690 DEBUG [org.modeshape.jcr.index.lucene.LuceneIndexProvider] (modeshape-reindexing-7-thread-1) Enabled index 'ns:bar' from provider 'lucene' in workspace 'default' after reindexing has completed {noformat} This is what I expected and queries produces expected results. It's fine - indexes were created and activated via configuration file. 2. Now I do not need indexing, so remove indexes section from configuration file.Only index-providers remains: {noformat} <index-providers> <index-provider name="lucene" classname="lucene" module="org.modeshape.index-provider.lucene" lockFactoryClass="org.apache.lucene.store.NoLockFactory" directoryClass="org.apache.lucene.store.RAMDirectory" codec="Lucene53"/> </index-providers> {noformat} Restarting WildFly gives me the following output: {noformat} 22:08:17,706 DEBUG [org.modeshape.jcr.RepositoryQueryManager] (modeshape-reindexing-6-thread-1) Performing full reindexing for repository 'test' and workspace 'default' ... 22:08:18,720 DEBUG [org.modeshape.jcr.index.lucene.LuceneIndexProvider] (modeshape-reindexing-6-thread-1) Enabled index 'ns:foo' from provider 'lucene' in workspace 'default' after reindexing has completed 22:08:18,720 DEBUG [org.modeshape.jcr.index.lucene.LuceneIndexProvider] (modeshape-reindexing-6-thread-1) Enabled index 'ns:bar' from provider 'lucene' in workspace 'default' after reindexing has completed {noformat} This is not what I expected and queries produces the same results as when indexes were defined in configuration file. So it doesn't matter whether indexes defined in configuration file or not - they are enabled till I remove index-providers at all. Is it a bug?	1	1	https://issues.jboss.org/browse/MODE-2581
15454	MODE-2568	[~rhauch]: I'm aware of the latest developments, my original plan was to try this using the latest/final 1.x (I think that's 1.0.9). As far as being faster than H2, IMO any "direct-disk" (i.e. non-JDBC) based storage _should_ be faster than going through JDBC. If you think {{MapDB}} is not worth pursuing (3.x does not have a final release yet, plus being a major rewrite does not give me the confidence of using it right of the bat) I can take a look at LevelDB and see if the batch writing mechanism actually works (if it does, we could use Level since no writes will overlap on a particular key) wdyt ?	0	0	https://issues.jboss.org/browse/MODE-2568
15455	MODE-2546	I agree that option 2 seems far more flexible, despite it being more complicated. My fear with #1 is that some databases might not really support `SELECT FOR UPDATE`, in which case to get ModeShape to work properly we might need to use an explicit locking column/table. I'd hate to have to do that, but option 2 gives us that flexibility. Can option 2 be implemented by relying upon the underlying database (e.g., `SELECT FOR UPDATE` within the transaction)? If so, option 2 gives us the best of both worlds.	1	1	https://issues.jboss.org/browse/MODE-2546
15456	MODE-2532	[~jkraushaar]: are you referring here to the ability of adding "transient sequencers" ? i.e. a sequencer that is only active for the lifetime of the VM but which is lost after a restart ? If yes, then we can consider this enhancement. If not, then the only way to add a persistent "dynamic" sequencer is via the configuration because it has to survive restarts.	1	0	https://issues.jboss.org/browse/MODE-2532
15457	MODE-2532	Would that really make sense though? IMO not really, since once you add a sequencer you'd want that to be present from there onwards, not only for the lifetime of that JVM. The current {{RepositoryConfiguration}} approach does allow "in-place updating" of the configuration of a running repository via [an update method|https://github.com/ModeShape/modeshape/blob/modeshape-4.4.0.Final/modeshape-jcr/src/main/java/org/modeshape/jcr/ModeShapeEngine.java#L561]. This is the same approach we're using from our AS kit. I understand the appeal of using annotations or a more OO based approach for configuration, but that would mean adding this support at a general repository level really, not just for sequencers. It does make sense, but it's a quite complicated and requires significant changes. If you or anyone else in the community wants to help and contribute, we're more than happy to review and integrate any submissions.	1	0	https://issues.jboss.org/browse/MODE-2532
15458	MODE-2518	This has already been discussed here: https://developer.jboss.org/message/943886 [~oleg.kulikov]: can you please take a look ? tx	1	1	https://issues.jboss.org/browse/MODE-2518
15459	MODE-2488	It seems that the Wildfly 9 Maven artifacts *are built using JDK 8*. This means that ModeShape's Wildfly subsystem cannot be updated to use Wildfly 9 *until ModeShape moves to Java 8*. Until that happens, the only thing that we can do is to try and check (and change if possible) if the Wildfly 8 subsystem is _binary compatible_ (i.e. will run) on Wildfly 9. [~rhauch] thoughts ?	1	1	https://issues.jboss.org/browse/MODE-2488
15460	MODE-2488	What Wildfly requires is unclear. [This page|https://docs.jboss.org/author/display/WFLY9/Getting+Started+Guide] suggests that Wildfly 9 only requires Java 7 to run, though it's not terribly clear since that page is in the Wildfly 9 documentation but the text on the page still refers to Wildfly 8. The [source code|https://github.com/wildfly/wildfly/tree/9.x] appears to require JDK 7 (and I can't find it in the top-level POM). We need to find out for sure. But, if Wildfly 9 does require Java 8 (source and target level), then we need to have a separate conversation about when ModeShape will require a minimum of Java 8 for **running**. Either way, we should be able to start requiring JDK 8 for building, even if we stick to source-level of 7 and target 7 for runtime. Have you tried this?	1	1	https://issues.jboss.org/browse/MODE-2488
15461	MODE-2482	Do we have test cases that verify this?	0	0	https://issues.jboss.org/browse/MODE-2482
15462	MODE-2458	Did you investigate the option of using/creating your own [Authentication/Authorization providers|https://docs.jboss.org/author/display/MODE/Custom+authentication+providers] ? That would be the preferred way to map your AD roles to one of the 4 roles required by the repository. These 4 roles are repository internal-roles which were not designed with the purpose of being configurable. Instead, we have the previously mentioned interfaces allowing you to plug in your custom security provider and tailor the permission checking in whichever way your application requires.	1	0	https://issues.jboss.org/browse/MODE-2458

#	Requirement	Questions	Answered	Successfully Answered	URL
15463	MODE-2452	{quote} What if the backup operation wrote to the data directory (perhaps a specific location that incorporated the timestamp that the backup was started) {quote} Since the REST service can be running in "any environment" (not just a container) maybe we should also support passing an optional {{location}} parameter which can be resolved as a system property - e.g. {{/backup?location=jboss.server.data.dir}}. If no such location is given as a parameter, we can simply use a "writable" default like {{java.io.tmpdir}}. The backup request can then return if successful something like {code:json} "location": "file:///<some_path>/<repository_name>_backup_timestamp" {code} The restore operation would then require this {{location}} parameter and read it when performing the restore. If not provided by the user, the restore request would fail with a {{400}}. wdyt ?	1	1	https://issues.jboss.org/browse/MODE-2452
15464	MODE-2452	Passing parameters in certainly makes things generic, but IMO that does not make them as useful or easy to use. I'm not convinced that it's worth exposing the option of writing a backup somewhere else: it complicates the behavior, and it may cause problems or performance issues if people specify locations like NFS directories. For the EAP/Wildfly subsystem, it would be better if the backup were smarter and simply did the right thing -- which IMO is writing out to the a directory within the data directory. And restore should then require the name of the backup, which can be used to determine the directory (within the data directory) where the backup exists. Can this relative location be set within the service's {{web.xml}} file? If not set, then perhaps backup is disabled (that's a decent option, IMO). The subsystem-based service would refer to the server's data directory? (IIRC, you can refer to system variables like {{jboss.server.data.dir}} within the {{web.xml}}; if not, then that might be an issue.) By putting the relative location within the {{web.xml}} file also allows other non-EAP/Wildflly installations to control where backups are written/read. Just my opinion.	0	0	https://issues.jboss.org/browse/MODE-2452
15465	MODE-2452	{quote} For the EAP/Wildfly subsystem, it would be better if the backup were smarter and simply did the right thing – which IMO is writing out to the a directory within the data directory. And restore should then require the name of the backup, which can be used to determine the directory (within the data directory) where the backup exists. {quote} How would this account for the difference between domain mode ({{jboss.domain.data.dir}}) and standalone mode ({{jboss.server.data.dir}}) ? In other words, there is no "unifying" data directory for both standalone and domain mode. {quote} Can this relative location be set within the service's web.xml file? If not set, then perhaps backup is disabled (that's a decent option, IMO). The subsystem-based service would refer to the server's data directory? (IIRC, you can refer to system variables like jboss.server.data.dir within the web.xml; if not, then that might be an issue.) {quote} AFAIK system property substitution in {{web.xml}} is not part of the EE spec, and [this issue|https://issues.jboss.org/browse/AS7-2931] seems to confirm that. So we can't really use it. Therefore I'm not sure I see how configuring the location in the {{web.xml}} file would solve this problem, particularly for JBoss AS. Of course, we can always go the easy route and use convention, meaning that the code can "hardcodedly" look in order at one of the following locations: {{jboss.domain.data.dir}}, {{jboss.server.data.dir}} and fallback perhaps to {{user.home}}, taking the first accessible one and using it as the root folder for backup/restore.	1	1	https://issues.jboss.org/browse/MODE-2452
15466	MODE-2452	{quote} Of course, we can always go the easy route and use convention, meaning that the code can "hardcodedly" look in order at one of the following locations: jboss.domain.data.dir, jboss.server.data.dir and fallback perhaps to user.home, taking the first accessible one and using it as the root folder for backup/restore. {quote} How about use the value in the {{web.xml}} if there (which is useful for non-EAP/Wildfly deployments), and otherwise try to resolve the well-known variables as you suggested?	0	0	https://issues.jboss.org/browse/MODE-2452
15467	MODE-2443	Let's try to clarify a bit the meaning of file-chooser capabilities. Here are possible options: - point to path on local file system; - point to path on server's file system; - any valid URI. These options allow us either backup/restore on the server side, backup/restore and download/upload to the local filsystem or even specify other remote storage. What do you think?	1	1	https://issues.jboss.org/browse/MODE-2443
15468	MODE-2422	You could reverse the subsqueries, but then makes it difficult to return {{artifact1}} properties in the select. I'm not even sure what the {{'element'}} represents -- is that a literal? Can you describe the node structure you're trying to query? (Perhaps that's better done on the forums.) This is very unlikely to be implemented, since it would require significant changes to the query engine. Non-correllated subqueries (what we have now) are very easily implemented compared to correlated subqueries.	1	1	https://issues.jboss.org/browse/MODE-2422
15469	MODE-2411	[~wesssel]: can you please be more clear what the issue is ? ModeShape does not use login modules, but rather security domains defined as {{security-domain}} elements in the EAP configuration. Whatever such a security domain contains, is out of the scope of ModeShape.	1	1	https://issues.jboss.org/browse/MODE-2411
15470	MODE-2411	[~wesssel], we are definitely not "bypassing" the security cache. We're literally using the JAAS API, as seen [here|https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/security/JaasProvider.java#L84]. Have you tried using the [JaasCredentials|https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr-api/src/main/java/org/modeshape/jcr/api/JaasCredentials.java] rather than a password? With this approach, you create the LoginContext and pass that in every time. This might give you the kind of control you need to make more use of the cache (if that's even possible).	1	1	https://issues.jboss.org/browse/MODE-2411
15471	MODE-2411	[~wesssel], sure your own LDAP login module would work and would be quick for you. But we'd love to have something similar so that our OOTB JAAS integration supports caching (if it is available on the classpath). Did support say how you're supposed to get a hold of the CacheableManager instance?	1	1	https://issues.jboss.org/browse/MODE-2411
15472	MODE-2410	[~rhauch]: this method was added as part of https://issues.jboss.org/browse/MODE-1286, but it is unused atm. and I don't understand its purpose. Does it make sense to keep it ?	1	1	https://issues.jboss.org/browse/MODE-2410
15473	MODE-2407	Yes, not looking below a given path would mean that fewer nodes are iterated. This would however place the "burden" on the client of making sure that nodes are stored in the correct places. Since this is a performance issue, have you looked at (profiled) where the time is actually spent ? (MapDB is the only index provider we have atm and we are aware of this performance issue: https://issues.jboss.org/browse/MODE-2387) The current indexing mechanism already allows index definitions to apply only to certain workspaces (by name only), meaning nodes not belonging to those workspaces will not be looked at: {code:javascript} "indexes" : { "index" : { "kind" : "value", "provider" : "invalid-provider", "nodeType" : "mix:title", "workspaces" : "<regex>" "columns" : "jcr:title(STRING)" } } {code} So one option would be to extend/implement this feature under the {{workspaces}} attribute and extend the semantic from a workspace name regex to a full-blown workspace path.	1	1	https://issues.jboss.org/browse/MODE-2407

#	Requirement	Questions	Answered	Successfully Answered	URL
15474	MODE-2407	{quote} (or even adding a new optional paths attribute if needed, if that makes the implementation easier) {quote} Yes, that's exactly what I had in mind. {quote} But I also agree with Horia Chiorean that profiling is the only way to fully understand where the time is spent. If the vast majority of the time is spent in MapDB, then when MODE-2387 is addressed will this still be a problem? {quote} OK, I'll try if playing with MapDB settings makes any (large) difference for me. By the way, would you have some ModeShape-specific tips on how to profile it easily? Still, imagine the case with millions of nodes in a repository and that I want to only index a few dozen nodes -- always sitting at the top of the tree by definition. Well, that's more or less my use case. Even if MapDB was über-fast, I think it would be a waste of CPU cycles to iterate over those millions of node... only to actually index a couple dozens of them.	1	1	https://issues.jboss.org/browse/MODE-2407
15475	MODE-2406	[~dalbani]: shouldn't 8.2 be compatible with 8.1 ? Have you tried deploying the 8.1 kit on 8.2 ?	1	1	https://issues.jboss.org/browse/MODE-2406
15476	MODE-2398	[~blafond], [~van.halbert], Would you care to comment on the priority of this, for both Modeshape and future support in Designer/Komodo?	1	0	https://issues.jboss.org/browse/MODE-2398
15477	MODE-2394	Is this a change that needs to be back ported to 3.8.x?	1	1	https://issues.jboss.org/browse/MODE-2394
15478	MODE-2394	Paul, In fixing the VDB sequencer, did u by chance fix the parsing to load the relational_schema (model) and relational_catalog (vdb) tables accordingly? I see vdb_model and vdb_virtualDatabase get updated, but I would have also thought the relational tables would have also.	1	1	https://issues.jboss.org/browse/MODE-2394
15479	MODE-2394	Hey [~van.halbert], I have created another PR which actually modifies the VdbSequencer and its accompanying classes to handle the changes to the Vdb schema that have already gone into the vdb.cnd. Is this what you are referring to or if not can you clarify please?	1	1	https://issues.jboss.org/browse/MODE-2394
15480	MODE-2387	In responce to the PR discussion on github, I guess this belongs here rather than there. It is not (a bit) slower, it's (in my case) nearly unusable. Infinispan is fast if configured correctly, but the MapDB backed Index is very slow because of file sync and it gets even worse with multiple threads running. So currently I only have one option: Patching Modeshape. A configuration setting that doesn't hurt anyone if it is disabled by default would help me not to patch. If you could tell which class boots up the index provider I could try to implement that myself so no work would have to be done on your side and we end up with a "hidden" conf setting that I may use as long as no other index provider is in place. Or maybe better, I could just create a fork of the LocalIndexprovider and use that? Maybe easier for now. Would that work, jsut clone the whole provider, put it in a differnt package and use that clone from the config?	1	1	https://issues.jboss.org/browse/MODE-2387
15481	MODE-2363	[~bes82]: what's the feature that's being requested here ? Shouldn't this be a question rather either on the ModeShape on ISPN forums around case-sensitive/insensitive storage in the Infinispan StringBasedJDBCStore ?	1	1	https://issues.jboss.org/browse/MODE-2363
15482	MODE-2363	But you are right, it's because of MySQL. But still I don't understand: hasNodeType(typeName) asks for a key that can't be found. But then when registerNodeType(type, true) tries to create the same key it suddenly can find the other key and complains. Shouldn't the behaviour be consistent throughout this two methods? It would also be nice if something about the case-sensitivity could be mentioned here: modeshape-examples/modeshape-jdbc-store-example/src/main/resources/infinispan-configuration.xml	1	1	https://issues.jboss.org/browse/MODE-2363
15483	MODE-2363	Could you please add this as xml comment to modeshape-examples/modeshape-jdbc-store-example/src/main/resources/infinispan-configuration.xml. {code} <!-- When using MySsql it's important to use a cs collation for the idColumn. Otherwise you might run into issues when defining case sensitive nodeTypes <stringKeyedTable prefix="ISPN_STRING_TABLE" createOnStart="true" dropOnExit="false"> <idColumn name="ID_COLUMN" type="VARCHAR(255) CHARACTER SET latin1 COLLATE latin1_general_cs"/> <dataColumn name="DATA_COLUMN" type="LONGBLOB"/> <timestampColumn name="TIMESTAMP_COLUMN" type="BIGINT"/> </stringKeyedTable> --> {code}	1	1	https://issues.jboss.org/browse/MODE-2363
15484	MODE-2363	That sounds like a great idea. [~bes82], would you mind opening a pull-request against https://github.com/modeshape/modeshape-examples? It can simply add an `infinispan-mysql.xml` file to the module that uses JDBC (e.g., it [this directory|https://github.com/ModeShape/modeshape-examples/tree/master/modeshape-jdbc-store-example/src/main/resources]).	1	1	https://issues.jboss.org/browse/MODE-2363
15485	MODE-2363	[~bes82], you're right about "{{_ci}}". Other resources on the web suggest using {{COLLATE=utf8_bin}} for collation in place of the default. WDYT?	1	1	https://issues.jboss.org/browse/MODE-2363
15486	MODE-2349	[~dalbani]: are you referring to the Wildfly (XML) configuration or the JSON configuration ? The JSON configuration is documented here: https://docs.jboss.org/author/display/MODE40/Query+and+search#Queryandsearch-Usingindexes. If you're referring to the Wildfly configuration, did you look at the {{modeshape.xsd}} that's installed with the WF kit ? That will always explain the significance of each attribute.	1	1	https://issues.jboss.org/browse/MODE-2349
15487	MODE-2349	Indeed, it *is* documented -- I'm sorry, how could I miss that page!? To answer to second question, yes, I was looking for documentation for an XML-based configuration in WildFly. But it seems that using the documentation for the JSON environment is just fine. However, I've had a brief look at it and there's some configuration attributes that are not mentioned in the XSD shipped with WildFly. Namely {{workspaces}} and {{synchronous}}. Should it be added the the XML schema definition?	1	1	https://issues.jboss.org/browse/MODE-2349
15488	MODE-2334	[~rstachowiak], ModeShape should not be using the TCCL, but instead when ModeShape needs to find arbitrary classes (e.g., sequencers, connectors, text extractors, index providers, etc.) it should be possible to specify the "classloader" in the configuration. ModeShape then delegates all classloading to the {{Environment}} instance that the {{RepositoryConfiguration}} instance optionally references. We may want to implement a new {{Environment}} class (perhaps subclassing {{LocalEnvironment}}) that is OSGi-specific and that maps "classloader" names to OSGi bundles. But, yes, using a named classloader in the config should be the preferred approach for finding security providers. Have you found others components loaded without using a named classloader?	1	1	https://issues.jboss.org/browse/MODE-2334
15489	MODE-2334	The Environment abstraction is actually the only place where we should be changing/using different class-loaders because when running in EAP/WF we have a special implementation that works in that case. So in the case of {{LocalEnviroment}}, ATM when we instantiate external components (like security providers) we use as a fallback loader {{getClass().getClassLoader()}} optionally passing in a list of Strings representing URI(s) coming from the {{classpath}} configuration attribute. These URIs (if present) will be passed into another URL-based CL: https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/LocalEnvironment.java#L132 So we could change this to use as the fallback loader the TCCL, but the question is why won't the classpath approach work in an OSGI enviroment ? Something like: {code} "security" : { "providers" : [{ "classname" : "servlet" "classpath": "my.jar" }] {code} where _my.jar_ can be any valid {{URI}} that is accessible at runtime by a {{URLClassLoader}}.	1	1	https://issues.jboss.org/browse/MODE-2334
15490	MODE-2334	[~rhauch]: how would we actually use this 3rd implementation ? The logic of {{Environment}} embedding is something that's done when creating a {{RepositoryConfiguration}}.	1	1	https://issues.jboss.org/browse/MODE-2334

#	Requirement	Questions	Answered	Successfully Answered	URL
15491	MODE-2334	[~hchiorean], that sounds great. I've already +1 the PR. But doesn't the {{JcrRepository}} class still need to be changed to use {{Environment}} for the security providers? If so, then that should be fixed as part of this issue. [~rstachowiak], if you think an OSGi-specific {{Environment}} class is beneficial and are willing to help us out with it, please log a new issue as a feature request.	1	1	https://issues.jboss.org/browse/MODE-2334
15492	MODE-2333	OK, I'll follow the steps you described then -- I'm still learning about the versioning thing, and I thought that checked-in nodes were set as read-only for good. Taken as a whole, coming from the traditional SQL world, I'm kind of surprised of the current behavior. Changing an SQL schema and its constraints is an all-or-nothing action. Either the database fully accepts the change or it doesn't if there's a mismatch with the actual data. But I can't think of a situation where the database wouldn't let me *read* some data because I have changed the schema inadvertently, say, a few weeks ago. -I see constraint definitions as requirements for what I *put* in the repository, not what I *read* from it.- In practice, how do I validate a modification of constraints? Do I have to make a offline copy of the repository, change the CND and then try to read all the properties of all the nodes, to be really sure that ModeShape won't throw an error out of the blue??	1	0	https://issues.jboss.org/browse/MODE-2333
15493	MODE-2332	After working with these changes in my app, I'm wondering if it would be cleaner to actually break JdbcMetadataConnector into two different classes: 1) JdbcConnector, extends ReadOnlyConnector - defines connectivity properties like dataSourceJndiName, username, password, connection pool settings, etc. 2) JdbcMetadataConnector, extends JdbcConnector - defines metadata-specific retrieval methods and properties like metadataCollector, metadataRetrievers, some of the init methods, etc. My extension is to retrieve data from a database and present it as nodes. It seems that extending the JdbcMetadataConnector would imply that my functionality is working the the DB metadata itself, which it's not. Though I wouldn't want to duplicate the DB connectivity structure that's currently in place. Separating these two aspects of the class would make for a cleaner extension for me. Would this separation make sense in other cases as well? Thanks! nikhil	1	1	https://issues.jboss.org/browse/MODE-2332
15494	MODE-2327	We have trimmed names ever since the project's inception, so changing that behavior might be a problem. Can you use the trimmed name but store the _actual file name_ in a separate property on those nodes that really need it?	0	0	https://issues.jboss.org/browse/MODE-2327
15495	MODE-2327	{quote} We have trimmed names ever since the project's inception, so changing that behavior might be a problem. {quote} I agree that changing the default behavior could have (unforeseen) consequences. It could be implemented as a new option then, to avoid risking problems with existing installations. {quote} Can you use the trimmed name but store the actual file name in a separate property on those nodes that really need it? {quote} Some uses case don't lend themselves easily to such workarounds: e.g. the WebDAV example I mentioned before. Plus storing the "original" name is a dedicated property introduces some (unwanted) complexity.	0	0	https://issues.jboss.org/browse/MODE-2327
15496	MODE-2327	By the way, I've noticed than {{String.trim()}} removes all leading / trailing characters with a code <= 0x20. So that's more than only the "space" characters that I was talking about. I know these are control characters, so somewhat rare, but they are not specifically mentioned in the JCR spec either. To sum things up, I maintain my point of view : the current implementation in ModeShape arbitrarily forbids a set of leading / trailing characters, for no good reason in my opinion. I just can't see why [0x00 - 0x20] characters are treated in a way different that all the other Unicode space characters for instance (e.g. NBSP). Apart from characters used in the JCR node name representation and syntax, *all* characters should be treated equally -- up to the user to prune any characters that he doesn't see fit. Using ModeShape as a file repository, this limitation on node names is precisely what introduces a "significant degree of complexity": why should I handle some filenames differently than others?? Compared to a conventional filesystem, have we ever seen one which forbids leading / trailing spaces?? The workarounds I've been presented with are error-prone and make me deviate from the standard implementation (e.g. via a specific name encoder). Plus, as I've mentioned before, accessing a repository through WebDAV (i.e. my use case for file synchronization) don't let me use this trick. Having said that, I leave the decision to your appreciation.	1	0	https://issues.jboss.org/browse/MODE-2327
15497	MODE-2323	Merged the PR into the {{master}} branch. This does not appear to fix the MODE-2280 issue, so do we want to backport this to the {{3.x}} (and {{3.8.x-prod-ip6.1}} branch? I'm not convinced that's necessary or even a good idea. [~hchiorean], thoughts?	1	1	https://issues.jboss.org/browse/MODE-2323
15498	MODE-2323	Not sure actually. Maybe port it ? I like very much the fact that we've simplified this logic a great deal.	1	1	https://issues.jboss.org/browse/MODE-2323
15499	MODE-2321	Reopening, because the change is incomplete: ModeShape also defines an explicit version of Apache POI, which is required by Tika. This change updated the Tika versions but didn't update the POI versions which means that now Tika (depending on the artifacts it's processing) can crash with: {code} 11:01:56,087 ERROR [org.modeshape.extractor.tika.TikaTextExtractor] (modeshape-text-extractor-32-thread-1) Error while extracting text : org.apache.poi.xssf.eventusermodel.XSSFSheetXMLHandler.<init>(Lorg/apache/poi/xssf/model/StylesTable;Lorg/apache/poi/xssf/model/CommentsTable;Lorg/apache/poi/xssf/eventusermodel/ReadOnlySharedStringsTable;Lorg/apache/poi/xssf/eventusermodel/XSSFSheetXMLHandler$SheetContentsHandler;Lorg/apache/poi/ss/usermodel/DataFormatter;Z)V: java.lang.NoSuchMethodError: org.apache.poi.xssf.eventusermodel.XSSFSheetXMLHandler.<init>(Lorg/apache/poi/xssf/model/StylesTable;Lorg/apache/poi/xssf/model/CommentsTable;Lorg/apache/poi/xssf/eventusermodel/ReadOnlySharedStringsTable;Lorg/apache/poi/xssf/eventusermodel/XSSFSheetXMLHandler$SheetContentsHandler;Lorg/apache/poi/ss/usermodel/DataFormatter;Z)V at org.apache.tika.parser.microsoft.ooxml.XSSFExcelExtractorDecorator.processSheet(XSSFExcelExtractorDecorator.java:190) at org.apache.tika.parser.microsoft.ooxml.XSSFExcelExtractorDecorator.buildXHTML(XSSFExcelExtractorDecorator.java:135) at org.apache.tika.parser.microsoft.ooxml.AbstractOOXMLExtractor.getXHTML(AbstractOOXMLExtractor.java:112) at org.apache.tika.parser.microsoft.ooxml.XSSFExcelExtractorDecorator.getXHTML(XSSFExcelExtractorDecorator.java:94) at org.apache.tika.parser.microsoft.ooxml.OOXMLExtractorFactory.parse(OOXMLExtractorFactory.java:112) at org.apache.tika.parser.microsoft.ooxml.OOXMLParser.parse(OOXMLParser.java:82) at org.apache.tika.parser.CompositeParser.parse(CompositeParser.java:244) at org.modeshape.extractor.tika.TikaTextExtractor$1.execute(TikaTextExtractor.java:143) at org.modeshape.jcr.api.text.TextExtractor.processStream(TextExtractor.java:89) [modeshape-jcr-api-4.1-SNAPSHOT.jar:4.1-SNAPSHOT] at org.modeshape.extractor.tika.TikaTextExtractor.extractFrom(TikaTextExtractor.java:133) at org.modeshape.jcr.TextExtractors$Worker.run(TextExtractors.java:171) [modeshape-jcr-4.1-SNAPSHOT.jar:4.1-SNAPSHOT] at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) [rt.jar:1.7.0_40] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) [rt.jar:1.7.0_40] at java.lang.Thread.run(Thread.java:724) [rt.jar:1.7.0_40] {code}	0	0	https://issues.jboss.org/browse/MODE-2321

#	Requirement	Questions	Answered	Successfully Answered	URL
15500	MODE-2316	[~bes82]: should this be an issue in JIRA or a topic in the [community|https://developer.jboss.org/en/modeshape] ?	1	1	https://issues.jboss.org/browse/MODE-2316
15501	MODE-2316	One more thing I didn't understand: After registering a new index the reindex thread logs that it is reindexing every node. During this time "something" constantly accesses / writes to the index directory/files. I don't know too much about the internals, but is this file access really needed? If I define an index on a nodetype that has no nodes yet, then every node has to be considered, I understand. But as soon as the node has been loaded and its nodeType is known to not match, couldn't there be an early return? If this is already the case than please excuse this comment, as said: I don't know that much about the internals, just looking at the file access and logs.	0	0	https://issues.jboss.org/browse/MODE-2316
15502	MODE-2311	The query plan structure changes frequently and is likely to continue to do so while the indexing system matures over the next few releases. Exposing in a public API the existing internal object structure is not an option. Creating a public structure explicitly for this is also far more trouble than its worth. Besides, most DBMSes expose the query plan in tabular form (again, which would require a lot of work). On the other hand, the amount of work necessary to create a simple, custom analyzer is quite low. Every line in the query plan is separate operation, and the indentation level describes the structure (if that's even necessary). Have a look at [this test code|https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/test/java/org/modeshape/jcr/ValidateQuery.java#L349] that analyzes the query plan using simple regular expressions. The index validation logic does not currently handle queries with joins, but that could easily be done with custom code (or with a pull-request).	0	0	https://issues.jboss.org/browse/MODE-2311
15503	MODE-2305	It seems {{3.10.1}} contains a valid fix, so this will be the version that will be brought in via the BOM in both {{3.x}} branches.	1	1	https://issues.jboss.org/browse/MODE-2305
15504	MODE-2304	What do you mean by: {quote} modified the configuration and then used a maven uberwar to bundle that and my own application war {quote} The explorer should be a generic webapp, in that it works with any ModeShape repository. What parts of it do you need to modify and why ?	1	1	https://issues.jboss.org/browse/MODE-2304
15505	MODE-2304	[~oleg.kulikov], you seemed to allude to a potential problem: is it possible to create a single JAR that embeds all of the resources (not just class files) for the repository explorer, so that when it is assembled into the WAR file all the GWT and non-GWT resources can be found correctly? I suspect this will be the most challenging part.	1	1	https://issues.jboss.org/browse/MODE-2304
15506	MODE-2304	If this is not technically possible (for GWT reasons) then IMO we have only 1 option and that is to only distribute the full WAR file. This means that we can/should get rid of the {{modeshape-web-explorer-war}} module since its content is identical to {{modeshape-web-explorer}}. The reason is that without a lightweight JAR, there is no "Bare WAR" concept that would work and/or be useful IMO. EDIT: actually, isn't there a way we could leverage JCA and still have the "Bare WAR" approach ? Have the webapp connect to another ModeShape repository via JCA ?	1	1	https://issues.jboss.org/browse/MODE-2304
15507	MODE-2304	{quote} Horia Chiorean, I agree with you regarding modeshape-web-explorer-war. JCA doesn't seem like a solution. As maximum we can expose JCR api but we need quite more (backup/restore, stats, etc) {quote} Doesn't the explorer use only the JCR API and ModeShape's public API? If so, a lightweight WAR should still be an option, since those API JARs should be publicly available via the JCA deployment.	1	1	https://issues.jboss.org/browse/MODE-2304
15508	MODE-2304	Here's what I see as three possible ways for users/developers to get the explorer app: # *Heavyweight WAR* - The purpose is to be a standalone WAR file that is easy to deploy and use, because it runs out of the box. It has everything and is completely self-contained: it has the explorer classes, all GWT-related artifacts, repository connector definition, ModeShape connector implementation, *and* all of the JCR API and ModeShape-specific JARs with a repository configuration file. # *Lightweight WAR* - The purpose is to be a WAR file that must be deployed on a container where ModeShape is running elsewhere in the container. This is similar to the WAR file for the JBoss environment, but in this case there's nothing specific about JBoss. It includes only the explorer-specific code: the explorer classes, all GWT-related artifacts, repository connector definition, and ModeShape connector implementation (which IIUC only uses the JCR API and ModeShape public API; if not, it should). It does *not* contain the JCR API or ModeShape-specific JARs (since these should be found elsewhere on the classpath), and since the WAR is not starting or running its own repository, it has no repository configuration file(s). # *Custom WAR* - This is custom built by our users, and it's a simple WAR application that (let's say) is Maven based. They have their custom code, but they also want the explorer functionality and/or the REST service and/or the WebDAV service and/or the CMIS web service. To get this all into their own custom app, they need to add Maven dependencies on the explorer, REST service, and WebDAV artifacts, which of course have dependencies on other libraries. Because they're assembling multiple end points, they have to build a custom {{web.xml}} that includes snippets for the explorer, REST service, WebDAV service, and/or CMIS service. It probably also has a dependency on ModeShape (public APIs and impls), provides one or more repository configurations, and manages the ModeShape engine; if not, then it's looking up the repository from JNDI and has access to the JCR API and ModeShape API. When their application is built, everything gets assembled into a single, runnable WAR file that is easily deployed. We currently build something akin to (1), but twice ({{web/modeshape-web-explorer}} and {{web/modeshape-web-explorer-war}}) and these are nearly identical. (I'm not sure what the difference is, actually. But they're both very large and include everything under the sun.) Option (2) is very similar to what {{deploy/jbossas/modeshape-jbossas-web-explorer-war}} is, though it's customized for the JBoss container. I'm not sure whether it's feasible to have this for other containers, since sharing classpath resources is really challenging. (For example, it used to be possible to do this with old Tomcat only by putting shared classpath resources on the global classpath. This is really a hack, and many admins won't want to do it. But for those that do, this might be the best way.) I'm not convinced we need this for non-JBoss containers. Option (3) lets somebody easy build their own WAR that uses an arbitrary combination of our web-based functionality. Yes, they have to manage their {{web.xml}} and it has to be correct, but theoretically this should be easy by simply depending upon multiple ModeShape artifacts. Presumably these are all JARs, and would be except for the explorer functionality. If we don't want to do (2) and want to do (1) and (3), then for (3) we want to provide an artifact to make it easy to reuse the explorer functionality. All of our other services are just JARs, and this works really well and is very conventional. But GWT throws a wrench into the gears and makes it challenging. My question is: how can we package the explorer into one or more artifacts so that we can achieve option (3) above? Can we do a JAR file and put everything in it? Should it be a JAR for our classes (including generated classes) but then with the GWT resources put into a ZIP file that has to be exploded into the custom web app codebase? What other options are there?	1	1	https://issues.jboss.org/browse/MODE-2304

#	Requirement	Questions	Answered	Successfully Answered	URL
15509	MODE-2304	{{http://localhost:8090/modeshape-explorer}} results in: {code} Caused by: java.lang.ClassNotFoundException: org.modeshape.common.logging.Logger at java.net.URLClassLoader$1.run(URLClassLoader.java:366) at java.net.URLClassLoader$1.run(URLClassLoader.java:355) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:354) at java.lang.ClassLoader.loadClass(ClassLoader.java:424) at java.lang.ClassLoader.loadClass(ClassLoader.java:357) at org.eclipse.jetty.webapp.WebAppClassLoader.loadClass(WebAppClassLoader.java:436) at org.eclipse.jetty.webapp.WebAppClassLoader.loadClass(WebAppClassLoader.java:389) at org.modeshape.web.server.JcrServiceImpl.<clinit>(JcrServiceImpl.java:76) at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeConstructorAccessorImpl.newInstance(NativeConstructorAccessorImpl.java:57) at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(DelegatingConstructorAccessorImpl.java:45) at java.lang.reflect.Constructor.newInstance(Constructor.java:526) at java.lang.Class.newInstance(Class.java:374) at org.eclipse.jetty.server.handler.ContextHandler$Context.createInstance(ContextHandler.java:2307) at org.eclipse.jetty.servlet.ServletContextHandler$Context.createInstance(ServletContextHandler.java:1164) at org.eclipse.jetty.servlet.ServletContextHandler$Context.createServlet(ServletContextHandler.java:1151) at org.eclipse.jetty.servlet.ServletHolder.newInstance(ServletHolder.java:976) at org.eclipse.jetty.servlet.ServletHolder.initServlet(ServletHolder.java:521) at org.eclipse.jetty.servlet.ServletHolder.getServlet(ServletHolder.java:422) at org.eclipse.jetty.servlet.ServletHolder.handle(ServletHolder.java:669) at org.eclipse.jetty.servlet.ServletHandler$CachedChain.doFilter(ServletHandler.java:1568) at org.modeshape.web.AuthFilter.doFilter(AuthFilter.java:61) at org.eclipse.jetty.servlet.ServletHandler$CachedChain.doFilter(ServletHandler.java:1548) at org.eclipse.jetty.servlet.ServletHandler.doHandle(ServletHandler.java:524) at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:143) at org.eclipse.jetty.security.SecurityHandler.handle(SecurityHandler.java:544) at org.eclipse.jetty.server.session.SessionHandler.doHandle(SessionHandler.java:221) at org.eclipse.jetty.server.handler.ContextHandler.doHandle(ContextHandler.java:1110) at org.eclipse.jetty.servlet.ServletHandler.doScope(ServletHandler.java:453) at org.eclipse.jetty.server.session.SessionHandler.doScope(SessionHandler.java:183) at org.eclipse.jetty.server.handler.ContextHandler.doScope(ContextHandler.java:1044) at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:141) at org.eclipse.jetty.server.handler.ContextHandlerCollection.handle(ContextHandlerCollection.java:199) at org.eclipse.jetty.server.handler.HandlerCollection.handle(HandlerCollection.java:109) at org.eclipse.jetty.server.handler.HandlerWrapper.handle(HandlerWrapper.java:97) at org.eclipse.jetty.server.Server.handle(Server.java:459) at org.eclipse.jetty.server.HttpChannel.handle(HttpChannel.java:280) at org.eclipse.jetty.server.HttpConnection.onFillable(HttpConnection.java:229) at	1	1	https://issues.jboss.org/browse/MODE-2304
15510	MODE-2304	[~hchiorean], I'm building now. Should I resolve this issue after merging?	1	1	https://issues.jboss.org/browse/MODE-2304
15511	MODE-2304	Could you please check this: I created a war file with an empty web.xml just as in your example and included the modeshape-web-explorer.jar As soon as I access Console.html I get a GWT popup containing this error: org.jboss.as.naming.NamingContext cannot be cast to org.modeshape.jcr.ModeShapeEngine Nothing in the server.log There is nothing else in the war file. Server is Wildfly 8	1	1	https://issues.jboss.org/browse/MODE-2304
15512	MODE-2304	I don't understand, why you write it will never work. It worked before, when I bundled the old 3.6 GWT app into the war. Without changing anything to the JCR access layer used in there. I looked at the mentioned war and that contains nothing special besides some security settings that I currently don't need and that should have nothing to do with the JCR layer. And it contains the jboss-deployment-structure.xml. I think this war is only suitable in the context of the whole kit. What I need it the https://github.com/ModeShape/modeshape/blob/master/web/modeshape-web-explorer-war that should work out of the box, correct? Thats the whole purpose of this war: ModeShape repository explorer that embeds ModeShape and can be deployed standalone And it does indeed start and the {code}shouldAccessInitialPage{code} test succeeds as well, but that test only tests if the app was deployed. After that the GWT app just presents a notice box which seems to be triggered by this server response: //EX[2,1,["org.modeshape.web.client.RemoteException/3886951464","jcr -- service jboss.naming.context.java.jcr"],0,7] I can't see anything being logged from the GWT app besides {code}Requested URI /modeshape-web-explorer/console/jcr{code}, so I have no idea why this Exception is triggered. Thats the main problem, nothing is beeing logged. In the old Version 3.6 I used there wasn't even an access to a repository until you provided JNDI name and workspace at the front page. That front page seems to be gone now and instead the error above is triggered.	1	1	https://issues.jboss.org/browse/MODE-2304
15513	MODE-2304	{quote} What I need it the https://github.com/ModeShape/modeshape/blob/master/web/modeshape-web-explorer-war that should work out of the box, correct? {quote} it will only work out-of-the-box if the container you're deploying into has something bound by default in JNDI under the {{jcr}} name. In other words: {code} InitialContext ic = new InitialContext(); engine = (ModeShapeEngine) ic.lookup("jcr"); {code} This works fine by default *only if you use the ModeShape Wildfly/EAP kit* which does the above binding behind the scenes when you startup the subsystem. EDIT: If you don't use the above (which IIUC is your case), then you need to make the above binding yourself (in your application). I agree that this is quite limiting ATM, so I'll open an issue to improve this for {{4.1}}.	1	1	https://issues.jboss.org/browse/MODE-2304
15514	MODE-2299	[~rhauch]: since we do have a {{2.8.1.GA}} version can you please create a {{2.8.2.GA}} version as well in JIRA ?	1	1	https://issues.jboss.org/browse/MODE-2299
15515	MODE-2277	When a node is checked out, all nodes below it are affected. Unfortunately, that means to determine if a node is checked out we have to walk all the way up the ancestors to the root to verify that none of them is checked out. IIRC, the logic starts with the node in question and walks up to the root, though it does stop when it finds a node that is checked out. I'm not really sure how we might mitigate this. Even if we had a cache of the node keys of the nodes that are currently checked out, we'd still have to walk up the ancestors to find the node keys of the ancestors. If that cache of checked out nodes contained the node keys and the paths (and we could maintain this in a cluster), then we'd have to know the path of the node. Again, that requires walking up the ancestors. -Perhaps one thing we could do is maintain state that records the number of checked out nodes. I would guess that for most repositories, locking is not used and the state would be 0. This would be super fast for a lot of cases (maybe even the majority), but as soon as you used locking it would be the same performance as what we have now.- Thoughts? Can you think of any alternatives?	1	1	https://issues.jboss.org/browse/MODE-2277
15516	MODE-2277	[~hchiorean], I merged the PR into the {{master}} branch, but it doesn't merge cleanly into {{3.x}}. Can you please create a new PR?	1	1	https://issues.jboss.org/browse/MODE-2277

#	Requirement	Questions	Answered	Successfully Answered	URL
15517	MODE-2265	Having investigate this a bit more, IMO there are only 3 viable options: 1) keep as is (nothing changes) 2) use MapDB and an on-heap structure. Even though off-heap may sound attractive, the reality IMO is that off-heap is very much dependent on the OS/JVM combination and this is still a bit of a "grey" area. Since the WS cache is a structure that's being used intensively, we can't really risk memory/Direct Buffer mishaps. Also, the WS cache should be "relatively" small & constant in size, since most performance tuning should come out of the main ISPN cache. 3) "reinvent the wheel" and come up with a LRU {{ConcurrentMap}} implementation. I don't think we should spend the effort of doing this since there are other libraries out there which already have this implementation (like google collections) but since we're already using {{MapDB}} IMO we should avoid bringing in yet another dependency. For example just the {{-collections}} part of Guava is 500k in size... So IMO it really comes down between (1) and (2). The question then is how do we compare the two ? One thing we can try is just looking at performance (i.e. how much time does it take in a given scenario to read/write entries into the cache). However I'm not really sure how relevant this is given that the churn is likely to very high and there are a lot of get/put operations that happen very atomically - for each node. Thoughts ?	1	1	https://issues.jboss.org/browse/MODE-2265
15518	MODE-2242	[~rhauch]: I understand changing this for {{master}} and {{3.8.x-prod}} but for {{3.x}} this is a completely non-backwards compatible change. Given that the {{anonymous-roles}} defaults change was not ported to {{3.x}}, do we really want this in {{3.x}} ? Regarding the webapps: if we're supposed to use the {{ApplicationRealm}} (which from what I understand is being used via the {{other}} security domain), I see no advantage in defining/changing a {{modeshape-security}} domain. The mechanism via which users will be able to register & map roles will still be the same. So IMO we should change the webapps security and map them to {{other}}. This change (overall) will impact *all* of our integration tests, since I'm not aware of any way of doing the "add-user" call from tests. Is there any easy way for us to adapt the tests & use (add) the Application Realm ? We shouldn't loose our ability to run the container-based tests IMO (there's always the possibility of overwriting the security domain, but ideally we should keep the config as close to the real one as possible)	1	1	https://issues.jboss.org/browse/MODE-2242
15519	MODE-2242	> Randall Hauch: I'm curious, what are the arguments for having a separate security domain that behaves like {{other}} ? This way, all ModeShape repositories and web applications use a ModeShape-specific security domain out-of-the-box. And out-of-the-box the users for ModeShape are managed with the {{add-user.sh}} script (and the "{{ApplicationRealm}}") just like user accounts for EAP. To change how everything is authenticated, simply change the "{{modeshape-security}}" security domain, and anything that uses the "{{other}}" default security domain, like other web apps, are unaffected. IOW, in this approach, you simply have to modify the existing domain. With the other approach, ModeShape repositories and web applications *and* other web applications use the built-in security domain "{{other}}", and to change ModeShape requires defining a new security domain, change each repository configuration (not a big deal) *and* alter the ModeShape web applications to use the new security domain (not straightforward). This is potentially a lot more work and more error prone.	1	0	https://issues.jboss.org/browse/MODE-2242
15520	MODE-2234	[~oleg.kulikov], what's the status on this?	1	1	https://issues.jboss.org/browse/MODE-2234
15521	MODE-2234	[~oleg.kulikov], MODE-2231 is scheduled for 4.1, whereas the others are scheduled for 4.0 and 3.x.	0	0	https://issues.jboss.org/browse/MODE-2234
15522	MODE-2233	[~oleg.kulikov], what's the status on this?	1	1	https://issues.jboss.org/browse/MODE-2233
15523	MODE-2233	[~oleg.kulikov], MODE-2231 is scheduled for 4.1, whereas the others are scheduled for 4.0 and 3.x.	0	0	https://issues.jboss.org/browse/MODE-2233
15524	MODE-2225	I'm looking at adding a test case for this issue. Is the preferred approach to add a test to the existing ModeShape test suit and issue a pull-request or to provide a standalone project that can be used to demonstrate the issue?	1	1	https://issues.jboss.org/browse/MODE-2225
15525	MODE-2225	I think I have found the issue. If you look at the AbstractNode.ChildNodeResolver constructor (line 3559) it sets the this.checkPermission flag to true only if checkPermission=true and AccessControl is enabled (see code below) {noformat} public ChildNodeResolver(JcrSession session, NodeKey parentKey, boolean checkPermission) { this.session = session; this.parentKey = parentKey; // we only need to check permissions if ACLs are enabled this.checkPermission = checkPermission && session.repository().repositoryCache().isAccessControlEnabled(); } {noformat} From what I can tell isAccessControlEnabled is only true if using ACLS. I did do a previous test with ACLS and child nodes where filtered correctly based on the configured ACLS. Would a appropriate fix be to also check permissions if the SecurityContext is an instance of AuthorizationProvider or AdvancedAuthorizationProvider?	1	1	https://issues.jboss.org/browse/MODE-2225
15526	MODE-2225	That change (see [here|https://github.com/ModeShape/modeshape/commit/22abf4509553a393b6eb046f4bcf397557f369b0#diff-76d801bf27b2288229aa7e262e2ddb6cR3562]) was made as part of MODE-2186. We do need to check the permissions (when called by a client), but I'm not sure why we decided to do this only when ACLs are enabled and not all of the time. So my suggestion would be to do this: {code} public ChildNodeResolver(JcrSession session, NodeKey parentKey, boolean checkPermission) { this.session = session; this.parentKey = parentKey; this.checkPermission = checkPermission; } {code} The session will actually perform the check using the appropriate AuthorizationProvider, AdvancedAuthorizationProvider, and ACL mechanism that are configured. Want to try this and run a {{mvn clean install -Pintegration}} to check for regressions? However, perhaps [~hchiorean] can respond with any context that he remembers. (MODE-2186 was about iteration not checking ACLs, so since permissions were not being checked maybe he wanted to add the check only when ACLs are used.)	1	1	https://issues.jboss.org/browse/MODE-2225
15527	MODE-2225	Those are all clustering related, and it's usually a symptom of the network configuration of the tests not matching your platform. For example, does your machine only use IPv6 or IPv4? If so, you might want to try one of these commands to see if they help: {code} mvn clean install -Pintegration -DpreferIpv4 {code} if your machine is configured to use (or only use) IPv4, or {code} mvn clean install -Pintegration -DpreferIpv6 {code} if your machine is configured to use (or only use) IPv6. The latter should be the default, since our build [prefers IPv6|https://github.com/ModeShape/modeshape/blob/master/modeshape-parent/pom.xml#L218], except when running the two integration module tests against Wildfly/EAP since Arquillian seems to prefer IPv4.	1	1	https://issues.jboss.org/browse/MODE-2225
15528	MODE-2225	[~ma6rl], you logged this against 4.0.0.Alpha3. Are you okay with us *not* backporting this to 3.x?	1	1	https://issues.jboss.org/browse/MODE-2225
15529	MODE-2215	[~rhauch]: why would we want to build our subsystem against Beta level API (EAP 6.3) ? I don't see any good reason to as long as the EAP 6.2.1.GA API is binary compatible with EAP 6.3.Beta. IMO we should only move to EAP 6.3 *when GA is available* - meaning Maven artifacts published in the techpreview repository.	1	1	https://issues.jboss.org/browse/MODE-2215

#	Requirement	Questions	Answered	Successfully Answered	URL
15530	MODE-2215	I thought the whole point of releasing the Beta was using the same components versions as EAP 6.3. Please note there are 2 different classes of components: 1) non EAP API related: Infinispan, JGroups, RestEASY etc. These are the ones which should be frozen by now in EAP 6.3 and they should have the same versions in the BOM. Updating to CR9 should align these versions for us. All of these components *must be available in a public Maven repo (via JBoss Nexus)*. 2) EAP API: components like {{jboss-as-controller}}, {{jboss-as-naming}} etc which our subsystem needs to integrate with. The versions of these is directly related to the EAP version & state. These are the ones I'm referring to as "beta-level" API. They are only available in "obscure" (and sometimes not public) Maven repositories which need to be tracked down. So if we were to change our subsytem API to depend on EAP 6.3.Beta API: a) we need to add this additional Maven repo to our settings (we just got rid of this custom Maven repo when we updated to EAP 6.2.1GA). If this is not a public repo, it means only people with VPN access can build the community version. b) we would *tie our 3.8.Final* community release with the EAP 6.3.GA release (afaik July is the best-case scenario): I doubt you'd want to release any earlier than EAP 6.3.GA because that means releasing a "beta-subsystem" Anybody integrating against ModeShape in EAP, *should not care what EAP API was used to build the subsystem*, as long as the subsystem works.	1	1	https://issues.jboss.org/browse/MODE-2215
15531	MODE-2215	I do agree that we need everything (including any EAP APIs) to be in a public repository, and this issue is predicated on that being possible. That's always an underlying assumption: all released community bits must depend on publicly-available bits. So if the APIs for EAP 6.3 Beta are not yet available in a public repo, then I agree that ModeShape 3.8.0.Beta1 be released to use EAP 6.2.1.GA. I also concur that if we upgrade to EAP 6.3 Beta for ModeShape 3.8.0.Beta1, then we are hitching ModeShape 3.8.0.Final to be based upon EAP 6.3 GA. IIRC, that was always the plan. But I'm open to release 3.8.0.Beta1 against EAP 6.2.1.GA (what we currently depend upon), if we're very confident that the resulting subsystem can be installed into both EAP 6.2.1.GA and EAP 6.3.Beta. The benefit is that this gives our community users a choice of which EAP version they want to use, while not limiting ourselves. Of course, there's another option here: make our next 3.8 release a FINAL rather than a BETA. The only reason I was suggesting a BETA was because I presumed we'd build against the EAP 6.3 Beta. If that's no longer true, is there really any point to calling this next release a BETA?	1	1	https://issues.jboss.org/browse/MODE-2215
15532	MODE-2215	{quote} Of course, there's another option here: make our next 3.8 release a FINAL rather than a BETA. The only reason I was suggesting a BETA was because I presumed we'd build against the EAP 6.3 Beta. If that's no longer true, is there really any point to calling this next release a BETA? {quote} I was under the impression that we would call it a Beta so that a) there is a release published in Nexus (we don't do snapshots) which can be used by other projects and b) it would give us a flexible window (time-wise) to get in more features & bug fixes until we release Final independent of EAP (e.g. MODE-2174) I was never under the impression that we're aligning our {{3.x}} with EAP 6.3. That being said, if ^^ is untrue, and we plan on making this alignment, then releasing Beta against EAP 6.3.Beta and the Final against EAP 6.3.GA sounds reasonable. In that case we should have another JIRA like you suggested above. So this is something you need to decide, I guess. From a compatibility perspective, I've deployed a subsystem built with EAP 6.2.1 on EAP 6.3.Alpha and EAP 6.3.Beta and it worked fine (meaning navigating repositories and nodes via the various webapps). This leads me to think that from an EAP API perspective, we are compatible (if we weren't, the repositories wouldn't start really).	1	1	https://issues.jboss.org/browse/MODE-2215
15533	MODE-2215	I originally _did_ anticipate that the 3.8 effort (e.g., the {{3.x}} branch) would be aligned with EAP 6.3, since that was primarily the codebase that will go into the JBoss Data Virtualization 6.1 product. This is why I tended to favor releasing 3.7.3.Final and 3.7.4.Final, rather than releasing an earlier version of 3.8. But now I see the advantage of releasing the 3.8.0.Final this week: it includes slight patch-level upgrades for Infinispan and JGroups useful to all users, and the EAP subsystem could be installed into EAP 6.2.x GA and 6.3 Beta (and we think GA). There's no reason to call this a "beta" release, since there really are just a handful of very small, low-impact, low-risk features and quite a few bug fixes. In fact, it's nearly the same and just as stable as 3.7.4.Final. We can then continue fixing critical issues on the {{3.x}} branch and cutting patch-level releases (e.g., "3.8.x.Final") as needed. We also will not issue any more 3.7.x patch releases. I'll update the issue to no longer upgrade EAP to 6.3 Beta, but to instead just upgrade the IP BOM. If that's the case, what is involved in that other than just changing the IP BOM version in our parent POM?	1	1	https://issues.jboss.org/browse/MODE-2215
15534	MODE-2215	@rhauch: why would this issue have the Fix Version {{4.0.0.Alpha4}} ?	1	1	https://issues.jboss.org/browse/MODE-2215

#	Requirement	Questions	Answered	Successfully Answered	URL
15535	MODE-2211	[~hchiorean], I should have been more clear about what would be cached in the {{CachedNode}}: only whether *that node* has an ACL child. It does *not* reflect anything about what those ACLs represent nor whether its descendants have any ACLs. So, when you said: {quote} ... regarding caching, consider the following scenario: there are 2 nodes A with its parent B. Node A does not have any ACLs while node B does. Initially, everyone has READ permissions for node B. Now, if there are 2 sessions: session 1 loads only Node A and keeps it for a while {quote} Okay, at this point the {{WorkspaceCache}} contains a {{LazyCachedNode}} for Node A, and when ACLs are checked this cached node might cache the fact that *it* does not have ACLs. That simply means that when session 2 attempts to read Node B and determine if it can from the ACLs, the session would: # load the {{LazyCachedNode}} for Node B into the {{WorkspaceCache}} # the {{AccessControlManagerImpl}} logic would attempt to determine if session 1 can read/write Node B: ## it looks at the {{SessionNode}} and then the {{LazyCachedNode}} for Node B to see if it has ACLs, and the {{LazyCachedNode}} would not already have the boolean flag and would then look for the ACL child. It would find it, so it would return "true" signifying that it does have ACLs. The {{AccessControlManagerImpl}} would then ask the {{CachedNode}} for the ACL child and would discover the permissions. ## Even if ACL was not on Node B, the {{AccessControlManagerImpl}} would walk up the tree and find the {{CachedNode}} for Node A, and the {{CachedNode}} would already have cached that it does not have ACLs. Thus the {{AccessControlManagerImpl}} logic would not have to try to read the (nonexistent) ACL child, and could simply walk up to the next parent and continue. {quote} while session 2 loads only Node B and at some point (while session 1 is active) removes all the permissions from B and saves. {quote} As soon as session 2 (which removed the permissions from Node B) is saved, the {{LazyCachedNode}} for node B in the {{WorkspaceCache}} will be purged via changes. So the second kind of caching (whether or not a {{CachedNode}} has ACLs) would always be in-sync with the persisted state. This still seems to work. {quote} If session 1 would've cached the fact that node A & node B are readable it wouldn't see the ACL changes made by B, making this a security hole IMO. The reason why this is important is that all the security checks are performed "up-top", in the JCR layers, not in the caching layer during {{WritableSessionCache#persistChanges}}. {quote} As I mention above, the second kind of caching would be fine, as it simply saves a little bit of time when trying to determine ACL privileges by caching which nodes do *not* have an ACL. As for the third improvement and keeping information in the session, I do agree we have to be careful here. For example, we might be able to at least know if a node is writable because it is found in the {{SessionCache}}. Yes, this conclusion might not be valid in some cases: this session modifies a node for which it does have privilege, another session modifies the ACLs and persists them so that this session no longer has permission, and this session tries to make more changes. However, as it stands now the first session would still be allowed to persist at least those changes it made before the other session changed the ACLs -- is that valid? Is it any worse to allow a session to modify any node any time *after* it has determined ACLs allow it to do so, even if those ACLs are changed at some later point but before the session saves its changes?	0	0	https://issues.jboss.org/browse/MODE-2211
15536	MODE-2208	[~rhauch]: after reading https://issues.jboss.org/browse/TEIID-2948, it appears the RBAC feature is around securing the EAP subsystem extension elements (resources/attributes). Since the ModeShape EAP subsystem is big (in terms of configuration options), what are the exact requirements around the various elements ? Also, I'm not entirely convinced this makes sense: the core ModeShape subsystem requires a properly configured Infinispan (and in some cases JGroups) subsystem. However, we have no control over the RBACC compliance of these 2 subsystems. Therefore, what's the purpose of securing a {{repository}} for example, when we have no control of the security of its cache ?	1	1	https://issues.jboss.org/browse/MODE-2208
15537	MODE-2168	Are there any changes in ModeShape that are required in order to support EAP 6.3 RBAC features? If not, I would suggest reading TEIID-2948 and the discussion that Ramesh had with Brian Stansbury.	1	1	https://issues.jboss.org/browse/MODE-2168
15538	MODE-2168	Probably, but that'd be better handled as a separate issue. [~vhalbert], do you want to log an issue for us?	0	0	https://issues.jboss.org/browse/MODE-2168
15539	MODE-2168	With MODE-2208 completed, can this JIRA be considered completed?	1	1	https://issues.jboss.org/browse/MODE-2168
15540	MODE-2162	[~hchiorean] Can you please guide me a bit how we can isolate Lucene classes with separated modules? Just hit into the problem with versions. The issues is that Elasticsearch uses older versions. The most fresh is 5.2.1	1	1	https://issues.jboss.org/browse/MODE-2162
15541	MODE-2162	[~oleg.kulikov]: I'm curious though, is this a runtime-thing only ? i.e. isn't ElasticSearch supposed to have its own API so that you don't interact directly with the Lucene API ? More precisely, shouldn't there be just a REST over HTTP API ? The other thing is I always thought Elasticsearch is really a service/remote server only. So there should be any Lucene up-front interaction required.	1	1	https://issues.jboss.org/browse/MODE-2162
15542	MODE-2162	[~hchiorean], [~rhauch], just hit to the potential dangerous place https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/query/engine/process/RestartableSequence.java#L187 In case of inaccurate {{ResultSearch}} implementation loop might be infinite. Would it be useful to limit it to some {{MAX}} value?	1	1	https://issues.jboss.org/browse/MODE-2162
15543	MODE-2159	[~hchiorean]I am still a bit confused with multiple column indexes. {{ProvidedIndex}} interface offers methods like {{add(NodeKey, Object)}} which are, in my opinion, still oriented for single column index. Would it be better to use map between string column name and value(s) as parameter? Like this {{void add(Map<String, Object> doc)}}	1	1	https://issues.jboss.org/browse/MODE-2159
15544	MODE-2159	[~hchiorean], additional parameter {{propertyName}} will be good me for as well. Should I wait for commit of your changes or just add same thing locally and later merge?	1	1	https://issues.jboss.org/browse/MODE-2159
15545	MODE-2135	The maintainer of MapDB is moving towards 2.0, so it's not clear when append-only support will be added -- unless it already has been in 1.0.6. [~hchiorean], can you check? If it's not there, we may want to either push this into a future release (e.g., 4.1) or mark it as WONTFIX. Thoughts?	1	0	https://issues.jboss.org/browse/MODE-2135

#	Requirement	Questions	Answered	Successfully Answered	URL
15546	MODE-2111	See also IRC transcript below from 5 December 2013: {quote} *barmintor*: rhauch: is it reasonable to write that ticket such that event filtering criteria are ordered in increasing cost (eventType, session locality, parent uuid or location, access privilege, nodeType as I understand it) *barmintor*: it's a shifting cost: If there *are* nodeType filters among others, then that is faster; if not, it might incur other costs *rhauch*: barmintor: thinking *barmintor*: rhauch++ *rhauch*: barmintor: I still think that all of the other filters operate against the Event without needing to load the node, so I think the node type filters should always be at the end since they are most expensive. *barmintor*: rhauch: I'm inclined to agree *barmintor*: but wanted someone else to think about it, too *rhauch*: barmintor: so the change would be to the order of the calls within this method: https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/JcrObservationManager.java#L991 Do you agree? *barmintor*: rhauch: the order there, but checking eventType first in https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/JcrObservationManager.java#L838 *barmintor*: since eventType checking is a bitwise op on a value in the event already *rhauch*: barmintor: so should the acceptBasedOnNodeTypeName(...) call just be moved out of https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/JcrObservationManager.java#L991 and into the appropriate event type(s)? *rhauch*: barmintor: actually, I'm not sure I understand why that filter is a function of the event type *barmintor*: rhauch: I meant that eventListenedFor is the cheapest filter, and it's called after all of the other filters *rhauch*: barmintor: ah, so that just moves to the front and the acceptBasedOnNodeTypeName(...) moves to the end? *barmintor*: rhauch: that's what I was thinking, yeah *rhauch*: barmintor: the acceptBasedOnUuid(...) probably would be pretty fast if we converted the String[] to a Set<String> in the constructor. Might be faster than the path-based one. {quote}	1	1	https://issues.jboss.org/browse/MODE-2111
15547	MODE-2111	[~barmintor], has there been any progress on a pull-request with the fix? Note that we want to release 3.7 early next week, so we'd like to get this addressed.	1	1	https://issues.jboss.org/browse/MODE-2111
15548	MODE-2111	branch in the msoffice-sequencer tests. I develop on a Mac, so... that's suspicious. See below: Tests run: 4, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 20.974 sec <<< FAILURE! shouldSequencePowerpointFiles(org.modeshape.sequencer.msoffice.MSOfficeMetadataSequencerTest) Time elapsed: 17.467 sec <<< FAILURE! java.lang.AssertionError at org.junit.Assert.fail(Assert.java:92) at org.junit.Assert.assertTrue(Assert.java:43) at org.junit.Assert.assertNotNull(Assert.java:526) at org.junit.Assert.assertNotNull(Assert.java:537) at org.modeshape.sequencer.msoffice.MSOfficeMetadataSequencerTest.shouldSequencePowerpointFiles(MSOfficeMetadataSequencerTest.java:167) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.rules.RunRules.evaluate(RunRules.java:18) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:47) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:50) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:222) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:234) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:133) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:114) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:188) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:166) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:86) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:101) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:74)	1	1	https://issues.jboss.org/browse/MODE-2111
15549	MODE-2109	A couple of questions: 1) why shouldn't there be any properties on the node types ? IMO something like the {{nt:unstructured}} node definition would also work. 2) {quote} Note that when using large collections, a repository's eviction policy should be sized larger than the total number of range buckets, Nrb, for all collections to be accessed quickly and efficiently. {quote} While this statement is correct, IMO it's largely academic: there will be a lot of other stuff in the cache, so the eviction policy size may not have any direct influence on the range buckets. From the cache's perspective, the range buckets are nothing more than regular entries. In other words, even with a larger eviction policy, all the buckets may be evicted if not used often compared to other nodes. 3) Adding & Removing a child should also mention that explicit locking needs to happen on the bucket, once it has been located. 4) Federation: {{mode:unorderedCollection}} *should not* be allowed to contain external children. This would create a "deadly mix" between segments and range buckets. 5) Removing the parent node: according to the JCR spec, when removing a node all its children should be removed as well. In this case, with millions of nodes, this is likely to take forever, considering that we need to load each child node and mark it as REMOVED for the {{WritableSessionCache}}. Actually, loading all those into memory might also produce OOM exceptions, so not sure what to do here. 6) In the previous case, locking should also be discussed: consider 2 concurrent sessions, one which removes the parent and one which adds a child. Without any kind of synchronization between the sessions (via cache locking) how can we guarantee a consistent result ?	1	1	https://issues.jboss.org/browse/MODE-2109

#	Requirement	Questions	Answered	Successfully Answered	URL
15550	MODE-2109	add residual property definitions, but that would force all nodes that use it to allow any properties. Someone couldn't extend {{mode:unorderedLargeCollection}}, for example, and limit the properties to some set they define; instead, there would always be the residual that would allow any property to be set on the node. By leaving off all property definitions and basically requiring someone to create a subtype with the kinds of properties they wanted, the developer is in full control of the properties. I'm not convinced the current approach is better. Yes, it gives more control to a repository designer, but it isn't terribly useful without creating a subtype. {quote} 2) "Note that when using large collections, a repository's eviction policy should be sized larger than the total number of range buckets, Nrb, for all collections to be accessed quickly and efficiently." While this statement is correct, IMO it's largely academic: there will be a lot of other stuff in the cache, so the eviction policy size may not have any direct influence on the range buckets. From the cache's perspective, the range buckets are nothing more than regular entries. In other words, even with a larger eviction policy, all the buckets may be evicted if not used often compared to other nodes. {quote} Yeah, my statement was pretty squishy. I was simply trying to make the point that the range buckets are accessed pretty randomly whenever children are accessed, and that for ideal child-access performance you probably want all of the range buckets to stay in the workspace cache (while leaving room for all the other nodes). Whether that can even be achieved is pretty iffy. It's also highly dependent upon how the clients access/read the children: fairly randomly vs frequently re-reads of the same recently-read children. On one hand, you want each range bucket to be small enough so that loading it from the cache store is faster. On the other, the larger they are, they larger percentage of the collection each range bucket owns and thus the more likely a given bucket will already be in the cache (under heavy read load). I think picking the number of buckets is going to be tricky. It would be nice if it were adaptable, but that would add a lot of complexity. {quote} 3) Adding & Removing a child should also mention that explicit locking needs to happen on the bucket, once it has been located. {quote} Agreed. {quote} 4) Federation: mode:unorderedCollection should not be allowed to contain external children. This would create a "deadly mix" between segments and range buckets. {quote} Hmm... I guess your right. I do think that it's perfectly acceptable ... a collection should be stored completely _within_ the repository. Besides, if in the future it is deemed absolutely critical, then we could consider removing this restriction. {quote} 5) Removing the parent node: according to the JCR spec, when removing a node all its children should be removed as well. In this case, with millions of nodes, this is likely to take forever, considering that we need to load each child node and mark it as REMOVED for the WritableSessionCache. Actually, loading all those into memory might also produce OOM exceptions, so not sure what to do here. {quote} Good catch. But perhaps it is no different than removing a large tree in one fell swoop - that, too, would probably fail with an OOM. So is that a more general problem we should consider? And if so do we need a separate JIRA? Do all the nodes need to be removed within the one transaction? For example, is it sufficient that the node being removed is removed in the transaction, and once that has completed then all descendants are _eventually_ (perhaps asynchronously) removed? We can't suspend the actual transaction, as then we'd have no way to roll back any internal transaction(s) should the outer one be rolled back. A general workaround is to just explicitly delete smaller portions of the subtree or collection, and then repeat this as many times as needed to delete the remaining portions. Yes, it would take forever, but it would work. {quote} 6) In the previous case, locking should also be discussed: consider 2 concurrent sessions, one which removes the parent and one which adds a child. Without any kind of synchronization between the sessions (via cache locking) how can we guarantee a consistent result ? {quote} I suspect you're referring to ISPN locking and not JCR locking. What if the code removing the parent collection attempts to lock all (actual and potential) range bucket entries in	1	1	https://issues.jboss.org/browse/MODE-2109
15551	MODE-2109	I see this ticket has been pushed further out into the future (to 4.3). Is that due to the complexity of the solution? the ModeShape priority roadmap? or lack of community interest? This performance issue is a significant pain-point for the [Fedora Repository|https://wiki.duraspace.org/display/FF/Fedora+Repository+Home] application.	0	0	https://issues.jboss.org/browse/MODE-2109
15552	MODE-2109	When searching through the "open" tickets in the ModeShape JIRA - Query: project = "MODE" AND status = Open ORDER BY votes DESC I see that this ticket is the most wanted feature from the community. If there is a consultant company (RedHat has been unresponsive) that you have worked with before, we would be interested in funding the implementation of this ticket.	0	0	https://issues.jboss.org/browse/MODE-2109
15553	MODE-2109	I have tested it with CMIS connector. I couldn't create document in folder(mode:unorderedSmallCollection). Does this feature supported for external connectors?	1	0	https://issues.jboss.org/browse/MODE-2109
15554	MODE-2093	[~oleg.kulikov], those screenshots look fantastic! Is there a way of supporting multiple repositories and their workspaces, perhaps somewhat akin to how GitHub shows a user/organization with multiple Git repositories?	1	1	https://issues.jboss.org/browse/MODE-2093
15555	MODE-2093	[~oleg.kulikov], what is the status of this?	1	1	https://issues.jboss.org/browse/MODE-2093

#	Requirement	Questions	Answered	Successfully Answered	URL
15556	MODE-2093	8.0.0.Final/standalone/configuration/standalone-modeshape.xml}}" to add the following line to the ModeShape subsystem configuration: {code} <webapp name="modeshape-explorer.war"/> {code} and then running with {{bin/standalone.sh -c standalone-modeshape.xml}}. The error (see below) prevents deployment. {code} 11:36:20,216 WARN [org.jboss.as.ee] (MSC service thread 1-16) JBAS011006: Not installing optional component org.modeshape.web.RestoreImportServlet due to an exception (enable DEBUG log level to see the cause) 11:36:20,239 ERROR [org.jboss.msc.service.fail] (MSC service thread 1-14) MSC000001: Failed to start service jboss.undertow.deployment.default-server.default-host./modeshape-explorer.UndertowDeploymentInfoService: org.jboss.msc.service.StartException in service jboss.undertow.deployment.default-server.default-host./modeshape-explorer.UndertowDeploymentInfoService: java.lang.ClassNotFoundException: org.modeshape.web.RestoreImportServlet from [Module "deployment.modeshape-explorer.war:main" from Service Module Loader] at org.wildfly.extension.undertow.deployment.UndertowDeploymentInfoService.createServletConfig(UndertowDeploymentInfoService.java:870) at org.wildfly.extension.undertow.deployment.UndertowDeploymentInfoService.start(UndertowDeploymentInfoService.java:242) at org.jboss.msc.service.ServiceControllerImpl$StartTask.startService(ServiceControllerImpl.java:1948) [jboss-msc-1.2.0.Final.jar:1.2.0.Final] at org.jboss.msc.service.ServiceControllerImpl$StartTask.run(ServiceControllerImpl.java:1881) [jboss-msc-1.2.0.Final.jar:1.2.0.Final] at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) [rt.jar:1.7.0_45] at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) [rt.jar:1.7.0_45] at java.lang.Thread.run(Thread.java:744) [rt.jar:1.7.0_45] Caused by: java.lang.ClassNotFoundException: org.modeshape.web.RestoreImportServlet from [Module "deployment.modeshape-explorer.war:main" from Service Module Loader] at org.jboss.modules.ModuleClassLoader.findClass(ModuleClassLoader.java:197) [jboss-modules.jar:1.3.0.Final] at org.jboss.modules.ConcurrentClassLoader.performLoadClassUnchecked(ConcurrentClassLoader.java:443) [jboss-modules.jar:1.3.0.Final] at org.jboss.modules.ConcurrentClassLoader.performLoadClassChecked(ConcurrentClassLoader.java:431) [jboss-modules.jar:1.3.0.Final] at org.jboss.modules.ConcurrentClassLoader.performLoadClass(ConcurrentClassLoader.java:373) [jboss-modules.jar:1.3.0.Final] at org.jboss.modules.ConcurrentClassLoader.loadClass(ConcurrentClassLoader.java:118) [jboss-modules.jar:1.3.0.Final] at org.wildfly.extension.undertow.deployment.UndertowDeploymentInfoService.createServletConfig(UndertowDeploymentInfoService.java:602) ... 6 more 11:36:22,244 INFO [javax.enterprise.resource.webservices.jaxws.servlet.http] (MSC service thread 1-1) WSSERVLET14: JAX-WS servlet initializing 11:36:22,299 INFO [org.wildfly.extension.undertow] (MSC service thread 1-1) JBAS017534: Registered web context: /modeshape-cmis 11:36:22,305 ERROR [org.jboss.as.controller.management-operation] (Controller Boot Thread) JBAS014613: Operation ("add") failed - address: ([("deployment" => "modeshape-explorer.war")]) - failure description: {"JBAS014671: Failed services" => {"jboss.undertow.deployment.default-server.default-host./modeshape-explorer.UndertowDeploymentInfoService" => "org.jboss.msc.service.StartException in service jboss.undertow.deployment.default-server.default-host./modeshape-explorer.UndertowDeploymentInfoService: java.lang.ClassNotFoundException: org.modeshape.web.RestoreImportServlet from [Module \"deployment.modeshape-explorer.war:main\" from Service Module Loader] Caused by: java.lang.ClassNotFoundException: org.modeshape.web.RestoreImportServlet from [Module	1	1	https://issues.jboss.org/browse/MODE-2093
15557	MODE-2093	It does work well, and I really like the new design. Very nice work, [~oleg.kulikov]. I do have a few comments, and I'd be interested in honest feedback from [~oleg.kulikov] on them: # The "path" link that shows the breadcrumb location of what is currently being deployed is a single link to the current view. That's great, but do you think it would make sense to break this path into separate links for each segment? So for "/jcr:system/mode:namespaces/nt", the "jcr:system" segment would link to "/jcr:system", the "mode:namespaces" segment would link to "/jcr:system/mode:namespaces", and "nt" would link to "/jcr:system/mode:namespaces/nt". This would make it much easier to jump up to an ancestor, yet there still is a link to the current page. # What would you think about moving the properties table below the child table? I'm not sure if that'd be better or worse, but it might save a few clicks. # How is a binary property shown? Can you link to see or download the content? How it's displayed might depend on the MIME type - some MIME types might be just displayed (maybe even with line numbers! :-D) # Can the links be displayed as standard blue (when not clicked, but if darker blue if previously-clicked) but without underline? They're still visible, but IMO the underling is pretty distracting nowadays. # Make the "Backup", "Restore", and "Download" more subtle, or even move to a page for administration. Also, can these be displayed only if user has admin role? Perhaps it'd be a bit easier to move to a separate page for administration, where we could also provide a bit more functionality in the future (like showing some metrics, etc.). Fixing to allow only administrators to see/use these buttons seems pretty important. # "Save" button should maybe move closer to where the content and edit buttons are. Not exactly sure when. It'd be fantastic if it enabled only when there's something to save. # Click on "Node types" and get a popup with "Note: 123" and occasionally get a "Note: 300". Otherwise seems to work. # Query language should default to "JCR-SQL2" # Query results show only one line for me, though I can scroll up and down via those buttons on the list. Not sure if we should handle these as part of this issue, or whether we should have separate issues for some/all of these. I could go either way.	1	1	https://issues.jboss.org/browse/MODE-2093
15558	MODE-2091	[~oleg.kulikov], have you made any progress on this? I'm not convinced this is an easy fix, but great if you can get some easy improvements by Monday Dec 16, 2013. (If not, please re-target to 4.0.)	1	1	https://issues.jboss.org/browse/MODE-2091
15559	MODE-2088	[~hchiorean], do you want to try to start this? The index system might very well benefit from this. As it currently stands, the global index writer for the repository would like process each of the incoming ChangeSet objects, and serially delegate to each of the index providers, each of which would serially delete to each of its indexes. I'm in the process of changing the IndexWriter interface to extend ChangeSetListener, so the bulk of all index updates happen by processing events. The exception is reindexing, which does not rely upon events whatsoever. We certainly could change the design so that each index provider, or rather each index, might register its own ChangeSetListener with the bus. That would mean that even in the short term the indexing system is relatively simple and (potentially) decoupled. With the current bus, each index would indeed be independent and decoupled (we use a separate worker thread for each listener). But each ChangeSetDispatcher (which the thread runs) still has a blocking queue. At the moment I'm worried about adding lots more listeners. (I think we'll be fine for a dozen or so indexes, but I suspect we'd have trouble with more than that. And if we decide to give each index in each workspace its own listener, then we have a permutation problem.) If we think it's worth it, then I'll start moving the index system over to using separate listeners for each index (and perhaps each index/workspace combination). Then if we implement a ring buffer, the index system shouldn't need to be changed.	1	1	https://issues.jboss.org/browse/MODE-2088
15560	MODE-2088	[~rhauch]: perhaps, but I'm not yet entirely sure how this would work/look like. I did look at the [disruptor|https://github.com/LMAX-Exchange/disruptor], but IMO it's not something that we can use as-is, so we'd have to implement something from scratch. At least a couple of issues the implementation would have to take into account: * support 1 writer/multiple readers from different threads * should there be a max size to the buffer ? If yes, what happens when events are received by the bus and exceed the current {{max}} ? If no, what's the difference between the ring buffer and an unbounded queue from which each listener takes (but doesn't remove) items ?	1	1	https://issues.jboss.org/browse/MODE-2088

#	Requirement	Questions	Answered	Successfully Answered	URL
15561	MODE-2088	Can you share the code for your favorite implementation? I'm curious to see how you're doing it. {quote} * preventing the main buffer from being overrun - whenever a new ChangeSet comes through and should be added to the buffer, it cannot be added until "there is a spot for it", meaning that there should be a gap of at least 1 between the current head of the buffer and the slowest reader. This means that the buffer has to always be on the lookout for the slowest reader and keep waiting (or sleep & retry) until a slot opens up. The slowest reader in my prototypes is determined by the global "min cursor" value, but that obviously sucks when there are lots of listeners. Sleep and retry is even worse for throughput. {quote} There should always be room for the buffer to add something, or else the ring buffer is not sized nearly large enough. For example, an {{Object[1000]}} is only 4K of memory, so it's probably worth it to have a sufficiently large buffer. Let's see how the first ChangeSet is added, which we'll assume is added at index 0. Until it is added, the listener threads do nothing (yes, we have to address that), but as soon as it's added the threads all process it as quickly as they can. When another is added, they all process it as quickly as they can. This continues, so after having added 100 ChangeSet, the majority of the listeners will already have processed nearly all of them (e.g., 998 of them), and a few might be a little slower at the moment (e.g., 96), though it's likely they may catch up. (Consider that 96 might have been a big ChangeSet, whereas 97-98 were both relatively small.) When a new listener is added, it has no need to process any of the previously-added ChangeSets, so in our example it would be ready to process the one at index 100, since 0-99 were all fired before the listener was added. {quote} * each reader thread has to "do something" if there is no new data in the buffer. This means that it can either keep checking its own cursor against the head of the buffer in a while loop (which effectively kills the CPU) or employ some sort of suspend/resume mechanism, via which the buffer tells it that there is "new data". This suspend/resume mechanism is already part of the JDK built-in queues and mimicking it with low level code seems to always perform worse. Overall, I'm inclined to say that a polling approach - the reader threads constantly polling the buffer waiting for data - offers the worst performance. Therefore, the other approach is to be able to notify the reader threads in some way but even so, the performance degradation is significant (even for a simple test of 1000 events and 10 listeners) compared to the current queued approach. {quote} I agree that this is perhaps the most challenging. There are multiple ways to do this (even with suspending the thread via wait), but we'd have to do a bit of research to figure out the fastest and most efficient way to do it. {quote} * removing consumed items from the buffer - the way I understand it, the buffer has to constantly look at the "free gap" between its latest head and the "slowest reader" (its cursor) and remove (or in my implementation nullify) all the elements in this gap. {quote} Sure, but isn't this "mind the gap" problem just comparing indexes? In theory this should never happen because nearly all listeners should be pretty close to the front of the ring buffer, not near the end. As I said above, if they're near the end, they are slowing down relative to the incoming ChangeSets and will (perhaps soon) be at the very end and will indeed block additions (assuming a fixed-size ring buffer). Something like this is a problem with this approach or with our current approach (in which case the listener's blocked queue fills up and ChangeSets are dropped. IMO, this condition means that either the ring buffer is too small (not leaving enough room for natural variation in the processing time of the ChangeSets) or the listener implementation is too slow (in which case it needs to be changed).	1	1	https://issues.jboss.org/browse/MODE-2088
15562	MODE-2088	Have you seen https://github.com/LMAX-Exchange/disruptor/wiki/Introduction ? Maybe it might be worth prototyping with the Disruptor to see what happens with their advanced implementation. If that doesn't even perform, then we can drop the whole thing immediately.	1	1	https://issues.jboss.org/browse/MODE-2088
15563	MODE-2086	with information/data already stored in the cache. The reason is that since https://issues.jboss.org/browse/MODE-1524, the "main" {{SchematicExternalizer}} extends {{org.infinispan.commons.marshall.Externalizer}} *and not* {{AdvancedExternalizer}}. This, in turn, means that any externalizer IDs are ignored completely by Infinispan. Instead, when serializing objects in the cache, Infinispan via JBoss Marshalling will write out *the fully qualified name (string)* of the externalizer class and will attempt to instantiate & use that upon deserialization. See https://github.com/jboss-remoting/jboss-marshalling/blob/master/river/src/main/java/org/jboss/marshalling/river/RiverMarshaller.java#L1415 and https://github.com/jboss-remoting/jboss-marshalling/blob/master/river/src/main/java/org/jboss/marshalling/river/RiverUnmarshaller.java#L1072 for the implementation details behind marshalling & unmarshalling. When attempting to change the package names, the binary incompatibility manifests itself with exceptions like: {code} org.infinispan.persistence.spi.PersistenceException: java.lang.ClassNotFoundException: org.infinispan.schematic.internal.SchematicEntryLiteral at org.infinispan.marshall.core.MarshalledEntryImpl.unmarshall(MarshalledEntryImpl.java:116) at org.infinispan.marshall.core.MarshalledEntryImpl.getValue(MarshalledEntryImpl.java:61) at org.infinispan.interceptors.CacheLoaderInterceptor.loadIfNeeded(CacheLoaderInterceptor.java:270) at org.infinispan.interceptors.CacheLoaderInterceptor.loadIfNeededAndUpdateStats(CacheLoaderInterceptor.java:335) at org.infinispan.interceptors.CacheLoaderInterceptor.visitGetKeyValueCommand(CacheLoaderInterceptor.java:105) at org.infinispan.commands.read.GetKeyValueCommand.acceptVisitor(GetKeyValueCommand.java:40) at org.infinispan.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:98) at org.infinispan.interceptors.EntryWrappingInterceptor.visitGetKeyValueCommand(EntryWrappingInterceptor.java:116) at org.infinispan.commands.read.GetKeyValueCommand.acceptVisitor(GetKeyValueCommand.java:40) at org.infinispan.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:98) at org.infinispan.interceptors.base.CommandInterceptor.handleDefault(CommandInterceptor.java:112) at org.infinispan.commands.AbstractVisitor.visitGetKeyValueCommand(AbstractVisitor.java:74) at org.infinispan.interceptors.locking.AbstractTxLockingInterceptor.visitGetKeyValueCommand(AbstractTxLockingInterceptor.java:60) at org.infinispan.interceptors.locking.OptimisticLockingInterceptor.visitGetKeyValueCommand(OptimisticLockingInterceptor.java:108) at org.infinispan.commands.read.GetKeyValueCommand.acceptVisitor(GetKeyValueCommand.java:40) at org.infinispan.interceptors.base.CommandInterceptor.invokeNextInterceptor(CommandInterceptor.java:98) at org.infinispan.interceptors.base.CommandInterceptor.handleDefault(CommandInterceptor.java:112) at org.infinispan.commands.AbstractVisitor.visitGetKeyValueCommand(AbstractVisitor.java:74) at org.infinispan.commands.read.GetKeyValueCommand.acceptVisitor(GetKeyValueCommand.java:40) at	0	0	https://issues.jboss.org/browse/MODE-2086
15564	MODE-2083	What do you consider a small cluster and a large cluster?	1	1	https://issues.jboss.org/browse/MODE-2083
15565	MODE-2077	Let's see what impact (1) has on the performance and go from there. If possible, can we do this while keeping (at least temporarily) the option of using the current JGroups-based lock implementation? This will make it easier to compare behaviors.	1	1	https://issues.jboss.org/browse/MODE-2077
15566	MODE-2066	Did you fork the https://github.com/ModeShape/modeshape repository? If you did, you can push your topic branch up there (don't make changes on master). See [these instructions|https://community.jboss.org/wiki/ModeShapeDevelopmentWorkflow].	0	0	https://issues.jboss.org/browse/MODE-2066

#	Requirement	Questions	Answered	Successfully Answered	URL
15567	MODE-2065	Should we support *any* version of Wildfly or is there a particular version we're targeting ? (i.e. I remember significant differences between AS 7.1 and AS 7.2).	1	1	https://issues.jboss.org/browse/MODE-2065
15568	MODE-2065	WildFly 8.0.0.CR1 was released Dec 21 and uses Infinispan 6.0.1.Final: http://wildfly.org/news/2013/12/21/WildFly8-CR1-Released/ https://community.jboss.org/thread/233347 stack trace showing such early interference with the WildFly boot suggests that the existing EAP6.x kit really isn't going to help those of us who want WildFly+ModeShape. WildFly is passing EE7 TCK 100%, so the next few weeks would the optimal window to test it with an early cut of a ModeShape4 with upgraded Infinispan, and get feedback to them before they go final. What do you think?	1	1	https://issues.jboss.org/browse/MODE-2065
15569	MODE-2065	@rhauch as mentioned in my previous comment, I still need to update the quickstarts. I was thinking we can use this issue & attach a 2nd PR. WDYT ?	1	1	https://issues.jboss.org/browse/MODE-2065
15570	MODE-2065	When playing around with clustering in ModeShape 4.0, we get the error mentioned in MODE-2188 when shutting down Wildfly (8.0.0.Final) {noformat} 17:04:05,973 ERROR [org.jgroups.jmx.JmxConfigurator] (ServerService Thread Pool -- 28) unregister MBean failed jgroups:type=protocol,cluster="modeshape-repository-clustered",protocol=FORK {noformat} NOTICE THE "protocol=FORK" message, which makes us think it's related to our programmatic creation of the JGroups FORK. Once this error occurs we are no longer able to get all of the Wildfly instances in my cluster to reconnect. The first instance connects but the second instance gets stuck attempting to register an MBean and remains that way until it times out. The full details of this issue are discussed on the ModeShape forum https://community.jboss.org/message/871880#871880	0	0	https://issues.jboss.org/browse/MODE-2065
15571	MODE-2061	In the testing I did it was actually the smaller files that showed increased performance using openssl not the larger. For the 64GB size file streaming vs openssl were basically the same. Did you find differently? My findings: stats for retrieving jcr:content node of 1.3GB and 64GB file: default modeshape secureHash - 12 sec/16.6 minutes cached openssl - 3.5 sec/8.1 minutes cached modeshape secureHash - 6.2 sec/8.1 minutes mocked - .002 sec But even with the openssl alternative, whether faster or slower, this issue - that touching jcr:content nodes recalculates the hash every time (that takes on the order of minutes to an hour for large file) still remains. Ideally, being naive to modeshape, it seems hashes of a full file should NOT be the BinaryKey identifier but I don't know the fallout of changing this. In the Fedora Commons Akubra module for example https://wiki.duraspace.org/display/FEDORA34/Configuring+Low+Level+Storage, the key is an MD5 of an ID rather than the full file, an approach I was trying to implement similarly with the "mocked" implementation of creating the SHA-1 hash from the file path rather than full file, but Randall expressed reservations about the fallout from this when files move around. One idea I just realized is perhaps for large files (above a threshold) to allow the option of a "disinterested hash", which rather than a full file hash (which takes to long), or a path hash (that may change), to calculate the hash on the first arbitrary KBs of a file, or the first and last KBs of a file (is that possible?) with the assumption that the probability of this being the same in more than one file is little to none. I'd be interested in your opinion about this approach.	1	1	https://issues.jboss.org/browse/MODE-2061
15572	MODE-2061	{quote} In the testing I did it was actually the smaller files that showed increased performance using openssl not the larger. For the 64GB size file streaming vs openssl were basically the same. Did you find differently? {quote} I didn't do a lot of in-depth testing, but I noticed during some of our unit tests that computed the SHA-1 for lots of very small to small files (e.g., mostly .class files in the '{{target}}' directory) that the results were noticeably slower when we always used {{openssl}}. For a few larger files, {{openssl}} was actually faster. So the results are more inline with my expectation that streaming small files and computing the SHA-1 in our code is actually faster that forking off a separate process, whereas streaming (very) large files in Java should be more expensive than a native {{openssl}}, which has access to more operating system optimizations. (That doesn't mean that {{openssl}} _uses_ all the available optimizations on all platforms.) You'll notice that my solution was to always use {{openssl}} (if available) for all files larger than {{n}}, where {{n}} defaults to 50kB but is entirely configurable. This approach kept the configuration as simple as possible while still allowing someone control over {{n}}, including setting {{n}} to a very large value if {{openssl}} is never to be used or a very small value if {{openssl}} should always be used. *Can you test the code in 'master' in similar ways to what you did before, and see if this change fixes your problem? If you can do this and let us know ASAP, because we plan to release 3.6 in the middle of next week (roughly October 16). Currently, there is nothing holding up the release.* {quote} Ideally, being naive to modeshape, it seems hashes of a full file should NOT be the BinaryKey identifier but I don't know the fallout of changing this. {quote} The whole point of the binary key is to have a unique identifier that is representative of the content but independent of location (and thus moves). Note that most of the requirements come from how ModeShape manages/stores binary values in its binary store, where a content-based key allows us to store only one copy of any given binary value. Now, the FileSystemConnector (or rather a subclass of it) may actually be able to use a different binary key, since the connector actually [stores the URI of the file inside the BinaryValue|https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/connector/filesystem/FileSystemConnector.java#L385]. We've not really investigated that, though it could very well be that this causes problems when copying or moving external content into internal content. Feel free to subclass the connector and try different options. Subclassing would also be a way to provide caching of binary keys. IMO, caching is not worth the complexity, especially when we can't use an _unbounded_ cache.	1	1	https://issues.jboss.org/browse/MODE-2061
15573	MODE-2061	{quote} This breaks the functionality of the getHash() and getHexHash() methods which are now returning this URI {quote} What exactly does "break" mean here ? In the case of {{UrlBinaryValue}} it may be true that "a true SHA1" isn't returned, but from looking at the code the main purpose of the {{getHash}} methods is to aid in binary values comparison (i.e. value1.equals(value2)). This should still hold when the string version of the URL is used. Yes, the Javadocs for the above methods aren't really correct in this case, but I'm not sure I see the functional problem.	1	1	https://issues.jboss.org/browse/MODE-2061
15574	MODE-2061	Thanks for the update Eric. How is the SHA1 used by the fixity service ? Is it used to make assumptions about the content (bytes) of the file itself ?	1	1	https://issues.jboss.org/browse/MODE-2061
15575	MODE-2048	Can we just add this as additional function for rest api of items instead of creating new handler? We just replace target url from "../items/<node_path>" to "../items/move".	1	1	https://issues.jboss.org/browse/MODE-2048

#	Requirement	Questions	Answered	Successfully Answered	URL
15576	MODE-2048	[~aangeles], unfortunately the "{{.../items/move}}" resource URI is does not follow the RESTful style because it puts the verb into the resource URI and because the same URI would apply to multiple resources. A RESTful approach would instead keep the URI referencing the resource (the node) that is being moved and uses the POST or PUT operation. The question is then what the body of the request looks like. My current thought is that this can be handled by our current API as follows. The client would do the following steps: # GET the source node (the node to be moved/shared), and specifically its identifier # GET the new parent node's representation # Add an entry for the source node in the correct location in the "children" nested document, keeping the same identifier that the node has in its original location. # PUT the updated representation of the new parent ModeShape (specifically the [ItemHandler.updateNode(...)|https://github.com/ModeShape/modeshape/blob/master/web/modeshape-web-jcr-rest/src/main/java/org/modeshape/web/jcr/rest/handler/ItemHandler.java#L613] method) would then behave as follows: # Make sure that the children in the representation match the children in the current node. Any nodes that don't exist need to be created, and any nodes that exist elsewhere (by ID) need to be either moved (if not shareable) or cloned (if shareable). WDYT? Thoughts, [~hchiorean]?	1	1	https://issues.jboss.org/browse/MODE-2048
15577	MODE-2046	After thinking about this today, I'm not convinced that this is that useful. For the most part, it is really not much code to create a session and ensure it is logged out in a finally block. So it really isn't saving much code. However, it does increase the complexity a fair amount. Should the function or the framework perform the save? What happens if there is an exception during the same; the transient state in the session is still there, and can be accessed to either discover or fix what was wrong. It's also more constraining to force the application code into this pattern, especially when the pattern requires jumping through to different methods. I'll leave this open for a little while in case people want to discuss and advocate for this. Otherwise, we'll close it with a REJECTED.	0	0	https://issues.jboss.org/browse/MODE-2046
15578	MODE-2040	Which version of JDK is this in ?	0	0	https://issues.jboss.org/browse/MODE-2040
15579	MODE-2040	Sounds good, I noticed there isn't a 4 branch? Is this something in the works or is it a bit further off? Perhaps there is somewhere where I could read about it?	0	0	https://issues.jboss.org/browse/MODE-2040
15580	MODE-2033	I agree that {{RepositoryConfiguration.validate()}} is probably not appropriate for the reasons you mention. I think the problem is that the code that calls {{ModeShapeEngine.deploy(RepositoryConfiguration)}} doesn't get any feedback of any problems encountered during deployment. Sure, the problems (really, warnings) are logged, but that doesn't really help the calling code. Obviously any _errors_ will prevent the repository from being deployed at all. But _warnings_ are just logged; is there a way of recording the deployment warnings/problems in the {{JcrRepository}} instance and making them available after the repository is deployed?	1	1	https://issues.jboss.org/browse/MODE-2033
15581	MODE-2021	org.modeshape.sequencer.image.ImageMetadataSequencerTest 13:57:20,129 INFO ModeShape version 3.6-SNAPSHOT 13:57:20,667 INFO ISPN000128: Infinispan version: Infinispan 'Delirium' 5.2.6.Final 13:57:21,862 INFO HSEARCH000034: Hibernate Search 4.2.0.Final 13:57:21,878 INFO HCANN000001: Hibernate Commons Annotations {4.0.1.Final} 13:57:27,024 INFO Image has IFD0 block information but is not one of the standard image types, marking it as TIFF Tests run: 6, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 7.295 sec <<< FAILURE! shouldExtractExifInformationFromJpeg(org.modeshape.sequencer.image.ImageMetadataSequencerTest) Time elapsed: 0.493 sec <<< FAILURE! org.junit.ComparisonFailure: expected:<...8-05-08T14:54:46.000[+03]:00> but was:<...8-05-08T14:54:46.000[-05]:00> at org.junit.Assert.assertEquals(Assert.java:125) at org.junit.Assert.assertEquals(Assert.java:147) at org.modeshape.sequencer.image.ImageMetadataSequencerTest.shouldExtractExifInformationFromJpeg(ImageMetadataSequencerTest.java:105) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.rules.RunRules.evaluate(RunRules.java:18) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:47) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:50) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:222) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:234) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:133) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:114) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:188) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:166) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:86) at	0	0	https://issues.jboss.org/browse/MODE-2021

#	Requirement	Questions	Answered	Successfully Answered	URL
15582	MODE-2018	administrator-defined indexes. When no such indexes are defined, the query engine is able to still answer the queries by "scanning" all nodes in the repository. This is like a regular relational database: all query functionality works (albeith slowly) even when no indexes are defined, though to improve performance simply define an appropriate index based upon the query or queries that are being used. All of ModeShape's query parsing, planning, and optimization steps are basically unchanged from the previous query system. There is one addition to the rule-based optimizer: a new rule looks at query plans and adds the potential indexes that might be of use in each access query portion of a query plan. Then, the query execution process (see below) chooses one of the identified indexes based upon the selectivity and cardinality. If no index is available for that portion of the query plan, then the query engine simply iterates over all queryable nodes in the repository. A new kind of component, called a "query index provider", allows the query engine to delegate various responsibilities around indexes to these providers. For example, a provider must provide an index planner that can examine the constraints that apply to an access query and determine if any of the provider's indexes can be used. When they are, ModeShape adds those indexes to the query plan. If the query engine uses one of those indexes, then provider must be able to return all of those nodes that satisfy the criteria as described earlier by its index planner. Finally, as ModeShape content changes, ModeShape will notify the index providers' of the changes so that they can ensure their indexes are kept up-to-date with the content. This means that a provider can implement the functionality using any kind of technology, and consequently, that ModeShape can begin to leverage multiple kinds of search and index technology within its query system. The ModeShape community anticipates having providers that use Lucene, Solr, and ElasticSearch. ModeShape will also likely come with a provider that maintains file-system based indexes. Additionally, providers can optionally support indexes on one or more properties. Thus, it will be possible to mix and match these providers, selecting the best technology for the specific kind of index. The new query engine does the execution in a very different way than the previous engine, which used Lucene to determine the tuples (that is, the values in each row) for each access query and that were then further processed and combined to form the tuples that were returned in the result set. The new engine instead uses a new concept of a "stream of node keys" for each access query: what actually implements that stream depends on many factors. A node sequence is an abstraction of a stream of "rows" containing one or more node keys. The interfaces are designed to make it possibly to lazily implement a stream in a very efficient manner. Specifically, a node stream is actually comprised of multiple "batches" of rows, and batches can be of any size. Consider when the engine findes no indexes are available for a certain access query. The engine simply uses a "node sequence" (or {{NodeSequence}}) implementation that returns in batches a row for each node in the repository. But if an access query involves a criteria on the path of a node, such as "{{... WHERE ISSAMENODE('/foo/bar') ...}}", then ModeShape knows that this query (or portion of a query) will have only one result, namely the node at "{{/foo/bar}}". ModeShape doesn't need an index to quickly find this node; it merely has to navigate to that path to find the one node that satisfies this query. ModeShape has several other optimizations, too: it knows when a query involves all children or descendants of a node at a given path, and can take this into account when optimizing and executing the query. All of these are handled with special {{NodeSequence}} implementations optimized for each case. For many access queries (i.e., part of a larger query), the engine will use one of the indexes identified by one of the providers. When this happens, ModeShape uses other {{NodeSequence}} implementations that utilize the underlying indexes to find the nodes that satisfy some of the criteria. The above describes how the engine uses a single {{NodeSequence}} instance for each each access query in a larger query. But how does the engine combine these to determine the ultimate query results? Basically, the	0	0	https://issues.jboss.org/browse/MODE-2018
15583	MODE-1986	What's the required timeframe for this?	1	1	https://issues.jboss.org/browse/MODE-1986
15584	MODE-1986	It seems to me that if a user (or Designer) knows the parser they want to use, they should be able to designate that parser and use the "parseUsing" method. On the other hand, if the user (or Designer) does not know what parser to use, it seems they would want the parsing framework to pick the best one. This is how the framework's "parse" method is intended to work. It sounds like there is a problem in the framework, or maybe in an individual parser, when the first parser fails, the subsequent parser is completing successfully when it shouldn't be (i.e., returning an incomplete, problem-filled, or incorrect node graph). In the case above, I assume you didn't know which parser you wanted to use so you used the generic "parse" method. If you would've received a fully parsed node graph (without knowing what parser was used), or received a chain of parser exceptions signifying no parser could handle the DDL, would you still have had an issue? Or maybe you just want the user to be able to pick the parser after the scoring since they may have a bit more knowledge about the DDL than the framework does. And sometimes there are ties in the scoring and I'm not sure how the framework works in that case. Regardless of all of that, it seems there is a problem, rightfully so, with the fact you received an AstNode with just UnknownStatement child nodes. The Teiid parser was written to throw an exception if the DDL does not parse. One or more of the other parsers apparently completely parse (no exceptions) even if they don't recognize the DDL.	1	1	https://issues.jboss.org/browse/MODE-1986
15585	MODE-1979	[~blafond], can you create a pull-request with the changes? (Get the latest from 'master', create a topic branch, commit your changes, push to your fork, then create a pull-request.) Also, do you have a test case? It'd be a big help, and we can merge the fix onto our '3.3.x' branch.	1	1	https://issues.jboss.org/browse/MODE-1979
15586	MODE-1960	So how can the memory-usage-might-go-crazy issue be tackled? How can I make sure that my _critical_ JCR server doesn't stall due to some weird, random MS Office file?	0	0	https://issues.jboss.org/browse/MODE-1960
15587	MODE-1955	So this JCR feature has always been very nebulous to me. Can you give some scenarios where/how this would be used?	1	1	https://issues.jboss.org/browse/MODE-1955

#	Requirement	Questions	Answered	Successfully Answered	URL
15588	MODE-1955	Your last idea about the connector exposing an area of "unfiled" content seems like the best approach. (Besides, using a central "{{/jcr:system/jcr:unfiled}}" and having content from multiple connectors there would be challenging.) Have you looked at how the Git connector works? Strictly speaking, Git's history really isn't hierarchical in any way (the log for a single branch can be very, very long), but it creates an artificial hierarchy for the different branches and tags, as well as a place to find all of the commits (with tree view of the file structure under each commit). We also designed the connector SPI to work fairly well with very large numbers of children by working with pages. For example, you don't want a node representing the "master" branch to contain a child node for every commit in that branch - doing so would make it extremely expensive to create the document for the "master" branch node. So, instead, the GitConnector looks up a small number of recent commits and adds them as the first page of children under a branch node. As the client navigates/iterates over the children, ModeShape will transparently fetch additional pages as necessary. We've created a specific {{[Pageable|https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/federation/spi/Pageable.java]}} interface that your {{Connector}} (or {{ReadOnlyConnector}} or {{WritableConnector}}) subclass and implement. If you do, then ModeShape can take advantage of this capability. Where paging tends to breaks down, however, is when the client "knows" (from a variety of ways) a particular commit ID and then asks for that commit under a particular branch. In this case, ModeShape knows that it needs to get a specific child of the parent, and rather than relying upon the is being (e.g., a child of the branch where the child key is known), what ModeShape actually does is translate that into a call to the "{{getChildNode(parentKey,childKey)}}" method of the connector. I think the Pageable facet of the Connector SPI will let you implement a connector that does have a separate "unfiled" area that is effectively a huge list, and will allow client applications to get a particular node in that "unfiled" area when they know which node they want (usually by getting a reference or identifier from some other content).	1	1	https://issues.jboss.org/browse/MODE-1955
15589	MODE-1955	Trying to add unfiled support Creating versionable nodes under /jcr:system/jcr:unfiled (as it is told in specification). "jcr:unfiled" added as nt:folder (regular node - not system) which lets us to manage it's children. When create regular (eg nt:file) nodes it works but for versionalbe ones we get duplicate key error. This session tried to create node '/jcr:system/jcr:unfiled/sjesjsj' (with key '589e942317f1e737769d21-5a53-4c6b-9dbd-458eaf4e3e96'), but a node with that key already exists. {code} javax.jcr.InvalidItemStateException: This session tried to create node '/jcr:system/jcr:unfiled/sjesjsj' (with key '589e942317f1e737769d21-5a53-4c6b-9dbd-458eaf4e3e96'), but a node with that key already exists. Caused by: org.modeshape.jcr.cache.DocumentAlreadyExistsException: 589e942317f1e737769d21-5a53-4c6b-9dbd-458eaf4e3e96 at org.modeshape.jcr.cache.document.WritableSessionCache.persistChanges(WritableSessionCache.java:1077) at org.modeshape.jcr.cache.document.WritableSessionCache.save(WritableSessionCache.java:543) at org.modeshape.jcr.JcrSession.save(JcrSession.java:1004) ... 45 more {code} As not versionable nodes works OK I assume error is caused by the fact that versionable node has corresponding VersionHistory having the same key in the same system workspace. Does it make sense to you? And how difficult it would be to make VersionHistory to have different id and being correctly linked to a node ?	0	0	https://issues.jboss.org/browse/MODE-1955
15590	MODE-1955	Yes, I think your analysis is correct. All of the '/jcr:system' branch is actually placed within a hidden 'system' workspace and then shared to each of the real workspaces. The version storage exists under '/jcr:system', and there will be a version history for a given node that has the same key as the versionable node (that exists in a separate workspace). By also putting the unfiled content here, then any unfiled nodes that are also 'mix:versionable' will result in ModeShape attempting to also create a tree with the same identifier under version storage in the same hidden workspace. I'm not sure that we can really change how version history is stored - there are existing repositories whose content would have to be migrated, and that's very undesirable. Instead, I think we have to find a way for the unfiled nodes to work without changing how the existing content is done. One option is to use a different workspace in the key, though that might mean we actually need to create a second hidden workspace for everything else related to NodeKeys to work properly. Can anyone else think of a different approach?	1	1	https://issues.jboss.org/browse/MODE-1955
15591	MODE-1955	I'm not 100% clear what the purpose of the {{jcr:unfiled}} area is, but should it be trying to store nodes with the same identifier (the id part of the key) as standard (in session created) JCR nodes ? In other words, if node A is a versionable node created in a JCR session, why is this node's identifier being duplicated in {{jcr:system/jcr:unfiled}} ? IMO we shouldn't start using a second hidden workspace, as that would only complicate things a great deal. As an alternative, can't a separate source key be used and an approach similar to the {{FederatedDocumentStore}} ?	1	1	https://issues.jboss.org/browse/MODE-1955
15592	MODE-1955	>I'm not sure that we can really change how version history is stored - there are existing repositories whose content would have to be migrated, and that's very undesirable Well, can this be configurable or something? Like add some switch and based on it enable unfiled + changed versionHistory storage. The "existing" production data issue should not be blocker for improvements I guess.	1	1	https://issues.jboss.org/browse/MODE-1955
15593	MODE-1955	{quote} > I'm not sure that we can really change how version history is stored - there are existing > repositories whose content would have to be migrated, and that's very undesirable Well, can this be configurable or something? Like add some switch and based on it enable unfiled + changed versionHistory storage. The "existing" production data issue should not be blocker for improvements I guess. {quote} We obviously have to be cognizant of maintaining backward compatibility. Generally speaking, we should avoid changing how data is persisted, unless there is a *really* good reason to do so. Besides, we've already described two possible approaches that should mitigate the duplicate IDs without affecting any existing data. So as long as at least one of those works, then there's no reason to (arbitrarily) change how the existing version histories are stored. Finally, it took a lot of time and effort to get the versioning functionality working, and I'm not enthused about changing some of the fundamental mechanisms, especially if doing so will break existing repositories.	1	1	https://issues.jboss.org/browse/MODE-1955

#	Requirement	Questions	Answered	Successfully Answered	URL
15594	MODE-1955	[~hchiorean] and I had a chat today, and we have some questions about this feature. Here is _everything_ that the JSR-283 specification has to say about unfiled content: {quote} 3.12 Unfiled Content Implementers that build a JCR repository on top of an existing repository that supports content objects which exist outside of a hierarchical structure may expose these objects as nodes below {{/jcr:system/jcr:unfiled}} (see §3.11 System Node). The hierarchical structure below {{/jcr:system/jcr:unfiled}} is implementation-dependent. JCR implementers may disallow discovery (listing) of the nodes beneath this folder. In such a case a call to {{Node.getNodes()}} on the {{jcr:unfiled}} node would throw a RepositoryException. JCR implementers may expose the nodes below {{jcr:unfiled}} to search through the query mechanism (see §6 Query).{quote} (The only other place unfiled content is mentioned is the associated Repository descriptor in Section 24.2.2.) It concerns me that in order to get this to work efficiently, we'd have to add a significant amount of code to actually disable the discovery (listing) of nodes under "{{/jcr:system/jcr:unfiled}}". Thus, ModeShape would likely treat the unfiled area as just regular content. Yes, it happens to be in the system area, but that actually makes it worse and harder to deal with. In other words, ModeShape would not treat unfiled content any differently than regular workspace nodes, and that this unfiled area would not really perform or scale any better than placing a bunch of nodes under a single "nt:unstructured" parent (or any other acceptable node type). So given this, why do we want to add this feature? What benefit does it offer beyond the existing capability to place lots of nodes under a single, regular parent node in a workspace?	1	1	https://issues.jboss.org/browse/MODE-1955
15595	MODE-1955	{quote} can you please explain why "/jcr:system/jcr:unfiled" is useful to a client beyond just putting nodes under a single normal parent? {quote} Consistency? Do not need hierarchy - go unfiled and get your parent node as null, need hierarchy - here is your root node? Sorry, maybe I'm misunderstanding the complexity of this change in modeshape, so for me either way is an equal effort :)	1	1	https://issues.jboss.org/browse/MODE-1955
15596	MODE-1918	Hi Oleg, do means this way? {code:xml} <ironjacamar> <connection-definitions> <connection-definition class-name="org.modeshape.jca.JcrManagedConnectionFactory" jndi-name="jcrJCA" enabled="true" use-java-context="true" pool-name="modeshape-Pool" use-ccm="true"> </connection-definition> </connection-definitions> <config-property name="repositoryURL"> file:/${jboss.server.data.dir}/ms-repo </config-property> </ironjacamar> {code}	1	0	https://issues.jboss.org/browse/MODE-1918
15597	MODE-1918	Dmitri, yes. You can find sample configuration by this link https://github.com/ModeShape/modeshape/blob/master/modeshape-jca/src/main/resources/META-INF/ironjacamar.xml Now you have some kind of problem with repository configuration (resource adapter was started and failed on repository start up) which is specified by URL. Can you please show your configuration file? Another reason might be a variable in the url, RA does not resolve properties so I am not sure about what is the actual path. Please try also to specify an absolute path to configuration just for check	1	1	https://issues.jboss.org/browse/MODE-1918
15598	MODE-1918	[~dimonv], the JSON configuration file you posted above is invalid: there is an extra closing brace before the "workspaces" field. {code:javascript} { "name" : "CIM", "storage" : { "transactionManagerLookup" : "org.infinispan.transaction.lookup.DummyTransactionManagerLookup" }, "security" : { "jaas" : { "policyName" : "modeshape-jcr" }, "anonymous" : { "username" : "<anonymous>", "roles" : ["readonly","readwrite","admin"], "useOnFailedLogin" : true } }, <---------- THIS LINE SHOULD NOT BE HERE ---------- "workspaces" : { "predefined" : ["ws1", "ws2"], "default" : "default", "allowCreation" : true }, "node-types" : ["nodetypes.cnd"] } {code} Such an error in the configuration file will prevent the repository from starting, which means that nothing will get bound in JNDI to "{{java:jcr/local/CIM}}" and the {{NameNotFoundException}} is logical. Please correct the JSON configuration file and try again, continuing to look at the log for any errors or problems. ModeShape should be outputting any such repository startup errors via ERROR log messages.	1	1	https://issues.jboss.org/browse/MODE-1918
15599	MODE-1918	[~dimonv], can you please give an update? May I ask why you're deploying ModeShape via the JCA Adapter approach when we have a special kit for AS7/EAP? The problem with _not_ using our special kit is that ModeShape requires a specific version of Infinispan, and using the kit is really the best way to make sure that the versions match. The other benefit of using the ModeShape kit for AS7/EAP (which simply unzips on top of AS7/EAP) is that you don't have to (and really should not) use the JCA adapter at all. That's because our kit installs a ModeShape-specific subsystem into the application server, and this subsystem is controlled by the normal AS7/EAP configuration mechanism. See our documentation on the [kit for AS7/EAP|https://docs.jboss.org/author/display/MODE/ModeShape+and+JBoss+AS7], including how to [configure the subsystem|https://docs.jboss.org/author/display/MODE/Configuring+ModeShape+in+AS7]. I'm considering marking this as either CANNOT REPRODUCE (if the aforementioned correction to the custom JSON configuration file works) or WONTFIX (since we have a special kit for AS7/EAP).	1	1	https://issues.jboss.org/browse/MODE-1918
15600	MODE-1918	Hello Randall, sorry for not commenting so long. Why am I deploying Modeshape as a JCA adapter? Because: - it is easier to deploy without changing application server - it is a proper JEE way to include a RAR file to EAR and deploy it - its easier to acquire JCR session over JNDI and connection factory - it is not necessary to change application server - every deployed application can use a different version of JCR implementation Regards, Dmitri	1	0	https://issues.jboss.org/browse/MODE-1918
15601	MODE-1901	I assume that you mean to include something like a [mode:plan] property in the response? Would this mean that you have to actually execute the query to get the plan back? That isn't always a great thing to do if the query you are trying to get an explanation of is causing problems by taking an hour to run. Have I misunderstanding this? The usual behaviour of EXPLAIN is to return a rowset containing the query plan so for example if I run the following explain then I get back two rows with little execution time. Since the result is in the same form as the result of any other query, the usual tools (via JDBC) work quite well with this. {code} mysql> explain select * from (select date,time,abs(id) as id from txn order by date desc) x order by time,abs(id); +----+-------------+-----------+------+---------------+------+---------+------+------+----------------+ | id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra | +----+-------------+-----------+------+---------------+------+---------+------+------+----------------+ | 1 | PRIMARY | <derived2> | ALL | NULL | NULL | NULL | NULL | 333 | Using filesort | | 2 | DERIVED | txn | ALL | NULL | NULL | NULL | NULL | 333 | Using filesort | +----+-------------+-----------+------+---------------+------+---------+------+------+----------------+ 2 rows in set (0.00 sec) {code} but if I run the actual query I get back 100s or 1000s of rows and it takes a lot of time to execute.	0	0	https://issues.jboss.org/browse/MODE-1901
15602	MODE-1868	That seems very reasonable. If it were possible to do this in a way that was reusable (for other connectors) of course we would choose that. But if it's not going to be part of the general contract for connectors, then it makes a lot of sense to do it as part of a specific implementation. If I understand the design arightly, we should be able to subclass JsonSidecarExtraPropertyStore and mix in node types in the *properties() methods. But JsonSidecarExtraPropertyStore is default visible. Would it be possible to make this a public class, or to insert an abstract class that it extends and which implements ExtraPropertiesStore? Then we could subclass it (or the abstract class) and add only a minimum of new logic.	1	1	https://issues.jboss.org/browse/MODE-1868

#	Requirement	Questions	Answered	Successfully Answered	URL
15603	MODE-1864	Had a more in-depth look at implementing this and before detailing some of the problems, the following need to be taken into account with regards to the current versioning implementation: 1. everything version related is stored in a predefined structure in the {{jcr:system}} workspace, which is "outside" the scope of any current client session. 2. the structure, node names & ids of the version history (for any given version node) are based on the SHA-1 of the node's {{NodeKey}}. 3. there was no way (until the API extension in {{4.0.0.Alpha1}}) to remove the version history of a node, even if that node was removed. Considering this, at least the following aspects/scenarios are impossible/way too complex to handle: * importing an XML via a session (not a workspace): according to the JCR spec it's possible to transiently import nodes into a session which will only be persisted on {{save}}. However, if a client were to call {{refresh}} those would have to be discarded. Because of 1. from above, it would impossible in this case to "transiently alter the system area". Therefore, any such feature would only realistically be possible for the "workspace-write" import operations, which isn't ok IMO. * importing an XML containing versions into a repository for a given node which already contained that node & versions: when importing nodes, if either the {{ImportUUIDBehavior.IMPORT_UUID_CREATE_NEW}} or {{ImportUUIDBehavior.IMPORT_UUID_COLLISION_REPLACE_EXISTING}} options are used, the key of the existing node or the key of the new node would effectively change. This means the hierarchy of the existing/incoming version history nodes would need to be changed to reflect this (see 2. from above). There is no straightforward way to accommodate this, because the process of initializing a version history is extremely complex. * importing an XML containing versions into a repository which doesn't contain the versionable node anymore, but still has its entire version history (see 3. from above). Since it's possible to restore(revive) a removed node from an existing version, how would the import behave in this case ? These are just a few aspects, but I'm sure a lot more would come up. Therefore, I think this feature is *not* realistically possible or *worth the effort*. As an alternative, the backup/restore process would restore the entire {{jcr:system}} area, including any versioning information. [~rhauch]: thoughts ?	1	0	https://issues.jboss.org/browse/MODE-1864
15604	MODE-1860	Do you think it's possible to change the location to use a Maven property?	1	1	https://issues.jboss.org/browse/MODE-1860
15605	MODE-1849	What about the dist zip we use for testing ?	1	1	https://issues.jboss.org/browse/MODE-1849
15606	MODE-1849	BTW, I noticed that we are putting our modules under "{{modules/system/layers/base}}". Should we be putting our modules in some location other than "{{base}}"? AS7-6344 seems to cover the changeover to this structure, and it references https://community.jboss.org/docs/DOC-47927 (which describes the basic concepts and responsibilities) and https://community.jboss.org/docs/DOC-48264 (more of an internal AS7 documentation about how to implement these layers).	1	1	https://issues.jboss.org/browse/MODE-1849
15607	MODE-1839	The JCR API doesn't provide a way of doing this. Calling "{{Version.remove()}}" or "{{VersionHistory.remove()}}" (which is the same {{remove()}} inherited from {{Item}}) is expressly forbidden and tested by the TCK. It is possible to remove all of the versions from the history using "{{VersionHistory.removeVersion(String versionName)}}", but it is not possible to remove the root version and the {{VersionHistory}} node. There are a couple of challenges to this. One, it's not really possible through the JCR API to find the version history for a versionable node that has already been deleted. That's because the only way to find the version history is through the "{{VersionManager.getVersionHistory(String path)}}", where "{{path}}" is the path of the versionable node. _*UPDATE:* I think it should be possible to pass an identifier path (e.g., a node's identifier surrounded with square brackets) to the {{VersionManager.getVersionHistory(String path}} method. If this doesn't work now, we should fix it._ Now, ModeShape implements this by looking up the versionable node by the path, obtaining the node key, and then finding the VersionHistory node given the history's key (which can be derived from the node key). So it would be possible to expose a ModeShape-specific API to obtain the history given a node identifier. Secondly, when a versionable node is removed, it can (theoretically) be restored using the {{VersionHistory}} node. Of course, getting the {{VersionHistory}} node without an existing versionable node is a challenge, as previously mentioned. Thirdly, there may be some regulatory reasons why the version history should not be removed. OTOH, there actually may be some regulatory or policy requirements for being able to remove all of the data associated with some node. But either way, no single policy for removing or keeping the history can be stipulated by ModeShape; instead, we can provide the tools and functionality so that the application can implement the desired behavior. Finally, there are several considerations to removing version history: # The version history for a node should be removed when the node is removed. This might be exposed with a ModeShape-specific public API to include version histories for all versionable nodes below some node being removed. For example, {{org.modeshape.jcr.api.Node.remove(boolean includeVersionHistories)}}. This is essentially a more thorough form of the standard {{Node.remove()}} that would work on subtrees as well. # Provide a ModeShape-specific API to obtain the version history for a (potentially-non-existant) versionable node given the versionable node's identifier, and to provide a way of destroying the entire version history via a ModeShape-specific API. I'm not sure how practical this is given that the caller must somehow know the identifier of the versionable node. Could this be called when the versionable node still exists (see next item)? _*UPDATE:* I think it should be possible to pass an identifier path (e.g., a node's identifier surrounded with square brackets) to the {{VersionManager.getVersionHistory(String path)}} method. If this doesn't work now, we should fix it._ # What happens when a versionable node is simply changed to no longer be versionable, but it is not removed. Essentially, the node would need to have "{{mix:versionable}}" explicitly in the list of mixins and not have it in any of the other node types. In this case, does JCR stipulate what happens to the version history when a versionable node is changed to no longer be versionable. If the version history is to be removed, then perhaps this is a good approach for removing the version history of a single node without requiring a non-standard public API. # Provide a ModeShape-specific API to reclaim/remove all version histories for versionable nodes that no longer exist. This could be done with a batch process, but does the caller need any feedback? We need to decide which of these (or other) aspects we want to deal with. We at least need to figure out what the specification and/or TCK tests expect for #3.	1	1	https://issues.jboss.org/browse/MODE-1839
15608	MODE-1839	Hello, I've found another behaviour that can be related to this jira. I'm trying to find a workaround to remove history in the way to remove the mixin mix:versionable. But I can't once that you use a versionable node I have an exception like this: {code} javax.jcr.nodetype.ConstraintViolationException: Cannot find a definition for the property named 'jcr:versionHistory' on the node at '/test1' with primary type 'nt:file' and mixin types: [mix:mimeType,mix:lastModified,mix:title] at org.modeshape.jcr.AbstractJcrNode.removeMixin(AbstractJcrNode.java:2534) at org.modeshape.VersioningAPITest.createVersioningNodesAndRemovingVersioningDelete(VersioningAPITest.java:209) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) {code}	0	0	https://issues.jboss.org/browse/MODE-1839
15609	MODE-1820	Are you forking/reusing another existing web app (like JCR Explorer)?	1	1	https://issues.jboss.org/browse/MODE-1820

#	Requirement	Questions	Answered	Successfully Answered	URL
15610	MODE-1820	[~oleg.kulikov], can you be a bit more specific by providing a URL? There are several projects ambiguously referred to as "JCR Explorer", including: * http://code.google.com/p/jackrabbitexplorer/ (and https://github.com/lucasponce/modeshape-tools) * http://www.jcr-explorer.org/screenshots.html We need to keep the original license information, and we need to verify that the licenses are all okay for us to use.	1	1	https://issues.jboss.org/browse/MODE-1820
15611	MODE-1820	@okulikov , you mentioned that you were going to rework this. Do you have an updated PR or status?	0	0	https://issues.jboss.org/browse/MODE-1820
15612	MODE-1820	Rebased and merged the [first pull-request|https://github.com/ModeShape/modeshape/pull/965] into the 'master' branch. Work still left to be done: # Add the webapp to the EAP kit (as discussed on IRC, and similarly to the other web apps) # Consider renaming the module to not use "`console`". We hope to have integration with the EAP/Wildfly console, whereas this application is more of a repository explorer. What about "`modeshape-explorer`"?	1	1	https://issues.jboss.org/browse/MODE-1820
15613	MODE-1820	+1 for explorer, but are we limiting this for use only with ModeShape as opposed to any JCR-compliant implementation ?	1	1	https://issues.jboss.org/browse/MODE-1820
15614	MODE-1818	IMO this should be part of 3.2, as it's an enhancement request. For 3.1.x it should suffice to bump up the default, hardcoded value.	1	1	https://issues.jboss.org/browse/MODE-1818
15615	MODE-1815	[~psiroky-redhat.com], what is the status of this effort? Are you waiting for anything in particular from us?	1	1	https://issues.jboss.org/browse/MODE-1815
15616	MODE-1815	[~rhauch]: the fix version is still 3.6, shouldn't it be 4.0 ?	0	0	https://issues.jboss.org/browse/MODE-1815
15617	MODE-1803	It is not straightforward to get the Session for which the connector method are being called, so using a ThreadLocal may be the best approach. Just remember that when the connector is called, the session will already be within the scope of a call to the Session, so watch out for any concurrency problems. Do you have to pass in the same session (because you need the transient state), or can it be any session?	1	0	https://issues.jboss.org/browse/MODE-1803
15618	MODE-1803	>Do you have to pass in the same session (because you need the transient state), or can it be any session? I guess the same session/transaction and yes the new ID may not be visible outside of it. It makes no sense to keep cross-reference record if there was a fail.	0	0	https://issues.jboss.org/browse/MODE-1803
15619	MODE-1801	Is this still necessary since we've introduced in 3.4.0.Final support for JCR ACLs? If not, then I'd like to close this.	1	1	https://issues.jboss.org/browse/MODE-1801
15620	MODE-1784	Changed the way ModeShape web apps are deployed via the AS7 kit: Instead of statically packaging web applications into the kit and copying them over in the AS7 deployments folder, the following changes were made: * the web applications were moved into a _deployments_ folder under the _org.modeshape_ AS7 module * the AS7 subsystem was extended, with {code}<webapp name="" exploded="" auto-deploy=""/> {code} directly under the _<subsystem>_ element ** _name_ represents the name of the folder or file inside the deployments folder ** _exploded_ (default true) indicates if the packaged web app is exploded or not ** _auto-deploy_ (default true) indicates if the web application should be deployed automatically, when the subsystem starts * based on the above, web applications from the deployments directory are deployed at runtime, when the subsystem starts	1	1	https://issues.jboss.org/browse/MODE-1784
15621	MODE-1777	Perhaps I'm doing it wrong, but I don't see a difference in which tests are run by command-line Maven for "{{@Test}}" methods with and without the "{{@SkipLongRunning}}" annotation. Should this be reopened?	1	1	https://issues.jboss.org/browse/MODE-1777
15622	MODE-1764	Or AS7 subsystem relies on the AS7 core (in this case the clustering subsystem) to inject a {{CacheContainer}} into our {{RepositoryService}}, which will then be retrieved by our environment. How can we change this to obtain an ISPN cache configured in AS7 using the 5.2 slot ?	1	1	https://issues.jboss.org/browse/MODE-1764
15623	MODE-1764	Do we still need to keep this open ? (the attached PR has been merged and we have separate issues for the remaining open problems)	1	1	https://issues.jboss.org/browse/MODE-1764
15624	MODE-1723	Just a status update ... Finished coding the composite binary store. Now all the repository content (children and properties) are visible in the RHQ component tree under the EAP node. Also added "Web Applications" content as a sibling to "Repositories." The only information is the name of the web app and if the content has been exploded. Is this something of value? Added a new repository properties grouping for the 5 document optimization properties. I've titled this grouping "Content Optimization." Is that a good title? Currently working on reading and writing the metrics window(s) data from/to the plugin data directory and then updating the RHQ metrics report when requested.	0	0	https://issues.jboss.org/browse/MODE-1723
15625	MODE-1713	The more I think about it, the more I agree that the document-level (i.e., low-level) changes would be much more difficult for a connector implementation to work with. As for the Semantic Changes, I agree that {{SessionNode}} should not directly implement the same SPI interface that connectors use. And that would mean essentially "copying" the changes in the {{SessionNode}} into the changes object, which would be safer all the way around (guarantees no leakage in either direction) {quote} One potential challenge I see here, is that modified values coming from the SessionNode (e.g. property values) may have a different interpretation on the "document layer", because on documents, values are transformed via the DocumentTranslator#valueToDocument call. This might be a more generic problem that's already present, but which we haven't encountered yet in tests. {quote} That's true, but if they use the {{DocumentReader}} and access the {{Property}} representation, those values auto-convert to the desired types. This does expose that (so far internal) API to the connector, but I think the "{{value}}" API is extremely useful in this regard, especially for the more arcane types like DATE. One issue I do expect to have, however, is that the changes can contain {{NodeKey}} values that may need to be reduced to {{String}}s, though we probably need to do this for only the matching source rather than do it blindly for all {{NodeKey}}s (as some {{NodeKey}} values might refer to content not included in the connector's content. Is that something we can/want to prevent?	1	1	https://issues.jboss.org/browse/MODE-1713
15626	MODE-1713	What about NodeKeys that refer to nodes outside of the connector's content? For example, in REFERENCE or WEAKREFERENCE properties. (I think it should be up to a connector to prevent this, unless we make a global decision that no connectors can ever have references to content outside of the source.)	1	1	https://issues.jboss.org/browse/MODE-1713
15627	MODE-1712	[~rhauch]: this seems to me as something quite complicated to implement and honestly I wouldn't know where to start. Also, the concept of "navigation" is a bit vague IMO. What's the driving factor/use case behind this requirement and why would we want this in 3.6 ?	1	1	https://issues.jboss.org/browse/MODE-1712
15628	MODE-1712	[~oleg.kulikov], what's the status on this? See the discussion on the [PR|https://github.com/ModeShape/modeshape/pull/992] for the latest.	1	1	https://issues.jboss.org/browse/MODE-1712

#	Requirement	Questions	Answered	Successfully Answered	URL
15629	MODE-1711	I see at least 2 different options here: 1) Define the projections as part of the repository JSON config file, within each source: {code} "externalSources" : { "readonly-files" : { "classname" : "org.modeshape.connector.filesystem.FileSystemConnector", "directoryPath" : "target/federation/files-read", "readonly" : true, "projections" : ["ws1:/node1/node2/root => /", "ws2:/node1/node2/file1 => /file1"] } } {code} Where the {{projections}} element represents an array of strings, where each string has the format {{workspaceName}} *:* {{repositoryPath}} *=>* {{externalPath}}. This format poses the following questions: * what if parts of the projection path do not exist, do we create those nodes ? If yes, what typing information do we use ? * what if parts of the projection path already exist, do we ignore the projection or just create the nodes for the missing segments ? * do we process the projections each time a repository starts ? 2) Define the projections as part of the initial content: {code} <?xml version="1.0" encoding="UTF-8"?> <jcr:root xmlns:jcr="http://www.jcp.org/jcr/1.0"> <node1/> <!-- Creates a projection node2/testFile -> / --> <node2> <testFile source-id="files-store" path="/"> <!--Any content would be ignored--> </testFile> </node2> </jcr:root> {code} Here an XML element with the {{source-id}} and {{path}} elements is treated as a projection which has the name/alias that of the XML element. For this to work, IMO we need some conventions in place: * both {{source-id}} and {{path}} are required for an element to be considered a projection. If either are missing, the node is discarded * any content (children) within a projection node is ignored Using initial content however, would mean projections are created only once, when initial import takes place on a "fresh" repository. Thoughts ?	0	0	https://issues.jboss.org/browse/MODE-1711
15630	MODE-1710	I'm trying to write unit tests for this new code, but having trouble finding a model. I can't seem to find any tests for the federation code in modeshape-jar/src/test/java. Is there somewhere else I should be looking?	1	1	https://issues.jboss.org/browse/MODE-1710
15631	MODE-1708	I think there are several ways we can specify how nodes should be excluded from the indexes (and thus not queryable): # Use a mechanism that uses node type information to determine which nodes are indexable. There is already something for full-text queries, but such a property would be ModeShape-specific. # Use a mechanism that encodes information within the document, but outside the normal JCR properties area. Connectors could automatically set this information. This would be similar to the {{cacheTTL}} mechanism currently used by the connectors to specify how long the node representations should be cached in a workspace. # Use a ModeShape-specific mixin to determine if a node should not be indexed. The advantages are that JCR clients can see this, that it can be defined in initial content, and it can easily be defined as part of node types (ensuring all nodes of that type are never indexed). But exposing to clients is also perhaps a disadvantage. # Use the "childrenInfo/count" value to determine if a node's children should not be indexed. For example, if the count is -1, then the children would never be index. Should we pick just one, or are there benefits to having multiple approaches? Note that {{RepositoryQueryManager}} already has [logic|https://github.com/ModeShape/modeshape/blob/master/modeshape-jcr/src/main/java/org/modeshape/jcr/RepositoryQueryManager.java#L343] to exclude the locks and version storage area from the indexes. Ideally we could eliminate that custom logic and use one of the mechanisms above.	1	1	https://issues.jboss.org/browse/MODE-1708
15632	MODE-1708	Yes, it would be very difficult to use option 2 in areas (e.g., system area, and by clients) that aren't normally using the Document representation directly. I think my preference is for option 3, simply because it's easily used everywhere by anyone. Plus, a node type definition can use the mixin as a supertype, which means JCR clients couldn't change the behavior of nodes of that type. Is it sufficient to have a single mixin that would signal that the node should not be indexed? Or do we need a second one that says "do not index my children"? IMO, the first would be sufficient for now, and if there's a need later we can always do it. We can even define the "do not index my children" as a supertype of "do not index me".	1	1	https://issues.jboss.org/browse/MODE-1708
15633	MODE-1708	I agree that for the time being, one mixin should be enough. How do you want to call it ? (e.g. mode:queryable)	1	1	https://issues.jboss.org/browse/MODE-1708
15634	MODE-1703	[~jpalka], can you please review the pull-request and (if possible) build ModeShape 3.0.1-SNAPSHOT and run your application with it?	1	1	https://issues.jboss.org/browse/MODE-1703
15635	MODE-1703	Jaroslaw, any more detail on the problems your having? Does this fix the issue (as I suspect it should)? I'm marking this as resolved, but can reopen if there still are issues.	1	1	https://issues.jboss.org/browse/MODE-1703
15636	MODE-1698	Based on the current description, this issue does not make sense to me: ModeShape 3.x and 4.x is/will be fully JCR (and TCK) compliant, meaning it will not alter any parts of the implemented JCR spec. Everything else is provided as an extension to the spec, meaning clients can always opt to *not use the extension*. Is there any good example how a {{strict}} mode would be beneficial ?	1	1	https://issues.jboss.org/browse/MODE-1698
15637	MODE-1695	The EJB related examples from the EAP quickstarts, make use of "dummy" web apps which basically contain JSF controllers and some front ends so that the user can "click trigger" the example. Do we really want this for our examples as well or is it enough to use "arquillian-driven" examples ? (the EAP quickstarts contain some of those as well)	1	1	https://issues.jboss.org/browse/MODE-1695
15638	MODE-1695	Added 4 different quickstarts, together with a distribution module: * _modeshape-servlet_ - is the porting of the modeshape-web-example module from the examples repository * _modeshape-ejb_ - is a quickstart that shows how to use ModeShape from 3 different types of EJBs: Stateful Singleton, Stateless CMT and Stateless BMT. * _modeshape-cdi_ is a quickstart that show how to use CDI together with _@Producer_ and _@Consumer_ methods to inject a ModeShape session * _modeshape-cli_ is a quickstart which contains a couple of CLI scripts that are used to setup/clear the repository. All quickstarts contain remote Arquillian based unit tests.	1	1	https://issues.jboss.org/browse/MODE-1695
15639	MODE-1688	If this is an issue specific only to the CMIS connector, IMO we should change the description of this issue. Also, looking the {{CmisConnector}} code, there isn't a configuration option which stores neither username nor password. From what I can tell, only SOAP service endpoints are configured. So how does password encryption apply - i.e. if there is a SOAP authorization endpoint, *it* should be responsible for handling encrypted passwords, not ModeShape.	1	0	https://issues.jboss.org/browse/MODE-1688
15640	MODE-1688	[~vasilievip]: are you referring to ongoing (not committed) work for the CMIS connector ? Based on https://github.com/ModeShape/modeshape/blob/master/connectors/modeshape-connector-cmis/src/main/java/org/modeshape/connector/cmis/CmisConnector.java I'm not sure I understand how you can set a username/password from the JSON configuration file. The other existing connectors: FS, DBMetadata and Git all function without configuring a user password.	1	0	https://issues.jboss.org/browse/MODE-1688
15641	MODE-1680	Ivan, what is rate of passed tests when you are running TCK?	1	1	https://issues.jboss.org/browse/MODE-1680
15642	MODE-1680	Can you give us any information about which tests caused which results? And can you give us a URL to the test source?	1	1	https://issues.jboss.org/browse/MODE-1680

#	Requirement	Questions	Answered	Successfully Answered	URL
15643	MODE-1680	Jackrabbit does support child axes in order by clauses, but it is non-standard. See [https://issues.apache.org/jira/browse/JCR-800]. [~oleg.kulikov], can you look into how much effort it would take to support it? We can then evaluate whether it's worth the effort. [~vasilievip], do you know why the CMIS<->JCR bridge implementation uses XPath and not the more powerful JCR-SQL2? Is there a way for the bridge to be enhanced?	1	1	https://issues.jboss.org/browse/MODE-1680
15644	MODE-1680	So, are you saying that there is an issue with apache chemistry? Here is the [apache cmis code which seems wrong|https://github.com/apache/chemistry-opencmis/blob/trunk/chemistry-opencmis-server/chemistry-opencmis-server-jcr/src/main/java/org/apache/chemistry/opencmis/jcr/impl/DefaultDocumentIdentifierMap.java#L33]	1	1	https://issues.jboss.org/browse/MODE-1680
15645	MODE-1598	Brian, thanks for the feedback. I agree that a command line tool would be valuable. Do you want to log an enhancement request? Scripting would be important, but can you state any specific JCR-related scenarios or capabilities or ideas you'd like to see in the tool? (I have several ideas, but I'd rather hear users' expectations and ideas first.)	0	0	https://issues.jboss.org/browse/MODE-1598
15646	MODE-1587	[~hchiorean], is the precondition for our AS7.2 is to have that version of AS7 _installed_ locally or to _build_ that version locally? If the latter, then the integration test's POM can still install the AS7.2 kit in the {{target}} directory.	1	1	https://issues.jboss.org/browse/MODE-1587
15647	MODE-1583	Node type functionality, at a minimum the ability to upload a CND file. And/or a JSON interface to create, update, retrieve node types?	0	0	https://issues.jboss.org/browse/MODE-1583
15648	MODE-1576	Changing the resolution to 'Cannot Reproduce Bug', since it has been shown to work.	0	0	https://issues.jboss.org/browse/MODE-1576
15649	MODE-1561	Is it possible to have the same mimetype configuration for the text extractor (included, excluded) in ModeShape3 (and json) as it is in ModeShape2 (and xml)?	0	0	https://issues.jboss.org/browse/MODE-1561
15650	MODE-1561	Implemented the writeLimit functionality for 3.x as well. However, our current version of Tika (1.2) has this issue: https://issues.apache.org/jira/browse/TIKA-1069 for which I've added a workaround.	1	1	https://issues.jboss.org/browse/MODE-1561
15651	MODE-1535	[~rhauch]: since we're moving away from HSearch, shouldn't this be closed ?	1	0	https://issues.jboss.org/browse/MODE-1535
15652	MODE-1516	One interesting thing regarding the sequencing configuration in AS7, is that the "classname" attribute from the AS7 configuration xml, gets translated to the "type" field from our sequencer configuration (org.modeshape.jboss.subsystem.AddSequencer#92) These would seem to indicate that the "classname" attribute from our JSON schema configuration cannot be set atm from an AS7 configuration file, meaning that no custom sequencers could be defined inside AS7. Wouldn't it be better to keep the same attribute names both in the AS7 config and our schema ? (e.g. type - type, classname - classname etc)	1	1	https://issues.jboss.org/browse/MODE-1516
15653	MODE-1516	Ok. Wouldn't it be better to name the AS7 xml attribute "type" for the sequencer so there's a 1-1 mapping ?	1	1	https://issues.jboss.org/browse/MODE-1516
15654	MODE-1516	org.modeshape.jcr.JcrRepository.login(JcrRepository.java:578) at org.modeshape.jcr.JcrRepository.login(JcrRepository.java:535) at org.modeshape.jcr.JcrRepository.login(JcrRepository.java:146) at org.modeshape.test.integration.RepositoryProvider.loginAndLogout(RepositoryProvider.java:86) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.jboss.as.ee.component.ManagedReferenceMethodInterceptorFactory$ManagedReferenceMethodInterceptor.processInvocation(ManagedReferenceMethodInterceptorFactory.java:72) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.invocation.InterceptorContext$Invocation.proceed(InterceptorContext.java:374) at org.jboss.as.weld.ejb.Jsr299BindingsInterceptor.doMethodInterception(Jsr299BindingsInterceptor.java:127) at org.jboss.as.weld.ejb.Jsr299BindingsInterceptor.processInvocation(Jsr299BindingsInterceptor.java:135) at org.jboss.as.ee.component.interceptors.UserInterceptorFactory$1.processInvocation(UserInterceptorFactory.java:36) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.invocation.WeavedInterceptor.processInvocation(WeavedInterceptor.java:53) at org.jboss.as.ee.component.interceptors.UserInterceptorFactory$1.processInvocation(UserInterceptorFactory.java:36) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.invocation.InterceptorContext$Invocation.proceed(InterceptorContext.java:374) at org.jboss.as.ejb3.concurrency.ContainerManagedConcurrencyInterceptor.processInvocation(ContainerManagedConcurrencyInterceptor.java:104) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.as.weld.ejb.EjbRequestScopeActivationInterceptor.processInvocation(EjbRequestScopeActivationInterceptor.java:82) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.invocation.InitialInterceptor.processInvocation(InitialInterceptor.java:21) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.invocation.ChainedInterceptor.processInvocation(ChainedInterceptor.java:61) at org.jboss.as.ee.component.interceptors.ComponentDispatcherInterceptor.processInvocation(ComponentDispatcherInterceptor.java:53) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.as.ejb3.component.singleton.SingletonComponentInstanceAssociationInterceptor.processInvocation(SingletonComponentInstanceAssociationInterceptor.java:53) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288) at org.jboss.as.ejb3.tx.CMTTxInterceptor.invokeInOurTx(CMTTxInterceptor.java:228) at org.jboss.as.ejb3.tx.CMTTxInterceptor.required(CMTTxInterceptor.java:304) at org.jboss.as.ejb3.tx.CMTTxInterceptor.processInvocation(CMTTxInterceptor.java:190) at org.jboss.invocation.InterceptorContext.proceed(InterceptorContext.java:288)	1	1	https://issues.jboss.org/browse/MODE-1516
15655	MODE-1493	Should this be changed in 3.0.x also?	0	0	https://issues.jboss.org/browse/MODE-1493
15656	MODE-1493	Assuming we'd be able to integrate the BOM, we'd probably do that only in 3.1 (master) so I don't think we need to worry about 3.0.x. However, the question is: do we still want 2.3.4 or 2.3.3 ?	0	0	https://issues.jboss.org/browse/MODE-1493

#	Requirement	Questions	Answered	Successfully Answered	URL
15657	MODE-1487	In addition to the VDB dependencies feature, the new VDB manifest (vdb.xml) includes another new feature. The VDB manifest now allows adding models to the VDB declaratively. In other words, for these models there will *not* be an "xmi" file in the VDB archive. These models are completely defined within the vdb.xml file by either (1) using DDL (for virtual models); or (2) using the "source" element to identify the connector (for source models). Since these declarative models do not have a resource associated with them, they will not be sequenced by the current model sequencer implementation. So any associated output nodes would have to be created by the VDB manifest sequencer. These declarative models will not have an "xmi:uuid" property. One solution would be to add a node type definition to the vdb.cnd that might look like this: // model name is the node name [vdb:declarativeModel] > mix:referenceable - mmcore:modelType (string) = 'PHYSICAL' autocreated < 'PHYSICAL','VIRTUAL' - vdb:visible (boolean) = 'true' autocreated - vdb:sourceTranslator (string) // only for PHYSICAL - vdb:sourceJndiName (string) // only for PHYSICAL - vdb:sourceName (string) // only for PHYSICAL - vdb:metadataType (string) = 'DDL' autocreated // only for VIRTUAL - vdb:modelDefinition (string) // only VIRTUAL All the above properties, except "vdb:metadataType" and "vdb:modelDefinition," are properties found in the existing "vdb:model" node type definition. However, since these declarative models will not have an "xmi:uuid" they cannot extend "vdb:model" since that definition requires the "xmi:uuid" property. Another solution might be to keep the common properties of declarative and XMI models in the "vdb:model" definition and create two separate definitions for declarative and XMI models. The VDB manifest's declarative models support Teiid's dynamic VDB feature but are not yet supported by Teiid Designer. Randall, should this issue address these declarative models? If so, would creating a node type definition like above be a good solution?	1	1	https://issues.jboss.org/browse/MODE-1487
15658	MODE-1487	{quote} Another solution might be to keep the common properties of declarative and XMI models in the "vdb:model" definition and create two separate definitions for declarative and XMI models. {quote} While I would prefer changing "{{vdb:model}}" to be a supertype for "{{vdb:declarativeModel}}" and "{{vdb:xmiModel}}", it does create a backward compatibility problem for anyone with a 2.x or 3.0 installation configured with a VDB sequencer. (The likelihood of 3.0 with a VDB sequencer is small, but a 2.x installation is more likely.) So perhaps we should move the common property definitions from the existing "{{vdb:model}}" to a new "{{vdb:abstractModel}}" that is referenceable, and to then make it the supertype of the existing "{{vdb:model}}" and a new "{{vdb:declarativeModel}}" node types. This isn't ideal, but it will give us backward compatibility. {quote} Randall, should this issue address these declarative models? If so, would creating a node type definition like above be a good solution? {quote} We can use this issue to track the support for declarative models, but we should probably rename the issue to describe that we're adding support for newer Teiid and Teiid Designer features (and enumerate them in the description). I don't see much point in artificially separating the work into separate issues -- using separate commits for logically-separate changes would still be useful, though. WDYT?	1	1	https://issues.jboss.org/browse/MODE-1487
15659	MODE-1487	The common VDB model properties can be moved into "vdb:abstractModel." I think making "vdb:abstracModel" a mixin not a node type definition would work. How does this look: [vdb:abstractModel] mixin - vdb:description (string) - vdb:visible (boolean) = 'true' autocreated - vdb:pathInVdb (string) - vdb:sourceTranslator (string) - vdb:sourceJndiName (string) - vdb:sourceName (string) // name is the node name (xmi models) [vdb:model] > vdb:abstractModel, mmcore:model - vdb:checksum (long) - vdb:builtIn (boolean) = 'false' autocreated + vdb:markers (vdb:markers) = vdb:markers copy // model name is the node name [vdb:declarativeModel] > nt:unstructured, mix:referenceable - mmcore:modelType (string) = 'PHYSICAL' autocreated < 'PHYSICAL','VIRTUAL' - vdb:metadataType (string) = 'DDL' autocreated - vdb:modelDefinition (string) As stated before, the declarative model cannot inherit from "mmcore:model" as it requires an "xmi:uuid." So the declarative model node type must also declare "mmcore:modelType" and "mix:referenceable."	1	1	https://issues.jboss.org/browse/MODE-1487
15660	MODE-1476	Are we renaming the repository as well ?	1	1	https://issues.jboss.org/browse/MODE-1476
15661	MODE-1472	Horia Chiorean <hchiorea@redhat.com> made a comment on [bug 814675|https://bugzilla.redhat.com/show_bug.cgi?id=814675] Where is the MissingMethodException generated from ? (this is not a java.lang exception and moreover, SLF4J 1.6.1 has the same API as 1.5.8 - used in SOA)	1	1	https://issues.jboss.org/browse/MODE-1472
15662	MODE-1472	Martin Vecera <mvecera@redhat.com> made a comment on [bug 814675|https://bugzilla.redhat.com/show_bug.cgi?id=814675] There are still other non-necessary dependencies in the library. What is preventing us from following the jbossall-client.jar approach? This is in its MANIFEST.MF: Class-Path: commons-logging.jar concurrent.jar ejb3-persistence.jar hi bernate-annotations.jar hibernate-commons-annotations.jar hibernate-c ore.jar hibernate-validator.jar hornetq-core-client.jar hornetq-jms-c lient.jar jboss-aop-client.jar jboss-appclient.jar jboss-aspect-jdk50 -client.jar jboss-client.jar jboss-common-core.jar jboss-deployers-cl ient-spi.jar jboss-deployers-client.jar jboss-deployers-core-spi.jar jboss-deployers-core.jar jboss-deployment.jar jboss-ejb3-common-clien t.jar jboss-ejb3-core-client.jar jboss-ejb3-ext-api.jar jboss-ejb3-pr oxy-clustered-client.jar jboss-ejb3-proxy-impl-client.jar jboss-ejb3- proxy-spi-client.jar jboss-ejb3-security-client.jar jboss-ejb3-vfs-sp i.jar jboss-ha-client.jar jboss-ha-legacy-client.jar jboss-iiop-clien t.jar jboss-integration.jar jboss-j2se.jar jboss-javaee.jar jboss-jsr 77-client.jar jboss-logging-jdk.jar jboss-logging-log4j.jar jboss-log ging-spi.jar jboss-mdr.jar jboss-messaging-client.jar jboss-profilese rvice-client.jar jboss-remoting.jar jboss-security-spi.jar jboss-seri alization.jar jboss-srp-client.jar jboss-system-client.jar jboss-syst em-jmx-client.jar jbosscx-client.jar jbossjts-integration.jar jbossjt s.jar jbosssx-as-client.jar jbosssx-client.jar jmx-client.jar jmx-inv oker-adaptor-client.jar jnp-client.jar netty.jar slf4j-api.jar slf4j- jboss-logging.jar xmlsec.jar	1	1	https://issues.jboss.org/browse/MODE-1472
15663	MODE-1472	Horia Chiorean <hchiorea@redhat.com> made a comment on [bug 814675|https://bugzilla.redhat.com/show_bug.cgi?id=814675] modeshape-client.jar should be used as a standalone client jar, not only inside AS but also outside (as a JDBC driver for example). Therefore, just adding the dependencies inside of MANIFEST.MF is not really an option. The logging dependency was optional, since logging works with just JUL logging. However, all the other packaged dependencies are needed to maintain the standalone functionality. What are the other dependencies which are conflicting atm (if you have stacktraces please post them as well). Thanks.	1	1	https://issues.jboss.org/browse/MODE-1472
15664	MODE-1462	Regarding the AS 7 subsystem: the code seems to be using MS logging & JBoss logging, but the pom.xml defines also dependencies to: log4j and commons-logging. Why are those dependencies needed ?	1	1	https://issues.jboss.org/browse/MODE-1462
15665	MODE-1461	Randall Hauch <rhauch@redhat.com> made a comment on [bug 813837|https://bugzilla.redhat.com/show_bug.cgi?id=813837] The last automated comment reflects that MODE-1461 now only deals with ModeShape 3.0 (e.g., 'master' branch), and not longer affects 2.x. Since there's no work to be done for 2.x, can we resolve this issue?	1	0	https://issues.jboss.org/browse/MODE-1461
15666	MODE-1459	It looks like this has been repeatedly de-prioritized. Do you know whether this will be actually worked on for 4.1?	1	0	https://issues.jboss.org/browse/MODE-1459

#	Requirement	Questions	Answered	Successfully Answered	URL
15667	MODE-1444	FAILURE! org.modeshape.jcr.value.binary.CassandraBinaryStoreTest Time elapsed: 0 sec <<< ERROR! java.lang.RuntimeException: java.net.SocketException: Address family not supported by protocol family: bind at org.apache.cassandra.net.MessagingService.getServerSocket(MessagingService.java:448) at org.apache.cassandra.net.MessagingService.listen(MessagingService.java:387) at org.apache.cassandra.service.StorageService.joinTokenRing(StorageService.java:533) at org.apache.cassandra.service.StorageService.initServer(StorageService.java:499) at org.apache.cassandra.service.StorageService.initServer(StorageService.java:397) at org.apache.cassandra.service.CassandraDaemon.setup(CassandraDaemon.java:309) at org.apache.cassandra.service.CassandraDaemon.init(CassandraDaemon.java:342) at org.apache.cassandra.service.EmbeddedCassandraService.start(EmbeddedCassandraService.java:52) at org.modeshape.jcr.value.binary.CassandraBinaryStoreTest.setUpClass(CassandraBinaryStoreTest.java:64) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:27) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:234) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:133) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:114) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:188) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:166) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:86) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:101) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:74) Caused by: java.net.SocketException: Address family not supported by protocol family: bind at sun.nio.ch.Net.bind(Native Method) at sun.nio.ch.ServerSocketChannelImpl.bind(ServerSocketChannelImpl.java:126) at sun.nio.ch.ServerSocketAdaptor.bind(ServerSocketAdaptor.java:59) at sun.nio.ch.ServerSocketAdaptor.bind(ServerSocketAdaptor.java:52) at org.apache.cassandra.net.MessagingService.getServerSocket(MessagingService.java:434) ... 30 more {code} 2)When running via IPV4 from my IDE: {code}	0	0	https://issues.jboss.org/browse/MODE-1444
15668	MODE-1444	Does Cassandra have support for variables inside their YAML file? If so, you could have one configuration file and use code to determine the variables. Might be easier; might not.	0	0	https://issues.jboss.org/browse/MODE-1444
15669	MODE-1440	Looks much better Randall... and all elements with defaults (as indicated in the schema) are not required. Correct?	1	1	https://issues.jboss.org/browse/MODE-1440
15670	MODE-1388	Targeting to 4.0. I'd really like to understand more than just the one use case. How is it useful to have some nodes disappear automatically after a period of time? If we have no new information, we'll probably close this as WONTFIX or REJECTED.	0	0	https://issues.jboss.org/browse/MODE-1388
15671	MODE-1365	Note that lots of {code} WARN HSEARCH000050: It appears changes are being pushed to the index out of a transaction. Register the IndexWorkFlushEventListener listener on flush to correctly manage Collections! {code} messages are being logged by Hibernate Search when we index nodes. See ISPN-1218 (which is marked as WONTFIX), but I think we need to address this.	0	0	https://issues.jboss.org/browse/MODE-1365
15672	MODE-1352	Isn't this a duplicate of MODE-1988 ?	1	0	https://issues.jboss.org/browse/MODE-1352
15673	MODE-1344	Can this be targeted to 3.1 (~Feb 2012), or does it need to be done by 3.0?	1	1	https://issues.jboss.org/browse/MODE-1344
15674	MODE-1344	That's going to be difficult to do and will involve some bigger enhancements to the sequencing framework in 2.7, since the 2.7 sequencing framework doesn't allow sequencers to create new node types. After all, the graph layer knows nothing about node types. It is possible, but just quite a bit of work (about 2-3 weeks total, compared with roughly 2-3 days in 3.0). An interim step might be to just be sure to read the extended properties and add them to the derived output, but not to create new node types. This means that the extended properties would appear in results (if included in the SELECT clause) but would not appear in the JDBC schema. Also, importing the models would not create the corresponding columns, but they could be modeled manually. Obviously this is less than ideal.	1	1	https://issues.jboss.org/browse/MODE-1344
15675	MODE-1310	Do we need to create an adapter layer for older sequencers, or should we just expect users to convert their custom sequencers?	1	1	https://issues.jboss.org/browse/MODE-1310
15676	MODE-1304	Since the PR has been merged, do we still need this open ?	1	1	https://issues.jboss.org/browse/MODE-1304
15677	MODE-1295	Updated code based on the above points as follows: 1. done 2. those tests are still excluded from modeshape 2.x. For 3.x. and JR all tests are being run 3. binary tests were added back, but will still be reported as failed until MODE-1308 is fixed 4. added warmup count setting. Default value of 5 5. done 6. added a sample report which generates a simple bar chart with the values for each test. We need to define exactly what other reports we want and how they should look. Also, modeshape-connector-benchmark it seems only does a println - nothing more - although there is some code commented out 7. TDB: where should I add this ? 8. that's definitively an option, and the reason why I avoided any modeshape specific dependencies in the api. I would say let's do it. 9. done - added a new test suite where anonymous sessions are created 10. done - added a new test suite where authenticated test session are created (however, the passed Credentials should match the ones with which the repo is configured)	1	1	https://issues.jboss.org/browse/MODE-1295
15678	MODE-1295	Is there a way to add the files in the new repo while preserving the history ?	1	1	https://issues.jboss.org/browse/MODE-1295
15679	MODE-1295	Merged the pull-request into the 'master' branch of the new 'modeshape-performance' repository. One last thing: can you update the 'README.md' file with a description of this framework and how to use it? For example, what each module represents, how to add other impls and/or versions, what commands to run, and where the results are.	1	1	https://issues.jboss.org/browse/MODE-1295
15680	MODE-1291	MODE-1581 recently added a 'getAllBinaryKeys()' method to the BinaryStore interface (needed for the backup and restore functionality). When I merged these changes, I simply stubbed out this method. [~tfromm], can you implement this method? Please see the new InfinispanUtil class for a utility method to find all the keys for a Cache (which should be enhanced with your technique of using the CacheLoader directly to get the keys for a local, persisted cache).	1	1	https://issues.jboss.org/browse/MODE-1291

#	Requirement	Questions	Answered	Successfully Answered	URL
15681	MODE-1286	/commit/e2e665fab36a1d35f1399e95433229a0ba27623f at org.modeshape.connector.git.GitConnector.getDocumentById(GitConnector.java:228) at org.modeshape.jcr.federation.FederatedDocumentStore.get(FederatedDocumentStore.java:230) at org.modeshape.jcr.cache.document.WorkspaceCache.documentFor(WorkspaceCache.java:142) at org.modeshape.jcr.cache.document.WorkspaceCache.documentFor(WorkspaceCache.java:155) at org.modeshape.jcr.cache.document.WorkspaceCache.getNode(WorkspaceCache.java:178) at org.modeshape.jcr.cache.document.AbstractSessionCache.getNode(AbstractSessionCache.java:171) at org.modeshape.jcr.cache.document.WritableSessionCache.getNode(WritableSessionCache.java:144) at org.modeshape.jcr.JcrSession.node(JcrSession.java:412) at org.modeshape.jcr.AbstractJcrNode$ChildNodeResolver.nodeFrom(AbstractJcrNode.java:3396) at org.modeshape.jcr.JcrChildNodeIterator.nextNode(JcrChildNodeIterator.java:113) at org.modeshape.jcr.AbstractJcrRepositoryTest.navigate(AbstractJcrRepositoryTest.java:292) at org.modeshape.connector.git.GitConnectorTest.shouldReadCommitSubgraph(GitConnectorTest.java:113) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:47) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:50) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:222) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.junit.runner.JUnitCore.run(JUnitCore.java:157) at com.intellij.junit4.JUnit4IdeaTestRunner.startRunnerWithArgs(JUnit4IdeaTestRunner.java:76) at com.intellij.rt.execution.junit.JUnitStarter.prepareStreamsAndStart(JUnitStarter.java:195) at com.intellij.rt.execution.junit.JUnitStarter.main(JUnitStarter.java:63) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at	1	1	https://issues.jboss.org/browse/MODE-1286
15682	MODE-1285	Can you explain why this would be helpful?	1	1	https://issues.jboss.org/browse/MODE-1285
15683	MODE-1252	Not sure I understand why REPOSITORY_NAME should be deprecated in 2.x: we could just change getRepository(Map parameters) to do an additional search in the Map for the REPOSITORY_NAME key and if that is present, use it instead of parsing the URL. If it's not present, just default to what the current code is. For 3.x, do we still need the JcrRepositoryFactory#getRepository(String jcrUrl, String repositoryName) method ?	1	1	https://issues.jboss.org/browse/MODE-1252
15684	MODE-1233	See MODE-1231 for the fix in the branch used in the 5.2 product branch.	1	1	https://issues.jboss.org/browse/MODE-1233
15685	MODE-1233	Can the fix not work for both branches?	1	0	https://issues.jboss.org/browse/MODE-1233
15686	MODE-1231	Is this only a problem when the sequence is being created?	1	0	https://issues.jboss.org/browse/MODE-1231
15687	MODE-1226	Perhaps the largest issue I can think of is deciding what the semantics should be with respect to children. Initially, having the client application be responsible for manually detaching all nodes they care about seems to make sense. But if a client detaches a parent node and all of its child nodes, would the detached parent node be able to find the child nodes via the 'getNodes()' and related methods?	0	0	https://issues.jboss.org/browse/MODE-1226
15688	MODE-1218	I'm not sure I really see the advantage of the {{ant start.demo}} or {{ant stop.demo}}. Seems really strange to have a script just to start the server _and_ Eclipse. Is this issue even needed anymore?	0	0	https://issues.jboss.org/browse/MODE-1218
15689	MODE-1205	Could you please clarify the difference between this and MODE-1107? Thanks!	0	0	https://issues.jboss.org/browse/MODE-1205
15690	MODE-1187	This feature is now supported in ModeShape 3.0 with the new JCR design. In fact, the nodes within the '/jcr:system' branch are already using non-UUID identifiers.	0	0	https://issues.jboss.org/browse/MODE-1187
15691	MODE-1157	We're approaching our Beta3 release date (hopefully late next week), and I'd like to see this merged into the 'master' branch. But frankly, I'm concerned about this feature getting in without simultaneously updating the Reference Guide. I have no doubt that your code review is good. IMO, once a feature is implemented, the unit/integration tests help us catch regressions in behavior, and without tests I'm just not confident that we won't mistakenly break this feature in the future when working on this class (and the Hibernate configuration code). Is there a reason we can set up an integration test to instantiate the JBoss Transaction Manager, register a JTA DataSource (e.g., HSQLDB), run our integration tests? See http://community.jboss.org/message/592630#592630 for the Maven dependencies used in Infinispan tests.	0	0	https://issues.jboss.org/browse/MODE-1157

#	Requirement	Questions	Answered	Successfully Answered	URL
15692	MODE-1157	Let me apologize up front for being so difficult about this change. My concern is that if we merge these changes into 'master' without the corresponding changes to make ModeShape JTA compatible (MODE-1123), we will actually cause some strange and IMO wrong behavior when ModeShape is used within (or with) user transactions. IIUC, prior to these changes it was not possible for ModeShape to use a JDBC DataSource that was set up as an XA resource. With these changes, it will be possible to do this, allowing the client (or more strictly speaking the JTA transaction) to dictate when changes are persisted in the database using JTA user transactions. For example, a client could do the following (or something equivalent within an EJB-land where the container manages transaction boundaries): # being the user transaction # make changes to a workspace using a session # save the session # commit the user transaction The crux of my concern is that ModeShape still thinks that the information is being saved to the store in #3 above, and will fire the change events. However, because the information is actually not persisted in the store until #4, we're firing events too early and we're updating the indexes too early. This might not be a (too much of a) problem in the nominal case, but actually can cause a lot of problems in two other situations. The first is when there are multiple saves on a single session within a single user transaction: # being the user transaction # make changes to a workspace using a session # save the session # make more changes to a workspace using a session # save the session # commit the user transaction At the end of Step 3, ModeShape will fire events (to all observers) that reflect the changes made in step 2. Then, at the end of Step 5, ModeShape will fire events (to all observers) that reflect the changes made in step 4. However, none of these changes were persisted in the store, and none will actually be visible to any other sessions. And because our query indexes are updated based upon events, we'll actually update our query indexes at step 3 and 5 above. Again, this leaks information out of the transaction, making it visible to other sessions and actually exposing incorrect data. Then what happens if the user transaction is rolled back rather than committed? # being the user transaction # make changes to a workspace using a session # save the session # **rollback the user transaction** At this point, after step 3 completes, ModeShape will fire the events and will (asynchronously) update the query indexes to reflect the changes made in step 2 and saved in step 3. But if the transaction is rolled back in step 4, the changes were not actually persisted in the store and thus the fired events and the query indexes are all flat out wrong. So if the above is accurate (and there's a fair possibility it is not), then it actually could be very problematic to merge these changes into 'master' without also making the rest of the changes for JTA support (MODE-1123). I'd be fine doing that, if the rest of the changes were going to happen in the same release. But it doesn't look like the rest of the MODE-1123 changes will be ready, as we're planning to release 2.5.0.Beta3 very soon (next week) and 2.5.0.Final very quickly after that (e.g., a week later). Please let me know if I'm missing something.	0	0	https://issues.jboss.org/browse/MODE-1157
15693	MODE-1150	BTW Randall... since it's JBoss.. do you plan to switch to JBoss Logging or stick to SLF4J ?	1	1	https://issues.jboss.org/browse/MODE-1150
15694	MODE-1132	Is this still targeted for 2.5.1.GA?	0	0	https://issues.jboss.org/browse/MODE-1132
15695	MODE-1122	Posted a pull request at https://github.com/ModeShape/modeshape/pull/64 that includes a proposed code change. If the change looks good, I can submit a new patch that also updates the reference guide to describe this behavior. Was the intent to also do this for JaasCredentials and SecurityContextCredentials?	1	1	https://issues.jboss.org/browse/MODE-1122
15696	MODE-1122	I think this should work the same way for all credentials. Do you think this behavior should be configurable? On one hand, it seems useful to always get anonymous access to the content if authentication fails. But on the other hand, you'll never know that your credentials were invalid, so an application that expects a user to log in will not be able to enforce authentication without completely disabling anonymous logins.	1	1	https://issues.jboss.org/browse/MODE-1122
15697	MODE-1116	We'll have to decide which of these should be made into OSGi bundles. All of them? Do the documentation modules really need to be OSGi-enabled?	1	1	https://issues.jboss.org/browse/MODE-1116
15698	MODE-1116	Do you want to create a pull-request from your GitHub fork and set the URL on this issue?	1	1	https://issues.jboss.org/browse/MODE-1116
15699	MODE-1098	Are these environment or system properties?	1	1	https://issues.jboss.org/browse/MODE-1098
15700	MODE-1091	This should be a pretty quick and easy fix. Targeting for 2.5.0.Beta1. Dmitri, would you mind specifying exactly what you changed with the Hibernate properties to specify the schema? Is it just the "hibernate.default_schema" property?	0	0	https://issues.jboss.org/browse/MODE-1091
15701	MODE-1081	What if we add those JARs in Nexus with an "unsigned" classifier ? Then ModeShape could use this classifier. What do you think ?	1	1	https://issues.jboss.org/browse/MODE-1081
15702	MODE-1075	I think that the MODE-456 fix addresses this, although it would require you to make your sequencer output path the same (or a subpath of) your sequencer input path. WDYT?	1	0	https://issues.jboss.org/browse/MODE-1075
15703	MODE-1075	Hey, I implemented sequencing as it was meant to be, different node for sequencing output, but now I find it complicates things and I rather have the output properties on the sequenced file. I'm not sure what the drawbacks could be though. Is there a resolution for this feature request ?	0	0	https://issues.jboss.org/browse/MODE-1075
15704	MODE-1051	Thanks, Edouard. Can you clarify for me what the problem or error actually is? The text above seems to imply that ModeShape does indeed make '/A/X' appear in the result set for your query, but '/B/X' does not. Am I reading that correctly? Is your test case failing, and if so on which line? (If you are testing with Jackrabbit, it'd be very interesting to note where your test case fails when running against Jackrabbit.)	1	1	https://issues.jboss.org/browse/MODE-1051
15705	MODE-1051	So the first query works (when using ModeShape) and the results are what the test case expects, but the second query does not work for ModeShape. Is that correct? As we've discussed in the discussion thread (see link above), I believe that this behavior matches the JCR 2.0 specification. Even though the specification gives a fair amount of room for implementations (and doesn't fully specify the expected behavior in a case like this), IMO the shareable nodes feature as described in the specification really isn't consistent with the query feature. Having said, to date we've chosen to implement ModeShape such that the shared nodes (created by the cloning process within a single workspace) are not visible to the query engine. Only doing this can ModeShape be consistent with the shareable nodes and query features as described in the specification. (Yes, this is in part due to the implementation details of these features, but that's what we have to work with at this time.)	0	0	https://issues.jboss.org/browse/MODE-1051
15706	MODE-1033	I'm wondering if the StreamSequencerAdapter should, as it looks at the existing content (to find the appropriate SNS index for the top-level nodes generated by the sequencing operation) should also look for and remove existing content in the output location that was also (previously) derived from the same node currently being sequenced. In other words, the new derived output would **replace* the content previously-derived from the same input. If so, should this always be done, or should this be driven by the specific sequencer configuration? (ATM, I'm leaning toward always doing this.)	1	1	https://issues.jboss.org/browse/MODE-1033
15707	MODE-1028	would it be possible to link the node back to the sequenced nodes via a correlation of related UUIDs?	1	1	https://issues.jboss.org/browse/MODE-1028

#	Requirement	Questions	Answered	Successfully Answered	URL
15708	MODE-1028	MODE-1033 does this indirectly via a PATH property on the derived (sequenced) information. We could store the UUID as a STRING property, but in order to use a REFERENCE (on the derived content or the sequenced file) we'd need one of them to be 'mix:referenceable'. I originally thought that when a file is deleted, we could easily query for 'mix:derived' with a 'mix:derivedFrom' property with a path matching the 'nt:file' node being deleted. Technically, this is very easy to do. The harder part, IMO, is the broader issue of where this logic should live? Should it live in the REST client (in the unpublish method)? Should it live in the JCR engine? Should it be implemented as an event listener (observer) that is added to the engine, perhaps defined/added in the configuration?	0	0	https://issues.jboss.org/browse/MODE-1028
15709	MODE-1027	should this maybe be incorporated in to MODE-888?	1	1	https://issues.jboss.org/browse/MODE-1027
15710	MODE-1019	My guess that it is easier to start of having it all in one file (inline). Is it hard to support both?	1	1	https://issues.jboss.org/browse/MODE-1019
15711	MODE-1011	Having one last difference, the DataBaseMetadata returned for getColumnType() is 'Name', but for the same column in the ResultSetMetaData getColumnType(), it returns 'String'. I noticed in the NodeTypeSchemata: types.put(PropertyType.NAME, typeSystem.getStringFactory().getTypeName()); For column types of 'Path', it's returned in both cases, and has a typeSystem.getPathFactory(). Should there be a typeSystem.getNameFactory() for 'Name' type and be return the 'Name' type in query results?	1	1	https://issues.jboss.org/browse/MODE-1011
15712	MODE-1011	TypeSystem does have this method: TypeFactory<?> getPathFactory(); Should that be removed and change the NodeTypeSchemata to use a string factory instead?	1	1	https://issues.jboss.org/browse/MODE-1011
15713	MODE-999	I thought the SLF4J JAR was to be excluded because it is already available on the global classpath. Is that true, Van?	1	1	https://issues.jboss.org/browse/MODE-999
15714	MODE-995	The properties shown represent all available descriptors for the JCR repository. Source is not exposed via a descriptor. Should source be added?	0	0	https://issues.jboss.org/browse/MODE-995
15715	MODE-943	[~rhauch]: this issue is almost 4 years old and I can barely find any information online about {{JSOP}}. Is there any benefit in keeping this open ? {{JSOP}} seems quite "dead" to me (and if it's not dead, then I don't really see it as a wide-spread standard)	0	0	https://issues.jboss.org/browse/MODE-943
15716	MODE-940	Should this read: implement SVN connector for 3.x ?	1	0	https://issues.jboss.org/browse/MODE-940
15717	MODE-939	The system workspace is not exposed through the JCR Workspace.getAvailableWorkspaceNames(), nor is it included in the list returned by the REST service. (This is true even when the SYSTEM_SOURCE_NAME is set to a workspace in the same source that the repository uses.)	0	0	https://issues.jboss.org/browse/MODE-939
15718	MODE-930	Changed the title to remove the notion of a User Defined Function (UDF), which I had originally suggested because I thought the use case just involved changing the results of a query. There are a couple of things to consider regarding how this might be implemented. First and foremost is whether the property is a "virtual" property or a "transformed" property. In my mind, a "virtual" property does not really exist, but is always created by using other properties of the node with, perhaps, external data. In other words, it looks like the "virtual" property exists in the persisted content, but it actually does not. A "transformed" property, on the other hand, probably does exist in the persisted content, but before it is returned it is transiently altered. Theoretically, this transformation might not require any other input, but it also may require external resources and/or other properties on the node. For example, the PropertyFactory approach might work for certain "transformed" properties that require the persisted values and optionally some external resource. However, since the PropertyFactory methods do not know about the node for which the property is being created, the factory could not possibly access other properties of the node, and thus would not work for the "virtual" case. Second, how should the resulting property values be treated within queries and search? Should they be excluded (non-queryable in property definition terms)? If it should be included, then should the underlying value or resulting value be indexed? Third, should be done at the lower-level graph API or within the JCR implementation layer. The graph API is far simpler, and does not address things like versioning. Plus, there really isn't a single concept in the graph API that might work for both "virtual" and "transformed" properties. As a result, my initial thought is that this might be better suited as a formal mechanism within the JCR implementation layer: it's higher-level; a single simple framework can likely handle all variations of "virtual" and "transformed" with a single mechanism; it can take advantage of but is not limited to the graph API (e.g., a custom PropertyFactory); and the result would likely be more straightforward and easier to understand. Thoughts?	1	0	https://issues.jboss.org/browse/MODE-930
15719	MODE-930	I didn't realize that the PropertyFactories don't know about the node for which the property is being created. This indeed defeats the purpose. (1) About the difference between virtual and transformed properties: Do you see implementation differences between the two? To me a virtual property might be the same as a readonly (null) transformed property.But I do think that the property itself should not contain any persisted data as you cannot ask for the 'raw' (untransformed) source. The only way to parameterize (is that english?) the transformation is to set other properties in the repository, possibly from the same parent node.So 'virtual' properties is probably the best name. (2) Considering the previous I suggest that these nodes are not indexed. (3) Looking at the JCR layer, the SessionCache is the best place to hook the property value handling at the JCR level?	0	0	https://issues.jboss.org/browse/MODE-930
15720	MODE-930	Should this be a blocker ?	1	1	https://issues.jboss.org/browse/MODE-930
15721	MODE-930	[~hchiorean], I don't think it should be a blocker. This kind of functionality just isn't defined in JSR-283. As useful as it would be, I think a lot of other features are higher priority. Do you see this as being useful for another required feature?	1	0	https://issues.jboss.org/browse/MODE-930
15722	MODE-916	Is this going to be committed to 2.2.x branch?	0	0	https://issues.jboss.org/browse/MODE-916
15723	MODE-899	Was this rolled back?	1	1	https://issues.jboss.org/browse/MODE-899
15724	MODE-899	Actually, I just saw that this was not made to the 2.2.x branch. Should it?	1	1	https://issues.jboss.org/browse/MODE-899
15725	MODE-888	Noticed IRestClient is deprecated, and from the comment, has been obsolete since 3.1. Should ModeShape Tools 3.3 be written to remove the IRestClient dependency and use its replacement?	1	1	https://issues.jboss.org/browse/MODE-888
15726	MODE-888	In that case is there any point in keeping MODE-1652 open ? Any new "REST client" would basically mean re-creating the module which we're about to delete here....	1	1	https://issues.jboss.org/browse/MODE-888
15727	MODE-887	Given that you can have more than one path for a particular result set if selectors are used, what is the suggested name for the property (or properties) containing the uri(s)? What about mode:uri for the default path and mode:uri-<selector name> if more than one selector was specified?	0	0	https://issues.jboss.org/browse/MODE-887

#	Requirement	Questions	Answered	Successfully Answered	URL
15728	MODE-887	Actually, I've recently seen more examples of using the "link" element in HTML/XML: {code} <link rel="node" href="http://example.com/resources/{repositoryName}/{workspaceName}/item/{path}"/> {code} See http://stackoverflow.com/questions/1271962/creating-hypermedia-links-in-a-custom-media-type and http://www.infoq.com/articles/subbu-allamaraju-rest Now, we're talking about JSON, so we can do something similar on each node, and it could be as simple as: {code} { "node" : "http://example.com/resources/{repositoryName}/{workspaceName}/item/{path}" } {code} That only includes the "view" URL and is implicitly a GET, but we could go so far as to explicitly define the "view", "edit" and "delete" URLs, too. {code} { "link" : [{ "type" : "application/json" "method" : "GET" "rel" : "view" "href" : "http://example.com/resources/{repositoryName}/{workspaceName}/item/{path}" } , { "type" : "application/json" "method" : "POST" "rel" : "edit" "href" : "http://example.com/resources/{repositoryName}/{workspaceName}/item/{path}" } , { "type" : "application/json" "method" : "DELETE" "rel" : "delete" "href" : "http://example.com/resources/{repositoryName}/{workspaceName}/item/{path}" }] } {code} I kinda merged some ideas from http://beautifulbeta.wikidot.com/json-feeds and Rails (see https://github.com/asplake/described_routes) Is that gross?	1	1	https://issues.jboss.org/browse/MODE-887
15729	MODE-887	Okay, regarding your questions: {quote} Given that you can have more than one path for a particular result set if selectors are used, what is the suggested name for the property (or properties) containing the uri(s)? What about mode:uri for the default path and mode:uri-<selector name> if more than one selector was specified? {quote} Yeah, that's interesting. That'd mean that certain columns go to one node, while other columns go to a different node. (There ca be more than just two, also.) I'll have to think about that more.	1	1	https://issues.jboss.org/browse/MODE-887
15730	MODE-881	Randall, can you try running this in your QA/Test environment? Its mostly working but one failed test, which I'm trying to determine is my pc or not.	0	0	https://issues.jboss.org/browse/MODE-881
15731	MODE-861	The patch looks great, although I can conceive of scenarios where this behavior would be seen as an information leak (e.g., if users' home directories were each mapped to a workspace). Could you please add an option that disables this behavior and then default it to false so that this behavior is exposed by default? Thanks!	1	1	https://issues.jboss.org/browse/MODE-861
15732	MODE-860	Wouldn't it be better to have this in the Teiid JOPR plugin, dedicated for finding VDBs that can be deployed? The plugin still could issue the queries using the JCR repository (via the JCR API, without needing any compile-time dependency on ModeShape) or using the RESTful service. Essentially, enter some criteria about the VDB name (wildcard), version (int range), uploaded date (date range), path (wildcard string), and/or description (full-text search or wildcard string), and it could return a table showing: VDB name (ideally hyperlinked), version, description, and path.	1	1	https://issues.jboss.org/browse/MODE-860
15733	MODE-827	I original developed a new JackrabbitXmlNodeTypeSource class that imported the XML content into a transient graph, read that Jackrabbit-oriented graph content and constructed an equivalent ModeShape CND graph structure in another transient graph that was then returned by the JcrNodeTypeSource.getNodeTypes():Graph method. I created several unit tests that used some XML files provided by Jozef, and got everything working. Total time was about 2 hours. But then I started thinking about it, and the inability to easily register node types via the NodeTypeManager has been bothering me for a while. The JcrNodeTypeSource mechanism isn't really very good, because it requires casting to the JcrNodeTypeManager implementation and using a proprietary method. What I've been trying to figure out is whether it's possible to create some loaders that basically created NodeTypeTemplate instances, and then allowed using the standard methods to register those new templates. This afternoon I took the opportunity to create a prototype of a new approach. I created a JcrNodeTypeFactory interface that defines a method to get an array of NodeTypeDefinition instances (remember that the template interfaces are subtypes of the definition interfaces). I then created the CndNodeTypeFactory and JackrabbitXmlNodeTypeFactory implementation classes. Both read their respective formats, and load the content into a graph with ModeShape's CND structure. Because they both operate in a similar way (the only difference is reading the file content and constructing the graph content; but processing the graph content to create the NodeTypeDefinition instances is common, as are a slew of loading methods), I pulled all of that common functionality into a new GraphNodeTypeFactory abstract base class. This new mechanism follows this pattern: CndNodeTypeFactory factory = new CndNodeTypeFactory(session); factory.load(file); // or stream or resource file if (!nodeTypeSource.getProblems().isEmpty()) { // Report problems } else { boolean allowUpdate = false; session.getWorkspace().getNodeTypeManager().registerNodeTypes(factory.getNodeTypes(), allowUpdate); } Note that there are actually a half-dozen "load(...)" methods that handle URLs, resource files, input streams, raw content, and even reading from a graph or subgraph. And the JackrabbitXmlNodeTypeFactory approach is almost identical - other than the class name and not quite as many "load(...)" methods. Yes, this still requires instantiating a ModeShape-specific class, but no other objects need to be cast from the standard JCR API interfaces. So I think this is a little cleaner. The CndNodeTypeFactory and JackrabbitXmlNodeTypeFactory both update the namespaces with those used by the node type definitions within the files, and they even support using different prefixes in the files than what may already be defined in the session. (Of course any new namespace used in the loaded node type definitions will be registered in the workspace using the files' namespace prefixes. And because they're registered with the workspace and not the session, those namespace registrations are available to the rest of the repository users.) Of course, other implementations will need to make sure they're registering new namespaces with the NamespaceRegistry in the session's workspace. Otherwise, the process of registering the node type definitions using the session's node type manager will result in errors due the inability to parse a prefixed name without a namespace registered for that prefix. Incidentally, I did discover a problem with importing XML that contains something like same-name-siblings. The XmlHandler was not keeping track of the SNS indexes when it was building the paths. (This is different than the JCR import, which does work fine.) A few changes to the XmlHandler class and everything works fine. I'm uploading a patch with these changes and asking for feedback on this approach. Is this better, worse, or no different than the older JcrNodeTypeSource approach? Is it worth these changes? (I also didn't want to commit this now, as Brian is working on a fix for MODE-810.)	1	1	https://issues.jboss.org/browse/MODE-827
15734	MODE-827	Guys, could you let me know ASAP (let's say this week) what is your decision regarding subject of this issue? I have to finish the work on Shaped Magnolia, so I need to know how to write (and what MS methods to use in) the ProviderImpl.registerNodeTypes(InputStream) method.	1	1	https://issues.jboss.org/browse/MODE-827

#	Requirement	Questions	Answered	Successfully Answered	URL
15735	MODE-824	I attached a patch (MODE-824_locks.patch) that's a first cut and providing the lock-related functionality in ManagedRepository. The existing ManagedLock class required some fields that were manifestly useful (e.g., expiration date, owning session) that did not exist in the ModeShapeLock class in the JCR layer. Since ModeShapeLock was intended to be valid over a long lifetime and some lock-related fields (e.g., expiration date, owning session) could change constantly and asynchronously, ModeShapeLock only held the immutable fields for the lock. Rather than try to hack these volatile fields into MSL, the patch adds a JcrLockSnapshot class and added a method to RepositoryLockManager to read all locks directly from the graph as a snapshot of the current lock state. Since the ManagedEngine doesn't connect to the repository through a JCR session, it seems essential that the lock snapshots be populated from a direct graph read anyway. Other than needing some Javadoc and a blessing, I am concerned that the patch adds the removeLock(String) method to JcrRepository, but doesn't provide any security checks in that method. Will the ManagedEngine run in some JAAS context that will allow for a security check? For that matter, the JcrRepository.getAllLocks() method also runs without a security check and can leak information (e.g., node UUIDs, user names, workspace names). Am I being paranoid?	0	0	https://issues.jboss.org/browse/MODE-824
15736	MODE-818	Isn't there a workaround for this? I thought the DDL sequencer, when processing this view, will already create a single AST for the whole view (including the definition expression). Is that not the case?	0	0	https://issues.jboss.org/browse/MODE-818
15737	MODE-813	Fantastic! Does this include the WebDAV WAR, too?	1	1	https://issues.jboss.org/browse/MODE-813
15738	MODE-813	Does that mean the user will need to crack open the JAR to modify the users/roles properties? A number of the other components place their users/roles properties files in server/{profile}/conf/props. Not sure if this is standard or convention but it's easier to access/modify.	0	0	https://issues.jboss.org/browse/MODE-813
15739	MODE-813	Can we change the managedConfigRepository.xml to not use the ModeShape URL as the default namespace? I think it's awkward that the repository name is "mode:repository".	1	1	https://issues.jboss.org/browse/MODE-813
15740	MODE-810	The patch was committed, so is there anything left to do on this issue?	1	1	https://issues.jboss.org/browse/MODE-810
15741	MODE-806	I believe the most RESTful style would be to define media types for each query language and each response format. Embedding the query language in the URL does seem easier on the surface (at least when thinking about interacting with it from a browser), but I think that's the wrong approach and doesn't fit in the REST style. I agree with your proposal to return a JSON-encoded representation of the QueryResult, since all of our query languages return that same object (with slightly differing columns for certain languages). Thus we probably need only a single response format MIME type ("application/jcr+query+results" works for me). To be consistent with the rest of the RESTful service's responses, the response format could be JSON with a container for the JSON fragments for each row in the QueryResult (including the URL for the node given by the Row's location). It's also be nice to have the query included. The suggestion to use limits and offsets is a very good one. This can be done with the JCR API without mucking with the query, by using Query.setLimit(long) and Query.setOffset(long). As for the media types for the query languages, I'd suggest something like "application/jcr+sql", "application/jcr+sql2", "application/jcr+xpath", and "application/jcr+search". Then a "doQuery(...)" method for each of these (which probably delegates to a single helper method that accepts the query language as a parameter). Does that work? WDYT?	1	1	https://issues.jboss.org/browse/MODE-806
15742	MODE-806	I briefly looked at the patch, and it looks great. This is fantastic functionality, and has been implemented well. Thanks! Do you have time to update the Reference Guide with the text in the previous comment and an example of a request and a (short) response?	1	1	https://issues.jboss.org/browse/MODE-806
15743	MODE-789	Can Chapter 4 Running the example applications, be renamed to something like: Example Application Setup? since you don't really run it there, and in 4.2 it indicates the next section is where the examples will be run.	1	1	https://issues.jboss.org/browse/MODE-789
15744	MODE-787	JCR-2663 and JCR-2648 are fixed, does it make sense to update these now or wait until all 4 are fixed?	1	0	https://issues.jboss.org/browse/MODE-787
15745	MODE-770	Overall, this is very good. A few comments: 1) The Getting Started guide talks about obtaining the RepositoryFactory instance by simply instantiating the JcrRepositoryFactory. While this is definitely allowed by the JCR 2 specification (and is in fact the minimum required supported behavior), the idiomatic JCR 2 approach seems to be to use the ServiceLocator pattern. And while that pattern does take more code, I think we should use it in the examples. The Reference Guide should probably start out by describing the ServiceLocator approach, but should also document simply instantiating the JcrRepositoryFactory. 2) The JcrRepositoryFactory methods are not all synchronized correctly. The "getRepository(...)" instance method is not synchronized, while the two "shutdown(...)" instance methods are. And the "getEngineFromConfigFile(...)" helper method is static and synchronized, making it synchronized on the class not the instance. Per IRC discussion, the helper method should be made an instance method, and all should be well. 3) I might suggest documenting in JcrRepositoryFactory that it supports the ServiceLocator pattern, including the reference to that class in the "META-INF/services/javax.jcr.RepositoryFactory", plus a simple code example that shows how to use it. (Yes, this is in the Getting Started and Reference Guides, but it's always good to have it in the JavaDoc.) Do these make sense? Should be committable once these are addressed.	1	1	https://issues.jboss.org/browse/MODE-770
15746	MODE-758	A separate RepositorySource implementation in the same project is probably the easiest, since it will have very different JavaBean properties than the existing source class. It might be possible to have a base class do most of the work, and the two subclasses simply define the JavaBean properties and the mechanism for creating the CacheContainer. Does that make sense?	1	1	https://issues.jboss.org/browse/MODE-758
15747	MODE-755	Attached patch that takes a first cut at a path-based repository that works with the new base repository. The patch also ports the FileSystem connector to make sure that the repository works in the "real world". There are still a few todos: 1. Fix the comments - they're all over the map and often wrong 2. Sort out exception handling in the file system connector (see below) 3. Review MapTransaction/PathTransaction to see if they have any common code that can be pushed into BaseTransaction. And some questions: 1. What's the new best practice for throwing exceptions that only occur during the commit? Right now they get swallowed in Connection. 2. Should we keep the prepare/commit framework in place or delete it for now under the principle of YAGNI? Any other feedback would be appreciated as well. Thanks!	0	0	https://issues.jboss.org/browse/MODE-755
15748	MODE-716	Attached a rough patch that implements the RepositoryProvider interface and attempts to lookup a Repositories object at a JNDI name provided through a servlet init parameter. If the object bound to the name is a mere Repository, this provider wraps it in a Repositories object. Is this what you had in mind? This code is completely untested and probably doesn't quite work, but I'd like to get some feedback before I clean it up.	1	0	https://issues.jboss.org/browse/MODE-716

#	Requirement	Questions	Answered	Successfully Answered	URL
15749	MODE-716	What if we reworked my initial patch so that the JNDI name of the repository could be injected more easily? That way, any JCR implementation that supported JNDI binding (which should be all of them, since the JCR TCK test gets its Repository reference through JNDI) could be used with this new RepositoryProvider. If we make the RepositoryFactory JNDI-bindable, then we still need a JndiRepositoryProvider for users who want to share their ModeShape repo across multiple applications. If we made the RepositoryFactory the JNDI binding point, then you might end up with one JNDI binding for the JcrEngine and another for the RepositoryFactory in that model. I may be missing something here though. What do you think about this approach? Sorry for the delay in responding.	1	0	https://issues.jboss.org/browse/MODE-716
15750	MODE-711	I thought this was available in 1.1.0Final as the "cacheConfigurationName" property on the connector as outlined in the reference document. Am I missing something?	0	0	https://issues.jboss.org/browse/MODE-711
15751	MODE-681	Should each type of jdbc driver (i.e., jndi, http, and file) be packaged separately or make one jar?	1	1	https://issues.jboss.org/browse/MODE-681
15752	MODE-681	Well, I take that back. I guess it depends on how big the JAR file is with all of the dependencies, and how that compares to other JDBC drivers. I don't think the JAR file will be too unreasonable, but if it is too big then we could package the JNDI and remote options as one JAR and the local option as another (larger) JAR. WDYT? How big is the all-encompassing JAR?	1	1	https://issues.jboss.org/browse/MODE-681
15753	MODE-681	Yikes. I guess the "file" option is all ModeShape JARs? I can see four options (perhaps there are others): 1) Make a single (monster) JAR. This doesn't seem very attractive for those that want to use the 'jndi' or 'http' mechanisms. 2) Make a single smaller JAR that contains everything needed for the 'jndi' and 'http' mechanisms, but that would not include the rest of the ModeShape JARs. If you want to use the 'file' mechanism, you'd be responsible for ensuring the ModeShape JARs are on the classpath. We could make this a bit easier if we provided a single, all-encompassing JAR for the ModeShape JCR, but at the moment it would entail using the 'modeshape-jcr-with-dependencies.jar' and any of the sequencer and connector '-with-dependency.jar' files needed by the particular config. That's not too bad of a workaround, IMO. A disadvantage of this option is that the driver still accepts the 'file' URLs even though not all classes might be there. 3) Make separate JARs: one for the 'jndi' and 'http' mechanisms, and a monster for the 'file' mechanism. 4) Split the JDBC module into two Maven projects/modules: one for the conventional 'remote' (e.g., 'jndi' and 'http' mechanisms), and a second module that relied upon the first and added the 'file' mechanism. This may not be practical, depending upon how easily it would be to subclass and override/add the file feature. But if it can be easily done, this may be an attractive option. This also makes each module fairly conventional, and the dependencies of each would match what they really are. An advantage of this option compared with option 3 is that this would one different Driver class that would accept the 'jndi' and 'http' URLs, and a different 'embedded' Driver class that would accept 'jndi', 'http', and 'file' URLs. The first is not practical, and IMO option 2 is much better than option 1. Option 3 and 4 are pretty similar, but use a different approach for making the separate JAR. Of 3 and 4, I think I like 4 better. So I think the choice is between 2 and 4. What's the practicality of option 4, given the current JDBC driver design?	1	1	https://issues.jboss.org/browse/MODE-681
15754	MODE-676	Why do you want to remove the "Each time <code>repository(String)</code> is called..." paragraph from the Getting Started guide? That functionality didn't go away. Perhaps the terminology about "load the existing definition" is not quite right? It also may be useful to use the "&version;" entity in place of the "1.0.0.Final" literals. I still think the default mode of the Java source file sequencer should be to produce the new content that matches what the Java class file sequencer produces. If anyone still wants the old format, they can set the "sourceFileRecorderClassName" property accordingly. Otherwise, the suggested changes look great!	1	1	https://issues.jboss.org/browse/MODE-676
15755	MODE-676	> Why do you want to remove the "Each time <code>repository(String)</code> is called..." paragraph from the Getting Started guide? That functionality didn't go away. Perhaps the terminology about "load the existing definition" is not quite right? Yes, but I believe that exact same note is about 3 paragraphs earlier. Apologies if there was a difference that I missed. > It also may be useful to use the "&version;" entity in place of the "1.0.0.Final" literals. That's a good call. Will the &version; entity match the version in Maven? > I still think the default mode of the Java source file sequencer should be to produce the new content that matches what the Java class file sequencer produces. If anyone still wants the old format, they can set the "sourceFileRecorderClassName" property accordingly. I am not comfortable making that call myself. I would not resist it either, although I still would like to hear from Serge on it. I'll get his $0.02 the next time that I see him. > Otherwise, the suggested changes look great! Thanks!	1	1	https://issues.jboss.org/browse/MODE-676
15756	MODE-676	>> Why do you want to remove the "Each time <code>repository(String)</code> is called..." paragraph from the Getting Started guide? >> That functionality didn't go away. Perhaps the terminology about "load the existing definition" is not quite right? > > Yes, but I believe that exact same note is about 3 paragraphs earlier. Apologies if there was a difference that I missed. You're right about the extra paragraph. My bad. >> It also may be useful to use the "&version;" entity in place of the "1.0.0.Final" literals. > >That's a good call. Will the &version; entity match the version in Maven? The "&version;" entity is defined in the "custom.dtd" file (which is not synchronized with Maven), and can simply be used in the DocBook source files.	1	1	https://issues.jboss.org/browse/MODE-676
15757	MODE-653	Well, what would the response be? Would you expect an XSL as well that maybe took the documents and formatted them?	0	0	https://issues.jboss.org/browse/MODE-653
15758	MODE-653	I was thinking more along the lines of a simple interface that would allow you to use a browser to: - see the available repositories (each with a link that does a GET on that repository) - see the available workspaces in a repository (each with a link that does a GET on that workspace, and showing links for the things you can do in that workspace, including browse and search) - browse and navigate (and eventually edit and remove) content under "items", with links for navigation and information about each node - submit queries and see results This is quite simple, and would merely allow more useful interaction with the REST service. WDYT?	1	1	https://issues.jboss.org/browse/MODE-653
15759	MODE-650	Ivan, can you please check the feature of import types from CMIS repository as we discussed and synchronization with bridge implementation regarding nt:file/nt:folder? Thanks!	1	1	https://issues.jboss.org/browse/MODE-650
15760	MODE-650	Oleg, reviewed changes and progress looks very good to me. Here are my findings: - cmis.cnd can be eliminated - CMIS (at least 1.0, I see in trunk they are added some "Item" for 1.1) has notion of "base type id" and "parent type id" for document/folder types definition. In case of mapping, cmis:document will be always base type id for all document types, but parent type id will vary from type to type (see my screenshots). So, when importing children you need to use parent type id to reflect hierarchy - Could you please elaborate about type definition cmis:repository (registerRepositoryInfoType)? What it is used for?	1	1	https://issues.jboss.org/browse/MODE-650
15761	MODE-650	Oleg, I made second patch from 31/Mar/13, could you please take a look?	0	0	https://issues.jboss.org/browse/MODE-650

#	Requirement	Questions	Answered	Successfully Answered	URL
15762	MODE-640	Do you think the factory should be logging AND throwing the exceptions? Seems to me that if it propagates an exception, it should not log it. This is usually a best practice, but this may need to be different because the engine may have multiple problems and because of the way it's part of JNDI and initialization. Thoughts?	1	1	https://issues.jboss.org/browse/MODE-640
15763	MODE-637	"org.modeshape.jboss.managed" (currently no subpackages). I've created a "Managed" object for each managed object listed in this JIRA except for the "Repositories" and "Connectors" groupings. Those will be created by JON/JOPR (per Ted). Each of the "Managed" objects will take a ModeShape delegate object in their constructor. Here is a first draft API for each of those managed objects. I'm still finalizing the method signatures and determining how to implement them. Please notice the "Questions" section for each object. Comments welcome. Thanks. 1. org.modeshape.jboss.managed.ManagedEngine Delegate class: JcrEngine Managed properties: Managed operations: 1. getRepositories() - returns Set<ManagedRepository> 2. getSequencingService() - returns ManagedSequencingService 3. getObservationBus() - returns ManagedObservationBus 4. getConnectors() - returns Set<ManagedConnector> Questions: 1. Is there a way to get the ObservationBus from the engine? 2. Would start() and shutdown() also be good managed operations? 3. Should managed connectors be cached or created each time getConnectors() is called? 4. Should managed repositories be cached or created each time getRepositories() is called? 2. Repositories -> Per Ted, this grouping will be handled by him 3. org.modeshape.jboss.managed.ManagedRepository Delegate class: JcrRepository Managed properties: 1. getQueryActivity() returns ? 2. getSaveActivity() returns ? 3. getSessionActivity() returns ? Managed operations: 1. listLocks() returns ? 2. listSessions(Comparable) returns ? 3. removeLock(Object lockToken) return boolean 4. terminateSession(Object sessionId) return boolean Questions: 1. Should descriptors be readonly managed properties? 2. Should options be readonly managed properties? 4. org.modeshape.jboss.managed.ManagedSequencingService Delegate class: SequencingService Managed properties: 1. getJobActivity() returns ? 2. getQueuedJobCount() return int Managed operations: 1. listQueuedJobs() returns ? Questions: 1. Should the properties of the Statistics be readonly managed properties? 2. Should the state (started, paused, etc.) be a managed operation? 5. org.modeshape.jboss.managed.ManagedObservationBus Delegate class: ObservationBus Managed properties: none identified in JIRA MODE-637 Managed operations: none identified in JIRA MODE-637 Questions: 1. Should hasObservers() be a managed operation? 2. Should shutdown() be a managed operation? 6. Connectors -> Per Ted, this grouping will be handled by him 7. org.modeshape.jboss.managed.ManagedConnectionPool Delegate class: RepositoryConnectionPool Managed properties: 1. getAvailableConnectionCount() returns int 2. getConnectionCount() returns int 3. getConnectionCreatedCount() returns int 4. getConnectionDestroyedCount() returns int 5. *getConnectionInUseCount() returns int 6. getConnectionInUseHighWaterMark() returns int 7. *getMaxSize() returns int 8. getMinSize() returns int Managed operations: 1. flush() Questions: 1. Should getKeepAlive]() be a managed property? 2. Should setKeepAlive(long, TimeUnit) be a managed operation? 3. Should getMaxFailedAttemptsBeforeError() be a managed property? 4. Should setMaxFailedAttemptsBeforeError(int) be a managed operation? 5. Should getRunState() be a managed operation? 6. Should setMaxFailedAttemptsBeforeError(int) be a managed operation? 7. Should setMaximumPoolSize(int) be a managed operation? 8. Should setPingTimeout(long, TimeUnit) be a managed operation? 9. Should setValidateConnectionBeforeUse(boolean) be a managed operation? 10. Should shutdown() be a managed operation? 11. Should prestartAllCoreConnections() be a managed operation? 12. Should prestartCoreConnection() be a managed operation? 8. org.modeshape.jboss.managed.ManagedConnector Delegate class: RepositorySource Managed properties: none identified in JIRA MODE-637 Managed operations: none identified in JIRA MODE-637 Questions: 1. Should getName() be a managed readonly property()? 1. Should getRetryLimit() be a managed operation? 2. Should setRetryLimit() be a good managed operation? 3. Should close() be a managed operation? 4. Should the RepositorySourceCapabilities	1	1	https://issues.jboss.org/browse/MODE-637
15764	MODE-637	What you describe sounds fantastic, and it's exactly along the lines of what I was hoping for. I'll try to answer your questions below (hopefully my organization will make sense): 1.1 Right now the ObservationBus is not exposed (better to not expose it until it's needed), so we'll have to add a way to get at it. We'll have to think about whether we want to expose public methods on JcrEngine, or perhaps have JcrEngine expose a "monitoring" object that has the accessor methods. (The latter option might simply make it more obvious that it's not for general use.) See also my answer for 5.1 1.2 "Start" and "Shutdown" might be good operations, as might a "Restart" (that shutsdown and starts the engine). 1.3 Offhand, the ManagedEngine should maybe just create new ManagedConnection objects each time, and should not hold onto them. That makes things much easier all around (no synchronization of the list of ManagedConnection, no removal, etc.), and I think this will work whether or not JOPR holds on to the managed objects. (If JOPR asks every time it needs them, then fine. But if JOPR caches the results, then those should be valid as long as JOPR "knows" when it should throw out its cached objects and refetch.) 1.4 Same as 1.3, but applied to ManagedRepository 3.1, 3.2, 3.3 (properties) These methods will all return an int, where the unit is something along the lines of "Sessions per minute". (Ted and I discussed other possible metrics, but they were always "instantaneous" values like "Open Sessions" that could very likely be not representative of other instants in time.) For now, the ManagedRepository will just have to fake up values. 3.1 (operations) listSessions will return an immutable list of (immutable) objects that each represent a session and that has the session's workspace name (String), username (String), internal session ID (string), and perhaps the time (DateTime) the session was created. This list of objects should be created each time the method is called. 3.2 (operations) listLocks will return an immutable list of immutable objects that each represent a lock in the repository and has the following fields: workspace name (String), session-based (boolean), locking session ID (string), lock expiration (DateTime), lock UUID (String), lock owner (String), deep (boolean). The actual locks are stored as nodes under "/jcr:system/mode:locks", so it should be pretty easy to read that information. 3.3 (operations) perfect. 3.4 (operations) perfect. 3.1, 3.2 (questions) Yes, descriptors and options should be read-only. Right now, changing an option requires restarting the engine. (Yeah, that's not ideal, but it's what it is at the moment. Until we get requests to make the options modifiable, we'll very likely leave it as it is.) 4.1 (properties) Job activity is an int (just like Session Activity). 4.2 (properties) Job queue is really just the list of runnables in the ExecutorService. We'll have to decide how to change this so we have the list of Futures. 4.1 (operations) We'll need an object that represents the Future. 4.1 (questions) yes, immutable 4.2 (questions) No need to manage the state at this time. 5.x Maybe it's not appropriate to expose the ObservationBus at this time. I think we should be very explicit about what we expose, and should not just expose something because it's there, and it sounds like maybe there's not much value in exposing it anyway. 7.1-7.4, 7.6-7.9 Shouldn't these just be managed properties? 7.5, 7.10-7.12 I don't think these need to be managed operations, since I'm not sure it makes sense to do these things to the pool (e.g., shutdown) without shutting down the engine. 8.1 Yes, getName() should be a read-only property 8.2 RetryLimit could be a writable property, couldn't it? 8.3 No, close() should not be a managed operation. (Doing that would make it unusable by the engine.) 8.4 Yes, these all can be managed, read-only properties. The remainder of the RepositorySource's bean properties can be writable properties. 8.5 Yes, ping() should be the only managed operation. What do you think?	1	1	https://issues.jboss.org/browse/MODE-637

#	Requirement	Questions	Answered	Successfully Answered	URL
15765	MODE-637	Delegate class: JcrEngine Managed properties: none identified in JIRA MODE-637 Managed operations: 1. getRepositories() - returns Set<ManagedRepository> 2. getSequencingService() - returns ManagedSequencingService 3. DELETED: getObservationBus() - returns ManagedObservationBus 4. getConnectors() - returns Set<ManagedConnector> 5. ADDED: start() returns void 6. ADDED: shutdown() returns void 7. ADDED: restart() returns void 8. ADDED: isRunning() returns boolean Questions: 1. RESOLVED: Is there a way to get the ObservationBus from the engine? NO. 2. RESOLVED: Would start() and shutdown() also be good managed operations? YES. 3. RESOLVED: Should managed connectors be cached or created each time getConnectors() is called? They should be created each time the the method is called. 4. RESOLVED: Should managed repositories be cached or created each time getRepositories() is called? They should be created each time the method is called. 2. Repositories -> Per Ted, this grouping will be handled by him 3. org.modeshape.jboss.managed.ManagedRepository Delegate class: JcrRepository Managed properties: 1. getQueryActivity() returns int 2. getSaveActivity() returns int 3. getSessionActivity() returns int 4. ADDED: getOptions() returns Map<String, String> 5. ADDED: getDescriptors() returns Map<String, String> Managed operations: 1. listLocks() returns returns List<ManagedLock> - sorted by owner 3. ADDED: listLocks(Comparator) returns List<ManagedLock> 2. ADDED: listSessions() returns List<ManagedSession> - sorter by user name 3. listSessions(Comparator) returns List<ManagedSession> 4. removeLock(String lockId) return boolean 5. terminateSession(String sessionId) return boolean Questions: 1. RESOLVED: Should descriptors be readonly managed properties? Yes. Descriptors should be readonly managed properties. Yes. Descriptors should be a managed readonly property. 2. RESOLOVED: Should options be readonly managed properties? Yes. Options should be readonly managed properties. 4. org.modeshape.jboss.managed.ManagedSequencingService Delegate class: SequencingService Managed properties: 1. getJobActivity() returns int 2. getQueuedJobCount() return int 3. ADDED: getNodesSequencedCount() return int 4. ADDED: getNodesSkippedCount() return int 5. ADDED: getStartTime() return long Managed operations: 1. listQueuedJobs() returns Collection<ManagedFuture> Questions: 1. RESOLVED: Should the properties of the Statistics be readonly managed properties? Yes they should be readonly managed properties. 2. RESOLVED: Should the state (started, paused, etc.) be a managed operation? No need to show state. 5. DELETED: org.modeshape.jboss.managed.ManagedObservationBus Delegate class: ObservationBus Managed properties: none identified in JIRA MODE-637 Managed operations: none identified in JIRA MODE-637 Questions: 1. Should hasObservers() be a managed operation? 2. Should shutdown() be a managed operation? 6. Connectors -> Per Ted, this grouping will be handled by him 7. org.modeshape.jboss.managed.ManagedConnectionPool Delegate class: RepositoryConnectionPool Managed properties: 1. getAvailableConnectionCount() returns int 2. getConnectionCount() returns int 3. getConnectionCreatedCount() returns int 4. getConnectionDestroyedCount() returns int 5. getConnectionInUseCount() returns int 6. getConnectionInUseHighWaterMark() returns int 7. getMaxPoolSize() returns int 8. getMinSize() returns int 9. ADDED: getKeepAliveTime() return long (nanoseconds) 10. ADDED: setKeepAliveTime(long) 11. ADDED: setMaximumPoolSize(int) 12. ADDED: getValidateConnectoinBeforeUse() return boolean 13. ADDED: setValidateConnectionBeforeUse(boolean) Managed operations: 1. flush() Questions: 1. RESOLVED: Should getKeepAliveTime]() be a managed property? Yes it should. 2. RESOLVED: Should setKeepAlive(long, TimeUnit) be a managed operation? It should be a managed property. 3. RESOLVED: Should getMaxFailedAttemptsBeforeError() be a managed property? Yes it should. 4. RESOLVED: Should setMaxFailedAttemptsBeforeError(int) be a managed operation? It	0	0	https://issues.jboss.org/browse/MODE-637
15766	MODE-637	Wanted to post an update to what I have done so far: - I've done a lot of work on the annotations for managed objects, properties, and operations. Still don't understand fully all what is required. - I've implemented all the methods not mentioned below. - I've stubbed in some dummy data for most things. So these managed classes can be used now before actually getting a real ModeShape engine. Here is the work left to do: - check-in existing code (waiting for 1.2 release first) - complete annotations (need to get help here) - Connection Pool 1. implement flush() returns boolean 2. how do you calculate connections destroyed count? 3. how do you calculate in-use high water mark? 4. how do you implement getMinSize() - Connector 1. need to find out how to reflectively add managed properties - Engine 1. how do you lookup engine in JNDI? - Repository 1. implement getQueryActivity() returns int 2. implement getSaveActivity() returns int 3. implement getSessionActivity() returns int 4. implement listLocks() returns List<ManagedLock> 5. implement listSessions() returns List<ManagedSession> 6. implement removeLock(String lockId) returns boolean 7. implement terminateSession(String sessionId) returns boolean - Sequencing Service 1. implement getJobActivity() returns int 2. implement getQueuedJobCount() returns int 3. should getStartTime() return a string-ified date instead of long? 4. implement listQueuedJobs returns Collection<ManagedFuture> 5. implement ManagedFuture Answers to above questions and help on how to implement the above methods is appreciated. Thanks!	0	0	https://issues.jboss.org/browse/MODE-637
15767	MODE-637	These new operations probably need to be on a managed repository, and we'll have to include the workspace name in the input. Having said that, do we have a way of viewing the available workspace names for a repository? If not, we may want to add that ability.	1	1	https://issues.jboss.org/browse/MODE-637
15768	MODE-637	Do we need a list of sequencers to display in JON/Jopr?	1	1	https://issues.jboss.org/browse/MODE-637
15769	MODE-637	We don't really need a list of sequencers, but we may want to provide a way to see the sequencer configurations. Each configuration defines a classname, classpath (tho that will be unused in our JBoss installation), and path expression [1]. Is there a way in JOPR to easily add these configurations, perhaps where the classname is really a choice list of named sequencers (and we specify the classname in the managed objects) and editable path expression strings? [1] http://docs.jboss.org/modeshape/latest/manuals/reference/html_single/reference-guide-en.html#path-expressions	1	1	https://issues.jboss.org/browse/MODE-637
15770	MODE-637	Randall... "Is there a way in JOPR to easily add these configurations, perhaps where the classname is really a choice list of named sequencers (and we specify the classname in the managed objects) and editable path expression strings? " By editable path expression strings, do you mean that is something you would like to persist for a sequencer? If we have a way to get all names sequencers and their properties, we can display them as a separate resource or maybe on the engine configuration page.	1	1	https://issues.jboss.org/browse/MODE-637
15771	MODE-612	The approach looks great. One question: it looks like in the SessionCache.JcrNodeOperations that the "compute(...)" implementation initializes the version history for the supplied node, but that in the Graph.save() method (where an anonymous class is used to call the "compute(...)" method) is done once for each child. Is this right? Won't this be expensive for large numbers of children (even when the parent and/or children are not versionable)?	1	1	https://issues.jboss.org/browse/MODE-612
15772	MODE-611	I'm willing to provide an implementation. to be honest, I think the current one should be renamed "connector-hibernate" would a target for the older JPA 1 or the newer JPA 2 be desired? I see some level of desire to expose criteria queries, and based on what I see you doing in the current connector, JPA 2 does automatically (deleting orphan children, for example)	0	0	https://issues.jboss.org/browse/MODE-611

#	Requirement	Questions	Answered	Successfully Answered	URL
15773	MODE-611	The MODE-1209 patch will switch us to doing just that (q.v., https://github.com/bcarothers/modeshape/commit/2505083f0e3a053b456c7889bb08ca5d092ce857#L4R1276). Hopefully, that will get committed to trunk today. We still have a lot of legacy settings on JpaSource itself though that map to implementation-specific functionality rather than parts of the JPA standard. For example, setting connection pool sizes, transaction isolation levels, etc. You could argue (correctly) that all of this can be set in persistence.xml, but if we only supported that method, we'd have a pretty big break in backwards compatibility. Also, it seems wrong not to be able to override some common settings at that level. I suppose that we _could_ add a GenericJpaAdapter that ignores all of the source settings (and adds warnings to that effect), allowing users to use any JPA implementation and control everything through the persistence.xml if they so chose. WDYT?	1	1	https://issues.jboss.org/browse/MODE-611
15774	MODE-580	We already have acquired all the resources for the new name: - we have a Magnolia project for the ModeShape JBoss.org site (http://www.jboss.org/modeshape), though it's not yet active as I haven't published it yet - we own the 'modeshape.org' domain is ours, and http://www.modeshape.org will soon redirect to http://www.jboss.org/modeshape, and we'll be able to use 'org.modeshape' for our package names. - we have a new JIRA, and I will be moving all of the DNA issues into the ModeShape JIRA. Our issues will be of the form "MODE-XYZ" plus a new SVN repository, new mailing lists, new blog site, new Twitter account, and new Ohloh project. - we have a new logo design and should soon have the images we need Here are the first steps: 1) The JBoss DNA SVN has already been locked down, and no more changes can be made to it. DONE! 2) All of the history from the DNA SVN has already been loaded into the ModeShape SVN. Basically, it looks exactly like the DNA repository did right after the 0.7 release. But at the moment, it is locked down as well, pending the next step. DONE! 3) I will move all DNA JIRAs (all of them, including those that are complete or incomplete) to the ModeShape JIRA. Any links to DNA issues will automatically redirect to the corresponding issue in the ModeShape JIRA. The numbers won't match, but JIRA has it's own way of doing things. 4) I've started renaming the projects, packages and classes. The projects will go from 'dna-...' to 'modeshape-...", and our package structure will go from "org.jboss.dna..." to "org.modeshape...". And any class that has "DNA" in the name will be renamed. I will then commit these changes to the ModeShape SVN. I will also move all of the JIRA issues from DNA to 5) I will then open up the ModeShape SVN, and you could switch your development environments over to use https://svn.jboss.org/repos/modeshape/trunk . But at this point, we still only want to change JavaDoc, comments, and the documents. After this is done, we'll release ModeShape 1.0.0.Beta1, which will essentially be the same thing as DNA 0.7 except with the new brand, new module names, and updated documentation. We'll publish the http://jboss.org/modeshape project, and we'll start publicizing everywhere we can. And, the code in SVN will be changed to be 1.0.0.GA-SNAPSHOT, and we can continue fixing issues, completing any part of the rebranding, and working towards 100% TCK runs and 1.0.0.GA! What can you do now? - Sign up for the new mailing lists at https://lists.jboss.org/mailman/listinfo/. As usual, we have modeshape-announce@..., modeshape-commits@..., modeshape-dev@..., modeshape-issues@..., and modeshape-users@.... - If you're on Twitter, follow @modeshape and retweet when the time comes - If you use an RSS reader, add our new blog at http://modeshape.wordpress.com/. No need to do anything if you're already subscribed to the JBoss Feed. - If you have any local changes, create some patches and edit them so that they'll work with the new project names and package structure - Go to the new IRC room, #modeshape on freenode.net - Relax for a few days while I do my magic with JIRA and SVN	1	1	https://issues.jboss.org/browse/MODE-580
15775	MODE-536	Serge, Is there a reason you suggested http://www.jboss.org/dna/downloads/eclipse-update-sites.html ? When we changed over to the new project site, the downloads were moved to http://downloads.jboss.org/jbossdna/ (for example http://downloads.jboss.org/jbossdna/0.6/jboss-dna-0.6-gettingstarted-examples.zip). JBoss.org automatically provides download metrics for artifacts on http://downloads.jboss.org, so it'd be useful to have the actual update site be at http://downloads.jboss.org/jbossdna/updatesite/<versionOrLatest>. Note that the last URL would have a readable HTML page that should point to the installation instructions (something akin to http://www.jboss.org/tools/download/update.html). This is the same pattern that JBoss Tools uses: a page at http://www.jboss.org/tools/download.html that describes the different update sites that are actually available at URLs like http://download.jboss.org/jbosstools/updates/stable/	0	0	https://issues.jboss.org/browse/MODE-536
15776	MODE-465	Has any work begun on this or is it still in the planning stage ?	1	1	https://issues.jboss.org/browse/MODE-465

#	Requirement	Questions	Answered	Successfully Answered	URL
15777	MODE-465	Actually, the work is pretty far along, and I'm actively working on it (now that 0.6 is out the door). I'm actually working on a (local) branch, and am hoping to commit what I have to trunk relatively soon. The design is based on the approach that JCR 2.0 has taken: multiple query languages are parsed into a single Abstract Syntax Tree (a "query model"), which is then executed by single engine. This means that 'dna-graph' will have a QueryEngine component parses string queries (given the name of the language, and a registered parser for that language that transforms the string query into an AST representation), plans and optimizes the query, constructs a processing plan, and delegates to a 'QueryProcessor' component to execute the processing plan and generate the results. These results are then returned to the caller of the QueryEngine. This design allows us to do a couple of different things. First, we can create multiple language parsers, yet have one engine. Second, the planning and optimization components can either be customized or replaced, so the whole query engine framework can be reused. Third, the query engine is completely independent of the execution, though the capabilities of the QueryProcessor can be taken into account during planning/optimization. We're actually providing a QueryProceessor implementation that uses Lucene indexes, and with this a way to populate the Lucene indexes by scanning the graph content and/or updating the indexes based upon events. We'll actually plug this query engine into the DNA JCR system in a couple of ways. Each connector will optionally be able to support querying its content (either by pushing down the queries to the underlying system, or by having a Lucene index and just instantiating the query engine within the connector). The federation connector will support queries by pushing down queries to the connectors being federated, and it will be able to manage a Lucene index for each connector that doesn't support queries. Then, since the JCR implementation will basically just implemented on top of the federated connector, it just has to delegate the queries down to the federated connector. Currently the following components are pretty complete (and have a lot of unit tests, though the tests are probably not sufficient): the QueryEngine, the AST classes, the query planner, the rule-based optimizer, and the QueryProcessor framework. The Lucene-based QueryProcessor implementation is at least 50% done and what remains is mostly integration. I'll be done with the JCR SQL2 parser (actually, the parser is complete and I have many unit tests, but only 50% of the parser is tested) in another day or so. What's left is: - commit the current state (I've been waiting to make sure all the pieces look like they'll fit together, and I'm probably no more than a few days away) - write an XPath parser - start integration testing of the query engine with the Lucene-based processor - add query support to the connector API - change some of our connectors to support queries (mostly just integrating the query engine and Lucene processor), though the federation connector will require some extra customization - add support for text extraction - expose through our JCR implementation - make sure the events work and are wired correctly, as well as set up mechanisms to manage the scanning and re-indexing of all content Another approach is to not yet put add query support to the connector API, but instead to just have our JCR implementation own the QueryEngine and Lucene-based processor. I'm not sure if this will save much time or effort (mostly because of federation), but it is an option. This is still a fair amount of work, but as you can see I think this design will be pretty powerful and will allow the connectors to really control how they handle queries. When it's committed to trunk, we'll be able to divvy up the tasks. WDYT so far? Yeah, this is a really rough description, but I do hope to document this more and blog about some of the details.	0	0	https://issues.jboss.org/browse/MODE-465
15778	MODE-455	a good approach - it provides master persisted lock information accessible to all distributed processes. The organization seems pretty good - /jcr:system/dna:locks/<lock_for_node_given_node_uuid>, where the lock node contains the information about the node being locked as well as the information about the owner of the lock. And the restriction to mix:referenceable nodes is obvious, since the nodes' UUIDs are used in the name of the DNA lock. Considering that this approach doesn't work with nodes that are not mix:referenceable, would it work to use the UUID for the lock (rather than the UUID of the locked node) as the node name? The UUID and/or path could then be kept in the lock info node (and because this node is locked, only the session with the lock token can modify the node and possibly move the node, and in that case the session could update the path in the lock info). This would make it more difficult to quickly look up a lock given a node, but isn't that the point of the in-memory map? The in-memory map could be changed to be keyed by Location, which would work by path and/or UUID and/or by other ID properties. I'm not sure whether this would even work, because I might be missing some aspects of the current design or failing to recognize all of the issues with this alternative. Thoughts? Back to the current patch. Here are a few comments: In WorkspaceLockManager: 1) As the patch defines the code, the DnaLock.isLive() method is getting backdoor access to the map outside of a lock scope, and this seems like an error. 2) I'm not sure whether the read-write lock is guarding only the in-memory map or also the persisted content. If it's the latter, then what about other sessions created in other processes? If the read-write lock is to guard only the in-memory map, then the map probably can be changed to be a concurrent map and the read-write lock removed entirely, which could very well have a significant performance impact since this method is called once for each node in the path of a node that is being modified. Also, is the map just providing an in-memory cache of the content in the repository? 3) The lock tokens appear to just be the string form of the locked node UUID. This is pretty easily forged, correct? Not sure that this is a bad thing, but it does give users a bit more rope to hang themselves. 4) The lockFor(UUID) method could be changed to accept an AbstractJcrNode, and thus would encapsulate a bit more of the WorkspaceLockManager internals. Of course, the lockFor(String token) would need to be changed a bit to reflect how the lockFor(UUID) worked. The rest of my comments are _very_ minor nits. In AbstractJcrNode: 1) lock() will be called a lot (for every modification method on any node or property) ... seems like it'd be useful to cache some of the references. Also (minor nit) seems like the code could be cleaned up to reduce duplication by using a do-while loop. In JcrRepository: 1) JcrRepository.login(...) calls JcrRepository.getLockManager() and passes the WorkspaceLockManager reference into the JcrWorkspace constructor. This additional constructor parameter really isn't necessary, since the JcrWorkspace constructor could get it's own WorkspaceLockManager from the JcrRepository that is passed as a parameter to the constructor. 2) Some of the use of the concurrent lockManagers map is not really correct. Basically, one should never do multiple checks within a single method, since the map can change from one map method to the next. So, for example: WorkspaceLockManager getLockManager(String workspaceName) { WorkspaceLockManager lockManager = lockManagers.get(workspaceName); if (lockManager != null) return lockManager; lockManagers.putIfAbsent(workspaceName, new WorkspaceLockManager(executionContext, this, workspaceName, locksPath)); return lockManagers.get(workspaceName); } should really be: WorkspaceLockManager getLockManager(String workspaceName) { // Check for an existing lock manager ... WorkspaceLockManager lockManager = lockManagers.get(workspaceName); if (lockManager != null) return lockManager; // Need to create one if one hasn't just been created ... lockManager = new WorkspaceLockManager(executionContext, this, workspaceName, locksPath); WorkspaceLockManager	1	1	https://issues.jboss.org/browse/MODE-455

#	Requirement	Questions	Answered	Successfully Answered	URL
15779	MODE-455	to the map outside of a lock scope, and this seems like an error. FIXED - switched to use of ConcurrentHashMap and removed the lock 2) I'm not sure whether the read-write lock is guarding only the in-memory map or also the persisted content. If it's the latter, then what about other sessions created in other processes? If the read-write lock is to guard only the in-memory map, then the map probably can be changed to be a concurrent map and the read-write lock removed entirely, which could very well have a significant performance impact since this method is called once for each node in the path of a node that is being modified. Also, is the map just providing an in-memory cache of the content in the repository? FIXED - Switched DnaLock to ThreadSafe (since it has no mutable internal state but does rely on the thread-safe workspace locks map), moved workspaceLocks to be a ConcurrentHashMap and deleted the manual locking 3) The lock tokens appear to just be the string form of the locked node UUID. This is pretty easily forged, correct? Not sure that this is a bad thing, but it does give users a bit more rope to hang themselves. FIXED - The lock tokens are now UUIDs themselves that are independent of the underlying node's UUID. 4) The lockFor(UUID) method could be changed to accept an AbstractJcrNode, and thus would encapsulate a bit more of the WorkspaceLockManager internals. Of course, the lockFor(String token) would need to be changed a bit to reflect how the lockFor(UUID) worked. NOT FIXED - There's a place where I need to walk up the graph from the current node to the root to make sure that none of the parents have a deep lock. I can get the underlying o.j.d.g.Node easier than the AbstractJcrNode. If I changed to using AbstractJcrNode as the argument, I'd have to do more materialization (albeit with the same number of data retrieves). I'm still open to changing this, but it doesn't seem worth the materialization penalty to me off the cuff. The rest of my comments are _very_ minor nits. In AbstractJcrNode: 1) lock() will be called a lot (for every modification method on any node or property) ... seems like it'd be useful to cache some of the references. Also (minor nit) seems like the code could be cleaned up to reduce duplication by using a do-while loop. NOT FIXED - I don't think that you can do any non-trivial lock caching (i.e., cache the results of lockFor(nodeUuid)) since that could be changed outside of the current session. I did cache the workspace().session().lockManager() calls. I think I see what you're saying about refactoring into a do-while, but the lock handling is different for a lock on the current node vs. one of its parents if lockIsDeep is false. I think it's slightly more efficient as is, but I may be missing something. In JcrRepository: 1) JcrRepository.login(...) calls JcrRepository.getLockManager() and passes the WorkspaceLockManager reference into the JcrWorkspace constructor. This additional constructor parameter really isn't necessary, since the JcrWorkspace constructor could get it's own WorkspaceLockManager from the JcrRepository that is passed as a parameter to the constructor. FIXED - Good point. The redundant constructor arg is no more. 2) Some of the use of the concurrent lockManagers map is not really correct. Basically, one should never do multiple checks within a single method, since the map can change from one map method to the next. So, for example: WorkspaceLockManager getLockManager(String workspaceName) { WorkspaceLockManager lockManager = lockManagers.get(workspaceName); if (lockManager != null) return lockManager; lockManagers.putIfAbsent(workspaceName, new WorkspaceLockManager(executionContext, this, workspaceName, locksPath)); return lockManagers.get(workspaceName); } should really be: WorkspaceLockManager getLockManager(String workspaceName) { // Check for an existing lock manager ... WorkspaceLockManager lockManager = lockManagers.get(workspaceName); if (lockManager != null) return lockManager; // Need to create one if one hasn't just been created ... lockManager = new WorkspaceLockManager(executionContext, this, workspaceName, locksPath); WorkspaceLockManager newLockManager = lockManagers.putIfAbsent(workspaceName, lockManager); if (newLockManager != null) { // Some other thread just added this lock	1	1	https://issues.jboss.org/browse/MODE-455
15780	MODE-455	I definitely like the changes to the graph layer, including the changes to the fluent graph API as well as the requests. It may be good to elaborate a bit in the JavaDoc for the connector capabilities and what 'supports locking' actually means. I think it's intuitive that if a connector "supports locking", it basically allows locks to be used on (at least some of) its content, and will either confirm or veto a lock requests. But it may be a bit more ambiguous as to what it means to not support locking. I'd guess this means that the connector is saying that locks will not be permitted on any of its content, and that if any lock requests do come down, they will be vetoed/rejected. Is this correct? And what will the behavior of our existing connectors be? IMO, the JPA, in-memory, JBoss Cache, and Infinispan connectors should certainly support locking, but since they are really "storage" connectors (meaning that the underlying systems are 'owned' by DNA and are being accessed by DNA) they have no need to do anything special with the lock requests, other than approve them. On the other hand, the file system connector should _not_ support locking (perhaps this can change when Java 7 enables interaction with native FS locks). And the SVN connector is read-only right now and thus should also not support locks, though when the connector is changed to support writes, the connector could also be changed to support locking by integrating with native SVN locks. Does this sound right? Probably should update the Reference Guide, too.	1	1	https://issues.jboss.org/browse/MODE-455
15781	MODE-454	Below are the TCK tests that are failing using my (not currently checked-in) implementation of the ObservationManager and why I think they are failing. One design note: I had to implement a check that makes sure Changes with a timestamp before the time a listener was registered do not get delivered to that listener. This was a problem because the events were created and the listener was being registered in one thread and the event notification was being handled in a different thread. I will write Jiras for (1) requests don't differentiate between setting a new property or changing an existing property, and (2) Changes need to include a session identifier so that the ObservationManager can handle the "noLocal" flag. ========== LockingTest testAddLockToNode - no PROPERTY_ADDED events received testRemoveLockFromNode - no PROPERTY_REMOVED events received NodeRemovedTest testMultiNodesRemoved - no NODE_REMOVED event for child nodes when a parent is removed NodeMovedTest testMoveWithRemove - fails with NPE (originalParent) in GraphSession.containsChangesWithExternalDependencies() NodeReorderTest testNodeReorder - no events received (need to research what happens when node.orderBefore(sourcePath, destinationPath) is called) PropertyAddedTest testSinglePropertyAdded - no events received since PROPERTY_ADDED changes are implemented as PROPERTY_CHANGED events testMultiPropertyAdded - no events received since PROPERTY_ADDED changes are implemented as PROPERTY_CHANGED events PropertyChangedTest testSinglePropertyChangedWithAdded - too many events as PROPERTY_ADDED changes are being broadcast as PROPERTY_CHANGED events AddEventListenerTest testNodeType - path is /testroot/node2/node3 but should be /testroot/node1/node3 (need to research this one) testPath - should not have path /testroot/node1 (so does isDeep mean not to include parent path?) testIsDeepFalseNodeAdded - path is /testroot/node1 and should be /testroot/node1/node2 (expected test results seems wrong??) testIsDeepFalsePropertyAdded - no events since PROPERTY_ADDED changes are sent as PROPERTY_CHANGED events testNoLocalTrue - still have to implement "noLocal" check in ObservationManager testUUID - no events received since no PROPERTY_ADDED changes are implemented as PROPERTY_CHANGED events WorkspaceOperationTest testCopy - no NODE_ADDED event generated for child node of copied parent node	1	1	https://issues.jboss.org/browse/MODE-454

#	Requirement	Questions	Answered	Successfully Answered	URL
15782	MODE-445	Overall, I think it turned out very nice. I only checked it closely in the JCR module (which led to me attaching this patch to hide three version-related classes), but it seemed to show about what I would have expected. On a related note, should we hide the JcrI18n class? I can't see anyone needing to access that directly from their application. I did get a three test failures when I ran "mvn -P assembly clean install" from modeshape-source (instead of modeshape-jcr). They all looked something like this: java.lang.AssertionError: Expected: is "JBoss, a division of Red Hat" got: "${project.parent.organization.name}" at org.junit.Assert.assertThat(Assert.java:502) at org.junit.Assert.assertThat(Assert.java:492) at org.modeshape.jcr.JcrRepositoryTest.testDescriptorValues(JcrRepositoryTest.java:484) at org.modeshape.jcr.JcrRepositoryTest.shouldProvideBuiltInDescriptorsWhenNotSuppliedDescriptors(JcrRepositoryTest.java:195) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.junit.internal.runners.TestMethod.invoke(TestMethod.java:59) at org.junit.internal.runners.MethodRoadie.runTestMethod(MethodRoadie.java:98) at org.junit.internal.runners.MethodRoadie$2.run(MethodRoadie.java:79) at org.junit.internal.runners.MethodRoadie.runBeforesThenTestThenAfters(MethodRoadie.java:87) at org.junit.internal.runners.MethodRoadie.runTest(MethodRoadie.java:77) at org.junit.internal.runners.MethodRoadie.run(MethodRoadie.java:42) at org.junit.internal.runners.JUnit4ClassRunner.invokeTestMethod(JUnit4ClassRunner.java:88) at org.junit.internal.runners.JUnit4ClassRunner.runMethods(JUnit4ClassRunner.java:51) at org.junit.internal.runners.JUnit4ClassRunner$1.run(JUnit4ClassRunner.java:44) at org.junit.internal.runners.ClassRoadie.runUnprotected(ClassRoadie.java:27) at org.junit.internal.runners.ClassRoadie.runProtected(ClassRoadie.java:37) at org.junit.internal.runners.JUnit4ClassRunner.run(JUnit4ClassRunner.java:42) at org.apache.maven.surefire.junit4.JUnit4TestSet.execute(JUnit4TestSet.java:62) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:140) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:127) at org.apache.maven.surefire.Surefire.run(Surefire.java:177) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.booter.SurefireBooter.runSuitesInProcess(SurefireBooter.java:345) at org.apache.maven.surefire.booter.SurefireBooter.main(SurefireBooter.java:1009)	1	1	https://issues.jboss.org/browse/MODE-445
15783	MODE-445	What is 'modeshape-source' ? This error typically occurs when the property files are not properly filtered before copying to the target/ folder.	1	1	https://issues.jboss.org/browse/MODE-445
15784	MODE-402	Can we agree on the limitation that all properties so defined must come before any child nodes?	1	1	https://issues.jboss.org/browse/MODE-402
15785	MODE-402	Brian, did you mean to say "there doesn't seem a way to add a property VALUE ..." ? The current graph API does allow you to set a property on a node, even if that node was created in a previous request (even if they're in the same batch). For example, I can create a node (maybe supplying a handful of properties), and then call setProperty(x).on(samePath) and that will work just fine. In fact, if setProperty is called right after the create, then the set property request even get merged into create request. The graph API currently doesn't allow adding a value to a property; this is the heart of DNA-379. The importer should already be keeping track of the path given the current stack, and you should always be able to set a property on the node (even after any child elements have just been finished). The only tricky part is that to handle multi-valued properties (i.e., where each value is represented by a different child element's content), then the importer might have to keep a map of the properties set via child elements (and overwrite them if additional child elements are read with additional values for a property), and set them only when the element is ended. If I'm correct, then it shouldn't matter whether the property child elements come before child elements, or in what order the property child elements appear. It just means that the importer has to maintain a stack of properties for each "incomplete node" (or do this on the existing stack). Waiting for DNA-379 would allow this to work without having to maintain the properties in the stack, since the processing of each child node (with content) would mean just adding that value to the property with the given name. I agree it would be easier. The only question is whether it's worth doing DNA-379 in 0.4.	1	1	https://issues.jboss.org/browse/MODE-402
15786	MODE-381	Applied both patches. The first hides the "dna:nodeDefinition" property from the JCR layer (meaning it might exist in the graph layer, but won't appear as properties of JCR nodes). The second patch adds to "nt:base" a property definition for a "dna:nodeDefinition" property that is protected and not mandatory (nor autocreated). The spec is not clear as to whether a JCR implementation can add properties to "nt:base" (the required supertype of all node types), the spec is clear that an implementation can add properties to any of the built-in node types through mixins. Therefore, the second patch actually adds this new definition to "nt:base" by having "nt:base" extend the new "dna:defined" mixin, which has the new "dna:nodeDefinition" property definition. Because the "dna:nodeDefinition" is protected, JCR clients will not be allowed to define another property with this name, nor will they be able to set the property (using a residual definition) to any value. All this validation is done through the normal process, requiring no code specific to "dna:nodeDefinition". Nicely done! This is a great demonstration on how node types, node definitions, and property definitions can be used to enforce a structure while having no additional validation code.	0	0	https://issues.jboss.org/browse/MODE-381

#	Requirement	Questions	Answered	Successfully Answered	URL
15787	MODE-363	I like the approach of using templates (ala 283). Is there a way to change the NodeTypeSource-approach to use the new registration mechanism? Seems like that might clean up things quite a bit. I also do like where you are going with refactoring of findPropertyDefinition to use Predicates. However, I'm not convinced it's better. One goal I had with the old design was that the logic was very obvious and "in your face" (not masked by a lot of method calls), and I deemed this as more beneficial than the ugliness of the duplication. So while the new approach gets rid of the duplication, it doesn't really keep the "in your face" logic. Plus, the new approach doesn't reduce the number of lines (385 new vs 373 old). Besides, I'd prefer for this change to be pulled out into a different patch (or even a different issue). Also, considering the size of the change to support registration and unregistration, would it be better to not address thread-safety in the same patch? Also, although using a thread-safe container would be beneficial, we made need to use an explicit ReadWriteLock if the thread-safe container doesn't offer the required level of atomicity. Requiring users to oder the types to be registered is a perfectly valid assumption, I think. At least for now, since it'd likely keep the implementation much more simplistic. One thing we're going to have to address is that we can't update existing JcrNodeType objects. So, if a supertype is changed (re-registered), we really need to rebuild all of the node types that extend it. Unfortunately this is the only way to truly have immutable objects.	1	1	https://issues.jboss.org/browse/MODE-363
15788	MODE-310	We need to use only SVN-compatible patches so that any users (including those without Eclipse workspaces) can apply the patches. (Although Eclipse can produce standard patch files for a project, it creates Eclipse-specific files when creating a patch from changes to multiple projects.) Creating patches is very easy. Let's assume we want to create a patch to capture local changes made against the latest revision of "trunk". Simply go to the "trunk" folder in your local working area, and issue this command: svn diff . > ~/DNA-000.patch where "DNA-000" represents the JIRA issue number that describes the changes (and where the patch will be uploaded). Note that the destination in the command creates the patch file in your home directory, but you could place it anywhere. Also note that 'diff' ignores binary files, so any binary files need to be uploaded separately (http://www.igorexchange.com/node/89#creating_patches) Now, someone else wants to apply the patch, so they download the patch file and apply the differences using the 'patch' utility: patch -p0 < ~/DNA-000.patch The "-p0" option makes sure that all files can be found.	0	0	https://issues.jboss.org/browse/MODE-310
15789	MODE-293	From Sergiy: My thoughts that we would want to anticipate wide range of customer's different deployment scenarios with minimum efforts, right? This is especially true with SOA-Platform and JBoss Enterprise Data Services Platform which includes JBoss ESB, JBoss jBPM, JBoss Rules and JBoss Enterprise Application Platform (JEE App server). So in addition to Embedded and MicroRepository we would certainly need "Enterprise" Repository that may be deployed as EAR into any JEE compatible server, and may be accessble through RMI in addition to HTTP, HTTP+REST, etc. IMO, we would still need to develop few "reference deployment packages" and I can participate in these efforts. Even more, we can integrate Jackrabbit transparently into web/app server - thanks to JackRabbit has JCA connector and JNDI, so DNA clients (REST API, web UI) will think they talking to JBoss DNA JCR. Jackrabbit also allows loading previously exported XML content into JCR repository - so we can test DNA predefined graph structure. That will also help with integrated and functional testing of our stack without waiting for DNA JCR implementation. For instance, REST API and web applications will require at least basic functional testing - I was thinking about using SOAP UI to test WS services and REST API. WAR packaging scenarios aree OK, but if you have multiple session facades (your business logic) plus multiple web apps (e.g. REST API, web admin UI based on Flex, web app based on RichFaces, web app based on XXX, etc) + shared code - the right approach would be to use EAR that was designed for that purpose and stores multiple WAR files, shared libs, session facade EJB3 artifacts all together. Actually this resembles MVC paradigm, right? Plus, remember that EBJ3 artifacts are just POJO with annotations - so that makes them simple and thin while benefiting from standard services like resource pooling, thread management, security and transaction propagations, JMX, etc. Same Stateless Session Bean could be accessed locally, remotely through RMI, and as WebService with absolutely no extra effors needed. On the other hand, JEE App servers getting smaller, becoming embedded like JBoss Embedded and small like Glassfish embedded. It's no longer monolithic 100Mb thing - we can prepare deployment with just few megs. Not to mention that nextJEE6 is talking about different profiles to support more lightweight scenarios. From Randall: I completely agree that DNA will can anticipate many different scenarios. But I also am concerned about over-engineering and adding complexity or intent where we don't need it yet. From Sergiy: What's the DNA overall deployment strategy? I've mentioned it here https://www.jboss.org/community/docs/DOC-12952 From Randall: There will be several. One is an embedded case, where an application wants to use a local repository and will access it through the JCR API. This repository may be federated from multiple sources, including the file system, and may use sequencers, allowing the application to have a repository with a mixture of the content it wants. A variation of this is the MicroRepository. This is also embedded, but uses just a few of the DNA artifacts plus very minimal dependencies. The idea is that an application could use the repository as the bootstrapping mechanism. Think of an OSGi repository, or Maven repository, or JBoss Microcontainer. Another case will be a shared resource, much like a database. I foresee this being most easily deployed as a .war file, or actually a choice of .war files that range from a simple RESTful service on top of multiple repositories, to including other web services and web applications). This is probably closest to your architectural diagram, but I'm not sure I know exactly what kind of technologies will be used - I presume we'll try to pick the best tools for the job, once we know more about the requirements. Also, I don't want to go too far with this, as there are plans within JBoss to use DNA inside the SOA-Platform and JBoss Enterprise Data Services Platform. It may be that the open-source DNA remains mostly a set of artifacts and maybe some tools, leaving the more "integrated stacks" of technologies for products (from JBoss and others).	1	1	https://issues.jboss.org/browse/MODE-293
15790	MODE-277	Isn't this a dupe of MODE-41?	1	0	https://issues.jboss.org/browse/MODE-277
15791	MODE-273	Wouldn't it be just as useful to sequence all of the content, and allow the user of the resulting information decide whether only the first line is used? One benefit is that a CSV importer in Tool X might be able to show the user some of the data, helping the user know whether particular columns are of interest.	1	1	https://issues.jboss.org/browse/MODE-273
15792	MODE-273	Would it be best to sequence the entire file (which could easily contain multiple megabytes of data) or just a sample (perhaps the first N rows)?	1	1	https://issues.jboss.org/browse/MODE-273
15793	MODE-266	Does the latest patch add enum and annotation support for the source version, too? That's awesome. Should we make the two sequencers produce the same content (by default)? And should we do that for 1.0? That seems to be a good thing to bring them together.	1	1	https://issues.jboss.org/browse/MODE-266

#	Requirement	Questions	Answered	Successfully Answered	URL
15794	MODE-235	databases as part of single server instance, Oracle uses single database but multiple schemas. We may review how database access tools like Aqua Data Studio (based on Java Swing and JDBC) intepreting such graphs. I'm talking about the root part of such graph - from root node to let's say table node. > For example: > <root database node>/<sub graph that difficult to map in generic way>/tables/... > <root database node>/<sub graph that difficult to map in generic way>/views/... > <root database node>/<sub graph that difficult to map in generic way>/users/... > > We need to think more how we will map subgraph highlighted above onto different db vendors concepts: > MS SQL/Sybase, etc: > <root database node>/databases/<db name>/user tables/<user name>.<table name> > <root database node>/databases/<db name>/system tables/sys.<table name> > <root database node>/databases/<db name>/views/<view name> > > > Oracle and others similar: > <root database node>/schema/<user name>/tables/<table name> > <root database node>/schema/<user name>/views/<view name> > > We may consider merging both styles together or having them separate (both approaches have advantages and disadvantages). I would prefer > generic, merged concept view - for benefits of DNA-JCR and REST toghether so clients don't have to change graph based on particular vendor. From Randall: Do we really think it's our job to come up with a single model? Why can't we simply expose what the database exposes? Like you said, it's mostly at the high-level (what I'd call the organizational level, above the tables/views/procedures/etc.). How about taking it one step at a time, and focusing on getting the functionality there to expose the same organizational structure that the database uses. At some point, we'll probably want to come up with a single approach, but even then I think we'd want to make it optional (specified in the metadata connector). I also think that we should use mixins for the type information (or maybe even node types that are extensible to allow any properties), and that means we could actually not have any nodes that exist to collect the types. This is much more along the lines of a So, Oracle's schema might show up as content: <db:schema jcr:name="Products"> <db:table jcr:name="PRODUCTS"> <db:column jcr:name="ID" ... /> <db:column jcr:name="NAME" ... /> ... <db:primaryKey ... /> ... </db:table> <db:table jcr:name="TABLES" jcr:mixins="db:system" ...> ... </db:table> <db:view ...> ... </db:view> <db:procedure> ... </db:procedure> </db:schema> I've shown that any system table/view/procedure could have an additional "System" mixin. The great thing about this approach (of using mixins rather than names in a hierarchy) is that you can more readily build queries to get the content in any structure you want. I've used an XML notation; for details about this notation, see the note somewhere after paragraph 5 in http://www.jboss.org/file-access/default/members/dna/freezone/docs/0.3/manuals/reference/html/repositories.html#dna-connector-federation. From Sergiy: > Anyway, for both styles we may add generic metadata subnode about RDBMS capabilities (that will hold all not specific to database object propertise like supportsMixedCaseIdentifiers(), storesUpperCaseIdentifiers(), supportsAlterTableWithAddColumn(), etc): > > <root database node>/capabilities From Randall: Yup. <db:capabilities supportsMixedCaseIdentifiers="true" db:supportsUpperCaseIdentifiers="true" ... /> From Sergiy: Sub nodes of table, view, etc nodes are pretty easy to model because they are similir for many databases. I was thinking about something similar to following (I guess we may have design sessions and may refactor it later): .../tables/<table name>/data [default primary type, so it can be omitted] table's data records .../tables/<table name>/data/record 1 .../tables/<table name>/data/record N From Randall: Let's table the data discussion for the moment. I've got very different ideas about how that will work, since this is too flat for tables with millions of rows. We may want to do something like this for a sampling of the data (something that's very useful when you're dealing with metadata, but even then there are enterprise security issues with this). From Sergiy: .../tables/<table name>/columns/<column name> .../tables/<table name>/constraints/<constraint name>	1	1	https://issues.jboss.org/browse/MODE-235
15795	MODE-235	I think we have a couple of choices when it comes to designing the hierarchy for database metadata. 1) Do not use node types at all, but instead use some other mechanism (e.g., hierarchy with predefined names to connote the types of content under it). This would be a difficult approach, as it relies too much on convention and is not flexible (e.g., what happens when a node has multiple types). It also is neglecting several of the primary features of JCR. 2) Use node types, where the node type is used to identify the "primary type" of the node. JCR only allows one primary type, and it is a constant for the lifetime of the node. Other aspects can be captured as mixins and/or property values. The node types and mixins would define the properties and children that are allowed. 3) Use a generic node type (e.g., "nt:unstructured") and do all the interesting things with mixins. Since the node type really allows anything, it doesn't matter that it is constant throughout the node's lifetime. Mixins describing the various facets ("system-ness", "table-ness", "view-ness", "schema-ness" etc.) could be added an removed at will. So if we rule out #1, then the choice is between #2 and #3. Normally, I'd be in favor of #3, but database metadata is fairly structured, and it's usually not important to change from one type to another. The node type could also restrict the types and number of children. Therefore, #2 seems like a better approach to me.	1	1	https://issues.jboss.org/browse/MODE-235

#	Requirement	Questions	Answered	Successfully Answered	URL
15796	MODE-233	the children of a node ... Children children1 = graph.get().children().at(path); Children children2 = graph.get().children().of(path); Children children3 = graph.get().children().of(path).matching(nameMatcher); Children children4 = graph.get().children().matching(nameMatcher).of(path); Children children5 = graph.get().children().having(propertyMatcher).of(path); // Get the values of a single property by iterating ... for (Object value : graph.get().property(name).at(path)) { } for (Object value : graph.get().property(name).on(path)) { } for (Object value : graph.get().property(name).of(path)) { } // Get the properties on a node by iterating ... for (Property property : graph.get().properties().at(path)) { } for (Property property : graph.get().properties().at(path).matching(nameMatcher)) { } for (Property property : graph.get().properties().at(path).matching(nameMatcher)) { } // Get a node (both properties and children) ... Node node1 = graph.get().node().at(path); Node node2 = graph.get().node().with(uuid); Node node3 = graph.get().node().having(uuid); Node node4 = graph.get().node().at(uuid); Node node5 = graph.get().node(path); Node node6 = graph.get().node(uuid); // Work with the properties and children of a previously obtained node... Children children6 = node1.getChildren(); Children children7 = node1.getChildren().having(propertyMatcher); Children children8 = node1.getChildren().matching(nameMatcher); Property prop7 = node1.getProperty(name); MultipleProperties props4 = node1.getProperties(); MultipleProperties props5 = node1.getProperties().having(propertyMatcher); MultipleProperties props6 = node1.getProperties().matching(nameMatcher); // Get a subgraph (nodes and properties) Subgraph subgraph1 = graph.get().subgraph("/a/b"); Subgraph subgraph2 = graph.get().subgraph("/a/b").ofDepth(3); Subgraph subgraph3 = graph.get().subgraph().at("/a/b").ofDepth(3); Subgraph subgraph4 = graph.get().subgraph().at(uuid).ofDepth(3); // Get nodes from within a subgraph ... Node node10 = subgraph1.getRoot(); Node node11 = subgraph1.getNode(relativePath); Node node12 = subgraph1.getNode(uuid); // We can add support for visitors, too // graph.get().subgraph(path).visit(myVisitor); // Then there are operations on a graph that result in changes to the graph. These consist of // copy, delete, move, create, and set. When using these, the graph records the commands but // delays executing (or committing) them against the source until the "commit()" method is called. // // Changes to the graph can be interspersed with "get()" commands. When this happens, it is // up to the graph implementation to define the behavior when the result of a get command // is requested (e.g., the get command is to be executed) while there are uncommitted changes. // Copy a portion of the content in the graph ... graph.copy().node("/a/b/").into("/x/y"); graph.copy().node("/a/b/").into("/x/y").and().get().node("/m/n"); graph.copyNode("/a/b/").into("/x/y"); graph.copy().node().at("/a/b").into("/x/y"); graph.copy().node().with(uuid).into("/x/y"); graph.copy().subgraph("/a/b").into("/x/y"); graph.copy().subgraph().at(uuid).into("/x/y"); graph.copySubgraph("/a/b/").into("/x/y"); graph.copy().children().of(uuid).into("/x/y"); graph.copy().children().matching(nameMatcher).on(uuid).into("/x/y"); graph.copy().children().of(uuid).matching(nameMatcher).into("/x/y"); graph.copy().children().of("/a/b").into("/x/y"); graph.copy().properties().of("/a/b").into("/x/y"); graph.move().node("/a/b/").into("/x/y"); graph.move().node("/a/b/").into("/x/y").and().get().node("/m/n"); graph.move().node().at("/a/b").into("/x/y"); graph.move().node().with(uuid).into("/x/y"); graph.move().properties().of("/a/b").into("/x/y"); graph.move().properties().of("/a/b").matching(nameMatcher).into("/x/y"); graph.move().children().of("/a/b").into("/x/y"); graph.move().children().of(uuid).into("/x/y"); graph.move().children().with(nameMatcher).on(uuid).into("/x/y"); graph.move().children().having(propertyMatcher).on(uuid).into("/x/y"); graph.move().children().matching(nameMatcher).on(uuid).into("/x/y");	0	0	https://issues.jboss.org/browse/MODE-233
15797	MODE-140	Renamed the "getEnvironment()" methods to "getExecutionContext()" to reflect the previous change in name of the ExecutionContext (was ExecutionEnvironment). Also renamed the "getProperties()" method of some of the SPI graph commands that returned an Iterator<Property> to now be "getPropertyIterator()", which is now consistent with the other graph commands. Made FederatedNode a concrete class (not an interface), moved it, and implemented additional interfaces so that it can be treated as a GetNodeCommand. Finally, made more changes to the FederationCommandExecutor implementation.	0	0	https://issues.jboss.org/browse/MODE-140
15798	MODE-138	The federation system will use JAAS to do authentication. This was implemented in the FederatedRepositorySource by making use of the ExecutionContextFactory.create(...) methods that will use JAAS LoginContext instances to authenticate the subject. Once an ExecutionContext (was called ExecutionEnvironment) is obtained by the FederatedRepositorySource, it uses that context to talk to the configuration repository and to create the FederatedRepository instance.	0	0	https://issues.jboss.org/browse/MODE-138
15799	MODE-53	The most obvious first approach for a URI design is to directly represent each node in the repository as a resource, and to map the repository operations to GET (read), PUT (update/create), POST (create), and DELETE (delete). For example, "/a/b/c/d" would identify the node "/a/b/c/d". A "GET /a/b/c/d" request would obtain a representation of the "/a/b/c/d" node. Perhaps the first challenge to this approach is whether to also treat properties as resources. If each property were treated as a resource, then the URI design would have to distinguish a node property named the same as a that of a child of the same node. For example, if a node /a/b/c in the previous example contained a child "d" and a property "d", then is "/a/b/c/d" referring to the node or the property? One option is to distinguish properties with a character ('@' comes to mind), although which character would depend on URL escaping requirements. It probably is possible to require this character only when there is a child with the same name as a property name (something that is not too frequent, I expect). The kind of resource to return (e.g., JSON, XML, HTML, etc.) should be driven by the "accept" HTTP header, which specifies the media ranges (mime types with optional parameters) that are being requested. Multiple mime types can be specified, with precedence rules and quality factor denoting the preferred order. See http://www.w3.org/Protocols/rfc2616/rfc2616-sec14.html By using the HTTP header, the URI no longer needs to deal with extensions or any other mechanism to denote the desired response format. And this means that the URI design can more closely match the resource names. Here's an example of a GET request to obtain a JSON representation of node "/a/b/c/d": GET /a/b/c/d Host: myhost.example.com Accept: application/json URI query parameters are useful to supply additional criteria and information about the representation to be returned, such as depth of the representation or reading just the properties or children. This seems to be an acceptable way of affecting the representation. What about operations, like move or copy? One choice is to use query parameters (e.g., PUT with a "?moveTo=/a/b/e" or POST with a "?copyOf=/a/b/e"). Another option would be to allow actions or operations in the representation that is PUT by the client. The following example is an HTTP request to move node "/a/b/c/d" to "/a/b/e" and simultaneously set the "jcr:title" property to "This is a new title": PUT /a/b/c/d Host: host.example.com Accept: application/json Content-Length: nnn Content-Type: application/json { "dna-action:moveTo": "/a/b/e" "jcr:title" : "This is the new title" }	1	1	https://issues.jboss.org/browse/MODE-53

#	Requirement	Questions	Answered	Successfully Answered	URL
15800	MODE-53	The URI design may need to make room for different repositories, workspaces, etc. The first segment of the URI (below the host or application URI) might specify the repository name, the second the workspace name, followed by the path to the node. For example: /repoA/myWorkspace/a/b/c/d would identify the node "/a/b/c/d" in the "myWorkspace" workspace of the "repoA" repository configured by the RESTful server. What about searching? Well, search could be done using the "/search" URI, making that a reserved name and unusable for the repository name in the configuration. That probably is pretty acceptable. What about a web application for users to view, navigate, edit the content through a browser? Well, its quite likely that this application would be in a different WAR (perhaps still deployed to the same server). However, there are two options if the web app is to be in the same WAR file: reserve the "/ui" URI for this (might be a good idea anyway), or simply use the HTTP accept header to specify that a client wants the (probably rich) HTML for a resource.	1	1	https://issues.jboss.org/browse/MODE-53
15801	MODE-48	There is some overlap between the JCR 1.0 and 2.0 API, but we may need separate API implementations (depending upon whether 2.0 is a superset of 1.0).	0	0	https://issues.jboss.org/browse/MODE-48
15802	MODE-34	Can this be done through URLs, or does this need access to the file system to get directory listings?	1	1	https://issues.jboss.org/browse/MODE-34
15803	MODE-34	Should this use WebDAV?	1	1	https://issues.jboss.org/browse/MODE-34
15804	MODE-31	the sequencer was whether the UUIDs used within the model files should be used as the "jcr:uuid" property value on the corresponding nodes. If we did do this, then resolution would be really easy, as model references could be mapped directly into JCR (weak) references, and JCR would magically do the dereferencing for us. But JCR (and ModeShape) uses "jcr:uuid" as a primary key, so there can be only one node with a particular UUID in a workspace. That means that if we went this route, we'd only be able to have one copy (or version) of a model in the repository at a time. This is the same constraint as we had with MMR and DTC, and IIRC there were quite a few cases where that proved to be a major limitation. So if we want to be able to upload a model any number of times into the repository (e.g., the model by itself and any VDBs containing a model), we need a different approach. If we instead capture the 'mmuuid' values as separate properties, then we can still have properties that use these XMI UUIDs as indirect references. This will require using JCR-SQL to resolve references (to the possibly multiple nodes, if there are multiple versions of a model in the repository), though when we do know a single node to which the reference resolves where the references resolve, we can actually cache/record this as a JCR weak reference property. In fact, a simple pattern of four properties allows us to capture the raw information (e.g., hrefs or mmuuids) and have slots to capture the resolved information (e.g., the names and weak references) if known. I've defined a simple naming convention for these properties. If the reference is single valued: {referenceName}Href - stores the href literal value from the XMI file {referenceName}XmiUuid - stores the XMI UUID to the referenced node, if the href had an embedded UUID (hrefs to data types and XSD components don't use UUIDs) {referenceName}Name - stores the name of the resolved node {referenceName} - stores the JCR weak reference to the resolve node If the reference has an upper bound greater than 1 (e.g., it is multi-valued), then the properties are slightly different: - {singularReferenceName}Hrefs - stores the href literal values from the XMI file - {singularReferenceName}XmiUuids - stores the XMI UUIDs to the referenced nodes, if the hrefs had embedded UUIDs (hrefs to data types and XSD components don't use UUIDs) - {singularReferenceName}Names - stores the names of the resolved nodes - {pluralReferenceName} - stores the JCR weak references to the resolve nodes References to model objects within the same model are easily resolved upon sequencing, and so we set all of the properties. However, references to objects in other models cannot be resolved at sequencing time (though I'm trying to make it so that sequencing a VDB causes the references within the VDB's models to all be resolved). Because we need the ability to upload a model into a repository more than once, the current sequencer uses the more complicated multi-property pattern. However, it'd be pretty easy to add another sequencer that uses the XMI UUIDs as the JCR UUIDs, but I'll wait on that since we wouldn't use that in production anyway. === Node Types === The sequencer follows JCR best-practices by defining all nodes to have a primary type of "nt:unstructured", meaning it's possible and valid for any node to have any property (with single or multiple values). However, it is still useful to capture the metadata about what that node represents, and so we use mixins to do this. For example, we have a "xmi:referenceable" mixin with a single "xmi:uuid" property, and that is patterned after the built-in "mix:referenceable" mixin. Since all model objects have mmuuids, all nodes will have this mixin. I've written a utility that reads in an ECore metamodel file and writes out a rough CND file containing a mixin node type for each metaclass, where that mixin contains a property definition for each metaclass attribute and a child node definition for each metaclass containment relationship. Each non-containment relationships are represented using the four-property pattern mentioned above. These automatically-generated CND files were then hand-modified to be slightly more usable. === Default Values === As you know, default values are not written by EMF into the XMI model files, so the sequencer has nothing to read in. Luckily, JCR has a way to deal with this. Because we're defining mixin node types	0	0	https://issues.jboss.org/browse/MODE-31
15805	MODE-28	Problem I see with sequencers is that they don't get anything close to a filename so guessing mimetype based on file extension is quite hard. That leads to other problems as without file extension, all e.g. OpenOffice documents are application/zip, all source files all text/plain, etc. Is there any other way how to get the currently sequenced node reference to search for other info?	1	0	https://issues.jboss.org/browse/MODE-28
15806	MODETOOLS-79	Did you want the choices to only be (1) include project - which is the current way, or (2) not include project? Or also maybe a (3) exclude project and enter a new root path? For instance, if there is an Eclipse project called "Test" with these files: /Test/a.txt /Test/b.txt /Test/folder/c.txt And a ModeShape publish area called "/files." The results would be: For (1) above (current way): /files/Test/a.txt /files/Test/b.txt /files/Test/folder/c.txt For (2) above: /files/a.txt /files/b.txt /files/folder/c.txt For (3) above with new path called "mypath": /files/mypath/a.txt /files/mypath/b.txt /files/mypath/folder/c.txt	1	1	https://issues.jboss.org/browse/MODETOOLS-79
15807	MODETOOLS-79	[~elvisisking]: do you think either (2) or (3) would require significant effort ? Based on this we'll have to see whether this can be included in the upcoming DV ER(s) or not.	1	1	https://issues.jboss.org/browse/MODETOOLS-79
15808	MODETOOLS-54	I created a new CND with a node type and property and there was no MS header. Is this the scenario? Could you have had some file template auto-inserted maybe?	1	0	https://issues.jboss.org/browse/MODETOOLS-54
15809	MODETOOLS-8	Randall, I was leaving that open until I resolved the update-site URL issue we discussed. This morning I talked to Mickael Istria about how we could change where the Jenkins job puts the update-site. Currently the job puts it here: /downloads_htdocs/modeshape/builds/staging/modeshape-tools-continuous/all/repo You mentioned you didn't like that path. What do you want that update-site URL for this job to look like? Mickael said I need to create a Jira and assign to Nick Boldt with the new path. Let me know the path you would like and I will write issue and get with Nick. Thanks!	0	0	https://issues.jboss.org/browse/MODETOOLS-8

#	Requirement	Questions	Answered	Successfully Answered	URL
15810	MODETOOLS-8	Although we're setting the location, isn't the rest of that path defined by the script? If so, I'd guess there were a reason JBT used "staging" and "all/repo". Perhaps that's because their release process involves copying/moving the code over? That's been my assumption: that the JBT folks had reasons for doing that. What happens after we have an update site for say 3.0 but then want to release 3.1? Shouldn't it get added into the existing update site somehow? If not and our job should just publish in the correct location, then we certainly can change do that (it does make things easier). Pulling this completely out of thin air, my initial guess might be these paths for our actual update sites: {{/downloads_htdocs/modeshape/tools/update/}}_{{\{stable|develop\}}}_{{/}}_{{\{helios|indigo|juno|etc\}}}_ This assumes that each build works with a single Eclipse platform version, and we therefore need to break out the Eclipse platform version in our URL structure. (I don't know when/if we'll start building for multiple platforms, but we should account for that possibility.) WDYT? We could possibly do without the "{{update}}" segment, but I'm not sure it's that big a deal to leave it and may even be useful to help users recognize it as an update site URL.	0	0	https://issues.jboss.org/browse/MODETOOLS-8
15811	MSC-125	Is it not just a matter of appending a list representation of the cycle array to the message?	1	1	https://issues.jboss.org/browse/MSC-125
15812	MSC-84	Can you be more specific? The API already allows for reporting of failure causes for services.	0	0	https://issues.jboss.org/browse/MSC-84
15813	NTA-94	Amos, You can make a start with the current docs. Currently it only deploys to a specific commit of WildFly (see docs). Can you figure out why it doesn't deploy to WildFly master and fix the issue? When I update the docs to show how to run the tool, you can check that that works on WildFly master.	0	0	https://issues.jboss.org/browse/NTA-94
15814	NTA-22	Why does the demo scenario create "two additional transactions"? I'd only expect to see one JTA transaction.	1	1	https://issues.jboss.org/browse/NTA-22
15815	NTA-10	Amos, Can you take a look at our other PR jobs and see if you can make one for NTA?	1	1	https://issues.jboss.org/browse/NTA-10
15816	NETTY-452	@Trustin: I just pulled it to master and 3.2. Can you close it?	1	0	https://issues.jboss.org/browse/NETTY-452
15817	NETTY-450	When implementing page/action cache in my web framework, I wanted to use HttpMethod as one of the elements of cache key. It's convenient if HttpMethod were serializable. But it's OK now. I don't want HttpMethod to be serializable any more. Thanks for your work with Netty.	0	0	https://issues.jboss.org/browse/NETTY-450
15818	NETTY-413	I would agree but the issue is: what is a message then ? What I mean is to count Written or Read Messages, you have to know how are those messages. For instance, take an HTTP answer, you can first try to write a first series of bytes with only the header, then make a second write with the body. Is it 2 messages or 1 ? Functionnaly it is 1, but technically, you made 2 writes... So my guess is that it could be difficult to have such counters in a standard way, even if I perfectly agree with you of the interest. So do you have any idea ?	1	0	https://issues.jboss.org/browse/NETTY-413
15819	NETTY-345	HttpChunkAggregator could merge the trailing headers into the aggregated HttpMessage. Would it be enough?	1	1	https://issues.jboss.org/browse/NETTY-345
15820	NETTY-336	Trustin, will this also deal with the problem in JBPAPP-6092 ? The problem there is that we run multiple Netty transports in a single JVM and the tread delimiter is static, so you can't have a prefix per transport there, because the thread delimiter class will be the one associated with the last transport started, in this case the Memcached one. I can't find a workaround for this either because when a thread is created, all I know is whether is a worker/master or whether it's a client or server one however, I have no extra information available to me. The thread naming delimiter should be a factor of the NettyChannelPipelineFactory or the ServerBootstrap, rather than being static.	1	1	https://issues.jboss.org/browse/NETTY-336
15821	NETTY-336	Is this still planned to be released?	0	0	https://issues.jboss.org/browse/NETTY-336
15822	NETTY-333	Hi Trustin.. the above branch already supports file upload. That's what i figured looking at the example for http server isn't it?	0	0	https://issues.jboss.org/browse/NETTY-333
15823	NETTY-333	Kevin, could you elaborate with me on if how it's different from Frederic's implementation? Maybe I am missing something.	1	0	https://issues.jboss.org/browse/NETTY-333
15824	NETTY-311	Not needed anymore.. IpFilter will do.	0	0	https://issues.jboss.org/browse/NETTY-311
15825	NETTY-307	@Trustin: Do you have any preferences if this should only get added to master ? I think we could add it to 3.2 branch too, as it is an isolated feature and so a "non risk" change.. WDYT?	1	1	https://issues.jboss.org/browse/NETTY-307
15826	NETTY-307	My guess is that there might be some duplication in Daniel's and Norman's work. Perhaps we should join the force and reduce any possible duplication? Norman, could you take a look into Daniel's pull request and see what we can do if you have some free time? Cheers, Trustin	1	0	https://issues.jboss.org/browse/NETTY-307
15827	NETTY-297	Hi Steven, I could make the framing code more strict to address this issue since the current implementation checks only the length field. Do you have an example packet that triggers this issue?	1	1	https://issues.jboss.org/browse/NETTY-297
15828	NETTY-257	This is a nasty issue.. and seems to boil down to the fact that CharsetEncoder and CharsetDecoder cannot compute the needed length for a given sequence of bytes, and thus it gets quite hard to know if the given operations will succeed (from the users perspective). If its to be implemented anyways then I suggest the following signatures: // Use this to copy into an existing array, returns number of chars written in dst public int getChars(int index, int length, char[] dst, int dstIndex, int dstMaxLength, String charsetName); // Use this to decode into a new array of needed size. CharsetEncoder and CharsetDecoder can produce this for you. public char[] getChars(int index, int length, String charsetName); // Write a number of chars to the ChannelBuffer, returns number of bytes written public int writeChars(int index, char[] from, int fromIndex, int fromLength, String charsetName); Regarding optimizations, then you could either create charset specific versions of these methods, ie. getASCIIChars, writeCharsInASCII. Or you could switch on the charset internally in each method. Decoding and encoding UTF-16 is quite easy to do naively. Same with US-ASCII except it can produce encoding errors. Im not sure its a good idea to handle anything else oneself unless its a strict subset of UTF-16.	1	0	https://issues.jboss.org/browse/NETTY-257

#	Requirement	Questions	Answered	Successfully Answered	URL
15829	NETTY-257	Perhaps the following might be enough? // Useful for reading fixed length strings // Returns the number of written chars int getString(int index, int length, char[] dst, Charset charset); int getString(int index, int length, char[] dst, int dstOffset, Charset charset); int getString(int index, int length, Appendable dst, Charset charset); // Useful for reading nul-terminated strings // Returns a status object that tells the number of read bytes and written chars (see below) DecodeResult getString(int index, int length, char[] dst, Charset charset, ChannelBufferIndexFinder terminatorFinder); DecodeResult getString(int index, int length, char[] dst, int dstOffset, Charset charset, ChannelBufferIndexFinder terminatorFinder); DecodeResult getString(int index, int length, Appendable dst, Charset charset, ChannelBufferIndexFinder terminatorFinder); + equivalent readString() methods (no need to specify index, readerIndex advances) // Returns the number of written bytes int setString(int index, char[] src, Charset charset); int setString(int index, char[] src, int srcOffset, int srcLength, Charset charset); int setString(int index, CharSequence src, Charset charset); int setString(int index, CharSequence src, int srcOffset, int srcLength, Charset charser); + equivalent writeString() methods (no need to specify index, writerIndex advances) final class DecodeResult { final int readBytes; final int writtenChars; } In case of overflow, IndexOutOfBoundsException could be raised. To get the safe number of required chars / bytes, you could use the following idiom: CharsetDecoder d = ...; int requiredChars = (int) ((double) numberOfBytesToDecode * d.maxCharsPerByte())); or CharsetEncoder e = ...; int requiredBytes = (int) ((double) numberOfCharsToEncode * e.maxBytesPerChar())); What do you think?	1	1	https://issues.jboss.org/browse/NETTY-257
15830	NETTY-257	Is the lack of EncoderResult on purpose? Im not sure its bad to support Appendable. You could always instanceof the dst and cast it to a CharBuffer to avoid making a copy. I guess the same could be done with the CharSequence from setString(). I guess thats it. It looks good to me.. except maybe a utility method that produces a String/char[] without giving a dst.	1	0	https://issues.jboss.org/browse/NETTY-257
15831	NETTY-257	Im not really sure what I was thinking about the EncoderResult.. About Appendable. What I tried to explain was that you could just allocate a new CharBuffer to decode into and then pass that to the Appendable instance. However, if the given Appendable is an instance of CharBuffer then just cast and pass that to the CharsetDecoder. Its a couple of days since I looked at it but Im pretty sure the append, and put methods on CharBuffer have the same behaviour so it should be sound to do. Does it make sense?	1	0	https://issues.jboss.org/browse/NETTY-257
15832	NETTY-257	That's a great idea (Appendable). Here's the revised signatures with much less number of methods: int getString(int index, int length, Appendable dst, Charset charset); DecodeResult getString(int index, int maxLength, Appendable dst, Charset charset, byte terminator); DecodeResult getString(int index, int maxLength, Appendable dst, Charset charset, ChannelBufferIndexFinder terminatorFinder); int readString(int length, Appendable dst, Charset charset); DecodeResult readString(int maxLength, Appendable dst, Charset charset, byte terminator); DecodeResult readString(int maxLength, Appendable dst, Charset charset, ChannelBufferIndexFinder terminatorFinder); int setString(int index, CharSequence src, Charset charset); int setString(int index, CharSequence src, int srcOffset, int srcLength, Charset charset); int setString(int index, Readable src, Charset charset); int writeString(CharSequence src, Charset charset); int writeString(CharSequence src, int srcOffset, int srcLength, Charset charset); int writeString(Readable src, Charset charset); final class DecodeResult { final int readBytes; final int writtenChars; } Internally, CharBuffer could be determined using the instanceof operator because CharBuffer implements both Readable and Appendable. However, I'm not really sure if the getters with ChannelBufferIndexFinder are really needed. Most encoded strings are terminated by one byte (e.g. NUL) except for UTF-16. I'd like to remove them if there is a common use case for using ChannelBufferIndexFinder to find a string terminator.	0	0	https://issues.jboss.org/browse/NETTY-257
15833	NETTY-246	I still think HTTP proxy connectivity should be separated into a new transport. For example: ClientSocketChannelFactory f = new HttpProxyClientSocketChannelFactory(new HttpTunnelClientChannelFactory(new NioClientSocketChannelFactory(...))); Is there any reason that the tunneling transport and the proxy connectivity should be combined into one transport?	1	1	https://issues.jboss.org/browse/NETTY-246
15834	NETTY-246	I've checked in a few small changes mostly related with code cleanup: * tabs -> spaces * private fields accessed by inner classes -> changed the access modifier to package-private * set missing svn:eol-style and svn:keywords properties - perhaps the autoprops configuration was not working when you were adding the files? I also added a couple 'FIXME' markers to source code regarding the unused variables. Could you take a look? Now I will replace java.util.logging with Netty's logging layer, which is also trivial. :)	1	1	https://issues.jboss.org/browse/NETTY-246
15835	NETTY-246	Sure. I do agree with your idea to improve the logging API. We could take care of the logging API problem as a separate issue. By the way, the log messages at WARN / ERROR level could be converted to the 'exceptionCaught' events because it is unusual to see such cases. What do you think?	1	1	https://issues.jboss.org/browse/NETTY-246
15836	NETTY-246	The cases where I can see things being logged at the WARNING level all relate to when an unknown message is received. This is likely to happen if a browser attempts to talk to the tunnel server, or if the client attempts to establish a tunnel with a normal web server that does not understand the request. It may be justified to throw exceptions up the pipeline if this occurs, but I don't know which exception. Is there a good, generic "connection failed" exception that's worth raising, that we could subclass to highlight the specifics? Some comments regarding the FIXME's you have added: HttpTunnelAcceptedChannelFactory - you have a FIXME here stating that the methods should not be exposed to users. The interface itself is default protected, so the interface will not be usable outwith the org.jboss.netty.channel.socket.httptunnel package. So, users cannot use these methods. Also, methods on interfaces are always public, whether you specify as such or not. The only thing you can do to narrow their availability is to change the interface to default protection instead of public. HttpTunnelClientWorkerOwner - similarly, this interface is not useable outwith the httptunnel package, so the FIXME seems invalid. HttpTunnelClientChannel - I agree with the FIXME in onBindRequest(). Perhaps we could bind the send channel to the specified SocketChannel, and bind the poll channel to an ephemeral port on the same interface?	1	1	https://issues.jboss.org/browse/NETTY-246
15837	NETTY-246	Rescheduling to ALPHA3 since ALPHA2 is supposed to be released today. Iain, could you let me know the progress of this issue? I didn't see check-ins recently. If you don't mind, I can fix them by myself. If you have local modifications, I'd suggest you to check in more often. Thanks!	1	1	https://issues.jboss.org/browse/NETTY-246
15838	NETTY-246	I notice that, when using OIO, proxy settings set by system properties will be applied automatically on the Socket used by Netty (new Socket() then uses sun's SocksSocketImpl that has built in support for HTTP and SOCKS proxies...). This is not the case when using NIO, however, whereas we need explicit support for HTTP and SOCKS put into Netty... Am I right?	1	1	https://issues.jboss.org/browse/NETTY-246
15839	NETTY-242	Hello, did you have the time to take a look at this contribution ? -- ron	1	0	https://issues.jboss.org/browse/NETTY-242

#	Requirement	Questions	Answered	Successfully Answered	URL
15840	NETTY-239	Does HttpRequest.getUri() return different values for the same HTTP request? HttpRequest.getUri() always return the second component of the request line in verbatim. For example, getUri() always returns '/jaxws-endpoint/endpoint/number1?wdsl' for the following request: GET /jaxws-endpoint/endpoint/number1?wsdl HTTP/1.0 ... Because the URI field is final internally, it is not supposed to change unless you did something with it in the getRequestPath(String) method or a different request was received. getUri() will return 'http://127.0.0.1:8871/jaxws-endpoint/endpoint/number1/' only if the request looks like the following: GET http://127.0.0.1:8871/jaxws-endpoint/endpoint/number1/ HTTP/1.0 ... Could you tell me what differences you've found?	1	1	https://issues.jboss.org/browse/NETTY-239
15841	NETTY-236	Is there a workaround for this? Right now this is a show stopper for me because I'm dealing with a legacy protocol that works in the described way...	1	1	https://issues.jboss.org/browse/NETTY-236
15842	NETTY-236	Thanks a lot for your time Trustin, I really appreciate it. I'm assuming then that currently there is no workaround for this scenario. Is that correct? I'm trying to do something like this on the client.... clientSocket.shutdownOutput(); Assert.assertEquals('0', (char) clientSocket.getInputStream().read());	0	0	https://issues.jboss.org/browse/NETTY-236
15843	NETTY-236	Is this feature in the pipe? It would be quite useful because it enables us to do a more orderly shutdown of the connection where we don't risk losing incoming traffic from the remote peer.	1	1	https://issues.jboss.org/browse/NETTY-236
15844	NETTY-236	Is now part of 4.0.0.x. You can now call Channel.shutdownInput() or Channel.shutdownOutput().	0	0	https://issues.jboss.org/browse/NETTY-236
15845	NETTY-155	I didn't realize that multicast already existed in netty. Where could I find the documentation on it? Thanks for letting me know, btw!	1	0	https://issues.jboss.org/browse/NETTY-155
15846	NETTY-153	Hi all, I'd like to request a further feature from RTSP protocol that would be nice to have in Netty. [RFC 2326|https://datatracker.ietf.org/doc/rfc2326/] for RTSP in its section 10.12 describes a method of interleaving binary data in RTSP stream. Those data may be sent from server to client or vise versa. The format is very simple: > Stream data such as RTP packets is encapsulated by an ASCII dollar > sign (24 hexadecimal), followed by a one-byte channel identifier, > followed by the length of the encapsulated binary data as a binary, > two-byte integer in network byte order. The stream data follows > immediately afterwards, without a CRLF, but including the upper-layer > protocol headers. Each $ block contains exactly one upper-layer > protocol data unit, e.g., one RTP packet. Although the existing RTSP codec is easy to extend to support interleaved data, it would be great to have such an extension provided natively. Is it realistic to have the extended codec in 3.2.0? Thanks.	1	1	https://issues.jboss.org/browse/NETTY-153
15847	NETTY-126	trustin, a couple of questions. // FIXME: Support both version 0 and 1 cookies (both for encoding and decoding) I coldn't extract the differences between 0 and 1, any chance you could shed some light? // FIXME: Encode all cookie fields, including domain, path, maxAge, secure, and comment. // FIXME: Decode all cookie fields, including domain, path, maxAge, secure, and comment Do you just want these added as fields to the cookie class and encoded.decoded if set by the user? // FIXME: CookieDecoder cannot assume that the first field is always the name-value pair. ok, what else could it be?	1	1	https://issues.jboss.org/browse/NETTY-126
15848	NETTY-126	> // FIXME: Support both version 0 and 1 cookies (both for encoding and decoding) > I coldn't extract the differences between 0 and 1, any chance you could shed some light? I think the difference is that version 0 cookies use 'expires' attribute and version 1 cookies use 'max-age' attribute, as described in the chapter 10 (HISTORICAL) of http://www.ietf.org/rfc/rfc2109.txt. I think we can simply convert 'expires' to 'max-age' and vice versa instead of adding the Expires attribute to the HttpCookie class. I also found that there's cookie version 2: http://www.ietf.org/rfc/rfc2965.txt. The version 2 adds more properties like 'Port', 'CommentURL', and 'Discard'. I think we need to support version 2, too. It's interesting that the version 2 cookie represents itself as version '1' cookie except that it uses a different header name. > // FIXME: Encode all cookie fields, including domain, path, maxAge, secure, and comment. > // FIXME: Decode all cookie fields, including domain, path, maxAge, secure, and comment > Do you just want these added as fields to the cookie class and encoded.decoded if set by the user? Yes. I added the fields for version 0 and 1 to the cookie class by myself, so what's left is to: * add the fields for version 2 * fix the decoder to understand version 0, 1, and 2 cookies. (HttpCookie.getVersion() should return 1 for version 2 cookies, as specified in the RFC.) * fix the encoder to be able to encode any version of cookies. (The desired version could be specified as an integer constructor parameter.) > // FIXME: CookieDecoder cannot assume that the first field is always the name-value pair. > ok, what else could it be? Check the example 5.1.3 of http://www.ietf.org/rfc/rfc2109.txt: Cookie: $Version="1"; Customer="WILE_E_COYOTE"; $Path="/acme" I think it can even be like this: Cookie: $Path="/acme"; $Version="1"; Customer="WILE_E_COYOTE";	0	0	https://issues.jboss.org/browse/NETTY-126
15849	NETTY-112	Alan, could you close this issue if this satisfies your need?	1	1	https://issues.jboss.org/browse/NETTY-112
15850	NETTY-112	On a server which sent the first two messages, and received the last, I see this: NioAcceptedSocketChannel(id: 513d764c-011f-1000-b629-0170773a61bd, /127.0.0.1:62046 => /127.0.0.1:8081) - (message: Welcome to darkstar.local!) NioAcceptedSocketChannel(id: 513d764c-011f-1000-b629-0170773a61bd, /127.0.0.1:62046 => /127.0.0.1:8081) - (message: It is Sat Feb 07 16:56:40 SAST 2009 now.) NioAcceptedSocketChannel(id: 513d764c-011f-1000-b629-0170773a61bd, /127.0.0.1:62046 => /127.0.0.1:8081) - (message: eee) It's a bit hard to distinguish between incoming and outgoing messages, am I doing something wrong?	1	1	https://issues.jboss.org/browse/NETTY-112
15851	NETTY-112	I think you are not using the latest revision. Could you try again with a fresh build?	1	1	https://issues.jboss.org/browse/NETTY-112
15852	NETTY-112	Trying with latest Netty 3.1.0.ALPHA4 from svn: NioAcceptedSocketChannel(id: 564dad4e-011f-1000-b6d9-018537fa90cb, /127.0.0.1:51527 => /127.0.0.1:8081) - (state: OPEN) NioAcceptedSocketChannel(id: 564dad4e-011f-1000-b6d9-018537fa90cb, /127.0.0.1:51527 => /127.0.0.1:8081) - (state: BOUND) NioAcceptedSocketChannel(id: 564dad4e-011f-1000-b6d9-018537fa90cb, /127.0.0.1:51527 => /127.0.0.1:8081) - (state: CONNECTED) org.jboss.netty.channel.DownstreamMessageEvent@f1235b - (write: Welcome to darkstar.local!) org.jboss.netty.channel.DefaultWriteCompletionEvent@8c3679 - (writtenAmount: 27) org.jboss.netty.channel.DownstreamMessageEvent@659db7 - (write: It is Sun Feb 08 16:32:27 SAST 2009 now.) org.jboss.netty.channel.DefaultWriteCompletionEvent@556d12 - (writtenAmount: 41) org.jboss.netty.channel.UpstreamMessageEvent@3a9bba - (received: eeee) Is this correct?	1	1	https://issues.jboss.org/browse/NETTY-112
15853	NETTY-112	Shall I checkout and try again? Another thing, with MINA if you added the logging handler at the beginning of the filter chain, you were shown hex representation of all characters. It's helped me a few times. Could/should this be done with Netty?	1	1	https://issues.jboss.org/browse/NETTY-112
15854	NETTY-109	It seems like the current implementation is supposed to work with the Local transport only. Can we make this work with any ChannelFactories?	1	1	https://issues.jboss.org/browse/NETTY-109

#	Requirement	Questions	Answered	Successfully Answered	URL
15855	NETTY-109	My findings so far: 1) ServletClientSocketChannelFactory itself could work with any ClientSocketChannelFactory. For now, it seems like it uses old socket I/O. Can we change this to accept any ClientSocketChannelFactory as a constructor parameter? For example: ServletClientSocketChannelFactory factory = new ServletClientSocketChannelFactory(new NioClientSocketChannelFactory(...)); or ServletClientSocketChannelFactory factory = new ServletClientSocketChannelFactory(new OioClientSocketChannelFactory(...)); 2) IIUC, your implementation should work any HTTP server that understands the HTTP tunneling protocol you designed. Therefore, it would be nice if we can replace 'Servlet' with something else. For example, ServletClientSocketChannelFactory could be renamed to HttpTunnelClientChannelFactory. Same applies to the package names; org.jboss.netty.channel.socket.servlet could be renamed to org.jboss.netty.channel.socket.http. 3) A ChannelFactory is supposed to create more than one Channel with different settings. Therefore, you need to create a subtype of SocketAddress and specify it when you initiate a connection attempt, rather than specifying it in the factory constructor. For example: ClientBootstrap b = new ClientBootstrap(new HttpTunnelClientChannelFactory(...)); b.connect(new HttpTunnelAddress("http://example.com/nettyServlet"); ... 4) URI is preferred to URL - HttpTunnelAddress could look like this: public class HttpTunnelAddress extends SocketAddress { public URI getUri() { ... } }	1	1	https://issues.jboss.org/browse/NETTY-109
15856	NETTY-108	The initial implementation looks fairly simple and OK. Considering most rough places will be cleaned up, there are a couple tiny things I've found: * Perhaps we don't need to depend on Executor at all? (i.e. firing an event immediately for minimal latency) * We might need to maintain an unbounded queue to prevent stack overflow though (the pinging-ponging case.) * The package should be org.jboss.netty.channel.invm rather than org.jboss.netty.channel.socket.invm. * You can extend AbstractChannelSink rather than implementing ChannelSink directly. * InvmServerChannel could extend AbstractServerChannel instead of AbstractChannel. * Invm -> InVm? I'm not sure this is better. Or should we replace it with 'Local'? (i.e. LocalChannelFactory) WDYT?	1	1	https://issues.jboss.org/browse/NETTY-108
15857	NETTY-108	* Perhaps we don't need to depend on Executor at all? (i.e. firing an event immediately for minimal latency) ok, but i guess this still has to be done in a separate thread? * We might need to maintain an unbounded queue to prevent stack overflow though (the pinging-ponging case.) Again, we just kick of a separate thread that pulls from the queue and calls fireMessageReceived * The package should be org.jboss.netty.channel.invm rather than org.jboss.netty.channel.socket.invm. ok * You can extend AbstractChannelSink rather than implementing ChannelSink directly. ok * InvmServerChannel could extend AbstractServerChannel instead of AbstractChannel. no probs * Invm -> InVm? I'm not sure this is better. Or should we replace it with 'Local'? (i.e. LocalChannelFactory) WDYT? actually i think i prefer local!	1	1	https://issues.jboss.org/browse/NETTY-108
15858	NETTY-108	>> Perhaps we don't need to depend on Executor at all? (i.e. firing an event immediately for minimal latency) > ok, but i guess this still has to be done in a separate thread? If it needs to be done in a separate thread a user can add ExecutionHandler (in org.jboss.netty.handler.execution) in the pipeline, so we don't need to do this in a separate thread. >> We might need to maintain an unbounded queue to prevent stack overflow though (the pinging-ponging case.) > Again, we just kick of a separate thread that pulls from the queue and calls fireMessageReceived This issue can be avoided by suppressing recursive calls by maintaining a ThreadLocal flag and event queue (unbounded). For example, take a look at here: http://tinyurl.com/9z2q89 (Check the member variable 'notifying.') Would this be fine for you? >> Invm -> InVm? I'm not sure this is better. Or should we replace it with 'Local'? (i.e. LocalChannelFactory) WDYT? > actually i think i prefer local! Cool! Let's go with 'Local' then. :) It would be somewhat confusing with 'ChannelLocal', http://tinyurl.com/8n4y5v, but I don't think it's a big deal.	0	0	https://issues.jboss.org/browse/NETTY-108
15859	NETTY-108	I have fixed several compatiblity issues with existing transports like NIO socket transport. It should behave more realistically. :) On the other hand, I simplified how a local transport is instantiated. LocalAddresses are managed internally when you bind / connect, so you don't need to assign a name when a ChannelFactory is created anymore. This is a pretty big change that might surprise Andy, but it seems like the local example works just fine. (I had to fix a bug in the example's PrintHandler though.) What's left is to: * ensure some event ordering (i.e. don't fire message event until channelConnected() is fired) * make LocalChannel.getParent() return the parent * ensure that a LocalServerChannel is notified with childChannelOpen() and childChannelClosed() * make LocalChannel.toString() returns a string with correct local/remote addresses ('null => null' for now)	0	0	https://issues.jboss.org/browse/NETTY-108
15860	NETTY-79	That's also an option I'm thinking about. However, I'm somewhat concerned about the performance implications of running idleness detection in I/O thread. Idleness detection is performed periodically, but it scans through the whole list of open channels which can be time consuming (i.e. O(N) operation.) Therefore, at least I want to make it run in a separate thread with lower priority. Does it make sense to you? :)	0	0	https://issues.jboss.org/browse/NETTY-79
15861	NETTY-79	I'm not very confident that it will be a very time consuming task. However, it is still true that it takes a constant amount of time. A separate thread approach adds another overhead of maintaining a map of all open channels in a thread-safe way, which is likely to cause contention. I'm thinking of adjusting idleness check interval dynamically from 1 second to 3 seconds depending on the number of connections. WDYT?	0	0	https://issues.jboss.org/browse/NETTY-79
15862	NETTY-79	Just finished the implementation of HashedWheelTimer: http://tinyurl.com/7ka4ey I used Hashed Timing Wheel which is described here: http://www.cse.wustl.edu/~cdgill/courses/cs6874/TimingWheels.ppt What's left is ChannelHandlers that utilizes HashedWheelTimer (i.e. ReadTimeoutHandler and WriteTimeoutHandler.)	1	1	https://issues.jboss.org/browse/NETTY-79
15863	NETTY-79	The current implementation notifies your handler with an ExceptionEvent (ChannelReadTimeoutException or ChannelWriteTimeoutException). The cost of creating an exception is zero because it's pre-instantiated. Would this work for you?	1	1	https://issues.jboss.org/browse/NETTY-79
15864	NETTY-79	I must admit that I am not terribly familiar with Nettty (yet), and there do not seem to be jUnit tests for this now, could you please give a brief example of how to setup the bootstrap (and constructing instances of the *TimeoutHandler classes, this is what I didn't get right now) to be able to both receive read and write timeouts?	1	1	https://issues.jboss.org/browse/NETTY-79
15865	OSJC-261	[~maxandersen] In follow-up to the email trail regarding the use of proxy with this client. I have learned that okhttpclient when run with the openshift tooling is initialized with an 'EclipseProxySelector' and only honors the network settings from the eclipse preferences. It does not appear to honor your system network proxy. Is this acceptable? We could probably modify if needed to honor both, but this seems unnecessary based on the concerns you had.	1	1	https://issues.jboss.org/browse/OSJC-261
15866	OSJC-261	Interesting - I thought SPDY was needed for upgrading the connection securely ? or is that only relevant for browsers and we can do more custom logic on our end ?	1	1	https://issues.jboss.org/browse/OSJC-261
15867	OSJC-250	Patch looks clear and simple from my view. I'll work on adding some API in openshift-tools' RSync class to make use of it once this is accepted by [~jcantrill]. [~adietish] did you test it? Does it work? ;)	1	1	https://issues.jboss.org/browse/OSJC-250

#	Requirement	Questions	Answered	Successfully Answered	URL
15868	OSJC-132	Hi [~chirino] sorry for the late reply, this went somehow unnoticed in my radar. I'm not sure about the time we have to implement this, we're on heavy pressure to get tooling/client-lib for v3 done by this summer. Is this an absolut must or is there even a lucky chance where you could contribute this?	0	0	https://issues.jboss.org/browse/OSJC-132
15869	OSJC-101	If I query the geargroups of an application via {code} curl -H "Accept: application/json; version=1.2" --user <USER>:<PASSWORD> https://openshift.redhat.com/broker/rest/domains/foogoo/applications/php/gear_groups -X GET | json_reformat | less {code} I get the following response: {code} { "api_version": 1.2, "data": [{ "additional_gear_storage": 0, "base_gear_storage": 1, ... {code} [~bparees]: Are these the 2 properties that you need? Do you have any pointers to what the numbers exactly mean?	1	1	https://issues.jboss.org/browse/OSJC-101
15870	OSJC-95	As far as I can tell the enum sizes are not actually used anywhere, except in the tests. Can we just remove the enums and rewrite the tests? Nothing in OSJC or JBossTools-OpenShift seem to be dependent on checking them (or else weird sizes would show up in the lists now anyways...)	1	1	https://issues.jboss.org/browse/OSJC-95
15871	OSJC-93	Per discussion in IRC today, we have a working build for this, which is coupled to your existing openshift component build: https://hudson.qa.jboss.com/hudson/job/openshift-java-client-master/ https://hudson.qa.jboss.com/hudson/job/jbosstools-3.3_trunk.component--openshift/ How will openshift component pick up the new client jar? Is the assumption that it'll see it on the trunk composite site [1] and therefore build against that? If so, then you'll want to add your new component's staging folder [2] and previous staging folder [3] to the composite site in SVN [4], then kick an rsync [5] to push the SVN change to [1]. [1] http://download.jboss.org/jbosstools/builds/staging/_composite_/trunk/ [2] http://download.jboss.org/jbosstools/builds/staging/openshift-java-client-master/all/repo/ [3] http://download.jboss.org/jbosstools/builds/staging.previous/openshift-java-client-master/all/repo/ [4] https://svn.jboss.org/repos/jbosstools/trunk/download.jboss.org/jbosstools/builds/staging/_composite_/trunk/ [5] http://hudson.qa.jboss.com/hudson/view/DevStudio/view/DevStudio_All/job/jbosstools-download.jboss.org-rsync-from-svn/ If not, please clarify.	1	1	https://issues.jboss.org/browse/OSJC-93
15872	OSJC-93	-uh, if people want to fetch github sources and just build the single non-osgi jar, they TOTALLY can.- -The "pollution", as you call it, of having this upstream plugin packaged up as an osgi jar, wrapped w/ a feature, wrapped with a site.xml, published to our _composite_ site, meant we could consume it like we do any other upstream plugin.- -Why is that bad? And why would you rather the overhead of manually committing embedded jars in SVN when you can just pull them from upstream via our existing Jenkins build infra? You've railed for years about the suck of Modeshape/Teiid having to commit binaries into SVN, yet unless I'm mistaken, that's exactly what you're proposing here.- Can we move the OpenShift component code to a second github repo, so that the work done here can be reused? If nothing else, it would allow us to start gradually moving components out of SVN and into git.	1	1	https://issues.jboss.org/browse/OSJC-93
15873	OSJC-93	Andre - isn't this already done and the issue can be resolved ? What is left?	1	1	https://issues.jboss.org/browse/OSJC-93
15874	OSJC-59	The work that I did for this issue has been listed as done by work within a commit by develepercorey on OSJC-53, however no references to the refactors appear in OSJC-53. Is this still something I can finish up?	0	0	https://issues.jboss.org/browse/OSJC-59
15875	OSJC-45	java-client the test-suite hits many 500 "internal error" on INT (34 currently and started happening on the 2013-04-10). Most of them happen when the client tries to list the available applications: {"data":null,"errors":{},"messages":[{"exit_code":-1,"field":null,"severity":"error","text":"\nReference ID: f0c6ada8cd6dd8f36cb3abc1ce9bc487"}],"status":"internal_server_error","supported_api_versions":[1.0,1.1,1.2,1.3,1.4],"type":null,"version":"1.0"} at com.openshift.internal.client.httpclient.UrlConnectionHttpClient.createException(UrlConnectionHttpClient.java:204) at com.openshift.internal.client.httpclient.UrlConnectionHttpClient.write(UrlConnectionHttpClient.java:176) at com.openshift.internal.client.httpclient.UrlConnectionHttpClient.post(UrlConnectionHttpClient.java:146) at com.openshift.internal.client.httpclient.UrlConnectionHttpClient.post(UrlConnectionHttpClient.java:142) at com.openshift.internal.client.RestService.request(RestService.java:146) at com.openshift.internal.client.RestService.request(RestService.java:98) ... 68 more Caused by: java.io.IOException: Server returned HTTP response code: 500 for URL: https://int.openshift.redhat.com/broker/rest/domains/1366023591076/applications at sun.net.www.protocol.http.HttpURLConnection.getInputStream(HttpURLConnection.java:1459) at sun.net.www.protocol.https.HttpsURLConnectionImpl.getInputStream(HttpsURLConnectionImpl.java:234) at com.openshift.internal.client.httpclient.UrlConnectionHttpClient.write(UrlConnectionHttpClient.java:174) ... 72 more The very same is not happening neither on STG nor on PROD. I was wondering if there's any particular reason for this, anything wrong on my side or is it STG not being stable? Another frequent error that keeps popping up here and there (on PROD and INT, very rarely on INT currently) is the following: com.openshift.client.OpenShiftEndpointException: Could not request https://openshift.redhat.com/broker/rest/user: Operation failed.Reason given: "OpenShift::AuthServiceException" ... at org.codehaus.plexus.classworlds.launcher.Launcher.main(Launcher.java:352) Caused by: com.openshift.internal.client.httpclient.HttpClientException: {"data":null,"messages":[{"exit_code":null,"field":null,"severity":"error","text":"OpenShift::AuthServiceException"}],"status":"internal_server_error","supported_api_versions":[1.0,1.1,1.2,1.3,1.4],"type":null,"version":"1.0"} at com.openshift.internal.client.httpclient.UrlConnectionHttpClient.createException(UrlConnectionHttpClient.java:204) at com.openshift.internal.client.httpclient.UrlConnectionHttpClient.get(UrlConnectionHttpClient.java:113) at com.openshift.internal.client.RestService.request(RestService.java:144) at com.openshift.internal.client.RestService.request(RestService.java:98) ... 67 more Caused by: java.io.IOException: Server returned HTTP response code: 500 for URL: https://openshift.redhat.com/broker/rest/user at sun.net.www.protocol.http.HttpURLConnection.getInputStream(HttpURLConnection.java:1459) at sun.net.www.protocol.https.HttpsURLConnectionImpl.getInputStream(HttpsURLConnectionImpl.java:234) at com.openshift.internal.client.httpclient.UrlConnectionHttpClient.get(UrlConnectionHttpClient.java:111) Thanks! André {quote} I started talking to Lili Nader which told me to look the operation references reported in the error up in the broker logs. I answered her that I could not do that since I had no root access on INT. We need to further follow up on this.	0	0	https://issues.jboss.org/browse/OSJC-45

#	Requirement	Questions	Answered	Successfully Answered	URL
15876	OSJC-44	pushed to master and rebased. verified on jenkins which now reports: {code} com.openshift.client.OpenShiftEndpointException: Could not request https://stg.openshift.redhat.com/broker/rest/api: java.net.ConnectException: Connection timed out at com.openshift.internal.client.RestService.request(RestService.java:106) at com.openshift.internal.client.RestService.request(RestService.java:91) at com.openshift.internal.client.RestService.request(RestService.java:72) at com.openshift.internal.client.AbstractOpenShiftConnectionFactory.getConnection(AbstractOpenShiftConnectionFactory.java:34) at com.openshift.client.OpenShiftConnectionFactory.getConnection(OpenShiftConnectionFactory.java:161) {code}	1	0	https://issues.jboss.org/browse/OSJC-44
15877	PLANNER-840	There is a trade-off between - adding more information to the trace log line - verbosity of the trace log line. What are the use cases to need those 2 additional information you need? (When in doubt, I 'd leave it out). There is 3) case btw: it's not just non doable moves that aren't mentioned, but also filtered moves/entities/values aren't mentioned (which would have performance implications if we track that).	0	0	https://issues.jboss.org/browse/PLANNER-840
15878	PLANNER-810	[~ederign] Datamodeler also can't copy-paste urls into a new browser, although the content of the url does seem to indicate that it's being trying to work, could this be an general uberfire problem?	1	0	https://issues.jboss.org/browse/PLANNER-810
15879	PLANNER-793	{quote}{Unit test needed for if a getter is overriden and the parent was autodetect but the child has an explicit annotation - and vica versa}{quote} What are the use cases for this? Would you for example want to override a {{getFactList()}} getter detected as a {{@ProblemFactCollectionProperty}} to {{@ProblemFactProperty}}? Anything else? I can only think of changing an entity property to problem fact property but that doesn't make sense to me, so I don't think it needs to be tested.	0	0	https://issues.jboss.org/browse/PLANNER-793
15880	PLANNER-790	1) I think this duplicates the open issue PLANNER-151 (which support not just XML, but like all XML have a java Config API alternative approach). 2) Covered already by EasyJavaScoreCalculator? 3) Covered already by the programmatic api (see docs chapter 4)? What do you think (after researching those pointers above)? Can this issue be closed as duplicate?	0	0	https://issues.jboss.org/browse/PLANNER-790
15881	PLANNER-763	[~jsniezek] Hey Joe, have you had a chance to look at the result of the "import example" dialog of Matej? Your feedback is welcome. Mine is above. The feature itself is great, this is just some polishing.	1	1	https://issues.jboss.org/browse/PLANNER-763
15882	PLANNER-707	The old repo's are here: https://github.com/guvnorngtestuser1?tab=repositories https://github.com/mbarkley/app-former-playground Don't remove them (would break 6.x) but clearly add a note that says "MOVED to https://github.com/droolsjbpm/kie-wb-playground"	0	0	https://issues.jboss.org/browse/PLANNER-707
15883	PLANNER-695	bq. Just heard of a case where they had non-unique id's on 2 different instances with different content of the same class. So that should definitely fail fast So this means it is forbidden to add two working objects with the same ID unless it is the same instance, right? Even if the two different instances are equal it should fail. It could also happen that the two instances have same ID, are equal and still there is some difference in their content (fields) which is not covered by equals(). This likely a bug in user's code but I guess we don't want to detect this case. On the other hand adding the same instance twice should pass or fail? If it's added it basically means no change but still there may be a reason to fail for example to detect a bug on a higher level. Lastly, how should this work with EQUALITY strategy? Adding the same instance twice is OK but adding two different instances that are equal should fail? This could detect bugs in user's equals() implementation.	0	0	https://issues.jboss.org/browse/PLANNER-695
15884	PLANNER-689	Agreed, if it doesn't change the 1 depot files (all current vrp files), which it doesn't I think. Would you be willling to send in a Pull Request for this change?	1	0	https://issues.jboss.org/browse/PLANNER-689
15885	PLANNER-667	Neither of these solutions takes into account potential nullity of the fields, nor do they allow distinguishing between natural ordering and user-provided ordering. In the past, writing Comparators was painful and people hated it. However, the Java 8 `Comparator` convenience methods are awesome. They handle all the corner cases and they are simple enough to use. Let's not over-engineer it, let's not bring annotation scanning into it, let's not introduce our own logic - let's have people simply write a Java 8 one-liner comparator.	0	0	https://issues.jboss.org/browse/PLANNER-667
15886	PLANNER-666	[~jlocker] Thanks for pointing this out Jiri. Yes, the change should be there, could you please a JIRA to track this? Thanks.	0	0	https://issues.jboss.org/browse/PLANNER-666
15887	PLANNER-664	This report looks valid. Can you also turn on DEBUG and TRACE logging and add that here? That will give me full confidence that it's a valid bug.	1	1	https://issues.jboss.org/browse/PLANNER-664
15888	PLANNER-664	I'm currently using Optaplanner 6.2. I have compiled a 6.2 version patched with your modifications. When I run the test program I attached to this issue with this 6.2 patched version, no hard rules are broken anymore but no process are placed. However 150 processes should be placed on the 75 hosts. Is it normal?	1	0	https://issues.jboss.org/browse/PLANNER-664
15889	PLANNER-659	[~ge0ffrey]: can you please check if there is now everything that should be : http://docs.jboss.org/optaplanner/release/6.4.0.Final/optaplanner-javadoc/index.html	1	1	https://issues.jboss.org/browse/PLANNER-659
15890	PLANNER-659	So it looks good. Is there a commit or PR with the changes to the assembly so all future versions are fixed too?	1	1	https://issues.jboss.org/browse/PLANNER-659
15891	PLANNER-642	Is your patch available on a github repo too (feel free to send it as a PR)? This bug doesn't surpise me - and we have to fix it of course - but the manner is difficult. *Workaround*: in the DistanceMeter, make it so the distance to/from any depot is 0, sot hey are sorted before the Customers (stops).	1	0	https://issues.jboss.org/browse/PLANNER-642
15892	PLANNER-642	[~ge0ffrey], thanks for your response. Created a pull request with the content of the patch (can you also delete the older attachment, it missed some code). Used the issue key in the PR message, so you may have to fix the automatic transition this issue made. As for the workaround: this will indeed work for now, but we have to make sure the distribution size maximum is far above the number of depots (in our own problem all the vehicles have their own starting location)	0	0	https://issues.jboss.org/browse/PLANNER-642
15893	PLANNER-620	3) My personal taste is a separate class, but I also fine with a *static* inner class. I've added some code above in the description. Would it be possible to generate that? 5) not a big fan of this, due to same reasons at 1: the comparator should really be compiled at build time, not runtime.	0	0	https://issues.jboss.org/browse/PLANNER-620
15894	PLANNER-607	This requires further discussion [~ge0ffrey] Thinking about it now, not sure whether we should use different model at all. Ideally we'd like to test whether the configuration works correctly as a whole. If we use different data model, it can get out of sync with rule definitions, value selectors, etc. wdyt?	0	0	https://issues.jboss.org/browse/PLANNER-607

#	Requirement	Questions	Answered	Successfully Answered	URL
15895	PLANNER-575	webexamples --- [WARNING] Error injecting: org.codehaus.plexus.archiver.jar.JarArchiver java.lang.ExceptionInInitializerError at jdk.internal.reflect.NativeConstructorAccessorImpl.newInstance0(java.base@9-ea/Native Method) at jdk.internal.reflect.NativeConstructorAccessorImpl.newInstance(java.base@9-ea/NativeConstructorAccessorImpl.java:62) at jdk.internal.reflect.DelegatingConstructorAccessorImpl.newInstance(java.base@9-ea/DelegatingConstructorAccessorImpl.java:45) at java.lang.reflect.Constructor.newInstance(java.base@9-ea/Constructor.java:453) at com.google.inject.internal.DefaultConstructionProxyFactory$1.newInstance(DefaultConstructionProxyFactory.java:86) at com.google.inject.internal.ConstructorInjector.provision(ConstructorInjector.java:108) at com.google.inject.internal.ConstructorInjector.access$000(ConstructorInjector.java:32) at com.google.inject.internal.ConstructorInjector$1.call(ConstructorInjector.java:92) at com.google.inject.internal.ProvisionListenerStackCallback$Provision.provision(ProvisionListenerStackCallback.java:115) at com.google.inject.internal.ProvisionListenerStackCallback$Provision.provision(ProvisionListenerStackCallback.java:133) at com.google.inject.internal.ProvisionListenerStackCallback.provision(ProvisionListenerStackCallback.java:68) at com.google.inject.internal.ConstructorInjector.construct(ConstructorInjector.java:90) at com.google.inject.internal.ConstructorBindingImpl$Factory.get(ConstructorBindingImpl.java:269) at com.google.inject.internal.InjectorImpl$2$1.call(InjectorImpl.java:1009) at com.google.inject.internal.InjectorImpl.callInContext(InjectorImpl.java:1066) at com.google.inject.internal.InjectorImpl$2.get(InjectorImpl.java:1005) at com.google.inject.internal.InjectorImpl.getInstance(InjectorImpl.java:1044) at org.eclipse.sisu.space.AbstractDeferredClass.get(AbstractDeferredClass.java:48) at com.google.inject.internal.ProviderInternalFactory.provision(ProviderInternalFactory.java:86) at com.google.inject.internal.InternalFactoryToInitializableAdapter.provision(InternalFactoryToInitializableAdapter.java:54) at com.google.inject.internal.ProviderInternalFactory$1.call(ProviderInternalFactory.java:70) at com.google.inject.internal.ProvisionListenerStackCallback$Provision.provision(ProvisionListenerStackCallback.java:115) at com.google.inject.internal.ProvisionListenerStackCallback$Provision.provision(ProvisionListenerStackCallback.java:133) at com.google.inject.internal.ProvisionListenerStackCallback.provision(ProvisionListenerStackCallback.java:68) at com.google.inject.internal.ProviderInternalFactory.circularGet(ProviderInternalFactory.java:68) at com.google.inject.internal.InternalFactoryToInitializableAdapter.get(InternalFactoryToInitializableAdapter.java:46) at com.google.inject.internal.InjectorImpl$2$1.call(InjectorImpl.java:1009) at com.google.inject.internal.InjectorImpl.callInContext(InjectorImpl.java:1066) at com.google.inject.internal.InjectorImpl$2.get(InjectorImpl.java:1005) at com.google.inject.internal.SingletonScope$1.get(SingletonScope.java:36) at	0	0	https://issues.jboss.org/browse/PLANNER-575
15896	PLANNER-551	Walter, Paulo's looking at the UF support for this (UF-288), once that's done - making this issue easier to fix, would you mind taking a look at this one?	1	1	https://issues.jboss.org/browse/PLANNER-551
15897	PLANNER-551	Hello Joe, Would you mind take a look and provide the corresponding icons for the "drools&jbpm", "advanced", and "persistence" tool windows? Regards, Walter.	1	1	https://issues.jboss.org/browse/PLANNER-551
15898	PLANNER-505	Geoffrey, I am trying to figure out the service API. I think the basic format would be: To get the list of the solvers: {code}GET /services/rest/server/containers/instances/<containerId>/solvers {code} To create a solver: {code}PUT /services/rest/server/containers/instances/<containerId>/solvers/<solverId> {code} To update a solver (to start a solver, i.e., call the solve() method, or to call terminateEarly()) {code}POST /services/rest/server/containers/instances/<containerId>/solvers/<solverId> {code} To get the solver status: {code}GET /services/rest/server/containers/instances/<containerId>/solvers/<solverId> {code} To get the best solution: {code}GET /services/rest/server/containers/instances/<containerId>/solvers/<solverId>/bestsolution {code} To dispose a solver: {code}DELETE /services/rest/server/containers/instances/<containerId>/solvers/<solverId>{code} I don't think it would be possible, due to the stateful nature of each solver, to create more than one instance of each solver configuration on the same container. Do you agree? I am not sure how to prevent it though, so it would be the user's responsibility.	1	1	https://issues.jboss.org/browse/PLANNER-505
15899	PLANNER-505	Hi Edson, In a multi-tenancy scenario, it's common to have multiple Solver instances of the same solver configuration running in parallel. Think hospital nurse rostering "as a service". "To get the solver status" - I presume this returns the best solution (which includes the best score). We can implement this adding Solver.addEventListener() that listens to new best solution events (which is on the solver thread) and stores it in an AtomicReference, so that GET rest call (which is on a servlet thread) can simply get the latest best solution from that AtomicReference and marshall it back. (Note the best solution is already "planning cloned" (see docs) by the solver itself, so most race conditions are impossible, but some with ProblemFactChanges (not supported in this phase) will be in the future). Let's meet up for a real-time discussion?	1	1	https://issues.jboss.org/browse/PLANNER-505
15900	PLANNER-505	Geoffrey, thinking further about the operations and which data they require, is this correct? * create a solver: all it requires is the name of the solver configuration file? Is it possible that we will have to support other parameters in the future? * start the solver: requires the input data. Anything else? If a solver is stopped later on, can it be re-started? Can I reuse the same solver with a different input data set? * check the status: what are the possible status a solver can be? IDLE, SOLVING, FINISHED? anything else? * get solution: this will return the best solution computed up to that moment, correct? * stop: I think you called this "early termination". Is any parameter necessary? * dispose: does not require any data.	1	1	https://issues.jboss.org/browse/PLANNER-505
15901	PLANNER-485	How about implementing some MRU cache, of a configurable size? Seems easier (and faster) than to try to create some decision logic around entities.	0	0	https://issues.jboss.org/browse/PLANNER-485
15902	PLANNER-485	how about use 2Dim implementation of [1]? it's efficient structure to store data in space. then we could find in log time good candidates for selection. Or am I missing something? [1] https://en.wikipedia.org/wiki/K-d_tree	0	0	https://issues.jboss.org/browse/PLANNER-485
15903	PLANNER-485	+1for the idea of also also support Most Recently Used cache strategies somehow, or maybe just a Most Nearby cache, and swapping the rest out to disk. Implementing this in a reliable, platform independent matter won't be easy. Where's the disk file? What happens in an openshift environment?	1	1	https://issues.jboss.org/browse/PLANNER-485
15904	PLANNER-485	Can we leave that to Infinispan? Let's not invent a cache where we already have one.	1	1	https://issues.jboss.org/browse/PLANNER-485

#	Requirement	Questions	Answered	Successfully Answered	URL
15905	PLANNER-478	[~jlocker] Would it be possible to check the performance regressions tests impact of these changes? especially on chained vars and tabu search. See commit on the right to see actual impact.	1	1	https://issues.jboss.org/browse/PLANNER-478
15906	PLANNER-467	We already have these distributions, as we're just piggy-backing on the kie distributions (but our stuff is masked out with the role plannermgmt). We need to look at how we'll expose this in docker. Shall we get a separate docker image for optaplanner which is basically the same as the drools one with an extra role?	0	0	https://issues.jboss.org/browse/PLANNER-467
15907	PLANNER-423	LookAndFeels to investigate: - Nimbus - Part of JDK IIRC? What is the minimum JDK version? We need to support 6. Yes, it's part of 6: https://docs.oracle.com/javase/tutorial/uiswing/lookandfeel/nimbus.html - Sustance http://insubstantial.github.io/insubstantial/substance/ - http://stackoverflow.com/questions/10134516/how-can-i-create-a-modern-looking-java-desktop-application - http://stackoverflow.com/questions/2934337/open-source-laflook-n-feel-for-java-swing-application - JGoodies. Has aged IIRC. - Which ones does IntelIJ support? See it's settings. Several of them are not open source (such as Alloy IIRC). The LaF will probably be a separate dependency jar, so - It needs to be in Maven Central (unless it's in the JDK) - It needs to be a compatible open source license (ASL, BSD, MIT) or maybe LGPL but definitly not GPL. See readme's dependency reqs list for the full req list. Also do a full build and test the runExamples.sh script. Normally that should all work out now with the new assembly xml's that put that new dependency jar in the correct folder, which should be /examples/binaries, not /binaries.	0	0	https://issues.jboss.org/browse/PLANNER-423
15908	PLANNER-411	Didn't test with any specific version (just an older master, before the uninitializedVariableCount change) and I get this exception when running the aggregator: {code} Exception in thread "main" java.lang.NullPointerException at org.optaplanner.benchmark.impl.result.SolverBenchmarkResult.determineTotalsAndAverages(SolverBenchmarkResult.java:310) at org.optaplanner.benchmark.impl.result.SolverBenchmarkResult.accumulateResults(SolverBenchmarkResult.java:289) at org.optaplanner.benchmark.impl.result.PlannerBenchmarkResult.accumulateResults(PlannerBenchmarkResult.java:288) at org.optaplanner.benchmark.impl.aggregator.swingui.BenchmarkAggregatorFrame.initPlannerBenchmarkResultList(BenchmarkAggregatorFrame.java:236) at org.optaplanner.benchmark.impl.aggregator.swingui.BenchmarkAggregatorFrame.init(BenchmarkAggregatorFrame.java:129) at org.optaplanner.benchmark.impl.aggregator.swingui.BenchmarkAggregatorFrame.createAndDisplay(BenchmarkAggregatorFrame.java:96) at org.optaplanner.examples.common.app.CommonBenchmarkApp.buildAndBenchmark(CommonBenchmarkApp.java:71) at org.optaplanner.examples.app.GeneralOptaPlannerBenchmarkApp.main(GeneralOptaPlannerBenchmarkApp.java:24) {code} The uninitializedVariableCount in SingleBenchmarkResult is null.. this could be fixed by using some sort of a default value.	1	1	https://issues.jboss.org/browse/PLANNER-411
15909	PLANNER-401	* The separate JVM requirement was dropped, see PLANNER-428 * Naming SubSingleBenchmarks ** Make a vote + proposals * Find a way for the user to set the number of benchmarks in the config (default to one), run it that many times. ** Make a vote the exact syntax and naming + proposals * Aggregate the results in the report: drop down menu in report to display individual runs, avg, med, maybe stddev (also display a violin (box) plot - where?). * Choose one of the statistics (average) and use that as the outcome of the single benchmarks -> Do not change the SingleBenchmarkResult interface, apart from renaming. * Merging SubSingleStatistics ** Make a vote (reference to meeting with Geoffrey) + proposals * Details & questions ** What about uninitialized subsingles? *** Handle them just like scores, add a link to PLANNER-405 where applicable. ** What about failed subsingles? *** Handle them just like singles, propagate up (overrides score). ** What about violin plot not supported OOTB in JFreeChart? *** [Violin plots vs box plots|http://www.cerebralmastication.com/2009/02/box-plot-vs-violin-plot-in-r/] *** Ask on their forums, create a JIRA to follow up on this. *** Use Box plot for now, migrate later, see PLANNER-429 and http://stats.stackexchange.com/questions/136160/jfreechart-boxplot-appearance	1	1	https://issues.jboss.org/browse/PLANNER-401
15910	PLANNER-401	repeated single benchmark runs? * A1: SubSingleBenchmark ** Pro: Falls into it's place in the hierarchy nicely ** Con: The name sounds weird ** *Votes*: +1 by [~ge0ffrey], [~oskopek] * A2: SingleBenchmarkRepetition ** Pro: The name is logical and concise ** Con: Doesn't fall into it's place in the hierarchy very intuitively ** *Votes*: +1 by [~ge0ffrey], +0 by [~oskopek] * A3: Rename SingleBenchmark to RepeatedSingleBenchmark (or anything else), rename SubSingleBenchmark to SingleBenchmark ** Pro: Could improve the readability/understandability considerably ** Con: A lot of renaming/rewriting in the code and docs ** Con: Will confuse users/developers for a while ** *Votes*: +0 by [~ge0ffrey], +0.5 by [~oskopek] * An: ... (further proposals welcome) *Vote B*: How do we name the configuration setting that lets the user configure the number of run repeats of single benchmarks? (prefixes depend on where we place the setting in the benchmark config) * B1: repeatCount ** *Votes*: +0 by [~ge0ffrey], [~oskopek] * B2: repeats ** *Votes*: -1 by [~ge0ffrey], [~oskopek] * B3: subSingleCount (adjusted to the result of Vote A) ** *Votes*: +1 by [~ge0ffrey], [~oskopek] * Bn: ... (further proposals welcome) *Vote C*: Where do we let the user configure the number of run repeats of single benchmarks? * C1: In solverBenchmark->problemBenchmarks ** Pro: can be inherited ** Pro: can be set per solverBenchmark ** Con: too deep? ** Con: cannot be set individually per singleBenchmark ** *Votes*: -0.5 by [~ge0ffrey], -1 by [~oskopek] * C2: In solverBenchmark ** Pro: can be inherited ** Pro: can be set per solverBenchmark ** Con: too high in the hierarchy, given it references single benchmarks? ** Con: cannot be set individually per singleBenchmark ** *Votes*: +1 by [~ge0ffrey], [~oskopek] * C3: In plannerBenchmark (globally) ** Con: cannot be inherited ** Con: cannot be set per solverBenchmark ** Con: too high in the hierarchy, given it references single benchmarks ** Con: cannot be set individually per singleBenchmark ** *Votes*: -1 by [~ge0ffrey], [~oskopek] * C4: (Somehow) per singleBenchmark with a default per solver benchmark ** Pro: can be inherited ** Pro: can be set per solverBenchmark ** Con: too low in the hierarchy, makes the configuration too granular and confusing? ** Pro: can be set individually per singleBenchmark ** *Votes*: -0.5 by [~ge0ffrey], [~oskopek] * Cn: ... (further proposals welcome) *Vote D*: How do we merge (Sub) SingleStatistics of the repeated single benchmark runs? * D1: We only look at the first subsingle and ignore the rest for the statistics ** Con: Arbitrary choice ** Con: Doesn't fit with the summary stats ** *Votes*: -1 by [~oskopek] * D2: Take the median subsingle. This is a real one and it's not an edge case. Median according to SingleBenchmarkRankingComparator (not user pluggable). If subsingle count is even, take the worse one of the 2 middle ones. ** Pro: Finds a _representative_ benchmark run that is also real at the same time ** Pro: Easier to implement ** Con: The median is not the average, so it will not be the same as the summary statistics ** *Votes*: +1 by [~oskopek] * D2b: Adjust summaries to use median ** Con: if hard broken or uninitialized in 1 subsingle, it isn't shown in the summary statistics ** *Votes*: -1 by [~ge0ffrey], [~oskopek] * D2c: Put median in dropdown (in addition to the average summary statistics (score)) ** Pro: Definitely better than not doing this ** *Votes*: +1 by [~oskopek] * D3: Average the single statistic points with some sort of moving average ** Pro: Is an average case, is indicative of the underlying individual run results ** Con: Isn't real time-lined data, the data doesn't tell the user much ** *Votes*: -0.5 by [~oskopek] * D4: [~jvahala]: Combine D2 with D3: Use the median, add a moving arithmetic average (or other sensible statistic), plot them on the same chart with the min/max, add distribution shading (frequency of values in a given area on the chart) ** Pro: Is quite descriptive about the average case, outliers and shows the distribution between them ** Con: Is the result worth the (implementation) cost? JFreeChart probably doesn't support shading in the way we want. ** Details: [in this comment|https://issues.jboss.org/browse/PLANNER-401?focusedCommentId=13105265&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-	1	1	https://issues.jboss.org/browse/PLANNER-401

#	Requirement	Questions	Answered	Successfully Answered	URL
15911	PLANNER-401	* Where do we let the user configure the number of run repeats of single benchmarks? ** +1 "C2: In solverBenchmark" Seems perfect to me ** -1 "C3: In plannerBenchmark (globally)" because doing it with C2 in <inheritedSovlerBenchmark> is the same ** -0.5 C4 and C1: overkill. YAGTNI. Or show me a use case to disprove this :) Where/when does a user want a difference in subsingle repeat count? * Between 2 solver configs if the first solver config doesn't rely on anything random (First Fit Decreasing, Custom initializer, etc) and the second does (Late Acceptance) * That's it. I see no reason to ever use a different repeat count per dataset.	0	0	https://issues.jboss.org/browse/PLANNER-401
15912	PLANNER-387	Are you sure the scoreDrl string is correct? That exception happens if it can't find the classpath resource from the classloader argument: https://github.com/droolsjbpm/optaplanner/blob/master/optaplanner-core/src/main/java/org/optaplanner/core/config/score/director/ScoreDirectorFactoryConfig.java#L384 Can you try doing a toString() of the classloader you're giving to build the SolverFactory. Does it equal "org.optaplanner.core [280]" or something else?	0	0	https://issues.jboss.org/browse/PLANNER-387
15913	PLANNER-387	are you sure that just fixing tests will resolve failing issue? ibm classloader is inconstistent with other classloaders in attributes and methods, I am not sure if methods which planner uses are same.	1	1	https://issues.jboss.org/browse/PLANNER-387
15914	PLANNER-372	Isn't this one resolved?	1	1	https://issues.jboss.org/browse/PLANNER-372
15915	PLANNER-368	[~jsniezek] I attached sample screenshots (termination.zip) of current implementation. Initially the tree is empty (termination_empty.png). Selecting "Add termination" option on the root/nested termination element displays a popup (termination_add_popup.png) where different termination types and their values can be picked. The results are added to the tree (termination_items.png). Nested terminations can be expanded/collapsed. All items but the root termination element can be removed. Removing nested termination element removes the whole subtree of its children. Could you please share your ideas from UX perspective?	1	0	https://issues.jboss.org/browse/PLANNER-368
15916	PLANNER-359	is specifying probability distribution optional here?	1	1	https://issues.jboss.org/browse/PLANNER-359
15917	PLANNER-359	and what if I change SORTED to RANDOM and add probability distribution? It still fail fast for me, but I don't see reason why.	1	1	https://issues.jboss.org/browse/PLANNER-359
15918	PLANNER-359	Which code snippets fail fast for you?	1	0	https://issues.jboss.org/browse/PLANNER-359
15919	PLANNER-359	I think the code above fails because the CH Placer can not work with never ending entity selectors, such as selectionOrder RANDOM. The rest of that config isn't relevant any more (including the nearby stuff). What does the error message state? But the valueselector with nearby selection itself, doesn't care what selectionOrder the entitySelector is using. That entity selectionOrder doesn't impact (as in "prohibits options in") the value selector.	0	0	https://issues.jboss.org/browse/PLANNER-359
15920	PLANNER-357	Richard also agrees with Jiri: "we have to define twice if we choose to have SimpleScore or HardScoreScore : once in the java class implements definition, and a 2nd time in the XML configuration file. Of course, an error is raised if there is a discrepency between the 2. I was wondering if it would be feasible to only define it once?" We should find a way to fix this. Proposal A) Remove the solver config line. The problem is that *the domain class defines the Score class, while the solver config (xml file usually) defines the ScoreDefinition instance*. A ScoreDefinition is the factory (and more) of Score. But Score instances (such as HardSoftScore) don't have a reference to their ScoreDefinition (such as HardSoftScoreDefinition) because they need to behave like super lightweight value classes (such as Long, Double, BigDecimal, LocalDate) when being serialized and exchanged. The real problem is that a ScoreDefinition _instance_ exists at bootstrap time, but _no_ Score instance exists as bootstrap time. Let me explain differently: Is it possible to determine a ScoreDefinition instance from looking at a Score instance? Yes. But we don't have a Score instance in the domain model at configuration time (there is no Solution instance yet, only a class). We do have a Score class (as a generic type even - but we can actually extract that despite erasure). Is it possible to determine a ScoreDefinition instance from looking at a specific Score class? No. For example, a BendableScore(Definition) has a hard and soft levelCount int. Those 2 ints are only known by the instances, not by the classes, therefor a Score class cannot unambiguously define the Score instance or ScoreDefinition instance.	1	1	https://issues.jboss.org/browse/PLANNER-357
15921	PLANNER-356	Do we want to do this?	1	1	https://issues.jboss.org/browse/PLANNER-356
15922	PLANNER-356	So, this is my idea of implementation: * All getters annotated with ```@ProblemFactCollectionProperty``` must return a collection type. All elements in these collections would be added among facts. * All getters annotated with ```@ProblemFactProperty``` will have their return values added among facts. * Support for ```getProblemFacts()``` will be removed altogether. Since the ```Solution``` interface is going away shortly after this is done (PLANNER-436), users will have to make compile-time changes to their code - and therefore will move to use the annotations. * This JIRA will not deal with the removal of ```Solution```. Instead, it will be limited to replacing the getter with annotations. What do others think?	1	1	https://issues.jboss.org/browse/PLANNER-356
15923	PLANNER-352	Quick question... what happens when I add the annotation on both the variable and the getter for the variable? And what if those two have different value range providers?	1	0	https://issues.jboss.org/browse/PLANNER-352
15924	PLANNER-343	That NoSuchMethodError is wierd. *Such an error only happens if the runtime classpath differs from the compilation classpath.* This line compiles, so during compilation it's correct. As of 6.3.0-SNAPSHOT, we're using reflections 0.9.10 {code} builder.addUrls(ClasspathHelper.forPackage(packageInclude, (ClassLoader[]) null)); {code} Are your tests using the classpath of 6.3.0-SNAPSHOT and therefore 0.9.10? Note that in 6.2, kie used an older reflections version, 0.9.9-RC1 or something.	1	1	https://issues.jboss.org/browse/PLANNER-343

#	Requirement	Questions	Answered	Successfully Answered	URL
15925	PLANNER-343	automatic domain scan should scan only classes in <packageInclude> right? or does it look in upper level? I have two nqueens domains and when I specify <packageInclude>org.jboss.qa.brms.planner.examples.nqueens</packageInclude> it gives me: {code}java.lang.IllegalStateException: The scanAnnotatedClasses (ScanAnnotatedClassesConfig([org.jboss.qa.brms.planner.examples.nqueens])) found multiple classes ([class org.jboss.qa.brms.planner.examples.nqueens.domain.NQueens, class org.jboss.qa.brms.planner.examples.nqueenscontplanning.domain.NQueens]) with a PlanningSolution annotation. at org.optaplanner.core.config.domain.ScanAnnotatedClassesConfig.loadSolutionClass(ScanAnnotatedClassesConfig.java:81) at org.optaplanner.core.config.domain.ScanAnnotatedClassesConfig.buildSolutionDescriptor(ScanAnnotatedClassesConfig.java:68) at org.optaplanner.core.config.solver.SolverConfig.buildSolutionDescriptor(SolverConfig.java:228) at org.optaplanner.core.config.solver.SolverConfig.buildSolver(SolverConfig.java:179) at org.jboss.qa.brms.planner.agile.AutoDetectSolutionDomainTest.testAutoNQueensDomainDetection(AutoDetectSolutionDomainTest.java:27){code} but "nqueenscontplanning" is different package and has nothing to do with "nqueens" package. If I specify package more strictly: <packageInclude>org.jboss.qa.brms.planner.examples.nqueens.domain</packageInclude> planner works.	1	0	https://issues.jboss.org/browse/PLANNER-343
15926	PLANNER-332	Question: What do we use the checkstyle plugin for? Maybe we can even just remove it? (although I doubt that)	1	1	https://issues.jboss.org/browse/PLANNER-332
15927	PLANNER-328	Is the workaround possible without extending the entire SolverConfig class? If I try to just inject my customScoreDirectorFactory into the XML I get an error saying SolverConfig expects only the scoreDirectorFactory tag there. I've tried modifying the SolverConfig and removing the XStreamAlias from the ScoreDirectorFactoryConfig field. Then aliasing that field to my class manually using stream but it will still only load the ScoreDirectorFactoryConfig class and not my custom one.	1	0	https://issues.jboss.org/browse/PLANNER-328
15928	PLANNER-328	I think I did it according to the instructions. I extended ScoreDirectorFactoryConfig, gave it an alias and put the alias into the XML file instead of scoreDirectorConfig. Is it possible that this is because ScoreDirectorFactoryConfig is not abstract nor an interface? I can post my code(sanitized) to the developer forum post I made here: https://groups.google.com/forum/#!topic/optaplanner-dev/3s-9kvYxP40 if that will help.	0	0	https://issues.jboss.org/browse/PLANNER-328
15929	PLANNER-328	I found the proper way to inject a child class into the configuration: {code:title=OptManager.java|borderStyle=solid} XStreamXmlSolverFactory solverFactory = new XStreamXmlSolverFactory(); solverFactory.aliasAttribute(SolverConfig.class, "scoreDirectorFactoryConfig", "flexScoreDirectorFactory"); solverFactory.alias("flexScoreDirectorFactory", FlexScoreLevelDirectoryFactoryConfig.class); solverFactory.addDefaultImplementation(FlexScoreLevelDirectoryFactoryConfig.class, ScoreDirectorFactoryConfig.class); {code} This does require modifying the existing XStreamXMLSolverFactory class to add aliasAttribute, alias and addDefaultImplementation methods. Can these be added to that class?	0	0	https://issues.jboss.org/browse/PLANNER-328
15930	PLANNER-325	This is a great start guys. A quick suggestion. Is there a way to hide Planner tab for field variables and add it dynamically only if the Entity is selected as Planning Entity ? Also, as an extension, even for the Entity, not keeping the planner tab by default. Perhaps have a checkbox in the top or something, that would enable or hide the tab. Last thing.. "Planner" in the tab is not sure if it is a right verbeage. Not mandatory we should come up with something. But worth thinking about a relevant title for that tab.	1	1	https://issues.jboss.org/browse/PLANNER-325
15931	PLANNER-322	When should solverStarted() be called? Do you require any additional data info besides that it's started? Would do this code cover your requirements? {code} public void run() { solverStarted(); solver.solve(); solverTerminated(); } {code} Reason I ask, is because inside the solver, the internal SolverLifecycleListener.solverStarted() event happens much later, after the initial solution has been scored and cloned. And it's even called multiple times if Solver.addProblemFactChange() is used. Same applies to SolverLifecycleListener.solverEnded().	1	1	https://issues.jboss.org/browse/PLANNER-322
15932	PLANNER-319	I am refactoring our model so that we can use AnchorChain/Inverse and a CustomShadowVar listener (to compute them all), but => PLANNER-252	1	1	https://issues.jboss.org/browse/PLANNER-319
15933	PLANNER-319	The 3th point is a separate issue, presuming it still reproduces: "When the fact-change API changed such a property, we also called before/afterVariableChanged, because we could not call before/afterFactChanged, because the changed field belongs to a planning entity." That should no longer be the case in 6.2.0.Final: before/afterFactChanged() should work on a planning entity to change a problem property of an entity (so without throwing an exception). Can you confirm on deny this with 6.2 or later?	1	1	https://issues.jboss.org/browse/PLANNER-319
15934	PLANNER-317	1) thanks for the idea to split Vehicle/VehicleEntity 2) my idea was to remove the consistence-check in OptaPlanner, so that OptaPlanner does never call the getter of the inverse shadow variable, only the setter then one can annotae the @InverseRelationShadowVariable on the Standstill interface, let the Vehicle stay a fact (no cloning), because only Customer would implement the setter setNextCustomer() properly. The Vehicle just provides a dummy-impl. that ignores the next customer. BTW: what do you mean with externalized model? Could not find this term in the manual. in 6.0 and 6.2	1	1	https://issues.jboss.org/browse/PLANNER-317
15935	PLANNER-315	Can the (B) annotated like bellow? Proposal D) {code} @CustomShadowVariable(source = {fire}, listener = FireToLavaAndSmoke.class) Lava lava; @CustomShadowVariable(source = {lava}) Smoke smoke; {code} If so, this looks simple. it defines smoke as a shadow variable and dependent to lava but no listener for this update.	1	1	https://issues.jboss.org/browse/PLANNER-315
15936	PLANNER-315	Just thinking about backwards compatibility. will we change it with words we are fixing a bug? because if not, there will be always way to implement bugged domain	1	1	https://issues.jboss.org/browse/PLANNER-315
15937	PLANNER-309	Increasing the limit as time passes is a good idea. It's similar to increasing the acceptedCountLimit for Tabu Search as time passes. And it's a bit similar to reheating of LA and SA. The problem is: at what event do we decide to increase such a limit? Do we hardcode a function against time or number of steps? Or do we go with something potentially smarter (also potentially worse) such as flat line detection etc. The trigger must definitely be pluggable, so follow a Strategy design pattern, so we can implement all these and do benchmarks to figure out the best config.	0	0	https://issues.jboss.org/browse/PLANNER-309
15938	PLANNER-282	Everything done! Nearly ;) What exact do you mean with: Adjust the kie-uber-distribution accordingly ?	1	1	https://issues.jboss.org/browse/PLANNER-282
15939	PLANNER-275	With some minor changes, ACC is improved to 670. I confirmed that ACC keeps the same with larger data set (e.g. instance05) while EasyJave gets slower. Sent a PR: https://github.com/droolsjbpm/optaplanner/pull/71 Cheers,	1	1	https://issues.jboss.org/browse/PLANNER-275
15940	PLANNER-255	Thanks for reporting. Which is the useless screenshot?	1	0	https://issues.jboss.org/browse/PLANNER-255
15941	PLANNER-255	Are you sure you're getting the exact same stacktrace? Your issue sounds like an entirely different issue to me (in which case, please open a new jira issue (and link this one)).	1	0	https://issues.jboss.org/browse/PLANNER-255

#	Requirement	Questions	Answered	Successfully Answered	URL
15942	PLANNER-255	This an upstream JDK issue, only for Mac. Can you submit a new issue in the mac port? https://java.net/jira/browse/MACOSX_PORT It looks like Mac JDK 1.7.0_71 again exposes this problem that was fixed in earlier release: https://java.net/jira/browse/MACOSX_PORT-774	0	0	https://issues.jboss.org/browse/PLANNER-255
15943	PLANNER-253	Should the J in Kjar be a capital?	1	1	https://issues.jboss.org/browse/PLANNER-253
15944	PLANNER-253	Wrt. "J" in "Kjar" - that depends... what does Drools use? Have they standardized? Is "Kjar" even the proper full name, or is it "Kie Jar", "Knowledge Jar" or whatever? +1 on "scoreKjarFile" and "scoreKjarGav". Cannot see the use for "scoreKBaseId".	1	1	https://issues.jboss.org/browse/PLANNER-253
15945	PLANNER-249	Isn't this implemented already? See doc section "4.3.5.2.3. ValueRangeProvider on the planning entity" http://docs.jboss.org/drools/release/latest/optaplanner-docs/html_single/index.html If it doesn't work for a specific case, pls update the title of this issue to reflect that.	0	0	https://issues.jboss.org/browse/PLANNER-249
15946	PLANNER-249	Wouldn't filters on moves/selectors do something very similar? You still have to evaluate the filters, so there is some overhead, but the overhead of filters can be kept pretty low... http://docs.jboss.org/drools/release/6.1.0.CR1/optaplanner-docs/html_single/index.html#filteredSelection	0	0	https://issues.jboss.org/browse/PLANNER-249
15947	PLANNER-245	Can we restrict this to optaplanner-benchmarker? Basically add a ProblemStatistic that registers listeners in optaplanner-core to find this stuff out and report it. I don't want to keep track of these statistics by default in any optaplanner-core run. That will slow things down. But I might want to keep track of them, if one of the metaheuristics demands them (similarly to constraintMatchEnabled)	1	1	https://issues.jboss.org/browse/PLANNER-245
15948	PLANNER-239	I stlll have the same problem described in this issue. Will that be implemented in 6.3?	1	1	https://issues.jboss.org/browse/PLANNER-239
15949	PLANNER-235	Martin, can you check these requirements? Requirements: - All terminations (except for terminateEarly) should put the Solver in the wait() state. The solve() method should not return. - terminateEarly() should not put the Solver in the wait() state and it should get it out of it if it's in it. The solve() method should return. So terminateEarly() is effectively the only way to kill the deamon properly. - If a problemFactChange is added, the Solver must get out of the wait state() - and should not re-enter it before all problemFactChanges are done and the solver has run (every time there are new). - If the Solver starts with no planning entities, it will naturally terminate quickly, then it should go into the wait() state: to wait for problemFactChange that add entities.	1	1	https://issues.jboss.org/browse/PLANNER-235
15950	PLANNER-220	Maybe the following bug could be related to this feature: The logarithmic scaling does not work if the problemScale is ambiguous and therefore set to -1L in ProblemBenchmarkResult.class. While creating a chart having problems with ambiguous scale JfreeChart is throwing a Runtime Exception: {quote} java.lang.RuntimeException: Values less than or equal to zero not allowed with logarithmic axis at org.jfree.chart.axis.LogarithmicAxis.autoAdjustRange(LogarithmicAxis.java:530) at org.jfree.chart.axis.NumberAxis.configure(NumberAxis.java:415) at org.jfree.chart.plot.XYPlot.configureDomainAxes(XYPlot.java:974) at org.jfree.chart.plot.XYPlot.datasetChanged(XYPlot.java:4539) at org.jfree.chart.plot.XYPlot.setDataset(XYPlot.java:1396) at org.optaplanner.benchmark.impl.report.BenchmarkReport.createScalabilityPlot(BenchmarkReport.java:527) at org.optaplanner.benchmark.impl.report.BenchmarkReport.writeBestScoreScalabilitySummaryChart(BenchmarkReport.java:302) at org.optaplanner.benchmark.impl.report.BenchmarkReport.writeReport(BenchmarkReport.java:202) at org.optaplanner.benchmark.impl.PlannerBenchmarkRunner.benchmarkingEnded(PlannerBenchmarkRunner.java:197) at org.optaplanner.benchmark.impl.PlannerBenchmarkRunner.benchmark(PlannerBenchmarkRunner.java:89) at de.akquinet.jbosscc.energyplanner.planner.runtime.benchmark.PlannerBenchmarkTest.testDifferentLocalSearchParametersMoreStatistics(PlannerBenchmarkTest.java:63) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:17) at org.junit.internal.runners.statements.FailOnTimeout$StatementThread.run(FailOnTimeout.java:74) {quote} Didnt find a way to disable logarithmic scaling via config, is there one?	1	1	https://issues.jboss.org/browse/PLANNER-220
15951	PLANNER-215	Does this solve the problem? http://docs.jboss.org/drools/release/latest/optaplanner-docs/html_single/index.html#reusingTheScoreCalculationOutsideTheSolver {code}for (ConstraintMatchTotal constraintMatchTotal : guiScoreDirector.getConstraintMatchTotals()) { String constraintName = constraintMatchTotal.getConstraintName(); Number weightTotal = constraintMatchTotal.getWeightTotalAsNumber(); for (ConstraintMatch constraintMatch : constraintMatchTotal.getConstraintMatchSet()) { List<Object> justificationList = constraintMatch.getJustificationList(); Number weight = constraintMatch.getWeightAsNumber(); ... } } {code}	0	0	https://issues.jboss.org/browse/PLANNER-215
15952	PLANNER-211	How to implement? See http://cs.stackexchange.com/questions/18467/non-uniform-random-distribution-how-do-i-get-a-random-between-100-and-180-that http://stackoverflow.com/questions/5314323/generate-a-random-number-with-max-min-and-meanaverage-in-java/6296366#6296366 http://commons.apache.org/proper/commons-math/javadocs/api-3.2/index.html	0	0	https://issues.jboss.org/browse/PLANNER-211
15953	PLANNER-210	I'm also interested in getting this issue fixed. Is there any work going on? If a create a PR on GH... will it get merged? Thanks a lot.	0	0	https://issues.jboss.org/browse/PLANNER-210
15954	PLANNER-210	Normally PR's are very welcome :) But because Random is about everywhere, it's a simple replace and I am currently planning to fix it next week, no need to bother. Feel free to spec out how you thing the solverconfig xml should look like if you want to switch to another random impl. How about? {code} <solver> <randomType>WELL_1024_A</> </solver> {code}	1	1	https://issues.jboss.org/browse/PLANNER-210
15955	PLANNER-210	Quick suggestion - can we, instead of <randomType>SOME_CONSTANT</randomType> let people provide an actual class name? That way, people can plug in whatever they like. (Should also help with debugging and testing.)	1	1	https://issues.jboss.org/browse/PLANNER-210
15956	PLANNER-210	Good point. Agreed that we want to support that scenario. Let's brainstorm about the solutions proposals on IRC?	1	1	https://issues.jboss.org/browse/PLANNER-210
15957	PLANNER-204	At what logging level do you want this? DEBUG or TRACE? -1 to do at logging level ERROR, WARN and even INFO: too much verbosity. DEBUG seems most appropriate?	0	0	https://issues.jboss.org/browse/PLANNER-204
15958	PLANNER-204	When do we output it? A) when the solver factory is created B) every time a solver is created	0	0	https://issues.jboss.org/browse/PLANNER-204

#	Requirement	Questions	Answered	Successfully Answered	URL
15959	PLANNER-202	Naming dilemma: Suppose we selected a capital near Brussels (so Paris, Amsterdam, London, etc). How do we call "Brussels"? A) Reference B) Original C) Origin D) Center E) ...	1	1	https://issues.jboss.org/browse/PLANNER-202
15960	PLANNER-194	Is this issue dead? We're cleaning up jira. If there's no activity on this issue in the next 3 months, I am going to reject it (based on my comment above).	0	0	https://issues.jboss.org/browse/PLANNER-194
15961	PLANNER-189	Hi Reinis. Interesting request :) I do have some concerns/doubts... A) I am not familiar with rule.getMetaData(). What kind of information does rule.getMetaData() gives us that we don't have already? Got some practical examples? Note that this info is already available: - ConstraintMatchTotal.constraintPackage is the DRL package - ConstraintMatchTotal.constraintName is the rule name - ConstraintMatch.justificationList is the list of causes that caused the rule to fire. My concerns/doubts: 1) I am mainly worried about bloating the ConstraintMatch or ConstraintMatchTotal objects. In the future, some algo's might store them in memory to make heuristic decisions based on the matches of previous steps. For example: prefer moves that resolve a ConstraintMatch which has been broken for all previous steps - even if they are worse that moves that don't resolve it. 2) I am also worried about the fact that the ConstraintMatchTotal objects outlive the RHS (they stick into ScoreHolder so the GUI can show them for example when you click on "Constraint Match".) Looks like rule.getMetaData() wraps an internal drools Map to make it unmodifiable for us. But what if drools itself modifies it later - could it "change" under our feet?	1	1	https://issues.jboss.org/browse/PLANNER-189
15962	PLANNER-189	Here's what getMetaData() does: http://docs.jboss.org/drools/release/6.0.0.CR3/drools-expert-docs/html_single/index.html#d0e4780 Basically I can declare some user-defined key-value pairs for some or all rules. Here are my scenarios: - after the planning I want to extract only "relevant" rules from ConstraintMatch collection (for instance only such I want to log or report to user if they're violated). This way I have my very own rule categorization; - eventually I want to group rules, because there are use cases where multiple drools rules map to one "logical" rule. Or to generalize, my intent is to be able to have my own taxonomy for rules in addition to package. Regarding your concerns - surely I can solve this differently (e.g. by grouping rules into packages) I wasn't simply aware of any implications. Also, according to what you write on ConstraintMatch design - aren't they intended to be generated AFTER planing by means of creating ScoreDirector and passing planning result to it so that it creates all the ConstraintMatch information? (am mentioning this because that way your concerns would have less relevance :)) On the other hand I do think I will need a continous planning very soon and in this case there is no "AFTER planing". Is it possible mebbe to discuss this a bit with drools guys to clarify what implications could there be? Another thought - rule meta data as its name implies is rather static information (similar to package or const name) and is unlikely to change (see also this drools link for the examples) because it is intended to enrich rule definition with some meta.	1	0	https://issues.jboss.org/browse/PLANNER-189
15963	PLANNER-167	Is this issue planned to be fixed in the near future? When not, how much work is it to add this feature?	1	0	https://issues.jboss.org/browse/PLANNER-167
15964	PLANNER-167	How would it behave? What are the use cases? Without those questions answered, this issue isn't defined well enough yet to consider starting implementation.	1	1	https://issues.jboss.org/browse/PLANNER-167
15965	PLANNER-167	*What are the use cases?* See attached image (simple use case). Each _Scheduling_ has an _Order_. The _Order_ has a _Runtime_. This _Runtime_ needs to be scheduled by one _Machine_ successively. Each day the _Machine_ has another capacity which is the _MachineCapacity_, sometimes the _MachineCapacity_ for a day is null --> e.g. on Saturday or Sunday. So depending on the optimized ("selected") days the number of days from the _Scheduling_ differs --> a _List<MachineCapacity>_ is needed. _@PlanningSolution: Solution_ ... returns a list of _Schedules_ (_@PlanningEntityCollectionProperty_) _@PlanningEnity: Scheduling_ ... has an _Order_ and a list of scheduled _MachineCapacities_ (_@PlanningVariable_) The number of _Schedules_ are equal to the number of _Orders_. *How would it behave?* Each optimization step the _Optimizer_ takes, depending on the selected algorithm, a number of _MachineCapacities_ and schedules this _MachineCapacities_ to the different _Schedules_. When each _Scheduling_ has enough _MachineCapacity_ (on the same Machine successively) to finish the scheduled _Order_ the _ScoreCalculator_ returns the hard constraints value 0. It is also necessary that a _Machine_ has only one _Scheduling_ at a specific moment.	1	0	https://issues.jboss.org/browse/PLANNER-167
15966	PLANNER-167	I changed the design so that I need a shadow variable, which is used for calculating the _EndDate_. This _EndDate_ is also used for the score calculation. I think it is no good design to use a shadow variable for calculating the score. But otherwise it would not be possible to solve this problem. Do you think it is good practice to use a shadow variable for calculating the score? I need to know the _EndDate_ because so I can check if another _Order_ is scheduled at the same date/time.	0	0	https://issues.jboss.org/browse/PLANNER-167
15967	PLANNER-110	Hi Michal, Are you sure this applies to drools-planner? If it's specifically aimed at the score rule calculation done by the RETE rule engine, the component should be drools-expert. If it's about how the drools-planner moves/steps should be optimized, then drools-planner is the correct component.	1	1	https://issues.jboss.org/browse/PLANNER-110
15968	PLANNER-103	approximately when will be resolved? :)	1	1	https://issues.jboss.org/browse/PLANNER-103
15969	PLANNER-100	The interface FeasibilityScore now exists, this termination can expect that that such a Score is being used. The other Terminations now have Tests (so it's better to copy one of those tests). The only open analyse question is how do we specify it in the TerminationConfig? My proposal of <hardAndSoftScoreFeasible> is outdated, especially now that we have the interface FeasibilityScore which is implemented by BendableScore and HardMediumSoftScore too.	1	1	https://issues.jboss.org/browse/PLANNER-100
15970	PLANNER-99	Why am I not able to assign this to myself? There's no such button anywhere.	1	1	https://issues.jboss.org/browse/PLANNER-99
15971	PLANNER-99	Click on the button assign and then in the dialog, assign to me?	1	1	https://issues.jboss.org/browse/PLANNER-99
15972	PLANNER-96	Copied from the change: Hi David, Thanks for sharing this change. Even though I don't want to change the BestSolutionEvent API like this (as it leaks the solver's "working" WorkingMemory), it does give me a very good insight on what you want to do. Furthermore, I 've recently specced out that the user provides a bestSolution and the workingMemory is a clone of that. So that means the bestSolution is never in the WorkingMemory. See the changes here: https://github.com/droolsjbpm/drools-planner/commit/5690541a567e5db36c6f2f5a66fda4f9c99d4086#diff-6 Next improvement in this will be to split up the SolutionDirector into SolutionDirectorFactory (holds RuleBase) and SolutionDirector (holds WorkingMemory) so the user can use the same SolutionDirectorFactory to create it's own, isolated SolutionDirector. That would mean you'd have to take the bestSolution and set it on your SolutionDirector (instead of sharing a WorkingMemory with the Solver). Can you live with that? It might be a small performance hit (but premature optimization is the root of all evil because not doing it would hinder multi-threaded solving in the future).	0	0	https://issues.jboss.org/browse/PLANNER-96

#	Requirement	Questions	Answered	Successfully Answered	URL
15973	PLANNER-94	For another issue, it should be possible to set the RuleBase instead of supplying the DRL references. Would that solve this issue? Maybe not because a service needs to be set on the WM, not the RuleBase which creates the WM. Update: It would not solve this issue because setting it on the rulebase doesn't work, because global variables need to be set on the WorkingMemory instances, not the RuleBase. However, every now and then planner needs to create a new WM (especially once we go multi-threaded/vm), currently only in debug mode I believe, so setting it on an individual WM isn't a good solution. Defining a map of globals to put on any WM that is created, might work.	1	1	https://issues.jboss.org/browse/PLANNER-94
15974	PLANNER-94	@Henrik That scenario makes sense. Here's solution proposal A: DroolsScoreDirectorFactory gets a Map<String, Object> globalVariablesToInjectMap for every entry in that map, DroolsScoreDirector injects them (key = global variable name, value = instance to inject) when creating a new working memory. ScoreDirectorFactoryConfig gets a Map<String, Object> droolsGlobalVariablesToInjectMap. Database connections in XML won't work well (because the connection will never get closed), so the use of droolsGlobalVariablesToInjectMap would call for the use of API configuration (see manual), which is fine. Chapter-Score_calculation.xml should document how to use it, preferably with a connection based example. It should clearly state that all global variables must be thread-safe. Open questions: Is the thead-safe requirement doable? In some cases they might need to hack it with ThreadLocals. Currently this is not a priority for me, but pull requests are welcome :)	1	1	https://issues.jboss.org/browse/PLANNER-94
15975	PLANNER-94	Thanks for a quick reply! I'll see what I can do. Proposal A is probably more reasonable, I can't see a reason for exposing the rest of WorkingMemory? I read through the stuff you mentioned, and this seems straightforward. Regarding thread safety: Should the globals possibly implement some interface to allow them to be notified when the solver session starts / ends? This could allow the globals to handle thread specific db-connection lifecycle stuff.	1	1	https://issues.jboss.org/browse/PLANNER-94
15976	PLANNER-94	Solution proposal D: Notifying the globals at the beginning and ending of workingMemory creation (not solve()) I don't prefer this proposal, because it requires the global instances to implement a custom interface, so you can't reuse JDCB's datasource directly etc. Solution proposal E: Notify a WorkingMemoryManager interface from DroolsScoreDirectorFactory every time a WM is created or destroyed. This is basically proposal B. I think I prefer this one. How to do proposal E: DroolsScoreDirector must delegate WM creation and destruction to DroolsScoreDirectorFactory DroolsScoreDirectorFactory can have a WorkingMemoryManager instance. Upon WM creation and destruction, DroolsScoreDirectorFactory calls WorkingMemoryManager. This gives WorkingMemoryManager an opportunity to add globals (and whatnot) to the WM during creation and cleanup during destruction.	1	1	https://issues.jboss.org/browse/PLANNER-94
15977	PLANNER-92	How does this relate to phasing? Phasing is the process where you run 10 minutes tabu search, 10 minutes simulated annealing, 10 minutes great deluge, 10 minutes tabu search, ... sequentially after one another on the same problem continue-ing on the results of the previous solver.	0	0	https://issues.jboss.org/browse/PLANNER-92
15978	PLANNER-83	Does branch and bound work with positive weighted constraints? Or only with strictly negative weighted constraints (= most use cases)? If not, we need a brute force solver too, for the use cases that use postive weighted constraints	0	0	https://issues.jboss.org/browse/PLANNER-83
15979	PLANNER-76	It's the homogeneous moves (as Java Currency in practice teaches us) that need to be spread on multiple threads. * One step thread (the thread which called solve()) ** Starting solution is created on the step thread (later we'll see if we can do multithreading support there too) ** It has a workingSolution (which is does already) ** multithreading starts from decideNextStep * N move threads ** Typically N = the number of CPU's (minus 1?) *** NOT the number of selected moves! ** Each has a work queue of Moves (= 1/N th of the selected moves) *** The stepThreads creates the Moves and distributes them among the moveThread workQueues. (A) *** The moveThreads can steal from the end of each other => work stealing queues **** probably use Java 6 implementations **DO NOT BREAK DELTA BASED SCORE CALCULATION!!! (See latest ref manual to understand why not!) *** Each moveThread has it's own workingSolution in a ThreadLocal *** A move from the workQueue of which the score and acceptChance needs to be calculated is done and undone on the moveThread's workingMemory. *** The stepMove (= winning move) is applied to the stepThread's workingSolution as well as on every moveThread's working solution **** So it needs to be cloned? ***** And applied on several solutions (B) (A) and (B) have problem to which I know no best answer at the moment: A Move holds a reference to an instance (for example Exam) of the stepThread's workingSolution (for example Examination). But when given to a moveThread it should execute on the instance of the moveThread's workingSolution, which is a different instance(actually a clone of the original Exam, but a different instance).	0	0	https://issues.jboss.org/browse/PLANNER-76
15980	PLANNER-63	Good idea. I often kill the Benchmarker process when I want all CPU resources available for my IDE, losing hours of benchmarking CPU time. Some things to pounder about: - Any solution for this issue and JBRULES-3477 will involve serializing benchmark results (including statistics) to disk. - best solution changed statistics data is now kept in memory for every solver for every file. Only at the end the graphs and csv's are made. Is this asking for a OOM for some future use cases? (yes) - Are the benchmarks being done in the correct order? Currently: per solver, then per file. If it's the other way around: per file, then per solver, we can output best solution changed statistics graphs sooner (and release that memory!) - "A new API is added to Planner that allows to resume the benchmark from the above intermediate file." I doubt the usability of this exact way. We can make it work automatically, like "<resumeFromLastIfThatIsUnfinished>true</>" (over-ridable by system property): -- Planner knows what the last benchmark in the benchmarkDirectory, because the benchmarks are in timestamped subdirs -- Problem is what if the benchmark config has changed meanwhile? We need some way to detect if any of the benchmarks already run are the exactly the same. But what if the drools-core jar or planner jar or any of the classes have been updated? Does it count as being the same benchmark? -- resumeFromLastIfThatIsUnfinished must be overridable by a system property, so you can just configure a run config that way without having to change the xml every time to enable/disable it.	1	1	https://issues.jboss.org/browse/PLANNER-63
15981	PLANNER-60	Could someone please add a brief note about this in the docs? I think the workaround is quite acceptable, but the docs are very misleading because they say "If null is a valid planning value, it should be included in the value range and the default way to detect uninitialized planning variables must be changed." (http://docs.jboss.org/drools/release/5.4.0.Beta2/drools-planner-docs/html/ch03.html)	1	1	https://issues.jboss.org/browse/PLANNER-60
15982	PLANNER-57	Could you comment an example of a use case that would benefit for global injection? I'd like to understand the real use cases for this to implement/test it better. If we reuse the same instance for all the threads, then that global instance needs to be thread-safe, is suspect to congestion and clustering will fail outright. (Note: multi-threading and clustering are design atm, not implementation.) So I 'd rather see a factory for such a global instance, but that means the global cannot output data after the Solver is done. Is that a problem?	1	1	https://issues.jboss.org/browse/PLANNER-57

#	Requirement	Questions	Answered	Successfully Answered	URL
15983	PLANNER-57	So an example is constraint costs. I'm looking at this for scheduling field workers, both what work, and what order, over a large amount of work and workers. For some sorts of work, being late is very expensive (for example where they are meeting a member of the public, and by law they have to be on time or there are financial penalties), where as for others the primary factor is travel time (when there are number of small pieces of work to do and travel costs dominate). These scaling factors between the different constraints need to be adjustable between runs, because they need to be adjustable per business area. As we discussed on IRC, these could be done by inserting facts, but that is going to slow and verbose for constantly looking them up as part of the rules. Make sense ?	1	1	https://issues.jboss.org/browse/PLANNER-57
15984	PLANNER-56	How to test this stuff: X) Run NurseRosteringBenchmarkApp with the program argument "sprint". This will take a few seconds => great for round-tripping during development. Y) Run CloudBalancingBenchmarkApp. This will take minutes or hours. It has numbers above 64k. Z) Run GeneralOptaPlannerBenchmarkApp. This will take an hour. This is a special corner case and might fail where others won't. W) Run another *BenchmarkApp. Play with <termination> in the *BenchmarkConfig.xml files to reduce round-trip time.	0	0	https://issues.jboss.org/browse/PLANNER-56
15985	PLANNER-54	IncrementalScoreDirector possibly suffers from this too. And CustomPhaseCommands in the mix, might this even more unreliable. Do we really want to support Solver caching? In hyperheuristics, I believe there will be only 1 phase too, so there it's irrelevant. In multi-phase setups (like packing santa's sleigh) and partitioning (divide and conquer, mapReduce) cases, it might be useful though.	0	0	https://issues.jboss.org/browse/PLANNER-54
15986	PLANNER-45	Is this issue even valid? I don't think drools recompiles the DRL every time. Note that multiple Solvers share the same KieBase if they come from the same SolverFactory.	0	0	https://issues.jboss.org/browse/PLANNER-45
15987	PLANNER-42	@Marcin Sounds like a bug in or misuse of Xstream. What happens if you comment out "xStream.setMode(XStream.ID_REFERENCES);" ?	1	1	https://issues.jboss.org/browse/PLANNER-42
15988	OVERLORD-134	Alright, unless you think otherwise, I think the layout is mostly there. Have at it!	0	0	https://issues.jboss.org/browse/OVERLORD-134
15989	PERFREPO-244	[~manovotn] Could you revisit, if this is still not possible? Test execution search via REST API is already implemented, editing test execution (also with values) was also implemented a few days ago, could you please try it? :)	1	1	https://issues.jboss.org/browse/PERFREPO-244
15990	PERFREPO-220	Hi [~manovotn], thanks for the information. Can I ask, how far did you get? This issue will probably be very important for us, since it's a blocker of https://issues.jboss.org/browse/PERFREPO-235 . Do you have any update on this?	1	1	https://issues.jboss.org/browse/PERFREPO-220
15991	PERFREPO-192	Hi Vojtech, thanks for the issue, the README should be updated. The master branch is development branch and the next release of the PerfRepo will require java 1.8 and EAP 6.4. It seems that you run maven with java 1.7. Try jdk 1.8. The WildFly is not supported, so there can be some issues. If it is possible try EAP 6.4	0	0	https://issues.jboss.org/browse/PERFREPO-192
15992	PERFREPO-192	[~vprusa] Thanks for reporting this issue! :) This should be definitely caused by running it on Java 7. Could you please verify that the issue was solved by upgrading to Java 8, so I could update the README file? Please feel free to report any other misunderstanding in README file.	1	1	https://issues.jboss.org/browse/PERFREPO-192
15993	PERFREPO-170	Adding the links is great and thank you! The second bullet is about the text included in the email. Here's an example from a recent email: {noformat} some of the alerts conditions failed. Alerts that failed: Alert name: Library mode Radargun WritesOnly/sec REPL TX improvement {noformat} It would help if the line with "Alert name:" included some more detail about the failure like: {noformat} some of the alerts conditions failed. Alerts that failed: Alert name: Library mode Radargun WritesOnly/sec REPL TX improvement. "ResponseTime was more than 10% different than last run" or "Expected result of x, but got y" {noformat} Sort of like the messages that our logged from a unit test assert failure? Not sure if this is possible, but I wanted to ask!	1	1	https://issues.jboss.org/browse/PERFREPO-170
15994	PERFREPO-170	This is definitely possible. Every alert has an "description" field (hehe, Field...) on it where I can describe, what that alert is guarding. I was thinking about adding it to the email, my only concern is that the mail could get long. Options: * directly include the description in the email * (currently implemented) the alert name is a link to special page in PerfRepo, where there is showed all information about the alert (condition, description, tags, link, etc - everything) Which one would you prefer (or maybe both)?	1	1	https://issues.jboss.org/browse/PERFREPO-170
15995	PERFREPO-153	Purposed library to use (demo): http://www.amcharts.com/demos/candlestick-chart/ It's pretty nice, allows to display as many values as needed, great zooming capabilities, etc. Would this suite your needs? [~jhladky] [~olichtne]	1	1	https://issues.jboss.org/browse/PERFREPO-153
15996	SECURITY-947	What version of WildFly / JDK?	1	0	https://issues.jboss.org/browse/SECURITY-947
15997	SECURITY-947	HI, in fact to be more clear, the problem is this exception : "LDAP: error code 4 - Sizelimit Exceeded" because MaxPageSize is setted in AD to 500 and sometimes we have more than 1000 groups, furthermore in the org.jboss.security there is i think 2 ways to configure it (tested the first not second) : in LdapExtModule.java (i recreate it and override it) : i add thos lines : first the variable : private static final String MAX_PAGE_SIZE = "maxPageSize"; then the code : /** * if there is more than 1000 groups in AD */ protected int maxPageSize = 500; String maxPSize = (String) options.get(MAX_PAGE_SIZE); if (maxPSize != null) { try { maxPageSize = Integer.parseInt(maxPSize); } catch (NumberFormatException e) { PicketBoxLogger.LOGGER.debugFailureToParseNumberProperty(MAX_PAGE_SIZE, this.maxPageSize); } } //set a maxPagesize ctx.setRequestControls(new Control[]{ new PagedResultsControl(maxPageSize, Control.CRITICAL) }); and in logging module you can add this configuration : <module-option name="maxPageSize" value="1000"/> i think it should be available directly in JBOSS classes Second solution, in same class : // Query for roles matching the role filter SearchControls constraints = new SearchControls(); constraints.setSearchScope(searchScope); constraints.setTimeLimit(searchTimeLimit); add this line (but i am not sure if it is working) constraints.setCountLimit(3500); i make a workaround with the first solution but a had to redeclare SecuritYActions.java and Decode.java because protected classes regards	0	0	https://issues.jboss.org/browse/SECURITY-947
15998	SECURITY-799	Patched NegotiationAuthenticator.java from 2.3.0.Final with SECURITY-793 changes. Works fine in Chrome and Firefox. In Internet Explorer it acts strange. When I open application, IE displays basic auth username/password prompt where I can put whatever I want (it does not matter) and just after it displays login form or basic auth window (depending on web.xml configuration) again which can finally be used to log in. Why is IE acting so strange?	1	1	https://issues.jboss.org/browse/SECURITY-799
15999	SECURITY-784	During fixing, could you also consider following points? * custom {{INITIAL_CONTEXT_FACTORY}} loaded from another module * fix the behavior in another LDAP login modules too (both PicketBox and JBoss Negotiation)	0	0	https://issues.jboss.org/browse/SECURITY-784
16000	SECURITY-784	Do we have any workaround while waiting the fix?	1	1	https://issues.jboss.org/browse/SECURITY-784
16001	SECURITY-784	Do we have any plan complete the fix for java 7?	1	1	https://issues.jboss.org/browse/SECURITY-784
16002	SECURITY-783	Peter, can you please take a look into Derek's patch?	1	0	https://issues.jboss.org/browse/SECURITY-783
16003	SECURITY-774	Peter, could you review and commit the attached patch please? Thanks.	1	1	https://issues.jboss.org/browse/SECURITY-774

#	Requirement	Questions	Answered	Successfully Answered	URL
16004	SECURITY-773	Peter, could you commit the attached patch please? Thanks.	1	1	https://issues.jboss.org/browse/SECURITY-773
16005	SECURITY-770	I don't have commit permissions to the svn repository. Peter, could you commit it please?	0	0	https://issues.jboss.org/browse/SECURITY-770
16006	SECURITY-761	Darran Lofthouse <darran.lofthouse@redhat.com> made a comment on [bug 1017974|https://bugzilla.redhat.com/show_bug.cgi?id=1017974] This issue has never attempted to address referrals. - Did mapping referrals work before this fix? - Did mapping referrals work before the fix that caused the regression that is fixed here? This issue is only about fixing the regression.	0	0	https://issues.jboss.org/browse/SECURITY-761
16007	SECURITY-761	Josef Cacek <jcacek@redhat.com> made a comment on [bug 1017974|https://bugzilla.redhat.com/show_bug.cgi?id=1017974] Russell, could you please describe the limitation in role search in Release Notes? Thanks.	0	0	https://issues.jboss.org/browse/SECURITY-761
16008	SECURITY-760	sgilda@redhat.com made a comment on [bug 1022679|https://bugzilla.redhat.com/show_bug.cgi?id=1022679] What is the current status of this bug? I'm still getting Jenkins build errors.	0	0	https://issues.jboss.org/browse/SECURITY-760
16009	SECURITY-745	Hi, this problems seems still to be around. I was searching for a replacement for the [GenericHeaderBasedAuthentication|https://developer.jboss.org/wiki/GenericHeaderBasedAuthentication] that works within the Undertow platform. In case there is one please point me to the correct place - I was only able to find tomcat based solutions... We are running Wildfly 9.0.0.Final and I am facing the exact same issue as the one described. Scenario: I wrote an Undertow AuthenticationMechanism that I use in a simple enterprise application with a web module and an ejb module. In the web module I have an secured path configured via web.xml where I expect an authenticated user. I also have included the Undertow Authentication Mechansim which registers correctly. When I access the page, I do get authenticated correcltly via the Undertow Authentication Mechanism and also receive a valid Subject and Principal (with according Roles) in the WebContainer. The servlet then accesses a local EJB which also has role-restrictions applied. Here I receive an authorization error since the Subject is null. The same scenario works just fine when I simply refresh the page and the JSESSIONSSO token is set. Would moving away from the Undertow Authentication Mechanism to a JASPIC Module solve this issue? Any help would be highly appreciated since this is the last thing that keeps us from moving away from WebSphere (there we are using a trust association interceptor (TAI) for header based authentication).	1	0	https://issues.jboss.org/browse/SECURITY-745
16010	SECURITY-744	Stefan - can you verify and incorporate?	0	0	https://issues.jboss.org/browse/SECURITY-744
16011	SECURITY-724	Peter - can you investigate and tell us all the interoperability problems we have here?	0	0	https://issues.jboss.org/browse/SECURITY-724
16012	SECURITY-724	I tried to create vault using 7.2.0-Alpha2 and use it in 7.1.1.Final with long alias and short alias and both working fine. subh aym, can you check it with latest AS7.2, please?	0	0	https://issues.jboss.org/browse/SECURITY-724
16013	SECURITY-671	Have you tried using the standard Kerberos login module in the second domain?	1	1	https://issues.jboss.org/browse/SECURITY-671
16014	SECURITY-671	What do you mean with "standard Kerberos login module"? Do you mean com.sun.security.auth.module.Krb5LoginModule? I can't use this directly in a security domain. Its LoginMethod always returns true, because it only puts the initialized Kerberos Ticket in the Subject as privat credentials. It has to be "surrounded" by JGSS-code. Please take a look at the classes I added to the Support Case. Here the standard Krb5LoginModule is used in combination with the JGSS API. There the Krb5LoginModule created the Kerberos-Ticket from username/passsword. But after this code is called it has to go to your SPNEGOLoginModule.innerLogin() method to accept the generated ticket. That's why I think my extensions should be integrated into the original SPNEGOLoginModule.	1	1	https://issues.jboss.org/browse/SECURITY-671
16015	SECURITY-476	Hi , we ran into similar problem similar problem with domain qualified principals and AdvancedLdapLoginModule while working with JBOSS EAP 4.3 , queries to AD failed since sAMAccountName doesn't include the domain part. We had to add the following code (denoted FIXME in the code snippet below) to findUserDN in org.jboss.security.negotiation.AdvancedLdapLoginModule to strip the DQN from the principals when required. The behavior is similar to "Use domain-qualified user names" available in Websphere global security. protected String findUserDN(LdapContext ctx) throws LoginException { if (baseCtxDN == null) { return getIdentity().getName(); } try { NamingEnumeration results = null; String filterPrincipalNameArg = getIdentity().getName(); // FIXME DOMAIN CUTOFF FIX - START if (isUseDomainQualifiedPrincipal) { // isUseDomainQualifiedPrincipal is a property exposed as a module option default is false Object[] filterArgs = {getIdentity().getName()}; } else { if (filterPrincipalNameArg != null && filterPrincipalNameArg.trim().length() > 0 && filterPrincipalNameArg.contains("@")) { filterPrincipalNameArg = filterPrincipalNameArg.substring(0, filterPrincipalNameArg.indexOf("@")); } Object[] filterArgs = { filterPrincipalNameArg }; } // FIXME DOMAIN CUTOFF - END // FIXME original: Object[] filterArgs = {getIdentity().getName()}; results = ctx.search(baseCtxDN, baseFilter, filterArgs, userSearchControls); if (results.hasMore() == false) { results.close(); throw new LoginException("Search of baseDN(" + baseCtxDN + ") found no matches"); } SearchResult sr = (SearchResult) results.next(); String name = sr.getName(); String userDN = null; if (sr.isRelative() == true) userDN = name + "," + baseCtxDN; else throw new LoginException("Can't follow referal for authentication: " + name); results.close(); results = null; if (log.isTraceEnabled()) { log.trace("findUserDN - " + userDN); } return userDN; } catch (NamingException e) { LoginException le = new LoginException("Unable to find user DN"); le.initCause(e); throw le; } } // FIXME DOMAIN CUTOFF - END // FIXME original: Object[] filterArgs = {getIdentity().getName()}; results = ctx.search(baseCtxDN, baseFilter, filterArgs, userSearchControls); if (results.hasMore() == false) { results.close(); throw new LoginException("Search of baseDN(" + baseCtxDN + ") found no matches"); } SearchResult sr = (SearchResult) results.next(); String name = sr.getName(); String userDN = null; if (sr.isRelative() == true) userDN = name + "," + baseCtxDN; else throw new LoginException("Can't follow referal for authentication: " + name); results.close(); results = null; if (log.isTraceEnabled()) { log.trace("findUserDN - " + userDN); } return userDN; } catch (NamingException e) { LoginException le = new LoginException("Unable to find user DN"); le.initCause(e); throw le; } } Since the Negotiation component was not a part of EAP in version 4.3 we had to implement this code , currently we are migrating to EAP v5 which has Negotiation 2.0.3 SP2/SP3 integrated thus we would like to have this FIX applied in the enterprise version for EAP 5. Thank you.	1	1	https://issues.jboss.org/browse/SECURITY-476
16016	SECURITY-448	Thats good to hear, where can I find details on this, it isnt in the documentation I have found? If this is the case shouldn't this ticket be resolved ?	0	0	https://issues.jboss.org/browse/SECURITY-448
16017	SECURITY-447	This seems like a pretty useful RFE and at some point we do need to start 'certifying' identity stores like MS AD - so this will be a useful enhancement to make that possible. My only concern is the availability of a test environment and the expertise to set it up and maintain it. If QE is comfortable with this then I vote that we commit it for EAP 5.1 as EAP 5.0.2 may be essentially rolled into EAP 5.1.	0	0	https://issues.jboss.org/browse/SECURITY-447
16018	SECURITY-345	Take a look?	0	0	https://issues.jboss.org/browse/SECURITY-345

#	Requirement	Questions	Answered	Successfully Answered	URL
16019	SECURITY-280	Our current implementation allows for more than one stack to be defined. What should getAppConfigurationEntry(String name) return when there are multiple stacks? Just return the first one?	1	1	https://issues.jboss.org/browse/SECURITY-280
16020	SECURITY-270	Darran, I may need to modify the NegociationAuthenticator to handle properly. If my implementation is good enough, I am willing to contribute along with other things. Have you an idea on how you want to do that already? Have you committed anything in the CVS repository or should I start from scratch? Cheers.	1	0	https://issues.jboss.org/browse/SECURITY-270
16021	SECURITY-244	Is this still an issue?	1	1	https://issues.jboss.org/browse/SECURITY-244
16022	SECURITY-242	(01:16:24 PM) marcus: anil: SECURITY-158 (01:16:30 PM) anil: marcus: ? (01:16:34 PM) marcus: anil: credential is not null (01:16:52 PM) marcus: anil: it gets populated in the getUsernameAndPassword method (01:17:39 PM) anil: the question is whether it gets added to the shared map (01:18:28 PM) anil: the thing is the password that was passed via the callback handler needs to be placed in shared map (01:21:02 PM) marcus: it gets added. if PasswordCallback returns anything, the credential is set, so is the password returned from getUsernameAndPassword (01:21:24 PM) marcus: anil: if you add the password to the shared map, you could be adding it hashed (01:21:42 PM) marcus: anil: which i don't think should be the case (01:30:25 PM) anil: marcus: I think the user is correct (01:30:35 PM) anil: see credential is the one that already existed in shared state (01:30:49 PM) anil: so it is either exists (if I am 2nd module) or null (01:31:11 PM) anil: password is what is returned by the handler and the JAAS contract says that that should be the pass that exists in the shared map (01:31:15 PM) anil: can u validate this? (01:41:37 PM) marcus: it is not null because in the getUsernameAndPassword method (called in login()) it is set to the value of the password callback (01:44:13 PM) marcus: the only way the credential is null is if PasswordCallback.getPassword() is null (01:44:31 PM) marcus: anil: i think the user is mistaken (01:44:34 PM) anil: this is why global variables are dangerous (01:45:40 PM) anil: marcus: I think in the end password and credential will contain the same information (01:45:46 PM) anil: right? (01:46:00 PM) anil: because of String pass = info[1] (01:46:01 PM) marcus: no... the password can be hashed if the options are included (01:46:21 PM) marcus: then it would store the hashed version of the password (01:46:43 PM) marcus: and not the plain text, or char[] (01:46:52 PM) marcus: which is the credential (01:48:00 PM) anil: marcus: I think u r correct (01:48:06 PM) anil: marcus: credential is the right password (01:48:11 PM) anil: marcus: that is unchanged (01:48:21 PM) marcus: anil: yes. i believe so (01:49:08 PM) anil: marcus: thanks for the second look. I suspected something fishy here	1	1	https://issues.jboss.org/browse/SECURITY-242
16023	SECURITY-141	Hi Darran, Is this ticket also related to falling back to a lesser WWW-Authenticate mechanism? (e.g. Digest or Basic). I have implemented a solution that refactors the NegotiationAuthenticator allowing the user to configure Basic fallback if they choose. In addition, the solution makes it pretty simple to incorporate Digest fallback as well. Would there be interest in this patch? If needed, I would be happy to create a separate ticket and provide the implementation. Thanks!	1	1	https://issues.jboss.org/browse/SECURITY-141
16024	SECURITY-141	Hi Darran, As you mentioned in your description that this ticket is specific to FORM authentication, I went ahead and created a ticket specific to supporting fallback to BASIC authentication, specifically SECURITY-448. I have attached the patch there so please let me know if you have any feedback. In addition, I think the provided patch could support FORM authentication as well with only a few modifications. I'm probably missing something, but I could only figure out how FORM authentication could address fallback case 1 in the comments. In case 1 the server gets a chance to respond to the user before the browser presents the 401 so we could definitely redirect them to the form instead of sending the 401. Did you have an idea for supporting case 2 with FORM authentication? Thanks again!	1	1	https://issues.jboss.org/browse/SECURITY-141
16025	SECURITY-128	Darran , can you please elaborate more on this one?	0	0	https://issues.jboss.org/browse/SECURITY-128
16026	SECURITY-89	Should we close this?	0	0	https://issues.jboss.org/browse/SECURITY-89
16027	SECURITY-36	Where is the patch? You have not attached it yet. What is the best way to test this?	1	1	https://issues.jboss.org/browse/SECURITY-36
16028	PLINK-747	Can you try adding {{org.picketlink.config}} to your dependencies ?	1	1	https://issues.jboss.org/browse/PLINK-747
16029	PLINK-691	Could not identify what may be happening from your logs. Can you please provide more details about your HA environment ? Also, how your PL configuration looks like ? Thanks. PS.: I'll be looking at that more deeply in April.	0	0	https://issues.jboss.org/browse/PLINK-691
16030	PLINK-683	This has been closed as Won't Fix - is there any plan to support this in future?	0	0	https://issues.jboss.org/browse/PLINK-683
16031	PLINK-652	Hey Mike, can you check it out again ? From upstream ...	1	0	https://issues.jboss.org/browse/PLINK-652
16032	PLINK-637	Which IDE are u using ? I can see the javadocs in Intelij though. And comments are there on each method of [1]. [1] https://github.com/picketlink/picketlink/blob/master/modules/base/api/src/main/java/org/picketlink/config/http/AbstractHttpSecurityBuilder.java	0	0	https://issues.jboss.org/browse/PLINK-637
16033	PLINK-612	Got it. I would suggest to leave this as is for now. We already have a pretty huge RFE backlog :) Are you ok with that ?	1	0	https://issues.jboss.org/browse/PLINK-612
16034	PLINK-575	Hi Ron, From a Apache DS perspective, can we stay with @Secures then ? Is this something missing in Apache DS ? Regards.	1	1	https://issues.jboss.org/browse/PLINK-575
16035	PLINK-571	Hey, That makes a lot of sense. Can you please give more details about your use case ? Maybe writing a quickstart + README or gist with your requirements. Would like to include this in 2.7.0.Beta2. Which should be the last Beta before we start CRs. Thanks.	1	1	https://issues.jboss.org/browse/PLINK-571
16036	PLINK-571	Hey Ron, Can you please test this out using latest SNAPSHOT version ? Best would be if you build upstream/master, 2.7.0-SNAPSHOT. Otherwise you may try to use a SNAPSHOT from JBoss Nexus, which I just deployed containing the necessary changes. In order to define the scope you just need to provide a configuration as follows: {code} public void onInit(@Observes SecurityConfigurationEvent event) { SecurityConfigurationBuilder builder = event.getBuilder(); builder .identity() .scope(WindowScoped.class); } {code} Regards.	1	1	https://issues.jboss.org/browse/PLINK-571
16037	PLINK-548	Can't you do that from a custom AttributeManager ? Maybe we need to pass some more context information when invoking it.	1	1	https://issues.jboss.org/browse/PLINK-548
16038	PLINK-548	Hey Dustin, Any news ? Can we close this issue ?	1	1	https://issues.jboss.org/browse/PLINK-548
16039	PLINK-548	You can close it. Should we need this functionality in the future we will look at writing a custom AttributeManager or engage you again.	1	1	https://issues.jboss.org/browse/PLINK-548
16040	PLINK-542	Have you tried to extend the SAML2EncryptionHandler and use the incoming request (and SAML message. eg: AuthnRequest) to decide whether encryption should be enabled or not ?	1	1	https://issues.jboss.org/browse/PLINK-542
16041	PLINK-535	That is something we've discussed and holding for some time. I think the best solution (but with a little more change) is improve the Query API to support that. Something like: {code} IdentityQuery<User> query = identityManager.createIdentityQuery(User.class); query.setParameter(CREATED_DATE).lessThan(date); query.setParameter(CREATED_DATE).greaterThan(date); query.setParameter(CREATED_DATE).between(date, anotherDate); {code} Or even review this API and provide a more meaningful fluent API. Wdyt ?	0	0	https://issues.jboss.org/browse/PLINK-535

#	Requirement	Questions	Answered	Successfully Answered	URL
16042	PLINK-514	unifiedpush-server/commit/079874789d0ecb7ee1f1930e1c449e0e589bb960 Running the build gives me these test failures: {code} SEVERE: CDI Beans module deployment failed org.apache.webbeans.exception.inject.DeploymentException: javax.enterprise.inject.UnsatisfiedResolutionException: Api type [org.picketlink.internal.CDIEventBridge] is not found with the qualifiers Qualifiers: [@javax.enterprise.inject.Default()] for injection into Field Injection Point, field name : eventBridge, Bean Owner : [null] at org.apache.webbeans.config.BeansDeployer.deploy(BeansDeployer.java:213) at org.apache.openejb.cdi.OpenEJBLifecycle.startApplication(OpenEJBLifecycle.java:187) at org.apache.openejb.cdi.ThreadSingletonServiceImpl.initialize(ThreadSingletonServiceImpl.java:162) at org.apache.openejb.cdi.CdiBuilder.build(CdiBuilder.java:43) at org.apache.openejb.assembler.classic.Assembler.createApplication(Assembler.java:799) at org.apache.openejb.assembler.classic.Assembler.createApplication(Assembler.java:612) at org.apache.openejb.assembler.classic.Assembler.createApplication(Assembler.java:608) at org.apache.openejb.testing.ApplicationComposers.before(ApplicationComposers.java:639) at org.apache.openejb.testing.ApplicationComposers.evaluate(ApplicationComposers.java:722) at org.apache.openejb.junit.ApplicationComposer$DeployApplication.evaluate(ApplicationComposer.java:64) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:271) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:70) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:50) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:238) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:63) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:236) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:53) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:229) at org.junit.runners.ParentRunner.run(ParentRunner.java:309) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:252) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:141) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:112) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:189) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:165) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:85) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:115) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:75) Caused by: javax.enterprise.inject.UnsatisfiedResolutionException: Api type [org.picketlink.internal.CDIEventBridge] is not found with the qualifiers Qualifiers: [@javax.enterprise.inject.Default()] for injection into Field Injection Point, field name :	1	1	https://issues.jboss.org/browse/PLINK-514
16043	PLINK-440	Did you get this working by annotating methods ? Please, consider the fix for PLINK-443.	1	1	https://issues.jboss.org/browse/PLINK-440
16044	PLINK-439	Last discussion about this issue. ========== Shane Bryzak: i think it's ok if we extend the @AttributeProperty annotation08:05:21 PM Shane Bryzak: so we have @IdentityStereotype(ROLE) on the class08:05:43 PM Shane Bryzak: then @AttributeProperty(stereotypeProperty = ROLE_NAME)08:06:01 PM Shane Bryzak: on the property field or method08:06:09 PM Pedro Igor: I see, actually, I did that after our last discussion ...08:06:32 PM Pedro Igor: but I really don't liked how we specify the property08:06:42 PM Pedro Igor: i think we're complicating ... i single annotation should do the trick ...08:07:00 PM Shane Bryzak: we shouldn't use a string value, it's not typesafe08:07:08 PM Pedro Igor: yeah, i know ..08:07:16 PM Shane Bryzak: and who knows, we might introduce a stereotype in the future for which there are multiple properties required08:07:43 PM Shane Bryzak: so a single annotation wouldn't be future-proof08:07:51 PM Pedro Igor: actually, we just create a new one ... 08:08:11 PM Pedro Igor: but i understand your point ..08:08:17 PM Pedro Igor: and I'm half way to make it working ..08:08:26 PM Pedro Igor: considering that "gaps" that i found and we discussed yesterday ..08:08:46 PM Shane Bryzak: yeah i still don't have an elegant solution for the relationship types...08:09:31 PM Pedro Igor: which is common for whatever solution we define ...08:09:34 PM Pedro Igor: ok08:09:37 PM Shane Bryzak: i see why the idea of interfaces is appealing to you08:09:50 PM Shane Bryzak: actually you know what..08:10:25 PM Pedro Igor: :)08:10:31 PM Shane Bryzak: we could use the stereotypes for the relationship types also08:10:41 PM Shane Bryzak: @RelationshipStereotype(APPLICATION_ROLE)08:10:54 PM Pedro Igor: yeah08:11:11 PM Shane Bryzak: that would solve the gaps issue, and be consistent with the other changes08:11:44 PM Pedro Igor: and also define which property is related with the role ?08:11:46 PM Shane Bryzak: right08:11:53 PM Shane Bryzak: which means i'm no longer happy with using @AttributeProperty...08:12:14 PM Shane Bryzak: hmm maybe we need a new annotation for that08:12:23 PM Pedro Igor: we can use that StereotypeProperty that you suggested in both cases ?08:12:58 PM Shane Bryzak: @RelationshipProperty(stereotype = ROLE)08:13:18 PM Pedro Igor: so we define a big enum with stereotype properties for identity types and relationships ?08:13:20 PM Shane Bryzak: either a single enum or two separate enums08:13:42 PM Shane Bryzak: if two enums, i think we need two annotations08:13:50 PM Shane Bryzak: maybe @StereotypeProperty is ok08:14:05 PM Pedro Igor: yeah08:14:07 PM Shane Bryzak: actually i like @StereotypeProperty, it's abstract enough to make sense 08:14:33 PM Pedro Igor: yeah08:14:37 PM Pedro Igor: so ..08:14:39 PM Shane Bryzak: so perhaps one big enum08:14:47 PM Shane Bryzak: @StereotypeProperty(RELATIONSHIP_ROLE)08:15:03 PM Shane Bryzak: @StereotypeProperty(IDENTITY_LOGIN_NAME) ==========	1	1	https://issues.jboss.org/browse/PLINK-439
16045	PLINK-437	Can you check again ? Sources and javadoc jars should be there.	1	1	https://issues.jboss.org/browse/PLINK-437
16046	PLINK-386	can i get a estimated timeframe on when work on this feature will start?	1	1	https://issues.jboss.org/browse/PLINK-386
16047	PLINK-386	Can you take a look at https://docs.jboss.org/author/display/PLINK/Unsolicited+Responses. This can help you setting up Unsolicited Responses using SAML v2.0. If you want SAML v1.1, the process is pretty much the same.	1	1	https://issues.jboss.org/browse/PLINK-386
16048	PLINK-358	Pedro, I started working on this. When I'm running a WildFly / EAP instance with the picketlink subsystem (the result of WFLY-2733), I can see two subsystems: - {{/subsystem=picketlink-federation}} - {{/subsystem=picketlink-identity-management}} But the code in branch PLINK-358 of the pickelink console [1] still uses the address {{/subsystem=picketlink/...}}. Do we need to update the resource addresses in the picketlink console? [1] https://github.com/picketlink2/picketlink-console/blob/PLINK-358/gui/src/main/java/org/picketlink/as/console/client/shared/subsys/model/Federation.java#L37	1	1	https://issues.jboss.org/browse/PLINK-358
16049	PLINK-349	Do you really need SPFilter ? Can you use the valves ? The reason is that SPFilter is very out dated and the valves (if you are using jbossweb or tomcat) are the best option. Regards.	0	0	https://issues.jboss.org/browse/PLINK-349

#	Requirement	Questions	Answered	Successfully Answered	URL
16050	PLINK-343	Will the enhancement be part of the 2.6.0 release? The picketlink-jetty-common project is not referenced from the parent binding pom https://github.com/picketlink/picketlink-bindings/blob/master/pom.xml#L44 and there are checkstyle errors compiling it due to it not having been refreshed in the last couple of weeks. I am eagerly awaiting a release for this feature so I don't have to use snapshots. Thanks!!!!	0	0	https://issues.jboss.org/browse/PLINK-343
16051	PLINK-327	Hi Pedro, I know you where running some more tests against other JPA implementations like EclipleLink and OpenJPA, based on this ticket. Was there any recent update ?	1	1	https://issues.jboss.org/browse/PLINK-327
16052	PLINK-319	Hi, I've also noticed the BDA issue when debugging the subsystem in WildFly. The behavior is exactly how you described, it works for the first BDA and fails for the second. Can you give us the link to your working EAR app ?	1	1	https://issues.jboss.org/browse/PLINK-319
16053	PLINK-319	Can you please take a look at the latest comments on WFLY-2689 ?	1	1	https://issues.jboss.org/browse/PLINK-319
16054	PLINK-313	can you elaborate as to exactly what version that is? We would really like to be able to use our existing saml auth to integrate with bluejeans.com here at redhat, and this is the primary blocker for that happening!	1	1	https://issues.jboss.org/browse/PLINK-313
16055	PLINK-313	Pedro - We are trying to validate this feature in PL2.5.3beta and are having trouble figuring the exact config syntax. Could you provide an example?	0	0	https://issues.jboss.org/browse/PLINK-313
16056	PLINK-313	Hi Danny, This looks like a bug when parsing the value of the Option. Can you try the patched jar which I've attached ? Thanks.	1	1	https://issues.jboss.org/browse/PLINK-313
16057	PLINK-313	Pedro, thank you for your reply. I deployed the attached .jar file but unfortunately the SAML response still looks the same as it did before. I'm not even sure what it's supposed to look like if both the assertion and complete response are signed. Can you verify that I've deployed the .jar in the correct manner (I posted the steps I took above)? Also, can you let me know which specific log entries would show up if both signatures were present and valid? Are there any other dependancies that I might not be using currently besides picketlink-core, picketlink-federation, and picketlink-jbas?	1	1	https://issues.jboss.org/browse/PLINK-313
16058	PLINK-313	Pedro, thanks. The above instructions were successful in helping me deploy 2.5.3. For this version, is there an RPM so that I can distribute these new .jar files without having to run the picketlink installer? I would rather not have it automatically deploy all the other .jar files or change my directory structure or config files.	1	1	https://issues.jboss.org/browse/PLINK-313
16059	PLINK-287	Hi Lincoln, We've a test case for that in our test suite. I`m not sure why you're not being able to retrieve the email: https://github.com/picketlink/picketlink/blob/master/modules/idm/tests/src/test/java/org/picketlink/test/idm/query/UserQueryTestCase.java#L143 I've also tried using the getResultCount() method and it is working. Can you give us the link to your application sources ? Regards. Pedro Igor	1	1	https://issues.jboss.org/browse/PLINK-287
16060	PLINK-273	Hey [~prabhalar], Can you provide more details about your use case and what exactly do you expect from IdP when sending AJAX request ? I understand that makes more sense a 401 than a redirect when request is an AJAX call. But is that all ?	0	0	https://issues.jboss.org/browse/PLINK-273
16061	PLINK-273	Hi All Was this made configurable? The description says "The request is to have a configuration switch ...".	1	0	https://issues.jboss.org/browse/PLINK-273
16062	PLINK-264	Is this still happening ? What about the solution used for PLINK2-107 ? Can we close this issue ?	0	0	https://issues.jboss.org/browse/PLINK-264
16063	PLINK-263	Is this fixed ?	0	0	https://issues.jboss.org/browse/PLINK-263
16064	PLINK-244	Hi Pedro, AeroGear was already updated to the final release. I know you guys have been working on documentation, but I have one question. Do we really need those entities? https://github.com/pedroigor/aerogear-aerodoc-backend/tree/PLINK-244/src/main/java/org/jboss/aerogear/aerodoc/model/entity	1	1	https://issues.jboss.org/browse/PLINK-244
16065	PLINK-244	Hi Pedro, I was digging into the sources and for some reason BasicModel.getAgent can't retrieve the agent (https://github.com/abstractj/aerogear-aerodoc-backend/tree/debug), I think is because we have a custom entity. How could we properly retrieve it? It will throw NPE on login. Another question, Sebastien mentioned that is necessary to retrieve and set the values again like he did on perfomLogin. Is that true?	1	1	https://issues.jboss.org/browse/PLINK-244
16066	PLINK-204	Hi Matthias, As we discussed, the filter is not re-authenticating the user if he is already authenticated. Even if the credentials were sent again. To avoid such confusion, I'll change the filter's logic to always check for the credentials even if the user is already authenticated. If they're present we always re-authenticate the user. Do you agree ? Thanks. Pedro Igor	1	1	https://issues.jboss.org/browse/PLINK-204
16067	PLINK-202	The following changes have been made: 1) SecurityContext has been renamed IdentityContext, and made an interface 2) DefaultIdentityManager has been renamed ContextualIdentityManager, and now implements the IdentityContext interface This change eliminates the weirdness we had around the 1:1 relationship between IdentityManager and SecurityContext, and also results in having to create one less object per IDM "session".	0	0	https://issues.jboss.org/browse/PLINK-202
16068	PLINK-197	Hi Mathias, Looking at the stack trace I noticed that the root cause is the column size being insufficient to hold the value. Are you using the IdentityObjectAttribute from picketlink-idm-schema ? If so, please try to use your own implementation for this class changing the column "value" length to a higher value. The default schema (pl-idm-schema) is only a sample mapping, you can always provide your own entities.	0	0	https://issues.jboss.org/browse/PLINK-197
16069	PLINK-154	If SCIM 2.0 isn't ready, can this implement SCIM 1.1?	1	0	https://issues.jboss.org/browse/PLINK-154
16070	PLINK-140	Deepali has provided some valuable feedback: "A few comments: Excellent overview in section 1. Version in following sentence needs to be fixed: "PicketLink 3.0 is licensed under the Apache License Version 2," As a developer, when would I use the identity management lib without the base module? This is a question that popped up while reading section 1. It wasn't easy to get to the PL Fed docs from the link, maybe a more specific link would help. E.g. https://docs.jboss.org/author/display/PLINK/PicketLink+Federation Even better would be to include those in the same doc as this one. In section 3.1, overview was helpful. After that, it would be good to understand the power of the library, rather that understand full details of each class and how they relate to each other. I found the snippets in section 3.3.7.1 and 3.3.8.2 to be most illustrative in understanding what I can do with these libs. I wonder if there is an easy way to start with that, or other usage examples and then cover details of the identity model and other details. This is just a thought, take my suggestion with a pinch of salt ;)"	0	0	https://issues.jboss.org/browse/PLINK-140

#	Requirement	Questions	Answered	Successfully Answered	URL
16071	PLINK-125	Feedback from Burr: 1) Start with an upfront arch diagram that demonstrates how the pieces come together. Where does IDM fit in a typical 3-tier web app arch? 2) What is an Agent? my speed reading did not seem to grok what "Agent" was, it is the returned type for getUser 3) Groups/relationships - I would like to see a more complex example worked in. Here is a potential hierarchy: Sales - North America -- East ---- Northeast ---- Southeast -- West -- Channels -- Federal ---- Defense ---- Civilian - Europe Middle East Africa (EMEA) - Asia Within these groups, you have sub-group level "administrators", those people who can modify permissions and roles. You have individual sales contributors who read/write to their customer's data but only read access to their peer's customer data. You have sales managers for a sub-group. 4) How to handle custom schema - this was part of Seam Security if I remember correctly 5) Forgot password - security questions (what was your high school mascot) 6) Captcha 7) Social login setup/configuration 8) X509 9) Kerboros (Windows/IE integrated login) 10) LDAP and AD integration Section 3.4.4 -what is the name of the default file on disk? That way I can go looking for it, see if it is editable using vi/notepad Section 3.6 - what if the identity provider is read-only? it is fairly common for the LDAP/AD to be read-only from a developers perspective, therefore adding users is harder. In many cases, my user identity is coming from Active Directory but the roles/permissions in a custom store for the application	0	0	https://issues.jboss.org/browse/PLINK-125
16072	PLINK-123	Running this command against the demo, I still get 401: {code:title=request.sh|borderStyle=solid} curl -v -I --user john:123 --digest "http://localhost:8080/aerogear-controller-demo/mycars" {code} The request output: {code:title=output.txt|borderStyle=solid} * About to connect() to localhost port 8080 (#0) * Trying ::1... * Connection refused * Trying 127.0.0.1... * connected * Connected to localhost (127.0.0.1) port 8080 (#0) * Server auth using Digest with user 'john' > HEAD /aerogear-controller-demo/mycars HTTP/1.1 > User-Agent: curl/7.24.0 (x86_64-apple-darwin12.0) libcurl/7.24.0 OpenSSL/0.9.8r zlib/1.2.5 > Host: localhost:8080 > Accept: */* > < HTTP/1.1 401 Unauthorized HTTP/1.1 401 Unauthorized < Server: Apache-Coyote/1.1 Server: Apache-Coyote/1.1 < Set-Cookie: JSESSIONID=FEa9MKlByC-KpFeOrFFecWRR.undefined; Path=/aerogear-controller-demo; HttpOnly Set-Cookie: JSESSIONID=FEa9MKlByC-KpFeOrFFecWRR.undefined; Path=/aerogear-controller-demo; HttpOnly < WWW-Authenticate: Digest realm="default",domain="/aerogear-controller-demo",nonce="MTM2Nzk0NjA2NTE0NDpmMTY2ODkwMC1hMjM3LTQ5ZjItOWYwNS03YTE5YTc4NTU5Mjk=",algorithm=MD5,qop=auth,stale="false" WWW-Authenticate: Digest realm="default",domain="/aerogear-controller-demo",nonce="MTM2Nzk0NjA2NTE0NDpmMTY2ODkwMC1hMjM3LTQ5ZjItOWYwNS03YTE5YTc4NTU5Mjk=",algorithm=MD5,qop=auth,stale="false" < Content-Type: text/html;charset=utf-8 Content-Type: text/html;charset=utf-8 < Content-Length: 958 Content-Length: 958 < Date: Tue, 07 May 2013 17:01:05 GMT Date: Tue, 07 May 2013 17:01:05 GMT < * Connection #0 to host localhost left intact * Issue another request to this URL: 'http://localhost:8080/aerogear-controller-demo/mycars' * Re-using existing connection! (#0) with host (nil) * Connected to (nil) (127.0.0.1) port 8080 (#0) * Server auth using Digest with user 'john' > HEAD /aerogear-controller-demo/mycars HTTP/1.1 > Authorization: Digest username="john", realm="default", nonce="MTM2Nzk0NjA2NTE0NDpmMTY2ODkwMC1hMjM3LTQ5ZjItOWYwNS03YTE5YTc4NTU5Mjk=", uri="/aerogear-controller-demo/mycars", cnonce="MTM2Nzk0", nc=00000001, qop="auth", response="24679f8db32aaa84d3ce8740632455fd", algorithm="MD5" > User-Agent: curl/7.24.0 (x86_64-apple-darwin12.0) libcurl/7.24.0 OpenSSL/0.9.8r zlib/1.2.5 > Host: localhost:8080 > Accept: */* > < HTTP/1.1 401 Unauthorized HTTP/1.1 401 Unauthorized < Server: Apache-Coyote/1.1 Server: Apache-Coyote/1.1 < Set-Cookie: JSESSIONID=TyBgSC+u-sVJWGnKYeBHqMrg.undefined; Path=/aerogear-controller-demo; HttpOnly Set-Cookie: JSESSIONID=TyBgSC+u-sVJWGnKYeBHqMrg.undefined; Path=/aerogear-controller-demo; HttpOnly * Authentication problem. Ignoring this. < WWW-Authenticate: Digest realm="default",domain="/aerogear-controller-demo",nonce="MTM2Nzk0NjA2NTE0Nzo5OTE3NDEzNS01ZWIwLTQzZDQtYjRhMi1jMWVmODBkNTUxMjE=",algorithm=MD5,qop=auth,stale="false" WWW-Authenticate: Digest realm="default",domain="/aerogear-controller-demo",nonce="MTM2Nzk0NjA2NTE0Nzo5OTE3NDEzNS01ZWIwLTQzZDQtYjRhMi1jMWVmODBkNTUxMjE=",algorithm=MD5,qop=auth,stale="false" < Content-Type: text/html;charset=utf-8 Content-Type: text/html;charset=utf-8 < Content-Length: 958 Content-Length: 958 < Date: Tue, 07 May 2013 17:01:05 GMT Date: Tue, 07 May 2013 17:01:05 GMT {code} The expected behavior would be 401 at first glance and 200 in the sequence right? I was running the same example with Node.js. Here comes the expected behavior: {code:title=output.txt|borderStyle=solid} s=12" * About to connect() to 127.0.0.1 port 3000 (#0) * Trying 127.0.0.1... * connected * Connected to 127.0.0.1 (127.0.0.1) port 3000 (#0) * Server auth using Digest with user 'bob' > HEAD /sessions.json?sEcho=2&iColumns=12 HTTP/1.1 > User-Agent: curl/7.24.0 (x86_64-	0	0	https://issues.jboss.org/browse/PLINK-123
16073	PLINK-98	There are some public methods on the JPA config that I'm not sure if we need to expose to users. Usually, these methods are only used by the JPA store. Eg.: getModelPropertyValue, setModelPropertyValue,getIdentityTypeDiscriminator. Should we move them to the impl module ?	1	1	https://issues.jboss.org/browse/PLINK-98
16074	PLINK-98	Sure, no problem. Can you take a look at the changes to the File store config ? https://github.com/pedroigor/picketlink/commit/a3b65692644e60f42e87b05b694ac9bc8df6fbd0 https://github.com/pedroigor/picketlink/commit/f52b0082b876f6493d8a9980864d657cac163301 If it is ok for you, I'll send a PR with all changes for all stores. Btw, what is the best way to initialize a identity store during the startup ?	1	1	https://issues.jboss.org/browse/PLINK-98
16075	PLINK-98	Shane, As we discussed, all configuration classes are now at the api module. I know you are now working with the refactoring of the store and context stuff, but I think we can close this issue. Can you confirm ?	0	0	https://issues.jboss.org/browse/PLINK-98

#	Requirement	Questions	Answered	Successfully Answered	URL
16076	PLINK-95	Adding your chat transcript on #picketlink for people interested on it: [22:24] abstractj (~abstractj@redhat/jboss/abstractj) joined the channel. [22:24] Topic is Only discussion on PicketLink (http://jboss.org/picketlink) [22:24] Set by asaldhan on July 6, 2010 1:59:56 PM GMT-03:00 [22:24] Mode is +cnt [22:53] asaldhan: abstractj: ping [22:53] asaldhan: abstractj: sbryzak: ping [22:53] sbryzak: asaldhan: pong [22:53] abstractj: asaldhan echo reply [22:54] asaldhan: abstractj: shane had some feedback on the session mgmt stuff [22:54] sbryzak: ah yes [22:54] abstractj: cool [22:54] asaldhan: abstractj: did you see the discussion we had 2 hours ago? [22:54] sbryzak: abstractj: could you describe the requirements for PLINK-95 ? [22:54] abstractj: asaldhan not really [22:54] jbossbot: jira [PLINK-95] Add session management support on PicketLink [Open (Unresolved) Feature Request, Blocker, Anil Saldhana] https://issues.jboss.org/browse/PLINK-95 [22:56] abstractj: sbryzak well following the discussions with Pedro this afternoon looks like PicketLink doesn't have session management support (looks like) [22:57] sbryzak: abstractj: are our assumptions about it correct though, that you guys need some kind of token-based solution for stateless apps? [22:58] abstractj: sbryzak token based solution is something quite simple, I did it using custom attributes from the IDM (is not something perfect, but it works) [22:58] abstractj: our concern is about session invalidation, expiration...etc [22:59] sbryzak: abstractj: and http sessions don't provide you with what you need? [23:01] abstractj: sbryzak well, we can work with HTTP sessions only, we just need to know which path to follow. Will picketlink support session management one day? Or we need to implement it ourselves? Or maybe use picketlink extensions? Is just a matter of say the correct path to follow to common web apps [23:02] qmx|away is now known as qmx [23:10] asaldhan: abstractj: if http sessions solves your use cases, you can stick with it. [23:10] asaldhan: abstractj: the web container as you know manages the sessions for you. [23:11] asaldhan: abstractj: DefaultIdentity in PL is session scoped. [23:13] abstractj: asaldhan 'kk so if I want to deal with session timeout, or invalidate some session, the way to go is implement it, correct? I'm checking because I don't want to reinvent the wheel. Will session management be supported on picketlink in the future or this is out of scope? [23:13] sbryzak: abstractj: sorry, had to take a skype call [23:14] asaldhan: abstractj: the http sessions are managed by the container. You can just use the settings for the web container, right? [23:14] sbryzak: abstractj: you know that the Identity bean is session scoped right? [23:14] abstractj: sbryzak relax! is 12 pm! [23:14] abstractj: asaldhan k [23:15] abstractj: sbryzak yup [23:16] sbryzak: abstractj: if you guys have some specific session management requirements, we can take a look at that and assess whether it belongs in picketlink [23:16] abstractj: sbryzak asaldhan so the correct way to go is: use the JBoss + session scoped beans for session management and implement ourselves? I'm double checking it, because I'm do lazy to do stuff that already exists :) [23:16] sbryzak: abstractj: by the way, i'd like to add an example of the typesafe authorization annotations to the aerogear demo if that's ok [23:16] abstractj: sbryzak feel free to do it, anytime [23:17] asaldhan: abstractj: unless you are not using http in any way, then worry about general session management (as in PBox5 and Shiro) [23:17] sbryzak: abstractj: cool.. is there any particular business method that would make sense to secure? [23:18] sbryzak: abstractj: maybe restrict Cars.getCars() to logged in users, or something like that? [23:18] abstractj: sbryzak feel free to change, delete...etc [23:19] abstractj: help is more than welcome [23:20] abstractj: asaldhan 'kk we will handle this, I was just checking because I don't want to duplicate efforts	1	0	https://issues.jboss.org/browse/PLINK-95
16077	PLINK-93	I don't have permission to change the assignee, so there we go. I'm testing the latest changes against https://github.com/abstractj/aerogear-controller-demo/commit/6940cdd084b7fde76f6ff5f4f86c4129e8239a7c And would appreciate if you guys could provide some example of usage with cURL. Currently I'm trying with curl -u "john:123" http://localhost:8080/aerogear-controller-demo/basic/ but not having too much success. Am I doing something wrong?	1	1	https://issues.jboss.org/browse/PLINK-93
16078	PLINK-93	Sorry Pedro, but I'm not following you on it. Could you please elaborate more? This JIRA is related to HTTP BASIC, not DIGEST (http://curl.haxx.se/docs/manpage.html#--digest) I don't have the same behaviour using Node.js.	1	1	https://issues.jboss.org/browse/PLINK-93
16079	PLINK-93	No problem Bruno. To use basic try that: curl --basic -b cookies.txt -c cookies.txt -u teste:teste "http://localhost:8080/aerogear-controller-demo/basic/" -v Regarding Node.js, can you please attach some example that i can use to test ?	1	1	https://issues.jboss.org/browse/PLINK-93
16080	PLINK-93	Thanks Pedro, I think that the result produced by this command will be the same take a look: curl --basic -b cookies.txt -c cookies.txt -u john:123 "http://localhost:8080/aerogear-controller-demo/basic/" -v Something really curious, I didn't get a chance to debug it on PicketLink, but when I try: curl -u john:123 "http://localhost:8080/aerogear-controller-demo/basic/" -v The response will be HTTP 404, but when I send a request with some invalid user: curl -u homer:simpson "http://localhost:8080/aerogear-controller-demo/basic/" -v The HTTP response will be 401. Is it expected? (I'm talking about the 404 response) If you want to try our demo: git clone https://github.com/abstractj/aerogear-controller-demo cd aerogear-controller-demo & git checkout picketlink_debug mvn clean package and deploy it For Node.js example is quite simple (https://github.com/jaredhanson/passport-http/tree/master/examples), but let me know if you need some help.	1	1	https://issues.jboss.org/browse/PLINK-93
16081	PLINK-84	Hi Bruno, here is what I did to try to simulate the issue: 1) Login using Username/Password 2) Back to the login page 3) Login using username/password/otp (the same goes to facebook authc) After these steps, the user is redirected to the index page. No authentication is performed in step 3 because the user is already authenticated. So, the user is never authenticate again even if using different credentials. This is done by the following code inside the authc endpoints: if (this.identity.isLoggedIn()) { return createResponse(authcRequest); } To overcome this misunderstanding I think we can try Anil's suggestion. What do you think ?	1	1	https://issues.jboss.org/browse/PLINK-84
16082	PLINK-84	Anil, looks like we're talking about different things here. What you're talking about is how to handle session timeouts, when the real issue is: - I have this method {code} @POST @Consumes(MediaType.APPLICATION_JSON) @Produces(MediaType.APPLICATION_JSON) public AuthenticationResponse login(final AuthenticationRequest authcRequest) { if (this.identity.isLoggedIn()) { return createResponse(authcRequest); } credential.setCredential(new OTPCredential(authcRequest.getUserId(), authcRequest.getPassword(), authcRequest.getOtp())); this.identity.login(); return createResponse(authcRequest); } {code} 1 - User will login successfully 2 - Next attempt to login I'll send another credential to the method .setCredential and invoke **again** identity.login(). Is another credential correct? So in theory, it should be validated. Am I wrong?	1	1	https://issues.jboss.org/browse/PLINK-84
16083	PLINK-84	Bruno, I agree that is another credential, but the user will not be re-authenticated because it is already authenticated. The execution stops at: return createResponse(authcRequest); The code you pointed out is not wrong if the requirement is "Users should not be able to authenticate again if they are already authenticated." We can have also another requirement that says "If an authenticated user try to authenticate again we should redirect him to the index page." For that we can change the TODO app. Another requirement I can see is "If an authenticated user try to authenticate again we should ask him if he wants to log out and log in with a different account.". Something like what Anil suggested but with the consent of the user. How do you see this ?	0	0	https://issues.jboss.org/browse/PLINK-84

#	Requirement	Questions	Answered	Successfully Answered	URL
16084	PLINK-81	Hi Bobby, Not sure if I understood your request correctly. Is that all about replacing the ServiceURL from SP's picketlink.xml and get its value from the original request ? The one that originated the AuthnRequest to an IdP ? And use the the "saveRestoreRequest" attribute to control this behavior ? Thanks.	1	1	https://issues.jboss.org/browse/PLINK-81
16085	PLINK-81	Hi [~blawrence], It is more clear. Thanks :) Accordingly with the specs, the AssertionConsumerServiceURL is intended to: "Specifies by value the location to which the <Response> message MUST be returned to the requester". That said, I think you should always use a value that corresponds to a protected URI in your application. Another important aspect of this behavior in PicketLink is that regardless the ACS value, PicketLink will always redirect the user to the original request once the SAML processing is done and he is authenticated at the SP. Any thoughts ? Cheers. Pedro Igor	1	1	https://issues.jboss.org/browse/PLINK-81
16086	PLINK-81	The scenario you described is clear enough to understand the motivation behind this issue. One thing we could do in this case is use system properties to specify the ServiceURL for SPs. So we can set different values when deploying the SP in different servers. But AFAIK, this is not your case. Your scenario is about a SP deployed in a single server and accessed in two different ways (directly on the server or via a reverse proxy). Support a "dynamic" ServiceURL/ACS from requests is something we need to take a closer look. I can think about some implications if we do that. Anyway, I think this issue is now clear as pure water. One last thing, how are you managing this situation today ? Thanks.	1	1	https://issues.jboss.org/browse/PLINK-81
16087	PLFED-343	Pull request sent. What I did is: - Improvement of SAMLEntityParser and SAMLEntitiesParser to filter whitespace characters (I had ClassCastException without this change) - Implementation of FileBasedEntitiesMetadataProvider so it can use metadata config file with "EntitiesDescriptorType" as root element. There is also FileBasedEntityMetadataProvider, which can handle metadata config file with single "EntityDescriptorType". New unit tests added for parsing. - AbstractIDPValve will read metadata for all SP if MetadataProviderType is available in picketlink.xml configuration - If SPSSODescriptorType from metadata contains AuthnRequestsSigned with value "false", then signature won't be verified for this SP and SAML2SignatureValidationHandler will ignore SAMLRequest from this SP. This adds possibility that validation from some SP is available but validation of other SP is disabled (For example Google Apps don't support signed SAMLRequest). If metadata are not provided, signatures are verified (same behaviour like before) - SAML2AuthenticationHandler will now save participant into IdentityParticipantStack according to this: --- If metadata are not provided, SAML2AuthenticationHandler will use same URL like original URL (Same behaviour like before) --- If metadata are provided and they contain "SingleLogoutService", then it will use URL provided by SingleLogoutService. This adds possibility to use different URL for logout then was URL for authentication (this is the case with Salesforce) --- If metadata are provided and they don't contain "SingleLogoutService", then participant is not registered into IdentityParticipantStack, which means that it will be ignored for Single Logout (This is needed for Google Apps for instance because Google Apps doesn't support SAML single logout profile)	0	0	https://issues.jboss.org/browse/PLFED-343
16088	PLFED-193	When there is no SAML Request/Response, after authentication, the IDP forwards to the index page that has links to service providers with the additional query param (?TARGET=url) as per the SAML 1.1 specification. When the user clicks on any of the links, IDP will send the user to the service provider with the saml assertion as SAMLResponse=xxx	0	0	https://issues.jboss.org/browse/PLFED-193
16089	PLFED-189	Does picket link have handler and/or login module for handling SAML 1.1 assertions from STS for ws-trust i.e. web services and not web application SSO. I can get to use the saml 1.1. token provider in picketlink-sts to issue saml 1.1 tokens but cannot find any handler/modules to validate it.	0	0	https://issues.jboss.org/browse/PLFED-189
16090	PLFED-172	Can you check in the changes?	1	1	https://issues.jboss.org/browse/PLFED-172
16091	PLFED-137	You should have opened a discussion thread in the picketlink forums rather than a JIRA issue. You should configure a different authenticator in the SP. Which authenticator exists in your context.xml?	1	1	https://issues.jboss.org/browse/PLFED-137
16092	PLFED-137	Ah sorry, it seemed like a missing feature so I put it in JIRA. I'm using Jetty and basing my work on the example here: http://community.jboss.org/wiki/PicketLinkSAMLSSOforWebContainers So, I'm loading the org.picketlink.identity.federation.web.handlers.saml2.SAML2IssuerTrustHandler, org.picketlink.identity.federation.web.handlers.saml2.SAML2LogOutHandler, and org.picketlink.identity.federation.web.handlers.saml2.SAML2AuthenticationHandler. The last one is the one that is reporting the error. Which authenticator would have done the job?	1	1	https://issues.jboss.org/browse/PLFED-137
16093	PLFED-105	RestEASY has some OAuth work done. Can you take a look at their work and see if we can reuse and advance that code to incorporate OAuth2?	1	1	https://issues.jboss.org/browse/PLFED-105
16094	PLIDM-4	Good to know. Thanks for the comment David! Is it present in 5.1.0.CR1? I was playing with it and 5.0.0.GA and not sure if I missed this behaviour. Actually the only issue I have with JUL is that it doesn't have log.fine("message", exception) style methods...	1	0	https://issues.jboss.org/browse/PLIDM-4
16095	PLINK2-87	Please, can you take a look at the following quickstart ? https://github.com/picketlink2/picketlink-quickstarts/tree/master/saml/employee-multi-protected It is about a service provider with multiple security-constraints, one for each protected folder/resources. The key configuration is how you configure the ServiceURL in picketlink.xml. You'll notice that we use a default entry point for the application using the "/sso/*" URI. This URI is also a protected resource and it is used as ServiceURL and will be used as the entry point for IDP-SP interactions. If the user has previously requested, let's say, /admin/index.jsp, after the login the user will be redirected back to this page with all the state saved (parameters, cookies, etc). Wondering if you can try it out and let us know if we need to consider something else. Thanks.	0	0	https://issues.jboss.org/browse/PLINK2-87
16096	PLINK2-87	Any feedback about the solution above ? Can we close this issue ?	0	0	https://issues.jboss.org/browse/PLINK2-87
16097	PLINK2-85	The IDP is now issuing Assertions with a AuthnStatement.SessionIndex. For now, the SessionIndex value is the same as the Assertion ID. As defined by the specs (saml-core-2.0-os.pdf): "In general, any string value MAY be used as a SessionIndex value. However, when privacy is a consideration, care must be taken to ensure that the SessionIndex value does not invalidate other privacy mechanisms. Accordingly, the value SHOULD NOT be usable to correlate activity by a principal across different session participants. Two solutions that achieve this goal are provided below and are RECOMMENDED: • Use small positive integers (or reoccurring constants in a list) for the SessionIndex. The SAML authority SHOULD choose the range of values such that the cardinality of any one integer will be sufficiently high to prevent a particular principal's actions from being correlated across multiple session participants. The SAML authority SHOULD choose values for SessionIndex randomly from within this range (except when required to ensure unique values for subsequent statements given to the same session participant but as part of a distinct session). • Use the enclosing assertion's ID value in the SessionIndex" Also, logout requests are always including the SessionIndex element.	1	1	https://issues.jboss.org/browse/PLINK2-85

#	Requirement	Questions	Answered	Successfully Answered	URL
16098	PLINK2-71	We have the need, in our project to reload dynamically the PDP policies. This feature is particuarly critical during the deployment period of out product, when policies are aligned frequently with our customer security model. Do you have a plan for this feature request ? Regards Vincent	1	0	https://issues.jboss.org/browse/PLINK2-71
16099	PLINK2-71	Vincent, we can look at this from multiple directions: 1) The servlets can be configured to create a new PDP (that would get new policies) periodically that can be configured. Saying, refresh the PDP every 2 minutes or so. 2) The individual policy locators such as ldap policy locator can be beefed up to look at aspects such as "modifyTimeStamp" on the ldap entry to refesh the policies that it has. Option1 is the simplest to do at the moment. Option 2 will take some extra time. What is your opinion?	0	0	https://issues.jboss.org/browse/PLINK2-71
16100	PLINK2-67	Can anyone give me a status update on this? Will the fix described in PLINK-146 work in the 2.5.x series? We have an SP that would very much like for us to be able to provide them assertions signed this way. Based on the description in PLINK-146 it looks like this will globally (ie. for all SP's) add a x509 signature block to the returned assertion in addition to the current RSA block? thanks	0	0	https://issues.jboss.org/browse/PLINK2-67
16101	PLINK2-51	Why is the content of the 2 conditions the same? {code} // try using the truststore alias from the service provider metadata. ServiceProviderType provider = this.spMetadata.get(serviceName); if (provider != null && provider.getTruststoreAlias() != null) { Certificate cer = this.trustManager.getCertificate(provider.getTruststoreAlias()); if(cer instanceof X509Certificate) { x509 = (X509Certificate)cer; } } // if there was no truststore alias or no PKC under that alias, use the KeyProvider mapping. if (x509 == null) { Certificate cer = this.trustManager.getCertificate(provider.getTruststoreAlias()); if(cer instanceof X509Certificate) { x509 = (X509Certificate)cer; } } {code} Our tests fails with NPE in the method: {code} 14:27:48,420 WARNING [org.apache.cxf.phase.PhaseInterceptorChain] (http-/127.0.0.1:8080-22) Application {urn:picketlink:identity-federation:sts}PicketLinkSTS#{urn:picketlink:identity-federation:sts}IssueToken has thrown exception, unwinding now: org.apache.cxf.interceptor.Fault: PLFED000010: Error obtaining public key for service: http://localhost:8080/jaxws-samples-wsse-policy-trust/SecurityService at org.apache.cxf.service.invoker.AbstractInvoker.createFault(AbstractInvoker.java:162) [cxf-api-2.7.6.redhat-2.jar:2.7.6.redhat-2] at org.apache.cxf.jaxws.AbstractJAXWSMethodInvoker.createFault(AbstractJAXWSMethodInvoker.java:213) at org.apache.cxf.service.invoker.AbstractInvoker.invoke(AbstractInvoker.java:128) [cxf-api-2.7.6.redhat-2.jar:2.7.6.redhat-2] at org.apache.cxf.jaxws.AbstractJAXWSMethodInvoker.invoke(AbstractJAXWSMethodInvoker.java:178) at org.apache.cxf.jaxws.JAXWSMethodInvoker.invoke(JAXWSMethodInvoker.java:68) at org.jboss.wsf.stack.cxf.JBossWSInvoker.invoke(JBossWSInvoker.java:129) at org.apache.cxf.interceptor.ServiceInvokerInterceptor$1.run(ServiceInvokerInterceptor.java:57) [cxf-api-2.7.6.redhat-2.jar:2.7.6.redhat-2] at java.util.concurrent.Executors$RunnableAdapter.call(Executors.java:439) [rt.jar:1.6.0_45] at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) [rt.jar:1.6.0_45] at java.util.concurrent.FutureTask.run(FutureTask.java:138) [rt.jar:1.6.0_45] at org.apache.cxf.workqueue.SynchronousExecutor.execute(SynchronousExecutor.java:37) [cxf-api-2.7.6.redhat-2.jar:2.7.6.redhat-2] at org.apache.cxf.interceptor.ServiceInvokerInterceptor.handleMessage(ServiceInvokerInterceptor.java:106) [cxf-api-2.7.6.redhat-2.jar:2.7.6.redhat-2] at org.apache.cxf.phase.PhaseInterceptorChain.doIntercept(PhaseInterceptorChain.java:271) [cxf-api-2.7.6.redhat-2.jar:2.7.6.redhat-2] at org.apache.cxf.transport.ChainInitiationObserver.onMessage(ChainInitiationObserver.java:121) [cxf-api-2.7.6.redhat-2.jar:2.7.6.redhat-2] at org.apache.cxf.transport.http.AbstractHTTPDestination.invoke(AbstractHTTPDestination.java:239) at org.jboss.wsf.stack.cxf.RequestHandlerImpl.handleHttpRequest(RequestHandlerImpl.java:93) at org.jboss.wsf.stack.cxf.transport.ServletHelper.callRequestHandler(ServletHelper.java:143) at org.jboss.wsf.stack.cxf.CXFServletExt.invoke(CXFServletExt.java:87) at org.apache.cxf.transport.servlet.AbstractHTTPServlet.handleRequest(AbstractHTTPServlet.java:286) at org.apache.cxf.transport.servlet.AbstractHTTPServlet.doPost(AbstractHTTPServlet.java:206) at javax.servlet.http.HttpServlet.service(HttpServlet.java:754) [jboss-servlet-api_3.0_spec-1.0.2.Final-redhat-1.jar:1.0.2.Final-redhat-1] at org.jboss.wsf.stack.cxf.CXFServletExt.service(CXFServletExt.java:135) at org.jboss.wsf.spi.deployment.WSFServlet.service(WSFServlet.java:140) [jbossws-spi-2.2.1.Final-redhat-1.jar:2.2.1.Final-redhat-1] at	0	0	https://issues.jboss.org/browse/PLINK2-51
16102	PBR-519	this commit (6b39d81) will break the support of Weld 1.1.13 (JBoss EAP-6.1.1) because in the Portal-Render-Phase the WeldPhaseListener don't get a RESTORE_VIEW Phase anymore and therefore no conversation-context was bound to current thread. Is this intended? Currently there is no way to configure around this...	0	0	https://issues.jboss.org/browse/PBR-519
16103	PBR-519	okay thanks another question, this changeset will break the JSR-329 spec too {quote} 5.2.6 Executing a Portlet Render Request .. ensure that all PhaseListeners listening on the before and after phase of the PhaseId.RESTORE_VIEW are called {quote} i know that JSR-329 is for JSF-1.2 Portlets. Is there a new spec (draft?) for JSF2-Portlets anywhere? which behavior i can expect from a JSF2-Portletbridge implementation? why can i skip the Portlet Integration Library with an CDI/1.1 conform implementation as you mentioned in the first comment?	0	0	https://issues.jboss.org/browse/PBR-519
16104	PBR-519	Hi Ken, Is there some new requirement or feature of CDI 1.1 that forces implementations (Weld, OpenWebBeans, CanDI) to get rid of the casts to HttpServletRequest? Or did the JBoss engineers simply decide to refactor the CDI 1.1 branch of Weld so that it no longer has the casts? Thanks! Neil	1	1	https://issues.jboss.org/browse/PBR-519
16105	PBR-495	Which browsers don't handle self closing tags properly? To my understanding all modern browsers can handle this fine, and is now the preferred method. Also, I'm not sure it's a good idea to be modifying code of anything that passes through the Bridge. If the code that includes the script is doing something that is not supported, ie. self closing tags, then it should be modified directly as opposed to getting the Bridge to modify it. Right now, I don't see a good reason for doing this.	1	1	https://issues.jboss.org/browse/PBR-495
16106	PBR-378	Wouldn't this be better to be included in the community project so that the regular build, or with a profile, could run the TCK? Seems preferable to waiting until a commit is made to find that the TCK is broken.	1	1	https://issues.jboss.org/browse/PBR-378
16107	PBR-378	Are you running the tests against Pluto or GateIn? Did you follow the advice in PBR-254 about which tests to de-activate as they won't work with other tests enabled?	1	1	https://issues.jboss.org/browse/PBR-378
16108	PBR-374	Also, is there a particular reason why the redirect is being done in a Phase Listener?	1	1	https://issues.jboss.org/browse/PBR-374
16109	PBR-307	This solution would break the createResourceURL function (create resourceURL pointing back to same portlet for a serverResourceRequest): {code:title=backingbean.java} static public String createResURL(String key) { FacesContext facesContext = FacesContext.getCurrentInstance(); String url; MimeResponse response = (MimeResponse) facesContext.getExternalContext().getResponse(); ResourceURL resourceURL = response.createResourceURL(); resourceURL.setResourceID(key); url = resourceURL.toString(); return url; } {code} additional, now it's impossible to get static-urls (cacheable for browsers) to images/css/.... or is there still a way to get the old (v2.2.0) behaviour?	0	0	https://issues.jboss.org/browse/PBR-307

#	Requirement	Questions	Answered	Successfully Answered	URL
16110	PBR-282	Hi, if the servlet is accessed through a portlet (resource request) than there is a portlet request and the bridge should access it correctly. Shouldn't it?	1	1	https://issues.jboss.org/browse/PBR-282
16111	PBR-267	How can I use this patch to get my richfaces-portlets working in liferay? Is this patch in the jboss-portletbridge svn? Can you provide a jar-file with this fix?	0	0	https://issues.jboss.org/browse/PBR-267
16112	PBR-254	Most failing tests thre UnsupportedOperationException... so I added this to AbstractExternalContext: private Bridge.PortletPhase getPortletPhase() { return (Bridge.PortletPhase) getRequestMap().get(Bridge.PORTLET_LIFECYCLE_PHASE); } @Override public String getRequestCharacterEncoding() { if (getPortletPhase() != null && request instanceof ClientDataRequest) { return ((ClientDataRequest) request).getCharacterEncoding(); } else { // other phases -- return null as per spec return null; } //return super.getRequestCharacterEncoding(); } This also seems to solve https://issues.jboss.org/browse/PBR-199, which is the same error. @Alex - can you double check this and make sure it is best approach? it is checked in as of r901 on trunk	1	1	https://issues.jboss.org/browse/PBR-254
16113	PBR-254	Many tests failed because of possible TCK issue, that threats URLs not started with application context as internals. Following this logic (and given Oracle's use of encodeActionURL to encode its equivalent of the OutputLink) then encodeAction should do a return portletResponse.encodeURL(target) if the url begins with a / and isn't (this application) context-path prefixed. I.e. we need the bridge code to give the portlet container a chance to encode the server/port (construct and absolute url) -- typically when the portlet container is remote. Sound about right? I know the spec doesn't say anything about this but it seems to fall out from this decision/interpretation. Does your code already do this? If not should it? -Mike- On 6/14/2011 1:29 PM, Michael Freedman wrote: > So what you are saying is that because ViewHandler.getResourceURL() is defined and it will always append the ContextPath then the spec/impl should assume that urls starting with '/' that don't contain a ContextPath must be out of the application (because all in application urls would have been channeled through getResourceURL? I.e. just using /xxx as a resource url without a context path doesn't work in servlet based JSF. Interesting -- okay I see your logic. Please file a JIRA against the TCK indicating the test url it should use ought to be "resources/myImage.jpg?javax.portlet.faces.BackLink=myBackLinkParam". If possible please make the needed changes yourself and verify (and then submit the jira with either the details of the code change or a patch file). > > I will file a bug against the RI concerning its not treating / as an external url -- and since we are changing the test/fixing a bug will just include in the 2.0.1 release. > -Mike-	1	1	https://issues.jboss.org/browse/PBR-254
16114	PBR-254	Was a little trigger happy to close this issue :) I will reopen it until we get full certification from Oracle. This is now tagged for certification and next release (QE). http://anonsvn.jboss.org/repos/portletbridge/tags/2.1.1.CR1/	0	0	https://issues.jboss.org/browse/PBR-254
16115	PBR-222	which plugins? we are still from jboss-parent pom here inheriting http://anonsvn.jboss.org/repos/portletbridge/tags/2.1.0.CR2.EPP51/parent/pom.xml and I haven't changed any versions in particular that I'm aware. But could have made changes to build with Maven 3.	1	1	https://issues.jboss.org/browse/PBR-222
16116	PBR-216	How to replicate this issue: - deploy the richfaces demo application into GateIn - load the portlet on one page and on another page but over wsrp - check the links at the top of the richfaces demo. In the non wsrp case they will point to external links, in the wsrp case they will point to a wsrp url (since the url is being encoded through wsrp).	1	1	https://issues.jboss.org/browse/PBR-216
16117	PBR-198	No longer an issue in the latest 2.0 CP05 patch. Just use portlet 2.0 public render params	1	1	https://issues.jboss.org/browse/PBR-198
16118	PBR-198	Yes, the public render parameter will work. But I am looking for a url parameter that works well with the page parameter defined in the Seam application. A url like http://localhost:8080/portal/portal/default/PortletParamTest?userUsername=demo1 still doesn't seem to work in portlet bridge 2.0. Am I missing anything here?	1	1	https://issues.jboss.org/browse/PBR-198
16119	PBR-198	Sure thing, keep in mind this is only for GateIn. If anybody needs this to work on JBoss Portal they can submit a patch with almost exact code as what I checked in. Gary - If you need this for EPP 4.3/JBoss Portal, can you submit a patch? otherwise I will resolve this issue...	0	0	https://issues.jboss.org/browse/PBR-198
16120	PBR-191	@Balazs - Can you verify and test the following snapshot in your liferay setup? https://repository.jboss.org/nexus/content/repositories/snapshots/org/jboss/portletbridge/portletbridge-api/2.0.1-SNAPSHOT/portletbridge-api-2.0.1-20100917.004919-1.jar https://repository.jboss.org/nexus/content/repositories/snapshots/org/jboss/portletbridge/portletbridge-impl/2.0.1-SNAPSHOT/portletbridge-impl-2.0.1-20100917.004942-1.jar	1	1	https://issues.jboss.org/browse/PBR-191
16121	PBR-188	Waldemar, Do you have a simple test war that illustrates the problem you describe? I could then use it in resolving the issue. Thanks Ken	1	1	https://issues.jboss.org/browse/PBR-188
16122	PBR-185	can you close this issue since PRP is not a problem and create a jira for doc update and another for archetype update?	1	0	https://issues.jboss.org/browse/PBR-185
16123	PBR-133	After I applied your changes to trunk, I am getting quite a few exceptions from other classes - mainly using classes from the richfaces 3.x codebase. Did you actually get this to compile? If so, what other changes did you make?	0	0	https://issues.jboss.org/browse/PBR-133
16124	PBR-119	Fix has been made on trunk. The artifact id naming convention was previously used for a workaround with cargo profiles and automatic deployment. But those have been removed from the archetype(s) for this next CR2 release. Snjezana - I tested the changes on JBT 3.0, can you confirm on 3.1 on your end?	0	0	https://issues.jboss.org/browse/PBR-119
16125	PBR-94	To be sure we are talking about the same subject : Is this feature related with the "A4J.AJAX.onExpired JavaScript function never called" problem ? This is discussed here : http://www.jboss.org/index.html?module=bb&op=viewtopic&t=144693	1	0	https://issues.jboss.org/browse/PBR-94
16126	PBR-58	what do you think of not creating any new @annotation, we could use the current @RaiseEvent and add some portlet option listing the names of the events which are relayed from internal only to the portal coordination. This would mean you don't need to add anything to seam, just add an auto-installing Seam component to the portletbridge which listens to configured events. This would be a standard Seam component needing access to the portlet configuration/option/properties and provide complete decoupling.	1	1	https://issues.jboss.org/browse/PBR-58
16127	PBR-58	Sounds even better. Have you started prototyping this at all?	1	1	https://issues.jboss.org/browse/PBR-58

#	Requirement	Questions	Answered	Successfully Answered	URL
16128	PBR-11	A small question about the created Archetypes, what is the functionality of the following lines in the *-objects.xml.... <instance> <instance-name>RichFacesEchoPortletInstance3</instance-name> <component-ref>richfacesEchoApp.RichFacesEchoPortlet3</component-ref> </instance> I am asking this because, these lines are coming with the Archetype but Portal 2.6.4 complains that instance definitions in *-objects.xml are deprecated but Archetype use it anyway. Are they specially necessary for RichFaces-Portal or it would function anyway if we live these lines out? The whole *-objects.xml looks like following. <deployments> <deployment> <if-exists>keep</if-exists> <parent-ref>tvpp</parent-ref> <instance> <instance-name>RichFacesEchoPortletInstance3</instance-name> <component-ref>richfacesEchoApp.RichFacesEchoPortlet3</component-ref> </instance> </deployment> <deployment> <parent-ref>tvpp</parent-ref> <if-exists>overwrite</if-exists> <page> <page-name>RichFacesEchoPortlet3</page-name> <window> <window-name>RichFacesEchoPortletWindow3</window-name> <instance-ref>RichFacesEchoPortletInstance3</instance-ref> <region>left</region> <height>0</height> <initial-window-state>maximized</initial-window-state> <properties> <property> <name>theme.renderSetId</name> <value>emptyRenderer</value> </property> </properties> </window> </page> </deployment> </deployments>	1	1	https://issues.jboss.org/browse/PBR-11
16129	PRESSGANG-76	Would a simple smoke Test that uses JDK's build in Transformer stuff a good start or do we need something completly different? {code} // jdk1.4.1 import javax.xml.transform.*; import java.net.*; import java.io.*; public class HowToXSLT { public static void main(String[] args) { try { TransformerFactory tFactory = TransformerFactory.newInstance(); Transformer transformer = tFactory.newTransformer (new javax.xml.transform.stream.StreamSource ("howto.xsl")); transformer.transform (new javax.xml.transform.stream.StreamSource ("howto.xml"), new javax.xml.transform.stream.StreamResult (new FileOutputStream("howto.html"))); } catch (Exception e) { e.printStackTrace( ); } } } {code} Alternatively, I can also look at just setting up pressgang-examples, and build that pdf and html as part of the build. If the xslt is invalid, that will just crash.	1	1	https://issues.jboss.org/browse/PRESSGANG-76
16130	PRESSGANG-70	Would we also need to create a ns version of the jbossorg-jdocbook-style plugin? At the moment, Publican still uses DocBook 4.5, so would we need to keep the non-ns version at 4.5 for cross-compatibility. Does the DTD/Schema declaration need to match the XSLT declaration? That's not to say that I shouldn't raise DocBook5 support as a future enhancement on the publican-list. I think I should release JDOCBOOKSTYLE-3, JDOCBOOKSTYLE-4, and this ticket as 1.2 now (branch the trunk for the release). I think moving the repo to the PressGang SVN would also be a good thing to do for the 1.2 release. I'll update the release milestones and you can go forth and implement the fix and branch the codebase. Depending on what other enhancements are implemented between this week and the next release (1.3 at this stage), it might be worth encouraging people to convert their CSS entensions to the ns version.	1	1	https://issues.jboss.org/browse/PRESSGANG-70
16131	PRESSGANG-69	Hi Geoffrey I really like the sound of Option 1. I had the same thought as you about how all the names are different. Yes, I've inherited a motley bunch of packages in their current form. In nautical terms, it's time to get them shipshape ;) In response to your proposed options: >Alternative 1: >* The parent pom has the id org.jboss.documentation-support:jbossorg-documentation-support >** All module pom.xml files must extend this parent pom.xml and declare nothing that is general to all modules (scm, distribution, ...). >* Each module gets the same groupId (regardless of it's former groupId): org.jboss.documentation-support >** Note that this already effectively breaks all pom level backward compatibility. I'd like to see "docbook-support" renamed to something like "jbossorg-jdocbook-styles" to make it a bit clearer. Seeing that I want to move these styles into PressGang, would it be appropriate to name the pom ID: org.jboss.PressGang:jbossorg-jdocbook-style >* The module artifactId's must be renamed: I agree totally; they are all over the place at the moment. Using the naming syntax as I understand it, and moving everything into PressGang, would these names be correct? >** hightlight => jbossorg-documentation-highlight org.jboss.jbossorg-jdocbook-styles:jdocbook-style-highlight >** jbossorg-fonts => jbossorg-documentation-fonts org.jboss.jbossorg-jdocbook-styles:jdocbook-style-fonts >** jbossorg-docbook-xslt => jbossorg-documentation-docbook-xslt org.jboss.jbossorg-jdocbook-styles:jdocbook-style-xslt >** jbossorg-jdocbook-style => jbossorg-documentation-jdocbook-style org.jboss.jbossorg-jdocbook-styles:jdocbook-style-CSS >** documentation => jbossorg-documentation-docs org.jboss.jbossorg-jdocbook-styles;jdocbook-style-userguide Do I understand how pom module naming works in my suggestions? If not please set me straight. Cheers J	1	1	https://issues.jboss.org/browse/PRESSGANG-69
16132	PRESSGANG-69	By changing the groupId's and artifactId's, we made backward incompatible (build level) changes, that's the reason why I think we should call a release with the new groupId's 2.0.0 instead of 1.2.0. But calling it version 1.2.0 is fine too. What I do feel stronger about is this: don't release without doing first "Move the contents of jboss_community_styles on directory up, rename the artifact jboss-community-styles to pressgang-parent" and "All the child modules must have a <parent> declaration to the pressgang-parent", as you 'll change the parent artifactId again by that action (= might be another backward incompatible (build level) change). "jaredmorgs: I don't think I want to publish the JDG along with the styles. The idea is to keep the styles and JDG separate, so we don't have to release them all in one large chunk. The JDG will be in states of development more than the styles." Good point. Still I 'd propose to make the directory name equal to the artifactId in the pom in that directory. As I see it, you 'll want the JDG to have a separate version number (and all the rest shares the same version number of course). Do you want the JDG pom to have as a parent the pressgang-parent pom? I would and do these things: - still rename it to pressgang-jboss-documentation-guide (both directory and artifactId) - overwrite the version number in pressgang-jboss-documentation-guide pom.xml (Don't overwrite the version number in the other child poms)	1	1	https://issues.jboss.org/browse/PRESSGANG-69
16133	PRESSGANG-69	So if I was building all the jboss style bundle artifacts from the parent-pom, how would I exclude the JDG from being packaged at the same time and released along with the style packages. What I want to avoid is having to complete tasks for the JDG before I can release the styles, or vice versa. This could be my ignorance of Maven packaging methods showing here.	1	1	https://issues.jboss.org/browse/PRESSGANG-69

#	Requirement	Questions	Answered	Successfully Answered	URL
16134	PRESSGANG-69	C) Looks like pressgang-jdocbook-style's pom.xml went missing when copying from the old repo to the new? Maybe similar for docs and docbook-support D) Every child should have <parent><groupid/artifactId/version/type></parent> reference to the pressgang pom. Just copy paste from https://svn.jboss.org/repos/labs/labs/jbossrules/trunk/drools-docs/drools-docs-planner/pom.xml . Since everything is inherited, you can delete these entities: groupId, version, organization, scm, repositories, build, distributionManagement, ... As for dependencies, you should use dependencyManagement in the pressgang pom to define the versions and just specify groupId and artifactId without version in the dependencies. E) I do agree that you don't want the JDG on the same level as the tools modules. so I 'd suggest this id's: groupId = org.jboss.pressgang (artifact-prefix: pressgang) parent module: pressgang Child modules of pressgang: - pressgang-documentation-guid - pressgang-tools Child modules of pressgang pressgang-tools: - pressgang-highlight - pressgang-fonts - pressgang-xslt - pressgang-xslt-ns - pressgang-jdocbook-style - pressgang-docs If you do A, B and C, I can do D and E (easier to do than to explain :).	0	0	https://issues.jboss.org/browse/PRESSGANG-69
16135	PRESSGANG-69	Hi, any news on this? Especially not executing A) is dangerous because it effectively forks the project. Want me to execute some of these?	0	0	https://issues.jboss.org/browse/PRESSGANG-69
16136	PRESSGANG-68	Is this fixed? If so, it's included in the 2.0.0 release. There was no 1.2 release (because of GAV changes to get some order into things, see pressgang svn)	0	0	https://issues.jboss.org/browse/PRESSGANG-68
16137	PRESSGANG-64	Have you verified that this works in both Maven and Publican?	1	1	https://issues.jboss.org/browse/PRESSGANG-64
16138	PRESSGANG-64	Hey Sean Do you have to specify any maven attributes when you publish, or does the publishing toolchain apply the correct formatting automatically? Do you need to make any pom.xml changes? Cheers J	1	1	https://issues.jboss.org/browse/PRESSGANG-64
16139	PRESSGANG-58	Hi Darrin I applied your patch, and verified that there are some <programlisting> without the role attribute applied. I was testing this on the JBoss Documentation Guide book, which has examples of this structure in the "Disallowed Attributes" section. (https://svn.jboss.org/repos/pressgang/) I'm not seeing the highlighting behaviour as described in your fix. The code remains unhighlighted. You can test this out using the 1.1.1.SNAPSHOT for jbossorg-jdocbook-style. If you are using the JBoss Docs Guide as the testbed, this is preconfigured in the POM. Would you be able to take another look and see what might be causing the issue? Cheers Jared	1	1	https://issues.jboss.org/browse/PRESSGANG-58
16140	PRESSGANG-58	Geoffrey I've got some information about how to get highlighting working in the JBoss Documentation Guide. The direct link to the info is here (Hudson-hosted document): http://hudson.jboss.org/hudson/job/JBoss_Documentation_Guide/lastSuccessfulBuild/artifact/JBoss_Documentation_Guide/target/docbook/publish/en-US/html_single/index.html#sgce-Structure_Guidlines_Code_Examples I'll raise a ticket to get this information updated, and have the latest jbossorg-jdocbook-style information added as well. I think the jDocBook Plugin Manual probably needs to be taken under the wing of PressGang, and perhaps integrated into the JDG? What do you reckon? Cheers Jared	1	1	https://issues.jboss.org/browse/PRESSGANG-58
16141	PRESSGANG-58	Tested this change on the drools docs. I replaced all <programlisting role="JAVA"> with <programlisting language="java">. I replaced all <programlisting role="XML"> with <programlisting language="xml">. >> I'm not seeing the highlighting behaviour as described in your fix. The code remains unhighlighted. Did you do the necessary replacing in the JBG?	1	1	https://issues.jboss.org/browse/PRESSGANG-58
16142	PRESSGANG-58	Because it uses xmlns:jhr="com.uwyn.jhighlight.renderer.Renderer" Does that mean that the highlight module in docbook-support is dead code? => No, looks like the highlight module extends the jhightlight artifact (see the resources property file)	1	1	https://issues.jboss.org/browse/PRESSGANG-58
16143	PRESSGANG-56	Did a few more experiments tonight. Firebug shows that firefox thinks there's text between the " " and the "<p>": " " (with maybe a \n). Possible path 1: Either we stop firefox from rendering that " " in xhtml, maybe by a html (likely) or CSS (unlikely) property. Possible path 2: Or we somehow tell docbook to not ident/pretty format the xhtml output, so the output is " <p>" It looks like hibernate is doing the second.	1	1	https://issues.jboss.org/browse/PRESSGANG-56
16144	PRESSGANG-53	PS: If this is not the correct project to make a jbossorg-jdocbook-style issue, my apologies. In that case, could you direct me where to move it to? It would be nice to centralize the documentation of both the jdocbook plugin and the jbossorg-jdocbook-style.	1	1	https://issues.jboss.org/browse/PRESSGANG-53
16145	PRESSGANG-52	(11:56:26 AM) sebersole: jmorgan: anyway, like i said i think we can get away with jhighlight (11:56:38 AM) sebersole: the issue seems to be how it is applied in the xslt (11:56:47 AM) jmorgan: Oh, OK. (11:56:51 AM) sebersole: yes it is limited in what it can highlight (11:57:11 AM) sebersole: is kate ported to qt4 yet? (11:57:17 AM) jmorgan: For our purposes however, XML and Java work just fine (for the most part, anyhow). (11:57:22 AM) jmorgan: Hmm, not sure again sebersole (11:57:37 AM) sebersole: so i'd say: (11:57:46 AM) sebersole: 1) if kate is qt4, and (11:58:12 AM) sebersole: 2) its highlighting library is ported to qt4, and (11:58:28 AM) sebersole: 3) there is a java language binding for qt4 (11:58:33 AM) sebersole: then we make use of all that (11:59:02 AM) sebersole: otherwise, its not worth the effort involved (11:59:11 AM) sebersole: and we just stick with what we got (11:59:25 AM) sebersole: (with fixes to xslt, hopefully)	1	1	https://issues.jboss.org/browse/PRESSGANG-52
16146	PRESSGANG-52	Yet another highlighting library: http://code.google.com/p/google-code-prettify/ Supports Java and xml and more: "bsh", "c", "cc", "cpp", "cs", "csh", "cyc", "cv", "htm", "html", "java", "js", "m", "mxml", "perl", "pl", "pm", "py", "rb", "sh", "xhtml", "xml", "xsl" Powers code.google.com and stackoverflow.com. Is written in JavaScript and renders at runtime in HTML... Threats: - Probably not applicable for code highlighting in PDF output. Given the code highlighting mess we're currently in, I might be able to live with this (Steve doesn't use PDF IIRC). - How does it react on a single html page with lots and lots of code fragments? How much slower does the page load in FF, Chrome and IE? This is critical.	0	0	https://issues.jboss.org/browse/PRESSGANG-52
16147	PRESSGANG-50	Hi Ondrej That sounds great. I wonder how we could get that into the PressGang styles? We would need to create a Maven artifact for that, wouldn't we? This is where I run into problems. I am not a developer (as you know) so what I can do is largely dependent on the skills I can get from the community. Would you know how to implement this into PressGang styles? Would you know someone who could lend a hand with this? Cheers J	1	1	https://issues.jboss.org/browse/PRESSGANG-50
16148	PRESSGANG-50	Hi Jared, could you provide some link to PressGang source so I could look where the syntax highlighter could go? Thanks.	1	1	https://issues.jboss.org/browse/PRESSGANG-50
16149	PRESSGANG-46	Hi Isaac Can you provide more information about where you are pulling the enterprise styles from? If I understand it correctly (and I am new to how these styles hold together), the jbossorg-jdocbook-styles are the base styles projects use. What each project does is then override certain aspects of the base CSS with project-specific style tweaks. Is this what has happened with the JBDS Maven build? If so, we can try and correct some of the issues in the base styles so everyone can benefit from them. However, I believe CSS overrides that projects implement are up to them to maintain. That being said, for those changes that projects feel improve the base styles, we need to move those into the base styles. This will result in a few positives: - take the pressure off projects to maintain extensive overrides. - improve the quality of the base jbossorg-jdocbookl-styles. We should investigate what CSS overrides are implemented in JBDS, and see what we can do to streamline them a bit for your team.	1	0	https://issues.jboss.org/browse/PRESSGANG-46

#	Requirement	Questions	Answered	Successfully Answered	URL
16150	PRESSGANG-43	I find I use replaceable in <programlistings> to highlight stuff users should change. Darrin, do you think there's a better way to do this. What method do you prefer to use? Do you find it puts garbage in the <programlisting> that will cause parse errors if users copy the code sample.	1	0	https://issues.jboss.org/browse/PRESSGANG-43
16151	PRESSGANG-42	Do you have any suggestions as to how this might be corrected in the styles, Sean? Have you tried to add the – and — entities to the .ent file? Does this allow jDocBook to render the em-dash and en-dash correctly?	1	1	https://issues.jboss.org/browse/PRESSGANG-42
16152	PRESSGANG-40	I am not using subtitle, just this {code} <book version="5.0" xsi:schemaLocation="http://docbook.org/ns/docbook http://www.docbook.org/xml/5.0/xsd/docbook.xsd http://www.w3.org/1999/xlink http://www.docbook.org/xml/5.0/xsd/xlink.xsd" xml:base="./" xmlns="http://docbook.org/ns/docbook" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xmlns:xs="http://www.w3.org/2001/XMLSchema" xmlns:xlink="http://www.w3.org/1999/xlink" xmlns:xi="http://www.w3.org/2001/XInclude" xmlns:ns="http://docbook.org/ns/docbook"> <info> <title>Drools Planner User Guide</title> <releaseinfo>Version &project.version;</releaseinfo> ... {code} But I see no difference in my html output: http://docs.jboss.org/drools/release/5.3.0.Final/drools-planner-docs/html/index.html and my single_html output: http://docs.jboss.org/drools/release/5.3.0.Final/drools-planner-docs/html/index.html Can you check if both of my pages have what you want yours to have?	1	0	https://issues.jboss.org/browse/PRESSGANG-40
16153	PRESSGANG-39	Hi Martin Have you tried using the highlight package from the new JBoss repos for the styles. They now reside in a different location. I've attached the pom.xml that we use to generate the JBoss Documentation Guide. What we can do with the highlighter package is limited. It is unfortunately not maintained by anyone. :(I know that Steve Ebersole has done a lot of work with Callouts in his Hibernate project. Try altering the dependencies in your project with those in the attached pom, and let me know how you go. Cheers Jared	1	1	https://issues.jboss.org/browse/PRESSGANG-39
16154	PRESSGANG-39	Hi Jared, can you please provide us with more info? Like which dependencies and plugins we should use and how to use the highlight package? I'm quite confused about it right now. Our pom.xml is located at https://github.com/weld/core/blob/1.1.0.Final/docs/reference/pom.xml. Thank you	1	1	https://issues.jboss.org/browse/PRESSGANG-39
16155	PRESSGANG-39	Hi Martin I can't see how you guys are implementing the base styles in your project pom. Usually you have to declare a styles package in your docs pom, but I can't see any mention of a styles dependency (which would include the highlighter package). Most projects declare the base styles in the pom, then override certain parts of them with their project-specific requirements. Such as the Weld Banner, and coloration. I've checked your pom, and I can't see any dependencies to the pressgang-jdocbook-style (formerly the jbossorg-jdocbook-style plugin dependency). The pressgang-jdocbook-style dependency contains callout numbers from 1 to 15. Is that the color that you saw in your code samples before they disappeared? I have also commented on WELD-826 to get more information. Cheers Jared	1	1	https://issues.jboss.org/browse/PRESSGANG-39
16156	PRESSGANG-39	Hmm, I just opened this page [1] up in FireFox, and the callouts look awesome to me. What browser are you using? [1] http://docs.jboss.org/weld/reference/1.1.0.Final/en-US/html_single/#d0e2802	0	0	https://issues.jboss.org/browse/PRESSGANG-39
16157	PRESSGANG-39	underscores in the area ID attribute values = no change # Removed full stops and tried CamelCase in area ID attribute values = no change I compared the code in the standard programlisting before the callout listing: {code} </p><pre xmlns="" xmlns:rf="java:org.jboss.highlight.XhtmlRendererFactory" class="XML"><!-- XML : generated by JHighlight v1.0 (http://jhighlight.dev.java.net) --> <faces-config version="2.0" xmlns="http://java.sun.com/xml/ns/javaee" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation=" http://java.sun.com/xml/ns/javaee http://java.sun.com/xml/ns/javaee/web-facesconfig_2_0.xsd"> </faces-config> </pre><p> {code} I then took a look at the code behind the programlistingco: {code} </p><div class="programlistingco"><pre xmlns="" xmlns:rf="java:org.jboss.highlight.XhtmlRendererFactory" class="XML"><!-- XML : generated by JHighlight v1.0 (http://jhighlight.dev.java.net) --> <web-app version="2.5" xmlns="http://java.sun.com/xml/ns/javaee" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation=" http://java.sun.com/xml/ns/javaee http://java.sun.com/xml/ns/javaee/web-app_2_5.xsd"<span	0	0	https://issues.jboss.org/browse/PRESSGANG-39

#	Requirement	Questions	Answered	Successfully Answered	URL
16158	PRESSGANG-39	It seems not ok. Could you please tell me which version 'maven-jdocbook-plugin' and 'jbossorg-docbook-xslt' should I use. I am now using '2.3.8' and '1.1.1': The document look like: <process id="paymentProcess"> ml_plain"> <actions> ml_plain"> <action id="step1" ref="prepareParameter"/> ml_plain"> <action id="step2" ref="insertJournal"/> ml_plain"> <action id="step3" ref="callUnionpay"/> ml_plain"> <action id="step4" ref="updateJournal" /> <action id="step5" ref="prepareResponse"/> </actions> </process> Here is my pom config: <groupId>org.jboss.maven.plugins</groupId> <artifactId>maven-jdocbook-plugin</artifactId> <version>2.3.8</version> <extensions>true</extensions> <dependencies> <dependency> <groupId>org.jboss</groupId> <artifactId>jbossorg-docbook-xslt</artifactId> <version>1.1.1</version> </dependency> <dependency> <groupId>org.jboss</groupId> <artifactId>jbossorg-jdocbook-style</artifactId> <version>1.1.1</version> <type>jdocbook-style</type> </dependency> </dependencies>	1	0	https://issues.jboss.org/browse/PRESSGANG-39
16159	PRESSGANG-39	Do you mean using 'pressgang-xslt' 2.x version ?	0	0	https://issues.jboss.org/browse/PRESSGANG-39
16160	PRESSGANG-36	I totally agree Geoffrey. What do you think we could do to substitute that highlighter?	1	1	https://issues.jboss.org/browse/PRESSGANG-36
16161	PRESSGANG-22	Hi Robb I 100% agree with you on this. I've seen this very issue in multiple places in the Enterprise versions of documentation. For example, in the EAP 5.1 Security Guide, there were heaps of "edit the jboss-service.xml file". Which one? Fixing this is a matter of going through each document and tying up the loose ends. We are doing this progressively in the enterprise documentation, but I'm not sure what is happening upstream. Our solution is putting the absolute file location from the JBoss installation directory. While making some sentence a little more verbose, it makes the file location crystal clear to the reader, and that after all is what user documentation is all about. Do you think it would be useful to include a section in the JBoss Documentation Guide about how to declare file locations in JBoss User Documentation (both enterprise and community)? I'm sure it wouldn't hurt. Cheers Jared	1	0	https://issues.jboss.org/browse/PRESSGANG-22
16162	PRESSGANG-21	Robb Does this relate to JBoss documentation, such as EAP? That is how I interpreted the ticket. The Middleware Docs team is fixing this issue as they discover it while fixing surrounding text for JIRAs, and when a guide is nominated for a cover-to-cover rebase. If the question relates to JBoss user guides, and you're not sure where to raise the bug against, please create a JIRA against the JBDOCS project. I can route it to the correct area for you. Cheers Jared	1	0	https://issues.jboss.org/browse/PRESSGANG-21
16163	PRESSGANG-16	Looks like there are 3 related repository's: - docbook-support: https://svn.jboss.org/repos/jbossas/projects/docs/docbook-support - pressgang (the JBoss documentation guide): https://svn.jboss.org/repos/pressgang/ - jdocbook plugin (the jdocbook plugin): https://svn.jboss.org/repos/maven/plugins/jdocbook/maven-jdocbook-plugin/trunk/ Not sure where the jdocbook plugin docs are hosted. Doesn't look to be in either of these. Is there any reason why these 3 repository's can't be consolidated?	1	1	https://issues.jboss.org/browse/PRESSGANG-16
16164	PRESSGANG-14	Hi Josh A bit of time has elapsed since this ticket was raised. Do you still think this information is accurate for our Audience positioning information? Cheers Jared	1	0	https://issues.jboss.org/browse/PRESSGANG-14
16165	PRESSGANG-7	Hi Steve Is maven-jdocbook-plugin 2.2.2 ready to roll yet? Just want to make sure I update the PressGang JBoss Documentation Guide with the correct version. Cheers Jared	1	1	https://issues.jboss.org/browse/PRESSGANG-7
16166	PRESSGANG-5	Hi Josh I wonder how relevant this will be now that Mark and Vlastimil have created the Project Documentation Editor? Because this procedure in unlikely to be followed by a community contributor, this may be something useful to add to our ECS Docspace rather than mention in the JBoss Docs Guide. What are your thoughts on this? Cheers Jared	1	0	https://issues.jboss.org/browse/PRESSGANG-5
16167	ABDERA-437	Abdera is in maintenance mode, there's not much active development going on but reported bugs do get fixed. I suppose we could/should do another release to catch up with the bugs that have been fixed in the trunk code, thats not happened mainly as users haven't been pushing for it - is a release important for you? How important is the httpclient upgrade to you - if we did a new Abdera release must it also have that httpclient version update?	1	1	https://issues.apache.org/jira/browse/ABDERA-437
16168	ABDERA-437	Thanks ant for your update! What we are concerning are: whether current Abdera or new Abdera to be adotped in our project is good enough to fix known issues after no release for over 3 years, whether new bug would be resolved in time if there is in future. For httpclient, since replaced by HttpComponents, will new bug(if any) about old httpclient be fixed if build several years ago is still kept? In my opinion, better to turn to latest HttpComponents to make sure all parts of Abdera are in maintenance and move forward, then users would be more willing to use it. Your thoughts?	0	0	https://issues.apache.org/jira/browse/ABDERA-437
16169	ABDERA-418	The definition of "Blocker" is: "Blocks development and/or testing work, production could not run". Can you explain why this issue meets that criteria, and in particular why you are unable to override the dependency in your project?	1	1	https://issues.apache.org/jira/browse/ABDERA-418
16170	ABDERA-287	Well that looks interesting James. What do you want to happen with this now - how would it be integrated into Abdera? (and just to remind, you do still have commit access here)	1	0	https://issues.apache.org/jira/browse/ABDERA-287
16171	ABDERA-287	Well, it's an update for the existing uri template impl that's already there. For now, unless there's major reason to get it into trunk, keeping the patch available in jira is sufficient I think. People who want to play around with it can patch their source. Fwiw, I've been playing around with a significant update to the i18n module, fixing bugs, improving performance, and finally getting around to the ICU support for the unicode stuff. When I'm done with that i'll submit the patch for that as well. https://issues.apache.org/jira/browse/ABDERA-287?page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel&focusedCommentId=13090162#comment-13090162] now - how would it be integrated into Abdera? http://tools.ietf.org/html/draft-gregorio-uritemplate-06) ---------- URI Templates specification. implementation. The URI Template spec has changed significantly. {.list*}{/list*}{?foo,var,hello}{&half,x,y}#{var}"); http://www.red.green.blue/red/green/blue?foo=That%27s%20right%21&var=value&hello=Hello%20World%21&half=50%25&x=1024&y=768#value	0	0	https://issues.apache.org/jira/browse/ABDERA-287
16172	ABDERA-278	Could you provide a patch of what you'd like the update manifest to look like?	1	1	https://issues.apache.org/jira/browse/ABDERA-278
16173	ABDERA-278	I've committed a fix to this to trunk and included it in the upcoming 1.1.2 release, it would be great if you could verify if this fixes problem for you. The trunk artifacts are published to the snapshot repo as version 1.2-SNAPSHOT or the 1.1.2 release candidate artifacts are at http://people.apache.org/~antelder/abdera/1.1.2-RC2/	1	1	https://issues.apache.org/jira/browse/ABDERA-278

#	Requirement	Questions	Answered	Successfully Answered	URL
16174	ABDERA-267	Abhishek's analysis is correct, and I would like to add some more explanations (wearing my Axiom developer hat): * Some background information about the isSetPrefixBeforeStartElement thing can be found in [1]. It is actually not Axiom that got it wrong, but a particular version of a particular StAX implementation. In Axiom 1.2.9, we isolated that code into what we call a StAX dialect so that it is only ever triggered when that particular StAX implementation is used. * On the other hand, the performance issue related to isSetPrefixBeforeStartElement was already addressed in 1.2.8 (see WSCOMMONS-325). So upgrading to that version instead of 1.2.9 may provide a quick fix for the issue. * In Axiom 1.2.9, the factories returned by StAXUtils are indeed immutable. The reason is that these factories are cached and therefore shared with other frameworks using Axiom. E.g. if Abdera and Axis2 are used in the same application, they would get the same factory instances from Axiom. Therefore we can't allow modifications of these factory instances, because this has the potential of causing subtle issues. In 1.2.9, we also introduced new features that allow to customize the factories in a controlled way, but for XMLOutputFactory instances, this is only implemented in the current trunk. * Since Abdera extends Axiom implementation classes (namely the LLOM implementation), it is indeed strongly coupled and therefore very sensitive to changes in Axiom. Probably the best way to avoid these issues is to have the Abdera trunk depend on the Axiom trunk so that incompatibilities are detected early by the Abdera build in Hudson. I will see what needs to be done in the Abdera and Axiom code to make them again play nicely together. [1] http://markmail.org/thread/37y3qjcfucb6756o	1	1	https://issues.apache.org/jira/browse/ABDERA-267
16175	ABDERA-267	Thanks both of you for for working on this issue, i've applied the patch to Abdera trunk in r1006114. We need to get this and a couple of other fixes out in an Abdera release but wont be able to do that till an Axiom release is done so Abdera can used a non-snapshot release. When do you think an Axiom release is likely to happen? (lets leave this jira open till the release is done)	1	1	https://issues.apache.org/jira/browse/ABDERA-267
16176	ABDERA-236	Pardon my complete lack of knowledge about OSGi matters, but would the following: - <packaging>jar</packaging> + <packaging>bundle</packaging> imply that JARs are not created anymore? I guess not.	1	1	https://issues.apache.org/jira/browse/ABDERA-236
16177	ABDERA-214	we already have a RSS extension, can you give us more details?	1	1	https://issues.apache.org/jira/browse/ABDERA-214
16178	ABDERA-214	Is there any progress on this JIRA ? I'm interested in consuming these in Apache Tuscany RSS Binding.	0	0	https://issues.apache.org/jira/browse/ABDERA-214
16179	ABDERA-188	Ok, I'll try to fix that diff problem and let you know. In the meantime, given that my implementation is a drop-in replacement to the current one, may you try to completely replace current source files with the ones contained into the V3 zip file? Thanks!	0	0	https://issues.apache.org/jira/browse/ABDERA-188
16180	ABDERA-188	I've tried to fix the diff file, but it requires too much manual modifications, which I think are error-prone. I think it's better to apply the diff as it is, skipping not found files, and then manually copy the following files from the V2 zip file: src/test/resources/atomResponse.xml src/main/java/org/apache/abdera/ext/opensearch/model/Query.java src/main/java/org/apache/abdera/ext/opensearch/model/IntegerElement.java src/main/java/org/apache/abdera/ext/opensearch/model/OpenSearchExtensionFactory.java And finally delete: src/main/java/org/apache/abdera/ext/opensearch/Query.java src/main/java/org/apache/abdera/ext/opensearch/IntegerElement.java src/main/java/org/apache/abdera/ext/opensearch/OpenSearchExtensionFactory.java What do you think?	0	0	https://issues.apache.org/jira/browse/ABDERA-188
16181	ABDERA-188	Hi guys, have you made any progress with this? Any chance to get this into trunk? Thanks!	0	0	https://issues.apache.org/jira/browse/ABDERA-188
16182	ABDERA-185	I introduced the NoOpTransaction in order to remove boiler-plate null checking and make code more readable: however, it is not strictly needed. Talking about the processor classes, do you want me to move them to another specific package and submit another diff file? I'm happy you appreciated my work: let me know how to proceed. If you'll commit my changes, I'll start working at improving OpenSearch support by using request processors. Thanks.	1	1	https://issues.apache.org/jira/browse/ABDERA-185
16183	ABDERA-182	There are other patch that adds support for using the requirements in the Route.match: https://issues.apache.org/jira/browse/ABDERA-139 Take a look at the patch and if you agree with it I could add them.	1	1	https://issues.apache.org/jira/browse/ABDERA-182
16184	ABDERA-182	I've looked at 139's requirements patch, the only issue i have with it is it doesn't let me do what i want with the requirements. It's a way to filter on your variables that are matched to path segments. (and it appears this is how it works in Python routes). However, I want to set a requirement on a url param being present, which is after the ? and so not part of the URI in the case of the patch in 139. This is why I put in generic accessor methods. Then I use the requirements during the resolve.	0	0	https://issues.apache.org/jira/browse/ABDERA-182
16185	ABDERA-182	This seems pretty non-controversial, since it doesn't change behavior, just makes some data accessible. Can I get a commit on this one? I'd like to use this feature in the project I'm working on and other workarounds have not panned out. thanks	0	0	https://issues.apache.org/jira/browse/ABDERA-182

#	Requirement	Questions	Answered	Successfully Answered	URL
16186	ABDERA-168	afin de traiter automatiquement un filtre par categories * pour les adapteurs qui n'implante pas le support des catégories. * * Si on ne le faisait pas, un adapteur ignorerait simplement les catégories * passées en paramètres. * * Un adapteur retournera un 400 automatiquement si le support des catégories * est à false, sans avoir à ce soucier de ce support dans le dao associé à l'adapteur. * * De plus, on en profite pour ajouter automatiquement les liens pour le paging. * * @see #getHandlesCategories() */ @Override protected void addFeedDetails(Feed feed, RequestContext request) throws ResponseContextException { // Y a-t-il au moins une categorie sur l'URL? if (request.getTarget().getParameter("category1") != null) { // Il y a des categories dans l'URI de la requete... if (!getHandlesCategories()) { // L'adapteur ne supporte pas les categories. On bypass le traitement normal. ServiceContext context = getServiceContext(request); context.getGenericErrors().reject(BaseErrorCode.REST_REQUEST_CATEGORIES_NOT_SUPPORTED.toString(), "Cet adapteur ne supporte pas les catégories."); checkForErrors(context, request); } } // TODO:RG Devrait utiliser un hook du genre afterGetEntries() plutôt // que de réimplanter complètement addFeedDetails(). // Ajouter automatiquement un lien next page pour le paging si le nombre d'entrées // dans le feed est un de plus que ce qui devait être retourné. // La méthode getServiceContext() a modifié cette valeur pour que l'adapteur // charge une entrée supplémentaire. Donc, si cette entrée est présente dans // le feed, cela indique qu'une prochaine page est disponible. feed.setUpdated(new Date()); Iterable<T> entries = getEntries(request); ServiceContext context = getServiceContext(request); int inflatedMaxResult = context.getMaxResults(); // On corrige l'info de paging dans le ServiceContext. context.setMaxResults(inflatedMaxResult - 1); boolean needNextPageLink = false; int collectionSize = 0; if (entries != null) { for (T entryObj : entries) { // Detection de la presence d'une prochaine page. ++collectionSize; if (collectionSize >= inflatedMaxResult) { needNextPageLink = true; // Ne pas ajouter cette dernière entrée. break; } Entry e = feed.addEntry(); IRI feedIri = new IRI(getFeedIriForEntry(entryObj, request)); addEntryDetails(request, e, feedIri, entryObj); if (isMediaEntry(entryObj)) { addMediaContent(feedIri, e, entryObj, request); } else { addContent(e, entryObj, request); } } } int currentStartIndex = context.getStartIndex(); // Ajout d'un lien premiere page? if (currentStartIndex > 1) { FeedPagingHelper.setFirst(feed, createPagingLink(request, 1)); } // Ajout d'un lien page précédente? if (currentStartIndex > 1) { int newStartIndex = currentStartIndex - context.getMaxResults(); if (newStartIndex < 1) { newStartIndex = 1; } FeedPagingHelper.setPrevious(feed, createPagingLink(request, newStartIndex)); } // Ajout d'un lien page suivante? if (needNextPageLink) { int newStartIndex = currentStartIndex + context.getMaxResults(); FeedPagingHelper.setNext(feed, createPagingLink(request, newStartIndex)); } // Un lien derniere page n'est jamais ajouté étant trop couteux à calculer. } /** * Génère un lien de paging pour insertion dans le feed par l'appelant. * * @param startIndex Le nouveau start-index. * @return Un lien absolue. */ private String createPagingLink(RequestContext request, int startIndex) { String uri = request.getResolvedUri().toString(); int separatorIndex = uri.indexOf("?"); if (separatorIndex != -1) { uri = uri.substring(0, separatorIndex); } StringBuilder link = new StringBuilder(uri); link.append("?").append(SearchConstants.START_INDEX).append("=").append(startIndex); for(String parameterName : request.getParameterNames()) { if (!parameterName.equals(SearchConstants.START_INDEX)) { String parameterValue =	1	0	https://issues.apache.org/jira/browse/ABDERA-168
16187	ABDERA-168	You could use a URI Template to do this by setting the template pattern and passing in the current uri parameters in the context with the appropriate changes for the page. e.g. something like: Template template = new Template("?{-join|&|a,b,c,d,page}"); Map<String,Object> context = new HashMap<String,Object>(); context.put("a", request.getParameter("a")); context.put("b", request.getParameter("b")); context.put("c", request.getParameter("c")); context.put("d", request.getParameter("c")); context.put("page", 2); String queryString = template.expand(context); you can then append the expanded querystring to the base request uri	0	0	https://issues.apache.org/jira/browse/ABDERA-168
16188	ABDERA-129	Applying this patch breaks the ant build. Can you take another look and make sure we can get a clean build	1	1	https://issues.apache.org/jira/browse/ABDERA-129
16189	ABDERA-106	The patch does not appear to line up with the current code in trunk and I don't know enough about spring to fix it up myself. Can you please take another look at the patch?	1	1	https://issues.apache.org/jira/browse/ABDERA-106
16190	ABDERA-104	I'm confused... getId() is for the feed ID. Do you mean getHref? I'm still a little fuzzy as to what you mean. Can you expand on the problem?	1	1	https://issues.apache.org/jira/browse/ABDERA-104
16191	ABDERA-104	Went back to RFC4287. I was missusing the getId() method. I now understand the difference between the atom:link@rel=self and the atom:id tags. However, I would like some pointers on providing feed ids... Do you recommend using a uuid, a guid or... where can I get pointers on the tag: constructs used in the unit tests? I can see the use of a feed id when building a generic adapter for a generic feed server and using the id to map into your feed table. However, I'm building an Abdera server over an existing application and the relation between the adapter and the database content is obviously hard coded in my adapters. I really have no need for the feed id. Should I just provide a filler value? Thanks again!	1	1	https://issues.apache.org/jira/browse/ABDERA-104
16192	ABDERA-97	I think this patch should work. Compiles .. all tests pass. Yesterday's break was due to not cleaning the build before compiling and running the tests. Can you please have a look at this.	1	0	https://issues.apache.org/jira/browse/ABDERA-97
16193	ABDERA-56	Was just wondering - are there still plans to integrate this into Abdera? (Not trying to rush you here, just don't want it to get out of date and fall off the roadmap)	1	0	https://issues.apache.org/jira/browse/ABDERA-56
16194	ACCUMULO-4673	Pull request has been merged. Is this done?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4673
16195	ACCUMULO-4658	[~dlmarion], can you add a brief description about this issue, to give the change some context?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4658
16196	ACCUMULO-4647	Can this be marked as resolved?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4647
16197	ACCUMULO-4644	Should this be closed?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4644
16198	ACCUMULO-4643	I was thinking about the use case of an iterator that does some type of transform. If not handled properly, this feature could lead to that iterator being reseeked with an untransformed key from an iterator below it. Where normally its only reseeked with transformed keys. For option #1, was the idea to add a interrupt boolean field to Key (like the delete boolean)? If so, then the iterator doing transformation would have to ensure it sets the interrupt field on the transformed key. For option 2, any iterator doing transformations should catch the exception, transform the key, and then thrown a new exception with the transformed key.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4643
16199	ACCUMULO-4643	[~ivan.bella] for your PR, what made you decide on option #2 over option #1?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4643
16200	ACCUMULO-4643	Whats the expected/desired behavior of this feature w.r.t. isolation? When an iterator yields and is later continued, will it be continued with the exact same data sources? Or could it be restarted with a different set of data sources (in which a compaction could have occurred or new data was added)?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4643

#	Requirement	Questions	Answered	Successfully Answered	URL
16201	ACCUMULO-4643	I created a [test suite|https://github.com/keith-turner/yield-test] and ran the following test. Test 1 : 128 tserver threads, 127 non-yielding long running scans, 1 thread doing 50 batch scans of 5 random ranges. Below are the times in ms for the batch scans. There were 5,000,000 key/values. Each long runnning scan was reading all. The total time for all long running scans was : 104 seconds {noformat} min: 30.0 max: 2395.0 mean: 882.3599999999999 standard deviation: 653.56063254159 {noformat} Test 2 : 128 tserver threads, 127 yielding long running scans, 1 thread doing 50 batch scans of 5 random ranges. The long running scans yielded every 10,000. Below are the times in ms for the batch scans. Total time for all long running scans was : 215 seconds {noformat} min: 428.0 max: 827.0 mean: 626.2800000000002 standard deviation: 104.54157263100298 {noformat} Test 3 : 16 tserver threads, 20 long running scans, 1 thread doing 50 batch scans of 5 random ranges. Below are the times in ms for the batch scans. There were 31,750,000 key/values. Each long runnning scan was reading all. The total time for all long running scans was : 174 seconds {noformat} min: 120.0 max: 548.0 mean: 210.54000000000002 standard deviation: 70.64345948031524 {noformat}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4643
16202	ACCUMULO-4643	Can this issue be closed, now that it has been merged into master? Is there more work to be done?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4643
16203	ACCUMULO-4642	bq. However, I think we can get there incrementally, by starting with eliminating the SystemConfiguration.getInstance() method. Are you referring to the AccumuloConfiguration object returned from a call to ServerConfigurationFactory.getSystemConfiguration() ? I can't find any SystemConfiguration class...	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4642
16204	ACCUMULO-4626	bq. When a single iterator tree references the same RFile blocks in different branches we sometimes get cache misses for one iterator even though the requested block is held in memory by another iterator. Have you figured out why this happens in the first place? If two iterators are reading from the same RFile, I would expect the block cache to contain the blocks for that file after the first iterator accesses it. Why are we getting the cache misses in the first place? Some kind of race condition where the block is not yet in the cache?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4626
16205	ACCUMULO-4626	Thanks for the explanation, Adam. bq. the sum of sizes of the referenced blocks across all of the concurrently running queries exceeds the 25% or so of the total cache that is reserved for single-use blocks That seems really aggressive in terms of eviction to me. Are we getting poor cache utilization because of that? Maybe there's some other characteristic which keeps more-often accessed blocks in cache? I'm trying to get a better understanding (admittedly without yet pulling up the code) about how this would affect more real-life workloads.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4626
16206	ACCUMULO-4620	bq. It would be great if all Accumulo modules were consolidate into a modules directory to put every module in one directory and reduce the number of directories in the root project directory What does this get us? I don't understand what value having fewer directories at the root of the repository provides.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4620
16207	ACCUMULO-4620	I don't understand the benefit here either. I see your PR is just moving everything into a modules/ directory. Can you explain what you are trying to fix and how this will accomplish that?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4620
16208	ACCUMULO-4620	bq. I haven't heard any actual reasons why this change will be disruptive. It is disruptive because it will cause the active developers to change how they access the filesystem. It's yet another thing to keep in mind as I'm doing command-line operations in Accumulo (based on the branch). My objection does not extend farther than I do not think it is confusing to have many modules, nor do I find the existing server directory confusing. I just disagree with the justifications/reasons for moving the modules in the first place. Anyways, I don't care to participate in argument/discussion about a directory -- my "-0" is all I really care to give. If others feel strongly that this should be done, you have the ability to proceed (my -0 is not a -1).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4620
16209	ACCUMULO-4614	Should that be a debug or info log? If it's just a one-shot load, INFO, might be okay. Otherwise, I'd suggest DEBUG.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4614
16210	ACCUMULO-4614	This has been merged. Is it done? Can this issue be closed, [~mikewalch]?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4614
16211	ACCUMULO-4609	Individual k/v pairs are not inspected during a bulk import. Data is not iterated over during this operation. If I understand your request correctly, it seems you're looking for the ability to perform some key/value filtering/transformation during first scan or major compaction after a file is bulk imported. is that correct? I think this is probably already possible with the existing scan/majc scopes. Are these insufficient?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4609
16212	ACCUMULO-4596	All the PRs have been merged. Is this issue complete or is there more subtasks to be done?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4596
16213	ACCUMULO-4584	[~elserj], mind taking a look at this? If we can remove these checks, any ideas on how to update the RpcWrapperTest to verify the correct behavior? The question is: is it better to wrap exceptions from oneway methods with TException like all the rest (effectively reverting ACCUMULO-4065) and have thrift drop them, or is it better to use reflection to figure out which methods are oneway and pass those through (keeping ACCUMULO-4065) unwrapped?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4584
16214	ACCUMULO-4584	bq. wrap exceptions from oneway methods with TException like all the rest (effectively reverting ACCUMULO-4065) and have thrift drop them I think the proper fix is to have Thrift handle this and for us to remove this work-around. However, having confidence that Thrift doesn't regress on this would be good (I'm cringing at having to debug this one again..). bq. any ideas on how to update the RpcWrapperTest to verify the correct behavior The only reasonable way I can think to try to verify this is to completely stub out the entire RPC pipeline. Along the lines of what is presently there in the test class now ({{foo_args}}, {{FakeServiceImpl}}, {{fake_fields}}), we could extend that to be a real end-to-end test that uses Thrift (instead of just mocking out the InvocationHandler). Maybe it would be easier in the long run to just write a more tightly-scoped unit test for RpcWrapper and create a new test that does a standard Client/Server setup for Thrift. We should be able to make a dumb IDL (with a oneway method that always throws an exception and a non-oneway method that returns a value), create a client and server, and then call the oneway method and the non-oneway method using the same connection. Does that make sense?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4584
16215	ACCUMULO-4532	I see a patch in the master branch for this. Is it done? Can it be closed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4532
16216	ACCUMULO-4531	I see some patches to the master branch for this issue. Is it done? Can this be closed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4531
16217	ACCUMULO-4530	More directly, is this one done?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4530
16218	ACCUMULO-4522	I see commits on 1.7, 1.8, and master branches. Is this done?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4522
16219	ACCUMULO-4510	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo-testing/pull/2#discussion_r97808927 --- Diff: core/src/main/java/org/apache/accumulo/testing/core/continuous/ContinuousIngest.java --- @@ -183,8 +177,9 @@ private static long flush(BatchWriter bw, long count, final int flushInterval, l return lastFlushTime; } - public static Mutation genMutation(long rowLong, int cfInt, int cqInt, ColumnVisibility cv, byte[] ingestInstanceId, long count, byte[] prevRow, Random r, - boolean checksum) { + static Mutation genMutation(long rowLong, int cfInt, int cqInt, ColumnVisibility cv, --- End diff -- Why was this reformatted? Does maven format?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4510

#	Requirement	Questions	Answered	Successfully Answered	URL
16220	ACCUMULO-4500	bq. we should limit the feature to be a mapping of names (type: String) to counts (type: signed Long) :coughs: That... is a histogram. Label (string) on x-axis, numeric value on the y-axis. What is this aversion to using the word "histogram"? I honestly don't understand the desire to avoid using that word. bq. I think it would be a good idea that when this information is exposed in the client API, it should be retrievable through a user-supplied aggregation/combiner function Can you sketch what you think this should look like? It's not entirely clear to me. bq. The reasoning for this is that client code doesn't normally deal with things at the granularity of files, but rather, the granularity of tablets, ranges, and tables Keith had made a good point elsewhere that, when files overlap tablets, we might double-count. That, combined with the counting (presently) not included on the IMM, these "counters" (winks) would not be 100% accurate. Personally, I think that's OK. bq. A summation function would probably be the most common, but certainly not the only useful aggregation function. Example of another aggregation function?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4500
16221	ACCUMULO-4500	bq. :coughs: That... is a histogram. I disagree. A histogram is a diagram... a particular kind of data visualization. Computer scientists overload the term to mean to mean the data structure to power such a visualization, and that's fine, but that's not even what we have here. A histogram has particular semantics where different values of a single variable are on one axis, and the frequency of those values are shown on the other. What we have here is more general than even that, because we're not necessarily referring to a single variable, nor do the magnitudes necessarily represent frequencies or anything like frequencies. Calling it a histogram implies semantics we don't necessarily need to impose. bq. Can you sketch what you think this should look like? It's not entirely clear to me. Something like: {code:java} public NamedCounters getCounters(Range range, Function<Long, Long> combiner); {code} bq. Keith had made a good point elsewhere that, when files overlap tablets, we might double-count. That, combined with the counting (presently) not included on the IMM, these "counters" (winks) would not be 100% accurate. Personally, I think that's OK. I agree. bq. Example of another aggregation function? Here's some: https://en.wikipedia.org/wiki/Aggregate_function I could think of others.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4500
16222	ACCUMULO-4500	Re passing in an instance of a Function, how can we run the function on remote JVMs? We could take the class name from the instance and use that to instantiate on remote JVMs. The following are two possible issues with taking the class name from the instance. * Anonymous inner classes (AIC) : If the user passes in Foo$1, it could be that the same AIC is Foo$2 on a remote JVM. This could happen when running a slightly different version of the code on the remote. Also if the Anonymous inner class is not static, could be tricky to instantiate on remote JVM. * Lambdas : Are not inner classes[1]. I have no idea if you can safely serialize[2] lambdas between class files generated from different versions of code??? I have wanted to look into what Spark does, but have not had a chance to dig into the details. I strongly suspect Spark can make certain assumptions based on the fact that all JVMs are running the exact same version of user code. In this case its reasonable to expect an AIC to have the same name across JVMs. I don't think we should assume Accumulo client and server processes are running the exact same version of user code. It seems given our circumstances that using a class name instead of an instance would be safer. I suppose we could just dissallow AICs (throw an exception). Not sure about lambdas. [1]: https://www.infoq.com/articles/Java-8-Lambdas-A-Peek-Under-the-Hood [2]: https://docs.oracle.com/javase/8/docs/api/java/lang/invoke/SerializedLambda.html	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4500
16223	ACCUMULO-4500	That was just a quick mock-up, not the actual suggested implementation. That mock-up wouldn't work anyway, because it's actually a list of Longs to map to a Long. But, then there's also the question... should this function be able to change behavior based on the counter name? Should it be able to emit new counters (like a many-to-many transformation) or just reduce? I don't know if we could do serialized lambdas or if we'd have to pass the class name, but you're right: The aggregation function will need to be available on both the client and the server side, because some aggregation might occur within the tserver, and some between tservers in the client. Maybe those should be two separate aggregation opportunities? Should our API only reflect in-cluster aggregation, and client code has to do its own cross-tserver aggregation because we'll just provide a stream? I don't know. I just know it'd be nice to put some serious thought into what any new API is going to look like... and based on a solid understanding of what features we're trying to provide at the most basic level.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4500
16224	ACCUMULO-4500	bq. A histogram has particular semantics where different values of a single variable are on one axis, and the frequency of those values are shown on the other. What we have here is more general than even that, because we're not necessarily referring to a single variable, nor do the magnitudes necessarily represent frequencies or anything like frequencies. Calling it a histogram implies semantics we don't necessarily need to impose. I really don't want to get into an argument over semantics. bq. What we have here is more general than even that, because we're not necessarily referring to a single variable I don't know this means. {code} public NamedCounters getCounters(Range range, Function<Long, Long> combiner); {code} I still don't understand what the {{Function<Long,Long>}} is accomplishing. Combination of numbers is (typically) addition. Why do we need a function here? Also, this doesn't let me fetch just the visibility data. If there are multiple types of counters stored in the rfiles, how do I refer to just one? How do I know which exist for some table? bq. Re passing in an instance of a Function, how can we run the function on remote JVMs? We could take the class name from the instance and use that to instantiate on remote JVMs This is 110% what I wanted to avoid. What was suggested as a *simple* database primitive is now being exploded into a "wholly generalized user-configured framework". I have no interest in working on that. If either of you do, I'm happy to help sketch out such a system, but I will lose all interest in building this simple feature if it's being stipulated that I have to build something so largely different to build the proposed *simple* feature.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4500
16225	ACCUMULO-4493	Do you know if it's possible to renew outside the shell process to get things going again for a long-running shell?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4493
16226	ACCUMULO-4493	bq. Do you know if it's possible to renew outside the shell process to get things going again for a long-running shell? It.. should be. I don't think I've ever tried to do this, but we should have access to the ticket cache (either in the default location or via the {{KRB5CCNAME}} environment variable), which should be all that we need. There is a call {{UserGroupInformation#reloginFromTicketCache()}}, but I'm not sure how (if at all) UGI would know about the ticket in the ticket cache (from the kinit). Might be some sort of disconnect since it's essentially found automatically by JAAS (instead of explicitly logged-in by the Accumulo shell). bq. Users should be able to launch the shell in a kerberos deployment using a keytab. So really, we're just adding another argument to the shell to let the user provide a keytab and then do a normal login+renewal? Makes sense.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4493

#	Requirement	Questions	Answered	Successfully Answered	URL
16227	ACCUMULO-4487	bq. I don't think we should try to find an Accumulo-specific build solution for this sort of thing, though +1 to a long-term upstream fix. How do you feel about a stop-gap until such an upstream feature could be committed/released, Christopher? My initial reaction was that I'd like to see us _not_ in this department, but I think the debt we would incur would be relatively small. [~cheng xu], have you reached out to the user@maven.apache.org list to see if there is some (perhaps, poorly documented) feature which could do what you've suggested here?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4487
16228	ACCUMULO-4481	in a related note, have we documented that upgrades from 1.5 and earlier to 1.8 and later isn't possible? I don't see anything in the 1.8.0 release notes nor the [1.8 user guide's section on upgrade issues|http://accumulo.apache.org/1.8/accumulo_user_manual#_upgrade_issues], for example.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4481
16229	ACCUMULO-4481	Isn't the LocalWALRecovery tool expressly mentioned in our user facing docs? we'll need to make sure those docs are clear about what version(s) the tool can be used in (IIRC this is already the case, but we should verify).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4481
16230	ACCUMULO-4479	{quote} This is prone to being overlooked and forgotten. We should have generic version-agnostic information in UPGRADING.md, and put more specific per-release notes in an UPGRADING section on the release notes for specific upgrade issues. At the very least, the version of this document for the current release does not need to contain upgrade instructions for upgrading to previous releases. {quote} counter-point: if we don't support direct upgrades from all older versions to all newer versions, having a dedicated upgrade doc makes it much easier for folks who have to go through multiple upgrades, e.g. to get from version 1.4 to 1.8. If we're going to rely on release notes, will that mean that each maintenance release will need to preserve an "upgrade considerations" section from the .0 release? Are we going to retroactively change the release notes for 1.5, 1.6, and 1.7 to include this information?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4479
16231	ACCUMULO-4479	Ok, I put a PR up for upgrade notes before I understood what Christopher meant. My concern here is that it was too easy to miss including upgrade notes for the 1.8.0. I am going to remove the UPGRADING.md file from all branches. All future releases should include upgrading notes in their release notes. I may back port upgrade notes to the 1.8.0 release notes as well.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4479
16232	ACCUMULO-4468	Hi [~wmurnane]. Thanks for the patch. bq. This functions correctly, of course, but one of the typical uses of this method is to compare adjacent rows to break them into larger chunks. For example, accumulo.core.iterators.Combiner repeatedly calls this method with subsequent pairs of keys. Do you have other examples of where this might be used in a tight loop? bq. This (marginally) improves the speed of comparisons in the relatively common case where only the last part is changing, with less complex code. How did you test this? What types of numbers did you see? {code} + case ROW: + if (!isEqual(row, other.row)) + return false; + break; default: throw new IllegalArgumentException("Unrecognized partial key specification " + part); } + return true; {code} Why not consolidate this to: {code} case ROW: return isEquals(row, other.row); {code}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4468
16233	ACCUMULO-4468	bq. Do you have other examples of where this might be used in a tight loop? I think there are lots of other examples in Accumulo itself; for example, in iterators.system.DeletingIterator and iterators.user.TransformingIterator, this method is called in a loop. In our use case, we're doing a similar thing: building a larger object out of multiple rows, by finding groups of rows which are equal under ROW_COLFAM. When each object is built from only a few rows, the CF equality comparison returns false pretty often (which is to be expected), but only after comparing row IDs, which are always the same in practice. bq. How did you test this? What types of numbers did you see? I haven't been able to install it on a cluster to test. The test suite does pass with this patch applied. I think it's a minor change; in the "rows are equal" case the same amount of work is done as with the existing code, although the parts are accessed in the opposite order. They're still compared mostly-in-order, as isEqual does, but the comment in that function was inspiration to try reversing the comparison order. Aside from performance, the code seems cleaner to me: there's no more repetition of e.g. the check of row equality. The bytecode (with Oracle javac 1.8.0_92) is substantially smaller: 389 bytes versus 167. bq. Why not consolidate this to: That's fine with me. I just wrote all the cases to look the same, instead of having a "special case" for the last comparison made. If some special work were required in the compare-equal case, it could go before the return statement.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4468
16234	ACCUMULO-4468	Hi [~wmurnane]. Thanks for the patch! After reviewing, I had a few comments: # This may speed things up only if the probability of not being equal is greater in the lower dimensions of the key than the higher ones. I'm not sure this is the case, as I don't have a strong sense of when this method is called, or how frequently. What analysis have you done to determine these relative probabilities? # Have you run any experiments to determine the performance differences in various use cases? # If this method is primarily used in user code, would a new method to compare in reverse order be better, to account for both cases? Perhaps it'd be better to optimize the Combiner code, rather than change the default behavior for all cases? # I think relying on fall-through behavior in switch statements should be avoided. It's prone to error, especially as code is refactored over time. I think it's better to avoid it than to suppress the warning. This may be a style choice, but it's a preference that the java compiler weighs in on (by making fallthrough a default compiler warning), and I'd prefer to avoid behavior which results in compiler warnings whenever possible.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4468
16235	ACCUMULO-4468	{quote} bq. Do you have other examples of where this might be used in a tight loop? I think there are lots of other examples in Accumulo itself; for example, in iterators.system.DeletingIterator and iterators.user.TransformingIterator, this method is called in a loop. In our use case, we're doing a similar thing: building a larger object out of multiple rows, by finding groups of rows which are equal under ROW_COLFAM. When each object is built from only a few rows, the CF equality comparison returns false pretty often (which is to be expected), but only after comparing row IDs, which are always the same in practice. {quote} Perhaps my concern didn't come across. From my perspective: I am concerned with a performance change that makes one case better. We need to understand if the case you outlined is "the norm" or "the exception". I was hoping you had context on this. {quote} bq. How did you test this? What types of numbers did you see? I haven't been able to install it on a cluster to test. The test suite does pass with this patch applied. I think it's a minor change; in the "rows are equal" case the same amount of work is done as with the existing code, although the parts are accessed in the opposite order. They're still compared mostly-in-order, as isEqual does, but the comment in that function was inspiration to try reversing the comparison order. Aside from performance, the code seems cleaner to me: there's no more repetition of e.g. the check of row equality. The bytecode (with Oracle javac 1.8.0_92) is substantially smaller: 389 bytes versus 167. {quote} At risk of being anti-social, I am -1 on any change for performance without numbers coming with it. There are many great tools out there to benchmark changes and don't necessarily require use of a cluster (it might actually be harder to test on a cluster). [JMH|http://openjdk.java.net/projects/code-tools/jmh/] and [Google Caliper|https://github.com/google/caliper] are two good tools for micro-benchmarking.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4468

#	Requirement	Questions	Answered	Successfully Answered	URL
16236	ACCUMULO-4468	Great stuff, [~wmurnane]! A few thoughts: bq. the original equals() copied to the new class is called customVanilla, and the original equals() in the original Key class is called standardEquals. The numbers to compare are really the two custom* ones; the standardEquals value is given just to show that it's in the same ballpark. I'm not grok'ing the difference between customVanilla and standardEquals. Why the variance in these two? Shouldn't they be essentially equivalent? bq. As a user of the API, I'd rather not have to think about equalsForward() versus equalsBackward(). I concur with you here. bq. I agree that any change should start with prejudice against it. However, I think the numbers above prove my case: when keys are presented in sorted order, which happens often in Accumulo, the proposed method of comparing is slightly but noticeably faster. The degree of improvement depends on the data, but it doesn't perform worse than the current solution in any case that I tested. Great, I hoped you didn't take my initial prejudice badly :). This is a great start. I'm curious trying to tweak a couple of other things. Sharing your project was super useful. Ultimately, if these numbers are as they appear (better in some cases, no worse in others), this is a great improvement. Expecting large contiguous blocks of keys where row or row+cf change very infrequently makes sense to optimize. It appears that {{compareTo(Object}} is also using a separate code path, so I don't think this would have a big affect on things like creating RFiles for bulk imports. I need to search through usages though.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4468
16237	ACCUMULO-4468	[~kturner], what do you think we should do here?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4468
16238	ACCUMULO-4456	org.apache.thrift.transport.TTransportException: java.net.ConnectException: Connection refused at org.apache.thrift.transport.TSocket.open(TSocket.java:226) at org.apache.accumulo.core.rpc.ThriftUtil.createClientTransport(ThriftUtil.java:320) at org.apache.accumulo.core.client.impl.ThriftTransportPool.createNewTransport(ThriftTransportPool.java:478) at org.apache.accumulo.core.client.impl.ThriftTransportPool.getTransport(ThriftTransportPool.java:410) at org.apache.accumulo.core.client.impl.ThriftTransportPool.getTransport(ThriftTransportPool.java:388) at org.apache.accumulo.core.rpc.ThriftUtil.getClient(ThriftUtil.java:139) at org.apache.accumulo.core.rpc.ThriftUtil.getClientNoTimeout(ThriftUtil.java:106) at org.apache.accumulo.core.client.impl.MasterClient.getConnection(MasterClient.java:70) at org.apache.accumulo.core.client.impl.MasterClient.getConnectionWithRetry(MasterClient.java:47) at org.apache.accumulo.core.client.impl.TableOperationsImpl.finishFateOperation(TableOperationsImpl.java:279) at org.apache.accumulo.core.client.impl.TableOperationsImpl.doFateOperation(TableOperationsImpl.java:341) at org.apache.accumulo.core.client.impl.TableOperationsImpl.doFateOperation(TableOperationsImpl.java:293) at org.apache.accumulo.core.client.impl.TableOperationsImpl.doTableFateOperation(TableOperationsImpl.java:1465) at org.apache.accumulo.core.client.impl.TableOperationsImpl.delete(TableOperationsImpl.java:650) at com.github.joshelser.MasterInteraction.run(MasterInteraction.java:55) at com.github.joshelser.MasterInteraction.main(MasterInteraction.java:72) Caused by: java.net.ConnectException: Connection refused at java.net.PlainSocketImpl.socketConnect(Native Method) at java.net.AbstractPlainSocketImpl.doConnect(AbstractPlainSocketImpl.java:345) at java.net.AbstractPlainSocketImpl.connectToAddress(AbstractPlainSocketImpl.java:206) at java.net.AbstractPlainSocketImpl.connect(AbstractPlainSocketImpl.java:188) at java.net.SocksSocketImpl.connect(SocksSocketImpl.java:392) at java.net.Socket.connect(Socket.java:589) at org.apache.thrift.transport.TSocket.open(TSocket.java:221) ... 15 more 2016-09-14 19:17:00,667 [zookeeper.ZooCache] WARN : Saw (possibly) transient exception communicating with ZooKeeper, will retry org.apache.zookeeper.KeeperException$SessionExpiredException: KeeperErrorCode = Session expired for /accumulo/332fbfc6-a17b-4104-8e7c-17f05cc485c3 at org.apache.zookeeper.KeeperException.create(KeeperException.java:127) at org.apache.zookeeper.KeeperException.create(KeeperException.java:51) at org.apache.zookeeper.ZooKeeper.exists(ZooKeeper.java:1045) at org.apache.accumulo.fate.zookeeper.ZooCache$2.run(ZooCache.java:319) at org.apache.accumulo.fate.zookeeper.ZooCache$2.run(ZooCache.java:295) at org.apache.accumulo.fate.zookeeper.ZooCache$ZooRunnable.retry(ZooCache.java:190) at org.apache.accumulo.fate.zookeeper.ZooCache.get(ZooCache.java:347) at org.apache.accumulo.fate.zookeeper.ZooCache.get(ZooCache.java:282) at org.apache.accumulo.core.client.ZooKeeperInstance.getInstanceID(ZooKeeperInstance.java:177) at org.apache.accumulo.core.zookeeper.ZooUtil.getRoot(ZooUtil.java:40) at	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4456
16239	ACCUMULO-4456	[~elserj] I'm not sure I fully understand what the replication service is doing, but it's possibly different. There seem to be 3 layers of service wrappers here. I didn't see that everywhere else. Does your comment imply that my fix was correct, or that I missed something?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4456
16240	ACCUMULO-4455	Just to clarify (since I asked), my intent was not specifically automated tests (since multi-threaded things are often hard to codify). I believe [~ivan.bella] said that he was primarily testing 1.6.x with [~phrocker] expressing some concerns about changes in newer versions of Accumulo. It would be nice that, for whatever methods were used to identify the problem and verify the fix, they are repeated for the other active releases (expressed in the fixVersions) that this change also landed on. [~ivan.bella], would you have the cycles to take this on?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4455
16241	ACCUMULO-4453	Hey [~dlmarion], it looks like you introduced these new logging messages in https://github.com/apache/accumulo/commit/2c88388995310 80adfdb33ba75f25f4cdd643963. Do you happen to recall why exactly you did this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4453
16242	ACCUMULO-4450	Did you want to fix this one for 1.7.3 and 1.8.1 also?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4450
16243	ACCUMULO-4446	[~lstav] how close is this? I'll leave it in 1.8.1 for now	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4446
16244	ACCUMULO-4446	[~lstav] yeah I saw it in there. Looked like it was close, do you think it is ready? Should it be pushed to 1.8.2?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4446
16245	ACCUMULO-4446	[~mjwall] as far as this issue is concerned, I think it is ready. After talking with [~elserj], he suggested creating an abstract class (which is in ACCUMULO-4552). One thing though, is that I branched the work from master, not from the 1.8 branch. Should I make the changes to the 1.8 branch too? Not familiar enough with changes for different versions.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4446
16246	ACCUMULO-4446	[~lstav] ah, ok. Typically, I think you should work on a ticket on the oldest branch it is assigned to. For this ticket that would be 1.7.3. When complete against, the 1.7 branch, we can merge it up into 1.8 then master. The alternative would be to merge it into master and then backport into 1.8 and 1.7, but I think starting at 1.7 is cleaner. This particular change would not be too back to backport, but can you create a new PR against 1.7 and reference https://github.com/apache/accumulo/pull/195. I'll get it merged up. Thanks man	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4446

#	Requirement	Questions	Answered	Successfully Answered	URL
16247	ACCUMULO-4443	[~ctubbsii], It appears to me that you changed the semantics of PerformanceTests with https://github.com/apache/accumulo/commit/be5c4334a7a6fda223c3213c66f4f9acda8e9f46 By having both categories, {{\@Category(\{MiniClusterOnlyTests.class, PerformanceTests.class\})}}, developers who don't want to run the PerformanceTests have to explicit specify them in the excludedGroups too. This is a change, no? I don't think the documentation reflects this change. If I'm not off my rocker, can you please fix the documentation?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4443
16248	ACCUMULO-4437	I agree. But, that leaves open the question about what this category is *for*. Is it just "not-mini-only", or is it "cluster tests which are not-mini-only".	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4437
16249	ACCUMULO-4437	bq. I agree. But, that leaves open the question about what this category is for. Is it just "not-mini-only", or is it "cluster tests which are not-mini-only". Yeah, in Accumulo I make the lazy assertion that "ITs" involve interaction with a running Accumulo in some form (all ITs are cluster tests) because that is the 99% case. I know we have exceptions. Perhaps the distinction would be good future-proofing.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4437
16250	ACCUMULO-4432	Is there a way to specify uncategorized tests? If so, maybe we don't need this. If not, then we could also probably do this with a checkstyle rule or maven plugin, rather than write our own. If we go that route, and one doesn't already exist, I'd be willing to create a maven plugin to do the check, so it could be used by multiple projects.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4432
16251	ACCUMULO-4432	bq. Is there a way to specify uncategorized tests? If so, maybe we don't need this. I am not aware of such a way. bq. If not, then we could also probably do this with a checkstyle rule or maven plugin, rather than write our own. If we go that route, and one doesn't already exist, I'd be willing to create a maven plugin to do the check, so it could be used by multiple projects. I am not sure if checkstyle would be sufficient to be able to analyze the type hierarchy (given the inherited annotation). I would be happy to use anything that already exists or you would like to create.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4432
16252	ACCUMULO-4432	Uh, nevermind. I see what you are doing now [~ctubbsii]. Do you want this in 1.8.1?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4432
16253	ACCUMULO-4423	Hey Dima! No pitchforks at all! We're happy to have help where offered. Presently, Maven still drives everything (which, personally, I like). The downside is that it does require a little bit of configuration instead of just doing something like picking up an hbase-site.xml and auto-bootstrapping itself off of the classpath. https://github.com/apache/accumulo/blob/master/TESTING.md#standalone-cluster We also, somewhere, have some integration that can run all of the Integration Tests we have presently via MapReduce. I'm... not sure where the documentation on that is. I think your approach would be nice for that. [~ecn] or [~ctubbsii], do you have anything you could point us to on the mapreduce IT execution?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4423
16254	ACCUMULO-4423	:) /me removes chainmail. Alright, let me dig around some more. I've been spending a bit of time in recent days playing around with what I found in that {{TESTING.md}} and did run into some complications when trying to point to my standalone cluster (like, at least on 1.7.2, an undocumented need to include a {{accumulo.it.cluster.clientconf}} system property). Do you happen to have any other documentation on getting tests running in standalone mode? Once I have my mind around it sufficiently, I'd love to try contributing something to make life easier for lazy engineers like myself (maybe starting with clarifying some docs, but then hopefully extending options for how people can run tests).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4423
16255	ACCUMULO-4423	bq. I've been spending a bit of time in recent days playing around with what I found in that TESTING.md and did run into some complications when trying to point to my standalone cluster (like, at least on 1.7.2, an undocumented need to include a accumulo.it.cluster.clientconf system property). Do you happen to have any other documentation on getting tests running in standalone mode? Hrm, [~mdrob] and I had chatted in IRC a few times on the matter too. Sounds like there's some bit-rot happening? I think I have this automated at $dayjob, let me cross-ref what is done there with the docs. This would be good to make sure is documented. bq. Once I have my mind around it sufficiently, I'd love to try contributing something to make life easier for lazy engineers like myself (maybe starting with clarifying some docs, but then hopefully extending options for how people can run tests). That'd be awesome! I'm going to go ahead and add you as a contributor so you can be autonomous on JIRA :)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4423
16256	ACCUMULO-4423	Are we interested in this for 1.7 and 1.8 branches too?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4423
16257	ACCUMULO-4419	bq. This ticket is to account for work done elsewhere in which I've made the compression pool configurable such that we either don't use the pool at all or use an adjustable pool based on commons-pool [~phrocker], can you give a quick overview as to why we want wouldn't want to use the pool or would want the adjustable pool as you described? Best as I can recall, the pooling limits the number of Compressors we have resident in memory to (de)compress data that we're reading/writing. Each of these compressors has some internal byte array which consumed a "large" amount of memory on heap, causing JVM issues. Is that appropriate background for this?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4419
16258	ACCUMULO-4419	Thanks, Marc! This is exactly what I was looking for. Some further thoughts: bq. Internally we've used the non pooled method since CodecPool did not trim the compressors when they weren't being used. In some cases a large number of sources will increase the memory foot print, so this was for memory usage; however, the ability to choose a factory was something I added to accommodate those who wanted to make their own choice. So, avoid the pool entirely to make sure when a "scan storm" happens, you don't hold on to compressors for any longer than absolutely necessary? bq. In testing I found significant performance increases by changing the 1K buffer higher. Neat, is this reproducible at all? Seems like some sort of harness would be useful in ever revisiting this kind of testing. bq. I didn't want to confuse the focus of the PR, but I did want to allow the CompressorFactory implementation and its configuration the ability to be updated. Understood and appreciated. Dynamic configuration is super helpful and (IMO) the norm now.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4419
16259	ACCUMULO-4407	I'm not clear. Are you suggesting just changing our code to chroot to the instance ID, or are you suggesting changing the layout of the ZK /accumulo area, or both? It sounds like just the former, but driven by the fact that the layout is confusing. FWIW, I opened ACCUMULO-1719 awhile back because I found the layout unnecessarily confusing. That issue is a bit old now, but simplification of the layout could compliment the use of chroot (it'd certainly be easier to change the layout, without risking breakage, if we were using chroot).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4407
16260	ACCUMULO-4383	My jenkins won't catch Java 6 problems, because I can't really get Java 6 to work with my version of Jenkins and Maven. (dependency jars have newer class files in them... polluting the classpath... hard to resolve). That's probably less of an issue with 1.7 branch and Java 7. Sooner we close out the 1.6 branch, the better for issues like this.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4383

#	Requirement	Questions	Answered	Successfully Answered	URL
16261	ACCUMULO-4376	bq. I like this idea, and it was something I was thinking about doing for the 2.0.0 API changes. It might be best if we coordinate this a bit, and decide what the next version of Accumulo will be, with these changes. I'd hate to add them in the old API, only to provide replacements in 2.0. It sounds like you haven't actually done anything yet for the 2.0 API re-write you've been working on. Why would we have to have replacements for it in 2.0? How would we future proof something that doesn't exist yet? :) I just don't want this to get bogged down in 2.0 hypotheticals. If there are concrete things we can do to make sure this will fit well into your redesign, please do share that info.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4376
16262	ACCUMULO-4376	Is the API design for 2.0 documented somewhere?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4376
16263	ACCUMULO-4376	There are a lot of important things in ACCUMULO-2589 but nothing related to how Keys and other structures are created. Are you talking about some specific task/comment of that JIRA?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4376
16264	ACCUMULO-4375	What version are you using? I see some Key constructors in the 1.7 branch that take byte[].	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4375
16265	ACCUMULO-4375	I would also like to add {{Range}} methods working on bytes, do you want me to do it in this JIRA or should I open another one?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4375
16266	ACCUMULO-4360	In 1.5.0 API calls and a shell command were added to cancel compactions, are you aware of these?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4360
16267	ACCUMULO-4358	[~ctubbsii] [~elserj] - can one of you two take a look? I feel like I did everything right, but rendering locally doesn't work so I have zero confidence in pushing this.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4358
16268	ACCUMULO-4357	For posterity, how did you identify the list of fixes? Do we have automated tooling for it? (Is it part of findbugs? checkstyle? coverity?)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4357
16269	ACCUMULO-4356	Would this mechanism also give the downstream user an easy way to see what license those dependencies are under?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4356
16270	ACCUMULO-4356	actually, could we have a [DISCUSS] on dev@ about this? it's a pretty big change to how we package the project.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4356
16271	ACCUMULO-4353	bq. If multiple tablet servers die in quick succession, such as from a rolling restart of the Accumulo cluster or a network partition, this behavior can cause a storm of reassignment and rebalancing, placing significant load on the master. In the case of a rolling-restart, isn't that more of an operational issue to solve (give Accumulo time to process the reassignments before restarting more nodes)? Can you expand a bit more on the specifics behind a network partition that you're trying to solve with this? Say, you lose a node? A rack? Half of your racks? Are we talking about a 5second interruption/delay? 30s? Minutes? bq. To avert such load, Accumulo should be capable of maintaining a steady tablet assignment state in the face of transient tablet server loss. Instead of reassigning tablets as quickly as possible, Accumulo should be await the return of a temporarily downed tablet server (for some configurable duration) before assigning its tablets to other tablet servers. I'm a little worried about this as a configuration knob -- I feel like it kind of goes against the highly-available distributed database which we expect Accumulo to be. When we don't reassign tablets fast, that is a direct lack of availability for clients to read data. bq. placing significant load on the master Can you expand on this some more? Given that assignment is arguably the most important thing for the Master to do, why are we concerned about letting the master do that as fast as it can (for the aforementioned reason)? Do we need to come up with a more efficient way for the master to handle the reassignment of many tablets? For example, I know that HBase has special logic to batch assignments together in one RPC call to regionservers in order to bring things online more quickly (I'm not sure if we have logic like that -- instead we just spam requests to a tabletserver to load a tablet until it does so). While seeing a pull request accompanying the issue reported, It seems a bit premature to me to see code without some discussion on what the problems are and how best to solve them.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4353
16272	ACCUMULO-4353	Are you attempting to design a mechanism that could be used to avoid re-balancing and have the master keep assignments where they were previously, knowing that servers will come back into operation? I only ask because I question why the load on the master is a problem. You will cause load since clients will persist for that length of time. Wont' you increase the number of thrift connections waiting since you may not re-balance for some time?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4353
16273	ACCUMULO-4353	bq. Can you expand on this some more? Given that assignment is arguably the most important thing for the Master to do, why are we concerned about letting the master do that as fast as it can (for the aforementioned reason)? Do we need to come up with a more efficient way for the master to handle the reassignment of many tablets? Reading through this, and bringing some first-hand experience, I don't think the issue is the the Master assigning tablets. It's the issue of tablet servers that are down for a short period of time. When a tserver goes down, the Master re-assigns the tablets. When the tserver comes back up, it goes through several rounds of balancing which could take a long time and cause a lot of churn. bq. I'm a little worried about this as a configuration knob – I feel like it kind of goes against the highly-available distributed database which we expect Accumulo to be. When we don't reassign tablets fast, that is a direct lack of availability for clients to read data. I don't see any harm done here as long as the default behavior is what happens today. Allowing an administrator to choose to delay tablet reassignment may not fit most use cases, but it could fit some. My 2 cents.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4353

#	Requirement	Questions	Answered	Successfully Answered	URL
16274	ACCUMULO-4353	bq. Are you attempting to design a mechanism that could be used to avoid re-balancing and have the master keep assignments where they were previously, knowing that servers will come back into operation? That is the idea, yes. There is definitely a tradeoff here. bq. If this is really about trying to make rolling-restarts better, I'd encourage a look at ACCUMULO-1454. As I mentioned before, I hadn't seen ACCUMULO-1454 before starting this. I've now looked at the discussion of that ticket. What I implemented was approximately what Christopher Tubbs and David Medinets were suggesting. I also read through Keith Turner's design proposal summary. I have some reservations with it: * It requires that each planned restart involves tablet servers changing ports. While the recent changes to Accumulo to support a narrow port range during port search would make this more plausible, it might still prove difficult to establish firewall rules for Accumulo. (Sean Busby raises this issue in the discussion). * What happens if a tablet is split after migration starts? It seems to me there might be a race condition here which would lead to incomplete migration between sibling tablet servers. Do we block assignment during the rolling restart, too? That seems seems like a cure worse than the problem. * Even barring those two concerns, I again raise the spectre of ops complexity. To transition a single server, I need to know (a) which port the "old" tserver was running on, and (b) which port the "new" tserver is running on. If I'm using some sort of dynamic port assignment (which I would need to unless I pointed the "new" tserver at an entirely different configuration), it could be non-trivial to gather these pieces of information. While the burden on the operator of a cluster of 5 tservers might not be significant, the burden on the operator of a cluster of 200 tservers might make this approach infeasible. And the non-triviality of determining the correct port migration mapping would also make the process difficult to robustly automate. bq. While seeing a pull request accompanying the issue reported, It seems a bit premature to me to see code without some discussion on what the problems are and how best to solve them. Ahh, my mistake then. As a new contributor to Accumulo, I still don't have a full grasp of the rules, either written or unwritten. My feeling from watching the list was that primary modus operandi was to present a (fully implemented) solution along with a proposed problem, and then to discuss the merits of the solution.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4353
16275	ACCUMULO-4345	bq. May want to bump Accumulo's gson version to 2.2.4. Any idea on how Gson uses their version strings (do they actually follow something like semver)?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4345
16276	ACCUMULO-4345	I've been using Gson 2.3.1 in Fedora for some time now with Accumulo's packaging there. I don't know what they guarantee with regard to their versioning, but newer versions don't seem to cause any issues.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4345
16277	ACCUMULO-4328	(about to be anti-social...) Looking at the pull request, you're targeting this for 1.6, [~dlmarion]? I'd have (at least one) reservation in pulling this into 1.6: This really goes against the spirit of our release guidelines. This looks like it would have a huge impact on users. It's hard to quantify this because we have no good guarantees defined on the shell scripts/configuration. However, given that 1.6 is close to EOL, this change doesn't make sense to me for 1.6 or 1.7. For a 1.8 or 2.0, I think this would be a nice feature. It would be nice to get some high-level documentation on how it's expected to work before looking at the changes though: * How do we track multiple processes in the shell scripts? * How do we differentiate multiple tservers on a host internally (inside the master)? * What configuration do users need to make to take advantage of this? Since you already have a patch put together for this, I'm guessing I'm mostly looking for user manual updates (instead of a design doc).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328
16278	ACCUMULO-4328	bq. This looks like it would have a huge impact on users. By default, it does nothing different than what happens today in 1.6 except for the pid in the log (which could be undone in the logger properties) bq. How do we track multiple processes in the shell scripts? I backported the pid support, so each tserver instance is named 'tserver-N' and it has an associated pid file. bq. How do we differentiate multiple tservers on a host internally (inside the master)? By hostname and port, just like today bq. What configuration do users need to make to take advantage of this? Take a look at the commented example in accumulo-env.sh, and they would need to set the tserver client port to zero.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328
16279	ACCUMULO-4328	bq. By default, it does nothing different than what happens today in 1.6 except for the pid in the log (which could be undone in the logger properties) It still breaks usability between 1.6.5 and 1.6.6. Things that would work in 1.6.6 will not in 1.6.5 which is the entire spirit of our compatibility statement. bq. the pid in the log (which could be undone in the logger properties) They don't have separate log files? That sounds nightmareish to debug :) {quote} bq. What configuration do users need to make to take advantage of this? Take a look at the commented example in accumulo-env.sh, and they would need to set the tserver client port to zero. {quote} Lots of places like to have stable ports (for security or monitoring reasons). Maybe there's a better approach to be had than resorting to random ports?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328
16280	ACCUMULO-4328	bq. It still breaks usability between 1.6.5 and 1.6.6. Things that would work in 1.6.6 will not in 1.6.5 which is the entire spirit of our compatibility statement. I don't understand how it breaks usability. Can you provide an example? bq. They don't have separate log files? Correct - not sure how you would do this with Log4J bq. Lots of places like to have stable ports (for security or monitoring reasons). Maybe there's a better approach to be had than resorting to random ports? Agreed, currently the tserver.client.ports property does not take a list or a range. I think that would be the solution here.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328
16281	ACCUMULO-4328	{quote} bq. It still breaks usability between 1.6.5 and 1.6.6. Things that would work in 1.6.6 will not in 1.6.5 which is the entire spirit of our compatibility statement. I don't understand how it breaks usability. Can you provide an example? {quote} This is a _feature_ which will magically appear in a 6th release of the 1.6 line. All five previous releases in the 1.6 do not have this feature. This is against what we try to do. Users who see this feature in 1.6.6 might think that they can also do it in 1.6.3. This is extremely confusion and against what the versions are supposed to represent. Think about your experiences with configuration properties suddenly appearing in a .Z (of an x.y.z version string) in Hadoop. How jarring would that be to you if 2.6.4 had something that 2.6.1 didn't? {quote} bq. They don't have separate log files? Correct - not sure how you would do this with Log4J {quote} I think this would be pretty important to figure out. Not having separate log files per daemon would be very strange to try to explain.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328
16282	ACCUMULO-4328	bq. How jarring would that be to you if 2.6.4 had something that 2.6.1 didn't? You mean like https://issues.apache.org/jira/browse/HDFS-7694? Yes, I went and looked for one :-)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328

#	Requirement	Questions	Answered	Successfully Answered	URL
16283	ACCUMULO-4328	Upfront, I'm neutral on this and could go either way - with the caveat that going from 1.6.x to 1.6.6 with this capability does not change any behavior if the feature is not enabled. And that would include needing to modify a current user script in any way. If I can use a 1.6.x config and run it with 1.6.6 without needing to change anything, well, maybe. I'm less worried about going backward, but we would want it documented clearly so that people know what to expect if the needed to roll back and were counting on this. Random ports bother me. If there is an option to assign a port, or default to a random one if you don't care, okay - if I can only use a random port, that seems to be a problem if we try to be more container friendly as well as port blocking considerations. The big question that I have is, even if you could do this, would you want to? Are there any performance improvements running multiple tservers vs increasing thread counts and memory on a single tserver? I don't think we are normally CPU bound, and multiple tservers are going to be contending for the same IO bandwidth. If we had a benchmark where this did show substantial improvement, then maybe its okay for the 1.6 line. If we don't have that or the time to evaluate it, then 1.7+ may be a better option so that we don't get into a case where - sure its possible - but don't ever do that because A, B ... Z in 1.6.6, use 1.7.x where we fixed that... .	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328
16284	ACCUMULO-4328	bq. ...with the caveat that going from 1.6.x to 1.6.6 with this capability does not change any behavior if the feature is not enabled. And that would include needing to modify a current user script in any way. I believe this to be the case, that you should not have to change anything. bq. Random ports bother me. If there is an option to assign a port, or default to a random one if you don't care, okay - if I can only use a random port, that seems to be a problem if we try to be more container friendly as well as port blocking considerations. I agree that random ports may be an issue, but its an issue that predates this JIRA. I feel that if we are going to modify tserver.port.client to take a list or range of values instead of 1 value for the purposes that have been stated, then that should be a separate JIRA. However, for the purposes of this issue, if someone wants to use this feature, then they are also choosing to deal with random ports. bq. The big question that I have is, even if you could do this, would you want to? I don't have any specific evidence to share, but in talking with [~kturner] he expressed an interest in this feature. Maybe he can give more specifics. I will say that one way to scale Accumulo is to add more tablet servers, and just because you have saturated one tablet server does not mean that you have saturated the disk or network bandwidth of the host on which it's running.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328
16285	ACCUMULO-4328	{quote} The big question that I have is, even if you could do this, would you want to? Are there any performance improvements running multiple tservers vs increasing thread counts and memory on a single tserver? I don't think we are normally CPU bound, and multiple tservers are going to be contending for the same IO bandwidth. If we had a benchmark where this did show substantial improvement, then maybe its okay for the 1.6 line. If we don't have that or the time to evaluate it, then 1.7+ may be a better option so that we don't get into a case where - sure its possible - but don't ever do that because A, B ... Z in 1.6.6, use 1.7.x where we fixed that... {quote} I think [~afuchs] did some sort of benchmark a while back which showed some promise, but that was, IIRC, related to contention in rolling WALs -- something that [~ecn] made better with ACCUMULO-3423 in 1.8. Good benchmarks to understand the cases when this would have performance benefits would be great documentation to include. bq. I agree that random ports may be an issue, but its an issue that predates this JIRA. Yes and no, IMO. This is formalizing the ability to run multiple TServers as a "feature" which means it needs to be figured out. bq. Would re-basing this PR against 1.7 be a sufficient compromise? I'm a little less worried against 1.7, but am still not be in favor of it. I'd like to hear what others have to say before I blindly -1 it. The only difference with my complaints with it landing in 1.7 instead of 1.6 is that we have 4 fewer point-releases on 1.7 than 1.6, but that's still two releases which would have different semantics. [~kturner], [~busbey], you both often have good strong opinions on compatibility -- a bit of your time to provide some opinions would be appreciated.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4328
16286	ACCUMULO-4328	Hadoop is a terrible example to look at, because our versioning there is a mess. Surprising changes regularly show up in double-dot releases and changes that can only be undone with a system rollback happen on single-dot releases. For example I had to argue against changes to the datanode's underlying filesystem layout on a double-dot release recently. I personally wouldn't want to see this in a 1.6 or 1.7 release for largely the same reasons Josh states. Changing basics around how we do process management on a double-dot release is tempting trouble. It's not this particular PR's problem, but we have essentially no testing around our scripts so any verification of how 'safe' changes like this are essentially come down to some subset of us working through them manually with ill-defined criteria. But I agree with Dave that this kind of change would not violate the letter of our versioning policy, so I wouldn't -1 on principle. If problems showed up after we had a 1.6.z release that contained it I would, however, then strongly advocate for reverting it in later 1.6 releases using the same 'not the public API' argument. Aren't we getting ready to have a 1.8 release soon? What's so pressing about this particular new feature that means we should introduce risk in our established lines rather than roll forward with more minor releases? The lack of user-facing docs on what can be expected from this feature, when folks should use it, how folks should use it, etc cause me concern. I get that we pretty regularly have an issue with these kind of "finishing touches" on feature in Accumulo, but it's particularly troublesome to me when these kinds of 'in-the-know' things show up in maintenance releases. I'd also strongly prefer having the different tserver instances log to distinct files, regardless of what version this lands in. I get that I can use post-processing to figure out which lines correspond to a given process, but I'm skeptical of log4j's ability to efficiently and safely have multiple processes share a log file. This should be doable with a log4j configs since we control the filename via a classpath config file. I don't have strong feelings on how we go about doing that, wether it's example config files that show making a config-per-expected-process (preferably referenced/explained in the user manual) or a parameter in a single log4j config that the launching script sets via a system property.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4328
16287	ACCUMULO-4325	What does this configuration look like? Where would it live?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4325
16288	ACCUMULO-4325	bq. Since it's not actually part of our project, which we intend to "release" to users, I'd prefer we keep it out of our development branches, if possible (not sure if it's possible or not... it's just my preference). How do you propose we configure how each branch gets tested then? Is that not a part of the codebase? It seems extremely silly to have a completely separate repository solely for a single shell script per branch.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4325
16289	ACCUMULO-4325	Do we actually need a personality per branch? I've been updating the HBase personality with the assumption that we wouldn't over there, and I think [~aw] has been doing the same for Hadoop. On the matter of jdk versions changing per branch, there's YETUS-369. HBase is waiting for that change, but Allen has a workaround documented there now that Hadoop uses.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4325

#	Requirement	Questions	Answered	Successfully Answered	URL
16290	ACCUMULO-4324	[~kturner] mentioned we needed to add a new DATA_VERSION to prevent things like 1.7 successfully coming up on a 1.8 instance. One ticket that I know requires this change is ACCUMULO-1124, which changed the RFile index. Anyone know of other tickets that should be listed as part of this? I am about to put a PR for this. Then I want to test an upgrade locally. Here is my plan. 1 - startup 1.7 2 - write some data, scan it 3 - stop Accumulo 4 - startup 1.8 on top of that hdfs/zk 5 - assure can read data, write some more data and scan it 6 - stop Accumulo 7 - try to restart 1.7 and verify it fails Are there any other things that should be done in testing an upgrade?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4324
16291	ACCUMULO-4314	[~kturner] - are you still on track to get this done soon? I see that ACCUMULO-1124 has been completed.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4314
16292	ACCUMULO-4314	I ran test with the changes in 1.7 for this issue using the same file I was testing the changes for ACCUMULO-1124 with. The total index size went from 6.9M to 3.6M. {noformat} $ accumulo rfile-info /accumulo/tables/2/default_tablet/A0000005.rf Reading file: hdfs://localhost:10000/accumulo/tables/2/default_tablet/A0000005.rf Locality group : <DEFAULT> Start block : 0 Num blocks : 20,041 Index level 1 : 4,140 bytes 1 blocks Index level 0 : 3,620,079 bytes 14 blocks First key : um:d:385:%03;%01;10.30.170.244>>o>/2954%af; data:current [] 4611686019157309597 false Last key : um:d:395:%03;%01;%ff; com.facebook>.www>s>/dialog/feed?app_id=90376669494... TRUNCATED data:current [] -6917529026891043602 false Num entries : 24,299,468 Column families : [data] Meta block : BCFile.index Raw size : 4 bytes Compressed size : 12 bytes Compression type : gz Meta block : RFile.index Raw size : 4,258 bytes Compressed size : 2,154 bytes Compression type : gz {noformat}	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4314
16293	ACCUMULO-4308	I just visually examined all the jars and they all appear to have a {{*.java}} file for every {{*.class}} file, so at least their source is currently available.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4308
16294	ACCUMULO-4306	Could you expand a little more on the likely usecases for this? I'm having trouble following and this seems like a pathway to easy security compromises. 1) In the case of HDFS not being exposed to an external network, wouldn't this change allow a user-submitted coprocessor executing server side to read arbitrary backing files? 2) In the case of a multi-tenant deploy of HDFS where Accumulo requires kerberos but other tenants don't, wouldn't lacking kerberos on HDFS mean that other tenants could read backing files? Why would Accumulo then have a requirement for kerberos for just its tenants? Wouldn't it also have the same problem as in #1? Perhaps a better reason for the same changes needed here: how do we handle authenticating the TServers to multiple HDFS instances in a multi-volume deployment? Right now, I suspect they'd have to all use the same keytab/principal?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4306
16295	ACCUMULO-4306	I think both of those points are really the same thing, which is "Can't a process running on the cluster that isn't Accumulo read backing files?" The answer is most likely yes, given a couple of circumstances: 1) Any user in the accumulo group will be able to read backing files. Authentication is a non-factor for any such process. Potentially using something like Ranger would allow stricter authorization policies to mitigate this. 2) Without HDFS authentication enabled, a user could masquerade as another user via webhdfs and that allows them to get around the authorization restriction mentioned above. In my use case, we have some other safeguards in place to mitigate this, such as cell level encryption that would require credentials external to both HDFS & Accumulo. For clarity's sake, when I say "user" I mean a user as HDFS would see it. I was not referring an Accumulo user, which is, by default, a synthetic account in ZooKeeper. As the architecture currently sits, only a subset of users (the Accumulo service principals, {{accumulo/HOST@REALM}}) interact with HDFS. Mostly I'm trying to get around the fact that "security" is either "on" or "off" for all of HDFS, and I have specific read paths through Accumulo for external clients that I'd prefer be strongly authenticated so they can get their proper authorizations. I was thinking about whether or not an encryption zone over Accumulo's data directory would help as well, but so long as a user can pretend to be another user, they'd be able to do #2 to decrypt the directory's contents. I believe that the multi-volume case *should* work correctly right now. So long as each NN process is running as its own service principal, SASL+GSSAPI should be able to handle negotiation tickets between a given TServer (as {{accumulo/$HOST@REALM}}) and the NN as if it were any other service. Would an update to our documentation/user manual that outlines the consequences of security configurations (both current and those as a result of this ticket) help sway you one way or the other? I think there's already gaps in our current capabilities now that are undocumented, and this would just add more unknown variables. Specifically you've mentioned reading backing files, but there are other concerns from Accumulo's perspective (such as user authorizations) that are a separate class of protection mechanisms which I'm also trying to consider. As as always, I appreciate the comments Sean :)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4306
16296	ACCUMULO-4306	Well, point one is no different than "out of the box" right now anyways. If you are allowed to deploy any code into Accumulo, it must be trusted code. Am I missing something in your example?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4306
16297	ACCUMULO-4306	{quote} Would an update to our documentation/user manual that outlines the consequences of security configurations (both current and those as a result of this ticket) help sway you one way or the other? I think there's already gaps in our current capabilities now that are undocumented, and this would just add more unknown variables. Specifically you've mentioned reading backing files, but there are other concerns from Accumulo's perspective (such as user authorizations) that are a separate class of protection mechanisms which I'm also trying to consider. {quote} Yes, this would help. HDFS without kerberos enabled is a pretty big red flashing light in my experience, so it would especially help me evaluate the delta we're talking about for likely my-first-cluster misconfigurations.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4306
16298	ACCUMULO-4195	While it seems trivial, designing this going particularly well for me. A naive solution just ends up with some {{FileOperationsOptions}} object which holds all possible parameters for all {{FileOperations}} methods. While this solves one problem (changes to {{FileOperations}} inputs no longer require global code changes), it introduces another: it's no longer obvious which methods support which parameter groups. e.g. "Why can't I specify a range when opening a reader for seeking?" "What does the set of column families mean when I'm opening a writer?" "How do I rate limit a {{getFileSize()}}?" A separate {{...Options}} object for each of the 5 primary method types would be possible, but would be a lot of boilerplate, and might prove difficult to maintain. My first attempt in code attempted to meet in the middle, and ended up enjoying the worst of both designs.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4195

#	Requirement	Questions	Answered	Successfully Answered	URL
16299	ACCUMULO-4195	GitHub user ShawnWalker opened a pull request: https://github.com/apache/accumulo/pull/95 ACCUMULO-4195 Added generalized configuration objects for RFile interaction Implemented a builder/fluent style of syntax, where each operation creates an operation object with both some methods for setting parameters and an execute() method to actually invoke the operation. I'm not really sure if this is in line with the improvement the reporter suggests. It does however address the only goal I can identify with that request: it should be possible to extend the interface `FileOperations` in a manner such as was done for ACCUMULO-4187 without requiring significant changes to unrelated code. Also, the automatic formatter hates me, and completely undid most of my eye-pleasing formatting. You can merge this pull request into a Git repository by running: $ git pull https://github.com/ShawnWalker/accumulo ACCUMULO-4195 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/accumulo/pull/95.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #95 ---- commit 8917a37a4134a04bb47fac8e14320ea6f1916ed4 Author: Shawn Walker <accumulo@shawn-walker.net> Date: 2016-04-19T19:54:31Z ACCUMULO-4195 Added generalized configuration objects for RFile interaction. Implemented a builder/fluent style of syntax, where each operation creates an operation object with both some methods for setting parameters and an execute() method to actually invoke the operation. ----	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4195
16300	ACCUMULO-4195	core/src/main/java/org/apache/accumulo/core/file/FileOperations.java --- @@ -48,38 +48,295 @@ public static FileOperations getInstance() { return new DispatchingFileFactory(); } + // + // Abstract methods (to be implemented by subclasses) + // + + protected abstract long getFileSize(GetFileSizeOperation options) throws IOException; + + protected abstract FileSKVWriter openWriter(OpenWriterOperation options) throws IOException; + + protected abstract FileSKVIterator openIndex(OpenIndexOperation options) throws IOException; + + protected abstract FileSKVIterator openScanReader(OpenScanReaderOperation options) throws IOException; + + protected abstract FileSKVIterator openReader(OpenReaderOperation options) throws IOException; + + // + // File operations + // + + /** - * Open a reader that will not be seeked giving an initial seek location. This is useful for file operations that only need to scan data within a range and do - * not need to seek. Therefore file metadata such as indexes does not need to be kept in memory while the file is scanned. Also seek optimizations like bloom - * filters do not need to be loaded. + * Construct an operation object allowing one to query the size of a file. * Syntax: * <pre> * long size = fileOperations.getFileSize().ofFile(filename, fileSystem, fsConfiguration).withTableConfiguration(tableConf).execute(); * </pre> */ public GetFileSizeOperation getFileSize() { return new GetFileSizeOperation(); } - public abstract FileSKVIterator openReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - RateLimiter readLimiter, AccumuloConfiguration tableConf) throws IOException; + /** * Construct an operation object allowing one to create a writer for a file. * Syntax: * <pre> * FileSKVWriter writer = fileOperations.openWriter() * .ofFile(...) * .withTableConfiguration(...) * .withRateLimiter(...) // optional * .withCompression(...) // optional * .execute(); * </pre> */ public OpenWriterOperation openWriter() { return new OpenWriterOperation(); } /** * Construct an operation object allowing one to create an index iterator for a file. * Syntax: * <pre> * FileSKVIterator iterator = fileOperations.openIndex() * .ofFile(...) * .withTableConfiguration(...) * .withRateLimiter(...) // optional * .withBlockCache(...) // optional * .execute(); * </pre> */ public OpenIndexOperation openIndex() { return new OpenIndexOperation(); } - public abstract FileSKVIterator openReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - RateLimiter readLimiter, AccumuloConfiguration tableConf, BlockCache dataCache, BlockCache indexCache) throws IOException; + /** * Construct an operation object allowing one to create a "scan" reader for a file. Scan readers do not have any optimizations for seeking beyond their * initial position. This is useful for file operations that only need to scan data within a range and do not need to seek. Therefore file metadata such as * indexes does not need to be kept in memory while the file is scanned. Also seek optimizations like bloom filters do not need to be loaded. * Syntax: * <pre> * FileSKVIterator scanner = fileOperations.openScanReader() * .ofFile(...) * .overRange(...) * .withTableConfiguration(...) * .withRateLimiter(...) // optional * .withBlockCache(...) // optional * .execute(); * </pre> */ public OpenScanReaderOperation openScanReader() { return new OpenScanReaderOperation(); } /** - * Open a reader that fully support seeking and also enable any optimizations related to seeking, like bloom filters. * Construct an operation object allowing one to create a reader for a file. A reader constructed in this manner fully supports seeking, and also enables any * optimizations related to seeking (e.g. Bloom filters). * Syntax: * <pre> *	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4195

#	Requirement	Questions	Answered	Successfully Answered	URL
16301	ACCUMULO-4195	core/src/main/java/org/apache/accumulo/core/file/FileOperations.java --- @@ -48,38 +48,295 @@ public static FileOperations getInstance() { return new DispatchingFileFactory(); } + // + // Abstract methods (to be implemented by subclasses) + // + + protected abstract long getFileSize(GetFileSizeOperation options) throws IOException; + + protected abstract FileSKVWriter openWriter(OpenWriterOperation options) throws IOException; + + protected abstract FileSKVIterator openIndex(OpenIndexOperation options) throws IOException; + + protected abstract FileSKVIterator openScanReader(OpenScanReaderOperation options) throws IOException; + + protected abstract FileSKVIterator openReader(OpenReaderOperation options) throws IOException; + + // + // File operations + // + + /** - * Open a reader that will not be seeked giving an initial seek location. This is useful for file operations that only need to scan data within a range and do - * not need to seek. Therefore file metadata such as indexes does not need to be kept in memory while the file is scanned. Also seek optimizations like bloom - * filters do not need to be loaded. + * Construct an operation object allowing one to query the size of a file. + * Syntax: + * <pre> + * long size = fileOperations.getFileSize().ofFile(filename, fileSystem, fsConfiguration).withTableConfiguration(tableConf).execute(); + * </pre> */ + public GetFileSizeOperation getFileSize() { + return new GetFileSizeOperation(); + } - public abstract FileSKVIterator openReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - RateLimiter readLimiter, AccumuloConfiguration tableConf) throws IOException; + /** + * Construct an operation object allowing one to create a writer for a file. + * Syntax: + * <pre> + * FileSKVWriter writer = fileOperations.openWriter() .ofFile(...) .withTableConfiguration(...) .withRateLimiter(...) // optional .withCompression(...) // optional .execute(); + * </pre> */ + public OpenWriterOperation openWriter() { return new OpenWriterOperation(); } + + /** + * Construct an operation object allowing one to create an index iterator for a file. + * Syntax: + * <pre> + * FileSKVIterator iterator = fileOperations.openIndex() .ofFile(...) .withTableConfiguration(...) .withRateLimiter(...) // optional .withBlockCache(...) // optional .execute(); + * </pre> */ + public OpenIndexOperation openIndex() { return new OpenIndexOperation(); } - public abstract FileSKVIterator openReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - RateLimiter readLimiter, AccumuloConfiguration tableConf, BlockCache dataCache, BlockCache indexCache) throws IOException; + /** + * Construct an operation object allowing one to create a "scan" reader for a file. Scan readers do not have any optimizations for seeking beyond their + * initial position. This is useful for file operations that only need to scan data within a range and do not need to seek. Therefore file metadata such as + * indexes does not need to be kept in memory while the file is scanned. Also seek optimizations like bloom filters do not need to be loaded. + * Syntax: + * <pre> + * FileSKVIterator scanner = fileOperations.openScanReader() .ofFile(...) .overRange(...) .withTableConfiguration(...) .withRateLimiter(...) // optional .withBlockCache(...) // optional .execute(); + * </pre> */ + public OpenScanReaderOperation openScanReader() { return new OpenScanReaderOperation(); } - public abstract FileSKVIterator openScanReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - RateLimiter readLimiter, AccumuloConfiguration tableConf, BlockCache dataCache, BlockCache indexCache) throws IOException; + /** - * Open a reader that fully support seeking and also enable any optimizations related to seeking, like bloom filters. + * Construct an operation object allowing one to create a reader for a file. A reader constructed in this manner fully supports seeking, and also enables any + * optimizations related to seeking (e.g. Bloom filters). + * Syntax: + * <pre> + *	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4195
16302	ACCUMULO-4195	core/src/main/java/org/apache/accumulo/core/file/FileOperations.java --- @@ -48,38 +48,295 @@ public static FileOperations getInstance() { return new DispatchingFileFactory(); } + // + // Abstract methods (to be implemented by subclasses) + // + + protected abstract long getFileSize(GetFileSizeOperation options) throws IOException; + + protected abstract FileSKVWriter openWriter(OpenWriterOperation options) throws IOException; + + protected abstract FileSKVIterator openIndex(OpenIndexOperation options) throws IOException; + + protected abstract FileSKVIterator openScanReader(OpenScanReaderOperation options) throws IOException; + + protected abstract FileSKVIterator openReader(OpenReaderOperation options) throws IOException; + + // + // File operations + // + + /** - * Open a reader that will not be seeked giving an initial seek location. This is useful for file operations that only need to scan data within a range and do - * not need to seek. Therefore file metadata such as indexes does not need to be kept in memory while the file is scanned. Also seek optimizations like bloom - * filters do not need to be loaded. + * Construct an operation object allowing one to query the size of a file. + * Syntax: + * <pre> + * long size = fileOperations.getFileSize().ofFile(filename, fileSystem, fsConfiguration).withTableConfiguration(tableConf).execute(); + * </pre> */ + public GetFileSizeOperation getFileSize() { + return new GetFileSizeOperation(); + } - public abstract FileSKVIterator openReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - RateLimiter readLimiter, AccumuloConfiguration tableConf) throws IOException; + /** + * Construct an operation object allowing one to create a writer for a file. + * Syntax: + * <pre> + * FileSKVWriter writer = fileOperations.openWriter() .ofFile(...) .withTableConfiguration(...) .withRateLimiter(...) // optional .withCompression(...) // optional .execute(); + * </pre> */ + public OpenWriterOperation openWriter() { return new OpenWriterOperation(); } + + /** + * Construct an operation object allowing one to create an index iterator for a file. + * Syntax: + * <pre> + * FileSKVIterator iterator = fileOperations.openIndex() .ofFile(...) .withTableConfiguration(...) .withRateLimiter(...) // optional .withBlockCache(...) // optional .execute(); + * </pre> */ + public OpenIndexOperation openIndex() { return new OpenIndexOperation(); } - public abstract FileSKVIterator openReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - RateLimiter readLimiter, AccumuloConfiguration tableConf, BlockCache dataCache, BlockCache indexCache) throws IOException; + /** + * Construct an operation object allowing one to create a "scan" reader for a file. Scan readers do not have any optimizations for seeking beyond their + * initial position. This is useful for file operations that only need to scan data within a range and do not need to seek. Therefore file metadata such as + * indexes does not need to be kept in memory while the file is scanned. Also seek optimizations like bloom filters do not need to be loaded. + * Syntax: + * <pre> + * FileSKVIterator scanner = fileOperations.openScanReader() .ofFile(...) .overRange(...) .withTableConfiguration(...) .withRateLimiter(...) // optional .withBlockCache(...) // optional .execute(); + * </pre> */ + public OpenScanReaderOperation openScanReader() { return new OpenScanReaderOperation(); } - public abstract FileSKVIterator openScanReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - RateLimiter readLimiter, AccumuloConfiguration tableConf, BlockCache dataCache, BlockCache indexCache) throws IOException; + /** - * Open a reader that fully support seeking and also enable any optimizations related to seeking, like bloom filters. + * Construct an operation object allowing one to create a reader for a file. A reader constructed in this manner fully supports seeking, and also enables any + * optimizations related to seeking (e.g. Bloom filters). + * Syntax: + * <pre> + *	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4195

#	Requirement	Questions	Answered	Successfully Answered	URL
16303	ACCUMULO-4195	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/95#discussion_r60946455 --- Diff: core/src/test/java/org/apache/accumulo/core/client/mock/MockTableOperationsTest.java --- @@ -225,7 +225,8 @@ private ImportTestFilesAndData prepareTestFiles() throws Throwable { fs.delete(tempFile, true); fs.mkdirs(failures); fs.mkdirs(tempFile.getParent()); - FileSKVWriter writer = FileOperations.getInstance().openWriter(tempFile.toString(), fs, defaultConf, null, AccumuloConfiguration.getDefaultConfiguration()); + FileSKVWriter writer = FileOperations.getInstance().openWriter().ofFile(tempFile.toString(), fs, defaultConf) --- End diff -- The "verbage" here feels a little strange to me (maybe because you tried to preserve the original `openWriter()` methods?). What do you think about something like: ``` FileOperations.getInstance().newWriterBuilder().ofFile(tempFile.toString(), fs, defaultConf) .withTableConfiguration(AccumuloConfiguration.getDefaultConfiguration()).build(); ``` Using the common "builder" and "build" terminology makes it a bit obvious how this works. "openWriter" would imply to me that something is being opened (that I would need to close) and I'm not sure what "execute" would be doing at all on some I/O handle. I know this is a little nit-picky, but if we're going to make the APIs better, I'd be in favor of making them as clear as possible to prevent ourselves from coming back here soon.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4195
16304	ACCUMULO-4195	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/95#issuecomment-214438022 @keith-turner how, if at all, does this play into your RFile API work? Have you and Shawn talked about this at all offline?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4195
16305	ACCUMULO-4192	bq. I think the removal of that TraceRunnable was fixing a bug Sorry, I wasn't suggesting that we should reinstate it, but that we should check for {{Trace.wrap}} in the appropriate places. Would that cause more problems, [~billie.rinaldi]? I'm trying to remember the details of ACCUMULO-3507.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4192
16306	ACCUMULO-4189	Github user ctubbsii commented on a diff in the pull request: https://github.com/apache/accumulo/pull/91#discussion_r60137579 --- Diff: examples/latest.md --- @@ -0,0 +1,6 @@ +--- +title: Accumulo Examples +redirect_to: https://accumulo.apache.org/1.7/examples/ --- End diff -- Should these be relative paths instead?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4189
16307	ACCUMULO-4187	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/90#discussion_r59489739 --- Diff: core/src/main/java/org/apache/accumulo/core/file/streams/PositionedOutputs.java --- @@ -0,0 +1,55 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.file.streams; + +import java.io.FilterOutputStream; +import java.io.IOException; +import java.io.OutputStream; + +import org.apache.hadoop.fs.FSDataOutputStream; + +public class PositionedOutputs { + public static PositionedOutputStream wrap(final OutputStream fout) { + if (fout instanceof FSDataOutputStream) { + return new PositionedOutputStream(fout) { + @Override + public long position() throws IOException { + return ((FSDataOutputStream) fout).getPos(); + } + }; + } else if (fout instanceof PositionedOutput) { + return new PositionedOutputStream(fout) { + @Override + public long position() throws IOException { + return ((PositionedOutput) fout).position(); + } + }; + } else { + return new PositionedOutputStream(fout) { --- End diff -- check for non-null?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4187
16308	ACCUMULO-4187	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/90#discussion_r59490427 --- Diff: server/base/src/main/java/org/apache/accumulo/server/util/ratelimit/SharedRateLimiterFactory.java --- @@ -0,0 +1,152 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.util.ratelimit; + +import com.google.common.collect.ImmutableList; +import java.util.List; +import java.util.WeakHashMap; +import java.util.concurrent.Callable; +import java.util.concurrent.atomic.LongAdder; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.util.ratelimit.GuavaRateLimiter; +import org.apache.accumulo.core.util.ratelimit.RateLimiter; +import org.apache.accumulo.server.util.time.SimpleTimer; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +public class SharedRateLimiterFactory { + private static final long REPORT_RATE = 60000; + private static final long UPDATE_RATE = 1000; + private static SharedRateLimiterFactory instance = null; + private final Logger logger = LoggerFactory.getLogger(SharedRateLimiterFactory.class); + private final WeakHashMap<String,SharedRateLimiter> activeLimiters = new WeakHashMap<>(); + + private SharedRateLimiterFactory() {} + + public static synchronized SharedRateLimiterFactory getInstance(SimpleTimer timer) { + synchronized (SharedRateLimiterFactory.class) { --- End diff -- Why not just `public static synchronized`?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4187

#	Requirement	Questions	Answered	Successfully Answered	URL
16309	ACCUMULO-4187	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/90#discussion_r59490490 --- Diff: server/base/src/main/java/org/apache/accumulo/server/util/ratelimit/SharedRateLimiterFactory.java --- @@ -0,0 +1,152 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.util.ratelimit; + +import com.google.common.collect.ImmutableList; +import java.util.List; +import java.util.WeakHashMap; +import java.util.concurrent.Callable; +import java.util.concurrent.atomic.LongAdder; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.util.ratelimit.GuavaRateLimiter; +import org.apache.accumulo.core.util.ratelimit.RateLimiter; +import org.apache.accumulo.server.util.time.SimpleTimer; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +public class SharedRateLimiterFactory { + private static final long REPORT_RATE = 60000; + private static final long UPDATE_RATE = 1000; + private static SharedRateLimiterFactory instance = null; + private final Logger logger = LoggerFactory.getLogger(SharedRateLimiterFactory.class); + private final WeakHashMap<String,SharedRateLimiter> activeLimiters = new WeakHashMap<>(); + + private SharedRateLimiterFactory() {} + + public static SharedRateLimiterFactory getInstance(SimpleTimer timer) { + synchronized (SharedRateLimiterFactory.class) { + if (instance == null) { + instance = new SharedRateLimiterFactory(); + // Update periodically + timer.schedule(new Runnable() { --- End diff -- Can we re-poll the configuration and reschedule the task to make the implementation responsive to dynamic configuration update? Having to restart the server to get a new configuration value sucks.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4187
16310	ACCUMULO-4187	server/base/src/main/java/org/apache/accumulo/server/util/ratelimit/SharedRateLimiterFactory.java --- @@ -0,0 +1,152 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.util.ratelimit; + +import com.google.common.collect.ImmutableList; +import java.util.List; +import java.util.WeakHashMap; +import java.util.concurrent.Callable; +import java.util.concurrent.atomic.LongAdder; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.util.ratelimit.GuavaRateLimiter; +import org.apache.accumulo.core.util.ratelimit.RateLimiter; +import org.apache.accumulo.server.util.time.SimpleTimer; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +public class SharedRateLimiterFactory { + private static final long REPORT_RATE = 60000; + private static final long UPDATE_RATE = 1000; + private static SharedRateLimiterFactory instance = null; + private final Logger logger = LoggerFactory.getLogger(SharedRateLimiterFactory.class); + private final WeakHashMap<String,SharedRateLimiter> activeLimiters = new WeakHashMap<>(); + + private SharedRateLimiterFactory() {} + + public static SharedRateLimiterFactory getInstance(SimpleTimer timer) { + synchronized (SharedRateLimiterFactory.class) { + if (instance == null) { + instance = new SharedRateLimiterFactory(); + // Update periodically + timer.schedule(new Runnable() { + @Override + public void run() { + instance.update(); + } + }, UPDATE_RATE, UPDATE_RATE); + // Report periodically + timer.schedule(new Runnable() { + @Override + public void run() { + instance.report(); + } + }, REPORT_RATE, REPORT_RATE); + } + return instance; + } + } + + public static SharedRateLimiterFactory getInstance(AccumuloConfiguration conf) { + return getInstance(SimpleTimer.getInstance(conf)); + } + + /** + * Lookup the RateLimiter associated with the specified name, or create a new one for that name. RateLimiters should be closed when no longer needed. + * + * @param name + * key for the rate limiter + * @param rateGenerator + * a function which can be called to get what the current rate for the rate limiter should be. + */ + public RateLimiter create(String name, Callable<Long> rateGenerator) { + synchronized (activeLimiters) { + if (activeLimiters.containsKey(name)) { + SharedRateLimiter limiter = activeLimiters.get(name); + return limiter; + } else { + long initialRate; + try { + initialRate = rateGenerator.call(); + } catch (Exception ex) { + throw new IllegalStateException(ex); + } + SharedRateLimiter limiter = new SharedRateLimiter(name, rateGenerator, initialRate); + activeLimiters.put(name, limiter); + return limiter; + } + } + } + + protected void update() { + List<SharedRateLimiter> limiters; + synchronized (activeLimiters) { + limiters = ImmutableList.copyOf(activeLimiters.values()); + } + for (SharedRateLimiter limiter : limiters) { + limiter.update(); + } + } + + protected void report() { + List<SharedRateLimiter> limiters; + synchronized (activeLimiters) { + limiters = ImmutableList.copyOf(activeLimiters.values()); + } + for (SharedRateLimiter limiter : limiters) { + limiter.report(); + } + } + + protected class SharedRateLimiter extends GuavaRateLimiter { + private final final	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4187

#	Requirement	Questions	Answered	Successfully Answered	URL
16311	ACCUMULO-4187	Github user ctubbsii commented on a diff in the pull request: https://github.com/apache/accumulo/pull/90#discussion_r59571367 --- Diff: core/src/main/java/org/apache/accumulo/core/file/FileOperations.java --- @@ -55,23 +56,24 @@ public static FileOperations getInstance() { */ public abstract FileSKVIterator openReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - AccumuloConfiguration tableConf) throws IOException; + RateLimiter readLimiter, AccumuloConfiguration tableConf) throws IOException; public abstract FileSKVIterator openReader(String file, Range range, Set<ByteSequence> columnFamilies, boolean inclusive, FileSystem fs, Configuration conf, - AccumuloConfiguration tableConf, BlockCache dataCache, BlockCache indexCache) throws IOException; + RateLimiter readLimiter, AccumuloConfiguration tableConf, BlockCache dataCache, BlockCache indexCache) throws IOException; /** * Open a reader that fully support seeking and also enable any optimizations related to seeking, like bloom filters. * */ - public abstract FileSKVIterator openReader(String file, boolean seekToBeginning, FileSystem fs, Configuration conf, AccumuloConfiguration acuconf) - throws IOException; + public abstract FileSKVIterator openReader(String file, boolean seekToBeginning, FileSystem fs, Configuration conf, RateLimiter readLimiter, --- End diff -- So, I've had some thoughts on this as well. A lot of our internal passing around of AccumuloConfiguration makes it easy to lose track of where the configuration object came from, making it harder to ensure we're not using it outside the intended scope (for example, if it represents a TableConfiguration). In many parts of our internal code, I would prefer we grab what we need from AccumuloConfiguration, combine it with any additional configuration appropriate for that specific context, and create a new context-specific configuration object to pass around within that context. If we were to do something like that, what would the context granularity be here? "file operations options", "file reading options", something else?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4187
16312	ACCUMULO-4187	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/90#discussion_r59578033 --- Diff: core/src/main/java/org/apache/accumulo/core/util/ratelimit/GuavaRateLimiter.java --- @@ -0,0 +1,51 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.util.ratelimit; + +/** Rate limiter from the Guava library. */ +public class GuavaRateLimiter implements RateLimiter { + private final com.google.common.util.concurrent.RateLimiter rateLimiter; + private long currentRate; + + public GuavaRateLimiter(long initialRate) { + this.currentRate = initialRate; --- End diff -- should there be a sanity check here to ensure non negative? or are there sufficient checks elsewhere?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4187
16313	ACCUMULO-4187	server/base/src/main/java/org/apache/accumulo/server/util/ratelimit/SharedRateLimiterFactory.java --- @@ -0,0 +1,152 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.util.ratelimit; +import com.google.common.collect.ImmutableList; +import java.util.List; +import java.util.WeakHashMap; +import java.util.concurrent.Callable; +import java.util.concurrent.atomic.LongAdder; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.util.ratelimit.GuavaRateLimiter; +import org.apache.accumulo.core.util.ratelimit.RateLimiter; +import org.apache.accumulo.server.util.time.SimpleTimer; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +public class SharedRateLimiterFactory { + private static final long REPORT_RATE = 60000; + private static final long UPDATE_RATE = 1000; + private static SharedRateLimiterFactory instance = null; + private final Logger logger = LoggerFactory.getLogger(SharedRateLimiterFactory.class); + private final WeakHashMap<String,SharedRateLimiter> activeLimiters = new WeakHashMap<>(); + private SharedRateLimiterFactory() {} + public static SharedRateLimiterFactory getInstance(SimpleTimer timer) { + synchronized (SharedRateLimiterFactory.class) { + if (instance == null) { + instance = new SharedRateLimiterFactory(); + // Update periodically + timer.schedule(new Runnable() { + @Override public void run() { instance.update(); } }, UPDATE_RATE, UPDATE_RATE); + // Report periodically timer.schedule(new Runnable() { @Override public void run() { instance.report(); } }, REPORT_RATE, REPORT_RATE); + } return instance; } } + public static SharedRateLimiterFactory getInstance(AccumuloConfiguration conf) { return getInstance(SimpleTimer.getInstance(conf)); } + /** * Lookup the RateLimiter associated with the specified name, or create a new one for that name. RateLimiters should be closed when no longer needed. * * @param name key for the rate limiter * @param rateGenerator * a function which can be called to get what the current rate for the rate limiter should be. */ public RateLimiter create(String name, Callable<Long> rateGenerator) { synchronized (activeLimiters) { if (activeLimiters.containsKey(name)) { SharedRateLimiter limiter = activeLimiters.get(name); return limiter; } else { long initialRate; try { initialRate = rateGenerator.call(); } catch (Exception ex) { throw new IllegalStateException(ex); } SharedRateLimiter limiter = new SharedRateLimiter(name, rateGenerator, initialRate); activeLimiters.put(name, limiter); return limiter; } } } protected void update() { List<SharedRateLimiter> limiters; synchronized (activeLimiters) { limiters = ImmutableList.copyOf(activeLimiters.values()); } for (SharedRateLimiter limiter : limiters) { limiter.update(); } } protected void report() { List<SharedRateLimiter> limiters; synchronized (activeLimiters) { limiters = ImmutableList.copyOf(activeLimiters.values()); } for (SharedRateLimiter limiter : limiters) { limiter.report(); } } protected class SharedRateLimiter extends GuavaRateLimiter { private final	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4187

#	Requirement	Questions	Answered	Successfully Answered	URL
16314	ACCUMULO-4187	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/90#discussion_r59585544 --- Diff: core/src/main/java/org/apache/accumulo/core/util/ratelimit/GuavaRateLimiter.java --- @@ -0,0 +1,51 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.util.ratelimit; + +/** Rate limiter from the Guava library. */ +public class GuavaRateLimiter implements RateLimiter { + private final com.google.common.util.concurrent.RateLimiter rateLimiter; + private long currentRate; + + public GuavaRateLimiter(long initialRate) { + this.currentRate = initialRate; --- End diff -- The docs for `tserver.compaction.major.throughput` specify using 0 for unlimited. Is specifying 0 or negative documented elsewhere?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4187
16315	ACCUMULO-4184	Github user ivakegg commented on a diff in the pull request: https://github.com/apache/accumulo/pull/93#discussion_r59746357 --- Diff: server/base/src/main/java/org/apache/accumulo/server/master/balancer/HostRegexTableLoadBalancer.java --- @@ -56,18 +61,24 @@ * table.custom.balancer.host.regex.pool.check=5m * Regex matching can be based on either the host name (default) or host ip address. To set this balancer to match the regular expressions to the tablet server * IP address, then set the following property: - * table.custom.balancer.host.regex.is.ip=true + * table.custom.balancer.host.regex.is.ip=true + * It's possible that this balancer may create a lot of migrations. To limit the number of migrations that are created during a balance call, set the following + * property (default 250): + * table.custom.balancer.host.regex.concurrent.migrations * */ -public class HostRegexTableLoadBalancer extends TableLoadBalancer { +public class HostRegexTableLoadBalancer extends TableLoadBalancer implements ConfigurationObserver { private static final Logger LOG = LoggerFactory.getLogger(HostRegexTableLoadBalancer.class); public static final String HOST_BALANCER_PREFIX = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex."; - public static final String HOST_BALANCER_OOB_CHECK = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex.oob.period"; + public static final String HOST_BALANCER_OOB_CHECK_KEY = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex.oob.period"; --- End diff -- How many oobs does a host balancer have to check? (in other words a more obvious property name might be nice).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4184
16316	ACCUMULO-4181	bq. The generic logger config, which I copied straight from the provided template, has the following: I'm sorry, but I'm not following. Can you specify in terms of the file: was the above what you had in monitor_logger.xml? Or generic_logger.xml? Are you saying that you (accidentally) configured the monitor to send updates to itself? Taking a glance at the example configs in the master branch, they look right to me. I just want to make sure we're not shipping stuff which is easy to mess up. I appreciate you helping out with the extra info.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4181
16317	ACCUMULO-4179	Github user madrob commented on a diff in the pull request: https://github.com/apache/accumulo/pull/98#discussion_r62910850 --- Diff: README.md --- @@ -113,3 +115,9 @@ with Accumulo 1.6.2 and 1.7.0 all API changes will follow [semver 2.0][12] [14]: examples/simple/src/main/java/org/apache/accumulo/examples/simple [15]: docs/src/main/resources/examples [16]: docs/src/main/asciidoc +[li]: http://img.shields.io/badge/license-ASL-blue.svg +[ll]: https://github.com/apache/accumulo/blob/master/LICENSE --- End diff -- Could this maybe link to our license on an apache.org host?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4179
16318	ACCUMULO-4179	Can this be closed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4179
16319	ACCUMULO-4178	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/87#discussion_r58382167 --- Diff: core/src/main/java/org/apache/accumulo/core/client/impl/TabletServerBatchWriter.java --- @@ -831,11 +814,16 @@ public void send(TabletServerMutations<Mutation> tsm) throws AccumuloServerExcep try { - long count = 0; - for (List<Mutation> list : mutationBatch.values()) { - count += list.size(); + + Set<Text> tableIds = new TreeSet<Text>(); + for (Map.Entry<KeyExtent,List<Mutation>> entry : mutationBatch.entrySet()) { + count += entry.getValue().size(); + tableIds.add(entry.getKey().getTableId()); } - String msg = "sending " + String.format("%,d", count) + " mutations to " + String.format("%,d", mutationBatch.size()) + " tablets at " + location; - Thread.currentThread().setName(msg); + StringBuilder msg = new StringBuilder(); --- End diff -- When java generates code for string concatenation, it uses StringBuilder. Is there any value in this case for using StringBuilder over concatenation?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4178
16320	ACCUMULO-4178	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/87#discussion_r58382980 --- Diff: core/src/main/java/org/apache/accumulo/core/client/impl/TabletServerBatchWriter.java --- @@ -16,37 +16,7 @@ */ package org.apache.accumulo.core.client.impl; -import java.io.IOException; --- End diff -- What did you use for import manipulation? I don't think the Accumulo project has declared how imports have been organized. However I think we do try to organize imports using Eclipse's default settings inorder to cut down on noise generated by different IDEs organizing imports in different ways. Can we just drop this portion of the patch? Did the code add any imports?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4178
16321	ACCUMULO-4177	FWIW, master is officially on minJdk=8. It would be neat to run a little benchmark to see what kind of a difference TinyLFU would make. Do you have any examples of tests that other projects used, [~ben.manes]?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4177
16322	ACCUMULO-4177	Thanks again, Ben, for the very thoughtful response. What do others think? I think this would be nice to investigate since we completely missed the BucketCache-train from HBase (and stayed on LRUMap this entire time).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4177
16323	ACCUMULO-4176	What about the original point of why you added this in the first place (ACCUMULO-4004)? The description for the property seems woefully sparse compared to the problem it was meant to work around. I assume this is something that you've run into in production Accumulo instances which is why I thought it would deserve documentation about why this property is important, the symptoms in which you might want to adjust it, and the implications of changing it. Am I blowing it out of proportion?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4176
16324	ACCUMULO-4176	The reason for the change is to work around issues with decommissioning in certain version of Hadoop. Do we have a section for things like that?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4176

#	Requirement	Questions	Answered	Successfully Answered	URL
16325	ACCUMULO-4173	server/base/src/main/java/org/apache/accumulo/server/master/balancer/HostRegexTableLoadBalancer.java --- @@ -0,0 +1,314 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.master.balancer; + +import java.net.InetAddress; +import java.net.UnknownHostException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; +import java.util.Set; +import java.util.SortedMap; +import java.util.TreeMap; +import java.util.regex.Pattern; + +import org.apache.accumulo.core.client.admin.TableOperations; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.data.impl.KeyExtent; +import org.apache.accumulo.core.master.thrift.TabletServerStatus; +import org.apache.accumulo.core.tabletserver.thrift.TabletStats; +import org.apache.accumulo.server.conf.ServerConfiguration; +import org.apache.accumulo.server.master.state.TServerInstance; +import org.apache.accumulo.server.master.state.TabletMigration; +import org.apache.thrift.TException; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +/** + * This balancer will create pools of tablet servers by grouping tablet servers that match a regex into the same pool and calling the balancer set on the table + * to balance within the set of matching tablet servers. + * Regex properties for this balancer are specified as: + * table.custom.balancer.host.regex.<tablename>=<regex> + * Periodically (default 5m) this balancer will check to see if a tablet server is hosting tablets that it should not be according to the regex configuration. + * If this occurs then the offending tablets will be reassigned. This would cover the case where the configuration is changed and the master is restarted while + * the tablet servers are up. To change the out of bounds check time period, set the following property: + * table.custom.balancer.host.regex.oob.period=5m + * Periodically (default 5m) this balancer will regroup the set of current tablet servers into pools based on regexes applied to the tserver host names. This + * would cover the case of tservers dying or coming online. To change the host pool check time period, set the following property: + * table.custom.balancer.host.regex.pool.check=5m + * Regex matching can be based on either the host name (default) or host ip address. To set this balancer to match the regular expressions to the tablet server + * IP address, then set the following property: + * table.custom.balancer.host.regex.is.ip=true + * + */ +public class HostRegexTableLoadBalancer extends TableLoadBalancer { + + private static final Logger LOG = LoggerFactory.getLogger(HostRegexTableLoadBalancer.class); + public static final String HOST_BALANCER_PREFIX = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex."; + public static final String HOST_BALANCER_OOB_CHECK = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex.oob.period"; + private static final String HOST_BALANCER_OOB_DEFAULT = "5m"; +	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4173
16326	ACCUMULO-4173	server/base/src/main/java/org/apache/accumulo/server/master/balancer/HostRegexTableLoadBalancer.java --- @@ -0,0 +1,314 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.master.balancer; + +import java.net.InetAddress; +import java.net.UnknownHostException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; +import java.util.Set; +import java.util.SortedMap; +import java.util.TreeMap; +import java.util.regex.Pattern; + +import org.apache.accumulo.core.client.admin.TableOperations; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.data.impl.KeyExtent; +import org.apache.accumulo.core.master.thrift.TabletServerStatus; +import org.apache.accumulo.core.tabletserver.thrift.TabletStats; +import org.apache.accumulo.server.conf.ServerConfiguration; +import org.apache.accumulo.server.master.state.TServerInstance; +import org.apache.accumulo.server.master.state.TabletMigration; +import org.apache.thrift.TException; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +/** + * This balancer will create pools of tablet servers by grouping tablet servers that match a regex into the same pool and calling the balancer set on the table + * to balance within the set of matching tablet servers. + * Regex properties for this balancer are specified as: + * table.custom.balancer.host.regex.<tablename>=<regex> + * Periodically (default 5m) this balancer will check to see if a tablet server is hosting tablets that it should not be according to the regex configuration. + * If this occurs then the offending tablets will be reassigned. This would cover the case where the configuration is changed and the master is restarted while + * the tablet servers are up. To change the out of bounds check time period, set the following property: + * table.custom.balancer.host.regex.oob.period=5m + * Periodically (default 5m) this balancer will regroup the set of current tablet servers into pools based on regexes applied to the tserver host names. This + * would cover the case of tservers dying or coming online. To change the host pool check time period, set the following property: + * table.custom.balancer.host.regex.pool.check=5m + * Regex matching can be based on either the host name (default) or host ip address. To set this balancer to match the regular expressions to the tablet server + * IP address, then set the following property: + * table.custom.balancer.host.regex.is.ip=true + * + */ +public class HostRegexTableLoadBalancer extends TableLoadBalancer { + + private static final Logger LOG = LoggerFactory.getLogger(HostRegexTableLoadBalancer.class); + public static final String HOST_BALANCER_PREFIX = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex."; + public static final String HOST_BALANCER_OOB_CHECK = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex.oob.period"; + private static final String HOST_BALANCER_OOB_DEFAULT = "5m"; +	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4173

#	Requirement	Questions	Answered	Successfully Answered	URL
16327	ACCUMULO-4173	server/base/src/main/java/org/apache/accumulo/server/master/balancer/HostRegexTableLoadBalancer.java --- @@ -0,0 +1,314 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.master.balancer; + +import java.net.InetAddress; +import java.net.UnknownHostException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; +import java.util.Set; +import java.util.SortedMap; +import java.util.TreeMap; +import java.util.regex.Pattern; + +import org.apache.accumulo.core.client.admin.TableOperations; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.data.impl.KeyExtent; +import org.apache.accumulo.core.master.thrift.TabletServerStatus; +import org.apache.accumulo.core.tabletserver.thrift.TabletStats; +import org.apache.accumulo.server.conf.ServerConfiguration; +import org.apache.accumulo.server.master.state.TServerInstance; +import org.apache.accumulo.server.master.state.TabletMigration; +import org.apache.thrift.TException; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +/** + * This balancer will create pools of tablet servers by grouping tablet servers that match a regex into the same pool and calling the balancer set on the table + * to balance within the set of matching tablet servers. + * Regex properties for this balancer are specified as: + * table.custom.balancer.host.regex.<tablename>=<regex> + * Periodically (default 5m) this balancer will check to see if a tablet server is hosting tablets that it should not be according to the regex configuration. + * If this occurs then the offending tablets will be reassigned. This would cover the case where the configuration is changed and the master is restarted while + * the tablet servers are up. To change the out of bounds check time period, set the following property: + * table.custom.balancer.host.regex.oob.period=5m + * Periodically (default 5m) this balancer will regroup the set of current tablet servers into pools based on regexes applied to the tserver host names. This + * would cover the case of tservers dying or coming online. To change the host pool check time period, set the following property: + * table.custom.balancer.host.regex.pool.check=5m + * Regex matching can be based on either the host name (default) or host ip address. To set this balancer to match the regular expressions to the tablet server + * IP address, then set the following property: + * table.custom.balancer.host.regex.is.ip=true + * + */ +public class HostRegexTableLoadBalancer extends TableLoadBalancer { + + private static final Logger LOG = LoggerFactory.getLogger(HostRegexTableLoadBalancer.class); + public static final String HOST_BALANCER_PREFIX = + Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex."; + public static final String HOST_BALANCER_OOB_CHECK = + Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex.oob.period"; + private static final String HOST_BALANCER_OOB_DEFAULT = "5m"; +	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4173
16328	ACCUMULO-4173	server/base/src/main/java/org/apache/accumulo/server/master/balancer/HostRegexTableLoadBalancer.java --- @@ -0,0 +1,314 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.master.balancer; + +import java.net.InetAddress; +import java.net.UnknownHostException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; +import java.util.Set; +import java.util.SortedMap; +import java.util.TreeMap; +import java.util.regex.Pattern; + +import org.apache.accumulo.core.client.admin.TableOperations; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.data.impl.KeyExtent; +import org.apache.accumulo.core.master.thrift.TabletServerStatus; +import org.apache.accumulo.core.tabletserver.thrift.TabletStats; +import org.apache.accumulo.server.conf.ServerConfiguration; +import org.apache.accumulo.server.master.state.TServerInstance; +import org.apache.accumulo.server.master.state.TabletMigration; +import org.apache.thrift.TException; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +/** + * This balancer will create pools of tablet servers by grouping tablet servers that match a regex into the same pool and calling the balancer set on the table + * to balance within the set of matching tablet servers. + * Regex properties for this balancer are specified as: + * table.custom.balancer.host.regex.<tablename>=<regex> + * Periodically (default 5m) this balancer will check to see if a tablet server is hosting tablets that it should not be according to the regex configuration. + * If this occurs then the offending tablets will be reassigned. This would cover the case where the configuration is changed and the master is restarted while + * the tablet servers are up. To change the out of bounds check time period, set the following property: + * table.custom.balancer.host.regex.oob.period=5m + * Periodically (default 5m) this balancer will regroup the set of current tablet servers into pools based on regexes applied to the tserver host names. This + * would cover the case of tservers dying or coming online. To change the host pool check time period, set the following property: + * table.custom.balancer.host.regex.pool.check=5m + * Regex matching can be based on either the host name (default) or host ip address. To set this balancer to match the regular expressions to the tablet server + * IP address, then set the following property: + * table.custom.balancer.host.regex.is.ip=true + * + */ +public class HostRegexTableLoadBalancer extends TableLoadBalancer { + + private static final Logger LOG = LoggerFactory.getLogger(HostRegexTableLoadBalancer.class); + public static final String HOST_BALANCER_PREFIX = + Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex."; + public static final String HOST_BALANCER_OOB_CHECK = + Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex.oob.period"; + private static final String HOST_BALANCER_OOB_DEFAULT = "5m"; +	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4173
16329	ACCUMULO-4173	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/83#issuecomment-208921993 > Any final comments? I'll take a look again since you asked :)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4173

#	Requirement	Questions	Answered	Successfully Answered	URL
16330	ACCUMULO-4173	server/base/src/main/java/org/apache/accumulo/server/master/balancer/HostRegexTableLoadBalancer.java --- @@ -0,0 +1,334 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.master.balancer; + +import java.net.InetAddress; +import java.net.UnknownHostException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; +import java.util.Set; +import java.util.SortedMap; +import java.util.TreeMap; +import java.util.regex.Pattern; + +import org.apache.accumulo.core.client.admin.TableOperations; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.data.impl.KeyExtent; +import org.apache.accumulo.core.master.thrift.TabletServerStatus; +import org.apache.accumulo.core.tabletserver.thrift.TabletStats; +import org.apache.accumulo.server.conf.ServerConfiguration; +import org.apache.accumulo.server.master.state.TServerInstance; +import org.apache.accumulo.server.master.state.TabletMigration; +import org.apache.thrift.TException; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + ++/** + * This balancer will create pools of tablet servers by grouping tablet servers that match a regex into the same pool and calling the balancer set on the table + * to balance within the set of matching tablet servers. All tablet servers that do not match a regex are grouped into a default pool. + * Regex properties for this balancer are specified as: + * table.custom.balancer.host.regex.<tablename>=<regex> + * Periodically (default 5m) this balancer will check to see if a tablet server is hosting tablets that it should not be according to the regex configuration. + * If this occurs then the offending tablets will be reassigned. This would cover the case where the configuration is changed and the master is restarted while + * the tablet servers are up. To change the out of bounds check time period, set the following property: + * table.custom.balancer.host.regex.oob.period=5m + * Periodically (default 5m) this balancer will regroup the set of current tablet servers into pools based on regexes applied to the tserver host names. This + * would cover the case of tservers dying or coming online. To change the host pool check time period, set the following property: + * table.custom.balancer.host.regex.pool.check=5m + * Regex matching can be based on either the host name (default) or host ip address. To set this balancer to match the regular expressions to the tablet server + * IP address, then set the following property: + * table.custom.balancer.host.regex.is.ip=true + * + */ +public class HostRegexTableLoadBalancer extends TableLoadBalancer { + + private static final Logger LOG = LoggerFactory.getLogger(HostRegexTableLoadBalancer.class); + + public static final String HOST_BALANCER_PREFIX = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex."; + + public static final String HOST_BALANCER_OOB_CHECK = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex.oob.period"; + + private static final String HOST_BALANCER_OOB_DEFAULT = "5m"; + +	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4173
16331	ACCUMULO-4173	server/base/src/main/java/org/apache/accumulo/server/master/balancer/HostRegexTableLoadBalancer.java --- @@ -0,0 +1,334 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.master.balancer; + +import java.net.InetAddress; +import java.net.UnknownHostException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; +import java.util.Set; +import java.util.SortedMap; +import java.util.TreeMap; +import java.util.regex.Pattern; + +import org.apache.accumulo.core.client.admin.TableOperations; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.data.impl.KeyExtent; +import org.apache.accumulo.core.master.thrift.TabletServerStatus; +import org.apache.accumulo.core.tabletserver.thrift.TabletStats; +import org.apache.accumulo.server.conf.ServerConfiguration; +import org.apache.accumulo.server.master.state.TServerInstance; +import org.apache.accumulo.server.master.state.TabletMigration; +import org.apache.thrift.TException; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + ++/** + * This balancer will create pools of tablet servers by grouping tablet servers that match a regex into the same pool and calling the balancer set on the table + * to balance within the set of matching tablet servers. All tablet servers that do not match a regex are grouped into a default pool. + * Regex properties for this balancer are specified as: + * table.custom.balancer.host.regex.<tablename>=<regex> + * Periodically (default 5m) this balancer will check to see if a tablet server is hosting tablets that it should not be according to the regex configuration. + * If this occurs then the offending tablets will be reassigned. This would cover the case where the configuration is changed and the master is restarted while + * the tablet servers are up. To change the out of bounds check time period, set the following property: + * table.custom.balancer.host.regex.oob.period=5m + * Periodically (default 5m) this balancer will regroup the set of current tablet servers into pools based on regexes applied to the tserver host names. This + * would cover the case of tservers dying or coming online. To change the host pool check time period, set the following property: + * table.custom.balancer.host.regex.pool.check=5m + * Regex matching can be based on either the host name (default) or host ip address. To set this balancer to match the regular expressions to the tablet server + * IP address, then set the following property: + * table.custom.balancer.host.regex.is.ip=true + * + */ +public class HostRegexTableLoadBalancer extends TableLoadBalancer { + + private static final Logger LOG = LoggerFactory.getLogger(HostRegexTableLoadBalancer.class); + + public static final String HOST_BALANCER_PREFIX = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex."; + + public static final String HOST_BALANCER_OOB_CHECK = Property.TABLE_ARBITRARY_PROP_PREFIX.getKey() + "balancer.host.regex.oob.period"; + + private static final String HOST_BALANCER_OOB_DEFAULT = "5m"; + +	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4173
16332	ACCUMULO-4171	Do we know how stable htrace's API is going to be as it moves forward from here on? Does htrace have an API statement, like semver? My concern is that now that Hadoop and Accumulo are both using HTrace, updating HTrace might further constrain the Hadoop versions we support (same with not updating it). If HTrace's API is not stable (or if Hadoop or ourselves use items which are not in HTrace's declared stable API), we may be limited to using only the exact version that Hadoop is using.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4171
16333	ACCUMULO-4171	Ah, that kind of makes sense. Do you recall why it seemed impossible to use? The only issue I encountered was that HDFS generates way too many spans, but we dealt with that in the Accumulo SpanReceiver by dropping the 0ms spans.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4171
16334	ACCUMULO-4171	Am I to read this that no Accumulo versions will work with Hadoop 2.6 and Hadoop 2.7?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4171

#	Requirement	Questions	Answered	Successfully Answered	URL
16335	ACCUMULO-4171	are we still interested in this for 1.9? (or some other version)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4171
16336	ACCUMULO-4171	What would alleviate those concerns though? e.g. I haven't seen anything in the HTrace community to suggest the issue I mentioned around upgrades has changed. Hadoop has a stated policy of not promising to avoid changes in dependencies across minor versions (or really any versions: [ref hadoop compat policy|http://hadoop.apache.org/docs/r2.7.2/hadoop-project-dist/hadoop-common/Compatibility.html#Java_Classpath]). We're already out of sync for tracing through to Hadoop for anything other than Hadoop 2.7.z. Do we need to go the shim route to try to handle tracing through to multiple versions? Or do we just bump our version along with Hadoop updates?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4171
16337	ACCUMULO-4165	It looks that the API you're envisioning will require some understanding of how locality groups are stored in RFiles. Have you considered omitting locality group support entirely, or just using a single default locality group if none are started before the first key is appended? For the factories, I'd assume the minimum information to provide is the filename? If so, should it default to "file://" if it begins with a "/"? For the parameters which take sizes, it'd be useful to be able to specify a string format, like "20M" instead of 1024*1024*20 bytes. Will this API be something that could be used internally to clean up some of our code which uses RFiles? (I hope so.)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4165
16338	ACCUMULO-4165	bq. For the factories, I'd assume the minimum information to provide is the filename? If so, should it default to "file://" if it begins with a "/"? If the default FileSystem is against the local filesystem, it seems reasonable to just always create a Path using that FileSystem (deferring to FileSystem to "localize" it, or tell us if it already has the wrong scheme?). bq. For the parameters which take sizes, it'd be useful to be able to specify a string format, like "20M" instead of 1024*1024*20 bytes. IMO, I think accepting a long representing bytes is fine, but :shrug:. I just think it's pretty easy for someone to make a constant in their code for certain numbers.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4165
16339	ACCUMULO-4165	bq. Have you considered omitting locality group support entirely, or just using a single default locality group if none are started before the first key is appended? Sorry I did not see these comments until now. Not going to drop support for locality groups. I think I will make it automatically start the default locality group. That should be a quick change.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4165
16340	ACCUMULO-4164	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/80#discussion_r56753645 --- Diff: core/src/main/java/org/apache/accumulo/core/file/blockfile/impl/SeekableByteArrayInputStream.java --- @@ -0,0 +1,132 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.file.blockfile.impl; +import java.io.IOException; +import java.io.InputStream; +/** + * This class is like byte array input stream with two differences. It supports seeking and avoids synchronization. + */ +public class SeekableByteArrayInputStream extends InputStream { + // make this volatile to ensure data set by one thread can be seen by another + private volatile byte buffer[]; + private int cur; + private int max; + @Override + public int read() { + if (cur < max) { + return buffer[cur++] & 0xff; + } else { + return -1; + } + } + @Override + public int read(byte b[], int offset, int length) { + if (b == null) { + throw new NullPointerException(); + } + if (length < 0 || offset < 0 || length > b.length - offset) { + throw new IndexOutOfBoundsException(); + } + if (length == 0) { + return 0; + } + int avail = max - cur; + if (avail <= 0) { + return -1; + } + if (length > avail) { + length = avail; + } + System.arraycopy(buffer, cur, b, offset, length); + cur += length; + return length; + } + @Override + public long skip(long requestedSkip) { + long actualSkip = max - cur; + if (requestedSkip < actualSkip) + if (requestedSkip < 0) + actualSkip = 0; + else + actualSkip = requestedSkip; + cur += actualSkip; + return actualSkip; + } + @Override + public int available() { + return max - cur; + } + @Override + public boolean markSupported() { + return false; + } + @Override + public void mark(int readAheadLimit) { + throw new UnsupportedOperationException(); + } + @Override + public void reset() { + throw new UnsupportedOperationException(); --- End diff -- Isn't this simple to implement? Is there a reason not to?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4164

#	Requirement	Questions	Answered	Successfully Answered	URL
16341	ACCUMULO-4164	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/80#discussion_r56753672 --- Diff: core/src/main/java/org/apache/accumulo/core/file/blockfile/impl/SeekableByteArrayInputStream.java --- @@ -0,0 +1,132 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.file.blockfile.impl; +import java.io.IOException; +import java.io.InputStream; +/** * This class is like byte array input stream with two differences. It supports seeking and avoids synchronization. */ public class SeekableByteArrayInputStream extends InputStream { // make this volatile to ensure data set by one thread can be seen by another private volatile byte buffer[]; private int cur; private int max; @Override public int read() { if (cur < max) { return buffer[cur++] & 0xff; } else { return -1; } } @Override public int read(byte b[], int offset, int length) { if (b == null) { throw new NullPointerException(); } if (length < 0 || offset < 0 || length > b.length - offset) { throw new IndexOutOfBoundsException(); } if (length == 0) { return 0; } int avail = max - cur; if (avail <= 0) { return -1; } if (length > avail) { length = avail; } System.arraycopy(buffer, cur, b, offset, length); cur += length; return length; } @Override public long skip(long requestedSkip) { long actualSkip = max - cur; if (requestedSkip < actualSkip) { actualSkip = requestedSkip; } cur += actualSkip; return actualSkip; } @Override public int available() { return max - cur; } @Override public boolean markSupported() { return false; } @Override public void mark(int readAheadLimit) { throw new UnsupportedOperationException(); } @Override public void reset() { throw new UnsupportedOperationException(); } @Override public void close() throws IOException {} public SeekableByteArrayInputStream(byte[] buf) { this.buffer = buf; this.cur = 0; this.max = buf.length; } public SeekableByteArrayInputStream(byte[] buf, int maxOffset) { this.buffer = buf; this.cur = 0; this.max = maxOffset; } public void seek(int position) { if (position < 0 || position >= max) throw new IllegalArgumentException("position = " + position + " maxOffset = " + max); this.cur = position; } public int getPosition() { return this.cur; } public byte[] getBuffer() { --- End diff -- Javadoc to be explicit that this is the actual array (changes to it will be represented in the InputStream? Do we want this do be `public`?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4164
16342	ACCUMULO-4164	I would love to see the perf test results for this change. Can you post them, [~kturner]?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4164
16343	ACCUMULO-4156	bq. I didn't see any special offsets for WAL files Sorry, I meant to say, definitively, that this doesn't exist. I think this was something I considered as a future improvement which would enable more responsive replication. bq. I think you'd need some marker for a WAL that lives through a flush so you don't do a double-insert incase of a failure after a flush. But a flush is just pushing the IMM to disk -- the records should already be recorded in the WAL by the time they make it into the IMM. Am I misunderstanding?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4156
16344	ACCUMULO-4147	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/74#issuecomment-188983450 @tristeng the proxy has fairly comprehensive ITs that excercise most of its API. There are located in `accumulo-test/src/main/java/org/apache/accumulo/test/proxy`. Would you be able to add test for these new namespace operations?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4147
16345	ACCUMULO-4147	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/74#issuecomment-189345114 > Any thoughts on where I should look? Looking at the test I have a strong suspicion that the reason the test is failing is because its using MockAccumulo. When things worked w/ your Python test client, was that proxy back a mini or real accumulo instance? I would discourage using Mock. Its been deperecated because it was not being well maintained as features (like namspaces) were added to Accumulo. I would also encourage you to add your namespace test to SimpleProxyBase. TBinaryProxyIT, TJsonProtocolProxyIT, etc all extend SimpleProxyBase but run configure a different protocol. All of the test in SimpleProxyBase are run using multiple thrift protocols. Also, the test in SimpleProxyBase use MiniAccumulo instead of MockAccumulo.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4147
16346	ACCUMULO-4147	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/74#discussion_r57719585 --- Diff: test/src/main/java/org/apache/accumulo/test/proxy/SimpleProxyBase.java --- @@ -836,6 +1020,22 @@ public void testExists() throws Exception { } @Test + public void testNamespaceExists() throws Exception { + try { + client.createNamespace(creds, namespaceName); --- End diff -- Also, will this test pass if not exception is thrown?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4147
16347	ACCUMULO-4147	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/74#issuecomment-205440868 Thanks @tristeng. I squashed and merge this, can you close this PR? I ran the following test w/ no problem before pushing. That caused some reformatting which was included in the commit I pushed. I also added `since 1.8.0` to the comment in the proxy.thrift file. ``` mvn clean verify -DskipITs mvn -Dit.test=".*proxy.*" -Dtest=foo -Dfindbugs.skip -Dcheckstyle.skip -DfailIfNoTests=false verify ```	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4147
16348	ACCUMULO-4132	It seems like your proposed change would be source compatible. However I am uncertain if its binary compatible. Will code compiled against {{TableOperations::addSplits(SortedSet<Text>)}} work with {{TableOperations::addSplits(Collection<Text>)}}? I suspect its not binary compat, but not positive. We can only make the changes if its binary compat. If not binary compat, another possible option is to add a method {{TableOperations::addSplits(Collection<Text>)}} and not remove {{TableOperations::addSplits(SortedSet<Text>)}}. Should be able to have both.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4132
16349	ACCUMULO-4111	Thanks for the patch. I will add you as a contributor in Jira and assign this issue to you. Two comment regarding the patch : * The class needs a {{@since 1.8.0}} javadoc tag * Would you be able to modify the patch to extend {{AbstractLexicoder}} and have a {{ protected Float decodeUnchecked(byte[] data, int offset, int len)}}} method? This is to follow the precedent set in ACCUMULO-3789.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4111
16350	ACCUMULO-4111	Thanks [~kturner] Was this the best way to deal with this process? If there is another way that is easier for you guys I'm happy to follow it next time. Could you put me down with the Satellite Applications Catapult (https://sa.catapult.org.uk/) and Europe/London or 0/+1 for the time zone.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4111

#	Requirement	Questions	Answered	Successfully Answered	URL
16351	ACCUMULO-4111	bq. Was this the best way to deal with this process? If there is another way that is easier for you guys I'm happy to follow it next time. Only thing I noticed was that you should feel free to leave "old" patches attached to the JIRA issue. The numbering scheme you were doing is recommended and can be helpful to go back and review old changes (in the face of heated discussion). You're also welcome to use pull requests on Github. We don't care too much either way. Thanks, again, for your contribution :)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4111
16352	ACCUMULO-4109	Is this something you'd like to do for 1.7.1, [~ctubbsii]? If so, any cycles to do it this week?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4109
16353	ACCUMULO-4109	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/64#issuecomment-173720610 Very nice stuff!! Left two comments that I thought of looking at the code. Have you run this through shellcheck? That would be a nice bar to set for ourselves.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4109
16354	ACCUMULO-4107	Are you thinking Hadoop metrics2-level reporting or something in the monitor (akin to the normal ingest rate graph)?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4107
16355	ACCUMULO-4104	What are you doing to colorize? I use a bash function in {{~/.bashrc}}: {code} function mvn() { trap 'echo -en "\e[m"' SIGHUP SIGINT SIGTERM command mvn "$@" 2>&1 | awk ' BEGIN { IGNORECASE=1 } /BUILD SUCCESS/ { print "\033[1;36m" $0 "\033[m" ; next } /^\[FATAL\]/ { print "\033[1;31m" $0 "\033[m" ; next } /^\[ERROR\]/ { print "\033[1;31m" $0 "\033[m" ; next } /^\[TRACE\]/ { print "\033[2;33m" $0 "\033[m" ; next } /^\[WARNING\]/ { print "\033[1;33m" $0 "\033[m" ; next } /^\[DEBUG\]/ { print "\033[2;34m" $0 "\033[m" ; next } /^\[INFO\]/ { print "\033[0;37m" $0 "\033[m" ; next } /^\[ALL\]/ { print "\033[1;37m" $0 "\033[m" ; next } { print } ' return ${PIPESTATUS[0]} ${pipestatus[1]} } {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4104
16356	ACCUMULO-4087	Using the "sticky" (previously chosen volume for that tablet) could work. I'm not a big fan of the sticky volume behavior. Another option, instead of using the previously chosen one, is to pull the table id from the {{KeyExtent}}, and re-choose with the appropriate chooser for that table (assuming the per-table chooser is configured as the general chooser). That might save some work examining the tablet's metadata, but at the risk of choosing another location (which is probably fine anyway). Of course, that only works if the {{FileUtils}} code is only being used for tablet- or table-scoped temp files, and not also used for general-purpose temp files (do we even have any such files? intermediate write-ahead logs? if we do, do they use this code?).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4087
16357	ACCUMULO-4086	Are we ready to merge this one in?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4086
16358	ACCUMULO-4086	I am needing this change for deployment onto our test system. Can we attempt to get this one wrapped up? I see there are now conflicts that have crept in. I will attempt to resolve those with another pull request into your branch.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4086
16359	ACCUMULO-4079	See ACCUMULO-2069. Almost two years to the day :-). I don't think we did anything, but I will double check when I get back in to the office to see if we changed something internally.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4079
16360	ACCUMULO-4078	Hi [~Matt Peterson]. Thanks for the feedback. I have some thoughts/remarks. bq. A new volume is added, with limited capacity, to be dedicated for a specific table. All other tables or files should be excluded from this new volume. How does the [PerTableVolumeChooser|https://github.com/apache/accumulo/blob/master/server/base/src/main/java/org/apache/accumulo/server/fs/PerTableVolumeChooser.java] fall short for what you'd like to do? I could see a static VolumeChooser implementation configured for the table in question that always chooses a specific volume. Every other table would use the default VolumeChooser (an implementation which ignores that special volume?). Maybe PreferredVolumeChooser already gives you the ability to limit the volumes you want to use (in a inclusive way, instead of exclusive). bq. It's not currently possible to specify preferred volumes for the write-ahead logs What's your intended effect from such a change? Remember that WALs are per-server, not per-table. bq. the failsafe chooser is the RandomVolumeChooser which will include the instance volume that needs to be excluded. It would be useful to have a configurable failsafe in this situation. By "failsafe" you mean a default chooser? I can see the value in mimicking how we do site-wide and per-table configuration values. When a per-table implementation doesn't exist (or is not relevant), it defers to the site-wide chooser. The approach we have now makes the provided implementations a bit more brittle since you inherit some default functionality (which I think is what you meant). bq. The volume chooser is called in FileUtils for temp directory creation but it could instead use the sticky dir to create the temp directory, not needing the volume chooser at all. I don't understand what you mean by "use the sticky dir". Why would we not want to use the volume chooser implementation when choosing a temp directory? This seems like an inconsistency since we do this for all other file-based operations. Also, are you planning to work on these changes? If so, let's get this ticket assigned to you :)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4078

#	Requirement	Questions	Answered	Successfully Answered	URL
16361	ACCUMULO-4078	Josh, Thanks for reviewing this ticket. {quote} What's your intended effect from such a change? Remember that WALs are per-server, not per-table. {quote} I'd like to exclude the WAL from the special volume for all tables. Let's assume that the special instance is appropriately sized for only that one table. The PerTableVolumeChooser will default to the random chooser for WAL files. The random chooser which will include all volumes, including the one that I want to exclude. {quote} I don't understand what you mean by "use the sticky dir". Why would we not want to use the volume chooser implementation when choosing a temp directory? This seems like an inconsistency since we do this for all other file-based operations. {quote} Sorry I realize now that the language was unclear. By "sticky dir" I was referring to the srv:dir entry for that tablet. The thought was that there is no need to choose a volume for the temp dir in the same way that Accumulo doesn't choose a volume each time it writes a new file for a tablet. It chooses a volume when first creating the volume and then continues to send its files there. {quote} How does the PerTableVolumeChooser fall short for what you'd like to do? I could see a static VolumeChooser implementation configured for the table in question that always chooses a specific volume. Every other table would use the default VolumeChooser (an implementation which ignores that special volume?). Maybe PreferredVolumeChooser already gives you the ability to limit the volumes you want to use (in a inclusive way, instead of exclusive). By "failsafe" you mean a default chooser? I can see the value in mimicking how we do site-wide and per-table configuration values. When a per-table implementation doesn't exist (or is not relevant), it defers to the site-wide chooser. The approach we have now makes the provided implementations a bit more brittle since you inherit some default functionality (which I think is what you meant). {quote} Yes, by "failsafe" I meant a default chooser, to be used when none is configured. Additionally, it would be used when a chooser is configured but cannot be used due to an error in its instantiation or configuration. There are cases with both the PreferredVolumeChooser and the PerTableVolumeChooser in which each will default to the RandomVolumeChooser, even if a site-wide default configuration is used. If the PerTableVolumeChooser is used as the general volume chooser and PreferredVolumeChooser is used as the site-wide table volume chooser, then this use case is nearly satisfied. The special table would be configured to include only the special volume. Most tables would inherit the site-wide configuration, which excludes the special volume. With this configuration, which I believe is the one you've suggested, there are a few cases when the special volume will be used: 1) .choose is called without a table id, resulting in files going to this special volume (e.g. WAL), (PreferredVolumeChooser and PerTableVolumeChooser behavior) 2) a table overrides the site-wide configuration but that override is misconfigured and fails, causing files to go to the RandomVolumeChooser (PerTableVolumeChooser behavior) PreferredVolumeChooser and PerTableVolumeChooser could use a configurable default, other than the RandomVolumeChooser. So there would be a general.volume.chooser and a general.volume.chooser.fallback.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4078
16362	ACCUMULO-4078	One last round of responses here on the parent :). Let's move discussion into the appropriate sub-tasks since Christopher took the time to make them {quote} bq. What's your intended effect from such a change? Remember that WALs are per-server, not per-table. I'd like to exclude the WAL from the special volume for all tables. Let's assume that the special instance is appropriately sized for only that one table. The PerTableVolumeChooser will default to the random chooser for WAL files. The random chooser which will include all volumes, including the one that I want to exclude. {quote} This would be an extreme overhaul of the WAL system. While it's easy to choose a different volume for a WAL, the point is that *all* tables hosted on a server share the same WAL file. There is no means to have a separate WAL for just one file. It's not impossible -- I just want to make sure you understand that this would be a serious undertaking. bq. Sorry I realize now that the language was unclear. By "sticky dir" I was referring to the srv:dir entry for that tablet. The thought was that there is no need to choose a volume for the temp dir in the same way that Accumulo doesn't choose a volume each time it writes a new file for a tablet. It chooses a volume when first creating the volume and then continues to send its files there. Gotcha, I understand where you were coming from now. Maybe it makes sense to have a srv:temp for a tablet? What's the net benefit of keeping a static temporary directory for a tablet? bq. If the PerTableVolumeChooser is used as the general volume chooser and PreferredVolumeChooser is used as the site-wide table volume chooser, then this use case is nearly satisfied. The special table would be configured to include only the special volume. Cool, I think this would be a nice addition. bq. PreferredVolumeChooser and PerTableVolumeChooser could use a configurable default, other than the RandomVolumeChooser. So there would be a general.volume.chooser and a general.volume.chooser.fallback. My first reaction is that it seems unnecessary to provide multiple "general*" volume chooser properties, but I need to read through your situations some more :). I think it would be best if we can continue to follow the global/site configuration and per-table/namespace overrides for VolumeChooser properties.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4078
16363	ACCUMULO-4078	bq. What's the net benefit of keeping a static temporary directory for a tablet? I'm not convinced there's even a benefit to having a static non-temporary directory for a tablet (srv:dir). Personally, I'd prefer choosing on every file creation (with as much information about the scope/environment as possible, so it can make a good decision based on current configuration vs. a past decision)... ACCUMULO-3090 bq. My first reaction is that it seems unnecessary to provide multiple "general*" volume chooser properties I agree. Some of the situations might be considered distinct "scopes" appropriate for a special config. However, many others might be behavior dependent on a particular configured global or per-table chooser, and it might be more appropriate to provide ACCUMULO-3384 to empower those various implementations.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4078
16364	ACCUMULO-4066	This one slipped in under my radar. Thought I'd give your changes a glance. 3x speed up is awesome! {code} - for (Condition cond : cm.getConditions()) { + // sort conditions inorder to get better lookup performance. Sort on client side so tserver does not have to do it. + Condition[] ca = cm.getConditions().toArray(new Condition[cm.getConditions().size()]); + Arrays.sort(ca, CONDITION_COMPARATOR); {code} To confirm, the server doesnt' rely on the sorted order, just hopes for it for performance reasons? I see a lot of changes in IteratorUtil (I assume to your point about loading iterators from the table config). How did this used to work? You had lots of new tests added for the other cases -- do we have good coverage for IteratorUtil already?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4066
16365	ACCUMULO-4066	bq. To confirm, the server doesnt' rely on the sorted order, just hopes for it for performance reasons? correct. bq. I see a lot of changes in IteratorUtil (I assume to your point about loading iterators from the table config). How did this used to work I needed parse table iter config once, cache that and then later merge condition iterators. I changed IteratorUtil to support this use case. Used to it would parse the table iterator config for each condition. I also modified the code to support caching class name to classes (Instead of going to the VFS classloader for each condition to load a class). Adding some test for IteratorUtil is a good idea.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4066
16366	ACCUMULO-4062	Is there a reason not to do this? What am I missing?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4062
16367	ACCUMULO-4062	{code}org.apache.accumulo.core.data.Mutation#hashCode{code} in master looks like this for me: {code} @Override public int hashCode() { return serializedSnapshot().hashCode(); } {code} what code are you looking at?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4062

#	Requirement	Questions	Answered	Successfully Answered	URL
16368	ACCUMULO-4062	Wouldn't this cause writes to arrive within a single tablet potentially in a different order then the user added them?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4062
16369	ACCUMULO-4062	Wouldn't computing the hash be kind of expensive (in comparison to List#add)? Couldn't the client just keep better track of the mutations they are generating?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4062
16370	ACCUMULO-4051	Is there a medium-term plan for syncing Accumulo and Thrift? The latter has gotten a lot of updates (security, compiler and library layers, etc) since 0.9.1, and needs a lot more, but it's hard to commit to making that happen on the Thrift end if we're pinned at an old version with no way forward.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4051
16371	ACCUMULO-4051	[~tlippy01], the problem for us is that it typically requires testing with a very close eye after upgrading Thrift. Historically, we've been bitten many, many times after upgrading the version of Thrift with performance regressions and more. As such, the worry from Accumulo's end is "will we have more issues?", but this isn't a blocker from doing it. We just need someone who can sit down, upgrade to 0.9.3, and run some of our tests over multiple nodes to check things out. In other words, I don't think we need multiple layers of planning yet; we just need someone to investigate.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4051
16372	ACCUMULO-4050	It's not clear to me what a "configuration container" means. Can one of you describe this in more detail, maybe include an example that would outline how the existing set of properties in ZK would appear with this suggested design? Similarly, what defines a snapshot WRT the data stored in ZK? Using ZK transaction id, maybe ZNode Version? https://zookeeper.apache.org/doc/r3.1.2/zookeeperProgrammers.html#sc_timeInZk. Any thought on how these snapshots could be used to enable follow-on work which enables us to build a client API that can guarantee that Accumulo servers have seen some configuration update?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4050
16373	ACCUMULO-4050	bq. I mean an object that is composed of the configuration (set of key value pairs) deserialized from ZK, as well as functions to perform the serialization/deserialization, and which can atomically swap out the reference to the complete configuration. This doesn't really help clarify the things that are vague to me :). Let me try to guess instead; maybe that will help describe what is not clear to me. A table would have multiple containers: one for the dynamic system-wide properties, one for the dynamic namespace-wide properties, and one for the dynamic table-specific properties? And these containers could be merged together (as per the currrent precedence of table > namespace > system)? Or, would the container be abstracting this hierarchy away? bq. By "snapshot", I mean an isolated immutable view of the entire set of configuration (k-v pairs) at a single point in time bq. I suppose we could incorporate the ZNode Version into the snapshot object or its containers somehow I think allowing a snapshot to have know what "time" it was created at would be all that we'd really need. Even if it's just a no-op now (always 0) or just {{System.currentTimeMillis()}}, we could easily swap in the necessary info from ZK down the road to piece together the logic to determine when every server has seen (at least) some update. bq. Any of that would definitely be follow-on work. Yes, completely. My only concern is that if you/Michael step up to do this work, let's look a few feet in front of us to make sure we at least don't make solving the client/server config disconnect harder :) What you have outlined so far sounds like a nice building block to fixing this stupid problem once and for all.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4050
16374	ACCUMULO-4049	How about: 1. Generalizing {{FirstKeyInRowIterator}} to a new class {{FirstPartialKeyInRowIterator}} (or perhaps {{StrippingIterator}}). 2. Modifying {{CountingIterator}} to act as a user iterator, being careful to return entries within the range to which it is seeked. It still just counts entries from its parent iterator. 3. Creating a user convenience class {{CountPartialKeysIterator}}. Internally it creates instances of the above two iterators; no new functionality; all this does is save the user from needing to add two iterator options. It accepts an option for the partial key: null (count all = no need for iterator #1), row (use both iterators), column family (use both iterators), column qualifier (use both iterators), column visibility (use both iterators), timestamp (same as count all), delete key (same as count all; we don't normally see deletes during scans). Client code would look like {code} BatchScanner bs; bs.setRanges(some_ranges); IteratorSetting is = new IteratorSetting(25, "my_row_counter", CountPartialKeysIterator.class); CountPartialKeysIterator.setPartialKey(PartialKey.ROW); bs.addScanIterator(is); long cnt = 0l; try { for (Map.Entry<Key, Value> entry : bs) { cnt += Long.parseLong(new String(entry.getValue().get())); } } finally { bs.close(); } return cnt; {code} I could look at this during the weekend if the plan looks good. This iterator should not be used at compaction time because it is not idempotent. The convenience class {{CountPartialKeysIterator}} is not a perfect solution. It would be better if we created a generic iterator that could load multiple iterators at runtime inside {{init}}, but this is out of this ticket's scope. We can set the default behavior of {{CountPartialKeysIterator}} to count all entries if no iterator option is set. I thought of naming iterator #3 {{CountPartsIterator}}, but this name could lead users to believe that it can count the number of unique column families. The iterator can only count the number of row+family partial keys; counting the number of unique column families requires more complexity and is less suitable for Accumulo's layout.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4049
16375	ACCUMULO-4045	bq. sessions would scale proportional to the number of Connector objects I had initially assumed that you meant a Connector would uniquely own a Session, but I'm not quite sure after reading this :) How do you envision the Session on the client-side mapping to the API? (e.g. would a single session be capable of "backing' multiple Connector objects?).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4045
16376	ACCUMULO-4045	How do we prevent unauthenticated connectors from piggy-backing on an existing session?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4045
16377	ACCUMULO-4045	bq. I had initially assumed that you meant a Connector would uniquely own a Session, but I'm not quite sure after reading this Right now, creating a Connector is where we store the user's credentials for future operations. We could create a session which could be used across connectors, but I was mostly thinking about an implementation which would have minimal impact on API and user experience. So, that'd basically mean one session per Connector. bq. How do you envision the Session on the client-side mapping to the API? (e.g. would a single session be capable of "backing' multiple Connector objects?). Essentially, a SessionToken would replace the PasswordToken in the ClientContext's Credentials.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4045
16378	ACCUMULO-4045	bq. We can also enforce a same-origin policy, to ensure that the session token is valid only for the client IP address which initiated the request for its construction. This provides additional mitigation against session hijacking. That's a nice perk I hadn't thought about on my own. bq. in future, I wouldn't mind forcing SASL for all connections This is a nice thought. Does the SASLTransport from thrift work with tframedtransport? (maybe that was the ssltransport which didn't work with tframedtransport) I don't remember anymore :)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-4045
16379	ACCUMULO-4040	Do you know what version they were running? And what the error was (i.e. what is the type of {{tablets}} on an IBM JVM)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4040
16380	ACCUMULO-4040	bq. IBM J9 JVM Do you know the specific version? https://www-01.ibm.com/support/knowledgecenter/SSYKE2_7.0.0/com.ibm.java.lnx.71.doc/user/java_jvm.html seems to show multiple targeted JREs via different "implementations" of the J9 jvm.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4040

#	Requirement	Questions	Answered	Successfully Answered	URL
16381	ACCUMULO-4039	Could you be a little more specific? The current RPC model is Half-Sync, Half-Async (HSHA). So requests and responses are handled asynchronously, but the RPC itself is executed by a shared thread-pool. This is a convenient paradigm to program against. Some things are tricky (like the service blocking on itself for meta services), but most requests are implemented as straight-line code: no callbacks, Futures and other bookkeeping. What limits are you pushing up against, and how would a different design solve it?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4039
16382	ACCUMULO-4022	I'm sorry, I was thinking of the NativeMap. My apologies. I am a newbie. I like these ideas! Thanks! One question, though. Would you envision the sync command being driven by the client or the read/write server?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4022
16383	ACCUMULO-4022	* Need to consider how to load balance scans between the read-write and read-only tablets. * Can there be more than one read-only copy of a tablet?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4022
16384	ACCUMULO-4014	Aside: are we considering killing accumulo-start and getting rid of this nonsense completely for 1.8?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-4014
16385	ACCUMULO-4004	bq. Some folks I'm working with have been wanting this to make replication more predictive. Definitely a nice bonus which gets around a quick of the replication impl. bq. It should be possible to manually roll WALs so that files on decommissioning datanodes are closed and the decommissioning process can complete [~ecn], do you have an explanation for why we even need to do this? I'd like to better understand why we need to change our code.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4004
16386	ACCUMULO-4004	[~dlmarion], are there any relevant sections in the book that should be updated after this change? Have you given this any thought -- it sounds like it would be important/useful for ops people to know about.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-4004
16387	ACCUMULO-3993	To confirm, this would just need to include bloomfilter in core and thrift for the server socket in server (as they both contain javadoc owned by the original copyright holder). The javascript stuff doesn't need anything done because they wouldn't contain anything in the javadoc jars. Anything I missed?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3993
16388	ACCUMULO-3993	Actually, this doesn't seem to be the case on hbase-common. Was that an HBase omission, or is there a reason that copied javadoc over there didn't require an updated NOTICE? bq. jars contain NOTICE file that incorrectly states the aggregate work is copyright the Accumulo project instead of the ASF I do see that the javadoc jars for rc1 (at least core) say "Copyright 2015 Apache Accumulo Project". So the above does apply as well.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3993
16389	ACCUMULO-3993	the issue back in HBase was that there were bundled works in the source/bin jar that didn't impact the javadocs, so the javadocs needed different files. Do all of our included works show up in the javadocs? i.e. I wouldn't expect the jquery / flot stuff to be present.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3993
16390	ACCUMULO-3993	bq. because it references jquery and that isn't present Wait, what? It's not? I thought that would've been brought in via src/main/resources. https://github.com/apache/accumulo/tree/1.5/server/src/main/resources/web/flot Are you sure you did your find correctly?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3993
16391	ACCUMULO-3991	bq. {{+server/src/main/java/org/apache/accumulo/server/util/TNonblockingServerSocket.java}} This file moves around across the branches. I'd prefer to see us just refer to the filename which should be unique. Would be nice to do that for the bloomfilters package, but a bit more difficult. Maybe enumerate those files instead of listing the package? What do you think? Otherwise, the patch lgtm.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3991
16392	ACCUMULO-3991	{quote} We should probably put the generic text of the MIT and BSD licenses in the LICENSE file, so we can just include them once for all bundled software to which that license applies, and move the copyright notices to the NOTICE file. I believe that's the correct way of utilizing these files. {quote} The current patch is done according to the [guidance on properly assembling license and notice files|http://www.apache.org/dev/licensing-howto.html]: permissive licensed works only in LICENSE, ASLv2 licensed works w/NOTICE file in NOTICE. It also only contains a single copy of each of the MIT and BSD-3 licenses referenced. {quote} Hey Sean, the files mentioned in ACCUMULO-3989 are based on old HBase versions of the same (thus the copyright notice in their headers). See http://s.apache.org/JXi for reference; that was a while ago, but some of it is still relevant. {quote} Ah. Should we include a "Portions from Apache HBase" notice in the NOTICE file? HBase didn't have a NOTICE file that included a copyright statement until very recently, so it's almost certain that the version these were copied from didn't include it; I think we'd be fine to just use as-is. Though perhaps a note about the history in the source file would be nice? {quote} bq. +server/src/main/java/org/apache/accumulo/server/util/TNonblockingServerSocket.java This file moves around across the branches. I'd prefer to see us just refer to the filename which should be unique. Would be nice to do that for the bloomfilters package, but a bit more difficult. Maybe enumerate those files instead of listing the package? What do you think? Otherwise, the patch lgtm. {quote} both excellent ideas, let me fix that up.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3991
16393	ACCUMULO-3990	bq. This might involve breaking out what we have in our top-level LICENSE/NOTICE files into those applicable for the specific jars. These would also be different from the javadoc jars (though they should be the same for the source jars) because the javadocs are probably all AL2.0 with no special notices. I thought this was clearly something that needs to be done for the Thrift and Javascript files based on ML discussion. Did I misunderstand?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3990
16394	ACCUMULO-3990	Applied [~busbey]'s addendum. Did the merge work up through 1.6, accounting for the break-out or server into server/*. Didn't verify jar contents in newer branches, but the necessary files should be in the right places now.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3990
16395	ACCUMULO-3982	Some points discussed with [~phrocker] off line: # the tablet servers used to clear the Locator cache after finishing a bulk import request. By switching to once-per-hour in ACCUMULO-3549, this is no longer true. # at larger scales, even once per hour is aggressive. Eventually you are re-reading the metadata table every few seconds, for no reason # the Locator needs more testing: for example, it always uses a single write lock to serialize lookups, which is a bad behavior for random tablet lookups. Locator does some simple read-ahead. Is it enough? Is it too much? # a configurable size-limited LRU cache makes much more sense if the concern is memory usage # TreeMap isn't exactly the most space-efficient store: [~kturner] thinks a 10-byte location entry requires 256 bytes. However, caching a million tablets would only use 0.25G which is not that much if you have a million-tablet Accumulo instance # the main issue [~phrocker] was investigating is ACCUMULO-3835, which has already been fixed in 1.8	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3982
16396	ACCUMULO-3982	bq. he Locator needs more testing: for example, it always uses a single write lock to serialize lookups, which is a bad behavior for random tablet lookups. Locator does some simple read-ahead. Is it enough? Is it too much? We've had a number of requests to get a Locator impl into the public API. Perhaps now would be a good time to refactor it into the API and make sure we can adequately test it? Would be happy to try to help get behind this.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3982

#	Requirement	Questions	Answered	Successfully Answered	URL
16397	ACCUMULO-3982	bq. It would be really nice if we could leverage guava's Caches here [Looks like|http://stackoverflow.com/questions/21966778/is-it-possible-to-iterate-over-a-guava-cache-in-order-of-insertion-access] Guava doesn't ever intend to support any sort of ordering on Cache. Maybe things have changed since that comment though. Generally speaking, what are the requirements/desirables for the locator? * Sorted iteration * Quick lookups * Quick updates * Constant memory WRT number of tablets * Avoid extra object creation/copy (except on cache updates) What else?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3982
16398	ACCUMULO-3982	bq. A quick investigation shows that doubles the data structure size overhead Makes me a little nervous -- doesn't the size of the data structures here affect both servers and clients? Granted, it's probably not a lot of memory total, but still makes me wonder. At first glance, something like Java's ConcurrentSkipListMap makes sense to me (and wouldn't require a lot of effort), but I haven't done enough reading yet to be sure there isn't something better. I wonder how the common access patterns the locator impl uses would affect the overall concurrency of the structure. For example, in the server-failure case, we would invalidate the locations of all tablets for a server which would likely be spread across the entire structure which would probably slow down as concurrency increases. Would being able to store the locations for a tablet close to the other tablets hosted by the same server be noticeably beneficial? bq. Adding and removing a million objects to both only takes ~5s. What did that look like size-wise? ~260MB like you said? Or is that 2*260MB maintaining a cache and a treemap?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3982
16399	ACCUMULO-3982	bq. It took about 0.5G. I also did experiments with a ConcurrentSkipListMap. The TreeMap was slightly faster than a SkipList. The SkipList used slightly less RAM (by a few percent). Cool, thanks for sharing. bq. it would be able to support different concurrent update strategies. I'm sure this would require more testing to see how concurrency pans out. bq. If we have both data structures, we can do our first look-ups on Cache, which will normally find the tablet, and that's a hash lookup, which is significantly faster. So that would be a nice benefit. Using two different data structures to maintain a single cache just comes across as a smell to me. I'm sorry, I know this is coming across as me harping on this approach. Because the locator is guaranteed to be used concurrently, this makes me worry even more about correctly keeping state in two different places over one. In a perfect world, the Locator impl is just holding a single copy of the data in the metadata table. While the Guava Cache interface is nice to use, is it going to be ultimately a more confusing implementation that have a single data structure to cache the metadata table?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3982
16400	ACCUMULO-3974	Thoughts on whether we should do this before we ship the next round of bug-fixes [~ctubbsii], [~ecn], [~kturner]? How critical does changing randomwalk to catch this strike you all?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3974
16401	ACCUMULO-3970	Thinking about the problem in a different way, might it be less intrusive to encode the rules for alternate views of data at ingest time and then push down the logic to deduplicate columns with multiple visibilities? I think this would prevent the need to inject down in the system visibility code and perform the merge/filter in "userspace iterators". For your DOB example, the table would actually contain both records: {noformat} (pt_id, demographic, pt_dob, PII_DOB) -> "1925-08-22" (pt_id, demographic, pt_dob, SHD_DOB) -> "1925" {noformat} At scan time, if the user has the ability to view this patient's personal data, they would scan with the PII_DOB visibility. Otherwise, they'd only have the SHD_DOB visibility. In the former case, you'd want to only present one value for demographic:pt_dob (not both) "1925-08-22". This lets you perform this filter on the server instead of unwinding it on the client. It does require logic in the driving application to assign the visibilities accordingly, but I think your original suggestion would also require this.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3970
16402	ACCUMULO-3970	Thanks for your comment, Billie. I see what you mean about not wanting to rely on the table admin to manage data visiblities. I guess my use case could be summarized as, "As a user of Accumulo, I want to be able to put sensitive data into the system, and have different users see different views of that data in accordance with organizational policy". Flipping it around, you're not relying on the data producer to get de-identification "right". In my mind it's a valuable feature, but I get that I be in a minority among Accumulo's users in wanting it.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3970
16403	ACCUMULO-3961	GitHub user rweeks opened a pull request: https://github.com/apache/accumulo/pull/42 Implementation of SeekingFilter, CfCqSliceFilter and CfCqSliceSeekingFilter OK let's try this again. Right bug id this time. Implementation for ACCUMULO-3961 including a filter which behaves similarly to the existing ColumnSliceFilter but can seek to the next potential match instead of explicitly rejecting every KV pair that falls outside the slice. Probably faster for small slices on large rows. Christopher switched the fix version to 1.8, I guess because adding new functionality should go into a minor release. So I'm targeting this against the master branch, because it doesn't look like a 1.8 branch exists yet. I know you guys require this to be rebased into a single commit. Is that something I can do on my side (I'm concerned about rebasing after I've pushed this to a public repo) or is that something you can do when you receive the pull request? You can merge this pull request into a Git repository by running: $ git pull https://github.com/phemisystems/accumulo ACCUMULO-3961 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/accumulo/pull/42.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #42 ---- commit a0c8d588109e2d8d0b3bdf862ec1d3e1e48c50e7 Author: Russ Weeks <rweeks@phemi.com> Date: 2015-07-23T03:53:01Z Initial commit of CfCqSlice code. Renamed files but Java class names are messed up. commit 9cc2c28251d445c86d7006bf1237ea2fd096b4e1 Author: Russ Weeks <rweeks@newbrightidea.com> Date: 2015-07-23T04:23:12Z Removing cyclic dependency on MiniAccumuloCluster by moving tests to test module commit 1100ba49ee09c97aadbe80429f54a13519a98552 Author: Russ Weeks <rweeks@newbrightidea.com> Date: 2015-07-23T05:48:12Z Reformatted and added comment to SeekingFilter commit e5ea222fa07ab8bb7211a12c87aa127146874f81 Author: Russ Weeks <rweeks@phemi.com> Date: 2015-07-24T00:20:08Z Slice tests passing commit f579acc4e60b7aa5cf36c59fad51ea36347204f0 Author: Russ Weeks <rweeks@phemi.com> Date: 2015-08-13T17:21:31Z Added class-level javadoc, fixed problem with OptionsDescriber commit dc2655622f71676012cf849a849e2e89fd936d26 Author: Russ Weeks <rweeks@phemi.com> Date: 2015-08-13T17:44:38Z tweaking the javadoc. ----	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3961

#	Requirement	Questions	Answered	Successfully Answered	URL
16404	ACCUMULO-3959	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/45#discussion_r37784571 --- Diff: core/src/main/java/org/apache/accumulo/core/client/BatchScanner.java --- @@ -16,19 +16,20 @@ */ package org.apache.accumulo.core.client; +import org.apache.accumulo.core.data.Range; + import java.util.Collection; import java.util.concurrent.TimeUnit; -import org.apache.accumulo.core.data.Range; - /** * Implementations of BatchScanner support efficient lookups of many ranges in accumulo. * BatchScanners are also appropriate for large, single ranges, as a BatchScanner will break those ranges up into separate RPCs provided the range spans more than one tablet and there are sufficiently many scan threads available. * Use this when looking up lots of ranges and you expect each range to contain a small amount of data. Also only use this when you do not care about the returned data being in sorted order. * If you want to lookup a few ranges and expect those ranges to contain a lot of data, then use the Scanner instead. Also, the Scanner will return data in sorted order, this will not. * Only use this when you do not care about returned data being in sorted order. --- End diff -- This was already broken before your patch, but I think javadoc need `<p>` markup for paragraphs. Not sure it will render as intended w/o it. Did you format these changes?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3959
16405	ACCUMULO-3959	Github user dhutchis commented on a diff in the pull request: https://github.com/apache/accumulo/pull/45#discussion_r37785744 --- Diff: core/src/main/java/org/apache/accumulo/core/client/BatchScanner.java --- @@ -16,19 +16,20 @@ */ package org.apache.accumulo.core.client; +import org.apache.accumulo.core.data.Range; + import java.util.Collection; import java.util.concurrent.TimeUnit; -import org.apache.accumulo.core.data.Range; - /** * Implementations of BatchScanner support efficient lookups of many ranges in accumulo. * BatchScanners are also appropriate for large, single ranges, as a BatchScanner will break those ranges up into separate RPCs provided the range spans more than one tablet and there are sufficiently many scan threads available. * Use this when looking up lots of ranges and you expect each range to contain a small amount of data. Also only use this when you do not care about the returned data being in sorted order. * If you want to lookup a few ranges and expect those ranges to contain a lot of data, then use the Scanner instead. Also, the Scanner will return data in sorted order, this will not. * Only use this when you do not care about returned data being in sorted order. --- End diff -- Correct, I see that the <p> tag is necessary from the online javadoc at http://accumulo.apache.org/1.7/apidocs/org/apache/accumulo/core/client/BatchScanner.html Will fix tonight when I return to my laptop. I don't think my editor (IntelliJ with the Eclipse code formatter plugin) adds the HTML tags automatically. On Mon, Aug 24, 2015 at 2:25 PM, Keith Turner <notifications@github.com> wrote: > In core/src/main/java/org/apache/accumulo/core/client/BatchScanner.java > <https://github.com/apache/accumulo/pull/45#discussion_r37784571>: > > * > > - * Use this when looking up lots of ranges and you expect each range to contain a small amount of data. Also only use this when you do not care about the > > - * returned data being in sorted order. > > - * > > - * If you want to lookup a few ranges and expect those ranges to contain a lot of data, then use the Scanner instead. Also, the Scanner will return data in > > - * sorted order, this will not. > > + * Only use this when you do not care about returned data being in sorted order. > > This was already broken before your patch, but I think javadoc need <p> > markup for paragraphs. Not sure it will render as intended w/o it. > > Did you format these changes? > > — > Reply to this email directly or view it on GitHub > <https://github.com/apache/accumulo/pull/45/files#r37784571>. >	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3959
16406	ACCUMULO-3952	Should there be an error if the script is unable to check the well-known directory for the jar in question, but it doesn't matter because all the necessary classes are in a jar in CLASSPATH already when the script is launched? How much of a problem is it that the jars aren't in the expected directory? Is it always a "fail loudly" situation, or only if CLASSPATH is empty when the script starts? How much checking can we actually do to ensure the classpath is properly set up before the JVM launches and it tries to load a class?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3952
16407	ACCUMULO-3952	bq. Should there be an error if the script is unable to check the well-known directory for the jar in question, but it doesn't matter because all the necessary classes are in a jar in CLASSPATH already when the script is launched? I didn't think we supported users setting CLASSPATH and then calling Accumulo scripts. If that does work, it's news to me and I would call it an undocumented/subject-to-change feature. bq. How much of a problem is it that the jars aren't in the expected directory? Is it always a "fail loudly" situation, or only if CLASSPATH is empty when the script starts? Very much of a problem? It's a fail very poorly situation if the log4j jar can't be found. I had to try very had to get a stacktrace. bq. How much checking can we actually do to ensure the classpath is properly set up before the JVM launches and it tries to load a class? I think it's very trivial to check the "start" classpath (slf4j and log4j) which is all I intended to check. Maybe accumulo-start too?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3952
16408	ACCUMULO-3952	bq. If that does work, it's news to me and I would call it an undocumented/subject-to-change feature. I fixed that with ACCUMULO-3685, because it was a long-standing issue, and I was correcting the bootstrap classpath issues. So, it has worked that way since at least 1.7.0 (but probably not before). I suppose it is an undocumented feature, but it's one I'd expect for a Java app. After all, I wouldn't exactly consider it normal behavior for a launch script to clobber-and-ignore standard user preferences/environment variables that affect the behavior of the thing it is launching, especially a "PATH" environment variable, where it is typical to append/prepend vs. clobber. bq. Very much of a problem? It's a fail very poorly situation if the log4j jar can't be found. I had to try very had to get a stacktrace. I agree it's definitely a problem. I suppose I was trying to ask how much of a problem that the script doesn't do the check vs. Java doing the check. Either way, the error needs to be more prominent. I'm just not sure to what extent we can preemptively check for required classes in the scripts, when the classes could be in different jars. I guess I was trying to ask whether we should consider these "required" locations, or "expected" locations, which would determine slightly different behavior and messages. bq. I think it's very trivial to check the "start" classpath (slf4j and log4j) which is all I intended to check. It's trivial to check for files with particular names, but it's less trivial to check for required classes being contained in those jars. So, this goes back to my question about "required" vs. "expected". Do we require that those specific classes be loaded from jars in those specific locations by those specific names? How much checking are you thinking is needed? At a certain point, Java's going to do a much more thorough checking than we could, for the requisite classes. Maybe it's just better to ensure that the CNFE is propagated better from Java, rather than implement our own checking? Or maybe we still want to do a minimal "sanity check" to see if the file with the expected name exists?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3952
16409	ACCUMULO-3952	I don't think it's wise to check for log4j in the environment. We've been slowly making our way toward pure slf4j-api consumption, and if log4j doesn't exist then the proper behaviour is "no logging" or possibly "print a warning to stdout/err" but never "fail quickly/loudly/somehow." What do you envision this verification doing, Josh?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3952

#	Requirement	Questions	Answered	Successfully Answered	URL
16410	ACCUMULO-3952	bq. I don't think it's wise to check for log4j in the environment. We've been slowly making our way toward pure slf4j-api consumption I don't understand the hesitation. The same goes for the slf4j jars. If they can't be found, we can't reasonably expect to run Accumulo at all. bq. What do you envision this verification doing, Josh? {{bin/accumulo}}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3952
16411	ACCUMULO-3952	Checking for slf4j is fine, since we advertise our dependency on it. I would like to believe that somebody can swap out our log4j jars and use logback or another implementation and Accumulo would still run mostly fine. Some functionality like monitor log forwarding might be missing, but it should still run. {quote} bq. What do you envision this verification doing, Josh? bin/accumulo {quote} What, not where. As in "if verify fails, then... stop process? print error? prompt for intervention?"	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3952
16412	ACCUMULO-3952	If it doesn't exist in the expected directory, but the classes can be loaded from the user-specified CLASSPATH, do you still want this error to occur? Do we really want it to be absolutely required to exist in this specific directory under all circumstances?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3952
16413	ACCUMULO-3952	Would CLASSPATH being set be a sufficient condition to disable the check? That's the only case I can imagine where the we can reasonably expect it to be okay that we didn't find the expected jar(s) in the expected location(s). I'm not sure there's much added value for providing another env variable to control this.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3952
16414	ACCUMULO-3948	Thanks for posting the WIP patches, Dave. Can you add them to reviewboard (or make a PR on github) when you get them ported to 1.8.0, please? It would be appreciated.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3948
16415	ACCUMULO-3948	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/51#discussion_r44564606 --- Diff: docs/src/main/asciidoc/chapters/administration.txt --- @@ -344,6 +344,49 @@ very sensitive to an excessive number of nodes and the sizes of the nodes. Appli which leverage the user of custom properties should take these warnings into consideration. There is no enforcement of these warnings via the API. +==== Configuring the ClassLoader + +Accumulo loads classes from the locations specified in the +general.classpaths+ property. Additionally, Accumulo will load classes +from the locations specified in the +general.dynamic.classpaths+ property and will monitor and reload them if they change. The reloading +feature is useful during the development and testing of iterators as new or modified iterator classes can be deployed to Accumulo without +having to restart the database. + +Accumulo also has an alternate configuration for the classloader which will allow it to load classes from remote locations. This mechanism +uses Apache Commons VFS which enables locations such as http and hdfs to be used. This alternate configuration also uses the ++general.classpaths+ property in the same manner described above. It differs in that you need to configure the ++general.vfs.classpaths+ property instead of the +general.dynamic.classpath+ property. As in the default configuration, this alternate +configuration will also monitor the vfs locations for changes and reload if necessary. + +===== ClassLoader Contexts + +With the addition of the VFS based classloader, we introduced the notion of classloader contexts. A context is identified +by a name and references a set of locations from which to load classes and can be specified in the accumulo-site.xml file or added +using the +config+ command in the shell. Below is an example for specify the app1 context in the accumulo-site.xml file: + +[source,xml] +<property> + <name>general.vfs.context.classpath.app1</name> + <value>hdfs://localhost:8020/applicationA/classpath/.*.jar,file:///opt/applicationA/lib/.*.jar</value> + <description>Application A classpath, loads jars from HDFS and local file system</description> +</property> + +The default behavior follows the Java ClassLoader contract in that classes, if they exists, are loaded from the parent classloader first. +You can override this behavior by delegating to the parent classloader after looking in this classloader first. An example of this +configuration is: + +[source,xml] +<property> + <name>general.vfs.context.classpath.app1.delegation=post</name> + <value>hdfs://localhost:8020/applicationA/classpath/.*.jar,file:///opt/applicationA/lib/.*.jar</value> + <description>Application A classpath, loads jars from HDFS and local file system</description> +</property> + +To use contexts in your application you can set the +table.classpath.context+ on your tables or use the +setContext()+ method on Scanner --- End diff -- What is the "context" that would be passed to `setContext(String)`?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3948
16416	ACCUMULO-3948	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/51#discussion_r44565049 --- Diff: server/tserver/src/main/java/org/apache/accumulo/tserver/tablet/ScanOptions.java --- @@ -99,4 +101,12 @@ public SamplerConfigurationImpl getSamplerConfigurationImpl() { public long getBatchTimeOut() { return batchTimeOut; } + + public String getContext() { + return context; + } + + public void setContext(String context) { --- End diff -- Is this actually used anywhere? Can we just make context final?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3948

#	Requirement	Questions	Answered	Successfully Answered	URL
16417	ACCUMULO-3948	test/src/main/java/org/apache/accumulo/test/functional/ScannerContextIT.java --- @@ -0,0 +1,227 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.test.functional; + +import static org.junit.Assert.assertEquals; +import static org.junit.Assert.assertFalse; +import static org.junit.Assert.assertTrue; +import static org.junit.Assert.fail; + +import java.util.Collections; +import java.util.Iterator; +import java.util.Map.Entry; + +import org.apache.accumulo.core.client.BatchScanner; +import org.apache.accumulo.core.client.BatchWriter; +import org.apache.accumulo.core.client.BatchWriterConfig; +import org.apache.accumulo.core.client.Connector; +import org.apache.accumulo.core.client.IteratorSetting; +import org.apache.accumulo.core.client.Scanner; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.data.Key; +import org.apache.accumulo.core.data.Mutation; +import org.apache.accumulo.core.data.Range; +import org.apache.accumulo.core.data.Value; +import org.apache.accumulo.core.security.Authorizations; +import org.apache.accumulo.core.util.CachedConfiguration; +import org.apache.accumulo.fate.util.UtilWaitThread; +import org.apache.accumulo.harness.AccumuloClusterHarness; +import org.apache.accumulo.minicluster.impl.MiniAccumuloClusterImpl; +import org.apache.hadoop.fs.FileSystem; +import org.apache.hadoop.fs.Path; +import org.hamcrest.CoreMatchers; +import org.junit.Assume; +import org.junit.Before; +import org.junit.Test; + +public class ScannerContextIT extends AccumuloClusterHarness { + private static final String CONTEXT = ScannerContextIT.class.getSimpleName(); + private static final String CONTEXT_PROPERTY = Property.VFS_CONTEXT_CLASSPATH_PROPERTY + CONTEXT; + private static final String CONTEXT_DIR = "file:///tmp"; + private static final String CONTEXT_CLASSPATH = CONTEXT_DIR + "/Test.jar"; + private static int ITERATIONS = 10; + + @Override + protected int defaultTimeoutSeconds() { + return 2 * 60; + } + + @Before + public void checkCluster() { + Assume.assumeThat(getClusterType(), CoreMatchers.is(ClusterType.MINI)); + MiniAccumuloClusterImpl mac = (MiniAccumuloClusterImpl) getCluster(); + } + + @Test + public void test() throws Exception { + // Copy the TestIterators jar to tmp + FileSystem fs = FileSystem.get(CachedConfiguration.getInstance()); + Path srcPath = new Path(System.getProperty("user.dir") + "/target/TestIterators-tests.jar"); + Path dstPath = new Path(CONTEXT_DIR + "/Test.jar"); + fs.copyFromLocalFile(srcPath, dstPath); + // Sleep to ensure jar change gets picked up + UtilWaitThread.sleep(5000); + try { + Connector c = getConnector(); + // Set the classloader context property on the table to point to the TestIterators jar file. + c.instanceOperations().setProperty(CONTEXT_PROPERTY, CONTEXT_CLASSPATH); + // Insert rows with the word "Test" in the value. + String tableName = getUniqueNames(1)[0]; + c.tableOperations().create(tableName); + BatchWriter bw = c.createBatchWriter(tableName, new BatchWriterConfig()); + for (int i = 0; i < ITERATIONS; i++) { + Mutation m = new Mutation("row" + i); + m.put("cf", "col1", "Test"); +	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3948
16418	ACCUMULO-3948	Github user dlmarion commented on a diff in the pull request: https://github.com/apache/accumulo/pull/51#discussion_r44656596 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/AbstractInputFormat.java --- @@ -484,7 +509,7 @@ public void initialize(InputSplit inSplit, JobConf job) throws IOException { if (null == authorizations) { authorizations = getScanAuthorizations(job); } - + String classLoaderContext = getContext(job); --- End diff -- None of the other variables in the method have the final modifier. Are you trying to raise the bar?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3948
16419	ACCUMULO-3948	Github user dlmarion commented on the pull request: https://github.com/apache/accumulo/pull/51#issuecomment-156186156 Can you point me to some tests that I need to modify?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3948
16420	ACCUMULO-3948	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/51#discussion_r44715330 --- Diff: core/src/main/java/org/apache/accumulo/core/client/ScannerBase.java --- @@ -242,4 +242,31 @@ * @since 1.8.0 */ long getBatchTimeout(TimeUnit timeUnit); + + /** + * Sets the name of the classloader context on this scanner. See the administration chapter of the user manual for details on how to configure and use + * classloader contexts. + * + * @param classLoaderContext + * name of the classloader context + * @throws NullPointerException + * if context is null + * @since 1.8.0 + */ + void setClassLoaderContext(String classLoaderContext); --- End diff -- What happens when a user sets a context that does not exists? When this happens, does the exception have an informative message that makes it easy for the user to determine what happened? Is there a test for this case?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3948
16421	ACCUMULO-3948	Github user dlmarion commented on a diff in the pull request: https://github.com/apache/accumulo/pull/51#discussion_r44968062 --- Diff: core/src/main/java/org/apache/accumulo/core/client/ScannerBase.java --- @@ -242,4 +242,31 @@ * @since 1.8.0 */ long getBatchTimeout(TimeUnit timeUnit); + + /** + * Sets the name of the classloader context on this scanner. See the administration chapter of the user manual for details on how to configure and use + * classloader contexts. + * + * @param classLoaderContext + * name of the classloader context + * @throws NullPointerException + * if context is null + * @since 1.8.0 + */ + void setClassLoaderContext(String classLoaderContext); --- End diff -- Can I get the server side configuration inside a ScannerImpl or TabletServerBatchReader? I don't see how, and if I did, do we want the client to get the server side configuration?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3948
16422	ACCUMULO-3948	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/51#discussion_r44980676 --- Diff: core/src/main/java/org/apache/accumulo/core/client/ScannerBase.java --- @@ -242,4 +242,31 @@ * @since 1.8.0 */ long getBatchTimeout(TimeUnit timeUnit); + + /** + * Sets the name of the classloader context on this scanner. See the administration chapter of the user manual for details on how to configure and use + * classloader contexts. + * + * @param classLoaderContext + * name of the classloader context + * @throws NullPointerException + * if context is null + * @since 1.8.0 + */ + void setClassLoaderContext(String classLoaderContext); --- End diff -- There are at least two ways to do this check. Do it when setClassLoader is called. Or do it when the thrift calls startScan and continueScan are made. An exception could be added to these thrift calls. How were you thinking of doing the check? For getting server side table props you can look into the implementation of `TableOperations.getProperties()`	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3948
16423	ACCUMULO-3944	I think I added this in ACCUMULO-514 when I noticed the system doing nothing for a bit of time when there was work to be done. I think the checks are usually just against in memory data structs and very quick (there may be cases where the checks go to hdfs). Were you seeing performance problems with this code?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3944

#	Requirement	Questions	Answered	Successfully Answered	URL
16424	ACCUMULO-3925	I think there are lots of places in the FateOps that throw ThriftTableOperationException (henceforth TTOE). But only a few were ignored. Did you go look at all of the TTOEs and only decide on those? I ask because another strategy could be to supply FATE a list of Exeception types not to warn on. For example ask FATE to never warn on a TTOE. This is not as clean as letting the FateOp decide.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3925
16425	ACCUMULO-3918	How is this different than between two files? You can't know exactly when the ZK props changed or when they propagated to the tserver. The fact that two locality groups share the same file says nothing about how those locality groups are expected to behave with respect to each other. I can understand the desire to have a consistent iterator stack for a single file... I just can't see how it would affect the user experience, because they don't know which data gets written to which file, or which locality groups are co-located in the same files anyway (we don't expose that to them, and it's not formally defined, nor is it guaranteed to not change). If this is considered a bug, then why isn't it considered a bug that this can happen between sequential compactions to different files (rather than just different locality groups)? I guess I'm just not understanding the use case. FWIW (not much), I typically advise people in situations where this matters, to take the table offline, change the table configuration, and then bring it back online.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3918
16426	ACCUMULO-3918	Well, once you have a consistent iterator stack for the whole file, you can log the iterator stack for the file (which would be follow on work). This would describe what iterator settings were in place when the file compacted. This is not the same as sequential compactions of different files because files are apparent to users. It's how users think about the data on disk. They would track that a major compaction happened and it outputted a file. I wouldn't expect the iterator stack to allow changes midway through a scan, why would I design things to change midway through a compaction. And telling people to offline their tables is horrible advise. That's a issue we should strengthen, not tell people to suffer through.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3918
16427	ACCUMULO-3918	bq. And telling people to offline their tables is horrible advise. That's a issue we should strengthen, not tell people to suffer through. Well, it's the only consistent thing we have at the moment... It's not ideal, but it does work. Can you ensure consistency for what you describe here given the race-conditions for ZK-propagation in general? The problem as you described it seems simple enough to fix, but I haven't looked to see if we'd just be kicking the can down the road.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3918
16428	ACCUMULO-3918	bq. you can log the iterator stack for the file (which would be follow on work) This sounds interesting for debugging purposes. Is there an existing ticket for this? Also, it seems you could just do the same thing per locality group, too. bq. files are apparent to users. It's how users think about the data on disk. I disagree. Files have nothing to do with the user's data model, and unless they are administering their own cluster, I don't know many users who even realize their is a relationshp between their data and the files stored on disk. Locality groups, on the other hand, at least have a loose relationship with the data model, because they are sets of column families. That's not to say they're not important... especially for performance and administration/debugging, of course. In any case, in your scenario, you're still going to have multiple current files for a tablet which have been potentially compacted with different iterator stacks. I don't see how that's fundamentally different than locality groups being compacted with different iterator stacks, unless you're only talking about full major compactions. Are you only concerned about those? (If so, that would go a long way towards my understanding your scenario.) bq. And telling people to offline their tables is horrible advise. That's a issue we should strengthen, not tell people to suffer through. I don't think it's "horrible advise[sic]". It's the best advice possible, given that it's pretty much the best guarantee we can give. I do agree it's insufficient, since there's not an online sync option available. I believe it's been proposed before to offer an online solution for ZK property sync, via the API... maybe an option to clear ZooCache for all the per-table properties.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3918
16429	ACCUMULO-3918	bq. We know that ZK has a defined window for propagation, Is that true? We know ZK has a defined window to maintain its session, but I don't think we have a defined window for any particular set of observers on a ZK watcher to complete and the ZooCache/TableConfiguration to be updated.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3918
16430	ACCUMULO-3913	What's the tangible benefit for users that you're going for here, Keith? I assume this is just one piece of a bigger picture. Took a quick glance at your branch, is there a reason that the Sampler isn't a specialization of SKVI? Seems very close to what Filter already does.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3913
16431	ACCUMULO-3913	bq. What's the tangible benefit for users that you're going for here, Keith? I assume this is just one piece of a bigger picture. Over the years I have heard multiple request for sampling. Recently someone was asking about the feasibility of supporting scanning the RFile indexes as an approximate sample. This was interesting idea, but it bugged me that different rfile indexes would have different subsets of keys. So if one RFile index had a key that was deleted in a later RFile, you would likely not see the delete. So I started thinking about pre-generating a sample that had a consistent set of keys. bq. is there a reason that the Sampler isn't a specialization of SKVI? Seems very close to what Filter already does. I have not put a lot of thought into this particular aspect. I mainly wanted to do a quick end to end prototype to see if there were any big gotcha's with the overall idea. One issue w/ using Filter for the sample function, may be that seek support is not needed.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3913
16432	ACCUMULO-3913	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/46#discussion_r38815363 --- Diff: core/src/main/java/org/apache/accumulo/core/client/ClientSideIteratorScanner.java --- @@ -297,4 +345,50 @@ public void setReadaheadThreshold(long batches) { } this.readaheadThreshold = batches; } + + private SamplerConfiguration getIteratorSamplerConfigurationInternal() { + SamplerConfiguration scannerSamplerConfig = getSamplerConfiguration(); + if (scannerSamplerConfig != null) { + if (iteratorSamplerConfig != null && !new SamplerConfigurationImpl(iteratorSamplerConfig).equals(new SamplerConfigurationImpl(scannerSamplerConfig))) { --- End diff -- Can the instantiation of the objects in the equals method be broken out of the if condition? There's a lot going on for a simple equality check that would be easier to follow if the impls are instantiated earlier IMO	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3913

#	Requirement	Questions	Answered	Successfully Answered	URL
16433	ACCUMULO-3913	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/46#discussion_r39076294 --- Diff: core/src/main/java/org/apache/accumulo/core/sample/impl/SamplerConfigurationImpl.java --- @@ -0,0 +1,181 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.sample.impl; +import java.io.DataInput; +import java.io.DataOutput; +import java.io.IOException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.LinkedHashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; + +import org.apache.accumulo.core.client.admin.SamplerConfiguration; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.tabletserver.thrift.TSamplerConfiguration; +import org.apache.accumulo.core.util.Pair; +import org.apache.hadoop.io.Writable; + +public class SamplerConfigurationImpl implements Writable { + private String className; + private Map<String,String> options; + + public SamplerConfigurationImpl(DataInput in) throws IOException { + readFields(in); + } + + public SamplerConfigurationImpl(SamplerConfiguration sc) { + this.className = sc.getSamplerClassName(); + this.options = new HashMap<>(sc.getOptions()); + } + + public SamplerConfigurationImpl(String className, Map<String,String> options) { + this.className = className; + this.options = options; + } + + public SamplerConfigurationImpl() {} + + public String getClassName() { + return className; + } + + public Map<String,String> getOptions() { + return Collections.unmodifiableMap(options); + } + + @Override + public int hashCode() { + return 31 * className.hashCode() + options.hashCode(); + } + + @Override + public boolean equals(Object o) { + if (o instanceof SamplerConfigurationImpl) { + SamplerConfigurationImpl osc = (SamplerConfigurationImpl) o; + return className.equals(osc.className) && options.equals(osc.options); + } + return false; + } + + @Override + public void write(DataOutput out) throws IOException { --- End diff -- How is this done with DataInput and Dataoutput?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3913
16434	ACCUMULO-3913	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/46#discussion_r39187156 --- Diff: core/src/main/java/org/apache/accumulo/core/sample/impl/SamplerConfigurationImpl.java --- @@ -0,0 +1,181 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.sample.impl; +import java.io.DataInput; +import java.io.DataOutput; +import java.io.IOException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.LinkedHashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; + +import org.apache.accumulo.core.client.admin.SamplerConfiguration; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.tabletserver.thrift.TSamplerConfiguration; +import org.apache.accumulo.core.util.Pair; +import org.apache.hadoop.io.Writable; + +public class SamplerConfigurationImpl implements Writable { + private String className; + private Map<String,String> options; + + public SamplerConfigurationImpl(DataInput in) throws IOException { + readFields(in); + } + + public SamplerConfigurationImpl(SamplerConfiguration sc) { + this.className = sc.getSamplerClassName(); + this.options = new HashMap<>(sc.getOptions()); + } + + public SamplerConfigurationImpl(String className, Map<String,String> options) { + this.className = className; + this.options = options; + } + + public SamplerConfigurationImpl() {} + + public String getClassName() { + return className; + } + + public Map<String,String> getOptions() { + return Collections.unmodifiableMap(options); + } + + @Override + public int hashCode() { + return 31 * className.hashCode() + options.hashCode(); + } + + @Override + public boolean equals(Object o) { + if (o instanceof SamplerConfigurationImpl) { + SamplerConfigurationImpl osc = (SamplerConfigurationImpl) o; + return className.equals(osc.className) && options.equals(osc.options); + } + return false; + } + + @Override + public void write(DataOutput out) throws IOException { --- End diff -- > RFile having a "thrift" dependency is a valid point, just not for worry about backwards compatibility IMO. agreed. I don't know of any issues w/ thrift serialization. > That's a smell to me. I have no stake besides the opinion that we should have one and not many serializations. :) It is screwy. What do you think about adding comments documenting the reasons for the two serializations?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3913

#	Requirement	Questions	Answered	Successfully Answered	URL
16435	ACCUMULO-3913	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/46#discussion_r39188675 --- Diff: core/src/main/java/org/apache/accumulo/core/sample/impl/SamplerConfigurationImpl.java --- @@ -0,0 +1,181 @@ + +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.sample.impl; + +import java.io.DataInput; +import java.io.DataOutput; +import java.io.IOException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.HashMap; +import java.util.LinkedHashMap; +import java.util.List; +import java.util.Map; +import java.util.Map.Entry; + +import org.apache.accumulo.core.client.admin.SamplerConfiguration; +import org.apache.accumulo.core.conf.AccumuloConfiguration; +import org.apache.accumulo.core.conf.Property; +import org.apache.accumulo.core.tabletserver.thrift.TSamplerConfiguration; +import org.apache.accumulo.core.util.Pair; +import org.apache.hadoop.io.Writable; + +public class SamplerConfigurationImpl implements Writable { + private String className; + private Map<String,String> options; + + public SamplerConfigurationImpl(DataInput in) throws IOException { + readFields(in); + } + + public SamplerConfigurationImpl(SamplerConfiguration sc) { + this.className = sc.getSamplerClassName(); + this.options = new HashMap<>(sc.getOptions()); + } + + public SamplerConfigurationImpl(String className, Map<String,String> options) { + this.className = className; + this.options = options; + } + + public SamplerConfigurationImpl() {} + + public String getClassName() { + return className; + } + + public Map<String,String> getOptions() { + return Collections.unmodifiableMap(options); + } + + @Override + public int hashCode() { + return 31 * className.hashCode() + options.hashCode(); + } + + @Override + public boolean equals(Object o) { + if (o instanceof SamplerConfigurationImpl) { + SamplerConfigurationImpl osc = (SamplerConfigurationImpl) o; + return className.equals(osc.className) && options.equals(osc.options); + } + return false; + } + + @Override + public void write(DataOutput out) throws IOException { --- End diff -- > It is screwy. What do you think about adding comments documenting the reasons for the two serializations? I think that's fine. I hadn't initially realized that this would be serialized in RFile. Until we consider something a bit more robust for serializing structures in RFile (I had talked about replacing the header with protobuf already), I think it's fine as-is.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3913
16436	ACCUMULO-3913	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/46#discussion_r39469628 --- Diff: core/src/main/java/org/apache/accumulo/core/compaction/CompactionSettings.java --- @@ -21,6 +21,7 @@ public enum CompactionSettings { + SF_NO_SAMPLE(new NullType()), --- End diff -- Could have static instances of these, but I don't think it's necessary since they're wrapped in the enum (e.g. we'd only have one instance anyways). Is that right?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3913
16437	ACCUMULO-3903	Is this work best done here or in HTrace libraries? Not sure I understand the boundaries well enough.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3903
16438	ACCUMULO-3903	bq. Is this work best done here or in HTrace libraries? Not sure I understand the boundaries well enough. (hopefully I don't get it wrong) Accumulo is a "SpanReceiver" which means that its an endpoint that receives Spans from clients that did timings. Clients include public API, server-side Accumulo functionality, HDFS, or your personal code. Accumulo has indexing functionality for writing those spans to a table (the Accumulo table). This is what backs the monitor. I think there is work going on in HTrace (via htraced) which is backed by LevelDB that is intended for small-to-medium based deployments. IMO, it makes sense to maintain our own SpanReceiver (since we have this scalable distributed database already over leveldb), but it does result in duplication of presentation efforts between Accumulo and htraced.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3903
16439	ACCUMULO-3903	bq. ...but it does result in duplication of presentation efforts between Accumulo and htraced. Does there have to be duplication? HTRACE-180 is trying to separate the display webapp. Could Accumulo use the HTrace UI running against Accumulo as a span store? (Would have to implement a few REST calls for the UI to pull on).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3903
16440	ACCUMULO-3903	bq. Does there have to be duplication? HTRACE-180 is trying to separate the display webapp Absolutely not, you're right -- that was a bad assumption of me to paint. bq. Could Accumulo use the HTrace UI running against Accumulo as a span store? Not sure how much work this would be. I'd have to take a look at htraced to see what it would actually take. Would be nice if we could just drop in Accumulo w/ the necessary REST service.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3903
16441	ACCUMULO-3903	I don't have a pressing need to have a solution right at this minute, so I would be interested in waiting for the development in Htrace webapp to stabilise. It seems like it would be a good approach to have one presentation. [~elserj], just wondering what is the status of the redevelopment of the monitor in (ACCUMULO-3005)? Do you think that this would work with displaying the Htrace GUI inside of the Accumulo monitor and still keep a consistent look and feel (i suppose that this could be addressed through css)? [~stack], would any work need to be done to the Htrace webapp to make it more suited for embedding it in other webapps? If so, would you like tickets created to cover this in Htrace Jira?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3903
16442	ACCUMULO-3903	[~Lonepine] bq. would any work need to be done to the Htrace webapp to make it more suited for embedding it in other webapps.... To be seen. Currently we just build a WAR. I've not tried deploying it in a 'container'. If you get to it before I, file issues against HTrace. What is there we know basically works for its what the htraced span receiver is putting up as UI.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3903
16443	ACCUMULO-3903	bq. just wondering what is the status of the redevelopment of the monitor in (ACCUMULO-3005)? Do you think that this would work with displaying the Htrace GUI inside of the Accumulo monitor and still keep a consistent look and feel (i suppose that this could be addressed through css)? [~Lonepine] it's essentially stalled. I was happy with the direction it was heading, but I haven't touched it in a long time. Would like to get back to it again. I'm not sure the level of effort to pull in two wars and spin things up that way. Look and feel can be something we address later on -- functionality is still first and foremost. If this is something you're interested in pursuing, let's try to firm up a plan that we think is the correct way forward and we can scope it out :)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3903
16444	ACCUMULO-3899	Are you up to date? Im pretty sure I fixed this already.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3899
16445	ACCUMULO-3899	I am not (probably a week behind), can you close this with a dupe against the fixing ticket?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3899

#	Requirement	Questions	Answered	Successfully Answered	URL
16446	ACCUMULO-3891	Cool, thanks. I figure that if we can include as much as possible w/o being a burden, that will simply interactions with it (less to think about 'will this page actually load in $deployment')	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3891
16447	ACCUMULO-3871	I see you committed a change to master which moved all of the integration tests from src/test/java to src/main/java. Can I ask why you did that? It seemed like the approach you were heading towards was creating a special fatjar with all of the stuff you needed to avoid libjars taking forever. It's not apparent to me why you needed to move the ITs to src/main as you could have just used the accumulo-test-$version-tests.jar. I'm concerned about not having our test code in the normal location to just support something that is not going to be generally used by everyone (most people won't have many nodes to run the ITs via mapreduce) Thanks.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3871
16448	ACCUMULO-3871	Thanks for the link. I missed that point from Eric's original point. Any suggestions on how can we prevent ITs from showing up against in test/src/test/java? Checkstyle rule? Developer section in the user manual? Email?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3871
16449	ACCUMULO-3858	Looks like you already made a similar change in master, [~ecn]. Was this specifically for the new WAL znodes? {code} commit 3fdd29f5222f9d1d32ca28b5ecf1d740a8d20f87 Author: Eric C. Newton <eric.newton@gmail.com> Date: Fri Apr 24 18:15:05 2015 -0400 ACCUMULO-3423 optimize WAL metadata table updates diff --git a/test/src/test/java/org/apache/accumulo/test/functional/WatchTheWatchCountIT.java b/test/src/test/java/org/apache/accumulo/test/functional/WatchTheWatchCountIT.java index 3b1dd2f..140fd59 100644 --- a/test/src/test/java/org/apache/accumulo/test/functional/WatchTheWatchCountIT.java +++ b/test/src/test/java/org/apache/accumulo/test/functional/WatchTheWatchCountIT.java @@ -52,7 +52,7 @@ public class WatchTheWatchCountIT extends ConfigurableMacIT { String response = new String(buffer, 0, n); long total = Long.parseLong(response.split(":")[1].trim()); assertTrue("Total watches was not greater than 500, but was " + total, total > 500); - assertTrue("Total watches was not less than 650, but was " + total, total < 600); + assertTrue("Total watches was not less than 675, but was " + total, total < 675); } finally { socket.close(); } {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3858
16450	ACCUMULO-3851	I thought client configuration was added in 1.7.0? Was it added in a 1.6.z?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3851
16451	ACCUMULO-3851	could we add it to the 1.6.0 release notes then? I don't see any mention of it there. I think it came in as a part of ACCUMULO-1009, but it's a big addition with a non-obvious relationship.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3851
16452	ACCUMULO-3842	As I see it, we have two main non-transient use cases for ZooKeeper: # Bootstrap information (primarily, ZooKeeper serves as an entry point for clients communicating with Accumulo). # Information which needs some degree of distributed consistency. A lot of the configuration falls into the second category (soon-ish consistency with ZK watchers and ZooCache). I think there's probably some stuff we could move over, but I also see other ways we can address some of those benefits you mention. * Loss of ZK configuration could be handled with some sort of snapshotting. * The number of watchers could be reduced greatly by changing our strategy for storing stuff. We could either have a single "watched" node (or one per table) and update that node when we make any other change. That would reduce the need to have as many watchers. We could also serialize table configuration a bit better (one node instead of many), which would also help address ACCUMULO-1568. I'm not sure how the change proposed would manifest the third benefit you mention (consistent updates of table props). Can you explain that, please? As I understand it, we use ZooKeeper, because it has watchers, which we can use to get consistency. I'm not aware of any similar mechanism with any alternatives.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16453	ACCUMULO-3842	Thanks for taking a read over this (very) bare outline so far. bq. I'm not sure how the change proposed would manifest the third benefit you mention (consistent updates of table props). Can you explain that, please? As I understand it, we use ZooKeeper, because it has watchers, which we can use to get consistency. I'm not aware of any similar mechanism with any alternatives. So, right now, we have eventually consistent configuration updates for tables. We don't know when the watchers will fire, but (IIRC) we know they will fire in the correct order and ever server will eventually see all updates. What we should really have to mimic the API we present is a strongly consistent means to update configurations. ZooKeeper doesn't keep us from accomplishing this. We would need to write code to actually get the strong consensus for ourselves. I know this is very hand-wavy at this point, but I think we're at the point where this is a problem we need to start thinking about because it's been a repeated problem for ourselves just in writing reasonable tests for Accumulo for ~2years now.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16454	ACCUMULO-3842	{quote} What we should really have to mimic the API we present is a strongly consistent means to update configurations. ZooKeeper doesn't keep us from accomplishing this. We would need to write code to actually get the strong consensus for ourselves. I know this is very hand-wavy at this point, but I think we're at the point where this is a problem we need to start thinking about because it's been a repeated problem for ourselves just in writing reasonable tests for Accumulo for ~2years now. {quote} Once it's finalized, could we leverage the recent HBase procedure store for this (ref HBASE-12439)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16455	ACCUMULO-3842	bq. Once it's finalized, could we leverage the recent HBase procedure store for this To the best of my understanding (which is lacking), I wasn't aware of procv2 solving this problem for HBase. Am I mistaken?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16456	ACCUMULO-3842	bq. ZooKeeper doesn't keep us from accomplishing this. We would need to write code to actually get the strong consensus for ourselves. Isn't this statement is true for HDFS also? Wether using zookeeper or HDFS, config needs to be cached on tservers for efficiency. For most operations config is not changing, do not want to make a synchronous read to ZK of HDFS before servicing each RPC. I think we need something at the API level to address this, regardless of implementation. bq. but I think we're at the point where this is a problem we need to start thinking about because it's been a repeated problem for ourselves just in writing reasonable tests for Accumulo for ~2years now. Yes, this is a big problem, lets fix it. I think ZK provides primitives that make adding an Accumulo API to wait for tservers to synch config to a certain point fairly easy to implement. I think a good place to start is thinking what the API solutions for this would look like. Some possible API changes : # Add something like {{waitForConfigSynch(long timeout, TimeUnit tu)}} to the API. If it times out could throw an exception w/ info about the tservers that did not update. # Change behavior of current methods that change config to wait for all tservers to sync (I don't like doing this, would rather leave the current behavoir). # Add a new method that allows setting a batch of key/value configs. This method would also wait for changes to sync across tservers. Something like {{updateConfiguration(Map<String, String> propsToSet, Set<String> propsToUnset, long timeout, TimeUnit tu)}}. Of course pass in class instead of 4 params. I kinda like 3. What other API changes could solve this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842

#	Requirement	Questions	Answered	Successfully Answered	URL
16457	ACCUMULO-3842	{quote} bq. ZooKeeper doesn't keep us from accomplishing this. We would need to write code to actually get the strong consensus for ourselves. Isn't this statement is true for HDFS also? Wether using zookeeper or HDFS, config needs to be cached on tservers for efficiency. For most operations config is not changing, do not want to make a synchronous read to ZK of HDFS before servicing each RPC. I think we need something at the API level to address this, regardless of implementation. {quote} I haven't put nearly enough thought into this (nor reading necessary literature), but my original thought was to use Accumulo tables as much as possible for the persistence (and wrap the access in some classes to make everything more natural). If we keep a notion of monotonically increasing versions for configs, a FATE op could wait for each server to report at least a minimal version. Discerning the "end" state might be difficult in the face of servers dying and starting...	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16458	ACCUMULO-3842	bq. Loss of ZooKeeper doesn't lose table configuration and users Is this a genuine problem for users? Keeping configuration information in ZK is pretty much ideal from my perspective.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16459	ACCUMULO-3842	{quote} bq. Loss of ZooKeeper doesn't lose table configuration and users Is this a genuine problem for users? {quote} IMO, we don't have enough rigor here to actually make me comfortable to tell someone that we have the means to backup and restore ZooKeeper. I know we have some tools, but I don't know how much this is actually used and tested. Ideally, it would be better if we can have a single way to backup data for an Accumulo cluster and later restore it (import/export table is a likely candidate for simplicity). Perhaps removing table configuration in ZK isn't a good idea, but I definitely think users information should not be stored in ZK.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16460	ACCUMULO-3842	bq. I definitely think users information should not be stored in ZK. Do you just mean the user password/authorizations/permissions database? Those are already pluggable and do not have to live in ZK. It would be a pain to put them inside Accumulo, creating a bit of circular dependency which would be prone to concurrency problems, but it could be doable. Do you have a strong reason you can express for changing the default behavior of storing this in ZK? Or is it just the lack of good comprehensive backup/restore tools which you've already mentioned (which seems to me to be an easier problem to solve)? From my perspective, ZK seems to be a relatively solid component. Because of that, it seems to me that burden is on any alternative to demonstrate a greater degree of reliability, scalability, or other benefit.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16461	ACCUMULO-3842	{quote} Do you have a strong reason you can express for changing the default behavior of storing this in ZK? Or is it just the lack of good comprehensive backup/restore tools which you've already mentioned (which seems to me to be an easier problem to solve)? {quote} One side of it is that I think the more we can store in tables, the easier it is for users to reason about. Having as much data as we can stored in Accumulo itself would be easier for users to say "how is X stored?". Another aspect is the scalability side. We know that making the problem fit in ZK's constraints is often an exercise to interacting with it at scale. We can likely do better scalability-wise in Accumulo itself. We probably aren't at that point where we _need_ to do this yet, but it's something to keep in mind. {quote} From my perspective, ZK seems to be a relatively solid component. Because of that, it seems to me that burden is on any alternative to demonstrate a greater degree of reliability, scalability, or other benefit. {quote} I'm pretty sure we already have solved the reliability and scalability problems in Accumulo itself. A bit of eating our own dogfood by leveraging a table. We're better suited at storing a larger category of data than ZooKeeper is (e.g. we can store larger blobs of data). Just to be clear: I think this has kind of broken down into two separate discussions already. One is the configuration consistency issue which I think we're all in agreement on needs to happen. The latter is my general opinion that we should start moving away from ZK as our "general not-table-related-data" store which deserves the continued discussion about what the actual observed benefits would be, short term and long term (and how, if at all, the two are related).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3842
16462	ACCUMULO-3841	What you propose might be possible. However, it's not just NativeMapIT which requires the native maps to be built. Many other tests probably rely on timing which only works reliably if the native maps are enabled (I suppose we could adjust these, as we find them). The dependency in the test pom on the native maps artifact would have to be moved to a profile (perhaps activated by the presence of the native maps). I'm not 100% sure the extent of this impact on our tests. Personally, I think it's reasonable to consider it a build requirement to have `make`, `g++`, and `javah` installed in order to build Accumulo. I am curious why the build proceeded at all if `make` failed, though. Was this in Jenkins or some other system where maven was executed with `--fail-at-end`? Otherwise, I'm not sure why accumulo-native would continue to be built.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3841
16463	ACCUMULO-3841	bq. I am curious why the build proceeded at all if `make` failed, though. Was this in Jenkins or some other system where maven was executed with `--fail-at-end`? Otherwise, I'm not sure why accumulo-native would continue to be built. Oh that might have been the case, actually. Thank you for debugging for me :P	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3841
16464	ACCUMULO-3840	That's a great idea and I'll work on it. I think there are a few other things that can go into the sub-class to help clean things up. One thing that I did notice is that OpTimer is only ever used at TRACE level (except for a test case) - even in a place or two where the comments say that for performance, logging should not be done. - That would seem to indicate that tracing would not be appropriate at that level and may for all of OpTimer usages. The reason that I didn't go further in this round with tracing is so that it will be done with consideration and consistently. The thing to do would be to get a feeling for where we need to measure client operations and then add them sort of top / down. Are you thinking nanoTime instead of currentTimeMillis? I thought about borrowing the Guava StopWatch / Ticker to approach to allow us to swap time sources, but then figured that it would be superseded by tracing and that nanos where unnecessary with the precision that we need and the fact that these are hardly ever used.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3840

#	Requirement	Questions	Answered	Successfully Answered	URL
16465	ACCUMULO-3840	bq. One thing that I did notice is that OpTimer is only ever used at TRACE level (except for a test case) - even in a place or two where the comments say that for performance, logging should not be done. - That would seem to indicate that tracing would not be appropriate at that level and may for all of OpTimer usages. Oh interesting. That's a fun comment. I imagine a bit of work to go into _good_ instrumentation. We have it sprinkled around, but could likely benefit from a planned visitation of the public API impl. bq. Are you thinking nanoTime instead of currentTimeMillis? I thought about borrowing the Guava StopWatch / Ticker to approach to allow us to swap time sources, but then figured that it would be superseded by tracing and that nanos where unnecessary with the precision that we need and the fact that these are hardly ever used. I remember reading something about how using nanoTime was better than currentTimeMillis for timings. Maybe it was https://blogs.oracle.com/dholmes/entry/inside_the_hotspot_vm_clocks {quote} If you are interested in measuring absolute time then always use System.currentTimeMillis(). Be aware that its resolution may be quite coarse (though this is rarely an issue for absolute times.) If you are interested in measuring/calculating elapsed time, then always use System.nanoTime(). On most systems it will give a resolution on the order of microseconds. Be aware though, this call can also take microseconds to execute on some platforms. {quote} The only other thing I'd request is to try to keep OpTimer/StopWatch improvements and consolidations as separate as possible from the log4j/slf4j switchover. We don't have a hard rule for one-change-per-commit, but it definitely helps if you can do that.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3840
16466	ACCUMULO-3840	Why is there a PR and something on RB?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3840
16467	ACCUMULO-3807	I agree with [~ctubbsii] having progress sit at %100 is worse than a moving indicator with funny looking numbers. Having numbers > than 100% does look odd, but as long as it is moving it reflects that recovery is working and will eventually (hopefully) complete. If it was clamped at 100% and then hung for any significant amount of time (maybe > 30 seconds), well I'm likely to start to panic that something else is wrong with the system. First, something happened that caused log recovery, and now log recovery is hanging.... With 1.4.x versions I seem to recall recovery taking minutes in "bad" cases. If the recovery indicator was at 100% for that amount of time it would certainly cause me concern. Maybe instead of a %, the number shown is # copy / sort operations? While it would be nice to know about how long things were going to take, if it can be estimated - otherwise it seems more important to provide feedback and reflect that things are still progressing - whatever form that is.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3807
16468	ACCUMULO-3795	Maybe {{table.scan.max.latency}} to keep some consistency with {{table.scan.max.memory}} instead of {{table.scan.flush.ms threshold}}? We already have nice parsing for {{TimeDuration}} in {{Property}}s. Don't need to specifically set ms. Sounds like a good addition though.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3795
16469	ACCUMULO-3793	bq. Do this by default in the start scripts. I'm not sure I'm a big fan of trying to change kernel parameters in our start scripts (nor sure how it would work as I assume you need root privileges). Wouldn't something like this be better served in a "System Configuration" section of the user manual which could also include things like disabling transparent huge pages, reducing swappiness, etc?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3793
16470	ACCUMULO-3774	Is anyone working on this one yet? I have some cycles I can start looking into this.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3774
16471	ACCUMULO-3774	It does not seem patch #1 covers the general case that the root tablet tries to write to itself after minc. Why should the root tablet should write to itself? [~elserj] saw another place where the root tablet writing to itself was causing problems.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3774
16472	ACCUMULO-3774	The problem (as I'm working through it): # at some point a tserver figures out that a WALog is no longer in use: * it is closed * no tablet has a references to it * the log is marked as unused to signal to the GC that the file can be removed # one of the likely times a log goes unused, is during a shutdown # at that point (but there are others) the metadata table is not available Some ideas: * mark the log *eventually* asynchronously * keep the log references in zookeeper only * use HDFS to track the logs, avoid the metadata table So, can zookeeper handle this? * a tablet server can theoretically ingest as fast as the WAL can write, which is the speed of HDFS * let's be an optimist and use a write rate of 128M/s * this allows for a log roll every 8 seconds on our magic hardware * let's assume a big cluster size: 10K nodes * each rollover updates zookeeper 3x: create, unused, deleted * 10K*3/8 = 3750 updates / sec * assuming 100 WALogs per server, and 500 bytes for each entry: 0.5G * based on published numbers (I found [this|http://wiki.apache.org/hadoop/ZooKeeper/Performance]), it looks like this would be easily do-able, even on quorum large enough to handle this magic cluster * I'm pretty sure the NN (or several of them) can keep up, since they are handling the WALs creation/deletion anyhow What are the benefits to each possible solution? Async: * async updates would require a recovery, probably for every "clean" shutdown Zookeeper: * we'll have to sync to avoid propagation delays (like knowing the full log set before recovery starts) * probably some other problems, zookeeper doesn't always act like I think it should NN: * using the NN to track files was not a particularly scalable approach, but it would be easy, and the multi-volume support allows us to abuse the NN a little more * given the current logging area {{/accumulo/walogs/host+port/UUID}} we would need some way for a restarted server to begin logging before all recovery has completed At this point, I'm leaning towards marking logs as unused by renaming them in HDFS, and storing them under the host+port+session. This eliminates the entry for the log from the metadata table altogether.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3774
16473	ACCUMULO-3774	Would it make things easier if we pinned a specific WAL _just_ for the root tablet? We could have a single place in ZK to track that and then specific logic for tservers. I'm not sure how much more complexity that would add though. Then we should use ZK+sync to make sure we always have a consistent view of the WAL for root (and have some great logging whenever we do something to it).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3774
16474	ACCUMULO-3771	Github user ctubbsii commented on the pull request: https://github.com/apache/accumulo/pull/35#issuecomment-99120050 There are two main problems with the Eclipse formatter: 1. It doesn't strictly adhere to the line length restrictions. 2. It adds a trailing space to indent empty javadoc lines. We've already accounted for the line length in the checkstyle rules by making it much larger than the Eclipse formatter. With this patch, I think we could probably just ignore trailing whitespace (and drop the corresponding checkstyle enforcement), since the formatter will remove the ones we mainly care about anyway (the javadoc ones are annoying, but there's not a lot of good options so long as Eclipse fails to fix that issue). We could just strip whitespace with `sed` after the formatter executes. One thing I'd like to see, though (for the purposes of a review) is just the pom changes, separate from any resulting formatting changes (I'm sure I can do this with some command-line Git Fu, though, so it's no big deal). I did see that the formatting significantly changed some lines that looked deliberately formatted to make it more readable. For those lines, we may wish to add the `@formatter:off` / `@formatter:on` tags to retain existing formatting. I'm also curious what happens with the `-Pthrift` profile activated. Will this plugin reformat the generated thrift code? Maybe that's okay, but it might be annoying. If the generated thrift code were in a separate package, it could be avoided more easily, though. I'm overall +1 to the idea, but it might need some tweaking. I especially like the idea that we could focus our style choices on readability, because writing to conform to a particular style doesn't matter any more (write in whatever style you like and the build will fix it).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3771

#	Requirement	Questions	Answered	Successfully Answered	URL
16475	ACCUMULO-3771	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/35#issuecomment-99123399 > One thing I'd like to see, though (for the purposes of a review) is just the pom change The PR has two commits. Commit 4b69fba has just pom changes. Is that what you are looking for? > It adds a trailing space to indent empty javadoc lines. Would you happen to know if this if fixed in later versions of Eclipse? This formatter plugin currently [relies on Eclipse 3.8.1](https://github.com/mblanchette/maven-java-formatter-plugin/blob/0.4/pom.xml#L66) code to do the formatting.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3771
16476	ACCUMULO-3771	Github user mikewalch commented on a diff in the pull request: https://github.com/apache/accumulo/pull/35#discussion_r29801880 --- Diff: maven-plugin/src/it/plugin-test/pom.xml --- @@ -113,6 +113,13 @@ </execution> </executions> </plugin> + <plugin> + <groupId>com.googlecode.maven-java-formatter-plugin</groupId> + <artifactId>maven-java-formatter-plugin</artifactId> + <configuration> + <configFile>${project.basedir}/../../../../contrib/Eclipse-Accumulo-Codestyle.xml</configFile> --- End diff -- What is LF?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3771
16477	ACCUMULO-3770	Where do these reports go? Are they rolled-over like the regular logs? The timeliness feature is nice, but that can be done in a separate SimpleTimer callback. Knowing the status of GC in a long-running, time-sensitive server is very important. I would be hesitant to remove it.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3770
16478	ACCUMULO-3770	Was talking with Bitweasil in IRC yesterday/today about this. He expressed some confusion about the GC logging we had. After digging, he found that we were actually logging the cumulative time spent in GC, not the time of the last cycle as expected. bq.Where do these reports go? Are they rolled-over like the regular logs? Typically, you'll get a {{gc.log\[\-date\]}} style file, but I believe this is also configurable. bq. The timeliness feature is nice, but that can be done in a separate SimpleTimer callback. Agreed. Timeliness check is very good and should be considered separate to GC logs. bq. Yeah, rollover is nice. Also does what java prints out have timestamps? Yes. Here are some examples from the logs that HDP creates for the secondary namenode: {noformat} 2016-09-14T09:57:00.476+0000: 54036.340: [GC (Allocation Failure) 2016-09-14T09:57:00.476+0000: 54036.340: [ParNew: 106050K->1130K(118016K), 0.0680896 secs] 125622K->20712K(1035520K), 0.0683270 secs] [Times: user=0.13 sys=0.00, real=0.07 secs] 2016-09-14T12:12:44.975+0000: 62180.839: [GC (Allocation Failure) 2016-09-14T12:12:44.975+0000: 62180.840: [ParNew: 106090K->1157K(118016K), 0.0503738 secs] 125672K->20742K(1035520K), 0.0505503 secs] [Times: user=0.10 sys=0.00, real=0.05 secs] 2016-09-14T14:51:46.898+0000: 71722.762: [GC (Allocation Failure) 2016-09-14T14:51:46.898+0000: 71722.762: [ParNew: 106117K->697K(118016K), 0.0705879 secs] 125702K->20662K(1035520K), 0.0707647 secs] [Times: user=0.14 sys=0.00, real=0.07 secs] {noformat} Another perk is that we can use standard tools for parsing these logs. Here is a nice little article that explains how to read these lines https://plumbr.eu/blog/garbage-collection/understanding-garbage-collection-logs	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3770
16479	ACCUMULO-3769	Huh. The work for YCSB Accumulo I did ended up at https://github.com/brianfrankcooper/YCSB Which looks like it's diverged?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3769
16480	ACCUMULO-3761	Thanks for the patch, [~rweeks]! {code} + kvBufSize += sourceTopKey.getSize() + sourceTopValue.getSize() + 128; {code} This was a curious addition that caught my eye. It looks like you copied this from {{TransformingIterator}} (given [~kturner]'s recommendation on the ML thread). Keith, is this just a constant to account for general object size overhead? Otherwise this looks really great! The test case is nice. One nit is that {{testEncodeAll}} doesn't have the JUnit test annotation, but that's something I can add before commit.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3761
16481	ACCUMULO-3761	[~rweeks] courtesy ping: any chance you can look into this failing test? Does it pass locally for you? This would be nice to add for 1.7.0, but isn't a big issue if it gets pushed to 1.7.1 (since our public API doesn't include iterators presently).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3761
16482	ACCUMULO-3761	Patch applied. Thanks, [~rweeks]. Can I add you to our contributors page w/ TZ and/or employer?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3761
16483	ACCUMULO-3751	Thanks [~medined], you pointed me to an idea. Column families are distinguished because they can be used for locality groups. Beyond that, the other fields are effectively laid out next to each other. Why distinguish between the parts of a Key that are considered indexable, versus the Value that is considered non-indexable? If we laid out in a single byte[] the {{row-colF-colQ-colVis-timestamp-delete-value}}, then we can sort on the value just as if we had a comparison by {{PartialKey.ROW_COLFAM_COLQUAL_COLVIS_TIME_DEL_VAL}}. The use of a column qualifiter is then a place to put information that sorts before column visibility and timestamp are considered. There's nothing preventing us from putting the source/parent field you suggest in the column qualifier position, value position or some other position (as long as we delimit the "fields" in the byte[] or track positions). When I describe Accumulo to friends, I tell them that "everything is a byte[]" that can be interpreted however you choose. This sounds like a greater realization of that motto. As for geographic data, the best data layout schemes for those are use things like Z-order curves and Hilbert curves, which seem hard to do in the current Accumulo. Were you thinking of something that would help with these, or perhaps something more elaborate?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3751
16484	ACCUMULO-3751	Out of curiosity, will you be at Accumulo Summit this week, [~denine99]? Was thinking that we could block out some time at the hackathon to brainstorm in person if you are. If you aren't, we could try to set up a video hangout (since there will at least be _some_ critical mass of devs there in person).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3751
16485	ACCUMULO-3720	You mentioned generating the report and optionally failing the build in the description and your last comment. What did you end up committing? Can you provide some details on how this does (or does not) affect the build process, please?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3720
16486	ACCUMULO-3720	Great, thanks for writing this up. I'm sure I'll be back to reference it. It sems like we're collecting a nice suite of automated tools that can be run at build time to help enforce good code at commit rather than release (findbugs, checkstyle, RAT, apilyzer). Tangential, it may be good to think about a developers section on the user manual to cover these (what do they do, what are failure conditions, how do you fix them, etc).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3720
16487	ACCUMULO-3720	Each plugin has its own documentation to link to (for most of those, the URL is documented in its own pom, as well as searchable in Google), if we do add such documentation. I don't know how much it's worth it to document the build plugins we're running, because the reasons for running them tend to be self-evident (e.g. the reasons we're using them is contained in the description of each), and I can't really think of what documentation would be useful beyond that and a link to the plugins' own documentation. Is there something beyond those things which you think might be valuable to put in said documentation? (BTW, if we were generating/publishing the maven site, a listing of our plugins, with links to their docs would be auto-generated. For example: http://code.revelc.net/zookeeper-maven-plugin/plugins.html)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3720
16488	ACCUMULO-3703	Fixes up the Environment to construct the correct AuthenticationToken. Doesn't address anything in the security module yet -- will deal with that in a separate ticket since it'll be a big rewrite.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3703
16489	ACCUMULO-3700	bq. configuration that maps namespaces to tablet servers What are your thoughts on defining tablet servers? Hostname regex or list? Something else?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3700

#	Requirement	Questions	Answered	Successfully Answered	URL
16490	ACCUMULO-3698	Chritopher, I was thinking about tight integration of Lucene with Accumulo. So that lucene can read and write to and from accumulo tables. I jsut want to use Lucene capabilities in Accumulo So that the over all system would be fast and scalable and does not limited by Lucene. Taking it as an example seems like you are using lucene independently to store index and Accumulo to store Documents. Is it the case?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3698
16491	ACCUMULO-3681	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/24#discussion_r26906707 --- Diff: core/src/main/java/org/apache/accumulo/core/client/impl/ScannerIterator.java --- @@ -92,13 +92,13 @@ public void run() { return; } } catch (IsolationException | ScanTimedOutException | AccumuloException | AccumuloSecurityException | TableDeletedException | TableOfflineException e) { - log.trace(e.getMessage(), e); + log.trace("{}", e.getMessage(), e); --- End diff -- General question: is it worth inserting the formatting string when there isn't any concatenation happening?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3681
16492	ACCUMULO-3665	You had mentioned to me in chat before that you wanted to figure this one out before you got the parent committed. Is this ready to be closed out now or do you still have more debugging to do here? Thanks.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3665
16493	ACCUMULO-3662	To be clear, the Makefile does the extracting from the environment. We could modify the Makefile to print a reasonable error message instead of build_native_library.sh. Under what circumstances have you seen this, though? When a JDK was not installed?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3662
16494	ACCUMULO-3662	I think we can add a quick check to verify {{$JAVA_HOME/bin/javah}} exists in the Makefile, right? That should be sufficient to determine if the JAVA_HOME points to a JDK. After that, we can just assume the JNI headers are in the include directory. We could check for the JNI headers directly, but that may not be as platform-independent as checking for the existence of javah. Do we need to do this check on Darwin also, or just Linux? I mean, can we assume Darwin has a JDK installed, or no?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3662
16495	ACCUMULO-3652	Can you give some details on this? I don't mind working on it if I understand what the goal is.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3652
16496	ACCUMULO-3652	There are several place in the core and shell that set the logging level during runtime which slf4j does not support. Example from o.a.a.core.client.mapred.AbstractInputFormat log.setLevel(logLevel); Does anybody have an idea how to refactor this with using slf4j? I've seen some hacks mentioned online but nothing looks enduring.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3652
16497	ACCUMULO-3652	I'm through all the modules except accumulo-test with the first (i.e. "easy") pass and should complete that tomorrow. Shall I create another issue to complete the hard parts of log4j replacement and/or client supplied logger?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3652
16498	ACCUMULO-3652	Will do. I'm almost done with accumulo-test so I'll wrap it up tomorrow and push the changes. BTW its going to be a huge change set; is a patch okay or should I get some credentials and push the branch to the repo and do a PR?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3652
16499	ACCUMULO-3652	Either patch or PR is fine. Are you familiar with ReviewBoard? If it's large and complex, that may be useful, too: https://reviews.apache.org/	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3652
16500	ACCUMULO-3652	OK, I've got rbt tools installed. How you guys create the diff file to upload to the review board?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3652
16501	ACCUMULO-3652	I was trying to figure out how to mark it "patch available" and used the Fixed status which marked it resolved. How did you mark it "patch available"?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3652
16502	ACCUMULO-3652	I've worked through the first round of the review board issues. Shall I update the patch attached to this issue, remove it all together, or wait till the review is over before updating it?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3652
16503	ACCUMULO-3652	I'm not sure what happened, but once again, the updated patch seems to contain a lot of unrelated changes from a merge. Do you have a branch on GitHub I can take a look at instead? Perhaps I can see if I can figure out how to make ReviewBoard happy?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3652
16504	ACCUMULO-3652	Looks like your patch includes a merge conflict file: {{server/tserver/src/main/java/org/apache/accumulo/tserver/tablet/Tablet.java.orig}} Are these changes necessary? {code} diff --git a/server/tserver/src/main/java/org/apache/accumulo/tserver/TLevel.java b/server/tserver/src/main/java/org/apache/accumulo/tserver/TLevel.java index 1c3e43f..7a15577 100644 --- a/server/tserver/src/main/java/org/apache/accumulo/tserver/TLevel.java +++ b/server/tserver/src/main/java/org/apache/accumulo/tserver/TLevel.java @@ -17,6 +17,7 @@ package org.apache.accumulo.tserver; import org.apache.log4j.Level; +import org.apache.log4j.Priority; import org.slf4j.Logger; public class TLevel extends Level { @@ -25,31 +26,31 @@ public class TLevel extends Level { public final static Level TABLET_HIST = new TLevel(); protected TLevel() { - super(Level.DEBUG_INT + 100, "TABLET_HIST", Level.DEBUG_INT + 100); + super(Priority.DEBUG_INT + 100, "TABLET_HIST", Priority.DEBUG_INT + 100); } static public Level toLevel(int val) { - if (val == Level.DEBUG_INT + 100) + if (val == Priority.DEBUG_INT + 100) return Level.DEBUG; return Level.toLevel(val); } static public void logAtLevel(Logger log, Level level, String msg, Object...objects) { switch(level.toInt()) { - case TLevel.DEBUG_INT: + case Priority.DEBUG_INT: log.debug(msg, objects); break; - case TLevel.ERROR_INT: - case TLevel.FATAL_INT: + case Priority.ERROR_INT: + case Priority.FATAL_INT: log.error(msg, objects); break; - case TLevel.INFO_INT: + case Priority.INFO_INT: log.info(msg, objects); break; - case TLevel.TRACE_INT: + case Level.TRACE_INT: log.trace(msg, objects); break; - case TLevel.WARN_INT: + case Priority.WARN_INT: log.warn(msg, objects); break; } {code} I think you got all of the occurrences of extra braces. For your next patch, can you also please rebase your multiples commits into a single commit? Multiple commits in the same patch files make it hard to review by eye. Thanks Bob, I think this is getting close.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3652

#	Requirement	Questions	Answered	Successfully Answered	URL
16505	ACCUMULO-3652	{code} @@ -139,7 +139,7 @@ public class TCredentialsUpdatingInvocationHandler<I> implements InvocationHandl try { clz = Class.forName(tokenClassName); } catch (ClassNotFoundException e) { - log.debug("Could not create class from token name: " + tokenClassName, e); + log.debug("Could not create class from token name: {} {}", tokenClassName, e); {code} Extra set of braces. {code} @@ -185,7 +185,7 @@ public class TabletStateChangeIterator extends SkippingIterator { boolean shouldBeOnline = onlineTables.contains(tls.extent.getTableId().toString()); if (debug) { - log.debug(tls.extent + " is " + tls.getState(current) + " and should be " + (shouldBeOnline ? "on" : "off") + "line"); + log.debug("{} is {} and should be {} line", tls.extent, tls.getState(current), (shouldBeOnline ? "on" : "off")); {code} Extra space between the third pair of braces and "line" {code} diff --git a/server/base/src/main/java/org/apache/accumulo/server/security/SecurityUtil.java b/server/base/src/main/java/org/apache/accumulo/server/security/SecurityUtil.java index e27a7e7..bf5ba8f 100644 --- a/server/base/src/main/java/org/apache/accumulo/server/security/SecurityUtil.java +++ b/server/base/src/main/java/org/apache/accumulo/server/security/SecurityUtil.java @@ -72,13 +72,13 @@ public class SecurityUtil { .... - log.error("Error logging in user " + principalConfig + " using keytab at " + keyTabPath, io); + log.error("Error logging in user {} using keytab at {} {}", principalConfig, keyTabPath, io); {code} Extra set of braces. {code} diff --git a/server/base/src/main/java/org/apache/accumulo/server/util/VerifyTabletAssignments.java b/server/base/src/main/java/org/apache/accumulo/server/util/VerifyTabletAssignments.java index c0c979b..3775fc8 100644 --- a/server/base/src/main/java/org/apache/accumulo/server/util/VerifyTabletAssignments.java +++ b/server/base/src/main/java/org/apache/accumulo/server/util/VerifyTabletAssignments.java @@ -126,7 +126,7 @@ public class VerifyTabletAssignments { try { checkTabletServer(context, entry, failures); } catch (Exception e) { - log.error("Failure on tablet server '" + entry.getKey() + ".", e); + log.error("Failure on tablet server ' {}. {}", entry.getKey(), e); {code} Extra set of braces {code} @@ -225,7 +226,8 @@ public class GarbageCollectWriteAheadLogs { } catch (FileNotFoundException ex) { // ignored } catch (IOException ex) { - log.error("Unable to delete wal " + path + ": " + ex); + //log.error("Unable to check for the existence of {} {}", swalog, ex); + log.error("Unable to delete wal {}",path , ex); {code} Nit: remove the commented code. {code} @@ -338,7 +340,7 @@ public class GarbageCollectWriteAheadLogs { return true; } } catch (InvalidProtocolBufferException e) { - log.error("Could not deserialize Status protobuf for " + entry.getKey(), e); + log.error("Could not deserialize Status protobuf for {} {}", entry.getKey(), e); {code} Extra set of braces {code} @@ -366,13 +368,13 @@ public class SimpleGarbageCollector extends AccumuloServerContext implements Ifa // atomically in one mutation and extreme care would need to be taken that delete entry was not lost. Instead of doing that, just deal with // volume switching when something needs to be deleted. Since the rest of the code uses suffixes to compare delete entries, there is no danger // of deleting something that should not be deleted. Must not change value of delete variable because thats whats stored in metadata table. - log.debug("Volume replaced " + delete + " -> " + switchedDelete); + log.debug("Volume replaced {} -> ", delete, switchedDelete); {code} Missing a pair of braces {code} @@ -1071,7 +1070,7 @@ public class Master extends AccumuloServerContext implements LiveTServerSet.Lis } catch (TTransportException e) { // ignore: it's probably down } catch (Exception e) { - log.info("error talking to troublesome tablet server ", e); + log.info("error talking to troublesome tablet server {}", e); } badServers.remove(server); tserverSet.remove(server); @@ -	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3652
16506	ACCUMULO-3643	Bump, didn't see any sort of ACK on this. When might you be able to fix these?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3643
16507	ACCUMULO-3633	[~elserj] some thoughts as I read over it. If you make a 2nd revision, would you mind putting it on RB? * init section : iterator config can come from other places than table config * Instantiation section : Could also mention per table classpath. Also could mention that one should not seek source in init() * Iterator design section : List is mentioned, could also mention Tree of iterators. Also could use Tree terminology in deepCopy section. Could mention deepCopy() should be called in init() * explain isolation, data sources do not change until top iterator returns key/value, could move [html isolation iterater documation|https://github.com/apache/accumulo/blob/1.6.2/docs/src/main/resources/isolation.html] to this new section in user manual * in addition to combiner and filter, could mention transforming iterator * next section : does it have to be a cached key value? * could explain why cross row operations are not recommended * hasTop section : java iterators have nothing similar to hasTop Some code like the following showing how tserver will call iterators for a scan may be useful. {code:java} List<KeyValue> batch; Range range = //range from client while(!overSizeLimit(batch)){ source = systemIterator() for(SKVI iter : iterators){ iter.init(source, opts, env) source = iter } //read a batch of data to return to client topIter = iterators.last() topIter.seek(range, ...) while(topIter.hasTop() && !overSizeLimit(batch)){ key = topIter.getTopKey() val = topIter.getTopValue() batch.add(new KeyValue(key, val) if(systemDataSourcesChanged()){ //code does not show isolation case, which will keep using same data sources until a row boundry is hit range = new Range(key, false, range.endKey(), range.endKeyInclusive()); break; } } } //return batch of key values to client {code}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3633
16508	ACCUMULO-3622	It looks like that will erase and rewrite everything under the user node for the root user. Can we quantify what will be destroyed by reseting the perms for system, namespace, etc? I think since we're talking about the root user those all get ignored maybe?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3622
16509	ACCUMULO-3621	Ok, I'm still perplexed, [~ctubbsii]. For STANDALONE, {{testDirList}} indexes ACCUMULO_HOME on the local node. Where were you getting /tmp from? Looks like I might have forgotten to update the comment in the test case. Perhaps that was throwing you? {code} case STANDALONE: dirListDirectory = ((StandaloneAccumuloCluster) getCluster()).getAccumuloHome(); {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3621
16510	ACCUMULO-3621	bq. Where were you getting /tmp from? From 1.6. Apparently 1.7 and later diverged from the old behavior.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3621
16511	ACCUMULO-3617	The same log statement which logs this message logs all the other incoming requests at the debug level also. It's in a "checkPermission" method, so it's trying to provide some auditing information for debugging permissions checks. Should we just move these all to trace or should we remove all of them? Or, should the unloadTablet message be specifically suppressed and the others shown?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3617
16512	ACCUMULO-3617	IMO, is there a purpose to auditing unloadTablet with the understand that the master will repeatedly spam it until the server actually responds to the request? Without more information, I don't think it makes sense to log it at all. If it actually included the extent in question, I could see it at TRACE. Given that it's in a checkPermission block, I think the intent was not to debug the fact that the server saw an unloadMessage, but more that it allowed an unloadMessage which, IMO, is useless.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3617
16513	ACCUMULO-3615	Is this another cance we could benefit from Curator and their PathChildrenCache?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3615

#	Requirement	Questions	Answered	Successfully Answered	URL
16514	ACCUMULO-3602	Left some commits on your most recent commit in that branch. Thanks for sharing, [~echeipesh]! Some bigger concerns that we'll need to talk through: * Resource management on the BatchScanner. Are we sure that it will get closed? I might have just missed it happening elsewhere (via ScannerBase). * All of the MapReduce classes/interfaces/etc under `org.apache.accumulo.core.client` are (painfully) considered to be in the public API. Since we're following semver, it's unlikely that we'll be able to make this change 1.6.3 as you've currently targeted. We'd have to go for 1.7.0. The mapred/mapreduce classes often fall into this gray area where we don't *really* want them all to be public API, but they got grandfathered in. We'll have to keep this in mind. * The duplication between RangeInputSplit and MultiRangeInputSplit is painful (I think you called this out to begin with). Maybe there's some more consolidation that can be done with RangeInputSplit? I'm worried about it because we've had a couple of bugs come out of RangeInputSplit due to the duplicative code and insufficient testing. As far as testing, you can try to take a look at AccumuloInputFormatIT.java. That has a good example of running a "mapreduce" job over a MiniAccumuloCluster. You could also check out the MiniMRYarnCluster class provided by Hadoop/YARN itself which should hopefully give you a more realistic test environment (as opposed to using the local MR runner). ExamplesIT.java also has some examples of running MR jobs against Accumulo which you could look at. In short, the more you can unit test, the better. A test or two to ensure high-level functionality would be good IMO.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16515	ACCUMULO-3602	Thank you for the review, I've made those adjustments. Just to spell out the issue, we can't change the RangeInputSplit interface because it would break semver, but by the same token adding MultiRangeInputSplit to the same package also breaks semver as it's not available in 1.6.2. This feels a little weird since the InputSplits are only ever used by the InputFormat, as I understand it. Can you see any way that this functionality making it into 1.6.3? If it's really not possible then the best approach is to have a universal InputSplit by combining functionality of the two and create appropriate reader based on split values and roll that into 1.7.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16516	ACCUMULO-3602	bq. Just to spell out the issue, we can't change the RangeInputSplit interface because it would break semver, but by the same token adding MultiRangeInputSplit to the same package also breaks semver as it's not available in 1.6.2 You got it. bq. This feels a little weird since the InputSplits are only ever used by the InputFormat, as I understand it. I hear you, but the concern for RangeInputSplit is that someone else could have extended it and your changes might break them. As for the addition of MultiRangeInputSplit, this is just strictly prohibited by semver as 1.6.3 should only contain bug-fixes. bq. Can you see any way that this functionality making it into 1.6.3? Sadly, I don't think so. Perhaps it's some consolation that 1.7 should happen in the next month or two? We've kind of been twiddling our thumbs on the matter. IMO, we should start pushing it out. bq. a universal InputSplit by combining functionality of the two and create appropriate reader based on split values and roll that into 1.7. Interesting, have you given thought into the logic that would control when either is used? I remember seeing some experiments where replacing the Scanner with a BatchScanner had very positive outcomes (obv you're one of them), but I wonder if it's always true? We should try be cautious in changing how the InputFormat works (even from 1.6 to 1.7).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16517	ACCUMULO-3602	Alright, 1.7 target it is then. I'll rebase on master. I'll code up a merged InputSplit to see what it would looks like. My intuition would be to be cautious and have a specific flag for the use of BatchScanner in cases when split has more than one range. There is still no way to support offline and client side scans using BatchScanner, so it doesn't cover all cases. The logic would be similar to what happens now except the branch will be on split property rather then split class. Should it be the default approach when possible? I'm not sure, probably yes. Locality would be tied to a tablet, which is hosted on one machine, you still can reap the benefits of several IO threads. The effort on querying the TabletLocator is spent in either case, so why not use it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16518	ACCUMULO-3602	bq. Alright, 1.7 target it is then. I'll rebase on master. Cool. Thank you, and again, sorry for the pain semver is causing. bq. Should it be the default approach when possible? I'm not sure, probably yes. Locality would be tied to a tablet, which is hosted on one machine, you still can reap the benefits of several IO threads. The effort on querying the TabletLocator is spent in either case, so why not use it. As long as the Split generation is still binning the ranges together on locality, I think this would be fine for most cases. There's still the possibility that tablets move the mappers read them, but that's nothing new. One case I thought of is that a Mapper could have special logic based on the fact that it will receive sorted Key-Value pairs which a BatchScanner-by-default would break. I'm not sure how common this is nor if it's a big concern. My cautious nature would be to say that it's safest to have 1.7 introduce configuration that allow use of BatchScanners and then 2.0 could switch that to the default? A nice enum could help make a setter which ensure API compat between 1.7 and 2.0 and let the default configuration change underneath. Haven't thought it through completely, just a first reaction. [~kturner] is really good at this stuff. Maybe he can weigh :)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16519	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27628662 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/AbstractInputFormat.java --- @@ -625,32 +659,52 @@ public float getProgress() throws IOException { } for (Map.Entry<KeyExtent,List<Range>> extentRanges : tserverBin.getValue().entrySet()) { Range ke = extentRanges.getKey().toDataRange(); - for (Range r : extentRanges.getValue()) { - if (autoAdjust) { - // divide ranges into smaller ranges, based on the tablets - RangeInputSplit split = new RangeInputSplit(tableName, tableId, ke.clip(r), new String[] {location}); - - split.setOffline(tableConfig.isOfflineScan()); - split.setIsolatedScan(tableConfig.shouldUseIsolatedScanners()); - split.setUsesLocalIterators(tableConfig.shouldUseLocalIterators()); - split.setMockInstance(mockInstance); - split.setFetchedColumns(tableConfig.getFetchedColumns()); - split.setPrincipal(principal); - split.setToken(token); - split.setInstanceName(instance.getInstanceName()); - split.setZooKeepers(instance.getZooKeepers()); - split.setAuths(auths); - split.setIterators(tableConfig.getIterators()); - split.setLogLevel(logLevel); - - splits.add(split); - } else { - // don't divide ranges - ArrayList<String> locations = splitsToAdd.get(r); - if (locations == null) - locations = new ArrayList<String>(1); - locations.add(location); - splitsToAdd.put(r, locations); + if (autoAdjust) { + // group ranges by tablet to be read by a BatchScanner + ArrayList<Range> clippedRanges = new ArrayList<Range>(); + for(Range r: extentRanges.getValue()) + clippedRanges.add(ke.clip(r)); + org.apache.accumulo.core.client.mapred.BatchInputSplit split = new org.apache.accumulo.core.client.mapred.BatchInputSplit(tableName, tableId, clippedRanges, new String[] {location}); + split.setFetchedColumns(tableConfig.getFetchedColumns()); --- End diff -- Can we set the common attributes on the splits in one place instead of both sides of the conditional? Does the type hierarchy make that simple?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602

#	Requirement	Questions	Answered	Successfully Answered	URL
16520	ACCUMULO-3602	core/src/main/java/org/apache/accumulo/core/client/mapreduce/AccumuloInputSplit.java --- @@ -0,0 +1,425 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.core.client.mapreduce; +import static java.nio.charset.StandardCharsets.UTF_8; +import java.io.DataInput; +import java.io.DataOutput; +import java.io.IOException; +import java.math.BigInteger; +import java.util.ArrayList; +import java.util.Arrays; +import java.util.Collection; +import java.util.HashSet; +import java.util.List; +import java.util.Set; +import org.apache.accumulo.core.client.ClientConfiguration; +import org.apache.accumulo.core.client.Instance; +import org.apache.accumulo.core.client.IteratorSetting; +import org.apache.accumulo.core.client.ZooKeeperInstance; +import org.apache.accumulo.core.client.mapreduce.lib.impl.ConfiguratorBase.TokenSource; +import org.apache.accumulo.core.client.mapreduce.lib.impl.InputConfigurator; +import org.apache.accumulo.core.client.mock.MockInstance; +import org.apache.accumulo.core.client.security.tokens.AuthenticationToken; +import org.apache.accumulo.core.client.security.tokens.AuthenticationToken.AuthenticationTokenSerializer; +import org.apache.accumulo.core.data.ByteSequence; +import org.apache.accumulo.core.data.Range; +import org.apache.accumulo.core.security.Authorizations; +import org.apache.accumulo.core.util.Base64; +import org.apache.accumulo.core.util.Pair; +import org.apache.accumulo.core.data.Key; +import org.apache.hadoop.io.Text; +import org.apache.hadoop.io.Writable; +import org.apache.hadoop.mapreduce.InputSplit; +import org.apache.log4j.Level; + +/** + * Abstracts over configurations common to all InputSplits. Specifically it leaves out methods + * related to number of ranges and locations per InputSplit as those vary by implementation. + * + * @see RangeInputSplit + * @see BatchInputSplit + */ +public abstract class AccumuloInputSplit extends InputSplit implements Writable { + protected String[] locations; + protected String tableId, tableName, instanceName, zooKeepers, principal; + protected TokenSource tokenSource; + protected String tokenFile; + protected AuthenticationToken token; + protected Boolean mockInstance; + protected Authorizations auths; + protected Set<Pair<Text,Text>> fetchedColumns; + protected List<IteratorSetting> iterators; + protected Level level; + + public abstract float getProgress(Key currentKey); + + public AccumuloInputSplit() { + locations = new String[0]; + tableName = ""; + tableId = ""; + } + + public AccumuloInputSplit(AccumuloInputSplit split) throws IOException { + this.setLocations(split.getLocations()); + this.setTableName(split.getTableName()); + this.setTableId(split.getTableId()); + } + + protected AccumuloInputSplit(String table, String tableId, String[] locations) { + setLocations(locations); + this.tableName = table; + this.tableId = tableId; + } + + private static byte[] extractBytes(ByteSequence seq, int numBytes) { + byte[] bytes = new byte[numBytes + 1]; + bytes[0] = 0; + for (int i = 0; i < numBytes; i++) { + if (i >= seq.length()) + bytes[i + 1] = 0; + else + bytes[i + 1] = seq.byteAt(i); + } + return bytes; + } + + public static float	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16521	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27629593 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapreduce/InputFormatBase.java --- @@ -296,6 +296,65 @@ protected static boolean isOfflineScan(JobContext context) { } /** + * Controls the use of the {@link org.apache.accumulo.core.client.BatchScanner} in this job. + * Using this feature will group ranges by their source tablet per InputSplit and use BatchScanner to read them. + * + * <p> + * By default, this feature is disabled. + * + * @param job + * the Hadoop job instance to be configured + * @param enableFeature + * the feature is enabled if true, disabled otherwise + * @since 1.7.0 + */ + public static void setBatchScan(Job job, boolean enableFeature) { + InputConfigurator.setBatchScan(CLASS, job.getConfiguration(), enableFeature); + } + + /** + * Determines whether a configuration has the {@link org.apache.accumulo.core.client.BatchScanner} feature enabled. + * + * @param context + * the Hadoop context for the configured job + * @since 1.7.0 + * @see #setBatchScan(Job, boolean) + */ + protected static boolean isBatchScan(JobContext context) { --- End diff -- Why JobContext here? And why not public?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16522	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27629761 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapreduce/RangeInputSplit.java --- @@ -179,81 +93,13 @@ public void readFields(DataInput in) throws IOException { if (in.readBoolean()) { localIterators = in.readBoolean(); } - if (in.readBoolean()) { - mockInstance = in.readBoolean(); - } - if (in.readBoolean()) { - int numColumns = in.readInt(); - List<String> columns = new ArrayList<String>(numColumns); - for (int i = 0; i < numColumns; i++) { - columns.add(in.readUTF()); - } - fetchedColumns = InputConfigurator.deserializeFetchedColumns(columns); - } - if (in.readBoolean()) { - String strAuths = in.readUTF(); - auths = new Authorizations(strAuths.getBytes(UTF_8)); - } - if (in.readBoolean()) { - principal = in.readUTF(); - } - if (in.readBoolean()) { - int ordinal = in.readInt(); - this.tokenSource = TokenSource.values()[ordinal]; - switch (this.tokenSource) { - case INLINE: - String tokenClass = in.readUTF(); - byte[] base64TokenBytes = in.readUTF().getBytes(UTF_8); - byte[] tokenBytes = Base64.decodeBase64(base64TokenBytes); - this.token = AuthenticationTokenSerializer.deserialize(tokenClass, tokenBytes); - break; - case FILE: - this.tokenFile = in.readUTF(); - break; - default: - throw new IOException("Cannot parse unknown TokenSource ordinal"); - } - } - if (in.readBoolean()) { - instanceName = in.readUTF(); - } - if (in.readBoolean()) { - zooKeepers = in.readUTF(); - } - if (in.readBoolean()) { - int numIterators = in.readInt(); - iterators = new ArrayList<IteratorSetting>(numIterators); - for (int i = 0; i < numIterators; i++) { - iterators.add(new IteratorSetting(in)); - } - } - if (in.readBoolean()) { - level = Level.toLevel(in.readInt()); - } - } - - @Override - public void write(DataOutput out) throws IOException { --- End diff -- Need to serialize offline here.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16523	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27629947 --- Diff: test/src/test/java/org/apache/accumulo/test/functional/AccumuloInputFormatIT.java --- @@ -146,6 +147,23 @@ public void testGetSplits() throws IOException, AccumuloSecurityException, Accum AccumuloInputFormat.setAutoAdjustRanges(job, false); splits = inputFormat.getSplits(job); assertEquals(ranges.size(), splits.size()); + + //BatchScan not available for offline scans + AccumuloInputFormat.setBatchScan(job, true); --- End diff -- Maybe expand this same test for isolated scans local iterators too?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16524	ACCUMULO-3602	[~echeipesh], do you think you'll have time to update your changes this week? We'd like to close down changes to 1.7.0.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602

#	Requirement	Questions	Answered	Successfully Answered	URL
16525	ACCUMULO-3602	Github user echeipesh commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27826065 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapreduce/AbstractInputFormat.java --- @@ -426,6 +442,7 @@ protected static ClientConfiguration getClientConfiguration(JobContext context) * the table name for which the scanner is configured * @since 1.6.0 */ + @Deprecated protected abstract void setupIterators(TaskAttemptContext context, Scanner scanner, String tableName, RangeInputSplit split); --- End diff -- Could this be implemented in terms of new setupIterators without breaking semver? Would save some ugliness.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16526	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27826429 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapreduce/AbstractInputFormat.java --- @@ -426,6 +442,7 @@ protected static ClientConfiguration getClientConfiguration(JobContext context) * the table name for which the scanner is configured * @since 1.6.0 */ + @Deprecated protected abstract void setupIterators(TaskAttemptContext context, Scanner scanner, String tableName, RangeInputSplit split); --- End diff -- I think we could have source compatibility with a switch from Scanner to ScannerBase, but I can't think of a way to handle AccumuloInputSplit instead of only RangeInputSplits. It should be straightforward to just call a new `setupIterator` method from the existing (now deprecated) method, no? Are you just trying to reduce the amount of code in here?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16527	ACCUMULO-3602	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27914201 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/AbstractInputFormat.java --- @@ -629,32 +676,37 @@ public float getProgress() throws IOException { } for (Map.Entry<KeyExtent,List<Range>> extentRanges : tserverBin.getValue().entrySet()) { Range ke = extentRanges.getKey().toDataRange(); - for (Range r : extentRanges.getValue()) { - if (autoAdjust) { - // divide ranges into smaller ranges, based on the tablets - RangeInputSplit split = new RangeInputSplit(tableName, tableId, ke.clip(r), new String[] {location}); - - split.setOffline(tableConfig.isOfflineScan()); - split.setIsolatedScan(tableConfig.shouldUseIsolatedScanners()); - split.setUsesLocalIterators(tableConfig.shouldUseLocalIterators()); - split.setMockInstance(mockInstance); - split.setFetchedColumns(tableConfig.getFetchedColumns()); - split.setPrincipal(principal); - split.setToken(token); - split.setInstanceName(instance.getInstanceName()); - split.setZooKeepers(instance.getZooKeepers()); - split.setAuths(auths); - split.setIterators(tableConfig.getIterators()); - split.setLogLevel(logLevel); - - splits.add(split); - } else { - // don't divide ranges - ArrayList<String> locations = splitsToAdd.get(r); - if (locations == null) - locations = new ArrayList<String>(1); - locations.add(location); - splitsToAdd.put(r, locations); + if (batchScan && supportBatchScan) { --- End diff -- Why check `supportBatchScan` here?. There is a check earlier that throws an exception `if (batchScan && !supportBatchScan)`.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16528	ACCUMULO-3602	Github user echeipesh commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27923743 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapreduce/lib/impl/InputConfigurator.java --- @@ -517,6 +517,74 @@ public static Boolean isOfflineScan(Class<?> implementingClass, Configuration co } /** + * Controls the use of the {@link BatchScanner} in this job. + * Using this feature will group ranges by their source tablet per InputSplit and use BatchScanner to read them. + * <p> + * By default, this feature is disabled. + * + * @param implementingClass + * the class whose name will be used as a prefix for the property configuration key + * @param conf + * the Hadoop configuration object to configure + * @param enableFeature + * the feature is enabled if true, disabled otherwise + * @since 1.7.0 + */ + public static void setBatchScan(Class<?> implementingClass, Configuration conf, boolean enableFeature) { + conf.setBoolean(enumToConfKey(implementingClass, Features.BATCH_SCANNER), enableFeature); + } + + /** + * Determines whether a configuration has the BatchScanner feature enabled. + * + * @param implementingClass + * the class whose name will be used as a prefix for the property configuration key + * @param conf + * the Hadoop configuration object to configure + * @return true if the feature is enabled, false otherwise + * @since 1.7.0 + * @see #setBatchScan(Class, Configuration, boolean) + */ + public static Boolean isBatchScan(Class<?> implementingClass, Configuration conf) { + return conf.getBoolean(enumToConfKey(implementingClass, Features.BATCH_SCANNER), false); + } + + /** + * Controls the use of the {@link BatchScanner} in this job. Using this feature will group ranges by their source tablet per InputSplit and use BatchScanner to read them. + * <p> + * By default, this feature is disabled. + * + * @param implementingClass + * the class whose name will be used as a prefix for the property configuration key + * @param conf + * the Hadoop configuration object to configure + * @param count + * number of threads to use with BatchScanner + * @since 1.7.0 + */ + public static void setBatchScanThreads(Class<?> implementingClass, Configuration conf, int count) { --- End diff -- In the current use case it makes the option unnecessary. Should it still remain for future refactors where `BatchInputSplit` gets initialized with ranges that span tablets? In theory each split could span tservers as well, but it's hard to imagine a reason for that. Perhaps reduce default to one thread for now?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602

#	Requirement	Questions	Answered	Successfully Answered	URL
16529	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27926355 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/AbstractInputFormat.java --- @@ -629,32 +676,37 @@ public float getProgress() throws IOException { } for (Map.Entry<KeyExtent,List<Range>> extentRanges : tserverBin.getValue().entrySet()) { Range ke = extentRanges.getKey().toDataRange(); - for (Range r : extentRanges.getValue()) { - if (autoAdjust) { - // divide ranges into smaller ranges, based on the tablets - RangeInputSplit split = new RangeInputSplit(tableName, tableId, ke.clip(r), new String[] {location}); - - split.setOffline(tableConfig.isOfflineScan()); - split.setIsolatedScan(tableConfig.shouldUseIsolatedScanners()); - split.setUsesLocalIterators(tableConfig.shouldUseLocalIterators()); - split.setMockInstance(mockInstance); - split.setFetchedColumns(tableConfig.getFetchedColumns()); - split.setPrincipal(principal); - split.setToken(token); - split.setInstanceName(instance.getInstanceName()); - split.setZooKeepers(instance.getZooKeepers()); - split.setAuths(auths); - split.setIterators(tableConfig.getIterators()); - split.setLogLevel(logLevel); - - splits.add(split); - } else { - // don't divide ranges - ArrayList<String> locations = splitsToAdd.get(r); - if (locations == null) - locations = new ArrayList<String>(1); - locations.add(location); - splitsToAdd.put(r, locations); + if (batchScan) { + // group ranges by tablet to be read by a BatchScanner + ArrayList<Range> clippedRanges = new ArrayList<Range>(); + for(Range r: extentRanges.getValue()) --- End diff -- Might this need to adhere to autoAdjust for correctness? I think it might be as simple as not clipping the ranges when autoAdjust is false.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16530	ACCUMULO-3602	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27973658 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/InputFormatBase.java --- @@ -315,7 +317,11 @@ protected static TabletLocator getTabletLocator(JobConf job) throws TableNotFoun protected abstract static class RecordReaderBase<K,V> extends AbstractRecordReader<K,V> { @Override - protected void setupIterators(JobConf job, Scanner scanner, String tableName, org.apache.accumulo.core.client.mapred.RangeInputSplit split) { + protected void setupIterators(JobConf job, ScannerBase scanner, String tableName, AccumuloInputSplit split) { --- End diff -- shouldn't the method being deleted be deprecated?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16531	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27974024 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/InputFormatBase.java --- @@ -315,7 +317,11 @@ protected static TabletLocator getTabletLocator(JobConf job) throws TableNotFoun protected abstract static class RecordReaderBase<K,V> extends AbstractRecordReader<K,V> { @Override - protected void setupIterators(JobConf job, Scanner scanner, String tableName, org.apache.accumulo.core.client.mapred.RangeInputSplit split) { + protected void setupIterators(JobConf job, ScannerBase scanner, String tableName, AccumuloInputSplit split) { --- End diff -- There's a since tag on AbstractRecordReader. Isn't that sufficient?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16532	ACCUMULO-3602	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27974211 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapreduce/InputFormatBase.java --- @@ -313,11 +374,34 @@ protected static TabletLocator getTabletLocator(JobContext context) throws Table protected abstract static class RecordReaderBase<K,V> extends AbstractRecordReader<K,V> { + /** + * Apply the configured iterators from the configuration to the scanner for the specified table name + * @param context the Hadoop context for the configured job + * @param scanner the scanner to configure + * @since 1.7.0 + */ @Override - protected void setupIterators(JobContext context, Scanner scanner, String tableName, org.apache.accumulo.core.client.mapreduce.RangeInputSplit split) { + protected void setupIterators(TaskAttemptContext context, Scanner scanner, String tableName, org.apache.accumulo.core.client.mapreduce.RangeInputSplit split) { --- End diff -- should this be deprecated instead of deleted?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16533	ACCUMULO-3602	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27976100 --- Diff: test/src/test/java/org/apache/accumulo/test/functional/AccumuloInputFormatIT.java --- @@ -146,6 +146,41 @@ public void testGetSplits() throws IOException, AccumuloSecurityException, Accum AccumuloInputFormat.setAutoAdjustRanges(job, false); splits = inputFormat.getSplits(job); assertEquals(ranges.size(), splits.size()); + + //BatchScan not available for offline scans + AccumuloInputFormat.setBatchScan(job, true); + + AccumuloInputFormat.setOfflineTableScan(job, true); + try { + inputFormat.getSplits(job); + fail("An exception should have been thrown"); + } catch (IOException e) {} + + AccumuloInputFormat.setOfflineTableScan(job, false); + + //BatchScan not available with isolated iterators + AccumuloInputFormat.setScanIsolation(job, true); + try { + inputFormat.getSplits(job); + fail("An exception should have been thrown"); + } catch (IOException e) {} --- End diff -- couldn't this fail because `setOfflineTableScan()` was called earlier? Therefore if `setScanIsolation()` does not cause a failure, the test may not detect this.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16534	ACCUMULO-3602	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r27979336 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/InputFormatBase.java --- @@ -315,7 +317,11 @@ protected static TabletLocator getTabletLocator(JobConf job) throws TableNotFoun protected abstract static class RecordReaderBase<K,V> extends AbstractRecordReader<K,V> { @Override - protected void setupIterators(JobConf job, Scanner scanner, String tableName, org.apache.accumulo.core.client.mapred.RangeInputSplit split) { + protected void setupIterators(JobConf job, ScannerBase scanner, String tableName, AccumuloInputSplit split) { --- End diff -- > There's a since tag on AbstractRecordReader. Isn't that sufficient? oic, yeah thats good.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16535	ACCUMULO-3602	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/25#issuecomment-91035731 > It seems like this change could encourage users to pass many ranges as configuration for the map reduce job. This could cause memory exhaustion for the job tracker. Agreed, but is this a new issue? If anything, it this these changes would reduce the total number of objects in memory (as a split could contain multiple ranges whereas previously it was 1:1). I remember hearing about solutions that use HDFS as a mechanism for supporting large numbers of ranges instead of the Configuration. Either way, I think the current changes are a good addition but just not in scope to address the bigger issue you brought up.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602

#	Requirement	Questions	Answered	Successfully Answered	URL
16536	ACCUMULO-3602	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/25#issuecomment-91038638 > Agreed, but is this a new issue? I think the intent of this PR is new. I think the intent is to efficiently support a map reduce job that reads from many small ranges. If we are going to do that, then lets do it in a way that scales better. > However, I still think that for many things, it's probably better to simply use an iterator with some filter criteria. This approach would work well for existing input format. The only drawback to this approach is it requires adding custom iterators to tservers classpaths. The batch scanner approach allows efficient reads of many small ranges from a tablet w/o custom iterators. Just need a custom function in the map reduce job to generate the ranges.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16537	ACCUMULO-3602	Github user ctubbsii commented on the pull request: https://github.com/apache/accumulo/pull/25#issuecomment-91038876 > Agreed, but is this a new issue? I wouldn't think so. I think it highlights a better implementation to satisfy the underlying use case than the one currently being proposed in this PR.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16538	ACCUMULO-3602	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/25#issuecomment-91041003 > > Agreed, but is this a new issue? > I think the intent of this PR is new. I think the intent is to efficiently support a map reduce job that reads from many small ranges. If we are going to do that, then lets do it in a way that scales better. Re-reading the issue description, yes, you are correct. However, what Eugene has already done is already an improvement over what currently exists, IMO. Assuming that the majority of Ranges that his Spark workflow creates are non-overlapping (a hunch given the subject matter -- very small boxes of interest), this still reduces the total memory use of splits and Ranges. The worst case would be numerous very large ranges. We're never going to have the ideal solution the first time around -- I just don't want this to stall because we come up with a different solution to an already solved problem. I'm happy to continue the discussion for dealing with large numbers of ranges efficiently, I'm just not convinced this is the place to do it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16539	ACCUMULO-3602	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/25#issuecomment-91049205 > If its not discussed now, then it will likely not be later. 1.7.0 will be released, and we will be supporting these API changes for a long time How do these changes preclude you from later adding in your proposed split generation code?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16540	ACCUMULO-3602	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/25#issuecomment-91049787 > How do these changes preclude you from later adding in your proposed split generation code? They do not. I am just thinking about maintenance over time. If the approach of passing a function is a better solution to the problem at hand, then only we only implement it and have less code to maintain.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16541	ACCUMULO-3602	Github user echeipesh commented on the pull request: https://github.com/apache/accumulo/pull/25#issuecomment-91641526 @keith-turner I don't think the range decomposition interface approach would handle my case without some kind of intermediate storage level, as I am reading it anyway. The Range -> List<Range> interface does not have enough parameters to determine the query plan. Also, if the query is user generated then a named class can not be present to encapsulate the parameters. From this standpoint it seems there are two options - Use some presence that is that would be populated by the query planner - Perhaps the Configuration object would still need a value to find the correct location - Expand the interface to to include some kind of `QueryConfiguration` interface that could be implemented by the user and would be passed in. The second approach sounds more interesting but it's hard to judge it since I don't think I have that problem. To clarify, what is the underlying concern, that the Configuration object gets too large and the (de)serialization becomes too expensive or not even possible? My particular use case is reflection on the fact that small ranges introduce too much overhead when they translate to mappers (in MR) and partitions (in spark). So it would be nice if there is some kind of minimum batch, tablet maps nicely to that. The only other addition to this problem, which is not explored in this PR, is to divide sequence of ranges per tablet server into "balanced", user configurable, number, of groups. Balanced in this case could mean something like: approximately equal number of possible row results. The idea would to be to match "x mappers per host" design. Sounds like it would be great, but I haven't tested anything like that yet, could be another PR if it works out?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16542	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r28554520 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/AbstractInputFormat.java --- @@ -629,32 +676,37 @@ public float getProgress() throws IOException { } for (Map.Entry<KeyExtent,List<Range>> extentRanges : tserverBin.getValue().entrySet()) { Range ke = extentRanges.getKey().toDataRange(); - for (Range r : extentRanges.getValue()) { - if (autoAdjust) { - // divide ranges into smaller ranges, based on the tablets - RangeInputSplit split = new RangeInputSplit(tableName, tableId, ke.clip(r), new String[] {location}); - - split.setOffline(tableConfig.isOfflineScan()); - split.setIsolatedScan(tableConfig.shouldUseIsolatedScanners()); - split.setUsesLocalIterators(tableConfig.shouldUseLocalIterators()); - split.setMockInstance(mockInstance); - split.setFetchedColumns(tableConfig.getFetchedColumns()); - split.setPrincipal(principal); - split.setToken(token); - split.setInstanceName(instance.getInstanceName()); - split.setZooKeepers(instance.getZooKeepers()); - split.setAuths(auths); - split.setIterators(tableConfig.getIterators()); - split.setLogLevel(logLevel); - - splits.add(split); - } else { - // don't divide ranges - ArrayList<String> locations = splitsToAdd.get(r); - if (locations == null) - locations = new ArrayList<String>(1); - locations.add(location); - splitsToAdd.put(r, locations); + if (batchScan) { + // group ranges by tablet to be read by a BatchScanner + ArrayList<Range> clippedRanges = new ArrayList<Range>(); + for(Range r: extentRanges.getValue()) --- End diff -- I think your intent is fine, it just sounds like we don't want to support `autoAdjust=true` when `batchScan=true`. I think that case is just nonsense and we could catch it in the earlier sanity check where you do the offline/isolate scan or localIterators check. Does that make sense?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16543	ACCUMULO-3602	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r28554794 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/InputFormatBase.java --- @@ -315,19 +352,8 @@ protected static TabletLocator getTabletLocator(JobConf job) throws TableNotFoun protected abstract static class RecordReaderBase<K,V> extends AbstractRecordReader<K,V> { @Override - protected void setupIterators(JobConf job, Scanner scanner, String tableName, org.apache.accumulo.core.client.mapred.RangeInputSplit split) { - List<IteratorSetting> iterators = null; - - if (null == split) { - iterators = getIterators(job); - } else { - iterators = split.getIterators(); - if (null == iterators) { - iterators = getIterators(job); - } - } - - setupIterators(iterators, scanner); + protected List<IteratorSetting> jobIterators(JobConf job, String tableName) { + return getIterators(job); --- End diff -- I'm not seeing where this is implemented. Am I just blind?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602

#	Requirement	Questions	Answered	Successfully Answered	URL
16544	ACCUMULO-3602	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r28560936 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/AbstractInputFormat.java --- @@ -629,32 +676,37 @@ public float getProgress() throws IOException { } for (Map.Entry<KeyExtent,List<Range>> extentRanges : tserverBin.getValue().entrySet()) { Range ke = extentRanges.getKey().toDataRange(); - for (Range r : extentRanges.getValue()) { - if (autoAdjust) { - // divide ranges into smaller ranges, based on the tablets - RangeInputSplit split = new RangeInputSplit(tableName, tableId, ke.clip(r), new String[] {location}); - - split.setOffline(tableConfig.isOfflineScan()); - split.setIsolatedScan(tableConfig.shouldUseIsolatedScanners()); - split.setUsesLocalIterators(tableConfig.shouldUseLocalIterators()); - split.setMockInstance(mockInstance); - split.setFetchedColumns(tableConfig.getFetchedColumns()); - split.setPrincipal(principal); - split.setToken(token); - split.setInstanceName(instance.getInstanceName()); - split.setZooKeepers(instance.getZooKeepers()); - split.setAuths(auths); - split.setIterators(tableConfig.getIterators()); - split.setLogLevel(logLevel); - - splits.add(split); - } else { - // don't divide ranges - ArrayList<String> locations = splitsToAdd.get(r); - if (locations == null) - locations = new ArrayList<String>(1); - locations.add(location); - splitsToAdd.put(r, locations); + if (autoAdjust) { + // group ranges by tablet to be read by a BatchScanner + ArrayList<Range> clippedRanges = new ArrayList<Range>(); + for(Range r: extentRanges.getValue()) --- End diff -- Should anything need to be added the `setBatchScan(...)` javadoc regarding `autoAdjust`?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16545	ACCUMULO-3602	Github user echeipesh commented on a diff in the pull request: https://github.com/apache/accumulo/pull/25#discussion_r28597474 --- Diff: core/src/main/java/org/apache/accumulo/core/client/mapred/AbstractInputFormat.java --- @@ -629,32 +676,37 @@ public float getProgress() throws IOException { } for (Map.Entry<KeyExtent,List<Range>> extentRanges : tserverBin.getValue().entrySet()) { Range ke = extentRanges.getKey().toDataRange(); - for (Range r : extentRanges.getValue()) { - if (autoAdjust) { - // divide ranges into smaller ranges, based on the tablets - RangeInputSplit split = new RangeInputSplit(tableName, tableId, ke.clip(r), new String[] {location}); - - split.setOffline(tableConfig.isOfflineScan()); - split.setIsolatedScan(tableConfig.shouldUseIsolatedScanners()); - split.setUsesLocalIterators(tableConfig.shouldUseLocalIterators()); - split.setMockInstance(mockInstance); - split.setFetchedColumns(tableConfig.getFetchedColumns()); - split.setPrincipal(principal); - split.setToken(token); - split.setInstanceName(instance.getInstanceName()); - split.setZooKeepers(instance.getZooKeepers()); - split.setAuths(auths); - split.setIterators(tableConfig.getIterators()); - split.setLogLevel(logLevel); - - splits.add(split); - } else { - // don't divide ranges - ArrayList<String> locations = splitsToAdd.get(r); - if (locations == null) - locations = new ArrayList<String>(1); - locations.add(location); - splitsToAdd.put(r, locations); + if (autoAdjust) { + // group ranges by tablet to be read by a BatchScanner + ArrayList<Range> clippedRanges = new ArrayList<Range>(); + for(Range r: extentRanges.getValue()) --- End diff -- Actually Auto-adjust is on by default, which would cause an exception as soon as batch scan is turned on. Should it be downgraded to a warning, since interaction between auto-adjust and batch scan is a "soft" change in the behavior?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16546	ACCUMULO-3602	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/25#issuecomment-94060780 Actually, compat is still broken with that setupIterators method. Do you want to fix that @echeipesh or should I just do when I apply the rest of your changes?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3602
16547	ACCUMULO-3599	I got to a point where all of the ITs are passing with a normal MAC and a MAC which is using Kerberos via Hadoop's MiniKdc. There were a few tests (those which tried to run MapReduce) which I just ignore when we're running with Kerberos on as I haven't figured out how I can actually get these working (do I just need to construct an inline Configuration with the necessary properties?). The short of it is that new ITs which only perform some task as the "root" user shouldn't have changed. Use the {{getConnector}} method and things should be fine -- the test harness will do the work for you depending on the environment. If the test uses multiple users, doesn't anything with user mgmt (e.g. {{createLocalUser}}), then things get a little different. Still a TBD to update the docs. A secondary goal of the above was to make sure that the changes made to support MAC w/ or w/o Kerberos also support a standalone cluster with Kerberos. I've already found an issue in the first test class I ran, so I anticipate more changes to come, but I wanted to get this in while I know it's in a good state. As always, if you notice new test failures, feel free to open a new issue and assign it to me.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3599
16548	ACCUMULO-3594	Can you clarify what you mean by "message classes"?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3594
16549	ACCUMULO-3594	bq. Alternatively, we could rely on Avro's generic serialization implementation (but it's a little slower). I'm not sure what this means. Can I lazy-internet you and ask for a "why avro" overview/blog/paper?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3594
16550	ACCUMULO-3592	What makes you say this? I'm suggesting we do this now so that it doesn't become a pain point later on. If we introduce the construct to move away from writables, it should allow the same support we have now, but guarantee a bit more backwards compatibility.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3592
16551	ACCUMULO-3591	Specifically the rolling-up of iterators and their options into a single ZNode? If we wanted to keep ZK readable (which I know has been stated as a big concern by others), we could use the JSON protocol for Thrift or protobuf-json for protocol buffers. Not directly related to your comment, just something related I wanted to mention before I forgot.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3591
16552	ACCUMULO-3591	Just for completeness, Avro also has a JSON serialization mode. Is there a reason folks prefer human-readable ZK instead of providing tools that can read a binary serialization?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3591
16553	ACCUMULO-3590	Is it possible to do a manual rolling upgrade from 1.6.2 to 1.72? Let say bring up 1.7 tablet servers then shutdown the 1.6? Or is data serialization currently a problem across the 2 versions?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3590
16554	ACCUMULO-3590	bq. The SystemToken is created using properties of 'instance' to create a hash. Some of these properties are created by default creating a unique hash between different versions. Oh, we added new "instance.*" properties to the configuration between 1.6 and 1.7? Bummer. We'll probably have to rethink how we compute the SystemToken in the future to account for this scenario. Probably another sub-task for this umbrella.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3590
16555	ACCUMULO-3590	Regarding this, what would the process for shutdown, upgrade, start up look like? My problem is that accumulo continually restarts the tablet server after a short interval (give or take 2 minutes) when using stop-here.sh. I don't expect to need to update accumulo that often, but would be nice to be able to do this one node at a time.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3590

#	Requirement	Questions	Answered	Successfully Answered	URL
16556	ACCUMULO-3590	bq. Regarding this, what would the process for shutdown, upgrade, start up look like? Exactly what you described. Some kind of shutdown step, upgrade of JARs on the given node, and restart of that node again. Movement of tablets hosted by the node being upgrade is the most difficult part to efficiently implement. bq. My problem is that accumulo continually restarts the tablet server after a short interval (give or take 2 minutes) when using stop-here.sh Huh? Accumulo does not automatically restart services out of the box. There are ways that you can configure the system to automatically restart stopped processes, but that's definitely not standard functionality.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3590
16557	ACCUMULO-3586	Do you know if the following would work? (I don't have a good way to test) {code} bin/bootstrap_config.sh -o -d target/example-configs/$s/standalone -s $s -j -v 2 >> $out 2>&1 {code}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3586
16558	ACCUMULO-3579	[~ctubbsii], can you take a look? {{mvn verify -Psunny}} passed for me.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3579
16559	ACCUMULO-3579	org/mortbay/jetty/servlet/Context at org.apache.hadoop.hdfs.server.namenode.NameNode$1.run(NameNode.java:400) at org.apache.hadoop.hdfs.server.namenode.NameNode$1.run(NameNode.java:395) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.Subject.doAs(Subject.java:415) at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1190) at org.apache.hadoop.hdfs.server.namenode.NameNode.startHttpServer(NameNode.java:395) at org.apache.hadoop.hdfs.server.namenode.NameNode.initialize(NameNode.java:337) at org.apache.hadoop.hdfs.server.namenode.NameNode.<init>(NameNode.java:569) at org.apache.hadoop.hdfs.server.namenode.NameNode.createNameNode(NameNode.java:1479) at org.apache.hadoop.hdfs.MiniDFSCluster.<init>(MiniDFSCluster.java:278) at org.apache.hadoop.hdfs.MiniDFSCluster.<init>(MiniDFSCluster.java:124) at org.apache.accumulo.test.AccumuloDFSBase.miniDfsClusterSetup(AccumuloDFSBase.java:68) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:24) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:27) at org.junit.runners.ParentRunner.run(ParentRunner.java:309) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:264) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:153) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:124) at org.apache.maven.surefire.booter.ForkedBooter.invokeProviderInSameClassLoader(ForkedBooter.java:200) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:153) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:103) Caused by: java.lang.ClassNotFoundException: org.mortbay.jetty.servlet.Context at java.net.URLClassLoader$1.run(URLClassLoader.java:366) at java.net.URLClassLoader$1.run(URLClassLoader.java:355) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:354) at java.lang.ClassLoader.loadClass(ClassLoader.java:425) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:308) at java.lang.ClassLoader.loadClass(ClassLoader.java:358) at org.apache.hadoop.hdfs.server.namenode.NameNode$1.run(NameNode.java:400) at org.apache.hadoop.hdfs.server.namenode.NameNode$1.run(NameNode.java:395) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.Subject.doAs(Subject.java:415) at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1190) at	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3579
16560	ACCUMULO-3567	Looks good, Keith. A couple of ideas: * Introduce what a TabletBalancer is when introducing the topic * How does someone use the GroupBalancer? Do they write their own implementation? Is it configuration based? * Another example could be tracking website traffic. Similar to geographic distribution, evenly spreading out the analytic records for multiple sites hosted by one organization (e.g. org1 => foo.com, bar.com, foobar.com, org2 => google.com, yahoo.com, etc). Same premise, just different application.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3567
16561	ACCUMULO-3565	Where can I find these dashboards?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3565
16562	ACCUMULO-3560	How can we create the Accumulo-equivalent of https://mms.mongodb.com or https://mongolab.com? These are hosted database solutions. mms.mongodb.com has an agent that will run on your local box to install mongo. mongolab.com is a front-end to AWS and other cloud providers.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3560
16563	ACCUMULO-3557	Who would be authorized to write here? If there are multiple instances sharing the same zookeeper (which I think is crazy-talk, but we support it) then it's not like a single instance secret is authoratative enough to make changes. Maybe the best we can do is just place an ACL on each individual znode and leave {{/instances}} open.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3557
16564	ACCUMULO-3557	bq. Restricting write access would be difficult, since the whole point of this name-to-id mapping paradigm was so somebody could easily re-initialize an instance with the same name, without clobbering an existing instance. In practice, how often is this actually used? I don't think I've ever done this in the years I've used Accumulo. I only see re-init'ing done when I'm destroying the instance. bq. And, the new instance could have a different instance secret, so what ACL do we use? Obviously, we would want to use the current instance.secret (since that's what the rest of the instance nodes would be acl'ed with). The problem would be knowing what the old instance.secret was to remove the old znode. The first thought I have is to have a password-prompt when we get an AuthFailed removing the old ZNode. The administrator running the reinitialization can enter whatever the old secret was. bq. The risks here are how it affects clients. Yes, that was my original concern. Users being redirected to a different Accumulo instances -- maliciously or otherwise. Anyone has the ability to change these nodes which means that anyone with access to the system can prevent users from normally accessing it by just nuking this node.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3557
16565	ACCUMULO-3557	bq. In practice, how often is this actually used? It *may* be useful for testing. However, I think it would be *more* useful to simply say something like "An instance was already created with that name. If you'd like to overwrite it, enter it's name here to confirm: " in order to conveniently remove the old one.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3557
16566	ACCUMULO-3557	Do we have a test for this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3557
16567	ACCUMULO-3548	{{KeyExtent}} is one that kind of worried me, because it's really not "well consumed". Does it make sense to pull a KeyExtent into the public API to just provide some stability? Or, should we try to make a better class which is well suited for public consumption in the public API which KeyExtents can be converted into?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3548

#	Requirement	Questions	Answered	Successfully Answered	URL
16568	ACCUMULO-3537	If we really need a separate property, I like [~kturner]'s suggestion of {{instance.volumes.replacements.relative}}. What do you think, [~elserj]?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3537
16569	ACCUMULO-3537	[~ctubbsii], what's the status here. Is this going to make 1.7.0?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3537
16570	ACCUMULO-3528	bq. we may want to wait until Hadoop bumps their dependency too. Do we know what version of Hadoop would contain this? 2.7.0? It's tough because we're either broken "now" or "later" WRT HDFS spans. Will our code still run against the old htrace and the new incubating htrace runtime dependencies (e.g. can the user swap in the htrace impl they want to use on the classpath w/o having to recompile the code)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3528
16571	ACCUMULO-3528	Overall it makes sense. Track the newer htrace stuff, and versions of HDFS which also use the new htrace libs will automatically start showing up in Accumulo. It encourages use of newer versions of HDFS and, hopefully, reduces our burden. Getting HDFS spans in Accumulo isn't a requirement, just a nice to have. Some quick observations: {code} + SpanReceiverBuilder builder = new SpanReceiverBuilder(wrapHadoopConf(conf)); + SpanReceiver rcvr = builder.spanReceiverClass(className.trim()).build(); + if (rcvr == null) { + log.warn("Failed to load SpanReceiver " + className); } + receivers.add(rcvr); + log.info("SpanReceiver " + className + " was loaded successfully."); } {code} Looks like you want an else after the {{if (rcvr == null)}} if-statement. This will print that it failed to load the spanreceiver, and then that it loaded it. {code} @Override public synchronized boolean offer(Runnable e) { - task.add((Comparable) e); + task.add((TraceRunnable) e); notify(); return true; } {code} In {{CompactionQueue}}, maybe add a {{Precondition}} check instead of a blind cast? At least we can give a better error if we ever get a Runnable that wasn't a TraceRunnable? I see in {{AsyncSpanReceiver}} and {{ZooTraceClient}} the {{configure(HTraceConfiguration)}} method has gone away. Was that nuked in o.a.htrace? The configure method was nice because it made the constructors do less "work" which typically makes it easier to test those classes with mocks.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3528
16572	ACCUMULO-3528	Is there more work to be done, or can we close this now?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3528
16573	ACCUMULO-3521	Can these be deprecated? OrIterator definitely has some bit rot. At least if we deprecate it could bring to light if they are used to not.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3521
16574	ACCUMULO-3521	[~ctubbsii] With the updates in PR #237, Josh's in https://github.com/apache/accumulo/pull/247 and the test I added in https://github.com/apache/accumulo-examples/commit/bec846ce3dec6b594e5e54a1349a99f5202dd91b can I mark this resolved? I think the other methods in (4.) have been removed.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3521
16575	ACCUMULO-3521	Sure. Any progress at all on this will have been useful. Can always re-run UCDetector again in the future.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3521
16576	ACCUMULO-3521	For the items removed in 2.0, are they marked deprecated in 1.8.x?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3521
16577	ACCUMULO-3520	1.7.0 is already released. Also, this is marked as "Fixed" rather than "Duplicate", but has a link indicating this is a duplicate of ACCUMULO-3209. I'm confused. Which ticket, and in which version was the work actually done?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3520
16578	ACCUMULO-3508	[~phrocker], I've added you as a contributor. Look forward to a patch -- thanks!	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3508
16579	ACCUMULO-3508	What's the status of this? Is a patch available? [~phrocker], are you still working on this?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3508
16580	ACCUMULO-3508	Should this be backported to 1.6.5 and/or 1.7.1, since it was originally reported against 1.6.0?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3508
16581	ACCUMULO-3490	The shell servlet integration seems like a separate, follow-on issue. What would it take to get this to work?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3490
16582	ACCUMULO-3490	bq. The shell servlet integration seems like a separate, follow-on issue Perhaps. My guess is that it would be very trivial to enable. bq. What would it take to get this to work? I haven't yet looked into what Hadoop presently does to enable this.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3490
16583	ACCUMULO-3479	I applied most of [~ctubbsii]'s attached patch in addition my changes. Thanks for the help!	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3479
16584	ACCUMULO-3477	bq. What worked really well was extending TThreadedSelectorServer, creating threads for ops I knew would be fast, and leaving things like a "startScan" or a "continue(Multi)Scan" handler to the thread pool. It required more coupling of the server to the processor/handler, but ultimately worked really well. Oh, cool. Thanks for sharing, [~phrocker]. I think coupling the processor impl to the server is kind of required for such an optimization (since we're essentially making decisions on how to run an RPC given the application-level characteristics of that RPC). Do you have any tangible numbers as to how well it worked, or was that just more "feeling"? Any thoughts on how we could measure the improvement of such a change?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3477
16585	ACCUMULO-3470	Also, shouldn't this also be tagged for 1.6 and 1.7?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3470
16586	ACCUMULO-3470	There are two remaining test classes that don't seem to have anything to do with Accumulo left in the accumulo-start module after this change. They seem to only be testing features of vfs hdfs support. Shouldn't those be moved over to be tests within commons-vfs instead? These two are VfsClassLoaderTest and ReadOnlyHdfsFileProviderTest.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3470
16587	ACCUMULO-3470	So the tests pass with the hadoop 2 profile, but not the hadoop 1 profile. The commons vfs2 HDFS dependency is Hadoop 2.6.0. It appears that the hadoop-1 profile was removed in Accumulo 1.7.0. What's the appropriate action here? Revert the commit in the Accumulo 1.6 branch so that it continues to work? I think the ReadOnlyHdfsFileProviderTest class can be removed in Accumulo 1.7.0 and beyond.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3470
16588	ACCUMULO-3464	Doesn't the resolution of "Done" address your issues of it not showing up in a release notes, [~ctubbsii]?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3464
16589	ACCUMULO-3423	Have all the ITs working now, except those related to replication. I talked to [~elserj] about the changes necessary. * the tablet server or master should inform replication when a WAL is closed * the GC should defer deleting WALs in use by replication * the tablet server should just let the GC clean up WALs	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3423

#	Requirement	Questions	Answered	Successfully Answered	URL
16590	ACCUMULO-3409	The RFC 6335 referenced in the description also defines the other two ranges of ports as: System (canonical; aka Well Known), and User (canonical; aka Registered), both of which are described as assigned by IANA. The "Ephemeral" ports are also termed Dynamic (canonical; aka Private), and their distinguishing trait is that they are not assigned by IANA, which is the case for us. While I don't necessarily see a problem changing the default ports for these services, I'm having a bit of trouble understanding the reasoning, since these ports are not registered with IANA and the designation of this range as "Private" seems to make it at least as appropriate as any other range. Further, that RFC also declares the expectation that assignment requests for numbers in the User Port range must provide an explanation as to why use of the Dynamic ports is not acceptable. Do we have any such justification, and/or is anybody planning on requesting IANA assignment for these services? Additionally, the only restriction, from what I can tell, about using the ports in the Dynamic range, is that the port number cannot be used as a service identifier... which is not an issue for us. We identify services through configuration, or registration in ZooKeeper, not through the numerical value of the port. In short, I think it'd be fine to change the numbers to something in the User range, but only if it is accompanied by a request to the IANA for port assignment, with justification and all. In lieu of that, I think using ports in the Dynamic range for defaults, as we've been doing, are fine.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3409
16591	ACCUMULO-3404	I'm all for reintroducing the class, but could you elaborate on the cause for the public API confusion? Do we point people to that class in our docs or something?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3404
16592	ACCUMULO-3404	Was that class instrumentation or the tracer service? If it's instrumentation, maybe it makes sense to keep deprecated even if it's not part of the public API, but if it's part of the tracer service, it should be safe to leave out. However, I'd rather we stick to our definitions going forward, and the public API does not define trace as part of it.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3404
16593	ACCUMULO-3402	what module were the classes in before? was it core?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3402
16594	ACCUMULO-3402	bq. what module were the classes in before? was it core? Yes.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3402
16595	ACCUMULO-3402	Commenting here instead of on the review, because it's more a design issue than a code issue. Ed's comment on ACCUMULO-1242 [here|https://issues.apache.org/jira/browse/ACCUMULO-1242?focusedCommentId=13812208&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13812208] becomes relevant again. Initially I solved it by moving the classes out of core, but that apparently caused other issues. ;) Can we file a follow on to address that and improve how we handle logging in the MR classes again? Otherwise, the patch LGTM (although admittedly, I only skimmed it, since much of it looked like classes were just added/deleted/moved).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3402
16596	ACCUMULO-3402	bq. Can we file a follow on to address that and improve how we handle logging in the MR classes again? Certainly. Thanks for bringing that up -- I had forgotten about it -- and for looking over things.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3402
16597	ACCUMULO-3402	[~elserj] Can you please explain the rational behind your shift in thinking? At one point you said the following, the switched to option 1. Basically can someone please ELI5 what the big win for users is for having a seperate M/R module? Lets call moving everything back to core option 3. bq. FWIW, the easiest resolution I can come up with short-term is a revert of the original changes (move the classes back to core/) and come up with a better solution moving forward. There is possibly an option 4. Break core into two modules core-mr and core-foo (not sure what to call it). Make the current core module depend on the two new modules. I don't really like option 4, but if there is some strong reason to put the MR code in its own module, option 4 may still maintain 1.6 pom compatability from a users perspective. [~elserj] and [~ctubbsii] it seems like you guys are leaning twoards option 1. I don't like this idea of deprecating the entire map reduce API because of historical churn in MR API. This churn occurred in 1.4, 1.5, and 1.6 see ACCUMULO-1533 (removed a lot of deprecated stuff in MR API), ACCUMULO-769, ACCUMULO-267. Of course as I said I don't understand what the win for users is w/ moving stuff into a seperate module, and that lack of knowledge makes me dislike option 1.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3402
16598	ACCUMULO-3394	This is a reasonable reorganization of some server-side utilities related to our RPC services. However, I would strongly prefer a different package name than {{thrift}}, which is a package name we typically use for generated thrift code. Is there a more appropriate package name that will be less likely to be confused with those we use for generated thrift code? If not, that's fine, too. I'll just have to improve my {{grep -v /thrift/}} pattern when I'm spelunking through our code. Maybe {{org.apache.accumulo.server.util.rpc}}?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3394
16599	ACCUMULO-3390	I've looked at the integration tests and I am unable to find where the Monitor is started. If someone can point to the file, I'd appreciate it. No rush. I think I've worked out how to start the Monitor thread but I'd like to compare my approach with the integration tests.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3390
16600	ACCUMULO-3384	"but there may be a use case for system-wide ones also." This caught my eye. Is there a use case for this or is this a boondoggle? The scenario you provided seems to be setting messages through ZK properties. Is this a need or just a nice to have ? Why can't we store this information in the monitor itself? a flat file? Seems like that would be a specific use case that is easy to address.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3384
16601	ACCUMULO-3381	Is trace even necessary to deprecate at all? Is trace part of the public API at all?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3381
16602	ACCUMULO-3375	So, there's a few reasons why I think we should not do this: # This could create a circular dependency (how do we authenticate users so they can interact with the metadata tables? who can scan the user table?) # There's no sense of consistency... and any real-time attempts to be consistent would kill performance: we'd have to scan this table every time we do any other operation: create table now triggers a table scan. # The user database is typically very small, and infrequently updated. The idea of using an Accumulo table to store this seems a bit overkill. Some of this might be a little bit easier with the system user, but this would be an enormous change, and I don't see a lot of value from it. I do see the potential to do it very badly that kills performance or breaks security. If the ZK implementation is not sufficient, we're probably better off improving the pluggable security APIs added in ACCUMULO-259. If anybody is interested in pursuing this, I'd like to review the design.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3375

#	Requirement	Questions	Answered	Successfully Answered	URL
16603	ACCUMULO-3375	Thanks for the thoughts, [~ctubbsii]. bq. This could create a circular dependency (how do we authenticate users so they can interact with the metadata tables? who can scan the user table?) My gut reaction was that the !SYSTEM user would continue to be "special", and perhaps not have an entry in said table. I think that would get around most of the cyclic issues. Other permission issues would still need to be sorted out. bq. There's no sense of consistency... and any real-time attempts to be consistent would kill performance: we'd have to scan this table every time we do any other operation: create table now triggers a table scan. Pretty sure our tables are consistent ;). Turning on the datablock cache for this table (like the metadata table) should alleviate most pain, but metrics would be needed. IsolatedScanners could also be used if we need strong consistency when scanning across a row. bq. The user database is typically very small, and infrequently updated. The idea of using an Accumulo table to store this seems a bit overkill. I think I thought the same thing initially, but how much of that is because of how we have implemented the user system to begin with? Consider a large organization with authentication via an external system like LDAP or Kerberos. Having a scalable store provided by us would be nicer for all those permissions/authorizations/etc. Need some more measurable numbers behind it.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3375
16604	ACCUMULO-3353	I don't know that I'd call this a bug. Maybe an improvement.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3353
16605	ACCUMULO-3338	Why did this commit strip useful comments from the template? Especially since the {{bootstrap_config.sh}} script explicitly looks for some of them, but was not also modified? Minimally, those comments should be re-added, unless there is sufficient reason to justify removing them. Also, under what circumstances are these {{/usr/hdp/current/*}} paths appropriate? I'm not familiar with that specific packaging, and I think there's possible risks for including too broad a classpath (as is, I tend to recommend people modify the example configs, and avoid the kinds of wildcards in the classpaths that we use in the example). (For that matter, I'm not necessarily comfortable with the {{/usr/lib/hadoop*}} ones either, because they seem to map to a very specific 3rd party packaging (and I don't want to get into the business of trying to support all possible 3rd party packaging schemes), but I'm not sure when those were introduced.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3338
16606	ACCUMULO-3338	bq. Why did this commit strip useful comments from the template A "close" comment followed immediately with a new "open" adds no value whatsoever in my opinion. That is why I removed the "close" comments. bq. Especially since the bootstrap_config.sh script explicitly looks for some of them, but was not also modified? I was not aware that comments were of any significance in generating the actual configs. I'm not going to feel bad for missing that, but thank you for bringing it to my attention. bq. they seem to map to a very specific 3rd party packaging Any packaging would be considered "very specific" since it's the only packaging that would be writing to such a location. These locations will ensure that Accumulo, out of the box with the *example* configurations, can run on top of a Hortonworks installation. Unless you can present some case where there are other jars that happen to also match these regexs, there is no risk. Again, these are *example* configurations which really aren't meant for production use. Users shouldn't be using these configs outside of testing and have custom configuration files which have a minimal set of classpath entries for their specific system.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3338
16607	ACCUMULO-3327	Did this ever happen?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3327
16608	ACCUMULO-3327	I meant more, is this something we should cherry-pick back to 1.7? It's not clear what the final change that was committed was.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3327
16609	ACCUMULO-3321	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/18#discussion_r20067167 --- Diff: core/src/main/java/org/apache/accumulo/core/client/impl/ThriftScanner.java --- @@ -115,12 +115,12 @@ public static boolean getBatchFromServer(Instance instance, Credentials credenti } catch (TApplicationException tae) { throw new AccumuloServerException(server, tae); } catch (TooManyFilesException e) { - log.debug("Tablet (" + extent + ") has too many files " + server + " : " + e); + log.warn("Tablet (" + extent + ") has too many files " + server + " : " + e); } catch (ThriftSecurityException e) { log.warn("Security Violation in scan request to " + server + ": " + e); throw new AccumuloSecurityException(e.user, e.code, e); } catch (TException e) { - log.debug("Error getting transport to " + server + " : " + e); + log.warn("Error getting transport to " + server + " : " + e); --- End diff -- This is tricky. In the case that a tserver fails, this would cause logging on the client. Ideally, the system will recover from a tserver failure and the client wouldn't even have to know that a failure happened. Might it make more sense to leave this at debug for the first N occurrences for a server and then elevate it to warn at the N+1th time?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3321
16610	ACCUMULO-3321	Github user busbey commented on the pull request: https://github.com/apache/accumulo/pull/18#issuecomment-62410572 Can we reuse the log message back off from the balancer warnings about "I can't balance because foo" ?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3321
16611	ACCUMULO-3321	Github user arielvalentin commented on the pull request: https://github.com/apache/accumulo/pull/18#issuecomment-62414640 @busbey I am afraid I am not sure what you are referring to. Would you be able to elaborate a bit for me? @joshelser I empathize with that concern because I don't have an instinct as to how much log output this will generate. Would you feel more comfortable with a more sophisticated pattern, where it logs something like "50 failures have occurred in the past N attempts"? I haven't familiarized myself with the code base enough to know if that is possible.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3321
16612	ACCUMULO-3321	Github user arielvalentin commented on the pull request: https://github.com/apache/accumulo/pull/18#issuecomment-62422047 @busbey Thanks for the hint. @joshelser had mentioned a different ticket [ticket|https://issues.apache.org/jira/browse/ACCUMULO-1268] that tries to address some design changes that the team wants to make to the client that would make it a bit more robust. I feel like the direction you are both preferring to head may fall under the umbrella of the "robust" client, e.g. sophisticated logging thresholds, monitoring, resiliency, etc... Am I reading too much into your comments?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3321
16613	ACCUMULO-3321	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/18#issuecomment-62422710 What are the pros and cons of putting the following in your log4j config vs this change? org.apache.accumulo.core.client.impl.ThriftScanner=DEBUG	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3321
16614	ACCUMULO-3317	Is this fix only on master so far? Edit: Never mind, I forgot that the branching scheme had changed.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3317
16615	ACCUMULO-3300	Can you see any problem with putting it in 1.6 as well?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3300
16616	ACCUMULO-3300	Instead of swapping out one hard coded value for another, can we expose this as a configuration to users?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3300
16617	ACCUMULO-3300	[~mdrob] Are you concerned that some people will not like this change? I would shy away from adding extra parameters for this due to the already high complexity of our configuration. I think the future of this feature ought to be migration to some sort of standard instead, like opentsdb. Can you think of a way to do this that reduces the complexity of core Accumulo?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3300

#	Requirement	Questions	Answered	Successfully Answered	URL
16618	ACCUMULO-3294	Thanks for the patch, [~parkjsung]. Some comments.. * Javadoc on the {{updateIterator}} method benefit from explicitly stating that this method is only intended for use when the iterator is already configured at the given scope. Otherwise, this javadoc is very nice. * {{TableIteratorConfigHelper.java}} needs some Javadoc. Class level (why does it exists? what's its purpose) and methods (what does the method do, what are the arguments required). * Does {{TableIteratorConfigHelper#transformOptions}} need to use a TreeMap? Is keeping these properties in sorted order important? If not, a HashMap would be better. Also relevant for {{TableOperationsHelper.java}}. * Would be nice to make sure the arguments to the constructor for {{TableIteratorConfigHelper}} are non-null. {{java.util.Objects#requireNonNull(Object)}} is nice for this. * Some basic comments on what each test is verifying in MockTableOperationsTest.java would be nice. * An integration test would also be appreciated that verifies this API in practice. e.g. set an iterator on a table, verify some query result, use this new API to change some configuration on the iterator and observe that the query results changed.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3294
16619	ACCUMULO-3294	The proposed change only updates a single iterator. A new API call that allows updating all of a tables iterators at once would be nice. Since [~ctubbsii] is working on atomic config changes, this new API could be made to atomically update all of a tables iterator settings at once. Looking at the current patch, it verifies an iterator with the given name is currently set on the specified scopes. What I am proposing would not do this. Why do you need this check?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3294
16620	ACCUMULO-3294	No substantial reason. My original reason was that since the method was called "update", I added a check so that it couldn't be used like a "set". It adds a little bit protection from surprises for clients since if I as a client (and perhaps its just me) am trying to "update" an iterator setting, I'm assuming that the iterator is there. If it suddenly disappears while I'm trying to update, then perhaps something happened that means I shouldn't update. It's not a spectacular solution to provide this protection since it's performing a read/modify/write with no locks so the iterator could still disappear. I don't feel too strongly about the check since there are use-cases for it not being there so if removing the check and making this behave more like a "set" is better aligned with other efforts, I can remove it. What would the new interface look like if we want this to update all of a tables iterators at once? A Collection<IteratorSettingScopePair>?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3294
16621	ACCUMULO-3294	bq. What would the new interface look like if we want this to update all of a tables iterators at once? A Collection<IteratorSettingScopePair>? Yeah something like that. Ilike the name {{ScopedIteratorSetting}}. So pass in a {{Collection<ScopedIteratorSetting>}}. {{ScopedIteratorSetting}} could extend or contain {{IteratorSetting}} and have scope info.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3294
16622	ACCUMULO-3284	Am I missing context from some off-list discussion on this ticket?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3284
16623	ACCUMULO-3284	I'm sorry, I don't understand. What's the basis for the -1? Or am I misinterpreting the closure? Why not just unassign yourself?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3284
16624	ACCUMULO-3268	Would it make sense to configure a BatchWriter with a {{RetryPolicy}}? Or did you envision some other solution?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3268
16625	ACCUMULO-3263	Can this be thought of more generally as a way to forcibly go from relative paths to absolute paths too?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3263
16626	ACCUMULO-3246	Let's not attribute too much wisdom to my comments here. Experiments should be done, and metrics exposed. But, the IMM / #activeTablets should be chunky. How chunky? I'm going with an HDFS block size. If HDFS block size is 100K, and there are ~10 tablets/server, let's go with 1M... maybe more since the data is compressed on the way out. What's the compression rate for your flush files? 10-1? 5-1? That would be good to know.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3246
16627	ACCUMULO-3236	How is this different from ACCUMULO-571? The whole idea seems antithetical to the idea of "clone", since the result will not be a clone of the original in any sense of the word "clone".	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3236
16628	ACCUMULO-3236	Right, we've identified two options to get what you want (union + delete + rename OR clone-ish/merge thing described here) and two options to get what I'd want (union OR clone + clone-ish/merge thing). For the irreversible merge you describe, it's one step your way, and a few primitive steps in my suggestion. For the reversible merge I describe, it's one step my way, or a few steps with what you describe. Both would be implemented using the optimized copy metadata file references, like clone does (plus your suggestion to insert extra split points to avoid chop). I'm not arguing against any of the requirements you've described, or the proposed internal implementation using the same clone mechanisms. The question I'm asking is: which makes the most sense as the API primitives? Personally, I'd go with the well-defined semantics of union and delete. I'm also pointing out that, from a user-facing perspective, this is not an improvement over a clone function... it's a brand new feature, with very different behavior from the clone operation.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3236
16629	ACCUMULO-3236	how about snapshot-import so that it lines up with bulk-import? also it gives a clear delineation on what gets included from Table B in Table A.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3236
16630	ACCUMULO-3236	I agree it's definitely more similar (semantically, not implementation-wise) to bulk import. The bulk import command is importDirectory(). Perhaps importFromTable(). Snapshot sort of implies that the table is consistent at a point in time, and I don't think we can guarantee that unless the source table is offline'd first. Is that going to be a prerequisite for this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3236
16631	ACCUMULO-3236	Has this discussion moved from "how should we do this" to "what should we call it?" AFAICT we all agree on: 1) snapshotting a table's metadata 2) merging in snapshot from 1 into another table's current state My use of "snapshot" isn't a preference-- historically I've always heard "clone" as being "just a snapshot." I do think "clone" is a composition of the snapshot'ing functionality, table creation, import, but I'm not inclined to make snapshotting a public facing API call for users.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3236
16632	ACCUMULO-3236	bq. Has this discussion moved from "how should we do this" to "what should we call it?" I think so. bq. I do think "clone" is a composition of the snapshot'ing functionality, table creation, import, ... I think "snapshot" implies some explicit view at a point in time of a table. Because of the way clone is implemented, I would probably describe the relationship between "clone" and "snapshot" somewhat the other way around. Clone is not so much a composition of a snapshot functionality. Rather, snapshotting can be achieved using clone, along with ensuring the table is offline (or at least, ingest is halted and the table is completely flushed). Clone by itself, can result in some very unexpected behavior if you assume it represents a snapshot in time of a table. For instance, a clone can pick up new data that was bulk imported to a tablet, but not older data that is waiting to be flushed... so even within a tablet, you can get gaps in time from clone, never mind the expectations for a snapshot view for an entire table.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3236
16633	ACCUMULO-3236	If we break wire compat, I'm -1 on inclusion in 1.6. I'd *really* like to maintain wire compat for 1.7, but am -0 if it is broken.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3236
16634	ACCUMULO-3236	What's the status here, [~vines]? Is this viable for 1.7.0?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3236
16635	ACCUMULO-3224	are we concerned about the fall back case where System.nanosTime passes through System.currentTimeMillis in a different set of units?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3224

#	Requirement	Questions	Answered	Successfully Answered	URL
16636	ACCUMULO-3224	[~busbey], do you mean on old kernels that don't have CLOCK_MONOTONIC support? No, I'm not concerned about that case. CLOCK_MONOTONIC is POSIX and (from what I understand) has been in the kernel since well before the 2.6.32 kernel in CentOS 6. And, in any case, that fallback on those very old kernels is no worse than what was there before this was patched. Do you have concerns about this case?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3224
16637	ACCUMULO-3210	Could we adapt [~bhavanki]'s work to isolate the random walk code for this? Maybe move it into a contrib?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3210
16638	ACCUMULO-3206	Copying from user list: It should suffice to list the number of entries for a table, tablet and tablet server. No need to worry about number of unique rows, number of unique column families, etc. By entry I mean number of (key,value)s. For load balancing, we care about how much physical data is on each tablet / tablet server. This is directly proportional to the number of entries, assuming that the key size and value size in bytes do not differ too drastically. If they do (say for raw documents of vastly different sizes), the **best measure is the size of the data in bytes** for each tablet / tablet server. I didn't suggest it because it doesn't look like Accumulo tracks it so it would involve a lot of new implementation and book-keeping, which could hamper performance. Accumulo does already track the number of entries for tables, tablets and tablet server. It's just hard to get to, relying on the format of the metadata table and accessing the non-public Monitor classes. Bringing it to the public API just looks like a matter of reworking the API and letting the client gather the information that the Monitor already does by connecting to each tablet server. Does that sound reasonable?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3206
16639	ACCUMULO-3206	Would it suffice to add his in the current major development line, or would you need it in 1.6?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3206
16640	ACCUMULO-3206	What kinds of use cases can this solve that the REST version of monitor stats in ACCUMULO-3005 doesn't? If we're still pushing to get that in (and I hope we are), couldn't we just add the missing bits (# bytes) to the info returned by that web service?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3206
16641	ACCUMULO-3204	[~ctubbsii], is there more to be done here? Can you update with the current state, please?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3204
16642	ACCUMULO-3197	As part of the removal of the confusing Accumulo configuration story, can you please make a point to be clear at the javadoc level about what API methods do exist and how they should be used? Given that I'm always confused about what methods I can actually use (albeit some of my confusion is related to the server-side AccumuloConfiguration methods), I imagine users are equally or more confused :)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3197
16643	ACCUMULO-3178	It was brought up in ACCUMULO-3181 that the VolumeChooser isn't strictly public API, even though it is selectable in the system-wide configuration properties, which is considered public API. I think it's the intention that the VolumeChooser (and other server-side pluggable components) should be in the public API, but the fact remains it is not currently. Currently, the examples use stuff not in the public API, but the goal of ACCUMULO-1641 is that the examples only use the public API at some point. I don't know what the implications for this are for this issue. Should the VolumeChooser be declared public API first? Given that it's pluggable, it may already be expected to be public API.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3178
16644	ACCUMULO-3178	bq. So, you typically have to run distcp first before you can bulk import. If you list both A and B in instance.volumes, but have a volume chooser only choose A. Now, you can bulk import files currently on B, but they'll get stored on A when they are compacted, eliminating the need to distcp. Couldn't the files you imported get deleted by A before B MajCs them? You would have to clone or offline the source table on A, no?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3178
16645	ACCUMULO-3178	What would delete them?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3178
16646	ACCUMULO-3177	bq. it allows any user with ALTER TABLE to change how the system manages hdfs internals This is already true for a great many HDFS parameters, such as replication, block size, durability, etc. The chooser is further removed from those things, since it at least keeps things at the Accumulo layer of abstraction rather than HDFS internals, because the focus of this issue is to specify a class to do the volume choosing for a particular table. It has little to do with HDFS... only Accumulo volumes. bq. namespaces property altering can be limited to a much smaller set of trusted individuals Table-level properties always override namespace-level properties, and it is presumed that more people will have permissions to alter a table's properties than will have namespace-level permissions, which it seems you are acknowledging by mentioning that it's common to have ALTER TABLE, so I'm not sure what you're suggesting here. Unless... are you suggesting that we create namespace-only properties, which cannot be set/overridden at the per-table level? Because that would be a *huge* change to re-work all the configuration internals. In any case, I think there's merit to having a first pass not require ACCUMULO-3176, by setting the parameter at the namespace level rather than at table creation. I don't think there's any reason to explicitly prevent a per-table property being set for this, though.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3177
16647	ACCUMULO-3177	[~hustjl22]: Can you rebase this patch onto the latest master? In particular, I think some of the calls to the choose method violate the goal of ACCUMULO-3181 and should be updated (some of the use of {{ServerConstants.get*()}} convenience methods seem to be forcing the chooser to choose paths, not volumes again). Additionally, I think some more problematic conflicts were introduced in {{Initialize.java}} since the last time the patch was updated, a few more {{choose}} methods exist, some substituting {{"/"}} with {{Path.SEPARATOR}}, and maybe a few more changes. Thanks for the work on this! Your contributions are appreciated.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3177
16648	ACCUMULO-3176	{quote} bq. Why not instead add the ability to wait for a set of properties to have been seen by all tservers? Because that really sucks for large numbers of tservers. Your single client now needs to communicate to every single server in the instance and ask them to confirm that they've seen a given property. The more realistic way to do it would be to introduce a new RPC call that forces all tservers to drop their ZooCache entries. However, the end result is that you still have to add new API calls to support this. {quote} Couldn't the RPC for setting configuration on a table tell all the tservers to drop cached entries related to said configuration change? that avoids adding new client RPC and solves the general race condition on table properties.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3176
16649	ACCUMULO-3176	bq. Couldn't the RPC for setting configuration on a table tell all the tservers to drop cached entries related to said configuration change Can you be more specific on how you see that working? I can't think of a simple way that wouldn't still be a broadcast to all servers. We could probably hide something inside the implementation for setProperty on {{TableOperationsImpl}} but someone will still have to reach out and talk to each tabletserver (whether client, master, or a tserver). A FATE operation could probably be used to wait for it to happen, but the master would just be doing the RPC then. The other option would be using ZooKeeper to broadcast that, but we've just looped onto the same problem we were trying to solve at that point :). We could potentially have the master delegate to other tservers to tell other tservers to drop their ZooCache for this table (akin to how bulk-loading delegation works) but that's rather complex and scary. Trying to make sure that all possible servers see an update seems more complex to me than just allowing properties to be provided at table creation time before the servers even know about that table.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3176

#	Requirement	Questions	Answered	Successfully Answered	URL
16650	ACCUMULO-3176	bq. Unless we disallow non-creation-time table properties, creation time properties still aren't a real solution. Do you mean that the update of a table property at some point in time (after it exists) is subject to the same race condition? That's definitely true. I'd speculate that the issue is possibly less heinous for iterators/combiners/constraints because it's likely that you only set at table creation but that's subjective. bq. we force a refresh of properties related to a table when it transitions from offline to online? It's not clear to me how this would work either. You could probably invalidate the table properties when a tablet is loaded on a tserver, which would encompass the offline->online transition. I think that could work. I don't like the story of "just offline and re-online your table to make sure configuration propagates". That's really nasty. Are we at least in agreement that allowing properties at table creation do mitigate the potential race condition? We would definitely need something more heavy-handed to properly fix this (or we'd have to do away with ZooCache...) that is probably out of scope for this issue specifically.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3176
16651	ACCUMULO-3176	{quote} Are we at least in agreement that allowing properties at table creation do mitigate the potential race condition? We would definitely need something more heavy-handed to properly fix this (or we'd have to do away with ZooCache...) that is probably out of scope for this issue specifically. {quote} No, because it only solves one race condition. It leaves people prone to needing to redo a table entirely because they didn't realize they had to set the property at creation time. Or worse, it leaves them unaware that they are in this state. {quote} It's not clear to me how this would work either. You could probably invalidate the table properties when a tablet is loaded on a tserver, which would encompass the offline->online transition. I think that could work. I don't like the story of "just offline and re-online your table to make sure configuration propagates". That's really nasty. {quote} It's not nasty if you can only set properties when it's offline. That would mean the story is "In order to ensure a consistent view of per-table configuration properties you can only alter them when the table is offline. When the table is brought online each tablet server will properly reflect the updated property." If I forgot to set something at creation time, I can follow a procedure to ensure I have a consistent set of properties. That might get more complicated if I need better assurances (like doing an offline MR job to verify the extant data is consistent with some constraint I want to add), but atleast there's a usable path that isn't "you have to create a new table and set the properties then."	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3176
16652	ACCUMULO-3147	bq. You just want to be able to specify to use a different (user-created) table Yeah, I wanted to focus less on the specifics of where the table came from and just make it dynamic. I don't want this to be an issue if we make a replication namespace and suddenly decide we want the replication table over the instead of the accumulo namespace (for example). bq. some other config, but ZK makes sense Yeah, the normal site+ZK configuration is what I meant. bq. I'd have preferred all the replication stuff sit outside Accumulo, as an add-on, and just provide the internal hooks necessary for that external system to operate. I will leave you to that to expose all of those scary internals out the clients :). I'm not convinced you can actually do this and properly track the replication status; however, you are more the free to investigate it yourself. bq. the problematic logic in the master and in the tests that arbitrarily create/delete/scan/alter permissions on the replication table How is it arbitrarily doing anything? There are a number of things that are always checked on the replication table and configured if not present. This doesn't change with what I suggested. The name of the table being used is simple not hard-coded. The assumptions on how that table must be configured are still the same regardless of the name. Yes, while some of the permissions granted are not entirely necessary, giving the root user permission to read/write/alter the table is 100% normal and expected. It would be nice if it's done up front (I'm assuming this is something you've done), but it doesn't mean we can't check it later to make sure that is still the case.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3147
16653	ACCUMULO-3147	bq. Yeah, I wanted to focus less on the specifics of where the table came from and just make it dynamic. I don't want this to be an issue if we make a replication namespace and suddenly decide we want the replication table over the instead of the accumulo namespace (for example). Describing it this way seems like an argument for making it external to Accumulo, as an add-on, not as. In my view, it's either a system-managed table or it's a utility table declared by some external, client system. As is, it seems like a hybrid. bq. I will leave you to that to expose all of those scary internals out the clients . I'm not convinced you can actually do this and properly track the replication status; however, you are more the free to investigate it yourself. I have no interest in even attempting to do that. My point was that this seems like it would have been a preferable implementation, because as is, replication seems to be a hybird user-system/baked-in system, and bloats Accumulo's internals. I'm not objecting to the point where I think we should revert (not seriously considering that anyway), but I do think it needs a lot more polishing before it can be released. bq. How is it arbitrarily doing anything? There are a number of things that are always checked.... The very idea of dynamic checking of the table's existence, and proper configuration, reveals the fact that the table is expected to be a user-managed table, not a system-managed one. This is the problem I have with the current implementation. It's treated as a user-managed table, but created at runtime as needed by the master, utilized by the system user for internal system features. It's tested as a user-table, also, being created by the root user in most tests, and deleted as needed for testing. Very little about how it's utilized reveals itself to be part of an internal system table. System tables are always present, have well-defined access permissions, exist in a non-conflicting system-managed namespace, have well-known table IDs, and have strict tests to ensure they behave this way. User tables can be created on the fly, exist anywhere except the system namespace, can be deleted by users, and have permissions that are appropriate to the external user/system that utilizes them. The trace table is a slight exception, because it existed before namespaces and somewhat blurs the boundaries, but even it is primarily treated as a user-controlled table by an external service, since the external tracer service creates the table, and uses credentials to do so, from a stored configuration file. But, the replication table is not created by an external service. It's created by the master, as needed, can conflict with user table names/config, which can block replication from working, does not have a fixed tableID, and is part of an internal feature. Don't get me wrong... I don't think this is unfixable (indeed, I'm trying to fix some of it). I think it just needs some more polishing before release. But, it's what I meant by "arbitrarily".	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3147
16654	ACCUMULO-3137	Work was logged on this issue. Is it complete?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3137
16655	ACCUMULO-3134	Since we can't predict what options everybody's custom compaction strategies will need, does it make sense to just accept a map for configuration?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3134
16656	ACCUMULO-3134	bq. Since we can't predict what options everybody's custom compaction strategies will need, does it make sense to just accept a map for configuration? My thinking was that these options would leverage a specific compaction strategy built into accumulo. If the compact command had an option for specifying a general compaction strategy w/ option map, then that could be mutually exclusive w/ these options. I think this approach will be less cumbersome to use, while still allowing generalization thats more cumbersome.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3134

#	Requirement	Questions	Answered	Successfully Answered	URL
16657	ACCUMULO-3124	This shouldn't be assigned to me. I didn't make this issue, nor do I consider it a blocker. It's probably a useful thing to do, but I don't have the time to do it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3124
16658	ACCUMULO-3115	Is this happening frequently? It's not a great test: trying to verify the durability model by performance. The test should be printing out the performance values: what are you seeing?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3115
16659	ACCUMULO-3110	The class level timeout set on these looks a lot tighter than the sum of the individual tests that were in them. Do these still pass when not disabled via timeout.factor?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3110
16660	ACCUMULO-3110	No, no it looks like you're correct and the [Timeout Rule is applied per-method|http://junit-team.github.io/junit/javadoc/4.11/org/junit/rules/Timeout.html]. I'm not sure why reading through things made me mis-remember it as applying a limit for all tests in the class. Is the javadoc on testsShouldTimeout clear enough on this or should I file a ticket to improve it?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3110
16661	ACCUMULO-3107	Where (all) are you proposing these changes be made? TableOperations, NamespaceOperations... anything else? Totally agree with you. I can't see any argument to returning an Iterable here. We already have a Map instantiated from the Thrift RPC. We should just pass that along.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3107
16662	ACCUMULO-3091	What are the use-cases that we need to hide user configurations from other users? Given that most things are already available to read in ZooKeeper, I want to make sure that there's well-defined reasons to try to implement this because the changes would be very widespread across the entire codebase.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3091
16663	ACCUMULO-3090	If a tablets dir is no longer in instance.volumes, then it should switch to a new volume. See {{VolumeUtil.decommisionedTabletDir()}}. Is this not working?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3090
16664	ACCUMULO-3090	One backup option starting w/ 1.5 is exporting tables. Can clone, offline, export.. leaving original table online while distcp copies exported files. Would need to write a little script to do all tables in an instance.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3090
16665	ACCUMULO-3090	bq. Are you referring to the replication feature, or something else? Yes, disaster recovery was the primary use case that drove replication and was created specifically to avoid the steps which you outlined.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3090
16666	ACCUMULO-3090	{quote} So, when Sean asked how a tablet spread across volumes would affect disaster recovery, I took it to mean backup and recovery. {quote} Correct, I meant recovering a given crashed cluster. Replication is needed for a real disaster recovery plan, but there will be plenty of folks who need to fix broken systems without having a hot-swappable backup. Even those with one would probably prefer to avoid copying as much as possible, esp those who have large enough clusters to warrant the complications of multiple volumes. For example, how does [recovering from a lost ZK quorum|http://accumulo.apache.org/1.6/accumulo_user_manual.html#zookeeper_failure] get impacted if the files for a given tablet are spread across volumes?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3090
16667	ACCUMULO-3089	I read through some of the first ticket. I will have to read through it all later. It sounds as if HSM gives you the ability to ask for a file on a particular medium or with some other characteristic (local or remote). Does this feature work within one namespace, or across namespaces?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3089
16668	ACCUMULO-3089	test/src/test/java/org/apache/accumulo/test/VolumeChooserIT.java --- @@ -0,0 +1,670 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.test; + +import static org.junit.Assert.*; + +import java.io.File; import java.util.HashMap; import java.util.Map; import java.util.SortedSet; import java.util.TreeSet; import java.util.Map.Entry; + import org.apache.accumulo.core.client.BatchWriter; import org.apache.accumulo.core.client.BatchWriterConfig; import org.apache.accumulo.core.client.Connector; import org.apache.accumulo.core.client.Scanner; import org.apache.accumulo.core.client.admin.TimeType; import org.apache.accumulo.core.conf.Property; import org.apache.accumulo.core.data.Key; import org.apache.accumulo.core.data.Mutation; import org.apache.accumulo.core.data.Range; import org.apache.accumulo.core.data.Value; import org.apache.accumulo.core.metadata.MetadataTable; import org.apache.accumulo.core.metadata.schema.MetadataSchema.TabletsSection.DataFileColumnFamily; import org.apache.accumulo.core.security.Authorizations; import org.apache.accumulo.minicluster.impl.MiniAccumuloConfigImpl; import org.apache.accumulo.test.functional.ConfigurableMacIT; import org.apache.hadoop.conf.Configuration; import org.apache.hadoop.fs.Path; import org.apache.hadoop.fs.RawLocalFileSystem; import org.apache.hadoop.io.Text; import org.junit.Test; + +/** + * + */ +public class VolumeChooserIT extends ConfigurableMacIT { + private static final Text EMPTY = new Text(); + private static final Value EMPTY_VALUE = new Value(new byte[] {}); + private File volDirBase; + private Path v1, v2, v3, v4; + + @Override + public void configure(MiniAccumuloConfigImpl cfg, Configuration hadoopCoreSite) { + // Get 2 tablet servers + cfg.setNumTservers(2); + + // Set the general volume chooser to the GeneralVolumeChooser so that different choosers can be specified + Map<String,String> siteConfig = new HashMap<String,String>(); + siteConfig.put(Property.GENERAL_VOLUME_CHOOSER.getKey(), org.apache.accumulo.server.fs.GeneralVolumeChooser.class.getName()); + cfg.setSiteConfig(siteConfig); + + // Set up 4 different volume paths + File baseDir = cfg.getDir(); + volDirBase = new File(baseDir, "volumes"); + File v1f = new File(volDirBase, "v1"); + File v2f = new File(volDirBase, "v2"); + File v3f = new File(volDirBase, "v3"); + File v4f = new File(volDirBase, "v4"); + v1f.mkdir(); + v2f.mkdir(); + v4f.mkdir(); + v1 = new Path("file://" + v1f.getAbsolutePath()); + v2 = new Path("file://" + v2f.getAbsolutePath()); + v3 = new Path("file://" + v3f.getAbsolutePath()); + v4 = new Path("file://" + v4f.getAbsolutePath()); + + // Only add volumes 1, 2, and 4 to the list of instance volumes to have one volume that isn't in the options list when they are choosing + cfg.setProperty(Property.INSTANCE_VOLUMES, v1.toString() + "," + v2.toString() + "," + v4.toString()); + + // use raw local file system so walogs sync and flush will work + hadoopCoreSite.set("fs.file.impl", RawLocalFileSystem.class.getName()); + + super.configure(cfg, hadoopCoreSite); + } + + // Test that uses two tables with 10 split points each. They each use the StaticVolumeChooser to choose volumes. + @Test(timeout = 60 * 1000) public void	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089

#	Requirement	Questions	Answered	Successfully Answered	URL
16669	ACCUMULO-3089	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/16#discussion_r18073259 --- Diff: core/src/main/java/org/apache/accumulo/core/client/admin/TableOperations.java --- @@ -103,6 +103,25 @@ void create(String tableName, boolean versioningIter, TimeType timeType) throws AccumuloException, AccumuloSecurityException, TableExistsException; + /** + * @param tableName + * the name of the table + * @param limitVersion + * Enables/disables the versioning iterator, which will limit the number of Key versions kept. + * @param timeType + * specifies logical or real-time based time recording for entries in the table + * @param properties + * initial table properties the user wants + * @throws AccumuloException + * if a general error occurs + * @throws AccumuloSecurityException + * if the user does not have permission + * @throws TableExistsException + * if the table already exists + */ + void create(String tableName, boolean limitVersion, TimeType timeType, Map<String,String> properties) throws AccumuloException, AccumuloSecurityException, --- End diff -- Why are we adding more constructors at all? It seems unrelated to the nature of the title.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16670	ACCUMULO-3089	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/16#discussion_r18073371 --- Diff: core/src/main/java/org/apache/accumulo/core/client/admin/TableOperations.java --- @@ -103,6 +103,25 @@ void create(String tableName, boolean versioningIter, TimeType timeType) throws AccumuloException, AccumuloSecurityException, TableExistsException; + /** + * @param tableName + * the name of the table + * @param limitVersion + * Enables/disables the versioning iterator, which will limit the number of Key versions kept. + * @param timeType + * specifies logical or real-time based time recording for entries in the table + * @param properties + * initial table properties the user wants + * @throws AccumuloException + * if a general error occurs + * @throws AccumuloSecurityException + * if the user does not have permission + * @throws TableExistsException + * if the table already exists + */ + void create(String tableName, boolean limitVersion, TimeType timeType, Map<String,String> properties) throws AccumuloException, AccumuloSecurityException, --- End diff -- That's a use case, but it isn't a necessity. Why can't the existing table property configuration methods be used?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16671	ACCUMULO-3089	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/16#discussion_r18073417 --- Diff: server/base/src/main/java/org/apache/accumulo/server/fs/VolumeChooser.java --- @@ -18,5 +18,7 @@ public interface VolumeChooser { + String choose(VolumeChooserEnvironment env, String[] options); --- End diff -- Plus, having duplicative methods like this is confusing for an implementation (when does choose1() get called and when does choose2() get called, why, etc). Maybe create some utility/helper method for implementations to get the current VolumeChooserEnvironment?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16672	ACCUMULO-3089	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/16#discussion_r18073730 --- Diff: server/base/src/main/java/org/apache/accumulo/server/fs/StaticVolumeChooser.java --- @@ -0,0 +1,68 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.server.fs; + +import org.apache.accumulo.core.conf.AccumuloConfiguration.AllFilter; +import org.apache.accumulo.core.conf.AccumuloConfiguration.PropertyFilter; +import org.apache.log4j.Logger; + +public class StaticVolumeChooser implements VolumeChooser { + private static final Logger log = Logger.getLogger(StaticVolumeChooser.class); + + public StaticVolumeChooser() {} + + @Override public String choose(VolumeChooserEnvironment env, String[] options) { + // Get the current table's properties, and find the preferred volumes property + PropertyFilter filter = new AllFilter(); + Map<String,String> props = new java.util.HashMap<String,String>(); + env.getProperties(props, filter); + String volume = props.get("table.custom.preferredVolumes"); + log.info("In custom chooser"); + log.info("Volumes: " + volume); + log.info("TableID: " + env.getTableId()); + if (volume != null) { + // If preferred volume is specified, make sure that volume is an instance volume. + boolean usePreferred = false; + for (String s : options) { + if (volume.equals(s)) { + usePreferred = true; + break; + } + } + // If the preferred volume is an instance volume return that volume. + if (usePreferred) + return volume; + // Otherwise warn the user that it is not an instance volume and that the chooser will default to randomly choosing from the instance volumes. + else + log.warn("Preferred volume, " + volume + ", is not an instance volume. Defaulting to randomly choosing from instance volumes"); + } + // If the preferred volumes property is not specified, warn the user, then use a RandomVolumeChoser to choose randomly from the given list of volumes + log.warn("No preferred volume specified. Defaulting to randomly choosing from instance volumes"); + return new RandomVolumeChooser().choose(options); --- End diff -- Lots of object creation and garbage collection. Couldn't this extend RandomVolumeChooser and call `super.choose(options)` when it can't find a static volume?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16673	ACCUMULO-3089	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/16#discussion_r18073784 --- Diff: server/base/src/main/java/org/apache/accumulo/server/fs/VolumeUtil.java --- @@ -261,7 +263,8 @@ private static String decommisionedTabletDir(ZooLock zooLock, VolumeManager vm, throw new IllegalArgumentException("Unexpected table dir " + dir); } - Path newDir = new Path(vm.choose(ServerConstants.getTablesDirs()) + "/" + dir.getParent().getName() + "/" + dir.getName()); + Path newDir = new Path(vm.choose(Optional.of(extent.getTableId().toString()), ServerConstants.getTablesDirs()) + "/tables/" + dir.getParent().getName() --- End diff -- You added an extra "tables" to the constructed `Path`. Was that intentional? Doesn't `ServerConstants.getTablesDirs()` give this to you already?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089

#	Requirement	Questions	Answered	Successfully Answered	URL
16674	ACCUMULO-3089	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/16#discussion_r18073825 --- Diff: server/base/src/main/java/org/apache/accumulo/server/util/FileUtil.java --- @@ -79,7 +81,7 @@ public Text getLastRow() { private static final Logger log = Logger.getLogger(FileUtil.class); private static Path createTmpDir(AccumuloConfiguration acuConf, VolumeManager fs) throws IOException { - String accumuloDir = fs.choose(ServerConstants.getTemporaryDirs()); + String accumuloDir = fs.choose(Optional.<String>absent(), ServerConstants.getTemporaryDirs()) + "/tables/"; --- End diff -- Why the extra "tables" in the path?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16675	ACCUMULO-3089	Jenna, thanks for getting involved in Accumulo! With the feedback on the current patch it looks like you're on your way to having a well formed implementation. Before too much more work goes into incorporating that feedback, can someone help me understand how the current implementation is addressing the concerns raised in my veto? This still looks like a lot of added complexity for marginal benefit, especially compared to a simpler approach like limiting particular namespaces to particular volumes.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16676	ACCUMULO-3089	[~busbey] wrote: {quote} Before too much more work goes into incorporating that feedback, can someone help me understand how the current implementation is addressing the concerns raised in my veto? This still looks like a lot of added complexity for marginal benefit, especially compared to a simpler approach like limiting particular namespaces to particular volumes. {quote} I replied to your concerns above. In particular, I explained that your "simpler approach" is practically identical to what is proposed. This is an implementation of precisely that sort of thing. You never replied with any disagreement, or explanation of how you thought they were different. Also, just a small point of order here: you can veto a commit, but you can't veto an idea or a contribution. As it stands, I see your "-1" as blocking this patch from being committed, but not the refinement, discussion, and review, of this contribution to the community. As long as your veto stands, it can't be commited to Accumulo, so maybe Jenna will have to host it somewhere else, but my hope is that you'll change your veto once you take a look at the change and realize it achieves precisely the sort of thing that you prefer "like limiting particular namespaces to particular volumes", especially since I believe your veto was based on a misunderstanding of this being scoped around heterogeneous storage.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16677	ACCUMULO-3089	Github user ctubbsii commented on a diff in the pull request: https://github.com/apache/accumulo/pull/16#discussion_r18100611 --- Diff: server/base/src/main/java/org/apache/accumulo/server/fs/VolumeChooser.java --- @@ -18,5 +18,7 @@ public interface VolumeChooser { + String choose(VolumeChooserEnvironment env, String[] options); --- End diff -- Right, this will break the interface. This is low risk, I think, though. The VolumeChooser isn't part of the public API (yet), is it? The alternative is that we can pass the environment with dependency injection to an annotated field, if it exists. If this is something that's a blocker, that's an option. There's also a slight behavioral change to consider here (improvement, I think) that may be relevant, if we're concerned about API. The original VolumeChooser interface did not choose between volumes.... it chose between paths. This patch changes things so that the API conforms to the user expectations that it will be provided with volumes from which to choose, which the framework will append any necessary paths to, if needed.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16678	ACCUMULO-3089	[~busbey] My question above was never answered. I'll paste it below so you don't have to search for it: bq. It sounds as if HSM gives you the ability to ask for a file on a particular medium or with some other characteristic (local or remote). Does this feature work within one namespace, or across namespaces? Personally, I don't understand how HSM is equivalent to Volume selection. It appears based on what I read that HSM would allow you to create a file either in memory, disk, or some other medium. Volumes are a mapping to an HDFS namespace. Does HSM allow you to specify that you want a file created in a particular namespace?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3089
16679	ACCUMULO-3089	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/16#discussion_r18106232 --- Diff: server/base/src/main/java/org/apache/accumulo/server/fs/VolumeChooser.java --- @@ -18,5 +18,7 @@ public interface VolumeChooser { + String choose(VolumeChooserEnvironment env, String[] options); --- End diff -- The user facing property general.volume.chooser allows user to set classes that implement the existing interface. That seems to put it in the public API? One option is to deprecate that property and introduce a new property that uses a new interface (don't base new interface on current one). If the old prop is set, could still use the old class/interface. If both are set, then error.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16680	ACCUMULO-3089	[~busbey] wrote: bq. Just as a simple example, limiting volumes to particular namespaces would not require changing when table properties are available. It would also not require changing the table creation api. Oh, I see. Yes, absolutely, that would be a simpler example. Certainly, that could be done, but I don't see how useful it would be. Without any environment, the only choosers possible are ones that have hard-coded decisions, which would have to be recompiled when the policy changes, or who rely on external/remote services for configuration. (This is true of the namespace example, for instance... because there is no way for the chooser to know which namespace it is making a decision for.) None of our other pluggable components have such limitations, and we've made it possible to configure each through environment (usually table properties, as in the case of iterators, constraints, etc.), so I don't see why this addition should be vetoed here. So, yes, that's a simpler example... but it's not an example which satisfies the criteria of this ticket, because this ticket is essentially about creating an example which can make smarter decisions, based on configuration. Is "there could be another implementation/example" your only remaining basis for your veto, or is there something inherently wrong with the idea of giving environment/configuration to the chooser? Also, how do you feel about the sub-tasks, the per-table volume chooser (similar to the per-table balancer), or the example preferred volume chooser? Does breaking it up like that help frame, change, or address your concerns in any way? To which of those two would your veto apply, or would it be for both? bq. For the record, my veto is not based on a misunderstanding. That seems to express a degree of confidence in one's understanding the other side of the conversation that I do not believe anybody could ever have. There is *always* a risk you've misunderstood somebody else's intentions (in this case, my intentions for this ticket I created). I understand this statement to mean that you do not believe you've misunderstood the intentions.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3089
16681	ACCUMULO-3089	bq. It doesn't seem that any progress is being made on resolving Sean's veto on this ticket. I suggest some other means of communication where back-and-forth is easier. Looking at the bylaws[1], this isn't a veto, it's an objection. The ASF community concensus building guide[2] appear to back this up. However, our bylaws state that a vote must take place since an objection has been raised. I'd like to prevent that by trying to understand the details behind the proposed alternative. If we can't get our questions answered here, how will changing the communication mechanism help? [1] http://accumulo.apache.org/bylaws.html#CTR [2] https://community.apache.org/committers/consensusBuilding.html	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3089

#	Requirement	Questions	Answered	Successfully Answered	URL
16682	ACCUMULO-3087	Have you tested this on OS X? I know hostname has different options in BSD/Mac.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3087
16683	ACCUMULO-3079	Nifty, did you use the code from ACCUMULO-3067 to see the performance improvements? Anything else that can be replicated?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3079
16684	ACCUMULO-3079	bq. I moved the synchronization to the filter, so the functionality is still there Right, I was just noting that you came to the same conclusion about being able to move it. bq. I did play around with removing the SyncronizedIterator altogether in the case where no client iterators are present Hrmm. I wonder how that works for compaction time iterators? Mostly a rhetorical question since you didn't include that change here. bq. I could subjectively say that the performance difference between not having the Synchronized iterator at all and moving its functionality into the VisibilityFilter was small enough to be hidden in the testing noise, but I don't have evidence to objectively state that. Ok, I didn't necessarily expect you to have specifics for that either. I brought it up mostly because I was curious. While having the synchronized "barrier" in the iterator stack does make isolate user iterators from mucking up the system iterators, I think I was really wondering if it even makes sense to "encourage" users to write multi-threaded iterators. I know I personally haven't come across any use cases where multiple threads are used effectively by an iterator. Would it make sense to remove the synchronized iterator(filter) from the default "pipeline" and wire up something through the client API and/or table configuration that can let users enable it. Not implying that this should be done for this ticket, it's just a parallel thought that has some relevancy here.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3079
16685	ACCUMULO-3079	Good question on the compaction scope changes. Since IteratorUtil is used in all scopes, the attached patch does remove the synchronization from the compaction scope. We don't generally do visibility filtering in the compaction scope, so I don't suppose the security vulnerability applies. Do you see any reason why we would need synchronization on the compaction scope? bq. I was really wondering if it even makes sense to "encourage" users to write multi-threaded iterators. I don't believe I have ever come across a use case in which multiple threads are needed (or useful) within an iterator either. From a security perspective, another approach to solving this problem would be to use the java SecurityManager to prevent iterators from using multiple threads.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3079
16686	ACCUMULO-3079	bq. Do you see any reason why we would need synchronization on the compaction scope? It's possible that a user configured their own crazy com.multi.threaded.MyIterator at compaction time which *could* need synchronization, but that's also a reach given that the use-case for multi-threading is a stretch anyways. bq. another approach to solving this problem would be to use the java SecurityManager to prevent iterators from using multiple threads. That's certainly stronger than my first thought of documenting somewhere "don't start threads in an iterator" :). Hopefully if a user tried to do this on their own, they'd realize that they might be doing something weird when they don't have the hooks to ensure 100% that any threads started also terminate.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3079
16687	ACCUMULO-3057	How does the shell do this for the ConfigCommand?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3057
16688	ACCUMULO-3045	I noticed you had a few methods in the commit that overrode the super class, but simply called super.method() (in particular, write() and readFields()). Is there some reason these don't just inherit from the parent class? Also, I'm a bit confused on the intent. Is the token intended to resolve on the client side to a PasswordToken? Or is it intended to resolve on the server-side after the CredentialToken is reconstructed from the RPC? (I only briefly glanced at the diff so far, but hoping to get some information on intent, before further evaluation). My current understanding is that this does the latter (resolving the actual credentials on the server-side). If that is the case, it seems like there may be a gap for a "PasswordResolver.getPasswordToken(CredentialProvider cp);" for client-side resolution (for instance, in the case of reading a password from the local filesystem or environment).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3045
16689	ACCUMULO-3045	bq. I noticed you had a few methods in the commit that overrode the super class, but simply called super.method() (in particular, write() and readFields()). Is there some reason these don't just inherit from the parent class? Ah, some cruft I could've cleaned up. I had thought I needed some extra members and, thus, had to add to the serialization, but ended up not needing to. I'll clean those up. bq. I'm a bit confused on the intent Ultimately the intent is to be able to use a CredentialProvider to authenticate against Accumulo. I'm not entirely positive what you mean about client side or server side, but I *think* the intent is server-side. The desired intent is to be able to seamlessly authenticate with Accumulo using a CredentialProvider (using the shell, a mapreduce job, really anything using a Connector). In practice, it comes off as very similar to what exists for presently for providing passwords to the shell from a file or an environment variable. Does that help?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3045
16690	ACCUMULO-3045	Ahh, yes, I understand you now. If it's not backed by some distributed "means", it's limited by the local resource. Yes, I agree with you 100%. It'll be interesting to see the types of implementations that come about (I'm not entirely sure what's in the works yet, myself) but, given that this is trying to be addressed at the "platform" (hadoop) level, I'm sure that we'll see some good implementations. I definitely liked the fact that I had to alter absolutely no security implementation code and could solely implement this by the creation of a new AuthenticationToken. I made a writeup in the user manual for ACCUMULO-2464; I could do the same for these changes (maybe add in a new subsection in the "Writing Clients" section, or similar). Are you looking for user info, developer info, both? Thoughts?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3045
16691	ACCUMULO-3042	I think we currently provide this functionality for things that call the hadoop command via bin/tool.sh. Presuming we add in this utility class and document how to use it, would we be able to remove that script?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3042
16692	ACCUMULO-3023	If we are going to do this automatically as a part of the 1.7.0 upgrade (read as the first time a 1.7.0 master is started on an instance?), what kind of fail-safes do we need? What happens if the user provides a bad value for the proposed configuration value from ACCUMULO-3537? Should we provide some tool that allows dry-runs? Make sure that the resulting fully-qualified path exists before writing it back to the metadata table?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3023
16693	ACCUMULO-3022	What exactly is the reasoning for this proposal? What problem is being avoided by limiting the length of table names?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3022
16694	ACCUMULO-3022	bq. What problem is being avoided by limiting the length of table names? Memory exhaustion of zookeeper servers. Maybe the issue should be more general. What other points in the Accumulo API can the user put arbitrary data into zookeeper?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3022
16695	ACCUMULO-3022	do namespace names go into zookeeper?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3022
16696	ACCUMULO-3022	How many, let's say, 200 character table names does it take to exhaust the Zookeeper memory and why would that be an Accumulo issue? Plucking a number out of the air, we could say that 2,500 table names of 200 characters each exhaust Zookeeper memory. If we arbitrarily limit table names to 100 characters then Zookeeper memory would be exhausted after 5,000 tables were created? It seems like table name size should be a guideline not an enforceable constraint. You have to know that six months after the constraint is added, someone will have a use case for longer table names and then we'll need to make the constraint configurable. Oh.. and does this need to have 'Major' priority?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3022
16697	ACCUMULO-3022	bq. do namespace names go into zookeeper? Yes, in the same way that table names do.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3022

#	Requirement	Questions	Answered	Successfully Answered	URL
16698	ACCUMULO-3021	What do you mean by 'very large split row'? I can think of several different definitions for this phrase.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3021
16699	ACCUMULO-3020	This issue has been marked "In Progress" for 9 months. What's the status?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3020
16700	ACCUMULO-3019	+1 on undeprecating {{ServerConfiguration}}. Sorry about the warnings. IMO a "ServerConfiguration" object should be some sort of configuration object, so I'm not in favor of having the factory implement it. I'm happy to work with you on figuring out some solution for the plugins. {{ServerConfiguration}} isn't part of the public API, though, is it? To what extent can we practically adjust it? I'm working on the static {{SiteConfiguration}} calls under ACCUMULO-2603. A complicating factor is that {{SiteConfiguration}} sort of belongs in o.a.a.server.conf and not the core, as [~ctubbsii] explained in ACCUMULO-2992, but I'm not planning to address that in the near term.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-3019
16701	ACCUMULO-3005	Is this at odds with ACCUMULO-1013?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3005
16702	ACCUMULO-3005	bq. Is this at odds with ACCUMULO-1013? I'm not really sure what a URL builder is, but I would imagine that this issue would subsume that issue.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3005
16703	ACCUMULO-3005	bq. are there specific advantages you've encountered that make the switch compelling? In using JaxRS resources over Servlets, there is definitely much more literature on configuring Grizzly with said resources, which is primarily how my decision was made so far. I haven't actually been able to find a technical difference (or even a comparison) between the two as to why you'd choose Jetty over Grizzly or vice versa (oddly..). I also haven't found any docs on how to configure Jetty to do the jaxrs stuff (as opposed to servlets) either. I did find that Grizzly publishes a WADL automagically which was pretty cool for me to at least validate that my resources were being loaded. I could see some benefit from that for consumers as well, but that's by no means a strong differentiation.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3005
16704	ACCUMULO-3005	You're not using Jackson directly, are you? I'd imagine you're just using it as the registered provider for {{@Produces("application/json")}}?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-3005
16705	ACCUMULO-3000	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/29#issuecomment-94579191 Overall, this looks nice. Thanks! Could you possibly add an integration test that verifies that you can use the new methods to start up a MiniAccumuloCluster instance? You could lift what MiniAccumuloClusterImpl already does to start ZooKeeper and then just do something simple like read from the Accumulo metadata table as a "can I actually connect to and interact with Accumulo".	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16706	ACCUMULO-3000	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r28747812 --- Diff: minicluster/src/main/java/org/apache/accumulo/minicluster/MiniAccumuloConfig.java --- @@ -108,6 +108,19 @@ public MiniAccumuloConfig setZooKeeperStartupTime(long zooKeeperStartupTime) { } /** + * Configure an existing ZooKeeper instance to use. Calling this method is optional. If not set, a new ZooKeeper instance is created. + * + * @param existingZooKeepers + * Connection string for a already-running ZooKeeper instance + * + * @since 1.7.0 + */ + public MiniAccumuloConfig setExistingZooKeepers(String existingZooKeepers) { --- End diff -- > Mini runner is user facing It is, but I don't see a place where we actually leak MiniAccumuloCluster/Config through that class. Did I miss something? If we want to be really picky (and these changes go on the impl), we should not even expose the new property via MiniAccumuloRunner. Is this what you were pointing to?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16707	ACCUMULO-3000	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r28750135 --- Diff: minicluster/src/main/java/org/apache/accumulo/minicluster/MiniAccumuloConfig.java --- @@ -108,6 +108,19 @@ public MiniAccumuloConfig setZooKeeperStartupTime(long zooKeeperStartupTime) { } /** + * Configure an existing ZooKeeper instance to use. Calling this method is optional. If not set, a new ZooKeeper instance is created. + * + * @param existingZooKeepers + * Connection string for a already-running ZooKeeper instance + * + * @since 1.7.0 + */ + public MiniAccumuloConfig setExistingZooKeepers(String existingZooKeepers) { --- End diff -- > It is, but I don't see a place where we actually leak MiniAccumuloCluster/Config through that class. Did I miss something? No, not implying anything leaks. > If we want to be really picky (and these changes go on the impl), we should not even expose the new property via MiniAccumuloRunner. Is this what you were pointing to? Just thinking that if we add it to mini runner, that we will need to continue to support that option in mini runner. I think its very similar to adding to API. I am not opposed or in favor of adding an option to mini runner or API. I don't know enough about the performance impact to have an opinion about adding to API or mini runner. I am really curious about the performance impact.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16708	ACCUMULO-3000	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r28800662 --- Diff: minicluster/src/main/java/org/apache/accumulo/minicluster/impl/MiniAccumuloConfigImpl.java --- @@ -313,6 +321,19 @@ public MiniAccumuloConfigImpl setZooKeeperStartupTime(long zooKeeperStartupTime) } /** + * Configure an existing ZooKeeper instance to use. Calling this method is optional. If not set, a new ZooKeeper instance is created. + * + * @param existingZooKeepers + * Connection string for a already-running ZooKeeper instance + * + * @since 1.7.0 + */ + public MiniAccumuloConfigImpl setExistingZooKeepers(String existingZooKeepers) { + this.existingZooKeepers = existingZooKeepers; --- End diff -- Should there a be a guava precondition check for null? How does the user unset/clear this setting?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16709	ACCUMULO-3000	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r28801819 --- Diff: minicluster/src/main/java/org/apache/accumulo/minicluster/MiniAccumuloConfig.java --- @@ -108,6 +108,19 @@ public MiniAccumuloConfig setZooKeeperStartupTime(long zooKeeperStartupTime) { } /** + * Configure an existing ZooKeeper instance to use. Calling this method is optional. If not set, a new ZooKeeper instance is created. + * + * @param existingZooKeepers + * Connection string for a already-running ZooKeeper instance + * + * @since 1.7.0 + */ + public MiniAccumuloConfig setExistingZooKeepers(String existingZooKeepers) { --- End diff -- If its going to go into the API, I think the following needs to be considered. * Need a getter in the API. What should the getter do/return if not set? * Do we want to allow the user to clear this setting, if so how does the user clear this setting? * Need to document in javadoc its relationship to `setZooKeeperPort()`. Should these be made mutually exclusive? Trying to set both would cause the 2nd set to throw an exception?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000

#	Requirement	Questions	Answered	Successfully Answered	URL
16710	ACCUMULO-3000	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r28803038 --- Diff: minicluster/src/main/java/org/apache/accumulo/minicluster/MiniAccumuloConfig.java --- @@ -108,6 +108,19 @@ public MiniAccumuloConfig setZooKeeperStartupTime(long zooKeeperStartupTime) { } /** + * Configure an existing ZooKeeper instance to use. Calling this method is optional. If not set, a new ZooKeeper instance is created. + * + * @param existingZooKeepers + * Connection string for a already-running ZooKeeper instance + * + * @since 1.7.0 + */ + public MiniAccumuloConfig setExistingZooKeepers(String existingZooKeepers) { --- End diff -- For the getter, would returing a guava Optional make sense? I have not used Optional, I have only read about it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16711	ACCUMULO-3000	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r30200397 --- Diff: minicluster/src/main/java/org/apache/accumulo/minicluster/MiniAccumuloRunner.java --- @@ -75,6 +75,7 @@ private static final String NUM_T_SERVERS_PROP = "numTServers"; private static final String DIRECTORY_PROP = "directory"; private static final String INSTANCE_NAME_PROP = "instanceName"; + private static final String EXISTING_ZOO_KEEPERS_PROP = "existingZooKeepers"; --- End diff -- Did we already address the addition here? MiniAccumuloRunner is also public API (like MiniAccumuloCluster). Seems odd to add it here. I don't think we have a good way to hide it though (like the MiniAccumuloClusterImpl), so maybe it's ok?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16712	ACCUMULO-3000	Github user ctubbsii commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r30266262 --- Diff: .gitignore --- @@ -23,3 +23,4 @@ /.pydevproject /.idea /*.iml +*.swp --- End diff -- What's this for?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16713	ACCUMULO-3000	Github user ctubbsii commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r30268333 --- Diff: minicluster/src/main/java/org/apache/accumulo/minicluster/MiniAccumuloRunner.java --- @@ -93,6 +94,7 @@ private static void printProperties() { System.out.println("#" + TSERVER_MEMORY_PROP + "=128M"); System.out.println("#" + ZOO_KEEPER_MEMORY_PROP + "=128M"); System.out.println("#" + JDWP_ENABLED_PROP + "=false"); + System.out.println("#" + EXISTING_ZOO_KEEPERS_PROP + "=localhost:2181"); --- End diff -- I realize you were probably just following what was there, but does it make sense to print these *before* the configuration options are parsed and set?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16714	ACCUMULO-3000	Github user ctubbsii commented on a diff in the pull request: https://github.com/apache/accumulo/pull/29#discussion_r30269047 --- Diff: pom.xml --- @@ -119,6 +119,8 @@ <accumulo.release.version>${project.version}</accumulo.release.version> <!-- bouncycastle version for test dependencies --> <bouncycastle.version>1.50</bouncycastle.version> + <!-- Curator version --> + <curator.version>2.7.1</curator.version> --- End diff -- Do we anticipate needing to control this on the command line, or elsewhere? We generally don't need properties for dependency versions (especially for test dependencies) unless the same version is used repeatedly for multiple artifacts (like jetty or slf4j) or if we anticipate swapping it out frequently during different build environments.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16715	ACCUMULO-3000	Github user ctubbsii commented on the pull request: https://github.com/apache/accumulo/pull/29#issuecomment-102496224 @joshelser You mean with regard to the public API? I don't see any semver violations there. It doesn't even appear to change any public API at all. It just adds an additional (optional) configuration item. Even if that were considered public API, it's an addition, so it should be fine (this is going on 1.8, right?).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-3000
16716	ACCUMULO-2997	Reading the comment, I'm not sure this is actually a bug. The comment implies that the race condition is not an issue. Should this ticket be closed as "Not a Problem"?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2997
16717	ACCUMULO-2992	How about an annotation, something like {{@ServerOnly}}? It may even be possible to have the compiler complain when it sees it used incorrectly: http://stackoverflow.com/questions/1752607/how-to-intentionally-cause-a-custom-java-compiler-warning-message Barring that, I'm fine with a comment. I'll suggest the class Javadoc for {{SiteConfiguration}}, but I'm open to whatever.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2992
16718	ACCUMULO-2984	is the plan for this to be another implementation of the MiniAccumuloCluster class ala Josh's work to have a cluster that spins up in yarn? or is the plan for it to be alterations on MiniAccumuloClusterImpl? I'd prefer the former if you're looking to run against an Accumulo config rather than e.g. just running against an existing HDFS/ZK.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2984
16719	ACCUMULO-2984	The latter. This is exactly like existing MAC, but just a case of A. providing an external ZK B. (may be optional) providing an external site.xml file and C. telling it not to try to initialize. Why would you prefer the former?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2984
16720	ACCUMULO-2981	Is there a checkstyle rule or something we could use to enforce not having a "_" identifier show up again?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2981
16721	ACCUMULO-2981	bq. Is there a checkstyle rule or something we could use to enforce not having a "_" identifier show up again? I do not know. However, to put this in context, we've had precisely one variable named "_" since the project began.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2981
16722	ACCUMULO-2981	-1. Why are we using a character that is difficult to type on traditional keyboards? What's wrong with T/F if we still want single character labels. Also, there were a bunch of other changes in here that should have been their own issue.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2981
16723	ACCUMULO-2981	[~mdrob] wrote: bq. Why are we using a character that is difficult to type on traditional keyboards? What's wrong with T/F if we still want single character labels. The underscore character was originally chose explicitly to make the table which documents the state change operations more readable. I chose a Unicode character which is just as valid as any other character in Java, which preserved the original readability intentions. I did try T/F and _X/__, but this choice looked the best to me. This character's usage is isolated to the array initialization, and should have minimal impact on future development. If you think another combination would be better, please feel free to change the variable name to your preferred character(s). bq. Also, there were a bunch of other changes in here that should have been their own issue. The only other changes in this commit were a type safety check of a generic method which prevented compilation on JDK8, and a change to the POM to make it easier for me to run a test build against 1.8. I wouldn't consider those two "a bunch". They were the minimum changes I needed to make in order to actually test the change. They're also pretty trivial. If you have an objection to either of those changes, feel free to open a new issue.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2981
16724	ACCUMULO-2976	Wouldn't this make more sense as part of the watchdogs people implement? Perhaps if it gets killed x times in y period, don't restart it.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2976
16725	ACCUMULO-2969	How do you plan on handling stale table mappings?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2969
16726	ACCUMULO-2969	can't we just make the thrift call whenever we need the mapping? and maybe a periodic poll for the table name tab completion?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2969
16727	ACCUMULO-2969	When accumulo began to support table renames, a lot of tablename -> tableid lookups were done in client code. {{Tables}} is used for those lookups. Given the asynchronous notification of zookeeper, we don't really know the true table state at the moment a change is made. How important is it to use zookeeper callbacks to keep the table list current?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2969

#	Requirement	Questions	Answered	Successfully Answered	URL
16728	ACCUMULO-2959	Last time I tried this, TravisCI couldn't handle the ITs. Are you just planning on enabling the unit tests?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2959
16729	ACCUMULO-2959	Could we achieve these same goals by updating our ASF Jenkins builds?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2959
16730	ACCUMULO-2959	I don't think so. I'm interpreting the goal as making it easier for a user with their own fork to set up automated builds. Does the ASF jenkins let anonymous clones take up resources on it?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2959
16731	ACCUMULO-2959	Also, just noticed that you're not listed on the [contributors page|http://accumulo.apache.org/people.html], [~mikewalch]. Would you like to be? (optionally with employer and/or timezone)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2959
16732	ACCUMULO-2951	I cannot reproduce. If it was related to thrift, could it have been fixed by ACCUMULO-2950?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2951
16733	ACCUMULO-2946	These two classes have two different interfaces, o.a.a.core.client.admin.TableOperations and o.a.a.core.client.admin.NamespaceOeprations with a lot of relative documentation. First thought: do you see collapsing the common methods into generic ones with <table> or <namespace> parameters? This may require the addition of a third interface that holds the generics and the original interfaces having those specific to their scope. Second thought: should a CommonOperationHelper class hold those methods that are needed by o.a.a.core.client.admin.NamespaceOperationsHelper and o.a.a.core.client.admin.TableOperationsHelper and use the boolean flag to set the scope as table or namespace respectively? I would mind work this issue with some guidance. Feel free to assign it to me and provide some implementation details.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2946
16734	ACCUMULO-2944	In an effort to make this painless for users, could we use the java -version from the $JAVA_HOME to figure out which JVM should be configured?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2944
16735	ACCUMULO-2944	Corey, frequently the default JAVA_HOME in the shell might not be set with the desired JVM since a number of different versions could be installed. I agree that it might be desirable to suggest the JVM installed on JAVA_HOME as a default selection if no other specific ones are selected. Should this be recorded as an enhancement in a separate JIRA?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2944
16736	ACCUMULO-2944	Github user busbey commented on the pull request: https://github.com/apache/accumulo/pull/9#issuecomment-47536342 Can you make sure you've pushed the updated version to your ACCUMULO-2944 branch? That will update this PR. Right now it looks like the previous 2 commits are still there.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2944
16737	ACCUMULO-2944	Github user busbey commented on a diff in the pull request: https://github.com/apache/accumulo/pull/9#discussion_r14414275 --- Diff: conf/templates/accumulo-env.sh --- @@ -51,10 +51,11 @@ test -z "$ACCUMULO_TSERVER_OPTS" && export ACCUMULO_TSERVER_OPTS="${POLICY} ${tS test -z "$ACCUMULO_MASTER_OPTS" && export ACCUMULO_MASTER_OPTS="${POLICY} ${masterHigh_masterLow}" test -z "$ACCUMULO_MONITOR_OPTS" && export ACCUMULO_MONITOR_OPTS="${POLICY} ${monitorHigh_monitorLow}" test -z "$ACCUMULO_GC_OPTS" && export ACCUMULO_GC_OPTS="${gcHigh_gcLow}" - test -z "$ACCUMULO_GENERAL_OPTS" && export ACCUMULO_GENERAL_OPTS="-XX:+UseConcMarkSweepGC -XX:CMSInitiatingOccupancyFraction=75 -Djava.net.preferIPv4Stack=true" +test -z "$ACCUMULO_GENERAL_OPTS" && export ACCUMULO_GENERAL_OPTS="${gcPolicy} -Djava.net.preferIPv4Stack=true ${jaxpDocBuilderFactory} ${randomNumGeneratorProviderOverride}" test -z "$ACCUMULO_OTHER_OPTS" && export ACCUMULO_OTHER_OPTS="${otherHigh_otherLow}" # what do when the JVM runs out of heap memory export ACCUMULO_KILL_CMD='kill -9 %p' +export GC_POLICY_ARGS="${gcPolicy}" --- End diff -- where does this exported variable actually get used?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2944
16738	ACCUMULO-2944	Github user haydenmarchant commented on the pull request: https://github.com/apache/accumulo/pull/9#issuecomment-47693949 I tried to comply with the contributor guidelines for the Pull Request. Can you be specific to what I am doing wrong?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2944
16739	ACCUMULO-2944	Github user busbey commented on a diff in the pull request: https://github.com/apache/accumulo/pull/9#discussion_r14696044 --- Diff: minicluster/src/main/java/org/apache/accumulo/minicluster/impl/MiniAccumuloClusterImpl.java --- @@ -259,7 +259,14 @@ private Process _exec(Class<?> clazz, List<String> extraJvmOpts, String... args) for (Entry<String,String> sysProp : config.getSystemProperties().entrySet()) { argList.add(String.format("-D%s=%s", sysProp.getKey(), sysProp.getValue())); } - argList.addAll(Arrays.asList("-XX:+UseConcMarkSweepGC", "-XX:CMSInitiatingOccupancyFraction=75", "-Dapple.awt.UIElement=true", Main.class.getName(), className)); + + String gcPolicyArgs = System.getenv("GC_POLICY_ARGS"); --- End diff -- For this to impact the current MAC uses in the ITs, won't we need to have this env variable set? should we be adding jvm-vendor activated profiles to the pom?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2944
16740	ACCUMULO-2944	Github user busbey commented on the pull request: https://github.com/apache/accumulo/pull/9#issuecomment-48436920 Overall patch is looking good. Could you rebase onto the current 1.6.1-SNAPSHOT and resolved the conflict? The commit message still doesn't start with the jira reference. Could you move it to the beginning of the summary line?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2944
16741	ACCUMULO-2944	Github user busbey commented on the pull request: https://github.com/apache/accumulo/pull/9#issuecomment-48513357 something has gone wrong in the rebase, since git now thinks all ~1750 commits are a part of the merge. would you mind if I just picked off your commit and did the cleanup to make it work on the current dev branch?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2944
16742	ACCUMULO-2944	It looks like a pull request was merged in for this ticket. Can this be closed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2944
16743	ACCUMULO-2943	This issue is marked "In Progress". What's the status? I see the pull request was closed, but it doesn't look like it was applied.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2943
16744	ACCUMULO-2942	Thanks for the patch Hayden! Would you like to be listed on [our contributors page|http://accumulo.apache.org/people.html]? If so, what would you like your employer listed as? What timezone?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2942
16745	ACCUMULO-2936	Would anyone object to creating a "Dos and Don'ts" development document to help address this? One thing I think that's going to become onerous is a testing plan that involves all of our tests for each implementation of the JDK we deem fit for Accumulo's use.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2936
16746	ACCUMULO-2936	[~bills] we should probably have this discussion on the dev@ list rather than here. Could you start it?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2936
16747	ACCUMULO-2934	Bill, did you test full ITs or a live install?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2934

#	Requirement	Questions	Answered	Successfully Answered	URL
16748	ACCUMULO-2933	bq. If the intention is to run performance tests to compare releases These comments, ACCUMULO-2368, and ACCUMULO-2766 have me desiring a test suite that times many different Accumulo operations. It would be so nice to have this inorder to detect to detect regressions. Accumulo has a lot of test for ensuring various operations are correct, but not much in the way of performance. Could time operations like the following and produce a report that could be compared w/ reports from other releases. * time creating tables * time deleting tables (w/ data, w/o data, w/ data in memory, w/ compactions running, w/ scans running) * time splitting (empty table, table w/ files) * time merging * time cloning (empty table, table w/ files) * time exporting a table * time importing a table * time bulk import * time group commit * etc	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2933
16749	ACCUMULO-2932	It would be preferable to prevent tablets that have different durability configurations from using the same WAL. Presuming that the user configuring things tunes down durability on some table because they need the performance gain. Should we really *always* sync the WALs involved in accumulo.metadata? I can see defaulting to it or even requiring that changing it be done at the table level (presuming we still ahve some system-wide default used for tables). But what if the user has stable backup power and isn't worried about other sources of sudden machine loss?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2932
16750	ACCUMULO-2932	bq. It would be preferable to prevent tablets that have different durability configurations from using the same WAL. Presuming that the user configuring things tunes down durability on some table because they need the performance gain. This would be much easier if a tablet could use multiple WALs (ACCUMULO-1083), but that's a difficult problem in itself. Specifying the explicit file name log in each tablet row does make this easier to track, but it's not required to do what it outlined here. bq. Should we really always sync the WALs involved in accumulo.metadata? I can see defaulting to it or even requiring that changing it be done at the table level (presuming we still ahve some system-wide default used for tables). But what if the user has stable backup power and isn't worried about other sources of sudden machine loss? Sure, that's a valid point.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2932
16751	ACCUMULO-2927	Yeah, I was trying to avoid all global ignores (notice that I prefixed the IDE ones with /), so nothing unusual (like IDE files outside of any actual module) would be hidden from users. Should I push or did you want to make changes first?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2927
16752	ACCUMULO-2916	Are there features that you would like to see in ZK 3.4? AFAIK, there are no serious bugs with 3.3.6 and it is very stable.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2916
16753	ACCUMULO-2913	up high enough) would be to find an invocation where we passed in a null SecurityErrorCode into AccumuloSecurityException (which is a lot of places). I'll try to fix the ultimately cause if/when I see it again. {noformat} Error Message org.apache.accumulo.core.client.AccumuloSecurityException: Error DEFAULT_SECURITY_ERROR for user root - Unknown security exception Stacktrace java.lang.RuntimeException: org.apache.accumulo.core.client.AccumuloSecurityException: Error DEFAULT_SECURITY_ERROR for user root - Unknown security exception at org.apache.accumulo.core.tabletserver.thrift.TabletClientService$startMultiScan_result$startMultiScan_resultStandardScheme.read(TabletClientService.java:8613) at org.apache.accumulo.core.tabletserver.thrift.TabletClientService$startMultiScan_result$startMultiScan_resultStandardScheme.read(TabletClientService.java:8590) at org.apache.accumulo.core.tabletserver.thrift.TabletClientService$startMultiScan_result.read(TabletClientService.java:8529) at org.apache.thrift.TServiceClient.receiveBase(TServiceClient.java:78) at org.apache.accumulo.core.tabletserver.thrift.TabletClientService$Client.recv_startMultiScan(TabletClientService.java:311) at org.apache.accumulo.core.tabletserver.thrift.TabletClientService$Client.startMultiScan(TabletClientService.java:291) at org.apache.accumulo.core.client.impl.TabletServerBatchReaderIterator.doLookup(TabletServerBatchReaderIterator.java:632) at org.apache.accumulo.core.client.impl.TabletServerBatchReaderIterator.doLookup(TabletServerBatchReaderIterator.java:592) at org.apache.accumulo.core.metadata.MetadataLocationObtainer.lookupTablets(MetadataLocationObtainer.java:181) at org.apache.accumulo.core.client.impl.TabletLocatorImpl.processInvalidated(TabletLocatorImpl.java:667) at org.apache.accumulo.core.client.impl.TabletLocatorImpl._locateTablet(TabletLocatorImpl.java:610) at org.apache.accumulo.core.client.impl.TabletLocatorImpl.binRanges(TabletLocatorImpl.java:281) at org.apache.accumulo.core.client.impl.TabletLocatorImpl.binRanges(TabletLocatorImpl.java:356) at org.apache.accumulo.core.client.impl.TabletLocatorImpl.processInvalidated(TabletLocatorImpl.java:660) at org.apache.accumulo.core.client.impl.TabletLocatorImpl._locateTablet(TabletLocatorImpl.java:610) at org.apache.accumulo.core.client.impl.TabletLocatorImpl.locateTablet(TabletLocatorImpl.java:440) at org.apache.accumulo.core.client.impl.ThriftScanner.scan(ThriftScanner.java:226) at org.apache.accumulo.core.client.impl.ScannerIterator$Reader.run(ScannerIterator.java:84) at org.apache.accumulo.core.client.impl.ScannerIterator.hasNext(ScannerIterator.java:177) at org.apache.accumulo.core.client.impl.ScannerIterator.next(ScannerIterator.java:214) at org.apache.accumulo.core.client.impl.ScannerIterator.next(ScannerIterator.java:48) at com.google.common.collect.Iterators.getOnlyElement(Iterators.java:325) at com.google.common.collect.Iterables.getOnlyElement(Iterables.java:271) at org.apache.accumulo.test.replication.CyclicReplicationIT.dataIsNotOverReplicated(CyclicReplicationIT.java:224) {noformat}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2913
16754	ACCUMULO-2913	Commit 643aaed80039f870a247f8c8e167152fd8b52b76 in accumulo's branch refs/heads/master from [~elserj] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=643aaed] ACCUMULO-2913 More replication test stabilization. Remove commented code, retry some checks for replication data a few times before failing the test, and remove the defaultTimeoutSeconds because it was overriding what was actually configured in the JUnit timeout annotation.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2913
16755	ACCUMULO-2913	Is there more work to be done on this? The sub-tasks are complete, and there are several commits already logged for this issue.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2913
16756	ACCUMULO-2905	bq. the core issue is that the memory manager is not being called because the mechanism that will trigger it only runs the manager once a second. Can you expand on this, [~ecn]? I see the block that [~vines] was talking about (250ms), but I don't see how the manager would only be invoked once a second. It appears to me that it will start running immediately after construction of the class and then every 250ms (as mentioned).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2905
16757	ACCUMULO-2905	What loop?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2905
16758	ACCUMULO-2905	bq. What loop? Pretty much the stuff after {noformat} if (startMinC) { {noformat}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2905

#	Requirement	Questions	Answered	Successfully Answered	URL
16759	ACCUMULO-2889	bq. We could just add all the hosted tablets to the log when it is new (common), and when new tablets are hosted (rare). Could also try to be a little smarter and keep some basic measure of "recently written to" tablets and update the log entry there instead of all tablets. That would help the write-once, read-many case that's common. bq. This would cause unnecessary recoveries in the case where a tablet is flushed or isn't otherwise using the log. I'm not entirely clear on what performing log recovery for a tablet using a WAL that doesn't have any updates for that tablet entails. Would we just be sorting and re-reading that WAL to not find any updates that need to be replayed? Do you think we would incur a noticeable penalty for doing this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2889
16760	ACCUMULO-2886	Very nice! Much better than before. Also, noticed the following included in your patch {code} --- a/pom.xml +++ b/pom.xml @@ -708,7 +708,7 @@ <plugin> <groupId>org.codehaus.mojo</groupId> <artifactId>exec-maven-plugin</artifactId> - <version>1.3.2</version> + <version>1.4.0</version> </plugin> <plugin> <groupId>org.eclipse.m2e</groupId> {code} Did you mean to commit that as well? Seems unrelated.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2886
16761	ACCUMULO-2883	bq. there may be room for some sort of permission model on metadata if we plan on exposing this more publicly. I was thinking about this problem this morning. Right now, anyone with a minor bit of cunning could read the metadata table and get the same information (because we can't restrict read access to it), right? Is that a fundamental problem we just have to accept? How could an extra permission layer on top help us increase the total security of the system? I'm not sure presently :)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2883
16762	ACCUMULO-2877	Could you reformat the patch so that it only contains the fix for ACCUMULO-2877? At the moment it also includes two previous commits.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2877
16763	ACCUMULO-2877	Is there a way we can fold this functionality into a part of the public API? I'm all for easing transitions (e.g. by leaving this in place for 1.5.x and 1.6.x), but I don't want us to get into the habit of leaving non-public API code in place because it's being used externally.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2877
16764	ACCUMULO-2877	bq. If the client API does not call CachedConfiguration, then the method is not necessary. We could re-open ACCUMULO-43. Would it make sense for public API methods that interact directly with HDFS ({{TableOperationsImpl#importDirectory}}) to accept a Configuration object? We could add that support moving forward, deprecate the version that relies on CachedConfiguration, and then have a clear time when we can remove CachedConfiguration. It looks like OfflineScanner could also benefit from the same approach (although we could create a Configuration with the relevant information from accumulo-site.xml too).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2877
16765	ACCUMULO-2866	Can you explain what exactly "makes sense"?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2866
16766	ACCUMULO-2866	The WAL should be a single block. A recovery file might be made into multiple smaller blocks... is that what you're thinking about?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2866
16767	ACCUMULO-2861	I was looking around in the code and found one of the things I have bumped into before. Table permissions, system permission, and namespace permission are encoded using the binary ordinals of the enum. So when dumping zookeeper a lot of base64 shows up instead of things like READ,WRITE, etc. In 1.6 the code that does this encoding is in {{org.apache.accumulo.server.security.handler.ZKSecurityTool}}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2861
16768	ACCUMULO-2861	bq. we'd get better performance by just creating a POJO and using jackson for serialization. Is that opinion or fact? The reason I was leaning towards thrift or protobuf is that we can use them to ensure that we're backwards compatible. Ensuring that with POJOs would be a bit more difficult and error-prone.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2861
16769	ACCUMULO-2855	A lookup attempt could be fine... I just find them to be problematic, because they make too many assumptions, and those assumptions may be completely invalid. Which NIC are you going to use? Which IP, if the NIC has multiple? Which hostname do you use if DNS reports multiple? What if the hostname known locally is not the one used publicly? Worse, we have to maintain the code that implements these assumptions, while artificially limiting users for whom these assumptions are invalid. And, the invalidity of these assumptions is possibly quite frequent, due to infrastructure like EC2, where users have separate private IP address space and public IP address space, and a DNS that they may have less control over. Of course, we could do both: allow explicit configuration and default to some heuristic... but then, we're adding a lot more complexity and potential for confusion for users that makes debugging harder. This is why I think software like httpd does it best when they just make it explicitly configurable (as in the Listen and ServerName directives in httpd.conf).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2855
16770	ACCUMULO-2855	bq. it might be better to extend MAC with a more generic cluster. I think Josh Elser was interested in doing something like that. I think for MAC's case, you'd still just want to provide a single hostname/addr to use for binding since its intent is to run solely on a single host. Things change in a lot of other ways too when you start thinking about using other deployment systems to programmatically run a cluster. bq. This is why I think software like httpd does it best when they just make it explicitly configurable (as in the Listen and ServerName directives in httpd.conf). I'd agree here. How about allowing a hostname/IP value to be used for MAC, default it to localhost so it continues to work as it does now, and then let the resolution happen as the network is configured? That would keep up from making any weird assumptions and satisfy Bill's original desire.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2855
16771	ACCUMULO-2855	bq. How about allowing a hostname/IP value to be used for MAC, default it to localhost so it continues to work as it does now, and then let the resolution happen as the network is configured? That's basically what I was thinking, too.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2855
16772	ACCUMULO-2851	We definitely wouldn't want to change how this functionality works, and, like Eric, I think removing them is the common case. However, I could see us adding something to the API which will not delete them if the user sets that option. Would that be satisfactory?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2851
16773	ACCUMULO-2844	Why make the change in the bug fix branches?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2844
16774	ACCUMULO-2844	I'd rather not debate if "master" is still offensive. I initially had thought Leader would work. When I started going through what the role does, I wasn't as sure. * Manage tablet assignments / balancing * Recovery coordination * Run FATE ops * Respond to Admin API commands * Handle restructuring ZK / write to HDFS stored metadata during upgrades Anything else it does?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2844
16775	ACCUMULO-2844	The biggest technical concerns I have are the backwards compat, which you plan to address (yay). I worry a bit about user confusion on these branches as well (I used a "slaves" file previously, but now I can use a "slaves" file or a "tservers" file, _why can I use either?_, _do they do different things?_, etc). I would rather the terminology be changed only in master, but I wouldn't -1 it if this was applied in earlier versions in a compatible manner. Conductor or Coordinator work for me as a new name for the Master. Something that doesn't imply that tservers must talk to the master before performing any action or similar.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2844

#	Requirement	Questions	Answered	Successfully Answered	URL
16776	ACCUMULO-2844	This is a good change to make for 1.7. Making the change in 1.5 and 1.6 seems like it would require a lot of extra work and testing (that would not be needed in 1.7) in order to ensure it does not break things for users. Also as [~bills] mentioned, the old names will still be present in 1.5 and 1.6 for compatibility reasons. I really like [~billie.rinaldi] suggestion for 1.5 and 1.6. bq. It may also be worth evaluating whether "master" is still offensive in the absence of "slaves". Renaming the slaves file to tservers will likely not cause users problems and it seems to solve the problem at hand. The term master by itself does not seem offensive, its a very widely used and uncontroversial term. While master by itself does not seem offensive, it also is not very accurate. Common uses are related to someone who is an expert. Renaming the slaves files leaves us with something thats just poorly named. Engineers often poorly names things. I really like a lot of the other names that have been suggested. As for renaming things in the master branch and not in 1.5-SNAPSHOT or 1.6-SNAPSHOT, git handles this really well. If the change were made in 1.5 and 1.6, would we post two sets of documentation on the web site? One for 1.6.0 and one for 1.6.1 using the different terminology?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2844
16777	ACCUMULO-2844	Haha, the timing is funny with both you revisiting this at the same time. I agree with both... this would be great for 2.0. Will look at the PR more closely tomorrow if I get a chance (unless it gets merged in before that... in which case, awesome). I did see that the PR didn't address it in the replication code. Josh, are you going to handle that part?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2844
16778	ACCUMULO-2844	bq. I did see that the PR didn't address it in the replication code. Josh, are you going to handle that part? Oh drat, I introduced more in there? I tried to use the terminology "peer". I'll have to grep. bq. I agree with both... this would be great for 2.0. Will look at the PR more closely tomorrow if I get a chance (unless it gets merged in before that... in which case, awesome). I just looked at the latest from Mike and think it looks fine to me. I will grep around to see what's there for replication, but am leaning towards just merging it in today.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2844
16779	ACCUMULO-2844	bq. I will grep around to see what's there for replication I see that some references sneaked in via the documentation. Why don't I just break out some sub-tasks and merge in Mike's stuff now. We can work on this in tandem.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2844
16780	ACCUMULO-2844	I think this is sufficiently done, and can be closed. Is there more to be done?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2844
16781	ACCUMULO-2844	bq. I think this is sufficiently done, and can be closed. Is there more to be done? Is there interest to also rename the "master" Accumulo process? Personally, I would like to see s/master/coordinator/ and s/garbage collector/janitor/, but I'm not sure what others think could/should be done for 2.0. In writing/speaking, I think both of these renames would remove a bit of ambiguity (e.g. does "master" refer to the "accumulo master" or a special physical node which different hardware than the other nodes? does garbage collection refer to the JVM or the accumulo "garbage collector" process?), and help ignore the previous debate of "\{offensiveness with the lack of slaves\}". There was concern about API compatibility also previously stated with such a rename. At a quick glance, I only see {{getMasterStats}} on the thrift RPC method side. As everyone knows, Thrift is great about letting us build cross-version compatibility between clients/servers. I would think that we could rename most instances, and leave those necessary for backwards compatibility there. I believe that we can change the general tone to something that provides greater clarity and reduces the likelihood to offend someone.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2844
16782	ACCUMULO-2842	What is the relationship between this setting and {{dfs.datanode.synconclose}}? If I set this to 'hflush' then do I need to worry about the datanode settings, or are they orthogonal?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2842
16783	ACCUMULO-2841	[~ctubbsii], if you're going to close a ticket for a patch that had a subset of the features mentioned, can you please open and link a follow on ticket and edit the description for what the current ticket actually provided. This patch only provides a place for information on a table. I imagine that it also works for namespaces, but there are no tests for it. And there's no ability to watch for changes.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2841
16784	ACCUMULO-2841	[~vines]: Yes, I can, and I would have had I recognized it to be the case that the patch only provided a subset. However, I did not interpret the work done to be a subset of the requested features, since I requested them, and all my expectations were satisfied. If they had not been, I would have commented accordingly on Jenna's pull request. There was only one other commenter (you), who I could have had any reasonable suspicion of having different expectations, but I did not interpret your comment back in June (phrased as *I'm wondering if we should...*), to represent a blocker or requirement for this issue, especially since you did not follow up in comments later about the original idea being metadata-oriented (which would have made your suggestion much, much, harder). That is the only feature I can imagine you mean was omitted, when you say "subset". I did, just now, create ACCUMULO-3121, in regards to your comment. While I didn't consider it within scope of the original identified issue here, I liked it enough to support a ZK solution to this issue so that future work could be done if somebody were interested in creating a ticket and pursuing that. However, I wouldn't have thought it to be my specific responsibility to create that issue just because the issue in which it was mentioned was closed by me, as you've implied. Anybody interested in that feature could have done that, but I've created it as requested, regardless. Since it was me, I hope the wording accurately represents your intentions/expectations for that follow-on work (because I didn't/can't know precisely what you had in mind when you first speculated about it in the comments above). Regarding namespace and permissions testing, all per-table config already works for namespaces, and that's tested separately, as well as tests for permissions to set properties. Those tests continue to pass after this patch. Because the implementation was a minor extension to an existing feature, in which those things are already tested, I thought the patch's provided test, which is more narrowly scoped to testing the provided change set, was sufficient. Had the implementation been a different one, in which those additional requirements were provided by a new mechanism, certainly, I'd have required tests for those new mechanisms before signing off on the patch. This brings up an interesting question: *is a patch for a new feature, which leverages existing features, expected to test the existing features because it relies on them to satisfy the new feature's requirements?* If the answer is "yes, always", it seems we could spend a lot of time writing redundant tests, which don't provide any additional meaningful code coverage. I think the answer is probably "sometimes", and I'd probably lean towards testing those existing features when they are used in a complex way, not covered by the existing test suite for those old features. I don't think that's the case here, though. This new feature isn't complex and doesn't change anything with regard to per-table configuration permissions or the use of per-table configuration items in namespace configs. Those features already are tested and work in the general case (regardless of which property is being set), so they should work in this case, also (setting a property with this new prefix). I don't think any additional tests for this are essential. If you still feel those additional tests are necessary, could you please create a sub-task, though?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2841
16785	ACCUMULO-2837	Dropping priority and changing from "Bug" to "Improvement", because it doesn't appear to represent an essential change. [~kturner], is there more to be done on this? All the sub-tasks appear to be complete. Was closing the parent simply overlooked, or did you want to add a sub-task to address the gzipping?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2837

#	Requirement	Questions	Answered	Successfully Answered	URL
16786	ACCUMULO-2827	This seems to be a pareto improvement on the previous HeapIterator, and it doesn't change any interfaces. Is anyone opposed to getting this into 1.5.2 and 1.6.1?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2827
16787	ACCUMULO-2827	The benchmark graph certainly looks promising and the changes seem straightforward enough. Have you considered running some before/after marks using continuous-ingest? Assuming we see an improvement in practice, I'd likely be in favor of including this in the 1.5 and 1.6 lines.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2827
16788	ACCUMULO-2827	[~parkjsung] bump ^^ -- Have you run any "end-to-end" benchmarks with these changes?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2827
16789	ACCUMULO-2827	Did you only do a single compaction using each of the old and new iterators, or is that a single sample intended to be representative of things? The most thorough thing might be to have a profiler hooked up and see how long the HeapIterator takes, so that we can ignore things like shell start-up time (which is often non-trivial).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2827
16790	ACCUMULO-2827	This was only a single compaction using each of the old and new iterators. Any more details on hooking up a profiler and seeing how long the HeapIterator takes? By profiler, do you mean something like jvisualvm? Do you want us to try and profile a major compaction as its running in a tserver vs a test harness? How would you like the profiler hooked up? Typically we've found that attaching profilers to the iterators greatly affect the performance. It should still show the benefit from this change though. Would those results be valid?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2827
16791	ACCUMULO-2827	It's interesting that all of the single column tests were slower. Is there a reason for this? Is there an ability to detect a single column in the RFile index and switch implementation at runtime?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2827
16792	ACCUMULO-2827	Yes, yes I was. [~kturner] - Looks like you pushed? Can the JIRA be marked as resolved?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2827
16793	ACCUMULO-2825	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/7#issuecomment-43547788 @ryaneleary I know you had mentioned that you wanted this abstraction point for a reason, but I don't recall what it was. Can you elaborate on the types of operations that this abstract class provides, some use case that the WholeRowIterator didn't sufficiently provide, and, if you have it, maybe an example class? That might help cement this for when/why a user would want to extend RowEncodingIterator	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2825
16794	ACCUMULO-2817	As for creating new issues, I wanted to see how the discussion here went before doing that. Was thinking of adding new decode method to the current implementations w/o changing the interface. Code like the following using the concrete implementations could still be written, which seems useful. In the future when we switch to Java8 we can modify the interface. What are the pitfalls of this approach? {code:java} LongLexicoder longLex = new LongLexicoder(); long l = longLex.decode(data, 5, 13); {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2817
16795	ACCUMULO-2817	Realized I need to add in bounds checking to all of the {{decode(byte[], int, int)}} impls. So, two design questions: 1. Bounds checking wouldn't be necessary for when {{decode(byte[])}} calls the other method, so would it be a good idea to do the following, or is this over-optimizing? {code:java} @Override public T decode(byte[] b) { return decodeUnchecked(b, 0, b.length); } /** @since 1.7.0 */ public T decode(byte[] b, int offset, int len) throws ValueFormatException { if (offset + len > b.length) { throw new ValueFormatException(String.format("offset + length %d exceeds byte array length %d", (offset+len), b.length)); } else { return decodeUnchecked(b, offset, len); } } /** @since 1.7.0 */ public T decodeUnchecked(byte[] b, int offset, int len) { ... } {code} 2. Do we agree it would be prudent to make a {{public abstract class AbstractEncoder<T> implements Encoder<T>}} that has the above boilerplate {{decode(byte[], int, int)}} and {{decode(byte[])}} methods?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2817
16796	ACCUMULO-2817	Rebased onto master at revision {{8f6558c}} This patch has my ideas on how abstract classes could be used. * Removed {{decode(byte[], int, int)}} from the Encoder (and thus also Lexicoder) interfaces, leaving them the same as on master. * Added {{org.apache.accumulo.core.client.lexicoder.AbstractEncoder}}. It implements {{decode(byte[])}} (required by Encoder interface), and {{decode(byte[], int, int)}} (not part of an interface). It then defines {{protected abstract T decodeUnchecked(byte[] b, int offset, int len)}}. This does not change the Encoder or Lexicoder interfaces. * Added {{org.apache.accumulo.core.client.lexicoder.AbstractLexicoder}}, which extends AbstractEncoder and also implements Lexicoder. * Refactored Lexicoders to extend AbstractLexicoder, and Encoders to extend AbstractEncoder. I think AbstractEncoder and AbstractLexicoder should be part of the public API since they may be useful to others, and do not cause any breaking changes. I looked through semver, and I don't think "any backwards incompatible changes are introduced to the public API," but feel free to correct me. *Question:* Do I add "@since" to any public methods, or only ones implementing an Interface, or all new methods?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2817

#	Requirement	Questions	Answered	Successfully Answered	URL
16797	ACCUMULO-2817	{code} diff --git a/core/src/main/java/org/apache/accumulo/core/replication/proto/Replication.java b/core/src/main/java/org/apache/accumulo/core/replication/proto/Replication.java index 030f4ec..72f706a 100644 --- a/core/src/main/java/org/apache/accumulo/core/replication/proto/Replication.java +++ b/core/src/main/java/org/apache/accumulo/core/replication/proto/Replication.java @@ -449,6 +449,17 @@ package org.apache.accumulo.core.replication.proto; throws com.google.protobuf.InvalidProtocolBufferException { return PARSER.parseFrom(data, extensionRegistry); } + public static org.apache.accumulo.core.replication.proto.Replication.Status parseFrom(+ byte[] data, int offset, int len) + throws com.google.protobuf.InvalidProtocolBufferException { + return PARSER.parseFrom(data, offset, len); + } + public static org.apache.accumulo.core.replication.proto.Replication.Status parseFrom(+ byte[] data, int offset, int len, + com.google.protobuf.ExtensionRegistryLite extensionRegistry) + throws com.google.protobuf.InvalidProtocolBufferException { + return PARSER.parseFrom(data, offset, len, extensionRegistry); + } public static org.apache.accumulo.core.replication.proto.Replication.Status parseFrom(java.io.InputStream input) throws java.io.IOException { return PARSER.parseFrom(input); {code} This is generated code. You don't need to be changing it. {{StatusCombiner}} will just have to be changed to call the necessary method on {{Status}} if no methods are exposed already. {code} try { Preconditions.checkNotNull(b, "cannot decode null byte array"); } catch (NullPointerException e) { throw new ValueFormatException(e); } try { Preconditions.checkArgument(offset >= 0, "offset %s cannot be negative", offset); } catch (IllegalArgumentException e) { throw new ValueFormatException(e); } try { Preconditions.checkArgument(offset + len < b.length, "offset + length %s exceeds byte array length %s", (offset + len), b.length); } catch (IllegalArgumentException e) { throw new ValueFormatException(e); } {code} You don't need to catch the Exceptions thrown by Preconditions. They're RuntimeExceptions -- just let them filter up. bq. Do I add "@since" to any public methods, or only ones implementing an Interface, or all new methods? Adding it to AbstractEncoder as you did is good. I think AbstractLexicoder is better off being in {{org.apache.accumulo.core.client.lexicoder.impl}} which would not need to have the annotation. Generally speaking, if you're introducing a new class into the [public API|https://github.com/apache/accumulo/blob/master/README.md#api], you can just add it at the class level and you don't need the method-level annotations. When new methods are added into that class later, {{since}} would be added on those methods. {code} /** * Calls {@link #decodeUnchecked(byte[], int, int)} as {@code decodeUnchecked(b, 0, b.length);}. * Does not check if b is null. */ @Override public T decode(byte[] b) { return decodeUnchecked(b, 0, b.length); } {code} Why not check that {{b}} is non-null? bq. I looked through semver, and I don't think "any backwards incompatible changes are introduced to the public API," but feel free to correct me. The concern is that a user who implemented their own Lexicoder against 1.6.x would suddenly stop working with 1.7.0. The way you currently did it looks fine to me (Lexicoder hasn't ultimately changed only the implementations in such a way that only adds a new method).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2817
16798	ACCUMULO-2817	bq. Hi folks, I'm interested in calling decodeUnchecked(byte[], off, len) as public API, since it removes the need to copy a subarray out from a larger array. Can we pull this discussion into another JIRA issue, [~dhutchis]? It would get confusing to track since this changeset was already released in 1.7.0.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2817
16799	ACCUMULO-2815	For comparison, do you have the numbers for SSL?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2815
16800	ACCUMULO-2815	That's okay. I was just thinking SSL in Accumulo. But, now that you mention it, I was kind of assuming HDFS/ZK were set up the same on all these scenarios (either with or without Kerberos). Are you suggesting that is not the case? That the HDFS settings are also configured with the above Kerberos settings (or none)?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2815
16801	ACCUMULO-2815	In other words, do these numbers reflect Accumulo's overhead with Kerberos, or do they represent the overhead of an entire system configured end-to-end with Kerberos?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2815
16802	ACCUMULO-2815	What do the headings mean?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2815
16803	ACCUMULO-2808	[~ctubbsii] - Do you need any more feedback on this before applying it?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2808
16804	ACCUMULO-2803	This looks good. Can you submit a patch generated using {{git format-patch}} so that the commit attributions are correct? Thanks!	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2803
16805	ACCUMULO-2799	I'm on to my second implementation of a "drain" operation which a client can invoke to know "when the data has been replicated" for a table. There are a couple of caveats here that make this difficult. * You might get a new WAL for a table even if no data has been ingested into it (WALs are per-tserver, not per-table) * How do you implement it with constant ingestion? Prevent the user from making a call that will never complete. (Is that intended?) What I've done for the present is take a snapshot of the current WALs in use by a table (using the metadata table and the 'log' colfam) and wait for all of those WALs to be fully replicated. This at least addresses both of the above points, but I'm still not positive if this is the "best" semantics for it. Thoughts, anyone else?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2799
16806	ACCUMULO-2797	Are there any non-English speaking people who could support this effort? I, unfortunately, am a mono-linguist.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2797
16807	ACCUMULO-2792	Vikram, can you check if this first change in 1.5 or in 1.6?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2792
16808	ACCUMULO-2789	[~bills] do you have more here that you're planning on doing? Do we want to add any higher-level documentation for this? I don't recall off the top of my head if we have docs elsewhere for continuous-ingest and randomwalk that we could add this to.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2789
16809	ACCUMULO-2789	Commit f0dcf1e5570c12098e716ec58ea2f4bc3321053b in accumulo's branch refs/heads/1.5.2-SNAPSHOT from Bill Slacum [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=f0dcf1e] ACCUMULO-2789 Added parameter to control the number of cells within a row When running the stress test, it is now possible to configure the writer to send mutations for rows that contain a random amount of column updates. This is controlled via the --min-row-width, --max-row-width and --row-width-seed parameters. I did some internal shuffling to make randomly generating numbers within a range more convenient. See the new RandomWithinRange class for details.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2789
16810	ACCUMULO-2789	Commit f0dcf1e5570c12098e716ec58ea2f4bc3321053b in accumulo's branch refs/heads/1.6.1-SNAPSHOT from Bill Slacum [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=f0dcf1e] ACCUMULO-2789 Added parameter to control the number of cells within a row When running the stress test, it is now possible to configure the writer to send mutations for rows that contain a random amount of column updates. This is controlled via the --min-row-width, --max-row-width and --row-width-seed parameters. I did some internal shuffling to make randomly generating numbers within a range more convenient. See the new RandomWithinRange class for details.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2789

#	Requirement	Questions	Answered	Successfully Answered	URL
16811	ACCUMULO-2789	Commit f0dcf1e5570c12098e716ec58ea2f4bc3321053b in accumulo's branch refs/heads/master from Bill Slacum [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=f0dcf1e] ACCUMULO-2789 Added parameter to control the number of cells within a row When running the stress test, it is now possible to configure the writer to send mutations for rows that contain a random amount of column updates. This is controlled via the --min-row-width, --max-row-width and --row-width-seed parameters. I did some internal shuffling to make randomly generating numbers within a range more convenient. See the new RandomWithinRange class for details.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2789
16812	ACCUMULO-2784	Does this ticket carry with it any particular design of experiments requirements? e.g. workload, iterations, etc?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2784
16813	ACCUMULO-2781	The stack trace is: {code:java} java.io.IOException: Can't get Master Kerberos principal for use as renewer at org.apache.hadoop.mapreduce.security.TokenCache.obtainTokensForNamenodesInternal(TokenCache.java:116) at org.apache.hadoop.mapreduce.security.TokenCache.obtainTokensForNamenodesInternal(TokenCache.java:100) at org.apache.hadoop.mapreduce.security.TokenCache.obtainTokensForNamenodes(TokenCache.java:80) at org.apache.hadoop.mapreduce.lib.output.FileOutputFormat.checkOutputSpecs(FileOutputFormat.java:142) at org.apache.hadoop.mapreduce.JobSubmitter.checkSpecs(JobSubmitter.java:456) at org.apache.hadoop.mapreduce.JobSubmitter.submitJobInternal(JobSubmitter.java:342) at org.apache.hadoop.mapreduce.Job$10.run(Job.java:1268) at org.apache.hadoop.mapreduce.Job$10.run(Job.java:1265) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.Subject.doAs(Subject.java:415) at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1491) at org.apache.hadoop.mapreduce.Job.submit(Job.java:1265) at org.apache.hadoop.mapreduce.Job.waitForCompletion(Job.java:1286) at org.apache.accumulo.core.client.mapreduce.AccumuloFileOutputFormatTest$MRTester.run(AccumuloFileOutputFormatTest.java:154) at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:70) at org.apache.accumulo.core.client.mapreduce.AccumuloFileOutputFormatTest$MRTester.main(AccumuloFileOutputFormatTest.java:163) at org.apache.accumulo.core.client.mapreduce.AccumuloFileOutputFormatTest.writeBadVisibility(AccumuloFileOutputFormatTest.java:195) {code}	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2781
16814	ACCUMULO-2762	bq. it could cause problems That's adequately, vaguely obtuse... Context?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2762
16815	ACCUMULO-2760	The intent of the test is to have the server's I/O hang, and have the master status monitoring fail. The master should remove the tserver's lock, which should kill it. We're seeing the tserver killed, but it seems to be happening because it is missing zookeeper deadlines, and not due to any action by the master. At least, not always.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2760
16816	ACCUMULO-2739	Urk, looks like {{Halt.halt(int, Runnable)}} is going to be a pain to deal with. I'm a little worried about changing code that was added to try to make these processes exit more gracefully. Since these are just running in their own threads (for the purposes of this issue), I would think that we have to ensure that Interrupting each process (thread) needs to then exit gracefully. Is there a better way to try to do this? (cc/ [~bhavanki])	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2739
16817	ACCUMULO-2739	[~elserj] do you know that spawning processes is whats making MAC slow? I am sure thats part of it, but is it the largest contributor? For example the interaction between master, root table and metadata table inorder to get first user tablet online could possibly be much slower than spawning processes. I looked into this type of thing a long time ago, but can not remember much. ACCUMULO-879 are an example of problems that have nothing to do w/ separate processes.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2739
16818	ACCUMULO-2739	Just for argument's sake, I'll mention Docker because it is supposed to start nodes in milliseconds. I know that people use Docker for Hadoop. Has anyone here used Docker?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2739
16819	ACCUMULO-2739	I'm not quite sure what you mean, [~medined]. Are you just stating that starting up fast is desirable? I don't think running Accumulo inside of docker containers is really relevant for what this outlines...	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2739
16820	ACCUMULO-2736	I stand corrected. If that's the case, then I'm not sure the distinction between "Fixed", "Done", and "Implemented"... unless we just want to use "Done" for tasks, and "Fixed" for bugs, and "Implemented" (did you verify that one?) for "New Feature" and "Improvement". If that's the case, then "Done" and "Implemented" *should* have fixVersions on them properly.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2736
16821	ACCUMULO-2732	I'd rather not see the README become a tutorial on maven, especially if it starts making assumptions about a user's build environment (maven version, maven's own resource requirements, or jvm execution paths, etc.). Even in build instructions, we have to make some assumptions that a user is familiar with the build tools or can debug build tool problems with documentation on the build tool itself. MAVEN_OPTS is not an accumulo build option, it's a maven configuration option, so I don't think it should go here. I agree with [~busbey] that this would be more suitable in a contributor guide (specifically, one on the website, or in a wiki).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2732
16822	ACCUMULO-2732	I'll pick this up since I suggested an avenue :) What else besides MAVEN_OPTS are some notable "gotchas" that we would need to include? Could point out the minimum maven version (defined in the pom), as well as supported/recommended JDK.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2732
16823	ACCUMULO-2732	I have a coworker who can reliably reproduce having the core/security/crypto/BlockedIOStreamTest fail on JDK6 with default maven surefire JVM settings, regardless of manipulating the MAVEN_OPTS (since those aren't actually used for forked surefire tests) and having JDK7 installed. The test always passes for me, I suspect because of a combination of default JVM max heap and because my system has no jdk6 (while hers does), resulting in my maven using a JRE7 for running tests even though it's targeting java 6. While specifying a max heap, I have all tests passing at Xmx 756MB while she requires Xmx 1GB [~elserj] (or anyone else) would you mind if she takes this over to contribute a pom fix for the surefire plugin configs? Or would you prefer a new ticket?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2732
16824	ACCUMULO-2731	This is marked as "Fixed", but no fixVersion or commit comments. Was this actually done?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2731
16825	ACCUMULO-2712	Can you give me a pointer as to where to track ASF-17?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2712
16826	ACCUMULO-2705	Shouldn't this just happen as a part of the upgrade process? How did the 1.5 -> 1.6 upgrade create the root table?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2705
16827	ACCUMULO-2705	SystemCredentials really doesn't seem to be the "correct" place for this either as they are explicitly meant for IPC. Should Accumulo internals have the ability to use the public API? I suppose I could try to circumvent things by hitting the {{TableOperations.create(String)}} implementation as a shortcut.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2705

#	Requirement	Questions	Answered	Successfully Answered	URL
16828	ACCUMULO-2689	Is it a good idea to be working around what appear to be bad system configuration issues? You already stated on ACCUMULO-2554 that you think it plausible to be this. Are there other benefits to starting ZK in process? Are you considering this a first step to making an in-process MAC?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2689
16829	ACCUMULO-2680	Before I up the timeout, did either of you notice anything that I might have missed? I saw it in the state where the first master killed itself, the second master was waiting for the lock, saw the "bogus" update that the test submits, but didn't acquire the master lock.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2680
16830	ACCUMULO-2680	[~elserj], what timeout factor were you using?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2680
16831	ACCUMULO-2680	but could you do that change in a follow on ticket? In case the current RC passes and 1.6.0 is closed out?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2680
16832	ACCUMULO-2657	Which Classpath does this property impact? Is it the running Accumulo instance?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2657
16833	ACCUMULO-2655	Also: {noformat} java.lang.Exception: Error running node Shard.xml at org.apache.accumulo.test.randomwalk.Module.visit(Module.java:286) at org.apache.accumulo.test.randomwalk.Framework.run(Framework.java:63) at org.apache.accumulo.test.randomwalk.Framework.main(Framework.java:122) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.apache.accumulo.start.Main$1.run(Main.java:141) at java.lang.Thread.run(Thread.java:744) Caused by: java.lang.Exception: Error running node shard.CompactFilter at org.apache.accumulo.test.randomwalk.Module.visit(Module.java:286) at org.apache.accumulo.test.randomwalk.Module.visit(Module.java:255) ... 8 more Caused by: java.lang.Exception: Saw unexpected document e900000000000000 doc: [] 1397147994361 false at org.apache.accumulo.test.randomwalk.shard.CompactFilter.visit(CompactFilter.java:87) at org.apache.accumulo.test.randomwalk.Module.visit(Module.java:255) ... 9 more {noformat}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2655
16834	ACCUMULO-2655	Does this arise because of multiple walkers running concurrent at the same time? I've noticed that the "ctt" tables it creates aren't namespaced by the walker host like other tests.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2655
16835	ACCUMULO-2655	[~ecn] did all walkers freeze of hang? Can you remove the nodes in the graph that causing hang and rerun?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2655
16836	ACCUMULO-2647	Thanks Jeffrey! Do you want to be listed on the [contributors list|http://accumulo.apache.org/people.html]? If so, how would you like your employer displayed (if at all)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2647
16837	ACCUMULO-2637	Well, looks like I'm not going to be able to do anything about automation at the moment since I still can't log into ASF jenkins (continuing ldap auth problems?) For everyone else, I assume this was how it was done back in the day? Did someone just occasionally deploy snapshots manually?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2637
16838	ACCUMULO-2630	Can you link to the dev list discussion? I was not able to find it after a cursory search, and think it would be good to include for posterity.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2630
16839	ACCUMULO-2627	[~mdrob], sorry this feel to the wayside. Do you have aspirations/cycles to update this before we ship out 1.7.0?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2627
16840	ACCUMULO-2627	Aspirations, yes. Cycles, no. If somebody else is able to pick up the work I had done and drive it home, that would be excellent. Are we allowed to add the Closeable interface to classes in a minor release? I'm not sure what impact that has on compatibility and am not particularly inspired to check the JLS on this.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2627
16841	ACCUMULO-2627	Looks like this doesn't apply after Eric's run-ITs-over-mapreduce. [~mdrob] plans to revisit this patch? Or should it just be closed as "Later"?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2627
16842	ACCUMULO-2613	Actually, to really drive down recovery time would we need to also ask to pin WALs for the !METADATA table? I think we would.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2613
16843	ACCUMULO-2613	That would be preferable, but couldn't a tablet server just recognize when it was adding a !METADATA related entry to the WAL and then request that it be pinned?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2613
16844	ACCUMULO-2613	How much gain would we get from caching metadata table files when we're already caching them in the tserver?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2613
16845	ACCUMULO-2613	even then, say we pinned all the WALs. how much space is that likely to be? less than 10G per node? that's not too bad on a hardware cluster.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2613
16846	ACCUMULO-2605	Commit 7ac4b54a76b85314a483378e01980c73f891a900 in accumulo's branch refs/heads/1.6.0-SNAPSHOT from [~ecn] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=7ac4b54] ACCUMULO-2605 list slf4j as a test dependency	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2605
16847	ACCUMULO-2603	This has been marked "In Progress" for over a year. What's the status of this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2603
16848	ACCUMULO-2594	Is that really relevant? It's garbage to have a class silently do the wrong thing whether or not we explicitly expose it to users. Edit: also, do we even have something that is in the public API that programmatically reports tablet locations?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2594
16849	ACCUMULO-2594	Are tableIds actually checked though? Despite the construction of a TabletLocator instance changing, I don't recall seeing this sanity check being added. The way that a TabletLocator is instantiated is unrelated to the original bug report.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2594
16850	ACCUMULO-2592	bq. That particular experimental annotation was specifically introduced for the configuration I wish you had stated that previously when I mentioned it before writing this up then :) Since you have opinions about this, I assume you also have an interest in how this looks? Such an API can be placed into the release notes for 1.6.0 for advertisement. What else do you want to discuss? Edit: Leveraging Hadoop's InterfaceStability and InterfaceAudience is likely a well suited approach here if you were to ask me. I'd assume they have a good grasp on things by now (not to mention the other projects which also leverage the annotations).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2592
16851	ACCUMULO-2592	I think it's probably a little late for 1.6.0 to begin discussing/documenting what we think "experimental" should mean in order to start annotating the API with that. What exactly does it mean to be experimental? The scope for the configuration annotations was very narrow and functional (as implementation detail to limit exposure of properties in the documentation and shell for features still in some stage of development). But, if we start using it more broadly, we're going to have to define the semantics of that annotation. Perhaps just a comment in the javadoc would be sufficient (which was my suggestion on the mailing list thread)? I'd prefer not to add more Hadoop dependencies just to borrow their annotations, and I'd also rather not blindly borrow the semantics of their annotations without at least some mailing list discussion to see if they are in line with our API goals. I'll create a sub-task for ACCUMULO-2589, to create annotations with well-defined semantics for the API, but for this, I think a simple javadoc comment is sufficient..	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2592
16852	ACCUMULO-2590	Do we need to make explicit that public inner classes are public API all the way down? i.e. if an inner class has an inner class, the inner-inner class will still be public, and so on.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2590

#	Requirement	Questions	Answered	Successfully Answered	URL
16853	ACCUMULO-2590	bq. Do we need to make explicit that public inner classes are public API all the way down? i.e. if an inner class has an inner class, the inner-inner class will still be public, and so on. That would be my expectation. Anything public or protected in an API class is part or the API. This would include methods, variables, classes, interfaces, and enums in the class. Did i miss anything?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2590
16854	ACCUMULO-2590	bq. Anything public or protected in an API class is part or the API. This would include methods, variables, classes, interfaces, and enums in the class. Did i miss anything? I think that's everything. I like that phrasing better as well.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2590
16855	ACCUMULO-2590	How's this suggested patch look for clarifying things?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2590
16856	ACCUMULO-2590	Would it be useful to also to make a statement about compatible between 1.X and 1.(x+1) releases. Deprecated methods may be dropped and methods may be added. Anything not deprecated should continue to work.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2590
16857	ACCUMULO-2589	How about we phrase it "more appropriate use of RuntimeExceptions", and debate the appropriateness of specific uses?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2589
16858	ACCUMULO-2589	Could we use subclasses of RuntimeException to allow for both [~mdrob]'s ease of failing and [~vines]'s desired ability to catch and recover (without have to do the current insanity of parsing strings)?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2589
16859	ACCUMULO-2589	I see that you are making progress on this (https://github.com/ctubbsii/accumulo/tree/ACCUMULO-2589), [~ctubbsii]. Can you give any sort of status here? Any possibility to get a WIP/feature branch up in Apache that we all can collaborate on? It would be good to make sure this doesn't start off in a bubble -- who knows, you might also get help from other devs too ;)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2589
16860	ACCUMULO-2589	bq. wired to the implementation in the core jar using java's ServiceLoader Interesting, I haven't looked at ServiceLoader at all before. A brief read-up seems like it makes a lot of sense for the future. bq. figure out how we want to do exception handling in the new API (there have been many suggestions) We should start this talk (again) sooner rather than later :). It'll be a doozie. bq. I'm also trying to find a sensible serialization practice that we can use on our "dumb" objects that doesn't depend on Hadoop libraries, but satisfies our serialization needs (in particular, credentials). This is interesting sounding, but I'm not 100% sure what you mean here. The thrift classes we have now? Do you have a separate issue for this already that we can start a discussion on? bq. As for uploading as a feature branch, I didn't want to spam the mailing lists with repeated rebase-ing, at least until it's a bit more stable and ready to contribute back to the project. I can collaborate on GitHub, though. I definitely don't want to discourage help from others... but I'd rather collaborate on GitHub until it's at least a little bit more stable and ready to contribute. My biggest regret in working on replication is how long I let it live in its own branch. This was the easy/lazy decision to make, but I'm convinced it's split me and the feature off from the rest of the community. If it lives in a bubble too long, it just becomes infeasible to bring the other devs up to speed. For replication, this was more "ok" (not saying it was acceptable) because it was an optional feature, but since the client API is of such importance, I would encourage you to start bringing things up for review. Assuming you have your own schedule for things you're working on, could you pull some code into a feature branch in Apache for things you have already worked out? For example, take the ServiceLoader code (or just look at the changes with a focus on the ServiceLoader): we can isolate the changes you've made so far for that, you can get some feedback on it, and then it would be one step closer to making it into master. Even if it changes again later (as pieces likely will), that's not the end of the world -- more than just you can help evolve the new implementation.Thoughts?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2589
16861	ACCUMULO-2589	How do you feel about about TimeUnit?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2589
16862	ACCUMULO-2573	Does something like Swift[1] make sense here? [1] https://github.com/facebook/swift	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2573
16863	ACCUMULO-2572	I meant as an example, not that we would actually hook SEP in, yes. Actually though. It's worth considering trying to build off of the replication support that's in HBase. Especially given the addition of security labels in HBase 0.98, it would be really nice if we got bi-directional HBase and Accumulo replication working. That would substantially lower the barrier for HBase users who wanted to try out Accumulo.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2572
16864	ACCUMULO-2567	Does anyone have objections to doing this?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2567
16865	ACCUMULO-2561	What are our options here? subclassing Error and throwing that? in that case we'd still run shutdown hooks, no? Do we have a design document somewhere that covers why we use the super-fast-stop-the-world version of exiting?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2561
16866	ACCUMULO-2555	I cannot duplicate this issue in my development environment or on my Jenkins environment. Attached is my best guess at a patch. [~ecn] - do you have any way to verify this before I push it?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2555
16867	ACCUMULO-2554	I am unable to reproduce this with Maven 3.0.4 or 3.1.1. Was there something else running on the port? Which OS is this failing on? Have you tried with 3.1.1 or newer Maven? Was this build running concurrently with another build that may have caused a port conflict? Was mvn executed with -T to do parallel builds across modules that both use mini?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2554
16868	ACCUMULO-2554	There shouldn't be any issues with 3.0.4. I guess this is just an issue that plagues mini in general (not being able to reserve a ZK port). I'm not sure if there's a good fix. How frequently do these happen for you? And, do you see any extra failures with the maven-plugin specifically on this (e.g. more reproducible in that module than the test module ITs)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2554
16869	ACCUMULO-2554	I don't recall ever seeing this either. I guess you have your work cut out to figure out what's happening :)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2554
16870	ACCUMULO-2554	I picked a random MAC dir from a 1.6.0-SNAPSHOT run yesterday. Looking at the following file : {noformat} /local/disk1/jenkins/workspace/accumulo16/test/target/mini-tests/org.apache.accumulo.test.DumpConfigIT_test/logs/ZooKeeperServerMain_1656483996.out {noformat} I see {noformat} 2014-04-16 17:10:10,077 [quorum.QuorumPeerConfig] INFO : Reading configuration from: /local/disk1/jenkins/workspace/accumulo16/test/target/mini-tests/org.apache.accumulo.test.DumpConfigIT_test/conf/zoo.cfg 2014-04-16 17:10:10,078 [server.ZooKeeperServerMain] INFO : Starting server {noformat} then I see {noformat} 2014-04-16 17:10:10,135 [server.NIOServerCnxn] INFO : binding to port 0.0.0.0/0.0.0.0:59228 {noformat} There are no lines in the log containing the string "dt_socket". I grepped every zookeeper out file for that string and did not see it. In the past DNS was flaky on my Jenkins machine and for some reason zookeeper would hang. I jstack the zookeeper server process and saw it was hung in java name resolution calls during initialization. Also I know there was very little in the zk out when this happened, but I can not remember exactly what. Have you been able to jstack?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2554
16871	ACCUMULO-2553	[~sonixbp] Does each reducer process a single group+range? When the map reduce job finishes, how are files associated with groups?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2553
16872	ACCUMULO-2552	Is this really a blocker for 1.6.0? Can we declare Volume support early access for 1.6.0 and then fix these things in a follow on version? Given the history around ACCUMULO-118 it looks like the Volume support change is going to take some soak time to be production ready, and I'd rather not hang up the 1.6.0 release on that happening.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2552

#	Requirement	Questions	Answered	Successfully Answered	URL
16873	ACCUMULO-2552	Found the following (potential) issues: 1. Defaults to fs.defaultFS/fs.default.name for FileSystem. Should use FileSystem from the Path instead of pulling the default. I think this has the potential to throw an error out of Hadoop about using the "Wrong FS" ./core/src/main/java/org/apache/accumulo/core/file/rfile/SplitLarge.java:56: FileSystem fs = FileSystem.get(conf); 2. Just remove the getAndWrap method since it doesn't appear to be used anywhere? ./server/base/src/main/java/org/apache/accumulo/server/trace/TraceFileSystem.java:815: return wrap(FileSystem.get(conf)); 3. Pulls the default filesystem and expects files to be located on it. Likely not of concern because the class itself isn't "generally" recommended. ./server/base/src/main/java/org/apache/accumulo/server/util/AddFilesWithMissingEntries.java:68: final FileSystem fs = FileSystem.get(conf); 4. Log message about the URI being used to find the instance_id is no longer accurate ./server/base/src/main/java/org/apache/accumulo/server/util/ZooKeeperMain.java:43: FileSystem fs = FileSystem.get(CachedConfiguration.getInstance());	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2552
16874	ACCUMULO-2552	[~ecn], you agree that the first two don't represent blockers for 1.6? Do we want to deprecate {{AddFilesWithMissingEntries}} for 1.6.0 to remove it in w/e is next?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2552
16875	ACCUMULO-2546	Isn't this likely just caused by the error in ACCUMULO-2494 not declaring a dep and instead getting commons-math via transitive? hadoop 2.3.0 switched to an incompatibly newer version of commons-math. I think this likely will be fixed as a part of the work [~mdrob] is already doing?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2546
16876	ACCUMULO-2545	What's the reasoning behind adding licenses to otherwise empty files?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2545
16877	ACCUMULO-2541	Something seems.. goofy here. How do we read 10 elements, 5 of which didn't have readhead, in less time than the Slow iterator is configured to run. Those numbers shouldn't be capable of being returned in less than 1 second (1000millis).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2541
16878	ACCUMULO-2538	Has not been failing for days.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2538
16879	ACCUMULO-2537	I have an auto-save action on files I edit in Eclipse that adds this. Is that sufficient, or do we want to enforce application of this globally?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2537
16880	ACCUMULO-2537	I'd rather see a single clean up commit against this ticket along with some kind of enforcement mechanism to make sure new contributions follow the rule. do we do any checkstyle in the build already?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2537
16881	ACCUMULO-2532	[~ecn], something else I noticed while reading through this code: {code:title=FileUtil.java:97-99} // try to reserve the tmp dir if (!fs.createNewFile(new Path(result, "__reserve"))) result = null; {code} It looks like this {{__reserve}} file is never used anywhere else. Do you recall what the intent of making it was?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2532
16882	ACCUMULO-2522	Commit 448a79e00c56684325078cfe257733a450e8f344 in accumulo's branch refs/heads/1.5.2-SNAPSHOT from [~ecn] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=448a79e] ACCUMULO-2522 fix timeout typo	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2522
16883	ACCUMULO-2514	bq. Can't switch over many of the System.out.printf's since we rely on the exact format of the string in stdout to process the test results and print pretty graphs. I fixed the others. Can't you {{log.info(String.format("%....", str, str))}}?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2514
16884	ACCUMULO-2503	Just saw this failure on my build server across both branches (1.6.0-SNAPSHOT, with hadoop1 and hadoop2, and master). {noformat} Tests run: 2, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 0.008 sec <<< FAILURE! testTimestamps(org.apache.accumulo.core.util.format.DateStringFormatterTest) Time elapsed: 0.007 sec <<< FAILURE! java.lang.AssertionError at org.junit.Assert.fail(Assert.java:86) at org.junit.Assert.assertTrue(Assert.java:41) at org.junit.Assert.assertTrue(Assert.java:52) at org.apache.accumulo.core.util.format.DateStringFormatterTest.testTimestamps(DateStringFormatterTest.java:47) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:17) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:26) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:271) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:70) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:50) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:238) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:63) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:236) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:53) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:229) at org.junit.runners.ParentRunner.run(ParentRunner.java:309) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:252) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:141) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:112) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:189) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:165) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:85) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:115) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:75) {noformat}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2503
16885	ACCUMULO-2503	Commit 55b17e476686e846cdfd519393ef506e8a35a234 in accumulo's branch refs/heads/1.6.0-SNAPSHOT from [~mdrob@cloudera.com] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=55b17e4] ACCUMULO-2503 fix time zone assertion	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2503

#	Requirement	Questions	Answered	Successfully Answered	URL
16886	ACCUMULO-2503	More failures (I'm assuming) from upstream changes {code} scanDateStringFormatterTest(org.apache.accumulo.core.util.shell.ShellTest) Time elapsed: 0.007 sec <<< FAILURE! java.lang.AssertionError: r f:q [] 1969/12/31 16:00:00.000 v present in test@fake t> scan -fm org.apache.accumulo.core.util.format.DateStringFormatter -st r f:q [] 1970/01/01 00:00:00.000 v was not true expected:<true> but was:<false> at org.junit.Assert.fail(Assert.java:88) at org.junit.Assert.failNotEquals(Assert.java:743) at org.junit.Assert.assertEquals(Assert.java:118) at org.apache.accumulo.core.util.shell.ShellTest.assertGoodExit(ShellTest.java:124) at org.apache.accumulo.core.util.shell.ShellTest.exec(ShellTest.java:99) at org.apache.accumulo.core.util.shell.ShellTest.exec(ShellTest.java:93) at org.apache.accumulo.core.util.shell.ShellTest.scanDateStringFormatterTest(ShellTest.java:233) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:17) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:26) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:27) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:271) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:70) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:50) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:238) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:63) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:236) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:53) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:229) at org.junit.runners.ParentRunner.run(ParentRunner.java:309) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:252) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:141) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:112) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:189) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:165) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:85) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:115) at org.apache.maven.surefire.booter.ForkedBooter.main(ForkedBooter.java:75) {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2503
16887	ACCUMULO-2501	Very sorry that this patch was sitting here idle for so long, [~rweeks]. Thank you very much for the contribution. I applied some minor formatting changes and also added a unit test that explicitly invoked deepCopy. It looked good to me so I applied it. I'll add you to our contributors page as well. Thanks again!	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2501
16888	ACCUMULO-2499	Is there any reason not to apply this patch to 1.6.0, since it's just a documentation update, and 1.6.0 hasn't been released yet?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2499
16889	ACCUMULO-2499	More details per [~elserj], Start with http://en.wikipedia.org/wiki/Ephemeral_port This is a pretty common convention (Jetty and InetSocketAddress are two good examples in the Java realm). When requesting to bind a socket, the port that you provide is typically used directly to state which port you want to bind to on the interface. When you provide '0' instead of some other valid port (e.g. 1-65535), convention states that an ephemeral port will be chosen instead. This ephemeral port is likely to be unique and not already bound. Thus, configuring ports to use '0' instead of some explicit value, should, in most cases, work around any issues of running multiple distinct Accumulo instances (or any other process which tries to use the same default ports) on the same hardware.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2499
16890	ACCUMULO-2499	Hi Al! there's trailing whitespace on both patches, and the commit message for the 1.6 specific patch says "CCUMULO-2499" instead of "ACCUMULO-2499." Could you put up a version 3 set of patches to take care of these things?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2499
16891	ACCUMULO-2499	Hi [~krinker]! Sorry I missed some bits earlier. Why include a ZooKeeper entry? We don't include e.g. port numbers related to HDFS. How about instead, we open with something like "In addition to needing access to ports associated with HDFS and ZooKeeper, Accumulo will use the following default ports." What's the advantage of including the ports used in the 1.6.0 unit tests? AFAIK, those should never come up in a deployment, right? [~ecn] any obvious missing entries?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2499
16892	ACCUMULO-2499	[~krinker] - Are the entries supposed to be the same for 1.5 and 1.6? If so, we can just merge them using a proper merge instead of relying on two patches	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2499
16893	ACCUMULO-2497	I need a little help. What does this actually provide, other then wrapping existing constructors inside another structure?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2497
16894	ACCUMULO-2497	A. I'm not great with mocking, so pardon my ignorance. But couldn't you just as easily mock out Instance instance = new ZooKeeperInstance(...) to sub in your special Instance? B. If this is for testing purposes, wouldn't this live solely in Test code, which means it could live in src/test/java...? I don't know where new ZooKeeperInstance is called in our standard code that would require this.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2497
16895	ACCUMULO-2497	Let me rephrase the second question. Where in the code do we, outside of tests, create ZooKeeper instances? HdfsZooInstance already has a singleton method for grabbing it, which I assume means that is easily mocked. So that just leaves ZooKeeperInstance, which, AFAIK, isn't actually used in our code.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2497
16896	ACCUMULO-2497	Thanks for the example [~bhavanki], helped me understand what you were thinking. It seems liike a useful change. What do you think about developing code that would use this patch in tandem? This patch could be held off until the test that use it are added in another ticket and both patches pushed at the same time. If the change is made by itself there is a chance it will never be used and would be orphaned code from the start. There is also the possibility that when you get into the details of writing the test you realize this should work differently.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2497
16897	ACCUMULO-2491	Commit 4a27da54696ed97b197ce0abd5bef0a2ccfc2b80 in accumulo's branch refs/heads/1.4.5-SNAPSHOT from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=4a27da5] ACCUMULO-2491 Correct table user state upon termination of Security randomwalk A partial backport of ACCUMULO-1162 to 1.4 - commit e93a11a only. As the Security randomwalk ends and drops the table user, it now also sets the user's existence flag to false.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2491
16898	ACCUMULO-2490	Mike, could you give an example of how this makes the API more cumbersome?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2490

#	Requirement	Questions	Answered	Successfully Answered	URL
16899	ACCUMULO-2490	Sure. {code} Pair<Integer,String> pair = new Pair<Integer,String>(10, "XX"); // Doesn't compile: // Type mismatch: cannot convert from Map.Entry<capture#1-of ? super Integer,capture#2-of ? super String> to Map.Entry<Integer,String> Entry<Integer,String> entry = pair.toMapEntry(); // does compile but forces the client to use wildcards. Entry<? super Integer,? super String> entry2 = pair.toMapEntry(); // Doesn't compile either: Entry<Number,CharSequence> entry3 = pair.toMapEntry(); {code}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2490
16900	ACCUMULO-2488	Thanks, that's good information. There are actually five tables, as it turns out. Now, the question is, what should this balance check test for? Should it test that the default balancer is working as expected? Or should it test that the outcome fits some general definition of "balanced"? For example, that no server should be more than _n_ standard deviations away from the mean.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2488
16901	ACCUMULO-2467	If the root cause of this ticket is "the website manual needs to be regenerated", then it looks like the 1.4 manual is also impacted. I noticed it doesn't have the client section on the proxy service, which is present in the user manual released for 1.4.4. New ticket? Expand this ticket?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2467
16902	ACCUMULO-2464	[~elserj]: What good is the hash instead of plaintext? At some point, the client software needs the *actual* private credentials, not something computed from them. I don't see how you can work around this, except to integrate with some interactive credential source (direct user prompt) or pass the buck to some other interactive credential source (gnome-keyring-daemon with an interactive prompt for unlocking the keyring, etc.) You could do something like what [Maven does|http://maven.apache.org/guides/mini/guide-encryption.html] to obfuscate passwords with an encrypted version, and put the decryption key in a separate file... but then you have to protect that other file instead of this one, so I don't see much value gained by such obfuscation. At best, it's no better than simply putting the plaintext password in its own file. And, for that, I like the kind of options that openssl provides for specifying a password source (which also formed the motivation for shell password options in ACCUMULO-1412). Something similar to ACCUMULO-1412 could be done for configuration, but essentially, these are all just variations on the same general theme of isolating sensitive configuration to a file and locking down file permissions to the user running the service.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2464
16903	ACCUMULO-2450	[~bhavanki], could you reinstate DeleteRowsTest as a unit test under src/server/src/test/java/org/apache/accumulo/server/master/ ?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2450
16904	ACCUMULO-2445	After using ByteBuffer a bit, I think it is extremely unfriendly to work with. I'm leaning towards something a bit simpler, more like our current ByteSequence. Of course, making it easily convertible to/from ByteBuffer would be helpful, so I'd want to do that. What do you guys think?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2445
16905	ACCUMULO-2439	Hi Joe! Thanks for the contribution. Could you post it to [the Apache review board|https://reviews.apache.org/] so it's easier to gather feedback?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2439
16906	ACCUMULO-2423	bq. How's about we ditch bash for something [...] like python or ruby I can't say I'm opposed to that. Granted, python and ruby would also have their own caveats about the version installed (python 2 to 3 being the biggest but still easy to work around).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2423
16907	ACCUMULO-2421	Commit 2b3202aee60a92c909d28d6376e73fc4cac1b370 in accumulo's branch refs/heads/master from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=2b3202a] Merge remote-tracking branch 'origin/1.6.0-SNAPSHOT' (for ACCUMULO-2421). Conflicts: core/src/main/java/org/apache/accumulo/core/client/impl/Tables.java	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2421
16908	ACCUMULO-2419	Question for the group: I plan to support specifying the thread pool size for {{SimpleTimer}} (default 1) with a configuration property, but callers of {{getInstance()}}, like {{TabletServer}} and {{Master}}, will need to pass in a configuration object for that to happen. Is it appropriate to update all of those callers within this ticket - say, as a separate commit - or is that better done under a separate ticket?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2419
16909	ACCUMULO-2394	I like the way this looks, but what's the gain over keeping it in ASF CMS directly as markdown? Wouldn't it meet all of the criteria for this issue? We'd have to export it for a release. I'm not sure that's more work than having to push updates out of the main repo to the website.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2394
16910	ACCUMULO-2394	bq. Wouldn't it meet all of the criteria for this issue? Probably -- the ASF CMS is just much more of a pain to work with IMO. It's also not a part of the source tree which I'm not a fan of either (you should get the release notes in the artifact you download just as you do the CHANGES). Thinking some more, it's not a requirement that the release notes are integrated into the main a.a.o website. We could link directly to the release notes content hosted elsewhere (just like CHANGES is now). Exporting isn't a huge concern to me because that's something that isn't difficult to do. Getting release notes in the release artifacts would arguably require just as much work (if not more since it has to be put in tarball, rpms and debs). Overall, I'm not sold on this approach, but I had to the time so I just did it (it literally didn't take more than about 30mins). Jekyll is so much easier to generally use, provides a lot of nice features in building something like this and is easier to get installed. Integration with everything else is a concern.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2394
16911	ACCUMULO-2394	Is this issue done?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2394
16912	ACCUMULO-2380	Commit 7059c767bfbcd2def59556ed63cdfe443f2d336b in accumulo's branch refs/heads/1.4.5-SNAPSHOT from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=7059c76] ACCUMULO-2380 Improve RunTests to produce useful mapper and job output When passed the -m option, run.py for functional tests will emit one line on standard output for each test that passes or fails. The output is suitable for processing when tests are run within a mapreduce job. The mapper for RunTests, which runs functional tests, is updated to pass -m to run.py. It also inserts the map task attempt ID into the custom output, so that testers can trace back any test to the task where it ran, to view its logs for example. The mapper also now uses counters to track the number of successes, failures, errors, etc. Note that since the job may attempt to re-run tests that error out, the total for these counters can exceed the total number of unique tests run.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2380
16913	ACCUMULO-2380	Commit 7059c767bfbcd2def59556ed63cdfe443f2d336b in accumulo's branch refs/heads/1.5.1-SNAPSHOT from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=7059c76] ACCUMULO-2380 Improve RunTests to produce useful mapper and job output When passed the -m option, run.py for functional tests will emit one line on standard output for each test that passes or fails. The output is suitable for processing when tests are run within a mapreduce job. The mapper for RunTests, which runs functional tests, is updated to pass -m to run.py. It also inserts the map task attempt ID into the custom output, so that testers can trace back any test to the task where it ran, to view its logs for example. The mapper also now uses counters to track the number of successes, failures, errors, etc. Note that since the job may attempt to re-run tests that error out, the total for these counters can exceed the total number of unique tests run.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2380

#	Requirement	Questions	Answered	Successfully Answered	URL
16914	ACCUMULO-2380	Commit 7059c767bfbcd2def59556ed63cdfe443f2d336b in accumulo's branch refs/heads/1.6.0-SNAPSHOT from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=7059c76] ACCUMULO-2380 Improve RunTests to produce useful mapper and job output When passed the -m option, run.py for functional tests will emit one line on standard output for each test that passes or fails. The output is suitable for processing when tests are run within a mapreduce job. The mapper for RunTests, which runs functional tests, is updated to pass -m to run.py. It also inserts the map task attempt ID into the custom output, so that testers can trace back any test to the task where it ran, to view its logs for example. The mapper also now uses counters to track the number of successes, failures, errors, etc. Note that since the job may attempt to re-run tests that error out, the total for these counters can exceed the total number of unique tests run.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2380
16915	ACCUMULO-2380	Commit 7059c767bfbcd2def59556ed63cdfe443f2d336b in accumulo's branch refs/heads/master from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=7059c76] ACCUMULO-2380 Improve RunTests to produce useful mapper and job output When passed the -m option, run.py for functional tests will emit one line on standard output for each test that passes or fails. The output is suitable for processing when tests are run within a mapreduce job. The mapper for RunTests, which runs functional tests, is updated to pass -m to run.py. It also inserts the map task attempt ID into the custom output, so that testers can trace back any test to the task where it ran, to view its logs for example. The mapper also now uses counters to track the number of successes, failures, errors, etc. Note that since the job may attempt to re-run tests that error out, the total for these counters can exceed the total number of unique tests run.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2380
16916	ACCUMULO-2378	Commit 389831727a8f6ab2fda98f279a5cba442081ca69 in accumulo's branch refs/heads/1.4.5-SNAPSHOT from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=3898317] ACCUMULO-2378 Replace LOGNAME usage in TestUtils.py with getpass When running functional tests under MapReduce, the LOGNAME environment variable is not available. This change switches TestUtils.py to use getpass.getuser(), which looks for a variety of environment variables to locate the current user, and which works more reliably.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2378
16917	ACCUMULO-2378	Commit 389831727a8f6ab2fda98f279a5cba442081ca69 in accumulo's branch refs/heads/1.5.1-SNAPSHOT from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=3898317] ACCUMULO-2378 Replace LOGNAME usage in TestUtils.py with getpass When running functional tests under MapReduce, the LOGNAME environment variable is not available. This change switches TestUtils.py to use getpass.getuser(), which looks for a variety of environment variables to locate the current user, and which works more reliably.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2378
16918	ACCUMULO-2378	Commit 389831727a8f6ab2fda98f279a5cba442081ca69 in accumulo's branch refs/heads/1.6.0-SNAPSHOT from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=3898317] ACCUMULO-2378 Replace LOGNAME usage in TestUtils.py with getpass When running functional tests under MapReduce, the LOGNAME environment variable is not available. This change switches TestUtils.py to use getpass.getuser(), which looks for a variety of environment variables to locate the current user, and which works more reliably.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2378
16919	ACCUMULO-2378	Commit 389831727a8f6ab2fda98f279a5cba442081ca69 in accumulo's branch refs/heads/master from [~bhavanki] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=3898317] ACCUMULO-2378 Replace LOGNAME usage in TestUtils.py with getpass When running functional tests under MapReduce, the LOGNAME environment variable is not available. This change switches TestUtils.py to use getpass.getuser(), which looks for a variety of environment variables to locate the current user, and which works more reliably.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2378
16920	ACCUMULO-2375	Please expand the description of this ticket. What do you feel is missing? I often use my IDE to run Accumulo and did not need any particular procedures to get started. Just add the dependency to a pom.xml file and start coding. If you want to run Accumulo in-process, you can use the various test techniques (MockAccumulo, Instamo, etc...).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2375
16921	ACCUMULO-2368	Can't the walogs be owned by the two new tablets?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2368
16922	ACCUMULO-2365	Is there a single point of contact at Ubuntu who could work with the charm developer and be responsive? I've seen integration efforts fail due to lack of response from vendors. Then the developer's time is simply wasted.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2365
16923	ACCUMULO-2357	Thanks for the explanation [~kturner]. Updates to numRead are not thread-safe. Is that not a concern here?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2357
16924	ACCUMULO-2357	bq. Updates to numRead are not thread-safe. Is that not a concern here? Right, its not. Its not a concern. Its not expected that multiple threads would be reading from the same iterator.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2357
16925	ACCUMULO-2353	Why file the ticket here and not in Hadoop-Common, [~dlmarion]?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2353
16926	ACCUMULO-2353	Looks like this was fixed in Hadoop 2.8. What's the disposition for this ticket?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2353
16927	ACCUMULO-2350	It's difficult to recommend an alternative to something that never had a well-defined behavior. Is that the case here? Did this ever have well-defined (intended) behavior?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2350
16928	ACCUMULO-2350	The case I saw is that {{instance.getConfiguration}} was replaced by {{new ServerConfiguration(instance)}}. If there's ambiguity, can't we still succiently make some sort of recommendation?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2350
16929	ACCUMULO-2350	what in the client API should be using a ServerConfiguration?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2350
16930	ACCUMULO-2349	Where would we log this message to?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2349
16931	ACCUMULO-2349	> http://logging.apache.org/log4j/1.2/faq.html#a3.6 {quote} Is there a way to get log4j to automatically reload a configuration file if it changes? Yes. Both the DOMConfigurator and the PropertyConfigurator support automatic reloading through the configureAndWatch method. See the API documentation for more details. Because the configureAndWatch launches a separate wathdog thread, and because there is no way to stop this thread in log4j 1.2, the configureAndWatch method is unsafe for use in J2EE envrironments where applications are recycled. {quote}	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2349
16932	ACCUMULO-2349	DomConfigurator is already used by the code to accomplish this. I'm still not sure how the current configuration files are a burden on anyone. I rarely even touch them. Given that the current logging configuration mechanism have been like they are now for years, I'm don't view them as such a bad thing. They work a given way, and don't require users to typically touch them. I've also said it once already elsewhere, but you cannot use log4j.properties to configure the log-forwarding as it currently stands as you cannot configure the AsyncAppender with it. Vikram, how are you suggesting that users use the log4j.configuration property to actually configure per-process logging? That sounds like a bunch of new environment variables? Did you have something else in mind?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2349
16933	ACCUMULO-2349	bq. We could also consider that if there is a file called monitor_log4j.properties in conf dir, it means user wants to use it for monitor only. Uhh.. how is that any different than what currently happens?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2349

#	Requirement	Questions	Answered	Successfully Answered	URL
16934	ACCUMULO-2349	[~vines] wrote: {quote}We offer individual configuration for services, how would we offer that with your proposed changes?{quote} Different directories for different configuration, by default, seems sensible to me: {code} /etc/accumulo/ ├── common.conf ├── master	├── java.env	├── log4j.properties	└── master.conf └── tserver ├── java.env ├── log4j.properties └── tserver.conf {code} The actual code should accept these as command-line arguments, though. The directory structure would just be for scripts. A script might look like: {code} # /etc/init.d/accumulo-tserver # ... more SysVInit stuff here CONF_DIR=/etc/accumulo/tserver [-r "${CONF_DIR}/java.env"] && source "${CONF_DIR}/java.env" # java.env contains CLASSPATH= and/or TSERVER_OPTS= lines [-r "${CONF_DIR}/tserver.conf"] && CONFIG="--config ${CONF_DIR}/tserver.conf" java $TSERVER_OPTS org.apache.accumulo.tserver.TServer $CONFIG # ... more stuff here {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2349
16935	ACCUMULO-2345	[~vickyuec] avoiding allocation of a Violations Object was intentional. The reasoning was that constraint checking was on the write path for every mutation and want constraint checking to have as little impact as possible when no constraints are violated. I am not sure if this was worthwhile, it does complicate the code. Do you know if this change makes a difference in performance? I wonder if there would be a difference between the versions calling the check function10,000,000 times in a standalone program for mutations that pass,.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2345
16936	ACCUMULO-2345	The patch introduces Mockito, while we already have Easymock available as a dependency. Could you convert the unit test to Easymock for consistency? I'd be fine if this led to a discussion about which framework the project should use, as long as we settle on just one.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2345
16937	ACCUMULO-2338	I am interested in taking this ticket on. Is anyone working on anything related that I should be trying not to conflict with?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2338
16938	ACCUMULO-2336	From an API perspective, one way to do this without changing internal details: {code:java} public class ApplicationTable extends org.apache.accumulo.client.api.Table { public ApplicationTable(String name, Properties applicationMetadata) { super(name); // create new table updateApplicationMetadata(applicationMetadata); } private void updateApplicationMetadata(Properties appData) { /* implementation */ } } {code} This depends on a good, extensible client API with usable Table objects that an application could extend.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2336
16939	ACCUMULO-2314	[~busbey], we were just chatting today about logs containing "DSF". Are these corrections right?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2314
16940	ACCUMULO-2307	bq. MAC flips shit if there ANYTHING in the directory prior. That was by design, later code expects the directories to be empty. Is there a better way to do this? At the very least the odd exception for "ssl" in MAC needs a comment and a constant. It would be best if it could be removed, it just seems like it could cause trouble.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2307
16941	ACCUMULO-2307	[~vines]. Thanks for following up. What precisely is wrong with the current options that MAC already has: the options to set the SSL params in MiniAccumuloConfig, so it can support anything for testing that the full Accumulo supports? The SslIT test that currently exists can simply continue to create temporary keystores in a separate directory with which to test MAC's ssl support against?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2307
16942	ACCUMULO-2300	That's a valid point -- is there a good annotation or some other way to advertise when we've made that conscious decision in certain cases?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2300
16943	ACCUMULO-2299	[~vickyuec], can you add a unit test to assert that an exception is thrown on malformed data?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2299
16944	ACCUMULO-2299	It would be better to use the {{expected}} parameter with the IllegalArgumentException on the JUnit Test annotation instead of the {code} try { do(); fail(); } catch { //pass } {code} The test case you wrote is subject to the serialization of the WholeRowIterator. I would think that it would be cleaner to serialize some key-values, truncate the byte array you pass to the decode method and test for a failure there. Then, you're not encoding serialization logic into your test case. Also, do we want to throw an IllegalArgumentException, or should it just throw an IOException that's already in the method signature? I'm leaning towards the latter at first thought, but I'm not entirely sure.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2299
16945	ACCUMULO-2290	Is the node removal command preferable to using stop-here.sh ?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2290
16946	ACCUMULO-2290	Hi Arshak! How's the work on this ticket coming along? Do you think you'll have something we can review soon? If not, would you mind if we pushed this out to a later release, perhaps aiming for 1.6.1?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2290
16947	ACCUMULO-2287	are you thinking we should modify our scripts to do the rolling OR remove impediments to configuring something like supervisord to do log rolling?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2287
16948	ACCUMULO-2287	How will this impact the accumulo_watcher.sh (ACCUMULO-3569) script?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2287
16949	ACCUMULO-2287	bq. How will this impact the accumulo_watcher.sh (ACCUMULO-3569) script? Reading that script, it appears to perform a loop of "start normally", "wait for it to exit", "check the output". This means that it would not actually perform any rotation of the logs. [~jvines@gmail.com] do you have opinions on the matter?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2287
16950	ACCUMULO-2273	Is there a greater goal here other than library preference?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2273
16951	ACCUMULO-2257	Attached patch where class is moved out. Also added a unit test. Question: Do we need to retain Environment.getAuthorizations? It's not being used anywhere.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2257
16952	ACCUMULO-2257	[~ecn] I wanted to add more unit tests and make code better (e.g. not having public fields getting modified directly). Do you think it's better if I make all these changes in master itself, and then move all classes out once 1.6 is released?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2257
16953	ACCUMULO-2256	What version is this with?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2256
16954	ACCUMULO-2253	Thanks Vikram! Would you mind rebasing against 1.4.5-SNAPSHOT so I can pull this through all the dev branches?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2253
16955	ACCUMULO-2239	Thanks, Alex! I took a quick look: * Can you eliminate trailing whitespace from lines where you find it? * Can you ensure every file ends with a newline? * The markup in the "overwrite previous ingest" commands for test1's README has an error. * Go ahead and place "TBD" or something as the content for test3's README. A follow-on JIRA can be for real content.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2239
16956	ACCUMULO-2229	I don't understand your suggestion. You're saying, keep them in the modules for their components, but they can be included in the output package right there? We don't produce separate packaged artifacts for each component, do we? The particular thing I'm suggesting is unmaveny is stuff like this in assemble/src/main/assemblies/component.xml: {{<directory>../server/tserver/src/main/packaging/etc/init.d</directory>}}. How are you proposing we resolve that? Or are you just saying that you think the cost of attaching it as an artifact outweighs the risk of having inter-sibling, filesystem path-based implied dependencies?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2229
16957	ACCUMULO-2219	Fate operations are executed in a thread pool. The size in configurable by setting master.fate.threadpool.size, it defaults to 4. Operations that do not lock the same table should be able to execute in parallel. If an operation tries to lock a table that already locked, it should yield and not tie up the fate thread. When you say things are backed up, do you know more about what was going on? Did you happen to look the output of the fate admin print command? It shows what operations have locks and are waiting to get locks.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2219

#	Requirement	Questions	Answered	Successfully Answered	URL
16958	ACCUMULO-2212	I was looking back at this, and this new {{o.a.a.server.zookeeper.ZooReaderWriterFactory}} does not appear to be used at all. Is that intended? Is this a bug? Should this class be deleted? Also, the problem originally reported (singleton with public constructor) seems to have shifted to this {{ZooReaderWriterFactory}} class instead of the original {{ZooReaderWriter}}. Since this has already shipped in 1.6.1, I suppose a new issue should be opened to deal with it, but before I do that, I'm kind of curious what the intentions were for this unused class.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2212
16959	ACCUMULO-2204	Do we have any way of verifying that 19 should not exist?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2204
16960	ACCUMULO-2204	bq. Do we have any way of verifying that 19 should not exist? Maybe make the last node point to the first node deleted. Write 19->29 and then delete 19. Its not perfect, but it adds an interesting variation to the test.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2204
16961	ACCUMULO-2180	Can we do this work as a contribution to Presto, then track what the Presto ticket is here? In general, we're trying to get these kinds of connectors baked in to the project that would be connecting. Accumulo's APIs are stable, so said connectors should be able to work across multiple Accumulo versions.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2180
16962	ACCUMULO-2180	I would be interested to know the status of this JIRA. Is this being done or not yet started?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2180
16963	ACCUMULO-2180	I recently read that Presto has no security (i.e. usernames). Did I read that correctly?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2180
16964	ACCUMULO-2180	I can add a blurb to the website when it gets released, [~elserj]. Should it go in {{projects.md}} or some other page?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2180
16965	ACCUMULO-2180	bq. I can add a blurb to the website when it gets released, Josh Elser. Should it go in projects.md or some other page? Awesome. Thanks in advance. That sounds like a good place to me. I'm going to go ahead and close this issue (since it was addressed in Presto, not here in Accumulo) and open an issue to make sure we don't forget about updating the website.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2180
16966	ACCUMULO-2178	[~elserj] If we add ranges for both cf and cq to the new Filter, then the ColumnSliceFilter becomes a special case of the new Filter where cf_start and cf_end are both null. Does that sound correct?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2178
16967	ACCUMULO-2160	Do we really want PasswordToken to return a copy on getPassword()? Will that violate it's "Destroyable"-ness in any way (maybe when it's overridden)? I'm not sure. Just want to make sure we're not doing something we don't want to do. I also wonder how many of these changes can be represented as warnings in Eclipse (or your IDE of choosing) and whether we should make them standards. Eclipse may even support some of these things (like stripping redundant modifiers off interfaces) as save-actions. Also, I noticed that many changes here added trailing whitespace (including indented blank lines). In TestLruBlockCache (perhaps elsewhere also), it should probably have used JUnit's Assert.assertArrayEquals(), for more informative failure messages, rather than assertTrue(Arrays.equals()) It's curious that Math.abs was inline'd into the class. What's the reasoning for that in the examples?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2160
16968	ACCUMULO-2160	Thanks for doing this. I'm a huge Findbugs fan. Some items that jumped out at me: * In several places you switched out {{Math.random()}} for {{Random.nextInt()}}. It may be worthwhile creating a single {{Random}} instance for all code to use, so that there aren't a bunch of them around. The downside is that {{Random}} isn't thread-safe, so a shared instance would have to be synchronized. I'll leave it to you whether that is an issue for performance ({{Math.random()}} had to synchronize anyway). * The change in the {{equals()}} methods of {{Condition}} and {{ConditionalMutation}} mean that their subclasses can now be considered equal to superinstances. That might matter, it might not. * I saw some stubbed {{hashCode()}} methods, which makes sense for this ticket, but we'll need to fix those. * The correction in {{Authorizations.contains(String)}} seems like it could be huge, since before it would always return false. Did that have any impact? * The correction in {{TabletServer.AssignmentHandler.run()}} also seems like it could have some significance in making {{recentlyUnloadedCache}} work better - maybe causing more {{UNLOAD_FAILURE_NOT_SERVING}} messages to be generated. (Looks like a good thing.)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2160
16969	ACCUMULO-2159	I don't think turning on or off is a good idea. What's the root of the issue for you-- that the values aren't human readable when scanning via the shell? That can be helped with the use of a formatter on the table. I think some better (or any, it appears) documentation would be better than getting rid of the value.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2159
16970	ACCUMULO-2149	Is this really a blocker?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2149
16971	ACCUMULO-2149	Would it be sufficient to have the verification MR job look at the job counters to verify that the number of data local map tasks are within some threshold of the total number of launched map tasks? Presumably, for release testing the continuous ingest test and the verification job will run on clusters with a node count >= 5.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2149
16972	ACCUMULO-2149	bq. Would it be sufficient to have the verification MR job look at the job counters to verify that the number of data local map tasks are within some threshold of the total number of launched map tasks? That sounds like a good check. Not as rigorous as I would like, but much better than nothing. If there is a bug in the Accumulo code and its giving the wrong location for a tablet, then M/R would report task as local that are not. For example if Accumulo was giving M/R the last location instead of the current location.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2149
16973	ACCUMULO-2148	I can verify that given a 1.4 instance, 1.6.0-SNAPSHOT properly says "Noooooooope." The error message is the generic "This version of Accumulo doesn't work with the version from your data" from server/base/src/main/java/org/apache/accumulo/server/Accumulo.java {code} if (dataVersion != ServerConstants.DATA_VERSION && dataVersion != ServerConstants.PREV_DATA_VERSION) { throw new RuntimeException("This version of accumulo (" + codeVersion + ") is not compatible with files stored using data version " + dataVersion); } {code} Does this upgrade just need to be manual? Would an upgrade from 1.5.1-rc2 be sufficient?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2148
16974	ACCUMULO-2148	As a part of this, I plan to update the README's section on upgrading. Right now it still gives the 1.4 -> 1.5 instructions. Besides the instructions about Fate that will be added as a part of ACCUMULO-2519, are there any other special considerations?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2148
16975	ACCUMULO-2148	[~elserj], that sounds like something that should be covered in [the user manual chapter on multi volume set ups|https://git-wip-us.apache.org/repos/asf?p=accumulo.git;a=blob;f=docs/src/main/latex/accumulo_user_manual/chapters/multivolume.tex;h=0a0e6fef06fdf8c233ef11b20618c190e2a719c1;hb=refs/heads/1.6.0-SNAPSHOT]. For one thing, I would expect to do such an upgrade in two steps # Upgrade 1.5.x -> 1.6.x and verify # move from single volume to multi volume and verify This looks, to me, like a distinct step unrelated to upgrades. Furthermore, as the user manual chapter points out, multivolume configurations are an advanced option only needed at a particular scale. It's not something we should be calling out in the README for common upgrades. Can you please file a follow on ticket?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2148

#	Requirement	Questions	Answered	Successfully Answered	URL
16976	ACCUMULO-2148	bq. I could maybe explain that the 1.4 to 1.5 upgrade process asynchronously handled migrating 1.4 write ahead logs and that 1.6 can no longer make retries? I think that would be a good idea. Perhaps a disclaimer that those successfully running 1.5 upgraded from 1.4 don't need to perform the verification of no local WALs, too. bq. I did mean "table metadata". We write into multiple tables as a part of the upgrade, not just into the reserved metadata table. (the new reserved root, for example) Oh I see. What about ".. in both ZooKeeper and the Accumulo system tables." instead? That would match nomenclature better I believe.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2148
16977	ACCUMULO-2143	What's the use case we're supporting by exposing any part of the ZooKeeper underpinnings? This smells like abstraction leakage.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2143
16978	ACCUMULO-2143	How is it abstraction leakage? What if I'm running a test and I want to verify the state of ZooKeeper after some operation(s)? (I think) My original intent of this was to keep myself from having to dig into "How do I make a ZooKeeper object from the ZooKeeper connection string". That's the biggest arguing point I see here.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2143
16979	ACCUMULO-2143	bq. (I think) My original intent of this was to keep myself from having to dig into "How do I make a ZooKeeper object from the ZooKeeper connection string". That's the biggest arguing point I see here. That's pretty easy, actually. {{ZooKeeper zk = new ZooKeeper(connectString, timeout, null);}} It's _even easier_ if we use Curator.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2143
16980	ACCUMULO-2143	bq. A MAC we use for Accumulo testing should be a different thing than the MAC we make part of the public API for end users. They have different audiences and maintenance costs (because we constrain ourselves wrt changing our public API). That's my point, they don't have to be separate entities. I don't see the addition of a convenience method on MAC limiting any user from using it to test against Accumulo's APIs. Are you saying that you think this addition should not even be available to a user?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2143
16981	ACCUMULO-2143	Josh, I started it but it wasn't much. It is at home, but something like public ZooKeeper getZookeeper() { return new Zookeeper(config.getZookeepers(), 30000, null); } I thought about creating a Zookeeper object as part of the start, and then having getZookeeper return that. The benefit would be the MAC could handle the close. But it would only be one object. I also thought about making an array of returned Zookeeper objects, one for every call to getZookeeper. This again would allow the MAC to close them on MAC.close. Don't worry about anything I worked on though, it was small.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2143
16982	ACCUMULO-2143	{quote} bq. A MAC we use for Accumulo testing should be a different thing than the MAC we make part of the public API for end users. They have different audiences and maintenance costs (because we constrain ourselves wrt changing our public API). That's my point, they don't have to be separate entities. I don't see the addition of a convenience method on MAC limiting any user from using it to test against Accumulo's APIs. Are you saying that you think this addition should not even be available to a user? {quote} Yes, I'm saying for users of the public API we should not expose this.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2143
16983	ACCUMULO-2143	So what is the current desire for this ticket? 1) Do nothing and close the ticket 2) Do what the ticket originally requests and implement a getZookeeper that returns a Zookeeper object in both the MiniAccumuloCluster (public API) and MiniAccumuloClusterImpl (internal) 3) Only implement getZookeeper to return a Zookeeper object in the MiniAccumuloClusterImpl 4) Something different, I am completely off If it is option 2 or 3, do you want MAC to close any returned Zookeeper objects automatically, or simply document they will need to be closed in the Javadoc? Thanks	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2143
16984	ACCUMULO-2143	Ok, so everyone is good with option 3, but having it return a ZooReaderWriter? Should every call to getZookeeper() return a new object or the same one? I'll see what I can do to close it when the MAC closes. Thanks	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2143
16985	ACCUMULO-2139	[~elserj], I like where you are going with this. The more testing we can automate the better. ACCUMULO-2036 can not easily be tested on a single machine, it requires multiple nodes. Therefore its not a good candidate for a integration test. We need a framework for writing and running test that require a cluster. Does anything like this exist? What about bigtop, does anyone know what it offers? Upgrade testing is a different animal, but I suppose it could be fully automated also. I made a start at this w/ {{test/system/upgrade_test.sh}}, but its not fully automated and its not very exhaustive.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2139
16986	ACCUMULO-2139	bq. The more testing we can automate the better Bingo. bq. What about bigtop, does anyone know what it offers? My (uneducated) presumption was that Bigtop would be more pain than it would be worth for our use of it, but that's a painfully uneducated assertion. I know [~jvines@gmail.com] and [~mdrob] have both done stuff with Bigtop, perhaps they could weigh in on how easy/hard this would be. Ultimately, I don't want to be bogged down with a bunch of other things -- I just want to write my Accumulo test and avoid the situation where I can't test what I want to test due to limitations of the testing harness itself. We want to test Accumulo here, not test Accumulo *and* it's functionality with everything else. Ideally, I'd like to not have to worry about setting up things like Hadoop/ZooKeeper for such a test, but leave them as a pre-req to provide fs.default.name/fs.defaultFS, the zk connection string, and whatever necessary env variables we need to bootstrap Accumulo. I dunno -- if tool exists that provides us what we need without a bunch of other limitations, we should use it.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2139
16987	ACCUMULO-2139	Cool! Thanks for the input, [~mackrorysd]. bq. currently Bigtop needs a "release" before it will build everything for you Is this really just some Maven GAV that can be pulled through Maven or Ivy, or something else? Ultimately, I think you hit the nail on the head with this note. While it would *also* be nice for us to catch changes upstream in ZooKeeper or Hadoop (or even help them find bugs!), I think that's a next step. Step one is really integration testing on semi-stable versions of Hadoop so we can identify if (when) we screwed up. Another thought here is that we could leverage something like Yarn through HOYA to automate the cluster deployment. It took me a while to realize this, but I think there are a few perks. # It aligns with (my) goal of treating Hadoop and ZooKeeper as provided # You can define cluster specifications (2 masters, 8tservers and 64 tracers) # It would run agnostic of *actual* resources available. I think that last point is actually a pretty cool one. As long as your have memory on the NodeManagers to actually service the requests for your cluster (you have enough resources to start each process), the number of nodes is irrelevant. While this may not be satisfactory for some of our tests (such as local mappers), but it still allows anyone to run the tests (we could mark tests that don't have the resources as "incomplete" or "not-applicable" given the resources).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2139
16988	ACCUMULO-2136	[~vickyuec] - How did you generate this patch? It looks like you used {{git diff}}, so I was able to apply it using {{patch -p1}} but not {{git am}}. Can you create another patch using {{git format-patch}} and upload it? Can you add javadoc to the {{Fate}} constructor to hint to the user about the need to call {{startTransactionRunners(int)}}? Nit: There is trailing whitespace on Fate.java:145. I'll assign the issue to you, thanks for taking this on!	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2136
16989	ACCUMULO-2132	bq. OMG thank you! Ha, you're welcome. Was also going to try to squash the "SCDynamicStore" warnings, but those didn't seem to get picked up like this -- didn't feel like digging into that one as well.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2132

#	Requirement	Questions	Answered	Successfully Answered	URL
16990	ACCUMULO-2126	{{mvn package -Dhadoop.profile=1 && mvn package -Dhadoop.profile=2}} fails, but I think that's unavoidable. Can you strengthen the comment {{+ <!-- Just to denote intention, won't work to activate profile -->}}, since it is now required, in order for the enforcer plug-in to validate? Everything else looks good!	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2126
16991	ACCUMULO-2114	In what way do you propose users interact with this value? The way it is currently implemented is similar to java classes in the default package. That is, the fully qualified table names for tables in the default namespace have no namespace prefix. Namespaces are not really first class citizens, like tables. So, can you clarify how you think this should change?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2114
16992	ACCUMULO-2114	Well, I was thinking that the "default" namespace could be the results of getConnector().namespaceOps().getDefaultNamespace(), and the connector-specific default namespace would be the results of getConnector(namespace).namespaceOps().getDefaultNamespace(). But, I can see how this could be non-intuitive for users to understand the answer to the question: "What is the *default* default namespace?". Alternatively, we could add a separate namespaceOps().getCurrentNamespace() or something. The merits of that can be debated later if we decide to pursue that feature. Other reasons for moving it to namespaceOps(): 1. Constants are harder to find if separated into a separate file, because users don't know to go looking there if they're working with a separate class. 2. Constants files tend to get bloated, which diminishes their value for consumers (although they still have value for refactoring). 3. Encouraging client use of constants, with compile-time binding, means that compiled client code could break if we change the value of the constant.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2114
16993	ACCUMULO-2113	[~jaredwinick] do you have a test that verifies that behavior? shouldn't the different applications be under different classloaders, which would give them different pools?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2113
16994	ACCUMULO-2113	Oof, that is a little brittle. On the other hand though, the best we can do is try to make sure it works in the environments which it was tested. Is the "Thrift Connection Pool Checker" thread name coming from us or from Thrift? This would be good to lift into a constant if we're setting it (probably a note for whomever integrates it into Accumulo).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2113
16995	ACCUMULO-2113	[~kturner] are you referring to a version that would require changes to the ThriftConnectionPool and ZooSession to make them able to shutdown their resources without these reflection and thread hacks? I wasn't planning on that in the near term, especially if there is a working, albeit ugly, way of cleaning up resources against an unpatched Accumulo 1.5.0. Would this non-reflection version be what actually gets added to 1.4.5, 1.5.1, 1.6.0? FWIW, I tested the example web application on Tomcat 6.0.37 with the same successful results.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2113
16996	ACCUMULO-2107	Is there a way to apply the standards partially to a file without being in Eclipse? I can apply across the whole file outside Eclipse, but that can lead to all sorts of extra edits.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2107
16997	ACCUMULO-2100	Why was the artifact re-versioned to 2.0.0? This could be very confusing if it doesn't match Accumulo's versioning (especially, the major version).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2100
16998	ACCUMULO-2100	[~elserj] I know we emailed back and forth about this, but now that I have I had a chance to digest it for a while I wanted to open up the discussion again. It seems like the reason for accumulo-instamo-archetype-1.4 and accumulo-instamo-archetype-1.5 etc is because of changes in the API for the MiniAccumuloCluster. For most projects, I would handle this with a wrapper around MiniAccumuloCluster to hid these changes. In fact, I already created a wrapper in some of my testing that extends the MiniAccumuloCluster to override close and clean up the the temp directory. Are the differences between MiniAccumulo version so complex a wrapper wouldn't work? When I think about it from a user perspective, I would want one archetype that allows you specify the version of accumulo to use. Something like {noformat} mvn archetype:generate -Daccumulo.version=1.5.1 {noformat} Two things I don't know if archetypes handle that would be needed 1 - A way to validate optional properties. It would be useful to specify a released version of the archetype supports version A, B and C and throw an error if you tried to use a different accumulo version. 2 - A way to include a specific version of the wrapper based on the accumulo version passed in. To me, here are benefits assuming those 2 items are possible. - Adding support for a new version of Accumulo would be adding new wrapper class instead of creating a new archetype. - It would allow for multiple version of minor release, like 1.4.4 and 1.4.3. I am sure users will stuck at certain releases behind the latest and would want to use the archetype to build new components.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2100
16999	ACCUMULO-2100	Nice -- makes sense. We've had some other discussions about how MAC should evolve. Instead of extending the minicluster itself, might it make more sense to add something new to the MiniAccumuloConfig? The changes you have might be a simple workaround for the short-term, but it would definitely be good to weigh in on the bigger picture for MiniAccumuloCluster (notably, starting a temporary instance, starting a temp instance from a real instance, and other things).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2100
17000	ACCUMULO-2100	bq. might it make more sense to add something new to the MiniAccumuloConfig? I also think this is the way to go. If we want MAC to be able to run out of a directory setup by a previous run, then add an option to MiniAccumuloConfig to support this.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2100
17001	ACCUMULO-2100	Sorry, I meant to answer that as well. I was planning to add a Wrapper to the MiniAccumuloConfig as well to handle the new options. I agree a better long term solution would be to roll this stuff into the MiniAccumuloCluster/Config. Putting a wrapper on both objects in the archetype adds capability for already released versions and would allow new design to be tested. Functionality could be moved to the MAC proper as it made sense. Does this sound reasonable?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2100
17002	ACCUMULO-2097	[~elserj] - can the sqoop connector take care of the brutally simple import use case?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2097
17003	ACCUMULO-2097	bq. can the sqoop connector take care of the brutally simple import use case? I don't know -- I haven't looked at that code yet. My comment was more meant in terms of schema design in Accumulo. Mapping csv data to a big-table database is probably the easiest "schema" for new users to understand.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2097
17004	ACCUMULO-2097	[~elserj] - I'm already working on the CSV/TSV part, even though there is not a JIRA for it yet. Do you want to wait for me to finish that so you can integrate it in, or do you have different plans?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2097
17005	ACCUMULO-2097	bq. I'm already working on the CSV/TSV part, even though there is not a JIRA for it yet. Do you want to wait for me to finish that so you can integrate it in, or do you have different plans? /me facepalm I just worked up a bunch of stuff last night too, haha. I was making something standalone and then going to see how/where it makes sense to merge it in. Lemme push my stuff to Github for now. Edit: https://github.com/joshelser/accumulo-csv I got most of the boilerplate stubbed out last night. I was planning on supporting rfile creation or mutation writing. I also planned to implement a rudimentary schema support to show how the non-fixed colfam set for a table can be used and tracked.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2097
17006	ACCUMULO-2084	Can you explain why clone table, or anything, needs to lock the namespace? Isn't locking on the fully qualified table name sufficient?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2084

#	Requirement	Questions	Answered	Successfully Answered	URL
17007	ACCUMULO-2079	[~mdrob]: This ticket was about adding restrictions to the argument checks to ensure that users don't do bad things to the built-in namespaces. I added those restrictions here, and *additionally* tried to make sure the error message sent back to the client is informative (rather than simply a TApplicationException, which we wrapped as an AccumuloException with no informative message). Perhaps I didn't go far enough, in ensuring these messages were presented well to the client... but that wasn't the primary goal of this ticket. I created a second ticket for that... and I fully back it. What I'm trying to understand by your "-1", is if there's something *bad* about the fix for *this* ticket, or if you're just annoyed it didn't go far enough (all the way back to the client). If it is something bad with *this* ticket, I can fix that here. If it's that it didn't go far enough... then this commit is fine as an incremental improvement, and I (we, whoever) can still do that work under a more tightly scoped ticket (which I created as ACCUMULO-2133). What I mean when I say I'm not forcing users to string parsing to get at the information... I mean that this fix did not attempt to make that current requirement any better. If they had to do that before, they still have to do that now... what I mean is that I did not add any additional requirement to do so, but I still made the improvement to restrict table creation in the system namespace. (Actually, I think I laid the groundwork for the improvements you're talking about in a future commit... it just wasn't done... yet.) The fix here did include documentation updates (minor javadoc), informative error strings, and additional tests... it just wasn't scoped to the API issue you're talking about... it was scoped to the restriction to create tables in namespaces. Certainly these overlap... but in some sense, all of our code overlaps... we have to draw the line somewhere in order to make incremental progress. Does that mean this commit should be reverted, or that it did not add value as an improvement, independent of that issue? In any case, if it's a priority, I can certainly begin work on Monday on ACCUMULO-2133. However, it was my understanding that there was some contention regarding exactly *how* the API changes should happen, and honestly, I'd rather avoid being at the center of that controversy, and just make progress in other areas, rather than throw myself under that particular bus. Either way, API changes or not... this change needed to happen... because we can't have users creating tables in the reserved namespace, or deleting built-in namespaces, etc.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2079
17008	ACCUMULO-2076	What about the new {{public ZooKeeperInstance(Configuration config)}} method introduced in 1.6? This was created for the sole purpose of passing in client side config. accumulo-site.xml contains server side config.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2076
17009	ACCUMULO-2076	Can clients and servers use a different set of zoo keepers? I imagine that those have to be the same between both, right? So what is the distinction between {{accumulo-site.xml}} and a client config at that point?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2076
17010	ACCUMULO-2076	bq. The accumulo-site.xml file used by servers has a secret that client should not see. This has come up a couple of times recently about how the accumulo-site.xml file isn't necessarily good to client consumption. Would it be better to think about how we can split up the logic into client and server xml files and make sure we get make decisions about how to structure said files from a segmentation of configuration params (for server and for public use) and for ease of API use.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2076
17011	ACCUMULO-2076	(Discussion taken from review board.) {quote} [~elserj] wrote: Wait, I just noticed that you have this RB marked as 1.4.5-SNAPSHOT (and your javadoc has @since 1.4.5, etc) but your original ticket is marked for 1.7.0. I do not think this should pushed back to 1.4.5 or 1.5.1. What is your intent? Also, I was just talking to [~bills] about creating Instances earlier by chance. Have you considered expanding this to also support HdfsZooInstance and MiniAccumuloInstance? This would also give us flexibility in (if we move this direction) to having threadsafe and non-threadsafe variants of the aforementioned Instances (also moving towards a fluent interface like [~bhavanki] mentioned). {quote} In theory, I'm fine with any fix version <= 1.7; in practice, I'm writing utilities and applications against 1.4 and 1.5 and I absolutely intend to use this in my own code, even if I have to refactor into an external utility class. Expanding this to HZI and MAI (and MockInstance?) is a good idea. I hadn't implemented those because I consider them relatively trivial after there is something for ZKI, and was more concerned about getting the first one right.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2076
17012	ACCUMULO-2076	Why does this have blocker priority?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2076
17013	ACCUMULO-2061	I do want to note that you are increasing the the configuration complexity for "common" users now. Currently, instance.dfs.uri doesn't need to be specified (and will be pulled out of hdfs-site.xml if on the classpath) which lets me make a new installation by only specifying the HDFS directory I want to use with instance.dfs.dir. Edit: Can I specify a path only in instance.volumes instead of URI which would default to the configuration from hdfs-site (defaultFS, is it?) to maintain backwards compatibility with configurations?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2061
17014	ACCUMULO-2061	Ok, that's what I typically expected when I hear relative paths in the context of 1.6, but when dealing with the configuration of these parameters, it wasn't clear to me. So what exactly needs to be decided here? Is deprecating instance.dfs.dir and instance.dfs.uri definitely what we want to do for 1.6? It looks like instance.volumes defaults to "" presently which doesn't jive with what Christopher said earlier about defaulting to "/accumulo" -- should this be updated? As of now, it looks like the "default" volume is still pulled only from instance.dfs.dir (will error if you try it through instance.volumes) and instance.volumes is used for additional volume definitions.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2061
17015	ACCUMULO-2061	[~elserj] wrote: {quote}So what exactly needs to be decided here?{quote} # instance.dfs.uri is certainly unnecessary outside of upgrades, only needed for resolving the previous relative paths. We can deprecate its use now... no deciding necessary. # instance.dfs.dir is needed to resolve the previous relative paths, but also (currently) as a (in my view) confusing extra element for each volume in instance.volumes... we need to decide if we want it to only be used for resolving the previous relative paths. Do we agree that the extra element adds little value as a separate property and is really only needed to resolve relative paths from previous versions?.... OR... do we think that it's useful to configure a common subdirectory with a separate property and that this warrants keeping the current behavior of joining two properties to make absolute paths? Personally, if we think there's any value in configuring a common subdirectory, i think we can support that feature by adding property interpolation in configuration parsing in the future. Example: {code} subdir = myApplication.accumulo-files instance.volumes = hdfs://nn1/${subdir},hdfs://nn2/${subdir} {code} However, while interpolation is a useful feature... I don't think it's a prerequisite for solidifying the behavior of absolute paths with a single property rather than two.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2061
17016	ACCUMULO-2060	Can the defaults be updated in a follow-on JIRA?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2060
17017	ACCUMULO-2060	bq. Can the defaults be updated in a follow-on JIRA? Hypothetically, but I don't see how that would make sense. I introduced what I consider a bug because now the AccumuloClassLoader is doing something different than what we advertise in the example configurations. Also, these changes were never released.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2060
17018	ACCUMULO-2052	Is this a documentation issue?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2052

#	Requirement	Questions	Answered	Successfully Answered	URL
17019	ACCUMULO-2050	Wow, that's rough -- but, giving some more feedback on the HDFS ticket would probably be good too :P I haven't looked at the viewfs stuff at all yet. It seems excessive to reimplement something whose functionality already (appears) to be in HDFS. Is delaying viewfs for 1.6 and planning for it in a 1.7 timeline reasonable? This might actually be mutually agreeable for a few other reasons (things that didn't make the 1.6 cut).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2050
17020	ACCUMULO-2039	The patch looks fine. One question for other committers: The patch eliminates a public method in ByteBufferUtil, since it was only used by the code being changed and could be refactored away. Is that elimination OK?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2039
17021	ACCUMULO-2039	ByteBufferUtil is not in our public API, so ti should be fine. We've been known to leak that abstraction, so I'd recommend checking the Accumulo contrib repos to make sure something else isn't impacted. I noticed this defensively copies the actual byte arrays. Do we have some quick way to check for a performance regression on this? I'm not sure how extensively we use getAuthorizationsBB, so I don't know how much overhead we're talking about adding. An alternative would be to use the byte arrays in place but call [asReadOnlyBuffer|http://docs.oracle.com/javase/6/docs/api/java/nio/ByteBuffer.html#asReadOnlyBuffer()] to make sure the generated BB's can't be modified. This won't work if whomever is using the BB needs to grab the underlying byte array, because read only buffer's don't allow access to them (I think Avro's serialization code does this).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2039
17022	ACCUMULO-2015	This patch will create gc/monitor/tracer files when initializing a standalone mode. It will also check for the existence of those files in start-all/here to avoid error messages from grep. Note that this is a change from current behaviour, hence I'm putting it up for review first.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2015
17023	ACCUMULO-2015	"standalone mode" is new to me, but, in my confusion, it looks like the concept doesn't exist in 1.5 or 1.6. Can you explain what standalone mode is supposed to be and why it doesn't exist in newer versions?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2015
17024	ACCUMULO-2015	"standalone mode" was a shortcut to launching a single node Accumulo instance in psuedo-distributed mode invoked by not specifying both the slaves and masters files in $ACCUMULO_HOME/conf. It was in the original code import, so there aren't tickets to look at for its original reasoning. The only [docs I've found for it are comments in the part of config.sh that set it up|https://git-wip-us.apache.org/repos/asf?p=accumulo.git;a=blob;f=bin/config.sh;h=485e259e232cbe0eb99180d1cafd6bc3a8e63488;hb=419aacc45279a3cd6b3b5bf61baf486f082a450a#l97]. AFAICT, it was to make it easier to spin up test instances. [~kturner] removed it from the 1.5.x line in [this commit|https://git-wip-us.apache.org/repos/asf?p=accumulo.git;a=commit;h=4b163a996e090550eb6a29e6d63eb0abb5b5853c] as a part of ACCUMULO-1071 (though I'm not sure if that ticket # was a typo). He would have to speak to what his reasoning was, there's not an explanation on the ticket or mailing lists at the time. If I had to guess, I'd say the presumption was that the addition of bootstrap (and minicluster) obviated the need. Since 1.5.x wasn't required to maintain default behavior, it was just removed instead of wrapped in some option.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-2015
17025	ACCUMULO-2010	When I was working ACCUMULO-1984, I did not see any compiler warnings pertaining to {{Closeable}}. Is this something IDE-specific?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2010
17026	ACCUMULO-2005	I will work on this ticket. Could someone with admin privs assign it to me. Thanks!	1	1	https://issues.apache.org/jira/browse/ACCUMULO-2005
17027	ACCUMULO-2005	The patch applies well to 1.4.5-SNAPSHOT. Unfortunately, argument handling changed substantially in 1.5.x, so the changes won't impact there. Could you checkout 1.5.1-SNAPSHOT and take a look at making your changes work with that version of the code? You'd just add another patch for that branch (e.g. ACCUMULO-2005.v4-1.5.1-SNAPSHOT.patch). If you don't think you'll have time, just drop a note here and we'll get someone to handle it when they have time.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2005
17028	ACCUMULO-2005	Hi Sean, Sorry, I've been late on emails. I've on-site with no internet access until now. Just caught up and saw this. Yes I'd like to go through rebasing exercise myself (to learn), but I probably won't be able to touch this til this weekend. Is that ok? If yes, please show me the steps. Thanks, Hung	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2005
17029	ACCUMULO-2005	Thanks for the contribution [~hpham]. How would you like to be listed on [the contributors page|http://accumulo.apache.org/people.html]? (namely organization and timezone, if either)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-2005
17030	ACCUMULO-1974	[~bhavanki], do you have the capabilities to rebase these changes on master easily? I'm getting failures in applying the patch: {noformat} error: patch failed: core/src/main/java/org/apache/accumulo/core/conf/Property.java:600 error: core/src/main/java/org/apache/accumulo/core/conf/Property.java: patch does not apply {noformat}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1974
17031	ACCUMULO-1960	Made the agitator script a bit more forgiving on how it needs to be run. Also defined an HDFS_USER and ACCUMULO_USER in continuous-env.sh, with them defaulting to the `whoami` in the example. 1) Doesn't require to be run as root, but will change users properly if you do 2) If you run it as yourself, and also run the datanode and accumulo processes as yourself, it won't bother changing users 3) If the user you run the agitator as isn't the same as the user for the process being agitated, the script will try to sudo. If you configured sudo correctly, it should work fine.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1960
17032	ACCUMULO-1957	I was talking offline with [~kturner]. Does anyone have a use-case for wanting to control write durability on the BatchWriter separately from table configuration?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1957
17033	ACCUMULO-1946	Yes, I would assume that it does slow down writes to some extent. I thought briefly about adding a note about this potentially being unnecessary if you're running with UPCs or redundant power supplies, but decided against it. Do [~kturner] or [~ecn] have an idea for the amount of penalty incurred? Looking at HBASE-5954, Lars reported a little under 50% penalty (if I'm reading that correctly, and those numbers didn't change in the end). Personally, I'd much rather warn that they should turn it on, and let someone turn it off if they're comfortable. What do you think about making a note about being able to turn it off when you have more robust hardware power failures? I don't necessarily want to say "we recommend this, but it's going to make it slower". I doesn't seem like the norm to accept data corruption for a relatively small improvement in performance, especially when we think in terms of commodity hardware. Thoughts?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1946
17034	ACCUMULO-1946	bq. Eventually, we may want to have these as table settings, or as mentioned in HBASE-5954, with each update. [~ecn], should we open a ticket? As you mentioned having these for ROOT and METADATA tables sounds important.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1946
17035	ACCUMULO-1946	[~elserj] why not document this in 1.4 also, now that it has hadoop 2 support? Is this config introduced in hadoop 2?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1946
17036	ACCUMULO-1946	What about the following situation, is it relevant? # Compaction occurs writes to file F1 and closes file F1 # Portions of file F1 on 3 datanodes are in memory (OS memory maybe) and not flushed to disk # All 3 datanodes F1 was written to are powered off	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1946
17037	ACCUMULO-1945	I would like to start contributing, could someone assign this to me?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1945
17038	ACCUMULO-1945	Hi Andrew! Could you add "help" to the list of valid commands on the usage line?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1945

#	Requirement	Questions	Answered	Successfully Answered	URL
17039	ACCUMULO-1945	No worries. Patch looks good, I'll apply it this evening. How would you like your name, company, and time zone to appear on the [contributors list|accumulo.apache.org/people.html], if at all?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1945
17040	ACCUMULO-1933	I added this convenience feature for memory units at the same time as the time units, which are definitely SI. So, I guess there's a choice we have to make... do we want to try to stick to SI for everything, or are memory units sufficiently distinct, to support a wider range of (non-fractional) unit declarations? Second, if we do think that memory units are sufficiently different, and to the extent that it makes sense to support the case-insensitive JVM precedent, how do we deal with "b" vs. "B" for bytes? [~ted.m]: since you've expressed a willingness to work on this, what are your thoughts?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1933
17041	ACCUMULO-1933	I don't think I can omit "B" and still be backwards compatible. If I leave "b" as is, then the code errors out currently. Hmm I think I have two options: 1. Support G, M, K, and B in all cases 2. Support G, M, and K in all cases but B only in upper case and add a clean error message for b What do you think. I'm open to both.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1933
17042	ACCUMULO-1933	Thanks Ted! Could you update the new test file to include an ASF license header? A few suggested improvements: * how about IllegalArgumentException instead of generic RuntimeException? * please break the two failure cases into their own tests that use JUnit's expected exception argument to hte Test annotation Also, please make sure you're creating the patch with the "git diff" args given [in our contributor guide|http://accumulo.apache.org/git.html#contributors] so that a committer can easily attribute the commit to you.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1933
17043	ACCUMULO-1933	Thanks Chris, What is the next steps?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1933
17044	ACCUMULO-1933	[~vines] - Thanks for rolling this in. Next time, can you also use the {{--signoff}} option of {{git am}}? It makes it easier to track who is doing what. Thanks!	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1933
17045	ACCUMULO-1910	The source for ./start/src/test/resources/ClassLoaderTestA/Test.jar (and B, C) is below. I have a question about packaging. Because this is a classloader test, was this created as a separate project so that it would not be in the accumulo namespace and is that how it went missing? TestCombinerX,Y are different, I need to check the others, but how do you want these to be packaged? (Stand-alone maven module, or maybe not as a source / class file but just added as a readme?) +++++++++ package test; public class TestObject implements Test { private int count = 0; public String hello() { return "Hello from testA"; } public int add() { return ++count; } }	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1910
17046	ACCUMULO-1910	bq. Because this is a classloader test, was this created as a separate project so that it would not be in the accumulo namespace and is that how it went missing? I'd have to look at the test cases that load these jars, but I don't think they were in their own namespace. I don't think we need to have anything in our source tree for them, but having the source files in the jar with the class file would be nice, IMO. So, a new jar for those which have no source is what I was thinking of. You could attach either the source Java files or a new jar (which contains the source). It's up to you. My gut reaction was that we need source (per ASF policy), but [~billie.rinaldi] pointed out on the mailing list that she doesn't think that to be the case. Regardless, it would be nice for people trying to understand this code to have access to all the source we use to test.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1910
17047	ACCUMULO-1910	[~etcoleman] - Did [~elserj] provide enough guidance, or do you have further questions on this? We'd love to see this submitted and committed.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1910
17048	ACCUMULO-1910	Thanks for the patch, Ed! When you say that the project version is set to 1.4.4, do you mean that the dependency version is set to 1.4.4? Because I don't see anything else that it might be. It looks like it compiled cleanly for me as applied, so I'm not really sure what the problem that you're describing is. After applying the patch and running {{mvn clean package}}, I spot checked {{./src/start/src/test/resources/ClassLoaderTestA/Test.jar}} and {{./src/start/target/test-classes/ClassLoaderTestA/Test.jar}}, and neither one had source. Should I have used a different maven goal to make it work? Minor nit: there is trailing white space in some of the files: {{Test.java}}, {{TestObject.java}}, and some (but not all) of the {{pom.xml}}.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1910
17049	ACCUMULO-1910	Yea... I think we're making this complicated for the sake of preserving compatibility in something that isn't part of the public API anyway. I strongly encourage the shortest path solution here, as the code was never _functionally_ wrong, just a slight bit opaque. Like [~elserj] alluded to earlier, a binary patch of swapping out one jar for another would likely have been sufficient. I added {{src/test-jars-src}} as another module to the main {{pom.xml}} and then saw the artifacts get built. I think it would be easiest for new users to grok fewer jars, since the source files will exist in the tree anyway, but feel free to disagree. My main concern at this point is making sure that the jars we are generating are the ones getting used by our tests, which I'm not convinced is the case right now. Can you create a new JIRA for the doc-gen issues that you are running into?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1910
17050	ACCUMULO-1910	I have no problem with just putting the jars into a patch. However, are there any issues with the manifest? Maven seems to insist that some of my info leak in - no problem with that if everyone is okay with it. With the current patch, it could be built on the jenkis server and then have only accumulo build artifacts. - (Seems cleaner to me.) Alternately, I could modify the build on my end, but sometimes it seems than me and maven are not the best of friends (i.e. suggestions more than welcome,) Or I could ditch maven and just jar cvf ... as long as no one wants to replicate the build...	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1910
17051	ACCUMULO-1907	[~ctubbsii] Seems like moving tables between namespaces could be a nice feature. For example a table could be created in one namespace and then "published" by moving it to another namespace. What are some examples of configuration changes that can cause a problem in table migration, and can we possibly address this by documentation?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1907
17052	ACCUMULO-1904	How would this simplify the API? Can you detail the use case(s)?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1904
17053	ACCUMULO-1904	Can the priority be lowered?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1904
17054	ACCUMULO-1904	[~medined] wrote: bq. How would this simplify the API? It would simplify the API by reducing the need to have all the convenience methods for creating a new table with pre-configured settings. If you could create the table initialized to an offline state, you could use the normal API to configure it before you brought it online. That simplifies a lot of the overloaded createTable cruft.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1904
17055	ACCUMULO-1898	I'm not sure if this aligns with what's been done >=1.5.0 as the NameNode information was removed from the monitor. What kind of confirmation are you looking for? Isn't authenticating with Accumulo (via the shell, for example) validation that your authentication system is working? I'm not a kerberos expert, but it seems like this would a post-installation check and not something I'd want/need to know as I continue to use the installation. Are you looking for the equivalent of "Security is OFF" for Accumulo as exists on the NN UI?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1898

#	Requirement	Questions	Answered	Successfully Answered	URL
17056	ACCUMULO-1898	There should be. Accumulo kerberos integration uses org.apache.hadoop.security.UserGroupInformation#loginUserFromKeytab which has {code} LOG.info("Login successful for user " + keytabPrincipal + " using keytab file " + keytabFile);{code} Check your log4j settings to make sure it's not be squelched?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1898
17057	ACCUMULO-1888	Commit 0e63755d5e3a9e441bae51d2f28402a3768820dc in branch refs/heads/1.6.0-SNAPSHOT from [~ecn] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=0e63755] ACCUMULO-1831 ACCUMULO-1888 use uuids to confirm WALog GC; remove WALog entries in the !METADATA table with the correct mutation	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1888
17058	ACCUMULO-1869	Hi [~busbey], the issue has been associated with the test component. Is that valid also for others modules? I mean, the others src/test/resources/log4j.properties?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1869
17059	ACCUMULO-1856	Would this be using the [metrics2 stuff from Hadoop Common|http://hadoop.apache.org/docs/stable/api/org/apache/hadoop/metrics2/package-summary.html] within Accumulo so we can then use the Ganglia integration from there? or just trying to do something similar to the Ganglia provider itself? (I'd prefer the former, but it means adding a dep on hadoop-common)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1856
17060	ACCUMULO-1853	Right, I'm not surprised. Something is wrong... can you jstack the processes and see where things are stopped (like I did over in ACCUMULO-1861)? If you want to post your logs, I can take a look.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1853
17061	ACCUMULO-1853	I noticed that. Will it be given more memory if I increase the maven opts? What do you set yours to?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1853
17062	ACCUMULO-1853	Is this done?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1853
17063	ACCUMULO-1834	Whatever the rules, I think they should complement the rules for committing patches submitted to JIRA. The more rigorous the rules for RB end up, the less they should be for patches submitted after review, so that committers don't have to exert the same effort twice. If a patch doesn't go through RB, it should face more scrutiny before commit. The Apache voting rules say 3 +1s for a code mod, but that seems high. I like 1 or 2 ship-its. Maybe if it gets only 1, the patch committer must not be the 1 ship-it. What changes don't need to go through RB? Documentation? Commenting?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1834
17064	ACCUMULO-1833	[~cmccubbin] I've been playing around with this. Can you give any insight into the type of workload you were running when you saw this? Even distribution of mutations across many tables? Many to one table and few to many tables? How many threads is "multiple"?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17065	ACCUMULO-1833	I did a plain test of ZooCache and it seemed okay. So now I'm wondering if the issue is related to the way this code gets the ZooCache. It's ultimately reusing the same ZooCache, but it goes through some hoops to retrieve it each time. Could you try out something like the attached patch and let us know if it has any effect on performance?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1833
17066	ACCUMULO-1833	bq. A simple way to avoid this extra work would be to make the cache in MultiTableBatchWriter use table names instead of table ids (the lookup in the cache could be done immediately and if a hash map was used would have constant time cost). But that change in behavior could break any existing code that relies on the current late binding of the table name to table id. Since the late binding is eventually consistent anyway, can't we make a time-sensitive cache (i.e. hold on to things for 5 seconds) that maintains the same semantics and still fixes this performance issue?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17067	ACCUMULO-1833	bq. Since the late binding is eventually consistent anyway, can't we make a time-sensitive cache (i.e. hold on to things for 5 seconds) that maintains the same semantics and still fixes this performance issue? That would at least alleviate the lock contention on ZooCache, but that's really just pushing the problem down a little farther. Perhaps it would push it far enough that clients would see it until they're at 100+ threads in one JVM? That'd probably be sufficient for a fix to this ticket. Digging through the code made it relatively apparent to me that there was a design decision that one ZooKeeperInstance (and Connectors from that Instance) will share a cache on top of ZK. If we want to talk about making changes to let the API be accessed more concurrently within one VM, I think we'd want to start at that point rather than try to address it at every usage of the caching.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17068	ACCUMULO-1833	bq. can't we make a time-sensitive cache (i.e. hold on to things for 5 seconds) that maintains the same semantics and still fixes this performance issue? That sounds good. Would need to ensure time going backwards is handled properly (just drop the cache entry). A "small" timeout could be determined empirically. If this approach looks promising, this should be marked a bug (it fixes a performance problem and does not change behavior or API) and fixed in earlier releases.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17069	ACCUMULO-1833	There's always Guava's cache which has a time-based expiration policy (http://docs.guava-libraries.googlecode.com/git/javadoc/com/google/common/cache/CacheBuilder.html). bq. Would need to ensure time going backwards is handled properly (just drop the cache entry). I was just thinking about a per-VM cache (at which point we only have to worry about the system clock drastically changing). Are you considering something larger than that?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17070	ACCUMULO-1833	bq. Are you considering something larger than that? Nope, just wanted to handle the system clock jumping backwards. Hopefully Guava's cache would handle this.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17071	ACCUMULO-1833	bq. Nope, just wanted to handle the system clock jumping backwards. Hopefully Guava's cache would handle this. Ok, cool. I'll try to get to this today. I'll try to peek at the implementation again, but I think it handles this ok. Should this go in 1.6.0 or >=1.6.1? I figured I'd put it in 1.5.1 either way.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17072	ACCUMULO-1833	bq. Should this go in 1.6.0 or >=1.6.1? I think its a performance bug fix and should go in 1.6.0.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17073	ACCUMULO-1833	bq. Should this go in 1.6.0 or >=1.6.1? I figured I'd put it in 1.5.1 either way. 1.5.1, 1.6.1, and 1.7.0 seem like the best option to me. I don't want to see 1.6.0 get held up and we've already drawn the line in the sand last week.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1833
17074	ACCUMULO-1833	I consider this a non-critical performance bug, and should be handled in a future release, whether it's all bugfix releases or just 1.7.0, I could go either way. If we do decide on bugfix releases and it's done this week, maybe I'll look the other way, provided it doesn't introduce any crazy issues. Is there a randomwalk test that uses the MTBW?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17075	ACCUMULO-1833	Ok, I'll be pushing some changes shortly to address Table cache invalidation. One thing that worries me (not sure if it should), but using an AtomicInteger *does* introduce the possibility of int overflow in a long-running process. Should we care about this? Is using an AtomicLong better? Thoughts on how best to work around that?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833
17076	ACCUMULO-1833	bq. I would be happier with an AtomicLong or supporting rolling Done. [~kturner] and/or [~ecn], can you look at these changes before I merge them into 1.5.1-SNAPSHOT? They're in the branch ACCUMULO-1833-caching.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1833

#	Requirement	Questions	Answered	Successfully Answered	URL
17077	ACCUMULO-1832	Is this truly a blocker for 1.6.0?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1832
17078	ACCUMULO-1832	How would you be in the same boat? This is a unique problem because of the fact that 1.6 stores full URIs, whereas 1.5 and earlier only stored relative paths. You could replace the namenode with a backup on a different server and things would just work (in theory, anyway).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1832
17079	ACCUMULO-1832	I like the idea of using a prefix (see [1]) instead of full URL. If there were a mapping of prefix to namenode address in the configuration, would that make rename operations simpler? It's like viewfs to some degree, but I think it allows greater flexibility in the long run as viewfs is limited to directories. In another release when we have a pluggable file placement strategy, then we could add weights or priority to the different prefixes. [1] https://issues.apache.org/jira/browse/ACCUMULO-118?focusedCommentId=13245562&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13245562	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1832
17080	ACCUMULO-1831	This should be governed by {{master.recovery.max.age}}. That is, we don't really GC recovery files, we just remove them when they have been sitting there for an hour (by default). Did you set this to some very low value?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1831
17081	ACCUMULO-1829	Bump. Could someone commit this? the follow on for docs (ACCUMULO-1835) is already in the 1.4.x branch claiming -f works.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1829
17082	ACCUMULO-1826	Thanks for keeping me honest, I was mainly just spitballing for the name. "accumulo version" or "accumulo about" or whatever sound just as reasonable to me. Don't they say that naming things is one of the hardest parts of computer science?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1826
17083	ACCUMULO-1826	I didn't find how to get all those using pure Java. Do you think it is a good idea to spin up a bunch of processes for the extra information? I guess we can just read some stuff out of {{/proc}} but I am super not-excited about hand parsing any of it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1826
17084	ACCUMULO-1824	Is this a documentation wish-list item that you'd like to see included somewhere? I'm not sure I understand the scope of this issue, aside from general frustration that may be had.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1824
17085	ACCUMULO-1824	Here's a [ticket | https://issues.apache.org/jira/browse/HBASE-5754] I posted to HBase's mailing list. I tried to include every detail I could think of so they could reproduce the problem: * versions of everything * reproducible, executable test case * documented strangeness: no problem at a million entries, more problems at a billion * multiple tests performed * feedback... I responded to the ticket with more information Is every user going to be able to do this? No. But most users can provide feedback, and logs. Perhaps the fix is something like: {noformat} $ ./bin/accumulo issue Please angryuser, describe what made you want to post a new issue? Describe, what unexpected behavior you first noticed: > shell is hung, can't scan my xyzzy table Accumulo 1.6.0-SNAPSHOT version abc12134... on hadoop 2.2.0, zookeeeper 3.4.0, ... you appear to have 3 node system, with 1.2G of debugging logs. I can deliver this to [dev@accumulo.apache.org]. Y/n? : y Note that the logs often contain small details about the data you are storing. If this is not a test system, we recommend not sending the logs. Do you want to send your logs? N/y? y {noformat}	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1824
17086	ACCUMULO-1817	Would this be using the [metrics2 stuff from Hadoop Common|http://hadoop.apache.org/docs/stable/api/org/apache/hadoop/metrics2/package-summary.html] within Accumulo so we can then use the Ganglia integration from there? or just trying to do something similar to the Ganglia provider itself? (I'd prefer the former, but it means adding a dep on hadoop-common)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1817
17087	ACCUMULO-1817	What's the level of effort for supporting both the new way and the old way with a scary warning message?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1817
17088	ACCUMULO-1817	bq. What's the level of effort for supporting both the new way and the old way with a scary warning message? That's a good question -- I haven't stumbled across anything that would allow for programmatic configuration of Metrics2 as a way to support the old mechanism. I'll poke a little bit, but my gut reaction is that it would be higher than I'd want to do :)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1817
17089	ACCUMULO-1806	Is this a dup of ACCUMULO-1658?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1806
17090	ACCUMULO-1804	Have a tar.gz need to upload..need to cancel	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1804
17091	ACCUMULO-1804	[~aarongmldt], thanks for the code! Digging through some of the source, it looks like you used the C++ interface for the Thrift proxy to connect R with Accumulo. That, combined with some of the API calls, it looks like raccumulo requires at least Accumulo 1.5.0 and an instance of the Thrift proxy running. Is that correct? It's been quite a while since I've used R; where do you see this fitting into the Accumulo community? Generally speaking, we have two avenues which integration code like this falls into. # Inclusion in core Accumulo codebase # A "contrib" project of Accumulo For #1, this would typically require a committer to sign up to ensure that the code is well-maintained as Accumulo itself grows. It is held to a certain level of testing and has a good expectation of working as expected since it would be released with Accumulo itself. For #2, a contrib project is a means for Accumulo to keep related, developed code near Accumulo. These projects typically follow their own schedule and aren't crucial to a release of Accumulo itself. To me, it seems like a contrib project is the best location for it at the moment. What do you think? Other committers? Do you intend to maintain and add additional functionality to raccumulo as people use it and find bugs or improvements? Thanks again. It's awesome to see contributions like this!	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1804
17092	ACCUMULO-1804	I love teh idea of this contribution, but I'm -1 (non-binding) on the pull request, because it places it in the main project repo. I'd like to see this in a contrib repository. That will take a ticket to INFRA to make a repo, and one of the Accumulo committers will have to do it. The current process for getting a new contrib in Accumulo is just an email to dev@ to get consensus on taking things on. Part of that discussion is the relative merit of having this as a contrib versus having it live in github. Aaron, I see you sent a brief message a few days ago to user@ about this. Could you (or Phil?) copy dev@ into that thread with a brief description of the motivation for adding it as a contrib? Aaron, could you please post your request to [Apache's review board|https://reviews.apache.org/] so that we can provide feedback via an Apache system?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1804
17093	ACCUMULO-1804	[~ctubbsii] Oh, I didn't event think about that. Good point. Do you remember what the ASF rules for non-ASF licensed code is? Can we even place this as a contrib if we wanted to? If it's a contrib, isn't it also subject to our release process as well (which the GPL would throw a wrench into)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1804
17094	ACCUMULO-1804	I installed R on an Ubuntu virtual machine. Then I installed RStudio and was able to run R scripts through my host browser. Then I downloaded the attached gz file. When I un'tarred the file, I saw no README or other indications how I should install or configure. What did I miss?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1804
17095	ACCUMULO-1804	This thread has started me on the path to create a Vagrantfile to run R and RStudio in a VM. Would you be interested in helping to extend that Vagrantfile to add the integration steps? Then we'd have a one step approach to see this integration working.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1804
17096	ACCUMULO-1802	[~afuchs], could you break those into separate tickets?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1802
17097	ACCUMULO-1801	What's the difference between this iterator and a stack that consists of a column family filter for the projection columns + the predicate input columns, the user predicate, and a column family filter for the projection columns?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1801

#	Requirement	Questions	Answered	Successfully Answered	URL
17098	ACCUMULO-1797	So in that case we just update the release notes to include some notion of how people should report said tests? Maybe as a part of casting a vote on the release? Should we still then require that at least one of the test runs be against each of the hadoop.profile versions (e.g. 2.0, 0.20)? Just to ensure we have some level of confidence when we say things like "works on Hadoop 2" ?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1797
17099	ACCUMULO-1797	bq. Should we still then require that at least one of the test runs be against each of the hadoop.profile versions (e.g. 2.0, 0.20)? I would veto a release that does not run through all of the test on at least one version of Hadoop and Zookeeper. This should shake out most pure Accumulo bugs. It would be nice to find and fix as many integration bugs as possible before release, but realistically we can only require what will actually be done. So what versions should we test (0.20, 1.2, 2.0,?) and who will test? I can run test against at least one version.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1797
17100	ACCUMULO-1797	Okay, then in the interests of not backing ourselves into a corner, we'll document that having a hadoop.profile doesn't imply that a release is officially supported against that profile. Instead the update to the release process will only require that the test set is run against whatever the default profile is? And then whichever versions come in with the release votes make up the supported set for the release? Or are all versions outside of the default profile just informational?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1797
17101	ACCUMULO-1797	When I said veto earlier, what I meant was "vote -1". bq. we'll document that having a hadoop.profile doesn't imply that a release is officially supported against that profile. Does "officially supported" == "we ran our test suite against a version and documented anything special"? bq. whichever versions come in with the release votes make up the supported set for the release? That sounds reasonable, assuming the reporter did not have to take an undocumented step to make the test run.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1797
17102	ACCUMULO-1797	bq. Perhaps making a more discrete testing plan for committers to follow on their "flavors of choice" (OS and Java, Hadoop, Zookeeper versions) would get committers more involved? Are you looking for something more than the testion section on the [relase page|http://accumulo.apache.org/governance/releasing.html]? If so I would be happy to provide more detail. bq. First off, I think that we should be testing Accumulo against each "stable" version of Hadoop at the time of that release – so, for 1.6.0, this equates to 1.2.1 and 2.2.0. I would love for release testing against 1.2.1 and 2.2.0 to be done, I am not against testing on multiple versions. Below are my thoughts : * Doing release testing takes a lot of time, If only one person decides to do this I don't its fair for the community to require anything of them. The community should be thankful for their contribution. I hope more people do release testing and I will be happy to help anyone who wants to tackle it. I suppose another way of looking at it, is that if we do require it then it may force someone to take it on inorder to get a release out? In the case that no one does this, then what do we do? * Whatever release testing we do should be documented. We have not really done this in the past. I think this information would be very useful. There was earlier discussion we were having where we trying to decide if it was ok to support one commercial flavor of Hadoop and not another. In this thread [~benson_basis] advised that we have the option to take whatever is contributed. If code, documentation, etc are contributed to support X and Y, but not Z then thats what we have. There is no need to wait for someone to contribute code and documentation for Z. Also we do not need to drop support for X and Y because no one contributed anything for Z.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1797
17103	ACCUMULO-1797	bq. Are you looking for something more than the testion section on the relase page? If so I would be happy to provide more detail. Nice, I actually didn't know that existed. There's a bit more we could write down in terms of "running functional tests" and how to invoke findbugs/pmd, but those are likely better suited as their own write-ups. In regards to testing on multiple versions of Hadoop, I'm happy to help out (if this wasn't painfully obvious already). The only limitation I see is, given that write-up, people sitting at home aren't going to be able to run the tests on a cluster. They might be able to run continuous-ingest/randomwalk, but that's a pretty low bar of catching bugs that I would expect to have been caught already by unit tests. Along the same lines of what Benson said, we test what we have cycles to or people volunteer to test, and that's what we have. We just need to make sure we meet the bar we've set in that release guidelines for the project.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1797
17104	ACCUMULO-1797	Bump. Can a committer take a look at this please?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1797
17105	ACCUMULO-1797	the PGP v GPG distinction isn't part of this change, it's the original text. Can we do that bit in a follow on?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1797
17106	ACCUMULO-1797	{quote}the PGP v GPG distinction isn't part of this change, it's the original text. Can we do that bit in a follow on?{quote} Sure... and, it's not that important. I had gotten a little confused between the dashes for the patch and the markdown syntax.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1797
17107	ACCUMULO-1796	Does the same problem happen with 1.5.x compiled with the Hadoop 1 profile and run on a Hadoop 2.2.0 cluster? If it does, I'd be much more inclined to do option #3. One of the complications with option #2 is that we don't appear to be doing it for any other of our releases, so it'll have more unknowns.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1796
17108	ACCUMULO-1796	I saw the patch on ReviewBoard. Do you have this patch as a git formatted patch?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1796
17109	ACCUMULO-1794	[~busbey] looking at the RB for this, the only thing still open is our discussion about reworking the agitator.pl to sync up with your changes here. Are you happy with the state up on RB, and then we can open a new issue to track reworking the (current) agitator? One thing I will note after ACCUMULO-1960, it is likely going to be desired that the script be smart enough to handle the case where a single user is running all components in the stack and not force them to start the agitator as root.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1794
17110	ACCUMULO-1794	Just applied locally and did a little bit of testing on 1.5.1 -- Overall looks good to me, but I did think of a couple of things that I'd like to see eventually (or at least get some feedback on the ideas). # Alternate between using haadmin and {{kill -9}}'ing the Namenode. We shouldn't see a difference here, but it would be nice to test coordinated failover and automatic failover # Some more validation before anything else: Can the user sudo to the hdfs admin user as they claim? Do the executables (hdfs, sudo) exist? Does the namespace provided exist (or can we find any namespaces if we're using all of them)? Can we find namenodes for the namespaces configured? The only other thing I'm curious about is when the script tries to choose an random namenode to make active, could we ever get in that block while ZFKC is in the middle of transition? In other words, is it possible to have no active namenodes while automatic failover is happening and we get an error because we try to force the transition?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1794

#	Requirement	Questions	Answered	Successfully Answered	URL
17111	ACCUMULO-1794	bq. 1. Alternate between using haadmin and kill -9'ing the Namenode. We shouldn't see a difference here, but it would be nice to test coordinated failover and automatic failover As I mentioned to [~kturner] in the review, doing this will require a heuristic. We can ask the HDSF admin tools for the hostname corresponding to the namenode id, but picking out the namenode process will be version dependent. I think he and I agreed that that sort of thing was better left to something like BigTop, since it attempts to work across projects. bq. 2. Some more validation before anything else: Can the user sudo to the hdfs admin user as they claim? opened as ACCUMULO-1982 about using sudo to users generally. bq. Do the executables (hdfs, sudo) exist? The existing tests for executability should cover this, no? Or are you looking for more specific error messages? bq. Does the namespace provided exist (or can we find any namespaces if we're using all of them)? Both of these cases are handled by the current error checking. the error message for the former is confusing (the message complains of a missing configuration value). bq. Can we find namenodes for the namespaces configured? This is covered in the current error handling. bq. The only other thing I'm curious about is when the script tries to choose an random namenode to make active, could we ever get in that block while ZFKC is in the middle of transition? In other words, is it possible to have no active namenodes while automatic failover is happening and we get an error because we try to force the transition? Yes, this is certainly possible. As things currently are, we'll simply log a message that this happened and try again the next time around. I couldn't think of anything else worth doing in that case. Note that it's also possible for an automatic failover to have changed which namenode is active while we are in the block that says to use the failover command. In that case, if there are only 2 namenodes we'll just do a no-op failover that says everything went fine. If Hadoop adds more than 2 namenodes per nameservice in the future, then I don't know what it will do but I know we'll log it and try again later. :)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1794
17112	ACCUMULO-1793	I don't think the merge to 1.5.1-SNAPSHOT worked as intended. I'm still seeing 2.0.4-alpha as the dependency in the top-level pom.xml Actually, I'm really confused now. The hadoop-1.0 profile in 1.4.5-SNAPSHOT now activates 0.20.203.0? Also, it looks like we also need to update the default hadoop dependency from 1.0.4 to 1.2.1? It was my understanding that 1.4 was going to keep the default 0.20.* hadoop dependency, and then it would make sense to have profiles that will automatically compile against 1.2.1 and 2.2.0. Did I miss discussion on this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1793
17113	ACCUMULO-1793	Right, but the hadoop-1.0 profile sets hadoop.version=0.20.203.0. Is 0.20.x binary compatible with hadoop-1? It just seems really weird to me that hadoop-1.0 uses a 0.20. And yes, I think 1.4 should also depend against 1.2.1 when using a hadoop-1 profile to maintain consistency across our other active branches. I had assumed neither hadoop-1.0 nor hadoop-2.0 profiles would be active by default. I can help here if you want/like.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1793
17114	ACCUMULO-1793	[~elserj], Hadoop 0.20 is part of the Hadoop 1.0 line. Since our examples in the README showed building against 0.23 by specifying the hadoop 2 profile and the version (and 0.23 is part of the Hadoop 2 lineage), it didn't seem odd me to have 1.4.x default the Hadoop 1 line to 0.20. But I'm pretty used to Hadoop versions being terrible at this point. The binaries built against 0.20.203.0 should be binary compatible with the Hadoop 1.x releases, but should not be binary compatible with the 2.x releases. Similarly, the binaries built against 2.2.0 won't be compatible with Hadoop 1.x releases (including 0.20.203.0). Unless we backport (or create) some fancy shim work under ACCUMULO-1796. Which part feels glued on? Having a third profile? Since 0.20 is a Hadoop 1 version, I'm not sure that we need to. If it makes the insanity of Hadoop versions easier to follow, I don't think it hurts. We could just change the name of the Hadoop 1 profile to Hadoop-0.20 and not have 1.4 speak to 1.x. As an aside, I think the patch would be better suited to ACCUMULO-1795 since that's the ticket that deals with building against 0.20.203 and non-Hadoop 2 compat.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1793
17115	ACCUMULO-1793	bq. Hadoop 0.20 is part of the Hadoop 1.0 line Blarg. Ok, this is just me being stupid. Blah blah, old hadoop versioning, blah blah. bq. Which part feels glued on? So, when we released 1.5, we started this practice of having the hadoop-1 and hadoop-2 profile which compiled against "good" versions of Hadoop for each line (read as: most recent "stable-ish"). Presently in 1.5 and beyond, this gives us hadoop-1->1.2.1 and hadoop-2->2.2.0. With the introduction of these profiles back in the 1.4 series, which previously only knew about 0.20.203.0, it's the same but different. Suddenly, if someone knows what we do in almost two major releases (1.5 and 1.6), and perform the same action (-Dhadoop.profile=1.0), they aren't going to compile against 1.2.1. It just doesn't sit right to me for 1) the 5th point release to suddenly have different build action than the previous 4, and 2) equivalently named build actions that do different things. That being said, I'm not even sure I like the patch I attached here to clarify things (pretty much why I didn't just commit it). I'm not sure what would be best. Like you said, with hadoop-0.20 being a part of "hadoop 1", but the "hadoop 1" profile using a version of hadoop in the 1.x series in 1.5, we're going to stomp on some approach we previously took. I haven't come up with anything that seems more straightforward yet.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1793
17116	ACCUMULO-1792	What's the reasoning for backporting ACCUMULO-600 as part of this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1792
17117	ACCUMULO-1792	Just that the discussion around it was in favor and removing it during this set of backports meant I didn't have to figure out getting the version left behind in 1.4 to run on Hadoop 2. Prefer it broken out?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1792
17118	ACCUMULO-1792	Could you specify what concerns you have back on [the thread discussing wikisearch in 1.4|http://mail-archives.apache.org/mod_mbox/accumulo-dev/201310.mbox/%3CCAGHyZ6Lw1W%2BWKO%3DN9yyxcRT2M_Y5UqXVTBFgf1UwX2c06y49-w%40mail.gmail.com%3E]? Our public api declaration is pretty specific about where people can rely on things not changing. FWIW, There is work to get wikisearch under 1.4 working on hadoop 2. A quick glance at the current contrib repo looks like it's find for Hadoop 2. It seems silly to backport those changes, particularly if the contrib repo works with a 1.4 cluster (which is the ask from Keith that I mentioned needing to answer).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1792
17119	ACCUMULO-1792	[~ecn], did you clean up your working directory before building? The line that is failing references o.a.a.core.Constants.METADATA_TABLET_COLUMN_FAMILY. That field exists in 1.4.x, but does not in 1.6.0-SNAPSHOT+. So it looks like your classpath is getting a 1.6.x version of that file and a 1.4.x version of o.a.a.server.util.Initialize. Could you try running `git clean -xdf` and then rebuilding and rerunning? Also note that as of the patch attached, building requires specifying both the hadoop profile and the hadoop version: {code} $> mvn -Dhadoop.profile=2.0 -Dhadoop.version=2.2.0 clean package && mvn -Dhadoop.profile=2.0 -Dhadoop.version=2.2.0 assembly:single -N {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1792
17120	ACCUMULO-1789	What did you change the timeouts to?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1789
17121	ACCUMULO-1789	How do I apply the scaling factor to the tests in 1.6, since they're not python based anymore?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1789

#	Requirement	Questions	Answered	Successfully Answered	URL
17122	ACCUMULO-1787	[~kturner] showed interest in this ticket. If you also think there is still interest, can you re-open? I can create a compaction strategy that supports both compression types.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1787
17123	ACCUMULO-1787	What release is this slated for?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1787
17124	ACCUMULO-1785	Bump. Could a committer take a look at pushing this?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1785
17125	ACCUMULO-1785	What branch is this patch applicable to?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1785
17126	ACCUMULO-1785	[~busbey] I applied the patch. There were merge conflicts when applying to 1.4.5-SNAPSHOT which I resolved. When I merged it into 1.5.1-SNAPSHOT there were also merge conflicts which I resolved. It merged cleanly from 1.5 to 1.6. Can you check the 1.4 and 1.5 conflict resolution?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1785
17127	ACCUMULO-1785	I just noticed that these changes also impacted how ACCUMULO_HOME is used. I'm not sure if it was intentional or not. Before these changes, bin/config.sh *always* calculated ACCUMULO_HOME based on the location of itself (config.sh). Since the changes for this issue, ACCUMULO_HOME is only calculated, when ACCUMULO_HOME is not already set or is set to the empty string. From an error-catching perspective, should we add something that does some validation of the provided ACCUMULO_HOME when we don't calculate it? I can easily see cases where the user mistypes an export of ACCUMULO_HOME or programmatically generates an incorrect path and suddenly none of the classpaths work.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1785
17128	ACCUMULO-1785	That was an intentional change. Can we file a follow on ticket for the verification of a set ACCUMULO_HOME?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1785
17129	ACCUMULO-1785	bq. That was an intentional change. Can we file a follow on ticket for the verification of a set ACCUMULO_HOME? Ok, I wasn't sure because things suddenly broke for me after this change, despite not setting ACCUMULO_VERIFY_ONLY. Opened ACCUMULO-1997.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1785
17130	ACCUMULO-1782	Is that record time? Thanks! *salute*	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1782
17131	ACCUMULO-1782	Bump! Can this ticket be resolved? Danke.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1782
17132	ACCUMULO-1778	I'm not sure I like the idea of exposing details of server configuration through the public API when it isn't needed for client operations. I'm not sure client code should have any knowledge that any monitor services even exist. At the very least, I'm concerned about this being a part of the public API. What if we replace the monitor with better monitoring solutions? What if the user isn't running a monitor? What if it is running multiple monitors? Would a utility class be better suited to this? Or maybe just put it in HdfsZooInstance rather than Instance itself? Will adding it to the Instance interface disrupt the API in undesirable ways? Is there a precedent for exposing a monitoring service within the service being monitored? It seems a bit awkward for a service to assume it's being monitored at all, never mind that it's monitored with a very specific type of monitoring tool (in this case, one that has a location attribute that it shares with the services it is monitoring).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1778
17133	ACCUMULO-1778	Can this ticket be closed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1778
17134	ACCUMULO-1773	Don't we need fully qualified URIs to handle multiple namenodes anyway?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1773
17135	ACCUMULO-1773	bq. Don't we need fully qualified URIs to handle multiple namenodes anyway? Any new file refs stored in the metadata table in 1.6.0 will use full qualified URIs. However relative paths can still exist in metadata from earlier versions. Unrelated to your question, the following situation is one that I was thinking of. Excluding relative paths, there can still be problems if configurations differ spatially and/or temporally # Tserver A writes a ref to metadata for hdfs://foo.com:123/accumulo/tables/9/default_tablet/0.rf # Later Tserver B writes a del marker to metadata for hdfs://foo:123/accumulo/tables/9/default_tablet/0.rf (its configured differently than Tserver A) # Accumulo GC deletes hdfs://foo:123/accumulo/tables/9/default_tablet/0.rf In the situation above a file thats referenced was deleted. The problem is that foo:123 != foo.com:123 even though they resolve to the same IP. If the Accumulo GC used /tables/9/default_tablet/0.rf instead of the fully qualified URIs, then it would not delete the file.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1773
17136	ACCUMULO-1773	bq. Don't we need fully qualified URIs to handle multiple namenodes anyway? The GC should create normalized relative paths for comparison. Absolute paths would still be stored in metadata table and used if the GC actually needs to delete something.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1773
17137	ACCUMULO-1772	Is this a strong requirement for ACCUMULO-118, or is this a nice utility that we will be needing at some point?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1772
17138	ACCUMULO-1771	Is this a dupe or abstraction of ACCUMULO-1766?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1771
17139	ACCUMULO-1771	bq. Is this a dupe or abstraction of ACCUMULO-1766? Its not a dupe of ACCUMULO-1766, but its caused by different code making the same assumption. The code treats the first configured URI specially. I am thinking the list of uris should be used as list of possible places to store new files. The configured list of uris does not represent all of the locations that Accumulo has stored files in the past. # User initializes with URI_1 in config, runs w/ this for a bit, then changes config to contain URI_2 and URI_3. # User initializes with URI_1 and URI_2, runs w/ this for a bit, then changes config to URI_0, URI_1, and URI_2 # User initializes with URI_1 and URI_2, user wants to reinit so they clear URI_1 (but forget to clear URI_2), User initializes again with URI_1 and URI_2 # User initializes instance 1 with URI_1. User initializes instance 2 with URI_2. User shuts down both instances and adds URI_2 to instance 1 config. In scenario 1 Accumulo will probablly not start because because it will look for the instance id and data version at URI_2. In secnario 2 Accumulo will look for instance id and data version at URI_0. In scenario 3 initialization may succeed even though URI_2 is dirty. Scenario 4 shows adding a dirty URI to an existing instance. A possible solution is the following. * init checks and initializes all URIs configured at the time it runs (upgrade would also need to do this) * master, tserver, etc check the set of URIs when they start and ensure instance id and data version are consistent across all. If server processes ever reload this prop, the check would need to be redone. * provide a utility to initialize a new URI. This utility would fail if the URI is not clean. The utility would have to be run before adding the URI to the Accumulo config. Another possible solution is to store the list of URIs in zookeeper and provide instance operations for mutating that list.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1771
17140	ACCUMULO-1771	What's the use case for running Accumulo with different kinds of volumes? That seems like it would be rife with consistency problems, and I would have a tough time recommending to users that they should stripe their data across HDFS and NFS, for example. Or even HDFS and MapR - the added complexity in that space is huge (as we're seeing from this and related issues). This JIRA may not be the place to have this conversation, but I'd definitely want to see us moving closer to what the rest of the Hadoop community is doing so that we can benefit from the 'free work' that it may provide.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1771
17141	ACCUMULO-1771	bq. We can move data version to zookeeper, and put zookeeper instance name in accumulo-site.xml to get the instanceId, rather than read from the hdfs path. How do you plan doing the migration to that? And does this mean that you would have to init and then distribute a new accumulo-site.xml to all of the nodes?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1771

#	Requirement	Questions	Answered	Successfully Answered	URL
17142	ACCUMULO-1769	The existing functional tests generate their own configurations to prevent clobbering known instances. Will this interfere?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1769
17143	ACCUMULO-1769	bq. The existing functional tests generate their own configurations to prevent clobbering known instances. Will this interfere? Would not want test to pick config from ACCUMULO_HOME/conf where instance.dfs.uri is defined and use the production accumulo dir. I am not sure if that could happen with this patch, buts that a concrete example of something that could possibly go wrong.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1769
17144	ACCUMULO-1769	Presuming ACCUMULO-1658 gets fixed, is the general case still a concern? As a part of that we can update the README to call out the need to have a configuration that doesn't point to your production information. How conservative do we want to be here? I was under the impression that the functional tests were already defensive in their generated configurations, since giving them access to any part of the production configuration is dangerous. The primary motivator for reading classpaths from accumulo-site.xml is so that we can cut down on how many places we have to maintain things like support jars for hadoop client. An alternative would be to switch to relying on `hadoop classpath`. The downside to this is that it's overly broad, given that for Hadoop 2 we only declare hadoop-client as needed.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1769
17145	ACCUMULO-1766	Seems like #1 is only bad in the case that an HDFS instance is removed, which can only happen if there is some migration path. #2 would probably be less vulnerable to breaking if the migration of namenode data is managed outside of Accumulo. Does anyone know if such a thing is possible. If we do go with #1 then we should have a script to handle the root tablet directory pointer migration and documentation around it.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1766
17146	ACCUMULO-1766	-1 for upgrade scripts. Do you remember nothing?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1766
17147	ACCUMULO-1761	So has this been addressed?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1761
17148	ACCUMULO-1758	Is this best implemented in Accumulo itself, or managed through whatever query application is doing the expiration? I don't think Accumulo has a concept of sessions.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1758
17149	ACCUMULO-1757	I was thinking about what the API change look like. Below is a possibility. It seems like you want the user to be able to pass a BatchWriterConfig XOR a BatchWriter. {code:java} class BatchDeleterConfig { public BatchDeleterConfig(BatchWriterConfig bwc){} public BatchDeleterConfig(BatchWriter bw){} public void setNumQueryThreads(); } class Connector { ... public abstract BatchDeleter createBatchDeleter(String tableName, Authorizations authorizations, BatchDeleterConfig config) ... } {code} Wouldn't you also want the ability to pass in a batch reader? The API above does not include that. Another approach may be allowing users to pass in a thread pool like ACCUMULO-1224	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1757
17150	ACCUMULO-1757	Isn't the binning/location cache static, though?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1757
17151	ACCUMULO-1755	Limiting binning to one thread may be unnecessary. Could bin in parallel and do a synchronous merge, similar to map reduce job w/ one reducer.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17152	ACCUMULO-1755	At a glance, wouldn't a simple queue that we could dump mutations on be a simple change (would just require some memory mgmt controls)? Clients can push onto the queue to buffer mutations as long as there is free memory. Heck, we could even let users pass in the queue (and totally ignore the memory control problem). The BW can then consume and bin mutations in batches from the queue. It seems like that might be more simple to me, but I might be missing some details too :)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17153	ACCUMULO-1755	bq. At a glance, wouldn't a simple queue that we could dump mutations on be a simple change (would just require some memory mgmt controls)? In hindsight, this is pretty much what MutationSet is already doing :)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17154	ACCUMULO-1755	Dave, thanks for resurrecting this ticket! Looks like your implementation is oriented towards the "don't hold a lock while binning" branch. This should work well for a multi-threaded client sharing a BatchWriter, but probably won't affect the performance of a single-threaded client. Do you have any perf tests you're using to evaluate how well your solution works? I don't remember whether any of our examples use a multi-threaded "mutation producer" pattern. One thing we might consider is whether we want to try to optimize the case of users writing simpler single-threaded clients and letting Accumulo handle all of the parallelization. That might lead us to an alternative implementation of moving the binning operation to another thread.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1755
17155	ACCUMULO-1755	Sounds like you have a good scenario to support the changes you're proposing. Do you have any performance measurements or profiling results you can share?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1755
17156	ACCUMULO-1755	No performance numbers, just seeing BLOCKED threads in a stack trace :-). I'll see what I can do about getting some performance numbers with and without my final patch. Do you think continuous ingest would be a good framework for this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1755
17157	ACCUMULO-1755	Github user dlmarion commented on a diff in the pull request: https://github.com/apache/accumulo/pull/75#discussion_r54292664 --- Diff: core/src/main/java/org/apache/accumulo/core/client/impl/TabletServerBatchWriter.java --- @@ -699,7 +701,16 @@ else if (Tables.getTableState(context.getInstance(), table) == TableState.OFFLIN } - void addMutations(MutationSet mutationsToSend) { + void queueMutations(final MutationSet mutationsToSend) { + queueThreadPool.submit(new Runnable() { --- End diff -- Are you suggesting re-using the existing ExecutorService?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17158	ACCUMULO-1755	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/75#issuecomment-189513944 > Addressed Keiths comments in the update. What about comment about exceptions?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17159	ACCUMULO-1755	Github user dlmarion commented on the pull request: https://github.com/apache/accumulo/pull/75#issuecomment-189523047 I added a catch clause to startProcessing to catch the RejectedExecutionException that can be thrown. Were you thinking of something else?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17160	ACCUMULO-1755	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/75#issuecomment-189526782 > I added a catch clause to startProcessing to catch the RejectedExecutionException that can be thrown. Were you thinking of something else? What you did is needed, but I was thinking of something else. If the call to `addMutations()` in the Runnable throws an unexpected exception, nothing will be done about that exception. The batch writer should transition to an error state. Thinking need to do something like the following. ```java void queueMutations(final MutationSet mutationsToSend) { queueThreadPool.submit(new Runnable() { public void run() { try{ addMutations(mutationsToSend); } catch (Exception e){ updateUnknownErrors(e); } } }); } ```	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1755

#	Requirement	Questions	Answered	Successfully Answered	URL
17161	ACCUMULO-1755	Github user keith-turner commented on a diff in the pull request: https://github.com/apache/accumulo/pull/75#discussion_r54319650 --- Diff: core/src/main/java/org/apache/accumulo/core/client/impl/TabletServerBatchWriter.java --- @@ -638,6 +648,21 @@ public MutationWriter(int numSendThreads) { queued = new HashSet<String>(); sendThreadPool = new SimpleThreadPool(numSendThreads, this.getClass().getName()); locators = new HashMap<String,TabletLocator>(); + binTimer.schedule(new TimerTask() { + @Override + public void run() { + MutationSet m = queue.poll(); + while (null != m) { + try { + addMutations(m); + } catch (Exception e) { + updateUnknownErrors("Error processing mutation set", e); + } + m = queue.poll(); + } + } + }, 0, 500); --- End diff -- Does this mean every 500ms check for something on the queue to bin? Whats the rational for this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17162	ACCUMULO-1755	Github user dlmarion commented on a diff in the pull request: https://github.com/apache/accumulo/pull/75#discussion_r54321264 --- Diff: core/src/main/java/org/apache/accumulo/core/client/impl/TabletServerBatchWriter.java --- @@ -638,6 +648,21 @@ public MutationWriter(int numSendThreads) { queued = new HashSet<String>(); sendThreadPool = new SimpleThreadPool(numSendThreads, this.getClass().getName()); locators = new HashMap<String,TabletLocator>(); + binTimer.schedule(new TimerTask() { + @Override + public void run() { + MutationSet m = queue.poll(); + while (null != m) { + try { + addMutations(m); + } catch (Exception e) { + updateUnknownErrors("Error processing mutation set", e); + } + m = queue.poll(); + } + } + }, 0, 500); --- End diff -- Well, in the previous revision I was adding a MutationSet to an ExecutorService that had an unbounded queue and 1 thread. That's not immediate. Can MutationSets be binned concurrently? If so, then I can change the timer from a fixed delay to fixed rate. Thoughts?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17163	ACCUMULO-1755	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/75#issuecomment-191289655 > which is a findbugs issue, and I don't see the issue that it's complaining about Can you reproduce it locally via `mvn verify -DskipTests -Dcheckstyle.skip`? I know the jenkins output can sometimes be... a little weird to parse for w/e reason.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1755
17164	ACCUMULO-1755	Thanks, Dave. Got any perf testing results that show how much this improved things?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1755
17165	ACCUMULO-1755	I took the test that I created and ran it against master and my feature branch with 1 to 6 threads. I didn't see much difference, but looking back at it now I think its because the test pre-creates all of the mutations and adds them as fast as possible. The test is really for multi-threaded correctness rather than performance. In the new code there is still a synchronization point when adding the binned mutations to the queues for the tablet servers. The send threads in the test (local mini accumulo cluster) must be able to keep up with adding of the binned mutations. I don't expect that to be the case in a real deployment. Good news - performance wasn't worse. I think a better test is to write a simple multi-threaded client that creates and adds mutations to a common batch writer. Then, time the application as whole trying to insert N mutations with 1 to N client threads. The previous implementation blocked all client threads from calling BatchWriter.addMutation(), meaning the clients could not do any work. In the new implementation the clients will be able to continue to do work, adding mutations, and even binning them in their own thread if necessary, before blocking. I'll see if I can re-test with this new approach in the next few days. Do you have a different thought about how to test this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1755
17166	ACCUMULO-1748	When the test then fails for me, should I be concerned? Does it indicate that whatever was done to make the pause unnecessary is not effective?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1748
17167	ACCUMULO-1748	Is this still an issue in 1.5.1?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1748
17168	ACCUMULO-1744	The merge command in the shell has the size option. Is this what you are looking for?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1744
17169	ACCUMULO-1741	Several of us mentioned in ACCUMULO-391 that an API change would be best for 1.6.0. Should old setter methods (like addIterator(), setInputTableName(), fetchColumns()) be deprecated or just removed altogether? I ask because I'd like to design a configuration object like TableQueryConfiguration that takes a table, iterator settings, ranges, and columns that can be passed into the InputFormat like so: {noformat} public static void setTableQueryConfigurations(JobContext job, TableQueryConfiguration... queryConfigs); {noformat}	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1741
17170	ACCUMULO-1737	Bit of a tangent, but what ordering guarantees does a BatchScanner have?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1737
17171	ACCUMULO-1729	[~vines], [~mberman] - Do you know which one? It would be nice to link the two issues. If something ever comes up with the other one and we need to do a (partial) revert, then it would be useful to know that it might adversely affect this JIRA as well.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1729
17172	ACCUMULO-1728	How would this work? What is the proposed mechanism for storing and checking this flag?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1728
17173	ACCUMULO-1728	I didn't have any particular mechanism in mind when I created this ticket. It was in response to your comment in an email thread: bq. You could implement a ZK flag that reports which mode (set by the Master) the entire instance is running in, but I think this should be used only as a fast-failure when constructing a secure client connector. Were you imagining any particular implementation? Also, I see this is a lower priority than SSL support in general, so I was feeling less pressure to get this into 1.6 than the patch on ACCUMULO-1009. Do you agree with that, or do you think this is a required feature for any SSL support to get accepted?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1728
17174	ACCUMULO-1722	[~ecn] [~ctubbsii] I am not familiar with namespaces or their permissions. They seem duplicative of system permissions, is that correct? There also seems to be an accumulo namespace and a "" namespace. Shall I skip the "" namespace? I think everything looks ok, but was wondering if you could check the syntax in the resulting files when you get a chance. Steps: 1. Apply patch 2. mkdir /tmp/backup 3. Run the following commands in the shell: createnamespace foo config -ns foo -s table.cache.index.enable=false createtable baz createtable foo.baz createuser me setauths -u me -s a,b,c,d grant Namespace.CREATE_TABLE -ns foo -u me grant Namespace.ALTER_TABLE -ns foo -u me grant Namespace.BULK_IMPORT -ns foo -u me createuser you grant Table.READ -t foo.bar -u you grant Table.READ -t foo.baz -u you 4. Run the following command on the command line: ./accumulo admin dumpConfig -a -d /tmp/backup 5. Optionally, you can run them separately: ./accumulo admin dumpConfig -s -d /tmp/backup ./accumulo admin dumpConfig -n -d /tmp/backup ./accumulo admin dumpConfig -t foo.baz -d /tmp/backup ./accumulo admin dumpConfig -u -d /tmp/backup	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1722
17175	ACCUMULO-1722	[~dlmarion] wrote: bq. They seem duplicative of system permissions, is that correct? Namespace permissions are not exactly duplicating system permissions. They also have some permissions that look like table permissions. bq. Shall I skip the "" namespace? No. The "" namespace is the default namespace. This is the one that is used for tables with no namespace qualifier. {{createtable foo}} would put {{foo}} in the default namespace. This is a quite common namespace to use, since it was logically the only one available prior to 1.6.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1722

#	Requirement	Questions	Answered	Successfully Answered	URL
17176	ACCUMULO-1717	How do you have Hadoop configured on the system(s) in question? It looks like Hadoop has defaulted to use the local filesystem. If the INFO statements below are not enough, I suppose there could be a more detailed message. I would not say that the Initialize should fail when the local filesystem is used as this is a product of the indirection supplied by the FileSystem object. Maybe a more detailed warning could be issued? There are components using the underlying Hadoop filesystem in local mode, specifically the MiniAccumuloCluster.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1717
17177	ACCUMULO-1713	It looks like those were both done by [~ecn]. Eric, are we good to close? Should the fix version include 1.4.5 and 1.5.1?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1713
17178	ACCUMULO-1709	Can't we handle all of these in Start?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1709
17179	ACCUMULO-1709	You mean remove the catch blocks from Master and TabletServer main, or rethrow the exceptions? I guess those catches were added for the extra logging. Do you think propagating the exception would be better than exiting immediately?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1709
17180	ACCUMULO-1709	[~stevel@apache.org]/[~billie.rinaldi@gmail.com], are you two happy with this issue? Can the JIRA be resolved?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1709
17181	ACCUMULO-1706	So is this wrapped? I see a commit a few days ago	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1706
17182	ACCUMULO-1705	I updated to 1.32 in the root pom.xml and everything seemed to keep working just fine. ITs all pass, shell and admin commands work fine. Not bothering to attach a patch because generating/applying it seems like a lot more work than a committer just bumping the version. [~mdrob], is there any reason you thought this would be a substantial task?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1705
17183	ACCUMULO-1695	Do we already do this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1695
17184	ACCUMULO-1691	Seeing ACCUMULO-2935 crop up makes me a bit worried about this change going into 1.6.1. Have other things silently changed that we haven't realized yet?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1691
17185	ACCUMULO-1690	Commit ba507fcdd1a77cd1edd441ffe954d74da2c433a0 in branch refs/heads/1.4.5-SNAPSHOT from [~ctubbsii] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=ba507fc] ACCUMULO-1690 Update rat plugin to 0.10	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1690
17186	ACCUMULO-1689	org.apache.accumulo.minicluster.MiniAccumuloClusterGCTest is failing for me in 1.5.1-SNAPSHOT. When I look in the GC logs for the unit test I see the following : {noformat} 2013-09-11 18:31:53,740 [gc.SimpleGarbageCollector] FATAL: org.apache.thrift.transport.TTransportException: Could not create ServerSocket on address 0.0.0.0/0.0.0.0:50091. org.apache.thrift.transport.TTransportException: Could not create ServerSocket on address 0.0.0.0/0.0.0.0:50091. at org.apache.thrift.transport.TNonblockingServerSocket.<init>(TNonblockingServerSocket.java:89) at org.apache.thrift.transport.TNonblockingServerSocket.<init>(TNonblockingServerSocket.java:68) at org.apache.thrift.transport.TNonblockingServerSocket.<init>(TNonblockingServerSocket.java:61) at org.apache.accumulo.server.util.TServerUtils.startHsHaServer(TServerUtils.java:215) at org.apache.accumulo.server.util.TServerUtils.startTServer(TServerUtils.java:278) at org.apache.accumulo.server.gc.SimpleGarbageCollector.startStatsService(SimpleGarbageCollector.java:421) at org.apache.accumulo.server.gc.SimpleGarbageCollector.run(SimpleGarbageCollector.java:212) at org.apache.accumulo.server.gc.SimpleGarbageCollector.main(SimpleGarbageCollector.java:164) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.accumulo.start.Main$1.run(Main.java:101) at java.lang.Thread.run(Thread.java:662) {noformat}	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1689
17187	ACCUMULO-1689	Are you running a local instance at the same time? I wonder if I should switch the GC to use "non-standard" ports? Or the random port stuff? Wasn't the random port stuff in 1.6? Answering my own question: yes, 1.6 ACCUMULO-1664	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1689
17188	ACCUMULO-1689	bq. Are you running a local instance at the same time? Maybe I was. I just tried running it again and it was ok.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1689
17189	ACCUMULO-1682	I've been playing around with swapping out the source iterator underneath the intersecting iterator for things that I want to query in ranges. I'm wondering if it would work to take your composite index approach and actually have the source iterator modify the key (to always be the lower bound of the range query if in fact the key is within the query range) before the key ever makes it to the intersecting iterator. In other words, let's say the term in question is AGE=>30 && AGE<=40. The intersecting iterator only needs to know that I care if AGE=30 and should tell each source iterator to begin there. If the source iterator knows about the range query AGE>=30 && AGE<=40, it can continue to iterate the keys that match but modify them to AGE=30 when getTopKey() is called. This way the intersecting iterator sees what it asked for and when the full document is retrieved, the actual value is pulled out. Excuse my ignorance if this is a big nono. I know that certain optimizations performed on scans make it very dangerous to modify keys. I can't foresee any conflicts with the order of the keys in the source iterator since I'm always rounding down and I'm scanning through increasing keys to start with. I'm assuming this has already been tried and doesn't work. Am I correct? UPDATE: Nevermind- I see better your original case about column qualifiers not being able to be returned in order unless they are somehow sorted. It works when you are querying a discrete value since the column qualifiers will be sorted already. Skipping through multiple different column families causes this not to be the case.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1682
17190	ACCUMULO-1677	Is it not possible to just load a native library off of the java classpath?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1677
17191	ACCUMULO-1675	[~ctubbsii] this really seems like a bug w/ the maven plugin. Would it be correct to say "its not correctly following maven conventions?" I switched it from improvement to bug. Feel free to switch it back if you think this is not the case.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1675
17192	ACCUMULO-1664	Does thrift let you get the local socket address after you open the port? The master does advertise its address in its lock.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1664
17193	ACCUMULO-1664	[~ecn], not sure I understand your comment. Are we hand editing generated sources now?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1664
17194	ACCUMULO-1654	Nice catch. Would it be possible to create a test that recreates this issue inorder to prevent regressions? Is there a sequence of operations (compactions, scans, splits) that will reproduce this issue? Do you think this "secret key handling strategy" deserves its own ticket? I'm thinking about the 1.6.0 release notes and wondering if it deserves its own line. It seems like this is a performance bug fix that does not change the encryption API or configuration? If its not a feature, its probably ok to lump it in with the other bug fix because. Bug fixes against unreleased features eventually just add noise to the release notes.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1654

#	Requirement	Questions	Answered	Successfully Answered	URL
17195	ACCUMULO-1654	Unfortunately I cannot think of a good way to create a test around this. We only found the bug when the OS and filesystem caching lined up *just so* such that one only got a few bytes back from reading an IV rather than the whole 16 bytes. Creating that kind of a boundary condition within a real-world scenario with any consistency as to become part of an automated suite.... I'm just not sure how I would go about it. Fiddle with a bunch of OS settings? This one was definitely wiley and took a while to appear during a continuous ingest run (many hours). There was no consistent set of actions that lead to its appearance. As far as the secret key handling strategy goes, I'd just consider it as part of the overall bug fix for the encryption stuff. No need for a separate line.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1654
17196	ACCUMULO-1654	I think what you would have to do to deterministicly replicate this is mock up a file system that just gives one byte at a time. That could be done as a simple wrapper that is inserted in some of the existing unit tests. Might be worth it if it is that easy, but I could go either way. Is this something a system like FindBugs would have caught?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1654
17197	ACCUMULO-1654	bq. Unfortunately I cannot think of a good way to create a test around this. I suspected this, but was not sure. So for release testing we probably need to run continuous ingest w/ and w/o encryption inorder to catch bugs like this. However this may be forgotten. Running random walk on a cluster may have caught this also. Maybe random walk needs the ability to generate a random config for Accumulo. This would allow random walk to run with encryption sometimes and sometimes without. Actually I suppose this random config generator could be use for continous ingest also. bq. Is this something a system like FindBugs would have caught? Seems possible. I suppose if you are calling read() w/o checking how much was read, that could be detected via static analysis. If find bugs does not have this check, it probably should.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1654
17198	ACCUMULO-1654	I applied the patch. I added the ticket number to the commit message. I just noticed that SecretKeyCache.ensureSecretKeyCacheInitialized() catches an IOException and only logs it, was that intentional?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1654
17199	ACCUMULO-1652	Is there accompanying documentation that should be updated with this? If we don't already, it would be nice to enumerate the environment variables that are configurable through $ACCUMULO_CONF_DIR/accumulo-env.sh	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1652
17200	ACCUMULO-1643	Are we specifying the *earliest* version we support, or the best version? I would like to update it, but I'm just the accumulo guy, and not the poor person sitting on top a heap of integration that happens to use accumulo.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1643
17201	ACCUMULO-1643	Best version, although I don't think it matters. AFAIK, compiling against 1.0.4 or 1.2.1 hasn't caused me any strife to date (it would also go against the whole semantic versioning...) From what I read, Hadoop typically said "if you're using 1.0.x, update to 1.1.x", and then "if you're using 1.1.x, update to 1.2.x", which would also imply to me not having any issues doing this.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1643
17202	ACCUMULO-1643	bq. will a 1.2.1 compiled version run fine against a 0.20.203? Building against 1.2.1 runs just fine against 0.20.205.0 locally. As the saying goes, it works on my machine.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1643
17203	ACCUMULO-1643	This was reopened for 1.5.1. Is it done now?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1643
17204	ACCUMULO-1641	the Wikisearch contrib is also guilty of this and should be fixed. Can someone add a component for "examples" so we can mark this properly?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1641
17205	ACCUMULO-1634	It means Ctrl-C will abort and will not forcibly shutdown the servers. Do you have a suggestion for a better message or a better way to interact users during shutdown?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1634
17206	ACCUMULO-1634	One thing I'd like to have clear in the message, is that Ctrl-C will only cancel the *forced* shutdown. It will not abort the managed shutdown that has already been attempted (and which may still succeed). Aside from that, my only concern is the lack of brevity in these messages. I'd prefer to keep script messages under 80 chars per line. How about separate messages for the forced shutdown case, and the successful shutdown + kill unresponsive servers case? Forced shutdown case: {code} Invalid password or unable to connect to the master Initiating forced shutdown in 15 seconds (Ctrl-C to abort) Initiating forced shutdown in 5 seconds (Ctrl-C to abort) {code} Normal shutdown case: {code} Accumulo shut down cleanly Utilities and unresponsive servers will shut down in 5 seconds (Ctrl-C to abort) {code}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1634
17207	ACCUMULO-1627	See [discussion on ACCUMULO-1734|https://issues.apache.org/jira/browse/ACCUMULO-1734?focusedCommentId=13780097&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13780097] for background. The method {{Mutation.equals(Mutation)}} is problematic for implementing {{equals()}} properly in a subclass. In short, for some object {{o}} of a subclass type {{SubMutation}}, a call to {{o.equals(m)}} where {{m}} is a {{Mutation}} will be bound to {{Mutation.equals(Mutation)}}, not {{SubMutation.equals(Object)}}. This is because the {{Mutation}} method has a more specific argument type, so the compiler prefers it over the proper one in the subclass. This is a problem for preserving symmetry for {{equals()}} as well as for checking additional fields in the subclass. {code:java} Mutation m = /* some mutation */; SubMutation sm = /* some other mutation */; sm.equals(m); /* <- bound to Mutation.equals(Mutation) */ {code} To remedy this I had thought to protect/rename {{Mutation.equals()}}, relegating it to a utility comparison method which it appears to be. However, as Keith suspected, this would break ABI compatibility. Code compiled against the old version would throw {{NoSuchMethodError}}. Another option is to have each {{Mutation}} subclass also override {{equals(Mutation)}}. It's messy but it should work. So, what is the best path according to the development policies for Accumulo? "Fix" the API at the expense of breaking the ABI, or introduce (modest) warts to work with what already exists? Thanks all.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1627
17208	ACCUMULO-1623	Janitor?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1623
17209	ACCUMULO-1621	When we dropped RPM/DEB packaging, I believe we also dropped all of our service scripts. Can this issue be re-scoped or closed out?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1621
17210	ACCUMULO-1612	Why is this a won't fix?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1612
17211	ACCUMULO-1602	Don't tablets already have fixed length id's? I thought this was what corresponded to the HDFS directory. Isn't the point of hashes that they are one-way and hard to reverse?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1602
17212	ACCUMULO-1602	bq. Don't tablets already have fixed length id's? I thought this was what corresponded to the HDFS directory. KeyExtent has a getUUID() method which basically hashes the extent. I think the monitor uses this. Older versions of Accumulo may have used this for the HDFS dir, newer versions use a counter from zookeeper. The tablets directory could be used instead of the hash, it would be shorter and should be unique. bq. Isn't the point of hashes that they are one-way and hard to reverse? The utility I was thinking of would scan the metadata table hashing each tablets extent until it finds the extent in question. If the tablets dir were used, then a special utility would not be needed. Could just grep the metadata table. There is one issue with using the tablet dir, it will not distinguish tablets that split. One of the children from a split will have the same tablet dir as the parent. Sometimes when debugging its important to distinguish the parent from the child. We could generate a unique id for each tablet and store this in a column in the metadata table. This allows the tablet to be found with grep in the Accumulo shell. The monitor could display this id.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1602

#	Requirement	Questions	Answered	Successfully Answered	URL
17213	ACCUMULO-1602	I am not a fan of a utility that, in essence, is running a brute force attack on the hash that we're trying to find. This also puts up a huge barrier to the usefulness of the naming. Why are we even worrying about this? Is there a max length to thread names that we need to fit in? What are the semantics for automatic expansion in a range compaction? Does it include the whole row? The whole tablet? Is this part of the API contract and specified in the docs, or is it simply a side-effect of the implementation?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1602
17214	ACCUMULO-1602	bq. Why are we even worrying about this? Is there a max length to thread names that we need to fit in? I look a the output of jstack alot and did not like the idea of a really long thread name. But thats just my preference. I looked at the current KeyExtent toString(). The max length it will print is around 160 to 170 bytes. Of course that could be modified to make it truncate more. Another thing to consider is that jstack output is shared for debugging purposes. If the jstack contained KeyExtents, then it possible it contains data from a table. A user may inadvertently share data from their table that they did not want to when sharing a jstack.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1602
17215	ACCUMULO-1602	bq. Another thing to consider is that jstack output is shared for debugging purposes. If the jstack contained KeyExtents, then it possible it contains data from a table. A user may inadvertently share data from their table that they did not want to when sharing a jstack. Don't we already have this "concern" with listscans? If you jstack a tserver running a scan, you have the extent in there. I personally don't think this is a big concern. If you have the *ability* to jstack/`kill -3` the tserver process and view the output, you likely can just print the rfile.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1602
17216	ACCUMULO-1602	bq. It was about the practice of sharing jstack output for debugging. If jstack contains KeyExtents, then jstack may need to be scrubbed of user data before sharing for debugging. Are you stating this as a concern about doing this or just stating it?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1602
17217	ACCUMULO-1602	bq. Are you stating this as a concern about doing this or just stating it? I think its an important consideration. It impacts how easy it is to share a jstack of Accumulo. This would be an incredibly useful feature, We have at least three options # Put table id in thread name # Put keyextent in thread name # Put tablet id in thread name I am in favor of 2 or 3. Considerations for #2 * easy to immediately use the jstack output * may have to be edited before sharing * user may not edit before sharing and leak info * more noise when examining jstack output Considerations for #3 * user can immediately share jstack output when seeking help, * have to take the tablet id and do a lookup somewhere to get the extent. * less noise when examining jstack output Personally I am in favor of #3. But I am not opposed to #2.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1602
17218	ACCUMULO-1601	Duplicate of ACCUMULO-1464. That ticket is marked with a fixVersion for 1.5.1 also, but this ticket isn't. Will this fix migrate easily to 1.5.x?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1601
17219	ACCUMULO-1599	Is there some magic that can automatically grab the test method name? Maybe munge the stack?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1599
17220	ACCUMULO-1599	Turns out we generate many gb of data in the mac directory in each test batch, so we certainly shouldn't save more than the last copy, and I'm not even sure it makes sense to do that. All I ever care about are the logs, so I'm thinking about cleaning everything else up after the test. Does that seem reasonable? Are there other subdirectories that should be nominated for saving?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1599
17221	ACCUMULO-1599	Is this ticket resolved?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1599
17222	ACCUMULO-1585	I'm signing up to do this, but I will take all the help I can get. I'll make it a priority to review/patch. If you can do it yourself, please reassign the ticket. Is there some approach we should take, similar to the way that datanodes self-identify to the namenode? That might be even better than the FQDN.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1585
17223	ACCUMULO-1585	[~nsfw], did you have a chance to do heavier testing on this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1585
17224	ACCUMULO-1581	Going to take a look at this but should it be limited by default? The JIRA references the behavior of the scan command which limits the output based on the showFewOpt and the default is not limited. For CLI consistency I would think we would want listscans to match the behavior of scan.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1581
17225	ACCUMULO-1578	bq. As a user, I would find it frustrating if I accidentally typoed my password and then proceeded type out a longer command, only to have it them come back as a failure. Are you referring to the shell when you say type out a command?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1578
17226	ACCUMULO-1576	How is this different from setting a different ACCUMULO_CONF_DIR for different runs? What is the benefit to having this as a separate variable?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1576
17227	ACCUMULO-1561	Superceded by ACCUMULO-2606 for 1.6.0 and later.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1561
17228	ACCUMULO-1557	Sounds great. I have a few more updates for the system tests -- are the continuous ingestion and random walk tests affected by these rewrites?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1557
17229	ACCUMULO-1551	I took a look at the changes. This API change builds around EntryConverter, which converts a Key+Value to an arbitrary Java object. As you mentioned this approach may not work well w/ heterogeneous data. Typo suffered from the same issue, however I think Typo was more rigid than EntryConverter. At first glance I think EntryConverter is an improvement over typo, it seems more flexible. I think some examples of using the API would make it easier to evaluate and understand it. Do you plan to address writing data? Does anyone know other APIs that abstract Accumulo's API? Typo and Gora were mentioned in the description. I have also seen [Accumulo-Fluent|https://github.com/Berico-Technologies/Accumulo-Fluent] When experimenting w/ Typo I took the approach of building a prototype API on top of the existing Accumulo API. I found this to be an easier way to explore the concept.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1551
17230	ACCUMULO-1551	[~ekohlwey]: I think the branches in github have diverged too much to make sense out of. Do you have an updated version of example code/static diff for your illustration?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1551
17231	ACCUMULO-1542	Commit 46b49658bce215fb8c74952a28491d7dddb929a7 in branch refs/heads/master from [~ecn] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=46b4965] ACCUMULO-1542 merge 1.5.1-SNAPSHOT to master	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1542
17232	ACCUMULO-1537	Chiming in here from the perspective of a "I'm new to Accumulo and I'm not sure how to make it build": When I'm coming into a new project that's Maven-based, generally the first thing I will do after pulling source is mvn clean install If this causes a build that takes 30+ minutes, I am pretty sure this isn't a great first-time developer experience (or a not-first-time developer experience). Might I suggest putting a few choice integration tests on the main "mvn clean install" pathway and hide the rest behind a profile?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1537
17233	ACCUMULO-1537	bq. Josh Elser and .. I think you mean "mvn verify". Nope, meant what I wrote which was essentially the same thing you elaborated upon. For day-to-day development, install is most likely not the lifecycle a user should be invoking. bq. If this causes a build that takes 30+ minutes, I am pretty sure this isn't a great first-time developer experience (or a not-first-time developer experience). Might I suggest putting a few choice integration tests on the main "mvn clean install" pathway and hide the rest behind a profile? I don't think we should be going against Maven conventions just to make a random first-timer have a "better" experience. In my opinion, if anything, we should go out of our way to encourage proper usage of Maven for the task at hand, which, again, is likely not `mvn install` as releases are available through maven-central (and wouldn't require a user to install the artifacts into their local repository).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1537

#	Requirement	Questions	Answered	Successfully Answered	URL
17234	ACCUMULO-1537	I'm having compile problems on trunk currently on the Testing module. The root cause is that some tests in ./test/src/main depends on code in ./test/src/test (MacTest is the base class of several and lives in ./test/src/test and is depended upon by tests in ./test/src/main) Moving MacTest to ./test/src/main requires junit be a main dependency as opposed to a testing dependency. With the patch attached, the Testing Modules builds. (currently running, will report in the morning).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1537
17235	ACCUMULO-1528	Maybe the VersioningIterator should return all values instead of non-deterministically picking one? Where are you thinking of finding the source identity and how would it be stored? On Fri, Jun 21, 2013 at 10:36 AM, Christopher Tubbs (JIRA)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1528
17236	ACCUMULO-1528	Do you mean when there are keys that are non-unique including timestamp?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1528
17237	ACCUMULO-1528	{quote}Do you mean when there are keys that are non-unique including timestamp?{quote} Which specific statement would you like me to clarify? {quote}If we go by file name/file timestamp...{quote} Acknowledged. I'm not proposing a solution, just identifying the problem.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1528
17238	ACCUMULO-1516	[~vines] - With the metadata table being split into two tables, I'm not sure how it would be sensible to keep that constant around. The code is still open for improvement, and I've been thinking about additional improvements under the parent ticket of this one (ACCUMULO-1481). For example, I've been working towards a cleaner API for metadata table operations (still designed for internal use). Do you have a comment on the MetadataServicer I added? I do plan on moving more code into that, and am open to feedback. However, I don't see how we can sensibly preserve the METADATA_TABLE_NAME constant, as any reasonable way somebody might have been trying to use this (despite it not being in the public API), will likely be broken by ACCUMULO-1481. I did try to preserve the Constants.NO_AUTHS as deprecated (and replaced it with Authorizations.EMPTY), because I thought that is something users would likely use. I'm also willing to deprecate other constants, as needed, but I don't see merit in preserving them all, since they were never part of the public API, and most aren't even useful outside of internal code.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1516
17239	ACCUMULO-1516	I was looking at this earlier today, and there is still a lot of znode stuff in the Constants file. Might it make sense to move that out to other files as well (likely in a follow on JIRA, since it would be out of scope for a root table subtask)?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1516
17240	ACCUMULO-1515	Is there any reason the README can't refer to the packaged user guide or the website for that information, though?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1515
17241	ACCUMULO-1515	The build.sh script does not seem focused on the person who wants to build from source and then deploy. Does it make sense to add the following? '--locally-deploy' option to perform "mvn package -DskipTests -DDEV_ACCUMULO_HOME=/var/tmp". '--make-deploy_artifacts' to perform "mvn package -P assemble". '--make-documentation' to perform "mvn package -P docs -DskipTests" With these options explicitly represented in build.sh (and a few comments) I can see the README file reduced in size.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1515
17242	ACCUMULO-1515	Right. Just to be clear, I didn't make the build.sh script. I only modified it to work better for creating and testing release candidates in 1.5.0. I'm not even sure I like it, but since it's there, I think I agree it could be modified to make it easier to build something ready for immediate deployment locally. However, I'm not entirely sure what the goal is. Is the goal to provide an easier mechanism for somebody unfamiliar with maven to build from source? Or, is the goal to document the correct maven command-line options to create the desired build artifacts, under the premise that the user is familiar with maven already (or is willing to become familiar with it as a prerequisite)? I think I'd almost rather just enumerate the profiles that are available on the command line, and explain how they affect the resulting artifacts. The nice thing about the script, though, is that it's easier to document once... and we just change the script, as needed, to activate different profiles if the necessary profiles change between releases.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1515
17243	ACCUMULO-1515	How about moving the maven profile documentation to docs/maven.html? The README file can point there. We can leave the build.sh as is. If there a few +1 votes for this change, I would not mind putting the initial draft together. I'm sure I'll miss details though. On Fri, Jun 14, 2013 at 11:17 PM, Christopher Tubbs (JIRA)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1515
17244	ACCUMULO-1515	[~bills] wrote: {quote}Would using to the README/INSTALL file convention a lot of projects have alleviate some of this concern? ....{quote} Possibly. The README could be very minimal and not change much, because the project description and links to more information won't change much. We'd have to clarify what the INSTALL file is, though. Moving it to a separate file doesn't necessarily help minimize doc maintenance or consolidate docs. Is INSTALL for the binary tarball, explaining how to run, or is it for the source tarball, explaining how to build, or both? At some point the contents... whether in README or INSTALL, have diminishing value, as we try to explain every possible scenario (building DEB/RPMs, building binary tarball from source, explaining all the maven profiles and options, etc.). If it were as simple as {{./configure && make}}, I'd be fine with that... because configure scripts and Makefiles in other projects often support making many different targets and the INSTALL files don't usually try to explain all of the different options. Unfortunately, we do, and that's where the value begins to diminish, becoming a tutorial on maven and instructions on every build option rather than a quickstart guide for 90% use case coverage. If we can reduce these instructions to a single one-line maven command, I'm all in favor. If we follow other examples, the INSTALL file should contain basic information about build dependencies and instructions on how to build the source in its entirety (PDFs, RPMs, DEBs, tarballs, everything) with default options (default hadoop profile, for instance). It should not be an attempt to explain every profile and build option, just as other projects don't attempt to explain every Makefile target or ./configure flag which is available. It's not too much to expect somebody diverging from the basic "build all" option to learn the build system and read the build configuration files.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1515
17245	ACCUMULO-1504	Added hashcode() and equals(Object obj) methods to the IteratorSetting class via eclipse method generator. They look good, do we need test cases?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1504
17246	ACCUMULO-1504	bq. do we need test cases? Yes, please.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1504
17247	ACCUMULO-1504	Hey Miguel, can you please follow the coding style for your new patch? I take it you're using eclipse given your comment, so you should be able to follow these instructions at the bottom of the page http://accumulo.apache.org/source.html	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1504
17248	ACCUMULO-1504	[~bluerain], thanks for your contributions. In the future, can you leave multiple versions of your patches attached to the tickets? It is nice to be able to trace history through comments and the patches.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1504
17249	ACCUMULO-1503	ACCUMULO-633 was the only issue addressed in the 1.3 branch since 1.3.6 was released. Since it's not a bug in 1.3, and since it would change behavior, I do not think it should even happen for 1.3.7, never mind warranting going through the release process.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1503
17250	ACCUMULO-1503	[~ctubbsii], is there still any work left to do on this issue?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1503

#	Requirement	Questions	Answered	Successfully Answered	URL
17251	ACCUMULO-1502	I would like to help out here. I have starting playing around with tracing and have seen the existing docs in docs/distributedTracing.html. I just have a few initial questions, technical and administrative. 1.) Should the documentation go in 1.5.1-SNAPSHOT? Seems like it would be reasonable to get into the 1.5 code 2.) Is the Administration chapter of the user manual the appropriate place? Maybe right after Monitoring? 3.) What server-side operations are traced by default? I have seen minor/major compactions and GC. 4.) Sampling rate is currently left explicitly to the user via CountSampler, correct? 5.) Is there a good way to trace a write without flushing the BatchWriter? I know the actual writes to the tablet servers are asynchronous from adding mutations so maybe it isn't possible. Thanks for any feedback that can be provided.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1502
17252	ACCUMULO-1502	Just added a patch to the Administration chapter of the docs. Does the User Manual HTML get generated somehow from the TeX? I also started examples at https://github.com/jaredwinick/accumulo-tracing-example that could be incorporated in with the other samples if desired. With respect to #5 above, I am unable to get the behavior you describe (tracing of writes without explicitly calling flush()) to work. I have attempted to make sure the asynchronous writes inside the BatchWriter have completed by setting the timeout very low, but still no luck.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1502
17253	ACCUMULO-1501	Hmm. Do you think the trace ttl should be a configurable property, or would it be better to pick a default value and make users change the property on the table manually?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1501
17254	ACCUMULO-1499	I have no objections. What's the benefit of moving it?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1499
17255	ACCUMULO-1499	Miguel, I took a look at your patch. Were you on an updated 1.6.0 branch when you made the change? The initialize() method in your patch seems to be missing some content. Here's the content that I have from 1.6.0: {code:java} /** * Set directories and fully populate site config */ MiniAccumuloConfig initialize() { if (!initialized) { libDir = new File(dir, "lib"); confDir = new File(dir, "conf"); accumuloDir = new File(dir, "accumulo"); zooKeeperDir = new File(dir, "zookeeper"); logDir = new File(dir, "logs"); walogDir = new File(dir, "walogs"); String classpath = System.getenv("ACCUMULO_HOME") + "/lib/.*.jar," + "$ZOOKEEPER_HOME/zookeeper[^.].*.jar," + "$HADOOP_HOME/[^.].*.jar," + "$HADOOP_HOME/lib/[^.].*.jar," + "$HADOOP_PREFIX/share/hadoop/common/.*.jar," + "$HADOOP_PREFIX/share/hadoop/common/lib/.*.jar," + "$HADOOP_PREFIX/share/hadoop/hdfs/.*.jar," + "$HADOOP_PREFIX/share/hadoop/mapreduce/.*.jar"; mergeProp(Property.INSTANCE_DFS_URI.getKey(), "file:///"); mergeProp(Property.INSTANCE_DFS_DIR.getKey(), accumuloDir.getAbsolutePath()); mergeProp(Property.INSTANCE_SECRET.getKey(), DEFAULT_INSTANCE_SECRET); mergeProp(Property.TSERV_PORTSEARCH.getKey(), "true"); mergeProp(Property.LOGGER_DIR.getKey(), walogDir.getAbsolutePath()); mergeProp(Property.TSERV_DATACACHE_SIZE.getKey(), "10M"); mergeProp(Property.TSERV_INDEXCACHE_SIZE.getKey(), "10M"); mergeProp(Property.TSERV_MAXMEM.getKey(), "50M"); mergeProp(Property.TSERV_WALOG_MAX_SIZE.getKey(), "100M"); mergeProp(Property.TSERV_NATIVEMAP_ENABLED.getKey(), "false"); mergeProp(Property.TRACE_TOKEN_PROPERTY_PREFIX + ".password", getRootPassword()); // since there is a small amount of memory, check more frequently for majc... setting may not be needed in 1.5 mergeProp(Property.TSERV_MAJC_DELAY.getKey(), "3"); mergeProp(Property.GENERAL_CLASSPATHS.getKey(), classpath); mergeProp(Property.GENERAL_DYNAMIC_CLASSPATHS.getKey(), libDir.getAbsolutePath()); mergePropWithRandomPort(Property.MASTER_CLIENTPORT.getKey()); mergePropWithRandomPort(Property.TRACE_PORT.getKey()); mergePropWithRandomPort(Property.TSERV_CLIENTPORT.getKey()); // zookeeper port should be set explicitly in this class, not just on the site config if (zooKeeperPort == null) zooKeeperPort = PortUtils.getRandomFreePort(); siteConfig.put(Property.INSTANCE_ZK_HOST.getKey(), "localhost:" + zooKeeperPort); initialized = true; } return this; } {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1499
17256	ACCUMULO-1499	Miguel, Would you like me to add you to the contributors page? Let me know if you want me to add an organization/timezone.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1499
17257	ACCUMULO-1499	[~sonixbp], what version of the thrift compiler are you using? It seems you have checked in slightly different thrift code than that generated by my version of thrift 0.9.0.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1499
17258	ACCUMULO-1499	bq. Would you like me to add you to the contributors page? Let me know if you want me to add an organization/timezone. I just added him, btw. I was chatting with him online and bugged him about his info :)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1499
17259	ACCUMULO-1498	bq. I'm attaching ctubbsii-git-suggestions.v1.txt with branching scenarios, based on conversations with Josh Elser and Keith Turner. I like most of this. Thanks for taking the time to write it down. My first concern was being crystal clear about when people can rebase/cherry-pick and when they must merge. Perhaps this is beating a dead horse though, "perform local development" should be sufficient. bq. I'm attaching ctubbsii-git-proposed-actions.v1.txt as a rough guide to the transition, assuming the previous attachment represents the target branching model. Questions: 1. [~billie.rinaldi@gmail.com] was correct when she said the site must stay in SVN to use the [CMS|http://www.apache.org/dev/cms.html#subversion] 2. Contribs move to their own repositories since they are separately versioned. See the number of Cordova repositories [here|https://git-wip-us.apache.org/repos/asf]. 3. Need to verify this one. Not sure if the infrastructure itself enforces ASF policy (cannot delete tags, cannot {{git push --force}} and rewrite history). Actions: 1. Need to remember to request this 4. Probably/Maybe? I'm not sure how robust git+svn clone is here. Verification/validation before the transition is complete would be necessary. 5. Not sure I follow this one. SVN branches will appear as git branches, and SVN tags as git tags. Verification/validation again. 6. Should pre-date 5 7. I hope to roll the existing source guide into the WIP doc I already started.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1498
17260	ACCUMULO-1498	{quote} Actions: 1. Need to remember to request this {quote} We could rename the branch ourselves if necessary, but probably they'd need to do it for us so it's the default branch when cloning. {quote} 4. Probably/Maybe? I'm not sure how robust git+svn clone is here. Verification/validation before the transition is complete would be necessary. {quote} Right now, I'm pretty sure the branches *don't* merge cleanly, so it would be very annoying if this made git branches not merge cleanly. If fixing it in SVN first will make git-svn and git happy later, we should do that. I'd really love to close the 1.3 branch prior to switching... it's one less branch we'd have to deal with. I'd also like to get this transition done *prior* to releasing 1.6.0, so we don't add another branch in the chain to deal with. {quote} 5. Not sure I follow this one. SVN branches will appear as git branches, and SVN tags as git tags. Verification/validation again. {quote} Right, but these are the long-running old versions branches we've discussed. This is related to number 4. The main point being that eventually, all the last tags in each series should merge forward cleanly. What we don't want to do is tag them without making them merge cleanly first... that would be *VERY* frustrating when doing bugfixes in older versions. {quote} 6. Should pre-date 5 {quote} Number 5 would create some more... unless you're including branch cleanup as part of the tagging process, then we're thinking similarly.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1498

#	Requirement	Questions	Answered	Successfully Answered	URL
17261	ACCUMULO-1495	Isn't that a completely separate way to just make the tabletserver "do less"? (sounds like its own JIRA issue if it's something you want to pursue..)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1495
17262	ACCUMULO-1494	Currently you can do bin/accumulo -add /path/to/example.jar:/path/to/example-dependency.jar com.example.ExampleMain --foo bar. Is that not sufficient?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1494
17263	ACCUMULO-1494	{quote}in the classpath local to the invocation, but [not] added to the classpath for any of the server processes{quote} How is this different from simply doing something like the following? {code} CLASSPATH=my.jar bin/accumulo shell {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1494
17264	ACCUMULO-1494	I will. Can you point me to information about the style formatter?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1494
17265	ACCUMULO-1494	[~bhavanki], is this issue resolved to your satisfaction?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1494
17266	ACCUMULO-1493	[~kbotzum], are you planning on submitting a patch for this that could be included in 1.6.0? I'm trying to weed out tickets that aren't going to be done for 1.6.0.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1493
17267	ACCUMULO-1491	What about firewalled networks? I imagine bundling them would be rather nice in that case. Maybe an alternative is to add a link on the monitor to the "most-up-to-date" version of the docs on a.a.o, but leave the release docs in place?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1491
17268	ACCUMULO-1491	this may sound crazy, but could we build an exported version of all our keep-in-one-place docs and include that in the binary dist? then have a config option in the monitor for pointing to the export vs live.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1491
17269	ACCUMULO-1490	What's the plan for the documentation? I'm concerned about making all the documentation only on the website, because that requires web access. Moving the html documentation pages to the website makes sense, and I might even be in favor of moving the user manual there. But it would be nice to keep minimal building information and examples packaged with our releases.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1490
17270	ACCUMULO-1490	Does anyone see value in using the Hadoop wiki [1] for documentation? It would certainly help as a place for keeping more up-to-date and dynamic documentation. The downside is that web access would be required, although it can generate printable pages. [1] http://wiki.apache.org/hadoop/	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1490
17271	ACCUMULO-1488	Do you think configuring the static inner class may throw some users? I think user will need to use a dollar sign for the last separator when configuring these, for example : {noformat} setiter -c org.apache.accumulo.core.iterators.user.BigDecimalCombiner$BigDecimalMaxCombiner ... {noformat} For 1.5.0 we have started adding since tags. The code should probably have a "@since 1.5.1,1.6.0" tag in the javadoc at the class level. Not sure if the inner classes will need it or if they will take it from the outer class. Also, there was some discussion about porting new features back to bug fix branches, and there was not really a consensus.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1488
17272	ACCUMULO-1488	What is the advantage of back-porting instead of contributing a jar that can be dropped into the lib/ext directory?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1488
17273	ACCUMULO-1488	{quote} What is the advantage of back-porting instead of contributing a jar that can be dropped into the lib/ext directory? {quote} I see none, and this would be my preferred method of supporting these in an older version, since they are completely modular and can be delivered in this way. At the *very* least, I think backporting a feature from a newer version when that newer version isn't even yet released (and therefore finalized in some way) creates an undue burden on *all* developers who could receive questions from users. If some users start having problems with this feature in 1.5.1 and later, or design code around it assuming it's available in 1.5.x, but have to migrate to 1.5.0 (due to a bug in 1.5.1, for instance), it will *not* be sufficient for me to tell them "[~afuchs] volunteered to take responsibility for that". Personally, I would like to see backported features (just features, not bugfixes) voted on for inclusion in a prior release line, precisely because of the additional support and maintenance risks.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1488
17274	ACCUMULO-1486	Are you sure we don't want to wait until the third RC of Accumulo 1.7.0 to solve this problem? We're kinda breaking with tradition here! ;-)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1486
17275	ACCUMULO-1483	Right, so it will include classes that aren't actually part of the published API if you must have something in the path to instantiate one of the public interfaces. e.g. instantiating ZooKeeperInstance requires that log4j be present, even though the client API doesn't make any claims about the log4j Logger. Making core take over this functionality is going to be ugly, and add some complexity to downstream people who don't actually want those transitives (since they'll need to do excludes). Better to have a different module, atleast for now. Do you have an exemplar user of this in mind?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1483
17276	ACCUMULO-1483	[~ctubbsii], what's the status? Please mark this as blocker if it needs to happen for 1.7.0 or update the fixVersion accordingly.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1483
17277	ACCUMULO-1482	I like the idea of making a system namespace and putting !METADATA and !!ROOT (assuming that also comes to fruition) into it. Assuming we do so, is there any more reason to put one to many !'s in front of METADATA (and ROOT)? I know [~mdrob] had some frustrations with jline2, history and ! being interpreted as a meta-character. From my understanding, the !'s were only to make them sort to the top of the monitor page for most people and be something unlikely for a user to want to use. Locking them into a system namespace seems to serve the same purpose then. Am I missing something else?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1482
17278	ACCUMULO-1482	{quote}is there any more reason to put one to many !'s in front of METADATA (and ROOT){quote} No, absolutely not. Implementing this feature would allow us to use normal table names for these system tables (we don't even have to make them upper-case to make them stand out). {quote}the !'s were only to make them sort to the top of the monitor page{quote} That's not the only reason. It's also that way for historical reasons, because, before tables had permanent IDs and could be renamed, table names (instead of IDs) were in the !METADATA table the !METADATA table's own tablets stored in the root tablet had to sort before other tables' tablets. (This is still a requirement, unless/until ACCUMULO-1481, but now applies to the table ID, which is why the tableID is now !0 for the !METADATA table).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1482
17279	ACCUMULO-1482	Does it makes sense to have a reserved "" namespace for internal use? Then we literally will have nothing reserved as a system namespace.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1482
17280	ACCUMULO-1481	Wait, what? Are we splitting and merging the root table now?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1481
17281	ACCUMULO-1468	Doesn't the maven-release-plugin automatically tag for you? Can it be configured to use signatures, or do we have to break automation to accomplish this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1468
17282	ACCUMULO-1468	bq. Doesn't the maven-release-plugin automatically tag for you? Can it be configured to use signatures, or do we have to break automation to accomplish this? Yes, it can do this automatically for you. I'm not sure if an option already exists that would do this automatically too. I'll look into it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1468
17283	ACCUMULO-1468	Should that document be merged with http://accumulo.apache.org/governance/releasing.html ?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1468
17284	ACCUMULO-1459	How do you know that a key will exceed a threshold until after you've passed it?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1459

#	Requirement	Questions	Answered	Successfully Answered	URL
17285	ACCUMULO-1459	bq. How do you know that a key will exceed a threshold until after you've passed it? I am thinking this can be pushed down to the deserialization code in rfile.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1459
17286	ACCUMULO-1458	The basic WordCount example itself can be used with a token file (just do -tf tokenfile rather than -p password), but I also added an example that doesn't use the ClientOpts argument parsing stuff. Also included a detailed explanation in README.mapred (and also found a small typo, forgot a -t in a command).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1458
17287	ACCUMULO-1454	This ticket make me laugh because my team at Toyrus.com developed rolling restart functionality back in 1999. We tracked the number of sessions assigned to each server using a load balancer. When a given server needed to be restarted, the load balancer changed the server's state to something like 'restarting' and did not send any sessions to it. The server would periodically check it's own status. If it saw 'rebooting' and it had zero sessions then it restarted itself. Is there some time delay before tablets are reassigned? Can the tserver restart within that window of time?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1454
17288	ACCUMULO-1454	[~medined] wrote: {quote}Is there some time delay before tablets are reassigned? Can the tserver restart within that window of time?{quote} I think that's the key to an elegant solution here: ensure a delay long enough for the tserver to come back and continue serving the tablets it had been, to avoid rebalancing the whole cluster, but not so long that a failure to come back would prevent re-assignment entirely.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17289	ACCUMULO-1454	bq. Maybe the 'restarting' status has a lease? Server must become responsive within a configurable time period. +1 bq. Currently when a tablet is closed, it interrupts running scans. Waiting for scans to finish is tricky, because it seems like you would not want to allow new scans to start. So while the scans running before close would see no delay, scans started after close will still see a delay. Would it make sense to "double host" tablets? Let existing scans finish on a tserver that is about to go down, meanwhile, load those same tablets elsewhere and point new scans to the new locations. At the end of the whole process, run one final balance to shake things out.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17290	ACCUMULO-1454	How do you plan to deal with port conflicts?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17291	ACCUMULO-1454	bq. How do you plan to deal with port conflicts? For any current version, the config could be modified to start up the "sidecar" processes with a port value of '0'. I've done pretty extensive testing in making sure we can bring up a cluster using completely dynamic ports.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17292	ACCUMULO-1454	I was thinking about use cases. * Administrator has 20 of 100 nodes with screwy java memory config, wants to fix those nodes w/ minimal impact. * Administrator has 100 of 100 nodes with screwy java memory config, wants to fix those nodes w/ minimal impact. * Administrator wants to upgrade cluster from 1.7.0 to 1.7.1 w/ minimal impact Are there any other important use cases? The first two are covered in the ticket description, I split them into two because one is a subset of tservers.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17293	ACCUMULO-1454	bq. If we expose a user level command to request that a tablet be moved to a given destination Do we necessarily care where they move? I don't have a good feel for the difference in cost between moving to a "sibling" tserver on the same node as opposed to a tserver on a completely different node. We also might just be fighting the balancer if we make it seem like the user has control over where a tablet is hosted. HBase has such a tool, no? Can we glean anything from their rolling upgrade support (good and bad)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17294	ACCUMULO-1454	bq. start a lot of new tserver instances This might be more difficult than it sounds. {{start-server.sh}} is fairly easy to manage, but {{stop-server.sh}} is pretty aggressive about just {{kill}}'ing the process on a host. I think we might have to expand on the shell scripts to really give an admin what they want. bq. migrate tablets <tserver instance src> <tserver instance dest> Would you consider this a FATE op that the master coordinates, or would you just add something directly to the TServer and let the client handle the coordination? Would that just bash ZooKeeper trying to migrate tablets off of ~200 tservers? (e.g. 200 tservers * 200 tablets)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17295	ACCUMULO-1454	bq. Would that just bash ZooKeeper trying to migrate tablets off of ~200 tservers? (e.g. 200 tservers * 200 tablets)? No more so than start-up. bq. Would you consider this a FATE op Yes: it needs to be part of the metadata processing state machine. Migrations are not presently a FATE op.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17296	ACCUMULO-1454	Is the upgrade case the same as the 'change configuration' use case? Why load/unload instead of move?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1454
17297	ACCUMULO-1453	One of the usual suspects for failed loads is walog recovery problems. The issue w/ RFiles mentioned in the issue description seems more tricky to isolate. If a tablet has a problematic file, the tablet will likely load successfully. Failure will occur when a scan of the rfile is attempted. When a tablet server fails, how do you know which tablet(s) caused the problem? Another approach to solving this issue may be to identify what can cause tablet server failure and try to defend against those. One possible cause could be key/value in a rfile that exceeds memory. This would be easy to defend against by making Accumulo refuse to load key/values that are too large. Another possible cause is an iterator that runs amok and consumes all memory. This is harder to defend against, ACCUMULO-1188 is one approach.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1453
17298	ACCUMULO-1452	Should the fix version for this include 1.4.5 and 1.5.1?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1452
17299	ACCUMULO-1444	[~sonixbp] I don't really understand what the title means. Can you expand here?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1444
17300	ACCUMULO-1444	My apologies for the late reply. Christopher is correct. On Tue, May 26, 2015 at 5:48 PM, Christopher Tubbs (JIRA) <jira@apache.org>	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1444
17301	ACCUMULO-1442	[~billie.rinaldi] - JLine2 is BSD Licensed. Google tells me that is Apache friendly, but I can't find an Official statement on it. Can you point me at something that definitive?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1442
17302	ACCUMULO-1442	Can somebody on a Mac test [^ACCUMULO-1442-v3.patch] and let me know if there are any issues?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1442
17303	ACCUMULO-1438	Umm, is this something we want to go moving around mid release cycle / this early into trunk? I thought we were trying to avoid tree conflicts.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1438
17304	ACCUMULO-1438	John, I'm sorry, I'm still getting used to our merge plan. Did I mess something up?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1438
17305	ACCUMULO-1438	That's what I was thinking as well. I should have confirmed it with the community first, but it seemed wonky to introduce that change in a future version and I don't want to split packages over different jars because that breaks OSGI environments. In the future, if I have a change to 1.5 and 1.6, should I wait to do the merge into 1.6? I'm still not completely clear on the plan there and I've noticed many others doing the merge from 1.5 into trunk. I need to add the new module to the assembly too so I reopened the ticket. Ill remove the rat plugin as well since there's no special includes to add. Sent from my iPad	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1438

#	Requirement	Questions	Answered	Successfully Answered	URL
17306	ACCUMULO-1438	bq. In the future, if I have a change to 1.5 and 1.6, should I wait to do the merge into 1.6? I'm still not completely clear on the plan there and I've noticed many others doing the merge from 1.5 into trunk. Its in some mailing list thread. We decided to hold off on major moves in trunk until after 1.5.1. The expectation is that we will find a lot of bugs that need to need to be fixed in 1.5 and merged forward. The plan is to apply bug fixes in 1.5 branch and merge them forward to trunk. We had a feature freeze on jan 25 for 1.5, but now its almost june and we still have not released 1.5. IMO we (myself included) need more discipline about making only necessary changes to 1.5. The goal should be to release 1.5, not continually improve it.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1438
17307	ACCUMULO-1426	Which documentation are you referring to?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1426
17308	ACCUMULO-1419	Hadoop has a 2.1.0-beta release, has anybody on our side done due diligence w.r.t API and code stability?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1419
17309	ACCUMULO-1419	Did you update the jenkins jobs?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1419
17310	ACCUMULO-1419	bq. Did you update the jenkins jobs? Nope, I didn't. Good call. Edit: Done	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1419
17311	ACCUMULO-1419	Can you point me to said warnings? I don't know what you're talking about.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1419
17312	ACCUMULO-1417	[~ecn] - looks like this made it into 1.5.0? Can we close the issue out?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1417
17313	ACCUMULO-1409	Will this add another bit of complexity? The single node accumulo instance is intended to what amount of entries? At what size does the having the native map truly affect performance? And to what degree? One billion entries ... Fifty Billion entries?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1409
17314	ACCUMULO-1409	bq. At what size does the having the native map truly affect performance? Its used to sort incoming data. Sorting data w/ C++ is faster than Java. A difference would be observed if you started writing to an empty Accumulo instance w/ native maps. bq. At what size does the having the native map truly affect performance? The performance of native maps should be invariant with data size. However sorting incoming data is only one part of what Accumulo does. Accumulo eventually has to do compactions, which have nothing to do with the native map. So I suspect you would see the largest performance increase from native maps when Accumulo is empty. I am thinking Amdahl's law kicks in as more data is written, compactions start and now native maps are only one part of the system thats processing data. It would be interesting to run experiments w/ and w/o native maps for 10M, 100M, 1B, and 10B entries.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1409
17315	ACCUMULO-1409	If bin/accumulo mini is used to start a Single Node Accumulo, the native maps should work out of the box (though the user will have to override the tserver.memory.maps.native.enabled property in the MAC's site config). Are we interested in seeing if native maps may help us in our pursuit to speed up the MAC during the automated integration tests?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1409
17316	ACCUMULO-1409	bq. Are we interested in seeing if native maps may help us in our pursuit to speed up the MAC during the automated integration tests? I suspect enabling them would not help for that case. Not much data is written. I think a lot of the time is spent in process startup and IPC snafus (i.e. accumulo init tries to talk to zookeeper before its fully up, it fails and then retries in a bit).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1409
17317	ACCUMULO-1409	bq. the only way I can guarantee the integration tests don't time out is if I enable the native maps with "-P native". Interesting, I wonder if this is because JavaGC is avoided and/or the native maps are faster at sorting. So native map support has recently been added to MAC. I have taken a quick look at the code, but not recently. It seems like native maps can be used when Accumulo itself uses MAC. However what about the case where an end user uses MAC? Can this be supported? It seems like supporting this would require putting jars maven repos that contain the shared libs.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1409
17318	ACCUMULO-1406	Should the command error/fail if the user specifies -scan and -all? This may indicate this user did not type what they were thinking (i.e. cut, paste, and edit a command adding the -scan option). I don't think commons cli can do this for us. AFAIK, can not specify an option as mutually exclusive w/ a group of options.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1406
17319	ACCUMULO-1405	bq. There's a lot of blocks already in place with the Instamo stuff. I would not say the blocks are in place at this point in time. Maybe after ACCUMULO-1378 is complete, instamo could leverage that work. When a lot of the tickets under ACCUMULO-1386 are complete and its easy to run a single node accumulo, how do we package that and make it available to users? Keep in mind that users will want to stop and start this instance, its not transient. Users will also want to configure it. Maybe maven is a way to do this, I am not really sure at this point. Seems like it would be something different than instamo, maybe a maven plugin for managing a single node accumulo instance. Will this work w/ native maps? Another option is creating a tarball, rpm, and deb. I think it would be useful to develop a list of the pros and cons of each approach. I agree w/ [~medined] we do want to be cognizant of the maintenance tail of whatever we create. We will have to continue to support in 1.7, etc.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1405
17320	ACCUMULO-1405	I just did an example assembly that copied all the dependencies (including provided) and packaged up the MAC as a shade jar. It was 32mb and that doesn't seem realistic to release inside the main tarball (which is... 11mb). I like the idea of having _developers_ pull down the dependencies on their own using a script in the tarball. We've also mentioned _non-developers_ using it to populate some data quickly for (possible) migration over to a fully-distributed cloud. In that latter case, can we guarantee that their system have maven installed or even connectivity to the internet? I know there must be downsides as well, but what if a user could do this: * untar accumulo-mini-1.6.0-bin.tar.gz * cd accumulo-mini-1.6.0 * ./bin/accumulo-mini -p propsFile	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1405
17321	ACCUMULO-1405	[~sonixbp], I have been thinking about a accumulo-mini-1.6.0-bin.tar.gz distro, from the perspective of a user. Why can't I run ./bin/start-all.sh and ./bin/stop-all.sh? Why don't I run init? Why do I configure tserver memory in a different place (props file vs accumulo-env.sh)? These are rhetorical questions. It seems that w/o working on the scripts (which I am not advocating for), this alternate distribution would present an entirely different experience. It also made me think about the monitor and gc, MAC does not start these. Do not really need to for unit test. Probably should start them for the single node use case. A maven plugin to start and stop a single node Accumulo instance and script in contrib to pull down deps into lib/deps seem like a good way to go.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1405
17322	ACCUMULO-1405	Now that you mention it, it doesn't. I would love to have a "mvn accumulo:run" mojo. That would be a HUGE win. Can we achieve this with your current StartMojo/StopMojo or can we plug in the MiniAccumuloRunner to kill two birds with one stone? Do we have a ticket for it yet?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1405
17323	ACCUMULO-1399	Personally, JEXL was the first thing that came to my mind. I was just thinking it would be neat if there was a feature that would allow you to create an alias that points to some file of commands. Maybe when creating the alias, you specify a file of commands and it copies it to some directory that can be recalled later. bq. http://www.beanshell.org/ Isn't that quite old and unsupported these days?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1399

#	Requirement	Questions	Answered	Successfully Answered	URL
17324	ACCUMULO-1399	{quote} The shell was originally designed to support pluggable commands. However, it was decided against making this a user-facing feature because of the confusion it might cause distinguishing between "official" features, and commands users might grow accustomed to on a particular customized installation (and because nobody had much interest in maintaining other users' commands that behave badly). {quote} since this was already thought of and decided against, should I abandon this idea? I'm thinking the answer might be yes, based on the above comment and the move to a CLI.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1399
17325	ACCUMULO-1398	Attached patch for review. Apply to trunk, here are the command options: dumpconfig -a dumpconfig -s dumpconfig -t t1,t2,t3 dumpconfig -s -t t1,t2,t3	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1398
17326	ACCUMULO-1397	[~mberman], are you still working on this? If so, feel free to assign it to yourself. Also, can you upload a patch file or open a review board when you are done? Having to get multiple {{.java}} files is a lot more cumbersome. Thank you!	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1397
17327	ACCUMULO-1395	[~John McNamee], would you be down for making some documentation updates for these changes? Removing the example configurations is a big change for users -- it would be good if we have something we can point them to that describes what was changed and how to use it. A new issue would be good to track that.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1395
17328	ACCUMULO-1395	Why are we suddenly adding new improvements on 1.6.0? Why did you commit this to 1.6.0, [~ctubbsii]?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1395
17329	ACCUMULO-1394	It does extend RegExFilter but it's in a different package than Filter. Does protected work in sub-packages?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1394
17330	ACCUMULO-1394	Is the link between jira and the code changes broken? On Thu, May 16, 2013 at 2:03 PM, Christopher Tubbs (JIRA)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1394
17331	ACCUMULO-1391	Is this a change to the MiniAccumuloConfig API or a completely new class altogether?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1391
17332	ACCUMULO-1390	Is this a duplicate of ACCUMULO-1380 or is this more specific to the features added in 1.6.0 that allow the MAC to be run as a single node development environment?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1390
17333	ACCUMULO-1390	bq. is this more specific to the features added in 1.6.0 that allow the MAC to be run as a single node development environment? Thats what I was thinking. Also, if its really easy to use, it may be used for more than development. I wonder if we will encounter the situation where a user starts with SNA then ask how to transition the data in it to a cluster.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1390
17334	ACCUMULO-1385	Changing global OS settings seems like an odd thing for an RPM to do. If a lot of RPMs did this, it seems like the order in which RPMs were installed could change a systems behavior. Is this a common things for RPMs to do? If so, is the installation order issue addressed in some way?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1385
17335	ACCUMULO-1385	bq. An example of an RPM that does add its own modular config to /etc/sysctl.d I did not even notice "/etc/sysctl.d/accumulo" in the summary. Putting it in /etc/sysctl.d resolves my concerns. Following this line of thought, should we up nproc and nofile in /etc/security/limits.d? It seems like adjustments to open file limits would need to apply to HDFS users more than the accumulo user, so that may be tricky.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1385
17336	ACCUMULO-1385	{quote}...should we up nproc and nofile in /etc/security/limits.d?{quote} It is my understanding that limits are security policies for user login sessions, and they don't affect daemons, so no, not unless we plan on these values affecting client code (in which case, they should be part of the client RPMs/DEBs).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1385
17337	ACCUMULO-1381	Would it be better if it took a URI instead of just a path that makes assumptions about places it could check? file:///path/to/file	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1381
17338	ACCUMULO-1380	[~ctubbsii] I'm finishing this up now. Is it too late to incorporate this into the RC?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1380
17339	ACCUMULO-1380	{quote}I'm finishing this up now. Is it too late to incorporate this into the RC?{quote} Too late for RC2, but if that fails, it can go in RC3. Don't let that stop you from contributing it to the 1.5 branch, though. It can go in 1.5.1 if RC2 passes a vote.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1380
17340	ACCUMULO-1378	The ability to get up and running with a single node Accumulo instance in a few minutes would be extremely useful. Setting up hadoop, zookeeper, and then Accumulo for a single node instance takes a while. MiniAccumulo is one option for achieving this. I think this will lead to adding complexity to MiniAccumuloCluster. For example people will want to be able to configure the amount of memory tservers use and make tservers use native maps. Maybe this is worthwhile because of the benefit. What other options are there?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1378
17341	ACCUMULO-1378	Also, is the goal really to completely replace Mock Accumulo?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1378
17342	ACCUMULO-1378	bq. Decomposing the core methods of the MAC into helper classes may give us more bang for our buck. [~sonixbp], are you proposing reorganizing the implementation of MAC or API changes? It seems we should rename this ticket to something like "Make it easy to run a single node Accumulo instance."	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1378
17343	ACCUMULO-1378	bq. are you proposing reorganizing the implementation of MAC or API changes? I'm referring to reorg of implementations for better code reuse- not API changes. You had me thinking about having a SingleNodeAccumulo that was different from the MiniAccumuloCluster. If we did that, many things would remain the same in both- processes get spun up, ports get configured, LogWriter gets created, shutdown hooks get registered, and vmargs optionally get attached to the processes for remote debugging. I've been teetering back and forth about this. On one hand, if the only difference between the two are the ability to re-use a directory/configuration, that doesn't seem to me like enough of a difference to separate them. On the other hand, if we want to treat one as a quick integration test environment and the other as a full development environment then we can define a broader set of features that differentiate the two. The latter seems like grounds for separation.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1378
17344	ACCUMULO-1378	bq. I've been teetering back and forth about this. [~sonixbp], could proceed with making changes to MAC to support single node Accumulo for now. Later on if we see good reasons for creating SingleNodeAccumulo, we can do that. bq. I'm ok with marking MockAccumulo deprecated, but not ok with removing it until MAC is "as fast". [~elserj], would you still be concerned about MAC's speed if ACCUMULO-1030 was implemented?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1378
17345	ACCUMULO-1378	bq. Josh Elser, would you still be concerned about MAC's speed if ACCUMULO-1030 was implemented? I think so. ACCUMULO-1030 seems to deal with a different subject, imo. To be clear, I'm just against _deleting_ MockAccumulo because I don't think we have an adequate replacement yet. There are known issues where MockAccumulo does things wrong; however, there are plenty of things it does right. As such, I don't see a reason to make other people _force_ to something new so soon after MockAccumulo was even introduced. I have no qualms against a strong recommendation to use MAC instead of Mock and a disclaimer that Mock may not be 100% accurate in implementation.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1378
17346	ACCUMULO-1377	Is this a plan to do the real fancy one mentioned on the mailing list or just fixing the issues with some buggy versions of browser from treating a link ending in tar.gz as a download?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1377

#	Requirement	Questions	Answered	Successfully Answered	URL
17347	ACCUMULO-1376	[~chentecaliente], functionally, the patch changes look good. A couple of small things: # What branch is this patch based on? # Make sure you use (or follow) the [codestyle guidelines|http://accumulo.apache.org/source.html]. Most notably, 2spaces instead of a tab character. # Could you add a test for this while you're making these changes? (1.5 ShellServerTest, 1.6.0+ ShellServerIT)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1376
17348	ACCUMULO-1376	Thanks for the comments Josh. 1) The patch is based on 1.6; however I imagine it can go back further. 2) I will take a look at the codestyle guidelines and make sure my IDE conforms to that. 3) I saw the ShellTest and didn't see a good way to add a test there, but did not see ShellServerTest/ShellServerIT so I didn't attempt to write a test there. I will write a test in ShellServerIT for 1.6, unless you think I should try to patch 1.5. Also, Is there documentation on the process to follow for submitting patches? I kind of just figured this out and I'm not sure if I did it properly. Thanks again Josh	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1376
17349	ACCUMULO-1375	bq. My last note - does thrift really need to be installed? IIRC, yes. You would need the language specific library for thrift installed to be able to use said thrift client code in that language (you need the python thrift library installed to use the python with the proxy)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1375
17350	ACCUMULO-1375	SVN#1478618. I updated the proxy/README file but feel it should be reviewed before this ticket is resolved. Can someone please take ownership of this ticket?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1375
17351	ACCUMULO-1375	bq. Can someone please take ownership of this ticket? got yer back.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1375
17352	ACCUMULO-1368	We may want to come to a decision about whether or not we want the MiniAccumuloCluster to be used for both integration tests AND as a quick-start development environment. Supplying a main method is simple but we should consider the packaging- is there going to be an option in accumulo-start's Main class? Will there be a different script for firing it up? Is it up to the user to jam this into a maven module with the necessary dependencies and call the main() method? Is a maven plugin going to be the sole caller of the method?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1368
17353	ACCUMULO-1368	I like how the proxy server main method takes one option, the name of a config file. I am thinking MAC main should do the same. bq. e an option in accumulo-start's Main class? I think this would be really nice to have. USers would not have to type the full class name. bq. Will there be a different script for firing it up? If a user can do "accumulo mini -p mini.props", would any other scripts be needed? bq. Is a maven plugin going to be the sole caller of the method? I think it makes sense to be able to call it via a maven plugin and the accumulo script.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1368
17354	ACCUMULO-1368	bq. Because we're relying on Hadoop & Zooekeeper's classpaths to provide many of the dependencies that we need to run the tserver, master, and init processes, using "accumulo mini" will only work if the two are installed. Good point, I was not thinking of that when saying we should add it to accumulo-start main. Now I understand your comment about packaging more fully. Are you thinking that launching via maven the only option, unless a single node accumulo package is created?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1368
17355	ACCUMULO-1368	bq. Are you thinking that launching via maven the only option, unless a single node accumulo package is created? I'd love to be able to run the MiniAccumuloRunner using the start.Main class. I can see that by omitting the commons and other dependencies from the lib directory must be saving a ton of space but it seems to be causing other issues as well- like some of the required dependencies missing from some versions of Hadoop. Are there any other reasons why we are not shipping many important dependencies with the assembled package? The current MiniAccumuloRunner is able to assume defaults so that a property file argument can be optional. It would be absolutely wonderful for running it to be as simple as 1) extract tarball 2) run {noformat}${accumulo_home}/bin/accumulo mini{noformat}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1368
17356	ACCUMULO-1367	Applied! Thanks for the patch, [~sonixbp]. Can you please do me a favor a double-check what I [committed|http://svn.apache.org/viewvc?view=revision&revision=1479227]? I had to massage your patch a little (that [~kturner] changin' packages), but it's probably fine.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1367
17357	ACCUMULO-1366	Can the problem be exposed by JMX? If so, the source (the iterator) is isolated from the reporting (i.e. aggregating)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1366
17358	ACCUMULO-1364	Did you determine if this was an HDFS bug?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1364
17359	ACCUMULO-1364	Your redundant file system wasn't redundant? Yeah. Oops. If we're misusing the dfs api, sure, we need to fix that; but, I don't see "Hadoop lost data!" as an Accumulo problem. Did you figure out what was happening at the dfs layer with the offending blocks?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1364
17360	ACCUMULO-1364	bq. My open question is whether we should hold the release of 1.5.0 until we can diagnose (or fail to replicate). My decision would be based on how much time is wanted to investigate. Obviously we're not there yet, but if this is going to be a few weeks or a month, I wouldn't want to wait. bq. our system clearly doesn't meet the level of reliability that we claim it meets I didn't think we advertise any redundancy above what Hadoop provides (in talking 1.5), so I don't really see that as false advertising (to put words in your mouth). By that measure, 1.4 was less redundant because we had only two copies of a WAL instead of the Hadoop default of 3. My view on this is that if you don't have the ability to say "oops, gotta reprocess" when faced with catastrophic, power failure, you should probably have a UPS to give you enough time to issue an orderly shutdown. Yes, it would be nice to figure out exactly what happened, but I know I don't have the knowledge/expertise to fully track the down the issue (nor the ability to look at your nodes to figure out what happened), and I'm not going to volunteer Eric (or anyone else, for that matter) to try and guess. Is someone other than just John looking into this (getting back to your 'should we delay 1.5' question).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1364
17361	ACCUMULO-1364	Couple of questions: 1) The logs that are corrupt, were they being written to actively when the power went out? 2) Are you testing with a current version of 1.5 or with the revision noted above? 3) Josh's question: how much time are you asking for?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1364
17362	ACCUMULO-1364	Looking through open HDFS tickets. Has anyone seen this [1]? It's not fixed, does that mean its still an issue, or maybe OBE? [1] https://issues.apache.org/jira/browse/HDFS-1944	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1364
17363	ACCUMULO-1364	Actually, the old loggers were thoroughly, though unintentionally, tested with power failure. At scale. More than once. In two weeks I could make this configurable, and users could make their own decision. But I would rather patch HDFS. I reproduced log failure on a VM this evening. The datanode just doesn't do anything to disk on a sync. [~billie.rinaldi], can you dig up an HDFS expert at Horton works who might know when/if this is fixed?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1364
17364	ACCUMULO-1364	Have we tried the same VM reset / plug pulling with 1.4?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1364

#	Requirement	Questions	Answered	Successfully Answered	URL
17365	ACCUMULO-1364	bq. Today, my client said (and I paraphrase) "Accumulo lost data? That's a blocker." No, Hadoop lost data. Taking what Eric says sounds like Accumulo closed a file in HDFS and data was not actually sync'ed to disk. bq. So, if you are using XFS and your system goes down hard, then you might lose data. "You mean having delayed write back might cause my data to not be written?" No filesystem is going to magically save you if your data wasn't actually written. How far are we going to go in regards to a user-space application on top of a user-space distributed filesystem to what happens in the kernel on one node to write bytes to disk (ignoring the fact that even that does not guarantee you have bytes on spindles and not just disk cache). I certainly don't hope this is going to go the way of "Accumulo requires ext3 with data=ordered". This discussion seems to be barking up the wrong tree to me. I'm glad Eric has been able to reproduce this at an HDFS layer (and I'm sure he'll work towards trying to get a fix there), but isn't a better solution here to make sure Accumulo has the means to work around this? The general consensus seems to be that people apparently think Accumulo should be resilient despite unreliable power. Perhaps I can shift discussion to that: are there things that Accumulo _could_ do on its own?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1364
17366	ACCUMULO-1364	Well, that makes me sad. At least we know now it's an issue throughout the chain and it's on our radars now. Out of curiosity, did it break the same way even though there were no appends? Or was it a different failure?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1364
17367	ACCUMULO-1364	What do you have in mind for testing? Just validating your results or something else?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1364
17368	ACCUMULO-1361	I noticed a problem with my patch anyways. Is this still something worth working on?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1361
17369	ACCUMULO-1360	Did this get fixed in 1.5.0?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1360
17370	ACCUMULO-1357	Thanks again [~sonixbp]! Looking at patch #1, I think enabling debugging should be a method on MAConfig. If its a method on MACluster then it should throw illegal state exception if called after start(). By putting it in the config object, it makes it more clear that it can not be enabled after start() it called. IMO ServerType should be in the mini package, its part of the MAC API. For setting memory, should we make that more strongly typed and less error prone? One way to do that is to use the Java TimeUnit convention. Maybe something like the following : {code:java} enum MemoryUnit { BYTE, KILOBYTE, MEGABYTE, GIGABYTE, etc } class MiniAccumuloConfig { . . . public MiniAccumuloConfig setMemory(ServerType, long memory, MemoryUnit unit){} } {code} Maybe this is not that useful if this usually ends up coming from a config file. If we do not do something more stronly typed, then we should at least validate the String passed into the set method immediately so that it fails fast.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1357
17371	ACCUMULO-1357	[~sonixbp]. I committed patch #2. I made some subsequent change in commit 1480705. Can you take a look at those? I noticed the methods relating to getting bytes were not public. When changing them to public, I thought that making them return a long would be more consistent with the setters than returning a string. I also change them names of the memory related getters and setters to make them more consistent. I was thinking the setDebug() and isDebug() names should be named somethng different. Something more specific that distiguishes from other types of debugging, like log4j debug level. Maybe something like setJDWPEnabled() and isJDWPEnabled(), those are much more precise, but kinda unsightly.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1357
17372	ACCUMULO-1354	Would this patch eliminate the need to do something like ACCUMULO-1402? If so, I'd be in favor of promoting this over that ticket. For the record, I have absolutely zero objection to using profiles for test dependencies, if that's what it takes.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1354
17373	ACCUMULO-1354	Oh, ok, I guess I could add the profiles back into the top-level pom. If they're there anyway, do you think I should put hadoop-tools and hadoop-distcp in them?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1354
17374	ACCUMULO-1350	Good point. Does shell --fake not currently work?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1350
17375	ACCUMULO-1349	Verified this with "mvn integration-test" Looks like it works well. Should we update jenkins to run "integration-test" phase by default?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1349
17376	ACCUMULO-1342	Why are giving this to meee? (Said as if I was the Doctor from Lost In Space)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1342
17377	ACCUMULO-1337	Can you please add some unit tests for this? And perhaps add comments in a javadoc explaining context - where and why would this iterator be helpful?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1337
17378	ACCUMULO-1337	Pushpinder, thanks for the patch. Some unit test would be nice. I do not think the continue case is correct. After an Accumulo client has received a batch of key values, it will reseek with the last key in the batch non-inclusive. So I am thinking the following code should not test the case 'sk.getColumnFamilyData().length() == 0'. The wholerowiterator always returns 0 len column families. This iterator does not do that. {code:java} if (sk != null && sk.getColumnFamilyData().length() == 0 && sk.getColumnQualifierData().length() == 0 && sk.getColumnVisibilityData().length() == 0 && sk.getTimestamp() == Long.MAX_VALUE && !range.isStartKeyInclusive()) { {code} Could the WholeRowItertor be generalized? Maybe create a more general version of the whole row iterator, and have the whole row iterator extend that for backwards compatibility.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1337
17379	ACCUMULO-1337	I justt ok a look at patch v1. * class javadoc should go before class, not before imports * maybe rename static function decodeRow() to decodeColumnFamily()? * should probably implement OptionDescriber... unless the configuring an iterator in the shell no longer requires this. I noticed the WholeRowIterator does not implement this also. * decodeRow()/decodeColumnFamily() should really have javadoc, since its user facing... this appears to be an issue w/ the WholeRowIterator * We have started adding since javadoc tags... this entire class could have a @since 1.6.0 tag in the javadoc * should the javadoc for filter change? seems its still oriented twoard rows? * might consider a unit test for deepCopy() Curious about the history behind testBug1() was this something you ran into w/ using the iterator w/ multiple tablets? At first the test threw me off because I did not notice the MultiIterator was narrowing the range. I looked at the diffs between the WholeRowIterator and WholeColumnFamilyIterator. The code for both are so short and the diffs significant enough that I am not sure if it makes sense to generalize the code. It seems like it would only make sense to do this if it makes maintenance easier and reduces bugs. If creating a generalized super class is more complex and makes maintenance harder, then its probably not worthwhile. What do you think?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1337
17380	ACCUMULO-1336	+1. Are you thinking of just the typo.encoder package with unit tests?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1336
17381	ACCUMULO-1336	bq. Are you thinking of just the typo.encoder package with unit tests? yeap, thats what I was thinking.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1336
17382	ACCUMULO-1336	Keith, do you have thoughts on which the lexicoders will live?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1336
17383	ACCUMULO-1336	bq. Keith, do you have thoughts on which the lexicoders will live? I was thinking somewhere under o.a.a.core.client, but do not have a specific sub package in mind.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1336
17384	ACCUMULO-1336	Would you have a problem with me taking this on?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1336

#	Requirement	Questions	Answered	Successfully Answered	URL
17385	ACCUMULO-1335	This could actually be a pretty simple change. I pulled the logic in the printDiskUsage() method up to a "getDiskUsage()" that returns the Map<TreeSet<String>,Long>. This isn't the prettiest thing to return in a client API. I supposed maybe a "DiskUsage" object that has accessors for the tables and sizes? What do you guys think?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1335
17386	ACCUMULO-1335	I took a quick look at the patch. It sounds like there may be a second patch, so I did not look at in detail. The javadoc below is incomplete. But if the return type will change, probably does not matter. {code:java} @return a map keyed by a set of table names where the value {code} Did you mean Configuration object? in the comment below? Check out org.apache.accumulo.core.util.CachedConfiguration. Eric created this a long time ago. He found that when a tablet server was loading lots of tablets, loading and parsing the hadoop config for each tablet was actually slowing down Accumulo startup. {code:java} // TODO: Refactor so that a connector isn't created each time. {code} Should the new method in the API throw TableNotFoundExcep?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1335
17387	ACCUMULO-1335	Just to make sure I'm following- A client using the java API should be able to call connector.tableOperations().getDiskUsage(Set<String> tables). The getDiskUsage() method in the TableOperationsImpl will call a getDiskUsage(String table) method in the TabletClientService. The TabletServer's thrift server implementation of TabletClientService.Iface will use the IDs of the tables, call out to the FileSystem.globStatus(new Path("/accumulo/" + tableId), and grab the size for the table(s) it's responsible for unless a failure occurs, at which point the tableOperations().getDiskUsage() method will pick a different tablet server and repeat the process for calculating the size. Is this what you are proposing?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1335
17388	ACCUMULO-1335	[~sonixbp] I just got a chance to look at your 2nd patch, it was nice thanks. I just committed the following changes. * return thrift connections in a finally block * do not retry on TApplicationException (extends TException) ... thrift will throw this if the server side code throws an unexpected Exception like a NPE.. not sure, but I think if a TException is thrown on the server side it will result in a TApplicationException on the client side... I should test that, because thats what I am expecting in the change I made * retry on TTransportException, this indicates a I/O error * moved the thrift calls from TabletServer to ClientServiceHandler .. there is not a great reason for this, it was just easier to reuse some existing code on the client side that randomly picks a tablet server. ClientServiceHandler exist for historical reasons. It used to be that the master or tablet server would service those methods. Then at some point we decided not burden the master w/ those methods. * checked if the user can read the tables in a more standard way. * One of the test you added went directly against the tablet server thrift API. I moved this to go against client API, because that seemed to test more code. Was there a specific reason for going straight to the tablet server?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1335
17389	ACCUMULO-1335	bq. Was there a specific reason for going straight to the tablet server? I could have just written the entire test in TableOperationsIT but it seemed reasonable to test the thrift service separately- just in case something small changes in that layer, it would be easier to isolate any problems that way. bq. retry on TTransportException, this indicates a I/O error I should have looked these up. I'm still getting used to thrift. I figured anything that threw an exception that's not an Accumulo-specific exception would warrant a retry- worst case is the exception occurs on all tablet servers and then we decide we can't fulfill the request. Probably not the most optimal approach.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1335
17390	ACCUMULO-1334	Hey Corey... can you merge your code with server/monitor/util/celltypes/NumberType.java and move it to core/util? That way we only have one mechanism for formatting large values. Plus, you can improve the code with your Math.log() trick.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1334
17391	ACCUMULO-1329	Was the README.mapred fixed? What specifically isn't correct?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1329
17392	ACCUMULO-1327	could we use the maven exec plugin to run the python markdown utility during a phase of our choosing?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1327
17393	ACCUMULO-1327	bq. could we use the maven exec plugin to run the python markdown utility during a phase of our choosing? Adding external dependencies into the build process makes me a sad developer. Even ones as benign as python.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1327
17394	ACCUMULO-1327	Does the python markdown stuff work with jython? if so we might be able to swing it with just maven plugins.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1327
17395	ACCUMULO-1321	Any update? Also, can you give details as to the configuration of the system? There should also be some output from the vfs related code in the tserver debug logs.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1321
17396	ACCUMULO-1321	I modified the unit test to use a jar w/ a class file thats different and it did not work. I have been deep in the VFS code trying to understand why. I think I have it figured out for the unit test. The VFS code creates a wrapper vfs filesystem for a jar/zip, which wraps a java.util.zip.ZipFile. Even after the file is deleted and recreated, vfs still references the old zip file (the java zip code keeps an open file handle to the zip file, so even though its deleted it can still read it). So the new classes in the new jar are not seen. I was trying to figure out how to clear this internal VFS cache, then I started thinking this behavior may be desirable. For classloaders created before the file changed, its nice for them to continue referencing the deleted zip file. This will result in consistent behavior, where the classloader always loads classes from one instance of a zip file and not different zip files over time. I think the solution is to completely recreate the VFS object when creating a new classloader. Older classloaders will reference older VFS objects which will reference deleted zip files. I need to do further work to understand how this solution works w/ something like HDFS where jar files are copied to a local tmp dir and then opened using java.util.zip.ZipFile	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1321
17397	ACCUMULO-1321	what are the remaining issues? - cache dir per VFS - ACCUMULO-599	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1321
17398	ACCUMULO-1321	Did that include rolling back the various work arounds that were committed? -- Cheers ~John	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1321
17399	ACCUMULO-1321	bq. Did that include rolling back the various work arounds that were committed? No, it just builds on the stuff you committed.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1321
17400	ACCUMULO-1321	Should we put in a ticket to the commons-vfs project about this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1321
17401	ACCUMULO-1321	bq. Should we put in a ticket to the commons-vfs project about this? What specific issue were you thinking of?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1321

#	Requirement	Questions	Answered	Successfully Answered	URL
17402	ACCUMULO-1321	rm old jars add new jar rename new jar in the ext directory Got the following stack trace {code} java.lang.ClassNotFoundException: blah.testIterator at org.apache.commons.vfs2.impl.VFSClassLoader.findClass(VFSClassLoader.java:181) at java.lang.ClassLoader.loadClass(ClassLoader.java:321) at java.lang.ClassLoader.loadClass(ClassLoader.java:266) at org.apache.accumulo.start.classloader.vfs.AccumuloVFSClassLoader.loadClass(AccumuloVFSClassLoader.java:102) at org.apache.accumulo.server.client.ClientServiceHandler.checkClass(ClientServiceHandler.java:259) at sun.reflect.GeneratedMethodAccessor184.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:616) at org.apache.accumulo.trace.instrument.thrift.TraceWrap$1.invoke(TraceWrap.java:63) at sun.proxy.$Proxy10.checkClass(Unknown Source) at org.apache.accumulo.core.client.impl.thrift.ClientService$Processor$checkClass.getResult(ClientService.java:2343) at org.apache.accumulo.core.client.impl.thrift.ClientService$Processor$checkClass.getResult(ClientService.java:2328) at org.apache.thrift.ProcessFunction.process(ProcessFunction.java:39) at org.apache.thrift.TBaseProcessor.process(TBaseProcessor.java:39) at org.apache.accumulo.server.util.TServerUtils$TimedProcessor.process(TServerUtils.java:156) at org.apache.thrift.server.AbstractNonblockingServer$FrameBuffer.invoke(AbstractNonblockingServer.java:478) at org.apache.accumulo.server.util.TServerUtils$THsHaServer$Invocation.run(TServerUtils.java:208) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1146) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at org.apache.accumulo.trace.instrument.TraceRunnable.run(TraceRunnable.java:47) at org.apache.accumulo.core.util.LoggingRunnable.run(LoggingRunnable.java:34) at java.lang.Thread.run(Thread.java:679) Caused by: org.apache.commons.vfs2.FileSystemException: Could not get attributes "jar:file:///cloud/accumulo/accumulo-current/lib/ext/blah.jar!/blah". at org.apache.commons.vfs2.provider.DefaultFileContent.getAttributes(DefaultFileContent.java:248) at org.apache.commons.vfs2.provider.DefaultFileContent.getAttribute(DefaultFileContent.java:285) at org.apache.commons.vfs2.impl.Resource.getPackageAttribute(Resource.java:87) at org.apache.commons.vfs2.impl.VFSClassLoader.definePackage(VFSClassLoader.java:244) at org.apache.commons.vfs2.impl.VFSClassLoader.defineClass(VFSClassLoader.java:215) at org.apache.commons.vfs2.impl.VFSClassLoader.findClass(VFSClassLoader.java:177) ... 21 more Caused by: java.lang.NullPointerException at org.apache.commons.vfs2.provider.jar.JarFileSystem.getAttributes(JarFileSystem.java:87) at org.apache.commons.vfs2.provider.jar.JarFileObject.doGetAttributes(JarFileObject.java:112) at org.apache.commons.vfs2.provider.DefaultFileContent.getAttributes(DefaultFileContent.java:241) ... 26 more {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1321
17403	ACCUMULO-1321	bq. I disagree, this bug is still extant. We just did work which made it less probable, and this is evidence in 1.5.0 that it's still here. Are you going to change the affects version to 1.5.0 and the fix version to 1.5.1? This practice in general seems really confusing to me, this bug is refed in the 1.5.0 release notes tagged in svn.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1321
17404	ACCUMULO-1321	bq. Well, this is very similar to that case though. Its confusing for someone trying to understand what was done when. Should we change the description of the ticket to "This ticket covers two similar cases, one of which was fixed in 1.5.0 and the other which was fixed in 1.5.1 ..."? bq. Source code changed, new jar, and instead of copying it in, it's a delete, move, rename. This sounds like a very nice concise description for a new ticket.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1321
17405	ACCUMULO-1314	Few different options for implementation, need feedback. # Switch from Property to Map # Provide expectations for classes (or interfaces) for each property. # How do we expose these class expectations in the CLI and xml based token generators? # Do we make our own DestroyableMap? # Do we enforce property classes in the Map, if custom, or in the Token? Or elsewhere?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1314
17406	ACCUMULO-1289	Some comments on ACCUMULO-1289.patch There are way more than diffs that would have I expected between ProxyServer in 1.4 and 1.5. This is may not be an issue with this patch, but I noticed it while looking at this patch. If this not resolved, I would be reluctant to release the proxy in 1.4. I do not think we should release it for 1.4 and not support it. In its current state, patching bug fixes to 1.4 and 1.5 would be painful and error prone. Can anyone provide insight into this issue? ProxyServer is missing apache header. I suppose this was an issue with the previous patch. I had to reformat SimpleTest in 1.4 and 1.5 to get it to compare nicely. After doing this I noticed the following : * Why not use getRandomProtocol() in 1.4 test? I have actually found proxy bugs in 1.5 by changing the protocol. * One test uses a constant of 10 and the other 100. Why the diff. * While comparing SimpleTest in the patch, I noticed some bugs in 1.5 that I think your patch fixes? I just checked in a fix for those to 1.5. r1469503 The patch changed thrift IDL, but does not include regenerated thrift code. I generated it using thrift 0.6.1 to experiment with the patch.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1289
17407	ACCUMULO-1281	Do we want to do it after every run or every Nth? Or perhaps this makes for the case that the !METADATA table should have a smaller in memory map, but that sounds like a bad idea.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1281
17408	ACCUMULO-1262	Is this worthwhile keeping open?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1262
17409	ACCUMULO-1262	bq. Is this worthwhile keeping open? I don't think so, especially since it doesn't look like anybody is interested in pursuing an internationalized format for displaying numbers.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1262
17410	ACCUMULO-1256	Do we need a special state? Could we just have trashed tables be offline tables with a particular naming convention?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256
17411	ACCUMULO-1256	bq. And perhaps all trash tables should be listed somewhere in the monitor, in a different table. I don't think it should be in the same table as Online and Offline tables. Yeah, maybe another page with a table showing how much space each table in the trash can is consuming and how many entries it has. I think space usage is the main thing a user would want to know about a table in the trash can. Is there anything else a user would like to know? This make me think that a user may want to delete a table in the trash can before it ages off. I suppose if delete table is called with the name of a table that is already in the trash can that it could go ahead and really delete it.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256
17412	ACCUMULO-1256	bq. Do we need a special state? Could we just have trashed tables be offline tables with a particular naming convention? Offline tables can be exported and map reduced over. I think we would not want to support this for tables in the trash can. I am thinking tables in the trash can should be completely invisible to the rest of the API, except for when the user explicitly ask for a list of tables in the trash. So if someone tries to export a table in the trash can, they should get an exception. Exporting a table in the trash would not be useful, the table could be deleted while a distcp is happening.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256

#	Requirement	Questions	Answered	Successfully Answered	URL
17413	ACCUMULO-1256	bq. Do we need a special state? Could we just have trashed tables be offline tables with a particular naming convention? Technically we could, but we already have a state model for the tables, it would be silly to break out of the model for a new feature. It could end up leading to odd state changes that we may inadvertently introduce down the road.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256
17414	ACCUMULO-1256	Presumably there will be a new configuration property indicating how long to wait before deleting the tables. Do you think this should be global or per-table? Which process will handle actually deleting the tables once they expire?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256
17415	ACCUMULO-1256	bq. Presumably there will be a new configuration property indicating how long to wait before deleting the tables. Do you think this should be global or per-table? Which process will handle actually deleting the tables once they expire? I think a global/master config option would be good. Maybe have the master do it. I have become very untrustworthy of time across nodes. For aging off fate operations, I made that happen within a master instance. An instance of the master keeps track of time for each complete FATE op that the client has not finalized. After 8 hours the master deletes it. If the master is restarted, then the stopwatch on each complete fate op starts over. If a master process does not run for 8 hours, it will never delete anything. This avoids the problem of time being off on different machines. It does not avoid the problem of time jumping on the node where the master instance is running. Could do the same thing for tables in trash, a running master keeps track of when each trash table should be deleted. If the timeout is one hour, then an hour after the master starts all tables that were in the trash when the master started would be deleted.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256
17416	ACCUMULO-1256	I can see a case where a user will create a table, delete the table, create again a table of the same name, and delete it. In your design both tables of name "x" will be preserved in the trash because they are given unique names. I like your point about being able to call "du" on the table in the trash but do you have any ideas surrounding any other means of a user being able differentiate between the tables in the client API? Are you thinking that trashed tables may still show up to some capacity in the monitor?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256
17417	ACCUMULO-1256	bq. I like your point about being able to call "du" on the table in the trash I just looked through TableOperations and did not see any other API calls that should work w/ a table in trash. If anyone else can think of any other exceptions, it would be nice to identify these now. bq. do you have any ideas surrounding any other means of a user being able differentiate between the tables in the client API? Do you mean if the user calls listTrash() and they see trash-foo-2013052201-01 and trash-foo-2013052202-01, how would they know which foo to undelete? I was thinking they would use the date. bq. Are you thinking that trashed tables may still show up to some capacity in the monitor? I think users will need an easy way to see what is in trash and how much space is used. This could be in the monitor, shell, or both. In the shell we could possibly add a flag to the du command to analyze space usage by tables in the trash. This would save the user the hassle of getting a list of tables in trash and calling du. What other capabilities users will need to manage this feature? It seems like a way to delete all tables in the trash would be useful, a way to empty the trash.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256
17418	ACCUMULO-1256	Should we spec out the shell changes that are needed and add those to the design doc?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1256
17419	ACCUMULO-1241	{quote}Most non-java files have Apache headers.{quote} Yes, those headers were mostly added after this ticket was created. However, the question about whether the header should include copyright still applies to non-java files. When you say "borrowed", do you mean, a derivative work or dependency? In the case of a derivative work, it seems that the derivative work is our work, and as long as we properly credit the original work, we can follow our own standards, rather than follow theirs. (Maybe our standards should include our header, followed by a brief comment explaining that this work is derived from the original source whose header is as follows, then the original header?) If, on the other hand, we're simply using their work as a dependency, then our packaging is wrong, and we should add their work as a proper dependency (in a separate package). The problem I'm seeing is that when bugs are fixed and features are added, I see everything under "org.apache.accumulo" as our own work, because I don't necessarily know the history of it to know why a license header contains a different copyright date or different formatting standards for the header that exists. I don't necessarily know, for instance, that the header is intending to convey that the original work, from which the current version is derived, had that copyright date, but that the derived work has a newer copyright date. Nor would I know if the original work is copied/pasted in its entirety, with only the package name and the visibility changed from protected to public. I think some of these concerns can be alleviated if we were more strict about packaging dependencies as dependencies, and adding more thorough comments to explain why the header from the original work was preserved (in addition to our own header, as suggested above, not in place of it). As for maintaining the NOTICE file's last copyright date, I think we can add that as a filtered maven resource, so it is automatically updated with the current year, so it doesn't get out of date, and the same goes for the assembly/pom.xml, which can use the same property to keep that up to date.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1241
17420	ACCUMULO-1241	bq. Yes, those headers were mostly added after this ticket was created. However, the question about whether the header should include copyright still applies to non-java files. The headers should not include copyright information unless the files are from third parties, in which case we should not change the headers or add our own. See http://www.apache.org/legal/src-headers.html "Third party" means non-ASF, so for the code we've borrowed from other ASF projects we could standardize the headers. But again, I personally like knowing exactly which files did not originate with Accumulo. bq. When you say "borrowed", do you mean, a derivative work or dependency? I believe the bloomfilter and blockfile/cache packages have minor modifications, and flot is unmodified?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1241
17421	ACCUMULO-1237	bq. incorporating Keith's Github version into the codebase I was not advocating for incorporating the github version into 1.4. My intent was to point out its existence as something that can be used w/ 1.4 today. Is this patch based on the github version? It seems like the API does not match the 1.5 version. Need to make the public methods match w/ 1.5. Also, the patch does not follow the 1.4 directory structure conventions.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1237
17422	ACCUMULO-1237	Keith, Originally, I had the MiniAccumuloCluster in a src/test module. Then I realized that 1.5 had a /test directory structure that included many of the same scripts in 1.4's /test. For that reason, I moved everything to /test. Should I just create src/test?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1237
17423	ACCUMULO-1237	bq. Should I just create src/test? I think either create src/test or put the code in src/server. In 1.4, a lot of the test code is in src/server. Howerver, we would not want to change the package name if its in src/server.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1237
17424	ACCUMULO-1237	Keith- If I move the MiniAccumuloCluster to src/server, are you saying you want me to keep the package org.apache.accumulo.test or change it to org.apache.accumulo.server.test? Seems like it should match the rest of the packages in that module, right?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1237

#	Requirement	Questions	Answered	Successfully Answered	URL
17425	ACCUMULO-1237	bq. If I move the MiniAccumuloCluster to src/server, are you saying you want me to keep the package org.apache.accumulo.test or change it to org.apache.accumulo.server.test? Seems like it should match the rest of the packages in that module, right? Wherever it goes, the package name needs to stay the same as 1.5, so code written against 1.4.4 will work w/ 1.5. I suppose a new module in src/test would be best. If this is done it should make the dependencies in user poms align more nicely. For example if a user had the following in their pom for 1.4, then they should be able to just change the version for 1.5. {code:xml} <dependency> <groupId>org.apache.accumulo</groupId> <artifactId>accumulo-test</artifactId> <version>1.4.4</version> </dependency> {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1237
17426	ACCUMULO-1237	I took a look at patch #2. * Unit test did not run : java.io.FileNotFoundException: /tmpminicluster-86bb2ad0-5ae9-479f-8fd1-f257d1d57553/conf/accumulo-site.xml (No such file or directory) * if you are using eclipse, can you organize imports? * source code formatting seems incorrect * why not just do "loggerProcess = exec(LogService.class);" with no try/catch? * would probably be good to take 'appendProp(fileWriter, "logger.sort.buffer.size", "50M", siteConfig)' from instamo, but use the constant :	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1237
17427	ACCUMULO-1237	[~sonixbp] a few comments on patch #3. I noticed the unit test was using slf4j. I think most of the Accumulo code uses log4j. For consistency, I think it would be better to stick w/ one logging package. I like the message you are logging about the dir its running in, that could be really useful for debugging. Also the formatting is still kinda off. I notice this because I am diffing the 1.4 and 1.5 code. There are a lot of white space diffs that add a lot of noise. When I reformat the patch #3, I get few diffs between 1.4 and 1.5. This is nice for merging bug fixes between branches. How are you formatting the code? I am curious what the difference is. The patch looks good otherwise.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1237
17428	ACCUMULO-1237	bq. I notice this because I am diffing the 1.4 and 1.5 code. I didn't realize that's what you meant when you said it wasn't formatted correctly. I noticed that I was indenting to 4 spaces instead of 2 and I fixed that in this latest patch. bq. How are you formatting the code? I set intellij to format to 2 spaces instead of 4 and then I auto-formatted the files. I'll copy 1.4 and make the small modifications again- there's only a few.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1237
17429	ACCUMULO-1229	Do you know how to get the list of current keepers from a single keeper? I've not yet figured it out.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1229
17430	ACCUMULO-1229	Do you know how to ask zookeeper for the list of replicas? I do not.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1229
17431	ACCUMULO-1229	Option 4 seems like it would result in the following sequence of operations # 1 zookeeper server is passed in # Use 1 zookeeper server to find location of master # Ask master for list of zookeepers # Use this list to create new zookeepers connection This seems a bit odd. Why is making a client list all zookeepers a burden? A client knowing of only one zookeeper server seems like a less reliable way to use Accumulo. Operating this way unnecessarily introduces a SPOF.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1229
17432	ACCUMULO-1229	Would there ever be a need for ZooKeeperInstance to only connect to the zookeepers the user passed (and no others)? I can not think of a use case at the moment. Want to avoid the situation where we cause problems for users when we were just trying to help them. Another option to consider is not changing ZooKeeperInstance at all and providing utility methods for obtaining a complete list of zookeeper servers from Accumulo. A problem with this option is that it does not work with existing user code. If ZookeeperInstance itself is changed then existing code could benefit.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1229
17433	ACCUMULO-1228	For the first proposal, I think it might be useful to add to the API the ability to exclude columns or column families. For the second proposal, however, allowing the client to specify details of locality groups means making locality groups a client-viewable feature, instead of a server-side optimization. Users don't "fetch a locality group", they query for "data identified by (in this case) a set of column families", regardless of how it is stored. Blurring the lines between storage optimization and the data model is confusing and risky, I think it's a bad idea. Further, it's very poorly defined... would you use the locality group as determined by the current configuration, or the locality groups specified in the RFiles on disk prior to the last locality group change? After all, you can get 10 different files with 10 different locality group characteristics... and this works fine if you ask for the data, but it would be... confusing and unpredictable from a user's perspective if you ask for a locality group. This is a bit like asking "give me all data in my files stored with snappy compression vs. gzip".	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1228
17434	ACCUMULO-1228	Why even have the locality group defined that way in the first place then? It seems brittle if we require a client to apply labels to sets of column families (that maybe have to be mutually exclusive across lables?) and then never use them. How does the system handle querying today if a locality group is inconsistent across RFiles? I'd more prefer to have a consistent API for creation and use. If we don't allow a user to query based on locality group name, then we shouldn't have them specify it in the first place. Likewise, if we require them to specify it at creation time, then they should be able to query using it. Having one without the other makes it kind of useless.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1228
17435	ACCUMULO-1228	It would be super useful to have the column family fetch/restrict utility through the client api. This can very positively impact the performance of a query, especially if you can exclude all columns which you know you don't want to look into. bq. Locality groups are an admin feature, for optimizing the data storage for more efficient queries Why do you call them an admin feature? In the simplest form, it's nothing more than a group of columns that we "set aside". Pulling out my BigTable paper, they specifically say "Clients can group multiple column families togther...". As such, a locality group is losely defined as some group of columns that are "accessed together", which in my mind, is something a client would be aware of when issuing a query. I suppose I disagree with you comparing them to a codec. True, the client doesn't care about the codec used to compress the data just as the client doesn't care about the distribution of files backing the table(t), but there are implications for a client to understand the layout of their data and clients can use that knowledge to more effectively retrieve it. As for the argument about a locality group changing and potentially not being propagated to disk yet, the locality group is nothing but a convenience shorthand, allowing a client to fetch multiple columns with a single call. Sure, you don't get the optimization, but the API is still much more simple for the user.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1228
17436	ACCUMULO-1228	bq. I would almost go as far to say that I expect locality groups to be immutable. I would argue that making this change (encouraging scans against locality groups) would in fact make locality groups immutable. Anyone who understands this new client API would be afraid to make changes to locality group config because they can not predict the consequences. Is this what we want? bq. "why did #3 happen?" Its a contrived example to illustrate the point that locality groups can be easily changed. A possible driver for this change would be a new and important use case for which a locality group change would be beneficial. But actually, they probably would not have made the locality group config change because of fear of breaking existing code. So the new use case would not benefit from Accumulo's ability to change locality group config on the fly. .	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1228

#	Requirement	Questions	Answered	Successfully Answered	URL
17437	ACCUMULO-1228	bq. A possible driver for this change would be a new and important use case for which a locality group change would be beneficial. But actually, they probably would not have made the locality group config change because of fear of breaking existing code. So the new use case would not benefit from Accumulo's ability to change locality group config on the fly. Someone is still making the conscious decision here to alter the locality group. They could've made a new locality group instead of destroying the original, no? Actually, I'm a little fuzzy here. Say we have a table with columns [A,B,C,D,E] and LG1=[A,B]. Say [~billie.rinaldi@gmail.com] makes LG2=[C,D]. Now [~kturner] changes LG1 from [A,B] to [A,B,C]. Would my "column search space" for specifying LG1 be [A,B,C,D] because of the overlap of LG1 and LG2? What happens here after a compaction? Would LG1 and LG2 share a common "file" that contains C or would LG1 "contain" all of LG2 and vice-versa because of the overlap? Am I missing the joke here?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1228
17438	ACCUMULO-1228	{quote}I still don't agree that there's an issue for us to be concerned over with client code breaking after someone *knowingly* redefines a locality group.{quote} You have to consider who is doing this modification. This type of modification seems to me to most likely occur when you have different data types occupying the same table, and query patterns change over time, and an admin changes locality configuration to improve performance for those types of queries. The most obvious such transition is from no locality groups (everything in the "default" locality group) to 2 locality groups (1 named, 1 default). Tailoring configuration to optimize for different query loads over time seems like a system administrative function to me, and I don't think that performing this action should break everybody's code that has been written with the assumption that the locality groups are static, nor do I see a substantial reason to encourage such assumptions by exposing that behavior in the public API. Locality groups do *not* represent the data model and should not be queried as such. They represent query patterns, and an optimization for changing load requirements. Now, it is a trivial matter to make the additional assumption that the query load is congruent with the data model for users who know and accept all the risks and implications of doing so... but that case can easily be satisfied by storing some concept of a ColumnFamilyGroup in the application code, and using that as a source for both the maintenance functions (configuring locality groups) *and* the query functions (constructing scanners). Another trivial option to satisfy the same use case is a very simple utility class that reads the current locality groups, and overloads their purpose to also represent a set of column families over which to query. Such a utility comes with risks, though, and it could be confusing to avail such a thing in the public API. I think this sort of thing is best left in an obscure package with lots of javadoc admonishments, or in user/application code. {quote}Scan profiles sound like they could be really useful.{quote} +10, and we're halfway there with IteratorSettings, and iterator profiles in the shell. Another observation: I think the mapreduce API for setting the Set of columns over which to query might actually be more intuitive than our main appending "fetchColumn" and "fetchColumnFamily" methods in Scanner. Allowing a fetch method that gets a set/group of columns/families might would make it even more trivial to have a one-liner that reuses locality group info (or other column family group info) to configure a scanner.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1228
17439	ACCUMULO-1224	Maybe the right answer is to have a Factory definition for creating threads that can be backed by a thread pool. It would implement isOpen and close methods. And any external implementations, even if they wrap a thread pool, would have to decide how to "close", whether that means to do a shutdownNow, or wait for threads to finish, etc. Where would a user specify the BatchConfig? In the Instance when it starts? Maybe it is just a Hadoop Configuration object that is wrapped with methods to add the thread pool factory class name, etc. Btw, this is probably not the right place to ask this, but how come the Connector class is a class and not an interface? Seems like there was probably a reason for this? (Instance too)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1224
17440	ACCUMULO-1224	bq. how come the Connector class is a class and not an interface? It's historical. It's on on its way to being an interface. First, we had to deprecate its constructors.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1224
17441	ACCUMULO-1223	Is this still valid?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1223
17442	ACCUMULO-1223	So, I actually don't even see slf4j-nop used in accumulo-start in master. I do, however, see slf4j-log4j12 as a transitive dependency of hadoop-client and hadoop-minicluster (it's in almost every hadoop artifact, actually... ; it'd be nice if Hadoop marked it as an optional dependency so it didn't propagate automatically.). Does this even matter, though, since it's just test output and it's being redirected to a file and safely ignored? (I guess that's sort of what I was getting at when I asked if it was still valid; I think this issue was created before we redirected console output in tests to a file) We could jump through hoops to exclude it from all our test dependencies, and add a test dependency on slf4j-nop, but I'm not sure what is to be gained from that (especially since we already have a direct dependency on log4j in accumulo-start; see ACCUMULO-3544).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1223
17443	ACCUMULO-1219	The procedure here is to copy another RFile into the missing file's place, correct? Are there any cases where you'd want to try and find a version of the RFile that work? Would you ever want to simply delete references to a file?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1219
17444	ACCUMULO-1219	bq. Would you ever want to simply delete references to a file? If you have some aggregation on that table where the result is important and needs to be accurate, is one I can think of right away.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1219
17445	ACCUMULO-1219	Josh does have it pretty much covered - I'll work up some text that refers to that section if the missing file is in the metadata table. Question: How does one create an empty rfile? I have one that has been passed down from my grandmother's second cousin's aunt and has been in the family for generations, but how would I create a new one? I thought I could just create an new table, but that does not seem to generate a file. Is it possible using the shell (using 1.5.1standard dist) with: create table sample1 insert row1 colf colq value1 delete row1 colf colq flush This gives me something like: \-rw\-r\-\-r\-\- 1 hdfs supergroup 199 2014-04-06 09:51 /accumulo/tables/2/default_tablet/F000000q.rf Or should it be done using AccumuloOutputFormat?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1219
17446	ACCUMULO-1219	[~ctubbsii], I don't see that method anywhere. Are you suggesting that such a method should exist? We're just looking for an easy way to document how an empty RFile could be created. The shell approach should work, but it "technically" isn't empty (because it still contains that delete). That's probably acceptable as a first stab with a caveat to make the delete for something bogus (not that you accidentally delete a record in the table you're trying to recover). I haven't found a high-level command that can create a purely empty RFile yet.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1219
17447	ACCUMULO-1219	In the interests of getting something together quickly that people can use without too many acrobatics, would people be fine with adding a CreateEmpty utility to go along with PrintInfo and SplitLarge?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1219

#	Requirement	Questions	Answered	Successfully Answered	URL
17448	ACCUMULO-1219	bq. In the interests of getting something together quickly that people can use without too many acrobatics, would people be fine with adding a CreateEmpty utility to go along with PrintInfo and SplitLarge? TBH, I wasn't even anticipating making something new for this. I assume an "almost empty" rfile would be sufficient -- the likelihood of the delete actually causing problems would be rather small assuming that people know "roughly" the format of the data that was in their table. I wouldn't be opposed to accepting such a utility if someone writes it up though.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1219
17449	ACCUMULO-1219	[~etcoleman], it looks like we're stalled on an RC for at least a little bit due to the inability to log into the ASF nexus. How can I help make sure this gets into 1.6.0?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1219
17450	ACCUMULO-1216	bq. However, it would be nice to go back and use SequenceFile. Why? What are the benefits?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1216
17451	ACCUMULO-1216	bq. the files are split-able Say I have a 1GB WAL sequence file, could I treat them as two 500MB files and replay them in parallel? I think the answer is no as the order of mutations applied is important, correct?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1216
17452	ACCUMULO-1216	bq. Say I have a 1GB WAL sequence file, could I treat them as two 500MB files and replay them in parallel? I think the answer is no as the order of mutations applied is important, correct? Right the order is correct. Having a splittable file would allow you sort those two 500MB chunks in parallel. Then you could do a merged read of the two sorted chunks.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1216
17453	ACCUMULO-1216	bq. It supports metadata, which we now have to do ourselves. There's some underlying machinery that can help with backward compatibility. Can you give some more details on these points?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1216
17454	ACCUMULO-1212	A simple fix would be to configure redirectTestOutputToFile=true [1] on the surefire-plugin. Personally, I prefer to be able to actually read the test times as they scroll by. The other simple fix would be to just add/fix the log4j configuration in src/test/resources. Either are possible; Do people have an opinion on the subject? I know this was brought up not to long ago, but I don't think we ever actually came to any consensus. [1] http://maven.apache.org/surefire/maven-surefire-plugin/test-mojo.html#redirectTestOutputToFile	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1212
17455	ACCUMULO-1212	bq. I somewhat want to see the normal test output in the build... I just don't want to see logging messages unless I need to (debugging). What do you define as normal? Anything going to System.out? Would you want to see logging above a certain level? I can nuke some of the really chatty stuff right now. There's not much point to even redirect that to files.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1212
17456	ACCUMULO-1203	Changing the interface will break formatters that users implemented for an earlier versions. If the interface had been an abstract class, we would have had more flexibility to deprecate and add. With the interface, it does not seem there is much that can be done. Does anyone object to making this change? I think if we are going to break it then we should make it an abstract class so we will have more flexibility in the future. Or maybe an interface and abstract class combo, encouraging users to extend the abstract class in the javadoc.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1203
17457	ACCUMULO-1203	[~kevin.faro] how do you feel about adding a formatter which formats the Long as the readable date string? At an initial glance, it seems to me that adding the option to `scan` overlaps a little with what the Formatters were intended to do. But, the next line of thought would be the ability to stack Formatters. Although, an easy workaround would to make an extension point from DefaultFormatter which does the long->String change and then users could extend it when they have custom Formatters. Thoughts?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1203
17458	ACCUMULO-1203	bq. how do you feel about adding a formatter which formats the Long as the readable date string? I like that option.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1203
17459	ACCUMULO-1203	I gave it a go. Was this in line with what you guys were thinking?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1203
17460	ACCUMULO-1203	bq. I gave it a go. Was this in line with what you guys were thinking? When Josh mentioned an extension point, I thought of DefaultFormatter having a method like the following. With that method DateStringFormatter could extend DefaultFormatter and only override that method. {code:java} protected String formatTimestamp(long ts){ return ""+ts; } {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1203
17461	ACCUMULO-1200	Hey Corey, I don't see it generating the python/ruby thrift code. Can you look into it? Everything else looks good.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1200
17462	ACCUMULO-1182	Will 1.5 be released without an up-to-date developer manual? On Tue, May 21, 2013 at 3:53 PM, Christopher Tubbs (JIRA)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1182
17463	ACCUMULO-1179	Is this supposed to use HDFSZooInstance to get the list of zookeepers?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1179
17464	ACCUMULO-1177	I was thinking about the design for all of the 1.6 walog changes. Currently each batch of mutations thats written to a walog has a sequence number attached to it. This sequence number relates the mutations to an instance of an in memory map. Currently this sequence number is recorded in start and stop minor compaction events in the walog. I'm thinking we can possibly dispense with these start and stop minc event in the walog and instead store the seq# in the metadata table. This could be done with the mutation that writes out a new minor compaction file to the metadata table to make it atomic. It could be stored in a new column. It would need to increase monotonically for the lifetime of a tablet. I think this will have two benefits. * For ACCUMULO-1083, I think this will greatly simplify minor compactions and recovery. I think it avoids having to group logs and consider each group in order at recovery time. We would not need to write start and stop minc events to all active walog groups (even if no mutations were written to the current walog of the group). * For this issue, I think it leads to the possibility of sorting less data at recovery time. We can analyze the metadata table and determine for each walog what (tablet, seq #) pairs are needed. Then we would only need to sort mutations where the (tablet, seq #) is > whats needed. The drawback of this approach is that the walog will be less self contained.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1177
17465	ACCUMULO-1177	bq. With this approach, the reader could perform in-memory sorts for all the tablets What is "the reader"? I can guess, but I am not exactly sure what you mean. bq. we can remove the define tablet entries in the logs. Now that we're using tableIds, they are short enough to use directly. A tablet id is composed of (tableID,endRow, prevRow). Define tablet in the walog shortens this entire thing. So it seems like it would still be useful if the rows are long?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1177
17466	ACCUMULO-1177	bq. What is "the reader"? The service that presently sorts the logs. That service would be replaced by a service that can read the WAL, and keep only the mutations that come after the id in the !METADATA table. It could sort the mutations in memory, flushing them out to files as needed. At the end of the file, it can add the files to the tablet and remove the WAL reference. Tablet servers should refuse to load a tablet that has WALs. bq. So it seems like it would still be useful if the rows are long? You're right.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1177
17467	ACCUMULO-1177	Marking this with fixVersion 1.7.0, since there's where the work logged against this issue have gone. Is this issue ready to close? I'm not sure, because there's still an outstanding sub-task. Do we want to move that sub-task off to its own top-level issue?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1177
17468	ACCUMULO-1175	Is this something different than ACCUMULO-420?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1175
17469	ACCUMULO-1175	Well crud, I didn't realize that. Does it apply for flush as well?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1175

#	Requirement	Questions	Answered	Successfully Answered	URL
17470	ACCUMULO-1175	bq. Does it apply for flush as well? Maybe. You have the option to flush before compacting. But there is no way to only apply it to flush. Is that what you want? {code:java} public void compact(String tableName, Text start, Text end, List<IteratorSetting> iterators, boolean flush, boolean wait) {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1175
17471	ACCUMULO-1167	Is there any additional work to be done on this outside of what was done for ACCUMULO-1313?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1167
17472	ACCUMULO-1166	Do the same release rules for TLPs apply to contrib projects (aka, this endeavor)? If so, I need to get my butt in gear with a released candidate and all that fun stuff so we can release this alongside 1.5.0. I didn't see anything obvious on http://www.apache.org/dev/release.html	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1166
17473	ACCUMULO-1166	I downloaded the trunk version. Then generated instamo-example. When I tried to run 'mvn install' I ran into a couple of problems. I needed to make log4j a top level dependency and select v1.2.16. Log4j v1.2.15 seems to have some issue in the maven repositories (see http://stackoverflow.com/questions/9047949/missing-artifact-com-sun-jdmkjmxtoolsjar1-2-1). There is also a v1.2.17 available. Once that dependency issue was resolved, I needed to add a dependency for accumulo-start because org/apache/accumulo/start/Main was not found. I'd love to get this working properly. What can I do to help?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1166
17474	ACCUMULO-1166	Ok, both should be working again now. I needed to account for the provide changes. Also, I have no idea why start, fate and trace weren't already included (and how they previously worked without them?). Also ran into some weirdness with ZK wanting to include jms, jmxri, and jmxtools (which you'd have to download yourself), so I just excluded them. Let me know if that works for you?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1166
17475	ACCUMULO-1165	What does this get us over properly integrating Kerberos support into our services? Kerberos is already the standard way that multi-service systems handle mutual authentication (both in general and specifically within Hadoop). It's already had far better analysis done than anything we come up with and operators that need secure set ups are more likely to be familiar with it (or to already have established a support path for dealing with it).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1165
17476	ACCUMULO-1160	What's left in the mapreduce API for this?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1160
17477	ACCUMULO-1159	I'm getting this on the shell {code} 2013-03-12 17:53:15,239 [trace.ZooTraceClient] DEBUG: Scanning trace hosts in zookeeper: /accumulo/750ccec4-8c71-47e5-abb8-c374a4ca8139/tracers 2013-03-12 17:53:15,239 [trace.ZooTraceClient] DEBUG: Trace hosts: [] 2013-03-12 17:53:15,276 [impl.ThriftTransportPool] TRACE: Creating new connection to connection to 172.16.101.54:9997:9997 2013-03-12 17:53:15,296 [impl.ThriftTransportPool] TRACE: Returned connection 172.16.101.54:9997:9997 (120000) ioCount : 90 2013-03-12 17:53:15,308 [shell.Shell] ERROR: java.lang.ClassCastException: org.apache.accumulo.core.security.handler.Authenticator$AuthProperty cannot be cast to java.lang.Comparable 2013-03-12 17:53:15,309 [shell.Shell] DEBUG: java.lang.ClassCastException: org.apache.accumulo.core.security.handler.Authenticator$AuthProperty cannot be cast to java.lang.Comparable java.lang.ClassCastException: org.apache.accumulo.core.security.handler.Authenticator$AuthProperty cannot be cast to java.lang.Comparable at java.util.TreeMap.compare(TreeMap.java:1188) at java.util.TreeMap.put(TreeMap.java:531) at java.util.TreeSet.add(TreeSet.java:255) at org.apache.accumulo.core.security.handler.ZKAuthenticator.getProperties(ZKAuthenticator.java:49) at org.apache.accumulo.core.util.shell.Shell.config(Shell.java:288) at org.apache.accumulo.core.util.shell.Shell.main(Shell.java:417) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:601) at org.apache.accumulo.start.Main$1.run(Main.java:99) at java.lang.Thread.run(Thread.java:722) {code}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1159
17478	ACCUMULO-1145	Would it help for now to use the AccumuloMiniCluster? I've had much success debugging iterators this way.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1145
17479	ACCUMULO-1136	Applied patch in r1452005 and added a note about the VFS classloader bq. In the documentation that mentions lib/ext, its probably worth mentioning the ability to load iterators from a URL. [~kturner] I assume you meant the VFS classloader? I know it's off by default, but was any documentation written for it? Just the basics: what it uses, how does it work with Accumulo, and how to configure it.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1136
17480	ACCUMULO-1134	Can't you already do that in {{monitor_logger.xml}}? Example available at http://svn.apache.org/viewvc/accumulo/trunk/conf/examples/1GB/native-standalone/monitor_logger.xml?revision=1382614&view=markup line 28 and 39.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1134
17481	ACCUMULO-1134	Is there a palace where user can specify configurations like these?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1134
17482	ACCUMULO-1132	What's your status on this ticket? I'll take it up if you don't have the cycles for it.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1132
17483	ACCUMULO-1131	What do you have in mind for examples?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1131
17484	ACCUMULO-1130	is there a reason you changed the Credential object from singular to plural in this patch?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1130
17485	ACCUMULO-1124	I was running Fluo's Webindex example on EC2 for a long period. After running the example I inspected some RFiles. Some of them had larger indexes than I expected. I suspect making a change mentioned in the ticket would reduce the index size. Below is info for an rfile that uses URLs from web pages in keys. I am going to experiment with generating shorter keys in the index for this file. This file was generated using 64K data blocks and 256K index blocks. {noformat} [centos@leader1 ~]$ accumulo rfile-info --histogram /accumulo/tables/7/t-0003uq7/A000rxoi.rf 2016-05-16 16:48:38,914 [rfile.PrintInfo] WARN : Attempting to find file across filesystems. Consider providing URI instead of path Reading file: hdfs://leader1:10000/accumulo/tables/7/t-0003uq7/A000rxoi.rf Locality group : notify Start block : 0 Num blocks : 0 Index level 0 : 0 bytes 1 blocks First key : null Last key : null Num entries : 0 Column families : [ntfy] Locality group : <DEFAULT> Start block : 0 Num blocks : 21,818 Index level 3 : 120,581 bytes 1 blocks Index level 2 : 451,008 bytes 2 blocks Index level 1 : 714,687 bytes 3 blocks Index level 0 : 6,915,137 bytes 25 blocks First key : um:d:385:%03;%01;10.30.170.244>>o>/2954%af; data:current [] 4611686019157309597 false Last key : um:d:395:%03;%01;%ff; com.facebook>.www>s>/dialog/feed?app_id=90376669494... TRUNCATED data:current [] -6917529026891043602 false Num entries : 24,299,468 Column families : [data] Meta block : BCFile.index Raw size : 4 bytes Compressed size : 12 bytes Compression type : gz Meta block : RFile.index Raw size : 120,754 bytes Compressed size : 21,719 bytes Compression type : gz Up to size count %-age 10 : 9292962 22.56% 100 : 14947371 74.88% 1000 : 59017 2.45% 10000 : 112 0.07% 100000 : 6 0.04% 1000000 : 0 0.00% 10000000 : 0 0.00% 100000000 : 0 0.00% 1000000000 : 0 0.00% 10000000000 : 0 0.00% {noformat}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1124

#	Requirement	Questions	Answered	Successfully Answered	URL
17486	ACCUMULO-1124	bq. I experimented with shortening keys in the index and that gave some nice improvements, but not as much as I expected. I realized that even with those changes, bad keys were still being placed in the index. I added code to keep statistics on key sizes and used those statistics to try to select keys that were <=AVG(keySize). I also excluded keys that were too big (greater than 3 std dev from the mean). I had the thought "how would we determine when index size is efficient" in the future (both evaluating the success of this change as well as identifying perf issues in the future). Did you give any thought about how we could expose this information more easily? Maybe we include some extra information in the file entry in metadata so that the master/monitor could easily aggregate/report on file statistics? Not suggesting it needs to happen now, but wondering your thoughts (since I assume you were doing all this investigation by hand).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1124
17487	ACCUMULO-1122	Can someone comment on this JIRA with some info on the current CMS situation? Or, in particular, what SVN path is the desired target for the 'mvn site' output in the CMS?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1122
17488	ACCUMULO-1122	Is this something that we still want? [~ctubbsii], I know you were working on related issues.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1122
17489	ACCUMULO-1102	Should the deleterange command exist in core/src/main/java/org/apache/accumulo/core/util/shell/commands? I am not seeing it. And my shell command does not seem to know about it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1102
17490	ACCUMULO-1098	Billie, do you have an HTML to PDF tool in mind?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1098
17491	ACCUMULO-1096	What is the convention?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1096
17492	ACCUMULO-1085	We ran an unintentional test the other day where we ran tdown on an actively ingesting cluster with lots of memory, and recovery of the tablets was too slow to wait for (would have taken many hours). I think that's a really good reason to try to get this into 1.5. Has anyone else experienced similar problems? bq. When walogs are recovered, they are immediately compacted. This is done to keep from running out of memory, in the case where multiple large in memory maps need to be recovered. Having multiple threads doing log recovery could possibly exhaust memory. The code that does recovery would need to defend against this in some way. Can we flush tablets that already have a lot buffered (i.e. those that are not being recovered)? Writing really small RFiles might be a source of inefficiency that is leading to slow recovery times, although I have no evidence to back that up.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1085
17493	ACCUMULO-1085	Stumbled across this one. Perhaps we could start running assignments in serial and, when we get far enough to know that there is no recovery to perform, we could kick off another assignment. Would help speed up the normal case, while being able to push the memory issues down the road.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1085
17494	ACCUMULO-1083	In the general scope of walog performance, I have a few questions that may or may not have been explored in development and testing. # Hadoop sync/appends have gone through several iterations since they were introduced into the core release. Does anyone have any info on the performance of appends in the various hadoop releases? I'm thinking [~billie.rinaldi] or [~tlipcon] may have some insight. # [~ecn]Can you provide some general insight on how the new waloggers are implemented? I have vague recollection of the incoming Mutations pooling with another thread constantly re-syncing. Does this achieve the same result as the group commits we had added recently? # General community question- do we have any standard benchmarks or performance charts for append, or even general HDFS performance, in regard to system configuration? Specifically filesystem use, system configuration for disks, mounting options, and the like? Thanks	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1083
17495	ACCUMULO-1083	Well, doesn't a factor of 1 still give some more durability since it'll be there if a tablet server dies from running out of memory or something like that?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1083
17496	ACCUMULO-1083	I will run the tests with replication factors of 2 and 1 to fill in the gaps on the table and report back. I wish that I could edit posts. Do I need special permission to do that or do you just have to be a committer?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1083
17497	ACCUMULO-1083	For 1.5, that table seems to indicate that adding n replicas of a log file results in a 1/(n + 1) effect on aggregate throughput. How would we make the logging process more concurrent-- would we add more log files per tablet? Per tablet server? Would adding more files actually do anything, since it seems like we're paying a penalty on a per replica basis, indicating that more raw files may decrease throughput?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1083
17498	ACCUMULO-1083	bq. [~ecn] Concurrent, as in, a sub-set of tablets use a log, or concurrent, as in, there are multiple log files available to log to, as described in the Big Table paper? I think this got lost in the chatter. bq. [~vines] Hadoop sync/appends have gone through several iterations since they were introduced into the core release. Does anyone have any info on the performance of appends in the various hadoop releases? I would be curious about the results from running the same test against 0.23 or 2.0 bq. [~billie.rinaldi] If people think it's important for everyone to be able to edit their own comments, we can ask. I'd have to agree it'd be nice for everyone to be able to edit their own message. Unless there's something to say otherwise. bq. [~brassard] It's also worth noting that the default Accumulo behavior is to use the default of Hadoop, so it's not necessarily set to 3 automatically. Given we used to run with 2 by default (and we didn't have widespread data-loss issues), we could have Accumulo do the setrep 2 by default for the WAL dir. But, I think this merits discussion which is parallel to the one we're already having. bq. [~bills] For 1.5, that table seems to indicate that adding n replicas of a log file results in a 1/(n + 1) effect on aggregate throughput. Sounds about right. But I'm not sure the picture isn't a bit larger than what we're looking at here. [~brassard], can you post relevant HDFS configurations (hdfs block size, datanode handler count, datanode xcievers, etc)? Given the beefiness of your nodes, performance seems low to me. Doing some hand-wavy math.. 15*8=120drives, ~100MB/s write per drive, gives you ~12GB/s write across your cluster. Even knocking that back to say 1.2GB/s write "actual".. that's still far above 5.7MB/s * 8nodes * 3replicas = 136MB/s. With 32 cores, I wouldn't think you're CPU bound. Yet it doesn't seem like you should be IO bound. I assume you're also memory heavy? Do you have stats from Ganglia or iostat that show IO performance? What about network performance? What filesystem are you using for the partitions hosting the data dirs? Without research, I also don't know how the TServer "finishing" writing a mutation aligns with the replicas writing. Does each datanode writing a replica need to complete the write before the TServer will return to the client? Does only one datanode need to finish? Lots of questions :D	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1083
17499	ACCUMULO-1083	bq. Does each datanode writing a replica need to complete the write before the TServer will return to the client? Yes, in the normal case. However multiple sync's and replications can run in parallel on the same file.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1083
17500	ACCUMULO-1083	A regular block write will return as soon as dfs.replication.min nodes have completed the operation (default 1). Are we certain that append behavior isn't the same?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1083
17501	ACCUMULO-1070	Great. I removed 1.5 from the fix version as we are passed the feature freeze date for 1.5.0. Are you also interested in auditing writes and bulk imports?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1070

#	Requirement	Questions	Answered	Successfully Answered	URL
17502	ACCUMULO-1070	I am still looking at the patch, but I have few comments. I will look at it some more tomorrow. * Did you intend to include changes to AccumuloApp as part of the patch? * the patch enables auditing by default, was this intended? * Will be logging binary data,this will be lossy. For example it will log the columns a user fetches. Could consider encoding base64 non-ascii characters, but I am not sure about this. * may be useful to audit credentials.getTokenClassName() in addition to credentials.getPrincipal() Phillip, I will add you as a contributor on jira and the Accumulo web site. Send me an email if you would like your org and timezone set on the web site.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1070
17503	ACCUMULO-1070	Thanks for looking over the patch Keith. We are back at work now, so should be able to make any changes necessary. I will update the patch to: - have the auditing disabled by default (i don't think that it should be on unless you want it) - The changes to AccumuloApp do not need to be updated, but it won't hurt if it does. - I will look into the logging of binary data. - We will add a section to the documentation explaining the audit feature and how to enable/disable. Can you explain what credentials.getTokenClassName() gives us that would be useful for auditing? Cheers Phil	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1070
17504	ACCUMULO-1070	bq. Can you explain what credentials.getTokenClassName() gives us that would be useful for auditing? [~ctubbsii] or [~vines] may be able to give you a more precises explanation. This is all in flux. Basically, it possible for a user to authenticate against a user database in zookeeper or a pluggable authentication system like ldap. I think the token class name may give a clue as to which the user is attempting to use.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1070
17505	ACCUMULO-1070	bq. Hey Keith can you please add me as a contributor on Jira too, thanks Sure. You worked on the patch w/ Philip, but I was not sure if you wanted an entry on the web page.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1070
17506	ACCUMULO-1070	Rob, is there a reason why you pulled all the printf constants to one place? That's not typical for the rest of the codebase, and it's nice to see the template next to point where it is expanded. Otherwise, it looks good.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1070
17507	ACCUMULO-1070	[~ptyoung] and [~robtallis] are you satisfied with the patches, as applied? Can we close the ticket? If not, please let me know what remains outstanding.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1070
17508	ACCUMULO-1052	Commits shouldn't be held just because a tablet is minor compacting, right? The tablet server has to run out of space in the write buffer across all tablets, or the tablets would have to be in a state of migration or going offline. Do you know what's causing the commit hold? Are minor compactions on tablets on this tablet server less efficient (this should be evident in the performance debug logs for that tablet server)? Could this be that the merging minor compactions are kicking in and reducing the compaction efficiency? If the number of files is growing unbounded then that seems to indicated the merging minor compactions are not kicking in. Is that a separate issue, or perhaps something related? How many files are referenced by the tablet in question?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1052
17509	ACCUMULO-1045	Is this handled by the -e option?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1045
17510	ACCUMULO-1045	bq. No. The shell has serious limitations. The intent is to obsolete it with more convenient utilities. Can you elaborate on this somewhat? I am aware of many of the limitations of the shell, but was under the impression that the thrift proxy was supposed to make writing python et. al. scripts a lot easier.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1045
17511	ACCUMULO-1041	At some point, regarding the changes to the security token stuff, the code in ConfiguratorBase.getToken() had the code reverted that I added for ACCUMULO-829, to read from the DistributedCache. Do you know if anything else got reverted aggressively that maybe should not have been?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1041
17512	ACCUMULO-1041	1. Looks good 2. This is primarily a convenience method for a lot of the testing frameworks. It's possible to work around, but would make a lot of the code more ugly. 3/4. Really like this idea, replace it with get property info or something which spits back sets of properties which can be used and descriptions. 5. Sounds good 6. I'm okay with spinning out the User Management stuff, as per other discussions. However, this is a bit of authorization work, but mostly Permission management. And I have an illogical liking of the name SecurityOperations, but I can't identify why. 7. Do you mean use a special token type so user's can't spoof the !SYSTEM principal? I'm okay with this too. 8. yup 9. yup 10. I thought I had migrated them over. Though I may have done it in one of the other iterations and gotten fed up and did things poorly until we were well cemented. 11. what items do you have in mind?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1041
17513	ACCUMULO-1041	First of all, was 1 supposed to be Instance.getConnector(String p, Props t)? Or was it actually supposed to be Connector...? Well the reason for the properties was for better handling in the cli and proxy. In the instance, it should basically be (Authenticator(Class.forName(getAutheticatorClass()).instanceOf()).login(properties). And perhaps we should expand the client side Authenticator interface to also have a get property sets which returns the valid pairings it requires. What sort of issues are you forseeing?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1041
17514	ACCUMULO-1041	bq. Any reconsiderations? If not, I would like to get started on this. Isn't Christopher working on this in the branch he created?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1041
17515	ACCUMULO-1031	Is there a utility to flag the malformed javadoc comments?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1031
17516	ACCUMULO-1030	I was actually just thinking about this today. Do you have any ideas about how to isolate configuration changes, existing tables, etc from different integration tests? It would be nice to spin up the single MAC instance in pre-integration-test, and not have to worry about what other integration tests are doing as I write a new test (I'm thinking of some sort of weird tests against !METADATA). Perhaps this is something that could be addressed down the road if it becomes an impediment. Regardless, I definitely like the general idea here.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1030
17517	ACCUMULO-1028	What issues of token serialization/deserialization would be introduced by this ticket?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1028
17518	ACCUMULO-1028	Agreed. I just don't know how that relates to this ticket, which is about the API, not serialization. I made another ticket that relates more to serialization (ACCUMULO-1027), to support a wider variety of tokens by only requiring them to know how to convert to a binary to be carried in a thrift object, and Eric made a ticket (ACCUMULO-1003) to question the API that determines the method by which tokens are actually serialized. Is there a way I can re-word this ticket, to make it more clear that it's about the API?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1028

#	Requirement	Questions	Answered	Successfully Answered	URL
17519	ACCUMULO-1028	I see. Certainly, (String, byte[]) would be least intrusive to the public API, and (Principal, byte[]) minimally more so. However, why should it take the already serialized token? That could complicate user code, because then they have to serialize the token (or rely on the system that provides the token to do it)... for *our* RPC calls. That doesn't make sense to me. We don't necessary need a serialized token... we just need to know how to serialize it, and the implementing class can define that so we don't have to. You said above that there's no need for them to follow a strict serialization/deserialization interface, but you're imposing one anyway (it's unavoidable), essentially "convertToBytes() and convertFromBytes()". Don't get me wrong... I've no problem with byte[] as the result of that imposition (in fact, I think it's necessary for ACCUMULO-1027)... I just don't like seeing arbitrary byte arrays in the public API... I'd rather do it behind the scenes... because a byte array could be anything, and you lose all compile time checks, and I don't think this is user-friendly for an API. We're developing in a high level language, and we can do better than expressing objects as non-expressive byte arrays in our API. You're right about the generics. It doesn't need to be a generic method. The reason why a generic method is because (I think) it's easier to express to the user what the method is expecting, but it could just as easily have been: {code:java} public Connector getConnector(Principal userPrincipal, AuthToken authToken); {code} They both achieve the same thing. (Unless we wanted to return a Connector<T> to support things like Connectors that automatically renew tokens that expire, or provide other token manipulation methods, like those in SecurityOperations, in which case, it would be better to have the generics. I'm not suggesting we do that, but the current implementation of ACCUMULO-259, without addressing ACCUMULO-1024 would benefit from it if the generic parameter is passed down to the Connector's SecurityOperations, so users can get compile-time checks against using inconsistent token types in the API). (Note also that I'm using "AuthToken" rather than "SecurityToken", because I think it better expresses its role, because "security" could mean anything.)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1028
17520	ACCUMULO-1026	It's important for users of our API to be able to automatically identify areas where their code might not work on future versions of Accumulo. @Deprecated ties into the standard convention for doing that analysis. If we replace @Deprecated with some other annotation then our users will lose the automatic warnings. Is there another way around this problem?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1026
17521	ACCUMULO-1024	Can we use Tomcat for inspiration? Too naive?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1024
17522	ACCUMULO-1024	I like it, [~ctubbsii]. Do you basically see it as a hardcoded ZKAuthenticator/Authorizor/PermHandler and then the Authent/Author/PermHandler have an option for nil? Or something else?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1024
17523	ACCUMULO-1024	What is your status on this ticket? I'm weary to do any 259 documentation until the dust clears. I would work on this, but you've had it marked in progress for a while. If you can give an update, I have some cycles to wrap on this.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1024
17524	ACCUMULO-1021	Do we want to make it 'secure' out of the box, or do we want it easy to use out of the box? And by easy to use, I mean avoiding any questions regarding security that the user may or may not understand/comprehend the repurcussions of. And I quote secure, because we can only go so far with an out of the box implementation so users may have an overextended sense of security because of this ticket.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1021
17525	ACCUMULO-1021	How does the "transparent encryption" feature in HDFS relate to this? IMO, I'd be much happier to rely on HDFS to encryption for us but I know HBase has also run into issues already with different semantics on some HDFS class implementations HBASE-13221.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1021
17526	ACCUMULO-1015	How much of this would we get with a combination of HDFS caching and HDFS short circuit reads?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1015
17527	ACCUMULO-1015	bq. How much of this would we get with a combination of HDFS caching and HDFS short circuit reads? Well, you'll still be incurring the overhead of the IPC to the datanode and reading the data from there, so I imagine it is still "significant" as opposed to completely in memory. I think it would likely depend on the level of implementation of an in-memory table. If we just keep the HDFS blocks in the tserver memory, that would be fairly "simple" and should still perform better, but I don't have any idea of actual numbers. That would be a good first step that someone could evaluate without implementing anything new in-memory 1. Load some data in a table 2. Warm the caches for that table 3. Test performance short-circuit reads 1. Turn off caches 2. Ingest data into a table 3. Compact to disk 4. Test performance	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1015
17528	ACCUMULO-1014	[~ecn], are you thinking about more complex operations like group, order and/or join too? I've seen some simple mockups that implement selection and projection using a basic REPL. Views could also be fairly easy to implement.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1014
17529	ACCUMULO-1014	How about we work to support [Phoenix|http://phoenix.incubator.apache.org/] instead of doing our own thing?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1014
17530	ACCUMULO-1014	So, Pig-0.13 and Hive-0.14 both have support for Accumulo. Trying to make Phoenix work would be another avenue for another query implementation. Do you want to keep this open, [~ecn]?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1014
17531	ACCUMULO-1013	Is there any way to quantify the number of hours this kind of change might take? If the change were broken into phases, what might those phases look like?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1013
17532	ACCUMULO-1013	Thanks. I did see the ticket for the Metrics2 integration, but have not looked yet at the implementation. Do you know if it is just exposing the metrics collected by the JMX MBeans, or a different set of metrics?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1013
17533	ACCUMULO-1013	bq. Do you know if it is just exposing the metrics collected by the JMX MBeans, or a different set of metrics? Yeah, same metrics we have exposed now. No new metrics were added, just the ability to use metrics2.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1013
17534	ACCUMULO-1012	AggregatingIterator is still in the code base, and has been deprecated since 1.4 - should it have been removed for 1.5? What other deprecated pieces should have been dropped in 1.5, but weren't?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1012
17535	ACCUMULO-1011	Each tablet could have a marker indicating if its a candidate for automatic merge. User created tablets would be marked ineligible. When a user created tablet automatically splits, the low child would be marked eligible (this is the system created tablet) the high child would left marked ineligible. User created tablets would only go away when the user initiates a merge. With this strategy automatic splits can be automatically merged. Is this what we want? Cloning would need to copy these markers.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1011
17536	ACCUMULO-1009	Another outstanding issue is that the mechanism by which mapreduce jobs are distributed to tasktrackers is a man-in-the-middle-able vector. I'm not sure how to address that besides making the location the client looks for truststores completely unconfigurable from the job config. At the moment, you can't inject novel trust along the wire, but if you have write access anywhere on the tasktracker node, you could point the tasks to a local truststore of your choosing. What do people think? Would keeping the trusted roots on HDFS be any better? Should you be able to configure where truststores are allowed to be loaded from? (although not through job config, of course) Better ideas?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009

#	Requirement	Questions	Answered	Successfully Answered	URL
17537	ACCUMULO-1009	Took a quick look at the patch. The patch adds a lot of constructors to ZooKeeperInstance. I'm thinking of just adding the following constructor : {code:java} /** * @param config provides all information needed to connect to an Accumulo instance. @see ClientConfiguration */ public ZooKeeperInstance(AccumuloConfiguration config){} {code} This would allow us maximum flexibility in the future w/o clutering the API as new features are added. The zookeepers, instance name, wether or not use ssl, etc. would be obtained from the config. This constructor would be the way to enable ssl. There is some code that pushes the client address and port across threads so its shows up in list scans, do you know if this is still working with the changes you made?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17538	ACCUMULO-1009	I also felt a little weird about the constructor bloat. Just having a single config constructor makes sense, but currently it'll be pretty awkward to use. Maybe I should also add a ClientConfigBuilder with some fluent methods for setting common config switches? I'll take a look at the client address listing code. Can you point me to a relevant entrypoint?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17539	ACCUMULO-1009	bq. Maybe I should also add a ClientConfigBuilder with some fluent methods for setting common config switches? To me that seems like a good way to go. It makes programmatically setting lots of optional configurations easier. The client side config objects constructor could possibly take the required options (or a file that must contain the required options). All the optional stull could be fluent methods are you suggested. bq. I'll take a look at the client address listing code. Can you point me to a relevant entrypoint? Its possible that org.apache.accumulo.server.util.TServerUtils.ClientInfoProcessorFactory may be setting this up. It sets the thread local clientAddress with the client info. The thread pool server has not been used in a while, so I am just curious if this still works.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17540	ACCUMULO-1009	WRT JSSE config, I wouldn't feel great about making the JSSE system properties be the only way to configure the accumulo client's SSL settings. It's JVM global, and accumulo clients may well make connections to multiple services or provide their own SSL server. Especially since I'm imagining the most common deployment will involve private roots, it seems overly restrictive to require the accumulo SSL config to be identical to the SSL config used across the entire client app. However, I do think it makes sense to optionally tell accumulo to use the JSSE config. I'll add that option. I'm working on separating AccumuloConfiguration from ClientConfiguration now, but they will have to cross paths at some point in the codebase, since accumulo services need to be able to make thrift connections to other accumulo services, and they will ultimately get their config from AccumuloConfiguration. Do you think it would make sense to have AccumuloConfiguration.getClientConfig()? Then, if we're in a context that does have an AccumuloConfiguration, we can access common code for creating connections, while contexts like ZooKeeperInstance need not know anything about it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17541	ACCUMULO-1009	I'm not rewriting any custom socket code...all the socket code is deep inside thrift. There's no way the JSSE properties can be respected automatically, since the thrift stack looks pretty different for an SSL server, so we will end up examining and branching based on the properties no matter how we do it. So are you ok with properties like {{javax.net.ssl.keyStore}} being defined in accumulo-site.xml and the Property enum? This is not a client-side setting; the tservers themselves need to know where to find the keyStore. I've also noticed there is already {{monitor.ssl.keyStore}} (and associated properties) used for configuring SSL for the monitor's webpage. This has different requirements from the thrift SSL settings, since there is only one monitor, but it would be good for it to be signed by a root that browsers are likely to recognize. This means that probably you want the monitor to have a cert cut from a "real" CA, but enforcing that requirement for every tserver will be unwieldy and not necessary for most deployments. Is there any particular reason that it's ok for the monitor to have custom SSL properties, but not the RPC connections? Or, if the monitor _should_ use JSSE properties, how should we resolve the different requirements between it and the RPC connections? Or do you think we should only be using JSSE properties for clients, and use our own settings for servers? In which case, which side do the thrift clients created by accumulo services fall on?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17542	ACCUMULO-1009	I've attempted to apply the patch to give it a try. The unit tests are failing with: {{{ Tests in error: testPerTableClasspath(org.apache.accumulo.minicluster.MiniAccumuloClusterTest): Unable to load client config from /tmp/junit8216417786686969359/conf/client.conf test(org.apache.accumulo.minicluster.MiniAccumuloClusterTest): Unable to load client config from /tmp/junit8216417786686969359/conf/client.conf }}} Can you update the patch against master and ensure that the unit tests run? I'm sure I just missed something trying to get the existing patch to work.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17543	ACCUMULO-1009	[~mberman] wrote: {quote}I'm not rewriting any custom socket code{quote} That's good at least, but... {quote}There's no way the JSSE properties can be respected automatically{quote} is simply not true. See TSSLTransportFactory.getServerSocket(int port) or the two param one (int, int)'s javadoc. Whether you use this method or the one that explicitly passes these properties, I don't care... so long as it falls back on the JSSE properties, the configuration looks similar to JSSE property names, and it supports all the features that the JSSE properties would if they were set as system properties. {quote}This is not a client-side setting{quote} The issue is applicable to both client-side configuration and server-side configuration. My previously expressed opinions were primarily for client-side configuration. For server-side, I've been thinking about better scoped configuration anyway (separate config files... not xml, and separate enums, for separate services). That would resolve some of these questions about where to put configuration and would address the monitor's limited ssl configuration as well. In the meantime, I think something like instance.javax.net.ssl.* or general.javax.net.ssl.* might be a decent compromise for server configuration. {quote}the tservers themselves need to know where to find the keyStore.{quote} Recall that part of my criticism of a keyStore-only based configuration was based on the idea that there are other JSSE properties allow you to configure an alternate provider as well (eg. javax.net.ssl.keyStoreProvider). {quote}This means that probably you want the monitor to have a cert cut from a "real" CA, but enforcing that requirement for every tserver will be unwieldy and not necessary for most deployments.{quote} How is this unwieldy? I concede that monitor may have different certs than tservers, but I don't agree that having a "real" CA for every tserver will be unwieldy or unnecessary. Having a CA that is trusted sign the certs used by the tservers is essential for clients communicating with the SSL-enabled instance... otherwise, they'll have to have a trustStore that includes the cert of every tserver... and *that* is unwieldy. Unless you were planning on making clients trust anything that the tserver throws at them... which would undermine the security of the feature. {quote}Is there any particular reason that it's ok for the monitor to have custom SSL properties, but not the RPC connections?{quote} No. See my suggestion above for the scoping of properties in separate properties files, and the compromise for prefixing them with a particular scope. The main point is I want users to understand at a glance that something like "monitor.javax.net.ssl.keyStore" is "javax.net.ssl.keyStore" scoped to "monitor". This gets murkier, if the property names are "monitor.ssl.keyStore", but we're stuck with that for now. The main point I want to make for the RPC is that I want to ensure we're moving forward towards a more intuitive and flexible configuration standards. It is not necessarily a blocker for the feature.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1009

#	Requirement	Questions	Answered	Successfully Answered	URL
17544	ACCUMULO-1009	{quote} >> There's no way the JSSE properties can be respected automatically is simply not true. See TSSLTransportFactory.getServerSocket(int port) or the two param one (int, int)'s javadoc. {quote} My point was, we at least need to figure out if we want to use a TSSLTransportFactory at all. And beyond that, the stack is still quite different since the THsHaServer doesn't support SSL. So I'm still confident "we will end up examining and branching based on the properties no matter how we do it" is true. {quote} I think something like instance.javax.net.ssl.* or general.javax.net.ssl.* might be a decent compromise for server configuration. {quote} This sounds fine to me. {quote} >> the tservers themselves need to know where to find the keyStore. Recall that part of my criticism of a keyStore-only based configuration was based on the idea that there are other JSSE properties allow you to configure an alternate provider as well (eg. javax.net.ssl.keyStoreProvider). {quote} Here I was just using keystore as an example property. Replace "keyStore" with "keyProvider" and my point still stands. But it sounds like this issue is addressed by your other answers anyway, so no big deal. {quote} >> This means that probably you want the monitor to have a cert cut from a "real" CA, but enforcing that requirement for every tserver will be unwieldy and not necessary for most deployments. How is this unwieldy? I concede that monitor may have different certs than tservers, but I don't agree that having a "real" CA for every tserver will be unwieldy or unnecessary. Having a CA that is trusted sign the certs used by the tservers is essential for clients communicating with the SSL-enabled instance... otherwise, they'll have to have a trustStore that includes the cert of every tserver... and that is unwieldy. Unless you were planning on making clients trust anything that the tserver throws at them... which would undermine the security of the feature. {quote} I think you've misunderstood me here. Absolutely the tservers' certs need to be signed by a known CA that can be trusted by clients. I believe there's value in the monitor's https cert being signed by a public root CA that is likely to be known about by web browsers. This is what I meant by a "real" CA, and this is the part that seems unwieldy (and usually unnecessary) for tserver certs. It's easier to distribute custom trusted roots to client applications than to end users' web browsers, and getting certs signed by public trusted roots is expensive, usually requires manual intervention, and is impossible if you want to do hostname verification for machines not on the internet. This is the distinction I was trying to draw between requirements for monitor vs. thrift certs. But, since it sounds like you're fine with a service qualifier before the JSSE-like property key, I believe that addresses my original concern.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17545	ACCUMULO-1009	[~ctubbsii] you have some very specific ideas about client side config that sound good. Do you think you would have time to put something together for 1.6 that [~mberman] could use for this feature? It seems to be a bit of dependency for this feature. I am just trying to work out how we can have minimal churn from a users perspective. Introducing this feature in 1.6 and a proper client side config in 1.7 is sub-optimal for users.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17546	ACCUMULO-1009	I just spent a bit looking at a patch. And I see a new one was posted recently, these may be against the older one. It would be nice to post multiple version of the patch and then I could state that my comments are against version 002 of the patch. * what does certutils do that keytool does not? or does it just make things easier? How do you expect the user to run this? "accumulo org.apache.accumulo.server.security.ssl.CertUtils" ? If this is the case, then this utility is not front and center and the user would really have to look for it. * should generate all print a disclaimer clearly stating any potential pitfalls/risk the user should be aware of? [~ctubbsii] would this address your concern? * Adding ssl options to MAC is done to test Accumulo itself. However this will also be available to Accumulo users using MAC. Is this an API you feel comfortable supporting for a few years? * I am not in favor of adding "public ZooKeeperInstance(String instanceName, String zooKeepers, boolean sslEnabled)". This is not a scalable approach for configuring Accumulo. For example when ACCUMULO-1268 is implemented we would not want constructors for : (..., sslEnabled), (...,sslEnabled,timeout), (...,timeout) . I think the constructors that take ClientConfiguration are sufficient and ClientConfiguration looks nice, like it. I am kinda torn about making client and sever config follow differnt models, but the new client config seems better. I still need to look at it in more detail. * what needs to be done for the proxy w.r.t to ssl? should ClientConfiguration be exposed to proxy api? should proxy server support ssl? maybe we should create sub-tickets Seems like we need test for the following scenarios : * Server requires clients have cert signed by trusted CA (test client w/ valid cert, w/o cert, w/ invalid cert) * Server requires SSL, client tries connecting w/o SSL * Server does not have a cert that client trust * Multiple ZooKeeperInstances in the same java process with different SSL config (is this possible? It looks like you included ssl options in the Key in ThriftTransportPool?) This could happen if a single process connected to multiple Accumulo instances	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17547	ACCUMULO-1009	[~vines]: We can make MAC easy to configure SSL for testing, regardless of whether the provisioning code is there or not. But what specific security configuration are you testing by providing a security switch that automatically provisions and configures MAC? The described situation is no different than configuring the jetty-maven-plugin with certs to test a webapp in Jetty with ssl... but you don't see Jetty with a "do security" configuration option that automatically chooses a security configuration to test against... you still have to provision certs (generating certs, using something like keytool-maven-plugin, if testing in maven, or by providing a specific cert), in order to test the app with that *specific* security configuration. Testing "security on" is next to meaningless. Testing a *specific* security configuration is far more valuable to users. So, I'm not convinced that users will get value from MAC having provisioning code. However, I *do* think there's value in documenting the provisioning and configuration of certs (with openssl or keytool) in an example. And... what would be *really* valuable is having this done automatically in an integration test baked-in to the next version of the Instamo maven archetype that supports SSL.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1009
17548	ACCUMULO-1009	Trying to catch up on the discussion (glad to see such in-depth discussion :D) As [~kturner] and [~_alexm] touched on, I think it may be good to take a step back and provision out the sub-tasks here. As a general statement, my gut agrees that we don't want to be in the "security provisioning" realm just for the monumental difficulties in doing it correctly. That being said, I think it would be prudent to have some sort of "basic" mechanism in which we can test things. The simplest approach to me would be to generate cert(s), keystore, local-CA, and w/e else we need to run "securely" for MAC, document how it was done, and then bundle that as a first go-around. A sub-task can be made to find a happy medium with what we could do automatically and what is best left up to the integrator/sys-admin? Testing security for the sake of "is it secure?" is likely without much gain, but there are definitely the edge-cases like [~kturner] pointed out which need testing. All that being said, trying to break down the larger wire encryption issue into some more tenable pieces is a good idea (plus so the next patch doesn't break the 4k line count :D)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1009

#	Requirement	Questions	Answered	Successfully Answered	URL
17549	ACCUMULO-1009	bq. So, you run accumulo init-ssl Where is this in the patch? I search the patch for "init-ssl" and find nothing. bq. Then, if you run init-ssl on another node, it copies the root to the local system, cuts a fresh private key off of it, and sticks both in the default locations If you have time I would really appreciate if you could post example commands for initializing a 3 node cluster using this process. bq. if you have connected successfully with some cert, and then in the same process you try to connect with a different cert, you could get a cached, connected transport I think it would be worthwhile addressing or at the very least opening a subticket to track the issue so its not forgotten. I'll open a subticket.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17550	ACCUMULO-1009	I'm assuming by "the recent major points" you're specifically talking about your own recent major points. Just clarifying because it sounds like you're asserting points that I don't believe are universally agreed upon. Some specific responses to new points that I haven't seen before upthread... bq. The proposed implementation, putting CA private certs in HDFS, is a very bad idea (it undermines the "authority" part of the term CA) Really? All it's saying is that anyone with the password to decrypt the key _is_ the authority. I don't believe there's anything fundamental about the concept of "authority" that requires they only work from one physical machine. As it is, if you have the instance secret, you can do whatever you want with the accumulo cluster. The root of trust has to start somewhere, and for the rest of accumulo, at the moment, that root of trust is the filesystem security on the site config. I don't see why it would be a problem for the same to be true for the SSL trust. bq. init-ssl would be useful, but if added, it can easily wrap keytool or openssl, rather than custom provisioning code Why is wrapping keytool or openssl better than wrapping bouncycastle? The interfaces are pretty much the same, except that BC can be pulled in as a maven dependency and doesn't require a brittle connection to tools separately installed on the system in varied versions and locations.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17551	ACCUMULO-1009	Ok, in an effort to unblock this changeset from the provisioning issue (which I'll file as a separate issue), I've attached a new patch where CertUtils is pulled out of server and minicluster has no capability to generate its own certs. I know [~ctubbsii] wants to switch to using keytool-maven-plugin for the provisioning in our integration tests, but it doesn't actually support cert generation without JDK 1.7, and my understanding is that we want to stay 1.6 compatible for the time being. (It's also not clear that it supports generating self-signed certs without external tools even then.) So, for now, I just moved CertUtils (and the corresponding bouncycastle deps) into a test-only scope. I've also included an \@Ignored test of our support for SSL config through JSSE system properties, however, at the moment it appears that setting JSSE system properties breaks ZK connections. This seems likely to be related to ZOOKEEPER-1554, but I'll do some more investigation and comment over there. So, I know [~ctubbsii] was pretty insistent that we support JSSE config, but should we include it even though we know it won't work, in anticipation of the ZK issues getting resolved in the future, or should we pull the support completely until there's a chance of it actually working?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17552	ACCUMULO-1009	bq. I haven't looked at every line of code yet. .... do you plan to? o.o	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1009
17553	ACCUMULO-1009	Sure, eventually. Don't let me be the hold-up, though. Why don't we drop it on ReviewBoard, so we can get multiple eyes on? [~busbey] seems to have been pushing to use that a bit more.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1009
17554	ACCUMULO-1009	Is this ready to go now? I have a patch, which I will attach, with attribution to [~mberman] which applies to 1.6.0-SNAPSHOT.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1009
17555	ACCUMULO-1008	I fixed up the tarball to lift the gen-* stuff out of proxy/target and put it in proxy/thrift. I also realized that the reactor was building proxy after assemble, so if you had a clean workspace, you wouldn't get anything but the java files for the proxy in your tarball. I did not address the blanket inclusion of **/*.py which causes some extra cruft to show up (e.g. core/target, trace/target, proxy/target etc) I'm not sure about the deb or the rpm though as I'm not on a debian/redhat based distro. Could you verify that both are actually packing the proxy stuff as intended, [~vines]? (I think you worked on that stuff most recently? If not, then whomever)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1008
17556	ACCUMULO-1008	So the deb still builds correctly for debian, but I'm not familiar with the proxy. How do I start it and what's the best way to test it?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1008
17557	ACCUMULO-1006	{{VisibilityParseException}} gets thrown in the code in three places in one method, none of which I'm sure are actually possible. The first two times are cases where we end up with an AND/OR node that has fewer than two children, but this is impossible to construct without running into a BadArgumentException in {{ColumnVisibilityParser.parse_(byte[])}}. The last case happens when the node type is not one of the types defined in the enum, which will only happen if somebody changes the code to make it possible. One _other_ way that this might happen is if somebody decides to extend {{ColumnVisibility}} because it is not {{final}}, but that has wire implications that I'm not prepared to think about yet. If Christopher is correct that VPE never gets thrown (and I think he is), then what is the intent of this JIRA?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1006
17558	ACCUMULO-1003	I understand that it's to provide alternative serialization methods, but we already have the extra level of indirection in the SecurityToken. If I want faster/better serialization, can't I just use a new class? And since hadoop already has a mechanism for substituting different classes for Writables, we can leverage that infrastructure for future improvements. For example, we've replaced the implementation of the objects written to the WAL between releases by updating the classes and providing the Hadoop library with the translation. This just looks overly complex. That's fine if it has a point.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1003
17559	ACCUMULO-1000	Could this be a family of iterators that implement the compare and set semantics, or is this meant to be resolved at insert time (AKA read-modify-write)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1000
17560	ACCUMULO-1000	From the mailing list: {quote} In the design document I suggested a method called putCondition() and if the value was null, thats mean that the user expected it to be absent. I don't really like that API for specifying expected absence. {quote} Could we do this as an additional method instead? relying on null for special API behavior smells like it'll lead to unexpected behavior if users mess up. {code} /** * Specify that a given cell must not exist */ void putConditionAbsent(CharSequence cf, CharSequence cq, ColumnVisibility cv) {} {code} And then putCondition can assert all non-null arguments.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1000
17561	ACCUMULO-1000	Would we want a conditional compare and delete analogues as well?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-1000
17562	ACCUMULO-1000	bq. Would we want a conditional compare and delete analogues as well? ContionalMutation will extend mutation, so the following would be supported. {code:java} ConditionMutation cm; cm1.putCondition(X); cm1.putDelete(Y); {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1000
17563	ACCUMULO-1000	Also, we should avoid leaking information about data being present at a level the user can't see. If visibility is really a part of the key, then as far as a user with only access FOOBAR is concerned, a cell with vis CATSDEE doesn't exist and the absence test should reflect that. Unless there's a different API already that lets a user say "Given these Auths, can I read everything in column X.Y.Z?"	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1000

#	Requirement	Questions	Answered	Successfully Answered	URL
17564	ACCUMULO-1000	It seems the implementation Keith gave is for compare and swap, not compare and set. This means that my pre-commit check is an equality operation, which may not be enough. One example that's always used is a bank account- when another person makes a charge to a bank account, they don't necessarily care that the account balance is the same as some initial check (ie, you said you have $100 but the bank says you have $90), I only care that the amount covered by the charged exists in the bank account (ie, this gumbo cost $50 and you have $90, you may pass go). I think we could hammer in this functionality by using an iterator that checks the precondition and returns a boolean- have you thought about how to add in that to {{ConditionalMutation}}?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1000
17565	ACCUMULO-1000	bq. One last thing, I am thinking of renaming putCondition() and putConditionAbsent() to putEqualityTest() and putAbsenceTest(). I initially found Test to be more confusing than Condition. What about ifExists(cf, cq, cv, val) and ifAbsent(cf, cq, cv)? That would provide a good semantic flow: cm.ifExists, cm.put. An alternative to ifExists could be ifMatches. I confess I don't have a full grasp of the use cases for this; would we ever want to test ifExists(cf, cq, cv)? Regarding the absence test with non-visible keys, maybe we should just let the test pass. Maybe the conditional writer should only allow writing visibilities that can be seen by the set of authorizations used for reading. You could use the same conditions with different authorizations to insert otherwise identical keys with different visibilities.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1000
17566	ACCUMULO-1000	bq. I think we could hammer in this functionality by using an iterator that checks the precondition and returns a boolean- have you thought about how to add in that to ConditionalMutation? My only thoughts about iterators so far is that the check will be done on the server side by reading the data through the iterator stack configured for the table. I suppose the ConditionalWriter could be configured to use iterators just like the scanner can. This made me think that ConditionalWriter could extend ScannerBase, took a look at the API most of the methods make sense except for iterator().	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1000
17567	ACCUMULO-1000	[~billie.rinaldi] and [~busbey] you both made comments about auths and vis, did the example w/ auth "A" and col vis "B" address your concerns? Want to make sure we are not talking past each other. bq. #1 seems most supportable, since it makes clear that you're only checking what you're allowed to see. As a bonus a change from required -> optional would be non-breaking, but not vice-vers good point. I am going to go w/ that option for now in my prototype. If anyone wants to see my client side prototype let me know. Once the discussion settles for a bit I will post another revision of the design doc.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1000
17568	ACCUMULO-1000	Can you put it up on the gitchubs or bitbucket?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-1000
17569	ACCUMULO-1000	bq. Ah, yeah that makes sense now. Presumably an IllegalArgumentException? When a batch of conditional mutations is passed to the conditional writer, can not throw an exception. Need to know the status of each one. I am thinking of adding something to the status enum, just not sure what to call the enum. bq. Can you put it up on the gitchubs or bitbucket? https://github.com/keith-turner/ACCUMULO-1000 [~busbey], I like the Cell proposal. I also agree w/ your arguments for not using it. bq. This might just be my bias against Object creation in Java. Thats a good bias to have (if only there were some safe way to allocate objects on the stack). Normally I would agree w/ this, but I am thinking in this case the operation (lock row on server side, read, maybe update) is so expensive that it possibly dwarfs this concern. However in the case where data is cached in the tserver its not as expensive. bq. the unwrapped version has the advantage of looking more like other methods used in Mutation. Agreed, consistency is very important. For me this is the main argument for not using the Cell+Test design. Changing the API for mutation is not really an option. [~bills] I just fully understood the significance of your earlier comment w.r.t API. Users will be able to do any check they like using iterators. Therefore we do not need to anticipate future mutation conditions besides equality and presence.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1000
17570	ACCUMULO-999	are we suggesting this would be a contrib?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-999
17571	ACCUMULO-999	[~kturner], would the goal of having Accismus as a contrib be to foster a community of its own (so that it can eventually stand as its own project) or is the idea to intertwine the development of Accumulo and Accismus such that at some future point they would merge? If the former, why not foster that community within the Incubator?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-999
17572	ACCUMULO-998	Just to clarify the scope of this feature, is it encryption at the database level, user level, table level, locality group, etc? It appears to me that creating an encryption codec and plugging it into the HDFS layer is what is being suggested for 980 and 981, which would be database level I think. Why is the implementation for 958 not the same?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-998
17573	ACCUMULO-998	Mike, thanks for the PDF-- very concise. In the known default properties table, between {crypto.secure.rng} and {crypto.default.key.stragegy.hdfs.uri}, the text is mangled for me in Chrome, Preview and Goodle Docs. Can you hook us up with a new version that has that info?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-998
17574	ACCUMULO-998	I have updated the encryption patch with a few improvements: 1. It works against trunk cleanly again. It had gotten stale. 2. We discovered a bug in our own internal testing where certain block caching situations would result in an encrypted stream not being correctly created. That bug has been fixed in this patch. How can we move forward on getting this patch applied?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-998
17575	ACCUMULO-998	Should this be fixVersion 1.6.0 and closed? What work on it remains?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-998
17576	ACCUMULO-997	Should this have also removed "logger" as an option?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-997
17577	ACCUMULO-996	What do people think about doing the following in the proxy instead of always passing user pass? {code} struct ProxyToken { long token; } ProxyToken authenticateUserPass(UserPass up); ProxyToken authenticateKerberose(KerberoseToken kt); i32 addConstraint (1:ProxyToken token, 2:string tableName, 3:string constraintClassName) {code} This would allow the proxy to use different authetication mechanisms. The token woudl be transient, generated using secure random, and time out when inactive.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-996
17578	ACCUMULO-996	bq. Another option is to leave the client api pre-259, and just make the server side have the pluggability, so we don't end up changing the authentication token types in quick succession. I spent the majority of my time on the server end, and unfortunately it's showing in the code. I still have not had a change to make a comprehensive review of this yet, its on my todo list. But what types of configurations were you hoping to support? Below are some possibilities. Do you feel the current API would support these? * Kerberose + LDAP * Username/pass + LDAP * Username/pass + Zookeeper * PKI + LDAP I just opened ACCUMULO-1009. Are there standard ways for doing PKI authentication other than SSL? If not then that may push PKI off til 1.6, but we would still want to think through the API implications as much as possible now. For example in the end we may want a connection factory that can produce authenticated (and possibly encrypted) connections to ACCUMULO servers. Users could supply this connection factory to Accumulo configured however they like.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-996
17579	ACCUMULO-993	What else is needed for this ticket?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-993

#	Requirement	Questions	Answered	Successfully Answered	URL
17580	ACCUMULO-980	Some comments on proposal V1. It seems like the IV would be transparent to RFile, it would just be encryption header information associated with a block. Just like each gzip block probably has some header. From RFiles perspective it just needs to be able to read and write blocks of data. When the encryption codec is not used, there is no per block IV. Does this sound correct? Taking this a step further, should encryption be pushed into BCFile? Currently RFile has no concept of compression, is just reads and write blocks of data to BCFile. BCFile handles compression and stores compression metadata like what codec to use for reading. Even RFiles own root meta block is stored as a regular BCFile meta block and compressed like everything else. Seems like modifying BCfile rather than RFile may be easier. I have already modified BCfile to support multi level indexes in 1.4. BCFile was copied because it was package private, but was not modified for a long time. Why is another interface needed? Why not use org.apache.hadoop.io.compress.CompressionCodec? Not saying we should or should not do this, but would like to hear your thoughts since you have looked into this. I see some things in the design doc that I suspect influence this decision, like needed to set Key and IV. While thinking about this I remembered the BigTable paper mentioned using two compression codecs in series. In the past we have not supported rolling upgrade from 1.x to 1.(x+1). Would only need to consider this if 1.6 supported it. Changes in the file format would be a small part of a larger effort to support rolling upgrade. Releases to date could always read a file produced by any previous version. So Accumulo 1.4 can read rfiles produced by any previous version of Accumulo. Is there any concern with storing unencrypted blocks in memory? The code currently caches uncompressed blocks (but still serialzed with RFile encoding) in memory. Would this be a concern in case these cached block are swapped out? Would we want to keep blocks encrypted in the cache and decrypt only as needed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-980
17581	ACCUMULO-980	bq. and another that is then used by the WAL file streaming, Is there some reason the WAL encryption could not use a compression codec?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-980
17582	ACCUMULO-980	Are you considering rekeying? I.e. encrypt new files key with new key and read old files key with an older key. Do you have any thoughts on key management? I was thinking that each tablet could have a key. The tablets key is encrypted with the table key. The tablet key is used to encrypt/decrypt file keys. A tablet never has to know the table key, it could ask a centralized service to decrypt its tablet key using the table key. Using this method, the table key would only need to be held in memory on one machine.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-980
17583	ACCUMULO-980	The design I'm seeing here is that each RFile will have its own encryption key embedded with in its own crypto header block. That key (call it an RFile encryption key) will be in turn encrypted to a different key handled by an outside entity (this is where that SecretKeyEncryptionStrategy class comes into play). Should that external key become compromised, then *presumably* the only thing that needs to be rekeyed is the RFile encryption key; however, I don't think HDFS works that way, because one can only write a file once, correct? If you really wanted to, you could "rekey" the RFile encryption key by creating a new copy of the RFile with its encryption key decrypted by the old key and re-encrypted to the new key. The rest of the RFile would not have to be decrypted and re-encrypted but the file would presumably have to be copied. If this were a grave concern, then the way to solve it would be to make the crypto block an external file to the RFile rather than something embedded inside it. The only problem with that structure is the possibility of the RFile becoming separated from its companion crypto block, and general untidiness putting lots of little files everywhere causes. Since these files generally aren't dealt with directly by users or administrators, maybe this is the direction to go?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-980
17584	ACCUMULO-980	bq. Actually, it will always do this, as each file has a separate key. I realize this, because each RFile has its own internal key encrypted with an external key. I was wondering about this external key, if you could have different ones for reading and writing. Need to establish some firm terminology for the different classes of keys, so discussion is clear. Maybe add something to the proposal naming each key type and defining it. Then that name could be referenced in discussion. bq. Overall, then, Keith, does the proposal make sense? Seems workable? What you are proposing for RFile sounds good. It would be nice to outline how you are thinking of integrating this rfile capability into Accumulo. I suppose these are some of the questions that I have been asking. When a major compaction happens, how does it get its keys? What does that interaction look like?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-980
17585	ACCUMULO-980	bq. Is there some reason the WAL encryption could not use a compression codec? Keith, could you expand on your suggestion that a compression codec could be used for WAL streaming encryption? The encryption would be much more reusable if we could do it all in a codec. What do people think about the encryption codec patch attached to HADOOP-9333? Could that be adapted for our purposes?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-980
17586	ACCUMULO-973	Integrated in Accumulo-Trunk-Hadoop-2.0 #3 (See [https://builds.apache.org/job/Accumulo-Trunk-Hadoop-2.0/3/]) ACCUMULO-973 Remove trivial warnings about imports, and added missing hashCode() implementations for objects that override equals(), to be nice to users. ACCUMULO-971 Fix some malformed javadoc return tags, and errors linking to hidden (protected) methods (I don't know why javadoc cares about these). (Revision 1434928) Result = FAILURE ctubbsii : Files : * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/BatchWriterConfig.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/impl/OfflineScanner.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/impl/TabletServerBatchReaderIterator.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/data/ColumnUpdate.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/iterators/user/IntersectingIterator.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/iterators/user/KeyTransformingIterator.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/util/shell/commands/InterpreterCommand.java * /accumulo/trunk/examples/simple/src/main/java/org/apache/accumulo/examples/simple/client/ReadWriteExample.java * /accumulo/trunk/fate/src/main/java/org/apache/accumulo/fate/TStore.java * /accumulo/trunk/server/src/main/java/org/apache/accumulo/server/logger/LogReader.java * /accumulo/trunk/server/src/main/java/org/apache/accumulo/server/security/SecurityUtil.java * /accumulo/trunk/start/src/main/java/org/apache/accumulo/start/classloader/AccumuloClassLoader.java * /accumulo/trunk/start/src/main/java/org/apache/accumulo/start/classloader/vfs/ContextManager.java * /accumulo/trunk/test/src/main/java/org/apache/accumulo/test/MiniAccumuloCluster.java * /accumulo/trunk/test/src/main/java/org/apache/accumulo/test/MiniAccumuloConfig.java	0	0	https://issues.apache.org/jira/browse/ACCUMULO-973

#	Requirement	Questions	Answered	Successfully Answered	URL
17587	ACCUMULO-973	Integrated in Accumulo-Trunk #645 (See [https://builds.apache.org/job/Accumulo-Trunk/645/]) ACCUMULO-973 Remove trivial warnings about imports, and added missing hashCode() implementations for objects that override equals(), to be nice to users. ACCUMULO-971 Fix some malformed javadoc return tags, and errors linking to hidden (protected) methods (I don't know why javadoc cares about these). (Revision 1434928) Result = SUCCESS ctubbsii : Files : * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/BatchWriterConfig.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/impl/OfflineScanner.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/impl/TabletServerBatchReaderIterator.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/data/ColumnUpdate.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/iterators/user/IntersectingIterator.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/iterators/user/KeyTransformingIterator.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/util/shell/commands/InterpreterCommand.java * /accumulo/trunk/examples/simple/src/main/java/org/apache/accumulo/examples/simple/client/ReadWriteExample.java * /accumulo/trunk/fate/src/main/java/org/apache/accumulo/fate/TStore.java * /accumulo/trunk/server/src/main/java/org/apache/accumulo/server/logger/LogReader.java * /accumulo/trunk/server/src/main/java/org/apache/accumulo/server/security/SecurityUtil.java * /accumulo/trunk/start/src/main/java/org/apache/accumulo/start/classloader/AccumuloClassLoader.java * /accumulo/trunk/start/src/main/java/org/apache/accumulo/start/classloader/vfs/ContextManager.java * /accumulo/trunk/test/src/main/java/org/apache/accumulo/test/MiniAccumuloCluster.java * /accumulo/trunk/test/src/main/java/org/apache/accumulo/test/MiniAccumuloConfig.java	0	0	https://issues.apache.org/jira/browse/ACCUMULO-973
17588	ACCUMULO-973	My IDE is showing no warnings. Can we wrap this ticket?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-973
17589	ACCUMULO-972	./docs/src/developer_manual/accumulo_developer_manual.tex:Read RPCs used in (8) are split into TODO ... ./docs/src/developer_manual/accumulo_developer_manual.tex:% TODO discuss how load and unload execute ./docs/src/developer_manual/accumulo_developer_manual.tex:% TODO: discuss minor compactions ./docs/src/developer_manual/accumulo_developer_manual.tex:% TODO: discuss major compactions ./docs/src/developer_manual/accumulo_developer_manual.tex:% TODO: discuss splits ./docs/src/developer_manual/accumulo_developer_manual.tex:% TODO: discuss garbage collection ./test/system/upgrade_test.sh:#TODO could support multinode configs, this script assumes single node config ./test/src/main/java/org/apache/accumulo/test/randomwalk/security/SecurityHelper.java: // TODO Auto-generated catch block ./test/src/main/java/org/apache/accumulo/test/randomwalk/security/WalkingSecurity.java: // TODO Auto-generated catch block ./test/src/main/java/org/apache/accumulo/test/randomwalk/concurrent/Compact.java: // TODO need to sometimes do null start and end ranges ./test/src/main/java/org/apache/accumulo/test/randomwalk/concurrent/Merge.java: // TODO need to sometimes do null start and end ranges ./test/src/main/java/org/apache/accumulo/test/randomwalk/shard/Search.java: // TODO breaking w/o reading all data causes batch reader to spew exceptions ./core/src/main/java/org/apache/accumulo/core/file/blockfile/cache/BlockCache.java: * Block cache interface. TODO: Add filename or hash of filename to block cache key. ./core/src/main/java/org/apache/accumulo/core/client/mapred/InputFormatBase.java: // TODO check that extents match prev extent ./core/src/main/java/org/apache/accumulo/core/client/mapreduce/InputFormatBase.java: // TODO check that extents match prev extent ./core/src/main/java/org/apache/accumulo/core/client/impl/OfflineScanner.java: // TODO share code w/ tablet ./core/src/main/java/org/apache/accumulo/core/client/impl/OfflineScanner.java: // TODO need to close files ./core/src/main/java/org/apache/accumulo/core/iterators/user/IntersectingIterator.java: // TODO: Should we change this to an ArrayList so that we can just add() ? ./core/src/main/java/org/apache/accumulo/core/iterators/user/IntersectingIterator.java: // TODO: Do I need to call new here? or can I just re-use the term? ./start/src/main/java/org/apache/accumulo/start/classloader/vfs/AccumuloVFSClassLoader.java: // TODO monitor time for lib/ext was 1 sec... should this be configurable? ./server/src/main/java/org/apache/accumulo/server/tabletserver/Tablet.java: // TODO this could hang causing other tablets to fail to load ./server/src/main/java/org/apache/accumulo/server/tabletserver/Tablet.java: // TODO check lastFlushID and lostCompactID ./server/src/main/java/org/apache/accumulo/server/tabletserver/Tablet.java: // TODO check seems uneeded now ./server/src/main/java/org/apache/accumulo/server/master/tserverOps/ShutdownTServer.java: // TODO move this to isReady() and drop while loop? ./server/src/main/java/org/apache/accumulo/server/master/tableOps/TableRangeOp.java: // TODO: add the entries only if there are no !METADATA table references ./server/src/main/java/org/apache/accumulo/server/master/tableOps/CompactRange.java: // TODO since not using tablet iterator, are there any issues w/ splits merges? ./server/src/main/java/org/apache/accumulo/server/master/tableOps/ExportTable.java: // TODO throw another type of exception? ./server/src/main/java/org/apache/accumulo/server/master/Master.java: // @TODO: maybe move this to Property? We do this in TabletServer, Master, TableLoadBalancer, etc. ./server/src/main/java/org/apache/accumulo/server/master/Master.java: // TODO this used to be TabletIterator... any problems with	1	0	https://issues.apache.org/jira/browse/ACCUMULO-972
17590	ACCUMULO-963	Why is this Won't Fix? I think BucketCache is still worth a look: http://www.n10k.com/blog/blockcache-showdown/	1	0	https://issues.apache.org/jira/browse/ACCUMULO-963
17591	ACCUMULO-962	It appears the the scanner that is pulling !METADATA table information is pulling in data beyond the keyextent requested. I tweaked the range computation, and added some additional debug: {noformat} 2013-01-15 15:59:12,338 [impl.ThriftScanner] DEBUG: Starting scan tserver=127.0.0.1:9997 tablet=!0;~;!0< range=[2;64af4f0d844d0e1c : [] 9223372036854775807 false,2;64af4f0d844d0e1c%00; : [] 9223372036854775807 false) ssil=[] ssio={} 2013-01-15 15:59:12,339 [tabletserver.TabletServer] DEBUG: ScanSess tid null !0 142 entries in 0.00 secs, nbTimes = [0 0 0.00 1] 2013-01-15 15:59:12,339 [util.MetadataTable] DEBUG: KV 2;64af4f0d844d0e1c;64ac083126e985b6 2;64af4f0d844d0e1c future:13c397780374fe2 [] 229776 false 127.0.0.1:9997 2013-01-15 15:59:12,339 [util.MetadataTable] DEBUG: KV 2;64af4f0d844d0e1c;64ac083126e985b6 2;64af4f0d844d0e1c srv:dir [] 29521 false /t-00007l9 2013-01-15 15:59:12,339 [util.MetadataTable] DEBUG: KV 2;64af4f0d844d0e1c;64ac083126e985b6 2;64af4f0d844d0e1c srv:lock [] 29521 false tservers/127.0.0.1:9997/zlock-0000000000$13c397780374c3b 2013-01-15 15:59:12,339 [util.MetadataTable] DEBUG: KV 2;64af4f0d844d0e1c;64ac083126e985b6 2;64af4f0d844d0e1c srv:time [] 29521 false M0 2013-01-15 15:59:12,339 [util.MetadataTable] DEBUG: KV 2;64af4f0d844d0e1c;64ac083126e985b6 2;64af4f0d844d0e1c ~tab:~pr [] 29521 false ^A64ac083126e985b6 2013-01-15 15:59:12,339 [util.MetadataTable] DEBUG: KV 2;64af4f0d844d0e1c;64ac083126e985b6 2;64b295e9e1b09682 loc:13c397780374f85 [] 209076 false 127.0.0.1:9997 2013-01-15 15:59:12,339 [util.MetadataTable] DEBUG: KV 2;64af4f0d844d0e1c;64ac083126e985b6 2;64b295e9e1b09682 ~tab:~pr [] 29516 false ^A64af4f0d844d0e1c {noformat} Why would a range of [2;64af4f0d844d0e1c, 2;64af4f0d844d0e1c%00) return a Key with row 2;64b295e9e1b09682 ?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-962

#	Requirement	Questions	Answered	Successfully Answered	URL
17592	ACCUMULO-958	Why do this for just walogs and not files? Would the cypherstreams handle sync/flush properly?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-958
17593	ACCUMULO-958	I was still under the impression that a ticket owner was responsible for reviewing and applying contributor patches. Did we change our methodology for this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-958
17594	ACCUMULO-958	For future record then, how long is appropriate? Because then it's worth mentioning for any future feature freezes that patches need to be in X time earlier to be brought it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-958
17595	ACCUMULO-958	Isn't there already an extension mechanism for FileSystem? Replace the FileSystem with a new kind of FileSystem which will de/encrypt data as it is read or written.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-958
17596	ACCUMULO-958	If this is going to make it in to the existing code (and stay), however polished it will be by the next release, I'd like to see it clearly marked as experimental, until it is available as a complete and coherent framework for encrypting table contents. So, I suggest moving the relevant classes into an "experimental" sub-package, and minimizing references to them in other code. I looked for a built-in "@Experimental" annotation, but couldn't find one, so we could create one for this sort of thing (but I still prefer the sub-package until it is no longer experimental). I do *not* think that they should be marked as "@Deprecated" because that implies a completely different point in the life cycle of the code (in fact, it implies the opposite end of that life cycle). That said, what exactly are the next actions, and the timeline for polishing this feature? From the previous comment, I gather "tests", and "tidiness" (which I interpret to mean QA refactorings, but not functional changes that incorporate critical feedback). Are there more anticipated actions that I've overlooked?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-958
17597	ACCUMULO-958	Since the interface has changed in this new patch, should we revert the first patch for 1.5?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-958
17598	ACCUMULO-958	Are you doing encryption before or after compression? It seems like that could be important. Presumably the data will compress better before encryption.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-958
17599	ACCUMULO-957	I like this. Can the commit of this ticket include putting the state diagram in the master comments in one of these?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-957
17600	ACCUMULO-956	I not sure there is a problem. The KeyTransformingIterator transforms only single keys, so you can't do any aggregation. If you had some kind of aggregating iterator on top of it on the stack, it would need to read enough keys from its source to do the aggregation. To produce your example, I think you'd have to scan all the cf: keys from the source to know it has what it needs for r cf: even [v2, v4]. Since that iterator would be scanning its source, there is no problem with the stack getting torn down and re-seeked. Or am I missing something?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-956
17601	ACCUMULO-956	The latest patch looks good. I was thinking of commiting and then making some further modifications. Was there anything else you wanted to do?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-956
17602	ACCUMULO-955	Integrated in Accumulo-Trunk-Hadoop-2.0 #5 (See [https://builds.apache.org/job/Accumulo-Trunk-Hadoop-2.0/5/]) ACCUMULO-955 Made BatchWriterConfig Writable, so it can be stored in a job's configuration in a human-readable way. Updated AccumuloOutputFormat to use it, and added unit tests for ACCUMULO-706 and ACCUMULO-955. Added an additional check for a reasonable minimum maxMemory value. (Revision 1435013) Result = FAILURE ctubbsii : Files : * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/BatchWriterConfig.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/mapreduce/AccumuloOutputFormat.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/util/shell/commands/DeleteCommand.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/util/shell/commands/InsertCommand.java * /accumulo/trunk/core/src/test/java/org/apache/accumulo/core/client/BatchWriterConfigTest.java * /accumulo/trunk/core/src/test/java/org/apache/accumulo/core/client/mapreduce/AccumuloOutputFormatTest.java * /accumulo/trunk/examples/simple/src/main/java/org/apache/accumulo/examples/simple/mapreduce/TeraSortIngest.java * /accumulo/trunk/server/src/main/java/org/apache/accumulo/server/test/continuous/ContinuousMoru.java	1	1	https://issues.apache.org/jira/browse/ACCUMULO-955
17603	ACCUMULO-955	Integrated in Accumulo-Trunk #647 (See [https://builds.apache.org/job/Accumulo-Trunk/647/]) ACCUMULO-955 Made BatchWriterConfig Writable, so it can be stored in a job's configuration in a human-readable way. Updated AccumuloOutputFormat to use it, and added unit tests for ACCUMULO-706 and ACCUMULO-955. Added an additional check for a reasonable minimum maxMemory value. (Revision 1435013) Result = SUCCESS ctubbsii : Files : * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/BatchWriterConfig.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/client/mapreduce/AccumuloOutputFormat.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/util/shell/commands/DeleteCommand.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/util/shell/commands/InsertCommand.java * /accumulo/trunk/core/src/test/java/org/apache/accumulo/core/client/BatchWriterConfigTest.java * /accumulo/trunk/core/src/test/java/org/apache/accumulo/core/client/mapreduce/AccumuloOutputFormatTest.java * /accumulo/trunk/examples/simple/src/main/java/org/apache/accumulo/examples/simple/mapreduce/TeraSortIngest.java * /accumulo/trunk/server/src/main/java/org/apache/accumulo/server/test/continuous/ContinuousMoru.java	1	1	https://issues.apache.org/jira/browse/ACCUMULO-955
17604	ACCUMULO-952	Could possible just make the command fail if it would create tablet with over a certain number of files. The user could override this with the -f option maybe.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-952
17605	ACCUMULO-948	We could change the log level on these message to debug. I think its warn now. Or maybe we can log a debug in some cases and warning in others. Curious what is triggering these for you. Is it an invalid config or an empty dir?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-948
17606	ACCUMULO-948	So the screwy config should be removed. That leaves us with the non-existant dir, should we warn about this? I am thinking we should. Whats the use case for having the config point to a non-existant dir?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-948
17607	ACCUMULO-936	Is this fixed by ACCUMULO-926?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-936
17608	ACCUMULO-935	I think the most appropriate place to stage a runnable Accumulo distribution in one's workspace (helpful for developers) is probably in the assemble/target directory. This would simplify much of the pom files that store artifacts in "../lib". Instead, the assemble project can simply depend on the modules it assembles, and use copy-dependencies to stage things in assemble/target/accumulo_home/lib. This eliminates a lot of additional configuration being done that assumes all modules are checked out and being built together, and removes complexity in plugin configurations to specify their output directories.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-935
17609	ACCUMULO-935	Won't this break the ability to run directly off of the classfiles? Or at least make it a giant headache to do so?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-935

#	Requirement	Questions	Answered	Successfully Answered	URL
17610	ACCUMULO-935	How do you use the new builds? Does it wipe out everything in the directory you specify and drop in the new files, so that you need to reconfigure the dev instance when you build?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-935
17611	ACCUMULO-927	What about a case in which I have my-code-1.0.jar deployed. I then find a problem in my code, cut a new version and deploy my-code-1.1.jar. Say in version 1.1, I have a class in which I added a new method that did not exist in 1.0. Wouldn't I have undefined behavior (my class in 1.0 or 1.1 could be used by Accumulo) if the 1.0 classes aren't removed completely? bq. I don't think it will be common that folks delete jars from HDFS without restarting Accumulo as the classes that are loaded from that jar will still be resident in the VM.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-927
17612	ACCUMULO-927	But your v1.0 class would still be loaded in the VM. I don't think there is a guarantee as to when classes get unloaded from the VM assuming that you have enabled those options. Your scenario would only work if the class was in a different package. You could however deploy the jar to a new location and set the table context to the new location. Sent from my Motorola ATRIX™ 4G on AT&T -----Original message----- From: "Josh Elser (JIRA)" <jira@apache.org> To: notifications@accumulo.apache.org Sent: Fri, Jan 4, 2013 02:54:12 GMT+00:00 Subject: [jira] [Commented] (ACCUMULO-927) Replace HDFS VFS FileSystem objects with the ones from Commons VFS [https://issues.apache.org/jira/browse/ACCUMULO-927?page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel&focusedCommentId=13543557#comment-13543557] Josh Elser commented on ACCUMULO-927: ---------- What about a case in which I have my-code-1.0.jar deployed. I then find a problem in my code, cut a new version and deploy my-code-1.1.jar. Say in version 1.1, I have a class in which I added a new method that did not exist in 1.0. Wouldn't I have undefined behavior (my class in 1.0 or 1.1 could be used by Accumulo) if the 1.0 classes aren't removed completely? bq. I don't think it will be common that folks delete jars from HDFS without restarting Accumulo as the classes that are loaded from that jar will still be resident in the VM. -- This message is automatically generated by JIRA. If you think it was sent incorrectly, please contact your JIRA administrators For more information on JIRA, see: http://www.atlassian.com/software/jira	1	0	https://issues.apache.org/jira/browse/ACCUMULO-927
17613	ACCUMULO-927	I take that back. You raise a valid point. In any case of a change to the monitored directory, the context classloader is recreated. I will look into this more to see if the deleted jar is ever removed from the local cache directory. Sent from my Motorola ATRIX™ 4G on AT&T -----Original message----- From: "Dave Marion (JIRA)" <jira@apache.org> To: notifications@accumulo.apache.org Sent: Fri, Jan 4, 2013 03:04:13 GMT+00:00 Subject: [jira] [Commented] (ACCUMULO-927) Replace HDFS VFS FileSystem objects with the ones from Commons VFS [https://issues.apache.org/jira/browse/ACCUMULO-927?page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel&focusedCommentId=13543564#comment-13543564] Dave Marion commented on ACCUMULO-927: ---------- But your v1.0 class would still be loaded in the VM. I don't think there is a guarantee as to when classes get unloaded from the VM assuming that you have enabled those options. Your scenario would only work if the class was in a different package. You could however deploy the jar to a new location and set the table context to the new location. Sent from my Motorola ATRIX™ 4G on AT&T -----Original message----- From: "Josh Elser (JIRA)" <jira@apache.org> To: notifications@accumulo.apache.org Sent: Fri, Jan 4, 2013 02:54:12 GMT+00:00 Subject: [jira] [Commented] (ACCUMULO-927) Replace HDFS VFS FileSystem objects with the ones from Commons VFS [https://issues.apache.org/jira/browse/ACCUMULO-927?page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel&focusedCommentId=13543557#comment-13543557] Josh Elser commented on ACCUMULO-927: ---------- What about a case in which I have my-code-1.0.jar deployed. I then find a problem in my code, cut a new version and deploy my-code-1.1.jar. Say in version 1.1, I have a class in which I added a new method that did not exist in 1.0. Wouldn't I have undefined behavior (my class in 1.0 or 1.1 could be used by Accumulo) if the 1.0 classes aren't removed completely? bq. I don't think it will be common that folks delete jars from HDFS without restarting Accumulo as the classes that are loaded from that jar will still be resident in the VM. -- This message is automatically generated by JIRA. If you think it was sent incorrectly, please contact your JIRA administrators For more information on JIRA, see: http://www.atlassian.com/software/jira -- This message is automatically generated by JIRA. If you think it was sent incorrectly, please contact your JIRA administrators For more information on JIRA, see: http://www.atlassian.com/software/jira	1	0	https://issues.apache.org/jira/browse/ACCUMULO-927
17614	ACCUMULO-925	Why can't it read the pid file dir out of the env file on the remote box?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-925
17615	ACCUMULO-925	Is there a reason why we can't apply this back to 1.6?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-925
17616	ACCUMULO-925	bq. Is there a reason why we can't apply this back to 1.6? Should probably have a discussion about it -- need to consider the impact on user automation around our scripts. We haven't really identified how the scripts should be covered by our compatibility statement (we have nothing about them now which is probably negligence). My gut reaction is that it wouldn't be worth it, but I'm not against someone doing the work to backport it (as long as we're sure we don't break scripts). It also would require the .out/.err file rotation (or at least splitting that logic out). IIRC, the hope was that usage of the existing start/stop-server and start/stop-all scripts didn't change, but I don't remember enough to say it is definitely compatible.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-925
17617	ACCUMULO-923	How do I programmatically set a table property? I want to set Property.TABLE_LOAD_BALANCER so that I can have a test balancer class which throws an exception in specific situations. I need to enable "configuration.getTableConfiguration(table).get(Property.TABLE_LOAD_BALANCER);" to return the specific mock class needed for each unit test. This would enable me to write unit tests ensuring the table name is in the exception messages.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-923
17618	ACCUMULO-918	Integrated in Accumulo-Trunk-Hadoop-2.0 #46 (See [https://builds.apache.org/job/Accumulo-Trunk-Hadoop-2.0/46/]) ACCUMULO-918 Added user iterator for visibilities. TODO: add a test. (Revision 1440118) Result = SUCCESS ctubbsii : Files : * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/iterators/system/VisibilityFilter.java * /accumulo/trunk/core/src/main/java/org/apache/accumulo/core/iterators/user/VisibilityFilter.java	0	0	https://issues.apache.org/jira/browse/ACCUMULO-918
17619	ACCUMULO-905	I was looking at the release notes for Java 7. Its up to update 10 now. I may have missed something, but the last release that had bug fix that seemed possibly concerning was u5. Some bug where the JVM would crash when PermGen was full, not sure if this bug was big deal or not. http://www.oracle.com/technetwork/java/javase/7u-relnotes-515228.html bq. There are known JDK 1.7 and zookeeper 3.4 issues. Can you elaborate?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-905
17620	ACCUMULO-905	I believe they were released a day or 2 ago, 6u39 and 7u13 - http://www.webupd8.org/2013/02/oracle-releases-critical-update-for.html Was that the exploit you're thinking of? Also, this is server software, browser plugins are something I'm not that concerned about.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-905

#	Requirement	Questions	Answered	Successfully Answered	URL
17621	ACCUMULO-905	For Accumulo 1.5, if you test it against Java 1.7 and find issues please open bugs. I agree w/ Benson dropping support for Java 1.6 is a serious decision. We can easily produce class files that work against Java 1.6 and Java 1.7 as long as we do not utilize new Java 1.7 features in Accumulo 1.6. While using new Java features would be nice, would any of these new features be necessary? From a security standpoint, a user can use Java 1.7 without us producing class files that only work with Java 1.7.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-905
17622	ACCUMULO-905	In short, no. But how long is an appropriate time to wait before moving on, since EOL seems to be the incorrect timeframe.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-905
17623	ACCUMULO-905	bq. But how long is an appropriate time to wait before moving on, since EOL seems to be the incorrect timeframe. For Accumulo 1.6.0 if you are interested in switching to Java 1.7, I propose you call a vote on the issue. If not for Accumulo 1.6.0 then we can reconsider the issue for Accumulo 1.7.0.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-905
17624	ACCUMULO-905	Update: The last security patch for Oracle JDK was 2013-04-16 (1.6.0_45). And, it looks like OpenJDK hasn't had a tag for 1.6 in about 8 months. OpenJDK was dropped from Fedora in version 16 (currently Fedora is 18, nearing 19, on ~6 month release cycles). I'm not sure what OpenJDK 1.6 support looks like in the next major RHEL/CentOS/Ubuntu releases, but I know that OpenJDK 1.7 has been supported since RHEL/CentOS 6.3 (~June 2012). I think the Accumulo 1.5 line will be stable enough to sustain users slow to adopt a newer version of Java for the foreseeable future. I don't think we need to hold back use of newer features for this any longer, personally. However, Keith is right... we should call a vote on the dev list.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-905
17625	ACCUMULO-905	Are there any statistics about how many groups are using 1.6 versus 1.7? What is the government using? On Wed, May 22, 2013 at 11:53 PM, Christopher Tubbs (JIRA)	0	0	https://issues.apache.org/jira/browse/ACCUMULO-905
17626	ACCUMULO-902	Do you have a use case? In the meantime, separate thread pools seem nearly ideal to handle a few fast minor compactions, and a queue of slower major compactions.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-902
17627	ACCUMULO-901	Does the master unload them all simultaneously or does it do it serially? If the former, then they shouldn't try to major compact before unloading. If the latter, perhaps we should introduce a majC halt as well.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-901
17628	ACCUMULO-898	We are talking about two distinct tasks now, right? First is to create an Storage implementation for Zipkin using Accumulo. Second is then to use Zipkin, and the Accumulo back-end from the first task, to instrument Accumulo code and store the traces to Accumulo. Am I following? I would like to help out somewhere if you think you could use it.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-898
17629	ACCUMULO-898	I got some quick and dirty code running using HTrace to log to Zipkin. Looks like they added a ZipkinSpanReciever to HTrace just the other week so this seemed like the easiest way to be able to push sample traces into Zipkin. I am sure it would be straightforward to convert the current Accumulo tracing code to push to a Zipkin collector too seeing as HTrace looks like they borrowed a lot of code from Accumulo tracing code. Have you by any chance started on the Storage implementation? I need to learn some Scala but as there are several Storage implementations to use as examples I guess it shouldn't be too hard. Let me know if you are interested in the HTrace example and I can share.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-898
17630	ACCUMULO-898	bq. The patch wouldn't apply cleanly with git am Is that due to whitespace errors? Those errors are in the generated thrift classes. Do we reformat those? bq. Since the SpanReceivers and related utilities are associated with the accumulo-tracer service, they should probably be moved there. I was worried about the dependency hierarchy there, since the jar containing the SRs will have to be added to the Hadoop classpath, and we probably don't want to be adding a bunch of jars. Though of course we don't really want there to be entire Accumulo modules with just a couple of classes in them either. Maybe it would be okay if dependencies on core and other modules don't creep into the SRs. bq. I'm not sure I understand the distinction in roles the SpanReceiverHost is supposed to serve, vs. the TraceUtil. SpanReceiverHost is kind of copied from Hadoop, from when we were planning to share the config, with the idea of switching over to using Hadoop's version at some point. But if we're not sharing config, there's no need to preserve it.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-898
17631	ACCUMULO-898	I've attached my current patch set, if you would like to compare. Can you re-assign this issue to yourself while you are working on it? Note that I got rid of the Hdfs wrapper classes, since HDFS has its own spans.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-898
17632	ACCUMULO-898	This is really cool. Does anyone know if someone is actively working on the Accumulo Storage implementation for Zipkin? I had started long ago (https://github.com/jaredwinick/zipkin/tree/master/zipkin-accumulo) but never came close to finishing.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-898
17633	ACCUMULO-895	Is there a reason that {{instance.dfs.dir}} requires the value to start with {{"/"}}?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-895
17634	ACCUMULO-895	bq. Is there a reason that instance.dfs.dir requires the value to start with "/"? I think the intent was to ensure it was an absolute path. It seems Path has an isAbsolute() method that handles windows paths. I will modify to use that.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-895

#	Requirement	Questions	Answered	Successfully Answered	URL
17635	ACCUMULO-895	Can run the {{ExampleAccumuloUnitTest}} now! Having issues with {{MapReduceExample}}, though. {noformat} ---- Running Mapred Against Accumulo UUID 1355244649903 81398c27-297e-4087-8ea8-ec0bdb976220 FLUSH 1355244659049 9144 24 1000000 1000000 FLUSH 1355244668456 9407 16 2000000 1000000 FLUSH 1355244680016 11560 27 3000000 1000000 FLUSH 1355244689021 9005 15 4000000 1000000 FLUSH 1355244703781 14760 13 5000000 1000000 [main] ERROR org.apache.hadoop.security.UserGroupInformation - PriviledgedActionException as:mike cause:java.io.IOException: Failed to set permissions of path: \tmp\hadoop-mike\mapred\staging\mike-101115261\.staging to 0700 Exception in thread "main" java.io.IOException: Failed to set permissions of path: \tmp\hadoop-mike\mapred\staging\mike-101115261\.staging to 0700 at org.apache.hadoop.fs.FileUtil.checkReturnValue(FileUtil.java:682) at org.apache.hadoop.fs.FileUtil.setPermission(FileUtil.java:655) at org.apache.hadoop.fs.RawLocalFileSystem.setPermission(RawLocalFileSystem.java:509) at org.apache.hadoop.fs.RawLocalFileSystem.mkdirs(RawLocalFileSystem.java:344) at org.apache.hadoop.fs.FilterFileSystem.mkdirs(FilterFileSystem.java:189) at org.apache.hadoop.mapreduce.JobSubmissionFiles.getStagingDir(JobSubmissionFiles.java:116) at org.apache.hadoop.mapred.JobClient$2.run(JobClient.java:856) at org.apache.hadoop.mapred.JobClient$2.run(JobClient.java:850) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.Subject.doAs(Unknown Source) at org.apache.hadoop.security.UserGroupInformation.doAs(UserGroupInformation.java:1093) at org.apache.hadoop.mapred.JobClient.submitJobInternal(JobClient.java:850) at org.apache.hadoop.mapreduce.Job.submit(Job.java:500) at org.apache.hadoop.mapreduce.Job.waitForCompletion(Job.java:530) at org.apache.accumulo.server.test.continuous.ContinuousVerify.run(ContinuousVerify.java:216) at org.apache.hadoop.util.ToolRunner.run(ToolRunner.java:65) at org.apache.accumulo.server.test.continuous.ContinuousVerify.main(ContinuousVerify.java:232) at org.apache.accumulo.instamo.MapReduceExample.run(MapReduceExample.java:43) at org.apache.accumulo.instamo.MapReduceExample.main(MapReduceExample.java:60) {noformat} Perhaps need to specify a mapred work directory in some conf before launching?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-895
17636	ACCUMULO-895	I figured out why it fails in Hadoop 2.0. The change below in Path makes it only perform Windows slash normalization if you're on a Windows system. I'm going to remove the Windows path test from PropertyTest until this get fixed. Think this is worth opening a Hadoop ticket? {noformat} static final boolean WINDOWS = System.getProperty("os.name").startsWith("Windows"); {noformat} {noformat} - if (path.indexOf("//") != -1) { - path = path.replace("//", "/"); - } - if (path.indexOf("\\") != -1) { - path = path.replace("\\", "/"); + path = StringUtils.replace(path, "//", "/"); + if (Path.WINDOWS) { + path = StringUtils.replace(path, "\\", "/"); {noformat}	0	0	https://issues.apache.org/jira/browse/ACCUMULO-895
17637	ACCUMULO-895	Can we not just use File.pathSeperator?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-895
17638	ACCUMULO-895	bq. These changes make MiniAccumuloCluster work on Windows. This is for Accumulo 1.4. Does anyone have an environment set up to build on Windows that can verify this patch fixes things? I'm a little scared to go down that path, personally.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-895
17639	ACCUMULO-894	Would it make sense to add this in an injectable way, with annotations? It seems like this is more future-proof, instead of changing the API. Alternatively it could be added as a final helper method to a base class, but I have a preference for annotations/injection for environment stuff like this.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-894
17640	ACCUMULO-894	Can somebody give an example of where this would be helpful?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-894
17641	ACCUMULO-891	Can you clarify how the API is making Authorizations seem immutable?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-891
17642	ACCUMULO-884	Do things break horribly if it's set in the client but not on the server? If not, we should also just set it in all of the examples.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-884
17643	ACCUMULO-884	Why was this bumped? I am under the impression that using it is just a documentation change? Sent from my phone, please pardon the typos and brevity.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-884
17644	ACCUMULO-884	bq. In my experiments on these solid state drives, enabling short-circuit reads more than doubled my read throughput! (TP measured in ops/s in a YCSB-derived read-only workload test.) If I remember correctly, I read somewhere that you won't see any benefit from shortcircuit reads until you actually get to hadoop-2 (maybe 0.23 too?). I'll see if I can find that information again... Interesting that you saw such a speedup. How influenced do you think your benchmark is by the YCSB workload itself? I did a bunch of testing early this year on spinning-disks with these parameters on 0.20 baseline and actually saw no performance gain trying to use the shortcircuit. I think I was also tweaking disabling checksums on local datanode reads. I think I had an even part read+write workload. JD had some info on HDFS-2246 but I'll leave it up to you to come to your own conclusions. Out of curiosity, in your read-only test, did you warm the Accumulo or OS caches before the test (or conversely, ensure they were cold)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-884
17645	ACCUMULO-884	bq. Interesting that you saw such a speedup. How influenced do you think your benchmark is by the YCSB workload itself? Our test was specifically designed to do uniformly random reads across a dataset that wouldn't come close to fitting in memory. Additionally, it was on solid state drives so we would assume that the relative overhead of pulling data through HDFS (as opposed to short-circuiting) would be large compared to a spinning disk scenario. bq. Out of curiosity, in your read-only test, did you warm the Accumulo or OS caches before the test (or conversely, ensure they were cold)? I didn't do anything special to warm/empty OS or Accumulo caches. Anecdotally, there was a pretty obvious warmup curve in read TP over the first few minutes of the test whenever Accumulo was not already warm with this table.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-884
17646	ACCUMULO-884	bq. Given that this requires that we run as the same user as the data node Is that the case? I thought it could be any user, just only one user.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-884
17647	ACCUMULO-884	What is the deliverable for this ticket? Is it a page of documentation on our web site? Is it part of the Accumulo configuration? Can we specify what tests need to be run as part of the closure criteria?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-884
17648	ACCUMULO-884	[~mridley], would you be willing to take on the documentation changes from this ticket? Your results would make a good blog post, either directly on the accumulo project blog or syndicated there.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-884
17649	ACCUMULO-869	Some comments about patch #1 * If isolation is enabled, should the parent classloader be the java system classloader, instead of the classloader that pulls stuff from $ACCUMULO_HOME/lib * Should isolation be configurable per context?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-869
17650	ACCUMULO-868	To change the directory from /tmp to something else, then we need to change line 273 in AccumuloVFSClassLoader from {noformat} vfs.setReplicator(new DefaultFileReplicator()); {noformat} and pass in a java.io.File instance that references a directory. Are we going to default this to something and also make configurable? What about $ACCUMULO_HOME/tmp ?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-868

#	Requirement	Questions	Answered	Successfully Answered	URL
17651	ACCUMULO-868	If the cache dir is not configured I think we should place it in the system tmp dir. There is no good way to do this in Java (until Java 7). Could do something like the following {code:java} cacheDir = new File(org.apache.commons.io.FileUtils.getTempDirectory(), "accumulo-vfs-cache"); {code} The nice thing about FileUtils.getTempDirectory() is that it returns something reasonable for Linux, Windows, and Mac. Is it safe for the multiple processes to use the same cache dir?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-868
17652	ACCUMULO-868	What do we do for the logs directory? Can we put the cache dir inside the logs dir?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-868
17653	ACCUMULO-868	bq. What do we do for the logs directory? Can we put the cache dir inside the logs dir? The default for walogs and logs is $ACCUMULO_HOME. User are encouraged to change both. Another problem I have seen with ACCUMULO_HOME is its on NFS sometimes, which is not good for logs or cache. System tmp dir will be local and writable. I share your concern about tmp filling up, I have seen this happen and when tmp dir is in OS partition it can be a disaster. So I think tmp has the best chance of working out of the box for a new user, but could cause problems in the long term on production systems. We could put the property in the example config (like walog config), encouraging user to set it to something appropriate for their system. Default it system tmp so it works out of the box for new users, but make system admins cognizant of this config for production systems. Comment in the example config should point out that this propery only matters if using vfs classloader.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-868
17654	ACCUMULO-853	What's with the changes in CachedBlockRead?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-853
17655	ACCUMULO-842	Is there a way to describe this task more concretely so that others might attempt it? Can the work be phased? Why is this a bug? Can you estimate the difficulty level?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-842
17656	ACCUMULO-842	Damon, Thanks for the patch, I have a few comments. * I am not keen on having an option to run if the master lock is held. I think modifications to Fate ops should only run if the master lock is not held. If there is a buggy fate op thats stuck in an infinite loop, deleting it from the ZooStore will not necessarily kill the thread in the master. In the case a user would still need to kill the master. I would not want the user to get the impression that this shell command will clean up state in the master process, it will not. Also, could get strange behavior in the master from concurrent modification of zookeeper. I think always killing the master process is the safest option. Killing the master for a bit should not impact reads and writes to Accumulo, but it will stall table operations. * seems to duplicates some code in AdminUtil, consolidate? * TxFailCommand.prepFail() is a nice improvement, I think this could be moved to fate.AdminUtil and have shell call AdminUtil * TxDeleteCommand.prepDelete() should always delete, even its NEW or IN_PROGRESS. Consider a buggy fate op thats in progress and always gets stuck in an infinite loop when it tries to run. Maybe just use fate.AdminUtil.prepDelete(). * consider using shellState.printLines(), it does pagination, in TxPrintCommand * javadoc for TxPrintCommand wrong * I think I am slightly in favor of having one shell command with options, less clutter in the shells list of commands. Was there a reason you chose to create three commands? * Did you use the Accumulo Eclipse formatter? The formatting seems mostly correct, but I am seeing some slight diffs like extra spaces in parametric types. Its odd. * May want to consider making the current Admin command log a warning saying its deperecated and point users to the shell commands. For testing, you can pick a random fate op and add an infinite loop to it :)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-842
17657	ACCUMULO-841	Are you sure? The xml file for security at one point set it. Sent from my phone, please pardon the typos and brevity.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-841
17658	ACCUMULO-841	Isn't the node counting necessary to support maximum hop?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-841
17659	ACCUMULO-839	Is anyone working on this issue or any other procedure which can enable search capabilities in Accumulo?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-839
17660	ACCUMULO-824	Where are the scripts located and where is the code to develop the packaging (the RPM)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-824
17661	ACCUMULO-818	Dave, can you please assign a ticket to yourself when you begin working it? Memory serves that a Jira ticket with an assignee implies that someone is working on it. Thanks!	1	1	https://issues.apache.org/jira/browse/ACCUMULO-818
17662	ACCUMULO-818	John, Are you referring to the api change in 818 or the shell command in 967? I'm thinking you mean the api change, but want to be sure. Dave	1	0	https://issues.apache.org/jira/browse/ACCUMULO-818
17663	ACCUMULO-813	Was this related to dat with a bad visibility? If so, ACCUMULO-360 may help avoid the situation. ACCUMULO-918 and ACCUMULO-844 will help users deal with this situation should it occur. Handling issues that arise higer in the iterator stack is tricky, but not impossible. For a corrupt rfile, where the problem is detected in the rfile code, handling the situation should be more straightforward. Want to understand the issue you ran into more before working on it.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-813
17664	ACCUMULO-804	Hadoop isn't a compiled dependency. Couldn't we keep the default with 1 for compilation time, but have a testing procedure of some sort which exercises both?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-804
17665	ACCUMULO-804	Should that be another dimension of the example configs? It would be better if we didn't make users have to reconfigure the classpath.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-804
17666	ACCUMULO-804	We already have 8 sets, so this would make it 16? Might be a little excessive. How does Accumulo respond to directories in general.classpaths that don't exist? If that doesn't cause an error, we could just add the 2.0 classpaths into the existing examples and they'd work for both.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-804
17667	ACCUMULO-803	looks good a few comments. * can you make this patch against trunk/1.5 instead of 1.4? I do not think this is something we would want to add to 1.4. * the user manual would need some mention of this new feature, in docs/src/user_manual/chapters/table_configuration.tex I think * in MockTable I think you should subtract from Long.MAX_VALUE instead of Integer.MAX_VALUE for consistency, even though its an int is being subtracted	1	1	https://issues.apache.org/jira/browse/ACCUMULO-803
17668	ACCUMULO-803	As a thought exercise, let's think about how far this might go. Since Combiners work on a column family, does that imply that we might want a different timestamp scheme on different column families? Should this turn into another plugin framework for assigning timestamps on tablet servers, or perhaps a timestamp data definition language for column families?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-803
17669	ACCUMULO-803	bq. does that imply that we might want a different timestamp scheme on different column families? Currently each tablets timestamp is stored in the metadata table. It tablets had to persist and indeterminate number of timestamps, then I do not think we could safely store that in the metadata table. Would need to store it somewhere else. Its nice storing the info in the metadata table because you can atomically update the timestamp on minor compaction and bulkimport by putting it in the same mutation. So where would this info be stored and how would it be atomically updated after minor compaction and bulk import? Would also have to be careful not to blow out memory on tservers. e.g. If each tablet on a tserver is keeping track of timestamps for many columns in memory. Would need to cache timestamp info from persisted store to avoid this problem.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-803

#	Requirement	Questions	Answered	Successfully Answered	URL
17670	ACCUMULO-803	bq. Since Combiners work on a column family, does that imply that we might want a different timestamp scheme on different column families? We might also consider setting an option on mutations to reverse system set time. This gives users the ability to target this in very tailored ways. Users could apply it to certain column family prefixes, to column qualifiers, etc. It gives a lot of flexibility w/o any server side code. Could add something like the following to mutation. {code:java} enum TimestampTransformations { NONE, REVERSE } class Mutation { //transformation that will be applied to system set timestamps public void setTimestampTransformation(TimestampTransformations tt); } {code} I think the only drawback with this is that it does not support setting system time for bulk insert. The current patch will support that.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-803
17671	ACCUMULO-803	I like the setting of timestamp transformation in the Mutation, but do you think people might want to set multiple different transformations for different column updates in the same mutation?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-803
17672	ACCUMULO-803	bq. I like the setting of timestamp transformation in the Mutation, but do you think people might want to set multiple different transformations for different column updates in the same mutation? I like it too, but the bulk import aspect is still tormenting me. I was being miserly with computation and storage when thinking about setting it at the mutation level. Eric just modified mutation so that the system timstamp is set once per mutation. It used to set the same timestamp for each column. I wanted to maintain this slight boost in ingest performance. I suppose we could still set the timestamp once permutation and just interpret it differently if a flag is set for the column when getTimestamp() is called. For bulk import, if the user request logical time, it will set the same timestamp for all keys in the file. We could make it reverse this timestamp for all keys in the file. But I have not thought of a simple method for do something more granular. Seems like the two are options are modify the Key written by bulk import or modify the special iterator that sets timestamps on bulk imported files when they are read. Don't really like this options. But we don't have to support anything more granular for bulk import.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-803
17673	ACCUMULO-803	bq. do you think people might want to set multiple different transformations for different column updates in the same mutation? One more consideration. This would add another set of put methods to mutation, which already has alot of put methods.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-803
17674	ACCUMULO-803	I agree, I don't feel comfortable allowing users to set column family ordering on a per mutation basis. With the API you propose, do you feel that fifo ordering would be best achieved by manipulating timestamps as is done the original patch? If so, does this introduce any strangeness in the relationship between typeType and fifoColumnFamilies? Could it make sense to use TimeType.MILLIS and a fifo column family? Perhaps yes. I feel a little uncomfortable with the fact that column families are dynamic but we would require the user to specify a set of column families at creation time if they want fifo behavior. This seems to run counter to the spirit of column families (and how they're used) in Accumulo. One way to solve this problem would be to add setColumnFamilyOrdering(String cf, boolean fifo) (or something like it) to TableOperations. This might not be a good idea because we run into the same problem we have with Mutations: the user could shoot themselves in the foot if they set the ordering on a column family to be different than the default for a table that already contains data. So, I have to admit that I lean closest to the original mechanism, but I could be biased because I'm the patch author :) Thoughts?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-803
17675	ACCUMULO-803	I feel bad about bumping this to 1.7 because it got bumped to 1.6 with a patch. [~alfarris@prodigy.net] or [~kturner], is it possible for you to do a revised patch for 1.6?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-803
17676	ACCUMULO-802	Took a quick look at v2 patch. Had the following comments. * maybe only delete namespaces if empty... could make attempt to delete non-empty namespace throw an error (could have a force option) * verify namespace props on server side.. for example configure iterator on namespace then try scanning/compacting... test this w/ and w/o table overriding prop... * test that namespace props survive across namespace rename (a good random walk test would do this) * need to test import and export table w/ namespaces * uncertain about automatic creation of namespaces at table creation time... this may hide errors in user code * maybe support cloning namespaces, a shortcut to clone all tables in a namespace + copy namespace props * it would be nice if the unit test covered all operations that can now take a namespace, for example get scanner, batchscanner, etc using namepsace+table	1	0	https://issues.apache.org/jira/browse/ACCUMULO-802
17677	ACCUMULO-802	A few additions to the v2 patch: # Don't keep both a {code}{ZKROOT}/namespaces/{namespaceId}/tables{code} and a {code}{ZKROOT}/tables/{tableId}/namespace{code} It makes things more complex when we have to maintain two structures for the same information, and will result in an inability to determine which is the authoritative source consistently if they diverge (we had this problem with tables in ZK and in the metadata table awhile back). I prefer the latter, because I think the table's own configuration should be the authoritative source of who has additional configuration for it. # Make sure the bytes written to ZooKeeper from strings (table namespace IDs, namespace names), are written as UTF8 bytes, and read back as UTF8 bytes. (For example, in Tables.getNamespace(Instance, String);); Since we don't have JDK7's UTF_8 constant, use our own: Constants.UTF8. # (Optional) In the master.thrift file, please ensure that params in new methods have sequential param IDs (1,2,3...). It's easier to read, and that makes it easier to add new params to thrift RPC methods that is backwards-compatible. Don't change old methods, though. # Avoid adding javadocs that don't add value (example: '@throws Exception' param without a comment) # If you add a TODO, please include the Accumulo JIRA number in the comment so it can be tracked (example: '// TODO security check (ACCUMULO-1479)') # Update the WIRE_VERSION to 3, and put the upgrade code in the if block. The data version was already incremented to account for the root table separation, so any upgrade checks based on the data version should work seamlessly. I'm adding [^ACCUMULO-802.v3.patch] with some of my recommendations.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-802
17678	ACCUMULO-790	When RFile was created not only was the file format considered, the end to end data process from client write to client read was also considered. The encoding that is utilized in RFile is also utilized to send data over the wire. RFile allowed much faster read than MapFile, but that speedup would not have mattered as much if the data could not be sent over the network faster. See Key.compress() and Key.decompress(). If these changes further speed up read rates, can we do something to speed up transfer rates? There may be nothing to do, but its something we can experiment with.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-790
17679	ACCUMULO-786	Looking at the commit just made, I think we should restructure the code to have mutation class in the server code. The server side Mutation could extend the client side mutation and add setSystemTimestamp(). The client code should not have setSystemTimestamp(). Also there is no reason for thrift to transfer the system time stamp that the user should never set. Should probably make the mutation members private instead of package private. I see this was done for testing, but can this be accomplished with the public API?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-786
17680	ACCUMULO-780	Are you proposing getting rid of all of the sample configs and having a script that ask some questions and generates an initial set of configs?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-780

#	Requirement	Questions	Answered	Successfully Answered	URL
17681	ACCUMULO-780	Can we build this functionality into the monitor? Maybe turn it into a management tool? Or are the security implications too dire? I'm picturing some kind of stacked bar chart where each process on a given server is one section of the stack. The stack being available RAM on the server.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-780
17682	ACCUMULO-780	If this is a proposed feature for 1.6.0, what are we going to do about ACCUMULO-1401 in 1.5.[0-9]+ fixes?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-780
17683	ACCUMULO-780	Thanks for the patch! We might need to discuss if Accumulo 1.7.0 will still have Hadoop 1 support, which could impact your patch, but that's best left for another day. A few other thoughts: bq. +\t-h --hlep\t print this help message\n" Typo. :) {quote} +echo $bin +echo $SOURCE {quote} Is this left over from debugging? If so, it should be removed. Is {{getopt}} cross-platform? I generally prefer to see that to custom parsing.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-780
17684	ACCUMULO-780	I think the heuristics are a good idea, but it seems like they can be addressed in (more specific) follow on issues, can't they? [~vines]: I think the 3GB example that existed prior already suffered from that same problem, didn't it? The patch provided doesn't add any new options, except the choice of which hadoop flavor to use, which was previously just a commented-out option. I think to satisfy this ticket, without creeping too far into other issues, it would be sufficient to add an "Other" memory option and prompt for additional size options. I did see that [~John McNamee] had added a link to the ReviewBoard above.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-780
17685	ACCUMULO-780	bq. I think the heuristics are a good idea, but it seems like they can be addressed in (more specific) follow on issues, can't they? I thought this was the initial intent of this ticket. Does that mean that the scope of this ticket is to substitute Java, ZooKeeper and Hadoop home env vars into a template? I have no problem if the work is broken down into pieces (in fact, it would probably be nice for everyone), but I don't think what is currently targeted is sufficient to remove {{conf/examples/*}} completely.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-780
17686	ACCUMULO-780	If that is the intent of this ticket, then [~vines] should not have resolved ACCUMULO-1395 as a duplicate of this one. Do we want to reopen that one, have [~John McNamee] attach his patch to that, and leave this one open for the smarter heuristics? Or, do we wish him to apply his patch as part of this ticket, and create a separate one for those smarter heuristics?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-780
17687	ACCUMULO-775	Is the scanning something that can be simulated? I'm thinking that a simulation might show the new technique has a prediced x% speedup.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-775
17688	ACCUMULO-765	Question: if I bring back org.apache.accumulo.core.zookeeper.ZooCache and make it a trivial subclass of org.apache.accumulo.fate.zookeeper.ZooCache, should I also bring back {code:java} public static String lookupInstanceName(org.apache.accumulo.core.zookeeper.ZooCache zooCache, UUID instanceId) {code} or is the new lookupInstanceName that takes the fate ZooCache OK, since the core ZooCache extending the fate ZooCache should work in the new method?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-765
17689	ACCUMULO-765	bq. or is the new lookupInstanceName that takes the fate ZooCache OK, since the core ZooCache extending the fate ZooCache should work in the new method? It seems like the new method would be ok with the proposed extension. It would be nice to test it. Could compile code against 1.4 that uses that method, and then run the compiled code against 1.5.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-765
17690	ACCUMULO-759	One other thought I had, was to use negative priorities. I thought this was kinda ugly. Negative numbers are currently not allowed. Could continue to disallow negative numbers for table iterators and allow them for scan time iterators. Negative priorities would always be interpreted on the server side as "after all table iterators". If negative numbers were used, we would want to do that in such a way that the user never entered a negative number. Could add the following to IteratorSetting : {code:java} public static final int MAX_TABLE_PRIORITY = Integer.MAX_VALUE; {code} The code example I gave above would change to the following. {code:java} //assume conn and scanner are initialized somewhow, just want to show their type Connector conn; Scanner scanner; int tableMax = IteratorSetting.MAX_TABLE_PRIORITY + 1; //this is effectively Integer.MIN_VALUE scanner.addScanIterator(new IteratorSetting(tableMax++, "foo1", ".org.bar.FooIter")); scanner.addScanIterator(new IteratorSetting(tableMax++, "foo2", ".org.bar.BarIter")); {code} The above is ugly, but I just want to show my thought process. I think the code below is much less offensive from a user perspective. It does something screwy with negative numbers behind the scenes, but that is hidden from the user. {code:java} class ScanIteratorSetting extends IteratorSetting { public ScanIteratorSetting(String name, String iteratorClass) super(Integer.MIN_VALUE, name, iteratorClass); } public ScanIteratorSetting(ScanIteratorSetting predecessor, String name, String iteratorClass) super(predecessor.priority+1, name, iteratorClass); } {code} So now the code would look like this. {code:java} //assume conn and scanner are initialized somewhow, just want to show their type Connector conn; Scanner scanner; ScanIteratorSetting is1 = new ScanIteratorSetting("foo1", ".org.bar.FooIter"); //comes after all table iterators scanner.addScanIterator(is1); ScanIteratorSetting is2 = new ScanIteratorSetting(is1, "foo2", ".org.bar.BarIter"); //comes after all table iterators and after foo1 scanner.addScanIterator(is2); {code} Can we make the code for chaining iterators more compact and intuitive?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-759
17691	ACCUMULO-759	What if we went with a more immutable form of Scanner, which was more representative of the data set? This would allow things like: {code:java} Scanner allTheThings = connector.createScanner(...); ScanIteratorSetting aggregateEverything = new ScanIteratorSetting(MyAggregatingIterator.class,properties); Scanner aggregatedThings = allTheThings.transform(aggregateEverything); // scan over the table with the iterator for(Entry<Key,Value> aggregate: aggregatedThings) { ... } // scan over the table without extra scan-time iterators for(Entry<Key,Value> thing: allTheThings) { ... } {code} I think this might be an easier way for users to conceptualize iterators.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-759
17692	ACCUMULO-759	When you say IteratorChain, do you mean the one from Apache Commons Collections? If not, I would look at this package. There are some pretty cool features.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-759
17693	ACCUMULO-755	bq. I am using this Ant build script to run FindBugs. I was unable to get FindBugs working properly with Maven. What problems did you have? Perhaps I/we can work through the issues you had with the maven findbugs plugin. I'd much rather only require developers have one build system installed.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-755
17694	ACCUMULO-755	@josh - I think that mvn findbugs:findbugs creates a set of 7 findbugs.xml files. One in each module directory. The XML is not readable. The The HTML report is generated for site commands only. Reading http://kohsuke.org/2011/08/23/my-epic-battle-with-findbugs-maven-plugin-and-how-i-utterly-lost/ was demoralizing. That's when I created the Ant build file. @john - I want to facilitate everyone using findbugs even an Accumulo novice. Why bother with an IDE when it is not needed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-755
17695	ACCUMULO-752	Thanks. Eric, can you elaborate on your suggestion? Shouldn't it be safe to assume that if the user wants a different file system that they will specify it in the file name using URI syntax? What if people begin leveraging the new method in old code where a conflict might occur? Under what conditions might that happen? How would you propose handling the current MockInstance constructor that doesn't take that argument?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-752

#	Requirement	Questions	Answered	Successfully Answered	URL
17696	ACCUMULO-752	So I'm still a little confused. On what filesystem, in your example, are /tmp/foo and /tmp/bar? Are you suggesting that the patch should limit interoperability to the local filesystem? Or that the default unqualified file system should always be the local filesystem and that people should have to use fully qualified paths if they want something else?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-752
17697	ACCUMULO-747	Is it asking for too much for all of them? Maybe enabling/disabling native maps should require a restart, but the rest of them look like they can be made configurable.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-747
17698	ACCUMULO-747	Where is the list of all properties? I expected to find it in in the user manual [1] but it didn't see one... [1] http://accumulo.apache.org/1.5/accumulo_user_manual.html#_configuration_2	0	0	https://issues.apache.org/jira/browse/ACCUMULO-747
17699	ACCUMULO-745	Please be more specific. Which utilities should read accumulo-site.xml? We should probably have one ticket for each to minimize the changes related to each.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-745
17700	ACCUMULO-745	{code} 11:40:36 runTest (simple.dynamicThreadPools.DynamicThreadPools) ..... ERROR:test.auto:This looks like a stack trace: Thread "org.apache.accumulo.server.test.TestIngest" died null java.lang.reflect.InvocationTargetException at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.accumulo.start.Main$1.run(Main.java:90) at java.lang.Thread.run(Thread.java:662) Caused by: java.lang.RuntimeException: org.apache.commons.cli.UnrecognizedOptionException: Unrecognized option: --debug at org.apache.accumulo.server.test.TestIngest.parseArgs(TestIngest.java:165) at org.apache.accumulo.server.test.TestIngest.main(TestIngest.java:256) ... 6 more Caused by: org.apache.commons.cli.UnrecognizedOptionException: Unrecognized option: --debug at org.apache.commons.cli.Parser.processOption(Parser.java:363) at org.apache.commons.cli.Parser.parse(Parser.java:199) at org.apache.commons.cli.Parser.parse(Parser.java:85) at org.apache.accumulo.server.test.TestIngest.parseArgs(TestIngest.java:162) ... 7 more {code}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-745
17701	ACCUMULO-739	Since there seems to be a difference of opinion, can the visibility label syntax be pluggable?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-739
17702	ACCUMULO-739	Bringing this back around... There's a specific need on the part of a few users to add . and / in order to support hierarchical structure in labels. Can we settle on implementing this feature using the current whitelist approach and open a new ticket to discuss the long term strategy?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-739
17703	ACCUMULO-736	Given that all data retrieval from Accumulo is returned to the client via a Java Iterator, I'm a little confused about the application of such a feature. If you have many columns in one row, you'll get a Key-Value pair for each Row-Column at which point it's an exercise for the client not to consume more than can fit in its memory. Other applications seem to be solvable through a custom Accumulo Iterator. For example, if you want to fetch all columns for a row in one call in batches of size N, this is easily done with an Accumulo Iterator without any core changes. Please assume I'm an HBase API novice. Can you provide some more concrete details? Thanks.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-736
17704	ACCUMULO-732	Have you seen the output from 'hexdump -C'? Would that format also be useful?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-732
17705	ACCUMULO-730	Question: do we need to maintain compatibility with hadoop 0.20.205.0? I'm about to check in a patch for ACCUMULO-730. I decided to use the plain local job tracker with MockAccumulo (instead of MiniAccumuloCluster) since it's a bit more straightforward and we haven't planned to deprecate MockAccumulo yet. Either way involves adding a couple of dependencies to the poms to get it to work for all of 1.0, 1.1, 0.23, and 2.0. The only one added to the top-level pom is commons-httpclient, which seems fine. To get it to work for 0.20.205.0 as well, we would have to add a dependency on jackson. I'm reluctant to do so, but if someone needs 0.20.205.0, I will.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-730
17706	ACCUMULO-730	Are the added dependencies only needed for tests? If so, they should probably just be in the test scope, and not the default scope. If they are needed beyond unit tests, then disregard.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-730
17707	ACCUMULO-722	David, I browsed the CFS site. The first question on that page asks about the tricky part where there's only one writer to a file. This is a key problem, since having a single writer to a write-ahead-log is key for proper failure conditions. In particular, we must ensure that the writes by a failing tablet server to its write-ahead log are denied while we use that file for recovery. The response "HDFS does not implement posix semantics" is true, but it understates the importance of this feature to the HBase and/or Accumulo WAL. And, it appears this is not an open-source solution, so I'm unable to test it without committing additional resources. Does anyone know if they support exclusive writer semantics?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-722
17708	ACCUMULO-722	among recent changes in HDFS: * HA NameNode (HDFS-3077) * Pluggable locality / failure tolerances (HDFS-385, HADOOP-8468) * Improving scalability of namespace operations (HDFS-5389) and block location mapping (HDFS-5711) Could we please forgo this ticket in favor of interested parties working on relevant bits within HDFS?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-722
17709	ACCUMULO-720	Is this the file? core/src/main/java/org/apache/accumulo/core/iterators/user/RegExFilter.java	0	0	https://issues.apache.org/jira/browse/ACCUMULO-720
17710	ACCUMULO-718	This is done. Does anyone think we should keep the old static methods around?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-718
17711	ACCUMULO-708	Cool stuff, Dave. I assume the usage of the HDFS path "/accumulo" is based on the default value for the "instance.dfs.dir" parameter? Also, have you thought about modifying the assemble profile to load lib correctly and then modify `accumulo init` to auto-load the default jars into HDFS? In other words, have you thought about automating the above process through the maven profile and Accumulo initialization given the "classloader.vfs.enabled" and "classloader.vfs.context.classpath.system" parameters in accumulo-site.xml?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17712	ACCUMULO-708	I scanned through the patch trying to get a sense of it at the high level. Tomorrow I will do a more in depth review. Have a few comments so far. * use general prefix for new config properties. general denotes its a configuration for all services (i.e. master, tserver, etc) * needs documentation. How does a user know this feature exists and what they can do with it? * probably should not use org/apache/commons/vfs2 package. If the commons project defines classes with same name and user loads Accumulo and Commons jar chaos could ensue. * I noticed there were a few accumulo site files introduced. Are these for testing? Some of them configured walog dirs which is no longer nesc in 1.5.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708

#	Requirement	Questions	Answered	Successfully Answered	URL
17713	ACCUMULO-708	bq. •use general prefix for new config properties. general denotes its a configuration for all services (i.e. master, tserver, etc) I'll include in another patch, I'll wait for your other comments and roll all the changes up into a new patch. bq. •needs documentation. How does a user know this feature exists and what they can do with it? Where should this be documented? User guide, readme, etc? bq. •probably should not use org/apache/commons/vfs2 package. If the commons project defines classes with same name and user loads Accumulo and Commons jar chaos could ensue. Understood. I know for a fact that these classes don't collide with ones in Commons VFS. Code has also been submitted to VFS for inclusion in a future release (VFS-442). bq. •I noticed there were a few accumulo site files introduced. Are these for testing? Some of them configured walog dirs which is no longer nesc in 1.5. These are there for tests. I believe I copied the 512MB example site file and modified for my tests. The walog property is not used and the test will not fail if the directory does not exist as the servers are not started during the tests.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17714	ACCUMULO-708	I applied the patch and have been unable to get the unit test to run. Keep seeing something like the following {noformat} Caused by: java.lang.NullPointerException at org.apache.hadoop.hdfs.MiniDFSCluster.startDataNodes(MiniDFSCluster.java:422) at org.apache.hadoop.hdfs.MiniDFSCluster.<init>(MiniDFSCluster.java:280) at org.apache.accumulo.test.AccumuloDFSBase.<clinit>(AccumuloDFSBase.java:66) ... 21 more {noformat} bq. Where should this be documented? User guide, readme, etc? Maybe add something to the users guide and add an example that shows off what can be done.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17715	ACCUMULO-708	Keith, are you getting other errors about directory permissions? see: http://search-hadoop.com/m/5Te0lu8g0c1/hbase+windows/v=plain	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17716	ACCUMULO-708	Still do not have the unit test running. It may be a configuration issue on my workstation, but its not the umask issue, still investigating. A few more comments on the code. * In AccumuloVFSClassLoader and AccumuloClassLoader the methods replaceEnvVars(), addUrl(), findAccumuloURLs(), etc. look very similar. Are these methods just copies? If not, what is the essence of the diffs? * Seems like unit test AccumuloContextClassLoaderTest would benefit from using two seprate jars. The same jar is loaded into two different context. Seems like it would be better to load two different jars with different classes. Then verify that each context contains only the expected classes and nothing else. * AccumuloReloadingVFSClassLoaderTest would also benefit from another jar. Test are needed for a new class coming into existence, a new version of an existing class coming into existence, and an existing class going out of existence. Moving forward on this we need to start thinking about having a table prop for configuring per table classloaders. If you would like help I can work on that after we get this checked in, just let me know. I was thinking of a scenario where we have the following two classloader contexts : * system class loader context : contains accumulo server and core jars and all deps in HDFS * table foo class loader context : contains user iterator jars in HDFS It seems like the two context above are siblings. Is this correct? Therefore the foo class loader will not have accumulo-core jars in itself or its parent class loaders. Will this cause problems when loading an iterator from the foo classloader?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17717	ACCUMULO-708	bq. I did run into the same not-helpful Exception Keith posted above and traced it back to the same issue Dave linked to. After I realized AccumuloDFSBase.java:25 was sending the logging to /dev/null I think my umask is correct {noformat} $ umask 0022 {noformat} AccumuloDFSBase.java line 25 is an import statement. Did you mean line 40 or line 49?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17718	ACCUMULO-708	bq. AccumuloDFSBase.java line 25 is an import statement. Did you mean line 40 or line 49? Uhh, my code still says line 25. The actual code I commented out was the following: {noformat} //Turn off the Jetty and MiniDFSCluster logging static { System.setProperty("org.apache.commons.logging.Log", "org.apache.commons.logging.impl.NoOpLog"); } {noformat} If your umask is 022, then you need to change the code from {noformat} conf.set("dfs.datanode.data.dir.perm", "775"); {noformat} to {noformat} conf.set("dfs.datanode.data.dir.perm", "755"); {noformat}	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17719	ACCUMULO-708	Josh, Please back out your change of setting dfs.datanode.data.dir.perm to 755. I believe this issue is fixed in HDFS-1560. Instead, please change the hadoop-test dependency in start/pom.xml to 1.1.0. I talked with Eric about changing the hadoop dependency from 0.20.205 to 1.1.0 for Accumulo 1.5. I talked with Keith yesterday and I believe he is past the directory permission issue, he is now running into another issue that has to do with the configuration of his computer. Keith, bq. •In AccumuloVFSClassLoader and AccumuloClassLoader the methods replaceEnvVars(), addUrl(), findAccumuloURLs(), etc. look very similar. Are these methods just copies? If not, what is the essence of the diffs? If I remember correctly, they are copies of methods in AccumuloClassLoader. I did not want to extend AccumuloClassLoader as I do not know its lifetime. bq. •Seems like unit test AccumuloContextClassLoaderTest would benefit from using two seprate jars. The same jar is loaded into two different context. Seems like it would be better to load two different jars with different classes. Then verify that each context contains only the expected classes and nothing else. I was testing that the same jar/class was in fact being loaded by two different classloaders, and were therefore different classes. I will add your test suggestions to my to-do list. All, I believe that the discussion around which contexts are supported is a separate discussion. I can think of several different scenarios in which contexts could be defined: per user, per table, per scan session, per minc/majc session, per application, etc. Regarding the heirarchy of the context classloader, I can see Keiths point above, and would also counter with a different scenario. Say for example that I wanted to run different versions (api compatible of course) of the Accumulo iterators on my table. The current setup would allow me to run a 1.5.0 tablet server and 1.5.1 iterators on my table. I can do this because the context classloader for my application is not a child of the system classloader. Changing the context classloader to be a child of the system classloader is a very small change.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17720	ACCUMULO-708	does it make sense to move the two open subtasks for this feature to the next release (not 1.5.0) and close this issue. The feature is in the 1.5.0 codebase. The two open tickets are not critical to 1.5.0.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-708
17721	ACCUMULO-706	This example sets the maxLatency to 2 minutes, but the code committed against this ticket sets the maxLatency to long.MAX_VALUE. I think we should have a default maxLatency of closer to 1 second or 1 minute. Otherwise people are going to have writers that sit around in long running processes and have data in their un-flushed buffers forever (I've seen this a couple of times), which is a source of confusion. Is there a counter-argument for making the default MAX_VALUE? Is there a performance problem with keeping it in the 1-100 second range?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-706

#	Requirement	Questions	Answered	Successfully Answered	URL
17722	ACCUMULO-706	bq. Is there a counter-argument for making the default MAX_VALUE? Is there a performance problem with keeping it in the 1-100 second range? No, not that I can think of. Not sure why I did that. Maybe to make it consistent with setTimeout(). I think 1 second would be too low, 1 or 2 minutes would be probably be an ok default. Seems like one second would cause small batches to be sent, which would be slower. I suppose if writing small batches were slower, then larger batches would build up. But it may oscillate. Seems like that could possibly cause unnecessary thrashing. Anyway, I think 2 minutes would be good. There was not a default previously, the user always had to specify something. If changed, would also need to update javadocs	1	0	https://issues.apache.org/jira/browse/ACCUMULO-706
17723	ACCUMULO-704	I think I may be missing something here... but I'll ask: Why does the build.sh script call pdflatex on the Accumulo user manual 4 times? http://svn.apache.org/viewvc/accumulo/trunk/docs/src/user_manual/build.sh?view=markup	0	0	https://issues.apache.org/jira/browse/ACCUMULO-704
17724	ACCUMULO-704	For review to see what people think: I moved the building of the user-manual from normal execution and placed it in its own profile that can be activated. I also updated the README. Thoughts? Does someone have recommendations on a more correct/better way to do this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-704
17725	ACCUMULO-704	Why can't the user-manual's build.sh file be run manually? Does it need to be invoked by maven?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-704
17726	ACCUMULO-703	Could you add some info to help about this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-703
17727	ACCUMULO-703	can you update docs/examples/README.bloom to use the new command... the example currently uses PrintInfo... should something be added to the user manual?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-703
17728	ACCUMULO-700	If we configure iterators using, say, comma-separated base64-encoded representations of IteratorSettings, a potential disadvantage is that the iterator information will no longer be readable in the job configuration. Does anyone need this to remain readable?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-700
17729	ACCUMULO-700	Will there be a tool that can convert the representations into human-readable form?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-700
17730	ACCUMULO-697	Does this work with client side iterators only? The top level iterator will usually be running server side, so you'll need to have an interface to set a serializer for some object at the top of the stack (and possibly another deserializer if you have custom iterators beneath it) and set a deserializer at the client level. At some point on the server side, the iterator will have to transform the user type into a {{byte[]}}, which is basically what a {{Value}} is. On the client side, the user will still have to supply code to deserialize the {{byte[]}}, which can already be accomplished with no code changes by using a library such as Google Collections/Guava or Apache Commons.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-697
17731	ACCUMULO-697	I like the concept, but does this go far enough? If Values aren't special, then are Keys special, and if so then why? Should we make our SortedKeyValueIterator implement Iterable<? extends Object> ? Then the bottom level iterator (RFile reader) would include KeyValue or Entry<Key,Value> objects, the top level iterator for scans would have to have objects that are serializable, and the top level iterator for compactions would have to implement Iterable<Entry<Key,Value>>. One of the problems we have with iterators now is that the Key and Value are accessed with separate methods, even though they're always read off of disk together. Splitting up the Key and Value on the server side is sort of arbitrary and could reduce our ability to parallelize iterators (if we ever decide that's something we want to do). Another problem is that SortedKeyValueIterator falls somewhere in between Java's Iterator and Iterable interfaces. SortedKeyValueIterator holds onto filters, aggregation parameters, etc. that make it act like a collection, and it keeps a pointer to somewhere in that collection like an Iterator. I think we should change SortedKeyValueIterator into more like an immutable collection, or a consistent, isolated, unchanging view of the data, and have it implement Iterable. That might open up opportunities for automating optimization of queries on the server side, or better support for built-in iterator tree definition languages.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-697
17732	ACCUMULO-697	I don't know why this is something we want to integrate into our API. A utility can easily be written to do whatever transformations you want on top of the scanner's iterator, so why push this down into our api? I think a change like this could only serve to make the api harder to use for little to no real gain.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-697
17733	ACCUMULO-697	What is K and V? Where are they defined? I know they mean key and value but I'd like to grok the fundamentals. Sorry for this basic question, and pardon me for asking it here, but I don't know what kind of web search would give me an answer.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-697
17734	ACCUMULO-697	bq. What is K and V? Where are they defined? I know they mean key and value but I'd like to grok the fundamentals. Sorry for this basic question, and pardon me for asking it here, but I don't know what kind of web search would give me an answer. http://en.wikipedia.org/wiki/Generics_in_Java	1	1	https://issues.apache.org/jira/browse/ACCUMULO-697
17735	ACCUMULO-693	[~medined], which version did you see this problem with? Can you provide more information on your environment?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-693
17736	ACCUMULO-677	Why should they depend on a different ACL? Grant/Revoke was intended to be an "ALTER USER" ACL, whereas Create User was intended to be a "CREATE USER" ACL, and this would *include* creating the initial authorizations. When you view separate it as "CREATE" and "ALTER" on the object "USER", it makes complete sense in an object oriented way. Separating them makes less sense, because it treats "CREATE USER" and "ALTER USER" as two completely independent actions, completely ignoring the common object you are manipulating ("USER"). If you implemented the above, then to create a fully functioning user, you'd have to have two separate permissions. I understand the desire to change the API to match this paradigm, if you were to desire to switch to it, but I personally think that leaving the "CREATE USER" and "ALTER USER" paradigm in place is better. That said... without deprecating or changing the "CREATE"/"ALTER" paradigm, you could add to the API a method to create a user without authorizations (unless that already exists).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-677
17737	ACCUMULO-671	Is this easier now that we have a root table?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-671
17738	ACCUMULO-670	Can you clarify how you think this code should behave?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-670
17739	ACCUMULO-670	[~kturner] - I believe I addressed this. Can you confirm before I close this ticket again?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-670
17740	ACCUMULO-652	This is still open. What are your plans for merging/closing this for 1.7.0?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-652
17741	ACCUMULO-650	So far I have added a desc command to the shell to display the descriptions, added a flag to tables to display the tables and the descriptions, added a flag to createtable to set the description. the desc command will either display the description for the table or it will update it using a flag. EX: {noformat} user@instance table> desc table [description of table] user@instance table> desc -u 'new description' table [new description] {noformat} Should I add the ability to set the description when importing a table or a directory? Also, should I clone the description of the table when the table is cloned?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-650
17742	ACCUMULO-650	Does a cloned table have any provenance? Maybe a cloned table should have "Cloned from XXXX at YYYY-MM-DD HH:MM:SS" appended to the end of the description? I haven't thought about this deeply. Is there any UUID associated with a table? On the other hand, since cloned tables are mutable, does it matter?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-650

#	Requirement	Questions	Answered	Successfully Answered	URL
17743	ACCUMULO-650	Since we're thinking. I've used clone in the following manner. I import a lot of stuff into one table. Then clone it so that other people can work with that data. Then I continue to work with the original table. My changes don't affect the other people. Dumb question time ... can clones be used like git for data? Can a clone be thought of as a tag or branch?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-650
17744	ACCUMULO-650	Kevin, can you add some unit tests?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-650
17745	ACCUMULO-650	I think the per table config route would be best. Ultimately its still sored in zookeeper, just in a more systematic way. This would involve adding an enum to o.a.a.core.conf.Property could still have API methods in table operation to set and get desc. These would be similar to methods for manipulating locality groups and iterator settings, which are also per table configs. Also, if you are still thinking of adding something to create table, I still think we should add a NewTableConfig method. Would still have an option on tables command to list descriptions. Would we still need the desc command? I do not really like the name, but I am uncertain about the functionality. As far as functionality, it may make things slightly simpler for the user rather than using the config command. Also having a specific shell command for setting it puts a layer of indirection between the implementation and the user.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-650
17746	ACCUMULO-650	Do we want to include this specific specialized per-table property in addition to the user-defined properties added in ACCUMULO-2841? I think 2841 pretty much satisfies all the use-cases. Any objections to closing?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-650
17747	ACCUMULO-638	Can this ticket be pushed to v1.6?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-638
17748	ACCUMULO-638	Is it OK if I assign this to myself? Looks like there hasn't been much activity since it was created. I'd be very grateful for an extra pair of eyes on [my work|https://github.com/rweeks/accumulo/tree/ACCUMULO-638], if anybody has some spare cycles. I guess it's complicated by the fact that most of the code changes will be in the Titan codebase, but there are a [few iterators|https://github.com/rweeks/accumulo/blob/ACCUMULO-638/core/src/main/java/org/apache/accumulo/core/iterators/user/SeekingFilter.java] that I think would be generally useful in Accumulo.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-638
17749	ACCUMULO-638	All yours Russ! Would you like me to set a target version? Or should we just go for whatever it lands on?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-638
17750	ACCUMULO-638	I'll try to make some time to give your changes a look. Would comments on Github be best?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-638
17751	ACCUMULO-626	Github user madrob commented on a diff in the pull request: https://github.com/apache/accumulo/pull/50#discussion_r44291578 --- Diff: iterator-test-harness/pom.xml --- @@ -0,0 +1,56 @@ +<?xml version="1.0" encoding="UTF-8"?> +<!-- + Licensed to the Apache Software Foundation (ASF) under one or more + contributor license agreements. See the NOTICE file distributed with + this work for additional information regarding copyright ownership. + The ASF licenses this file to You under the Apache License, Version 2.0 + (the "License"); you may not use this file except in compliance with + the License. You may obtain a copy of the License at + + http://www.apache.org/licenses/LICENSE-2.0 + + Unless required by applicable law or agreed to in writing, software + distributed under the License is distributed on an "AS IS" BASIS, + WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + See the License for the specific language governing permissions and + limitations under the License. +--> +<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd"> + <modelVersion>4.0.0</modelVersion> + <parent> + <groupId>org.apache.accumulo</groupId> + <artifactId>accumulo-project</artifactId> + <version>1.8.0-SNAPSHOT</version> + </parent> + <artifactId>accumulo-iterator-test-harness</artifactId> + <name>Apache Accumulo Iterator Test Harness</name> + <description>A library for testing Apache Accumulo Iterators.</description> + <dependencies> + <!-- TODO Don't force downstream users to have JUnit --> + <dependency> + <groupId>junit</groupId> + <artifactId>junit</artifactId> + <!--optional>true</optional--> + </dependency> + <dependency> + <groupId>log4j</groupId> + <artifactId>log4j</artifactId> + </dependency> + <dependency> + <groupId>org.apache.accumulo</groupId> + <artifactId>accumulo-core</artifactId> + </dependency> + <dependency> + <groupId>org.apache.hadoop</groupId> + <artifactId>hadoop-client</artifactId> + </dependency> + <dependency> + <groupId>org.reflections</groupId> + <artifactId>reflections</artifactId> + </dependency> --- End diff -- Do we need to add this to the NOTICE?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-626
17752	ACCUMULO-626	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/50#issuecomment-155167375 For iterators that can be run at major compaction, how can we exercise delete keys?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-626
17753	ACCUMULO-626	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/50#issuecomment-155168971 > For iterators that can be run at major compaction, how can we exercise delete keys? Very good question! The completely mocked-out IteratorEnvironment is a sign that I haven't thought past the scan case. I'm not even sure if the ColumnFiltering iterator with SortedMapIterator is even a 100% valid in-memory reproduction of how it works in a TabletServer. I had left input entirely up to the user, so I'm not sure of a simple way to test deletions without potentially messing up some expected table structure. There are some tricky edge cases at compaction that I haven't yet considered how to encapsulate. Combiners on non-full compactions are another good example.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-626
17754	ACCUMULO-626	Github user joshelser commented on the pull request: https://github.com/apache/accumulo/pull/50#issuecomment-163673419 Drat, another thing I just thought of: this is meant to be user facing, and I think it should eventually hit public API. It might be premature to put it there immediately (since this hasn't been vetted by anyone but myself), but does anyone have thoughts on a path to "graduate" into public API status?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-626

#	Requirement	Questions	Answered	Successfully Answered	URL
17755	ACCUMULO-626	Github user joshelser commented on a diff in the pull request: https://github.com/apache/accumulo/pull/50#discussion_r47255171 --- Diff: iterator-test-harness/src/main/java/org/apache/accumulo/iteratortest/testcases/DeepCopyTestCase.java --- @@ -0,0 +1,63 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to you under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.iteratortest.testcases; + +import java.io.IOException; +import java.util.Collections; +import java.util.TreeMap; + +import org.apache.accumulo.core.data.ByteSequence; +import org.apache.accumulo.core.data.Key; +import org.apache.accumulo.core.data.Value; +import org.apache.accumulo.core.iterators.SortedKeyValueIterator; +import org.apache.accumulo.iteratortest.IteratorTestInput; +import org.apache.accumulo.iteratortest.IteratorTestOutput; +import org.apache.accumulo.iteratortest.IteratorTestUtil; +import org.apache.accumulo.iteratortest.environments.SimpleIteratorEnvironment; + +/** + * Test case that verifies that an iterator can use the generated instance from {@code deepCopy}. + */ +public class DeepCopyTestCase extends OutputVerifyingTestCase { + @Override + public IteratorTestOutput test(IteratorTestInput testInput) { + final SortedKeyValueIterator<Key,Value> skvi = IteratorTestUtil.instantiateIterator(testInput); + final SortedKeyValueIterator<Key,Value> source = IteratorTestUtil.createSource(testInput); + try { + skvi.init(source, testInput.getIteratorOptions(), new SimpleIteratorEnvironment()); + SortedKeyValueIterator<Key,Value> copy = skvi.deepCopy(new SimpleIteratorEnvironment()); --- End diff -- Thanks for taking a look, Dylan! > call deepCopy in the middle of an iteration, in addition to the beginning after init() but before any next() calls Yeah, there are lots of good tests we can do around deepCopy. I'll see if I can add one for this specific case. Please feel free to send a PR if you definitely want to see this in the initial changeset :) > Some iterators may do this in order to "save" the iteration at a certain point. But that's not something the "system" does, is it? I can see the value, but I'm not sure if that's something that would happen through "normal" Accumulo use.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-626
17756	ACCUMULO-626	Github user dhutchis commented on a diff in the pull request: https://github.com/apache/accumulo/pull/50#discussion_r47256202 --- Diff: iterator-test-harness/src/main/java/org/apache/accumulo/iteratortest/testcases/DeepCopyTestCase.java --- @@ -0,0 +1,63 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to you under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.accumulo.iteratortest.testcases; + +import java.io.IOException; +import java.util.Collections; +import java.util.TreeMap; + +import org.apache.accumulo.core.data.ByteSequence; +import org.apache.accumulo.core.data.Key; +import org.apache.accumulo.core.data.Value; +import org.apache.accumulo.core.iterators.SortedKeyValueIterator; +import org.apache.accumulo.iteratortest.IteratorTestInput; +import org.apache.accumulo.iteratortest.IteratorTestOutput; +import org.apache.accumulo.iteratortest.IteratorTestUtil; +import org.apache.accumulo.iteratortest.environments.SimpleIteratorEnvironment; + +/** + * Test case that verifies that an iterator can use the generated instance from {@code deepCopy}. + */ +public class DeepCopyTestCase extends OutputVerifyingTestCase { + @Override + public IteratorTestOutput test(IteratorTestInput testInput) { + final SortedKeyValueIterator<Key,Value> skvi = IteratorTestUtil.instantiateIterator(testInput); + final SortedKeyValueIterator<Key,Value> source = IteratorTestUtil.createSource(testInput); + try { + skvi.init(source, testInput.getIteratorOptions(), new SimpleIteratorEnvironment()); + SortedKeyValueIterator<Key,Value> copy = skvi.deepCopy(new SimpleIteratorEnvironment()); --- End diff -- Ah I see Josh-- is the purpose to test only system iterators, as opposed to generic tests to check user-defined iterators? That makes more sense. Some users may define iterators that expect a column family, whereas all the tests seek with `Collections.<ByteSequence> emptySet(), false`. There may be many other places where a user does something "non-normal" intentionally. If we make this public API, we should state that user iterators do not necessarily need to pass all the IteratorTestCases, or at least that failing an IteratorTestCase does not imply a bug.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-626
17757	ACCUMULO-626	Github user keith-turner commented on the pull request: https://github.com/apache/accumulo/pull/50#issuecomment-163715371 > but does anyone have thoughts on a path to "graduate" into public API status? Well currently iterators are not declared as public API.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-626
17758	ACCUMULO-625	There has not, and it is not slated for 1.5. Do you have thoughts on what you would like to see?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-625
17759	ACCUMULO-625	Currently there's a UniqueColumn map/reduce example (org.apache.accumulo.examples.simple.mapreduce.UniqueColumns). Ideally I would like to take that example and turn it into an iterator. Due to the tear up and tear down factor, is it safe to assume i can only perform this operation per row bases? So the iterator would find all the distinct column in a single row and the client would have to perform a final distinct for all the rows?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-625
17760	ACCUMULO-622	How do we determine where to make changes? When will we know all changes are complete?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-622
17761	ACCUMULO-618	Does this matter? We also now have "instamo", an unbranded module like cloudtrace. Should we also brand that artifact as "accumulo-instamo" and cloudtrace as "accumulo-trace"? All these are pretty trivial changes, but I can think of at least one benefit: searching in Maven repositories indicates that the artifacts are associated with the Accumulo software stack (vs. simply being produced by the Accumulo group, which is indicated by the groupId).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-618
17762	ACCUMULO-618	bq. Should we also brand that artifact as "accumulo-instamo" and cloudtrace as "accumulo-trace"? I think accumulo-trace is good. bq. searching in Maven repositories indicates that the artifacts are associated with the Accumulo software stack Hopefully instamo does not end up any maven repos. Its just meant to serve as an example accumulo maven project.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-618
17763	ACCUMULO-617	That should be the default behavior? Is this not the case and/or did it change in 1.5?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-617

#	Requirement	Questions	Answered	Successfully Answered	URL
17764	ACCUMULO-617	Will ACCUMULO-1479 satisfy this ticket? Can I close this as a duplicate?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-617
17765	ACCUMULO-600	Agree that its complex. Have you considered moving it out of the build entirely so that you don't have to keep it up to date. By moving it out (to GitHub as an example) it could be maintained on its own schedule.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-600
17766	ACCUMULO-600	Does anyone know if moving WikiSearch to Github would preclude us from including it as part of future releases? I know there's ASF policy saying that code for ASF projects must be kept in ASF's SVN repository -- does that extend to contrib parts of a release?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-600
17767	ACCUMULO-600	Hi Adam, and Keith: any reason to move it to Github if it's currently within the ASF Accumulo project? I guess I'm not seeing the logic about removing it in general? Won't that splinter the development?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-600
17768	ACCUMULO-600	Hi Josh: {quote} I think the biggest issue has been that most of us who have been a part of the wikisearch development are not currently Accumulo developers (in other words, without svn commit access). {quote} Sorry to hear that -- how much contribution would you say you've been doing? {quote} Correct me if I'm wrong, anyone, but I believe Github was suggested as an alternative to the ASF SVN if it was decided that those who will be maintaing the wikisearch example will stay disjoint from those who are Accumulo committers. {quote} Do you have any context on this? Like an email thread on list where this was discussed?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-600
17769	ACCUMULO-600	{quote} Sorry to hear that – how much contribution would you say you've been doing? {quote} The few patches I've made have been coordinated with Eric Newton. Some on jira, some off. {quote} Do you have any context on this? Like an email thread on list where this was discussed? {quote} No, most of this was discussed offline. I'd encourage those who fall into this category to chime back in with their opinion rather than me asserting an opinion for them. :)	1	0	https://issues.apache.org/jira/browse/ACCUMULO-600
17770	ACCUMULO-600	Billie, since you did the work to pull out of trunk and into contrib, can we close this? We can make new ticket(s) to address the TODOs you left.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-600
17771	ACCUMULO-588	David, the descriptions look good. Did you miss mapred? Also, I have a script that copies these READMEs into the Accumulo site and treats them as markdown (see http://accumulo.apache.org/1.4/examples/). I was able to modify this script to place the links to the individual examples next to your descriptions, so it will look good when we publish the 1.5 example.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-588
17772	ACCUMULO-588	Thanks for the fix-up. I'm fairly sure I did the mapred writeup. However, it probably disappeared when I reformatted. http://accumulo.apache.org/1.4/examples/ does not show the brief blurbs about each example. Would it be OK to add a .md extension to the READMEs file to indicate that markdown syntax can be used? Your script might need updating. Would that change affect anything else?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-588
17773	ACCUMULO-587	I finally got a chance to look into this. I don't think anything special needs to be done for {{MultiTableBatchWriterImpl}} other than adding the {{finalize()}}. The object that actually needs to be closed (the {{TabletServerBatchWriter}}) is private and not shared. The {{BatchWriter}} implementation returned by the {{getBatchWriter(String table)}} method is a non-static inner class and therefore all references to any {{BatchWriter}} created by the {{MultiTableBatchWriterImpl}} must be lost before {{finalize()}} could be called. The only wrinkle with {{MultiTableBatchWriterImpl}} is the {{close()}} method may throw a {{MutationsRejectedException}}, so that must be handled in the {{finalize()}} method. I'm not sure if you'd want to turn that into a logged 'non-exception' or a {{RuntimeException}}; I could see the argument for either way (although I'd lean toward the {{RuntimeException}}). Regardless of the choice for handling the {{MutationsRejectedException}}, the memory/thread leak would be avoided because {{TabletServerBatchWriter}}'s {{close()}} method closes the thread pool in a finally block even if it throws the {{MutationsRejectedException}}. So I think just adding a method like this to {{MultiTableBatchWriterImpl}} provides the needed check: {code} @Override protected void finalize() { if (!closed) { log.warn(MultiTableBatchWriterImpl.class.getSimpleName() + " not shutdown; did you forget to call close()?"); try { close(); } catch (MutationsRejectedException mre) { log.error(MultiTableBatchWriterImpl.class.getSimpleName() + " internal error.", mre); // ... AND/OR ... throw new RuntimeException("Exception when closing " + MultiTableBatchWriterImpl.class.getSimpleName(), mre); } } } {code}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-587
17774	ACCUMULO-581	It looks like the following commands are covered at least in some fashion: * authenticate * config * createtable * createuser * debug * delete * deleteiter * deletetable * dropuser * flush * getauths * grant * help * insert * passwd * revoke * scan * setauths * setiter * table * tablepermissions * tables * user * userpermissions * users * whoami Which leaves these: * -?- * -about- * -addauths- * -addsplits- * -bye- * -classpath- * -clear- * -clonetable- * -cls- * -compact- * -constraint- * -deletemany- * -deleterows- * -deletescaniter- * -deleteshellter- * -deleteuser- * -droptable- * -du- * -egrep- * -execfile- * -exit- * -exporttable- * -formatter- * -getgroups- * -getsplits- * -grep- * -help- * -history- * -importdirectory- * -importtable- * -info- * -interpreter- * -listcompactions- * -listiter- * -listscans- * -listshelliter- * -maxrow- * -merge- * -notable- * -offline- * -online- * -ping- * -quit- * -renametable- * -setgroups- * -setscaniter- * -setshelliter- * -sleep- * -systempermissions- * -trace-	1	1	https://issues.apache.org/jira/browse/ACCUMULO-581
17775	ACCUMULO-580	Is the suggestion to add a new method to the BatchScanner interface to set the batch size?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-580
17776	ACCUMULO-569	Does an accumulo shell exposed through the monitor (ACCUMULO-196) accomplish this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-569
17777	ACCUMULO-564	What is the status of hadoop 0.23 relative to hadoop 1.1 and hadoop 2? It seems like it's the worst of both worlds: alphaness of 2 without the HA stuff that's in the beta (1.1). Perhaps it's reasonable to say no hadoop 0.23 support?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-564
17778	ACCUMULO-550	I noticed yesterday when working on ACCUMULO-652 that the index entries at level 1 and higher are still interspersed between level 0 blocks with the current technique. Is there value in keeping indexes at a given level greater than 0 close to each other, or is that overkill?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-550
17779	ACCUMULO-532	Can this be closed or resolved?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-532
17780	ACCUMULO-532	I'm looking at this a bit closer, and was going to update this contrib to use Hama 0.6.0 release (latest release as of this writing) in the process of implementing ACCUMULO-769. However, I noticed a problem. Is there a reason BSPJob does not extend Job, and BSPJobContext does not extend JobContext? It seems it would be very convenient if it did.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-532
17781	ACCUMULO-527	The indentation of a few statements in shell.setInstance is incorrect, making it hard to read. Should the code error if multiple options are given? For example if -z, -zi, and -zh are given. I suppose the existing code did not error if -z and --fake were specified. I tried running the new unit test w/ eclemma and it showed no coverage in Shell... I wonder if this is because powermock.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-527
17782	ACCUMULO-527	[~bhavanki] I just pushed your commits and a commit of mine where I changed the unit test. Can you take a look at that change? Thats the first time I have used PowerMock, its interesting, seems to tightly couple the test to the implementation rather than the expected behavior. Does anyone have any issue w/ adding the dependencies? I reasoned that since they were scoped at the test level there was no need to seek consensus.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-527

#	Requirement	Questions	Answered	Successfully Answered	URL
17783	ACCUMULO-527	[~bhavanki], is this issue resolved?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-527
17784	ACCUMULO-527	It is. Keith had added some updates on top of mine and they are fine with me. I'm satisfied. [~kturner], are you good with it?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-527
17785	ACCUMULO-511	Before this was added listscans already suffered from the issue of too many columns. I am thinking we need to do something to help users select columns. I suppose user cans pipe the output to awk or cut. Maybe thats good enough? Are there other options to consider beside the following? * add options to list scans for selecting columns to print * do nothing and let users use awk or cut	0	0	https://issues.apache.org/jira/browse/ACCUMULO-511
17786	ACCUMULO-511	I see the words "too many columns" above, and I can't but think that perhaps we should just treat data from listscans as though it were results from a table scan. We already know how to deal with formatting output, filtering columns, etc. from any other real table. The problem is that the "listscans" isn't a real table... but it would be trivial to treat it as one, and get all the features that come with tables as a result. Why shouldn't we be able to get a virtual table view of metadata like this?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-511
17787	ACCUMULO-508	Doesn't the AccumuloInputFormat already do this? I'm pretty sure there's an autoAdjust feature that merges overlapping ranges, splits ranges on tablet boundaries, and then assigns them. This is done by default. If this feature is turned off, the ranges are given to mappers exactly as they were given to the job: 1 range per mapper.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-508
17788	ACCUMULO-507	Good point. Do you think we should make this change in 1.4 branch?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-507
17789	ACCUMULO-507	See ACCUMULO-826. Rolled back changes that created a file 1.4 : r1400976 1.5. : r1401004. What should we do for 1.5?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-507
17790	ACCUMULO-507	I'm tempted not to fix this since it seems impossible to satisfy both our requirements: not leaving large files around and not killing the job when the process exits. Is it too onerous to adjust the mapred.user.jobconf.limit property to allow for jobs with large numbers of ranges?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-507
17791	ACCUMULO-507	Can you clear the ranges from the conf after the splits are generated? Not sure how it's done now, but individual tasks should be able to reconstruct what they need from the split metadata, and then there's no reason to pass around all the ranges to each mapper.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-507
17792	ACCUMULO-507	I think this fix was reverted for the wrong reason. ACCUMULO-826 found that a MapReduce job would fail if the process that started the job was killed. This was an issue because we were writing the user's password to a file that was being deleted on exit. Whenever a new map task is kicked off it needs to read the password, so it was trying to read a nonexistent file. But the ranges don't need to be read by each map task, they only need to be accessed once when getSplits is called, which happens before the job is actually submitted. Thus it shouldn't matter if the file containing the ranges is deleted in the middle of a job -- if the process exits before the job is actually submitted, the job will fail, but that seems OK to me. The other issue pointed out in ACCUMULO-826 is valid, that the file was being written to the file system, added to the distributed cache, then read directly from the file system. The ranges file shouldn't have been added to the distributed cache at all, since it's not needed by the slave nodes. However, there may be little point in re-applying this fix if the mapred.user.jobconf.limit applies to the whole job submit directory. Using the ranges file method might effectively halve the size of the job submit directory, but you could still hit the limit if you had enough ranges. I guess I'll try to verify this is true. Does anyone have opinions about this issue?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-507
17793	ACCUMULO-507	bq. Does anyone have opinions about this issue? I do not have a problem with reverting the revert if the job keeps running. What about random files being left around? Was the file place somewhere that M/R will eventually clean up if the process that created it does not?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-507
17794	ACCUMULO-507	Happily, I have verified that my initial impression was incorrect: the size limit is on the job configuration file only. bq. What about random files being left around? It looks like the files are created in the user's home directory. I'm not sure how FileSystem.deleteOnExit works, but since the ACCUMULO-826 bug was about the files being deleted too soon, I'm inclined to think it's effective. This is what the javadocs say: "Mark a path to be deleted when FileSystem is closed. When the JVM shuts down, all FileSystem objects will be closed automatically. Then, the marked path will be deleted as a result of closing the FileSystem." I'll look at recreating a patch. For 1.5.0?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-507
17795	ACCUMULO-507	deleteOnExit works if the process is killed w/ SIGTERM. It will not work if the nodes dies or the process is killed w/ SIGKILL. A job submit directory was mentioned, I was wondering if the file would be placed there and if it would be periodically cleaned up by the job tracker. But the file was placed in user home dir, so that would be left in case of unclean shutdown. Could the file be placed somewhere that the jobtracker does cleanup?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-507
17796	ACCUMULO-507	I was wondering about that. The job submit directory is an example of such a temporary directory, but that wouldn't be created until after the user has called setRanges and submitted the job. Maybe I should revisit deleting the ranges from the job configuration when getSplits is called, after the ranges are read. Is there ever a case where a job would be submitted more than once? If so, this change would cause a failure on the second submit.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-507
17797	ACCUMULO-504	I am looking for offending log messages. I ran the following and did not see anything obvious. {noformat} grep "log\." ./src/server/src/main/java/org/apache/accumulo/server/master/Master.java | grep "erver" {noformat} I saw the following, but these log messages are already commented out. {noformat} // log.info("Current servers " + currentTServers.keySet()); // log.info("Current servers " + currentTServers.keySet()); {noformat} Other code besides Master.java could log. Does anyone know of log messages that log all of the servers?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-504
17798	ACCUMULO-502	David: did Keith's response answer your question, is the accompanying Javadoc addition sufficient, and can this ticket be closed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-502
17799	ACCUMULO-498	Note that it is possible to have a Key with a visibility that is not a valid ColumnVisibility (see ACCUMULO-193). How should we handle calls to getColumnVisibilityParsed() in this case? Just let it throw the error? I don't think we should parse the CV unless someone asks for it, but I guess we could store it once someone asks for it, and drop what is stored if the CV changes.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-498
17800	ACCUMULO-498	Stumbled across this ticket and saw it's been sitting with a patch. Does anyone have anything else to weigh in for/against it or any possible adjustments which need to be made for it?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-498

#	Requirement	Questions	Answered	Successfully Answered	URL
17801	ACCUMULO-490	I have been thinking about this very problem for quite some time. What is the suggested approach, Thrift? I have some JRuby classes that wrap an earlier version of Accumulo's Java API and had hoped to find time to port that to the latest version. It would then be fairly easy to add a Rails app on top of that and provide a RESTful interface. The Rails app could easily run on the JVM via JRuby. I envisioned an API like this: GET table/colf/colq PUT table/colf/colq with the data being the value for a new record POST table/colf/colq with data being the value to update a record GET table/colf?startcolq=abc&endcolq=def DELETE table/colf/colq We could return xml, json, or whatever by adding a .json or .xml to the request. Handling authentication might take some work. Need to provide a way to give the user column visibilities. I have some generic ideas there as well. I haven't worked out in my mind yet how to handle paging. Expanding that, we could then provide client libraries in different languages, like the MongoDB drivers. If we provided a Javascript one, we would need to support JSONP, Dojo's window.name transport or some other hack to get around Javascript's SOP restriction so you could POST, PUT and GET from a different host. I have no idea how this approach would compare to a Thrift service. Whatever the design, I would love to help work on this.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-490
17802	ACCUMULO-490	Hey David, Interesting class. What exactly is it doing besides opening a socket? I would have expected to see it look at the incoming request and then either retrieve data or store data. Maybe I didn't understand the intent of your ticket. I can certainly appreciate simplicity. What I have in mind is more like http://wiki.apache.org/hadoop/Hbase/HbaseRest. It will take some effort to make it useful and performant. Looks this is still a ways off if ACCUMULO-482 needs to be fixed first. I'll start a separate project on github with my ideas and post a link when I have something useful. Thanks	0	0	https://issues.apache.org/jira/browse/ACCUMULO-490
17803	ACCUMULO-490	Does the proxy provided in ACCUMULO-482 satisfy your needs for this?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-490
17804	ACCUMULO-483	Does this boil down to a "randomize the location of all tablets in a range" operation? How would you test the effectiveness of this action?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-483
17805	ACCUMULO-482	It was a little hard to understand the problem the way you stated it, but I think I understand. You mean any part of the keyvalue pair might be sensitive and need to be protected and the !METADATA table might expose some of those. So you're saying if a proxy is created you don't want !METADATA entries to be sent to it, because the way most instances of Accumulo are deployed involves the Accumulo client being run on a protected machine, i.e. a machine on the security perimeter that the user can't just examine the memory of. But I don't think METADATA entries need to be exposed outside the trusted accumulo client in order for a proxy to work ... one could simply ferry scan parameters and the resulting keyvalue pairs to and from the client (for queries). The proxy could also be configured to simply disallow reading directly from the !METADATA table. Would that still leave the possibility of sensitive elements of KeyValue pairs being leaked?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-482
17806	ACCUMULO-482	Have there been any updates on this (Sapan/Jim)? Talking to devs from other start-ups makes me think this feature is more important than I originally thought. Not having a thrift proxy API or a built-in REST API is a real barrier to non-java/non-JVM devs using/adopting Accumulo. This barrier is not there for Cassandra and HBase since they both have Thrift APIs. I think many devs in our community discount this since Java is so dominant, but outside the government, esp. in smaller start-ups there is a real aversion to Java, instead favoring python, Ruby, Node JS, etc.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-482
17807	ACCUMULO-482	As I have been researching ways to make Accumulo accessible to other languages, I ran across JCC[1] used by PyLucene[2] and I thought it was interesting. Anyone have experience with this? Would it be an option for exposing Accumulo API to Python? [1] : http://lucene.apache.org/pylucene/jcc/index.html [2] : http://lucene.apache.org/pylucene/	1	1	https://issues.apache.org/jira/browse/ACCUMULO-482
17808	ACCUMULO-482	I think the continued use of PKeyValue obfuscates what is going on here: {noformat} struct PMutation { 1:binary row 2:list<PKeyValue> cells; 3:list<PKeyValue> deletedCells; } {noformat} Seems that I would expect {noformat} struct PMutation { 1: binary row 2: map<PColumnUpdate, binary> cells; 3: set<PColumnUpdate> deleteCels; } {noformat} Why would a PColumnUpdate would be confusing to create directly? It would be confusing to me to get a runtime error because I specified the row as something different in the PKeyValue.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-482
17809	ACCUMULO-482	bq. On the one hand, introducing a PColumnUpdate object violates 1.3 If the set of data objects is too small, it will make the API harder to use. How does introducing PColumnUpdate make it harder to use? Can you give an example?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-482
17810	ACCUMULO-482	bq. Make it maintainable I am thinking the maintenance headaches in the future will come from the many table, instance, and security operation methods exposed via the thrift API and less so from the the read and write methods. But maybe not. Does the thrift API only provide a way to create a batch scanner? No way to create a regular scanner?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-482
17811	ACCUMULO-482	bq. Does the thrift API only provide a way to create a batch scanner? No way to create a regular scanner? There's currently only a way to create a batch scanner. bq. I am thinking the maintenance headaches in the future will come from the many table, instance, and security operation methods exposed via the thrift API ... I can easily remove methods from the API if we think they are not needed for this interface. Let's discuss if any of the ones I implemented won't be needed on a regular basis.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-482
17812	ACCUMULO-482	More considerations * How do we let users specify that they want to let system set timestamp? * make it easy to write what you scan... writing PKey supports this * make it easy to make column updates to a row atomic... writing list<PKeyValue> really works against this * how easy is it to write code against the API and how does code that uses the API look ... I think this is really important... being forced to repeat the same row for a mutation works against this Thinking about this, we could really minimize structure overlap by making PKeyValue that was composed of PKey, Value. And then PKey composed of a row and PColumn. I am thinking that trying to minimize the structs like this makes it more cumbersome to write the code. {noformat} updates = [new PColumnValue(new PColumn("cf","cq1","cv",7),"val1"), new PColumnValue(new PColumn("cf","cq2","cv",7),"val2")] deletes = [new PColumn("cf","cq3","cv",8)] Mutation m = new Mutation("row", updates, deletes); {noformat} If we had PColumnValue that was flat, I think the following code is easier to write than the code above. But this makes it more cumbersome to write code that scans data and writes it to another table. {noformat} updates = [new PColumnValue("cf","cq1","cv",7,"val1"), new PColumnValue("cf","cq2","cv",7,"val2")] deletes = [new PColumn("cf","cq3","cv",8)] Mutation m = new Mutation("row", updates, deletes); {noformat} I do not think there will be any confusion when a user looks at these flat data structs. I suppose I think users will start with the functions and not the data structs. At least thats what I would do. For example I would look at the write method and see that it takes a mutation, then I would look at mutation to see what it needs.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-482
17813	ACCUMULO-482	bq. How do we let users specify that they want to let system set timestamp? I just modified PKey to make timestamp an optional value. If it is unset then the system will set the timestamp. I believe this is the best way to specify this so we don't have any "magic timestamp values" (0, -1, etc) that users may actually want to use for some reason.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-482

#	Requirement	Questions	Answered	Successfully Answered	URL
17814	ACCUMULO-482	bq. I just modified PKey to make timestamp an optional value Interesting, not a thrift feature I am familiar with. In a few of the popular languages, is the generated code for dealing with an optional field easy to write? Does this complicate reading data in anyway? When reading data, we would always expect the timestamp to be set, so it would be odd to have to constantly check if its set when reading.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-482
17815	ACCUMULO-482	bq. In a few of the popular languages, is the generated code for dealing with an optional field easy to write? Not exactly sure about all languages. In Java optional fields aren't in the constructor and need to be set via a method call. bq. Does this complicate reading data in anyway? When reading data, we would always expect the timestamp to be set, so it would be odd to have to constantly check if its set when reading. Well, I copy whatever is in the key coming out of accumulo into the response value. That value cannot be null since it is a primitive. So on a scan the timestamp will always be set. (This happens in Util.toThrift(Key)). In Java the timestamp is still a primitive so it won't throw an execption if you try to get an unset timestamp value, i suspect it will be 0 since it is appears to be an uninitialized long primitive. I think if you try to access an unset optional non-primitive, it will return null. There are also isSet* methods generated for optional values that tell you if the value is unset (as opposed to set to null).	0	0	https://issues.apache.org/jira/browse/ACCUMULO-482
17816	ACCUMULO-482	Is there anything else that we're still expecting to see accomplished for the 1.5.0 release in regards to the thrift proxy? It looks ready to go from where I'm standing!	0	0	https://issues.apache.org/jira/browse/ACCUMULO-482
17817	ACCUMULO-482	I looked at the proxy/README today and some things were unclear to me. When I compile Accumulo from the top level, isn't the proxy submodule included? Why would it need to be built separately? In order to build the proxy module from within the proxy directory, I needed to run 'mvn install' at the top level because of a dependency. The proxy.properties file which needs to be edited is not the one in the build directory, it's the one in the installation directory. Also, the 'accumulo proxy' command should be run from the installation directory. I also got a chance to look at the proxy/python/README file. Perhaps it's just that I know very little about python but "PYTHONPATH=path/to/generated/api:path/to/thrift/libs python TestClient.py" does not look right. Should there be a ';' before the python command? Will python people know what the 'path/to/generated/api' and 'path/to/thrift/libs' are?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-482
17818	ACCUMULO-477	What does deprecated since 1.4.1 imply? Does that mean we would remove the methods in 1.5? Also, the javadocs are incorrect. Syntax for methods is Classname#methodName (or just #methodName if it's the current class), and there's a link to a method that doesn't exist (setOptions presumably should be addOptions).	1	0	https://issues.apache.org/jira/browse/ACCUMULO-477
17819	ACCUMULO-473	Keith, when are you proposing doing the binary search? On the block after the block has been read into memory or before? I would advocate keeping block sizes small enough to be read into memory and if a binary search is done, do it on the block in memory. This is in order to restrict lookups and scans to one disk seek.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-473
17820	ACCUMULO-473	I experimented with dynamically indexing cached rfile blocks and the results were promising. For my experiments I created a Rfile with 100K entries and a 50 byte value. I ran the attached Test.java program. This program seeks every key in the file backwards, forwards, and randomly. The rfile block size was 100k and the file had 69 blocks, so each block had around 1.4K entries. ||Test||Lookup Rate w/o transient index||Lookup Rate w/ transient index|| |Backwards| 8.5K lookup/sec | 86K lookup/sec | |Random | 8.5K lookup/sec | 75K lookup/sec | |Forwards | 452K lookup/sec | 500K lookup/sec | For the backwards test I took the last three numbers, the first two were slower as a result of the JIT not kicking in. For these test there was a 9x to 10x speedup. For each block of 1.4K keys an index of 31 keys was eventually created, then it stopped indexing (I added some prints to the code to get this info). The code does not create indexes where the keys in the block are closer than 32 keys. These test show that seeking forwards is really fast, the code was already optimized for this. So if an iterator can seek forwards it should, but this change will help with seeking anywhere. I am not sure why seeking forwards is so much faster. My current theory is that alot more object creation is done when seeking backwards or randomly. I have not confirmed this. Even when I adjust dynamic index code to create larger indexes, its still does not approach the forward seek speed.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-473
17821	ACCUMULO-466	Could also make the Accumulo file output format do something similar. It would throw an error if you try to write a really large key. If the user really wants to do this, they can configure the output format to accept larger keys. Should this be a 1.5 fix?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-466
17822	ACCUMULO-461	The ConstraintLoader and ConstraintChecker are in the server module. Do we think they should be moved to the client module? It doesn't seem like a good idea to make a separate loader and checker for the client.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-461
17823	ACCUMULO-458	Probably the test.Test interface should be commented as well. How does the TestTemplate file get turned into the ClassLoaderTestA/Test.jar (and the other two jar files)?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-458
17824	ACCUMULO-457	Benson, would you like to remove that src directory, or would you prefer for me to give it a try? Should this patch be applied before removing src, or would you make a different patch once it's gone?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-457
17825	ACCUMULO-456	Does it make sense to build on top of an existing utility like Sqoop? I don't know how Sqoop could be expanded to consider split points and iterators. Maybe there is something else?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-456
17826	ACCUMULO-455	Does tool.sh only work when the class you're running is a Tool?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-455
17827	ACCUMULO-452	I had another thought on this: locality groups are good for features that are in a relatively constant, low cardinality set with a fairly dense distribution across the primary partitioning dimension. Also, queries must be aligned with the locality group frequently enough to amortize the cost of that partitioning. This means that the current column family-based locality groups only really help when cells in sorted order frequently oscillate between locality groups. I want to say that this type of feature tends to be something that is explicitly modeled based on how the user wants to query their data. If the user decides to put this information in the row or the column qualifier, could they just as easily put it into the column family? By the way, expressions like John mentions in ACCUMULO-164 help to groups high cardinality features into a low cardinality set of groups, so I think we're on the same page there. Partitioning based on the timestamp is an interesting consideration. In this case, you would want a small number of ranges of timestamps to be "active" (not aged off yet) at any one time. Timestamps are a bit special, though, because they tend to be inserted in increasing order. Instead of using the locality group mechanism, we might achieve better performance by modifying the major compaction selection algorithm to avoid merging files that have very different timestamp ranges. Keeping track of timestamps on a per-file or per-block basis would support bulk filtering, and would be as (or more) efficient than locality groups. Might this be another approach to consider? Like Aaron, I think we need some more details on envisioned scenarios in which more generic locality groups would be useful before we jump too deeply into implementing them.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-452

#	Requirement	Questions	Answered	Successfully Answered	URL
17828	ACCUMULO-452	I am thinking that keeping a min/max time stamp per file may satisfy some use cases but not all. It would certainly be helpful. The compaction algorithm may need to be modified as Adam suggested to make it more effective. The way major compaction currently works in Accumulo older data will eventually end up in the largest file. If your goal is to avoid this file under certain circumstances, then the user has no explicit control over that. Also if you want to age off older data, you will probably still need to read this entire file to do that. If they want to scan the last 6 months of data for example and the largest file overlaps this time range but only 10% of the data in the file matches the range, then a lot of data needs to be filtered. Does HBase do anything special to deal with case. Why limit locality groups to only column families? * Increases model complexity. I think this is true. I think the complexity of the locality group model is not increased. If you understand partitioning on column families, you will easily understand the concept of partitioning on any part of the key. It certainly does increase the complexity of the big table model as a whole though. It would certainly give users more rope to hang themselves. Personally I am not opposed to this. * Increases code complexity. I do not think this is true. This would actually simplify the code and make this functionality much easier to test in isolation. I have found this with iterators, they dramatically decreased the complexity of the scan code. When iterators were first introduced, the scan loop was starting to get fairly complex. This seems a lot cleaner than customizing the current code to meet needs. OF course, end users may not care about the complexity of the accumulo source code. They just want it to solve their problems. * There are no compeling use cases. These must exist. I think the original time based locality group is one, is their a better simpler way to achieve this? That would remove this use case. The HBase design is simpler in terms of the model, but the code sounds more complex. Also this model does not give the user explicit control w/o allowing them to configure the compaction process in some complex way.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-452
17829	ACCUMULO-452	Even if users use the server-provided timestamps or their own, the timestamp still falls after the row and column, and is used the same way: to limit values after the rows and columns have been identified. To me it seems as if this happened, as a little play: BigTable Guys: look you can physically partition your data automatically using the rows! Users: Great! That works, but maybe I want an additional, secondary partitioning? BG: hmm. ok, how about you can also partition on the column family? It's the next item in the hierarchy, doesn't add too much complexity, pretty straightforward. Just specify them into groups called locality groups and I think we can keep this under control. Users: Yay! You guys rock! BG: You're welcome. Other users: Hey, locality groups are cool, but can I partition on column qualifiers? BG: why are rows and column families insufficient? OU: well, I don't know, I just really like to slice things every way possible. BG: sigh .. Yet other users: Wait, what about timestamps? You know what's more general than partitioning on a few elements of the data model? Partitioning on ALL the elements of the data model! So sweet. More general means more better! BG: I'm quitting to go work at Facebook.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-452
17830	ACCUMULO-452	Users do want this capability. They keep asking for it. We do turn around them and tell them to sort their data differently. They don't always like that answer. The intent of this is to meet a user need. Storing temporal information in the column family is a possibility. It would work well for some cases, like having two locality groups one thats the current month and another thats everything else. You put the month in the column family and reconfigure the locality groups every month. However, if you would like something like LG1 = < day old, LG2 = < month old, LG3 = < year old this would not be possible w/ the current locality group implementation. However ACCUMULO-164 may make this possible. Store time to the day in the column family. John pointed out one problem w/ this, its hard to automatically determine that patterns match disjoint sets. I need to think through ACCUMULO-164 some more and see what the possible gotchas are. If you have to duplicate the data in the timestamp into your column family to accomplish your goals, does this indicate a problem with the model? It do not think its clean, but its ok w/ me.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-452
17831	ACCUMULO-448	Is the concept of 'diameters' explained somewhere? I think this is the first time I have heard the term in relation to Accumulo.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-448
17832	ACCUMULO-443	Does this mean that we should take the targetJDK=1.6 from the pmd configuration?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-443
17833	ACCUMULO-443	I don't think so. What is the harm in leaving it there?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-443
17834	ACCUMULO-442	The comments look great. They lead me to my next question. The script does not delete old generated Java files. I don't know much about Thift. Could it be possible for the templates to have changed so that some generated Java files are no longer needed? How would they be managed?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-442
17835	ACCUMULO-442	Is there any reason why this script should not delete the contents of the src/trace/src/main/java/cloudtrace/thrift directory for each run of the script? Is anything in that directory other than thrift-generated files?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-442
17836	ACCUMULO-441	I am down to one last javadoc warning which says: {{noformat}} [WARNING] javadoc: warning - Multiple sources of package comments found for package "org.apache.commons.collections" {{noformat}} Any idea how to resolve this message? Also, I have only changed javadocs comments. The code compiles and the tests run with no errors. Is there anything stopping me from committing to the trunk tomorrow night? 42 files were changed.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-441
17837	ACCUMULO-429	Perhaps check the user you're running as (like not root) Do we want to do PID level process checking? In addition to existence of certain directories, we also want to check writability.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-429
17838	ACCUMULO-429	Is there some standard to judge when to use /usr/bin or /usr/local/bin or even /usr/local/accumulo as a symlink to /usr/local/accumulo-1.4?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-429
17839	ACCUMULO-429	I cleaned up the symlink checking for ACCUMULO-1071 and added a couple new features. The patch also checks dependencies in a consistent way. Has anyone used (or a fan/hater) of monit (http://mmonit.com/monit/)? I could put together a set of monit configurations for the accumulo processes.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-429
17840	ACCUMULO-429	+! for monit. The more utlities we can preconfigure the better. Would you put the scripts in the contribs directory or somewhere else?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-429
17841	ACCUMULO-421	Is this something that can be easily simulated with [HDFS Attributes|http://blog.cloudera.com/blog/2014/06/why-extended-attributes-are-coming-to-hdfs/]?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-421

#	Requirement	Questions	Answered	Successfully Answered	URL
17842	ACCUMULO-420	This feature needs to work correctly in the case of concurrent compaction request. If two user request compaction of a table at around the same time and one passes iterators and one does not, then iterators should be applied to all relevant tablets for the user that passed iterators. Also, tablet servers can decide to compact all files in a tablet at any time. This another form of concurrency that must be considered. The current design is that there is a compaction counter in zookeeper. A compaction operation will increment this. As each tablet is compacted it records what the compaction counter was when it started. Compaction operations wait for the compaction counters for each relevant tablet in the metadata table to reach at least its count. The following is a possible new design. compact FATE op # set compaction iterators in zookeeper # increment compaction counter in zookeeper # wait until all tablets in compaction range have at least this operations compaction counter # remove compaction iterators in zookeeper TServer compact all operation # get compaction iterators from zookeeper # get compaction counter in zookeeper # compact with current compaction iterators # write compaction counter to metadata table Considerations * setting compaction iterators in zookeeper must accomodate concurrent request * the iterators are intentionally read from zookeeper before the counter. The iterators are written before the counter. Does zookeeper guarantee that iterators written before the counter will always be seen? If not there is a race condition. * If two users submit compaction request with iterators, this design could run compactions with both users iterators settings. Does anyone see an issue w/ this? * Iterators from one user compaction request could be applied to multiple compactions of a single tablet.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-420
17843	ACCUMULO-418	Would you like to see this in 1.4.x? If so, then we could possibly create a reader that uses the current file format and reads the index at the end of the file. For 1.5 we can change the file format to avoid reading the index.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-418
17844	ACCUMULO-418	I took a look at how sequence file and sequence file input format handles this. It writes a 128 bit sync marker between records in the file. The sync marker is unique per file. It reads the beginning of the file to read the sync marker and other info. Then for the HDFS block you are interested in it start scanning looking for the 128 bit sync marker. Once it finds that it keeps reading records until it passes the block boundary. So sequence file does some non local reads for the beginning fo the file. In 1.4 w/ the multilevel rfile index, reading the index may involve more than reading the end of the file. If the index is more than one level, then the nodes of the tree are sprinkled throughout the file.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-418
17845	ACCUMULO-415	Currently scan iterators are stored as per table config as shown below. {noformat} root@test14 ctt_001> config -t ctt_001 -f iterator.scan ---------+---------- SCOPE | NAME | VALUE ---------+---------- table | table.iterator.scan.vers ..... | 20,org.apache.accumulo.core.iterators.user.VersioningIterator table | table.iterator.scan.vers.opt.maxVersions .. | 1 ---------+---------- {noformat} One way to support iterator profiles is to use a different prefix in the per table config. For example table.iterator.profile.<prof name>.scan.<iter name>. The Scanner could have a method added to it that allows a user to select a profile to use. This is backwards compatible because it just adds a method to the scanner. Table operations API would need to be adjusted to support profiles, hopefully we can do this in a backwards compat way too.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-415
17846	ACCUMULO-415	Could this be accomplished entirely in client code, if iterator settings were serializable to a client config?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-415
17847	ACCUMULO-411	Once we are finished imporving the javadocs for 1.4 we need to put them on the web site. Seems like I remember reading about a way to put a tarball of javadocs in Aapache CMS when we were first setting up the web site, but I do not remember where.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-411
17848	ACCUMULO-407	Does this property file change need to be done on all nodes in the cluster? Do you think it would be hard to have log4j read from an hdfs file? Also note that Accumulo might be running in an environment in which the files are all 'locked down' (i.e. unchangeable). In the past, I have written a small web interface to change the log levels. The web application writes the log4j property file.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-407
17849	ACCUMULO-407	All nodes already need to have a valid local Accumulo installation (which includes the conf directory), so no, I don't see a point to having it read from HDFS currently. There could be some benefit from exposing something in the monitor. I'd like to see this rolled into the upcoming 1.4.0 release, so I didn't want to make a big change. I would assume that you have some sort of tool to copy a new log4j file out already installed on a cluster (e.g. pscp, pdcp). Same goes for making a read-only conf directory writable to do the copy. All that being said, I don't think anyone would object to you writing something to procedurally alter the log level either. Please have at it.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-407
17850	ACCUMULO-407	Would it make sense to set the default log level for org.apache.accumulo to INFO in generic_logger.xml and monitor_logger.xml as well? I suspect there's no need to generate these logs by default any longer if DEBUG level can be turned on dynamically.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-407
17851	ACCUMULO-407	Found some issues- + String propertyLogConfig = String.format("%s/conf/log4j.properties", System.getenv("ACCUMULO_HOME"), application); + if (new File(logConfig).exists()) { + PropertyConfigurator.configureAndWatch(propertyLogConfig, 5000); line 1 should not have the application argument line 2 should be new File propertyLogConfig To put aside these minor foibles, there are a few questions we need to ask ourselves in this implementation. 1. Do we want to be watching both the *_logger.xml AND the log4j.properties? 2. Do we want to provide a mechanism to pick up the introduction of an application specific logger.xml file after start up? 3. Do we want to continue to have, by default, 2 different logs for each application, or can we just have 1? As for me, 1. Not if we don't have to. I believe the log4j.properties is only used as a default until we start watching the *_logger.xml. 2. Not for 1.4, but definitely in the future. 3. Unsure, but not for 1.4	1	1	https://issues.apache.org/jira/browse/ACCUMULO-407
17852	ACCUMULO-407	{quote} 1. Do we want to be watching both the *_logger.xml AND the log4j.properties? {quote} Yes, because log4j.properties provides a mechanism for influencing the logging configuration for all applications regardless of whether they have a specific *_logger.xml or not. {quote} 2. Do we want to provide a mechanism to pick up the introduction of an application specific logger.xml file after start up? {quote} Would this involve picking up the new application specific configuration file and disabling the watch on the generic configuration file? This would involve implementing a custom config file watcher because log4j doesn't seem to have hooks to cancel a watch initiated by configureAndWatch. An alternative would be to watch both a generic and application specific logger, but I suspect that could quickly get confusing. {quote} 3 . Do we want to continue to have, by default, 2 different logs for each application, or can we just have 1? {quote} How about this: by default the A2 (debug) appender is commented out and default log level is set to INFO. This would produce 1 log per application. If someone wants to re-enable DEBUG logging they can uncomment the appender and set the levels to DEBUG to get the current behavior with 2 files per application OR they do whatever they they want.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-407
17853	ACCUMULO-403	It seems like the decisionIterator shouldn't be restricted to the same column families that are returned by the RowFilter. Is the suggested usage that the subclass reseek in the acceptRow method if it wants different columns?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-403

#	Requirement	Questions	Answered	Successfully Answered	URL
17854	ACCUMULO-397	I have not looked at any accumulo code .. and I apologize for making this suggestion blind ... but should static methods be avoided so make testing easier?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-397
17855	ACCUMULO-391	The patch is looking pretty good. An additional thing you'll need to do is update the fetchColumns method so that you're fetching columns for particular tables. Also, how would you feel about getting rid of TableKey and leaving this as an InputFormat<Key,Value>? I suggest this because the RangeInputSplit already contains the table name, and a Mapper can access it through ((RangeInputSplit) context.getInputSplit()).getTableName(). It's somewhat awkward, but the advantage of keeping InputFormat<Key,Value> is that you can have Mappers that don't care which table they're running over and can be used with either the single table or multi-table input format. If we want to make it easier to grab the table name, we could add a public static method that pulls it from a Context, or from an InputSplit. A separate thing I want to do is get rid of the AccumuloIterator and AccumuloIteratorOption configuration objects and just make IteratorSetting Writable so it can be used directly. I'll open another ticket about that, though.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-391
17856	ACCUMULO-391	I've addressed all the issues on this thread- is everyone okay with me merging this in?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-391
17857	ACCUMULO-391	Also it looks like TypedValueCombiner, TTimeoutTransport, and core/src/test/resources/log4j.properties shouldn't be included in the patch. Which IDE are you using? Something appears to be adding a space between a method name and the parenthesis, e.g. "getInstance (" which doesn't match our code style. Also, it may be your IDE that has changed the AccumuloInputFormatTest to import InputFormatBase.addIterator and other static methods instead of just InputFormatBase, changing calls to this method from AccumuloInputFormat.addIterator to just addIterator. I found this to be confusing.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-391
17858	ACCUMULO-391	bq. What happens if I configure iterators for a table that I did not set as an input? I was trying to stay away from enforcing the order in which the user calls the static configuration methods but maybe this is unavoidable. I could always throw an exception in the initialize() method but that would be too late.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-391
17859	ACCUMULO-391	What do you think about taking the union of all the tables configured as input tablets, or with ranges, iterators, or fetched columns?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-391
17860	ACCUMULO-391	bq. What do you think about taking the union of all the tables configured as input tablets, or with ranges, iterators, or fetched columns? I think this may be easier to use but more error prone. There are more places where you can mistype a table name. A fluent class for configuring tables to read would be another alternative. The class could serialize itself to job or a static method could take job and varargs of these table config classes. This completely eliminates the problem from the API, the API does not allow setting configuration for a table you are not reading and the table name is only specified once (in the constructor of the config object).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-391
17861	ACCUMULO-391	bq. There's nothing tying us to the current method of configuring jrrbs What are "jrrbs"? bq. I wouldn't worry about having things configured for tables that aren't going to be scanned. I can't see a bug happening from it that wouldn't be caught with minimal exercise/testing. I am pretty sure people will end up spending lots of time trying to figure out why their scan time iterators are not being executed because they mistyped a table name. I agree that they will most likely hit this while testing, but its still wasting a developers time.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-391
17862	ACCUMULO-391	[~sonixbp] I do not want to impede progress, this will be a very nice feature to have for 1.6.0. I would like to see the API improved and I can spend time on that. How can we move forward? One possible way we could proceed is that you iron out all of the issues identified except for configuring tables not read, and I can pick up from there. I can work in the branch you created (you could rebase it first if you wanted a more concise commit history).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-391
17863	ACCUMULO-391	Looks good to me. Nitpicking: can you run the formatter over TableQueryConfig and InputConfigurator before pulling into master? It looked like those were still a little goofy. A broader general question: InputFormatBase (mapred and mapreduce) both have a protected static method getTabletLocator method that is tagged with a {{@since 1.5.0}} whose signature was changed with these changes. Do we just update the javadoc from since 1.5.0 to 1.6.0 because it wasn't a part of the "public API"? Do we have any general practice for javadoc tags on internal API changes?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-391
17864	ACCUMULO-391	That's a great question Josh. Originally, I changed the @since tags to 1.6.0 but since I only changed a few lines in the method, I couldn't justify the change to the javadoc. It would be nice if we could define a general practice for this. On another note- are we planning on deprecating the existing legacy methods for adding input to a single table? I'll add those in before I merge as well but I wanted to make sure that's what everyone wanted.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-391
17865	ACCUMULO-391	[~kturner], [~bills], and any other interested parties, there didn't seem to be a final consensus on whether or not the current single-table API methods should be deprecated in favor of a new multi-table configuration input method. Can I merge over what I have without deprecating the methods for now? If everyone in this thread is comfortable with the API changes, should we just deprecate them?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-391
17866	ACCUMULO-391	[~sonixbp], can you clean up the warnings in trunk? I know you're not finished, but it's driving me a little crazy.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-391
17867	ACCUMULO-390	Would it make sense to create a deb package specifically for a standalone vm (like ubuntu on virtualbox)? This deb file could have the tiny configuration set up automatically.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-390
17868	ACCUMULO-390	Is that supposed to say 512MBB?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-390
17869	ACCUMULO-378	This sounds really cool. I looked at the HBase doc, it seems like it replays the walogs on the slave cluster through the client API. Where you thinking of doing this for all tables, or just select tables? What are your thoughts on replicating user and table metadata in zookeeper? What are your thoughts on enabling replication for existing data? (we clould clone the table, copy its existing data, and replicate new walogs created after the clone operation). How are you thinking of handling bulk imported data? (could possible copy to slave and bulk import on their also, this could be a FATE operation initiated by the bulk import FATE operation). What are your thoughts on replicating split and merge operations on the master cluster? I am wondering how much we can leverage FATE to make this easier and more reliable.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-378
17870	ACCUMULO-378	I need a bit of clarification- are you adapting the WAL to log to HDFS via appends or are you working on a mechanism to shove the logs into HDFS once they are complete?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-378
17871	ACCUMULO-378	I would like to collaborate w/ you on this. It seems like a starting point might be a design doc. Would you mind putting together a design doc detailing your thoughts on this? Any other suggestions on how we could collaborate? We could also meet at the meetup (http://www.meetup.com/Accumulo-Users-DC/events/45491582/) if you are in this area.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-378

#	Requirement	Questions	Answered	Successfully Answered	URL
17872	ACCUMULO-378	[~kturner], [~sapanbshah42], did you guys ever come up with any sort of design document? Looking back at the last chatter, we were still in a localfs WAL capability which is a bit out of date considering current Accumulo support :) Given [~rmutyala]'s question on dev@a.a.o about this, any interest in thinking about this as a major 1.7 feature? It will give something to think about while testing 1.6.0 :)	1	1	https://issues.apache.org/jira/browse/ACCUMULO-378
17873	ACCUMULO-376	- turns debug logging on or off listscans - list what scans are currently running in accumulo. See the accumulo.core.client.admin.ActiveScan javadoc for more information about columns. trace - turns trace logging on or off -- Exiting Commands ---------- bye - exits the shell exit - exits the shell quit - exits the shell -- Help Commands ---------- about - displays information about this program help - provides information about the available commands info - displays information about this program ? - provides information about the available commands -- Iterator Configuration ---------- deleteiter - deletes a table-specific iterator deletescaniter - deletes a table-specific scan iterator so it is no longer used during this shell session listiter - lists table-specific iterators setiter - sets a table-specific iterator setscaniter - sets a table-specific scan iterator for this shell session -- Permissions Administration Commands -- grant - grants system or table permissions for a user revoke - revokes system or table permissions from a user systempermissions - displays a list of valid system permissions tablepermissions - displays a list of valid table permissions userpermissions - displays a user's system and table permissions -- Shell Execution Commands ---------- execfile - specifies a file containing accumulo commands to execute history - Generates a list of commands previously executed -- Shell State Commands ---------- authenticate - verifies a user's credentials cls - clears the screen clear - clears the screen notable - returns to a tableless shell state sleep - sleep for the given number of seconds table - switches to the specified table user - switches to the specified user whoami - reports the current user name -- Table Administration Commands -------- clonetable - clone a table config - prints system properties and table specific properties createtable - creates a new table, with optional aggregators and optionally pre-split deletetable - deletes a table droptable - deletes a table du - Prints how much space is used by files referenced by a table. When multiple tables are specified it prints how much space is used by files shared between tables, if any. offline - starts the process of taking table offline online - starts the process of putting a table online renametable - rename a table tables - displays a list of all existing tables -- Table Control Commands ---------- addsplits - add split points to an existing table compact - sets all tablets for a table to major compact as soon as possible (based on current time) flush - flushes a tables data that is currently in memory to disk getgroups - gets the locality groups for a given table getsplits - retrieves the current split points for tablets in the current table merge - merge tablets in a table setgroups - sets the locality groups for a given table (for binary or commas, use Java API) -- User Administration Commands --------- createuser - creates a new user deleteuser - deletes a user dropuser - deletes a user getauths - displays the maximum scan authorizations for a user passwd - changes a user's password setauths - sets the maximum scan authorizations for a user users - displays a list of existing users -- Writing, Reading, and Removing Data -- delete - deletes a record from a table deletemany - scans a table and deletes the resulting records deleterows - delete a range of rows in a table. Note that rows matching the start row ARE NOT deleted, but rows matching the end row ARE deleted. egrep - searches each row, column family, column qualifier and value, in parallel, on the server side (using a java Matcher, so put .* before and after your term if you're not matching the whole element) formatter - specifies a formatter to use for displaying table entries grep - searches each row, column family, column qualifier and value in a table for a substring (not a regular expression), in parallel, on the server side importdirectory - bulk imports an entire directory of data files to	1	0	https://issues.apache.org/jira/browse/ACCUMULO-376
17874	ACCUMULO-361	Corey, would that effectively be a short cut to "clone and export" on a list of tables?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-361
17875	ACCUMULO-359	How would one track the cluster size, memory availability, machine power, and other tasks running during the testing? It seems like any given job has a lot of factors that affect performance.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-359
17876	ACCUMULO-359	Can we defer this to an external tool, like YCSB?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-359
17877	ACCUMULO-351	Are the options for compression codecs pluggable? Seems like it would be nice to just have a compressor/decompressor interface and tell people that they can just throw implementations in lib/ext. Does Accumulo do much with the compressors other than toss them over to HDFS for getting codecs?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-351
17878	ACCUMULO-351	bq. Are the options for compression codecs pluggable? Not really w/ Accumulo's copy of BCFile. See org.apache.accumulo.core.file.rfile.bcfile.Compression.Algorithm. Would be nice to change this or see if a newer version of BCFile has.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-351
17879	ACCUMULO-348	Could we batch the splits to tablet servers instead of sending them individually?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-348
17880	ACCUMULO-343	In for a penny ... I've run through all of the examples and the user manual but I managed to misunderstand what was happening with the setauths command and visibility labels. Can examples using no visibility labels (i.e., no -l option on the insert) be added to README.visibility for completeness? If there are no visibility labels is that equivalent to "public" so that anyone can see it? Please add something like "Accumulo does not keep sort of list of visibilities. Furthermore there is no definition of visibilities, so there is no such thing as defined and undefined visibilities. This might lead to an issue with misspelled visibility labels. Keep this in mind when debugging a misbehaving application." somewhere in the README.visibility file. Can you add a few scan examples to the "Inserting data with visibilities" section? For example, {noformat} # Clear authorizations test> setauths -c test> scan # no records are displayed because no authorizations are set and all records in the test table have visibility labels. test> scan -s broccoli 21 21:56:40,121 [shell.Shell] ERROR: java.lang.RuntimeException: org.apache.accumulo.core.client.AccumuloSecurityException: Error BAD_AUTHORIZATIONS - The user does not have the specified authorizations assigned test> setauths -s apple test> scan # no records are displayed because the one record in the table needs both "apple" and "carrot" authorizations due to the visibility label. test> setauths -s apple,carrot test> scan florida george:age [(apple&carrot)|broccoli|spinach] 50 test> setauths -s broccoli test> scan florida george:age [(apple&carrot)|broccoli|spinach] 50 test> setauths -s apple,carrot,broccoli # If you have several authorizations, you can use the scan command to display only specific authorizations. test> scan -s broccoli florida george:age [(apple&carrot)|broccoli|spinach] 50 # Of course, the visibility logic is still applied, therefore you must specify both apple and carrot test> scan -s apple,carrot florida george:age [(apple&carrot)|broccoli|spinach] 50 {noformat} Is there any way to set the VisibilityConstraint as a default constraint for all tables?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-343
17881	ACCUMULO-339	Looks good. Thanks for the ticket and the patch. Sapan, I've added you to our contributors list in JIRA. David, I'm adding you as well due to your sheer volume of reported issues. This will allow you to have more control over tickets.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-339

#	Requirement	Questions	Answered	Successfully Answered	URL
17882	ACCUMULO-303	I was initially poking around with the idea of being able to set multiple different formatters on a table, but eventually backed it out for the simplicity of a single formatter per table. I also have some concern regarding formatter "logic" being in the createtable, formatter, and config commands. I'd be interested in any feedback as to what you think a good strategy is moving forward. Does the formatter command still need to exist if all it's really doing is acting as a wrapper around a config call? Is it beneficial to preserve some backwards compatibility of the formatter command? Does anyone actually use the formatter command now?? :) I can look into adding a functional test as well.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-303
17883	ACCUMULO-303	If you are up to it, could you create a MockShell that would accept lines of input and provide lines of output to test this? It would make testing the shell a unit test, which would be much faster/frequent.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-303
17884	ACCUMULO-303	Yes, you would get a ClassNotFoundException if you tried to set a formatter that doesn't exist on the Accumulo classpath. Additionally, if there was one previously set on the table, you would, I think, see the same ClassNotFoundException when you try to switch to that table in the shell. This could very easily happen if someone upgrades Accumulo and forgets to copy the jars from the old Accumulo installation's lib/. Is there a desired benefit to falling back onto the DefaultFormatter if the specified formatter cannot be loaded (ClassNotFoundException, ClassCastException, etc)? Having a missing formatter class set *will* mess up all the scans until removed or the classpath being fixed; however, (Batch)Scanners won't be affected.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-303
17885	ACCUMULO-303	Sorry for the delay on this. Please poke around on this and let me know if I missed anything we discussed above. Things covered in ACCUMULO-303-bugfixes-with-mock-shell.patch: * Added argument to formatter option in CreateTableCommand * Removed any caching of formatter inside of Shell (will always check ZK) * Will fall back to DefaultFormatter if the configured Formatter cannot be loaded * Added option to ScanCommand to allows users to provide a formatter of their choice to scan with * Reworked Shell, created MockShell, and created FormatterCommandTest which uses MockShell	0	0	https://issues.apache.org/jira/browse/ACCUMULO-303
17886	ACCUMULO-291	I do not understand this ticket. Is this related to calling the init() method on an iterator, or something else?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-291
17887	ACCUMULO-290	Does the proxy provide essentially the same functionality requested by this ticket?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-290
17888	ACCUMULO-274	True. Is the main() method needed in FileDataIngest? If it stays, to me, that means that separate documentation is needed to show people how to use it.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-274
17889	ACCUMULO-274	I know that I am a nudge but could someone put this example in context. Why would someone *want* to store file data into Accumulo instead of directly into HDFS? Or into a super file?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-274
17890	ACCUMULO-263	What would we be checking here? The existence of the table?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-263
17891	ACCUMULO-261	Setting the number of entries can also be very useful. You can already set the batch size in bytes on a per-table basis through the table.scan.max.memory parameter. Do we need to be able to set it per scan?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-261
17892	ACCUMULO-259	Thought I made a ticket for this, but closest I can find is the comment in ACCUMULO-238 . This is step one. Step two is to do as you described- seperate the 'authority provider' from the 'authenticator'. This way we can let users authenticate against one system that may or may not be able of holding credentials. So, as you said, they could authenticate against one system and then use zookeeper for their various credentials. This still leaves us with the issue of coming up with a versatile interfaces that provides enough information for most providers. Also, should this be a seperate ticket, or just one large operation? And I agree on the errors.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-259
17893	ACCUMULO-259	What would be the reason for two tickets? Are you thinking of making some initial changes in 1.4 and doing more in 1.5?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-259
17894	ACCUMULO-259	Well, that's the question. How much do we want to do and when? 1.4 is nearing release and I worry that this change is too big, but I have a feeling that we want to get this pluggable interface in place (primarily due to the release of hadoop 1.0 which has kerberos support). Unless you have an objection, I'm proposing doing the full split (authenticator, authority provider, and proper utilization of these interfaces) for 1.4 so we have them available ASAP.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-259
17895	ACCUMULO-259	The functional tests are failing with: {noformat} java.lang.NullPointerException at org.apache.accumulo.core.client.ZooKeeperInstance.getSecurityTokenClass(ZooKeeperInstance.java:326) at org.apache.accumulo.core.util.shell.Shell.updateUser(Shell.java:937) at org.apache.accumulo.core.util.shell.Shell.config(Shell.java:279) at org.apache.accumulo.core.util.shell.Shell.main(Shell.java:374) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.accumulo.start.Main$1.run(Main.java:97) at java.lang.Thread.run(Thread.java:662) {noformat}	1	0	https://issues.apache.org/jira/browse/ACCUMULO-259
17896	ACCUMULO-259	SunnyDay ran fine, and I did a test and had the exact same tests failing in the branch and in the trunk the other day. Which tests are failing for you?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-259
17897	ACCUMULO-247	How do you install pdflatex on Ubuntu 11.10?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-247
17898	ACCUMULO-247	How do you install pdflatex on Ubuntu 11.10?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-247
17899	ACCUMULO-244	Eric, I think you had mentioned possibly using iptables to accomplish this. Is that still what you are thinking?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-244
17900	ACCUMULO-244	This ticket is over a year old. Is this still a concern?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-244
17901	ACCUMULO-241	We have a few goals for the language used in security labels (expressions made up of operators and authorizations): they should be easy to read by human and by computer, and they should be unambiguous, the Boolean logic operators should be easily distinguished from the atomic authorizations, labels should be backwards compatible forever, and the language should be extensible to anything we might want to do with it in the future. To support backwards compatibility while leaving room for extension, we originally reserved all non-alphanumeric characters and only allowed alphanumeric characters within authorizations. When our users asked for '_', '-', and ':' for use in authorizations, we added those to the white list. Moving to a black list approach is a bit more limiting to extensibility, but I think it can be done while preserving the possibility of adding future capabilities. Supporting escaping of reserved characters might be another option, but that might reduce the human readability. The big question is what do we want to do with cell-level security in the future? I think we probably want to support "not" at some point, so probably '!' and '~' should be reserved. If we do want to support escaping, we should probably reserve '\' or '#' and ';'. It has been hinted that we might want to support something like regular expressions, so '*', '?', '[', ']', '+', ... How about variable substitution, with '%' or '$'? Maybe it would be better to keep a white list for now?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-241

#	Requirement	Questions	Answered	Successfully Answered	URL
17902	ACCUMULO-241	Doh! For some reason I kept not seeing the word 'visibility'. Now I understand what labels are being discussed. Would it make sense to use syntax like NOT(FOO) instead of !FOO in labels. Using function-based syntax (if that terminology makes sense) reduces the need for fiddly single-character flags and should be more readable. It also leaves the door open for an extensible mechanism whereby keywords can be attached to user-supplied functionality.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-241
17903	ACCUMULO-241	I suspected the default encoding on the build server was different, but I was not sure. Does anyone know if there is a way to run something on the build server w/o committing? I will make changes to allow the encoding to be explicitly specified and see if that helps.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-241
17904	ACCUMULO-227	That makes sense - but then are you calling for an increase in the granularity of isolation (i.e. finer grained, down at the cell level) or the opposite? How does the mutation-level isolation disallow mutations with multiple identical cells? I was just looking at the isolation that Accumulo provides currently and it needs to be clearer to users. If this is largely an internal change that's fine, but if not we should make sure the isolation model makes sense for most users.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-227
17905	ACCUMULO-212	Maybe this ticket should be closed, it is not well defined. Whats actually being asked for? Is it the hooks? I do not understand why hooks would be called coprocessors? Accumulo already has some hooks/plugins (Iterators, Constraints, Balancer, Memory Manager, Authentication). If additional hooks are wanted, then why not open a ticket for each one? Maybe this ticket could be reopened in the future as a parent ticket for additional hooks. Also, the first two items mentioned in the ticket already exist.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-212
17906	ACCUMULO-199	Is this still something you want to do? I'm not sure I see the value of creating packages for a specific environment when we can create generic packages for generic environments, and provide ways to easily configure that environment. I'm not sure it makes much sense to manipulate modify our build profiles for users to rebuild. If anything, it'd be better to create an entirely separate project, which pulls in all the relevant accumulo dependencies, and repackages, with the desired configuration, if a user wished to do that.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-199
17907	ACCUMULO-196	Here are some instructions for enabling shell access in the monitor page. Generate keystore and truststore: keytool -genkey -alias <serverAlias> -keyalg RSA -keypass <keyPass> -storepass <storePass> -keystore keystore.jks keytool -export -alias <serverAlias> -storepass <storePass> -file server.cer -keystore keystore.jks keytool -import -v -trustcacerts -alias <serverAlias> -file server.cer -keystore cacerts.jks -keypass <keyPass> -storepass <storePass> Enter "yes" when prompted "Trust this certificate?" Set the following properties in accumulo-site.xml: monitor.ssl.keyStore - /path/to/keystore.jks monitor.ssl.keyStorePassword - keyPass monitor.ssl.trustStore - /path/to/cacerts.jks monitor.ssl.trustStorePassword - storePass (Re)Start the monitor and there will be a link for the shell. You can use it in mock mode, or log in as an Accumulo user.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-196
17908	ACCUMULO-164	John made the comment offline that determining if a set of patterns matches disjoint sets of column families may not be possible. I think this is may be true for regular expressions. However, it may be easy to determine this automatically with limited wildcarding. If only prefix wildcards were allowed, it seems like the following algorithm would ensure they are disjoint. {noformat} boolean isDisjoint(Set<String> prefixes){ while(prefixes.size() > 1){ String shortestPrefix = removeShortestString(prefixes); for(String prefix : prefixes){ if(prefix.startsWith(shortestPrefix)){ return false; } } } return true; } {noformat} Does this seem correct? For suffixes, startsWith() would be replaced with endsWith(). So maybe we can handle all prefix wildcards or all suffix wildcards. Can we verify anything else is disjoint? I do not think so. The following wildcards could match overlapping sets. {noformat} *a* *b* {noformat} And so could the following. {noformat} foo* *bar {noformat} So even though the literal parts of the above wildcards are unique, they can still match overlapping data.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-164
17909	ACCUMULO-164	A prefix wildcard is a simple form of a range. One other thing we may want to consider is simply having ranges. So the candidates for this ticket would be prefix wildcards, suffix wildcards, and/or ranges. These can be all be automatically verified as disjoint. Is allowing this type of setting in locality group configuration worthwhile w/o changing the scanner API? If we decided to allow prefix wildcards, then it would be easy to create a locality group with millions of actual column families. However there is no way on the client side to ask for everything in that locality group without enumerating all possible column families. It may not be possible to enumerate all possible columns families, therefore it may not be possible to read an entire locality group. To remedy this the client API and iterator API would need to be changed to allow specification of column prefixes (or ranges, or column suffixes).	1	1	https://issues.apache.org/jira/browse/ACCUMULO-164
17910	ACCUMULO-164	Well put, [~ctubbsii]. I agree with everything you say. A bit off to the side: I think one of the best lesser-realized "features" of Accumulo falls into the fact that column families are not fixed. To this extent, I think your final point about creating this functionality in an "expert" mindset further enhances the ability to support unlimited column families. The feature in its simplest sense does conflict with the original idea behind locality groups, but is that always "bad"? I'm not sure, but it's definitely different. Do you have any ideas on how to present such a feature that would avoid steering the common user toward it? Is healthy warning/documentation sufficient?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-164
17911	ACCUMULO-164	{quote} The feature in its simplest sense does conflict with the original idea behind locality groups, but is that always "bad"? I'm not sure, but it's definitely different. {quote} We've already extended the original idea behind locality groups, by allowing users to specify more than one column family for a locality group. And, I think that is definitely not "bad" ("good", even). This is just an easier way to select multiple families to put in a locality group, based on a common characteristic (like common prefix). However, I question why something like "common prefix" should be a desirable selection mechanism for multiple families in the first place. Not only are (in the case of the common prefix) these data naturally grouped locally without any use of locality groups, it's not clear to me that something like "common prefix" is the most sensible way to group related families in the general case. I'm not sure there *is* a general case, though. Perhaps len < 4 is more useful than identifying a common prefix for some users? Further, the only application for this, that I can think of, is when users introduce variability into the family that allows the number of distinct families to grow continuously (which, I think can be, and should be, done in the qualifier instead). So, I personally see little benefit to it (at least, for the common prefix case; though full regexes or suffixes would certainly have greater benefit). Maybe the most useful, and general, thing we could do to provide users the most utility to select families for a locality group, is to allow users to inject a user-defined hash function (maybe in JEXL?) to bin families into discrete localities by the arbitrary method of their choosing? {quote} Do you have any ideas on how to present such a feature that would avoid steering the common user toward it? Is healthy warning/documentation sufficient? {quote} If implemented, I think documentation should be sufficient to address all of my concerns. The main thing is just make it clear that the feature is used to *select* multiple column families, so that it's not implied that families with variability *are* the same "family". The API treats non-equal families as distinct, and that's how we should discuss them.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-164

#	Requirement	Questions	Answered	Successfully Answered	URL
17912	ACCUMULO-156	The trie used to be used for the visiblity filter parsing. We ended up switching it to a hashmap shortly before it arrived on the apache SVN, hence it's continued existence for Keith's toy project. Also, what do you mean the acceptance test?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-156
17913	ACCUMULO-143	Hi Brian, Did you have a chance to work on that? We are considering HCatalog->Hive->Accumulo and the work you are doing on Hive->Accumulo sounds very promising.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-143
17914	ACCUMULO-142	can this ticket be closed?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-142
17915	ACCUMULO-140	Would it be worthwhile adding fadvise like support to HDFS, where it can be given hints about the application plans to access a file?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-140
17916	ACCUMULO-139	Is it possible to make the compression used pluggable? I thought at one point this was the case - people could choose between LZO, Gzip, etc ..	1	0	https://issues.apache.org/jira/browse/ACCUMULO-139
17917	ACCUMULO-138	{{/do-component-build}} uses {{-Dhadoop.version=0.20.205.0}} to select a specific Hadoop version. Not only is this a very old version, it's at odds with the Bigtop "build everything together" view. Is there a way to to pull this out, to have it pick up the central bigtop-managed versions?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-138
17918	ACCUMULO-138	[~vines] I'm interested in adding accumulo to bigtop. Can I assign this jira to myself and post a patch? Is that okay with you?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-138
17919	ACCUMULO-138	In keeping with our general policy of "integrate within external projects where possible", this should be done as a part of the relevant BIGTOP ticket (BIGTOP-1175). [~gkesavan], could you please work with the assignee of that ticket to move forward? If no one objects, I plan to close this ticket as wontfix in favor of BIGTOP-1175.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-138
17920	ACCUMULO-136	Any movement on this topic? Should this ticket be resolved?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-136
17921	ACCUMULO-132	If the problem is slow tablet servers, then another way to deal with this is to detect and decommission slow tablet servers. If the problem is a slow datanode, many of the proposed solutions may still result in slow reads. In this case a solution may be to detect and decommission slow data nodes. When decommissioning datanodes and tablet servers, would need to ensure the whole cluster is not decommissioned :) Another source of slow reads I have seen is at the network level. Once I observed an entire rack running slow. After manually removing all nodes in this rack, everything sped up.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-132
17922	ACCUMULO-118	Why was this switched to a blocker? What is it blocking?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-118
17923	ACCUMULO-118	bq. It's time to start thinking about 1.6, and I intend for this to be part of that release. My goal is a working branch in 2 months. Do you plan to post a design doc outlining your approach to tackling this?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-118
17924	ACCUMULO-118	How does this related to ACCUMULO-722? Seems like if we implemented ACCUMULO-722 then we wouldn't need this, but this is a much easier ticket, right?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-118
17925	ACCUMULO-118	bq. Would avoding storing direct pointers to namenodes in the metadata table be sufficient to satisfy this? Always have a level of indirection like viewfs? Yes, a simple direct mapping is what I would like to see. Using a hash function is not simple or direct and I don't see why its needed. If I have a backup of the files from ns1 on ns2 and vice versa (think hot backup) and ns2 fails, then I can just change the mapping so that both ns1 and ns2 point to the same Hadoop HDFS instance. No update of the metadata table needed.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-118
17926	ACCUMULO-118	I was looking at some docs on viewfs. If possible, I am thinking we should not do anything that would preclude using viewfs. It seems like if URIs were supported for tablet dirs and files (along with a way to choose a tablet dir) that this would almost be enough to support viewfs. {noformat} 1;m srv:dir viewfs://clusterX/accumulo1/tables/abc 1;m file:viewfs://clusterX/accumulo1/tables/abc/F0000002.rf [] 196,1 1< srv:dir viewfs://clusterX/accumulo2/tables/abc 1< file:viewfs://clusterX/accumulo2/tables/abc/F0000003.rf [] 196,1 {noformat} If we want to further develop our own indirection layer, then maybe we should define our own URI prefix. Something like ans://. How independent should this URI be? Something like ans://<namespace name>/<path> would assume that you know where to look <namespace name> up. If the URI were like ans://<zookeepers>+<instance id>+<namespace name>/<path> then it would be more self contained. I do not think its necessary to make it self contained, its for internal use and would be translated by as needed. I was thinking about how bulk import will work in this federated world. Below is one way this could work. * Client calls import dir w/ /foo1 * Accumlo client code uses local config to convert /foo1 to URI hdfs://nn1/foo1 * hdfs://nn1/foo1 is passed to Accumulo server code via thrift * Accumulo server code looks at URI to determine where to move to, determines it has accumulo dir hdfs://nn1/accumulo. * moves files in hdfs://nn1/foo1 to hdfs://nn1/accumulo/tables/abc * Replaces hdfs://nn1/accumulo/tables/abc with ans://ns1/accumulo/tables/abc * Does bulk import of files in ans://ns1/accumulo/tables/abc Is this how this should work? The scenario above implies that Accumulo needs a dir on each namenode and way of mapping URIs to the appropriate Accumulo dir. Need to wor through this scenario w/ viewfs also.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-118
17927	ACCUMULO-118	bq. This suite of tickets is really starting to concern me. I am also concerned, I am worried about administrators having a really bad experience (bricked Accumulo instance). In hindsight I think this feature was merged in before it was complete and should not have been merged in. What can we learn from this to make development go more smoothly in the future? Personally I think if a new feature is incomplete (in terms of test, usability, documentation, etc), breaks other features of Accumulo, or introduces new problems it should not be merged in. I did not realize all of the problems absolute paths could cause until well after it was merged in. In hindsight I think the design of this feature should have started with administrative use cases. Maybe that would have brought to attention the problems w/ absolute paths much earlier. bq. I would like to start the conversation about marking this item and items around it as @Experimental I wish it were that simple. Wether you use multiple namenodes or not, Accumulo will use absolute paths in 1.6 and administrators do not have a reliable way to manage those paths. I have a fix for ACCUMULO-1832 in the works, I should be done w/ that very soon. I think that fix will address my major concern w/ this issue.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-118
17928	ACCUMULO-118	My comment about voting on design documents was a possible solution to a problem. The problem is how can we get peer reviewed designs? How can we motivate each other to spend a good chunk of time critically thinking about our designs. This issue had a design document and I remember reading and thinking "makes sense". I do not remember investing much time trying to find flaws with it, I think I just wanted to get back to whatever I was immersed in at the time. My thinking is that calling for a vote on the dev list after writing a design doc for a feature seems like a simple thing to try w/o any additional overhead because we have the dev list in place.	1	0	https://issues.apache.org/jira/browse/ACCUMULO-118
17929	ACCUMULO-115	Is this something we still want to do at all, with all of the test framework work being done?	0	0	https://issues.apache.org/jira/browse/ACCUMULO-115
17930	ACCUMULO-115	We currently use both EasyMock and PowerMock. [~jesse_yates], are you interested in rebasing your test for TabletServerBatchDeleter onto a current branch (namely 1.4.6-SNAPSHOT)?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-115
17931	ACCUMULO-87	grid_or2 is pretty good, although it might be a little hard to read. Aaron, can you make rearrange the letters into a line to see how it looks?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-87
17932	ACCUMULO-87	Aaron, Can you try creating starburst using a rounded rectangle? I would like to see how it looks w/o the corners.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-87

#	Requirement	Questions	Answered	Successfully Answered	URL
17933	ACCUMULO-87	Also - is the idea that the community votes for the best logo 'concept' which could be reworked a bit, to make it pretty, etc., or will the vote determine the golden master PNG file to be used going forward?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-87
17934	ACCUMULO-87	I committed an initial version of the logo display page to site/trunk/accumulo/misc/logo_contest_20111111.html. Can somebody make that available on the accumulo page? I'm about to hop on a plane.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-87
17935	ACCUMULO-87	Here's a site where you can look at all of the logos, sized evenly(ish), and grouped under categories that I made up: http://incubator.apache.org/accumulo/misc/logo_contest_20111111.html What do you think about this for the voting: 1. In the first voting round, everybody votes for as many groups of logos as they like, taking names from the page linked above (e.g. +1 "Greek God", +1 "Bond-O", ...). 2. We will rank the groups by popularity and take the top N for a second voting round. 3. In the second voting round, everybody votes for one group, and the top vote-getter wins. Shall we say 23:59:59 EST on Wednesday, 2011-11-16 for the first round deadline, and 23:59:59 EST on Friday, 2011-11-18 for the second round deadline?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-87
17936	ACCUMULO-84	A couple of requirement questions. 1. If the Cloudbase WAL fails to log the mutation (e.g. out of disk space), the requirement is for the WAL to continue in "READ-ONLY" mode as opposed to stopping as it does now. By continuing without WAL logging, can the Table still be recovered? 2. After the WAL enters READ-ONLY mode, is the requirement to start logging again once disk space becomes available? Or is the requirement that once the WAL enters READ-ONLY mode, it will always be in READ-ONLY mode until it is stopped and restarted? JimZ	1	0	https://issues.apache.org/jira/browse/ACCUMULO-84
17937	ACCUMULO-81	Is there an standard eclipse format for the accumulo source or for apache in general?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-81
17938	ACCUMULO-77	Can you be more specific about which errors are obscured (preferably in the form of a unit test, says ecn)?	1	0	https://issues.apache.org/jira/browse/ACCUMULO-77
17939	ACCUMULO-76	Can you clarify which tests should be converted, or close this ticket? Some tests may be able to be converted to integration tests using MiniAccumuloCluster.	0	0	https://issues.apache.org/jira/browse/ACCUMULO-76
17940	ACCUMULO-49	mlockall_agent sounds pretty cool. We have experimented w/ mlock. We wrote some JNI code, but the user had up the ulimit for it to work. One issue we ran into is that thread stacks seem to be preallocated when mlock is enabled. Without mlock, memory is only allocated on the stack as it is used. So w/ the default 1M stack, each thread would allocate 1M of mem when it was created. Used pmap to verify this. Have you had any issues w/ this?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-49
17941	ACCUMULO-49	ah - interesting. No, I haven't done any real testing of the above agent. Do you mean that, even if threads started and then stopped, the memory would be left allocated and thus "leak" over time? Or just that you have to set the limit to at least heapsize + maxThreads*1M?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-49
17942	ACCUMULO-19	Thanks Scott, I patched the 1.4 branch, and was able to build a .deb and install it on my Ubuntu laptop. I had to install Cloudera cdh3u2 to satisfy the "hadoop" dependency. Shouldn't the package depend on default-jre instead of jdk? Is this a difference between debian and ubuntu? I naively tried to run the resulting install, and as expected, it failed horribly. I created ACCUMULO-210 to build on this package for a better user install experience. If you wanted to take on this ticket, that would be great!	1	1	https://issues.apache.org/jira/browse/ACCUMULO-19
17943	ACCUMULO-19	Thanks Scott, I patched the 1.4 branch, and was able to build a .deb and install it on my Ubuntu laptop. I had to install Cloudera cdh3u2 to satisfy the "hadoop" dependency. Shouldn't the package depend on default-jre instead of jdk? Is this a difference between debian and ubuntu? I naively tried to run the resulting install, and as expected, it failed horribly. I created ACCUMULO-210 to build on this package for a better user install experience. If you wanted to take on this ticket, that would be great!	1	1	https://issues.apache.org/jira/browse/ACCUMULO-19
17944	ACCUMULO-14	Oh, I had no idea that there were functional tests. Those are definitely a solid tool. However, I still think that having a set of integration tests (before functional tests) is a reasonable goal. By developing a pipeline of tests, developers will have a shorter iteration time when building new features with more confidence that they aren't breaking anything. Side note: are we running these functional tests on the nightly build? I couldn't see where we were. This is something important that actually needs to be run as part of *a* build cycle to make sure we aren't breaking anything. Integration tests would be great for this on jenkins since we don't need to spin up hadoop or zk. On top of the utility for developers, users should still get a good way to test things, to make sure that they are writing 'good' functionality (e.g. new iterators). A more 'real' mock instance would go a long way to help that.	1	1	https://issues.apache.org/jira/browse/ACCUMULO-14
17945	ACCUMULO-14	Jesse, I took a quick look at the patch. It was hard to see what changes were made to the tablet server and master to achieve the goals of this ticket. There were formatting changes, many additions of final, additions of @Override. All of this noise makes it hard to know what changes where made to these very important classes. If you think things in the master and tablet should be marked final, I think that work should be done in a separate ticket w/ separate patches. If you are you using eclipse, are you using formatter mentioned at http://incubator.apache.org/accumulo/source.html? Would it be possible to submit a patch that did not have so many extraneous changes? Or point out what the critical change to the master and tablet server were?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-14
17946	ACCUMULO-14	I think instamo at least has to have the Apache pom as its parent to get maven repository deployment to work. Keith, is there a reason you didn't want to make the examples pom the parent of instamo (so it could inherit version numbers for its dependencies)?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-14
17947	ACCUMULO-14	bq. Keith, is there a reason you didn't want to make the examples pom the parent of instamo. Yes, I wanted to it to be a self contained example that a developer could use to quickly get up and running with Accumulo. I do not want to deploy it. An instamo jar in maven repos would have no use. I added it as a child of the examples pom just so it would automatically be pulled into eclipse, should I remove that?	1	1	https://issues.apache.org/jira/browse/ACCUMULO-14
17948	ACCUMULO-1	Commit 2c9dee325c6407b6ce7221521ba536ec0f633061 in branch refs/heads/master from [~ecn] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=2c9dee3] ACCUMULO-1.5.1-SNAPSHOT need to store location without full path	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1
17949	ACCUMULO-1	Commit cfc0e18ec40648d3a75eca7c57058dbf0ed34811 in branch refs/heads/1.4.5-SNAPSHOT from [~ecn] [https://git-wip-us.apache.org/repos/asf?p=accumulo.git;h=cfc0e18] ACCUMULO-1.5.1-SNAPSHOT need to store location without full path	1	1	https://issues.apache.org/jira/browse/ACCUMULO-1
17950	ACE-414	The agent current doesn't keep track (in a persistent way) whether a DP is completely downloaded or not. And due to ACE-413, it simply restarted the download of a DP when the agent was restarted. With ACE-413 now fixed, it no longer does this, and simply will install the DP *once it has verified with the server that it indeed has the complete DP downloaded*. I'm a little hesitant to built in all kinds of completion checks to reduce the number of HTTP calls upon startup. I do see a functional requirement to allow offline DPs to be installed as a kind of bootstrapping mechanism for new agents, but this is also what ACE-162 is about (albeit described in a more cryptic way). IMO, implementing that issue solves this problem as well. [~wilsib], do you agree?	1	1	https://issues.apache.org/jira/browse/ACE-414

#	Requirement	Questions	Answered	Successfully Answered	URL
17951	ACE-394	Currently the Workspace implementation is located in the REST client. The WebUI client has a similar functionality. Do you want to consolidate this as well? E.g. having a project assembling the ACE JAVA Client API & implementation and the 3 projects mentioned above are the user/consumer of such a project?	1	0	https://issues.apache.org/jira/browse/ACE-394
17952	ACE-390	Can we reuse the new Maven functionality of Bnd(tools) for this?	1	0	https://issues.apache.org/jira/browse/ACE-390
17953	ACE-390	Applied the patch from [~brampouwelse] in rev#1727101. We now only create POM files for the agent, do we need to create them for other bundles as well? [~marrs] WDYT?	0	0	https://issues.apache.org/jira/browse/ACE-390
17954	ACE-388	2 cents: On 1) I think this should not be limited to downloads. Any webservice should be able to defer an agent request. On 2) What will be our default implementation (Eg. Configurable max concurrent requests)? It should indeed be pluggable; a use-case discussed is lazy deployment store calculation. In this case the logic should issue a retry-after for targets request a version list where the calculated deployment store is out of date.	1	0	https://issues.apache.org/jira/browse/ACE-388
17955	ACE-382	Under ACE-347 we have added FeedBackHandler as a control delegate and the FeedbackChannel interface (see below) This allows a custom controller to lookup and write to channels as well as deciding to send the feedback to the server. In addition I propose; * What channels are available is a configuration aspect that must match channels on the servers * What the local log maximum size is before truncating is a configuration aspect (see ACE-381) * DefaultController will simply send feedback every sync iteration for all available channels Does this cover the "control all aspects" as mentioned in the description of this issue? {code} public interface FeedbackHandler { /** Returns the feedback channels names */ List<String> getChannelNames(); /** Returns the feedback channel for a name */ FeedbackChannel getChannel(String name); } {code} {code} public interface FeedbackChannel { /** Synchronizes the feedback with the current server. Ensures the server has at least as much feedback data as we do. */ void sendFeedback() throws RetryAfterException, IOException; /** Logs a new message to the channel.*/ void write(int type, Map<String, String> properties) throws IOException; } {code}	0	0	https://issues.apache.org/jira/browse/ACE-382
17956	ACE-381	Ronald, the patch does not apply to trunk. I tried it with: patch -p1 <file Could you take a look and supply an updated patch?	0	0	https://issues.apache.org/jira/browse/ACE-381
17957	ACE-380	Configuring multiple serverURLs was implemented under ACE-379. The (default) DiscoveryHandler will use these in order and (if not disabled) check if it can connect to them before returning. Thus if 2 serverURLs are configured the first will be used as long as it is available. Once the first become unavailable it tries the second etc. Two notes: * This allows load of a large set of agents to be balanced by configuration. It is not a round-robin implementation per agent. * This only checks if a serverURL can be connected to. It does not check/interpret statuscodes (eg. the "too busy". Thus the default implementation does not cover "if the first is busy, go on to the next", a custom DiscoveryHandler could. Should we incorporate some checking/interpreting of statuscode in the check? I am not convinced. WDYT? [~marrs] [~jajans]	1	0	https://issues.apache.org/jira/browse/ACE-380
17958	ACE-377	I've updated the patch to work with r1527533 (merging StatefulTargetRepositoryTest). However, the LogIntegrationTest does not compile, because of a package constraint on org.apache.ace.log . Could you have another look at that?	0	0	https://issues.apache.org/jira/browse/ACE-377
17959	ACE-375	Following the analysis, could you please start working on the implementation? I will post the analysis soon. :)	1	1	https://issues.apache.org/jira/browse/ACE-375
17960	ACE-347	I'm currently still evaluating ACE in total and especially the MA. Our environment is very maven centric. Therefore it's not very easy to understand and handle the current project structure and the built bundle artifacts and the run configurations. Therefore I agree to your idea of redesigning the ACE MA very much. I would also prefer your project structure instead of spreading the services over multiple projects and assembling them again to a single jar. I also found your launcher implementation very interesting and inspiriting. What do you think about implementing nearly all stuff regarding the artifact management outside of the OSGi container. I mean: within or as part of this launcher?? And providing a service as mentioned by you into the OSGi container? This would have several benefits: * separating artifact management (the artifact distribution part) from the deployment / installation * allowing to even update the felix version etc. * stopping and (re)starting the framework, even with different settings/properties (e.g. in case of failures: restarting with deleted bundle cache) I (we) have several conceptual ideas how such an interface could look like. How can I provide my suggestions? Thx Wilfried	1	1	https://issues.apache.org/jira/browse/ACE-347
17961	ACE-347	Thanks for the feedback guys! [~wilsib] I can see the value of totally separating the MA from the deployment and, as you say, being able to update the container as well or even running different containers. Another approach we discussed before is to let the MA fork an entirely separate process so that, even when the deployment crashes or exits, the MA lives on and could restart. It would potentially also allow updating the MA itself without affecting the deployment... However, there are some cases where we want to punch holes in the isolation. First, the DeploymentAdmin spec itself requires us to deploy Resource Processors in the deployment and these at least require deployment admin packages. Second, strictly speaking we should raise appropriate events during deployment through eventadmin (there are some spec issues here and therefore I do not really care too much for this one). Finally, I would like it to be possible for the deployment to talk to MA services (eg. write something in a logstore or trigger/authorize an update) and I would like the MA to be extensible/customizable (preferably through services).. The again, as Marcel states, the MA should also be able to deploy as a bundle in a standard container.. are we talking about one MA that can do it all or different forms of the smae core logic? I'm not sure yet :) [~marrs] Agreed.. I can start stripping/simplifying more pending final deliberations on the discussion above. But getting down to 100k will require some drastic measures :) Wrt eliminating CM/EA note that only the interface packages are in there. Do you mean even decoupling from these interfaces?	1	1	https://issues.apache.org/jira/browse/ACE-347
17962	ACE-343	Hi Bram, I dont think it is good to specify the version in "packageinfo", because mostly it's just pick up the version you specified for the packages in the appropriate Bndtools view. Unfortunality, Bndtools automatically creates a "packageinfo" file with "version=1.0" as soon as you export the package. I faced the same problem in the past, but prefer to use "package-info.java" to add javadoc for my packages as well. Maybe we should do the same? So as I said, I would still change versions for packages inside the Bndtools view and not in the "packageinfo" file. What do you think? Yes it is a workspace patch :) And I used subclipse to checkout the projects, because it is easier to just let it also create a patch, and had never problems until now :P The fact that "svn:propset" are inlcuded is, because I had to add "generated" folders to the ignoring list. Or not?	1	1	https://issues.apache.org/jira/browse/ACE-343
17963	ACE-343	Ok got it ;) Learned something new. Didnt know that it is also working if putting it to the "packageinfo" file. But should Bndtools not responsible to synch the versions between "packageinfo" and the version you put in for the exported package in the "bnd.bnd" file?! Another question: Can we put javadoc comments on package level into that file as well? (never did it before, just used .java or .html) :) With respect to patch files, haven't thought that there is a difference between the generated files and I will keep that in mind and try to not use subclipse.	1	1	https://issues.apache.org/jira/browse/ACE-343
17964	ACE-343	BTW: Will provide another patch soon ;)	1	1	https://issues.apache.org/jira/browse/ACE-343

#	Requirement	Questions	Answered	Successfully Answered	URL
17965	ACE-335	To start with your last question, you can attach a patch by pressing "More Actions" and then choosing "Attach Files". What kind of classes could not be found before you used a specific version of osgi.core, because I assumed that all OSGi releases so far have been backward compatible. Or was something even older than 4.1.0 being used? The changes so far make sense, Having the repository on the local drive is just a speed optimization and ideally we should still support both. However, we can do this one step at a time. Let's first make sure we migrate and can run the server and targets, plus all our tests, and then worry about the continuous build and how to setup everything for releases. Timing wise, it makes sense to do this migration now, since everybody will upgrade to Bndtools 2 now anyway.	0	0	https://issues.apache.org/jira/browse/ACE-335
17966	ACE-333	Do you want to use the sub-bundle feature from Bndtools?	1	1	https://issues.apache.org/jira/browse/ACE-333
17967	ACE-333	Got it. The only thing you have to be careful is if someone wants to use Maven for the Bundle build. Sub-bundles makes a headless build more difficult. But at this discussion: https://groups.google.com/forum/?hl=en&fromgroups=#!searchin/bndtools-users/sub$20bundles/bndtools-users/R3LnHluaTaM/kXuaBV2VxEMJ you explained a good strategy to solve that. Did you ever tried it?	1	1	https://issues.apache.org/jira/browse/ACE-333
17968	ACE-333	Great! Shall I take steps for this issue? Would write a list of components I see and their sub-bundles as a comment and after we discussed and adjusted that structure if needed, I can start to map it the current trunk. What do you think?	1	1	https://issues.apache.org/jira/browse/ACE-333
17969	ACE-333	Sure. So lets talk about it next week?	0	0	https://issues.apache.org/jira/browse/ACE-333
17970	ACE-316	Going towards R5 repositories seems an interesting direction. What would be the effects from a functional perspective?	0	0	https://issues.apache.org/jira/browse/ACE-316
17971	ACE-316	Comitted in r1468366. The OBR is now in control of the location on disk an this is reflected in the API. A create is will return an 201 with a location header identifying where the resource is located. The layout is based upon bsn/version (eg. org/apache/bla/org.apache.bla-1.0.0.jar). The default strategy for determining bsn/version is based upon the bundle manifest headers. If the resource is not a bundle a fallback to the now optional filename parameter is attempted expecting something like <bsn>(-<version>)?(.<extension>)?. If a bsn but no version is found the version defaults to "0.0.0". A complicating factor proved to be the VelocityArtifactPreprocessor that also uploads processed templates. As i duplicates the OBR upload logic and made assumption on that it could assume filename I had to make some behavioral changes there. As for configuration resources we must use the filename strategy mentioned above it does the following; Assuming the template is of form <bsn>-<version>.<extension> a generated variant will be of form <bsn>.<targetID>-<targetVersion>.<extension>. Questions: * OBR upload code is duplicate over the ArtifactRepositoryImpl and the VelocityArtifactPreprocessor passing the OBR around as a string uri. Should we not introduce a client service and/or will that be adressed when we introduxe real repository services? * Should the VelocityArtifactPreprocessor (mis)use the obr? We can not assume that that will work with an actual (upstream) repository service in the future and it is creating a lot of garbage in there. Why not keep those generated variants local?	1	0	https://issues.apache.org/jira/browse/ACE-316
17972	ACE-302	I definitely don't object to adding the lists to Nabble. Would others be interested in this as well?	1	1	https://issues.apache.org/jira/browse/ACE-302
17973	ACE-302	Since no one seems to be interested, what do you think [~marrs] to close this one?	1	0	https://issues.apache.org/jira/browse/ACE-302
17974	ACE-301	Is this still an issue? At present the build source/target is at 1.6 and I have no trouble building with this setting or even 1.7 :) My jdk: {code} "1.7.0_17" Java(TM) SE Runtime Environment (build 1.7.0_17-b02) Java HotSpot(TM) 64-Bit Server VM (build 23.7-b01, mixed mode) {code}	1	0	https://issues.apache.org/jira/browse/ACE-301
17975	ACE-289	what about "after->configureAdditionalServices" and "before->configureProvisionedServices"?	1	0	https://issues.apache.org/jira/browse/ACE-289
17976	ACE-289	Please can someone review and apply that patch that we can close this issue :)	1	0	https://issues.apache.org/jira/browse/ACE-289
17977	ACE-249	Can you take a look and fix this?	1	1	https://issues.apache.org/jira/browse/ACE-249
17978	ACE-232	Attached a simple patch that seems to do the job. Needs testing. Note that I had to do a few import/export tweaks that may or may not make sense eg. excluding condperm 1.1	0	0	https://issues.apache.org/jira/browse/ACE-232
17979	ACE-217	Could you take a look at this?	1	0	https://issues.apache.org/jira/browse/ACE-217
17980	ACE-217	2) The artifacts column shows all artifacts in the current artifact-repository and you can remove them at will using the 'X'. I can see how tree-views and filters may help, but there is a separate issue for that at ACE-319 1) The 'Repository view' in the "Add bundles" dialog is filtered against the current state of the artifact-repository. An initial approach is to allow deletion of these resources. There are some considerations though; * The fact that a resource is not referenced from the current artifact-repository doe NOT mean it is not referenced from older revision in the deployment-repository. Thus deleting a resource could break deployment packages for target at old/rollback versions. I see two ways to address this (at some point); First we could check all deployment revisions and state of targets and disallow deleting resources within some reasonable limit. However, this is probably pretty complex, expensive and hard to communicate to the user. The second approach is to introduce resource caching on the server so that it is no dependent on the user-managed / upstream repository when constructing a deployment package. I discussed this with Marcel before and IMHO that would be a far more robust approach that will also prepare us for multiple (upstream) repositories. * OBR management is somewhat scatted around the code base. In this case the OBR upload is encapsulated by the ArtifactRepositoryImpl and it does ot support delete. As proposed in ACE-341 the repository should become a first-class service allowing us to decouple repository management. Again something that will prepare is for multiple (upstream) OBRs where some will be manageable and some wont. This also relates to the R5 discussion at ACE-314. In the scope of this issues I propose to; * Allow deletes of resource in the OBR view because, as Paul argues, the OBR is private and are assumed to know what your are doing. * For the time being implement the delete functionality in the webui code, pending a OBR solution that covers the broader issues. WDYT?	0	0	https://issues.apache.org/jira/browse/ACE-217
17981	ACE-211	From that page; {quote} Binary Distributions What applies to canonical source distributions also applies to all redistributions, including binary redistributions: Any redistribution must obey the licensing requirements of the contents. The best way to do that will likely depend on the binary packaging form. When assembling binary distributions, it is common to pull in and bundle additional dependencies which are not bundled with the source distribution. These additional dependencies must be accounted for in LICENSE and NOTICE. As a result, the LICENSE and NOTICE files for a binary distribution may well differ from those in the source distribution it was built from. In any case, the principle remains the same: LICENSE and NOTICE must exactly represent the contents of the distribution they reside in. {quote} I interpret this as " we must maintain and include a LICENSE/NOTICE per bundle as well. Some may use a generic variant, some may be specific. Is that correct? The current build does not include any LICENSE/NOTICE into the bundles.	1	1	https://issues.apache.org/jira/browse/ACE-211
17982	ACE-194	Cool! That would be an interesting feature to add. Is there any progress on this to report?	0	0	https://issues.apache.org/jira/browse/ACE-194

#	Requirement	Questions	Answered	Successfully Answered	URL
17983	ACE-191	Three comments after reviewing this: Firstly, the ServerType (enum) class seems to be missing from this patch. Secondly, right now, you more or less "hardcode" the server types by explicitly creating an enum for them. That means that every time we want to add a new server type, we need to modify this code. I think we should try to make the server types configurable. Finally, there is a hardcoded link to an S3 file in the codebase right now. I understand that's where you now "parked" the GF zip file, but I think we should not hardcode such links. By solving remark #2 I'm hoping this one goes away as well. WDYT?	1	1	https://issues.apache.org/jira/browse/ACE-191
17984	ACE-191	I agree this shouldn't be hardcode. After discussing some options today I changed the amazon launcher to have more configuration options. With these changes, the code is completely generic as before, but can now be configured to download extra files from either the ACE obr or any external url. This is sufficient to bootstrap Glassfish, and I expect most other servers too. WDYT?	0	0	https://issues.apache.org/jira/browse/ACE-191
17985	ACE-191	The current composition of the ace-target-devserver always includes de amazon-nodelauncher. This component, on startup, always connects directly to amazon, with in my case, incorrect credentials. This gives a nasty and lengthy stacktrace in the logfile: 2011.10.27 14:24:29 WARNING - Bundle: org.apache.ace.nodelauncher.amazon - Invocation of 'start' failed. - org.jclouds.rest.AuthorizationException: POST https://ec2.us-east-1.amazonaws.com/ HTTP/1.1 -> HTTP/1.1 401 Unauthorized at org.jclouds.aws.handlers.ParseAWSErrorFromXmlContent.refineException(ParseAWSErrorFromXmlContent.java:117) (....) Is this as it should be ? I can image Ace being used without deployment to Amazon.... Is the idea to have a separate assembly for deployment to Amazon and for non-Amazon deployment, or will there be one Ace assembly for all kinds of deployments ? In the latter case, it might not be desirable to always connect directly to Amazon on the startup ?	1	0	https://issues.apache.org/jira/browse/ACE-191
17986	ACE-185	Thanks, Marcel. No problem about the empty blocks. I'm fully aware of different styles between people and between projects. Having a style guide around implies that it should be respected, thus my contributing the checkstyle file. Common sense is good and a zero warning policies evil (I know a place where some people really want that (and a zero FindBugs policy, too)). Interpreting the style guide, I've set many checks to produce "errors", rather than "warnings", since I did not know how seriously the style guide is to be taken inside the project. I encourage the Ace team to tweak those settings to their liking. My contribution is just a starting point. I suggest to leave at least the "tab" check and the "file header" check on "error" level. Placing the ZIP on the style guide page would not have been my first choice, though. I already mentioned that I had trouble finding that page. I fear that most people will miss it. Personally, I'd have added the two files to https://svn.apache.org/repos/asf/incubator/ace/trunk. There you're bound to stumble over it.	1	1	https://issues.apache.org/jira/browse/ACE-185
17987	ACE-180	Issue has been resolved for a long time, waiting for potential confirmation from the reporter. Closing it now.	0	0	https://issues.apache.org/jira/browse/ACE-180
17988	ACE-146	These features will probably be subsets of the total assemblies we use now for these different targets. We can leave out Pax Runner that currently bootstraps a target and loose the existing shell. Other than that, it's probably straightforward to come up with these features. Do we want to create a new module 'ace-karaf' as a home for these descriptors?	1	1	https://issues.apache.org/jira/browse/ACE-146
17989	ACE-136	Has long been moved to branches/legacy-ant so it's no longer part of the normal build.	0	0	https://issues.apache.org/jira/browse/ACE-136
17990	ACE-106	Is there a current version somewhere?	1	0	https://issues.apache.org/jira/browse/ACE-106
17991	ACE-103	Implemented, can someone else test this and give me some feedback?	0	0	https://issues.apache.org/jira/browse/ACE-103
17992	ACE-102	Issue has been resolved for a long time, waiting for potential confirmation from the reporter. Closing it now.	1	1	https://issues.apache.org/jira/browse/ACE-102
17993	ACE-101	Given the latest sources, this would now be possible. You can disable the default discovery service by adding "-Dorg.apache.ace.discovery.property=disabled" to the MA startup, and plug in your own discovery service by adding the MA-option "bundle=my.own.discovery.Activator" (given that the custom discovery bundle is added to the class path). For example: [localhost:~/]$ java -cp org.apache.ace.launcher-0.8.1-SNAPSHOT.jar:my.own.discovery-0.0.1.jar -Dorg.apache.ace.discovery.property=disabled bundle=my.own.discovery.Activator identification=MyTarget Does this cover this issue? If so, can it be closed?	0	0	https://issues.apache.org/jira/browse/ACE-101
17994	ACE-100	Issue has been resolved for a long time, waiting for potential confirmation from the reporter. Closing it now.	0	0	https://issues.apache.org/jira/browse/ACE-100
17995	ACE-94	This is somwhat strange, can you give a hint: which goals do you run ? e.g. "clean install" ?? Also, can you provide a full log? I am not fully aware that we do use 2.0.x incompatible plugins - but may be. Besides that, i would rather recommend running with Maven 3, as i don't see the benefit of using outdated Maven versions. Its a tool, its free and you get everything but old strange issues with 2.0.x versions. Its true, corporates cannot move with any Maven release cycle, but on the other hand, tools get better over time, you lose less time running into corner-issues, get better error messages, faster builds. -> I know, you just point to adding the "prerequisites" configuration. Just wanted to get behind the reason of keeping 2.0.x alive.	0	0	https://issues.apache.org/jira/browse/ACE-94
17996	ACE-94	> This is somwhat strange, can you give a hint: which goals do you run ? e.g. "clean install" ?? > Also, can you provide a full log? Attached > I know, you just point to adding the "prerequisites" configuration. Just wanted to get behind the reason of keeping 2.0.x alive. It is not by choice! I have to keep some software compatible with it. I switch between 2.0.11, 2.1.1. en 3.0 all day. My env just happend to be set to 2.0.11 when I tried to build Ace. So be my guest and add prerequisites 3.0. I would do the same for my package if the boss would let me ;)	0	0	https://issues.apache.org/jira/browse/ACE-94
17997	ACE-79	Angelo, could you take a look at these?	1	1	https://issues.apache.org/jira/browse/ACE-79
17998	ACE-79	A couple of tweaks I needed to to to get tests going under mvn... 1) Hack IntegrationTestBase by adding empty public cleanup() to get things to compile. Do you have a local modification somewhere? 2) Add explicit surefire plugin configuration doing two things: 2a) Tweak pax-runner config to run on integration-test phase instead of test phase so you will anoy less people with all this testing :) 2b) The parent brings in a testng dep which messes up the testrunner somehow. Disable it in the execution by specifying a non existing testNGArtifactName Now I still have some test failing but I'm not sure thats a mvn problem as I get the same exceptions in eclipse? Tests in error: foo(org.apache.ace.it.server.AssemblyConfigureTest) initializationError0(org.apache.ace.it.server.MinimalGatewayTest)	1	1	https://issues.apache.org/jira/browse/ACE-79
17999	ACE-78	This task indeed applies to the unit tests. Porting the integration tests is ACE-79. If you're confident all are ported, you can close the issue. Great work!	1	1	https://issues.apache.org/jira/browse/ACE-78
18000	ACE-74	dont we have that already ?	1	0	https://issues.apache.org/jira/browse/ACE-74
18001	ACE-71	Depending INFRA has just been resolved. Can someone with enugh rights create a CI build for both, the ant build as well as the maven build ? Once this is done, it would be great to have the maven artifacts being deployed somewhere (original reason of the hudson move).	0	0	https://issues.apache.org/jira/browse/ACE-71

#	Requirement	Questions	Answered	Successfully Answered	URL
18002	ACE-69	I've checked in some changes for this, but I'm not sure of the environment that the webapp is supposed to live in. Assuming it's a bundle, but there's some other strangeness that follows if the project type isn't 'war', so it's probably going to take a few extra steps to get the bundle manifest generated with the war. What is the typical deployment container for an OSGi-based webapp?	1	1	https://issues.apache.org/jira/browse/ACE-69
18003	ACE-69	Hi Brian, yeah, about the "packaging=war" in ace-webui: well, usually you can deploy wars right away with paxweb or new R4.2 (enterprise) web container impls. Though i've seen that the ant based web ui is packaged as a normal bundle, also setting their own osgi headers. So, i would do the same in our build: make it a normal bundle, put resources into src/main/resources instead src/main/webapp and try to tweak things the way the original looks like for now Sure, in the very end we could end up with a single war, too but then we change things already in this transition phase. wdyt ?	1	0	https://issues.apache.org/jira/browse/ACE-69
18004	ACE-67	Toni, are you still working on this? Currently, this is not a blocking issue for the Maven conversion, but I'm leaving it open in case you still intend to do this. If not, feel free to close it.	1	1	https://issues.apache.org/jira/browse/ACE-67
18005	ACE-62	Brian, I tried: git push git@github.com:topping/Apache-Ace-Mavenization.git But I got! ERROR: Permission to topping/Apache-Ace-Mavenization denied to jbonofre. fatal: The remote end hung up unexpectedly Could you give write permission on the repo for my github user (jbonofre) ? Thanks JB	1	1	https://issues.apache.org/jira/browse/ACE-62
18006	ACE-62	Hi Jean-Baptiste, I've added you and Tony. Can you confirm that you have access? Cheers, Brian	1	0	https://issues.apache.org/jira/browse/ACE-62
18007	ACE-62	Hi gang, sorry I've been AWOL, I'm just getting a new client squared away from a narrowly-avoided catastrophe. I'm not much of a git user, Toni, do you need to push something? I am cloned off github, and if I look at http://github.com/topping/Apache-Ace-Mavenization, I don't see any of your changes. I'm grateful you did the import and can manage that though, because I have no experience with this svn<->git coexistence thing yet. Otherwise, your approach looks great to me. Maybe we should use this jira issue for module synchronization so we don't step on each other. Question about source moves: Under SVN, the history is kept with a move, but will that history be kept in git or will it need to be done again manually in svn once we have the POM set finished? Just curious, seems like we can manage it in any event. Anyway, I am checked out, looking forward to putting in some time on this now that this client has some tasks to do on their end. Cheers, Brian	1	0	https://issues.apache.org/jira/browse/ACE-62
18008	ACE-61	New exception (used in previous patch) to simplify error handling in rest servlets. Once in a while, this exception should be (if finally accepted) staged up to a higher package. Will then probably be used by other ace rest servlets.	1	0	https://issues.apache.org/jira/browse/ACE-61
18009	ACE-56	How to use latest patch: - apply from ace root, you will get a new subproject called devtools - just src are included - compile-dependencies: java5, osgi core api - runtime dependencies: equinox 3.5+ or felix 2.0+ There is a unclear behaviour when using felix, which has to be debugged from a felix point of view (unrelated to ace). However both should work. I used equinox (for the time being).	1	0	https://issues.apache.org/jira/browse/ACE-56
18010	ACE-56	How relevant is this patch? It's been around for a while, and probably could be updated to work with the new Maven build, but mostly I use a remote debugger whenever I want to debug an ACE node (server or target). Moving to OSGi 4.2 and using the new launcher API should make this patch a lot easier though. Toni, do you still want to pursue this?	1	1	https://issues.apache.org/jira/browse/ACE-56
18011	ACE-46	There are three directions we could go in here: 1) Adding a feature to the "choose your bundle" dialog to import artifacts from Nexus. Each imported bundle is then copied to the "internal" OBR of ACE. 2) Adding a feature to the "choose your bundle" dialog to import artifacts from Nexus without copying them to the "internal" OBR of ACE. 3) Using Nexus as the "internal" OBR of ACE. I would prefer going with option 1, or maybe 2. 3 is more work since we would then have to support "uploading" bundles to Nexus too (and the "internal" OBR also supports adding different types of artifacts). A more advanced version of 1) would be a version where, if you select a bundle in Nexus, it: a) wraps it into a bundle if it is not already one (probably by using Pax URL), or b) analyze dependencies and present you with options to import those too, maybe even automatically create a group for all those dependencies. WDYT?	1	1	https://issues.apache.org/jira/browse/ACE-46
18012	ACE-33	I noticed that Pax Web 0.6.0 can find a port itself; might that help us in this case? see http://wiki.ops4j.org/display/paxweb/Pax+Web+-+0.6.0 , http://issues.ops4j.org/jira/browse/PAXWEB-126	1	0	https://issues.apache.org/jira/browse/ACE-33
18013	ACE-32	Patch adds two new things to the "ant package" phase for dev-server. I just went for one particuar target to let you cross check which direction we go and what to expect. 1. core/deploy/target/dev-server/paxrunner/ A folder with just a script that calls the original pax runner at runtime. just call "./run.sh" on the folder stated above. 2. core/deploy/target/dev-server/standalone this calls pax runner at build-time (inside ant) and let it construct a particuar platform output. The final folder contains routly (or should at least) an output that very much looks like the current target output. So, you also can call "run.sh" here. Look inside to see that there is no pax runner involved anymore. Benefit when compared to the current approach: framework independence Where is the configuration ? No real user configuration is done in the build scripts. All configurations are read from core/conf/<target>/ For Pax Runner we have two configurations there: --> platform.setup There you specify the normal pax runner options. This includes the selection of the framework (felix,equinox,knopflerfish,concierge) and its versions. --> platform.properties additional properties that are needed to be configured as "environment" properties (-D) like http.port and others.	1	0	https://issues.apache.org/jira/browse/ACE-32
18014	ACE-32	thanks for applying! I think now is also a good time to get rid of (now unused) checked in libraries like felix. Having them still in leads to the assumption that thats being used (which it is - at least shoud - not). wdyt ?	0	0	https://issues.apache.org/jira/browse/ACE-32
18015	ACE-29	Committed quite a big refactoring this morning. This should now work. Does anybody want to validate that it actually does (beyond "my machine" ;))?	0	0	https://issues.apache.org/jira/browse/ACE-29
18016	ACE-27	This patch adds a working bundle columns, which filters out all other artifacts than bundles. Next up: bundle creation.	1	1	https://issues.apache.org/jira/browse/ACE-27
18017	ACE-27	This patch adds the possibility to import bundles from an OBR running on the same server. Furthermore, some minor layout changes.	1	1	https://issues.apache.org/jira/browse/ACE-27
18018	ACE-23	Also, wouldn't it be better to group the related tasks for the GWT based web UI under a common WebUI Task ? That way we can see the relationship and also see your idea a bit better what is missing to satisfy the initial goal.	1	0	https://issues.apache.org/jira/browse/ACE-23
18019	ACE-23	@Angelo: Pax web extender does not exist anymore. Its now part of pax web: org.ops4j.pax.web-extender:pax-web-ex-war -> org.ops4j.pax.web:pax-web-extender-war But 0.7.0 has not been released yet. Should i bounce alin ??	1	0	https://issues.apache.org/jira/browse/ACE-23
18020	ACE-23	@Toni Right, I got this jar from the pax-web website. It would be nice if we could get a release, could you do the honors for that?	1	0	https://issues.apache.org/jira/browse/ACE-23
18021	ACE-20	About the off-the-shelf Felix: The Felix downloads website still pointed to 1.8.0. I built this framework from the svn tag, and we have some custom properties in the config.properties file. About the MultiFrameworkStarter: do you know whether something like Pax Runner can be made to run multiple frameworks in the same VM? That would solve the same problem.	1	1	https://issues.apache.org/jira/browse/ACE-20

#	Requirement	Questions	Answered	Successfully Answered	URL
18022	ACE-20	Yep i am actually using just Pax Runner/Exam to start Ace stuff.. but i am not 100% sure about what the MuliFrameworkStarter is about.. what is it good for ?	0	0	https://issues.apache.org/jira/browse/ACE-20
18023	ACE-19	@Marcel, are you doing this as part of the svn import or should i provide a patch ?	0	0	https://issues.apache.org/jira/browse/ACE-19
18024	ACE-19	Would be good if we also can get rid of the "liq" naming ? Or does it have any meaning in ace ? I would replace it with "core" if we stay with the current structure.. (server,liq/core,gateway). Same is true for liq referals in build.xml (project name) Patch not necessary i guess ;)	1	0	https://issues.apache.org/jira/browse/ACE-19
18025	ACE-18	Is it normal that i cannot assign tasks to myself ?	1	0	https://issues.apache.org/jira/browse/ACE-18
18026	ACE-18	@Marcel, what is about this one, do we have concrete idea on what is need in addition for this? Ideas ? You mentioned the automatic versioning thing from DevCon, but i don't see the short term releationship to this (simple) maven export.. hints? Would help out here if i would know "what to do" ;)	1	0	https://issues.apache.org/jira/browse/ACE-18
18027	AMQ-6585	[~chris@die-schneider.net], Do you have any thoughts on the best way to go about this? I think one issue is going to be that the spring-xbean bundle still has a dependency on Spring 3.2.x and there is no updated bundle for Spring 4.x so we may need to repackage it.	1	1	https://issues.apache.org/jira/browse/AMQ-6585
18028	AMQ-6585	I think it would be better to move the whole camel code into the camel repository. Not sure if the camel plugin would then still be feasible though. How important is the camel plugin for activemq users? I personally never used it and also do not know anyone who used it.	0	0	https://issues.apache.org/jira/browse/AMQ-6585
18029	AMQ-6585	I'm not really sure how important it is, I've never used it either. [~tabish121] and [~gtully], what are your thoughts on this?	1	1	https://issues.apache.org/jira/browse/AMQ-6585
18030	AMQ-6581	It's always tough to pick the right log level for some messages as the level really depends on the use case. In the case of expired messages, sometimes they are expected and other times it might indicate a larger issue that needs troubleshooting as in your case. It seems like logging expired messages at level WARN is too high to me even though, generally speaking, most messages would probably be expired on the broker side before dispatch so this scenario shouldn't happen too often. Would the INFO level be acceptable in your use case?	1	1	https://issues.apache.org/jira/browse/AMQ-6581
18031	AMQ-6571	Thank you for the quick turnaround - much appreciated. Is there a rough timeline for 5.15.0 vs. 5.14.4 (and other 5.14.x)?	1	1	https://issues.apache.org/jira/browse/AMQ-6571
18032	AMQ-6503	You need to have an ICLA on file to be granted access to modify the wiki. https://www.apache.org/dev/new-committers-guide.html#cla Have you submitted one yet?	0	0	https://issues.apache.org/jira/browse/AMQ-6503
18033	AMQ-6428	Can you squash the commits and rebase? After a quick glance this looks pretty good and I saw that you commited a similar patch to Artemis so it makes sense to have matching test utilities in this project as well.	1	1	https://issues.apache.org/jira/browse/AMQ-6428
18034	AMQ-6428	I tried squashing my two commits, but I think I screwed something up - the PR is showing a bunch of other commits from other people. I'm not exactly sure how to fix this other than closing the current PR and opening a new one. Is there an easier way?	1	1	https://issues.apache.org/jira/browse/AMQ-6428
18035	AMQ-6408	Commit 8e994366b3c5e7540296d80a12ae2569010743cb in activemq's branch refs/heads/activemq-5.14.x from [~tabish121] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=8e99436] https://issues.apache.org/jira/browse/AMQ-6408 Update to latest proton-j 0.14.0 (cherry picked from commit 1595912bbf3225c8d1a92cf9a0d54746756c6c12)	1	1	https://issues.apache.org/jira/browse/AMQ-6408
18036	AMQ-6403	Commit 5a874816b74a485822bcb60e1cf7fc550644d62a in activemq's branch refs/heads/master from [~gtully] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=5a87481] AMQ-6403 - add indexDirectory attribute to kahadb plist impl - settable via broker service tempDataStore	1	1	https://issues.apache.org/jira/browse/AMQ-6403
18037	AMQ-6403	Commit f0d898cc9d7cbeea19ee54334e504ac52ef854cc in activemq's branch refs/heads/activemq-5.14.x from [~gtully] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=f0d898c] AMQ-6403 - add indexDirectory attribute to kahadb plist impl - settable via broker service tempDataStore (cherry picked from commit 5a874816b74a485822bcb60e1cf7fc550644d62a)	1	1	https://issues.apache.org/jira/browse/AMQ-6403
18038	AMQ-6397	This change makes sense as the value should apply to both sends/receives. One thing I did notice is that with this change the timeout period will now be cached permanently since it will be set only once on the client creation which is then cached. The way it is currently constructed, the timeout is set on every send so in theory the timeout could be changed later via the setter. I don't really think this is an issue as the original intention was probably to configure the setting once and not change it later. What do you think [~tabish121] ?	1	1	https://issues.apache.org/jira/browse/AMQ-6397
18039	AMQ-6362	Commit a65f5e7c2077e048a2664339f6425d73948d71ce in activemq's branch refs/heads/master from [~gtully] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=a65f5e7] AMQ-6362 - merge duplicated code - fix regression in AMQ4889Test and ExceptionListenerTest	1	0	https://issues.apache.org/jira/browse/AMQ-6362
18040	AMQ-6349	I already have an edited version of the web console on my desktop. Is there a way to upload it?	1	1	https://issues.apache.org/jira/browse/AMQ-6349
18041	AMQ-6339	Commit 34d7b0bfcb3a3d61ffe961f53ed5824d54aac232 in activemq's branch refs/heads/master from [~tabish121] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=34d7b0b] https://issues.apache.org/jira/browse/AMQ-6339 Update to latest Netty 4.0.x release	1	1	https://issues.apache.org/jira/browse/AMQ-6339

#	Requirement	Questions	Answered	Successfully Answered	URL
18042	AMQ-6301	but was temporarily disabled because of the Karaf issues) The problem is now that spring dm is embedded, it avoids the Karaf issue but the OsgiBundleXmlApplicationContext still has a dependency on Spring and doesn't work properly with Spring 4. Specifically I'm seeing a NoSuchMethodError because it relies on a deprecated method that was removed. Spring DM probably needs to go away entirely for the correct fix. {noformat} java.lang.NoSuchMethodError: org.springframework.core.CollectionFactory.createConcurrentMap(I)Lorg/springframework/core/ConcurrentMap; at org.springframework.osgi.util.internal.ClassUtils.<clinit>(ClassUtils.java:164) at org.springframework.osgi.context.support.TrackingUtil.getService(TrackingUtil.java:160) at org.springframework.osgi.context.support.OsgiBundleXmlApplicationContext.lookupNamespaceHandlerResolver(OsgiBundleXmlApplicationContext.java:243) at org.springframework.osgi.context.support.OsgiBundleXmlApplicationContext.createNamespaceHandlerResolver(OsgiBundleXmlApplicationContext.java:201) at org.springframework.osgi.context.support.OsgiBundleXmlApplicationContext.access$000(OsgiBundleXmlApplicationContext.java:60) at org.springframework.osgi.context.support.OsgiBundleXmlApplicationContext$1.run(OsgiBundleXmlApplicationContext.java:124) at java.security.AccessController.doPrivileged(Native Method)[:1.8.0_60] at org.springframework.osgi.context.support.OsgiBundleXmlApplicationContext.loadBeanDefinitions(OsgiBundleXmlApplicationContext.java:119) at org.springframework.context.support.AbstractRefreshableApplicationContext.refreshBeanFactory(AbstractRefreshableApplicationContext.java:129) at org.springframework.context.support.AbstractApplicationContext.obtainFreshBeanFactory(AbstractApplicationContext.java:605) at org.springframework.context.support.AbstractApplicationContext.refresh(AbstractApplicationContext.java:509) at org.springframework.osgi.context.support.AbstractDelegatedExecutionApplicationContext.access$301(AbstractDelegatedExecutionApplicationContext.java:69) at org.springframework.osgi.context.support.AbstractDelegatedExecutionApplicationContext$1.run(AbstractDelegatedExecutionApplicationContext.java:186) at org.springframework.osgi.util.internal.PrivilegedUtils.executeWithCustomTCCL(PrivilegedUtils.java:85) at org.springframework.osgi.context.support.AbstractDelegatedExecutionApplicationContext.normalRefresh(AbstractDelegatedExecutionApplicationContext.java:182) at org.springframework.osgi.context.support.AbstractDelegatedExecutionApplicationContext$NoDependenciesWaitRefreshExecutor.refresh(AbstractDelegatedExecutionApplicationContext.java:89) at org.springframework.osgi.context.support.AbstractDelegatedExecutionApplicationContext.refresh(AbstractDelegatedExecutionApplicationContext.java:175) at org.apache.activemq.osgi.ActiveMQServiceFactory.updated(ActiveMQServiceFactory.java:102) at org.apache.felix.cm.impl.helper.ManagedServiceFactoryTracker.updated(ManagedServiceFactoryTracker.java:159)[7:org.apache.felix.configadmin:1.8.8] at org.apache.felix.cm.impl.helper.ManagedServiceFactoryTracker.provideConfiguration(ManagedServiceFactoryTracker.java:93)[7:org.apache.felix.configadmin:1.8.8] at org.apache.felix.cm.impl.ConfigurationManager$UpdateConfiguration.run(ConfigurationManager.java:1753)[7:org.apache.felix.configadmin:1.8.8] at	1	1	https://issues.apache.org/jira/browse/AMQ-6301
18043	AMQ-6301	Can we target this issue for ActiveMQ 5.14? I would hate to see it get released and have OSGi compatibility with Spring 4.X still remain broken.	1	1	https://issues.apache.org/jira/browse/AMQ-6301
18044	AMQ-6274	There is already a property called originalDestination that is part of a Message. It should be set by the CompositeDestinationBroker so I don't think this is necessary. Can you see if that works for you? We try not to add headers to messages if we don't have to so that there are no conflicts with user defined headers. Also, in general, it's best to include a test or two when submitting a patch to verify that behavior works as intended. But in this case I think what you are trying to accomplish should already exist by using the property I pointed out above.	1	1	https://issues.apache.org/jira/browse/AMQ-6274
18045	AMQ-6274	The basic use case is wiretapping JMS Destinations to feed copies of the messages to another system for auditing purposes. The composite destinations work great for the wiretap - I can use a pattern to specify a group of destinations and where to forward copies. The problem is tracing the data in the audit system - I need a way to tell where the message originally came from. I worked around not having this feature in the past by putting my own JMS User Property on every message when I published it, but that proved to be error prone (i.e. it wouldn't get done once in a while for new destinations). Adding this type of functionality to the broker will eliminate those user-errors. Does that makes sense?	1	1	https://issues.apache.org/jira/browse/AMQ-6274
18046	AMQ-6274	OK - I updated the PR. Will this make it into 5.13.4? The customer I'm doing this for is using 5.13.2 right now and they need this - I'd rather not make them wait until 5.14 if at all possible.	1	1	https://issues.apache.org/jira/browse/AMQ-6274
18047	AMQ-6239	Thanks for the patch. This makes a lot of sense to me, no real reason I can see to create an array list when we can just use iterators. The only thing would be is if for some reason we wanted the iterator to not see future modifications. (since items are copied into lists, any future adds or removes while using the iterator wouldn't be seen). What do you think [~tabish121]?	1	1	https://issues.apache.org/jira/browse/AMQ-6239
18048	AMQ-6239	By the same logic, should we handle the situation of the underlying list getting extended concurrently? If so, then the "if (!orderedPendingList.isEmpty())" optimization (which was already there) in the ctor should be removed. I wonder if we are contemplating a use case that's not actually required..	0	0	https://issues.apache.org/jira/browse/AMQ-6239
18049	AMQ-6239	Would it be possible the change goes into the next release? The issue wastes an amazing amount of CPU, as can be seen for example by attached JFR.	0	0	https://issues.apache.org/jira/browse/AMQ-6239
18050	AMQ-6192	Using the content length is a simple way to check if it is a bytes message or should be converted using a transformer. That being said, I don't see why the logic couldn't be changed to first check for a transformation header and to honor that if it exists before going the byte message route. If that header is specifically set it seems like it would be best to honor that regardless of whether or not the content-length is set. (as you pointed out the spec says it should include it anyways) [~tabish121], What do you think about this change?	1	1	https://issues.apache.org/jira/browse/AMQ-6192
18051	AMQ-6184	We need the best of both worlds, depending on the use case. The default is a SynchronousQueue as before. I added:{code} org.apache.activemq.transport.nio.SelectorManager.workQueueCapacity{code} when > 0, it will swap out the SynchronousQueue for a capacity limited queue with the specified workQueueCapacity In this way, a limited pool (one that is at maximumPoolSize) can queue up some work before getting rejected and falling back to use the calling thread.	0	0	https://issues.apache.org/jira/browse/AMQ-6184
18052	AMQ-6182	The security dept ask me to disable "TRACE method" on this port 61614. I saw that it's disable for default on port 8161 where are the webapps admin, hawtio, etc. (at org/eclipse/jetty/webapp/webdefault.xml for jetty webserver there is a <security-constraint>) Is possible disable it on port 61614 too ? Please could someone tell me how to do that ...	1	1	https://issues.apache.org/jira/browse/AMQ-6182
18053	AMQ-6182	Hi Christopher, I hope this is the version https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=473b328 which disables http trace method. Could you please help me to figure it out. Thanks in advance	1	1	https://issues.apache.org/jira/browse/AMQ-6182

#	Requirement	Questions	Answered	Successfully Answered	URL
18054	AMQ-6147	Did a quick test of Qpid JMS against a 5.14.0-SNAPSHOT build that sets the 'setEmitFlowEventOnSend' to false and the zero prefetch integration test 'JmsZeroPrefetchTest.testTwoConsumers' fails every time so there's some work needed to understand the nature of the change in behavior.	1	1	https://issues.apache.org/jira/browse/AMQ-6147
18055	AMQ-6135	Is there anything else I need to do to get this accepted?	1	0	https://issues.apache.org/jira/browse/AMQ-6135
18056	AMQ-6135	Is there a timeframe for getting this merged, rejected or at least commented upon?	1	0	https://issues.apache.org/jira/browse/AMQ-6135
18057	AMQ-6119	Can you indicate where the documentation was that you refer to, so that it can be updated? ActiveMQ has required at least Java 7 for a while now.	1	1	https://issues.apache.org/jira/browse/AMQ-6119
18058	AMQ-6118	When you say you tried to set it in the activemq file, do you mean activemq.bat? Because I think you can set ACTIVEMQ_SSL_OPTS in activemq.bat for Windows and it should get passed the java command.	1	1	https://issues.apache.org/jira/browse/AMQ-6118
18059	AMQ-6118	Thank you Christopher. Do you mean to set it like this in activemq.bat: set ACTIVEMQ_SSL_OPTS="-Dcom.sun.security.enableCRLDP=true -Docsp.enable=true -Docsp.responderURL=http://myOCSP-url" When I echo the ACTIVEMQ_SSL_OPTS, I get exactly my configured line. Could you please tell me when I have to execute the activemq.bat? Also I asked myself when does Activemq check the CRL via OCSP (when Client is connecting or earlier)? Sorry for that many questions but I got no luck on the mailinglists. many thanks in advance, Marko	1	0	https://issues.apache.org/jira/browse/AMQ-6118
18060	AMQ-6118	Okay, I've set the OCSP command in the activemq.bat and started the broker via activemq-admin.bat. Now my JVM has this arguments stored. I'm able to see it through jconsole. My Client, which is revoked by the ocsp responder, connects successfully as before - without any log entry. Also I can't see any try to connect my ocsp responder via http port 80 in my Firewall logs. Any ideas how to get it work? What could it be? (And how does it work with linux?) I've attached a screenshot showing my jconsole output with the jvm arguments successfully applied.	1	0	https://issues.apache.org/jira/browse/AMQ-6118
18061	AMQ-6118	Hey Christopher, thanks for your opinion. The post you linked in above, is from me. Could you tell me how to raise this request into "needs_review" section? Or how to create a feature request? I spent some time on an alternative idea with the static crl list. I could bring in an idea for adding the crl list loadup to the runtime configuration. So it's possible for the users to load it manually without broker restart. That's an alternative way to get client certificates managed. It would be very glad, when this feature would be reviewed in total.	0	0	https://issues.apache.org/jira/browse/AMQ-6118
18062	AMQ-6118	Wow, thank you! Dejan you're the man. Could you introduce me into the Windows Distribution? I'm confused setting the ACTIVEMQ_SSL_OPTS. Where do I have to set all the configurations in Windows Distribution? Thanks for reply. much regards	0	0	https://issues.apache.org/jira/browse/AMQ-6118
18063	AMQ-6106	Thanks for the pull request...this is a good start but the main thing missing would be a way to configure the BrokerService before calling start. Maybe having an initialization method that gets called before starting the broker would be a good idea. Each test typically has a different broker configuration and being able to do something like the following in a unit test would be necessary to make the Rule usable. {code:java} @Rule public EmbeddedActiveMQBroker broker = new EmbeddedActiveMQBroker() { @Override public void init() { brokerService.setKeepDurableSubsActive(false); } }; {code} [~tabish121], Do you think this is a good spot for this which is currently a new subproject called activemq-junit under tooling? It seems like a good place to me because then it can be included in other modules for testing like activemq-http or activemq-unit-tests as a testing dependency.	1	1	https://issues.apache.org/jira/browse/AMQ-6106
18064	AMQ-6106	Would it be possible to get this incorporated into the upcoming 5.11.x/5.12.x/5.13.x releases? I know it works with these because I started using it with 5.11.x. There was the issue with AdvisoryMessages not working when a memory store was used, but that's been resolved and that was the only issues I've seen with this. The reason I ask is I've used this quite a bit in my Camel tests and I've had to copy it into several source structures - I'd really like to start using this version instead.	1	0	https://issues.apache.org/jira/browse/AMQ-6106
18065	AMQ-6100	Can we close this one out now [~dejanb]?	1	0	https://issues.apache.org/jira/browse/AMQ-6100
18066	AMQ-6099	Also, I'm not convinced you'd ever have more than 2 versions at once; you'd have to have messages that were stuck in the store for longer than the lifetime of a particular version, and only when the from version, intermediate version, and to version all used different versions of OpenWire. I guess it's possible, but it doesn't seem very likely. Either way, though, once you can handle two versions, is the code to support three actually that complicated? Is there even any code at all, once you have support for two versions?	1	1	https://issues.apache.org/jira/browse/AMQ-6099
18067	AMQ-6007	I wrote a little test and I was able to verify this behavior. It looks like this is a result of a rewrite done with the SchedulerBroker. If a scheduledJobId is already set on message send, the broker won't schedule the message and will send it immediately. This was done to prevent duplicates from being scheduled when running a network of brokers and to prevent overwriting a scheduled job. This jobId needs to be preserved on the message for things like the management destination. In your specific scenario, I think the best thing to do is probably to configure camel to exclude that property when re-sending. Will this work for you? I will add some documentation to the scheduler page to highlight this so that users know to make sure that property is not set when sending a scheduled message either manually or some other process such as Camel doing it.	1	1	https://issues.apache.org/jira/browse/AMQ-6007
18068	AMQ-6001	Commit 789419ab9d87b735688adb5843427fb4deaea590 in activemq's branch refs/heads/activemq-5.12.x from [~tabish121] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=789419a] https://issues.apache.org/jira/browse/AMQ-6001 Refill credit at a slightly faster rate. (cherry picked from commit bbcd93803207f22d6c549bc79dd7d0c28e16be26)	1	1	https://issues.apache.org/jira/browse/AMQ-6001
18069	AMQ-5962	Commit caac4f784ff93de71b776cce19804907dda868ca in activemq's branch refs/heads/master from [~tabish121] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=caac4f7] https://issues.apache.org/jira/browse/AMQ-5962 http-client v4.5.1 now latest.	1	1	https://issues.apache.org/jira/browse/AMQ-5962
18070	AMQ-5958	Commit ad00e1fe52bba3efab77f528712953fa8e9666ca in activemq's branch refs/heads/master from [~tabish121] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=ad00e1f] https://issues.apache.org/jira/browse/AMQ-5958 Update commons pool2 to latest release: 2.4.2 as of now. This closes #139	1	1	https://issues.apache.org/jira/browse/AMQ-5958
18071	AMQ-5958	Commit 3ffbaf9d896f3bf16274438cdf53aac88bca26e9 in activemq's branch refs/heads/activemq-5.12.x from [~tabish121] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=3ffbaf9] https://issues.apache.org/jira/browse/AMQ-5958 Update commons pool2 to latest release: 2.4.2 as of now. This closes #139	1	1	https://issues.apache.org/jira/browse/AMQ-5958
18072	AMQ-5949	Understood. Could we instead adjust the log level statements for failed deliveries from INFO to WARN?	0	0	https://issues.apache.org/jira/browse/AMQ-5949
18073	AMQ-5941	[~ben_nisbet] I have made this change on master which is 5.13-SNAPSHOT. Would you be able to test this out and verify the change works for you? I don't have a Weblogic instance to test on. I can also backport this to the 5.12 branch if needed.	1	1	https://issues.apache.org/jira/browse/AMQ-5941

#	Requirement	Questions	Answered	Successfully Answered	URL
18074	AMQ-5915	[~gtully], The main issue with the RuntimeConfigurationPlugin is that a lot of times I'll configure a broker just using the BrokerService and not the xml config. I've been looking at how the Runtime plugin works to add virtual destinations and I'd probably use a similar strategy that the plugin uses (ie call reapplyInterceptor() in AbstractRegion to reapply the virtual destinations) but I was thinking about adding addVirtualDestination/removeVirtualDestination methods to BrokerService so I could programmatically modify them without having to use xml config. What do you think?	1	1	https://issues.apache.org/jira/browse/AMQ-5915
18075	AMQ-5889	http://activemq.2283324.n4.nabble.com/nio-buffer-memory-problem-td4712706.html could the changes in this ticket lead to the behavior in this problem?	1	0	https://issues.apache.org/jira/browse/AMQ-5889
18076	AMQ-5879	Pretty sure it would log the stack trace if the log level is set to debug, many folks find the full stack trace on those error annoying when they are at warn level as there are valid cases where the transport can throw out an error such as connection drop. Can you peek again at the code and make sure just turning on debug level logging isn't enough?	0	0	https://issues.apache.org/jira/browse/AMQ-5879
18077	AMQ-5867	[~gtully] I have the need to support Networkconnector (Network connector needs to use SSL client authentication). This is not working because the sslcontext is being shared between the serverfactory and client factory and there is no way to tell the java ssl api to use a specific key for the server socket. I am thinking of the following changes Enhance the SSLContext to enable two stores one for server another one for other uses namely Networkconntector. <sslContext> <sslContext keyStore="file:${activemq.conf}/broker.jks" keyStorePassword="test" trustStore="file:${activemq.conf}/broker.ts" trustStorePassword="test" serverKeyStore="file:${activemq.conf}/broker-server.jks" serverKeyStorePassword="test" serverTrustStore="file:${activemq.conf}/broker.ts" serverTrustStorePassword="test" /> If the server store is not configured use the default store. This preserves the backward compatibility. Enahnce the sslContext class to handle both stores appropriately. Let me know your thoughts Thanks -MM </sslContext>	1	1	https://issues.apache.org/jira/browse/AMQ-5867
18078	AMQ-5777	The latest patch looks good, one thing I think you ought to consider though. Now that you've simplified the checks for action and header length I think the value of the StompFrameSize inner class in StompWireFromat is pretty small and the whole thing could be replaced by a single long value in StompWireFormat (reset on return from unmarshal) and also in StompCodec which does currently reset a class wide object on codec reset. This avoids the need to keep calculating the length on every byte read from the socket stream and if kept as a class variable in StompWireFormat it avoids the garbage of creating a the size object on each unmarshal call. What do you think? Given that StompFrameSize is passed to a couple methods you might keep it and just have it be an accumulator or just use an AtomicLong.	1	1	https://issues.apache.org/jira/browse/AMQ-5777
18079	AMQ-5774	How would you differentiate the use case for this vs the current maxFrameSize setting that is available on the transports now?	1	1	https://issues.apache.org/jira/browse/AMQ-5774
18080	AMQ-5774	The use case is similar in that the max size would be set to prevent a DOS attack. However, {{maxFrameSize}} is specific to OpenWire and a max message size is a more generic setting that can apply to any protocol. What about other formats like AMQP, MQTT, STOMP, etc....would {{maxFrameSize}} still apply with those other formats?. Also, what is the relation of an OpenWire frame to a message? Is the frame size always going to be the same size as the message or are the values different? (ie can a message be chunked up into more than one frame) It seems more logical to me to apply the limit to a specific message size than to specify it for a specific protocol.	1	1	https://issues.apache.org/jira/browse/AMQ-5774
18081	AMQ-5751	[~xabhi], this issue is something that is useful to me as well so I went ahead and implemented this and submitted a pull request. If this commit is accepted then this feature can be enabled by setting a property on a destination policy, called "maxDestinations". This will limit the number of created destinations to that number all destinations under that policy. For example: {code:xml} <policyEntry queue="q.>" maxDestinations="500" memoryLimit="1mb"> <!-- The maximum number of queues for this policy will be limited to 500 --> </policyEntry> {code} Can a commiter review my pull request and let me know if any changes should be made to it? Thanks, Chris	1	1	https://issues.apache.org/jira/browse/AMQ-5751
18082	AMQ-5748	Hi Christopher, this is a substantial contrib and it looks good :-) My first thought was how different would the be from mKahaDB with a store per destination and the existing store size metric. I guess that would be more coarse, but it would be per destination. I think that would need a tweak on mKahaDb to isolate the store size but that would be trivial in comparison. How would that map with your use case? On the change to messageStore, if the intent is some metric that is accessed frequently, maybe introduce a destinationStatistic that is backed and initialised by the index. It would mean that the calculation is always up to date and there will be not store block on the index iterations. Am I correct in thinking this is a precursor to doing flow control based on total message size, if that is the case it extending the destination statistic to encompass this would make sense. That is the pattern for messageCount.	1	1	https://issues.apache.org/jira/browse/AMQ-5748
18083	AMQ-5748	[~gtully] and [~tabish121], I have updated my pull request to add a new statistics object, {{MessageStoreStatistics}}. This bean is now used to track both the message count and message size of the store (and can be used for future stuff too). I updated {{getMessageCount}} to now be non-blocking as it delegates to the statistics bean instead. Can one of you take a look and let me know what you think? I believe the counters should be updated in the proper locations (and recovered properly on restart) but I could have missed something. Right now the bean is only updated in Kaha and the Memory message store. Would it be possible to get the final version of this patch into 5.12? I also noticed that a unit test was failing, the testRecoveryAfterCorruptionCheckSumExistingIndex method in JournalCorruptionEofIndexRecoverTest. I looked at for a little while but I haven't figured out what is wrong with it yet. I'll keep looking at it but if you notice anything obvious before I figure it out let me know. Thanks, Chris	1	1	https://issues.apache.org/jira/browse/AMQ-5748
18084	AMQ-5748	Yeah that makes sense. Did any other tests fail for you? If you want I can go back through and lower the values in all the tests to where I know it will always work versus trying to guess. For example, I could lower the messageSize test assertion that you pointed out to "> 20000". Since 20 messages are being sent and the payload is 1000 bytes it will always be at least 20000 no matter what happens in the future.	1	1	https://issues.apache.org/jira/browse/AMQ-5748

#	Requirement	Questions	Answered	Successfully Answered	URL
18085	AMQ-5748	- Store COMMIT FAILED: java.lang.NullPointerException at org.apache.activemq.store.memory.MemoryMessageStore.removeMessage(MemoryMessageStore.java:99) at org.apache.activemq.store.memory.MemoryMessageStore.removeMessage(MemoryMessageStore.java:86) at org.apache.activemq.store.memory.MemoryTransactionStore$4.run(MemoryTransactionStore.java:375) at org.apache.activemq.store.memory.MemoryTransactionStore$Tx.commit(MemoryTransactionStore.java:107) at org.apache.activemq.store.memory.MemoryTransactionStore.commit(MemoryTransactionStore.java:272) at org.apache.activemq.transaction.LocalTransaction.commit(LocalTransaction.java:70) at org.apache.activemq.broker.TransactionBroker.commitTransaction(TransactionBroker.java:252) at org.apache.activemq.broker.MutableBrokerFilter.commitTransaction(MutableBrokerFilter.java:117) at org.apache.activemq.broker.TransportConnection.processCommitTransactionOnePhase(TransportConnection.java:499) at org.apache.activemq.command.TransactionInfo.visit(TransactionInfo.java:100) at org.apache.activemq.broker.TransportConnection.service(TransportConnection.java:334) at org.apache.activemq.broker.TransportConnection$1.onCommand(TransportConnection.java:188) at org.apache.activemq.transport.ResponseCorrelator.onCommand(ResponseCorrelator.java:116) at org.apache.activemq.transport.MutexTransport.onCommand(MutexTransport.java:50) at org.apache.activemq.transport.vm.VMTransport.iterate(VMTransport.java:270) at org.apache.activemq.thread.PooledTaskRunner.runTask(PooledTaskRunner.java:133) at org.apache.activemq.thread.PooledTaskRunner$1.run(PooledTaskRunner.java:48) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at java.lang.Thread.run(Thread.java:745) 2015-07-07 18:37:36,937 [/localhost#61-2] - WARN Service - Async error occurred: javax.transaction.xa.XAException: STORE COMMIT FAILED: Transaction rolled back at org.apache.activemq.transaction.LocalTransaction.commit(LocalTransaction.java:75) at org.apache.activemq.broker.TransactionBroker.commitTransaction(TransactionBroker.java:252) at org.apache.activemq.broker.MutableBrokerFilter.commitTransaction(MutableBrokerFilter.java:117) at org.apache.activemq.broker.TransportConnection.processCommitTransactionOnePhase(TransportConnection.java:499) at org.apache.activemq.command.TransactionInfo.visit(TransactionInfo.java:100) at org.apache.activemq.broker.TransportConnection.service(TransportConnection.java:334) at org.apache.activemq.broker.TransportConnection$1.onCommand(TransportConnection.java:188) at org.apache.activemq.transport.ResponseCorrelator.onCommand(ResponseCorrelator.java:116) at	0	0	https://issues.apache.org/jira/browse/AMQ-5748
18086	AMQ-5736	I would like to see this feature implemented as well but for now my project has worked around it by using Virtual Destinations to share subscriptions. ActiveMQ Artemis (the proposed name for the HornetQ donation) already supports JMS 2.0 so if Artemis becomes ActiveMQ 6 at some point in the future then ActiveMQ will get JMS 2.0 support at that point. Can someone from the PMC comment on whether or not there are any plans to add JMS 2.0 support to the 5.x broker? I'm assuming the answer is no since the current focus seems to be on getting Artemis to be the next generation of ActiveMQ.	0	0	https://issues.apache.org/jira/browse/AMQ-5736
18087	AMQ-5734	Dejan - can u peek at https://fisheye6.atlassian.com/changelog/activemq-git?cs=f5283a904589eb4a9344fe42885cdacf92de68dd i found that the retained message (body size ==0) was not being accounted for in the test. Can you verify that that message should be dispatched?	1	0	https://issues.apache.org/jira/browse/AMQ-5734
18088	AMQ-5706	That will be more difficult to achieve than it seems because durable topic subscribers share a single instance of the message. Flipping the redelivered flag for one would effect all of the rest. I guess the use case is valid for durable subs - but it would require some major surgery. Would virtual topics work in your use case?	1	0	https://issues.apache.org/jira/browse/AMQ-5706
18089	AMQ-5696	I'd like to see a set of messaging tools as well, but wouldn't want to pollute the activemq-client jar which should probably stay as slim as possible. What about a activemq-msgtool that uses activemq-client? It could combine simple put / get w/ the performance features? Other products have a simple msgput, msgget type cli that make it handy for admins to quickly assess brokers and broker networks	0	0	https://issues.apache.org/jira/browse/AMQ-5696

#	Requirement	Questions	Answered	Successfully Answered	URL
18090	AMQ-5692	{code:xml} <transportConnector name="stomp" uri="stomp://0.0.0.0:61613?transport.defaultHeartBeat=10000,0&transport.useKeepAlive=true&trace=true"/> {code} - Run consumer {code} perl ./AMQ-5692.pl consumer {code} - Run producer {code} perl ./AMQ-5692.pl producer {code} After consuming around 8000 messages, the consumer hangs and won't report any more messages consumed. Take a broker thread dump and observe two stuck threads as in (1) below. Despite the configured stomp heart-beat defaultHeartBeat=10000,0 the connection does not time out after 10 seconds and no keep-alive is sent either. Have not managed to reproduce the issue with the ActiveMQ stomp client library. (1) {code} "ActiveMQ BrokerService[localhost] Task-4" daemon prio=5 tid=0x00007fe9ca087800 nid=0x650b runnable [0x00000001192d9000] java.lang.Thread.State: RUNNABLE at java.net.SocketOutputStream.socketWrite0(Native Method) at java.net.SocketOutputStream.socketWrite(SocketOutputStream.java:113) at java.net.SocketOutputStream.write(SocketOutputStream.java:159) at org.apache.activemq.transport.tcp.TcpBufferedOutputStream.flush(TcpBufferedOutputStream.java:115) at java.io.DataOutputStream.flush(DataOutputStream.java:123) at org.apache.activemq.transport.tcp.TcpTransport.oneway(TcpTransport.java:176) at org.apache.activemq.transport.stomp.StompTransportFilter.sendToStomp(StompTransportFilter.java:98) at org.apache.activemq.transport.stomp.StompSubscription.onMessageDispatch(StompSubscription.java:103) at org.apache.activemq.transport.stomp.ProtocolConverter.onActiveMQCommand(ProtocolConverter.java:870) at org.apache.activemq.transport.stomp.StompTransportFilter.oneway(StompTransportFilter.java:62) at org.apache.activemq.transport.AbstractInactivityMonitor.doOnewaySend(AbstractInactivityMonitor.java:304) at org.apache.activemq.transport.AbstractInactivityMonitor.oneway(AbstractInactivityMonitor.java:286) at org.apache.activemq.transport.MutexTransport.oneway(MutexTransport.java:68) at org.apache.activemq.broker.TransportConnection.dispatch(TransportConnection.java:1419) at org.apache.activemq.broker.TransportConnection.processDispatch(TransportConnection.java:938) at org.apache.activemq.broker.TransportConnection.iterate(TransportConnection.java:984) at org.apache.activemq.thread.PooledTaskRunner.runTask(PooledTaskRunner.java:133) at org.apache.activemq.thread.PooledTaskRunner$1.run(PooledTaskRunner.java:48) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at java.lang.Thread.run(Thread.java:744) "ActiveMQ Transport: tcp:///192.168.178.28:60902@61613" daemon prio=5 tid=0x00007fe9ca086800 nid=0x3b0f waiting on condition [0x0000000118f2d000] java.lang.Thread.State: WAITING (parking) at sun.misc.Unsafe.park(Native Method) - parking to wait for <0x00000007c0ecc2f0> (a java.util.concurrent.locks.ReentrantLock$NonfairSync) at java.util.concurrent.locks.LockSupport.park(LockSupport.java:186) at	1	1	https://issues.apache.org/jira/browse/AMQ-5692
18091	AMQ-5692	Hi, It looks to me the soWriteTimeout way is just for this kind of issues and should be used. The inactivity monitor shouldn't do anything to terminate the connection when it is in writing status. The reason is that any connection in writing status means it is 'busy working', not inactive. What do you think? Howard	1	0	https://issues.apache.org/jira/browse/AMQ-5692
18092	AMQ-5678	I guess this is very similar to the AMQ-5522 case. Great idea though	1	1	https://issues.apache.org/jira/browse/AMQ-5678
18093	AMQ-5662	[~northlander], Do you mind doing a rebase on this and updating your PR now that 5.13.1 is out (there were some fixes done on the webconsole recently). This seems like a good candidate to be merged into a 5.14.0 release.	1	0	https://issues.apache.org/jira/browse/AMQ-5662
18094	AMQ-5659	thanks for the patch. One thought. Would it not be simpler to just break out of the loop on an exception? {code} } catch(IOException e) { break; // ensure we don't loop }{code}	1	0	https://issues.apache.org/jira/browse/AMQ-5659
18095	AMQ-5603	Is there any estimate for when this feature will be designed and implemented? Thanks -Bill	0	0	https://issues.apache.org/jira/browse/AMQ-5603
18096	AMQ-5603	[~ceposta] I extended scope to have a preallocate_async option - so that we fire a job to preallocate the next data file when we rotate. This seems to reduce latency jitter on publish but it needs validation on real hardware. On the kernel_copy strategy - I did not see any point to deleting/recreating the template file so that is changed. Can you peek at that? In an effort to keep this safe, the sync is very explicit in journal. The whole class was synchronized which had to change to allow contention on the nextWriteId between the appender and the async task.	1	1	https://issues.apache.org/jira/browse/AMQ-5603
18097	AMQ-5603	Commit 231679d936e986a00930a96cc859b0a5b6d77486 in activemq's branch refs/heads/master from [~gtully] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=231679d] https://issues.apache.org/jira/browse/AMQ-5603 - don't async preallocate to keep match data file count	1	1	https://issues.apache.org/jira/browse/AMQ-5603
18098	AMQ-5573	Can you please document this option on http://activemq.apache.org/tcp-transport-reference.html also. I had to dig a lot into google and ActiveMQ nabble to find this little option. I have just one question, since it is a transport connection will this affect the TCP and NIO transport connections also? Because for them we already have maxFrameSize Thanks, Abhi	1	0	https://issues.apache.org/jira/browse/AMQ-5573
18099	AMQ-5558	[~dejanb], There's a user thread about this change here http://activemq.2283324.n4.nabble.com/activemq-all-jar-tight-coupling-with-log4j-td4710022.html in regards to adding the log4j confuration to the activemq-all jar. Do you think it makes sense to provide the log4j conf file separately? I know it's kind of nice to also provide the config file for a default logging configuration. There's probably not really a good answer as some people would want it included and some people would want the config separate to customize it.	1	0	https://issues.apache.org/jira/browse/AMQ-5558
18100	AMQ-5554	I think there are 2 separate issues, this is the simplest one. Steps to reproduce: Plain vanilla Karaf (2.3.9 or 2.4.1). Install *only* the activemq-blueprint feature (otherwise we end up in the 2nd scenario, a bit more complex). ActiveMQ 5.11.0 behaves the same way. {code} karaf@root> features:chooseurl activemq 5.10.1 Adding feature url mvn:org.apache.activemq/activemq-karaf/5.10.1/xml/features karaf@root> features:install activemq-blueprint karaf@root> features:list | grep activemq [installed] [5.10.1] activemq-client activemq-core-5.10.1 ActiveMQ client libraries [installed] [5.10.1] activemq activemq-core-5.10.1 ActiveMQ broker libraries [uninstalled] [5.10.1] activemq-broker-noweb activemq-5.10.1 Full ActiveMQ broker with default configuration [uninstalled] [5.10.1] activemq-broker activemq-5.10.1 Full ActiveMQ broker with default configuration and web console [uninstalled] [5.10.1] activemq-camel activemq-5.10.1 [uninstalled] [5.10.1] activemq-web-console activemq-5.10.1 [installed] [5.10.1] activemq-blueprint activemq-5.10.1 {code} From the console deploy the broker-blueprint.xml (attached to the deploy directory. It will work. However, using a custom plugin (uncomment one of the two lines) gives different errors. Arguably, using the spring namespace is not a smart idea and shouldn't even be attempted.	0	0	https://issues.apache.org/jira/browse/AMQ-5554

#	Requirement	Questions	Answered	Successfully Answered	URL
18101	AMQ-5552	I agree with some of the comments, I am mostly against this proposal. Will try to explain why. I agree for instance that running tests for hours unreliably gives a bad first impression. Quite horrible, I'd think, and not only first impression, second and third too. This proposal only addresses releases from the title of this issue. Time is not a problem as this is a one time excercise. Any entity/company who uses internal builds have no problem running it overnight, however long it takes, as long as the result is pretty much guaranteed. For us as a community is it enough to say that it's a good release it passed the smoke testing? I would say not, the ASF guidelines say the same thing. If the PMC votes an official release, we guarantee that it actually works, which is demonstrated by the tests that pass. Those that don't should be fixed or removed (if they serve no purpose). I understand the need to be pragmatic. However, even if we run the tests with -fae, if tests fail in a <module>, tests in all other modules that depend on it will not be run. Therefore failing tests may hide other failing tests, very hard or impractical to discover. There is another practical reason. At some point, I introduced a regression while merging to a branch. One of the best ways to identify the source is to use git bisect. Given the number of commits I needed some 7 iterations. Now with activemq, where it takes some 10 mins for a build skipping tests, plus running the test(s) you're interested, it would take at least 15-20 min per iteration. While a significant chunk of time, at least it's feasible. If the tests you start with are unreliable, it's a much, much harder effort. There are other arguments, community related. I don't know what would constitute resolution of this issue. In my opinon certainly not something that would be at odds with the ASF guidelines for releases, or at least they way I understand them.	0	0	https://issues.apache.org/jira/browse/AMQ-5552
18102	AMQ-5537	I hope I was not too rude by specifying Fix Version as 5.11. Where is the place or the mechanism to discuss such enhancement requests?	0	0	https://issues.apache.org/jira/browse/AMQ-5537
18103	AMQ-5509	Using the code above to create a test program, I get the following reported via log4j: {noformat} 2015-01-08 15:47:38,890 INFO [com.artnaseef.RunBroker.main()] org.apache.activemq.store.SharedFileLocker - Database /tmp/amq5509/data/testAmq5509Broker/scheduler/lock is locked... waiting 10 seconds for the database to be unlocked. Reason: java.io.IOException: Failed to create directory '/tmp/amq5509/data/testAmq5509Broker/scheduler' {noformat} Is a similar message not logged, or are different results expected?	1	0	https://issues.apache.org/jira/browse/AMQ-5509
18104	AMQ-5503	Regarding failed ActiveMQ-PR-Build #20 (https://builds.apache.org/job/ActiveMQ-PR-Build/20/console): I have no idea whats going wrong here - i did not touched any of the files the job is complaining about. I did a resync by the following procedure: {code} git remote add upstream git@github.com:apache/activemq.git git fetch upstream git checkout trunk git merge upstream/trunk git push {code} What can i do?	1	1	https://issues.apache.org/jira/browse/AMQ-5503
18105	AMQ-5471	If this is just a service for hosting and serving XML files, why can't you just use an existing COTS product? Apache, for example, or one of the various products that allow you to host Maven repos? What justifies re-inventing the wheel here?	1	1	https://issues.apache.org/jira/browse/AMQ-5471
18106	AMQ-5471	If you just want to server static content, httpd, nginx are fine. If the content is more dynamic it starts to get more complicated. Sure you can map the http resource to a directory that is a git repository and then have hooks or some other external mechanism that deals with the fs content. I thought about maven repos, but I think they'd be overkill for this. If you want to use this service to provide not just configuration but also status of the brokers in the nob, in which case no cots product would help you. I don't see this as innovation in the wheel industry. Is there a simpler solution I miss?	1	1	https://issues.apache.org/jira/browse/AMQ-5471
18107	AMQ-5398	What is the impact of making the following change <broker xmlns="http://activemq.apache.org/schema/core" brokerName="localhost" dataDirectory="${activemq.base}/data" destroyApplicationContextOnStop="true"> <plugins> <discardingDLQBrokerPlugin dropAll="true" dropTemporaryTopics="true" dropTemporaryQueues="true" /> </plugins> </broker> What I really want is to put a policy whereby messages in DLQ older than 5 days are automatically deleted.	1	0	https://issues.apache.org/jira/browse/AMQ-5398
18108	AMQ-5383	Is it worth to invest that on ActiveMQ 6? there is some support required for reference countings on the server for temporary queues on the anonymous subscriptions... I'm not sure how difficult it would be to impleent that on ActiveMQ 6 (current master). Maybe it could be a result of the merging between HornetQ and ActiveMQ which we will be focusing now.	0	0	https://issues.apache.org/jira/browse/AMQ-5383
18109	AMQ-5378	I submitted the ICLA at 19/Sep/13. Is it necessary to submit a new one?	1	1	https://issues.apache.org/jira/browse/AMQ-5378
18110	AMQ-5378	Why is the service wrapper not needed anymore?	0	0	https://issues.apache.org/jira/browse/AMQ-5378
18111	AMQ-5378	Hi Marc, this still doesn't work properly on osx {code}mkdir: illegal option -- n usage: mkdir [-pv] [-m mode] directory ... INFO: Using default configuration (you can configure options in one of these file: /etc/default/activemq /Users/dejanb/.activemqrc) INFO: Invoke the following command to create a configuration file bin/activemq setup [/etc/default/activemq | /Users/dejanb/.activemqrc] INFO: Using java '/Library/Java/JavaVirtualMachines/jdk1.7.0_12.jdk/Contents/Home/bin/java' ERROR: '-n /Users/dejanb/workspace/oss/mq/assembly/target/apache-activemq-5.11-SNAPSHOT/bin/..//bin/activemq.jar' does not exist, define ACTIVEMQ_HOME in the config{code} Is there a real reason to change all that logic that deals with ACTIVEMQ_HOME? The current one works fine for years.	1	1	https://issues.apache.org/jira/browse/AMQ-5378
18112	AMQ-5378	Hi Marc, there's still problem with the proposed solution on osx. Also, I think we're adding more complexity to the already complex script, which already gave us problems in the past (as it's not easy to test it properly on all platforms). I'd like to propose a simpler solution that should still give us the desired results. It's presented in the attached patch (AMQ-5378.patch). The new file is added in the bin directory called "env" which contains default configuration. With this we have a simple file where users can tune jvm and other runtime options. This file is added to the list of files configuration files and will be executed by default. Setting up a global configuration is then just a matter of copying this file to the users folder or /etc. With this we simplify the main shell script and get the desired functionality without touching the other logic which has been stable for a while now. Also, we can maybe remove "setup" operation as it's just a matter of copying this script to the appropriate place. I haven't touch setup yet, just wanted to discuss general principle first. What do you think of this approach?	1	1	https://issues.apache.org/jira/browse/AMQ-5378
18113	AMQ-5378	Hi Dejan, thanks for spending time in the init-script infrastructure. I´m sure we are not far from a working solution. Which problems do you still experience on macosx? Can you send me the output of the problematic command? (send the file to ms-activemq@256bit.org) {code} /bin/sh -x ./activemq <the-problemtic-option> 2>&1|tee debug-the-nasty-initscript-improvement-of-marc.txt {code} It's a good idea to have default configuration in the distribution directory. I agree - that is more userfriendly. I will have a look for your patch at monday. (i tried to apply to my fork at https://github.com/breuninger-ecom/activemq, but it seems that you created it based on the original, outdated and buggy init script)	1	1	https://issues.apache.org/jira/browse/AMQ-5378
18114	AMQ-5378	BTW. Do you have a wiki account? I couldn't find it by your name.	1	1	https://issues.apache.org/jira/browse/AMQ-5378

#	Requirement	Questions	Answered	Successfully Answered	URL
18115	AMQ-5378	[~dejanb] No, i do not own a wiki account. How can i get one? (...where can i find the wiki login) Have a look at my patches - i fixed various problems i the old script: - the setup option was broken"bstat," - the "bstat" does not throw exception anymore if the broker is not running - the values of the default configuration creation were improved - it is now possible to chain various symlinks to the installation/the initscript (also classpath issues) /etc/init.d/activemq --> /opt/activemq/current/bin/activemq /opt/activemq/current/ --> /opt/activemq/apache-activemq-5.10.0/ - options like create, encrypt and decrypt are now working if the broker is not running - the help output was improved - if there is more than one defaults file, only the first one is loaded I think it would be a huge step back to restart with the inital state. I thought about the default configuration. I think we should not do this, because this renders automatic calculation of the current step to be useless. The creation of the old defaults file was broken, therefore strange if/then/else constructs were contained in the file created by the "setup" option. A working init script is a important thing for system administrators, i think it would be a good idea to get the changes upstream in a fast way. (Probably it's good idea to write a small integration test suite for the init script)	1	1	https://issues.apache.org/jira/browse/AMQ-5378
18116	AMQ-5378	Hi Marc, I cherry-picked and committed most of your changes. What I didn't commit is: - default configuration - I still think sourcing the separate script is the way to go - checkrunning changes - the commands like browse, bstat, are not limited to only working with local broker (you can use --jmxurl to connect them to remote ones), so it's not a job of a script to check if broker is running locally. We need to improve implementation of these commands, which should be done under a new Jira - changing the name of pid files - I know you mentioned multiple instances as the reason for this change. I'm not sure if you played with "create" command which should create a new "base" for multiple brokers if you want to do that. That should make those instance shell scripts also obsolete, as we can just copy "env" script to the new bin directory. I think we should focus on improving "create" functionality as the solution for multi-tenancy. - I think the test script is a great idea, but didn't include it for now, as there's a problem with it. There's general issue with using "sed -i" and "sed ~s" on osx (and probably other platforms), so we should avoid it. We should continue working on the good test script - I also removed setup command from the script. Thanks for your contribution so far. Please continue to provide fixes and improvements against the current trunk. But please, keep you patches focused on a single fix/improvement. It will make it easier to test and merge. I'll try to include some of the QA folks from RedHat to test the script on different platforms and verify it. For the future (especially for 6.x), I'd like us to get more in line of the Apollo and Karaf scripts. As for the Cygwin, as you managed to make it work, I'd leave it for now and see how to best support windows users in the future.	1	1	https://issues.apache.org/jira/browse/AMQ-5378
18117	AMQ-5344	Is there any reason why the suggested patch (https://github.com/apache/activemq/pull/45) has never been merged to trunk?	1	0	https://issues.apache.org/jira/browse/AMQ-5344
18118	AMQ-5294	Have made the necessary changes and will send a pull request. The logic that has been implemented attempts to make the configuration as minimal and sane as possible without breaking existing workflows. The logic works by deferring the sampler behaviour to client (producer/consumer) settings by default. If the client asks to send/receive messages by count, both the duration of the samplers and the ramp up/down times are ignored; this is due to the non-deterministic nature of the completion time. If the client asks send/receive messages by time, the default behaviour for a sampler is: * use the client duration as its duration, unless explicitly overridden (e.g. tpSampler.duration=15000) - then it's the user's call * set the ramp up time to 0, and the ramp down time to 0 A user may override the default behaviour by either: * setting the rampUpTime/rampDownTime - this is the highest level override * setting two new properties rampUpPercent/rampDownPercent. These properties cannot exceed 100 in total, and boundary checks are in place.	1	0	https://issues.apache.org/jira/browse/AMQ-5294
18119	AMQ-5272	Ok, I can take a stab at it when I have some spare time. Be advised that any patch will necessarily be somewhat partial, as the small solution presented here is only intended to handle our own STOMP requirement (not MQTT), will not provide documentation, and will not necessarily fit neatly into the larger conceptual/architectural framework of ActiveMQ (with which I am less than familiar). Will this be valuable? It was my hope that this example would illuminate the logic and necessity and motivate the development of a fully fleshed out first party solution along the same lines. I cannot commit to developing such a solution myself unfortunately.	0	0	https://issues.apache.org/jira/browse/AMQ-5272
18120	AMQ-5210	Is anyone working on this issue. Or anyone has idea, how to solve this issue. Due to this issue, we are not able to use to activeMQ - network broker and finally we are not able to scale it	0	0	https://issues.apache.org/jira/browse/AMQ-5210
18121	AMQ-5192	Hi [~tabish121], Thanks for your update on this one. Are we correct in thinking that this would only tell us how many times a particular job had fired? We're looking for a count of messages scheduled for each delivery, per queue. Thanks	1	0	https://issues.apache.org/jira/browse/AMQ-5192
18122	AMQ-5192	Hi Tim, What if we implemented the scheduler persistence with JDBC, would it then be possible to get the count of messages scheduled for re-delivery per queue? See AMQ-5238.	0	0	https://issues.apache.org/jira/browse/AMQ-5192
18123	AMQ-5191	And btw where in the user manual do you intended this should be added? eg can you provide a link etc?	0	0	https://issues.apache.org/jira/browse/AMQ-5191
18124	AMQ-5157	Did anyone get a chance to look at this ? Thanks, Anuj	1	0	https://issues.apache.org/jira/browse/AMQ-5157
18125	AMQ-5137	While it may make sense to have true be the default, this is a significant change. Anyone using multiple consumers distributed across a NOB in order to load-balance will find themselves with a big change between 5.9 and 5.10 as none of their cross-network consumers are sharing load from local producers. Depending on the topology, that could - in the worst case - cut down total consumption by a large amount. Especially if anyone has a topology with few local consumers... I recommend not changing the default. It's configurable. Is there a real need here?	0	0	https://issues.apache.org/jira/browse/AMQ-5137
18126	AMQ-5114	The title of this JIRA refers to AbortSlowAckConsumerStrategy, whereas the commit was in AbortSlowConsumerStrategy, so it affects both classes. When I first read the title, I assumed that the bug didn't apply to my 5.8.0 broker, since AbortSlowAckConsumerStrategy wasn't added until 5.9.0 and the Affects Version is 5.9.0. Looking at the commit, it looks like this would affect both slow consumer strategies, and that it would affect and pre-5.9.0 versions; can we update the title and Affects Version to clarify its scope, to avoid misleading future readers?	0	0	https://issues.apache.org/jira/browse/AMQ-5114

#	Requirement	Questions	Answered	Successfully Answered	URL
18127	AMQ-5091	ActiveMQ standalone has this problem with the 1.2 jolokia {code} WARN | Error for /api/jolokia java.lang.IllegalAccessError: tried to access method org.json.simple.JSONValue.escape(Ljava/lang/String;Ljava/lang/StringBuffer;)V from class org.json.simple.JSONObject at org.json.simple.JSONObject.toJSONString(JSONObject.java:116) at org.json.simple.JSONObject.toJSONString(JSONObject.java:101) at org.json.simple.JSONObject.toJSONString(JSONObject.java:108) at org.jolokia.http.AgentServlet.handle(AgentServlet.java:295) at org.jolokia.http.AgentServlet.doGet(AgentServlet.java:249) at javax.servlet.http.HttpServlet.service(HttpServlet.java:693) at javax.servlet.http.HttpServlet.service(HttpServlet.java:806) at org.eclipse.jetty.servlet.ServletHolder.handle(ServletHolder.java:652) at org.eclipse.jetty.servlet.ServletHandler.doHandle(ServletHandler.java:445) at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:137) at org.eclipse.jetty.security.SecurityHandler.handle(SecurityHandler.java:521) at org.eclipse.jetty.server.session.SessionHandler.doHandle(SessionHandler.java:227) at org.eclipse.jetty.server.handler.ContextHandler.doHandle(ContextHandler.java:1044) at org.eclipse.jetty.servlet.ServletHandler.doScope(ServletHandler.java:372) at org.eclipse.jetty.server.session.SessionHandler.doScope(SessionHandler.java:189) at org.eclipse.jetty.server.handler.ContextHandler.doScope(ContextHandler.java:978) at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:135) at org.eclipse.jetty.server.handler.HandlerCollection.handle(HandlerCollection.java:154) at org.eclipse.jetty.security.SecurityHandler.handle(SecurityHandler.java:521) at org.eclipse.jetty.server.handler.HandlerCollection.handle(HandlerCollection.java:154) at org.eclipse.jetty.server.handler.HandlerWrapper.handle(HandlerWrapper.java:116) at org.eclipse.jetty.server.Server.handle(Server.java:367) at org.eclipse.jetty.server.AbstractHttpConnection.handleRequest(AbstractHttpConnection.java:486) at org.eclipse.jetty.server.AbstractHttpConnection.headerComplete(AbstractHttpConnection.java:926) at org.eclipse.jetty.server.AbstractHttpConnection$RequestHandler.headerComplete(AbstractHttpConnection.java:988) at org.eclipse.jetty.http.HttpParser.parseNext(HttpParser.java:640) at org.eclipse.jetty.http.HttpParser.parseAvailable(HttpParser.java:235) at org.eclipse.jetty.server.AsyncHttpConnection.handle(AsyncHttpConnection.java:82) at org.eclipse.jetty.io.nio.SelectChannelEndPoint.handle(SelectChannelEndPoint.java:628) at org.eclipse.jetty.io.nio.SelectChannelEndPoint$1.run(SelectChannelEndPoint.java:52) at org.eclipse.jetty.util.thread.QueuedThreadPool.runJob(QueuedThreadPool.java:608) at org.eclipse.jetty.util.thread.QueuedThreadPool$3.run(QueuedThreadPool.java:543) at java.lang.Thread.run(Thread.java:744) {code}	1	1	https://issues.apache.org/jira/browse/AMQ-5091
18128	AMQ-5077	some thoughts - without having looked at the particulars. for async producers - are you setting a producerWindow? that should allow pending sends to accumulate broker side. I think the root problem is ack contention for the index lock - some sort of ack batching would help there. one other thought - for composite dests - we could send to each dest in parallel - I don't think we do atm.	1	1	https://issues.apache.org/jira/browse/AMQ-5077
18129	AMQ-5077	Hi, We weren't setting a producer window before, so whatever's default was used. Have tried some quick tests now - setting abnormally high/low sizes have no effect (1k, 10k, 100k, 1mb, 10mb). Is this used if flow control is disabled? Client side, we've seen calls blocked at the network level, rather than any AMQ specific code - i think regardless of the window size, it's not been hit at the time the network call is blocked. I've been using mostly NMS clients for the testing as we'd thought the problem was initially in the NMS library - i'll get a java test set up next week and go through the code path it takes and where the time is spent waiting. I think we can either say that singular acks contending on the index lock is a problem, or that we write the messages to the index too quickly in the first place - if we just want to handle the ideal case, where consumers are able to ACK the message near-instantaneously, then we could remove the need to write the messages and then remove them in short order from the index just by delaying the writes by the roundtrip time from consumer to broker. It might reduce durability, but we've already accepted that tradeoff with other settings - so it makes sense to me to try and optimise away the index writes/contention, rather than attempt to batch up the acks, which might cause less contention but probably guarantees that the messages are written in the first place. Thanks, Richard	1	0	https://issues.apache.org/jira/browse/AMQ-5077
18130	AMQ-5077	Attached is an XML config which i think shows the issue more clearly. There is a topic with routings set up to 5 queues: TopicRouted.5.Embedded There is an embedded bean which uses a VM transport to publish from TopicRouted.5 to TopicRouted.5.Embedded If i publish to either, eventually the messages will end up on the correct queues 1-5. However, if i publish to TopicRouted.5, i'm only limited by the broker receive rate, and eventually the memory limit on the topic (~13k published/second until this point), while if i publish to TopicRouted.5.Embedded i'm limited by the speed at which the broker is able to persist messages to the queues (~3k/sec published). While i understand that this is likely to lead to message loss if the broker were to restart while the embedded router had messages pending, it would also allow us to scale a lot better with bursts of messages. Is this expected behaviour? Do you know of any way to make the broker behave like this by default? Thanks, Richard	1	0	https://issues.apache.org/jira/browse/AMQ-5077

#	Requirement	Questions	Answered	Successfully Answered	URL
18131	AMQ-5077	Leaving aside the question of message loss on failover, We have 2 goals here: - getting the maximum possible throughput to consumers - Never blocking/delaying a producer until the JMS hits an understood/visible limit (memory/diskstore). I need to come up with a better test case to see how a larger ProducerWindowSize affects this, but for the moment i don't believe it's working as we would want it to. Currently the limit on the rate which we're able to deliver messages from producers to consumers is the speed at which the JMS can write/remove messages from the index & diskstore. This happens in such a way that producers block on the send() call. org.apache.activemq.store.kahadb.KahaDBStore: public Future<Object> asyncAddQueueMessage(final ConnectionContext context, final Message message) public void removeAsyncMessage(ConnectionContext context, MessageAck ack) My reading of the code is that messages can be dispatched before the store task has completed, and if the ACK arrives before the store completes, then the store operation is cancelled. This also implies that a message could be delivered without being written to disk. I'm not sure at what point in this process the producer receives the ACK. If the consumer were quick enough to receive, process and ACK the messages in question, then we'd optimise away the need to ever write to the diskstore, and not have an issue here. However, our SAN seems to be fast enough, combined with network latency, to ensure that the disk writes are nearly always in progress before the ACK arrives. In this case, all the work to write/remove messages from the diskstore & index, and the synchronisation overheads of doing this, happen on the NIO worker threads. This delays the producers in a way that isn't visible to the producer. Whether sending messages synch or async, all the producer code sees is that calling send() takes longer. My understanding is that increasing the producer send window would allow it to keep more messages in flight before it has received ACKs for them - but would not help when it's blocked at the network level. I'll see if i can come up with a more specific test case that shows the effect of varying the producer send window. What i think we're looking for is some option along the lines of ConcurrentDispatchThenStoreIfNeeded - first dispatch the message, then wait for a timeout period, and then only persist the message to the diskstore incurring the disk/synchronisation penalties if an ACK doesn't arrive on time. This would be for a low (<100ms?) time, respect memory limits on the broker for total messages in flight, and would allow the producer send rate to scale with the slowest consumer receive speed, rather than the sum of all queue writes possible on the JMS. Current behaviour: Producer -> broker with topic -> queue routings -> consumers: - Producer is blocked by speed at which broker can write to all queues. Consumers receive messages at speed JMS can write. Queue write limit is global. Producer -> broker with embedded routing bean -> consumers (router waits for send call to complete before acking the message received): - Producer is able to write up to producer window size. Embedded bean is blocked by broker write speed - consumers receive messages at speed JMS can write. Queue write limit is global. Ideal situation: Producer -> broker with topic -> queue routings and ConcurrentDispatchThenStoreIfNeeded set: Producer is able to write up to producer window size. Broker is able to dispatch to consumers at consumer receive rate limit, only writing to disk if consumers slow down. Does that make sense, or do you think i'm misunderstanding the issue here? Thanks, Richard	1	0	https://issues.apache.org/jira/browse/AMQ-5077
18132	AMQ-5077	producer catch it in code like: {noformat} "main" prio=10 tid=0x000000000bc3b000 nid=0x4109 runnable [0x0000000041ebe000] java.lang.Thread.State: RUNNABLE at java.net.SocketOutputStream.socketWrite0(Native Method) at java.net.SocketOutputStream.socketWrite(SocketOutputStream.java:113) at java.net.SocketOutputStream.write(SocketOutputStream.java:159) at org.apache.activemq.transport.tcp.TcpBufferedOutputStream.flush(TcpBufferedOutputStream.java:115) at java.io.DataOutputStream.flush(DataOutputStream.java:123) at org.apache.activemq.transport.tcp.TcpTransport.oneway(TcpTransport.java:176) at org.apache.activemq.transport.AbstractInactivityMonitor.doOnewaySend(AbstractInactivityMonitor.java:304) at org.apache.activemq.transport.AbstractInactivityMonitor.oneway(AbstractInactivityMonitor.java:286) at org.apache.activemq.transport.TransportFilter.oneway(TransportFilter.java:85) at org.apache.activemq.transport.WireFormatNegotiator.oneway(WireFormatNegotiator.java:104) at org.apache.activemq.transport.failover.FailoverTransport.oneway(FailoverTransport.java:658) - locked <0x000000050f60c5e8> (a java.lang.Object) at org.apache.activemq.transport.MutexTransport.oneway(MutexTransport.java:68) at org.apache.activemq.transport.ResponseCorrelator.oneway(ResponseCorrelator.java:60) at org.apache.activemq.ActiveMQConnection.doAsyncSendPacket(ActiveMQConnection.java:1321) at org.apache.activemq.ActiveMQConnection.asyncSendPacket(ActiveMQConnection.java:1315) at org.apache.activemq.ActiveMQSession.send(ActiveMQSession.java:1853) - locked <0x000000050f60c668> (a java.lang.Object) at org.apache.activemq.ActiveMQMessageProducer.send(ActiveMQMessageProducer.java:289) at org.apache.activemq.ActiveMQMessageProducer.send(ActiveMQMessageProducer.java:224) at org.apache.activemq.ActiveMQMessageProducerSupport.send(ActiveMQMessageProducerSupport.java:269) {noformat} My understanding of flow control & the producer window size: Client side: - window size is set. - Before each send, current size of all messages in flight is checked to see if window is exceeded. - if producerWindow.waitForSpace() doesn't block, then the message is sent. - After the message is sent, the producer in flight size is incremented by the message size (and decremented when the ack is received). Broker side: - Each queue has a memory limit set, as well as overall memory limit and disk store limit. - For each message dispatched for a given queue, each of these limits is checked. - if any limit is set and sendFailIfNoSpace is set to true, the producer should get an exception sent back. In none of my tests have i caught any thread stuck inside the flow control handling logic. In all cases they're inside network code - producer side as above, broker side in something like: {noformat} "ActiveMQ NIO Worker 29" daemon prio=10 tid=0x000000001775d000 nid=0x6a0d runnable [0x000000004473d000] java.lang.Thread.State: WAITING (parking) at sun.misc.Unsafe.park(Native Method) - parking to wait for <0x0000000638b62430> (a org.apache.activemq.store.kahadb.KahaDBStore$StoreQueueTask$InnerFutureTask) at java.util.concurrent.locks.LockSupport.park(LockSupport.java:186) at java.util.concurrent.FutureTask.awaitDone(FutureTask.java:425) at java.util.concurrent.FutureTask.get(FutureTask.java:187) at	1	0	https://issues.apache.org/jira/browse/AMQ-5077

#	Requirement	Questions	Answered	Successfully Answered	URL
18133	AMQ-5077	Hi, I'm happy that those 2 changes would help the producer in not getting blocked by the JMS. It is also likely to improve throughput to the consumers, but i think we could make further improvements there. A new executor using Concurrent store and dispatch is likely to result in more messages in flight to the consumers - but i'm still concerned that the performance of this (this being the ultimate queue writes/sec the broker can achieve) will be hard to determine, as it will be a function of how quick the underlying disk store is as well as the average roundtrip time for the consumer to receive, process and ACK each message. If we have a large queue of messages being written to the diskstore, and relatively quick consumers, then we can optimise away the disk writes - but i'm not sure that this is visible at the moment. I would still be interested in i guess some form of delay queue for the diskstore writes - a configurable property for a minimum delay to wait before writing messages through to the index/diskstore, which you could benchmark against your expected consumer ACK roundtrip time to determine if you expect to be able to optimise away the disk writes completely under most situations. If this delay is small enough, and the producer window doesn't get decremented till either the diskstore write completes or the consumer ACK arrives, we could still have some resilience against message loss with this in place. Do you think that would be a useful option, or cause more problems than it could solve? Thanks, Richard	1	0	https://issues.apache.org/jira/browse/AMQ-5077
18134	AMQ-5030	Is there any reason why Camel cannot be upgraded to 2.13.1 instead of 2.12x?	0	0	https://issues.apache.org/jira/browse/AMQ-5030
18135	AMQ-5024	retrieve the message meta data as opposed to the actual payload which can be retrieved only as needed? Thank you Ashwin ------ ERROR | Failed to page in more queue messages java.lang.OutOfMemoryError: GC overhead limit exceeded at java.util.HashMap.createEntry(HashMap.java:869) at java.util.HashMap.addEntry(HashMap.java:856) at java.util.HashMap.put(HashMap.java:484) at org.apache.activemq.broker.region.cursors.OrderedPendingList.addMessageLast(OrderedPendingList.java:56) at org.apache.activemq.broker.region.Queue.doActualDispatch(Queue.java:1927) at org.apache.activemq.broker.region.Queue.doDispatch(Queue.java:1839) at org.apache.activemq.broker.region.Queue.pageInMessages(Queue.java:2003) at org.apache.activemq.broker.region.Queue.iterate(Queue.java:1491) at org.apache.activemq.broker.region.Queue.wakeup(Queue.java:1709) at org.apache.activemq.broker.region.PrefetchSubscription.acknowledge(PrefetchSubscription.java:416) at org.apache.activemq.broker.region.AbstractRegion.acknowledge(AbstractRegion.java:426) at org.apache.activemq.broker.region.RegionBroker.acknowledge(RegionBroker.java:408) at org.apache.activemq.broker.BrokerFilter.acknowledge(BrokerFilter.java:77) at org.apache.activemq.broker.TransactionBroker.acknowledge(TransactionBroker.java:287) at org.apache.activemq.broker.BrokerFilter.acknowledge(BrokerFilter.java:77) at org.apache.activemq.broker.MutableBrokerFilter.acknowledge(MutableBrokerFilter.java:87) at org.apache.activemq.broker.TransportConnection.processMessageAck(TransportConnection.java:508) at org.apache.activemq.command.MessageAck.visit(MessageAck.java:236) at org.apache.activemq.broker.TransportConnection.service(TransportConnection.java:329) at org.apache.activemq.broker.TransportConnection$1.onCommand(TransportConnection.java:184) at org.apache.activemq.transport.ResponseCorrelator.onCommand(ResponseCorrelator.java:116) at org.apache.activemq.transport.MutexTransport.onCommand(MutexTransport.java:50) at org.apache.activemq.transport.vm.VMTransport.iterate(VMTransport.java:241) at org.apache.activemq.thread.PooledTaskRunner.runTask(PooledTaskRunner.java:129) at org.apache.activemq.thread.PooledTaskRunner$1.run(PooledTaskRunner.java:47) at java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1145) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:615) at java.lang.Thread.run(Thread.java:724) WARN | /admin/browse.jsp java.lang.OutOfMemoryError: GC overhead limit exceeded at com.sun.el.parser.SimpleCharStream.<init>(SimpleCharStream.java:281) at com.sun.el.parser.SimpleCharStream.<init>(SimpleCharStream.java:287) at com.sun.el.parser.ELParser.<init>(ELParser.java:1484) at com.sun.el.lang.ExpressionBuilder.createNodeInternal(ExpressionBuilder.java:139) at com.sun.el.lang.ExpressionBuilder.build(ExpressionBuilder.java:188) at	0	0	https://issues.apache.org/jira/browse/AMQ-5024
18136	AMQ-5009	i updated the release page - the point of the single -all jar is just ease of use. But there is a maven ease of use and ease of exclusion use case for sure - maybe we can have both the -all.pom and the artifact. Is that possible?	0	0	https://issues.apache.org/jira/browse/AMQ-5009
18137	AMQ-5008	Hi. Why there is no feedback here? It is very important feature for "security freaks". Patch needs additional work as it uses label jumps (technique known from such languages like BASIC known from 1960s). But with a little tuning and some unit test it might be very useful and a kind of green light for "security freaks"- most often financial sector like banks. Of course once this functionality will be done, there will be a problem of knowing that certificated used to establish connection was compromised. So once established using not yet compromised certificate, client app can still work as long as will not disconnect. But this is second side of the problem. Greetings Piotr Klimczak	1	0	https://issues.apache.org/jira/browse/AMQ-5008
18138	AMQ-4940	Before I start any work -- would a patch for this (upgrading Jetty to the latest stable release) be welcomed? I just don't want to waste my time if it's something that isn't desired.	1	1	https://issues.apache.org/jira/browse/AMQ-4940
18139	AMQ-4940	Steve, how far did you get with this? I think the old version of jetty is causing issues with java 8 so was interested in trying a patch. Just found AMQ-5517 adding support for Jetty 9.	1	0	https://issues.apache.org/jira/browse/AMQ-4940
18140	AMQ-4903	Does this work with Karaf 2.x also?	1	1	https://issues.apache.org/jira/browse/AMQ-4903
18141	AMQ-4888	I guess since you turn off persistence you assume the broker is memory only? Yeah maybe the disk check should not happen if its not persistent? Not sure if we have improved this in 5.9? Can you try with that release?	0	0	https://issues.apache.org/jira/browse/AMQ-4888
18142	AMQ-4888	Well, it's more that it issues an ERROR, but then somehow works anyway (should that message rather be a WARN?). And perhaps the default size of 50GB is a bit high. I understand temp space may be used for swapping messages, perhaps I need to use VM cursor to avoid potential disk use. I also realize that when creating a Broker just with new BrokerService(), I have no idea how its default configuration looks like.	0	0	https://issues.apache.org/jira/browse/AMQ-4888
18143	AMQ-4885	Agreed. Fernando, can you clean up the patch? not sure how the package.html patches are necessary for this jira... as well as the binary diffs on the .class files...	0	0	https://issues.apache.org/jira/browse/AMQ-4885
18144	AMQ-4867	I am also using Master/Slave configuration for ActiveMQ broker and faced the similar issue. Is there a way to initialize the plugin for failover broker ?	0	0	https://issues.apache.org/jira/browse/AMQ-4867

#	Requirement	Questions	Answered	Successfully Answered	URL
18145	AMQ-4867	what is this plugin doing? A slave broker blocks on creating the persistence adapter which is typically locked by the master. When the persistence adapter starts then it starts the broker, pluigns and transport/network connectors. The idea is that a client can only connect to a broker when it is started and has exclusive access to the store. So any plugin will not be able to do much b/c the broker can't do much. You could init your plugin as a spring bean and have it created when the context is loaded but maybe a plugin is not the way to go. Please provide some more context on what you are trying to achieve and why two instances of your plugin logic need to coexist.	1	0	https://issues.apache.org/jira/browse/AMQ-4867
18146	AMQ-4867	Can you wire the monitoring heartbeats to a JMX operation? JMX mbean server is started before we start the persistence adapter so you can query the brokerView Mbean. The isSlave attribute will reflect the current state of the failover broker.	1	0	https://issues.apache.org/jira/browse/AMQ-4867
18147	AMQ-4848	GitHub user zmhassan opened a pull request: https://github.com/apache/activemq/pull/28 https://issues.apache.org/jira/browse/AMQ-4848 - added NetworkConnectors... Fixed activemq:list command and it now displays the following output in karaf: ```javascript JBossA-MQ:karaf@root> activemq:list TransportConnectors = {openwire=tcp://localhost:38009?maximumConnections=1000} BrokerName = amq-broker TransportConnectors = {openwire=tcp://localhost:61616?maximumConnections=1000} BrokerName = amq ``` You can merge this pull request into a Git repository by running: $ git pull https://github.com/zmhassan/activemq AMQ4848 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/activemq/pull/28.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #28 ---- commit 40b1909256522ccccbcf6e3547eae48039d097f7 Author: zmhassan <zak.hassan1010@gmail.com> Date: 2014-06-17T20:24:43Z https://issues.apache.org/jira/browse/AMQ-4848 - added NetworkConnectors to activemq:list command ----	1	1	https://issues.apache.org/jira/browse/AMQ-4848
18148	AMQ-4848	GitHub user zmhassan opened a pull request: https://github.com/apache/activemq/pull/59 https://issues.apache.org/jira/browse/AMQ-4848 - Karaf update added Netw... This was a fix that I sent pull request #28 which was from Jun 17, 2014. I fixed activemq:list command and it now displays the network connectors available for the broker on karaf. Can anyone review this patch. If the fix is good then can this be merged? ![screen shot 2015-02-11 at 4 14 05 pm 2](https://cloud.githubusercontent.com/assets/1269759/6156952/a8b3eae2-b20b-11e4-8cd1-7ae9bca809d0.png) Also I've updated the pax exam unit test so the integration test does not fail. ![screen shot 2015-02-11 at 4 26 39 pm 2](https://cloud.githubusercontent.com/assets/1269759/6156968/cfaef740-b20b-11e4-9dd7-1fe04a2e84b9.png) You can merge this pull request into a Git repository by running: $ git pull https://github.com/zmhassan/activemq AMQ-4848-master Alternatively you can review and apply these changes as the patch at: https://github.com/apache/activemq/pull/59.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #59 ---- commit 669f6078bf3c7d25ea346917e70913ca6472c06b Author: zmhassan <zak.hassan1010@gmail.com> Date: 2015-02-11T21:31:44Z https://issues.apache.org/jira/browse/AMQ-4848 - Karaf update added NetworkConnectors to activemq:list command ----	1	1	https://issues.apache.org/jira/browse/AMQ-4848
18149	AMQ-4831	I think the current implementation is ok. I agree that on the presentation part it would be the job of the console to display something appropriate and aesthetically pleasing. The mbean should supply an accurate value as it may be used by machines as well, not only human consumption. As far as jconsole goes, it presents the raw data it gets as expected. I'd suggest closing this as a won't fix, or move it to need-review. wdyt?	0	0	https://issues.apache.org/jira/browse/AMQ-4831
18150	AMQ-4831	Typically, I like to see formatting handled entirely on the UI front, and that applies here. Also, the display looks acceptable to me. With that said, tracking average message size as a real number seems excessive. Even for very small message sizes (say 1-10 bytes), it doesn't seem to add much value (would you monitor and have a threshold at 1.5 bytes?). So, changing it to an integer or long seems viable.	1	0	https://issues.apache.org/jira/browse/AMQ-4831
18151	AMQ-4825	Why should there be OpenShift/Fuse-specific logic in the ActiveMQ code? Why not use AMQ_PORT environment variable?	0	0	https://issues.apache.org/jira/browse/AMQ-4825
18152	AMQ-4819	Sounds good Matt, do you wanna take a stab at implementing this?	0	0	https://issues.apache.org/jira/browse/AMQ-4819
18153	AMQ-4819	Sure thing. I'll work up a patch in the next day or so. Can you check out my patch for AMQ-4818? I want to make sure I'm going through the motions properly. Thanks	1	1	https://issues.apache.org/jira/browse/AMQ-4819
18154	AMQ-4818	Yeah, its just add'l fields to existing attributes, but a unit test is always a good idea. I'll look to extend/add. Are the unit tests for the 'master' branch out of Git building for you? I'm getting errors.	1	0	https://issues.apache.org/jira/browse/AMQ-4818

#	Requirement	Questions	Answered	Successfully Answered	URL
18155	AMQ-4762	This is what you see when starting ActiveMQ in standalone mode {code} davsclaus:/opt/apache-activemq-5.9-SNAPSHOT$ bin/activemq console INFO: Using default configuration (you can configure options in one of these file: /etc/default/activemq /Users/davsclaus/.activemqrc) INFO: Invoke the following command to create a configuration file bin/activemq setup [/etc/default/activemq | /Users/davsclaus/.activemqrc] INFO: Using java '/Library/Java/JavaVirtualMachines/jdk1.7.0_25.jdk/Contents/Home/bin/java' INFO: Starting in foreground, this is just for debugging purposes (stop process by pressing CTRL+C) Java Runtime: Oracle Corporation 1.7.0_25 /Library/Java/JavaVirtualMachines/jdk1.7.0_25.jdk/Contents/Home/jre Heap sizes: current=1004928k free=994439k max=1004928k JVM args: -Xms1G -Xmx1G -Djava.util.logging.config.file=logging.properties -Dhawtio.authenticationEnabled=false -Dcom.sun.management.jmxremote -Djava.awt.headless=true -Djava.io.tmpdir=/opt/apache-activemq-5.9-SNAPSHOT/tmp -Dactivemq.classpath=/opt/apache-activemq-5.9-SNAPSHOT/conf; -Dactivemq.home=/opt/apache-activemq-5.9-SNAPSHOT -Dactivemq.base=/opt/apache-activemq-5.9-SNAPSHOT -Dactivemq.conf=/opt/apache-activemq-5.9-SNAPSHOT/conf -Dactivemq.data=/opt/apache-activemq-5.9-SNAPSHOT/data Extensions classpath: [/opt/apache-activemq-5.9-SNAPSHOT/lib,/opt/apache-activemq-5.9-SNAPSHOT/lib/camel,/opt/apache-activemq-5.9-SNAPSHOT/lib/optional,/opt/apache-activemq-5.9-SNAPSHOT/lib/web,/opt/apache-activemq-5.9-SNAPSHOT/lib/extra] ACTIVEMQ_HOME: /opt/apache-activemq-5.9-SNAPSHOT ACTIVEMQ_BASE: /opt/apache-activemq-5.9-SNAPSHOT ACTIVEMQ_CONF: /opt/apache-activemq-5.9-SNAPSHOT/conf ACTIVEMQ_DATA: /opt/apache-activemq-5.9-SNAPSHOT/data Loading message broker from: xbean:activemq.xml INFO | Refreshing org.apache.activemq.xbean.XBeanBrokerFactory$1@60356eb0: startup date [Tue Oct 01 16:24:25 CEST 2013]; root of context hierarchy INFO | PListStore:[/opt/apache-activemq-5.9-SNAPSHOT/data/localhost/tmp_storage] started INFO | Using Persistence Adapter: KahaDBPersistenceAdapter[/opt/apache-activemq-5.9-SNAPSHOT/data/kahadb] INFO | Apache ActiveMQ 5.9-SNAPSHOT (localhost, ID:davsclaus.air-50551-1380637466269-0:1) is starting INFO | Listening for connections at: tcp://davsclaus.air:61616?maximumConnections=1000&wireFormat.maxFrameSize=104857600 INFO | Connector openwire Started INFO | Listening for connections at: amqp://davsclaus.air:5672?maximumConnections=1000&wireFormat.maxFrameSize=104857600 INFO | Connector amqp Started INFO | Apache ActiveMQ 5.9-SNAPSHOT (localhost, ID:davsclaus.air-50551-1380637466269-0:1) started INFO | For help or more information please see: http://activemq.apache.org WARN | Store limit is 102400 mb, whilst the data directory: /opt/apache-activemq-5.9-SNAPSHOT/data/kahadb only has 75851 mb of usable space - resetting to maximum available disk space: 75851 mb INFO | Welcome to hawtio 1.2-M19 : http://hawt.io/ : Don't cha wish your console was hawt like me? ;-) INFO | Starting hawtio authentication filter, JAAS authentication disabled INFO | Using file upload directory: /opt/apache-activemq-5.9-SNAPSHOT/tmp/uploads INFO | jolokia-agent: Using access restrictor classpath:/jolokia-access.xml INFO | Web console type: embedded INFO | ActiveMQ WebConsole initialized. INFO | Initializing Spring FrameworkServlet 'dispatcher' {code}	1	1	https://issues.apache.org/jira/browse/AMQ-4762
18156	AMQ-4759	We need an OSGi bundle for it. Currently 1.4.4 is the last version we have http://repo2.maven.org/maven2/org/apache/servicemix/bundles/org.apache.servicemix.bundles.xstream/	1	1	https://issues.apache.org/jira/browse/AMQ-4759
18157	AMQ-4738	@Deven - have a peek at https://issues.apache.org/jira/browse/AMQ-4682 - which allows runtime mods to the xml configuration. does that meet your use case?	1	1	https://issues.apache.org/jira/browse/AMQ-4738
18158	AMQ-4701	Hi, I am also facing a similar problem. The number of threads are not decreasing after using NIO transport. Can somebody tell in what scenarios this kind problem occurs? Thanks, Sam	1	0	https://issues.apache.org/jira/browse/AMQ-4701
18159	AMQ-4701	I have provided my configuration and test results. Can someone from ActiveMQ developers team help me with this. As Timothy said it reduces clients for large number of clients, So can you please specify what should be this large value. I have tried with 20 clients. After how many clients nio starts reducing threads ? Thanks, Anuj	1	0	https://issues.apache.org/jira/browse/AMQ-4701
18160	AMQ-4694	JNA appears to be (at least based on what distribution says) dual licensed under LGPL and ASL from version 4.0.0 onward. Does Apache preclude the use of open source work that isn't exclusively ASL? The version of JNA referenced in the patch is 3.5.1, but version 4.0 could be used instead.	1	1	https://issues.apache.org/jira/browse/AMQ-4694
18161	AMQ-4693	Hi ActiveMQ team, Could you please let me know if there is a plan to add GSS API based Kerberos Authentication to ActiveMQ TCP/STOMP connectors anytime soon. Could you please also provide rough estimate on timelines/version when this will be available? We use ActiveMQ and would like to expand its usage to authenticated message exchanges. Unfortunately, the simple text based user/password authentication supported currently by ActiveMQ will not be enough for our use cases and we want to use single sign on based authentication. Thanks!!	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18162	AMQ-4693	Hi Timothy, Thanks for the quick response. Can you please atleast tell me what are the rough timelines when ActiveMQ team could take this up? Would this be available in ActiveMQ-6.0 over next few months or is this too far away? Since Apache Qpid already supports Kerberos, I was thinking that code from Qpid can be ported to ActiveMQ without much effort? Is this too much of work? Are there any challenges in implementing kerberos for activemq? I see various ppl are interested in this in various communities. Thanks, Sam	1	1	https://issues.apache.org/jira/browse/AMQ-4693
18163	AMQ-4693	Also, how do you suggest somebody implement Kerberos support in ActiveMQ? Can it be done via ActiveMQ plugins? Since a user writing plugin does not have handle of the actual socket/connection, it may be difficult to add kerberos protocol/support via a broker plugin. Even if socket is accessible from broker plugin, activemq will use the socket for its own communication protocol and user will not be able to superimpose kerberos protocol on top of that connection stream parallely. Also, its not clear how will server authenticate itself to the user in such cases. I think kerberos authentication support will have to be inbuilt in activemq and cannot be done via plugins. This is the main problem due to which i cannot implement kerberos via plugin unless this is added into core activemq code. Thanks, Sam	1	1	https://issues.apache.org/jira/browse/AMQ-4693
18164	AMQ-4693	Thanks for the pointer. The problem is how does an ActiveMQ client (both JMS/STOMP) pass kerberos credentials to the ActiveMQ broker. I don't see any customized headers/fields in the connect/handshake call that clients can use to pass the kerberos credential/token. Clients also don't have access to the socket created by ActiveMQ library, so kerberos credentials cannot be written directly on the socket. Also, since ActiveMQ has its connect/handshake own protocol and message formats, how would client's kerberos credentials be passed to broker while activemq is doing its down protocol -- this seems complicated to overlap the two. Also the other way around, how would the broker authenticate itself to the clients? -Sam	1	1	https://issues.apache.org/jira/browse/AMQ-4693
18165	AMQ-4693	Hi Timothy, Please give the rough timelines when can we expect kerberos authentication in ActiveMQ. I can see many people are interested in this issue and actually if kerberos authentication can be added into the next coming version, it will be a great help. And as Sam suggested that Qpid already supports Kerberos, it won't be much overhead to port the code into ActiveMQ. Thanks, Anuj	1	0	https://issues.apache.org/jira/browse/AMQ-4693

#	Requirement	Questions	Answered	Successfully Answered	URL
18166	AMQ-4693	Let's step back a moment. In order to support modern authentication methods such as Kerberos, you need the capability for the client and server to exchange an arbitrary number of messages at the beginning of a session in order to set up a security context. That performs authentication, and provides a context which the peers can use protect subsequent messages. Is there a way in ActiveMQ whereby we can plug in to do this? It's been suggested to me that there is not.	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18167	AMQ-4693	Hi, As per my understanding, in order to add Kerberos authentication for ActiveMQ client and broker to mutually authenticate themselves to each other (using Kerberos or GSS API), we need to have a common channel using which client and broker could arbitrarily exchange some messages (tokens) before the GSS Security Context is established. Only after two way exchange of messages (tokens) over a common channel, the client and broker can authenticate themselves to each other. However, in order to exchange such messages initially, we would need common channel for client/broker to communicate. With the current plugin/interceptor support provided in ActiveMQ, there is no way for client and server to mutually exchange some messages before they are considered authenticated to each other. I only see one way passing of user name and password from client to broker which isn't sufficient to establish GSS security context and mutual authentication. It looks like there has to be some additional mechanism for doing this two way authentication (and exchange of tokens) before any other other communication can take place. So, it does not look like we can add kerberos based authentication to activemq with the current plugin support. Could ActiveMQ developers please confirm my understanding here? Thanks, Deepak	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18168	AMQ-4693	Interesting discussion! I agree that without a mechanism for mutual exchange of messages initially, GSS security context cannot be eastablished and kerberos authentication cannot take place. Could the ActiveMQ developer team comment on Deepak's and Richard's queries and confirm kerberos authentication cannot be added with the currently available plugin support? I guess all of us are stuck unless these queries are answered. Thanks, Sam	1	1	https://issues.apache.org/jira/browse/AMQ-4693
18169	AMQ-4693	Alright folks ! So can we conclude that adding a security plugin to do kerberos based authentication is NOT an option here? Tim, Christian -- Can we broaden the scope here and try add an SASL framework? The out-of-the-box implementation should provide GSSAPI based kerberos authentication and others like BASIC, PLAIN, DIGEST, MD5 could be added later. I understand that most ActiveMQ devs are too occupied in other stuff but I think its time that ActiveMQ be made more secure & robust in this regard since some of the other MQ products already support these things.	0	0	https://issues.apache.org/jira/browse/AMQ-4693
18170	AMQ-4693	Can somebody from ActiveMQ developer team confirm my understanding in the previous update ? Thanks, Anuj	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18171	AMQ-4693	Hi Gary, ActiveMQ team, Now that ActiveMQ-5.9 is out, I just wanted to check and confirm that Kerberos/SASL based authentication is really being planned for the next ActiveMQ release (v 5.10). This will be tremendously helpful. Could you please provide a rough timeframe on when this feature will be available? Thanks, Deepak	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18172	AMQ-4693	Hi Gary, Did you get a chance to look at my previous post? Could you or somebody else from ActiveMQ team please provide a rough timeline if possible? Thanks a ton!! -Deepak	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18173	AMQ-4693	We are also eagerly awaiting this feature. This will really take ActiveMQ's usage and use cases to next level and will be useful for many people in the community and industry. Can ActiveMQ developers please prioritize this feature since this seems like basic stuff which should be supported? =John=	1	1	https://issues.apache.org/jira/browse/AMQ-4693
18174	AMQ-4693	What's the status on this based on the current voting? Can somebody from the ActiveMQ team please respond to this? Can we assume this is definitely a part of ActiveMQ-5.10? Given that Apache products like Qpid support kerberos based authentication, this shouldn't be too hard to add, no? Note that SASL based authentication should work for Openwire as well as STOMP. -steve	0	0	https://issues.apache.org/jira/browse/AMQ-4693
18175	AMQ-4693	Hi Gary, Now that we are past January, is there any plan to look at this sometime soon? Thanks! -Deepak	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18176	AMQ-4693	Gary, ActiveMQ Gods, When do we plan to start on this? Please let us know if this is on track for ActiveMQ-5.10 release. Could you please provide a rough timeline on when this feature could be available. That will be really helpful! Thanks! -Deepak	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18177	AMQ-4693	I agree that this may be a significant work to implement. However since several people have expressed interest in this feature and the benefits of adding this feature would also be significantly high, can we atleast get a rough timeline on when ActiveMQ experts/developers may start looking at this? I think this feature would take ActiveMQ's usage in the industry to the next level and make it useful in a bunch of scenarios where it currently cannot be used due to lack of strong security mechanisms. Since this was/is originally marked for ActiveMQ-5.10 release, can we still assume that this will be available in ActiveMQ-5.10? Are you implying that this feature will not be available in ActiveMQ-5.10 at all? Getting some clarity on timeline would atleast help a lot of people to make proper choice of message queue implementation that they want to use in their projects. Many Thanks! -Deepak	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18178	AMQ-4693	Hi ActiveMQ team, Did you get a chance to look at this? Could you please provide a rough timeline for this based on current priorities and project pipeline? Can this be taken up sometime soon given that this will be useful for a lot of use cases in general and will take ActiveMQ usage to the next level? Thanks, Deepak	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18179	AMQ-4693	Hi Ryan! Thanks for feedback! As this solution is built on top of RFC2712, it is using TCP and SSL protocol with Kerberos tokens handshake. As failover works for plain SSL, then it should work for Kerberos too I believe. I have performed some tests: 1. Amended org.apache.activemq.network.NetworkFailoverTest to use krb5 protocol instead of tcp- PASSING 2. Crafted my own test with broker1 is starting, connecting 3 consumers with failover protocol, broker1 stopping, starting broker2, sending 3 messagaes. As a result consumers are reconnecting from broker1 to broker2 and consuming messages properly. Do you remember what exactly failed in your scenario? Haven't tested it yet on 2 different machines. Do you remember whether your problem occurs when testing failover on one machine? What I haven't tested yet is token expiration and renewal, which might cause the problem you described above. Will try to test it soon and will be back with results. I am not keen in Kerberos, so any help is welcome :) Thanks in advance for feedback. Piotr Klimczak	1	1	https://issues.apache.org/jira/browse/AMQ-4693
18180	AMQ-4693	Hi Piotr, Really excited to see this getting implemented! Is this new feature scheduled for the upcoming 5.11 release? Thanks, Deepak	1	0	https://issues.apache.org/jira/browse/AMQ-4693
18181	AMQ-4693	[~nannou9] I was inspired by ur PR to implement same in artemis - https://issues.apache.org/jira/browse/ARTEMIS-1264 however it appears the solution is in a bit of a java bubble and maybe a time warp. One I peeked out of the java realm I found the OpenSSL discussions and noted the evolution from RFC2712. It *is* too good to be true that any TLS client can get kerberos authentication for free. It comes at the cost of having an insecure channel. The use case is Kerberos SSO over an insecure channel, have you found that useful in practice?	1	1	https://issues.apache.org/jira/browse/AMQ-4693

#	Requirement	Questions	Answered	Successfully Answered	URL
18182	AMQ-4693	Hi [~gtully]. When I was looking for possible solutions, I have studied some other projects and also got inspired by one of Apache project which had it already implemented (don't remember which one unfortunately). My assumption was that if this solution was already once accepted by apache committers to the other project, then it should be good enough for ActiveMQ too, but indeed it looks like things moved on since then. At this point of time (late 2014), I have found that RFC2712 was the simplest implementation for this JIRA as other approaches required quite a lot of work as per comments made by others. Haven't used it anywhere in any project, as did it purely for community and self learning of kerberos. Kind regards	1	1	https://issues.apache.org/jira/browse/AMQ-4693
18183	AMQ-4693	[~gtully] As RFC2712 (in fact JSSE) is now past, the only options we have as I understand is GSS API. Question now- plain GSS or via SASL? C world will rather uses SASL if we are looking for C libs compatibility. Also, while JSSE came at almost no extra cost as it was over SSL, GSS/SASL are bit more difficult to implement (but definitely doable) as it requires sending tokens over stream pipe, instead of just stream. Surprisingly NIOSSLTransport doesn't seem to be far away from what GSS needs, as it also does (un)wrapping, handshaking etc but over SSL. So implementing GSS might not be that tricky as I initially thought, but I might be missing sth as I haven't been doing anything for ActiveMQ for last 2.5 years- bad :(I know. WDYT? Does it make sense? Personally I would love to implement it. Just not sure if my RL work will allow me as I am spending a lot of time in overtime in last months. Might give it a go during the weekend. Unless all hands should now go to Artemis? Cheers.	1	1	https://issues.apache.org/jira/browse/AMQ-4693
18184	AMQ-4679	Hi Mr. Timoty Bish! Thanks for your reply. It allowed me to move forward a bit. The download of ActiveMQ moved forward significantly. But there is another problem. The download stopped a one transitive dependency artifact: org.apache.qpid:proton-jms:jar:0.3.0-fuse-2 I searched for this artifact in the repository search engine and there is a difference between the version ActiveMQ is using and the one in the Maven repository. Look org.apache.qpid:proton-jms:jar:0.3.0-fuse-2 org.apache.qpid:proton-jms:jar:0.3 Do you see the difference? Here is my full error display > Downloading org.apache.qpid:proton-jms:pom:0.3.0-fuse-2 > Building SomeLocalApplication > Downloading org.apache.qpid:proton-jms:jar:0.3.0-fuse-2 > Buildr aborted! > RuntimeError : Failed to download org.apache.qpid:proton-jms:jar:0.3.0-fuse-2, tried the following repositories: > http://repo1.maven.org/maven2/ > https://repository.apache.org/content/repositories/snapshots/ Can you find a solution to this discrepancy? I suspect that you might have to update library org.apache.qpid:proton-jms:jar:0.3.0-fuse-2 to the most current one available. Also, I strongly suggest that you test the full download of ActiveMQ with the coordinate you give on your website. You might have to delete your own local Maven to do a full test.	1	0	https://issues.apache.org/jira/browse/AMQ-4679
18185	AMQ-4652	Since Java 8 is out now and ActiveMQ is still built with Java 7. When can we expect ActiveMQ compiled with java 8 ? Thanks, Anuj	1	0	https://issues.apache.org/jira/browse/AMQ-4652
18186	AMQ-4644	Can you find the valid SQL That informix can use? It is possible to define custom SQL configurations for different database vendors. So we can have an informix, but would need to valid SQL to use.	0	0	https://issues.apache.org/jira/browse/AMQ-4644
18187	AMQ-4621	I managed to fix the xml configuration issue, however this case still fails with the new strategy. Are you able to provide any insight into why this is still failing? Note this the same test case as on AMQ-4533.	1	1	https://issues.apache.org/jira/browse/AMQ-4621
18188	AMQ-4621	Should we reopen this issue? Tim mentioned to me via email he made some adjustments to the fix for this issue, but I can't find them in the Apache ActiveMQ repo, git://git.apache.org/activemq.git.	1	1	https://issues.apache.org/jira/browse/AMQ-4621
18189	AMQ-4594	I just followed Christian's instructions on his blog post http://www.christianposta.com/blog/?p=315 using the latest 5.9-SNAPSHOT https://repository.apache.org/content/groups/snapshots/org/apache/activemq/apache-activemq/5.9-SNAPSHOT/ and it seems hawtio is correctly included; so I guess this issue is almost closed ;). Though its worth leaving open until just before the release; so we can ensure we upgrade to the latest hawtio release. (e.g. 1.2 GA should be out by then).	0	0	https://issues.apache.org/jira/browse/AMQ-4594
18190	AMQ-4594	I'm building from source (commit 3e58a5a4ebd4d311f6d82065f602ace911318038), using maven 3.05 and encountering this same exception. Is there a workaround for the time being to be able to demo hawtio to co-workers?	1	0	https://issues.apache.org/jira/browse/AMQ-4594
18191	AMQ-4567	I suppose the original idea was for read/write access to JMX to be a admin-priviledged function... Should we enhance that to enforce authn authz at the jmx level with the thought users might be using JMX?	1	0	https://issues.apache.org/jira/browse/AMQ-4567
18192	AMQ-4565	Is this problem related to or even the same as AMQ-4356?	1	0	https://issues.apache.org/jira/browse/AMQ-4565
18193	AMQ-4553	Attached first attempt... Gary/Tim, can you review?	1	1	https://issues.apache.org/jira/browse/AMQ-4553
18194	AMQ-4552	Hi Claus, do you mind if I assign to me and work on this Jira ?	1	0	https://issues.apache.org/jira/browse/AMQ-4552
18195	AMQ-4492	Hello [~chris@die-schneider.net],[~jbonofre], I have been seeing all the messages in the Karaf dev list about Karaf 3 so I am getting all excited (Karaf 3 started to make me think of "Waiting For Godot"...:)). I see that activemq 5.9 is soon to come out so I would like to upgrade to that as well. My apologies if I am not understanding here - What exactly is the problem? I can happily run activemq 5.8.0 without Karaf (in an eclipse workspace with bndtools - equinox + gogo)...though I don't have access to the activemq commands (which are Karaf specific). Why should Karaf 3.0.0 cause a problem with gogo package bindings? Surely Karaf should export the same gogo packages as native felix gogo does? If it doesn't won't it break other OSGi features that may use gogo? I would hate activemq 5.9.0 not to be able to work in Karaf 3.0.0. just because of the Karaf command-line (especially since it will likely be months for 5.10.0). The command-line commands, are not essential, are they (I don't require them to run activemq 5.8.0 with just equinox and gogo)?	0	0	https://issues.apache.org/jira/browse/AMQ-4492
18196	AMQ-4492	What is the status of this? Is there a new Karaf release such as 3.0.1 that helps or is there some work needed on ActiveMQ side?	1	0	https://issues.apache.org/jira/browse/AMQ-4492
18197	AMQ-4433	Here's a new patch (with tests!) that addresses some of the missing validation: https://github.com/qerub/activemq/commit/22776a2960231305674c70e795daec80db02ab69 https://github.com/qerub/activemq/commit/22776a2960231305674c70e795daec80db02ab69.patch What do you think? (I've forgotten exactly which cases were too hard to validate, but it includes the transport.* sub-parameters.)	1	0	https://issues.apache.org/jira/browse/AMQ-4433
18198	AMQ-4399	Is there more to this ticket? Or can it be closed?	1	1	https://issues.apache.org/jira/browse/AMQ-4399
18199	AMQ-4374	I'm guessing since ActiveMQ already has this feature that you don't want us to add it again as that would be redundant, and it'd be redundant too. Is this for Apollo?	1	0	https://issues.apache.org/jira/browse/AMQ-4374
18200	AMQ-4365	My understanding is that when using multiple kahaDB (mKahaDB) persistence adpaters, we only need one lease-database-locker for all the adapters, right? The xsd requires that we define one for each kahaDB. Does it mean we end up with multiple locker-threads checking and/or updating the lock-table? How does it work with mKahaDB really?	1	0	https://issues.apache.org/jira/browse/AMQ-4365

#	Requirement	Questions	Answered	Successfully Answered	URL
18201	AMQ-4365	Thanks Gary. So you mean that at the moment we better not to use the Paul's lease-locker implementation if we're using mKahaDB? Well right now we are in a kind of no-lock-no-luck situation, meaning that we couldn't make the master/slave setup to work with (m)kahaDB for one of our customers, because of the filesystem used (they're not using NFSv4 as recommended). Perhaps we need to reconsider using mKahaDB and getting back to a single kahaDB with amq-db-lease-locker. Is there any plan from your side to come up with a solution in the next release?	1	0	https://issues.apache.org/jira/browse/AMQ-4365
18202	AMQ-4365	With this change in place does the broker handle the two scenarios where the database is either offline at start-up or goes offline at at some point after start-up? If so, how? If not, what modifications are needed for the broker to cope with these scenarios? I am trying to get an idea for the nature and magnitude of any required changes. Ideally the broker should attempt to (re)connect indefinitely in both scenarios.	1	0	https://issues.apache.org/jira/browse/AMQ-4365
18203	AMQ-4358	Matt, for "time" are you looking for a timestamp or the amount of time in ms? FirstEnqueueTime = amount of time it took to do the first enqueue since the broker has been up, LastEnqueueTime = amount of time it took to do the last enqueue that the destination knows about? etc..	1	0	https://issues.apache.org/jira/browse/AMQ-4358
18204	AMQ-4354	Which statistics are you referring to?	1	1	https://issues.apache.org/jira/browse/AMQ-4354
18205	AMQ-4354	All the stats for non-transitive objects under org.apache.activemq. Broker stats, destination stats, transport/network connector stats, etc.. (connection stats wouldn't make any sense b/c they are transitive). Not persisting statistics results in weird scenarios where after a restart you have "4,232" pending messages, but "0" enqueued messages. Admins could then rely on "resetStatistics" to clear them, instead of a broker restart	0	0	https://issues.apache.org/jira/browse/AMQ-4354
18206	AMQ-4341	Attaching a patch that: i) removes dependency flag from xbean-spring ii) removes depedenct flag from jms spec iii) removes reference to the activemq-http bundle (since its part of the uberbundle) iv) Creates a new integration test that tests the installation of the activemq broker feature against obr.	1	0	https://issues.apache.org/jira/browse/AMQ-4341
18207	AMQ-4341	Instead of removing the dependency flags from xbean-spring and the JMS specification, wouldn't it be better to ensure the activemq-osgi bundle imports those required packages? That would trigger the installation of those bundles too, but it would also avoid installing a second version of xbean-spring if a suitable one is already available.	1	0	https://issues.apache.org/jira/browse/AMQ-4341
18208	AMQ-4341	@loannis, this looks great. thanks for the patch. Peeking at org.apache.activemq.karaf.itest.ObrFeatureTest#testBroker that is disabled, any idea why that gets a clash of ports on 8181? on spring-dm in the activemq feature. Do u recall why that is required?	1	1	https://issues.apache.org/jira/browse/AMQ-4341
18209	AMQ-4341	Is this import version range correct? > + org.springframework.osgi*;version="[3,4]";resolution:=optional, I would assume spring osgi is version 1.2.1 being used. So it should be org.springframework.osgi*;version="[1.2,1.3)";resolution:=optional, And there is imports of sun packages + sun.misc*;resolution:=optional, + sun.nio*;resolution:=optional, Is that correct?	1	0	https://issues.apache.org/jira/browse/AMQ-4341
18210	AMQ-4341	@Claus, I think that you are right about the org.springframework.osgi version range. The sun packages are needed since some of the internal packages in the activemq-osgi make use of it. @Gary, its pretty safe to enable the test. Regarding the spring-dm feature, maybe we could get rid of it (we have the imports optional anyway). I think that the only case where its acutally needed is activemq-camel, which again is optional. Does this make sesne?	0	0	https://issues.apache.org/jira/browse/AMQ-4341
18211	AMQ-4311	What store are you assuming for this? Kahadb? Jdbc? Memory? Can you put together a test that shows PFC won't block? message size for pending messages in a tx should be counted against memory usage... any time we do a message.incrementReferenceCount(), the usage memory will be updated (if it's the first reference). In org.apache.activemq.broker.region.Queue#doMessageSend we check if transaction (context.isInTransaction) and then we do message.incrementReferenceCount()... In org.apache.activemq.broker.region.Topic#doMessageSend we increment the reference regardless of whether it's in a tx. You might be on to something, but can you give more details?	1	0	https://issues.apache.org/jira/browse/AMQ-4311
18212	AMQ-4305	When testing this - can i use activemq-rar-5.9-SNAPSHOT against an ActiveMQ 5.7.0 server or do i need to install a snapshot distribution of the server as well? (I'm working together with Mikael Falkvidd who commented above).	0	0	https://issues.apache.org/jira/browse/AMQ-4305
18213	AMQ-4303	Can we resolve this now?	0	0	https://issues.apache.org/jira/browse/AMQ-4303
18214	AMQ-4292	Hi Kevin, I just tried to stick in the Thread.sleep(100) bit, but that did not help me reproduce the issue. Are your still running into this?	1	0	https://issues.apache.org/jira/browse/AMQ-4292
18215	AMQ-4204	Are you using the JMS bridge of ActiveMQ http://activemq.apache.org/jms-to-jms-bridge.html Or are you using the activemq-camel component and letting Camel be the bridge?	0	0	https://issues.apache.org/jira/browse/AMQ-4204
18216	AMQ-4188	Claus did you get anywhere on this one?	1	1	https://issues.apache.org/jira/browse/AMQ-4188
18217	AMQ-4165	The functionnality provided by this mode had its issues (such the one i stumbled upon https://issues.apache.org/jira/browse/AMQ-3688) But it was exposed as a valid option in the documentation until recently (and not completly removed yet) ; cutting it like that is harsh. None of the few "alternatives" can handle the same use cases, and are expecting a very different set of requirements (such as a database or a shared filesystem). Is there any plan to complete the amq options in this area or is the project dropping the ball on HA ?	1	0	https://issues.apache.org/jira/browse/AMQ-4165
18218	AMQ-4146	Hiram, where can we document this? I wonder if you can add a note about this change/new feature in ActiveMQ 5.8 at http://activemq.apache.org/selectors.html	1	0	https://issues.apache.org/jira/browse/AMQ-4146
18219	AMQ-4144	Is this one good now?	1	0	https://issues.apache.org/jira/browse/AMQ-4144
18220	AMQ-4125	@Gary the two store adapters are now marked as deprecated in the trunk code, should we close this out and open a separate ticket to remove the activemq-amq-strore module in v5.9?	0	0	https://issues.apache.org/jira/browse/AMQ-4125
18221	AMQ-4120	Dont we have a SQL server adaptor now? Can you check AMQ 5.9 release?	0	0	https://issues.apache.org/jira/browse/AMQ-4120
18222	AMQ-4120	This still happens in 5.13.2. Is it safe to ignore the warnings?	0	0	https://issues.apache.org/jira/browse/AMQ-4120
18223	AMQ-4093	seems like this should be categorized as severe unless there is a known workaround. how are others dealing with this issue? tested briefly using the multiple kahadb persistence adapter approach to create a separate store for each queue...this seems to work, but may not scale for large number of queues. I also tried this scenario with AMQ 5.9/leveldb and get similar behavior (few messages in single queue block message store cleanup, etc)...this was unexpected.	0	0	https://issues.apache.org/jira/browse/AMQ-4093
18224	AMQ-4085	Do you have a unit test to demonstrate your issue? In some cases this is a non-recoverable error so having the failover transport spin on it wouldn't make any sense.	1	1	https://issues.apache.org/jira/browse/AMQ-4085

#	Requirement	Questions	Answered	Successfully Answered	URL
18225	AMQ-4085	Could you tell me where is the logic or what's the class and method in the broker to close the connection after a short delay. I have checked the source code in the following URL but found no delay closure is involved, http://svn.apache.org/repos/asf/activemq/tags/activemq-5.6.0/activemq-core/src/main/java/org/apache/activemq/broker/TransportConnection.java(processAddConnection) http://svn.apache.org/repos/asf/activemq/tags/activemq-5.7.0/activemq-core/src/main/java/org/apache/activemq/broker/TransportConnection.java(processAddConnection) Or maybe this is only fixed in fusesource releases?	0	0	https://issues.apache.org/jira/browse/AMQ-4085
18226	AMQ-4082	Would it also be possible to do some sort of graceful shutdown (eg as we do in Camel). To allow amble time for inflight messages to complete, and to deny new messages from coming in etc.	0	0	https://issues.apache.org/jira/browse/AMQ-4082
18227	AMQ-4079	Can you point out what you needed to do to make this happen in the Camel build case? The current ActiveMQ source builds as a result of the source plugin in the Apache parent pom from what I can tell and takes the name from the root pom which is the oddly named activemq-parent one.	1	1	https://issues.apache.org/jira/browse/AMQ-4079
18228	AMQ-4053	commons-beanutils-1.8.3.jar commons-beanutils-core-1.8.0.jar - ought to be same version number jsp-2.1-6.1.14.jar jsp-api-2.1-6.1.14.jar - the api is possible in the jsp itself servlet-api-2.5-20081211.jar servlet-api-2.5-6.1.14.jar - should remove the 1st JAR spring-oxm-1.5.10.jar spring-oxm-tiger-1.5.10.jar - I think this is part of spring 3 now stax-api-1.0-2.jar stax-api-1.0.1.jar stax2-api-3.0.2.jar - gotta figure out what to keep xpp3-1.1.4c.jar xpp3_min-1.1.4c.jar - is the _min needed? And in lib/web we have some JSP JARs from glassfish. Wonder if we should use the same JARs as in lib/optional ? core-3.1.1.jar - what is this JAR?	0	0	https://issues.apache.org/jira/browse/AMQ-4053
18229	AMQ-4034	When you add the AMQP transport in the activemq.xml like: {code} <transportConnector name="amqp" uri="amqp://0.0.0.0:5672?maximumConnections=1000&wireformat.maxFrameSize=104857600"/> {code} And then you connect to it via AMQP, the following exception is generated: {code} Exception in thread "ActiveMQ Transport Server Thread Handler: amqp://0.0.0.0:5672?maximumConnections=1000&wireformat.maxFrameSize=104857600" java.lang.NoClassDefFoundError: org/apache/qpid/proton/jms/JMSVendor at org.apache.activemq.transport.amqp.AmqpTransportFilter.<init>(AmqpTransportFilter.java:52) at org.apache.activemq.transport.amqp.AmqpTransportFactory.compositeConfigure(AmqpTransportFactory.java:44) at org.apache.activemq.transport.TransportFactory.serverConfigure(TransportFactory.java:246) at org.apache.activemq.transport.amqp.AmqpTransportFactory.serverConfigure(AmqpTransportFactory.java:56) at org.apache.activemq.transport.tcp.TcpTransportServer.handleSocket(TcpTransportServer.java:415) at org.apache.activemq.transport.tcp.TcpTransportServer$1.run(TcpTransportServer.java:348) at java.lang.Thread.run(Thread.java:722) Caused by: java.lang.ClassNotFoundException: org.apache.qpid.proton.jms.JMSVendor not found by org.apache.activemq.activemq-osgi [54] at org.apache.felix.framework.BundleWiringImpl.findClassOrResourceByDelegation(BundleWiringImpl.java:1460) at org.apache.felix.framework.BundleWiringImpl.access$400(BundleWiringImpl.java:72) at org.apache.felix.framework.BundleWiringImpl$BundleClassLoader.loadClass(BundleWiringImpl.java:1843) at java.lang.ClassLoader.loadClass(ClassLoader.java:356) {code}	1	1	https://issues.apache.org/jira/browse/AMQ-4034
18230	AMQ-4032	I maybe misunderstand... but it seems to me that JMX connections through the RemoteJMXBrokerFacade already supports failover.. {code} String jmxUri = properties.getProperty("jmxUri"); Properties p = System.getProperties(); p.setProperty("webconsole.jmx.url", jmxUri); brokerProxy = new RemoteJMXBrokerFacade(); SystemPropertiesConfiguration spc = new SystemPropertiesConfiguration(); brokerProxy.setConfiguration(spc); {code} I am able to failover back and forth between the AMQ nodes and the RemoteJMXBrokerFacade in the client connects to the other node on the next call. Is there something I am missing?	0	0	https://issues.apache.org/jira/browse/AMQ-4032
18231	AMQ-4027	Is there no ending dot in Oracle? I wonder if it's not better to do a startsWith test instead? And test for starting with "Oracle". ?	0	0	https://issues.apache.org/jira/browse/AMQ-4027
18232	AMQ-4024	Gary, shall we create a configuration switch for that? If yes, are transport addresses exchanged in any other functionality? I'm pondering whether to keep the switch as part of the rebalancing options, or whether it makes sense to promote it to the <broker /> level? e.g. {code} <broker name="abc" networkAddressStrategy="fqdn|hostname|ipaddr"> ... </broker> {code}	0	0	https://issues.apache.org/jira/browse/AMQ-4024
18233	AMQ-4009	We're going to need a bit more information here in order to help you out. Can you elaborate please?	0	0	https://issues.apache.org/jira/browse/AMQ-4009
18234	AMQ-3978	Can we bump up the Priority of this issue, this becomes a major problem in a high volume setup. A single message in DLQ effects the entire system, is there a known work around for this issue?	1	1	https://issues.apache.org/jira/browse/AMQ-3978
18235	AMQ-3978	With AMQ-6203 being closed, how far are we from having decent log compaction (as described in AMQ-3978 or the other related Jira issues)?	1	0	https://issues.apache.org/jira/browse/AMQ-3978
18236	AMQ-3944	Test Joram JMS conformacne cases added in rev 1365608. Use JoramJmsTest to execute them. It currently has 4 tests commented out since they are failing. Need to look deeper to see why.	1	1	https://issues.apache.org/jira/browse/AMQ-3944
18237	AMQ-3933	can you attach your test or describe the nature of it?	0	0	https://issues.apache.org/jira/browse/AMQ-3933
18238	AMQ-3917	@Timothy Can you please give some overview how you fixed this? Do you use the 'id' itself to identify a subscription uniquely? Is 'activemq.subscriptionName' header still there or is it removed? We have some hacks in our code and getting more understanding of your fix would help us remove those hacks when we move to ActiveMQ 5.7.0. Thanks in advance!!	0	0	https://issues.apache.org/jira/browse/AMQ-3917
18239	AMQ-3914	can u validate that the existing: org.apache.activemq.store.jdbc.adapter.TransactJDBCAdapter will work happily with 4.0. You need to specify {{adapter="org.apache.activemq.store.jdbc.adapter.TransactJDBCAdapter"}} in the xml config for jdbc. If that is the case, all that is needed is a copy of: http://svn.apache.org/viewvc/activemq/trunk/activemq-core/src/main/resources/META-INF/services/org/apache/activemq/store/jdbc/microsoft_sql_server_jdbc_driver_3_0?view=log to ..microsoft_jdbc_driver_4_0_for_sql_server	1	1	https://issues.apache.org/jira/browse/AMQ-3914
18240	AMQ-3914	Thanks Petr. I have fixed this. Does the file naming look correct to you? {code} davsclaus:(svn)activemq[trunk:1442509]/$ svn commit --message "AMQ-3914: MS SQL JDBC driver 4.0 fixed. Thanks to Petr Nídl for the details." Adding activemq-jdbc-store/src/main/resources/META-INF/services/org/apache/activemq/store/jdbc/lock/microsoft_jdbc_driver_4_0_for_sql_server Adding activemq-jdbc-store/src/main/resources/META-INF/services/org/apache/activemq/store/jdbc/microsoft_jdbc_driver_4_0_for_sql_server Deleting activemq-jdbc-store/src/main/resources/META-INF/services/org/apache/activemq/store/jdbc/microsoft_sql_server_jdbc_driver_4_0 Committed revision 1442517. {code}	1	1	https://issues.apache.org/jira/browse/AMQ-3914
18241	AMQ-3901	Can you provide some more information here like broker versions, what the client's are doing around the time you see this etc.	0	0	https://issues.apache.org/jira/browse/AMQ-3901

#	Requirement	Questions	Answered	Successfully Answered	URL
18242	AMQ-3894	A useful feature! XML tidbit, it's <redeliveryPolicyEntries> Also, when setting sendToDlqIfMaxRetriesExceeded="false", a message gets silently dropped when max retries is reached. Expected, and if so, worth documenting? I'd have expected the message stays in the queue.	1	0	https://issues.apache.org/jira/browse/AMQ-3894
18243	AMQ-3888	Can't that be another artifact? Something like activemq-jee-specs ? It's better to add one dependency (to run on Tomcat; which would conflict anyway with the jetty dependnecy today) than having to exclude 8 artifacts or define them one by one with 'provided' scope.	0	0	https://issues.apache.org/jira/browse/AMQ-3888
18244	AMQ-3886	Nicholas, do you work on a patch for this?	0	0	https://issues.apache.org/jira/browse/AMQ-3886
18245	AMQ-3883	Hello Freeman, Would it not be better to compare the getName() method of Principal interface? For example: public boolean equals(Object o) { if (this == o) { return true; } if(o instanceof Principal) { if (((Principal) o).getName().equals(this.getName())) { return true; } } if (o == null || getClass() != o.getClass()) { return false; } final GroupPrincipal that = (GroupPrincipal)o; if (!name.equals(that.name)) { return false; } return true; } This way, compare will not only work with any other JAAS implementation.	1	0	https://issues.apache.org/jira/browse/AMQ-3883
18246	AMQ-3876	There is some overlap with trace=true&startLogging=false&dyanmicManagement=true tcp transport options. These allow u to enable (via jmx) trace logging on a connection, of all of the operwire commands. Accessing the log remotely I guess is still a difficulty but a using a remote appender would sort that. did you find some limitations with the existing TcpTransport trace support?	0	0	https://issues.apache.org/jira/browse/AMQ-3876
18247	AMQ-3876	I've had trouble in the past configuring transport tracing, and was never able to set it up dynamically. Can you walk me through how to use the Transport tracing to address the following scenario (which we just recently had in production)? - NMS Client reports that they are not receiving messages - Via VisualVM, we found the Subscription has increasing Enqueue and Dequeue counts - Now we want to tap that connection to see just what is "crossing the wire" I am concerned about the complexity of configuring loggers, but I admit there's much there with which I am not knowledgable. Such as remote loggers. Only yesterday did I learn how to dynamically change logger levels via the log4j MBean. One big consideration is keeping this tracing separate from other log activity to both (a) keep from generating large volumes of noise in the normal logs and (b) to keep the tracing easy to read. (BTW - I'm an old UNIX/C developer at heart. This prototype is modeled on truss/strace/tcpdump)	1	0	https://issues.apache.org/jira/browse/AMQ-3876
18248	AMQ-3861	And is there something you want to do on the AMQ side? If not we should resolve this ticket.	1	1	https://issues.apache.org/jira/browse/AMQ-3861
18249	AMQ-3833	Gary, can this be moved to 5.8 release?	0	0	https://issues.apache.org/jira/browse/AMQ-3833
18250	AMQ-3779	Hi Gary, I Was looking for exactly the same functionality for one of our new use cases. Was happy to found a request but seems like this has been long forgotten. Can this be taken up for the next version please? Thanks, Bhanu	1	0	https://issues.apache.org/jira/browse/AMQ-3779
18251	AMQ-3723	Have you tried with a newer release of xbean? There is a 3.11.1 release http://search.maven.org/#artifactdetails%7Corg.apache.xbean%7Cxbean-spring%7C3.11.1%7Cbundle	1	1	https://issues.apache.org/jira/browse/AMQ-3723
18252	AMQ-3693	I have attached a patch for this work, but please note that the patch was created against Jetty 7.6.0.RC4. The changes will need to use the full Jetty 7.6.0 release when that release is made available. Is it possible to get a new version of ActiveMQ cut when Jetty 7.6.0 is released so that this CVE can be addressed right away?	0	0	https://issues.apache.org/jira/browse/AMQ-3693
18253	AMQ-3693	Applying the new patch and running locally, access to the AMQ web console produces large numbers of the following exception on the broker: {noformat} WARN | /admin/styles/prettify.css;jsessionid=16o74xc6vb13s1dc6cpufjy8de java.lang.NullPointerException at org.eclipse.jetty.security.SecurityHandler.handle(SecurityHandler.java:514) at org.eclipse.jetty.server.session.SessionHandler.doHandle(SessionHandler.java:227) at org.eclipse.jetty.server.handler.ContextHandler.doHandle(ContextHandler.java:1031) at org.eclipse.jetty.servlet.ServletHandler.doScope(ServletHandler.java:406) at org.eclipse.jetty.server.session.SessionHandler.doScope(SessionHandler.java:186) at org.eclipse.jetty.server.handler.ContextHandler.doScope(ContextHandler.java:965) at org.eclipse.jetty.server.handler.ScopedHandler.handle(ScopedHandler.java:117) at org.eclipse.jetty.server.handler.HandlerCollection.handle(HandlerCollection.java:149) at org.eclipse.jetty.security.SecurityHandler.handle(SecurityHandler.java:499) at org.eclipse.jetty.server.handler.HandlerCollection.handle(HandlerCollection.java:149) at org.eclipse.jetty.server.handler.HandlerWrapper.handle(HandlerWrapper.java:111) at org.eclipse.jetty.server.Server.handle(Server.java:349) at org.eclipse.jetty.server.AbstractHttpConnection.handleRequest(AbstractHttpConnection.java:449) at org.eclipse.jetty.server.AbstractHttpConnection$RequestHandler.headerComplete(AbstractHttpConnection.java:910) at org.eclipse.jetty.http.HttpParser.parseNext(HttpParser.java:634) at org.eclipse.jetty.http.HttpParser.parseAvailable(HttpParser.java:230) at org.eclipse.jetty.server.AsyncHttpConnection.handle(AsyncHttpConnection.java:76) at org.eclipse.jetty.io.nio.SelectChannelEndPoint.handle(SelectChannelEndPoint.java:609) at org.eclipse.jetty.io.nio.SelectChannelEndPoint$1.run(SelectChannelEndPoint.java:45) at org.eclipse.jetty.util.thread.QueuedThreadPool.runJob(QueuedThreadPool.java:599) at org.eclipse.jetty.util.thread.QueuedThreadPool$3.run(QueuedThreadPool.java:534) at java.lang.Thread.run(Thread.java:662) {noformat}	0	0	https://issues.apache.org/jira/browse/AMQ-3693
18254	AMQ-3692	Can you provide a bit more documentation how you would use this?, eg how you would configure this in the broker XML file.	1	1	https://issues.apache.org/jira/browse/AMQ-3692
18255	AMQ-3692	ActiveMQ provides karaf commands to check the status of brokers etc. Can you not use that as well? I would assume these commands can tell you that the broker is not available as well.	1	1	https://issues.apache.org/jira/browse/AMQ-3692
18256	AMQ-3687	Thomas, do you work on patches?	0	0	https://issues.apache.org/jira/browse/AMQ-3687
18257	AMQ-3644	Have you raised this issue with Apache Derby? People very often uses prepared statements for the same SQL queries to be executed.	1	0	https://issues.apache.org/jira/browse/AMQ-3644
18258	AMQ-3634	Hi Dejan Are you also handling this case here? http://activemq.2283324.n4.nabble.com/KahaDB-detecting-corrupt-index-file-td2541943.html	1	0	https://issues.apache.org/jira/browse/AMQ-3634
18259	AMQ-3623	I looked at the patch and I just noticed that the appender is added to the logging system but it is not removed when the test has completed. Since the appender references the test case then the test case instance will never be reclaimed until the logging system is removed and that is when the application ends. Would you consider adding the code that remove the added appender when the test has completed.	1	0	https://issues.apache.org/jira/browse/AMQ-3623

#	Requirement	Questions	Answered	Successfully Answered	URL
18260	AMQ-3621	An update for those following this issue that might not be following the ActiveMQ dev mailing list: I made quite a bit of progress over the weekend. Here is the result of the latest effort: https://github.com/lhazlewood/activemq/tree/trunk/activemq-shiro Change log is here: https://github.com/lhazlewood/activemq/commit/be2100ea2c5a421cd21a4c3ae158c65e6e14490f It works as follows: The ShiroPlugin installs 2 BrokerFilters at the moment: - A SubjectFilter which is responsible for constructing and associating a Subject instance with Connections. Downstream filters perform actual security checks. - An AuthenticationFilter (downstream from the SubjectFilter) that enforces client authentication as necessary. It supports system and anonymous connections as well if desired. I still have to create an AuthorizationFilter (that would be downstream from the AuthenticationFilter) that would allow access control checks (can you subscribe to a destination or not, etc). All of this code is tested with 100% class, method and line coverage. Comments, suggestions and feedback are most welcome.	1	1	https://issues.apache.org/jira/browse/AMQ-3621
18261	AMQ-3621	Hi Claus, The initial implementation is complete and has 100% test coverage for all functionality. The only thing remaining is to finish up the documentation: http://activemq.apache.org/shiro.html How would you guys like to receive the code contribution? Can we get it in to the 5.10 release? As an Apache PMC Chair, I already have CLAs on file, so everything should be clear/ready to go.	1	1	https://issues.apache.org/jira/browse/AMQ-3621
18262	AMQ-3621	Actually, I recant that last statement :) There are some minor architectural changes I want to clean up first before they make it into the ActiveMQ repository. When are you guys releasing 5.10? I'd like to get it in that release if possible.	1	1	https://issues.apache.org/jira/browse/AMQ-3621
18263	AMQ-3621	Les, I've committed this to trunk. Suggestions: Can you please finish documenting the Authentication and Authorization sections toward the bottom (marked "work in progress" at the moment). Maybe add referring links (or details if you have it working) about how to integrate with some of the other system Shiro supports (like LDAP, certs?) For the authorization part would be great to explain the verbs allowed (create, remove, wildcards, etc) Can you also give a note about the default behavior of vm:// connectors being treated as 'system' user? Otherwise, everything looks great. Thanks for all the unit tests, that definitely makes life easier :) Keep an eye out on the next nightly builds that include this commit. Please review that nightly and verify that it's to your liking. Commit: To https://git-wip-us.apache.org/repos/asf/activemq.git e7e317d..f9451e5 trunk -> trunk	1	0	https://issues.apache.org/jira/browse/AMQ-3621
18264	AMQ-3621	Has the docs been updated? Is there more work on this?	0	0	https://issues.apache.org/jira/browse/AMQ-3621
18265	AMQ-3621	Is there fore example an activemq-shiro feature to install in Karaf?	0	0	https://issues.apache.org/jira/browse/AMQ-3621
18266	AMQ-3615	Hi Timothy Bish, Pls help me ! My application use JMS for sending message. My client sends a request to the server. While the server is working with my request, my client stops (network problem) before the server finishes. When the server is finished, it'll return to the client, but the server can't see the client which sent the request to server, at which point, the server will return an exception in log file. The exception is : javax.jms.InvalidDestinationException: Cannot publish to a deleted Destination: temp-queue://ID:PC0092-49463-1344231871819-0:0:9 [^] And now I use jms.watchTopicAdvisories for fix this problem, And it's ok, I can't see this problem in the log file. My question is : what are you doing with exception in this case :) ? Sorry If it's stupid question :) Thanks, Dung	0	0	https://issues.apache.org/jira/browse/AMQ-3615
18267	AMQ-3573	Isn't there an API in JDK6+ now that you can ask for disk space?	0	0	https://issues.apache.org/jira/browse/AMQ-3573
18268	AMQ-3569	Gary whats the status of this? Can it be moved to 5.8 if there is still work to do?	1	0	https://issues.apache.org/jira/browse/AMQ-3569
18269	AMQ-3564	Does this work with multiple slaves? (quite a usual scenario)	1	0	https://issues.apache.org/jira/browse/AMQ-3564
18270	AMQ-3541	Tim- There isn't currently a unit test for the TimeStampingBrokerPlugin. I'd be more than happy to write one. Is there a current plugin unit test that you suggest I should use as a starting point? Thanks, Matt Pavlovich	1	1	https://issues.apache.org/jira/browse/AMQ-3541
18271	AMQ-3491	This issue was referred to by Gary Tully in http://activemq.2283324.n4.nabble.com/5-5-Downloads-tp3897983p3898090.html. I cannot find a single download for any version of ActiveMQ on any mirror site or the primary site. Were all versions pulled due to this issue? Is there a timeframe for the software to be available to download agin? Thanks!	1	1	https://issues.apache.org/jira/browse/AMQ-3491
18272	AMQ-3491	First I have to say I really appreciate the decisive action on this unfortunate issue - you guys did the right thing. I have a suggestion for the future (and for all ASF projects, really). Would you consider running something like Black Duck Protex or equivalent as part of pre-release testing - to assure users that there are no license incompatability issues? Maybe running the checks just for major releases at least? {color:red}*If every ASF project could access some license-checking tool*{color} (and use it as a pre-release check) it would add great value to the ASF brand - assuring users that *all reasonable steps* have been taken to avoid any licensing difficulties.	0	0	https://issues.apache.org/jira/browse/AMQ-3491
18273	AMQ-3476	I will try to write a testcase however with the current test cases I run into a build error, using maven 3 and this command line from the activemq-core folder: mvn test -Dtest=org.apache.activemq.transport.stomp.Stomp11Test.java Result: Running org.apache.activemq.transport.stomp.Stomp11Test Tests run: 15, Failures: 0, Errors: 15, Skipped: 0, Time elapsed: 3.042 sec <<< FAILURE! Results : Tests in error: testConnect(org.apache.activemq.transport.stomp.Stomp11Test): Error creating b ean with name 'org.apache.activemq.xbean.XBeanBrokerService#0' defined in class path resource [org/apache/activemq/transport/stomp/stomp-auth-broker.xml]: Insta ntiation of bean failed; nested exception is java.lang.ExceptionInInitializerErr or (... many similar messages) Are additional configuration parameters necessary to run the tests? Many thanks in advance	0	0	https://issues.apache.org/jira/browse/AMQ-3476
18274	AMQ-3476	Looks like your classpath is not right, how are you running the tests? From inside and IDE?	0	0	https://issues.apache.org/jira/browse/AMQ-3476
18275	AMQ-3453	Gary, thanks for the quick reply. I remember doing this some time back but could not infer much from it. Can you please let me know what kind of patterns I can search for in the log file to find out such destinations.	1	1	https://issues.apache.org/jira/browse/AMQ-3453
18276	AMQ-3453	James / Gary, Can you let me know what kind of pattern (any key words) I should be searching for from the log file. I believe it would be logging a lot of things once I start the broker in DEBUG or TRACE mode. I have spent a lot of time trying to resolve this, so looking for some help in directly identifying the issue.	1	1	https://issues.apache.org/jira/browse/AMQ-3453
18277	AMQ-3411	svn add activemq-console/src/test/java/org/apache/activemq/console Must be done or the tests will be lost. This patch adds: " --user <user> Set the username to use connecting", " --password <password> Set the password to use connecting", " --factory <Class> Set connection factory that should be used for creating connections", " Must be a subclass of ActiveMQConnectionFactory", " --passwordFactory <Class> Set class which will do password lookups", " Must be a subclass of PasswordFactory", And fixes the AmqMessagesUtil to use the connection factory so username/password are used for the connections.	0	0	https://issues.apache.org/jira/browse/AMQ-3411
18278	AMQ-3410	What happened here is the patch didn't result in the avn add: svn add activemq-console/src/test/java/org/apache/activemq/console	0	0	https://issues.apache.org/jira/browse/AMQ-3410

#	Requirement	Questions	Answered	Successfully Answered	URL
18279	AMQ-3394	can you provide a test case that captures this use case? currently a redispatch of an in-flight message is tied to consumer close. I guess you want that tied to a timer? the abort slow consumer policy is one way to tie the death of a consumer to a timer. is that close to what you want?	1	1	https://issues.apache.org/jira/browse/AMQ-3394
18280	AMQ-3394	Hi Gary, thanks for the inquiry. I'm not sure co-mingling slow consumer policy with this semantics will achieve the same, but I need to study it some more. The test case or rather, workflow, would go something like this: 1. A "work" message is queued. 2. The "work" message is de-queued by consumer A. It is now off the queue (or hidden) and the "work" is being performed by consumer A. However, the broker has triggered the message's TTL (or some other time indicator) specified by the producer who published it. Say 2 minutes for this message. 3. Consumer A is set to INDIVIDUAL_ACKNOWLEDGE and has not yet acknowledged the message because it is attemping to process the "work" from it. If successful, it will ACK the message. 4. Consumer A suffers a power outage while processing the message and never ACKS it. 5. 2 minutes expire and the original "work" message re-appears on the queue (presumably by the broker). 6. Consumer B receives the or9iginal "work" message, processes the work successfully and ACKS the message after that (within 2 minutes). The message is gone forever now. All are happy. Does this make sense? The idea for this comes from Amazon's Simple Queue Service which provides this semantics.	0	0	https://issues.apache.org/jira/browse/AMQ-3394
18281	AMQ-3388	How about working up a patch file and attaching with grants to apache?	0	0	https://issues.apache.org/jira/browse/AMQ-3388
18282	AMQ-3383	Would it make more sense here to use a byte array instead of a string?	1	1	https://issues.apache.org/jira/browse/AMQ-3383
18283	AMQ-3337	This creates a problem when producer isn't registered for any particular destination {code}java.lang.NullPointerException at org.apache.activemq.broker.jmx.ManagedRegionBroker.createObjectName(ManagedRegionBroker.java:674) at org.apache.activemq.broker.jmx.ManagedRegionBroker.addProducer(ManagedRegionBroker.java:271) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.advisory.AdvisoryBroker.addProducer(AdvisoryBroker.java:133) at org.apache.activemq.broker.CompositeDestinationBroker.addProducer(CompositeDestinationBroker.java:56) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.security.AuthorizationBroker.addProducer(AuthorizationBroker.java:168) at org.apache.activemq.broker.MutableBrokerFilter.addProducer(MutableBrokerFilter.java:99) at org.apache.activemq.broker.TransportConnection.processAddProducer(TransportConnection.java:507) at org.apache.activemq.command.ProducerInfo.visit(ProducerInfo.java:105) at org.apache.activemq.broker.TransportConnection.service(TransportConnection.java:306) at org.apache.activemq.broker.TransportConnection$1.onCommand(TransportConnection.java:179) at org.apache.activemq.transport.ResponseCorrelator.onCommand(ResponseCorrelator.java:116) at org.apache.activemq.transport.TransportFilter.onCommand(TransportFilter.java:69) at org.apache.activemq.transport.vm.VMTransport.dispatch(VMTransport.java:121) at org.apache.activemq.transport.vm.VMTransport.start(VMTransport.java:136) at org.apache.activemq.transport.TransportFilter.start(TransportFilter.java:58) at org.apache.activemq.transport.TransportFilter.start(TransportFilter.java:58) at org.apache.activemq.broker.TransportConnection.start(TransportConnection.java:896) at org.apache.activemq.broker.TransportConnector$1$1.run(TransportConnector.java:217) at java.lang.Thread.run(Thread.java:636){code}	0	0	https://issues.apache.org/jira/browse/AMQ-3337
18284	AMQ-3328	Hi Team, Can anyone kindly review this patch? Thanks Freeman	1	0	https://issues.apache.org/jira/browse/AMQ-3328
18285	AMQ-3322	Is this useful ? Shall I improve the patch in some way?	1	1	https://issues.apache.org/jira/browse/AMQ-3322
18286	AMQ-3297	Can someone review this patch? Thanks in advance. Freeman	1	0	https://issues.apache.org/jira/browse/AMQ-3297
18287	AMQ-3247	append a new patch(AMQ-3247-new.patch), from this patch we get a chance to set DEFAULT_JMX_URL/jmxUser/jmxPassword from system properties, which enable us configure those default value in ESB container, but still not change current logic in activemq-console. Can someone review this patch? Thanks Freeman	1	1	https://issues.apache.org/jira/browse/AMQ-3247
18288	AMQ-3205	Can you grant license to ASF for the attached file.	0	0	https://issues.apache.org/jira/browse/AMQ-3205
18289	AMQ-3189	Did you try the option Gary mentioned?	1	0	https://issues.apache.org/jira/browse/AMQ-3189
18290	AMQ-3174	Hi Adrian, just a question, why did you removed parallel threads in consumer tool? Regards, Dejan	1	1	https://issues.apache.org/jira/browse/AMQ-3174
18291	AMQ-3164	Couldn't this be solved by using apache jsvc? There is a patch to support commons-daemon start/stop: https://issues.apache.org/jira/browse/AMQ-2968 Unfortunately, I don't use the bundled start/stop scripts, so I haven't submitted patches for this: This page has info about how to set up activemq using commons-daemon under windows: http://blog.bigrocksoftware.com/2010/10/07/commons-daemon-procrun-as-a-java-service-wrapper-for-activemq/	1	0	https://issues.apache.org/jira/browse/AMQ-3164
18292	AMQ-3135	Fixed with svn revision 1103673 I added a clientIp property that's set on the broker side (when connection is made) and have a value of remote socket address (it has ip and port combination). I think this should be enough. Can you give it a whirl and see if it works for you.	1	1	https://issues.apache.org/jira/browse/AMQ-3135

#	Requirement	Questions	Answered	Successfully Answered	URL
18293	AMQ-3135	This seems to break the network of broker feature :) 2011-05-30 14:49:21,813 [98.xx.xx:49408] INFO TransportConnection - Started responder end of duplex bridge dev3-dev1@ID:dev3.xx.net-59818-1305794003658-0:1 2011-05-30 14:49:21,814 [ctiveMQ Task-11] WARN Service - Async error occurred: java.lang.NullPointerException java.lang.NullPointerException at org.apache.activemq.broker.jmx.ManagedRegionBroker.createObjectName(ManagedRegionBroker.java:674) at org.apache.activemq.broker.jmx.ManagedRegionBroker.addProducer(ManagedRegionBroker.java:271) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.advisory.AdvisoryBroker.addProducer(AdvisoryBroker.java:133) at org.apache.activemq.broker.CompositeDestinationBroker.addProducer(CompositeDestinationBroker.java:56) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.broker.BrokerFilter.addProducer(BrokerFilter.java:93) at org.apache.activemq.security.AuthorizationBroker.addProducer(AuthorizationBroker.java:168) at org.apache.activemq.broker.MutableBrokerFilter.addProducer(MutableBrokerFilter.java:99) at org.apache.activemq.broker.TransportConnection.processAddProducer(TransportConnection.java:507) at org.apache.activemq.command.ProducerInfo.visit(ProducerInfo.java:105) at org.apache.activemq.broker.TransportConnection.service(TransportConnection.java:306) at org.apache.activemq.broker.TransportConnection$1.onCommand(TransportConnection.java:179) at org.apache.activemq.transport.ResponseCorrelator.onCommand(ResponseCorrelator.java:116) at org.apache.activemq.transport.TransportFilter.onCommand(TransportFilter.java:69) at org.apache.activemq.transport.vm.VMTransport.dispatch(VMTransport.java:121) at org.apache.activemq.transport.vm.VMTransport.start(VMTransport.java:136) at org.apache.activemq.transport.TransportFilter.start(TransportFilter.java:58) at org.apache.activemq.transport.TransportFilter.start(TransportFilter.java:58) at org.apache.activemq.broker.TransportConnection.start(TransportConnection.java:896) at org.apache.activemq.broker.TransportConnector$1$1.run(TransportConnector.java:217) at java.lang.Thread.run(Thread.java:636) 2011-05-30 14:49:21,816 [m://localhost#2] INFO DemandForwardingBridgeSupport - Network connection between vm://localhost#2 and tcp:///88.xx.xx.181:49408 shutdown due to a local error: java.lang.NullPointerException	1	1	https://issues.apache.org/jira/browse/AMQ-3135
18294	AMQ-3135	So AMQ-3337 should be fixed on trunk ... Can you test it out now?	1	1	https://issues.apache.org/jira/browse/AMQ-3135
18295	AMQ-3135	That's OK, it's a connection of the bridge to the local broker. What does the *outbound* connection info says?	1	1	https://issues.apache.org/jira/browse/AMQ-3135
18296	AMQ-3121	Ivan's patch at least makes recovery for inbound messaging plausible, but as Gary points out leaves a connection leak. I don't recall how amq gets a connection for delivering inbound messages. If this connection is cached, can we arrange to use the connection obtained for transaction recovery for inbound message delivery?	1	0	https://issues.apache.org/jira/browse/AMQ-3121
18297	AMQ-3097	Just from a JMS spec point of view, is it allowed to add support for extensions such as function calls? Despite it could be a nice feature.	1	0	https://issues.apache.org/jira/browse/AMQ-3097
18298	AMQ-3024	JobSchedulerStore assumes that it is directory based. By non-Kaha persistence I would assume non file-based persistence. Is there a separate issue for that?	0	0	https://issues.apache.org/jira/browse/AMQ-3024
18299	AMQ-3017	hmm.. works for me on linux .. what os ?	1	1	https://issues.apache.org/jira/browse/AMQ-3017

#	Requirement	Questions	Answered	Successfully Answered	URL
18300	AMQ-3017	OS X, this is the exception {code} java.lang.ClassCastException: sun.net.www.protocol.file.FileURLConnection cannot be cast to java.net.HttpURLConnection at org.apache.activemq.blob.DefaultBlobUploadStrategy.uploadStream(DefaultBlobUploadStrategy.java:49) at org.apache.activemq.blob.DefaultBlobUploadStrategy.uploadFile(DefaultBlobUploadStrategy.java:43) at org.apache.activemq.blob.BlobUploader.upload(BlobUploader.java:51) at org.apache.activemq.command.ActiveMQBlobMessage.onSend(ActiveMQBlobMessage.java:173) at org.apache.activemq.ActiveMQSession.send(ActiveMQSession.java:1738) at org.apache.activemq.ActiveMQMessageProducer.send(ActiveMQMessageProducer.java:231) at org.apache.activemq.ActiveMQMessageProducerSupport.send(ActiveMQMessageProducerSupport.java:241) at org.apache.activemq.blob.FilesystemBlobTest.testBlobFile(FilesystemBlobTest.java:72) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at junit.framework.TestCase.runTest(TestCase.java:168) at junit.framework.TestCase.runBare(TestCase.java:134) at org.apache.activemq.CombinationTestSupport.runBare(CombinationTestSupport.java:90) at org.apache.activemq.CombinationTestSupport.runBare(CombinationTestSupport.java:96) at junit.framework.TestResult$1.protect(TestResult.java:110) at junit.framework.TestResult.runProtected(TestResult.java:128) at junit.framework.TestResult.run(TestResult.java:113) at junit.framework.TestCase.run(TestCase.java:124) at junit.framework.TestSuite.runTest(TestSuite.java:232) at junit.framework.TestSuite.run(TestSuite.java:227) at org.junit.internal.runners.JUnit38ClassRunner.run(JUnit38ClassRunner.java:79) at org.apache.maven.surefire.junit4.JUnit4TestSet.execute(JUnit4TestSet.java:59) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:115) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:140) at org.apache.maven.surefire.Surefire.run(Surefire.java:109) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.booter.SurefireBooter.runSuitesInProcess(SurefireBooter.java:290) at org.apache.maven.surefire.booter.SurefireBooter.main(SurefireBooter.java:1017) {code}	1	1	https://issues.apache.org/jira/browse/AMQ-3017
18301	AMQ-3017	desktop:~/workspace/activemq$ mvn test -Dtest=FilesystemBlobTest [INFO] Scanning for projects... [INFO] Reactor build order: [INFO] ActiveMQ [INFO] ActiveMQ :: KahaDB [INFO] ActiveMQ :: JAAS [INFO] Openwire Generator [INFO] ActiveMQ :: jmdns 1.0 [INFO] ActiveMQ :: Core [INFO] ActiveMQ :: RA [INFO] ActiveMQ :: Pool [INFO] ActiveMQ :: Camel [INFO] ActiveMQ :: Console [INFO] ActiveMQ :: Optional [INFO] ActiveMQ :: Jar Bundle [INFO] ActiveMQ :: File Server [INFO] ActiveMQ :: Blueprint [INFO] ActiveMQ :: Apache Karaf [INFO] ActiveMQ :: RAR [INFO] ActiveMQ :: Run Jar [INFO] ActiveMQ :: Spring [INFO] ActiveMQ Tooling [INFO] ActiveMQ :: Memory Usage Test Plugin [INFO] ActiveMQ :: Performance Test Plugin [INFO] ActiveMQ :: StartUp Plugin [INFO] ActiveMQ :: Web [INFO] ActiveMQ :: Web Demo [INFO] ActiveMQ :: XMPP [INFO] ActiveMQ :: Web Console [INFO] ActiveMQ :: Assembly [INFO] -- [INFO] Building ActiveMQ [INFO] task-segment: [test] [INFO] -- [INFO] [enforcer:enforce {execution: default}] [INFO] [bundle:cleanVersions {execution: cleanVersions}] [INFO] Setting property: classpath.resource.loader.class => 'org.codehaus.plexus.velocity.ContextClassLoaderResourceLoader'. [INFO] Setting property: velocimacro.messages.on => 'false'. [INFO] Setting property: resource.loader => 'classpath'. [INFO] Setting property: resource.manager.logwhenfound => 'false'. [INFO] [remote-resources:process {execution: default}] [INFO] ---------- [INFO] Building ActiveMQ :: KahaDB [INFO] task-segment: [test] [INFO] ---------- [INFO] [enforcer:enforce {execution: default}] [INFO] [bundle:cleanVersions {execution: cleanVersions}] [INFO] [remote-resources:process {execution: default}] [INFO] [resources:resources {execution: default-resources}] [INFO] Using 'UTF-8' encoding to copy filtered resources. [INFO] Copying 2 resources [INFO] Copying 3 resources [INFO] [compiler:compile {execution: default-compile}] [INFO] Nothing to compile - all classes are up to date [INFO] [resources:testResources {execution: default-testResources}] [INFO] Using 'UTF-8' encoding to copy filtered resources. [INFO] Copying 1 resource [INFO] Copying 3 resources [INFO] [compiler:testCompile {execution: default-testCompile}] [INFO] Nothing to compile - all classes are up to date [WARNING] DEPRECATED [systemProperties]: Use systemPropertyVariables instead. [INFO] [surefire:test {execution: default-test}] [INFO] Surefire report directory: /home/maurer/workspace/activemq/kahadb/target/surefire-reports [INFO] ---------- [INFO] Building ActiveMQ :: JAAS [INFO] task-segment: [test] [INFO] ---------- [INFO] [enforcer:enforce {execution: default}] [INFO] [bundle:cleanVersions {execution: cleanVersions}] [INFO] [remote-resources:process {execution: default}] [INFO] [resources:resources {execution: default-resources}] [INFO] Using 'UTF-8' encoding to copy filtered resources. [INFO] skip non existing resourceDirectory /home/maurer/workspace/activemq/activemq-jaas/src/main/resources [INFO] Copying 3 resources [INFO] [compiler:compile {execution: default-compile}] [INFO] Nothing to compile - all classes are up to date [INFO] [resources:testResources {execution: default-testResources}] [INFO] Using 'UTF-8' encoding to copy filtered resources. [INFO] Copying 4 resources [INFO] Copying 3 resources [INFO] [compiler:testCompile {execution: default-testCompile}] [INFO] Nothing to compile - all classes are up to date [WARNING] DEPRECATED [systemProperties]: Use systemPropertyVariables instead. [INFO] [surefire:test {execution: default-test}] [INFO] Surefire report directory: /home/maurer/workspace/activemq/activemq-jaas/target/surefire-reports [INFO] ---------- [INFO] Building Openwire Generator [INFO] task-segment: [test] [INFO] ---------- [INFO] [enforcer:enforce {execution: default}] [INFO] [bundle:cleanVersions {execution: cleanVersions}] [INFO] [remote-resources:process {execution: default}] [INFO] [resources:resources {execution: default-resources}]	1	1	https://issues.apache.org/jira/browse/AMQ-3017
18302	AMQ-3003	@Gary Have you tested this. I have tried to implementing this, but it seems that the message still goes back to the DLQ of the topic.	1	1	https://issues.apache.org/jira/browse/AMQ-3003
18303	AMQ-2922	Our first stab at implementing mutliple stores, with example configuration file. Are we on the right track with this approach and is this something that could be included in a future release?	0	0	https://issues.apache.org/jira/browse/AMQ-2922
18304	AMQ-2922	Is there a way to configure all DLQs to use one persistence adapter, rather than having to list them out one-by-one? DLQs get processed offline, and so separating them from "live" queues would be a great configuration option.	1	0	https://issues.apache.org/jira/browse/AMQ-2922
18305	AMQ-2922	We have a standard of creating DLQs with the suffix .DLQ. Example Sample.Queue.DLQ Tried the below wild card with filtered kaha db. Doesnt seem to work. <filteredKahaDB queue=">.DLQ"> Is this supported? How could we solve this.	0	0	https://issues.apache.org/jira/browse/AMQ-2922

#	Requirement	Questions	Answered	Successfully Answered	URL
18306	AMQ-2911	Here is the patch for the 5.5 SNAPSHOT (10th Sept 2010)	1	1	https://issues.apache.org/jira/browse/AMQ-2911
18307	AMQ-2875	Hi Adrian- I definitely agree that there is room for improvement on the networkConnector/network of brokers front to make it more usable, specifically in the large-scale scenarios you describe. A couple of comments: 1. Any improvement or modification to networkConnector-type configuration would be greatly enhanced by having the configuration be dynamically updatable and persisted. (ie.. save JMX/REST config changes to the activemq.xml file) 2. What do you think about supporting the "host" portion of the JMS url for routing? It seems to me that what you are describing might be achieved by supporting that, instead of tracking advisories up and down nodes. Having the ability to support routing a message to "kiosk46" when specifying the destination as "queue://kiosk46/MY_QUEUE" and client sending message to "brokerA" (HQ) might be easier to implement and manage than a number of network/tree connectors. This would also make migrations from commercial platforms that currently support this, such as MQ series, more straightforward. Thoughts? In my personal experience from implementing ActiveMQ at a global organization with 1000's of nodes is that the number of queues can become problematic in a straight-forward network-of-brokers model. Also, if all the endpoints have unique queues, it pushes the complexity to the application developers which creates a negative developer experience (imho). [Preferred-- Apps deployed to kiosk46 know to connect to queues and don't need to know anything about its current host] Host: Kiosk46 - queue:///ORDERS_Q - queue:///TICKETS_Q - queue:///INVENTORY_Q Host: Kiosk48 - queue:///ORDERS_Q - queue:///TICKETS_Q - queue:///INVENTORY_Q [Versus-- Apps need to consider their location when creating queue names] Host: Kiosk46 - queue:///K46_ORDERS_Q - queue:///K46_TICKETS_Q - queue:///K46_INVENTORY_Q host: Kiosk48 - queue:///K48_ORDERS_Q - queue:///K48_TICKETS_Q - queue:///K48_INVENTORY_Q Thanks, Matt Pavlovich	1	0	https://issues.apache.org/jira/browse/AMQ-2875
18308	AMQ-2834	Is it expected that BrokerService will log an error about temporary store limit and max journal file size if your broker is not configured as persistent?	1	1	https://issues.apache.org/jira/browse/AMQ-2834
18309	AMQ-2826	So any suggestions on how to proceed? The maven profile for now could simply prevent the tests from being run... Not sure what type of infrastructure is available in the CI environment, as far as having the ability to stand up a single node cassandra and zookeeper instance to run the tests agains at some point... How is the JDBC Store tested against all the different vendors that are supported?	1	1	https://issues.apache.org/jira/browse/AMQ-2826
18310	AMQ-2826	Hi Scott, Has there been any movement on QSandra in the last two years? Do you know of anyone using it in production? Or is it a now defunct project? Any information would be helpful! Thanks, Les	1	0	https://issues.apache.org/jira/browse/AMQ-2826
18311	AMQ-2821	Without this improvement, is an inactive queue reference (without any producer nor consumer) propagated to all brokers in a network of brokers topology ? Or does this reference finish to disappear after the non simultaneous restart of all brokers ? Eric-AWL	1	0	https://issues.apache.org/jira/browse/AMQ-2821
18312	AMQ-2716	{quote} target <---------- sender \ \/ temp reply queue ----------> receiver {quote} Now I close the resource in following order. {quote} receiver.close(); replyQueue.delete(); sender.close(); session.close(); connection.close(); {quote} does AMQ not follow the javax.jms.Connection definition? Even in the javadoc from ActiveMQConnection#close there is the same description... {quote} Closing a connection causes all temporary destinations to be deleted. {quote} But it seems that we must delete temp queue explicitly	1	1	https://issues.apache.org/jira/browse/AMQ-2716
18313	AMQ-2702	rev 935952 introduced activemq-spring module and moved a few classes there. Incidentally activemq-pool module doesn't have spring dependency any more. The main question is what should we do with all tests in activemq-core (and other modules), that depends on Spring (and preventing us to do a better separation). IMHO we should move them either in this newly-created activemq-spring module or activemq-all module or even previously created sys test project. It will require some refactoring and probably make debugging a bit harder, but I think it's the move we have to make.	0	0	https://issues.apache.org/jira/browse/AMQ-2702
18314	AMQ-2702	Now that this class is in the separate activemq-spring module, would it be possible to reintroduce the {{InitializingBean}} and {{DisposableBean}} interfaces on the {{PooledConnectionFactoryBean}} (cfr. AMQ-2702.diff) so it can be used in a plain Spring context without requiring the use of {{<context:annotation-driven/>}} to init/destroy the bean?	0	0	https://issues.apache.org/jira/browse/AMQ-2702
18315	AMQ-2668	individual store limits may not be possible as the store is shared, but possibly a precent store usage usage metric would work. request queue could block on 70% full and response queue could block on 100% full. This approach is already uses by the cursors via the cursorMemoryHighWaterMark. Do you think that would satisfy your use case?	1	1	https://issues.apache.org/jira/browse/AMQ-2668
18316	AMQ-2664	This is probably a good idea but... may not be sufficient for some QA tools that uses the classes (e.g. FindBugs). Does the activemq-core require every version of OpenWire bundled in? Could they be different jars? I may be totally wrong... Benoit	1	0	https://issues.apache.org/jira/browse/AMQ-2664
18317	AMQ-2636	Thanks for committing our initial patch! I will shortly be attaching a second patch (against svn revision 921822 of the ActiveMQ trunk) that adds in support for setting the Type Of Service bits as well as the Differentiated Services bits. The patch will also cleans up some tabs and incorrect spacing that I mistakenly let slip into the previous patch as well as does some refactoring on the TransportUriTests. When would it be appropriate for me to add new information about this new Transport option to the existing documentation (http://activemq.apache.org/tcp-transport-reference.html)?	1	1	https://issues.apache.org/jira/browse/AMQ-2636
18318	AMQ-2636	This patch fixes an issue that was discovered by running our automated wireshark test scripts where I was assuming the wrong two bits for the ECN bits, and so only multiples of 4 were working on machines which use 0 for the default ECN. (tested against trunk) The code for the automated wireshark test scripts can currently be found here: http://maljub01.svnrepository.com/comp190/trac.cgi/browser/scripts/dscp_test Is there some way we can either include the automated wireshark test scripts somewhere in the ActiveMQ trunk (maybe as an example, since it involves a Producer that communicates with the ActiveMQ broker) or put it somewhere on the ActiveMQ website so that future uses can benefit from it?	0	0	https://issues.apache.org/jira/browse/AMQ-2636

#	Requirement	Questions	Answered	Successfully Answered	URL
18319	AMQ-2591	I think the heart of the issue here is that security can be tricky with activemq (at least as far as locking down queues). The idea for easy security was to use the browser session id as the queue name... like "A653D78C432AEEC38D142ACCD2703BAF.queue" If the connection to the server was in SSL... there would be a very low risk of man in the middle attacks. Then to get your private messages you would just subscribe up to this session id queue name. Since it is very unlikely anyone could guess your session id... it would be very unlikely that anyone could guess your queue name. The idea would be that unless they knew your session id... they would have no idea which queue to subscribe to to eavesdrop on your messages. Along with this... we look at write only queues or topics for the server to listen to. That way people can publish messages upstream with their session id attached... however no body but the server could read them. The point was brought up however... that AMQ lets you wildcard subscriptions. In this case... they could listen to all queues and would have no trouble finding your "private" queue. I think the general hope is for a way to make a queue private without having to have an authentication mechanism for each individual user subscribing to queues. Our system for example has thousands of users all needing one private queue to receive "secret" messages on. It seems a little daunting to give each of these users unique passwords to access their queue. If there was a way to stop wildcarding... are there other holes people see with this approach?	0	0	https://issues.apache.org/jira/browse/AMQ-2591
18320	AMQ-2485	The master needs to be restarted and the slave attached again. A slave cannot play catch up with an already started and active master. there is a waitForSlave option on a master that can facilitate the synchronized start of the master. Would this suffice?	0	0	https://issues.apache.org/jira/browse/AMQ-2485
18321	AMQ-2476	yea, this is long over due. Just wonder about sandbox rather than a sub module off trunk. Think a sub module on trunk and a default profile that does not execute tests may be best as it will be part of the default compilation. what is the benefit of using the sandbox?	1	0	https://issues.apache.org/jira/browse/AMQ-2476
18322	AMQ-2456	Hi, could a commiter have a look at this patch and see if it's ok to commit. We've been running it in production for a while now and it shows no problems (and is useful). cheers, James.	1	1	https://issues.apache.org/jira/browse/AMQ-2456
18323	AMQ-2456	Hi, I use this plugin but i have a problem with a ssl connection. When the connection is open with the option duplex=false, all work correctly, ssl and tcp. But with the option duplex=true, the plugin doesn't use the ssl configuration. Do you have an explanation? Thanks! Rémy lesiak	1	1	https://issues.apache.org/jira/browse/AMQ-2456
18324	AMQ-2456	Hi Remy, what version of ActiveMQ ? Also, can you send me the relevant connector section of the activemq.xml for both brokers. cheers, James.	1	1	https://issues.apache.org/jira/browse/AMQ-2456
18325	AMQ-2455	This issue is certainly not fixed for me in 5.4.1. Is there any additional configuration I need to do? Source and config for my bridge setup to OpenMQ is here : http://github.com/centraldesktop/openmq-spring-bean	0	0	https://issues.apache.org/jira/browse/AMQ-2455
18326	AMQ-2444	so it states the three conditions, any of which can occur depending on the application. But if any occur, redelivery based on the current redelivery policy will ensue. Any configured redelivery policy will effect all three conditions. How would you word it differently?	1	1	https://issues.apache.org/jira/browse/AMQ-2444
18327	AMQ-2438	To be very precise, I dont have a failover trasport configured so I dont have a backup broker. If my activemq server dies or there is a network outage, when I restart my server will the server automatically recover to guarantee zero message loss or duplication in event of server failure?	0	0	https://issues.apache.org/jira/browse/AMQ-2438
18328	AMQ-2438	Here is the scenario I've tested. 1. started with single broker1 2. clients connected using failover protocol failover:(tcp://broker1:61616) 3. database used by the broker1 goes offline 4. broker1 stops 5. database comes back online After this broker1 never comes back. I am using RAR deployed to JBoss, is there any specific configuration need to be done for broker-config.xml?	1	0	https://issues.apache.org/jira/browse/AMQ-2438
18329	AMQ-2412	I'll check that out, what's the failure you're seeing? That test passes here for me...	1	1	https://issues.apache.org/jira/browse/AMQ-2412
18330	AMQ-2395	Does org.apache.activemq.store.jdbc.Statements#setTablePrefix not suffice?	1	0	https://issues.apache.org/jira/browse/AMQ-2395
18331	AMQ-2390	I just noticed this recently myself, so thanks for creating this issue, Dejan. Another item that added many more files to the ActiveMQ distribution is the inclusion of Dojo for the web console. Is it really necessary to include all of these files?	1	0	https://issues.apache.org/jira/browse/AMQ-2390
18332	AMQ-2386	Support for Byte type messages seems to be broken. If I send a message with a content-length header and any NULL characters I get the following exception: 2010-05-25 12:28:54,652 | DEBUG | Transport failed: java.io.EOFException | org.apache.activemq.broker.TransportConnection.Transport | ActiveMQ NIO Worker java.io.EOFException at java.io.DataInputStream.readFully(DataInputStream.java:180) at java.io.DataInputStream.readFully(DataInputStream.java:152) at org.apache.activemq.transport.stomp.StompWireFormat.unmarshal(StompWireFormat.java:146) at org.apache.activemq.transport.stomp.StompWireFormat.unmarshal(StompWireFormat.java:60) at org.apache.activemq.transport.stomp.StompNIOTransport.serviceRead(StompNIOTransport.java:121) at org.apache.activemq.transport.stomp.StompNIOTransport.access$000(StompNIOTransport.java:49) at org.apache.activemq.transport.stomp.StompNIOTransport$1.onSelect(StompNIOTransport.java:73) at org.apache.activemq.transport.nio.SelectorSelection.onSelect(SelectorSelection.java:94) at org.apache.activemq.transport.nio.SelectorWorker$1.run(SelectorWorker.java:119) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908) at java.lang.Thread.run(Thread.java:619)	1	0	https://issues.apache.org/jira/browse/AMQ-2386
18333	AMQ-2381	How do I update online doc in this regard? Should I?	0	0	https://issues.apache.org/jira/browse/AMQ-2381
18334	AMQ-2365	We intentionally don't add temporary destinations - a lot of broker behaviour works by assuming that the method won't return any. Why do you want the temporary destinations ? - we could provide a work around for your specific use case	0	0	https://issues.apache.org/jira/browse/AMQ-2365
18335	AMQ-2352	Hi, when you use webconsole-properties.xml you should configure your console using system properties. System properties configuration currently allows you to use -Dwebconsole.jms.user and -Dwebconsole.jms.password properties which will be used in created connection factory. As a demonstration I started 5.3.0 broker using conf/activemq-security.xml $ bin/activemq xbeanconf/activemq-security.xml and successfully connected using mvn jetty:run -Dwebconsole.type=properties -Dwebconsole.jms.url=tcp://localhost:61616 -Dwebconsole.jmx.url=service:jmx:rmi:///jndi/rmi://localhost:1099/jmxrmi -Dwebconsole.jms.user=system -Dwebconsole.jms.password=manager Does this works for you? If not, can you clarify a bit more on what are you trying to achieve? Regards Dejan	1	1	https://issues.apache.org/jira/browse/AMQ-2352
18336	AMQ-2352	hi, Where can I get the xbeanconf/activemq-security.xml to give it a try? As far as I remember I could start up using jetty and connect to it using webconsole, but I couldn't send a message to the queues, because the system properties ${webconsole.jms.userName} and ${webconsole.jms.password} do not automatically map to the connectionFactory bean. I hope this helps, Philippe	1	0	https://issues.apache.org/jira/browse/AMQ-2352

#	Requirement	Questions	Answered	Successfully Answered	URL
18337	AMQ-2352	context org.mortbay.jetty.plugin.Jetty6PluginWebAppContext@da3772{/act ive,C:\ontwikkeling\galileoworkspace\amaris-amq\patched\activemq-web-console-5.3.0.war} org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'brokerQuery' defined in ServletC ontext resource [/WEB-INF/webconsole-properties.xml]: Error setting property values; nested exception is org.springframe work.beans.NotWritablePropertyException: Invalid property 'configuration' of bean class [org.apache.activemq.web.RemoteJ MXBrokerFacade]: Bean property 'configuration' is not writable or has an invalid setter method. Does the parameter type of the setter match the return type of the getter? at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.applyPropertyValues(AbstractAuto wireCapableBeanFactory.java:1279) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.populateBean(AbstractAutowireCap ableBeanFactory.java:1010) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.doCreateBean(AbstractAutowireCap ableBeanFactory.java:472) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory$1.run(AbstractAutowireCapableBea nFactory.java:409) at java.security.AccessController.doPrivileged(Native Method) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.createBean(AbstractAutowireCapab leBeanFactory.java:380) at org.springframework.beans.factory.support.AbstractBeanFactory$1.getObject(AbstractBeanFactory.java:264) at org.springframework.beans.factory.support.DefaultSingletonBeanRegistry.getSingleton(DefaultSingletonBeanRegis try.java:222) at org.springframework.beans.factory.support.AbstractBeanFactory.doGetBean(AbstractBeanFactory.java:261) at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:185) at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:164) at org.springframework.beans.factory.support.DefaultListableBeanFactory.preInstantiateSingletons(DefaultListable BeanFactory.java:429) at org.springframework.context.support.AbstractApplicationContext.finishBeanFactoryInitialization(AbstractApplic ationContext.java:728) at org.springframework.context.support.AbstractApplicationContext.refresh(AbstractApplicationContext.java:380) at org.apache.activemq.web.WebConsoleStarter.createWebapplicationContext(WebConsoleStarter.java:59) at org.apache.activemq.web.WebConsoleStarter.contextInitialized(WebConsoleStarter.java:43) at org.mortbay.jetty.handler.ContextHandler.startContext(ContextHandler.java:548) at org.mortbay.jetty.servlet.Context.startContext(Context.java:136) at org.mortbay.jetty.webapp.WebAppContext.startContext(WebAppContext.java:1239) at org.mortbay.jetty.handler.ContextHandler.doStart(ContextHandler.java:517) at org.mortbay.jetty.webapp.WebAppContext.doStart(WebAppContext.java:466) at org.mortbay.jetty.plugin.Jetty6PluginWebAppContext.doStart(Jetty6PluginWebAppContext.java:124) at org.mortbay.component.AbstractLifeCycle.start(AbstractLifeCycle.java:50) at org.mortbay.jetty.handler.HandlerCollection.doStart(HandlerCollection.java:152) at org.mortbay.jetty.handler.ContextHandlerCollection.doStart(ContextHandlerCollection.java:156) at	1	0	https://issues.apache.org/jira/browse/AMQ-2352
18338	AMQ-2352	Hi Philippe, are you sure you don't have 5.2.0 left somewhere in the classpath? I just tried the following and it worked well: 1. Download the 5.3.0 web console war (http://repo2.maven.org/maven2/org/apache/activemq/activemq-web-console/5.3.0/activemq-web-console-5.3.0.war) 2. Copy it to clean jetty install webapps folder 3. Start secured ActiveMQ 5.3.0 (bin/activemq xbean:conf/activemq-security.xml) 4. start jetty (java -Dwebconsole.type=properties -Dwebconsole.jms.url=tcp://localhost:61616 -Dwebconsole.jmx.url=service:jmx:rmi:///jndi/rmi://localhost:1099/jmxrmi -Dwebconsole.jms.user=system -Dwebconsole.jms.password=manager -jar start.jar) 5. use web console at http://localhost:8080/activemq-web-console-5.3.0/ Regards Dejan	1	1	https://issues.apache.org/jira/browse/AMQ-2352
18339	AMQ-2352	The scenario you describe works on my local activemq. However when I try to validate on our deployed activemq here, that also has authentication on jmx, I fail to specify the username and password. Previously (on 5.2.0) I used: -Dwebconsole.jmx.role=username -Dwebconsole.jmx.password=password Apparently in 5.3.0 it is changed to: -Dwebconsole.jmx.user=username -Dwebconsole.jmx.password=password But still I fail to login. Caused by: javax.security.auth.login.FailedLoginException: Invalid username or password at com.sun.jmx.remote.security.FileLoginModule.attemptAuthentication(FileLoginModule.java:432) at com.sun.jmx.remote.security.FileLoginModule.login(FileLoginModule.java:279) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at javax.security.auth.login.LoginContext.invoke(LoginContext.java:769) at javax.security.auth.login.LoginContext.access$000(LoginContext.java:186) at javax.security.auth.login.LoginContext$5.run(LoginContext.java:706) at java.security.AccessController.doPrivileged(Native Method) at javax.security.auth.login.LoginContext.invokeCreatorPriv(LoginContext.java:703) at javax.security.auth.login.LoginContext.login(LoginContext.java:575) at com.sun.jmx.remote.security.JMXPluggableAuthenticator.authenticate(JMXPluggableAuthenticator.java:173) at sun.management.jmxremote.ConnectorBootstrap$AccessFileCheckerAuthenticator.authenticate(ConnectorBootstrap.ja va:181) at javax.management.remote.rmi.RMIServerImpl.doNewClient(RMIServerImpl.java:212) at javax.management.remote.rmi.RMIServerImpl.newClient(RMIServerImpl.java:179) at sun.reflect.GeneratedMethodAccessor131.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at sun.rmi.server.UnicastServerRef.dispatch(UnicastServerRef.java:294) at sun.rmi.transport.Transport$1.run(Transport.java:153) at java.security.AccessController.doPrivileged(Native Method) at sun.rmi.transport.Transport.serviceCall(Transport.java:149) at sun.rmi.transport.tcp.TCPTransport.handleMessages(TCPTransport.java:466) at sun.rmi.transport.tcp.TCPTransport$ConnectionHandler.run(TCPTransport.java:707) at java.lang.Thread.run(Thread.java:595)	1	0	https://issues.apache.org/jira/browse/AMQ-2352
18340	AMQ-2352	Hi Dejan, I am using AMQ 5.3.0 as embedded broker in Tomcat which works fine. For the admin console, I took the patch for SystemPropertiesConfiguration but the activemq-web-console-5.3.0.war still doesn't work. I pass the JAVA_OPTS="-Dwebconsole.type=properties -Dwebconsole.jms.url=tcp://localhost:61616 -Dwebconsole.jmx.url=service:jmx:rmi:///jndi/rmi://localhost:1099/jmxrmi -Dwebconsole.jmx.user=system -Dwebconsole.jmx.password=manager" as stated in documentation. Am I missing something else for this to work in tomcat. Is the admin webapp tested for embedded broker in tomcat ? Thanks Bhupi	1	0	https://issues.apache.org/jira/browse/AMQ-2352

#	Requirement	Questions	Answered	Successfully Answered	URL
18341	AMQ-2352	Hi Bhupi, there shouldn't be any difference between tomcat-embedded and standalone brokers. Are you sure the patch is applied correctly? You can try switching properties of user and password fields to see if the issue is still there. Also, if your broker is secured (and not jmx), you should use -Dwebconsole.jms.user=system -Dwebconsole.jms.password=manager properties.	0	0	https://issues.apache.org/jira/browse/AMQ-2352
18342	AMQ-2352	I downloaded the webconsole again from this location. http://repo2.maven.org/maven2/org/apache/activemq/activemq-web-console/5.3.0/activemq-web-console-5.3.0.war The fix is not included in that version. Where should I look to download a version that includes the fix? Thanks	0	0	https://issues.apache.org/jira/browse/AMQ-2352
18343	AMQ-2343	Hi, this is an old issue, but I can't find if it was already fixed. How is the progress of it?	1	0	https://issues.apache.org/jira/browse/AMQ-2343
18344	AMQ-2330	Could you add the patch ?	0	0	https://issues.apache.org/jira/browse/AMQ-2330
18345	AMQ-2324	Just wanted to add some more context. This means the Spring Integration {{JmsOutboundGateway.sendAndReceive}} will fail in a multi-broker environment as well. Should this be considered a bug or is it really an Improvement? (In other words, is there a workaround?)	0	0	https://issues.apache.org/jira/browse/AMQ-2324
18346	AMQ-2318	Is there an issue for StoreUsage and TempUsage ?	0	0	https://issues.apache.org/jira/browse/AMQ-2318
18347	AMQ-2245	I see the same stack trace today in 5.3.1 after trying to force start a BrokerService that was stopped. Is there a way to replace the ErrorBroker?	0	0	https://issues.apache.org/jira/browse/AMQ-2245
18348	AMQ-2139	Hi I'm working on servicemix 3 and I had some problems with load testing. I have multiple servicemix instances interconnected together with activemq brokers. After long investigation, and JProbe analyses I found that the pb is on the TcpTransport. So I tried your proposition and it seems to work fine (with a new TcpTransportFactory and a PooledTcpTransport! But I think it's not the best solution when using some big messages because there are only one socket used to communicate between two servicemix containers. My questions are : Is it necessary to have a MutexTransport due to the OpenWireFormat marshall method is already serialized ? Do you think it is possible to create a pool of sockets in the Tcp transport to build a network of brokers ? Thanks	0	0	https://issues.apache.org/jira/browse/AMQ-2139
18349	AMQ-2129	Seeing these now.. not 100% sure if ti's related to the patch: {code} RemoteConnection error: java.lang.IndexOutOfBoundsException: Index: 5, Size: 5 java.lang.IndexOutOfBoundsException: Index: 5, Size: 5 RemoteConnection error: java.lang.IndexOutOfBoundsException: Index: 6, Size: 5 at java.util.ArrayList.RangeCheck(ArrayList.java:546) at java.util.ArrayList.get(ArrayList.java:321) at org.apache.activemq.dispatch.AbstractPooledDispatcher.chooseDispatcher(AbstractPooledDispatcher.java:132) at org.apache.activemq.dispatch.PriorityDispatcher$2$1.execute(PriorityDispatcher.java:89) at org.apache.activemq.flow.FlowController.checkUnblockSources(FlowController.java:424) at org.apache.activemq.flow.FlowController.dispatchOverflowQueue(FlowController.java:341) at org.apache.activemq.flow.FlowController.onUnthrottled(FlowController.java:300) at org.apache.activemq.flow.FlowController$1.onUnthrottled(FlowController.java:83) RemoteConnection error: java.lang.IndexOutOfBoundsException: Index: 7, Size: 5 at org.apache.activemq.flow.AbstractLimiter.notifyUnThrottleListeners(AbstractLimiter.java:41) at org.apache.activemq.flow.SizeLimiter.remove(SizeLimiter.java:68) at org.apache.activemq.flow.RemoteConnection$WindowLimiter.remove(RemoteConnection.java:308) at org.apache.activemq.flow.RemoteConnection$WindowLimiter.onProtocolMessage(RemoteConnection.java:323) at org.apache.activemq.flow.RemoteConnection.onCommand(RemoteConnection.java:92) at org.apache.activemq.transport.InactivityMonitor.onCommand(InactivityMonitor.java:210) at org.apache.activemq.transport.TransportSupport.doConsume(TransportSupport.java:84) at org.apache.activemq.transport.tcp.TcpTransport.doRun(TcpTransport.java:234) at org.apache.activemq.transport.tcp.TcpTransport.run(TcpTransport.java:216) at java.lang.Thread.run(Thread.java:613) {code}	0	0	https://issues.apache.org/jira/browse/AMQ-2129
18350	AMQ-2106	Really need this. Can this be started plz!	0	0	https://issues.apache.org/jira/browse/AMQ-2106

#	Requirement	Questions	Answered	Successfully Answered	URL
18351	AMQ-2105	directory - When I start the application, I get the following error and the web admin doesn't work: [org.mortbay.log] jetty-6.1.9 2009-06-11 16:11:20,157 WARN [org.mortbay.log] Can't reuse /tmp/Jetty_0_0_0_0_8161_admin__admin__9l5s6p, using /tmp/Jetty_0_0_0_0_8161_admin__admin__9l5s6p_36110 2009-06-11 16:11:20,429 WARN [org.mortbay.log] Failed startup of context org.mortbay.jetty.webapp.WebAppContext@1958cc2{/admin,file:/usr/local/activemq/webapps/admin/} java.util.zip.ZipException: error in opening zip file at java.util.zip.ZipFile.open(Native Method) at java.util.zip.ZipFile.<init>(ZipFile.java:114) at java.util.jar.JarFile.<init>(JarFile.java:133) at java.util.jar.JarFile.<init>(JarFile.java:97) at org.mortbay.jetty.webapp.TagLibConfiguration.configureWebApp(TagLibConfiguration.java:168) at org.mortbay.jetty.webapp.WebAppContext.startContext(WebAppContext.java:1217) at org.mortbay.jetty.handler.ContextHandler.doStart(ContextHandler.java:510) at org.mortbay.jetty.webapp.WebAppContext.doStart(WebAppContext.java:448) at org.mortbay.component.AbstractLifeCycle.start(AbstractLifeCycle.java:39) at org.mortbay.jetty.handler.HandlerCollection.doStart(HandlerCollection.java:152) at org.mortbay.component.AbstractLifeCycle.start(AbstractLifeCycle.java:39) at org.mortbay.jetty.handler.HandlerWrapper.doStart(HandlerWrapper.java:130) at org.mortbay.jetty.Server.doStart(Server.java:222) at org.mortbay.component.AbstractLifeCycle.start(AbstractLifeCycle.java:39) at org.mortbay.jetty.xbean.JettyFactoryBean.run(JettyFactoryBean.java:47) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.invokeCustomInitMethod(AbstractAutowireCapableBeanFactory.java:1413) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.invokeInitMethods(AbstractAutowireCapableBeanFactory.java:1374) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.initializeBean(AbstractAutowireCapableBeanFactory.java:1334) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.doCreateBean(AbstractAutowireCapableBeanFactory.java:473) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory$1.run(AbstractAutowireCapableBeanFactory.java:409) at java.security.AccessController.doPrivileged(Native Method) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.createBean(AbstractAutowireCapableBeanFactory.java:380) at org.springframework.beans.factory.support.AbstractBeanFactory$1.getObject(AbstractBeanFactory.java:264) at org.springframework.beans.factory.support.DefaultSingletonBeanRegistry.getSingleton(DefaultSingletonBeanRegistry.java:221) at org.springframework.beans.factory.support.AbstractBeanFactory.doGetBean(AbstractBeanFactory.java:261) at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:185) at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:164) at org.springframework.beans.factory.support.DefaultListableBeanFactory.preInstantiateSingletons(DefaultListableBeanFactory.java:429) at org.springframework.context.support.AbstractApplicationContext.finishBeanFactoryInitialization(AbstractApplicationContext.java:729) at	0	0	https://issues.apache.org/jira/browse/AMQ-2105
18352	AMQ-2105	I downloaded activemq-rpm_standalone-project.tar.bz2 unzipped it, built MAVEN_OPTS=-Xmx800m mvn clean install -Dmaven.test.skip=true and it created an RPM [INFO] Installing /opt/junk/activemq-rpm_standalone-project/target/activemq-rpm-5.3-1-SNAPSHOT.rpm to /root/.m2/repository/org/apache/activemq/activemq-rpm/5.3-SNAPSHOT/activemq-rpm-5.3-SNAPSHOT.rpm # ls -l target/activemq-rpm-5.3-1-SNAPSHOT.rpm -rw-r--r-- 1 root root 25722620 Dec 21 09:55 target/activemq-rpm-5.3-1-SNAPSHOT.rpm The RPM says "not relocatable", is there an option available to change this?	0	0	https://issues.apache.org/jira/browse/AMQ-2105
18353	AMQ-2105	I have not used the latest releases of the unix-maven-plugin, but I don't think it currently supports relocatable. Looking at [1], it seems straight forward to add support to unix-m-p. Can you add a jira for it? As far as I can tell ActiveMQ seems "relocatable ready", so it seems doable. In my opinion relocatable is a minor convenience feature, not a blocking issue for adding rpm support to ActiveMQ... [1] http://www.rpm.org/max-rpm/ch-rpm-reloc.html	0	0	https://issues.apache.org/jira/browse/AMQ-2105
18354	AMQ-2105	Hi, thank you for wrining the specfile - i take a quick view at your RPM (just inspected by using cpio): Some ideas/hints from a sysadmin's view: - rpm creation should be automated Maybe you can add activemq to SUSE build service to create RPMs for several plattforms. (see: http://de.opensuse.org/Build_Service) Automatic creation of a daily snapshot based on svn-trunk would be a cool idea. - ActiveMQ persistence data is currently located in /var/cache/activemq - this looks not very LSB compliant. Maybe /var/lib/activemq would be a better way - /etc/activemq should only contain one single configuration (all other activemq-*.xml configurations should be placed in /usr/share/activemq/example/conf) - Basic configuration for the init-script should be moved to /etc/default/activemq - ActiveMQ comes with the java wrapper: /usr/share/activemq/bin/linux-x86-64/libwrapper.so /usr/share/activemq/bin/linux-x86-64/wrapper => it seems that this is the main reason to have separate rpm-archives for different platforms => in my opinion it would be a good idea to deliver activemq as a "noarch" package by using a init-script which do not need a binary helper. (see also my implementation https://issues.apache.org/activemq/browse/AMQ-2453, unluckily i got no feedback what is needed to get it in activemq-svntrunk) => in my opinion it would be a good idea to deliver activemq as a "noarch" package by using a init-script which do not need a binary helper. (noarch packages can i.e. be used on linux-x86/x86_64/s390/arm/.. and also on systems which support rpm like aix/pseries,...) - ActiveMQ should not run as root-user as default I'm also interested in having distribution packages of activemq - if you like i can contribute to the specfiles.... Regards Marc	1	0	https://issues.apache.org/jira/browse/AMQ-2105
18355	AMQ-2105	Hi James, thanks for the great work on this. Can you attach the sources and grant a license for inclusion in Apache project? I'm a bit busy at the moment, but will look at this and related issues (including making start script better) soon. What I'd do as well is to create a PGP signature for it, so its source can be verified (as it is the case with other distributions). Marc, all improvements are welcomed. It'd be great to make ActiveMQ a first-class citizen on Linux.	1	0	https://issues.apache.org/jira/browse/AMQ-2105

#	Requirement	Questions	Answered	Successfully Answered	URL
18356	AMQ-2105	{quote} * rpm creation should be automated Maybe you can add activemq to SUSE build service to create RPMs for several plattforms. (see: http://de.opensuse.org/Build_Service) Automatic creation of a daily snapshot based on svn-trunk would be a cool idea. {quote} Yes, those were just some temp rpms We build it automatically on our Koji (same as fedora) buildserver. Latest version at : [https://koji.afroditi.hellasgrid.gr/koji/packageinfo?packageID=492] {quote} * ActiveMQ persistence data is currently located in /var/cache/activemq - this looks not very LSB compliant. Maybe /var/lib/activemq would be a better way {quote} Yes, probably. Now fixed. {quote} * /etc/activemq should only contain one single configuration (all other activemq-*.xml configurations should be placed in /usr/share/activemq/example/conf) {quote} This seems to be a problem with the upstream tarball creation. In the SVN repo they're in a separate directory and get merged during the build process. I'll submit a patch against the assembly descriptor to change this. {quote} * Basic configuration for the init-script should be moved to /etc/default/activemq {quote} The init.d script loads from /etc/sysconfig/activemq which is the RHEL/fedora equivalent of the debian /etc/default. I currently don't create a /etc/sysconfig/activemq by default. I probably should. {quote} * ActiveMQ comes with the java wrapper: /usr/share/activemq/bin/linux-x86-64/libwrapper.so /usr/share/activemq/bin/linux-x86-64/wrapper => it seems that this is the main reason to have separate rpm-archives for different platforms => in my opinion it would be a good idea to deliver activemq as a "noarch" package by using a init-script which do not need a binary helper. (see also my implementation https://issues.apache.org/activemq/browse/AMQ-2453, unluckily i got no feedback what is needed to get it in activemq-svntrunk) => in my opinion it would be a good idea to deliver activemq as a "noarch" package by using a init-script which do not need a binary helper. (noarch packages can i.e. be used on linux-x86/x86_64/s390/arm/.. and also on systems which support rpm like aix/pseries,...) {quote} Well the different platforms tend to have different rpms anyway since there are differences in how the different platforms lay things out and what is provided for init.d scripts (e.g. start-stop-daemon vs runuser and /etc/default vs /etc/sysconfig. Also we we move to a 'real' rpm built for Fedora from source it'll be x86_64 anyway since it must put the .so files for gcj into the package. Tomcat5 on RHEL uses jakarata-commons-daemon for startup which also has a small binary component. I know there are some licencing problems with JWS which means we can't a newer version - I'll look at a patch to use jakarta-commons-daemon instead and see if it's better. {quote} * ActiveMQ should not run as root-user as default {quote} the rpm creates an 'activemq' user and modifies the init.d to use it by default. Thanks for the feedback, and any modifications you have to make it work on non-RHEL rpm platforms would be welcome ! The latest rpm (5.3.2-3) is here : [https://koji.afroditi.hellasgrid.gr/koji/buildinfo?buildID=623] and I'll attach the tarball of the scripts to this ticket now.	0	0	https://issues.apache.org/jira/browse/AMQ-2105
18357	AMQ-2092	just wondering when this occurs, would it make sense to spit out the toString() of the object? Is it something got to do with how TransportConnection is shutdown?	1	0	https://issues.apache.org/jira/browse/AMQ-2092
18358	AMQ-2040	Is this for browsing messages from the webconsole?	1	0	https://issues.apache.org/jira/browse/AMQ-2040
18359	AMQ-2018	Are you seeing errors when deploying ActiveMQ to the dmServer? I've just deployed both 5.1 and 5.2 to the SpringSource dmServer with any issues.	0	0	https://issues.apache.org/jira/browse/AMQ-2018
18360	AMQ-2006	woo, can you reattach your patch and give it the apache love. that is be sure and indicate that you wish to assign rights to apache.	1	1	https://issues.apache.org/jira/browse/AMQ-2006
18361	AMQ-1978	But do you really want to have the full text in a toString() value? I'd limit it to 60 chars or so.	0	0	https://issues.apache.org/jira/browse/AMQ-1978
18362	AMQ-1978	A hard limit of 60 could be good, better may be something of the form "abc....xyz" so the beginning and the end of the text are visible. need for a limit seems to be reflected in the tests, I see a bunch TextMessage related failures like : ---------- Test set: org.apache.activemq.openwire.v1.ActiveMQTextMessageTest ---------- Tests run: 1, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 0.046 sec <<< FAILURE! testGenerateAndReParsingIsTheSame(org.apache.activemq.openwire.v1.ActiveMQTextMessageTest) Time elapsed: 0.02 sec <<< ERROR! java.lang.OutOfMemoryError: Requested array size exceeds VM limit at java.lang.AbstractStringBuilder.<init>(AbstractStringBuilder.java:44) at java.lang.StringBuffer.<init>(StringBuffer.java:92) at org.apache.activemq.util.MarshallingSupport.readUTF8(MarshallingSupport.java:320) at org.apache.activemq.command.ActiveMQTextMessage.getText(ActiveMQTextMessage.java:85) at org.apache.activemq.command.ActiveMQTextMessage.toString(ActiveMQTextMessage.java:152) at java.lang.String.valueOf(String.java:2615) at java.lang.StringBuilder.append(StringBuilder.java:116) at org.apache.activemq.openwire.DataFileGeneratorTestSupport.testGenerateAndReParsingIsTheSame(DataFileGeneratorTestSupport.java:97) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at junit.framework.TestCase.runTest(TestCase.java:154) at junit.framework.TestCase.runBare(TestCase.java:127) at junit.framework.TestResult$1.protect(TestResult.java:106) at junit.framework.TestResult.runProtected(TestResult.java:124) at junit.framework.TestResult.run(TestResult.java:109) at junit.framework.TestCase.run(TestCase.java:118) at junit.framework.TestSuite.runTest(TestSuite.java:208) at junit.framework.TestSuite.run(TestSuite.java:203) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585) at org.apache.maven.surefire.junit.JUnitTestSet.execute(JUnitTestSet.java:210) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.executeTestSet(AbstractDirectoryTestSuite.java:135) at org.apache.maven.surefire.suite.AbstractDirectoryTestSuite.execute(AbstractDirectoryTestSuite.java:160) at org.apache.maven.surefire.Surefire.run(Surefire.java:81) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:585)	1	1	https://issues.apache.org/jira/browse/AMQ-1978

#	Requirement	Questions	Answered	Successfully Answered	URL
18363	AMQ-1928	Hi, Could you please let us know what is the time line of providing the exception for clients? If I interpret "Fix Version/s: 5.4.0" correctly then, this means to be fixed in 5.4.0. What would be an approximate date for the release? We see there still no proper exception for client in apache-activemq-5.3.1, when setting this <transportConnector name="openwire" uri="tcp://0.0.0.0:61616?maximumConnections=2"/> after sometime client just bails out with javax.jms.JMSException: Wire format negotiation timeout: peer did not send his wire format. Will it work for STOMP as well? Thank you, Konstantin	1	0	https://issues.apache.org/jira/browse/AMQ-1928
18364	AMQ-1919	FWIW the current code just swaps SESSION_TRANSACTED to be AUTO_ACKNOWLEDGE which seems even worse?	1	0	https://issues.apache.org/jira/browse/AMQ-1919
18365	AMQ-1894	Attached patch (configuration.patch) is a proposal for new configuration files for 5.3.0. First of all it trims down the default configuration as Bruce proposed. I also set kahaDB as a default persistence mechanism for 5.3.0. Next we have camel and jetty configurations pulled out into separate configuration files activemq-camel.xml and activemq-jetty.xml so that they can be easily imported into other configuration files. Currently inclusion of both of these files are commented in default configuration and I guess the first question is whether we want to keep it that way. activemq-demo.xml is pretty much the same as the current configuration file, but with kahaDB and producer flow control explicitly turned on. I created some more configuration examples. They are all located assembly/src/sample-conf. I renamed and modified some of the existing ones (network examples - thanks Andreas) and added some more. All files with activemq-*.xml pattern from that folder are included into conf folder of the distribution. There are some other files there, that I think are outdated, so we should probably remove them (another debating question) Here is the list of new configuration examples: activemq-jdbc.xml - demonstrates jdbc storage activemq-security- demonstrates secured broker activemq-command - demonstrates command agent and xmpp protocol activemq-static-network-broker1.xml and activemq-static-network-broker2.xml - demonstrates static network setup activemq-dynamic-network-broker1.xml and activemq-dynamic-network-broker2.xml - demonstrates multicast network setup shared file system master slave can be demonstrated using the default configuration file. What would be good to have as well: activemq-scalability.xm and activemq-throughput.xml which will show how to tune broker for these two use cases. I'll start working on them now. Please post your comments on these. If there are no major objections, we can start by committing this patch and working further on polishing them up. More comments in the files and wiki documentation would be great. Cheers	0	0	https://issues.apache.org/jira/browse/AMQ-1894
18366	AMQ-1888	is /schemas/ relevant past 5.0? /schema/ (note singular) has the necessary files for 5.x /schemas/ is a link to the m1 repo, I guess we could copy the relevant files to the m2 repo in a similar location but I wonder who references /schemas/ ?	1	0	https://issues.apache.org/jira/browse/AMQ-1888
18367	AMQ-1853	It is nice that there will finally be done something about this problem, eventhough I am not sure exactly that it is the same problem as I have been struggling with. This issue talks about the problem occuring if just on message is to be redelivered. The problem I have, only completely blocks delivery of new fresh messages, when there is more than maxMessagesPerSessions * maxSessions messages waiting for redelivery on the same consumer at the same time. I am using MDBs in Glassfish v2.1 as the consumer - maybe that makes the difference. For those with access - see more on https://fusesource.com/issues/browse/DEV-2279. Will this fix also solve my problem? Regards, Per Steffensen	1	0	https://issues.apache.org/jira/browse/AMQ-1853
18368	AMQ-1853	I see a problem. I downloaded AMQ 5.6-SNAPSHOT, configured it and my ConnectionFactory (wrapped) bean looks like this: <bean id="jmsConnectionFactory" class="org.apache.activemq.pool.PooledConnectionFactory" destroy-method="stop"> <property name="connectionFactory"> <bean class="org.apache.activemq.ActiveMQConnectionFactory"> <property name="brokerURL"> <value>${jms.broker.urls}</value> </property> <property name="redeliveryPolicy"> <bean class="org.apache.activemq.RedeliveryPolicy"> <property name="maximumRedeliveries" value="4"/> <property name="initialRedeliveryDelay" value="3000"/> <property name="useExponentialBackOff" value="true"/> <property name="backOffMultiplier" value="2"/> </bean> </property> <property name="nonBlockingRedelivery" value="true" /> </bean> </property> </bean> The problem is that if 'nonBlockingRedelivery' is set to TRUE AFTER RedeliveryPolicy property, redelivering takes much more time than specified in RedeliveryPolicy object properties. On the other hand - while 'nonBlockingRedelivery' is put BEFORE declaration od RedeliveryPolicy - it takes under consideration only 'maximumRedeliveries' property - but with no delay and multiplier. That makes whole redelivery thing totally wrong - i got defult 6 requests one by one. Is this normal and expected behaviour? How can I achieve taking redeliveryPolicy properties to run?	0	0	https://issues.apache.org/jira/browse/AMQ-1853
18369	AMQ-1853	Is the URL parameter for this nonBlockingRedelivery=true, jms.nonBlockingRedelivery=true, jms.redeliveryPolicy.nonBlockingRedelivery=true, or something else entirely? (*Chris*)	1	0	https://issues.apache.org/jira/browse/AMQ-1853
18370	AMQ-1832	2ying Do you think your patch can solve my problems AMQ-1825 ?	0	0	https://issues.apache.org/jira/browse/AMQ-1832
18371	AMQ-1832	I think your solution is bad. How can I stop Master? Stopping working master means that my business process will crashed... What's the problem if it is just a network problem, not server crashing? In this case slave must automaticaly reconnect to master and replicate state, it mustn't shutdown. Data replication must be independent from Command architecture... Propogating commands to slave is bad architecture. In fact, absense of command means absence of data replication.	0	0	https://issues.apache.org/jira/browse/AMQ-1832
18372	AMQ-1804	Reopening ticket due to xbean name mapping conflict that may cause problems for some users. I propose either removing the xbean mapping from the org.apache.activemq.ActiveMQConectionFactory class or removing the org.apache.activemq.spring.ActiveMQConnectionFactory class. Do we need both of these classes?	1	0	https://issues.apache.org/jira/browse/AMQ-1804
18373	AMQ-1757	It seems we missed September. What needs to be done now and how can I help?	1	0	https://issues.apache.org/jira/browse/AMQ-1757
18374	AMQ-1732	Would it be possible to add an ack mode to acknowledge all messages since last ack'ed message (but not all those delivered), sort of like a sliding window protocol?	1	0	https://issues.apache.org/jira/browse/AMQ-1732
18375	AMQ-1689	BTW.. not 100% sure how much this will help. So lets say that sync call is taking a long time to complete.. the client writes will start gettting batched up in the next batch.. then by the time the sync completes the next write batch it gets will be > 1 since several calls should have batched up.	0	0	https://issues.apache.org/jira/browse/AMQ-1689
18376	AMQ-1676	Hi Joseph when you attached the patch you did not click on the option that states you grant the ASF a ASL 2.0 license on your patch. We can't commit this change until you re-attach and click that option.	0	0	https://issues.apache.org/jira/browse/AMQ-1676

#	Requirement	Questions	Answered	Successfully Answered	URL
18377	AMQ-1660	Perhaps it would be better if we try to work around this problem with postgress so that you don't have to restart your broker. (and if you did I would recommend you just use a cron job to do that). Wonder if we use 2 lock statements in 2 different transactions. After some time we commit transaction 1 and get the lock again. Then we do the same with the 2nd transaction. and keep alternating. This should avoid the old transaction problem. We just need to deal with new brokers coming online and getting 1 of the locks while the master is re-newing the lock. So a salve need to release that lock fast if it can't get the 2nd lock. Make sense?	0	0	https://issues.apache.org/jira/browse/AMQ-1660
18378	AMQ-1657	The documentation explicitly does state that a configuration of multiple brokers, which are aware of each other, must be uniquely named. Is this an issue in 5.3? It is very easy to do this, dynamically or hard-coded, so I don't know whether new handling is merited here although I could easily argue that elegant or helpful handling is not a bad idea.	0	0	https://issues.apache.org/jira/browse/AMQ-1657
18379	AMQ-1614	The fix provided in patch is not 100% complete. PooledConnectionFactory provides following 3 constructors: 1. PooledConnectionFactory() 2. PooledConnectionFactory(ActiveMQConnectionFactory connectionFactory) 3. PooledConnectionFactory(String brokerURL) The New class AmqJNDIPooledConnectionFactory fixes JNDI problem but though this class extends PooledConnectionFactory, it is provided with support for zero argument constructor. Can you please provide all 3 constructors for AmqJNDIPooledConnectionFactory class and release patch again?	1	0	https://issues.apache.org/jira/browse/AMQ-1614
18380	AMQ-1609	The question is what would an example STOMP frame look like with type info in it?	0	0	https://issues.apache.org/jira/browse/AMQ-1609
18381	AMQ-1587	Please apply at your leisure. BTW, I'm trying to update (i.e. resize) the images on the Confluence related to this patch. I am unable to delete overwrite my previous uploaded images. Do I need better karma for this? :) I posted a thread about this below. see: http://www.nabble.com/Confluence-Image-Removal-Replacement-td15701101s2354.html	1	1	https://issues.apache.org/jira/browse/AMQ-1587
18382	AMQ-1587	patch applied in rev 645580. Thx! Not sure what the problem is with confluence.. could you rename and upload the images? Perhaps that would be the simple work around?	1	0	https://issues.apache.org/jira/browse/AMQ-1587
18383	AMQ-1567	Thanks. New version of the patch works great. Comment on the actual functionality of the patch: would it be possible that if a transform were not supplied on a stomp subsription, that the message would automatically be translated back to text per the supplied transformation? At the moment, map messages / object messages come back with an empty body if the subscriptions has not supplied a transform. They should probably translate the message backwards to it's closest approximation of the original sent text.	0	0	https://issues.apache.org/jira/browse/AMQ-1567
18384	AMQ-1567	This is the behavior of the legacy transformer. I have on my todo list to add the transformation-error header in case that specified transformer is not found and we can handle somehow this situation too. I also want to make a cache of translators, so we don't have to make one so often. I've planed to make another patch for these functionalities, but I can do it in this one as well (probably in the next few days). I'm not sure what transformation should be used to transform map and object messages to text by default. Do you have any suggestions? Maybe to use xml transformation by default? As for the Jettison driver, can you provide some examples (test cases) so I can reproduce it. I think this is the easiest way to do it, and of course we can fix Jettison bugs and consistency issues (it is a simple Stax driver).	0	0	https://issues.apache.org/jira/browse/AMQ-1567
18385	AMQ-1516	Please submit a patch if you think it'll help. Wanna create a patch against trunk?	1	0	https://issues.apache.org/jira/browse/AMQ-1516
18386	AMQ-1513	Rob, can you commit this please?	1	1	https://issues.apache.org/jira/browse/AMQ-1513
18387	AMQ-1513	is the intention of the patch to split the options between producer and consumer? seems like the patch was partially applied but there is no say 'Consumer' options section, just options. If you have a moment, would it be possible to capture another diff against trunk of exactly what you want?	1	1	https://issues.apache.org/jira/browse/AMQ-1513
18388	AMQ-1493	Hey Dejan, when you attached the patch you did not select the "ASF Granted License" option. Could you just confirm real quick that you are contributing this to the ASF under the ASL 2.0 lic? Thanks!	0	0	https://issues.apache.org/jira/browse/AMQ-1493
18389	AMQ-1492	Can we resolve this now Claus?	1	1	https://issues.apache.org/jira/browse/AMQ-1492
18390	AMQ-1477	With this change, my MulticastDiscoveryAgent will only get "localhost" as the broker name no matter what value I set in broker's activemq.xml Is this a bug? My application is broken when upgrading from 4.1.1 to 5.x @@ -398,7 +353,7 @@ if (selfService != null) { String payload = getType(); payload += started.get() ? ALIVE : DEAD; - payload += DELIMITER + brokerName + DELIMITER; + payload += DELIMITER + "localhost" + DELIMITER; payload += selfService; try { byte[] data = payload.getBytes();	0	0	https://issues.apache.org/jira/browse/AMQ-1477
18391	AMQ-1476	Why don't you reconfigure the logging, so that only errors from the org.apache.activemq.store.jdbc.DefaultDatabaseLocker category are logged? I think that's the fastest way to get rid of the problem, it doesn't even require a recompilation. The addition of the line log4.logger.org.apache.activemq.store.jdbc.DefaultDatabaseLocker=WARN to the log4j.properties in the conf folder of ActiveMQ should do the job (I didn't test this).	1	0	https://issues.apache.org/jira/browse/AMQ-1476
18392	AMQ-1470	Any idea if this will be fixed for the 4.x release? If so, what does the timeline look like?	0	0	https://issues.apache.org/jira/browse/AMQ-1470
18393	AMQ-1377	hi Dan!! i am using your latest updated amq.js file with cook's prototype file and eliminate the amq_prototype_adapter.js.. the problem is that when i want to call AjaxServlet with the following code Ajax.Request(uri, { method: 'post', postBody: body, onSuccess: org.activemq.Amq.endBatch}); then the AjaxServlet does not get called? Could u plz tell me any solution and review my attached amq.js? Thanks	1	0	https://issues.apache.org/jira/browse/AMQ-1377
18394	AMQ-1377	hi dan!! Could u please explain me vat does following means? var sendPoll = function() { // Workaround IE6 bug where it caches the response // Generate a unique query string with date and random var now = new Date(); var data = 'timeout=' + timeout * 1000 + '&d=' + now.getTime() + '&r=' + Math.random(); var options = { method: 'get', data: data, success: pollHandler, error: pollErrorHandler}; adapter.ajax(uri, options); }; u'hv already set 25 sec that the long-polling socket will stay connected. does it mean u r going to set 25000 secs for timeout parameter which seems quite unrealistic.. iam confused.. by using ur file the following scenerio occurs: 1. ActiveMQ and tomcat running on my machine and if i disabled network connection from my machine the pooling doesnt break, js continously send request for message to MessageListenerServlet 2. but other machines running same application through my IP stop sending sending requests to servlet means polling breaks here, and on the establishment of network again js become able to send request to servlet and response also come back to activeMQ but pooling doesnt occur so that js fail to get response from servlet.. hope i explain the possible scenario.? Thanks	1	0	https://issues.apache.org/jira/browse/AMQ-1377

#	Requirement	Questions	Answered	Successfully Answered	URL
18395	AMQ-1377	hi dan!! Could u please explain me vat does following means? var sendPoll = function() { // Workaround IE6 bug where it caches the response // Generate a unique query string with date and random var now = new Date(); var data = 'timeout=' + timeout * 1000 + '&d=' + now.getTime() + '&r=' + Math.random(); var options = { method: 'get', data: data, success: pollHandler, error: pollErrorHandler}; adapter.ajax(uri, options); }; u'hv already set 25 sec that the long-polling socket will stay connected. does it mean u r going to set 25000 secs for timeout parameter which seems quite unrealistic.. iam confused.. by using ur file the following scenerio occurs: 1. ActiveMQ and tomcat running on my machine and if i disabled network connection from my machine the pooling doesnt break, js continously send request for message to MessageListenerServlet 2. but other machines running same application through my IP stop sending sending requests to servlet means polling breaks here, and on the establishment of network again js become able to send request to servlet and response also come back to activeMQ but pooling doesnt occur so that js fail to get response from servlet.. hope i explain the possible scenario.?	1	0	https://issues.apache.org/jira/browse/AMQ-1377
18396	AMQ-1377	The timeout parameter takes milliseconds, so the multiplication by 1000 converts to seconds. It is not 25000 seconds, but 25000 milliseconds, which is 25 seconds. The other parameters ('d' and 'r') are just added to make the query string unique. IE6 has a bug where it will not refetch the exact query it has seen before over GET and instead shows you the cached result (regardless of any no-cache directives from the server). I'm not sure what your problem is. Like I said, this has been working fine in my app for nearly 6 months. You will need to debug your own unique scenario, I suggest using the firebug plugin, setting breakpoints in the scripts and seeing what is happening. The only thing that comes to mind is that if you are actually terminating the Tomcat server, that is a very different thing than a client having intermittent connectivity. My solution is really just to handle the network cable getting unplugged or the wifi being flaky, it is not designed to handle a server restart. When you restart the server, the session that your client had (and any topic subscriptions) are probably no longer valid, and so the client may need to make a new request for the topic..... maybe something like this: //Un-Listen from topic org.activemq.Amq.removeListener("yourTopic", "topic://Your.Topic"); //One second delay to wait for the Un-Listen to happen //Then Re-Listen setTimeout(function(){ org.activemq.Amq.addListener("yourTopic", "topic://Your.Topic", yourHandler); }, 1000);	1	0	https://issues.apache.org/jira/browse/AMQ-1377
18397	AMQ-1377	hi Dan!! I have added your's modified amq.js files in my application, its work quite fine. Thanks. I want to know one thing is there any impact of pooling on activeMQ performance? Actually ActiveMQ got disconnected after giving the following exception at console "Could not accept connection : Channel was inactive for too long" and also i got this exception in application log: "Uncategorized exception occured during JMS processing; nested exception is javax.jms.JMSException: Could not connect to broker URL: tcp://localhost:61616. Reason: java.net.ConnectException: Connection refused; nested exception is java.net.ConnectException: Connection refused" I wonder due to pooling ActiveMQ becomes slower than Ajax client, is it true or my misconception? Thanks	0	0	https://issues.apache.org/jira/browse/AMQ-1377
18398	AMQ-1377	Is it on purpose that amq.js starts with a poll, even before you have the chance of setting up a listener? This usually means a client initially has to wait "timeout" seconds without having a chance of receiving something.	1	1	https://issues.apache.org/jira/browse/AMQ-1377
18399	AMQ-1377	Arjan, I created an example that successfully used the prototype adapter. I then tried to use your new dojo adapter but did not have much success. Whenever I add a listener it does the POST operation but the polling terminates shortly after. The current poll finishes; but a new one does not start. It seems that the dojo.xhrGet does not get a chance to do its callback to keep the polling alive. Did you have any issues with your dojo adapter?	0	0	https://issues.apache.org/jira/browse/AMQ-1377
18400	AMQ-1377	Dejan any advice on why I could not get this to work? Thanks	1	1	https://issues.apache.org/jira/browse/AMQ-1377
18401	AMQ-1377	is there an eta for the 5.4 release that will contain this fix?	1	0	https://issues.apache.org/jira/browse/AMQ-1377
18402	AMQ-1377	Nathanael, there's no hard date on 5.4 release, we're working hard on making it happen in the next couple of months. Mark, can you post your client code and configuration to the user forum? Cheers	1	0	https://issues.apache.org/jira/browse/AMQ-1377
18403	AMQ-1361	This needs to be rolled back - can the patches be resubmitted but with Apache licence headers please ?	0	0	https://issues.apache.org/jira/browse/AMQ-1361
18404	AMQ-1281	This fix wouldn't happen to all fix the queue browser problem in the Web Console? Even when using the Web Console within ActiveMQ, if all of the destinations are password protected, the Web Console will crash horribly if you try to browse a queue, with the following error. I think your patch won't help, because it seems that the queue browser is using JMS, not JMX. org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'queueBrowser' defined in ServletContext resource [/WEB-INF/webconsole-embedded.xml]: Instantiation of bean failed; nested exception is org.springframework.beans.BeanInstantiationException: Could not instantiate bean class [org.apache.activemq.web.QueueBrowseQuery]: Constructor threw exception; nested exception is javax.jms.JMSException: User name or password is invalid. Caused by: org.springframework.beans.BeanInstantiationException: Could not instantiate bean class [org.apache.activemq.web.QueueBrowseQuery]: Constructor threw exception; nested exception is javax.jms.JMSException: User name or password is invalid. Caused by: javax.jms.JMSException: User name or password is invalid. at org.apache.activemq.util.JMSExceptionSupport.create(JMSExceptionSupport.java:46) at org.apache.activemq.ActiveMQConnection.syncSendPacket(ActiveMQConnection.java:1211) at org.apache.activemq.ActiveMQConnection.ensureConnectionInfoSent(ActiveMQConnection.java:1293) at org.apache.activemq.ActiveMQConnection.createSession(ActiveMQConnection.java:294) at org.apache.activemq.web.SessionPool.createSession(SessionPool.java:97) at org.apache.activemq.web.SessionPool.borrowSession(SessionPool.java:80) at org.apache.activemq.web.QueueBrowseQuery.(QueueBrowseQuery.java:41) ...	0	0	https://issues.apache.org/jira/browse/AMQ-1281
18405	AMQ-1281	Seems like this patch has been committed? Should this one be closed?	1	0	https://issues.apache.org/jira/browse/AMQ-1281
18406	AMQ-1273	We have exact same problem in our application at the moment, i'm trying an approach where i am starting the broker in the separate thread. any other tips i can use? is there a patch available for this feature?	1	0	https://issues.apache.org/jira/browse/AMQ-1273

#	Requirement	Questions	Answered	Successfully Answered	URL
18407	AMQ-1241	Since I don't want to be the mean guy I actually did something about it :) The attached patch allows 3 modes of operation, chosen via a VM-Flag (webconsole.type) * jndi (default): Gets a connection factory from "jms/connectionFactory" and the jmx-connection-string (f.e. service:jmx:rmi:///jndi/rmi://localhost:1099/jmxrmi) from "jmx/url" * properties: Uses the system properties webconsole.jms.url and webconsole.jmx.url for the urls of the JMS resp. JMX connectors of the broker * invm: Uses BrokerRegistry.findFirst() for the broker (fixes the current bug with custom named brokers) All modes support the following (well, the invm-mode doesn't really^^): * jmx-Failover: Syntax: <url1>, <url2>, <url3>, ... (f.e. "service:jmx:rmi:///jndi/rmi://serverA:1099/jmxrmi, service:jmx:rmi:///jndi/rmi://serverB:1099/jmxrmi") that chooses the first available broker * jms-Failover: Standard behaviour of the connectionfactory, no code changed Other things changed: * Fixed the handling of the 'query'-object in the request-context of spring (was bugged before, they wouldn't get cleaned up at the end of the request, leading to leftover open consumers) * removed any direct usage of vm://localhost since that often leaded to not-wanted broker startups * explicit initialization of the WebClient class (so it'd use the right ConnectionFactory) * Changed the WebClient to use the JMS-classes instead of the ActiveMQ specific ones (removes problems with application-server-provided session/connection pooling) Things that don't work yet or are missing: * QueueBrowserServlet: This class does own connection handling and creating. I couldn't find out why it'd do that, but I didn't dare doing a that deep change. So atm the RSS/Atom feeds are broken in the web-console * running the war inside the broker (former standard-configuration) doesn't work due to classloader issues (specific: BrokerRegistry is twice in the classpath). I'd suggest to reconfigure the jetty classloading or to remove the active-mq jars from the 'in-broker-war'. * the patch misses the change to the startup-scripts to set the system-property "webconsole-type" to invm * documentation (apart from this here) I'd be happy to discuss this change with you guys and get some input about the changes I made. I don't look at them at finished now, but as a step in the direction. Thanks and greetings Mario	1	1	https://issues.apache.org/jira/browse/AMQ-1241
18408	AMQ-1238	Hi Hiram, Can you take a look at this? I'm not a linux/debian expert so I can't verify it. :)	1	0	https://issues.apache.org/jira/browse/AMQ-1238
18409	AMQ-1237	- The idea is to use one "conf/credentials.properties" file which will define activemq.username and activemq.password properties used to create a connection to a broker from console - By default, the connection is always created with system/manager credentials, but that can be changed at one place for all connections - TODO: add the similar support for command agent and camel context (check if there's any other possible connections) so that broker can start all its services when secured	1	1	https://issues.apache.org/jira/browse/AMQ-1237
18410	AMQ-1204	The class activemq-core/src/main/java/org/apache/activemq/store/jdbc/adapter/SybaseDBCAdapter.java was created to resolve this issue. However there is a TYPO in the name of the class: it should be SybaseJDBCAdapter instead of SybaseDBCAdapter Otherwise, it cannot be automatically discovered. Additionally, this class is not included in the latest 4.2.0 snapshot. Could you please include?	1	1	https://issues.apache.org/jira/browse/AMQ-1204
18411	AMQ-1191	What is the current status of this issue? I would like to use Master/Slave failover together with a MS SQL 2005 installation. From reading different issues linking to this problem I am not getting a clear picture if this is or will be resolved.	0	0	https://issues.apache.org/jira/browse/AMQ-1191
18412	AMQ-1191	modified the database locker code to iterate on a result set if one is present, committed in r787990 The logic is conditional on getMetaData returning null if the lock statement will not result in a ResultSet, some don't. Can you verify that this works with your driver?	1	1	https://issues.apache.org/jira/browse/AMQ-1191
18413	AMQ-1157	Hello, I downloaded yesterday the 5.0-Snapshot soruce and debugged (cause I run into the same problem) ... even if above mentioned changes are applied still another change is missing: Class: AbstractRegion.java Method: start() Here you find this: ConnectionContext context = new ConnectionContext(); context.setBroker(broker.getBrokerService().getBroker()); context.getBroker().addDestination(context, dest); Also here the SecurityContext need to be added: context.setSecurityContext(SecurityContext.BROKER_SECURITY_CONTEXT); ConnectionContext context = new ConnectionContext(); context.setBroker(broker.getBrokerService().getBroker()); context.setSecurityContext(SecurityContext.BROKER_SECURITY_CONTEXT); context.getBroker().addDestination(context, dest); Can this be made soon?	1	0	https://issues.apache.org/jira/browse/AMQ-1157
18414	AMQ-1127	Is this still an issue do you think Tim?	1	0	https://issues.apache.org/jira/browse/AMQ-1127
18415	AMQ-1091	So, shouldn't this result in a download option for source on the AMQ 4.1 download page?	1	0	https://issues.apache.org/jira/browse/AMQ-1091

#	Requirement	Questions	Answered	Successfully Answered	URL
18416	AMQ-1082	for has indeed been deleted. There's still plenty of free space on the disk. ERROR DataManagerImpl - Looking for key 1 but not found in fileMap: {2=data-TopicSubscription-2 number = 2 , length = 33554006 refCount = 12755, 4=data-TopicSubscription-4 number = 4 , length = 33554158 refCount = 17435, 6=data-TopicSubscription-6 number = 6 , length = 20247053 refCount = 10538, 3=data-TopicSubscription-3 number = 3 , length = 33553411 refCount = 17037, 5=data-TopicSubscription-5 number = 5 , length = 33552256 refCount = 16291} ERROR ListContainerImpl - Failed to get value for offset=102, key=(-1, -1, 0), value=(1, 2001, 2022), previousItem=-1, nextItem=204 java.io.IOException: Could not locate data file data-TopicSubscription-1 at org.apache.activemq.kaha.impl.data.DataManagerImpl.getDataFile(DataManagerImpl.java:121) at org.apache.activemq.kaha.impl.data.SyncDataFileReader.readItem(SyncDataFileReader.java:60) at org.apache.activemq.kaha.impl.data.DataManagerImpl.readItem(DataManagerImpl.java:130) at org.apache.activemq.kaha.impl.container.ListContainerImpl.getValue(ListContainerImpl.java:841) at org.apache.activemq.kaha.impl.container.ContainerValueCollectionIterator.next(ContainerValueCollectionIterator.java:48) at org.apache.activemq.broker.region.cursors.FilePendingMessageCursor.next(FilePendingMessageCursor.java:186) at org.apache.activemq.broker.region.TopicSubscription.dispatchMatched(TopicSubscription.java:347) at org.apache.activemq.broker.region.TopicSubscription.acknowledge(TopicSubscription.java:191) at org.apache.activemq.broker.region.AbstractRegion.acknowledge(AbstractRegion.java:312) at org.apache.activemq.broker.region.RegionBroker.acknowledge(RegionBroker.java:401) at org.apache.activemq.broker.TransactionBroker.acknowledge(TransactionBroker.java:181) at org.apache.activemq.broker.BrokerFilter.acknowledge(BrokerFilter.java:74) at org.apache.activemq.broker.BrokerFilter.acknowledge(BrokerFilter.java:74) at org.apache.activemq.broker.MutableBrokerFilter.acknowledge(MutableBrokerFilter.java:88) at org.apache.activemq.broker.TransportConnection.processMessageAck(TransportConnection.java:456) at org.apache.activemq.command.MessageAck.visit(MessageAck.java:179) at org.apache.activemq.broker.TransportConnection.service(TransportConnection.java:258) at org.apache.activemq.broker.TransportConnection$1.onCommand(TransportConnection.java:164) at org.apache.activemq.transport.ResponseCorrelator.onCommand(ResponseCorrelator.java:95) at org.apache.activemq.transport.TransportFilter.onCommand(TransportFilter.java:65) at org.apache.activemq.transport.vm.VMTransport.syncOneWay(VMTransport.java:99) at org.apache.activemq.transport.vm.VMTransport.oneway(VMTransport.java:86) at org.apache.activemq.transport.MutexTransport.oneway(MutexTransport.java:44) at org.apache.activemq.transport.ResponseCorrelator.oneway(ResponseCorrelator.java:60) at org.apache.activemq.ActiveMQConnection.asyncSendPacket(ActiveMQConnection.java:1165) at	1	0	https://issues.apache.org/jira/browse/AMQ-1082
18417	AMQ-1082	Hi Albert, Is this still the case ? I'm sure this was fixed a week or so ago cheers, Rob	1	0	https://issues.apache.org/jira/browse/AMQ-1082
18418	AMQ-1023	The change to TopicSubscription looks ugly. Is there some reason that this information is not updated for all destinations. Also couldn't singleDestination simply replaced by a null check?	0	0	https://issues.apache.org/jira/browse/AMQ-1023
18419	AMQ-1008	James.Strachan wrote: > As for the temporary queue/topic, the JMS[1] spec states that temporary > queues/topics may be used by any session using that connection. This seems >> sensible and should probably be a mechanism embraced by STOMP. If connecting >> to an ActiveMQ server, then the destinations /temporary/queue/foo, or >> /queue/temp/foo would be nice. >Yeah. I guess the decision needs to be, does the client make up some >name and have the broker auto-create it, or should there be an >explicit CREATE/DESTROY type verbs on the STOMP protocol. An explicit client request to CREATE temporary queue/topic where the server creates and return via CREATED frame would have been be ideal. One possible server side implementation could associate the temp queue/topic to the session, and automatically deletes it when session ends which is consistent JMS side behavior. I suggested this a while back but the consensus was not to change the stomp protocol, rather extend it for ActiveMQ. Thus, I have proposed a simple temp queue/topic naming convention which doesn't require the Stomp protocol change or extension: destination: /tempqueue/Foo or destination: /temptopic/Bar This temp destination must be globally unique name as in the case of JMS side. And the server somehow associates the temp queue/topic to the session, and deletes it without a client request when the session ends. Mostly a temp queue is created so that it can be used for replyTo. >I guess the simplest thing that could possibly work is just to use a >naming convention; that /temporary/queue/* (or /tempqueue/*) indicates >temporary queues. Though the complication then arises that a client >may wish to specify the reply-to destination (more here >http://incubator.apache.org/activemq/stomp.html) - which if there was >a client side alias (/tempqueue/foo) that would need to get switched >to the globally unique queue name. I prefer /tempqueue/Foo, /temptopic/Bar type convention which quite similar to regular queue/topic: /queue/Foo, /topic/Bar >I'm not sure how useful the SESSION is; as folks can use a transaction >for a session? A stomp connection can be viewed as an implicit session. Having multiple session within a stomp connection may introduce complexity -Sileshi	1	0	https://issues.apache.org/jira/browse/AMQ-1008
18420	AMQ-984	a few years later but BrokerService has an isSlave() method, Explosing that readonly attribute would give the same information I guess. It will only be true while waiting for the lock, when the lock is acquired it will be false. Would this suffice?	0	0	https://issues.apache.org/jira/browse/AMQ-984
18421	AMQ-982	Hi, There is no authenticated connection documentation for embedded broker either. Following information may help. ---------- How to enable authentication for embedded broker? - Download activemq-jaas-x.x.x.jar and add to the CLASSPATH - Create conf directory and add to the CLASSPATH - Put login.config into conf directory - Put users.properties and groups.properties into conf\org\apache\activemq\security - Use following code for initializing the embedded broker // Server initialization BrokerService broker = new BrokerService(); broker.addConnector("tcp://localhost:61616"); broker.setPlugins(new BrokerPlugin[] { new JaasAuthenticationPlugin () }); broker.start();	1	0	https://issues.apache.org/jira/browse/AMQ-982
18422	AMQ-982	Thanks for this Ali, its awesome! BTW the entire website is a wiki so you can edit or add any information there directly... http://activemq.apache.org/how-does-the-website-work.html http://activemq.apache.org/how-do-i-edit-the-website.html do you fancy adding your documentation changes directly to the wiki? You'll then get more credit for your work as you'll be seen as the contributor to the documentation. You can attach files to wiki pages too	1	0	https://issues.apache.org/jira/browse/AMQ-982
18423	AMQ-943	Hey Dejan, when you attached the patch you did not select the "ASF Granted License" option. Could you just confirm real quick that you are contributing this to the ASF under the ASL 2.0 lic?	0	0	https://issues.apache.org/jira/browse/AMQ-943

#	Requirement	Questions	Answered	Successfully Answered	URL
18424	AMQ-912	Hi Sepand, I reviewed the patch and I noticed that the standard apache licence header is not on the new files included in the patch. Could you please add the headers and create a new diff? Thanks, Hiram	1	1	https://issues.apache.org/jira/browse/AMQ-912
18425	AMQ-912	Hi Sepand, After I apply the latest patch, there's a build failure in the Qdox stuff when I run mvn package. What did you have to do to resolve that before? Any ideas? Thanks, -Kelly java.lang.ArrayIndexOutOfBoundsException: 0 at org.apache.xbean.spring.generator.QdoxMappingLoader.loadElements(QdoxMappingLoader.java:147) at org.apache.xbean.spring.generator.QdoxMappingLoader.loadNamespaces(QdoxMappingLoader.java:107) at org.apache.xbean.spring.generator.QdoxMappingLoader.loadNamespaces(QdoxMappingLoader.java:101) at org.apache.xbean.maven.XBeanMojo.execute(XBeanMojo.java:153) at org.apache.maven.plugin.DefaultPluginManager.executeMojo(DefaultPluginManager.java:412) at org.apache.maven.lifecycle.DefaultLifecycleExecutor.executeGoals(DefaultLifecycleExecutor.java:534) at org.apache.maven.lifecycle.DefaultLifecycleExecutor.executeGoalWithLifecycle(DefaultLifecycleExecutor.java:475) at org.apache.maven.lifecycle.DefaultLifecycleExecutor.executeGoal(DefaultLifecycleExecutor.java:454) at org.apache.maven.lifecycle.DefaultLifecycleExecutor.executeGoalAndHandleFailures(DefaultLifecycleExecutor.java:306) at org.apache.maven.lifecycle.DefaultLifecycleExecutor.executeTaskSegments(DefaultLifecycleExecutor.java:273) at org.apache.maven.lifecycle.DefaultLifecycleExecutor.execute(DefaultLifecycleExecutor.java:140) at org.apache.maven.DefaultMaven.doExecute(DefaultMaven.java:322) at org.apache.maven.DefaultMaven.execute(DefaultMaven.java:115) at org.apache.maven.cli.MavenCli.main(MavenCli.java:256)	1	0	https://issues.apache.org/jira/browse/AMQ-912
18426	AMQ-908	Here are the Geronimo ones: [GeronimoUserPrincipal|http://svn.apache.org/viewvc/geronimo/server/trunk/modules/geronimo-security/src/main/java/org/apache/geronimo/security/realm/providers/GeronimoUserPrincipal.java?view=markup] [GeronimoGroupPrincipal|http://svn.apache.org/viewvc/geronimo/server/trunk/modules/geronimo-security/src/main/java/org/apache/geronimo/security/realm/providers/GeronimoGroupPrincipal.java?view=markup] I'm sure Glassfish has some too, though I don't know where. Perhaps Harmony as well? My assumption was that we should look for a constructor with a single String, and if we don't find that, look for an empty constructor and a setName method that takes a single String, and if we don't find that, throw an Exception. If someone complains that they got the exception, then I guess we'll look at their principal classes to determine the next fallback position. :)	0	0	https://issues.apache.org/jira/browse/AMQ-908
18427	AMQ-892	The JMX RMI port can be set in ActiveMQ 4.x using the JMX system properties. See the [How can I monitor ActiveMQ? page|http://activemq.org/site/how-can-i-monitor-activemq.html#HowcanImonitorActiveMQ-AdvancedJMXConfiguration] under the Advanced JMX Configuration heading. If the com.sun.management.jmxremote.port system property is set within the SUNJMX environment variable, the ActiveMQ start script will pick it up and pass it to the JVM.	0	0	https://issues.apache.org/jira/browse/AMQ-892
18428	AMQ-877	I tried to apply this patch to the current trunk and it broke test cases. I wonder if trunk has moved along a bit since your patch. Could you try create another patch please and we can try again? Do the test cases all work for you with your patch applied?	0	0	https://issues.apache.org/jira/browse/AMQ-877
18429	AMQ-868	I might be missing something but the only CMS code I know of is under the activemq-cpp space in SVN which is current and tagged release work. Are you referring to the CPP clients that exist in sandbox?	0	0	https://issues.apache.org/jira/browse/AMQ-868
18430	AMQ-855	> e.g. if we did implement prefetch of zero - which means don't deliver a message to a consumer at all - > unless they perform a consumer.receive() - even then, the consumer could then hang/deadlock and never > actually acknowledge or process the message. If consumer hangs and message is not processed is a different problem. Even if messages are lost in such case, nobody would blame AMQ. If you kill consumer and re-despatch, it would be brilliant. But if my consumers do not hang, crash or burn, can I just get my messages? James, could you please re-open this issue? I spend 1 month working on this project and now I will have to throw it all away and do explaining with my manager. I cannot apply my own fix, you will not fix it. What am I supposed to do now?	1	1	https://issues.apache.org/jira/browse/AMQ-855
18431	AMQ-855	James, I added implementation of pulling consumer whereby pullMessage() returns a response. Could you please browse through it and give me a feedback? Patches are for file versions PrefetechSubscription.java 1.15 and 1.22 ActiveMQMessageConsumer.java Current implementation does not work for me, because when recieve(long) call times out the prefetch window is increasing and we are back to where we started. Thanks a lot.	1	1	https://issues.apache.org/jira/browse/AMQ-855
18432	AMQ-855	'm going to start looking into it today. I may not use the DataResponse to return the message right away. I think I would still prefer to use the normal message dispatching flow and logic. What I'm thinking is that on timeout, or when receiveNoWait() is called, and no message has arrived to be dispatched, we dispatch a message to the consumer letting it know that no message was available. Then it can unblock. The nice thing about this is that you don't get timing issues between the client and server where the client times out but the broker actually dispatched the message.	1	1	https://issues.apache.org/jira/browse/AMQ-855
18433	AMQ-838	I haven't looked through the code for this at all yet, but maybe an option to force the deletion would be enough? Can someone at least provide a stack trace for this problem?	1	0	https://issues.apache.org/jira/browse/AMQ-838
18434	AMQ-827	That's kind of silly. :) What's wrong with suffixes, which are pretty pervasive at this point? :) Heck, anyone accustomed to configuring a Java heap knows that 100m means 100MiB. I think it's a lot more intuitive.	0	0	https://issues.apache.org/jira/browse/AMQ-827
18435	AMQ-826	Thanks James! How can I see the files? I saw test resources on SVN (properties, ldif, xml) but no Java files. What is the next step in actually using this module? How do I build AMQ with this? I'll try to solve ApacheDS issue you reported. I'll be using it on my project so I'll work with it anyway.	1	0	https://issues.apache.org/jira/browse/AMQ-826

#	Requirement	Questions	Answered	Successfully Answered	URL
18436	AMQ-826	Hi, I've been busy with another module. Based on Sepand's contribution, I made LDAPCertificateLoginModule that takes certificate from SSL connection and verifies it against LDAP. I'll commit it together with unit test. Now, about testing with Apache directory: What (I understood) you asked for (and I volounteered) is the ability of unit tests to start (an embedded instance of) directory server loladed with testing data and perform unit tests against it. Right now, I start manually ApacheDS and perform unit tests but we want that automated, right? If that is the case, I believe I can make that easily (I already started ApacheDS as embedded). But, if I understood you correctly, you would like to utilize existing libraries from AMQ project and submit only the difference, is this correct? I'll try to find the differences, too, but my knowledgde of Maven is extremly poor so I won't be able to produce some kond of '.pom' or whatever appropriate. I'll just submit the jars and the code, would that be acceptable?	1	0	https://issues.apache.org/jira/browse/AMQ-826
18437	AMQ-826	I finished improvement together with unit test (running on external ADS). There are two assumptions I want to confirm: 1. Composite destinations ACL set of a composite destination is a union of ACL sets of each particular destination. I deduced this from code (DefaultAuthorizationMap) and implemented the same although it does not seem logical to me. Intersection of sets would be more appropriate, I believe. Should I implement the intersection or leave the union? 2. Wildcard destinations Wildcards are given in authorization policy source (xml map or ldap or...) to allow creation (primarily) of a destination in a certain namespace. Wildcard is ">" which means any destination. This meaning is unlimited in depth i.e. "ActiveMQ.Advisory.>" will suffice both for "ActiveMQ.Advisory.Connection" (> = Connection, same level) and for "ActiveMQ.Advisory.Queue.ABC123" (> = Queue.ABC123, one level more). Is this correct? Should I leave it as it is or restrict ">" to the same level only?	0	0	https://issues.apache.org/jira/browse/AMQ-826
18438	AMQ-826	Thanks for wildcard link. I did not implement '*', I'll finish it as well. Is it possible to have kind of regular expression like STOCKS.PRICE.NYSE.*BM ? Regarding composite destinations, I would like your attention: Union of ACLs means that if a user has privilege on at least one destination, all destinations will allow operation. Intersection of ACLs means that if a user lacks privilege on at least one destination, no destination will allow operation. I'll produce a test to verify this but my point is that current implementation of union is a security leak (if my understanding is correct). Suppose that a guest user wants to read from a destination not authorized for guests, say destination USERS.SECRET. A guest may create a destination in GUEST space with all necessary privileges, say GUEST.ALLOW. Now, the user creates a composite destination (GUEST.ALLOW, USERS.SECRET) and attempts an operation: Case UNION: as operation is permitted on GUEST.ALLOW it is sufficient for composite destination; operation is performed on both destinations in spite of the fact that user is not authorized for the other. Case INTERSECTION: as operation is NOT permitted on USERS.SECRET no operation is attempted on composite destination. Now, maybe I got it wrong but the method 'getXXXXXACLs()' in DefaultAuthorizationMap is pretty clear - it adds all ACLs from all entries...	0	0	https://issues.apache.org/jira/browse/AMQ-826
18439	AMQ-826	was able to run this as it is but I wasn't able to run it in activemq-core and activemq-jaas eclipse projects. I tried to add libraries but I always ended up with exception, both in core and jaas modules: Exception in thread "main" org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'configuration' defined in file [C:\ActiveMQ\activemq\activemq-jaas\server.xml]: Instantiation of bean failed; nested exception is org.springframework.beans.BeanInstantiationException: Could not instantiate bean class [org.apache.directory.server.configuration.MutableServerStartupConfiguration]: Constructor threw exception; nested exception is java.lang.IncompatibleClassChangeError Caused by: org.springframework.beans.BeanInstantiationException: Could not instantiate bean class [org.apache.directory.server.configuration.MutableServerStartupConfiguration]: Constructor threw exception; nested exception is java.lang.IncompatibleClassChangeError Caused by: java.lang.IncompatibleClassChangeError at org.apache.directory.server.core.authn.AuthenticationService.<clinit>(AuthenticationService.java:66) at org.apache.directory.server.core.configuration.StartupConfiguration.setDefaultInterceptorConfigurations(StartupConfiguration.java:161) at org.apache.directory.server.core.configuration.StartupConfiguration.<init>(StartupConfiguration.java:88) at org.apache.directory.server.configuration.ServerStartupConfiguration.<init>(ServerStartupConfiguration.java:65) at org.apache.directory.server.configuration.MutableServerStartupConfiguration.<init>(MutableServerStartupConfiguration.java:43) at sun.reflect.NativeConstructorAccessorImpl.newInstance0(Native Method) at sun.reflect.NativeConstructorAccessorImpl.newInstance(Unknown Source) at sun.reflect.DelegatingConstructorAccessorImpl.newInstance(Unknown Source) at java.lang.reflect.Constructor.newInstance(Unknown Source) at org.springframework.beans.BeanUtils.instantiateClass(BeanUtils.java:82) at org.springframework.beans.BeanUtils.instantiateClass(BeanUtils.java:59) at org.springframework.beans.factory.support.SimpleInstantiationStrategy.instantiate(SimpleInstantiationStrategy.java:52) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.instantiateBean(AbstractAutowireCapableBeanFactory.java:639) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.createBeanInstance(AbstractAutowireCapableBeanFactory.java:625) at org.springframework.beans.factory.support.AbstractAutowireCapableBeanFactory.createBean(AbstractAutowireCapableBeanFactory.java:380) at org.springframework.beans.factory.support.AbstractBeanFactory$1.getObject(AbstractBeanFactory.java:245) at org.springframework.beans.factory.support.DefaultSingletonBeanRegistry.getSingleton(DefaultSingletonBeanRegistry.java:141) at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:242) at org.springframework.beans.factory.support.AbstractBeanFactory.getBean(AbstractBeanFactory.java:156) at org.springframework.beans.factory.support.DefaultListableBeanFactory.preInstantiateSingletons(DefaultListableBeanFactory.java:290) at org.springframework.context.support.AbstractApplicationContext.refresh(AbstractApplicationContext.java:348) at org.springframework.context.support.FileSystemXmlApplicationContext.<init>(FileSystemXmlApplicationContext.java:89) at org.springframework.context.support.FileSystemXmlApplicationContext.<init>(FileSystemXmlApplicationContext.java:74) at	0	0	https://issues.apache.org/jira/browse/AMQ-826
18440	AMQ-816	Hi, would this include also tackling the problem that having a persistent consumer on one broker he won't be forwarded the messages that were persisted in some other? I am thinking about the following case: we have a network of brokers, say brokerA and brokerB. A consumer subscribes with a unique id client1:subscription1 to brokerB. Therefore, if offline, messages will be persisted on brokerB. After a while, brokerB goes down, and the same consumer fails over to brokerA. He's able to receive all new messages, but will not receive persisted messages when brokerB comes back to life. What I would like to see is basically a consumer be recognized as unique over an entire network of brokers and not just on one broker, so messages persisted in one broker could also be forwarded whenever a durable subscription would appear somewhere else on the network. Thank you! daniel	1	0	https://issues.apache.org/jira/browse/AMQ-816
18441	AMQ-805	Hi Tom, do you have a test case?	1	0	https://issues.apache.org/jira/browse/AMQ-805

#	Requirement	Questions	Answered	Successfully Answered	URL
18442	AMQ-805	Does this make sense as a way to deal with the problem? Obviously before committing this we should make the interval configurable, and probably make it exponential. I am also still concerned about thread safety, but that is a separate issue.	0	0	https://issues.apache.org/jira/browse/AMQ-805
18443	AMQ-795	For this plugin we can just have a specific admin role (ie. tempAdmin) for the temporary destinations. Then we can grant by default everyone admin roles for temp destinations and still be able to remove them for certain users. Since temp destination names are generated at runtime we'll need an approach so we can map all temp destinations to a specific role. I was thinking of just adding a setter method inside the authorizationMap to set the authorizationentry for temporary destinations : ie. ..... <authorizationMap> <authorizationEntries> <authorizationEntry queue="QUEUE.USERS.>" "read="admins" write="admins" admin="admins"/> <authorizationEntry topic="TOPIC.USERS.>" read="admins" write="admins" admin="admins" /> </authorizationEntries> <tempAuthorizationEntry tempAdmin="tempAdmins" /> </authorizationMap>\	0	0	https://issues.apache.org/jira/browse/AMQ-795
18444	AMQ-794	I tried this in the broker config and by defaulting advisorySupport to false in BrokerService. I hard coded isAdvisorySupport( ) to return false... Should these still be created?: 2006-07-19 17:22:09,179 [localhost:61616] INFO TransportConnection - Created new TransportConnection. Broker:a01 2006-07-19 17:22:09,201 [127.0.0.1:39964] INFO ActiveMQDestination - Creating destination: topic://ActiveMQ.Advisory.TempQueue 2006-07-19 17:22:09,201 [127.0.0.1:39964] INFO ActiveMQDestination - Creating destination: topic://ActiveMQ.Advisory.TempTopic 2006-07-19 17:22:09,201 [127.0.0.1:39964] INFO RegionBroker - Adding consumer for destination: topic://ActiveMQ.Advisory.TempQueue,topic://ActiveMQ.Advisory.TempTopic 2006-07-19 17:22:09,201 [127.0.0.1:39964] INFO AbstractRegion - Adding consumer: ID:photon.duncllc.com-39963-1153347728978-1:0:-1:1 2006-07-19 17:22:09,204 [127.0.0.1:39964] INFO AbstractRegion - Adding subscription to destination:Topic: destination=topic://ActiveMQ.Advisory.TempQueue,topic://ActiveMQ.Advisory.TempTopic, subscriptions=0	0	0	https://issues.apache.org/jira/browse/AMQ-794
18445	AMQ-788	Fixing in trunk revision 418486 Should we also do this for the 4.0.2 release?	0	0	https://issues.apache.org/jira/browse/AMQ-788
18446	AMQ-730	Which are the parts of the javadoc which you think most need improvement? http://incubator.apache.org/activemq/maven/activemq-core/apidocs/	0	0	https://issues.apache.org/jira/browse/AMQ-730
18447	AMQ-729	Can you reproduce this on 4.0.1?	0	0	https://issues.apache.org/jira/browse/AMQ-729
18448	AMQ-729	FWIW, I tried to execute the test case but on Windows machine (which obviously changes everything) and with 4.0.1. It all worked. Has anyone else been able to reproduce this behaviour with 4.0.1?	0	0	https://issues.apache.org/jira/browse/AMQ-729
18449	AMQ-702	Adrian, Do you known when you will get around to fixing this?	1	0	https://issues.apache.org/jira/browse/AMQ-702
18450	AMQ-688	How can we collaborate on this work? The described functionality is crucial to ActiveMQ's success within our organization.	1	0	https://issues.apache.org/jira/browse/AMQ-688
18451	AMQ-660	The statement that is being executed is: DELETE FROM ACTIVEMQ_MSGS WHERE (EXPIRATION<>0 AND EXPIRATION<?) OR ID <= (SELECT min(ACTIVEMQ_ACKS.LAST_ACKED_ID) FROM ACTIVEMQ_ACKS WHERE ACTIVEMQ_ACKS.CONTAINER=ACTIVEMQ_MSGS.CONTAINER) Johan, does MaxDB support subselects?	1	1	https://issues.apache.org/jira/browse/AMQ-660
18452	AMQ-660	OK, I got an answer back from one of the MaxDB devs. Apparently, the issue is a known bug, that is fixed in the latest 7.5 version of it (and assume in the 7.6 branch, too). Thus, long term there is no issue here. However, since I'm running it all on Debian Sarge, I'm relying on the Debian team to potentially backport the fix into their stable package version, or run their testing package of MaxDB instead. Until I can look into that, is the ACTIVEMQ_MSGS_EIDX index critical to the functioning of activemq, or can I drop it for now and still expect reasonable performance?	1	1	https://issues.apache.org/jira/browse/AMQ-660
18453	AMQ-660	OK, I've manually dropped the offending index, and it seems not to be recreated by adding the <createSchemaStatements/> element to the config as explained above (were I to deploy activemq on a new machine using the same MaxDB version, I assume that I'd need first to start the broker without a createSchemaStatements tag to have it create the full schema, and then manually drop the insex, insert the element in activemq.xml and restart, correct?). However, adding the deleteOldMessages attribute to the statements element gives this error when starting the broker: ERROR: java.lang.RuntimeException: Failed to execute start task. Reason: org.springframework.beans.factory.BeanDefinitionStoreException: Line 26 in XML document from class path resource [activemq.xml] is invalid; nested exception is org.xml.sax.SAXParseException: The value of attribute "deleteOldMessagesStatement" associated with an element type "statements" must not contain the '<' character. ERROR: java.lang.Exception: org.springframework.beans.factory.BeanDefinitionStoreException: Line 26 in XML document from class path resource [activemq.xml] is invalid; nested exception is org.xml.sax.SAXParseException: The value of attribute "deleteOldMessagesStatement" associated with an element type "statements" must not contain the '<' character. Removing the less than character fixes the issue, but obviously makes the statement wrong. Is there some way to rewrite the less than character in activemq.xml? Finally, so far I'm still using the statements attributes that were needed to get MaxDB to operate with MaxDB in the first place (binaryDataType ="LONG BYTE" longDataType ="INTEGER" stringIdDataType="VARCHAR(250) ASCII"). Can I remove them, now that the snapshots starting on 3/29 contain automatic support for MaxDB, or are they still needed?	1	1	https://issues.apache.org/jira/browse/AMQ-660
18454	AMQ-660	A further piece of input: I changed the deleteOldMessagesStatement to be "DELETE FROM ACTIVEMQ_MSGS WHERE ID = (SELECT min(ACTIVEMQ_ACKS.LAST_ACKED_ID) FROM ACTIVEMQ_ACKS WHERE ACTIVEMQ_ACKS.CONTAINER=ACTIVEMQ_MSGS.CONTAINER) (I know that just dropping the < character probably makes the statement not very useful, but wanted to move past the error I reported above), and let the broker run for a while (with the ACTIVEMQ_MSGS_EIDX index dropped). I then get this message regularly: WARN JDBCPersistenceAdapter - Old message cleanup failed due to: com.sap.dbtech.jdbc.exceptions.JDBCDriverException: SAP DBTech JDBC: Column index 1 was not found. Did my change of the delete statement cause the dropped index to be needed, or do I need to make other changes to the statements used?	1	1	https://issues.apache.org/jira/browse/AMQ-660
18455	AMQ-660	Hiram, Using the deleteOldMessagesStatement you specify above does make the "Column index 1 was not found" go away, and the example producer and consumer still work. I just have one remaining question: The offending ACTIVEMQ_MSGS_IDX index still seems to be automatically created on startup, even though I have the empty createSchemaStatements element in activemq.xml like this: <persistenceAdapter> <journaledJDBC journalLogFiles="5" dataDirectory="../activemq-data" dataSource="#maxdb-ds"> <statements> <statements deleteOldMessagesStatement="DELETE FROM ACTIVEMQ_MSGS WHERE <>? AND ID <= (SELECT min(ACTIVEMQ_ACKS.LAST_ACKED_ID) FROM ACTIACTIVEMQ_ACKS WHERE ACTIVEMQ_ACKS.CONTAINER=ACTIVEMQ_MSGS.CONTAINER)"/> <createSchemaStatements /> </statements> </journaledJDBC> </persistenceAdapter> Do I need to nest the statements elements in some other way to get the createSchemaStatements to have the desired effect?	1	1	https://issues.apache.org/jira/browse/AMQ-660

#	Requirement	Questions	Answered	Successfully Answered	URL
18456	AMQ-657	Hi Kevin, could you re test this issue against 4.0 final? The failover transport should bre reaisnt those interupted and resumed events now.	1	1	https://issues.apache.org/jira/browse/AMQ-657
18457	AMQ-657	Hiram, I will test this out. Just applied the 4.0 (incubator) release to my test system the other day. Btw: when will 4.0 get out of incubator?	1	1	https://issues.apache.org/jira/browse/AMQ-657
18458	AMQ-656	Mats, thank you for that update and the patch. I have run the unittests with Broker 4.0 on Windows XP. The unittest are ok, but however there is an "unknown command" when closing the connection after receiving the Shutdown-message. Here is the last part of the output when running with uri: "tcp://localhost:61616/trace=true" . . . Send/receive a text message to a queue synchronously Sending command: cmd.id = 35, corr.id = -1, type = SESSION_INFO Received command: cmd.id = 32, corr.id = 35, type = RESPONSE Sending command: cmd.id = 36, corr.id = -1, type = CONSUMER_INFO Received command: cmd.id = 33, corr.id = 36, type = RESPONSE Sending command: cmd.id = 37, corr.id = -1, type = PRODUCER_INFO Received command: cmd.id = 34, corr.id = 37, type = RESPONSE Sending command: cmd.id = 38, corr.id = -1, type = ACTIVEMQ_TEXT_MESSAGE Received command: cmd.id = 35, corr.id = -1, type = ACTIVEMQ_MSG_DISPATCH Received command: cmd.id = 36, corr.id = 38, type = RESPONSE Sending command: cmd.id = 39, corr.id = -1, type = ACTIVEMQ_MSG_ACK Sending command: cmd.id = 40, corr.id = -1, type = REMOVE_INFO Received command: cmd.id = 37, corr.id = 40, type = RESPONSE Sending command: cmd.id = 41, corr.id = -1, type = REMOVE_INFO Received command: cmd.id = 38, corr.id = 41, type = RESPONSE Sending command: cmd.id = 42, corr.id = -1, type = REMOVE_INFO Received command: cmd.id = 39, corr.id = 42, type = RESPONSE [OK] Sending command: cmd.id = 43, corr.id = -1, type = REMOVE_INFO Received command: cmd.id = 40, corr.id = 43, type = RESPONSE Sending command: cmd.id = 44, corr.id = -1, type = SHUTDOWN Received command: cmd.id = 41, corr.id = -1, type = SHUTDOWN ERROR: Unknown command: ♂ Received exception = 'Unmarshal failed; unknown data structure type 204, at c:\programme\incubator-activemq-4.0\openwire-cpp-bugfix\src\main\cpp\activemq\protocol\openwire\openwireprotocol.cpp line 144' ERROR: Received a broker exception: Unmarshal failed; unknown data structure type 204, at c:\programme\incubator-activemq-4.0\openwire-cpp-bugfix\src\main\cpp\activemq\protocol\openwire\openwireprotocol.cpp line 144 Exiting read loop due to exception: Unmarshal failed; unknown data structure type 204, at c:\programme\incubator-activemq-4.0\openwire-cpp-bugfix\src\main\cpp\activemq\protocol\openwire\openwireprotocol.cpp line 144 --------- Where does this additional stuff come from? Thanks for your help Heinrich	1	0	https://issues.apache.org/jira/browse/AMQ-656
18459	AMQ-648	Hi, You could configure the management context in the broker config using the following code snippet: <broker useJmx="true"> <managementContext> <managementContext connectorPort="2011" connectorPath="/mypath"/> </managementContext> .... </broker> You could do the same thing for POJOs as well, AFAIK. BTW, as a minor note, the console doesn't support it yet if you change the domain name. :) Is this what you are looking for? Regards, Adrian Co	1	0	https://issues.apache.org/jira/browse/AMQ-648
18460	AMQ-637	Sounds fine with me :) BTW is the idea that the socket will only be open for the duration of the input stream, so that the transfer is independent of the underlying JMS connection? We do need to chunk the message on the broker side to avoid it having to load the whole stream into RAM; is this OK with you?	0	0	https://issues.apache.org/jira/browse/AMQ-637
18461	AMQ-629	Adiran, Is this still broken?	0	0	https://issues.apache.org/jira/browse/AMQ-629
18462	AMQ-608	Hiram, Thanks. I will use svn diff in the future. So, what is the time frame of 4.0.1? Like I commented in the other issue, when will 4.0 get out of incubator?	1	1	https://issues.apache.org/jira/browse/AMQ-608
18463	AMQ-596	what's the change in configuration to use this feature? I didnt see it in the docs...	1	0	https://issues.apache.org/jira/browse/AMQ-596
18464	AMQ-591	How does this relate to bug AMQ-775? http://issues.apache.org/activemq/browse/AMQ-775	0	0	https://issues.apache.org/jira/browse/AMQ-591
18465	AMQ-517	Nathan, can we resolve this issue now?	1	0	https://issues.apache.org/jira/browse/AMQ-517
18466	AMQ-498	I'm unable to crash the transport by doing this with 5.1.0, Hiram. Is this still a problem in 5.1 or have the commands been changed to handle this correctly?	0	0	https://issues.apache.org/jira/browse/AMQ-498
18467	AMQ-498	Fyi, this change broke us. In the future you should try to make changes that are backwards compatible. If they are not, they should be explicitly mentioned in the release notes. In this case, how many people really need it? It could have been that if the maxFrameSize is not explicitly set, then frame size checking would be disabled - as before. I will say, it is a great feature for those who have publicly exposed brokers. Thanks!	1	0	https://issues.apache.org/jira/browse/AMQ-498
18468	AMQ-486	Hi James, added only basic implementation for Referenceable. I'm not sure if what is needed is the whole jndi stuff. (i.e. creating the JNDIStorableInterface interface, JNDIBaseStorable abstract class, ObjectFactory, etc.). Let me know if you require the whole thing. :) Thanks! Is the object factory needed? I'm not that familiar yet with jndi, so I wasn't sure. :)	1	1	https://issues.apache.org/jira/browse/AMQ-486
18469	AMQ-486	Ported jndi implementation of references from 3.x to 4.x (just copied it actually). I noticed that the initial context does not use references, but stores the whole object, should this be change to just use the object reference?	0	0	https://issues.apache.org/jira/browse/AMQ-486
18470	AMQ-388	Richard, as a start could you let us know what the JDBC driver name is? e.g. {code} connection = ... connection.getMetaData().getDriverName() {code} Also what is the closest JDBC language to it? i.e. whats the most efficient way of dealing with BLOBs via JDBC - using getBytes()? Or streams/blobs?	0	0	https://issues.apache.org/jira/browse/AMQ-388
18471	AMQ-378	Is this still relevant? Does the slow consumer functionality not cover this?	1	0	https://issues.apache.org/jira/browse/AMQ-378
18472	AMQ-370	Is the sleep time formula correct anyway ? It would be nice to support different backoff strategies. I'd prefer to have the redelivery time computed like this: sleepTime = initialRedeliveryTimeout * pow(backoffIncreaseRate, deliverCount) truncated to some maximum value.	1	0	https://issues.apache.org/jira/browse/AMQ-370
18473	AMQ-360	Is there any documentation on how the total ordering of topics across the network will be achieved? Can you point us to any white papers on the issue. We would like to know if this feature will be in ActiveMQ 5.3.0 and when the release date for 5.3.0 be. Thanks Jim....	0	0	https://issues.apache.org/jira/browse/AMQ-360
18474	AMQ-340	This sounds like WSN root then maps to a broker instance. Since a JVM can run more than 1 broker, the WSN layer just needs to route requests to different brokers based on the root. What do you think?	0	0	https://issues.apache.org/jira/browse/AMQ-340
18475	AMQ-321	Do we know if this is still the case in 4.0?	1	0	https://issues.apache.org/jira/browse/AMQ-321
18476	AMQ-318	How to configure the checkpointInterval of the JournalPersistenceAdapter in 5.3.1 release? We are planning to use high performance Journal + Mysql permanent persistent . So our config is: <persistenceAdapter> <journaledJDBC dataDirectory="${activemq.base}/broker" dataSource="#mysql-ds" /> </persistenceAdapter> But we are also planning to control the checkpointInterval, so puzzling how to configure it. Thanks lot.	1	0	https://issues.apache.org/jira/browse/AMQ-318

#	Requirement	Questions	Answered	Successfully Answered	URL
18477	AMQ-318	Looks like this is not possible in latest releases. Could you please let me know how to configure the Journal's checkpointinterval in the latest releases? Is it possible? Or any workaround? We are using 5.4.0 and the config we are using is: <persistenceFactory> <journalPersistenceAdapterFactory journalLogFiles="5" dataDirectory="${basedir}/activemq-data" dataSource="#mysql-ds /> </persistenceFactory> Any help would be really appreciated. Thank you. Krishna	1	0	https://issues.apache.org/jira/browse/AMQ-318
18478	AMQ-305	SET AUTOCOMMIT TO OFF is no longer supported in Postgre8.1 When i start the broker with JDBCPersistence the following error occurs Nov 17, 2007 2:21:27 PM org.apache.activemq.store.jdbc.JDBCPersistenceAdapter createAdapter INFO: Database driver recognized: [postgresql_native_driver] Exception in thread "main" java.io.IOException: ERROR: SET AUTOCOMMIT TO OFF is no longer supported at org.apache.activemq.util.IOExceptionSupport.create(IOExceptionSupport.java:43) at org.apache.activemq.store.jdbc.TransactionContext.getConnection(TransactionContext.java:62) at org.apache.activemq.store.jdbc.TransactionContext.begin(TransactionContext.java:147) at org.apache.activemq.store.jdbc.JDBCPersistenceAdapter.start(JDBCPersistenceAdapter.java:156) at com.glenwood.Queue.StartBroker.createPersistenceAdapter(StartBroker.java:49) at com.glenwood.Queue.StartBroker.main(StartBroker.java:28) Caused by: java.sql.SQLException: ERROR: SET AUTOCOMMIT TO OFF is no longer supported at org.postgresql.core.QueryExecutor.execute(QueryExecutor.java:131) at org.postgresql.jdbc1.AbstractJdbc1Connection.ExecSQL(AbstractJdbc1Connection.java:482) at org.postgresql.jdbc1.AbstractJdbc1Connection.ExecSQL(AbstractJdbc1Connection.java:461) at org.postgresql.jdbc1.AbstractJdbc1Connection.setAutoCommit(AbstractJdbc1Connection.java:953) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(Unknown Source) at java.lang.reflect.Method.invoke(Unknown Source) at org.postgresql.ds.common.PooledConnectionImpl$ConnectionHandler.invoke(PooledConnectionImpl.java:360) at $Proxy0.setAutoCommit(Unknown Source) at org.apache.activemq.store.jdbc.TransactionContext.getConnection(TransactionContext.java:58) ... 4 more I am using postgresql-8.1-409.jdbc3.jar with Postgresql 8.1 ActiveMq.xml follows <persistenceAdapter> <journaledJDBC journalLogFiles="5" dataDirectory="../persistencedata" dataSource="#postgres-ds"/> </persistenceAdapter> <bean id="postgres-ds" class="org.postgresql.ds.PGPoolingDataSource"> <property name="serverName" value="1.1.1.180"/> <property name="databaseName" value="merge_emr_billing"/> <property name="portNumber" value="5432"/> <property name="user" value="pcare"/> <property name="password" value="pcare007"/> <property name="dataSourceName" value="postgres"/> <property name="initialConnections" value="1"/> <property name="maxConnections" value="10"/> </bean>	0	0	https://issues.apache.org/jira/browse/AMQ-305
18479	AMQ-269	Hi Ramzi We're hopefully gonna code freeze 3.1 on friday (22nd July) to get ready for 2 weeks of system testing prior to the 3.1 release. If you get this issue resolved by then, great it can go out in the 3.1 release - if not, I'll edit this issue to say that the fix version is 3.2. Good luck :)	1	1	https://issues.apache.org/jira/browse/AMQ-269
18480	AMQ-244	matthew.clark@despatch.net reported on the list: Hi there, just thought I would point out that the parseBoolean method has been used in the Producer/Consumer example. Obviously this method was only introduced in Java 1.5 so breaks the examples for 1.4.x users. Does this indicate a shift to Java 1.5 or was this a slip of the hand?	0	0	https://issues.apache.org/jira/browse/AMQ-244
18481	AMQ-197	Dominic, Fritz, Could you attach you server's log file to this incident. Would make fixing this problem easier.	1	1	https://issues.apache.org/jira/browse/AMQ-197
18482	AMQ-197	Is this still an issue? There's been a fair amount of work tidying up support for SQL Server - could you double check this is still an issue in 3.0-SNAPSHOT?	0	0	https://issues.apache.org/jira/browse/AMQ-197
18483	AMQ-196	I've removed the sleep() calls and used a semaphore now instead. I'm hoping this resolves any issues with Hermes as well - could someone verify please? By all means reopen this issue if you still have an issue.	0	0	https://issues.apache.org/jira/browse/AMQ-196
18484	AMQ-186	Yes, I will create a patch to add this functionality. I think I found another problem though. While the soTimeout property is exposed. The "read loop" (run() method) is mishandling it. I found the following code: if (changeTimeout) { socket.setSoTimeout(soTimeout); } if (changeTimeout) { socket.setSoTimeout(0); } This will set the read timeout, then immediately turn if off again before actually performing the read. Without the CVS history I am not able to determine when this problem was introduced. Does anyone have any knowledge of when the soTimeout property support and if the second call to socket.setSoTimeout() was just a coding mistake or intentional?	1	0	https://issues.apache.org/jira/browse/AMQ-186
18485	AMQ-183	Some notes.. For a Derby database, when a create table is issued when the table already exists, SQLException.getSQLState() will return "X0Y32" (see LANG_OBJECT_ALREADY_EXISTS_IN_OBJECT in Derby's SQLState class). If we were to use the getTables(catalog,schemaPattern,tableNamePattern,types[]) method in the java.sql.DatabaseMetaData class, couldn't it return more than one table if the table was create under two different schemas (e.g. created by two different users)? Do we have enough information available to pass to getTables(..) to reliably detect if the table has already been created?	0	0	https://issues.apache.org/jira/browse/AMQ-183
18486	AMQ-183	Well, that's a common problem. Suppose the table does exist in another schema, how does AMQ know if the DBA meant for it to use that one or to create another in this schema? The "attempt create" model currently used assumes one configuration. Another issue with using metadata is the presence of synonyms - there may not be a table with that name, but there may be a synonym that refers to a different one. The main reason I opened this issue is that there is no real benefit to the user in having the stack trace dumped and it just adds confusion; I think just logging a message would be an improvement. More advanced check on metadata would be useful but I agree that handling all the different scenarios possible in dbms land are hard (we dealt with many of the same things wanting to create tables for CMP EJBs).	0	0	https://issues.apache.org/jira/browse/AMQ-183

#	Requirement	Questions	Answered	Successfully Answered	URL
18487	AMQ-122	The new supportJMSPriority configuration attribute is not supported in 4.0-M4. When I try start my broker with that attribute set I get the following error: ERROR: java.lang.RuntimeException: Failed to execute start task. Reason: org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'org.apache.activemq.xbean.XBeanBrokerService' defined in class path resource [activemq.xml]: Error setting property values; nested exception is org.springframework.beans.NotWritablePropertyException: Invalid property 'supportJMSPriority' of bean class [org.apache.activemq.xbean.XBeanBrokerService]: Bean property 'supportJMSPriority' is not writable or has an invalid setter method: Does the parameter type of the setter match the return type of the getter? ERROR: java.lang.Exception: org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'org.apache.activemq.xbean.XBeanBrokerService' defined in class path resource [activemq.xml]: Error setting property values; nested exception is org.springframework.beans.NotWritablePropertyException: Invalid property 'supportJMSPriority' of bean class [org.apache.activemq.xbean.XBeanBrokerService]: Bean property 'supportJMSPriority' is not writable or has an invalid setter method: Does the parameter type of the setter match the return type of the getter? I am using a 4.0-M4 downloaded on 9 March. Also - to use the priority, is it necessary to set the priority on both the producer and the messages?	0	0	https://issues.apache.org/jira/browse/AMQ-122
18488	AMQ-122	I am using activeMQ 4.0.1 version and i added "jmsPriority" enabled in my broker's config file, <broker name="myBroker" useJmx="true" supportJMSPriority="true"> .... </broker> When i started the broker on command line, I got the following error: $ bin/activemq ACTIVEMQ_HOME: c:\activemq401 Loading message broker from: xbean:activemq.xml ERROR: java.lang.RuntimeException: Failed to execute start task. Reason: org.springframework.beans.factory.BeanCreationException: Error creating bean with name 'myBroker' defined in class path resource [activemq.xml]: Error setting property values; nested exception is org.springframework.beans.NotWritablePropertyExcept ion: Invalid property 'supportJMSPriority' of bean class [org.apache.activemq.xb ean.XBeanBrokerService]: Bean property 'supportJMSPriority' is not writable or h as an invalid setter method: Does the parameter type of the setter match the ret urn type of the getter? ERROR: java.lang.Exception: org.springframework.beans.factory.BeanCreationExcept ion: Error creating bean with name 'myBroker' defined in class path resource [ac tivemq.xml]: Error setting property values; nested exception is org.springframew ork.beans.NotWritablePropertyException: Invalid property 'supportJMSPriority' of bean class [org.apache.activemq.xbean.XBeanBrokerService]: Bean property 'suppo rtJMSPriority' is not writable or has an invalid setter method: Does the paramet er type of the setter match the return type of the getter? Is this problem happening only on 4.0.1 version of activeMQ ? is it fixed in 4.2 ? Please help and clarify how i can enable priority for messages, producers and consumers. - Sreenivas Sama	0	0	https://issues.apache.org/jira/browse/AMQ-122
18489	AMQ-122	Hi, I'm having the same issue. In my case, I cannot simply use a message selector because I need to have a fixed amount of consumers. I have a maximum of six consumers allowed, and they are used to process both low and high-priority messages. I would like to be able to use my six listeners to the maximum of their capacity instead of dedicating a certain number to low-priority messages and the rest to high-priority messages... Could someone please point me to a patch or fix this issue ? Thanks ! Philippe	1	0	https://issues.apache.org/jira/browse/AMQ-122
18490	AMQ-122	Is this issue really fixed - as in having added code that solves the problem? Or is it rather resolved as in "Wont-fix" - because there is a suggested workaround in the resequencer pattern?	0	0	https://issues.apache.org/jira/browse/AMQ-122
18491	AMQ-116	How about, if we added an AdminMBean so that you could shut down the broker by calling the 'stop' method in a JMX console? We should have a JMX-JMS bridge soon as well; so management commands could be called by sending a message as well. So then we'd have both options - using a JMX console and by sending a message to a management queue/topc/	1	1	https://issues.apache.org/jira/browse/AMQ-116
18492	AMQ-107	Vey nice! It would be cool if resolved issues contained a link the associated documentation when appopriate. Like in this case, how is max attempts set? Can it be set per queue? Can it be changed at runtime? Can a unique deadletter queue be specified per delivery queue?	0	0	https://issues.apache.org/jira/browse/AMQ-107
18493	AMQ-103	Per spec, ObjectMessages are allways serialized on send so I think to original intent of this issue is actually a non-issue. What was an issue is if two ear/wars were deployed each with his own private copy of activemq and the user wants to have the apps talk to each other without going over tcp/ip. Simple scenario is that webapp 1 is deployed with an embeded broker and webapp 2 wants to connect to it using a vm transport. Since the ActiveMQ classes are not == due to being loaded in different classloaders the current vm transport will throw class cast exceptions it tries to talk to the broker. Solution: use the activeio:vmpipe://localhost transport. The activeio vmpipe moves mesage in low level packets that are easy to bridge from one classloader to another. Yes, ActiveMQ packets do need to serialize down the byte[]s but the network stack overhead is eliminated and you also get a direct invocation on the broker like the current vm transport works.	0	0	https://issues.apache.org/jira/browse/AMQ-103
18494	AMQ-88	Could you give us some more details to go on - e.g. how are you running these commands? Is the broker still running when the consumer tool starts up? What command line arguments are you using? Both the producer and consumer tools are working fine for us over here. e.g. using the example instructions... http://activemq.codehaus.org/Examples or just running things directly in an IDE. Maybe a full dump of the command shell in which you're running the ConsumerTool would help?	1	1	https://issues.apache.org/jira/browse/AMQ-88
18495	AMQ-88	The log you attached looks like the broker's output, not the output of the ConsumerTool. From the broker's perspective, this looks like the consumer tool is terminating and the socket is closing (this could just be a spurious log message which we fixed in 1.1). Does the ConsumerTool receive a message and terminate? Maybe things are working correctly and this is just a spurious exception message that we should be eating (which we are now doing in 1.1 I think)? If not - could you attach the output of the ConsumerTool. If you've already attached the full output, then something strange is going on, since the consumer doesn't throw any exceptions that it can't connect, or has lost connection, yet the broker does which is strange.	1	1	https://issues.apache.org/jira/browse/AMQ-88
18496	AMQ-88	Ya James, you are right that the log which i sent is for the Broker and the ConsumerTool and even this is true that ConsumerTool is getting terminated normally without throwing any exception. But now the question is y does not it pickup the messages when the same are present in the queue and the MessageBroker is running, because i do not find anything wrong with the ConsumerTool's code. Secondly James, what did i observe when i ran ProducerTool(MessageBroker is running) and put some messages, the size of the file "Store.db" under "ActiveMQ" directory does not increase. But when i stop the Broker, that file gets updated and its size increases!! It seems there are many things which are going wrong :)	1	1	https://issues.apache.org/jira/browse/AMQ-88
18497	AMQ-65	Wanna take a stab at this?	0	0	https://issues.apache.org/jira/browse/AMQ-65
18498	AMQ-62	Where is Java serialization currently used? Is it when sending a JMS message that you chose to send serialized objects?	0	0	https://issues.apache.org/jira/browse/AMQ-62

#	Requirement	Questions	Answered	Successfully Answered	URL
18499	AMQ-50	What problem are you having? I've been using the maven rar plugin for a couple of months on activemq with no problems. The source code stays in src/java just like any other module, only the ra.xml etc goes in src/rar/META-INF	1	1	https://issues.apache.org/jira/browse/AMQ-50
18500	AMQ-49	Are we ready now? :)	1	1	https://issues.apache.org/jira/browse/AMQ-49
18501	AMQ-47	Good point. How much should we even try to cope with? Client version < server verion. Server version < client version. And Dito for server to server communications. Should we try to resolve this in 2.0? I think at a minimum we need to allow future servers to cope with old versions of the messages in it's persistence stores.	0	0	https://issues.apache.org/jira/browse/AMQ-47
18502	AMQ-45	What good is a CCI implementation compared to using jms interfaces? Are you intending to suggest that a jca 1.0 outbound-only adapter implementation would be useful?	0	0	https://issues.apache.org/jira/browse/AMQ-45
18503	AMQ-22	Are you refering to providing some visualization tools into the types of messages that are flowing through the broker?	0	0	https://issues.apache.org/jira/browse/AMQ-22
18504	AMQ-21	Could you explain what you mean by this a little more?	0	0	https://issues.apache.org/jira/browse/AMQ-21
18505	AMQ-15	This seems to be more dup checking than reliability. Or is this meant to encompass implementing one of the HTTPR specs?	1	0	https://issues.apache.org/jira/browse/AMQ-15
18506	AMQ-11	Is this sorted now? Can we mark it as fixed?	0	0	https://issues.apache.org/jira/browse/AMQ-11
18507	AMQNET-422	I can compile that code to .NET 2.0 at least in VS 2012 and 2013. Are you using an old compiler? And the only change to None DTC that some code was extracted from Dispatch into a new method DispatchToAsyncListener. Is this really a too big change?	1	0	https://issues.apache.org/jira/browse/AMQNET-422
18508	AMQNET-403	do you have a build?	1	0	https://issues.apache.org/jira/browse/AMQNET-403
18509	AMQNET-403	Hi, we tested it, looks good. url parameter flag is: transport.priorityBackup=true Do you have a paypal account (personal/team) which you can use to accept payment?	0	0	https://issues.apache.org/jira/browse/AMQNET-403
18510	AMQNET-327	Hi, Eddie. Thanks for the feedback. I fully understand that some are unable to upgrade to VS2010. However, would it be too much trouble for you to use the NAnt scripts for building? You will still be able to perform any necessary debugging in your own project from the resulting build output. The only thing you won't be able to do is build from within the IDE, unless you install the free VS2010 Express. I guess I'm looking to find out what you won't be able to do and how much of an inconvenience it will be. There is a certain amount of inconvenience in maintaining multiple IDE build projects, which are themselves redundant to the NAnt scripts.	1	1	https://issues.apache.org/jira/browse/AMQNET-327
18511	AMQNET-290	Very good, I'm really interested in testing this: What's the simplest way ? Via Fuse support ?	1	1	https://issues.apache.org/jira/browse/AMQNET-290
18512	AMQNET-290	Hi, Does the current implementation support both local and distributed transactions ? Another point, if I understand correctly, in distributed mode NMS will be enlisted as durable. But in local mode, NMS will be enlisted as volatile. Am I correct ? If yes, what are the consequences ? Is the message lost ? Regards	1	1	https://issues.apache.org/jira/browse/AMQNET-290
18513	AMQNET-290	Thanks for answer. I understand that enlisting as a durable resource manager is difficult. My only concern is about the side effects. I don't understand what's happening when the code using NMS crashes in the middle of a transaction ? Do we lose the message ? Is the transaction aborted ? Thanks again for the snapshot, I should be able to test it next week.	1	1	https://issues.apache.org/jira/browse/AMQNET-290
18514	AMQNET-271	Instead of a callback interface style of programming, wouldn't it be more .NET-like to use an event driven architecture? For example: The IMessageProducer interface would support a PreSendMessageEvent that could be used to register an event handler for processing the message. Likewise, the IMessageConsumer interface would support a PreDispatchMessageEvent event. This would make implementations cleaner and simpler, by creating a separation of concerns. The implementor would not be forced to implement both transforms, since they could create event handlers individually in case they only want to handle pre-dispatch events and don't care about pre-send events. Just trying to understand this feature a bit more.	1	1	https://issues.apache.org/jira/browse/AMQNET-271
18515	AMQNET-271	Events allow for multiple subscribers which you definitely don't want in this case. Plus I always find the event / delegate model in .NET to be a bit annoying as I have to read a book every time I want to use them. The purpose of this feature is to allow client to configure a producer or consumer to transform a message into something they understand, for instance a client could use to send say an IStreamMessage which they know a Stomp client is listening on an have the MessageProducer apply the transform on the fly to convert it to a TextMessage encoded in a way the Stomp client can process. Or it could be used just as a means of content enrichment. I don't think that you'd want to multicast that given that each receiver could return their own converted message. Which one would you end up sending? We could of course provide a default implementation that overrides both and return null (null right now in my code just means use the original), then the user can use that to allow them to only implement the transform method they wanted. Plus this could add some extra code like helpers to copy standard NMS headers etc from the old to the new message. I had also intended to use this interface to create a generic NMS message Transformer that the client could use to transform NMS messages commands from one client to the command type that's needed in another. For example, the NMS.ActiveMQ client could accept ITextMessage from a MSMQ client and internally transform it into the NMS.ActiveMQ ITextMessage implementation. If there's some slick .NETy way of doing all this then feel free to comment, just make sure it works on the .NETCF 2.0 framework so it can run in NMS.Stomp :) Just wanted to get the ball rolling on something so nothing set in stone right now, comment away....	1	1	https://issues.apache.org/jira/browse/AMQNET-271
18516	AMQNET-271	Yeah, those look good. Although, I would change the naming somewhat to ProducerTransformerDelegate and ConsumerTransformerDelegate. The reason for this is you can have a cleaner name for the instance variable. For example: {code:title=snippet.cs} interface IMessageConsumer { ... ConsumerTransformerDelegate ConsumerTransformer; ... } interface IMessageProducer { ... ProducerTransformerDelegate ProducerTransformer; ... } {code} While I see how the JMS spec is attempting to support a resilient protocol, this change seems to be a departure from the current programming model. Isn't this something that should be handled internally to the provider implementation itself? This seems to be exposing a bit of the lower-level message handling up to the application layer. At what point should the application layer need to deal with message formats? Wouldn't the application layer just create a message in whatever logical format it needs and then send it?	1	1	https://issues.apache.org/jira/browse/AMQNET-271
18517	AMQNET-263	I think you must be. I just successfully ran all of the unit tests on Mono 2.6.4 on Windows 7. It seems like there was the normal number of pass/fail results, but I would have to double check to be sure. However, the tests all did run. My final result numbers were: "Passed: 416, Failed: 26, Errors: 3, Time: 568.95" I do notice that there are some visual drawing anomolies of the NUnit interface when running under Mono, and if I click on the Text Output tab, an error dialog pops up. How are you running the tests, and what version of Mono are you using? I am also running the .NET 3.5 build of the tests. Update: I ran the tests again on .NET CLR, and the results are slightly different: "Passed: 418, Failed: 24, Errors: 3" There are two tests that fail on Mono, but pass on .NET. I'm not sure which they are or if that might be expected or not.	1	0	https://issues.apache.org/jira/browse/AMQNET-263

#	Requirement	Questions	Answered	Successfully Answered	URL
18518	AMQNET-262	After researching and thinking on this topic, I am beginning to question the use of the GAC with NMS. One of the primary strengths of the architecture of NMS with broker specific provider implementations is that you can do simple file copies to install the assemblies. This allows for side-by-side installation of different versions of NMS. What is the requirement to install NMS and the provider assemblies into the GAC? I can't see how this has any benefit, but I can see how it has several downsides. Perhaps explaining how installing NMS+Providers into the GAC is a "good thing" would be helpful. I'm beginning to think we are trying to find a solution to a problem that doesn't exist.	1	1	https://issues.apache.org/jira/browse/AMQNET-262
18519	AMQNET-245	Just double checked and the broker will set the "timestamp" field so adding that should be fine. You'd need to add in the read of the TimeStamp header in the StompWireFormat as well I'd think. Can you work up a complete patch with the changes needed and include a Unit Test to validate that the functionality is working as expected. That would help us out a lot since I won't have time to get to that this week.	1	1	https://issues.apache.org/jira/browse/AMQNET-245
18520	AMQNET-239	Hi, Timothy, At http://msdn.microsoft.com/en-us/library/bb738067.aspx?ppud=4#SSLwiththe.NETCompactFramework microsoft states: {quote}The .NET Compact Framework has built-in support for SSL.{quote} How does that fit together? Thanks, Markus	1	1	https://issues.apache.org/jira/browse/AMQNET-239
18521	AMQNET-239	Sorry, i was pretty unclear. As Markus assumed i was talking about the Implementation of SSL-support for the Compact Framework. Unfortunately I'm not able to share the patch I've used to accomplish this goal. Don't get me wrong, I'm not selfish and I don't want to slave-drive the Apache.NMS project - it's just licensing issues. In spite of all that i want to outline what we have done. - Build bouncycastle (use the latest source code version from CVS - the released version 1.6.1 doesn't build for target netcf-2.0) - Derive a SslStream class from the compact framework stream Stream class - SslStream is acting as a 'proxy' to the encrypted in-/out-streams provided by the bouncycastle library (be aware of the async calls!). Implement SslStream similar to the CLI Implementation of Stream (http://www.123aspx.com/rotor/RotorSrc.aspx?rot=42052). - Add the initialization of the SSL-Stream to the SslTransport class. Create a NetworkStream around your socket. Use the bouncycastle TlsProtocolHandler to encrypt the network stream. Use your implementation of SslStream to wrap the Input-/OutputStream of the TlsProtocolHandler. No offense.	1	0	https://issues.apache.org/jira/browse/AMQNET-239
18522	AMQNET-237	Can you add some more info to this, your broker config and a sample app showing how your use case would be good.	1	1	https://issues.apache.org/jira/browse/AMQNET-237
18523	AMQNET-236	Should this item be Resolved with the change made to throw the exception?	1	0	https://issues.apache.org/jira/browse/AMQNET-236
18524	AMQNET-226	If I'm understanding this correctly then it looks ok to me. The Destination allows you the option of overriding the Prefetch policy settings on the Connection which should be considered the defaults for all Consumers. If you want a specific destination to have some other prefetch value than you set the "?consumer.prefetchSize=5" option on the destination and it overrides the defaults. What was the behaviour you were expecting?	1	1	https://issues.apache.org/jira/browse/AMQNET-226
18525	AMQNET-224	Why don't you just create your own implementation of IRedeliveryPolicy and set that on the Connection object?	1	1	https://issues.apache.org/jira/browse/AMQNET-224
18526	AMQNET-206	I have a question concerning your unit tests: The session is in AcknoledgementMode "AutoAcknowledge": In that case the messages shoud automatically disappear from queue after receiption, or am I wrong ?! Anyway that's the behavior I'd expect?!	0	0	https://issues.apache.org/jira/browse/AMQNET-206
18527	AMQNET-206	After some research on the web, it seems that your implmentation is right. QueueBrowser does never remove messages from queue. So, I guess the best method to acknowledge messages is to store their JMSMessageId's while browsing and "consume" them with a MessageConsumer in a second round... Or does anybody knows a better way of achieving this (Because with that solution all messages to be acknowledged are downloaded twice...)?	0	0	https://issues.apache.org/jira/browse/AMQNET-206
18528	AMQNET-205	Tim, I found the following on MSDN: "The GZipStream class uses the gzip data format, which includes a cyclic redundancy check value for detecting data corruption. The gzip data format uses the same compression algorithm as the DeflateStream class." What do you think about using GZipStream instead of DeflateStream class? Reference: http://msdn.microsoft.com/en-us/library/system.io.compression.deflatestream.aspx	1	0	https://issues.apache.org/jira/browse/AMQNET-205
18529	AMQNET-201	That is because MSMQ is naturally one way, and that's because the way I understand it you use MSMQ, ActiveMQ, etc to apply load leveling functionality. Meaning, you off load the processing because it is generally too long for a client to wait for a response and you can load balance the processing as well as level it out as a response is not real time. Maybe I don't understand the customer driven motivation you mention. Why would they want to switch from request/reply done through WCF to use ActiveMQ? Why not just continue using the request/reply through WCF and increase the timeout values? If you are concerned about reliability of message transfer then you can enable reliability through the SOAP 1.2 endpoints. If you are concerned with ability to load level messaging, then why would you want to do request/reply with ActiveMQ over of point-to-point (one-way) or pub/sub? Yes ActiveMQ implements request/reply, but with .NET we have WCF to do it whereas not sure if other languages have such a technology.	0	0	https://issues.apache.org/jira/browse/AMQNET-201
18530	AMQNET-199	Can we document somewhere what this feature does? What does it do?	1	0	https://issues.apache.org/jira/browse/AMQNET-199
18531	AMQNET-198	I cannot use this workaround... the only open port is 61616... so I cannot access web console... I gave a try to the patch proposed in issue AMQNET-206... and it works! Could you consider to include it in release 1.2.0, please?	1	0	https://issues.apache.org/jira/browse/AMQNET-198
18532	AMQNET-159	Thanks Jim! Would you recommend just pulling from trunk to pick up this change? Any implications with picking up some of the completed items on 1.2 at all??? We're looking to go to production soon with the NMS provider and ActiveMQ.	1	0	https://issues.apache.org/jira/browse/AMQNET-159
18533	AMQNET-159	I reviewed the changes and found one difference. In FailoverTrasport.doConnect() if(_asyncConnect) { // now wait for transport to be populated, but timeout eventually allDone.WaitOne(_asyncTimeout, false); } you are passing false into an overloaded WaitOne method. I was not using this method. I wasn't passing in false just _asyncTimeout. Is there a specific reason this was done? If so, how come? If not, what are the implications of this? Also, would it make sense to make the async connection the default behavior?	1	0	https://issues.apache.org/jira/browse/AMQNET-159
18534	AMQNET-159	Okay. Thanks Timothy. Last question is whether to make this the default behavior then. Should it be opt in or opt out? Originally I thought opt in so as not to modify the default behavior of NMS ActiveMQ, but I don't see why anyone wouldn't want the provider to be faster connecting to the master broker when using the failover protocol.	1	0	https://issues.apache.org/jira/browse/AMQNET-159

#	Requirement	Questions	Answered	Successfully Answered	URL
18535	AMQNET-137	Regarding redelivery policy, currently client acknowledgement mode will not redeliver a message. There is no support of redeliveries. Would it be stretching NMS too far to allow to configure for a mode that "does" redeliver messages when the session is in client acknowledgement mode? If I understand correctly the redelivery logic is on the provider side and not the broker so perhaps this is possible. Straying too far from the java implementation wouldn't be a good thing though so would like everyone's opinion... Here is a discussion in the forum regarding the situation. http://www.nabble.com/Transactions-and-redelivery-td25915300.html We have coded a solution to redelivery optimistically or pessimistically with client acknowledgement mode. Optimistic will redeliver immediately before processing additional messages in the queue (like transactional acknowledgement mode). The message will not get ack'd until it exhausts all possible redeliveries. Pessimistic will simply clone the message, placing it at the end of the queue, and ack the current message. It will embed some metadata into the message and only redeliver (clone the message) up to the max redelivery count on the redelivery policy. Thoughts?	1	1	https://issues.apache.org/jira/browse/AMQNET-137
18536	AMQNET-129	It is possible to delete durable subscribers. Check the following API: ISession.DeleteDurableConsumer(). As far as deleting queues/topics, I will need to check on that. Just as background, what is the motivation for wanting to delete a queue? I also assume that you would want to delete a topic as well. If it is simply a matter of keeping things cleaned up, that I would ask if a temporary queue or temporary topic might better suit your implementation requirements?	1	1	https://issues.apache.org/jira/browse/AMQNET-129
18537	AMQNET-126	It appears that WCF does not support the request/reply model "The request-reply model is indeed inadequate for a queued transport (e.g. you can't block a thread while waiting for the response). You can correlate messages at the application level, typically using a durably stored correlation indicator (e.g. purchase order number). To process responses you would need a response queue and a service that would process messages sent to that queue." (http://social.msdn.microsoft.com/Forums/en-US/wcf/thread/2fb0aafd-078f-4638-bb99-372cbf0158fd/ for more details)	0	0	https://issues.apache.org/jira/browse/AMQNET-126
18538	AMQNET-105	If 1.0 isn't out yet, why should this deprecated method remain in 1.1? Why not schedule this fix for 1.0 with deprecation and then remove it for 1.1?	1	0	https://issues.apache.org/jira/browse/AMQNET-105
18539	AMQNET-100	Could you elaborate more on this request? What restrictions are you referring to? Could you show an example of the style you are looking for?	1	0	https://issues.apache.org/jira/browse/AMQNET-100
18540	AMQNET-97	This issue is in a fixed state, but should an enhancement issue be created to actually implement the QueueBrowser functionality in the ActiveMQ provider?	1	0	https://issues.apache.org/jira/browse/AMQNET-97
18541	AMQNET-93	Just to clarify, do you plan to mirror the current list of standard JMS exceptions in NMS? As for non-standard exceptions, perhaps wrap it in UnknownVendorException to the caller has an indication to do something special? Just for completeness the list of exception in EMS that don't map are AuthenticatoinExcetpion, CannotProceedExceptin, CommunicationException, ConfigurationException, InvalidNameException, ServiceUnavailableException.	0	0	https://issues.apache.org/jira/browse/AMQNET-93
18542	AMQNET-79	Awesome! What is the status of the contribution? Is it already part of Spring.NET? Feel free to post out to the Dev mailing list, or you can e-mail me directly if you like. I'll try and help out as I can.	1	1	https://issues.apache.org/jira/browse/AMQNET-79
18543	AMQNET-79	I have a (very!) basic implementation of WCF now : datagram only (i.e. IInputChannel/IOutputChannel), no request-response, no duplex, and sessionless. Which branch should I check it in to?	1	0	https://issues.apache.org/jira/browse/AMQNET-79
18544	AMQNET-42	I think this may be possible now that you have added support for STOMP 1.1. What do you think?	1	0	https://issues.apache.org/jira/browse/AMQNET-42
18545	AMQNET-27	I've added an ExceptionListener event on the IConnection which should serve a similar purpose to ExceptionListener in JMS. Does this seem an OK solution to you?	0	0	https://issues.apache.org/jira/browse/AMQNET-27
18546	AMQNET-26	replicate the functions in C#. Below are my comments as to the steps or actions required to fix this problem. If others can comment or correct my comments, then we can get started on this. If others are willing to help with the coding or testing, that would be great. All that I state below is what I have gleaned from the Java code. There is some interpretation of the code that may be in error. If so, then those with sufficient knowledge should correct my interpretation. A quick jump through the stack: We start with a connection factory derived from ActiveMQConnectionFactory() (C# ConnectionFactory()). This connection factory returns an IConnection which, as part of the process, creates an ITransport using a TransportFactory(). In Java, the resource information allows for various implementations of ITransport which includes FailoverTransport(). Different from C#, the Java implementation breaks out the logic for creating a transport so that it can be overridden by subclasses if desired. Not sure if we need that or not. In the TransportFactory, it all starts with the URI associated with the connection. The "failover:(URI1, URI2, ..., URIn)" syntax is passed to the TransportFactory.connect() to create the associated ITransport object. Looking at TransportFactory.connect() it uses findTransportFactory() to determine which class to load based upon the location passed. The URI.getScheme() method returns "failover" which is used to load the failover transport factory (org.apache.activemq.transport.failover.FailoverTransportFactory.java). This is based upon the META-INF information which redirects "failover" to that class. At this point we have a choice. We can allow for dynamic loading of classes (which can be done - not too hard) allowing for a similar run-time loading of appropriate classes to support various ITransport interface classes, or we can punt and do something easy for a compile-time loading of classes. My simple vote is to do a compile-time reference to the class which we can move to run-time loading at a future date. I have done run-time loading of classes, so if desired I can explain how I did it, but do we want that level of complexity at this point? Tabling that issue, the fundamental questions now reside around the differences between TransportFactory() and FailoverTransportFactory(), followed by the differences between Transport() and FailoverTransport(). The Java version of TransportFactory includes support for composite connections. This is used to support reliable and HA connections. Basically, composite connections represent a list of URIs associated with a connection. Most of the extended functionality of the transports is handled through subclassing. Want a ITransport that does X, then create a new subclass of Transport() or something that implements ITransport() and now you have a new feature X. However, there is a client-side and server-side for feature X (e.g. SSL). For SSL support, the client must understand how to connect properly, and the server needs to understand how to accept the connection properly. Thus some of the classes that support ITransport() also support TransportServer() which handles the server-side of things (I think this is correct). Not all features require this to work properly. For instance, the failover: capability is a client-only feature that does not require server-side coding. If we are to round-out all the neat features of the Java-based interfaces to be supported here with .NET, then we need to carefully support this dual concept of client- and server-side coding for the features. I think that is the level of complexity that is implied by all the features in the Java-based code. Essentially it breaks down into two sets of supported interfaces (Transport and TransportServer) which are required depending upon the specific feature. The rest of the feature code seems to be dependent upon the features being implemented. OK - so what does this mean for our implementation? I would suggest the following: 1. We create a Failover directory for our files with an implied "Failover" namespace extension. 2. Create a FailoverTransportFactory.cs and FailoverTransport.cs classes that will implement our new features (details to follow soon). 3. Explicitly reference this new namespace and based upon the schema passed either call TcpTransportFactory() or FailoverTransportFactory(). Using this ITransportFactory create the appropriate ITransport class. 4. As we add new features, we can then decide to continue with this	1	1	https://issues.apache.org/jira/browse/AMQNET-26

#	Requirement	Questions	Answered	Successfully Answered	URL
18547	AMQNET-26	Jim, Assuming I get somewhere with this how do I get the modified sources up to the system (or you)? I can do an SVN EXPORT and let you figure out where new files have been added, but would like to know what your preferred method is. I am still resolving some compile-time errors induced by my translation of the Java equivalent code to C#. If we had a sandbox SVN repository, I could check in my current snapshot for others to help. Jim, is something like that possible? If not, I hope to do some of my first tests later today. David, thanks for the offer. I would like your assistance if we can resolve a couple of these issues mentioned above. Assuming I get to the testing stage soon, it may be that your assistance with testing will be most valuable.	1	1	https://issues.apache.org/jira/browse/AMQNET-26
18548	AMQNET-26	I've just did an SVN UPDATE on my project tree and saw that there were many modifications performed by Jim in his preparation for a release. I was trying to merge his changes with mine but I cannot build the vendor portion (the NUnit libraries are missing). Jim, can you update me as to how to build this now? I can build the vendor .Net 2.0 library, but the test library refuses to build due to a lack of NUnit library. I presume if I get by this then the top-level project will build correctly. After a merge I can then generate a HEAD-level patch for us to use. David, please take a look at the patch and see if you can start seeing what I was doing. Or you may want to wait until I can get these changes updated to the HEAD version of the software.	1	1	https://issues.apache.org/jira/browse/AMQNET-26
18549	AMQNET-26	Clean checkout of HEAD does not build for me. I am missing the NUnit libraries, so all the TEST libraries cannot be built. Is there a simple way to build all without the TEST libraries being built? Or can we get the NUnit libraries included with the project? Or get instructions on what to do? Thanks, -Allan	1	1	https://issues.apache.org/jira/browse/AMQNET-26
18550	AMQNET-26	OK - I got it to compile. I now understand what Jim was doing in his reorganization of Apache.NMS. It used to be that the Apache.NMS portion was fully loaded under Apache.NMS.ActiveMQ as an external link. Now, only a fixed, deployed version of Apache.NMS is referenced by Apache.NMS.ActiveMQ. Because of that, and because that deployed version did not include .NET 2.0 CF, my builds were failing. Also, the nunit.framework.extensions.dll is needed by the Apache.NMS.ActiveMQ project. This required a change in the nant.build (patch attached). So, to get it to work for me I had to check out the Apache.NMS project (trunk) and build the the project (no problems there). I wanted to create the "deploy" directories but the only target that would do so is "deploy-all". This also does an SVN COMMIT which is not what I want to do. So I created a new target called "deploy-local" which is called by "deploy-all" so that I could do what was needed without the SVN COMMIT. Is there a better way to do this? After copying the files from Apache.NMS to Apache.NMS.ActiveMQ, I could now build the Apache.NMS.ActiveMQ project after making the nunit change referenced above (patch attached). From here, I shall attempt to add my previous patches for failover support. Since the project is in two pieces, I'll need to break up my previous patch into smaller pieces that can be applied to each separately.	1	1	https://issues.apache.org/jira/browse/AMQNET-26
18551	AMQNET-26	Since I had a hand in this code, would you mind posting the complete URL string that you are using? Also, when does the failure occur? During ConnectionFactory(url)? CreateConnection()? The changes in URISupport should remain (mostly) compatible with the previous version. But there was some code changed in the CreateConnection() call that dealt with URL processing. Having this additional information would go a long way to understanding and fixing this problem.	1	1	https://issues.apache.org/jira/browse/AMQNET-26
18552	APLO-288	Did you try: http://activemq.apache.org/nms/ which has a STOMP library Apache.NMS.Stomp.	0	0	https://issues.apache.org/jira/browse/APLO-288
18553	APLO-280	Hi. Does Apollo supports message groups with advanced settings like a TTL for a message group or sequence reset (Like Active MQ does, With the JMSXGroupSeq = -1 thing)? Thanks.	0	0	https://issues.apache.org/jira/browse/APLO-280
18554	APLO-267	I need to send messages between JMS and STOMP (JMS <-> TOPIC <-> STOMP (WebsocketClient)), but it is not working. I did try with ActiveMQ 5.8 and Apollo 1.6 with the current WebSocket-Client from Apollo 1.6. What i did find out: * Content-Length as described in http://activemq.apache.org/stomp.html !2013-09-04 12_07_39-Apache ActiveMQ-Stomp.png! ** but how do you remove the content-length header property? Parameter in the connection!? * Workaround for STOMP -> JMS in Ruby: http://www.juretta.com/log/2009/05/24/activemq-jms-stomp/ * Stackoverflow: http://stackoverflow.com/questions/13401374/activemq-to-apollo-transition-openwire-to-stomp-protocol-configuration * Problem is reproduceable, Java class and Websocket-Client is attached to this ticket and they are very simple. Should i workaround and try to remove the header property content-length from JMS side (TextMessage) and STOMP to be able to communicate?	0	0	https://issues.apache.org/jira/browse/APLO-267
18555	APLO-262	Patch to include: [addition] ProtocolFactoryTest [modified] OpenwireProtocolHandler [modified] OpenwireParellelTest [addition] openwire-features.md The openwire-features file can be viewed by running jetty:run in the apollo-website module and navigating to http://localhost:8080/documentation/openwire-features.html Let me know how we should include it in the main docs: Should we split out the STOMP docs into its own file? and link to both stomp and openwire from the main doc file (user-manual.md)	1	1	https://issues.apache.org/jira/browse/APLO-262
18556	APLO-257	Lionel, Is this something that happens regularly? Any other hints (machine specs, etc) for how to create this? I've been trying and cannot recreate it yet (using the stomp benchmark from APLO-176)	1	1	https://issues.apache.org/jira/browse/APLO-257

#	Requirement	Questions	Answered	Successfully Answered	URL
18557	APLO-257	We ran the same benchmark scenario on the same hardware with the latest snapshot (apollo-99-trunk-20120929.031419-121) with two changes: - more memory: -server -Xmx16384m -Xms4096m -XX:-UseBiasedLocking - using the default slow_consumer_policy="block" We still see memory exceptions, however, these are "GC overhead limit exceeded". Here are some stack traces: 2012-10-01 18:49:51,934 | WARN | stackref=13a1c2bf7a8 java.lang.OutOfMemoryError: GC overhead limit exceeded at org.fusesource.hawtdispatch.internal.util.TimerHeap.executeReadyTimers(TimerHeap.java:79) at org.fusesource.hawtdispatch.internal.TimerThread.run(TimerThread.java:144) 2012-10-01 18:50:14,710 | WARN | stackref=13a1c2bf7a9 java.lang.OutOfMemoryError: GC overhead limit exceeded at java.util.concurrent.ThreadPoolExecutor.addThread(ThreadPoolExecutor.java:671) at java.util.concurrent.ThreadPoolExecutor.addIfUnderMaximumPoolSize(ThreadPoolExecutor.java:721) at java.util.concurrent.ThreadPoolExecutor.execute(ThreadPoolExecutor.java:657) at org.fusesource.hawtdispatch.package$RichExecutor.execute(hawtdispatch.scala:171) at org.fusesource.hawtdispatch.package$RichExecutorTrait$class.apply(hawtdispatch.scala:68) at org.fusesource.hawtdispatch.package$RichExecutor.apply(hawtdispatch.scala:169) at org.apache.activemq.apollo.broker.security.JaasAuthenticator.authenticate(JaasAuthenticator.scala:66) at org.apache.activemq.apollo.stomp.StompProtocolHandler$$anonfun$on_stomp_connect$3$$anonfun$apply$mcVsp1.applymcVsp(StompProtocolHandler.scala:1107) at org.fusesource.hawtdispatch.package$$anon$3.run(hawtdispatch.scala:357) at org.fusesource.hawtdispatch.internal.SerialDispatchQueue.run(SerialDispatchQueue.java:96) at org.fusesource.hawtdispatch.internal.pool.SimpleThread.run(SimpleThread.java:77) 2012-10-01 18:50:48,827 | WARN | stackref=13a1c2bf7aa java.lang.OutOfMemoryError: GC overhead limit exceeded at java.util.Arrays.copyOfRange(Arrays.java:3065) at org.fusesource.hawtdispatch.transport.AbstractProtocolCodec.read(AbstractProtocolCodec.java:307) at org.fusesource.hawtdispatch.transport.TcpTransport.drainInbound(TcpTransport.java:657) at org.fusesource.hawtdispatch.transport.TcpTransport$6.run(TcpTransport.java:538) at org.fusesource.hawtdispatch.internal.NioDispatchSource$3.run(NioDispatchSource.java:226) at org.fusesource.hawtdispatch.internal.SerialDispatchQueue.run(SerialDispatchQueue.java:96) at org.fusesource.hawtdispatch.internal.pool.SimpleThread.run(SimpleThread.java:77) 2012-10-01 18:51:08,070 | WARN | stackref=13a1c2bf7ab java.lang.OutOfMemoryError: GC overhead limit exceeded 2012-10-01 18:51:38,771 | WARN | stackref=13a1c2bf7ac java.lang.OutOfMemoryError: GC overhead limit exceeded It would be good if Apollo could run happily in its 16GB without causing memory related exceptions...	1	1	https://issues.apache.org/jira/browse/APLO-257
18558	APLO-256	Hi Ricardo, Could you provide details about the paths you've installed Apollo to and the command that your trying to execute that fails?	0	0	https://issues.apache.org/jira/browse/APLO-256
18559	APLO-244	I need to get better hardware. What kind are you running on?	1	1	https://issues.apache.org/jira/browse/APLO-244
18560	APLO-244	Are you running the consumer on a different box or are you still running it on /w the producers?	1	1	https://issues.apache.org/jira/browse/APLO-244
18561	APLO-244	Hi Lionel, Could you attache the ActiveMQ config that your using for the test run? Would like to also test it for comparison as you have done.	1	1	https://issues.apache.org/jira/browse/APLO-244
18562	APLO-196	Hi Pablo, How about we define a new more powerful Filtering interface. How does this look to you? /** * A Protocol filter can filter frames being sent/received to and from a client. It can modify * the frame or even drop it. */ abstract class ProtocolFilter2 { /** * Filters a command frame received from a client. * returns None if the filter wants to drop the frame. */ def filter_inbound[T](frame: T):Option[T] /** * Filters a command frame being sent client. * returns None if the filter wants to drop the frame. */ def filter_outbound[T](frame: T):Option[T] }	1	1	https://issues.apache.org/jira/browse/APLO-196

#	Requirement	Questions	Answered	Successfully Answered	URL
18563	APLO-177	Here it is... ==> apollo.log <== 2012-08-27 08:30:05,161 | DEBUG | created queue: test4udp | 2012-08-27 08:30:05,163 | DEBUG | producer route attaching to consumer. | 2012-08-27 08:30:05,168 | WARN | java.lang.NullPointerException | 13966c41ce1 ==> stacktrace.log <== 2012-08-27 08:30:05,169 | WARN | stackref=13966c41ce1 java.lang.NullPointerException at org.apache.activemq.apollo.stomp.StompUdpProtocol$$anon$1$$anonfunorgapache$activemq$apollo$stomp$StompUdpProtocol$$anon$$updated_headers$2.apply(StompUdpProtocol.scala:205) at org.apache.activemq.apollo.stomp.StompUdpProtocol$$anon$1$$anonfunorgapache$activemq$apollo$stomp$StompUdpProtocol$$anon$$updated_headers$2.apply(StompUdpProtocol.scala:204) at scala.collection.Iterator$class.foreach(Iterator.scala:660) at scala.collection.JavaConversions$JIteratorWrapper.foreach(JavaConversions.scala:573) at scala.collection.IterableLike$class.foreach(IterableLike.scala:73) at scala.collection.JavaConversions$JListWrapper.foreach(JavaConversions.scala:615) at org.apache.activemq.apollo.stomp.StompUdpProtocol$$anon$1.org$apache$activemq$apollo$stomp$StompUdpProtocol$$anon$$updated_headers(StompUdpProtocol.scala:204) at org.apache.activemq.apollo.stomp.StompUdpProtocol$$anon$1$$anon$2.delivery(StompUdpProtocol.scala:151) at org.apache.activemq.apollo.broker.protocol.UdpProtocolHandler$UdpProducerRoute$$anon$2$$anonfun$$init$$1.apply(UdpProtocol.scala:254) at org.apache.activemq.apollo.broker.protocol.UdpProtocolHandler$UdpProducerRoute$$anon$2$$anonfun$$init$$1.apply(UdpProtocol.scala:254) at org.apache.activemq.apollo.broker.Sink$$anon$1.passing(Sink.scala:56) at org.apache.activemq.apollo.broker.SinkMapper$class.offer(Sink.scala:94) at org.apache.activemq.apollo.broker.Sink$$anon$1.offer(Sink.scala:55) at org.apache.activemq.apollo.broker.OverflowSink.drain(Sink.scala:159) at org.apache.activemq.apollo.broker.OverflowSink$$anonfun$1.apply$mcV$sp(Sink.scala:135) at org.fusesource.hawtdispatch.package$$anon$3.run(hawtdispatch.scala:357) at org.apache.activemq.apollo.broker.MutableSink$$anonfun$downstream_$eq$2.apply(Sink.scala:205) at org.apache.activemq.apollo.broker.MutableSink$$anonfun$downstream_$eq$2.apply(Sink.scala:203) at scala.Option.foreach(Option.scala:197) at org.apache.activemq.apollo.broker.MutableSink.downstream_$eq(Sink.scala:203) at org.apache.activemq.apollo.broker.protocol.UdpProtocolHandler$UdpProducerRoute.on_connected(UdpProtocol.scala:261) at org.apache.activemq.apollo.broker.DeliveryProducerRoute$$anonfun$connected$1.applymcVsp(Router.scala:222) at org.fusesource.hawtdispatch.package$$anon$3.run(hawtdispatch.scala:357) at org.fusesource.hawtdispatch.internal.SerialDispatchQueue.run(SerialDispatchQueue.java:96) at org.fusesource.hawtdispatch.internal.pool.SimpleThread.run(SimpleThread.java:77)	1	1	https://issues.apache.org/jira/browse/APLO-177
18564	APLO-177	The NPE was getting generated due to the 'add_user_headers' option. This should now be fixed in build 289 [1]. What did you expect the user header to get set to? You would need to have an authenticated user. Luckily, you can add a login/passcode header to the message to authenticate each UDP message. [1]: https://builds.apache.org/job/ActiveMQ-Apollo-Deploy/289/console	1	1	https://issues.apache.org/jira/browse/APLO-177
18565	APLO-177	Are you using a default login.config ??	1	1	https://issues.apache.org/jira/browse/APLO-177
18566	APLO-176	No, those are NOT really harmless. The in memory queue swaps blocks of messages out to disk but keeps track of the number of messages and bytes in that are contained in the block. At a later time when the a message from the block is needed, the block is loaded back up from the store and we compare the number of messages and bytes that we were able to load from the store to the expected size of the block. If there is a difference the queue adjust it's internal counters to account to the change. Why, does it do this? This was put in place so that a queue could tolerate if a 3rd party changed the message store directly (For example purged a store of all the messages in a queue). Is this normal in your case? No. I would NOT expect the LevelDB store and BDB stores to be inconsistent since they don't support external management. Now that hard part.. Whats the easiest way I can reproduce this? What kind of messaging scenarios is this running? Do all your brokers do this?	1	1	https://issues.apache.org/jira/browse/APLO-176
18567	APLO-176	Hi Lionel, Have you guys seen this anymore?	1	1	https://issues.apache.org/jira/browse/APLO-176
18568	APLO-176	Hi Richard, What version are you using?	1	1	https://issues.apache.org/jira/browse/APLO-176
18569	APLO-169	If the keystore is configured with two or more certificates, how does the broker pick which to use as the broker's cert?	1	0	https://issues.apache.org/jira/browse/APLO-169
18570	APLO-166	Once the OpenWire protocol if fully implemented and documented :) When will that be? Since Apache is a volunteer organization, it's hard to say when it will be completed. I can just tell you what the current state its. I would classify the current state of the OpenWire Protocol as being alpha/beta quality. It does not net support XA transactions and not all use cases are fully unit tested. Since the initial focus will be to to just make it feature and backward compatible, we probably will not be focusing on optimizing the OpenWire performance until thats done. ActiveMQ development will NOT stop. It has a large user based and we will continue to incrementally improve that code base. Apollo is jus providing a fresh new platform to do some things better.	1	1	https://issues.apache.org/jira/browse/APLO-166
18571	APLO-158	STOMP over websockets has been supported since the 1.1 release. Most modern browsers support WebSockets, should we even take the time to implement AJAX interface?	1	0	https://issues.apache.org/jira/browse/APLO-158
18572	APLO-131	Hi Lionel, I just tried manually testing durable subs against the latest build and i did not see any errors. Have you tried deleting the data directory (there have been some schema changes recently)? Is it reproducible? How?	1	1	https://issues.apache.org/jira/browse/APLO-131
18573	APLO-93	Here's a patch, the current openwire tests still work. However I notice there's an occasional NPE: Exception in thread "hawtdispatch-DEFAULT-1" java.lang.NullPointerException at org.apache.activemq.apollo.broker.protocol.HeartBeatMonitor.schedual_check_writes(HeartBeatMonitor.scala:43) at org.apache.activemq.apollo.broker.protocol.HeartBeatMonitor$$anonfun$start$1.applymcVsp(HeartBeatMonitor.scala:71) at org.fusesource.hawtdispatch.package$$anon$4.run(hawtdispatch.scala:347) at org.fusesource.hawtdispatch.internal.TimerThread$2.run(TimerThread.java:153) at org.fusesource.hawtdispatch.internal.SerialDispatchQueue.run(SerialDispatchQueue.java:83) at org.fusesource.hawtdispatch.internal.pool.SimpleThread.run(SimpleThread.java:61) Think it would be safe to break out of the schedual_check_read/schedual_check_write loop in this case? I suspect this is just because the connection is torn down before the monitor has checked the connection.	1	1	https://issues.apache.org/jira/browse/APLO-93

#	Requirement	Questions	Answered	Successfully Answered	URL
18574	APLO-28	the page has been started.. http://activemq.apache.org/apollo/versions/1.0-SNAPSHOT/website/documentation/extending-guide.html We should also document how to add custom stores, protocols, and transports.	1	0	https://issues.apache.org/jira/browse/APLO-28
18575	APLO-6	Hey Sam, Where do you think would be the best place to configure that option? A a broker level? Per connector? Per client?	1	1	https://issues.apache.org/jira/browse/APLO-6
18576	APLO-6	Hi Hiram, thanks for the fix. As I'm not familiar with JVM, can you give an example of how this option is set please?	1	0	https://issues.apache.org/jira/browse/APLO-6
18577	ARTEMIS-1185	Github user franz1981 commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1295#discussion_r119025093 --- Diff: artemis-cli/src/main/java/org/apache/activemq/artemis/cli/commands/tools/journal/StatJournal.java --- @@ -0,0 +1,288 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.activemq.artemis.cli.commands.tools.journal; +import java.io.File; +import java.io.FileOutputStream; +import java.io.PrintStream; +import java.util.Arrays; +import java.util.Comparator; +import java.util.concurrent.TimeUnit; +import java.util.concurrent.atomic.AtomicLong; +import java.util.concurrent.locks.LockSupport; +import io.airlift.airline.Command; +import io.airlift.airline.Option; +import org.apache.activemq.artemis.cli.commands.ActionContext; +import org.apache.activemq.artemis.cli.commands.Configurable; +import org.apache.activemq.artemis.core.config.impl.FileConfiguration; +import org.apache.activemq.artemis.core.io.buffer.counters.FlushSampleFlyweight; +import org.apache.activemq.artemis.core.io.buffer.counters.Profiler; +import org.apache.activemq.artemis.utils.collections.MulticastSample; +@Command(name = "stat", description = "Provide performance statistics around the journal usage") --- End diff -- It is a weird name...Do you have something in mind? It is a journal sampling profiler and the ideas on how it works are taken from: - https://linux.die.net/man/1/perf-record - https://github.com/jvm-profiling-tools/honest-profiler But in the future I'm planning to provide a single CLI command that could collect perf data using a filter of some kind or different CLI commands for macro profiling areas ...it is something I want to discuss on the dev list...	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1185
18578	ARTEMIS-1185	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1295 what I am missing: - doc how to use this... (is that internal only)? - how to read the file... I am a bit confused with the file format... Also, it would be nice to be allowed to configure where the file would be generated.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1185
18579	ARTEMIS-1185	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1295 @franz1981 I was just looking from the POV of the code.. TimedBuffer->profiler (whatever the name is) is instantiating an output file. shouldn't it be configurable for the running server? the CLI to read it could also get the file from the configuration if you're running from the instance.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1185
18580	ARTEMIS-1185	Github user franz1981 commented on the issue: https://github.com/apache/activemq-artemis/pull/1295 @clebertsuconic It is something I think can over-complicate the environment setting for a user and could potentially makes the mechanism less efficient too. That's why: - the file is not a "real" file on Linux (it uses [/dev/shm](https://www.cyberciti.biz/tips/what-is-devshm-and-its-practical-usage.html)) to avoid potential costs on page swapping/writeback of the dirty region, hence the position/name could not be bound to a fixed configured position - the file need to be zeroed on broker start and to make this behaviour safer&simpler to be implemented I prefer to let the file system do it for me, creating a new file with `pid-date of creation` info on the name (+ File::deleteOnExit), hence a fixed file name is not an option here too wdyt?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1185
18581	ARTEMIS-1185	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1295 @franz1981 the only reason I didn't merge this is that this would be a hidden feature... no docs.. can you add a doc chapter about this?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1185
18582	ARTEMIS-1185	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1295 @franz1981 you are misunderstanding me.. the CLI won't ever run concurrently with the server.. it's an issue on your CLI.. you are holding the lock on the server.. hence you're not able to capture it.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1185
18583	ARTEMIS-1185	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1295 @franz1981 I tried.. cannot run .. because of the lock... lets just talk over tomorrow..	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1185
18584	ARTEMIS-1185	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1295 I think this could be made a bit more generic.. So I have moved the file creationg to MultiCastBuffer. This could be extended to other measures.. E.G. Message sending and consuming... etc.. etc... So, right now the logic for missing measures is inside FlushSampleFlyweight.. while it should / should be at the buffer... I still think we should name this PipedBuffer.. you could create many pipes. .one for the journal, one for messages.. .etc.. etc... it's a pipe of information.. where you can have multiple receivers... if you create a chain of information (I'm calling pipe.. we can find other names) for messages.. it could be there... Or maybe we could make this generic for not just the journal.. .and still have a single File.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1185
18585	ARTEMIS-1179	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/1288 Clear the caches when the session is closed?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1179
18586	ARTEMIS-1179	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1288 When a session is closed it's thrown away right, and you make a new session. Would that not be redundant?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1179
18587	ARTEMIS-1179	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1288 oh yeay sevntu not available again.... [ERROR] Unresolveable build extension: Plugin org.apache.felix:maven-bundle-plugin:3.3.0 or one of its dependencies could not be resolved: Could not find artifact org.apache.felix:maven-bundle-plugin:jar:3.3.0 in sevntu-maven (http://sevntu-checkstyle.github.io/sevntu.checkstyle/maven2) @ On a serious note, it seem this tool in the repo, isn't reliable, this isn't the first time I've seen this issue or others have. How much real value is it adding?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1179
18588	ARTEMIS-1174	That's an easy change.. doesn't even need a JIRA... why don't you just send a PR? or it needs some extra work?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1174

#	Requirement	Questions	Answered	Successfully Answered	URL
18589	ARTEMIS-1163	Github user klausboeing commented on the issue: https://github.com/apache/activemq-artemis/pull/1269 @jbertram I made the adjustments as needed. How will the ARTEMIS-1163 ticket be processed? Thanks for listening.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1163
18590	ARTEMIS-1162	Github user franz1981 commented on the issue: https://github.com/apache/activemq-artemis/pull/1268 @michaelandrepearce Nope, no problem for me...but I'm not 100% sure is the best way to do it: changing the new one to be like the original one is a matter of 1 line change and a boolean flag to be put in the right spot. Instead of having: ``` if (checkpoint || currentPendingSyncs >= estimatedOptimalBatch) { ``` Need to be changed into: ``` if (checkpoint || (smartCoalescing && currentPendingSyncs >= estimatedOptimalBatch)) { ``` Finished. Now you have one implementation that cover the 2 behaviours. wdyt?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1162
18591	ARTEMIS-1157	Can you work up a test (perhaps based on one of the examples shipped with Artemis) which demonstrates the issue?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1157
18592	ARTEMIS-1157	Based on the examples is a bit difficult for us. ArtemisMqCrashDemoClient.java is the client code. Client is consuming the keystore from: src/test/resources/artemis-key.jks. Now, if the client makes a second connection from the same factory, he runs in the stacktrace provided as comment in the file. broker.xml is the server. We are using a shared-store and inside the <cluster-connections> we reference one live- and backup connector. The configured keystorepath defined on <connector name="netty-live-connector"> is getting propagated to the client but of course, at the path: /etc/pki/tls/cert/node1.client.jks is no keystore on the client-side. We know, that for this client code provided in ArtemisMqCrashDemoClient.java we could resolve it by using: -Dorg.apache.activemq.ssl.keyStore=src/test/resources/artemis-key.jks but what should we do if the same JVM needs to connect to a different broker with a different keystore?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1157
18593	ARTEMIS-1157	bq. Imagine a broker sets in the cluster-connector - lets say one or two cipher suites with the property: enabledCipherSuites=TLS_AES_WITH... Now if the clients JRE - for whatever reason - does not support that cipher suite (JRE to old, disabled in JRE security settings, JRE from vendor that does not support it at all...) In that case, the initial connection from the client will work, but as soon as a failover occurs, client connection fails. Couple of things: - What's the use-case where you'd want to set {{enabledCipherSuites}} on a connector and _not_ set it on the corresponding acceptor? - What's the use-case where different cluster nodes would support different cipher suites. I understand both of these use-cases are technically possible, but at this point they don't make sense to me so I'm not inclined to consider them worth supporting. bq. Another thing is: clientA -> HA-brokerA -> Bridge -> HA-brokerB -> clientB. This is a simple setup but consists out of 4 independent components and each component is high-available so in the most basic case this will result in 8 Boxes (VM's, containers, whatever...). I'm confused about this use-case. I assume the "4 independent components" here are clientA, HA-brokerA, HA-brokerB, & clientB. I don't understand how that results in 8 boxes since typically only the brokers would consist of a live/backup pair. Also, is the "Bridge" a cluster bridge (which is automatically configured by the internal cluster manager) or a bridge defined in broker.xml explicitly? bq. Configuring all this systems in a consistent way in large organisations can be challenging. Because if just one component is configured the wrong way, initially everything works but as soon as a failover occurs, the system stops working. As long as the brokers are configured consistently (which is the recommendation), for example using the same value for enabledCipherSuites on their acceptors, then any client that can't connect to the backup will also fail to connect to the live. bq. I just want to highlight here, that the current concept can be "complicated/intransparent"... The SSL use-case here certainly isn't the most elegant, but given the broker's clustering design I can't think of anything better at this point (and I'm not convinced I need to at this point). bq. This will also require a lot of testing every time you change something. I think this is a bit of an overstatement. As stated before, the recommendation is for all cluster nodes to be configured equivalently. If that recommendation is followed then I think you can avoid most of the issues you've identified here. Transparent/automatic functionality like fail-over and connection load-balancing assumes equivalent configuration between nodes. I believe this is the case for most software that provides similar functionality.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1157
18594	ARTEMIS-1157	> I understand both of these use-cases are technically possible, but at this point they don't make sense to me so I'm not inclined to consider them worth supporting. You are right, it actually doesn't make sense to me, if I now read it again, I would have the same question: why should you do this. > As long as the brokers are configured consistently (which is the recommendation), for example using the same value for enabledCipherSuites on their acceptors, then any client that can't connect to the backup will also fail to connect to the live. Also here. Maybe we do something, that isn't best practice. To sum this up: At the end of the day, the client needs the keystore (the file) present on the client machine to connect (to live or backup). This file needs to be at the same location where the HA-broker is getting it from (except system property overwrite). If the password for the keystore should not be transmitted over the network, we overwrite it, again, with system properties on both ends and do not specify it in the connector definition (client and broker, that should work right?). We were somehow fixed to the idea, that client and broker decide on it's own, where to find it's own keystore/password combo. Thanks for your patience ;-) Best, Philipp	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1157
18595	ARTEMIS-1156	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1263 @franz1981 just updated to add in the journal compactor. I ran out of time today, to deploy and run this with a profiler on our servers to verify that this hasn't pushed the auto boxing further down or else where, i will do this tomorrow. In particular can you look at the way i have handled addAll for the NonBlockingHashSetLong, is there a better way? Also for the HashMap that had long keys in the journal which is not concurrent, I've used the LongHashMap from netty.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1156
18596	ARTEMIS-1156	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1263#discussion_r115794851 --- Diff: artemis-journal/src/main/java/org/jctools/maps/NonBlockingHashSetLong.java --- @@ -0,0 +1,105 @@ +/* + * Licensed under the Apache License, Version 2.0 (the "License"); + * you may not use this file except in compliance with the License. + * You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICLongNSLong-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTILongS OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.jctools.maps; + +import java.io.Serializable; +import java.util.AbstractSet; +import java.util.Iterator; + +/** + * A simple wrapper around {@link NonBlockingHashMapLong} making it implement the + * {@link java.util.Set} interface. All operations are Non-Blocking and multi-thread safe. + */ +public class NonBlockingHashSetLong extends AbstractSet<Long> implements Serializable { --- End diff -- Isn't this class available on any jctools package? Can't we just reuse the library instead of copy & paste?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1156

#	Requirement	Questions	Answered	Successfully Answered	URL
18597	ARTEMIS-1156	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1263 @michaelandrepearce Oh.. never mind.. I can ammend during checkout.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1156
18598	ARTEMIS-1156	Github user franz1981 commented on the issue: https://github.com/apache/activemq-artemis/pull/1263 @michaelandrepearce Got it, there is another effective solution to this: build an interner, similar to a StringInterner, but with a different hashed logic in order to avoid the JVM duplicating autoboxed instances each time. Few lines of code to explain it better: ``` //it will retain in LRU fashion the last 64 Long instances LongInterner longCache = new LongInterner(64); long id; //instead of: put(id , connection); //interned instance Long cachedId = longCache.intern(id); put(cachedId, connection); ``` wdyt? The caching/escape analysis are different things AFAIK: - caching is performed by the library on long->Long, and put by the compiler when autoboxing (eg [Long::valueOf](http://grepcode.com/file/repository.grepcode.com/java/root/jdk/openjdk/8u40-b25/java/lang/Long.java#837)) - [escape analysis](http://docs.oracle.com/javase/7/docs/technotes/guides/vm/performance-enhancements-7.html#escapeAnalysis): newly created Long instances no longer exists on the heap but being unpacked on the stack if used temporary, but need to call by hand new Long(long), bypassing the cached logic and do not retain it (seems not our case anyway) eg http://www.rationaljava.com/2015/12/how-long-does-it-take-jvm-to-effect.html	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1156
18599	ARTEMIS-1156	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1263 So we have: 1. primitive collections 1. id's as objects (Long in this case) 1. flywheel/interning caching pattern For 1: we have tried, looks good as long as we just do hot paths, can very quickly over complicate our code, or start bringing negative effects. For 2: I haven't had a chance to see how much work, but this maybe simpler/cleaner imho. Though does lend its self to people in future accidentally introducing the boxing issue again. For 3: Should work nicely, but only upto a predefined cache size. e.g. what occurs if cache size is N and i end up with N+1000 longs, we end up in same position as with Long.valueOf. Also some complexities introduced, how big should the cache be? how does someone tune it, i run a very large vm, but fred runs a small one, this works for an app that you deploy just in your own DC, but I'm weary of setting a magic cache size where it gets deployed in many differing env's. For this reason, id personally favour looking into option 2, then keeping a low scoped option 1.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1156
18600	ARTEMIS-1156	Github user franz1981 commented on the issue: https://github.com/apache/activemq-artemis/pull/1263 Have you taken a look to http://fastutil.di.unimi.it ?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1156
18601	ARTEMIS-1156	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1263 @michaelandrepearce merged... did some tweaks to preserve history on TypedProperties.. all merged now... Awesome work on this!	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1156
18602	ARTEMIS-1151	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1256#discussion_r115269492 --- Diff: tests/unit-tests/src/test/java/org/apache/activemq/artemis/tests/unit/core/journal/impl/TimedBufferTest.java --- @@ -205,162 +203,4 @@ public int getRemainingBytes() { } } - - /** - * This test will verify if the system will switch to spin case the system can't perform sleeps timely - * due to proper kernel installations - * - * @throws Exception - */ - @Test --- End diff -- why did you remove this test? it was important when I added it. I would rather keep a method public and a test than remove a test to make a method private.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1151
18603	ARTEMIS-1142	Just add a method in DivertConfiguration setTransformer(Transformer aTransformer); in addition or current setTransformerClassname(String s); Is it clear?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1142
18604	ARTEMIS-1142	Makes sense. although the transformer is instantiated at ActiveMQServerImpl, right? Take a look at the newly added Plugin.. maybe that would also help you?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1142
18605	ARTEMIS-1142	I will tale à look on ActiveMQServerImpl. But I don't understand your sentence "Take a look at the newly added Plugin". What plugin are you talking? If you have some links?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1142
18606	ARTEMIS-1141	or a discuss thread? how this can be used?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1141
18607	ARTEMIS-1141	I ran the built with tests with no problems. Did I miss something ? I missed the ArtemisFeatureTest, but the feature could not be installed, so the test is definitely not working very well to detect failures. To the "how can that be used", the build ensures that you can't break the feature installation. Try removing an important bundle from a feature and you'll see the build will break. If a new package dependency is added which cause a new import to be added to one of the bundles, the build will also fail fast.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1141
18608	ARTEMIS-1141	[~gnt] thank you so much for this.. it would have saved me a day of work last week if this was already implemented.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1141
18609	ARTEMIS-1131	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1232#discussion_r113757529 --- Diff: docs/user-manual/en/perf-tuning.md --- @@ -206,6 +206,14 @@ tunings won't apply to JDKs from other providers (e.g. IBM or JRockit) of your queues and the size and number of your messages. Use the JVM arguments `-Xms` and `-Xmx` to set server available RAM. We recommend setting them to the same high value. + + When under periods of high load, it is likely that Artemis will be generating + and destroying lots of objects. This can result in a build up of stale objects. + To reduce the chance of running out of memory and causing a full GC + (which may introduce pauses and unintentional behaviour), it is recommended that the + max heap size (`-Xmx`) for the JVM is set at least to 5 x the `global-max-size` of the broker. --- End diff -- Can you change the calculation for global-max-size then to be 1/5 of the -Xmx? Look on ConfigurationImpl	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1131
18610	ARTEMIS-1129	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1272 Re the issue. We are able to run against remote broker fine. are you running a remote client or invm client? The invm parts aren't shaded as if you're running invm you would have the core server classes anyhow. , which is why as normal dependency these aren't there. Also what version is broker?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1129
18611	ARTEMIS-1129	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1272 Is there any way to add tests ?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1129
18612	ARTEMIS-1129	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/1272 Maybe use it in the examples?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1129
18613	ARTEMIS-1129	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1272 @jbertram all of them?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1129

#	Requirement	Questions	Answered	Successfully Answered	URL
18614	ARTEMIS-1129	Github user gemmellr commented on the issue: https://github.com/apache/activemq-artemis/pull/1272 Would an alternative be just offering a separate client binary convenient download, rather than just the single all-inclusive archive? Or documenting a maven command to gather all the current jars needed for a given client release? Those wouldn't really be any harder to actually use than the single jar in the end, but would avoid needing new uber jars that could end up being a little fragile given all the shading config used to create them and going much less tested than the main modules. Setting that aside, I notice there are various profiles specifying modules in the pom, and I wondered if it is sufficient to only add these new ones in the one place as done, or if they also need added to some of the other profiles? Not sure exactly what they all do.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1129
18615	ARTEMIS-1129	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1278 @michaelandrepearce Shade is generating a few files at the sub-modules directories (see the changes). this was failing: ``` mvn -Ptests test ``` I only ran the examples before.. and the build is now broken without this. do you think these changes are ok? I'm concerned also that the release would generate these files on the source profile. Something to be checked also.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1129
18616	ARTEMIS-1129	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1278 @clebertsuconic def agree that it should be git ignore , as it is generated at build time. are you saying by adding this exclusion in rat list, the tests break, or without adding it it breaks? i am just running the build with the -Ptests profile now, waiting for it to complete locally. I know there is the setting dependencyReducedPomLocation which allows you to move it to another directory, problem is this breaks any remaining build process and marked buggy in the shade plugin docs so i don't think we should look at that.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1129
18617	ARTEMIS-1129	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1278 @clebertsuconic I'm happy with this change. I have a few more tidy up bits, been working on, as have been exploding the jars to check whats copied in and some bits like notice and license and a load of crud from jgroups, whilst still waiting the -Ptest to complete... shall i raise a pr to your branch or direct to master?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1129
18618	ARTEMIS-1124	Github user mtaylor commented on the issue: https://github.com/apache/activemq-artemis/pull/1217 @franz1981 two comments here. 1. The network time out should also be used on the page store. 2. The default timeout is 20seconds but the docs says 30s. I'll merge this as it is as the doc is a minor fix and page can be added later. Can you follow up with a patch for this. Thanks	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1124
18619	ARTEMIS-1116	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/1206 Couple of things: - The commit message should follow the format outlined in the [Hacking Guide](https://github.com/apache/activemq-artemis/blob/master/docs/hacking-guide/en/maintainers.md#commit-messages). - You have a fair number of checkstyle violations. See more [here](https://builds.apache.org/blue/organizations/jenkins/ActiveMQ-Artemis-PR-Build/detail/ActiveMQ-Artemis-PR-Build/2607/pipeline). You run the checkstyle stuff locally by using 'mvn -Pdev install'. - I wonder if this feature would be better if it was more general so that it could be used with any login module rather than isolated only to the LDAP login module. What do you think?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1116
18620	ARTEMIS-1116	Github user sjhiggs commented on the issue: https://github.com/apache/activemq-artemis/pull/1206 That's fine, I was assuming "from" was the external/LDAP role and "to" is the internal AMQ role, do you see it the opposite?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1116
18621	ARTEMIS-1116	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/1206 > I was assuming "from" was the external/LDAP role and "to" is the internal AMQ role, do you see it the opposite? I see the "from" was the external role and "to" is the internal Artemis role as you said. > ...the from will probably need a different delimiter because the role name may be something like a DN. Whatever makes sense from an implementation perspective is fine with me. I was just going on the fact that the original commit used commas as a delimiter.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1116
18622	ARTEMIS-1110	This looks done. Can you confirm?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1110
18623	ARTEMIS-1096	[~jbertram], can global-max-size have a computed default (or at least a safe default), so that an explicit -1 remains a legal way to disable the limits?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1096
18624	ARTEMIS-1096	[~jbertram], one more question. Regarding your first paragraph - I was under the impression that max-size-bytes=-1 at the address level would defer to the global limit, and you would eventually trigger paging or blocking this way. Is that correct? I think I saw the broker behave this way in my testing.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1096
18625	ARTEMIS-1096	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1196 Please.. don't merge this	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1096
18626	ARTEMIS-1087	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1166 Why don't you use the properties, and allow it through URI?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1087
18627	ARTEMIS-1077	Is this done?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1077
18628	ARTEMIS-1073	Github user gemmellr commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1138#discussion_r108394942 --- Diff: artemis-cli/src/main/resources/org/apache/activemq/artemis/cli/commands/etc/broker.xml --- @@ -61,10 +61,13 @@ ${ping-config.settings}${journal-buffer.settings}${connector-config.settings} <global-max-size>100Mb</global-max-size> <acceptors> + <!-- useEpoll means: it will use Netty epoll if you are on a system (Linux) that supports it --> + <!-- amqpCredits: The number of credits sent to AMQP producers --> + <!-- amqpLowCredits: The server will send the # credits specified at amqpCredits at this low mark --> --- End diff -- Should there maybe be some coverage of these options in the documentation? Do we think enough people will tweak them that they should be in the default configuration file?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1073
18629	ARTEMIS-1073	Is this finished?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1073
18630	ARTEMIS-1070	[~romain.manni-bucau] do you have Artemis integrated with MDBs already deployed somewhere? do you have some guide on how to bootstrap it? The difficult part is testing it. writing the code is the easy part.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1070
18631	ARTEMIS-1070	I couldn't find the blog... do you have a link?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1070
18632	ARTEMIS-1060	Can you elaborate on the use-case here?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1060
18633	ARTEMIS-1056	Is the reference on b819026dfcdcd4a1228774ba87dcae2e94b13a9b correct? It seems like it probably isnt for this JIRA, and the log message doesnt match an commit comment in the JIRA.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1056
18634	ARTEMIS-1052	Is this done, [~clebertsuconic]?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1052

#	Requirement	Questions	Answered	Successfully Answered	URL
18635	ARTEMIS-1045	Is this complete?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1045
18636	ARTEMIS-1042	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1098#discussion_r106419741 --- Diff: artemis-core-client/src/main/java/org/apache/activemq/artemis/core/client/impl/TopologyMemberImpl.java --- @@ -129,6 +131,22 @@ public String toURI() { return "tcp://" + host + ":" + port; } + public URI toBackupURI() { + TransportConfiguration backupConnector = getBackup(); + if (backupConnector == null) { + return null; --- End diff -- Instead of returning null, can't we return the other members on the list? I foresee a lot of cases where users don't have backup servers on the cloud.. and just have the infra-structure reconnecting them from data. Can you provide other nodes as a list here?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1042
18637	ARTEMIS-1008	GitHub user andytaylor opened a pull request: https://github.com/apache/activemq-artemis/pull/1059 ARTEMIS-1008 - Allow web server host to be set on create https://issues.apache.org/jira/browse/ARTEMIS-1008 You can merge this pull request into a Git repository by running: $ git pull https://github.com/andytaylor/activemq-artemis ARTEMIS-1008 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/activemq-artemis/pull/1059.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #1059 ---- commit 2b369f9ef96adc19c31aa7cc2c6e6376c515b43f Author: Andy Taylor <andy.tayls67@gmail.com> Date: 2017-03-01T10:19:49Z ARTEMIS-1008 - Allow web server host to be set on create https://issues.apache.org/jira/browse/ARTEMIS-1008 ----	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1008
18638	ARTEMIS-1008	Commit 2b369f9ef96adc19c31aa7cc2c6e6376c515b43f in activemq-artemis's branch refs/heads/master from [~andytaylor] [https://git-wip-us.apache.org/repos/asf?p=activemq-artemis.git;h=2b369f9] ARTEMIS-1008 - Allow web server host to be set on create https://issues.apache.org/jira/browse/ARTEMIS-1008	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1008
18639	ARTEMIS-1008	Commit 2b369f9ef96adc19c31aa7cc2c6e6376c515b43f in activemq-artemis's branch refs/heads/master from [~andytaylor] [https://git-wip-us.apache.org/repos/asf?p=activemq-artemis.git;h=2b369f9] ARTEMIS-1008 - Allow web server host to be set on create https://issues.apache.org/jira/browse/ARTEMIS-1008	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1008
18640	ARTEMIS-1002	Github user mtaylor commented on the issue: https://github.com/apache/activemq-artemis/pull/1053 @clebertsuconic The difference is that the DEFAULT pooled is platform specific on whether to use direct buffers or not. We specifically switch this to false in the PartialPooledByteBufAllocator. Could this be an issue here?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1002
18641	ARTEMIS-1002	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1053 @mtaylor Oh.. right... I actually read it org.apache.activemq.core.client.. something like that... that's the one class I saw it.. but I confused the usage.. my bad!	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1002
18642	ARTEMIS-995	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1044#discussion_r102959366 --- Diff: artemis-junit/src/test/java/org/apache/activemq/artemis/junit/MultipleEmbeddedActiveMQResourcesTest.java --- @@ -38,6 +38,11 @@ static final String ASSERT_RECEIVED_FORMAT = "Message should have been received from %s"; static final String ASSERT_COUNT_FORMAT = "Unexpected message count in queue %s"; + static { + ThreadLeakCheckRule.addKownThread("MemoryPoolMXBean notification dispatcher"); --- End diff -- Isn't there a way to kill these threads?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-995
18643	ARTEMIS-995	[~eduda] do you have more coming before the release? should I close this? We can keep this open and if you send more JIRAs on this type of fix.. we can keep this open until the release if you have more coming.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-995
18644	ARTEMIS-994	In situations like this where changes are really only to gain performance I like to quantify the use-case and goals as much as possible. I assume you've done benchmarking for your use-case and determined that the current implementation does not meet your performance goals. Can you provide details on your specific use-case, how you tested performance, what numbers you received, and what numbers you hope to achieve with these changes?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-994
18645	ARTEMIS-994	[~jbertram] I/O is I/O.. there's nothing you can do to improve performance on the IO itself.. you have a physical medium somehow, right? Either a DISK, a SSD, a Network sending waves, whatever you do you can't go faster than physical barriers.. What you can do is use less CPU and generate less Garbage on the process... I have seen the implementation when Norman was still at RedHat writing it, and we have good indications that this will happen. If we can use less CPU and generate less Garbage that will generate better throughput on the system, and there is only one way to measure that.. we implement the thing and use benchmarks to measure the gain. Francesco has a few written on his github that we can reuse.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-994
18646	ARTEMIS-994	{quote} ...use benchmarks to measure the gain. {quote} This is essentially what I'm asking about, although you're talking about retro-active benchmarking. I'm asking about pro-active benchmarking - the kind of benchmarking that would allow one to identify a performance bottleneck for a particular use-case. In any case, were going to want to perform benchmarking to understand what kind of gains this change will make and what use-cases it will benefit the most. I'm not naive enough to believe it will be the best option for any and all use-cases. If we can acquire data on this front "for free" from someone who's already done it (e.g. [~michael.andre.pearce]) then so much the better. Do you think this is a waste of time, [~clebertsuconic]?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-994
18647	ARTEMIS-994	[~michael.andre.pearce], do you have information on how to reproduce the stack-trace detailed in your earlier comment?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-994
18648	ARTEMIS-994	[~norman] I notice going back through git history it once had @author as you in the code. Also i guess you're very familiar with the native netty epoll ;) writing that also. Do you have any comments on the above, and the issue we hit? Is what I've done to get around the issue sensible?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-994
18649	ARTEMIS-994	[~clebertsuconic] thanks, i have removed this class now, and my branch is uptodate. Doing a mvn clean install, succeeds since removing this i assume this runs all the tests you mention. What should we do next, shall i make PR?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-994
18650	ARTEMIS-994	A PR is good... we may work beyond it... but it's a good start.. it would preserve your authorship.. and we may go beyond it with config options.. etc...	1	1	https://issues.apache.org/jira/browse/ARTEMIS-994

#	Requirement	Questions	Answered	Successfully Answered	URL
18651	ARTEMIS-994	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1093 I am running the testsuite, and if everything is ok I should merge it tomorrow... I have 50% of tests running and so far so good.. will know in 2 hours... BTW: can you run these commands please? ```sh git pull -rebase upstream master git push origin -f {your branch name} ``` and then: ```sh git rebase -i upstream/master ``` And squash these committs into one, with a description that matches it? I could do this myself, but I would rather have you finding the best commit name once you squash them.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-994
18652	ARTEMIS-994	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1093 @michaelandrepearce I don't think you really need three parameters for this... just one would do... - useEpoll... the other two you introduced could just use the same parameter, being the same semantic.. being just epoll. That way you would configure to use epoll or not with a single switch. If you need to configure different settings.. than you can just update the values.. it gets easier on users IMHO. WDYT?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-994
18653	ARTEMIS-994	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1093 @clebertsuconic I notice this didn't get merged still. Did I miss understand and you expected me to make the configuration changes? if so no worries, just let me know.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-994
18654	ARTEMIS-994	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1093 @clebertsuconic i just did a complete fresh clone and rebuild of this PR/My branch which i rebased the other day. It still built ok for me. Should i rebase again? Maybe some commit in the last few days might have broken something?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-994
18655	ARTEMIS-994	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1093 Many thanks Clebert. Had a brief look over tonight. Seems good. Will spend a bit more time tomorrow and also take the chance to update my fork. Sent from my iPhone > On 24 Mar 2017, at 00:47, clebertsuconic <notifications@github.com> wrote: > > > Please review master ? > > — > You are receiving this because you were mentioned. > Reply to this email directly, view it on GitHub, or mute the thread. >	0	0	https://issues.apache.org/jira/browse/ARTEMIS-994
18656	ARTEMIS-948	Couple of questions: # What is triggering the creation of the destination? Is it a producer sending a message or consumer being created? # If JConsole doesn't list the destination how do you know it is being created? # Do you want "topic" semantics? If so, why are you using a destination prefixed with "jms.queue." (which will have queue semantics)? # Do you have a reproducible test-case you could provide?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-948
18657	ARTEMIS-948	You should use the user's list to ask question, only open JIRAs for confirmed bugs please. Can we close this?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-948
18658	ARTEMIS-922	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/982#discussion_r98759149 --- Diff: tests/integration-tests/src/test/java/org/apache/activemq/artemis/tests/integration/addressing/AddressingTest.java --- @@ -224,37 +224,30 @@ public void testMulticastRoutingBackwardsCompat() throws Exception { @Test public void testPurgeOnNoConsumersTrue() throws Exception { - SimpleString address = new SimpleString("test.address"); SimpleString queueName = SimpleString.toSimpleString(UUID.randomUUID().toString()); - // For each address, create 2 Queues with the same address, assert both queues receive message - boolean purgeOnNoConsumers = true; - Queue q1 = server.createQueue(address, RoutingType.MULTICAST, queueName, null, true, false, Queue.MAX_CONSUMERS_UNLIMITED, purgeOnNoConsumers, true); - + server.createQueue(address, RoutingType.ANYCAST, queueName, null, null, true, false, false, false, false, 1, true, true); + assertNotNull(server.locateQueue(queueName)); ClientSession session = sessionFactory.createSession(); - session.start(); - - ClientConsumer consumer1 = session.createConsumer(q1.getName()); - consumer1.close(); - - assertFalse(server.queueQuery(queueName).isExists()); + ClientProducer producer = session.createProducer(address); + producer.send(session.createMessage(true)); + session.createConsumer(queueName).close(); --- End diff -- is this really synchronous? don't you need a Wait here?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-922
18659	ARTEMIS-922	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/982#discussion_r98766157 --- Diff: tests/integration-tests/src/test/java/org/apache/activemq/artemis/tests/integration/addressing/AddressingTest.java --- @@ -224,37 +224,30 @@ public void testMulticastRoutingBackwardsCompat() throws Exception { @Test public void testPurgeOnNoConsumersTrue() throws Exception { - SimpleString address = new SimpleString("test.address"); SimpleString queueName = SimpleString.toSimpleString(UUID.randomUUID().toString()); - // For each address, create 2 Queues with the same address, assert both queues receive message - boolean purgeOnNoConsumers = true; - Queue q1 = server.createQueue(address, RoutingType.MULTICAST, queueName, null, true, false, Queue.MAX_CONSUMERS_UNLIMITED, purgeOnNoConsumers, true); - + server.createQueue(address, RoutingType.ANYCAST, queueName, null, null, true, false, false, false, false, 1, true, true); + assertNotNull(server.locateQueue(queueName)); ClientSession session = sessionFactory.createSession(); - session.start(); - - ClientConsumer consumer1 = session.createConsumer(q1.getName()); - consumer1.close(); - - assertFalse(server.queueQuery(queueName).isExists()); + ClientProducer producer = session.createProducer(address); + producer.send(session.createMessage(true)); + session.createConsumer(queueName).close(); --- End diff -- I know.. but will delete the queue be synchronous as well? no executors on the implementation whatsoever?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-922
18660	ARTEMIS-905	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/945 Does this yield any measurable performance gains? I'm curious about any benchmarks performed.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-905
18661	ARTEMIS-905	Github user ig-michaelpearce commented on the issue: https://github.com/apache/activemq-artemis/pull/945 Out of interest why JCTools over some other libs with similar goals? Agrona, Kolobroke, GS collections, fastutil, HPPC & Trove	1	1	https://issues.apache.org/jira/browse/ARTEMIS-905
18662	ARTEMIS-905	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/945 +1 on this. only thing is: make sure you run a full testsuite before running this.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-905
18663	ARTEMIS-905	Github user ig-michaelpearce commented on the issue: https://github.com/apache/activemq-artemis/pull/945 @franz1981 if all reference to ConcurrentHashSet are replaced, can the class be removed in org.apache.activemq.artemis.utils?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-905
18664	ARTEMIS-904	Github user bennetelli commented on the issue: https://github.com/apache/activemq-artemis/pull/944 Can't find any checkstyle warnings. Where are you looking for it? The Jenkins build failed because of an error in LDAP stuff which I didn't touch with my new markdown documentation.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-904

#	Requirement	Questions	Answered	Successfully Answered	URL
18665	ARTEMIS-904	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1231 @jbertram / @bennetelli nothing works after this.. you can't even start a server. What was changed here that broker it? I may have to revert it if not fixed by tomorrow. i have a release to make.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-904
18666	ARTEMIS-904	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/1231 I don't understand. What specifically isn't working now? I ran the PR build myself which starts the broker quite a few times for different tests and everything worked.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-904
18667	ARTEMIS-898	Github user michaelandrepearce commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1242#discussion_r114443703 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/ActiveMQServerImpl.java --- @@ -1808,6 +1822,33 @@ public void callPostQueueDeletionCallbacks(final SimpleString address, } @Override + public void registerBrokerPlugins(final List<ActiveMQServerPlugin> plugins) { + configuration.registerBrokerPlugins(plugins); + } + + @Override + public void registerBrokerPlugin(final ActiveMQServerPlugin plugin) { + configuration.registerBrokerPlugin(plugin); + } + + @Override + public void unRegisterBrokerPlugin(final ActiveMQServerPlugin plugin) { + configuration.unRegisterBrokerPlugin(plugin); + } + + @Override + public List<ActiveMQServerPlugin> getBrokerPlugins() { + return configuration.getBrokerPlugins(); + } + + @Override + public void callPlugins(ActiveMQPluginRunnable pluginRun) { + if (getBrokerPlugins().size() > 0) { -- - End diff -- Can this if statement be made before the runnable class is intsatiated? To avoid object creation if no plugins	1	1	https://issues.apache.org/jira/browse/ARTEMIS-898
18668	ARTEMIS-898	[~cshannon] Thank you for implementing this feature! Your interface expands on what was available in the 5.x line, and I think plugin writers will benefit from the lifecycle approach (beforeX / afterX, etc). Use of default methods a huge plus as well. I did a gap analysis comparing the 5.x plugin API w/ your new Artemis plugin API and found a couple gaps. Would it be helpful for me to comment them here or would you prefer a follow-on ticket?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-898
18669	ARTEMIS-898	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1258 did you think about using ServiceLocator?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-898
18670	ARTEMIS-898	Github user cshannon commented on the issue: https://github.com/apache/activemq-artemis/pull/1258 No, what do you mean?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-898
18671	ARTEMIS-898	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1258 Could I suggest this feature goes to 2..1.1 release rather than holding up 2.1.0 tag? If there's still design and implementation details being fleshed out.?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-898
18672	ARTEMIS-878	Github user mtaylor commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/924#discussion_r92656323 --- Diff: artemis-cli/src/main/java/org/apache/activemq/artemis/cli/commands/queue/UpdateQueue.java --- @@ -31,7 +31,7 @@ String name; @Option(name = "--deleteOnNoConsumers", description = "whether to delete when it's last consumers disconnects") - Boolean deleteOnNoConsumers = null; + String deleteOnNoConsumers = null; --- End diff -- There's an underlying problem with this approach. The presence of an argument represents true and the non presence represents false. But what if the field is optional? Does that imply false? Specifically if you want to update say a queue by setting "deleteOnNoConsumers" to a different value that is already set, how do we distinguish whether or not the user wants to set it to false, or the user does not want to update it? IMO it's better to be explicit about options like this. If we really want to keep the presence as an option then perhaps we can add another flag, enableDeleteOnNoConsumers, IMO this is less usable than just adding an option for "true" or "false" to be passed in.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-878
18673	ARTEMIS-878	Github user franz1981 commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/924#discussion_r92684194 --- Diff: artemis-cli/src/main/java/org/apache/activemq/artemis/cli/commands/queue/UpdateQueue.java --- @@ -31,7 +31,7 @@ String name; @Option(name = "--deleteOnNoConsumers", description = "whether to delete when it's last consumers disconnects") - Boolean deleteOnNoConsumers = null; + String deleteOnNoConsumers = null; --- End diff -- Yes, I'm reusing this branch :+1: Right now, I'm searching for an already existing example on CLI commands to look on how it is implemented (and use the same name pattern) as a reference, but i haven't found it yet, do you know where it is?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-878
18674	ARTEMIS-878	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/924#discussion_r92696163 --- Diff: artemis-cli/src/main/java/org/apache/activemq/artemis/cli/commands/queue/UpdateQueue.java --- @@ -31,7 +31,7 @@ String name; @Option(name = "--deleteOnNoConsumers", description = "whether to delete when it's last consumers disconnects") - Boolean deleteOnNoConsumers = null; + String deleteOnNoConsumers = null; --- End diff -- look at Create, allow-anonymous | --require-login (Remember when I talked about being creative ? :)) anyways.. I was looking into this, and I found a better way to address this, by using inputs on certain defaults.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-878
18675	ARTEMIS-873	[~jmesnil], I considered this when I implemented it, but the main impetus for this was parity with 5.x which uses the 1024 multiple. I didn't want a disparity in actual behavior given the same abbreviation. That said, if the community agrees with your suggestion I'm happy to implement it. Would you mind raising a \[DISCUSS\] thread on the user mailing list and making your proposal?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-873
18676	ARTEMIS-873	I'm starting to agree with Jeff on this. What about timeouts... if you define 3K Seconds it should be 3000 seconds and not 3 * 1024 seconds. What should we do.. K should mean 1024 on bytes, and 1000 on other values? it sounds horrible that way. Time for a DISCUSS thread?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-873
18677	ARTEMIS-869	I'm not clear on how "having the delivery mode explicitly part of the API" is actually more user friendly than simply using the existing setting. Can you elaborate on that point?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-869
18678	ARTEMIS-865	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/896 nice one.. thanks	1	1	https://issues.apache.org/jira/browse/ARTEMIS-865
18679	ARTEMIS-843	The JMS API from Oracle is licensed with the CDDL which requires some special care when including in an Apache product. My guess is that's why it wasn't used originally. In any case, can you elaborate on what you mean by "things become complicated because of that"? To my knowledge ActiveMQ has always used JMS specs from Geronimo (many Java EE application servers have their own versions of the specs).	0	0	https://issues.apache.org/jira/browse/ARTEMIS-843

#	Requirement	Questions	Answered	Successfully Answered	URL
18680	ARTEMIS-843	In my case its netbeans platform environment which does not like when packages being imported by more than one module. And it does not help in any way, it just complains. To find out which deps have those packages, check versions, find the version, which suits all the JMS impls, find repos and register them, exclude this jms dep, include mine, retest all its users - I ment all these fun stuff I had to find out what geronimo is and check if it is actually THAT jms, or what? If its that jms then why it was repackaged? - it was not obvoius at all And geronimo seems dead since 2013	0	0	https://issues.apache.org/jira/browse/ARTEMIS-843
18681	ARTEMIS-843	I suppose strange is relative. As I indicated before, the Geronimo project serves as an umbrella for Java EE specs for the Apache community. Therefore most, if not all, Apache projects that need Java EE specs use the ones from Geronimo. That list includes Artemis, of course, as well as ActiveMQ 5.x, Apollo, Camel, Qpid JMS, etc. Back when I worked on HornetQ it used the JMS spec from JBoss. I'd say this practice is fairly common in larger communities like Apache and JBoss.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-843
18682	ARTEMIS-839	I am pretty sure I had this somewhere else. The only issue to do this would be when transactions are involved. Which journal would you keep the transactions? I guess we could still have the main journal acting as the TX coordinator and have IDs registered.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-839
18683	ARTEMIS-839	What's the status on this? We have found we are maxing out at 50k msg/sec (durable) with ssd, we tried raiding but the raid over head was actually detremental. As such we cannot benefit by having multiple disks to get more performance. Where as if we could have multiple journal dirs then we could obviously make use because we just mount the disks direct and get total durable throughput higher. This would be akin to some high performance databases do like aerospike etc.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-839
18684	ARTEMIS-839	[~jbertram] we're running durable persistent messages in sync send over many clients, ramping up the clients as it progresses, what we note is essentially the disk is bottle neck, surprise surprise. Because of the disks we have the raid stack overhead actually degrades the performance we're getting. This is exact reason why other tech like the one i noted have started to look to support JBOD disks for their persistent areas and itself stripe the data it needs to write. Also aside from performance improvements, this would give a functional benefit of flow isolation, that you can segregate flows so that very latency sensitive flows could goto one disk and less sensitive can goto another (or other isolation reasons), so that at the disk level they can't impact one another. [~clebertsuconic] you note you had some initial work in this area, is this something you could share?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-839
18685	ARTEMIS-830	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/873 Can you just close it? I tried to close it myself and it didn't work.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-830
18686	ARTEMIS-824	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/867 @lulf build failure on this PR is accurate (apparently). Can you fix it? Try ``` mvn -Pdev install -DskipTests=true ``` on your local box.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-824
18687	ARTEMIS-824	Github user lulf commented on the issue: https://github.com/apache/activemq-artemis/pull/867 @clebertsuconic I'm not able to reproduce the issue. Both with and without skipTests=true works for me. From the jenkins test run it looks like it timed out, and I'm not able to find any test results. Is it possible to rerun the build?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-824
18688	ARTEMIS-822	Are you talking about the synchronized block in the write method of JournalImpl where the writes are enqueued?. If so, some sort of sync will probably always be needed to compute the correct offset for each record and checking for a rollover. But it's worth looking at to see if the time in the sync block could be reduced or rewritten in a way to reduce contention. As a reference, in the ActiveMQ 5.x broker, KahaDB does something very similar where records are enqueued and batched together under a lock and then a single thread writes the batches to the journal file.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-822
18689	ARTEMIS-821	I thought that specifying relative delay instead of absolute time is a way to deal with the possibility of clock skew between clients and server. Another illusion shattered. :| This is yet another strange thing related to this. Section 7.9. of JMS 2.0 Spec: "Note however that the setDeliveryDelay method on Message cannot be used to set the delivery delay of a message." What?!? Why they even put the method there, then, if it does not do what it is supposed to do?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-821
18690	ARTEMIS-819	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/864#discussion_r84759299 --- Diff: pom.xml -- @@ -128,8 +128,8 @@ <project.build.sourceEncoding>UTF-8</project.build.sourceEncoding> <project.reporting.outputEncoding>UTF-8</project.reporting.outputEncoding> - <maven.compiler.source>1.7</maven.compiler.source> - <maven.compiler.target>1.7</maven.compiler.target> + <maven.compiler.source>1.8</maven.compiler.source> --- End diff -- Do we really want this in 1.5, or target for 1.6?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-819
18691	ARTEMIS-815	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/964 @dejanb, where are you on the test effort? Also, can you amend your commit message to follow the 50/72 format discussed in the [Hacking Guide](https://github.com/apache/activemq-artemis/blob/master/docs/hacking-guide/en/code.md#commitMessageDetails)?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-815
18692	ARTEMIS-802	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/850 Just for consistency's sake, can you remove the square brackets from around the JIRA? Back in the old days those were required for tracking commits, but they aren't needed with current tools. Also, can you add a simple test to ensure the new functionality is working as expected and that if it changes inadvertently in the future we'll catch it?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-802
18693	ARTEMIS-796	So, can I close these as won't fix, given we already have support for management?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-796
18694	ARTEMIS-796	We already have Management Notifications in Artemis. If you see any missing events, please add specific JIRAs for the ones you think we should add. You mentioned producers.. even though we don't have a Producer's server object.. some protocols will have internal Producers.. and when core-protocol is sent we could add some control on destinations... can you add a JIRA for that ?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-796

#	Requirement	Questions	Answered	Successfully Answered	URL
18695	ARTEMIS-793	artemis-server-osgi/src/main/java/org/apache/activemq/artemis/osgi/OsgiBroker.java --- @@ -1,223 +1,252 @@ -/* - * Licensed to the Apache Software Foundation (SF) under one or more - * contributor license agreements. See the NOTICE file distributed with - * this work for additional information regarding copyright ownership. - * The ASF licenses this file to You under the Apache License, Version 2.0 - * (the "License"); you may not use this file except in compliance with - * the License. You may obtain a copy of the License at - * - * http://www.apache.org/licenses/LICENSE-2.0 - * - * Unless required by applicable law or agreed to in writing, software - * distributed under the License is distributed on an "AS IS" BASIS, - * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. - * See the License for the specific language governing permissions and - * limitations under the License. - */ -package org.apache.activemq.artemis.osgi; - -import java.io.File; -import java.lang.management.ManagementFactory; -import java.util.ArrayList; -import java.util.Dictionary; -import java.util.Enumeration; -import java.util.HashMap; -import java.util.HashSet; -import java.util.Hashtable; -import java.util.List; -import java.util.Map; -import java.util.Set; - -import org.apache.activemq.artemis.api.core.Interceptor; -import org.apache.activemq.artemis.api.core.TransportConfiguration; -import org.apache.activemq.artemis.core.config.FileDeploymentManager; -import org.apache.activemq.artemis.core.config.impl.FileConfiguration; -import org.apache.activemq.artemis.core.remoting.impl.netty.TransportConstants; -import org.apache.activemq.artemis.core.server.ActiveMQComponent; -import org.apache.activemq.artemis.core.server.ActiveMQServer; -import org.apache.activemq.artemis.jms.server.config.impl.FileJMSConfiguration; -import org.apache.activemq.artemis.spi.core.protocol.ProtocolManagerFactory; -import org.apache.activemq.artemis.spi.core.security.ActiveMQJAASSecurityManager; -import org.osgi.framework.BundleContext; -import org.osgi.framework.ServiceRegistration; -import org.osgi.service.component.ComponentContext; -import org.osgi.service.component.annotations.Activate; -import org.osgi.service.component.annotations.Component; -import org.osgi.service.component.annotations.ConfigurationPolicy; -import org.osgi.service.component.annotations.Deactivate; -import org.osgi.util.tracker.ServiceTracker; - -@SuppressWarnings({"unchecked", "rawtypes"}) -@Component(configurationPid = "org.apache.activemq.artemis", configurationPolicy = ConfigurationPolicy.REQUIRE) -public class OsgiBroker { - - private String name; - private String configurationUrl; - private String rolePrincipalClass; - private Map<String, ActiveMQComponent> components; - private Map<String, ServiceRegistration<?>> registrations; - private ServiceTracker tracker; - - @Activate public void activate(ComponentContext cctx) throws Exception { final BundleContext context = cctx.getBundleContext(); final Dictionary<String, Object> properties = cctx.getProperties(); configurationUrl = getMandatory(properties, "config"); name = getMandatory(properties, "name"); rolePrincipalClass = (String) properties.get("rolePrincipalClass"); String domain = getMandatory(properties, "domain"); ActiveMQJAASSecurityManager security = new ActiveMQJAASSecurityManager(domain); if (rolePrincipalClass != null) { security.setRolePrincipalClass(rolePrincipalClass); } String brokerInstance = null; String karafDataDir = System.getProperty("karaf.data"); if (karafDataDir != null) { brokerInstance = karafDataDir + "/artemis/" + name; } // todo if we start to pullout more configs from the main config then we // should pull out the configuration objects from factories if available FileConfiguration configuration = new FileConfiguration(); if (brokerInstance != null) { configuration.setBrokerInstance(new File(brokerInstance)); } FileJMSConfiguration	1	1	https://issues.apache.org/jira/browse/ARTEMIS-793
18696	ARTEMIS-793	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/839 There is one change that is wrong here. You are using the ProtocolTrackerCallback to admin DataSources. Can you separate this into a separate class? DataSourceTrackerCallback? And instead of returning a true or false, can you return the name of the datasources? Also: I coudln't comment on the change here given the merge mess you have on this branch now. Can you execute these commands please? ``` git pull --rebase upstream master git push origin -f your-branch-name ```	1	1	https://issues.apache.org/jira/browse/ARTEMIS-793
18697	ARTEMIS-793	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/839#discussion_r83920218 --- Diff: artemis-server-osgi/src/main/java/org/apache/activemq/artemis/osgi/ProtocolTrackerCallBack.java --- @@ -25,4 +25,6 @@ void addFactory(ProtocolManagerFactory<Interceptor> pmf); void removeFactory(ProtocolManagerFactory<Interceptor> pmf); + + void setDataSourceDependency(boolean dataSourceDependency); --- End diff -- what about renaming this to: addDataSource(WhateverInformationForTheDatasource); the callback could then decide what to do. the best would be to have a class extending ProtocolTrackerCallback: ServerTrackerCallback extends ProtocolTrackerCallback And at the caller we would check an instanceof. Would that be possible with OSGI? BTW: thanks for rebasing... at least I can comment on these changes now.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-793
18698	ARTEMIS-793	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/839 @graben I have merged, after I reverted some of the bad code I made you do. Can you check please? And sorry about that.. I really thought the ProtocolTracker was externally used.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-793
18699	ARTEMIS-790	Is this done?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-790
18700	ARTEMIS-786	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/835 looks nice... Let me run the whole testsuite before merging though.. I want also to run the examples. Have you done it? I'm a bit concerned with the examples as this will touch configurations.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-786
18701	ARTEMIS-786	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/835 @jbertram I will let you handle it then... please run the whole testsuite, including examples?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-786
18702	ARTEMIS-786	Github user gaohoward commented on the issue: https://github.com/apache/activemq-artemis/pull/835 @jbertram I move the hash factory api into the same PasswordMaskingUtil to reduce confusion, so that both 2-way encryption and one-way hashing APIs are in the same util class. But I'm not sure what's the benefit of bring in Jasypt. The PasswordMaskingUtil provides a general API which enables user to use customized password encoder/decoder like Jasypt, if he really want it. What do you think?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-786
18703	ARTEMIS-786	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/835 Is there a command to generate an encoded password so that you can put that into your configuration? I didn't see anything related to that.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-786
18704	ARTEMIS-786	Github user gaohoward commented on the issue: https://github.com/apache/activemq-artemis/pull/835 good idea. If you mean just doing the masking that could be done easily. If you mean also update the broker.xml with the command, I think it'll need some thinking and should be done in a separate JIRA. what do you think?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-786
18705	ARTEMIS-781	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/854 Ohhh... this is Artemis-781 branch. didn't realize that. should we be doing PRs on it? or just commit straight? I was doing straight committs when I was working with Howard on OpenWire.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-781

#	Requirement	Questions	Answered	Successfully Answered	URL
18706	ARTEMIS-780	[~martyntaylor] Would it make sense to put this out on confluence? Is there one for artemis, or perhaps a directory on the activemq one?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-780
18707	ARTEMIS-780	Items to consider: 0. Pre-configuring destinations is a big brain drain, so anything that can be client-driven is a win. Also, protocol specific handlers could perform the admin operations on-demand. For example: session.createDurableSubscriber(...) The JMS handler create the subscription on the behalf of the client. 1. Separate topic and queue namespaces.. in JMS topic:///foo != queue:///foo. The addressing will need some sort of way to separate the two during naming collisions. 2. In ActiveMQ 5.x, AMQP and STOMP handled the addressing by using queue:/// and topic:/// prefixes. I don't think that is necessarily a bad thing, but something to consider b/c we need to support #1 3. As far as destination behaviors, how about using uri parameters to pass provider (Artemis) specific settings on-the-fly? For example: in AMPQ the address could be topic:///foo?type=nonSharedNonDurable etc.. same for MQTT, STOMP, etc. There is precedence in using uri parameters to configure the Destination in JMS as well. IBM MQ has session.createQueue("My.Queue?targetClient=1") Note: AMQP supports options as well, so that could be used as well. However, uri's tend to be better for externalizing configuration management. 4. Destination name separation b/w protocol handlers. STOMP and MQTT like "/" and JMS likes "." as destination name separators. Is there any thought to having a native destination name separator and then have protocol-specific convertors? 5. Fully qualified destination names that include a broker name. Other providers support fully-qualified destination names in JMS following the format: queue://$brokerName/$queueName. Adding this would go a long way to supporting migration of current applications without having to change client-code. Note: This would probably impact cluster handling as well, so perhaps in phase 1 there is just a placeholder for supporting a broker name in the future?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-780
18708	ARTEMIS-762	Github user tabish121 commented on the issue: https://github.com/apache/activemq-artemis/pull/814 Was the commit of the log file intentional?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-762
18709	ARTEMIS-756	Once more...Is this done, [~johndament]?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-756
18710	ARTEMIS-756	[~johndament] how is this? Now that I know OSGi a bit better I'm not sure the example makes sense.. Maybe docs, sure... you know OSGi better than I do.. so what you think about this?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-756
18711	ARTEMIS-756	Github user johnament commented on the issue: https://github.com/apache/activemq-artemis/pull/857 So there's a couple of things I don't like right now - need some input from you guys. - The artemis maven plugin seems to need to redeclare dependencies. Is there a way to do this in a profile? - Do I need a `broker.xml` or can that be inherited somehow?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-756
18712	ARTEMIS-756	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/857 That is the finalize closing the connection that was left open. Is that a issue? do you have anything to be called when the osgi is going down and close the connection?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-756
18713	ARTEMIS-756	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/861#discussion_r84735576 --- Diff: artemis-rest/src/main/java/org/apache/activemq/artemis/rest/Jms.java --- @@ -0,0 +1,174 @@ + +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.activemq.artemis.rest; + +import javax.jms.BytesMessage; +import javax.jms.JMSException; +import javax.jms.Message; +import javax.jms.ObjectMessage; +import javax.ws.rs.core.MediaType; +import javax.ws.rs.ext.MessageBodyReader; +import javax.ws.rs.ext.Providers; +import java.io.ByteArrayInputStream; +import java.lang.reflect.Type; + +import org.apache.activemq.artemis.rest.util.HttpMessageHelper; +import org.jboss.resteasy.core.Headers; +import org.jboss.resteasy.spi.ResteasyProviderFactory; +import org.jboss.resteasy.util.GenericType; + +public class Jms { --- End diff -- Is this on purpose?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-756
18714	ARTEMIS-743	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/787#discussion_r79860068 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/QueueConfig.java --- @@ -0,0 +1,242 @@ + +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.activemq.artemis.core.server; + +import org.apache.activemq.artemis.api.core.SimpleString; +import org.apache.activemq.artemis.core.filter.Filter; +import org.apache.activemq.artemis.core.filter.FilterUtils; +import org.apache.activemq.artemis.core.paging.PagingManager; +import org.apache.activemq.artemis.core.paging.cursor.PageSubscription; + +public final class QueueConfig { + + private final long id; --- End diff -- The id is assigned when the queue is persisted. Why this is part of the config options?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-743

#	Requirement	Questions	Answered	Successfully Answered	URL
18715	ARTEMIS-743	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/787#discussion_r79860213 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/ActiveMQServerImpl.java --- @@ -1290,10 +1291,7 @@ protected ServerSessionImpl internalCreateSession(String name, SessionCallback callback, OperationContext context, boolean autoCreateJMSQueues) throws Exception { - return new ServerSessionImpl(name, username, password, validatedUser, minLargeMessageSize, autoCommitSends, autoCommitAcks, preAcknowledge, configuration.isPersistDeliveryCountBeforeDelivery(), - xa, connection, storageManager, postOffice, resourceManager, securityStore, managementService, this, configuration.getManagementAddress(), - defaultAddress == null ? null : new SimpleString(defaultAddress), callback, context, autoCreateJMSQueues ? jmsQueueCreator : null, - pagingManager); + return new ServerSessionImpl(name, username, password, validatedUser, minLargeMessageSize, autoCommitSends, autoCommitAcks, preAcknowledge, configuration.isPersistDeliveryCountBeforeDelivery(), xa, connection, storageManager, postOffice, resourceManager, securityStore, managementService, this, configuration.getManagementAddress(), defaultAddress == null ? null : new SimpleString(defaultAddress), callback, context, autoCreateJMSQueues ? jmsQueueCreator : null, pagingManager); --- End diff -- this line gets too big.. can you break it?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-743
18716	ARTEMIS-737	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/789 I think it would also be good to have a clear explanation of why this approach is superior to the approach already adopted in the test-suite. Additional questions: 1. What expectation is there to adopt this style of testing? Should all new tests use this style? 2. Do you plan on converting all existing test-cases to this style? If not, I'd be concerned with having 2 methods of doing the same thing introducing confusion.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-737
18717	ARTEMIS-737	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/789 I guess this is intended for users to use on their tests? What was the intention, to refactor Artemis' testsuite, or to provide rules for users running their own tests?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-737
18718	ARTEMIS-737	Github user hqstevenson commented on the issue: https://github.com/apache/activemq-artemis/pull/789 I'd be happy to add an example, but I'm not sure exactly what you're looking for. All of the unit tests show the basics of how you embed a sever into your test, and how you use the server afterwards. Is there a specific type of test you'd like to see? BTW - once this is released, I'm going to contribute a "example-camel-artemis" to the Camel project - that will show how to use the embedded resources with Camel.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-737
18719	ARTEMIS-737	@clebertsuconic -I had an idea - I noticed that none of the examples have a "unit test". I know you can run them with some the maven infrastructure that is available in the examples, but I thought it would be nice to have a unit test for an individual example. So I updated the "Bridge" example (examples/features/standard/bridge) and added a unit test using the EmbeddedJMSResource - can you give it a look? Note that I had to change the broker.xml files to get rid of a security exception I was seeing with the test - not really sure what's going on there. But if you like the idea, I could update a few more of the examples to include this type of unit test.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-737
18720	ARTEMIS-717	Can you quantify the "high burden" that the broker is under with a large number of groups? I'm mainly thinking of hot spots in the code where a lot of time is spent (e.g. stuff you might find by attaching a profiler), but an application-level quantification would be fine as well.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-717
18721	ARTEMIS-717	How long do you plan to use a particular group? For example, if you start the year with a group named 'X' do you expect to end the year with that same group 'X' or will 'X' not be used after a certain point? I assume that when you say "doesn't expire mappings" you mean that any group ever created will need to exist in perpetuity. Also, how often will a group be used on average in your use-case?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-717
18722	ARTEMIS-714	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/771 @jmesnil there was some changes I made to JDBC that are conflicting with this. Do you mine looking at them? Also.. this is showing 3 commits.. shouldn't you squash them and push -f to your branch?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-714

#	Requirement	Questions	Answered	Successfully Answered	URL
18723	ARTEMIS-714	org.apache.activemq.artemis.tests.integration.client.LargeMessage If you run any JDBC test, there's a NPE happening: ```java [main] 15:27:39,377 WARN [org.apache.activemq.artemis.core.server] AMQ222010: Critical IO Error, shutting down the server. file=NULL, message=null: java.lang.NullPointerException at org.apache.activemq.artemis.core.persistence.impl.journal.JDBCJournalStorageManager.init(JDBCJournalStorageManager.java:89) [:] at org.apache.activemq.artemis.core.persistence.impl.journal.AbstractJournalStorageManager.<init>(AbstractJournalStorageManager.java:208) [:] at org.apache.activemq.artemis.core.persistence.impl.journal.JournalStorageManager.<init>(JournalStorageManager.java:97) [:] at org.apache.activemq.artemis.core.persistence.impl.journal.JDBCJournalStorageManager.<init>(JDBCJournalStorageManager.java:47) [:] at org.apache.activemq.artemis.core.server.impl.ActiveMQServerImpl.createStorageManager(ActiveMQServerImpl.java:1799) [:] at org.apache.activemq.artemis.core.server.impl.ActiveMQServerImpl.initialisePart1(ActiveMQServerImpl.java:1930) [:] at org.apache.activemq.artemis.core.server.impl.LiveOnlyActivation.run(LiveOnlyActivation.java:61) [:] at org.apache.activemq.artemis.core.server.impl.ActiveMQServerImpl.start(ActiveMQServerImpl.java:447) [:] at org.apache.activemq.artemis.tests.integration.client.LargeMessageTest.testBufferMultipleLargeMessages(LargeMessageTest.java:1655) [:] at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) [rt.jar:1.8.0_73] at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) [rt.jar:1.8.0_73] at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) [rt.jar:1.8.0_73] at java.lang.reflect.Method.invoke(Method.java:497) [rt.jar:1.8.0_73] at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) [junit-4.11.jar:] at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) [junit-4.11.jar:] at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) [junit-4.11.jar:] at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:17) [junit-4.11.jar:] at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:26) [junit-4.11.jar:] at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:27) [junit-4.11.jar:] at org.junit.rules.ExternalResource$1.evaluate(ExternalResource.java:48) [junit-4.11.jar:] at org.junit.rules.TestWatcher$1.evaluate(TestWatcher.java:55) [junit-4.11.jar:] at org.junit.rules.ExternalResource$1.evaluate(ExternalResource.java:48) [junit-4.11.jar:] at org.junit.rules.ExternalResource$1.evaluate(ExternalResource.java:48) [junit-4.11.jar:] at org.junit.rules.TestWatcher$1.evaluate(TestWatcher.java:55) [junit-4.11.jar:] at org.junit.rules.RunRules.evaluate(RunRules.java:20) [junit-4.11.jar:] at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:271) [junit-4.11.jar:] at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:70) [junit-4.11.jar:] at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:50) [junit-4.11.jar:] at org.junit.runners.ParentRunner$3.run(ParentRunner.java:238) [junit-4.11.jar:] at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:63) [junit-4.11.jar:] at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:236) [junit-4.11.jar:] at org.junit.runners.ParentRunner.access$000(ParentRunner.java:53) [junit-	1	1	https://issues.apache.org/jira/browse/ARTEMIS-714
18724	ARTEMIS-679	I meant.. you wouldn't be able to start a server without a replica.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-679
18725	ARTEMIS-670	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/705 I think this gets a bit confused for users... Why don't you add the following commands: create-queue create-topic delete-queue delete-topic You can have abstract classes for what is common among the classes. It will be the same implementation almost, but the user would have more control on what he wants to be done. The way this is structure, you are using a create-queue command to delete it. I don't think it's too intuitive to the user.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-670
18726	ARTEMIS-670	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/705 @gaohoward when I wrote the delete-topic up here was just as an example of separating the commands. I didn't mean to be literal. Do you actually need to differenticate delete-topic from delete-queue or it could just be delete?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-670
18727	ARTEMIS-659	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/694 @johnament tests are good here. can we merge this? or you have other plans to this dependency on your next work on CDI?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-659
18728	ARTEMIS-658	Can you elaborate on the use-case here a bit? Perhaps it should be obvious, but it isn't obvious to me.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-658
18729	ARTEMIS-658	The only thing I see at that link is a few notes about logistics. There's nothing about the general use-case. Am I missing something?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-658
18730	ARTEMIS-656	I believe it's true that Artemis doesn't currently do hostname verification for SSL/TLS connections, and I aim to correct that. However, I'm a bit confused by the description. It says, "Artemis doesn't configure a trust manager when connecting bridges..." As I understand it, bridge connections are treated just like any other client connection which connects to a Netty acceptor and a trust manager will be created for that acceptor assuming the it was configured with a valid trustStorePath and trustStorePassword. The same is true for the Netty connector which the bridge uses. Is this what you're talking about? In any case, I'm not clear on how the trust manager relates to hostname verification. Also, I don't understand the phrase, "...there's no attempt to verify that the hostname of the target broker matches the one that triggered the connection." Unless I'm thinking about this wrong, the name of the host that triggers the connection will always be different from the name of host of the target broker unless the two are actually running on the same host (which would be unusual). Therefore, they shouldn't match. Right?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-656
18731	ARTEMIS-642	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/672#discussion_r72420954 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/QueueImpl.java --- @@ -2971,6 +2975,7 @@ public void run() { else if (consumerRate < threshold) { RemotingConnection connection = null; RemotingService remotingService = ((PostOfficeImpl) postOffice).getServer().getRemotingService(); + ActiveMQServer server = ((PostOfficeImpl) postOffice).getServer(); --- End diff -- The code now casts twice. Should this be in the interface? Should the casting be done once? Can server be extracted for the line above?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-642

#	Requirement	Questions	Answered	Successfully Answered	URL
18732	ARTEMIS-642	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/672#discussion_r72421069 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/QueueImpl.java --- @@ -2983,6 +2988,18 @@ else if (consumerRate < threshold) { if (connection != null) { ActiveMQServerLogger.LOGGER.slowConsumerDetected(serverConsumer.getSessionID(), serverConsumer.getID(), getName().toString(), connection.getRemoteAddress(), threshold, consumerRate); if (policy.equals(SlowConsumerPolicy.KILL)) { + if (connection instanceof CoreRemotingConnection) { + CoreRemotingConnection coreconnection = (CoreRemotingConnection)connection; + + Channel channel1 = coreconnection.getChannel(1, -1); --- End diff -- maybe `clientChannel` or something similar? `channel1` doesn't give much description.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-642
18733	ARTEMIS-642	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/672#discussion_r72421190 --- Diff: artemis-core-client/src/main/java/org/apache/activemq/artemis/core/protocol/core/impl/wireformat/DisconnectConsumerWithKillMessage.java --- @@ -0,0 +1,58 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ +package org.apache.activemq.artemis.core.protocol.core.impl.wireformat; +import org.apache.activemq.artemis.api.core.ActiveMQBuffer; +import org.apache.activemq.artemis.api.core.SimpleString; +import org.apache.activemq.artemis.core.protocol.core.impl.PacketImpl; +public class DisconnectConsumerWithKillMessage extends PacketImpl { + private SimpleString nodeID; + public DisconnectConsumerWithKillMessage(final SimpleString nodeID) { super(DISCONNECT_CONSUMER_KILL); this.nodeID = nodeID; } public DisconnectConsumerWithKillMessage() { --- End diff -- Is this constructor needed? If not, then the field can be final.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-642
18734	ARTEMIS-642	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/672#discussion_r72686692 --- Diff: artemis-core-client/src/main/java/org/apache/activemq/artemis/core/protocol/core/impl/ActiveMQSessionContext.java --- @@ -759,6 +769,12 @@ protected void handleReceiveProducerFailCredits(SessionProducerCreditsFailMessag handleReceiveProducerFailCredits(message.getAddress(), message.getCredits()); } + protected void handleReceiveSlowConsumerKillMessage(DisconnectConsumerWithKillMessage message) { + if (message.getNodeID() != null) { + setKilled(true); --- End diff -- since its always called with true, would it make sense to have a `kill()` method instead?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-642
18735	ARTEMIS-642	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/672#discussion_r72686789 --- Diff: artemis-core-client/src/main/java/org/apache/activemq/artemis/core/protocol/core/impl/RemotingConnectionImpl.java --- @@ -387,4 +389,15 @@ public void setClientID(String cID) { public String getClientID() { return clientID; } + + @Override + public void sentKilledMessage(SimpleString nodeID, Version version) throws Exception { --- End diff -- should probably be `sendKillMessage`?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-642
18736	ARTEMIS-642	Github user bayern39 commented on the issue: https://github.com/apache/activemq-artemis/pull/672 @clebertsuconic Hi Clebert, I know how to get client version,but i do not know how to get supported version, can you tell me more detail？Thanks	1	1	https://issues.apache.org/jira/browse/ARTEMIS-642
18737	ARTEMIS-642	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/672#discussion_r72964880 --- Diff: artemis-core-client/src/main/java/org/apache/activemq/artemis/core/protocol/core/impl/RemotingConnectionImpl.java --- @@ -387,4 +388,15 @@ public void setClientID(String cID) { public String getClientID() { return clientID; } + + @Override + public void killMessage(SimpleString nodeID) throws Exception { + if (clientVersion < DisconnectConsumerWithKillMessage.VERSION_INTRODUCED) { + throw new Exception(String.valueOf(clientVersion)); --- End diff -- Can this be a more specific ActiveMQ Exception? Even if its just `ActiveMQException`	1	1	https://issues.apache.org/jira/browse/ARTEMIS-642
18738	ARTEMIS-628	Github user jbertram commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/661#discussion_r72120144 --- Diff: README.md --- @@ -44,4 +44,7 @@ If you are trying to copy the examples somewhere else and modifying them. Consid # if trying to modify the 'topic' example: cd examples/jms/topic && mvn dependency:list +### The HornetQ Donation --- End diff -- What's this for?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-628
18739	ARTEMIS-627	[~gemmellr] Sure, thanks Robbie. I can certainly see how the current behaviour is a good middle ground between performance and protecting the broker. I just wanted to make sure we don't just shut this down and ignore it because it can hurt performance. I did some testing and it's quite easily to exhaust the brokers memory with a single dodgy client. I'll get the doc updated and we can address this as a feature request later down the line. I'm in the process of adding BLOCK for AMQP, we could easily add reject to the BLOCK policy. Using amqp:resource-limit-exceeded, I guess it's a little more important with AMQP since the flow control credits represent number of messages (vs number of bytes as in CORE) so it could be possible to go way over the limit quite quickly. Do you have any thoughts on what would be best in teh AMQP scneario? Thanks	1	1	https://issues.apache.org/jira/browse/ARTEMIS-627
18740	ARTEMIS-607	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/617 Could you squash these two commits together?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-607
18741	ARTEMIS-605	Can you elaborate on the use-case here a bit? If you want to prevent users from sending rogue management messages you can simply not grant them the {{manage}} permission.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-605

#	Requirement	Questions	Answered	Successfully Answered	URL
18742	ARTEMIS-601	The current implementation idea for ARTEMIS-66 is to write changes made via the management API back to the XML so the broker.xml would still be the master and changes would survive restarts. Your runtime configuration wouldn't be different from the configuration file. I assume all your management tools are all going to need to change to support Artemis' new XML format anyway. Could they not change to execute a simple shell/perl/python script. I'm not terribly familiar with Puppet. Is it only able to modify text files? If it's built for managing deployments and runtime configuration I'd be surprised if it didn't support JMX since that is the standard for managing Java applications.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-601
18743	ARTEMIS-601	Justin, this is the first time I read about "Artemis' new XML format". Is this documented/discussed somewhere? Of course, management tools could be extended with scripts talking JMX that try to find the differences between the expected configuration and the existing configuration and fire up the the right API calls. This however would move the complexity from one place (the broker) to multiple places (all the various management tools used in the wild). Instead of having one single implementation, evolving along with the broker, we would have different implementations, each with their own set of bugs, that all need to be kept in sync with changes in the broker. ActiveMQ's approach (periodically reload the configuration file) is far superior and goes beyond only security settings. It would be a pity to lease this feature in Artemis.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-601
18744	ARTEMIS-600	It's not clear to me what from the list of "Drawbacks in current design" are actually consequences of the design or just bugs in the implementation (e.g. #4 and #5 in that list sound like bugs to me, and #3 sounds like a misconfiguration). IMO, most clustered message grouping use cases don't make a lot of sense in the first place. Clustering is a way to increase performance by distributing load across multiple brokers. However, message grouping necessarily decreases performance since message consumption must be serialized. If consumption throughput is low (due to message grouping serialization) then it's likely users will want to impose flow control on producers to avoid paging unconsumed messages to disk thereby reducing producer throughput. The need for a cluster in this scenario is dubious, but I digress. [~ryeats], I'm not sure what you mean by "fully supported." Can you elaborate on that? The currently message grouping implementation is working for a fair number of users. I would expect any bugs to be fixed.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-600
18745	ARTEMIS-600	[~ryeats], at this point I can't confirm anything as the JIRA description contains no real information about what exactly is failing and in what use-cases the failures occur. IMO, this would have been better as a thread on the mailing list rather than a JIRA. [~mnovak], can you open individual JIRAs for drawbacks 3, 4, and 5 and explain the relevant use-case, what specifically isn't working, etc. Reproducers are always welcome.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-600
18746	ARTEMIS-592	Martyn, here is a description of our use case. All our clients use STOMP. We have a producer sending messages to a single topic, for instance {{jms.topic.foo}}. We have two different clients (let's call them {{A}} and {{B}}) who both want to receive all these messages. No problemo, each client creates its own durable subscription on the topic and this materializes as two different queues inside Artemis. However, we want to separate these clients and make sure one client cannot read the messages in the other client's queue. How can this be expressed as a {{<security-setting>}} element? We succesfully use ActiveMQ 5.x for this setup. We use http://activemq.apache.org/virtual-destinations.html and we assign different security settings to the different queues attached to the same topic ({{/topic/foo}}), for instance {{/queue/Consumer.A.VirtualTopic.foo}} and {{/queue/Consumer.B.VirtualTopic.foo}}.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-592
18747	ARTEMIS-592	Thanks for the explanation Lionel. I think I understand the problem, but just to ensure we're on the same page, let me iterate back what I think the problem is. There are two STOMP clients, A and B, which both subscribe to add jms.topic.foo. Which looks something like: {code} / ----------> [Ma,Ma,Ma] (Sub Q Client A) JMS Producer ----------> jms.topic.foo / \ \ ----------> [Mb,Mb,Mb] (sub Q Client B) {code} The thing you are trying to prevent, is client A setting up a subscription to jms.topic.foo and potentially consuming client B messages? I've looked into the STOMP code and I believe there is a bug that causes this problem. Subscriptions queues are not properly created, meaning two subscriptions to the same address could receive each others messages. Is this what you are seeing? If so would fixing the bug resolve this issue?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-592
18748	ARTEMIS-592	The only scenario is which a test client could interfere with the production flow in the way you describe is if the production client disconnects, the test client subsequently connects, is then authenticated, then pretends to be the production client by passing in the production client id. Is this the scenario you are trying to avoid? ActiveMQ Artemis does have a very different internal model of addressing, queues and routing and there may well be cases in which ActiveMQ 5.x needed locking down that doesn't really apply to Artemis. If you are on IRC I could go over the model with you, it might be useful to know regardless? #apache-activemq	1	0	https://issues.apache.org/jira/browse/ARTEMIS-592
18749	ARTEMIS-592	Github user mtaylor commented on the issue: https://github.com/apache/activemq-artemis/pull/701 @jbertram I'm not sure about using the queue name to control access. It requires users to understand how subscriptions work internally and how the queue name is constructed, which might be different across protocols. Perhaps a neater approach could be to allow all users of a particular role to share subscriptions, controlled by an address setting. e.g. User defines a role "production", all users within this role are able to share subscriptions between themselves. Many roles could be created that are able to share subscriptions, "production", "test", "application-x", only users within the role are able to share the subscription. Production Role: user1, user2 Test Role: user3, user4 user1 and user2 can share the subscription, user3 and user4 can also share a subscription, user1 and user3 can not share a subscription. We could also create a new security-setting to limit which roles are able to share subsctipions and for which addresses, something like "sharedSubscriptionGroup" (or better name), which specifies which roles are allowed to share subscriptions. e.g. ```xml <security-setting match="jms.topic.news.us.#"> <permission type="createDurableQueue" roles="user"/> <permission type="send" roles="us-user"/> <permission type="shareSubscription" roles="production,test"/> </security-setting> ``` We already store which user created the queue "the queue owner", we could use this information to determine whether or not to allow another user to share this subscription. You can get the queue owner roles and the current user roles to make the comparison and do a check on the security setting. Does this make sense? Any thoughts?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-592
18750	ARTEMIS-592	Github user mtaylor commented on the issue: https://github.com/apache/activemq-artemis/pull/701 @clebertsuconic sure, I think that similar feature is virtual topics, which solve the same problem as shared subscriptions. The way it works in ActiveMQ 5.x is that the user specifies the virtual topic queue to allow the shared subscription. This JIRA is with respect to virtual topics, so the first thing we need to do is actually an equivilent (shared subscription) support. That's when this JIRA starts to become a problem. Once we allow users to share subscriptions, how can we control access to the shared subscription. Queues are just an implementation detail.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-592

#	Requirement	Questions	Answered	Successfully Answered	URL
18751	ARTEMIS-592	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/701 I am not sure <permission type="shareSubscription" roles="production,test"/> would really be that user friendly as you defined. it could be confusing actually. you already have permissions to createQueue, createDurableQueue and consume. shareSubscriptions is exactly that... and it makes confusing to add an extra setting here. Artemis is a multi-protocol broker.. we are making other protocols more mature and your definitions of shareSubscription there may not fit all the protocols.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-592
18752	ARTEMIS-592	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/701 There is one concept on core called createSharedQueue, which is a temporary queue that will exist as long as there is an user connect into it. shared=true might confuse with that approach, unless this operation was specifically in regard to this type of shared queue. is that the case?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-592
18753	ARTEMIS-592	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/701 @jbertram any one... is this about Temporary Shared Queues? or anything in particular. if this is just about shared temp queues, then Martyn's suggestion could make sense.. if not.. it makes it even more confusing.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-592
18754	ARTEMIS-592	IMHO, it is very important to define precisely how Artemis handles durable STOMP subscriptions. Since they are durable, they are related to some kind of queues somewhere. Can we send messages directly to these queues? With Apollo and ActiveMQ 5.x, we can. Do we allow multiple connections to the same queue (see @clebertsuconic question "Wouldn't be easier if we just failed more than one connection with the same client-id?")? With Apollo and ActiveMQ 5.x, we can and messages will be split between all consumers, like with regular queues. Can we define different security settings (this is the initial scope of this Jira issue)? With Apollo and ActiveMQ 5.x, we can. Since Artemis supports multiple protocols, the next question is how much of this is (mostly) protocol independent and ends up in the Core and how much is protocol specific. From a user point of view, the more gets protocol independent, the better. I would prefer to have the same security settings regarding which protocols the clients end up using. Justin's proposal to concatenate the address and queue names looks fine to me.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-592
18755	ARTEMIS-592	"Can we send messages directly to these queues? With Apollo and ActiveMQ 5.x, we can." In STOMP this is possible. But not just any client has access to the queue. The STOMP implementation will prepend durable subscriptions with the client-id. So subscribing to "foo" with client-id "123" allows you to consume from the queue "123.foo". This is how the STOMP implementation behaves. "Do we allow multiple connections to the same queue (see @clebertsuconic question "Wouldn't be easier if we just failed more than one connection with the same client-id?")? With Apollo and ActiveMQ 5.x, we can and messages will be split between all consumers, like with regular queues." We do. Similar to how we handle durable subscriptions (explained above). It's possible to have more than 1 STOMP client with the same ID connect to the broker. Thus, all clients with ID 123 are able to subscribe to "123.foo". This has some limitations, a client with ID 123, is not able to "share a durable subscription queue" with client ID 456. This behaviour of sharing queues based on the client-id field isn't supported across all protocols. In order to do this we may have to break the spec of some protocols (which may or may not be acceptable) MQTT is an example of where this is explicitly outlined in the spec as not supported. TBH I'm not sure if the shared durable subscription ability (through client-id) was intentional or was a side effect of the implementation. "Can we define different security settings (this is the initial scope of this Jira issue)? With Apollo and ActiveMQ 5.x, we can." Yes this fix allows you to do that with STOMP. "Since Artemis supports multiple protocols, the next question is how much of this is (mostly) protocol independent and ends up in the Core and how much is protocol specific." Justin's point, which I think is a good one. Is that by securing queues, it's possible to apply it across protocols. The slight concern I had was that protocol implementations use queues to expose the desired protocol features. Using security settings at the queue level, requires exposing the underlying implementation details of each protocol. I'm now feeling OK with this if we can somehow standardise the internal naming schemes across all protocols. I believe STOMP durable subscriptions are <client-d>.<subscription-id>, MQTT is similar <client-id>.<subscription-address>. "From a user point of view, the more gets protocol independent, the better. I would prefer to have the same security settings regarding which protocols the clients end up using. Justin's proposal to concatenate the address and queue names looks fine to me." Yes, to make this non protocol specific we could perhaps support something similar to virtual destinations or standardise on the naming scheme.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-592
18756	ARTEMIS-592	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/750 Are you sure about ARTEMIS-592? ARTEMIS-592 is about fine grain access control.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-592
18757	ARTEMIS-590	I think this idea has merit. However, I'd like to keep the system properties to a minimum as they shouldn't be the normal way a connector is configured. The SSL related system properties only exist because that's the way Java has historically configured keystore-related details. The ideal way to configure this is with a transport constant so you could use something like {{tcp://127.0.0.1:61616?sslEnabled=true;useDefaultSSLContext=true}}. If this approach is taken then the question arises as to which settings take precedence since this new constant and the existing keystore-related constants as well as the system properties are mutually exclusive. Can you clarify how you envision this working?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-590
18758	ARTEMIS-587	Github user johnament commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/662#discussion_r72125621 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/management/impl/ActiveMQServerControlImpl.java --- @@ -1360,6 +1359,130 @@ public String listProducersInfoAsJSON() throws Exception { } @Override + public String listConnectionsAsJSON() throws Exception { + checkStarted(); + + clearIO(); + try { + JSONArray array = new JSONArray(); + Set<RemotingConnection> connections = server.getRemotingService().getConnections(); + for (RemotingConnection connection : connections) { --- End diff -- Can this wait until the json change?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-587
18759	ARTEMIS-584	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/621#discussion_r69741235 --- Diff: docs/user-manual/en/security.md --- @@ -10,6 +10,13 @@ long. To change this period set the property `security-invalidation-interval`, which is in milliseconds. The default is `10000` ms. +To assist in security auditing the `populate-validated-user` option exists. If this is `true` then +the server will add the name of the validated user to the message using the key `_AMQ_VALIDATED_USER`. +For JMS and Stomp clients this is mapped to the key `JMSXUserID`. For users authenticated based on +their SSL certificate this name is the name to which their certificate's DN maps. If `security-enabled` +is `false` and `populate-validated-user` is `true` then the server will simply use whatever user name +(if any) the client provides. This option is `false` by default. + --- End diff -- only thing missing is adding this to configuration-index.md, can you do it?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-584

#	Requirement	Questions	Answered	Successfully Answered	URL
18760	ARTEMIS-584	Github user mikehearn commented on the issue: https://github.com/apache/activemq-artemis/pull/621 Is there a reason this is disabled by default? Performance, maybe? It seems pretty useful, it's just odd that it's off by default. It makes me wonder what the catch is.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-584
18761	ARTEMIS-582	Couple of questions: * The cited MCollective documentation indicates that orphaned single-use queues will "clog up ActiveMQ's available memory." Can you confirm that the behavior here in Artemis is similar to that of ActiveMQ 5.x? In other words, do orphaned single-use queues clog up Artemis' available memory? If so, can you quantify the clogging (e.g. X number of orphaned single-use queues in Y bytes of heap causes high GC or OOME)? * Artemis supports the automatic deletion of JMS queues (see [the docs|http://activemq.apache.org/artemis/docs/1.3.0/queue-attributes.html] for more details on that). Can MCollective be configured to use JMS queues rather than core queues for the single-use queues that it needs?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-582
18762	ARTEMIS-582	Not knowing the internals of Artemis, I have no idea how much unused destinations cost in terms of memory. MCollective is just one example of an application using temporary destinations, the rate of creation of unused destinations will vary a lot with the application. "Can MCollective be configured to use JMS queues rather than core queues for the single-use queues that it needs?" Currently, no. MCollective comes with two connectors: ActiveMQ and RabbitMQ (see https://github.com/puppetlabs/marionette-collective/tree/master/lib/mcollective/connector). Since Artemis is not backward compatible with ActiveMQ 5.x, a third connector would need to be written for Artemis. Finally, it is not clear to me how automatic deletion of topics and queues work in Artemis. When would they be deleted? Immediately or after an "inactivity timeout"? In my STOMP tests, both Core and JMS addresses and queues have been created and some of these (did not check which ones) stayed after several hours and several broker restarts, triggering warnings such as: {code} 2016-06-20 13:20:03,107 [org.apache.activemq.artemis.core.server] WARN AMQ222205: OutOfMemoryError possible! There are currently 400 addresses with a total max-size-bytes of 4,194,304,000 bytes, but the maximum memory available is 764,411,904 bytes. {code}	1	0	https://issues.apache.org/jira/browse/ARTEMIS-582
18763	ARTEMIS-582	Couple of things... * There's no need to know any internals of Artemis or ActiveMQ for simple comparative analysis. All that's necessary is a standardized test used for both brokers and a way to measure the results which could be as simple as watching the memory graph in VisualVM or JConsole. * Artemis implements OpenWire so it should (theoretically) be backward compatible. Although looking more closely at MCollective it appears it's using STOMP which will almost certainly suffer some incompatibilities between brokers. * The documentation I cited previously describes when an auto-create queue or topic will be deleted. * Can you provide a simple test to reproduce what you're seeing? I wrote a quick test in the Artemis test-suite to ensure that auto-created queues are deleted as expected and it passed without issue.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-582
18764	ARTEMIS-582	I'm having troubles isolating the part that triggers this problem. Is there some debugging that I could enable on the broker side to provide useful insight on what's going on there?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-582
18765	ARTEMIS-581	Makes sense... Although I would always fail when that disk limit is hit. It would make understanding and configuration easier for users. Controlling what to do in case of a disk full, would give users too many options and would complicate user's decision. Would anyone be contrary to always fail the Send or Transaction in case the configured disk limit is hit?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-581
18766	ARTEMIS-581	Github user mtaylor commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/749#discussion_r77493191 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/ServerSessionImpl.java --- @@ -1249,6 +1256,14 @@ public RoutingStatus send(final ServerMessage message, final boolean direct) thr @Override public RoutingStatus send(final ServerMessage message, final boolean direct, boolean noAutoCreateQueue) throws Exception { + + // If the protocol doesn't support flow control, we have no choice other than fail the communication + if (!this.getRemotingConnection().isSupportsFlowControl() && pagingManager.isDiskFull()) { + ActiveMQIOErrorException exception = ActiveMQMessageBundle.BUNDLE.diskBeyondLimit(); + this.getRemotingConnection().fail(exception); + throw exception; --- End diff -- I'm not sure if killing the connection is the right thing to do here? You may also be killing consumers on the connection... Can't we just start dropping messages once this threshold is reached?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-581
18767	ARTEMIS-581	Github user mtaylor commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/749#discussion_r77524240 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/ServerSessionImpl.java --- @@ -1249,6 +1256,14 @@ public RoutingStatus send(final ServerMessage message, final boolean direct) thr @Override public RoutingStatus send(final ServerMessage message, final boolean direct, boolean noAutoCreateQueue) throws Exception { + + // If the protocol doesn't support flow control, we have no choice other than fail the communication + if (!this.getRemotingConnection().isSupportsFlowControl() && pagingManager.isDiskFull()) { + ActiveMQIOErrorException exception = ActiveMQMessageBundle.BUNDLE.diskBeyondLimit(); + this.getRemotingConnection().fail(exception); + throw exception; --- End diff -- Why are these the only 2 solutions? Surely we want to keep consumers alive so we can gracefully recover from the situation? Killing the connections could kill all the consumers too making it difficult to recover once the memory and disk limit is reached. DROP doesn't necessarily mean the messages are lost, the client should choose an appropriate qos level and retry later. Is it possible to configure it to go from PAGING to DROP once the disk limit is reached?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-581

#	Requirement	Questions	Answered	Successfully Answered	URL
18768	ARTEMIS-581	Github user mtaylor commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/749#discussion_r77524869 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/ServerSessionImpl.java --- @@ -1249,6 +1256,14 @@ public RoutingStatus send(final ServerMessage message, final boolean direct) thr @Override public RoutingStatus send(final ServerMessage message, final boolean direct, boolean noAutoCreateQueue) throws Exception { + + // If the protocol doesn't support flow control, we have no choice other than fail the communication + if (!this.getRemotingConnection().isSupportsFlowControl() && pagingManager.isDiskFull()) { + ActiveMQIOErrorException exception = ActiveMQMessageBundle.BUNDLE.diskBeyondLimit(); + this.getRemotingConnection().fail(exception); + throw exception; --- End diff -- What do you mean by silently fail? If the client chooses an appropriate QoS level (or protocol equivilent) it'll never get an ack and thus as far as the client is concerned the broker never received it. On Mon, Sep 5, 2016 at 2:49 PM, clebertsuconic <notifications@github.com> wrote: > In artemis-server/src/main/java/org/apache/activemq/artemis/ > core/server/impl/ServerSessionImpl.java > <https://github.com/apache/activemq-artemis/pull/749#discussion_r77524662> : > > @@ -1249,6 +1256,14 @@ public RoutingStatus send(final ServerMessage message, final boolean direct) thr > > @Override > > public RoutingStatus send(final ServerMessage message, final boolean direct, boolean noAutoCreateQueue) throws Exception { > > + > > + // If the protocol doesn't support flow control, we have no choice other than fail the communication > > + if (!this.getRemotingConnection().isSupportsFlowControl() && pagingManager.isDiskFull()) { > > + ActiveMQIOErrorException exception = ActiveMQMessageBundle.BUNDLE.diskBeyondLimit(); > > + this.getRemotingConnection().fail(exception); > > + throw exception; > > there is no way to tell the producers on these protocols that the message > failed. Keep coming would endup on a drop and the producer may thing it was > sent ok. It would be a message loss without communicating the failure to > the producer. > > Is it possible to configure it to go from PAGING to DROP once the disk > limit is reached? > you can already configure global-max-size to DROP. > I *strongly* advise against silently fail. (I have a strong opinion on > this subject). I won't silently fail a message send on purpose here. > — > You are receiving this because you were mentioned. > Reply to this email directly, view it on GitHub > <https://github.com/apache/activemq-artemis/pull/749/files/f84d9d105901253d079ec8dd42b888f15f356354#r77524662>, > or mute the thread > <https://github.com/notifications/unsubscribe-auth/AAOCy2c6jvpra-IFJzF4wUJTfzmGqJqhks5qnB3PgaJpZM4J0MjB> > . >	1	1	https://issues.apache.org/jira/browse/ARTEMIS-581
18769	ARTEMIS-579	Looking at the code, it seems that the functionality is present (via {{reload=true}}) but not documented. Please document it. Also, {{TextFileCertificateLoginModule.java}} contains: {code} * The user and role files' locations must be specified in the * org.apache.activemq.jaas.textfiledn.user and * org.apache.activemq.jaas.textfiledn.role properties respectively. NOTE: This * class will re-read user and group files for every authentication (i.e it does * live updates of allowed roles and users). {code} Could you please confirm that Artemis will *not* re-read user and group files for every authentication but only when the files change? In this case, the comment in the code should be changed...	0	0	https://issues.apache.org/jira/browse/ARTEMIS-579
18770	ARTEMIS-573	The existing documentation contains "Any other identifiers used in core filter expressions will be assumed to be properties of the message". However, the documentation does not specify what an identifier is. Can an identifier contain a dash (like {{foo-bar}})? Note: in STOMP, it is fairly common to use dashes in header names and, from memory, Apollo does support identifiers with dashes.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-573
18771	ARTEMIS-573	Looks like the JMS client enforces the {{Character.isJavaLetterOrDigit}} check, but there is nothing prohibiting the use of a '-' character in a core filter. Can you provide the test that is failing for you?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-573
18772	ARTEMIS-560	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/962#discussion_r96415467 --- Diff: artemis-jdbc-store/src/main/java/org/apache/activemq/artemis/jdbc/store/file/JDBCSequentialFile.java --- @@ -88,7 +88,12 @@ public boolean isOpen() { @Override public boolean exists() { - return isCreated; + if (isCreated) return true; + try { + return fileFactory.listFiles(extension).contains(filename); + } catch (Exception e) { --- End diff -- This may bite us, can you at least log an error here?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-560
18773	ARTEMIS-550	Github user mattrpav commented on the issue: https://github.com/apache/activemq-artemis/pull/875 Has there been mention of the transactedSend and concurrentSend options? These are super useful for improving reliability and scalability when using Virtual Topics.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-550
18774	ARTEMIS-550	Artemis has another concept known as "diverts." With diverts, you can effectively subscribe a durable queue to a topic, so that the queue receives all messages that were sent to the topic and holds them until its consumer is online. My understanding is that, in practice, this provides the same functionality as (and possibly more flexibility than) virtual topics. Could your use case be covered by diverts instead? https://activemq.apache.org/artemis/docs/1.4.0/diverts.html	0	0	https://issues.apache.org/jira/browse/ARTEMIS-550
18775	ARTEMIS-550	Github user mattrpav commented on the issue: https://github.com/apache/activemq-artemis/pull/875 @graben re: Transacted: if the PostOffice is transacted by default, that's great. The option in 5.x felt more like a "shore-up" to get to transacted. I don't see a compelling reason to support non-transacted broker-side message movement. A note in the migration guide would suffice. re: Concurrency: I see this having value, and agree it should be a separate JIRA/PR. re: Selector/Filter: that's great that you have it on your radar. You mentioned filter, so I feel obligated to clarify-- this selector is determined by the _consumer_ using the standard JMS selector available in the API. The broker then store these to properly filter messages if the consumer goes offline. Does Artemis' filter support this behavior, or is filtering broker-side configuration? Thanks!	0	0	https://issues.apache.org/jira/browse/ARTEMIS-550
18776	ARTEMIS-550	Github user andytaylor commented on the issue: https://github.com/apache/activemq-artemis/pull/875 @graben Could you explain why virtual topics need to be added. As far as I was aware in ActiveMQ its just to provide shared consumers and queue like semantics on a topic which Artemis does by default?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-550
18777	ARTEMIS-550	Github user rlodge commented on the issue: https://github.com/apache/activemq-artemis/pull/875 @andytaylor How does Artemis provide the same functionality by default (without configuration for each "virtual topic / queue" being added to the broker's configuration)? This is one of the only (maybe the only) feature keeping us from attempting to migrate to Artemis from ActiveMQ. I've looked through the documentation and haven't been able to find a similar feature other than diverts, which seem to need broker configuration and broker restart, which would be a dealbreaker, since our environment requires clients to to be able to dynamically create the queue (not just queue-like) destinations on the fly without broker reconfiguration or restart.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-550

#	Requirement	Questions	Answered	Successfully Answered	URL
18778	ARTEMIS-550	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/875 @rlodge you can have as many consumers you like on a QueueSubscription with Artemis. You can limite the client-id if you like, but you can also bypass that. The only use case I can see for artemis with Virtual topics is migration case... users migrating to artemis without changing code.. but I want to make sure I understand the requirements. What is needed? Route the virtual queue to another existing core queue? if that's the case it's pretty simple implementation. (I guess you would have most of what is needed on this PR)?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-550
18779	ARTEMIS-550	Github user andytaylor commented on the issue: https://github.com/apache/activemq-artemis/pull/875 @rlodge so I think all that is doable without having to add the notion of a virtual topic. With the addressing changes this should be even easier to add to the Openwire protocol manager. one thing you could explain is does "we want to be able to broadcast a message to an unknown number of different subscribers" mean to different addresses or just one address like myTopic?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-550
18780	ARTEMIS-550	Github user rlodge commented on the issue: https://github.com/apache/activemq-artemis/pull/875 @mtaylor What I'm hearing is that basically, following changes in 780, a client would be able to send to an address "RecordUpdateNotifications", other clients could subscribe to "RecordUpdateNotifications::SystemARecordUpdateNotifications" and it would act essentially like virtual topics do now; all "RecordUpdateNotifications::*" subscriptions would read from queues coming from the single address? If so, do OpenWire clients need to use a specific queue / address naming scheme, or would it happen automatically somehow (and could it be mapped from the VirtualTopic.* naming scheme)? It's not a terribly large problem for us to go through and change our queues from VirtualTopic.* to a slightly different naming scheme, though we'd rather change nothing about the clients; it would be more difficult to somehow change the API's we were calling...	0	0	https://issues.apache.org/jira/browse/ARTEMIS-550
18781	ARTEMIS-550	Github user andytaylor commented on the issue: https://github.com/apache/activemq-artemis/pull/875 @graben It sounds different, what is it a pool of, connections?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-550
18782	ARTEMIS-550	[~graben] Are you saying that it is not possible to implement these things using the new address model? I agree that the existing naming conventions i.e. using "::" is different to ActiveMQ 5.x, and there are gaps wrt how things are configured. For example using prefixing. The JIRA I added above is intended to track gap in the naming schema, like for example using prefixes instead of FQQN. What I don't want to do, is just clone the way ActiveMQ 5.x does virtual topics, as it doesn't fit in well with the Artemis architecture. Instead, I plan on providing generic solutions for things like address aliasing that solve part of the problem, whilst being an overall useful feature for other use cases.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-550
18783	ARTEMIS-548	Github user jbertram commented on the issue: https://github.com/apache/activemq-artemis/pull/653 There's a handful of Stomp subscriber tests running without issues. Can you clarify a bit on what is broken. If you can outline a specific use-case then I can write a test for it.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-548
18784	ARTEMIS-529	Is there any chance you'd be willing to send a PR for this?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-529
18785	ARTEMIS-529	Is there any chance you'd be willing to learn more about Artemis' internals in order to send a PR for this?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-529
18786	ARTEMIS-529	[~jbertram] Hi Justin, do you know where we are with this issue? I was of the understanding that all ActiveMQ 5.x JAAS modules were available in Artemis? If not are there any other outstanding?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-529
18787	ARTEMIS-517	Are you wanting to be able to execute this on the backup server or on the live server or both? {{org.apache.activemq.artemis.core.server.impl.SharedNothingBackupActivation#isRemoteBackupUpToDate}} already exists for the backup. Of course, the traditional management mechanisms won't be available until the backup actually becomes live. For the live server, {{org.apache.activemq.artemis.core.server.impl.Activation#getReplicationManager}} exists already which (if not null) can be used to invoke {{org.apache.activemq.artemis.core.replication.ReplicationManager#isSynchronizing()}}. If that returns {{false}} then you know the live and backup have synced up.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-517
18788	ARTEMIS-517	SharedNothingBackupActivation & Activation are in impl packages and are not considered public API. Besides I find that clunky to have to look at 3 different objects to get this information. Could this be exposed by a single method on the ActiveMQServer interface instead?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-517
18789	ARTEMIS-517	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/513#discussion_r62853199 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/ActiveMQServerImpl.java --- @@ -674,6 +674,26 @@ public void setJMSQueueDeleter(QueueDeleter jmsQueueDeleter) { this.jmsQueueDeleter = jmsQueueDeleter; } + @Override + public boolean isInitialReplicationSynchronizationComplete() { --- End diff -- suggestion: wouldn't isReplicaSync be a better name?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-517
18790	ARTEMIS-508	Checked! I'm not sure i prefer the vanilla's single Thread Executor as first-class choice to reduce the contention/need of synchronization of the journal; for few reasons: 1) it use internally a LinkedBlockingQueue: - it does not support in the right way batch offering due to the lock on the submit of new elements on it, implying a sort "OS queue effect" that lead to impredicible latencies, thus it will make even less predictable what is trying to cure - it creates new nodes without recycling in any way,this could lead to a lot a bad effects on latencies and on the overall application behaviour due to the pressure on the garbage collector (eg Full GC pauses + http://psy-lob-saw.blogspot.it/2016/03/gc-nepotism-and-linked-queues.html) - the on-heap Java objects it uses are sensible to Card Marking issues and pointer chasing issues (less cache friendly!) 2) creates a lot of garbage submitting new Runnable for each write request (+ new Future<T>,FutureTask etcetera) 3) it doesn't order properly the write requests (cannot be made fair on the requests scheduling without pain) Thus i see only two opportunities of improvements: 1) single threaded (like my current implementation): reduce features but make it fast as hell (single writer principle is ever a good idea IMHO and less branches let it more predictable and faster) 2) multi-threaded: using an approach similar to this https://groups.google.com/forum/#!topic/mechanical-sympathy/5jMIVNf7zXA or with the single threaded version plus a lock-free ringBuffer that feed an event loop (that could use the single-threaded version of the SequenceFile) What do you think about it?Do you see other possibilities?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-508
18791	ARTEMIS-508	I'm not sure To be honest... The synchronization we have as I remember is just to enqueue on the writes on the Sequential Files.. Look at the TimedBuffer... A lot of these are just waiting the callback through the operation context, it's based on notifications. But I'm definitely open for improvements. Why don't we start with including your MappedSquentialFile in the codebae and then look for possibilities for improvements? We could even do a google hangout or any sort of video meeting to look over where to make improvements. We could maybe do that after mid next week?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-508

#	Requirement	Questions	Answered	Successfully Answered	URL
18792	ARTEMIS-508	I'm very happy if the mapped journal could became part of the codebase!!! Perfect with Hangout! Do you have any preference about day and/or time of day? In this way i could ask you directly details to understand better the context or the assumpions of the usage of the Journal....	1	1	https://issues.apache.org/jira/browse/ARTEMIS-508
18793	ARTEMIS-487	Can you explain more.... Do I have to write a new method or new class...	1	0	https://issues.apache.org/jira/browse/ARTEMIS-487
18794	ARTEMIS-473	I've created new jira for the description - ARTEMIS-679 - Activate most up to date server from master-slave(live-backup) pair. If split brain happens then there is not much Artemis can do about it. Still it can recover from quite common cases. Basically 3 situation can happen when split brain happens (=master and slave are active at the same time): a) Clients do not loose connection to master and stay connected to master. b) Clients loose connection to master and failover to backup. c) Clients loose connection to master and slave at same time. They're trying to reconnect to master-slave pair. I believe that for situations a) and b) Artemis can recover when network is reconnected. In the moment when master and slave notice that they're active at the same time, they will check who has external (no in-vm) connections. Server without external client connections will restart. Only server with the clients has the up-to-date journal. Option c) is problematic as clients can connect to master or slave so in this case there is nothing Artemis can do. wdyt?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-473
18795	ARTEMIS-444	Github user andytaylor commented on the pull request: https://github.com/apache/activemq-artemis/pull/425#issuecomment-199269101 Can you improve this so the log message makes sense and to actually use the method. so something like 'LIBAIO is installed but /path/somewhere does not support AIO'.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-444
18796	ARTEMIS-444	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/425#discussion_r56841836 --- Diff: artemis-journal/src/main/java/org/apache/activemq/artemis/core/io/aio/AIOSequentialFileFactory.java --- @@ -114,6 +116,36 @@ public static boolean isSupported() { return LibaioContext.isLoaded(); } + public static boolean isSupported(File journalPath) { + if (!isSupported()) { + return false; + } + + File aioTestFile = new File(journalPath, AIO_TEST_FILE); + try { + int fd = LibaioContext.open(aioTestFile.getAbsolutePath(), true); + LibaioContext.close(fd); + aioTestFile.delete(); + } catch (Exception e) { + // try to handle the file using plain Java + // return false if and only if we can create/remove the file using + // plain Java but not using AIO + try { + if (!aioTestFile.exists()) { + if (!aioTestFile.createNewFile()) return true; --- End diff -- Isn't that a bit convoluted?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-444
18797	ARTEMIS-439	Can this be resolved? It appears to be done.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-439
18798	ARTEMIS-418	Based on my perusal of the ActiveMQ 5.x code it looks like MinMessageSize, AverageMessageSize, and MaxMessageSize are based on all the messages that have been sent to the queue rather than the messages that are currently in the queue. Is that correct? I couldn't find any documentation specifically explaining these metrics.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-418
18799	ARTEMIS-418	[~jbertram] This is resolved as incomplete? Is it partially fixed?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-418
18800	ARTEMIS-413	I'm not terribly familiar with the STOMP implementation, but I see several tests in {{org.apache.activemq.artemis.tests.integration.stomp.StompTest}} related to durable topic subscriptions (e.g. {{testDurableSubscriberWithReconnection()}}) which suggests it is supported. Can you give that a look?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-413
18801	ARTEMIS-409	What does the body of the response say in this situation?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-409
18802	ARTEMIS-405	Also, some attributes are documented as {{N/A}}. For instance we have in the Core: {code} org.apache.activemq.artemis:address=*,brokerName=*,module=Core,name=*,serviceType=Queue,type=Broker: ========== Attributes: Address java.lang.String [ro], "address this queue is bound to" ConsumerCount int [ro], "number of consumers consuming messages from this queue" DeadLetterAddress java.lang.String [ro], "dead-letter address associated with this queue" DeliveringCount int [ro], "number of messages that this queue is currently delivering to its consumers" Durable boolean [ro], "whether this queue is durable" ExpiryAddress java.lang.String [ro], "expiry address associated with this queue" Filter java.lang.String [ro], "filter associated with this queue" FirstMessageAge java.lang.Long [ro], "age of the first message in milliseconds" FirstMessageAsJSON java.lang.String [ro], "first message on the queue as JSON" FirstMessageTimestamp java.lang.Long [ro], "timestamp of the first message in milliseconds" ID long [ro], "ID of this queue" MessageCount long [ro], "number of messages currently in this queue (includes scheduled, paged, and in-delivery messages)" MessagesAcknowledged long [ro], "number of messages acknowledged from this queue since it was created" MessagesAdded long [ro], "number of messages added to this queue since it was created" Name java.lang.String [ro], "name of this queue" Paused boolean [ro], "whether the queue is paused" ScheduledCount long [ro], "number of scheduled messages in this queue" Temporary boolean [ro], "whether this queue is temporary" {code} But in JMS: {code} org.apache.activemq.artemis:brokerName=*,module=JMS,name=*,serviceType=Queue,type=Broker: ========== Attributes: Address java.lang.String [ro], "N/A" ConsumerCount int [ro], "number of consumers consuming messages from this queue" DeadLetterAddress java.lang.String [ro], "dead-letter address associated with this queue" DeliveringCount int [ro], "N/A" ExpiryAddress java.lang.String [ro], "expiry address associated with this queue" FirstMessageAge java.lang.Long [ro], "age of the first message in milliseconds" FirstMessageAsJSON java.lang.String [ro], "first message on the queue as JSON" FirstMessageTimestamp java.lang.Long [ro], "timestamp of the first message in milliseconds" MessageCount long [ro], "N/A" MessagesAdded long [ro], "N/A" Name java.lang.String [ro], "N/A" Paused boolean [ro], "N/A" RegistryBindings [Ljava.lang.String; [ro], "N/A" ScheduledCount long [ro], "number of scheduled messages in this queue" Selector java.lang.String [ro], "selector for the queue" Temporary boolean [ro], "N/A" {code} This gives the impressions that the {{MessageCount}} is not used in JMS but only in the Core. However, in my tests {{MessageCount}} seems to be correct in JMS too. So, what does {{N/A}} mean here?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-405

#	Requirement	Questions	Answered	Successfully Answered	URL
18803	ARTEMIS-374	Github user jbertram commented on the pull request: https://github.com/apache/activemq-artemis/pull/362#issuecomment-178012905 I need some way to differentiate messages that are put on the head of the queue by the scheduler and those that aren't because they need to be treated differently (notice the new 'if' block in addHead()). A message property seemed the simplest way to do that, but I'm open to other possibilities. Do you have any suggestions here?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-374
18804	ARTEMIS-374	Github user jbertram commented on the pull request: https://github.com/apache/activemq-artemis/pull/362#issuecomment-178061415 I'm not following. From what I can tell in addHead() (where I need to do the check) Rreference.getScheduledDeliveryTime() will always return 0 so there's no way to differentiate scheduled messages from non-scheduled messages by using this property. Can you clarify your suggestion?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-374
18805	ARTEMIS-373	[~martyntaylor] I'd like to know more details about this improvement. Is it still valid?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-373
18806	ARTEMIS-373	Is the java test related here is JDBCJournalTest.java ?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-373
18807	ARTEMIS-373	[~martyntaylor] At first, when I run JDBCJournalTest.java, testCallBacks and testReadJournal passed whereas testInsertRecords failed. Then, I remove the Thread.sleep() line in testInsertRecords and after that it also passed. So, Is that it or what to do next? Thanks.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-373
18808	ARTEMIS-355	As I understand the documentation a single JMS bridge is closely tied to a single source/target destination pair. Just to make sure I understand you correctly. Are JMS bridges really a viable temporary solution if you need seamless duplex connectivity for all JMS destinations across all nodes of a broker network?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-355
18809	ARTEMIS-347	This looks complete. Can you resolve it, [~clebertsuconic]?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-347
18810	ARTEMIS-339	Is this done?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-339
18811	ARTEMIS-322	Hi Clebert, I'm sorry about that, I'm new one day on this. I will try to explain better, I'm trying to implement and Junit test as a method above. When I'm using the ActiveMQ Artemis core lib, I was debugging the code because every time I'm getting the error code: {panel:title=My Exception|borderStyle=dashed|borderColor=#ccc|titleBGColor=#F7D6C1|bgColor=#FFFFCE} javax.jms.JMSException: There is no topic with name topic at org.apache.activemq.artemis.jms.client.ActiveMQSession.createTopic(ActiveMQSession.java:373) at com.redhat.messaging.topicagents.Publisher.create(Publisher.java:52) at com.redhat.messaging.ArtemisRedHatTopicTestCase.testSendTextMessageToDestination(ArtemisRedHatTopicTestCase.java:29) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at junit.framework.TestCase.runTest(TestCase.java:176) at junit.framework.TestCase.runBare(TestCase.java:141) at junit.framework.TestResult$1.protect(TestResult.java:122) at junit.framework.TestResult.runProtected(TestResult.java:142) at junit.framework.TestResult.run(TestResult.java:125) at junit.framework.TestCase.run(TestCase.java:129) at junit.framework.TestSuite.runTest(TestSuite.java:252) at junit.framework.TestSuite.run(TestSuite.java:247) at org.junit.internal.runners.JUnit38ClassRunner.run(JUnit38ClassRunner.java:86) at org.eclipse.jdt.internal.junit4.runner.JUnit4TestReference.run(JUnit4TestReference.java:50) at org.eclipse.jdt.internal.junit.runner.TestExecution.run(TestExecution.java:38) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.runTests(RemoteTestRunner.java:459) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.runTests(RemoteTestRunner.java:675) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.run(RemoteTestRunner.java:382) at org.eclipse.jdt.internal.junit.runner.RemoteTestRunner.main(RemoteTestRunner.java:192) {panel} As you can see, I created a topic called topic on my code bellow: {code:borderStyle=solid} // create a Session session = connection.createSession(false, Session.AUTO_ACKNOWLEDGE); // create the Topic to which messages will be sent Topic topic = session.createTopic("topic"); {code} I debbuged and I would like a help to fix it, on package org.apache.activemq.artemis.core.protocol.core.impl.ChannelImpl who package org.apache.activemq.artemis.core.protocol.core.impl.ActiveMQSessionContext on line 243 has a piece of code {code} public ClientSession.AddressQuery addressQuery(final SimpleString address) throws ActiveMQException { SessionBindingQueryResponseMessage_V2 response = (SessionBindingQueryResponseMessage_V2) sessionChannel.sendBlocking(new SessionBindingQueryMessage(address), PacketImpl.SESS_BINDINGQUERY_RESP_V2); return new AddressQueryImpl(response.isExists(), response.getQueueNames(), response.isAutoCreateJmsQueues()); } {code} this response value is not comming with values, and I understand should return that my topic exists. Because I'd created a temporary topic as I told you before.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-322

#	Requirement	Questions	Answered	Successfully Answered	URL
18812	ARTEMIS-322	I appreciate the willingness to ensure this is a configurable option. I am curious what the reasonable Use Case is for _enforcing_ Topic's to be created at an Administrative level? Or a better question is, what is the Use Case for a product *failing* if a Topic has not been created via the Administrative Console? This very question had resulted in a lengthly debate back in the day when JBoss MQ was just getting started. The specification states that the creation, deletion and management of Topic's are out of the JMS scope. The API and specification provides for a TemporaryTopic for 'reply scenarios', however, that topic can only be used by the TopicConnection which created it. Meaning any sort of TopicConnection pool becomes untenable. There were also issues with Selectors so that wasn't usable either. Then we had to deal with the Use Case of creating unique, per-client, durable queue's on an as-needed basis. For example, a client is launched which receives a live stream of events. That client is meant to be long running and always up to date. The idea would be to create a Durable Subscriber on a brand new Unique Queue which receives the events for that Client. The client could then specifiy any filtering criteria when registering for the events. Due to Client's having access to different sets of Events based on ACL's, or potentially being slightly off from one another it isn't possible to have a single _known_ queue ahead of time. And again, TempTopic's cannot be durable and are stuck to that TopicConnection. The client crashes, gets disconnected, or uses a connection pool and TempTopic's cannot be used. In today's age of Messaging systems all these problems aren't really issues anymore. Routes and Exchanges in AMQP solve so many problem. Modern JMS implementations do very well with JMSReplyTo, Selector's and allowing on-demand queue's/topic's to be created. I bring all this up because at the end of the day it literally took 3 lines of code to change JBoss MQ (which i did on my own fork of the code) to allow for the creation of dynamic queues. And the merge requests were denied as the specification didn't account for the administration of queues and topics. The bigger issue was that JBoss MQ was ALSO lacking any sort of reasonable API and CLI interface to administer the queues. So to make an application work a person had to log into the administrator console and add every single Topic and Queue along with any permissions, if needed, before the application would work. Again, thank you for taking the time to take into consideration the idea that there are valid use cases for wanting to dynamically create topics and queues. Making it configurable provides the best of both. JBoss MQ (aka JBoss Messaging) was the ONLY Broker to enforce this restriction making it all the more incredulous. Now that we have a protocol AND specification of AMQP, that seems to be the way forward. It handles exponentially more scenarios then the JMS API. Thank you all for the work you do.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-322
18813	ARTEMIS-252	What's the difference between that and move? How retry would work with topics? the move operation would put it back to all the subscriptions.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-252
18814	ARTEMIS-252	I like the idea yeah. I just wanted to clarify what to do on the topic case. If you resend back to the address all subscriptions will get the message. Do we need to send back to the single subscription where it came from? If u send a PR I can complete with the topic case if needed. There is hack that could be used to limit a core internal queue or topic subscriptions	1	1	https://issues.apache.org/jira/browse/ARTEMIS-252
18815	ARTEMIS-252	I guess question is what would be the semantic on activemq 5? Would it deliver back to the failed subscription only?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-252
18816	ARTEMIS-245	That's all that's needed? Why don't you send a Pull Request instead of a patch?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-245
18817	ARTEMIS-245	Is this patch sufficient to get the RA working with TomEE?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-245
18818	ARTEMIS-245	that's all it's needed? why don't you send a Pull request at https://github.com/apache/activemq-artemis/pulls	0	0	https://issues.apache.org/jira/browse/ARTEMIS-245
18819	ARTEMIS-245	The JavaDoc for javax.resource.spi.DissociatableManagedConnection (which you've implemented in your patch) says, "An implementation of this interface must support the lazy connection association optimization." It's not clear to me that the Artemis RA supports this optimization. Can you confirm that it does? Otherwise this change wouldn't be appropriate.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-245
18820	ARTEMIS-235	This issue is unresolved and marked as to be fixed for 1.1.1. However, 1.2.0 has already been released... Could it be targeted for 1.3.0 instead?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-235
18821	ARTEMIS-233	tests/integration-tests/src/test/java/org/apache/activemq/artemis/tests/integration/mqtt/imported/MQTTTest.java --- @@ -1001,91 +1005,93 @@ public void testCleanSession() throws Exception { outstanding task to add cross protocol support. This task should rework these tests. The tests are included here and commented out to ensure ActiveMQ and Artemis tests are in sync. */ - // @Test(timeout = 60 * 1000) - // public void testSendMQTTReceiveJMS() throws Exception { - // doTestSendMQTTReceiveJMS("foo.*"); - // } + @Test(timeout = 60 * 1000) + public void testSendMQTTReceiveJMS() throws Exception { + doTestSendMQTTReceiveJMS("foo.*", "foo/bar"); + } - // public void doTestSendMQTTReceiveJMS(String destinationName) throws Exception { - // final MQTTClientProvider provider = getMQTTClientProvider(); - // initializeConnection(provider); - // - // // send retained message - // final String RETAINED = "RETAINED"; - // provider.publish("foo/bah", RETAINED.getBytes(), AT_LEAST_ONCE, true); - // - // ActiveMQConnection activeMQConnection = (ActiveMQConnection) cf.createConnection(); - // // MUST set to true to receive retained messages - // activeMQConnection.setUseRetroactiveConsumer(true); - // activeMQConnection.start(); - // Session s = activeMQConnection.createSession(false, Session.AUTO_ACKNOWLEDGE); - // javax.jms.Topic jmsTopic = s.createTopic(destinationName); - // MessageConsumer consumer = s.createConsumer(jmsTopic); - // - // // check whether we received retained message on JMS subscribe - // ActiveMQMessage message = (ActiveMQMessage) consumer.receive(5000); - // assertNotNull("Should get retained message", message); - // ByteSequence bs = message.getContent(); - // assertEquals(RETAINED, new String(bs.data, bs.offset, bs.length)); - // assertTrue(message.getBooleanProperty(RetainedMessageSubscriptionRecoveryPolicy.RETAINED_PROPERTY)); - // - // for (int i = 0; i < NUM_MESSAGES; i++) { - // String payload = "Test Message: " + i; - // provider.publish("foo/bah", payload.getBytes(), AT_LEAST_ONCE); - // message = (ActiveMQMessage) consumer.receive(5000); - // assertNotNull("Should get a message", message); - // bs = message.getContent(); - // assertEquals(payload, new String(bs.data, bs.offset, bs.length)); - // } - // - // activeMQConnection.close(); - // provider.disconnect(); - // } + public void doTestSendMQTTReceiveJMS(String destinationName) throws Exception { + final MQTTClientProvider provider = getMQTTClientProvider(); + initializeConnection(provider); + // TODO As with other tests, this should be enabled as part of the cross protocol support with MQTT. - // @Test(timeout = 2 * 60 * 1000) - // public void testSendJMSReceiveMQTT() throws Exception { - // doTestSendJMSReceiveMQTT("foo.far"); - // } + // send retained message + final String address = "jms/queue/" + mqttAddress; + final String RETAINED = "RETAINED"; + // TODO As with other tests, this should be enabled as part of the cross protocol support with MQTT. - // public void doTestSendJMSReceiveMQTT(String destinationName) throws Exception { - // final MQTTClientProvider provider = getMQTTClientProvider(); - // initializeConnection(provider); - // - // ActiveMQConnection activeMQConnection = (ActiveMQConnection) cf.createConnection(); - // activeMQConnection.setUseRetroactiveConsumer(true); - // activeMQConnection.start(); - // Session s = activeMQConnection.createSession(false, Session.AUTO_ACKNOWLEDGE); - // javax.jms.Topic jmsTopic = s.createTopic(destinationName); - // MessageProducer producer = s.createProducer(jmsTopic); - // - // // send retained message from JMS - // final String RETAINED = "RETAINED"; - // TextMessage sendMessage = s.createTextMessage(RETAINED); - // // mark the message to be retained - //	1	1	https://issues.apache.org/jira/browse/ARTEMIS-233

#	Requirement	Questions	Answered	Successfully Answered	URL
18822	ARTEMIS-233	tests/integration-tests/src/test/java/org/apache/activemq/artemis/tests/integration/mqtt/imported/MQTTTest.java --- @@ -1001,91 +1005,93 @@ public void testCleanSession() throws Exception { outstanding task to add cross protocol support. This task should rework these tests. The tests are included here and commented out to ensure ActiveMQ and Artemis tests are in sync. */ - // @Test(timeout = 60 * 1000) - // public void testSendMQTTReceiveJMS() throws Exception { - // doTestSendMQTTReceiveJMS("foo.*"); - // } + @Test(timeout = 60 * 1000) + public void testSendMQTTReceiveJMS() throws Exception { + doTestSendMQTTReceiveJMS("foo.*", "foo/bar"); + } - // public void doTestSendMQTTReceiveJMS(String destinationName) throws Exception { - // final MQTTClientProvider provider = getMQTTClientProvider(); - // initializeConnection(provider); - // - // send retained message - // final String RETAINED = "RETAINED"; - // provider.publish("foo/bah", RETAINED.getBytes(), AT_LEAST_ONCE, true); - // - // ActiveMQConnection activeMQConnection = (ActiveMQConnection) cf.createConnection(); - // // MUST set to true to receive retained messages - // activeMQConnection.setUseRetroactiveConsumer(true); - // activeMQConnection.start(); - // Session s = activeMQConnection.createSession(false, Session.AUTO_ACKNOWLEDGE); - // javax.jms.Topic jmsTopic = s.createTopic(destinationName); - // MessageConsumer consumer = s.createConsumer(jmsTopic); - // - // // check whether we received retained message on JMS subscribe - // ActiveMQMessage message = (ActiveMQMessage) consumer.receive(5000); - // assertNotNull("Should get retained message", message); - // ByteSequence bs = message.getContent(); - // assertEquals(RETAINED, new String(bs.data, bs.offset, bs.length)); - // assertTrue(message.getBooleanProperty(RetainedMessageSubscriptionRecoveryPolicy.RETAINED_PROPERTY)); - // - // for (int i = 0; i < NUM_MESSAGES; i++) { - // String payload = "Test Message: " + i; - // provider.publish("foo/bah", payload.getBytes(), AT_LEAST_ONCE); - // message = (ActiveMQMessage) consumer.receive(5000); - // assertNotNull("Should get a message", message); - // bs = message.getContent(); - // assertEquals(payload, new String(bs.data, bs.offset, bs.length)); - // } - // - // activeMQConnection.close(); - // provider.disconnect(); - // } + public void doTestSendMQTTReceiveJMS(String destinationName) throws Exception { + final MQTTClientProvider provider = getMQTTClientProvider(); + initializeConnection(provider); + // TODO As with other tests, this should be enabled as part of the cross protocol support with MQTT. + // @Test(timeout = 2 * 60 * 1000) + // public void testSendJMSReceiveMQTT() throws Exception { + // doTestSendJMSReceiveMQTT("foo.far"); + // } + // send retained message + final String address = "jms/queue/" + mqttAddress; + final String RETAINED = "RETAINED"; + // TODO As with other tests, this should be enabled as part of the cross protocol support with MQTT. + // public void doTestSendJMSReceiveMQTT(String destinationName) throws Exception { + // final MQTTClientProvider provider = getMQTTClientProvider(); + // initializeConnection(provider); + // + // ActiveMQConnection activeMQConnection = (ActiveMQConnection) cf.createConnection(); + // activeMQConnection.setUseRetroactiveConsumer(true); + // activeMQConnection.start(); + // Session s = activeMQConnection.createSession(false, Session.AUTO_ACKNOWLEDGE); + // javax.jms.Topic jmsTopic = s.createTopic(destinationName); + // MessageProducer producer = s.createProducer(jmsTopic); + // + // // send retained message from JMS + // final String RETAINED = "RETAINED"; + // TextMessage sendMessage = s.createTextMessage(RETAINED); + // // mark the message to be retained	1	1	https://issues.apache.org/jira/browse/ARTEMIS-233
18823	ARTEMIS-233	tests/integration-tests/src/test/java/org/apache/activemq/artemis/tests/integration/mqtt/imported/MQTTTest.java --- @@ -1001,91 +1005,93 @@ public void testCleanSession() throws Exception { outstanding task to add cross protocol support. This task should rework these tests. The tests are included here and commented out to ensure ActiveMQ and Artemis tests are in sync. */ - // @Test(timeout = 60 * 1000) - // public void testSendMQTTReceiveJMS() throws Exception { - // doTestSendMQTTReceiveJMS("foo.*"); - // } + @Test(timeout = 60 * 1000) + public void testSendMQTTReceiveJMS() throws Exception { + doTestSendMQTTReceiveJMS("foo.*", "foo/bar"); + } - // public void doTestSendMQTTReceiveJMS(String destinationName) throws Exception { - // final MQTTClientProvider provider = getMQTTClientProvider(); - // initializeConnection(provider); - // - // send retained message - // final String RETAINED = "RETAINED"; - // provider.publish("foo/bah", RETAINED.getBytes(), AT_LEAST_ONCE, true); - // - // ActiveMQConnection activeMQConnection = (ActiveMQConnection) cf.createConnection(); - // // MUST set to true to receive retained messages - // activeMQConnection.setUseRetroactiveConsumer(true); - // activeMQConnection.start(); - // Session s = activeMQConnection.createSession(false, Session.AUTO_ACKNOWLEDGE); - // javax.jms.Topic jmsTopic = s.createTopic(destinationName); - // MessageConsumer consumer = s.createConsumer(jmsTopic); - // - // // check whether we received retained message on JMS subscribe - // ActiveMQMessage message = (ActiveMQMessage) consumer.receive(5000); - // assertNotNull("Should get retained message", message); - // ByteSequence bs = message.getContent(); - // assertEquals(RETAINED, new String(bs.data, bs.offset, bs.length)); - // assertTrue(message.getBooleanProperty(RetainedMessageSubscriptionRecoveryPolicy.RETAINED_PROPERTY)); - // - // for (int i = 0; i < NUM_MESSAGES; i++) { - // String payload = "Test Message: " + i; - // provider.publish("foo/bah", payload.getBytes(), AT_LEAST_ONCE); - // message = (ActiveMQMessage) consumer.receive(5000); - // assertNotNull("Should get a message", message); - // bs = message.getContent(); - // assertEquals(payload, new String(bs.data, bs.offset, bs.length)); - // } - // - // activeMQConnection.close(); - // provider.disconnect(); - // } + public void doTestSendMQTTReceiveJMS(String destinationName) throws Exception { + final MQTTClientProvider provider = getMQTTClientProvider(); + initializeConnection(provider); + // TODO As with other tests, this should be enabled as part of the cross protocol support with MQTT. + // @Test(timeout = 2 * 60 * 1000) + // public void testSendJMSReceiveMQTT() throws Exception { + // doTestSendJMSReceiveMQTT("foo.far"); + // } + // send retained message + final String address = "jms/queue/" + mqttAddress; + final String RETAINED = "RETAINED"; + // TODO As with other tests, this should be enabled as part of the cross protocol support with MQTT. + // public void doTestSendJMSReceiveMQTT(String destinationName) throws Exception { + // final MQTTClientProvider provider = getMQTTClientProvider(); + // initializeConnection(provider); + // + // ActiveMQConnection activeMQConnection = (ActiveMQConnection) cf.createConnection(); + // activeMQConnection.setUseRetroactiveConsumer(true); + // activeMQConnection.start(); + // Session s = activeMQConnection.createSession(false, Session.AUTO_ACKNOWLEDGE); + // javax.jms.Topic jmsTopic = s.createTopic(destinationName); + // MessageProducer producer = s.createProducer(jmsTopic); + // + // // send retained message from JMS + // final String RETAINED = "RETAINED"; + // TextMessage sendMessage = s.createTextMessage(RETAINED); + // // mark the message to be retained	1	1	https://issues.apache.org/jira/browse/ARTEMIS-233
18824	ARTEMIS-209	Can you test this again with 2.0?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-209
18825	ARTEMIS-209	[~chug] Tim fixed something on master. can you try master?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-209
18826	ARTEMIS-208	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/132#issuecomment-134000025 Let me the one merging this ?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-208

#	Requirement	Questions	Answered	Successfully Answered	URL
18827	ARTEMIS-208	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/132#issuecomment-134293839 This fix will only work on tests.... on a default server, the acceptor will be at 0.0.0.0, and you will be sending tcp://0.0.0.0:port to the client. It's also wrong to the fact that the same protocolManager will server multiple acceptors in case you have multiple ports configured. on that case you will be sending the BrokerInfo for a different port or IP. Do you have the IP the user used to connect anywhere on the connection? So, you would have to create a brokerInfo per request.. instead caching a single one... for that reason this fix is not valid!	1	1	https://issues.apache.org/jira/browse/ARTEMIS-208
18828	ARTEMIS-201	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/245#issuecomment-158999892 @jbertram think the following: An user sets 100 destinations.. with 100M cache... the memory would be able to fit only 10 paged destinations.. The user has monitoring in place... whenever a consumer drops and page increases he will take action and restart the client... that means... he would make sure he wouldn't get to the 100 paged message using monitoring tools... That's a valid scenario.. with your messages users will assume they must have memory available at all times.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-201
18829	ARTEMIS-201	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/245#issuecomment-159134310 Can you at least say "when you have more than X destinations in page mode" on your log message? X = Memory Size / Page MAX-Size... just as an approximation	1	1	https://issues.apache.org/jira/browse/ARTEMIS-201
18830	ARTEMIS-201	I did this per user's request... This can go away with 1.5.0 now, since you can setup global-max-size. We can set this to INFO or change the logs for sure. What do you suggest as a message?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-201
18831	ARTEMIS-200	Github user gaohoward commented on the pull request: https://github.com/apache/activemq-artemis/pull/163#issuecomment-139489885 Hi clebert, it's strange the protonTest failed in this build. I don't think my changes affect that. My local build is all fine. Anyway can you take a look of my new code changes? Thanks	1	1	https://issues.apache.org/jira/browse/ARTEMIS-200
18832	ARTEMIS-193	[~ataylor284] https://github.com/apache/activemq-artemis/blob/master/tests/integration-tests/src/test/java/org/apache/activemq/artemis/tests/integration/openwire/SimpleOpenWireTest.java#L478 It's working, what doesn't work?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-193
18833	ARTEMIS-163	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/98#issuecomment-125996627 What was done as part of this: - The native layer is now a simplified native implementation. In theory much easier to implement. - The native layer won't call any mallocs for io_submit though, get_iocb is a function that will have a cirtcular cache of objects that will be reused on writes. we used to malloc before every write submit - I have moved stuff around to make it easier to understand (for instance separated packages for nio and aio under the journal) - merged AIOCallback and a useless class IOSubmit into a single IOCallback - added a test on create to sync test performance - The callbacks are also reused now through a circular log Everything else stayed the same, the whole testsuite is passing and I have executed extensive tests on this (through manual testing)	1	1	https://issues.apache.org/jira/browse/ARTEMIS-163
18834	ARTEMIS-149	We already have notifications which is the same thing basically. Do we need to implement anything here?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-149
18835	ARTEMIS-149	Is that an open wire thing, or an activemq thing? We can probably refactor the test to use the new API?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-149
18836	ARTEMIS-149	Github user gaohoward commented on the pull request: https://github.com/apache/activemq-artemis/pull/64#issuecomment-119778194 Can you take a look @clebertsuconic ? Thanks.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-149
18837	ARTEMIS-149	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/64#issuecomment-119791408 I need your help with something.. right now I can't include the activemq-tests on the PRs because they are using invalid javadocs.. do you think you could fix that? remove the javadocs?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-149
18838	ARTEMIS-149	Github user gaohoward commented on the pull request: https://github.com/apache/activemq-artemis/pull/64#issuecomment-119795768 do you mean remove the javadocs in every java file in amq5 tests? I can do that.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-149
18839	ARTEMIS-146	Why it can't depend on the JMS stuff? I don't see how open wire would work without JMS.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-146
18840	ARTEMIS-146	Github user gaohoward commented on the pull request: https://github.com/apache/activemq-artemis/pull/55#issuecomment-116439289 I don't think this has to do with any JMS queues and JMX objects. It's only for openwire. Why it is not valid?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-146
18841	ARTEMIS-146	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/55#issuecomment-116444946 First, you are changing semantics of a production of a production server because the test is not adapting... it's wrong. Second, when a JMS Object is created it needs to also create JMS objects for the JMS Created Queue, that should happen through the JMS Manager. You are cutting this directly into a core-queue, it's also wrong. Why can't you fix the test and add the JMS Manager?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-146
18842	ARTEMIS-146	Github user gaohoward commented on the pull request: https://github.com/apache/activemq-artemis/pull/55#issuecomment-116448885 OK I get your point. The reason I didn't want to bring in JMS manager is because openwire is not supposed to depends on JMS manager. It should depends only on the core server. Like jms manager, the openwire should only be a facade layer around core. isn't it?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-146
18843	ARTEMIS-146	Github user gaohoward commented on the pull request: https://github.com/apache/activemq-artemis/pull/55#issuecomment-116913275 Hi Clebert, I updated the branch, can you take a look? Thanks Howard	1	1	https://issues.apache.org/jira/browse/ARTEMIS-146
18844	ARTEMIS-128	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/9#issuecomment-106791273 This PR seems dependent on #8, why don't you cherry-pick this commit into that branch and have a single PR?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-128

#	Requirement	Questions	Answered	Successfully Answered	URL
18845	ARTEMIS-127	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/8#issuecomment-106826140 Howard, your PR doesn't build because of these issues: /work/apache/six-merge/artemis-selector/target/generated-sources/javacc/org/apache/activemq/artemis/selector/hyphenated/SimpleCharStream.java:223: error: [DepAnn] Deprecated item is not annotated with @Deprecated public int getColumn() { ^ (see http://errorprone.info/bugpattern/DepAnn) Did you mean '@Deprecated'? /work/apache/six-merge/artemis-selector/target/generated-sources/javacc/org/apache/activemq/artemis/selector/hyphenated/SimpleCharStream.java:232: error: [DepAnn] Deprecated item is not annotated with @Deprecated public int getLine() { ^ (see http://errorprone.info/bugpattern/DepAnn) Did you mean '@Deprecated'? /work/apache/six-merge/artemis-selector/target/generated-sources/javacc/org/apache/activemq/artemis/selector/strict/SimpleCharStream.java:223: error: [DepAnn] Deprecated item is not annotated with @Deprecated public int getColumn() { ^ (see http://errorprone.info/bugpattern/DepAnn) Did you mean '@Deprecated'? /work/apache/six-merge/artemis-selector/target/generated-sources/javacc/org/apache/activemq/artemis/selector/strict/SimpleCharStream.java:232: error: [DepAnn] Deprecated item is not annotated with @Deprecated public int getLine() { ^ (see http://errorprone.info/bugpattern/DepAnn) Did you mean '@Deprecated'? Note: Some input files use unchecked or unsafe operations.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-127
18846	ARTEMIS-127	Github user gaohoward commented on the pull request: https://github.com/apache/activemq-artemis/pull/8#issuecomment-107281546 The code is only for compiling and running the tests to get the result. after that do a 'mvn clean' will get rid of all those codes.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-127
18847	ARTEMIS-127	Github user gaohoward commented on the pull request: https://github.com/apache/activemq-artemis/pull/8#issuecomment-107915237 Hi Clebert, I got the PR build passed finally. Can you take a look at it? Thanks Howard	1	1	https://issues.apache.org/jira/browse/ARTEMIS-127
18848	ARTEMIS-127	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/8#issuecomment-107951267 @gaohoward I don't understand this... what's under ./activemq5-unit-tests... it seems the whole activemq5 broker implementation. You even have a ProtocolManagerFactory for the OpenWireProtocol there. Why do you need a whole broker implementation for tests? you could just depend on the maven package.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-127
18849	ARTEMIS-127	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/8#issuecomment-107992912 There must be a better way. Why don't you just implement BrokerService properly instead of depending on the whole activemq5?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-127
18850	ARTEMIS-127	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/8#issuecomment-109303356 I'm good either way. The good thing about having the tests committed is it would be easier to maintain them on the long term. Howard had a whole broker copy on the commit here that's why we asked the change. @gaohoward I'm taking a day off today, I can look on this on monday. You can keep working on this. BTW: You had a fix as part of this PR. is it lost or did you squash it? don't forget to update JIRA please.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-127
18851	ARTEMIS-127	Github user thiagokronig commented on the pull request: https://github.com/apache/activemq-artemis/pull/8#issuecomment-110427482 @clebertsuconic please update the parent version pom: <version>1.0.0-SNAPSHOT</version> Should we have some kind of rule in the build to enforce that a child pom always conform with his parent, to at least generate a warning?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-127
18852	ARTEMIS-127	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/8#issuecomment-110828541 I think we should keep Error Prone on the compilation.. unless it starts to bother people. So far I think it's useful. can you send the PR?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-127
18853	ARTEMIS-119	start it for some reason: I built the server, create the server, ran it and got an exception ```sh cd artemis-distribution/target/apache-artemis-1.0.1-SNAPSHOT-bin/apache-artemis-1.0.1-SNAPSHOT/bin ./artemis create /tmp/myserver --user a --password a --allow-anonymous y cd /tmp/myserver/bin ./artemis run ``` and I got this: ``` 23:29:35,747 INFO [org.apache.activemq.artemis.integration.bootstrap] AMQ101005: Using broker configuration: xml:/tmp/lixo/etc/bootstrap.xml 23:29:36,202 INFO [org.apache.activemq.artemis.integration.bootstrap] AMQ101000: Starting ActiveMQ Artemis Server 23:29:36,235 INFO [org.apache.activemq.artemis.core.server] AMQ221000: live Message Broker is starting with configuration Broker Configuration (clustered=false,journalDirectory=../data/journal,bindingsDirectory=../data/bindings,largeMessagesDirectory=../data/large-messages,pagingDirectory=../data/paging) 23:29:36,273 INFO [org.apache.activemq.artemis.core.server] AMQ221013: Using NIO Journal 23:29:36,337 INFO [org.apache.activemq.artemis.core.server] AMQ221043: Protocol module found: [artemis-server]. Adding protocol support for: CORE 23:29:36,343 INFO [org.apache.activemq.artemis.core.server] AMQ221043: Protocol module found: [artemis-amqp-protocol]. Adding protocol support for: AMQP 23:29:36,357 INFO [org.apache.activemq.artemis.core.server] AMQ221043: Protocol module found: [artemis-hornetq-protocol]. Adding protocol support for: HORNETQ 23:29:36,360 INFO [org.apache.activemq.artemis.core.server] AMQ221043: Protocol module found: [artemis-openwire-protocol]. Adding protocol support for: OPENWIRE 23:29:36,489 INFO [org.apache.activemq.artemis.core.server] AMQ221043: Protocol module found: [artemis-stomp-protocol]. Adding protocol support for: STOMP 23:29:36,995 INFO [org.apache.activemq.artemis.core.server] AMQ221003: trying to deploy queue jms.queue.DLQ 23:29:37,064 INFO [org.apache.activemq.artemis.core.server] AMQ221003: trying to deploy queue jms.queue.ExpiryQueue 23:29:37,193 INFO [org.apache.activemq.artemis.core.server] AMQ221020: Started Acceptor at 0.0.0.0:61616 for protocols [CORE,AMQP,HORNETQ,STOMP,OPENWIRE] 23:29:37,197 INFO [org.apache.activemq.artemis.core.server] AMQ221020: Started Acceptor at 0.0.0.0:5445 for protocols [HORNETQ,STOMP] 23:29:37,199 INFO [org.apache.activemq.artemis.core.server] AMQ221020: Started Acceptor at 0.0.0.0:5672 for protocols [AMQP] 23:29:37,202 INFO [org.apache.activemq.artemis.core.server] AMQ221020: Started Acceptor at 0.0.0.0:61613 for protocols [STOMP] 23:29:37,204 INFO [org.apache.activemq.artemis.core.server] AMQ221007: Server is now live 23:29:37,204 INFO [org.apache.activemq.artemis.core.server] AMQ221001: Apache ActiveMQ Artemis Message Broker version 1.0.1-SNAPSHOT [nodeID=ec9eebc2-0a69-11e5-8817-db258e48d4cb] Exception in thread "main" java.lang.NoClassDefFoundError: javax/servlet/http/HttpSessionIdListener at org.eclipse.jetty.server.session.SessionHandler.<clinit>(SessionHandler.java:54) at org.eclipse.jetty.servlet.ServletContextHandler.newSessionHandler(ServletContextHandler.java:241) at org.eclipse.jetty.servlet.ServletContextHandler.getSessionHandler(ServletContextHandler.java:334) at org.eclipse.jetty.servlet.ServletContextHandler.relinkHandlers(ServletContextHandler.java:180) at org.eclipse.jetty.servlet.ServletContextHandler.<init>(ServletContextHandler.java:157) at org.eclipse.jetty.webapp.WebAppContext.<init>(WebAppContext.java:253) at org.eclipse.jetty.webapp.WebAppContext.<init>(WebAppContext.java:201) at org.apache.activemq.artemis.component.WebServerComponent.deployWar(WebServerComponent.java:100) at org.apache.activemq.artemis.component.WebServerComponent.configure(WebServerComponent.java:59) at	1	1	https://issues.apache.org/jira/browse/ARTEMIS-119

#	Requirement	Questions	Answered	Successfully Answered	URL
18854	ARTEMIS-98	Github user HornetQBot commented on the pull request: https://github.com/apache/activemq-6/pull/206#issuecomment-93590710 Can one of the admins verify this patch by saying "ok to test"?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-98
18855	ARTEMIS-97	Github user jbertram commented on the pull request: https://github.com/apache/activemq-6/pull/205#issuecomment-92938159 My understanding is that these properties weren't change due to compatibility with older HornetQ clients. Have you tested these changes with HornetQ clients to ensure there are no regressions?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-97
18856	ARTEMIS-97	Github user HornetQBot commented on the pull request: https://github.com/apache/activemq-6/pull/205#issuecomment-92938168 Can one of the admins verify this patch by saying "ok to test"?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-97
18857	ARTEMIS-97	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-6/pull/205#issuecomment-92983909 We will need to fix the testsuite before being able to merge this one... @jgenender you used your master for this... this PR will probably be open for a few days... Can you make sure you open a new branch if you send another PR non related to this? I will need a few days before I can merge this mainly because of the testsuite.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-97
18858	ARTEMIS-93	Are you interested that I help with the OSGi enabling?	1	0	https://issues.apache.org/jira/browse/ARTEMIS-93
18859	ARTEMIS-93	Github user cschneider commented on the pull request: https://github.com/apache/activemq-artemis/pull/254#issuecomment-160649094 Sorry.. forgot to use the profile again. I will make sure I run it from now on. Btw. The checks only seem to take 10 seconds for the whole build. Wouldn't it make sense to include them in mvn install? In this case I excluded the org.apache.activemq.artemis.cfg from the rat check as I think it does not make sense to add a license to a 3 line cfg. Is this ok?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-93
18860	ARTEMIS-92	[~jmesnil] What do you suggest? We could have an add transformer method on the Registry you mentioned, and whenever a class with that name is found we could just use that one. Any idea?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-92
18861	ARTEMIS-92	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/29#issuecomment-112476725 looks good.. will wait the build to merge it	1	1	https://issues.apache.org/jira/browse/ARTEMIS-92
18862	ARTEMIS-89	Github user mtaylor commented on the pull request: https://github.com/apache/activemq-6/pull/191#issuecomment-82264396 @clebertsuconic This was probably mentioned earlier, but this seems like a considerable change. Can we hold off merging until after the release. Thanks.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-89
18863	ARTEMIS-80	Hi Andy, is the "old" config with acceptors/connectors still present? I'm asking because of compatibility reasons and easy migration. Looks like that both of the ways of configuration will be needed. Thanks, Mirek	0	0	https://issues.apache.org/jira/browse/ARTEMIS-80
18864	ARTEMIS-79	[~daniel.pocock] can you update this with your PR again for this? It seems we lost this over the renames.. but I had the impression this was merged already. Can you please let me know what's the state on this one here since I lost track due to the renames? I believe this one is done!	1	0	https://issues.apache.org/jira/browse/ARTEMIS-79
18865	ARTEMIS-43	Can I get a second set of eyes on this. We may have the wrong license. We seem to have: http://www.apache.org/legal/src-headers.html but from: http://www.apache.org/licenses/LICENSE-2.0.html#apply the expectation is something else. I wonder which is correct? Also, I ran:{code}mvn org.apache.rat:apache-rat-plugin:0.11:check{code} on https://git-wip-us.apache.org/repos/asf/activemq-6.git and there are 811 files that don't have a license. From http://www.apache.org/legal/src-headers.html - all human readable files should have a copyright header but there are exceptions, if the content is too simple to mandate a copyright - http://www.apache.org/legal/src-headers.html#faq-exceptions I think the xml config from the examples and starts scripts from the distro do fall under that category. But all of the documentation does not. The documentation in .md will need a license header me thinks, hence the reopen.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-43
18866	ARTEMIS-43	Isn't this one done, Howard?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-43
18867	ARTEMIS-35	Whatever we do here keep it generic and pluggable	0	0	https://issues.apache.org/jira/browse/ARTEMIS-35
18868	ARTEMIS-19	Github user clebertsuconic commented on the pull request: https://github.com/apache/activemq-artemis/pull/13#issuecomment-108526486 I see that you changed a lot of tests.. but I didn't see any new tests validating the added semantic? did I miss anything there?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-19
18869	ARTEMIS-13	I'm not a fan of auto-delete. I would probably mark them as transient because they're not registered on the server, but I wouldn't necessarily delete them (unless an MDB pointing to one would count as a constant consumer, hence never deleting it). Should it persist between restarts? Hmmm... not sure of activemq's behaviour here previously.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-13
18870	ARTEMIS-13	Is this issue trying to resolve the same issue as https://issues.jboss.org/browse/HORNETQ-302 (Consider allowing creation of destinations via the pure JMS API)? Please disregard my off topic chatter below otherwise. :) If it is, you probably want to look into how ActiveMQ 5.x does it. AFAIK, in AMQ5 the JMS api Session.createQueue() methods does not communicate with the server, so the queues are created only when the first message is received by the broker. This saves a small amount of overhead in the createQueue method call. Nice, IMO! (Yes, you might argue that repetetive calls to createQueue is an antipattern, but it happens anyway.) https://github.com/apache/activemq/blob/trunk/activemq-client/src/main/java/org/apache/activemq/ActiveMQSession.java#L1264 AMQ5 users also expect the queues to be durable and non-temporary, i.e not removed when unsubscribed. I do not recall if you can configure thesese caracteristincs though.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-13
18871	ARTEMIS-13	Ok, I understand the distinction between this and HORNETQ-302. Still, I ask that you consider the rest of my comment as relevant. Especially my point about not making them transient or auto-deleting destinations created this way. Autodeleting destinations, say created in a certain name space, seems like a separate functionality. Generally I do not see why destinations created automatically, once they are created, should be treated any different to destinations created in any other way (administratively, declaratively, using the native AMQ6-api, etc). That's not a distinction that exists in AMQ5. When this issue has been resolved, I would very much like to see that Session.createQueue in the AMQ6 api is a no-op, as far as server communication goes. Should I open a separate issue for that?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-13
18872	ARTEMIS-13	Could you please update the documentation accordingly to reflect this behavior? I have seen it used in [ARTEMIS-218|https://issues.apache.org/jira/browse/ARTEMIS-218] issue.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-13
18873	ARTEMIS-13	[~michalxo], I'm not clear on what documentation needs to be added. The [original commit for this issue|https://github.com/apache/activemq-artemis/commit/e293d80f08de58e4ce2bafd9d63c1a4185e079a9] included documentation which describes the functionality (in queue-attributes.md). What specifically is lacking in your opinion?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-13
18874	ARTEMIS-12	What exactly is required for this task? Is this just for documentation purposes or do we want to change ActiveMQ 6 to be more like 5.x?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-12

#	Requirement	Questions	Answered	Successfully Answered	URL
18875	ARTEMIS-11	Is there anything left to do on the distribution or can this be resolved?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-11
18876	ARTEMIS-8	We need to verify what are our boundaries for compatibility? going forward or back and forth. Should we have newer clients talking to older servers? We are intending to inclined to older clients talking to newer servers... if anyone thinks of any hard requirement for back compatiiblity on internal protocols (open-wire, core)... should raise a voice now. We will always have AMQP as an option for back and forth compatibility.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-8
18877	ARTEMIS-6	The following third party deps have been identified as having Cat X licenses. Please take these into consideration when performing this task. CatX Licenses:. Artifact: org.jboss.jboss-transaction-spi License: LGPL Notes: Will be removed as part of: https://issues.apache.org/jira/browse/ACTIVEMQ6-6 Artifact: org.jboss.jbossts.jts.jbossjts-jacorb License: LGPL Notes: Will be removed as part of: https://issues.apache.org/jira/browse/ACTIVEMQ6-6 Artifact: org.jboss.ironjacamar.ironjacamar-core-api License: LGPL Notes: Will be removed as part of: https://issues.apache.org/jira/browse/ACTIVEMQ6-6 Artifact: jboss.jbossts.jbossts-common License: LGPL Notes: Apparently JMS Tests need this. Needs some investigation. Artifact: org.jboss.security.jbosssx Lisence: LGPL Notes; Both JMS Tests and JBoss Integration are dependent on this. Needs some work to pull out dependent tests and intvestigation into integration module. Do we want to keep this hanging around? Artifact:>org.jboss.security.boss-security-spi Lisence: LGPL Artifact: org.jboss.ejb3.jboss-ejb3-ext-api License: LGPL Notes: javaee examples require this. Do we want javaee examples to hang around here. These may be best suited outside of activemq repository as they appear to be somewhat Jboss specific.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-6
18878	ARTEMIS-6	Isn't this done?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-6
18879	ARTEMIS-4	just leave this task with me for now... there are a bunch of other things I'm having to touch to fix the native... (variable names.. paths on the poms... etc).	1	1	https://issues.apache.org/jira/browse/ARTEMIS-4
18880	AMQCPP-253	What does your Socket.h look like with the changes to TcpSocket.cpp? I keep getting the following error: Error 744: "decaf/net/TcpSocket.cpp", line 32 # The operand of a static_cast must be a pointer to or lvalue of a complete class; the actual type was apr_socket_t *. : socketHandle(static_cast< SocketHandle >((apr_socket_t*)INVALID_SOCKET_HANDLE)), I've tried it with and without the (apr_socket_t*) cast in TcpSocket.cpp. Also tried it with/wo the changes to Socket.h where INVALID_SOCKET_HANDLE was changed from int to apr_socket_t* and it's intialization is placed outside of the class declaration. My environment: bash-2.04$ uname -srvm HP-UX B.11.23 U ia64 bash-2.04$ aCC --version aCC: HP aC++/ANSI C B3910B A.05.52 [Sep 05 2003]	1	1	https://issues.apache.org/jira/browse/AMQCPP-253
18881	AMQCPP-253	Timothy, Compiler settings were the same for the examples and the the lib build. I was able to manually build the SimpleProducer.cpp example and run it, but it segfaults somewhere in pthread_create. Wolfgang, is there any more information about your environment that you could provide? For example: Is it PA-RISC, or IA-64? Are you building 32-bit or 64-bit versions of the libraries? What version of the apache portable runtime (APR) are you using? Also, are you able to run any of the example programs or "make check"?	1	1	https://issues.apache.org/jira/browse/AMQCPP-253
18882	AMQCPP-253	Brian, Is there any particular reason why you use so ancient version of the aCC compiler? I wonder if these changes are necessary to compile using the latest version of aCC. $ uname -srvm HP-UX B.11.23 U ia64 $ aCC --version aCC: HP C/aC++ B3910B A.06.23 [May 18, 2009]	1	0	https://issues.apache.org/jira/browse/AMQCPP-253
18883	AMQCPP-250	I've checked in the complete but not very well tested InactivityMonitor class along with changes to enable its usage. Unit testing looks good so far, however more tests need to be added and some real testing with cable pulls etc needs to be done. Users are welcome to check out SVN rev. 881016 and give it a try. Test reports welcome.	0	0	https://issues.apache.org/jira/browse/AMQCPP-250
18884	AMQCPP-214	Can you please tell me the time line of 3.0 release? When shall it happen? We can organize ourselves based on that release. Thanks, Sahay	1	0	https://issues.apache.org/jira/browse/AMQCPP-214
18885	AMQCPP-203	Timothy, Your patch differs from my working copy only in "if (buffer) delete [] buffer, buffer = NULL", i.e. in the freeing of the buffer during destruction. I wonder how will this prevent other threads (recv) from writing into an already unallocated memory (which is, if I understood you well, what seems to be going on)... (Scratching head)	1	1	https://issues.apache.org/jira/browse/AMQCPP-203
18886	AMQCPP-202	Would you agree that we have resolved this by starting / stopping the broker between each test cycle? You've not seen any other hangs since then correct?	1	1	https://issues.apache.org/jira/browse/AMQCPP-202
18887	AMQCPP-202	Morning! The night test is still running (on #657 now) (without CMSTemplate) !! I'll let it a few hours more to be sure that the problem relies on this particular test; btw, there were no segmentation faults either. On the other hand, AMQCPP-199 stopped on the 75th pass - I'll connect gdb now to see what's going on there, and post the log file at AMQCPP-199. For the double thread, I'll have to take a deeper look at the code to figure out a way to detect it and break the execution. BTW, which APR version are you running? We should be seeing the same locks...	1	1	https://issues.apache.org/jira/browse/AMQCPP-202
18888	AMQCPP-189	Are there any consumers running to get the messages from the broker?	1	1	https://issues.apache.org/jira/browse/AMQCPP-189
18889	AMQCPP-177	This sounds fine, why don't you go ahead and make the changes to the studio projects and commit them for 2.2	1	1	https://issues.apache.org/jira/browse/AMQCPP-177
18890	AMQCPP-177	Can you raise this issue on the developer list so we can see if we can get that sorted out.	0	0	https://issues.apache.org/jira/browse/AMQCPP-177
18891	AMQCPP-140	Patch looks promising. Might want to think a bit about these warnings when building on Windows though: 1>..\src\main\activemq\network\SSLSocket.cpp(58) : warning C4244: 'argument' : conversion from 'activemq::network::Socket::SocketHandle' to 'int', possible loss of data 1>..\src\main\activemq\network\SSLSocket.cpp(95) : warning C4244: 'argument' : conversion from 'activemq::network::Socket::SocketHandle' to 'int', possible loss of data The offending lines look like this: SSL_set_fd(ssl, socketHandle); On Windows, SocketHandle is a Winsock2 SOCKET, which seems to be a typedef of UINT_PTR, which I don't think can safely be cast to an int (if it even makes sense to cast it). Can OpenSSL work with Winsock2 sockets? Maybe there's an alternative function to SSL_set_fd that should be called in this case?	0	0	https://issues.apache.org/jira/browse/AMQCPP-140
18892	AMQCPP-140	OpenSSL support implemented in trunk. Tested on Redhat and Windows. Hopefully others with access to different platforms will report their results. Notes: * SSL transport in ActiveMQ-CPP is specified using the transport URI prefix 'ssl' * The key and trust store are set as System properties using the decaf library methods as follows. To set the location of the client key file (Must be in PEM format) call: decaf::lang::System::setProperty("decaf.net.ssl.keyStore", "<path/filename>"); optionally a password can be provided as: decaf::lang::System::setProperty("decaf.net.ssl.keyStorePassword", "password"); To set the location of the trusted certificate for the broker, or the trusted root certificate for all your brokers the location of the file (again in PEM format) is set via: decaf::lang::System::setProperty("decaf.net.ssl.trustStore", "<path/filename>"); These methods should be called before you create your first ConnectionFactory.	0	0	https://issues.apache.org/jira/browse/AMQCPP-140

#	Requirement	Questions	Answered	Successfully Answered	URL
18893	AMQCPP-124	I'm not sure that the Major and Minor versions should come from the client build versions. I was thinking they would just be protocol-specific version info. For example, if a connection is using openwire v2, the major version would be 2. So basically, the connector would hard-code the version info for whatever it supports. Does this make sense?	1	0	https://issues.apache.org/jira/browse/AMQCPP-124
18894	AMQCPP-124	I'm fine with that. I'm not quite sure what the motivation was for this particular part of the JMS API anyway. Does anyone else want to chime in? Albert?	1	0	https://issues.apache.org/jira/browse/AMQCPP-124
18895	AMQCPP-115	Looks like the only place that Properties is used in CMS is in the Destination API, where it is read-only. Perhaps we could change the getProperties method to simply return a std::vector< std::pair< std::string, std::string > >? It might not be as nice from the user perspective, but it might be an easy way to eliminate Properties as a dependency of CMS. What do u guys think?	1	1	https://issues.apache.org/jira/browse/AMQCPP-115
18896	AMQCPP-115	I'll contact the author. Is the ASF's policy on this stuff written up somewhere? As far as I can tell, this macro being licensed under the GPL has no implications for the licensing of the rest of AMQCPP.	0	0	https://issues.apache.org/jira/browse/AMQCPP-115
18897	AMQCPP-115	Still trying to contact the ax_cflags_option author. One more thing to think about with regards to DLLs and shared libraries: we have some code in header files. Could this cause problems? Here I'm thinking along the lines of some internal structure changing without affecting the DLL/shared library "interface", but the user then has some AMQCPP code linked into their executable that expects the old structure. If I have this right, one should simply be able to upgrade the shared library in this case, without requiring the user to recompile their program. Meanwhile, I've been testing pyactivemq with the shared library on Linux and it seems to work without problems.	0	0	https://issues.apache.org/jira/browse/AMQCPP-115
18898	AMQCPP-115	I've got a visibility macro in trunk that seems to be working. I've moved both this macro and the addrinfo check into separate .m4 files in an attempt to clean up the configure.ac file. Is there anything else to be done for this issue, or can we finally put it to rest?	1	1	https://issues.apache.org/jira/browse/AMQCPP-115
18899	AMQCPP-102	source bundle has been created and the vote has been called. Closing this issue.	0	0	https://issues.apache.org/jira/browse/AMQCPP-102
18900	AMQCPP-100	has any progress been made on this?	1	0	https://issues.apache.org/jira/browse/AMQCPP-100
18901	AMQCPP-93	I also have been doing some profiling of the activemq-cpp-2.0.1 tag and I am finding that the DataInputStream class is fairly inefficient and can be greatly improved upon. Most of the inefficiencies are in the readString and readFully methods. readString appends to the string object one character at a time, which is very inefficient. Also, most methods are calling readFully for one byte. It would be better to create a different method to read one byte which would then call the read method the returns the unsigned char. I can probably work on a patch but I have never work with an open source project before. How do I proceed?	1	1	https://issues.apache.org/jira/browse/AMQCPP-93
18902	AMQCPP-93	Gerald, very nice! Would it be possible for you to decouple your benchmark from any "private" code and attach it here?	1	1	https://issues.apache.org/jira/browse/AMQCPP-93
18903	AMQCPP-93	Timothy, I have carefully looked at BufferedInputStream and ByteArrayInputStream read(unsigned char*, size_t) and they have NO LOGIC to return -1. These methods currently always throw IOException which means that the contract is never followed in the first place. Since -1 is never returned and IOException is always thrown, DataInputStream::readFully can never throw a EOFException. The unit tests cases don't test these conditions. I have a small test program which can be walked through with the debugger if you wish. Where should we go from here? According to the current logic, unsigned char read() and size_t read(unsigned char*, size_t) act exactly the same for reading a single byte. We will always hit a major performance problem until we can stop calling a virtual function for each character in a string. The other unmarshalling read statements can probably be improved upon by calling readFully for more than 1 character, but the only way I can think of that we can improve string is to add a readString pure virtual function in the InputStream class which then all derived classes need to implement. Please research and comment.	1	1	https://issues.apache.org/jira/browse/AMQCPP-93
18904	AMQCPP-93	If I understand correctly, this was with Openwire? Did you perhaps try the test with Stomp? We're currently using Stomp due to Openwire's performance issues.	0	0	https://issues.apache.org/jira/browse/AMQCPP-93
18905	AMQCPP-40	Which standard specifies "unix" or "__unix__" as something that is defined on all UNIX systems? I'd like to read up on this. Thanks!	1	0	https://issues.apache.org/jira/browse/AMQCPP-40
18906	AMQCPP-36	> I don't follow this rule. Do you mind if the patches move member functions to .cpp files? As I see it, headers shouldn't include implementation details. > One person's extra whitespace is another persons useful code separation. What I meant by stray whitespace is that there are lines that have 3 or 4 spaces at the end where they serve no purpose. It's the kind of whitespace that Eclipse's "Remove trailing whitespace" removes and the kind that Checkstyle warns you about. I'll update the Linux build files in a new patch. Meanwhile, here's a signedness fix from StompFrame.h.	1	1	https://issues.apache.org/jira/browse/AMQCPP-36
18907	AMQCPP-36	Yes, I do mind. I prefer to leave small functions in the headers of little classes like this, if they get inlined by the compiler or not its fine by me. Simply by nature of C++ the headers have implementation detail, and for that matter most people shouldn't need to be in that code as the CMS interfaces define the core functionality that most users of AMQCPP will be using. The question to ask is does the code work, if yes is the answer, then making all these trivial changes is just pulling time away from implementing the current enhancement under development.	1	1	https://issues.apache.org/jira/browse/AMQCPP-36
18908	AMQCPP-31	Would a better solution be to use the autoconf-style #ifdefs for headers and methods, like we've done elsewhere in the code?	1	0	https://issues.apache.org/jira/browse/AMQCPP-31
18909	AMQCPP-16	1) has been implemented in trunk. This required changes to DataInput/OutputStreams and ByteArrayInput/OuptutStreams to facilitate the new BytesMessage API. I also update the InputStream and OutputStream APIs to remove "const" where not necessary. Tim, give this a try in windows (regarding the const changes) ... gcc is very giving on this particular issue :) Also, could you comment on any of the others that you have implemented in trunk? Thanks.	1	1	https://issues.apache.org/jira/browse/AMQCPP-16
18910	AMQCPP-8	patch has been applied to trunk ... not sure if a Fix Version applies - are we including the activemq-cpp code in releases?	0	0	https://issues.apache.org/jira/browse/AMQCPP-8
18911	AIRAVATA-2384	Assigned this to you [~bharvadl]. Is this and https://issues.apache.org/jira/browse/AIRAVATA-2380 the same? If so close one of them as its a duplicate.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-2384
18912	AIRAVATA-2370	[~eroma_a], This should be fixed now. Can you test again?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-2370
18913	AIRAVATA-2369	[~marcuschristie] and [~tilaks26] Why have we removed User settings from gateway provider ? It would have been useful for them in case of needing email change, etc.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2369
18914	AIRAVATA-2340	Hello [~eroma_a] , I am interested in Ticket about Airavata Data Management . As my imagine , this should be a website to share the resource between users. How can I discuss further about the idea ? Please nicely let me know. Thank you. Regards, Minh	0	0	https://issues.apache.org/jira/browse/AIRAVATA-2340

#	Requirement	Questions	Answered	Successfully Answered	URL
18915	AIRAVATA-2335	Sent following email to [~scnakandala] {quote} Supun, I'm installing Airavata on Jetstream and ran into a problem when the sharing registry gets initialized. I'm getting this exception in the logs 2017-03-06 22:11:43 \[main] ERROR o.a.a.a.s.h.AiravataServerHandler - Error occured while reading airavata-server proper org.apache.thrift.TApplicationException: getDomain failed: unknown result at org.apache.airavata.sharing.registry.service.cpi.SharingRegistryService$Client.recv_getDomain(SharingRegistr at org.apache.airavata.sharing.registry.service.cpi.SharingRegistryService$Client.getDomain(SharingRegistryServ at org.apache.airavata.api.server.handler.AiravataServerHandler.initSharingRegistry(AiravataServerHandler.java: at org.apache.airavata.api.server.handler.AiravataServerHandler.<init>(AiravataServerHandler.java:120) I think the problem is that the domain default gateway domain doesn't exist and getDomain returns null, but Thrift doesn't support null returning methods. This explains why I get a thrift exception but there is no logged exception from the sharing registry itself. The problem starts at this line \[1]. What would be the best way to go about fixing this? My first thought was to add a method like 'doesDomainExist' to the sharing registry. Another, probably quicker idea would be to use getDomains() and check if it is empty or doesn't contain the default gateway domain. Thanks, Marcus \[1] https://github.com/apache/airavata/blob/a700073d6fddd9407cb3ed49ab31c5e9465291b4/airavata-api/airavata-api-server/src/main/java/org/apache/airavata/api/server/handler/AiravataServerHandler.java#L133-L133 {quote}	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2335
18916	AIRAVATA-2318	[~smarru] how can we implement this with little time ? Can we do something with defining it as a output file to bring back in App interface but as a optional output? meaning it exists bring back if not existing ignore.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-2318
18917	AIRAVATA-2283	PRE_REQUISITES ========== 1. embeded derby server 2. RabbitMQ on 149.165.228.91 ASSUMPTIONS ========== 1. we have already built airavata and tarball is present in distribution 2. No other processes are blocking airavata ports How to build ========== 1. mvn clean install in test-suit 2. It will pick AiravataIT and run stated test cases How to run ========== 1. Make sure you have proper airavata-server.properties inside test-suit also 2. Embeded derby server needs to be enables in main distribution (set start.derby.server.mode=false) What it does? ========== 1. It will extract airavata build from distribution and deploy it with embeded derby server 2. Automatically creates and tests gateway 3. Automatically creates and tests computer resource 4. Automatically creates and tests storage resource 5. Automatically creates and tests application 6. Automatically launches application and tests result	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2283
18918	AIRAVATA-2283	GitHub user ajinkya-dhamnaskar opened a pull request: https://github.com/apache/airavata/pull/91 Airavata integration test Jira: AIRAVATA-2283 Integration test suit to test airavata experiment lifecycle It creates resources locally and uses local provider to test simple echo experiment. PRE_REQUISITES ========== 1. embeded derby server 2. RabbitMQ on 149.165.228.91 ASSUMPTIONS ========== 1. we have already built airavata and tarball is present in distribution 2. No other processes are blocking airavata ports How to build ========== 1. mvn clean install in test-suit 2. It will pick AiravataIT and run stated test cases How to run ========== 1. Make sure you have proper airavata-server.properties inside test-suit also 2. Embeded derby server needs to be enables in main distribution (set start.derby.server.mode=false) What it does? ========== 1. It will extract airavata build from distribution and deploy it with embeded derby server 2. Automatically creates and tests gateway 3. Automatically creates and tests computer resource 4. Automatically creates and tests storage resource 5. Automatically creates and tests application 6. Automatically launches application and tests result You can merge this pull request into a Git repository by running: $ git pull https://github.com/ajinkya-dhamnaskar/airavata airavata-integration-test Alternatively you can review and apply these changes as the patch at: https://github.com/apache/airavata/pull/91.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #91 ---- commit ce359dc0f086dcc5719b2ec42237e7f1d505b135 Author: Ajinkya Dhamnaskar <adhamnas@umail.iu.edu> Date: 2016-12-18T18:44:43Z modules/test-suite/multi-tenanted-airavata/src/main/resources/airavata-server-stop.sh Register gateway, compute resource, storage resource, application and launch application commit b46aa6e483d4e5e93e3d521b40ffa9ca103c35a9 Author: Ajinkya Dhamnaskar <adhamnas@umail.iu.edu> Date: 2016-12-18T18:46:52Z modules/test-suite/multi-tenanted-airavata/src/main/resources/test-framework.json Integration test config and server stop script commit cd2abca6996573ee6696a9bec4a2e3d3a64e41e5 Author: Ajinkya Dhamnaskar <adhamnas@umail.iu.edu> Date: 2016-12-18T18:48:11Z Adding test framework prperties commit 6c4ac5a128b8ea2dced763862f77188e0793754a Author: Ajinkya Dhamnaskar <adhamnas@umail.iu.edu> Date: 2016-12-18T18:49:05Z code cleanup commit c072479683dc11bb5ac28a278a4e5a2d175eeaee Author: Ajinkya Dhamnaskar <adhamnas@umail.iu.edu> Date: 2016-12-18T18:50:02Z file for future reference commit 4d43fc48aef971c5122940f30332d1cf30bc135e Author: Ajinkya Dhamnaskar <adhamnas@umail.iu.edu> Date: 2016-12-18T18:51:28Z adding pom commit 1e310f3ad6c85bfbdec0dca8a4e0761b08e6b847 Author: Ajinkya Dhamnaskar <adhamnas@umail.iu.edu> Date: 2016-12-18T18:52:24Z changing pom and airavata properties in main distribution commit 3f1b73bcaea1df1af76bb2f5b70d82d6f7c0d0f5 Author: Ajinkya Dhamnaskar <adhamnas@umail.iu.edu> Date: 2016-12-18T18:53:00Z adding readme ----	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2283
18919	AIRAVATA-2225	[~smarru], Why do we even need to expose this feature in the UI? Why not just always set 'Enable Auto Scheduling' to true and hide it from the user? I'd really like to not even expose 'Enable Auto Scheduling' to the user. My guess is the user will have no idea what it means.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2225
18920	AIRAVATA-2225	[~smarru], Does this feature need application support? Wouldn't the application need to be resumable to be autoscheduled? Maybe whether to autoschedule should be in application configuration. I'm thinking we can * not show 'Enable Auto Scheduling' to user * in Orchestrator, check if application supports auto scheduling and if it does just automatically enable auto scheduling I'm just trying to figure out where in the system this knowledge should be kept.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2225
18921	AIRAVATA-2178	[~shameera], I'm working on changing the experiment creation page in PGA to allow a user to select their own compute resource account for running the job. Have you thought about what changes we need to make to the data models so that PGA can convey that the user wants to use their own compute resource account? I'm not sure where it would best go, but I'm thinking some sort of flag on the ComputationalResourceSchedulingModel.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-2178
18922	AIRAVATA-2174	[~eroma_a], I've updated dev seagrid. I've done some testing, but can you also test the experiment sharing? It should allow you to share during execution and after execution.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2174
18923	AIRAVATA-2170	[~scnakandala], Looks like whenever a user saves a shared project, that user becomes the owner of the project. Is that by design?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2170

#	Requirement	Questions	Answered	Successfully Answered	URL
18924	AIRAVATA-2164	Pull request for code changes: https://github.com/apache/airavata-php-gateway/pull/23 Supun, can you review the pull request when you have a chance?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2164
18925	AIRAVATA-2158	[~anujbhan], I also need some way to set the description on a credential when creating it. That is, where the current API allows {noformat} string generateAndRegisterSSHKeys (1: required security_model.AuthzToken authzToken, 2: required string gatewayId, 3: required string userName) {noformat} I need instead something like {noformat} string generateAndRegisterSSHKeys (1: required security_model.AuthzToken authzToken, 2: required string gatewayId, 3: required string userName, 4: optional string description) {noformat} Can you add that on to this issue or would you prefer I create a separate issue for it?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-2158
18926	AIRAVATA-2117	Here is the discussion summary: * Is this user story a direct descendant to Use-Case "Resource & job management for user's with third party Allocations" in https://issues.apache.org/jira/browse/AIRAVATA-2108?focusedCommentId=15512194&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-15512194 * Will the user be given the privilege to update this preferences or the gateway admin does that ? * What Data will the Gateway need ? username, password, sshURL * Who will configure the SSH key, user or the Gateway ? * In order to provide this feature, An optional endpoint will be exposed out of Airavata API, which will update the Userprofile data modal and respective changes to Gfac and PGA User Interface.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2117
18927	AIRAVATA-2117	[~anujbhan], For the UserResourceProfile model, what is the userId? Is it the username, email address, or something else?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2117
18928	AIRAVATA-2117	Thanks Anuj. I have another question. In PGA, I tried to query to see if a user has an existing UserResourceProfile {code} public static function get_or_create_user_resource_profile() { $userId = Session::get('user-profile')['username']; $gatewayId = Session::get('gateway_id'); $userResourceProfile = Airavata::getUserResourceProfile(Session::get('authz-token'), $userId, $gatewayId); return $userResourceProfile; } {code} However when I try this I get {noformat} Airavata \ API \ Error \ AiravataSystemException HELP Error while retrieving user resource profile. More info : Error while retrieving user resource profile. More info : javax.persistence.NoResultException: Query "SELECT p FROM UserResourceProfile p WHERE p.userId =:param0 AND p.gatewayID =:param1" selected no result, but expected unique result. {noformat} I think this is because [the backend code calls getSingleResult()|https://github.com/apache/airavata/blob/50cf71e44431e1dce3653df25790443d4a77ec0f/modules/registry/registry-core/src/main/java/org/apache/airavata/registry/core/app/catalog/resources/UserResourceProfileResource.java#L161-L161]. Is there a better method to use to figure out if a UserResourceProfile already exists?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2117
18929	AIRAVATA-2117	[~anujbhan], Should getUserResourceProfile return null if there is no record instead of throwing an exception? The alternative I guess is that I catch the AiravataSystemException, but I don't like that because then the code would be assuming that the AiravataSystemException was thrown because a UserResourceProfile doesn't exist.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-2117
18930	AIRAVATA-2100	Hi Lahiru, How many instances do you recommend to setup only for log analysis in production with enabling HA?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-2100
18931	AIRAVATA-2070	Hi Lahiru, Could you please explain, reason for this approach instead of sending directly to logstash?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-2070
18932	AIRAVATA-2068	Could you explain how airavata instances publish messages, it is for one partition or different partitions? Do we have different partitions for gfac, orchestrator ,api .... components? Is there any diagram explain your approach?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-2068
18933	AIRAVATA-2011	The fix here seems to be disabling edit capabilities for users with read permission only, is that correct?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-2011
18934	AIRAVATA-1987	Hi Gourav, can you send a pull request by adding this into airavata code base?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1987
18935	AIRAVATA-1987	Hi Suresh, can you direct me where to insert this module?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1987
18936	AIRAVATA-1956	I think now is a good time to do the upgrade, if it is quick. I'm looking at the documentation and the latest Laravel is 5.4: https://laravel.com/docs/master/upgrade However, [starting in Laravel 5.3, PHP 5.6.4 or higher is required|https://laravel.com/docs/master/upgrade#upgrade-5.3.0]. I'm wondering how easily or not we'll be able to upgrade our servers where PGA runs to PHP 5.6.4 or higher.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1956
18937	AIRAVATA-1956	I can log in to PGA but now I get this error: {noformat} TSocket: Could not connect to tls://gw56.iu.xsede.org:9930 {noformat}	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1956
18938	AIRAVATA-1956	[~marcuschristie], most likely its SSL related. Did you try to connect to a non-secure Airavata and test it?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1956
18939	AIRAVATA-1948	When Summary and Detail were initally thought of, they were based on which user role is looking at it. A Gateway Admin gets to see Experiment Details/Summary. All users get to see only Experiment Summary. Are we moving away from this and deciding to show Details to everyone? Then taking this action would make sense, otherwise not.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1948
18940	AIRAVATA-1948	So in Experiments' context, do we want all user types be able to see a detailed view?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1948
18941	AIRAVATA-1894	Do we need to make a front end web UI for workflows editor using front-end technologies like javascript, html or css for adding drag and drop features? And we can use Google's GSON library for JSON implementation in workflow.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1894
18942	AIRAVATA-1893	Hi Anuj, Can you create a child page under https://cwiki.apache.org/confluence/display/AIRAVATA/GSoC+2016 and add your proposal there? * You will need create an account first on wiki and tell us the user name so we can give you required permissions. * Can you also start the proposal in gsoc dashboard and choose external link for the draft?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1893
18943	AIRAVATA-1880	tested and verified in PGA (http://149.165.156.11/) in JS Can set read only flag to true or false in application interface. When this flag is set for URI, no difference at the experiment create, edit and clone. When this is set to true for String, Integer or float user cannot change the value given by the gateway admin at creation, edit or clone of an experiment	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1880
18944	AIRAVATA-1847	In the first look - - Like Eroma, I also like the previous simpler single-headed elephant's logo more over the three-headed versions because of simplicity. - i am not inclined towards having black color on the logo, mainly because to me it doesn't represent an opensource foundationn. It suits more to a stylish/classy brand like a car or a watch company, but that is just my opinion. - I like the use of font in 'APACHE AIRAVATA' in the three-headed RED logo. - I don't know why, but the three-headed logos with trunks in the same direction do not resonate with me. Can we get an explanation behind the making of the logos to get a deeper understanding about them of why they are the way they are? ---------- I'm not a graphic designer and I understand very little about philosophy of logo design, but I have a couple of suggestions that I would like to add with examples that we can try if we have time - We could try flat colors used now-a-days like - http://www.flatuicolorpicker.com/ We could try giving a material feel to the logo like this one - https://cdn.colorlib.com/wp/wp-content/uploads/sites/2/2013/10/hudddle-logo.png ----------	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1847

#	Requirement	Questions	Answered	Successfully Answered	URL
18945	AIRAVATA-1847	[~eroma_a] [~ndoshi] How about you consult with Suresh and you can come back with some concise feedback. I would be willing to do any of this but I feel like you all should align your opinions and then I can make changes. I used types of flat colors on the logo. What explanation is needed for the logos? I want to emphasize Airavata and the elephant. What do you mean by material feel? Thanks, Blair Carlisle	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1847
18946	AIRAVATA-1847	Nipurn, Shivika, Its greatly impressive the way you are able to embed the gateways concept into the logo. It will be better to frame Airavata as a general purpose distributed systems framework and not constrain it to one use case. Since distributed is hard to illustrate beyond few connected things like [1] [2] [3], can you also make one attempt on not worrying any of these and suggest a logo and branding which just helps to promote the words Apache Airavata (or only Airavata) without necessarily have a thoughtful meaning? [1] - http://www-fp.pearsonhighered.com/assets/hip/images/bigcovers/0321263545.jpg [2] - https://openit.com/wp-content/uploads/2015/01/network-concept-600x450.jpg [3] - http://www.igi-global.com/Images/Covers/9781466626478.png	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1847
18947	AIRAVATA-1792	Github user machristie commented on the issue: https://github.com/apache/airavata-php-gateway/pull/52 @scnakandala when you have a chance, can you please review and merge this PR? Thanks.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1792
18948	AIRAVATA-1716	What you mean by duplicate with other jobs? We are not storing the sub-stringed job id. We are just doing it when we run the query. I tested it with all the resources and it works fine. Please pass the example where its not working and I am happy to test.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1716
18949	AIRAVATA-1663	Chathuri, just curious, how are we going to start all services in one go?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1663
18950	AIRAVATA-1649	Hi Abhishek, the proposal is reasonable with details. Can you immediately submit it in melange - https://www.google-melange.com/gsoc/homepage/google/gsoc2015, make it public and send a link on this JIRA? We could suggest improvements there.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1649
18951	AIRAVATA-1649	Hi Abhishek, Yes I can view your proposal now. This is a comment to both you and [~priyanshu-patra]. Both your proposals are looking identical and I had to spend time to get the differences out. Can you both in your respective proposals clearly distinguish your deliverables. Thrift is a polygot framework and you can see the full list of supported languages at - http://thrift.apache.org/lib/ Airavata currently supports a subset of these as client SDK's (Java, PHP, C++, Python). Can you make sure you have the initial deliverable of contributing good examples to register applications, create and monitor executions into Airavata Client SDK repo? Once you understand Airavata well enough through these examples, then you can integrate these into GenApp.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1649
18952	AIRAVATA-1649	All right sir, and as you suggested in the second part of earlier comment- we will modify our proposals by adding an initial deliverable of making some examples which will be added to the Airavata Client SDK repo by writing a simple program,getting to run via Airavata,defining it to genapp and then encoding the steps required to run any program via airavata and put that in genapp. Does that sound good?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1649
18953	AIRAVATA-1649	Hello Sir, I didn't actually understand your last comment. Can you please explain me.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1649
18954	AIRAVATA-1646	In regards to this project, I would like to understand the scope of the work. I am interested in identifying what are the best ways to incorporate Cassandra into the project. Here are some questions I have particularly towards metadata and provenance: - Do we have access to the apache thrift data model currently in use by Airavata? If so, can we modify this model? - What other object store technologies are you interested in (Cassandra and MongoDB)? - How will the metadata be used? Depending on metadata usage it can affect which technologies and which features of that specific technology we should enable. - What are some examples of meta data is being stored? Is the data structured or unstructured? - What kind of provenance data is being stored? - What kind of queries would you expect to be run on the provenance data? - Do we need look into Apache Storm for querying streaming data? - Will we receive accounts on NSF XSEDE clusters for this project? Thanks, Doug	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1646
18955	AIRAVATA-1646	I am really enthusiastic about the option to explore big data technologies (concepts and implementations) for Airavata Data Management needs. I have some questions: Is information such as Application Characteristics available? For example, are applications Read/Write/IO/CPU/memory intensive? Also, is one of the outcomes of this project a recommendation for a software pipeline that shows which big data technologies should be combined? Do we need look into Apache Storm for querying streaming data? And, what is the data size and how many nodes can we run data on? Finally, where will these experiments be run? Locally on my own machine? Pankaj	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1646
18956	AIRAVATA-1646	Hi Doug, Please see the responses embedded below: Do we have access to the apache thrift data model currently in use by Airavata? If so, can we modify this model? --I consider this project as a exploratory, so yes we could branch the master and have you modify the thrift data models. You can look at them here - https://github.com/apache/airavata/tree/master/airavata-api/thrift-interface-descriptions What other object store technologies are you interested in (Cassandra and MongoDB)? --It will be premature to state a preference. The key thing here is to understand the problem well enough and make a recommendation if relational databases are good, or if key-value or column, document or graph databases can better address Airavata metadata needs. How will the metadata be used? Depending on metadata usage it can affect which technologies and which features of that specific technology we should enable. --This is a very open ended question. I will hope you can propose a project keeping in mind you will need to explore this answer in interactions with airavata community. What are some examples of meta data is being stored? Is the data structured or unstructured? --Currently all the metadata is very structured. An example would be to look the experiment model. User requests an experiment which will get executed on remote resources, in the process transforms data. The metadata capture also included states of simulation or data analysis tasks. Once you run sample experiments, this will be more clearer. What kind of provenance data is being stored? --Currently very minimal to none. Basic information like user provided metadata, resources used to compute, job dimensions. A big missing piece is to collate provenance of input data and augment provenance of generated data with application details and simulation/analysis configurations. What kind of queries would you expect to be run on the provenance data? --this will be very subjective to the data domain. An example could be, query for all radar assimilation data which have a quality score of 5. We could find more concrete pointers. Do we need look into Apache Storm for querying streaming data? -- Not right away, but I could foresee some usage. For instance, if we have to process metadata extraction from all the archived data, I could see storm helping to run such a topology. We could also employ a storm cluster to shred deep data from all input requests. Again, we need to adapt with the usecases a bit here. Will we receive accounts on NSF XSEDE clusters for this project? --Yes we could get you access to various clusters including XSEDE if absolutely needed by the project.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1646

#	Requirement	Questions	Answered	Successfully Answered	URL
18957	AIRAVATA-1646	Hi Pankaj, Yes we could application characteristics. Typically (or currently) airavata executed applications are one or more of CPU/Memory/IO intensive. Most of the time they are CPU intensive. I am not sure what you mean by pipeline in this context. Can you elaborate? The goal is to explore if big data tools can solve/enhance airavata usecases. So this could lead to some open ended nature of this project. But to concretely help you started, may be you should consider the following usecase: Assume Airavata is populated with lots of previously run simulations. analyze the data and when a scientist is configuring a new application, can we help guide it based on previous success/failures?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1646
18958	AIRAVATA-1646	Hi Chris, I fully agree we need to bring this to surface again. There were couple of academic projects (one undergraduate team, one masters student) who explored using OODT for Airavata data management needs, need to check on that. Rishi Verma and few of us also discussed some synergies [1]. How about we find time at ApacheCon and make plans so we can pursue either GSoC, or summer of code or other opportunities to make progress? [1] - http://airavata.markmail.org/thread/5uabiaceuj2eqayl	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1646
18959	AIRAVATA-1636	Hello, I've got several questions about implementing a benchmark framework. First, as Airavata is available for installation on private machines / clusters rather than being centralized to one global installation, does it have a mechanism built in to record data from all experiments that it has executed? For example, if an organization is using Airavata to manage and run experiments, will any future benchmark framework be able to access all previous records (including things like experiment type, input size, computation resource, and run time) to prime its decision engine? Or do users have to manually record this data until the benchmark framework is installed and able to record it? Are all the computation resources running a uniform scheduler? Or is there some common API to access their scheduling queue? If the benchmark framework predicts that resource A will execute an experiment in 3 hours while resource B will compute it in 6 hours, but resource A is allocated for the next 6 hours, it would be ideal to schedule the experiment for resource B instead. I'm wondering if this can be used to help inform the decision. All execution on the resources is *not* performed through Airavata, correct? For example, other users may choose to run their experiments through a terminal or other applications? If all scheduling is done through Airavata, information relevant to the previous question could be provided by the benchmark framework itself. My first idea for such a framework is to provide a set of known parameters and historic execution data to a machine learning algorithm or framework and provide extra information to the user as they schedule an experiment. If this has been discussed as a possibility already, is there a link to the discussion? Thanks, John	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1636
18960	AIRAVATA-1635	1. Is there any documentation that describes about the API functions of Middleware Service [2] and Airavata SDK? 2. Do I have to setup a local SCiGap (Airavata) instance or is there a existing setup?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1635
18961	AIRAVATA-1635	In GridChem-Client [1] build and bin folders also have pushed to the repository. Is it necessary to have them in repo? This will lead to huge list of changes once the project is built locally. [1] https://github.com/SciGaP/sha2-GridChem-client	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1635
18962	AIRAVATA-1635	Hi all, Suresh : Ok. I removed previous forks in my repo and these are new forks https://github.com/DImuthuUpe/GridChem-Client https://github.com/DImuthuUpe/GridChem-Middleware-Service I will be working on my forks until I improve builds. Hopefully we may be able to convert ant builds to maven builds. Then I'll send you a pull request. I tried to run GridChem-2.0.0.jar using java -jar GridChem-2.0.0.jar but it gave following error. How can I resolve this? It seems like some libraries are not added to the classpath. Exception in thread "main" java.lang.NoClassDefFoundError: org/apache/commons/cli/CommandLineParser at java.lang.Class.getDeclaredMethods0(Native Method) at java.lang.Class.privateGetDeclaredMethods(Class.java:2451) at java.lang.Class.getMethod0(Class.java:2694) at java.lang.Class.getMethod(Class.java:1622) at sun.launcher.LauncherHelper.getMainMethod(LauncherHelper.java:494) at sun.launcher.LauncherHelper.checkAndLoadMain(LauncherHelper.java:486) Caused by: java.lang.ClassNotFoundException: org.apache.commons.cli.CommandLineParser at java.net.URLClassLoader$1.run(URLClassLoader.java:366) at java.net.URLClassLoader$1.run(URLClassLoader.java:355) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:354) at java.lang.ClassLoader.loadClass(ClassLoader.java:423) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:308) at java.lang.ClassLoader.loadClass(ClassLoader.java:356) ... 6 more Sudhakar: Thank you. I'll look into that and let you know in case of issue. I may need some details of how to build and start GridChem-Middleware-Service. I can see some targets in build.xml but couldn't figure out which target to invoke. There are some shell scripts inside the project (restart_tomcat7.sh, restart.sh, run_make.sh, startup_tomcat7.sh, startup.sh, test.FTD.sh). What is the correct script(s) that are needed to start the server?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1635
18963	AIRAVATA-1635	What is the method to run GridChem client? Is it just running GridChem-2.0.0.jar in bin folder or is there a specific way to run it? And there are several targets in build file that can be invoked. What is the target I need to invoke in order to build this again?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1635
18964	AIRAVATA-1635	Suresh/Sudhakar, Can I have access to a working system of GridChem as we have discussed earlier? Then I'll be able to get familiar with its use cases.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1635
18965	AIRAVATA-1635	I tried GridChem client by putting some jobs on the servers and getting outputs. Input files that we push has a format like this. %chk=water.chk %nprocshared=1 %mem=500MB #P RHF/6-31g* opt pop=reg gfinput gfprint iop(6/7=3) SCF=direct Gaussian Test Job 00 Water with archiving 0 1 O H 1 0.96 H 1 0.96 2 109.471221 Is this format some standard way of passing jobs or specific format for GridChem? Does client directly pass this file to middleware or parse this and pass only necessary data?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1635
18966	AIRAVATA-1635	Hi Suresh, I have almost finished porting Create/Edit Experiment interface of GridChem client. I need to test it with applications like Gaussian or GAMESS. Some resources can be found form [1] and [2] but they do not explain how to register those applications in locally deployed Airavata Server. What are the steps I need to follow in order to install them in my local machine? Or can I get an already hosted server with these applications registered for my testing purposes? [1] https://cwiki.apache.org/confluence/display/AIRAVATA/Script+Example+-+GridChem+Gaussian [2] https://cwiki.apache.org/confluence/display/AIRAVATA/Gaussian+Input+examples	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1635
18967	AIRAVATA-1624	Hi Hasini, Please refer to [1] for some brief discussion on this topic. In Scenario 2 of your document (figure 3) does wso2IS currently have such user store extension or will it have to be written? What will be the requirement of Gateway User Store? I am assuming most of these include a custom mysql user management or some use bundled postgre database like with drupal/wordpress. Suresh [1] - http://markmail.org/thread/vveb33oefxdbecxq	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1624
18968	AIRAVATA-1624	Hi Hasini, Some of your points in the discussion seem to indicate you assume Airavata will be deployed as a stand alone. But we are intending for it to be deployed as a tenanted platform. Will your design change if a single deployment has to serve all three use cases (gateways with diverse user management scenarios)? Suresh	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1624

#	Requirement	Questions	Answered	Successfully Answered	URL
18969	AIRAVATA-1624	Hi Hasini, Can we also consider a scenario 4 (which is a slight variation of Scenario 2 and 3) where gateway and Airavata needs to communicate via a persistent token (instead of a per user OAuth token)? Typically in a OAuth approach user id hand holding the trust and allowing Gateway and Airavata to trust his credentials. In this scenario, the gateway takes care of user authentication (either with custom user store or through federated identity) and performs some actions on the gateway portal. Subsequently, some user actions will require Airavata interactions at which point gateways send user information after authenticating and authorizing through a one time pre-issued gateway API token (similar to the developer tokens widely used). Airavata also supports Desktop applications (and in future mobile apps) in addition to the Web based gateways. Since your proposed implementation is OAuth centric, will all scenarios work for desktop and mobile apps?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1624
18970	AIRAVATA-1624	Hi Suresh, I was trying to edit the wiki page you have linked above, however it seems that I do not have permission to edit that. Can you please grant me the permission to edit the wiki page of the proposal at https://cwiki.apache.org/confluence/display/AIRAVATA/%5BGSoC+Proposal%5D+Securing+Airavata+API? My username for wiki is: hasinig. Thanks & Best Regards, Hasini.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1624
18971	AIRAVATA-1624	Hi Hasini, Thanks for the clarifications. I am planning to implement secure airavata client using PHP for the php gateway. Can you give me the instructions on how to develop the php client to use the IS OAuth services. Are there any existing client libraries available that I can use and what is your recommendation? Thanks Supun	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1624
18972	AIRAVATA-1421	What's the status of this? There have been some commits. Please assign yourself and update the status.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1421
18973	AIRAVATA-1420	Supun, the patch seems dubious because of its size (53MB). Can you only submit a patch with source code, also, make sure to remove .idea directories and all their contents? Also, remove all binaries (classes, jars..) and log files. Lastly, did you make sure the code has a maven build and all dependencies are coming from maven central? (make sure you did not add any custom repositories in the pom file).	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1420
18974	AIRAVATA-1415	Is this fixed? We have commits.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1415
18975	AIRAVATA-1397	Why is there a 'server-allowed-file-size' variable in pga_config.php? Is this just supposed to be the same as 'upload_max_filesize' in php.ini? If so, we don't need that since we can get the value from php.ini. Currently we have a bit of JavaScript that checks to see if files selected for upload are greater than the minimum of post_max_size and 'server-allowed-file-size'. However, * I would propose getting right of 'server-allowed-file-size' and instead look at the minimum of upload_max_filesize and post_max_size * it wasn't set up for the optional input files but I've fixed that locally That catches this problem earlier, client side. Here are some other improvements that I think can be made # Display the max upload file size in the form so that users (and administrators) have some clue as to what the max file size is. # Determine server side if the upload_max_filesize has been exceeded and redisplay the experiment creation form with an error message. # Determine server side if the post_max_size has been exceeded. Not sure what can be done in this case since it looks like $_POST is empty. We can redisplay the form with an error but the values provided by the user will be completely empty. #* See also http://stackoverflow.com/a/2133726 With the client side detection, I'm not sure how important #2 and #3 are. Since there isn't much that can be really be done with #3, I don't think it is worth it. I think #2 is worth it however, since you can't always rely on client side JavaScript.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1397
18976	AIRAVATA-1389	Can someone take this on?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1389
18977	AIRAVATA-1380	Currently the web browser back button works fine in most scenarios. Only place it does work is when the user tries to navigate through experiment creation, cloning and editing with browser back button user can navigate to browse experiments and browse projects respectively from experiment summary and Project summary. Shall we leave it as it is?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1380
18978	AIRAVATA-1380	Going back from experiment summary page to create experiment seem to be creating issues. The CPU, node related infromation is missing. Can we disable this button? [~scnakandala] lets discuss this issues and if possible we can push it to next version since we are going to do a next gen PGA anyway	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1380
18979	AIRAVATA-1374	Hi Nipun, These patches did not get applied before. Can you please retake these patches with latest code and re-submit?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1374
18980	AIRAVATA-1361	In Experiment -->Browse gave time to and from and also state COMPLETED. Only 3 experiments got retrieved. where as there should be 100+ Time given 01/15/2016 9:20 AM - 01/15/2016 3:20 AM	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1361
18981	AIRAVATA-1257	Here is the data model struct ExperimentSummary { 1: required string experimentID, 2: required string projectID, 3: optional i64 creationTime, 4: required string userName, 5: required string name, 6: optional string description, 7: optional ExperimentStatus experimentStatus, 8: optional list<ErrorDetails> errors } Is this ok to go forward ?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1257
18982	AIRAVATA-1190	We need to know the limits of input validation. How much validation will the Thrift clients do? What happens if the client gets unexpected values (string instead of int, for example)? Does the client throw an exception?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1190
18983	AIRAVATA-1176	Is this working correctly after the recent SDK updates?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1176
18984	AIRAVATA-1176	Experiment ID is now passed in via the command line. Works properly with updated SDK. Only updates the experiment's name - is that OK for the tutorial?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1176
18985	AIRAVATA-1171	Added command-line arguments for user and experiment name. Is this sufficient for the tutorial?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1171
18986	AIRAVATA-1160	Did this get applied?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1160
18987	AIRAVATA-1159	Was this patch applied?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1159
18988	AIRAVATA-1140	What is the status of this?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1140
18989	AIRAVATA-1043	Commit 8a615ec049d4c4f8421cf8ac20046e3a6744da9d in airavata's branch refs/heads/master from [~raminder] [https://git-wip-us.apache.org/repos/asf?p=airavata.git;h=8a615ec] Save JobDetails and use task details. AIRAVATA-1043, AIRAVATA-1042	1	1	https://issues.apache.org/jira/browse/AIRAVATA-1043
18990	AIRAVATA-1008	This is relating to current GFac interface : org.apache.airavata.gfac.cpi.GFac IMO it would make sense to change, public boolean submitJob(String experimentID) throws... to public void submitTask(String TaskID) throws... Functions to add cancelTask(...) But do we need a restartTask(...)? Do we have a realtime notification on GFacTask status changes? This is important for the workflow interpreter since it is inefficient for the interpreter to poll the registry and wait for status changes to trigger its node executions.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1008
18991	AIRAVATA-1006	Hi Marlon, Suresh, Please give feedback to [1]. Will this be too abstract ? Someone who programs against this should only know about the thrift data models. He does not have to know about the registry structure. [1] https://git-wip-us.apache.org/repos/asf?p=airavata.git;a=blob;f=modules/registry/registry-cpi/src/main/java/org/apache/airavata/registry/cpi/Registry.java;hb=25cc2514	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1006

#	Requirement	Questions	Answered	Successfully Answered	URL
18992	AIRAVATA-1006	Hi Chathuri, This is indeed very intuitive. One a flip side we can look at critically and say its too abstract and high level, but I find the abstraction very simple and neat. When data models change adding removing fields, this CPI will almost require no changes what so ever. Now this will burden all components to focus on airavata thrift data models and adapt to changes. But that may anyway require and it alleviates form changes to registry cpi as well. Its also very clean in the sense that internal organization of the data and the relations between them, whether structured, unstructured or semi-structured is now a registry implementation detail. + 1 for moving foreword. How will the update work? lets say if there are some optional values filled in and a request comes to update the object with a values for new parameters and for some previously set parameters, will all of them be updated? How will the transactions be enforced when two non co-ordinating components are trying to update the same obejct? Suresh	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1006
18993	AIRAVATA-1006	Also, will it be useful to have a getSupportedRegistryDataTypes? I like the isExist, but wondering if it will be useful to get a fill list.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-1006
18994	AIRAVATA-1005	As the CPI's are evolving, integrating them into various components is warranting initiating the components, setting the constructors and especially non-blocking calls creating lot of objects in memory. How about we let the thrift services penetrate to these components and move foreword with CPI level thrift service? I will assume lazy consensus and prototype the registry CPI as a service, meanwhile lets please discuss the alternatives.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-1005
18995	AIRAVATA-994	Thanks for the feedback Raman. How about adding these two table to keep workflow inputs and outputs ? CREATE TABLE WORKFLOW_INPUTS (EXPERIMENT_ID VARCHAR(255) NOT NULL, KEY VARCHAR (255) NOT NULL, VALUE VARCHAR (255), PRIMARY KEY (EXPERIMENT_ID, KEY), FOREIGN KEY (EXPERIMENT_ID) REFERENCES EXPERIMENT_METADATA(EXPERIMENT_ID) ON DELETE CASCADE) CREATE TABLE WORKFLOW_OUTPUTS (EXPERIMENT_ID VARCHAR(255) NOT NULL, KEY VARCHAR (255) NOT NULL, VALUE VARCHAR (255), PRIMARY KEY (EXPERIMENT_ID, KEY), FOREIGN KEY (EXPERIMENT_ID) REFERENCES EXPERIMENT_METADATA(EXPERIMENT_ID) ON DELETE CASCADE)	1	1	https://issues.apache.org/jira/browse/AIRAVATA-994
18996	AIRAVATA-994	Thanks Chathuri. The new model look good to me. Should we call it WORKFLOW_* or EXPERIMENT_*. We need to record, which workflow/application to invoke in one of the models.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-994
18997	AIRAVATA-992	What are the intended functionalities from this API? The JIRA AIRAVATA-991 says "Provides access to client queriable Airavata configuration data". What I'm interested is what client applicable Airavata configurations are we going to expose via this?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-992
18998	AIRAVATA-992	Saminda, this is a good question and I thought about a need other day but could recollect now. I was intending some configurations like what type of resources for Airavata service communicate to? as an example this will return the providers from gfac like EC2, GRAM and so on. Do you have any concrete needs from this block of API? As a general rule, it will be better for us not to add too many api methods and data we foresee, but do not have concrete use now. Its better to base all of API based on what we know now. I agree with API community thats its a API development is an art and not an engineering (hence the title craft and not design) and I am myself a bad artists. So please crtique and change liberally.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-992
18999	AIRAVATA-991	Hi Suresh, Currently we have data models inside airavata-api-stubs and airavata-data-models which leads to redundancy and confusion. From airavata components, we use data models inside airavata-data-models project. Instead of having two places, can we have all the generated data models in airavata-data-models ?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-991
19000	AIRAVATA-985	Thanks, Sachith for the contributions. The wiki pages look good. The sample code is constructed as unit tests, however, which may make it hard for new developers easily adapt to their usage. Is there a better way to present this? Also one other cleanup is needed: the patch includes an ItelliJ file (gateway-developer-guide.iml) that you should delete.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-985
19001	AIRAVATA-965	I guess one of the reasons to have 2.1 is to set the bouncycastle 1.49 as a dependency. How much time do you think it is required to test the jglobus release? and when do you want to release jglobus 2.1?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-965
19002	AIRAVATA-965	Can i create a branch with Jglobus and unicore changes and then rollback the version of Jglobus? We can continue to test both the provider and can request a jglobus release after enough testing. It will be easy to merge changes to trunk is this case.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-965
19003	AIRAVATA-963	Nice work Sachith. What new integration tests are you planning to add?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-963
19004	AIRAVATA-963	Hi Sachith, How should I apply the patch ? Do I have to apply all the three files or do I have to only apply integration_tests_refactor_update3.patch ?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-963
19005	AIRAVATA-949	for Sha validation for file in `find . -type f -not -name '*.asc' -not -name '*.md5' -not -name '*.sh' -not -name '*.sha`` do gpg --print-md SHA512 ${file} \| diff - "${file}"".sha" | echo "testing : $file"".sha" done For MD5 validation gpg --print-md md5 ${file} \| diff - "${file}"".md5" | echo "testing : $file"".md5" I can't access the page So I posted the code here. By adding them all together we can automate the whole process with one script.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-949
19006	AIRAVATA-944	Hi Saminda, Shall I create a document which depicts how to use Airavata API to create and use gateways to submit a simple job? It's not mentioned anywhere and it is very useful. We need How-to complete documentation on how we can do the above.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-944
19007	AIRAVATA-921	Hi Vijeyendra, This is a big patch and I see too many directories and files within the root directory. This is first time I am seeing the monitoring contribution and unable to review it. Can you please explain this patch a bit?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-921
19008	AIRAVATA-921	Hi Vijayendra, I committed this path, but did not know how to use it. Can you please add a README and also wiki documentation about this module? -- Suresh	1	0	https://issues.apache.org/jira/browse/AIRAVATA-921
19009	AIRAVATA-916	I did(apply Shameera's patch), but that doesnt seem to fix it. I am kind of confused here. Why does the client side of this API need to show all these methods to build SOAP headers. From how I understand it, this should just show the methods which the user will directly interact with (the get/list and set methods).	1	1	https://issues.apache.org/jira/browse/AIRAVATA-916
19010	AIRAVATA-916	Hi [~shameera] No , even I was having the same issue ! [~subho] I didn't see any service change in this patch, where have you added the file?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-916
19011	AIRAVATA-906	Hi Shameera, Thanks for this patch. I am wondering why some schemas are using draft-03 and some draft-04, is there any compatibility reason? I see some of the optional values are not explicitly called out for. So should we make sure we add "required": false for all optional values? Lastly, what is the json schema equivalent of unbounded arrays?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-906
19012	AIRAVATA-899	Saminda, can you please take a look at how to make it work on 0.8? I am attaching the patch and i tried to make most of the changes but still not working.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-899
19013	AIRAVATA-899	Hi Raman, Can you check the attached patch ("patch-0.8.diff") with the 0.8 release tag?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-899

#	Requirement	Questions	Answered	Successfully Answered	URL
19014	AIRAVATA-898	well. [INFO] Clustering has been disabled [ERROR] A ClassNotFoundException error occurred while validating the message builder org.apache.airavata.json.StaxonJSONBuilder org.apache.axis2.deployment.DeploymentException: A ClassNotFoundException error occurred while validating the message builder org.apache.airavata.json.StaxonJSONBuilder at org.apache.axis2.deployment.DescriptionBuilder$2.run(DescriptionBuilder.java:306) at org.apache.axis2.java.security.AccessController.doPrivileged(AccessController.java:132) at org.apache.axis2.deployment.DescriptionBuilder.findAndValidateSelectorClass(DescriptionBuilder.java:291) at org.apache.axis2.deployment.DescriptionBuilder.processMessageBuilders(DescriptionBuilder.java:228) at org.apache.axis2.deployment.AxisConfigBuilder.processMessageBuilders(AxisConfigBuilder.java:754) at org.apache.axis2.deployment.AxisConfigBuilder.populateConfig(AxisConfigBuilder.java:218) at org.apache.axis2.deployment.DeploymentEngine.populateAxisConfiguration(DeploymentEngine.java:703) at org.apache.axis2.deployment.WarBasedAxisConfigurator.<init>(WarBasedAxisConfigurator.java:157) at org.apache.axis2.transport.http.AxisServlet.initConfigContext(AxisServlet.java:525) at org.apache.axis2.transport.http.AxisServlet.init(AxisServlet.java:443) at org.apache.catalina.core.StandardWrapper.initServlet(StandardWrapper.java:1228) at org.apache.catalina.core.StandardWrapper.loadServlet(StandardWrapper.java:1147) at org.apache.catalina.core.StandardWrapper.load(StandardWrapper.java:1043) at org.apache.catalina.core.StandardContext.loadOnStartup(StandardContext.java:4957) at org.apache.catalina.core.StandardContext$3.call(StandardContext.java:5284) at org.apache.catalina.core.StandardContext$3.call(StandardContext.java:5279) at java.util.concurrent.FutureTask$Sync.innerRun(FutureTask.java:303) at java.util.concurrent.FutureTask.run(FutureTask.java:138) at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:895) at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:918) at java.lang.Thread.run(Thread.java:680) Caused by: java.lang.ClassNotFoundException: org.apache.airavata.json.StaxonJSONBuilder at java.net.URLClassLoader$1.run(URLClassLoader.java:202) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:247) at java.lang.Class.forName0(Native Method) at java.lang.Class.forName(Class.java:171) at org.apache.axis2.util.Loader.loadClass(Loader.java:261) at org.apache.axis2.util.Loader.loadClass(Loader.java:229) at org.apache.axis2.deployment.DescriptionBuilder$2.run(DescriptionBuilder.java:298) ... 20 more After adding that you will not get the above error. But you will get the below error. [ERROR] de/odysseus/staxon/json/stream/JsonStreamFactory org.apache.axis2.deployment.DeploymentException: de/odysseus/staxon/json/stream/JsonStreamFactory at org.apache.axis2.deployment.AxisConfigBuilder.processMessageFormatters(AxisConfigBuilder.java:742) at org.apache.axis2.deployment.AxisConfigBuilder.populateConfig(AxisConfigBuilder.java:249) at org.apache.axis2.deployment.DeploymentEngine.populateAxisConfiguration(DeploymentEngine.java:703) at	1	1	https://issues.apache.org/jira/browse/AIRAVATA-898
19015	AIRAVATA-897	Hi, Here i have attached patch for this, please review and commit. Thanks, Shameera.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-897
19016	AIRAVATA-894	Hi devs, I have attached patch for this, Please review and commit. Thanks, Shameera.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-894
19017	AIRAVATA-894	Hi Shameera, I am sorry, I missed to apply this patch and proceed with others and it now conflicts. Can you please re-take the patch from latest sandbox area? Suresh	1	1	https://issues.apache.org/jira/browse/AIRAVATA-894
19018	AIRAVATA-893	Sanjaya, thank you for paying attention to this lingering issue. The patch looks good. One concern I have is duplication of the schema files in the distribution. Any time we duplicate the schemas, we risk keeping them in sync. Can you instead copy these files as a assembly task from the commons location? Let me know if you need help with. You can look at the distribution module on how the config files are copied, you can use similar ant tasks to do the same here.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-893
19019	AIRAVATA-893	Hi Sanjaya, Thanks for taking initiative to fix this. Nice work on preparing the patch. I have some suggestions which you can improve on it. Apart from the improvement Suresh has mentioned do you think you can update the patch to have a order of preference of selecting online/offline automatically? The reason is that unless we educate users of this setting in the property file ("offlineModeEnabled") users may not know that they are required to update it when they loose or gain Internet connectivity. So my suggestion is let, offlineModeEnabled=auto|false|true and when "offlineModeEnabled" is "auto" (or null) check if the schema url is accessible over Internet and if not use the local schema files which you've setup. Also inorder to dynamically copy the schema files during the build time use the maven dependency plugin to retrieve the "airavata-gfac-schema-utils" jar (which will be built before any of the distributions) and extract it to a temporary location and copy the schema files using an ant task in the maven ant plugin. (Let me know if you need any help doing this) Notice that the "modules/rest/webapp" and "modules/integration-tests" will also require the schema files.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-893
19020	AIRAVATA-893	Thanks Saminda! It is a great and convenient improvement to support auto. Following is my understanding about auto implementation. Please clarify if your understanding is different. Each time when a request comes, application tries to connect to online schema. If it fails to connect online schema, validates schema using offline. It's a combination of both online and offline validation where it falls back to offline validation in case of connection failure in online validation. I have modified the build to copy the schema from resources/schema in commons/gfac-schema module. I think "airavata-gfac-schema-utils" also copied the schema from the same location. Is there any particular reason such as possible post processing to copy it from the jar file ? I'll update the "modules/rest/webapp" and "modules/integration-tests" to use the same mechanism. I might need some help to test these modules. I have tested the solution only with XBaya and Airavata Server modules.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-893
19021	AIRAVATA-891	changed scaffolding to allow for bootstrap 3 [~sanchitagarwal108@gmail.com] : I have changed most of your code to conform to bootstrap3 except for globusFileTransfer.html and launchWorkflow.html. Could you make the changes for these files?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-891
19022	AIRAVATA-889	Sanchit, Can you please submit a svn patch? Based on the svn client you are using comand line or eclipse you will be able to google and find instructions how to take a patch.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-889
19023	AIRAVATA-889	Sanchit, index.html has a conflict, can you do a svn update, resolve conflicts and take a patch again?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-889
19024	AIRAVATA-889	Hi suresh I think due to the zero byte file, index file got corrupted, can you please apply EI_indexfile.patch to correct this.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-889
19025	AIRAVATA-889	Sanchit, can you please do a svn update and retake the patch? It conflicts after I committed Subho's patch.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-889
19026	AIRAVATA-889	Sanchit, this is much better now, but still has a conflict in app/views/exec.launchWorkflow.html. Can you do a svn up on this file and retake the patch?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-889

#	Requirement	Questions	Answered	Successfully Answered	URL
19027	AIRAVATA-889	Sanchit, I applied this patch now, but it complained on app/scripts/controllers/executionInterfaceCntl.js and app/scripts/controllers/HostDescriptor.js. Can you take an svn up and add any missing files into a new patch.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-889
19028	AIRAVATA-889	Sanchit, Can you please submit an updated patch for views taken from latest gsoc13 folder?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-889
19029	AIRAVATA-889	Sanchit, I applied the patch, but looks like the HostDescriptor.js is missing in svn. Can you add take an svn update and verify again? I have seen the RegistryDescriptors and HostDescriptor.js is being toggled back and forth, which is the one you intend to use? The controllers have RegistryDescriptors.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-889
19030	AIRAVATA-880	Hi Viknes, The patch looks good for the current requirements you have. Do you identify a notion of user roles or groups in your requiremnets?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-880
19031	AIRAVATA-880	That is correct. Right now we assume the user is admin and they can do everything (authoring descriptors/workflows/experiments). But as I understand you are not trying to retrieve a list of users who has "admin rights" right? Just all the users or the users who can run experiments? Could you breifly define the portal users who can access your dashboard and what users a portal user can querie to get data in the current design? And if you have any future directions to the design that would also be helpful if anyone wants to do follow up developement of your work.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-880
19032	AIRAVATA-870	Sry I do not understand the CORS filter to review this. Chathuri/Amila, can you please review this patch?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-870
19033	AIRAVATA-859	i noticed data is still going to Gram_Data table. Its not getting recorded to GFac_Job_Data. Do i need to change any configuration?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-859
19034	AIRAVATA-858	Hi Saminda, Can you please explain bit more about this issue. Regards Lahiru	1	0	https://issues.apache.org/jira/browse/AIRAVATA-858
19035	AIRAVATA-819	Heshan I agree this is a needed capability. For API usage, this will be trivial, because the client which calls the API can read such kind of a file and pass the comma seeprated values. How do you propose this XBaya? Add a new helper component?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-819
19036	AIRAVATA-814	What is needed to be done here? [~smarru]	0	0	https://issues.apache.org/jira/browse/AIRAVATA-814
19037	AIRAVATA-803	Are you still getting this error in the trunk?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-803
19038	AIRAVATA-798	Hi Suresh, Thank you for the clarification. I will look at the frameworks which you have highlights. Also I worked in this kind of UI related thing in my internship, so this will be a interesting task for me. BTW I have another small problem. As I understand XBaya communicates with Airavata Server in the back-end while we do UI related tasks. Do we have to concentrate on those functionality while implementing the UI composer? Thanks AndunSLG	1	0	https://issues.apache.org/jira/browse/AIRAVATA-798
19039	AIRAVATA-798	Hi, I have been studying this topic (in view of making a proposal at GSoC'13) and peeking into some of the Airavata code that might be relevant. I see a solution to the problem as decomposed into the following - 1. An interpreter for WSDL, so the code which we write integrates directly with the API which XBaya is using right now. There is freedom to choose to do this either purely in Javascript(frontend) or with a Java/Scala based backend built with something similar to the http://www.playframework.com/. 2. The actual HTML5 based frontend which is used to design and deploy workflows. As Andun pointed out, there are some nice alternatives to pick from (jQueryUI and jsPlumb seems a good alternative), or maybe something that uses the HTML5 Canvas itself like RaphaelJS (http://raphaeljs.com/) 3. Add any new calls to the Airavata API as required by the new frontend. Also side - by - side look at methods of expressing the workflow slightly differently from what it is now using a series of messages based off WSDL. To maybe a more abstract(JSON/REST maybe?) which makes design of the frontend a little easier. So, the backend from (1) can probably translate from the WSDL representation to the new one and back again. Am I missing something here? Or can some more features be added to this?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-798
19040	AIRAVATA-798	Hi, I think the purely JavaScript based implementation of Workflow designer will be good since we can maintain it without adding other complex dependencies. In my internship I have worked in WSO2(www.wso2.com). I have specially worked in there ESB which is built on top of Synapse. That ESB have a nice UI with this kind of a implementation which is built using purely jQuery. All the enterprises service integrations, message flows are designed using this sophisticated UI. This is a sample of that UI http://docs.wso2.org/wiki/download/attachments/16846268/31.png?version=1&modificationDate=1349302422000. They have the same approach Airavata Server and XBaya GUI communicates. Front-end UI communicate using API provided by back-end. So because of that exposure I think the approach is good. Using RapahelJS or KineticJS or YUI will be a handy choice to make our life easy. Suresh what do you think? What will be the long term plan to move the UI of Airavata to web based version? Thanks AndunSLG	1	0	https://issues.apache.org/jira/browse/AIRAVATA-798
19041	AIRAVATA-798	Hi Suresh, I have digged deep in to the source code via debugging to see how the XBaya create Workflows, persists them local file system, persists them in registry. I understood following things, - All the workflow design happens on top of the org.apache.airavata.xbaya.ui.graph.GraphCanvas - The workflow data represented via the org.apache.airavata.workflow.model.wf.Workflow (I think this is the DAG representation which you are talking about in Airavata Workflow Architecture in [1]) - Local file system persistences, retrieving done via org.apache.airavata.xbaya.core.generators.WorkflowFiler - Registry based persistence, retrieving has done via the REST client org.apache.airavata.rest.clientUserWorkflowResourceClient So If I focus on the task of implementing Web based Workflow Composer for Airavata, IMO we can reuse all three other than 1st one. 1st one have to be replace with a JavaScript based implementation. When that handles design part all the other things related to workflow can be done reusing the others. Since WorkflowFiler is a XBaya class, we have to refactor it to bring inside this web based implementation Is that approach is appropriate or do we have to come from the scratch for this web based implementation ? (Because you have mentioned that we are in the initial step to move to a web based UI for Airavata) Thanks AndunSLG [1] - http://people.apache.org/~smarru/papers/airavata-gce11.pdf	1	0	https://issues.apache.org/jira/browse/AIRAVATA-798

#	Requirement	Questions	Answered	Successfully Answered	URL
19042	AIRAVATA-770	I tried the patch and following are the challenges. 1. BESProvider is dependent on cog-jglobus-1.8.0-bc146 which is not the released version of JGlobus. GramProvider is dependent on cog-jglobus-1.8.0 and if i change the dependency to bc146 version then the logs are not coming for GramProvider. Strange part is BESProvider logs are fine. 2. To test BESProvider, i left the build to use cog-jglobus-1.8.0-bc146,. a). There is a dependency on CA certificate. Shahbaz can you provide location to download the CA certificate for Test and Juelich Unicore server. b). After adding the CA cert, i am able to submit the job fine but the job failed. 2013-02-26 07:07:28,066 [main] INFO BESProvider - Job https://zam1161v01.zam.kfa-juelich.de:8002/INTEROP1/services/BESActivity?res=1ec91f99-599c-4762-8cfb-30ed0b3cca71 is Failed. It would be great if the provider provide more details on the failure cause. 3. I tried to remove the dependency on specific version of cog and use the released version. I used cog-jglobus-1.8.0 version. Doing that i am getting following error. I tried to add dependency on released versions of bouncycastle but that did not help. Can you share the changes made to cog? That will help us to discuss solution to solve version problem. Caused by: java.io.IOException: Can not load the PEM private key: java.io.IOException: Can not load the PEM private key: no input data (empty source?) at eu.emi.security.authn.x509.impl.CertificateUtils.internalLoadPK(CertificateUtils.java:229) at eu.emi.security.authn.x509.impl.CertificateUtils.parsePEMPrivateKey(CertificateUtils.java:210) at eu.emi.security.authn.x509.impl.CertificateUtils.loadPEMKeystore(CertificateUtils.java:363) at eu.emi.security.authn.x509.impl.PEMCredential.<init>(PEMCredential.java:110) at eu.emi.security.authn.x509.impl.PEMCredential.<init>(PEMCredential.java:93) at org.apache.airavata.gfac.provider.impl.BESProvider.initSecurityProperties(BESProvider.java:201) ... 28 more	1	1	https://issues.apache.org/jira/browse/AIRAVATA-770
19043	AIRAVATA-770	Shahbaz, I tried to apply the patch but the provider failed while executing initSecurityProperties(jobExecutionContext) method. Attached the error logs file. I tried to remove the hadoop dependencies to see of this is ssl library issue but i am still getting the same error. Can you please take a look?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-770
19044	AIRAVATA-770	Hi Shahbaz, Does unicore submission really need all the extensions in eu.unicore.jsdl.extensions. can we use the build time generated jsdl beans and not use any unicore extensions? We probably should add the jsdl SMPD extension and get away right? - Thanks Suresh.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-770
19045	AIRAVATA-759	NOTICE and LICENSE files exist in the META-INF of the gsissh library jar. Do i need to add anything else?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-759
19046	AIRAVATA-716	I thought you need to have configuration for mysql in the authenticator.xml. I resolved the issues by adding mysql configuration to the authenticator.xml. Sorry for misunderstanding. Hence re-opening the issue.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-716
19047	AIRAVATA-657	Hi Shahbaz, thank you very much for the contribution, this is very useful. Can you please change the package name of the client only to org.apache.airavata.sample.bes or something like that and resubmit the patch? Thanks, Suresh	1	1	https://issues.apache.org/jira/browse/AIRAVATA-657
19048	AIRAVATA-657	Hi Shahbaz, Can we integrate this with jglobus 2.0 key store style GSI proxies and trusted certificates? Also, can we remove dependencies to - de.fzj.unicore.* packages? Since the goal for Airavata will be add standard BES/JSDL, unicore specific implementations will hinder with other BES providers. Is there any reason you will need to use the security package from eu.unicore.security instead of jaav X509?? We can eliminate lot of dependencies if we align these proxy generation (and key store style authentication) with standard java security. Thanks, Suresh	1	1	https://issues.apache.org/jira/browse/AIRAVATA-657
19049	AIRAVATA-656	Thanks for the pointer Shahbaz. I am planning to change the airavata sample to use xmlbeans as well. Can you please attach some sample JSDL's which will help to make sure I get them right? Please include MPI job samples as well.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-656
19050	AIRAVATA-656	Hi Shahbaz, I removed dependencies to jsdl-xmlbeans-2.1.0.jar and wsag4j-jsdl-types and directly added the OGF schemas of jsdl, jsdl-posix, jsdl-hpcpa and jsdl-spmd to airavata xmlbeans generation. This will ease us from adding a third party repo and associated license requirements. We still need to verify the legal rules of adding OGF schemas to the source repo, and I asked Apache Legal for verification - https://issues.apache.org/jira/browse/LEGAL-166. I understand from a practical stand point you mention Precommands, Postcommands, ResourceRequests will require custom extensions to JSDL. May be at a later point, we will revisit this. But for now, do we need any of these extensions? Can we get away with requiring attributes from - https://unicore.svn.sourceforge.net/svnroot/unicore/jsdl-xmlbeans/trunk/src/main/schema/jsdl-unicore.xsd and work straight of JSDL? This will help Airavata as it can create a generic BES/JSDL client. Thanks, Suresh	1	1	https://issues.apache.org/jira/browse/AIRAVATA-656
19051	AIRAVATA-651	Hi Suresh, Is this still valid issue? , As i can see, workflow interpreter module has been removed from build. Thanks, Shameera.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-651
19052	AIRAVATA-649	Patch applied.. thanks amila for the contribution. Lahiru	1	0	https://issues.apache.org/jira/browse/AIRAVATA-649
19053	AIRAVATA-630	It's working on Firefox, but not in Chrome. I applied the patch to trunk. Thanks for the contribution Viknes.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-630
19054	AIRAVATA-599	Can you explain a bit more about this.. right now users cna configure which registry they want tot connect by configuring repository.properties file inside bin directory and start xbaya again. Since this is a temporary solution until we comeup with registry service where you don't have to touch registry database credentials in XBaya GUI level.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-599
19055	AIRAVATA-599	I was looking how can i connect to remote server registry from the Xbaya interface as in 0.4. I can live with property file for now. A suggestion i have is in current case, we can connect to registry while loading Xbaya and load the objects to avoid user confusion.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-599
19056	AIRAVATA-599	How is the DB connection handled in case of JNLP? i need to add the connection parameters in JNLP file?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-599
19057	AIRAVATA-568	What happens if two clients send the same topic? This will be a bad feature to support, rather we should recommend the clients to query the server to return a topic id. More over, the API should discourage having a separate topic id, the experiment id should be unique across a deployment and for all tracking purposes including event topic, we should start using the experiment id.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-568
19058	AIRAVATA-568	I would say thats too complicated to support interoperability with external execution systems. I would even argue, two airavata instances should also not check for uniqueness across them. Its not about default vs advanced option. My concerns are mainly, this is a bad design choice to let the clients send a unique id to a server which has serious implications if the id is not unique.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-568
19059	AIRAVATA-568	What you are talking about should be handled using workflow instance id and that need to be unique for every run but not the topic id? Topic belongs to messaging system and we should not use that to provide uniqueness in orchestration system.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-568
19060	AIRAVATA-562	The sql script should be auto generated by JPA provider. It does not make sense to use JPA otherwise. We want to build REST service interface over these service and may have web interface also in future. Why not we use Spring Framework to configure JPA and REST services?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-562

#	Requirement	Questions	Answered	Successfully Answered	URL
19061	AIRAVATA-562	Sorry I didn't notice the patch much earlier. I created a new module to incorparate the new AiravataRegistry API JPA implementation @ [1]. Can you please resubmit the patch for that module? Code that goes in to [2] will become obsolete once we have the jpa implementation of the AiravataRegistry. (The [2] was created to provide a jpa provenance registry for the existing jackrabbit registry). Again, sorry for all the rework involved. 1. https://svn.apache.org/repos/asf/incubator/airavata/trunk/modules/registry/airavata-jpa-registry/ 2. https://svn.apache.org/repos/asf/incubator/airavata/trunk/modules/commons/provenance-registry/	1	1	https://issues.apache.org/jira/browse/AIRAVATA-562
19062	AIRAVATA-562	Saminda, Can you remove mysql driver dependency from the pom?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-562
19063	AIRAVATA-562	Chathuri, Thanks for the contribution, I was just reviewing this patch, it looks good and now aligns well with modules Saminda suggested. Just one change, within src you have redundant java folders, can you get rid of one and move everything up? Its like this - src/main/java/java/org instead of src/main/java/org.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-562
19064	AIRAVATA-561	Hi Amila, Did you intend the patch to be applied to trunk? Since the trunk is on the verge of a release preparation, I only added it to the sandbox area [1]. Also I had to skip the pom since there is a no master pom for sandbox. I also did not add in any h2 binary files, can you point them in the zip files so I can add them manually. Thanks, Suresh [1] - https://svn.apache.org/repos/asf/incubator/airavata/sandbox/airavata-registry-rest	1	0	https://issues.apache.org/jira/browse/AIRAVATA-561
19065	AIRAVATA-561	Patch applied.. thanks amila for the contribution.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-561
19066	AIRAVATA-560	Hi Amila, thanks for the contribution. Do you intend this to be applied in the trunk? I haven't seen a feature list yet for next release (0.5), but unless you plan to incorporate this into next release, I suggest to first having this in a sandbox and integrate it with trunk right after the release. I will let the 0.5 RM Saminda make a call on this though.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-560
19067	AIRAVATA-546	I skimmed through the code, its very nice. Currently the provenance (workflow produced data) is bottlenecking, so having this persistent solution will be great. Do you intend to test this with both Derby and mysql?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-546
19068	AIRAVATA-541	Hi Bhathiya, After the patch was applied there is a test failure in the registry-api module. It seems the following change causes the issue, Index: commons/registry-api/src/main/java/org/apache/airavata/registry/api/impl/AiravataJCRRegistry.java ========== --- commons/registry-api/src/main/java/org/apache/airavata/registry/api/impl/AiravataJCRRegistry.java (revision 1373906) +++ commons/registry-api/src/main/java/org/apache/airavata/registry/api/impl/AiravataJCRRegistry.java (working copy) @@ -177,7 +177,9 @@ Node node = serviceNode.getNode(serviceId); Property prop = node.getProperty(XML_PROPERTY_NAME); result = ServiceDescription.fromXML(prop.getString()); - } catch (Exception e) { + } catch (PathNotFoundException e) { + return null; + } catch (Exception e) { throw new ServiceDescriptionRetrieveException(e); I suppose you wanted to sync the two methods to be the same before you got rid of the "getServiceDesc" method? I managed to run the tests successfully after reverting the above change only. Can you verify the issue & give another patch to the existing code please.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-541
19069	AIRAVATA-541	The logic of returning null seems to be to point out that the requested service description is not available in the registry... IMO this is primitive... It is much cleaner to throw an Exception. Users can catch it and figure-out the reason why the service description was not retrieved (inner exception can give the exact reason). WDYT?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-541
19070	AIRAVATA-531	Milinda, Could you please update this jira once the patch from gsoc is applied to the trunk (after the 0.5 release).	1	0	https://issues.apache.org/jira/browse/AIRAVATA-531
19071	AIRAVATA-518	Hi Marlon, I would like to try on this task. There are few things i wanted to clarify. 1. Since Airavata now uses Git for version control i am bit unclear how to use svn checkin numbers in order to get versions. 2. I can see that Airavata.Client exposes API method getAPIVersion() which can be used to get version. Also there is a class called airavataAPIConstants in airavata-api-stubs package which also contains AIRAVATA_API_VERSION. Can't we utilize one of these two in displaying the version number? Thank you very much! Ishara	1	0	https://issues.apache.org/jira/browse/AIRAVATA-518
19072	AIRAVATA-517	Heshan, Looks like the jars are not being signed. com.sun.deploy.net.JARSigningException: Found unsigned entry in resource: http://localhost/~smarru/lib/xpp3_xpath-1.1.4c.jar at com.sun.javaws.security.SigningInfo.getCommonCodeSignersForJar(SigningInfo.java:411) at com.sun.javaws.security.SigningInfo.check(SigningInfo.java:276) at com.sun.javaws.LaunchDownload.checkSignedResourcesHelper(LaunchDownload.java:1741) at com.sun.javaws.LaunchDownload.checkSignedResources(LaunchDownload.java:1505) at com.sun.javaws.Launcher.prepareResources(Launcher.java:1281) at com.sun.javaws.Launcher.prepareAllResources(Launcher.java:634) at com.sun.javaws.Launcher.prepareToLaunch(Launcher.java:336) at com.sun.javaws.Launcher.prepareToLaunch(Launcher.java:208) at com.sun.javaws.Launcher.launch(Launcher.java:125) at com.sun.javaws.Main.launchApp(Main.java:451) at com.sun.javaws.Main.continueInSecureThread(Main.java:283) at com.sun.javaws.Main$1.run(Main.java:116) at java.lang.Thread.run(Thread.java:680) Suresh	1	0	https://issues.apache.org/jira/browse/AIRAVATA-517
19073	AIRAVATA-513	I have created a sample workflow to test "if-condition" and attached it here. But It is not working properly. Can anyone help me to correct the mistake I have done when designing this workflow. ?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-513
19074	AIRAVATA-513	Hasitha, not sure if I followed your question. Did you mean how to verify the T/F within the code or visually? The control connection (indicated with dotted line as opposed to data path solid lines) is the true path, and the bottom edge is connected to the false path. The config allows you to control the number of input edges and specify a expression using xpath.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-513
19075	AIRAVATA-477	Have these been addressed in separate Jiras? If so, we should link and close.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-477
19076	AIRAVATA-473	HI Raminder, What do you mean by this issue. Can you please tell how to reproduce this. Regards Lahiru	1	0	https://issues.apache.org/jira/browse/AIRAVATA-473
19077	AIRAVATA-473	Can you try to test adding a local provider in for-each loop? This was while back and i think generic provider was not able to call local provider service.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-473
19078	AIRAVATA-440	Did this get fixed?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-440
19079	AIRAVATA-372	Does this still happen? Since jackrabbit is not being used now?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-372
19080	AIRAVATA-371	commit #1328989 should fix the problem. It works fine on my machine: Windows 7 SP1 without Cygwin. Ross, could you verify it. Thanks,	1	0	https://issues.apache.org/jira/browse/AIRAVATA-371

#	Requirement	Questions	Answered	Successfully Answered	URL
19081	AIRAVATA-357	Exploration of Apache Whirr based Hadoop deployment on clouds in the first phase will provide us with the required knowledge about the complexity of the task and problem associated with moving from local clusters to cloud platforms. I think as the next phase we should consider about moving part of the job to cloud based on the increasing computing resource requirements. This will allow us to understand overhead introduced by this approach and complexities involved like resuming partially completed jobs. Finally we could explorer the area of bursting the computation into computational grids. I think that there will be a requirement for some heuristics based algorithm for selecting which computing resource to choose and when to burst into cloud. But I think we should implement the simplest case first and then go for the complex scenarios. I would like to know where this component will reside in Airavata architecture? Will this be a activity in a workflow or workflow activity + some daemon which handles provisioning of cloud resources and job? I'm bias towards the later approach, but I would like to know what is the most natural way when compared with the other Airavata components?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-357
19082	AIRAVATA-357	That sounds good and I'm impressed about the GSoc proposal you submitted.. Following are couple of comments regarding the steps you proposed in the proposal, that you would have to think through.. 1. What's the complexity of partial migration of Hadoop jobs and in what level are you planning on doing it (eg: adding cloud worker instance to the existing Hadoop cluster vs starting a new master in the cloud to migrate part of the data as a separate job vs ....)? 2. What kind of complexity/sophistication you are aiming with of the heuristics based resource selection algorithms? and how this needs to be different to the existing grid oriented solutions to cater for the clouds? What ideas/algorithms can we borrow from those existing systems? I'm bit optimistic about the step 3 you propose with regards to the limited time frame we have, but I'm fine with it given the research nature of that component. IMHO any contributions for step 3 would be very valuable, even if we weren't able to finish implementing that step completely during the GSOC.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-357
19083	AIRAVATA-356	Good patch Pulasthi. We haven't been actively maintaining WS-Notification, I quickly ran the test case but did not look deep enough. Will review it further and apply it over the weekend. I understand other samples of messenger are still using ant build so you went might have followed it with ant as well. Can you consider contributing a maven build for this sample so we can have a samples profile in ws messenger build?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-356
19084	AIRAVATA-345	hi, Im Sujan. I need to send proposal for gsoc 2012 for AIRAVATA-345. how can i do it? please let me know the steps what should i do ? Thank u email: hardsujan@gmail.com	1	0	https://issues.apache.org/jira/browse/AIRAVATA-345
19085	AIRAVATA-345	Here I wanted to tell u what I understand about requirements. In this dash board there are many jobs has to do as a sequential flow . there is expecting a automatic work flow to do this jobs as a flow. It need build.xml files ,ant , web server to share resources like zip files and etc. we need to create automatic system for send data via that web server and those data we can send as zip file . in the zip files there should be xml files with data. then we can make a tool for doing this job as a work flow . I need to send a proposal for this event. is it posible let me knw . here my Email address : hardsujan@gmail.com	1	0	https://issues.apache.org/jira/browse/AIRAVATA-345
19086	AIRAVATA-345	Thank you suresh for your email. according to AIRAVATA-345, what i understood is We have to create a tool to get the daily information about workflow of the application. say for an example If we consider there is a workflow to get the project details from each project of a company. These each project have daily updates , defects, redundancies , line of codes per class and etc. we can get these details by creating a tool. And also Using DBUnit testing we can check whether is this work flow running correctly. I would like to join Gsoc 2013 jira group. Let me know more information about project or tool. Im with my finla year in University of Moratuwa. Thank you.	1	0	https://issues.apache.org/jira/browse/AIRAVATA-345
19087	AIRAVATA-345	---------- Running org.apache.airavata.security.configurations.AuthenticatorConfigurationRe aderTest C:\airavata-0.6\modules\security\target\test-classes\. [main] WARN org.apache.airavata.security.util.PasswordDigester - Password hash m ethod is not configured. Setting default to plaintext. [main] WARN org.apache.airavata.security.util.PasswordDigester - Password hash m ethod is not configured. Setting default to plaintext. C:\airavata-0.6\modules\security\target\test-classes\. [main] WARN org.apache.airavata.security.util.PasswordDigester - Password hash m ethod is not configured. Setting default to plaintext. Tests run: 2, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 0.529 sec <<< FA ILURE! Running org.apache.airavata.security.userstore.JDBCUserStoreTest [main] WARN org.apache.airavata.security.util.PasswordDigester - Password hash m ethod is not configured. Setting default to plaintext. Tests run: 1, Failures: 0, Errors: 0, Skipped: 0, Time elapsed: 3.301 sec Running org.apache.airavata.security.userstore.LDAPUserStoreTest Tests run: 1, Failures: 0, Errors: 0, Skipped: 1, Time elapsed: 0.001 sec Running org.apache.airavata.security.userstore.SessionDBUserStoreTest Tests run: 2, Failures: 0, Errors: 0, Skipped: 0, Time elapsed: 0.718 sec Results : Failed tests: testInit(org.apache.airavata.security.configurations.Authenticat orConfigurationReaderTest) Tests run: 6, Failures: 1, Errors: 0, Skipped: 1 [INFO] ---------- [INFO] Reactor Summary: [INFO] [INFO] Airavata ..... SUCCESS [4:15.142s] [INFO] Airavata Commons ..... SUCCESS [0.067s] [INFO] Airavata GFac Schema Utilities ..... SUCCESS [2:53.496s] [INFO] Airavata Common Utilities ..... SUCCESS [3:15.216s] [INFO] Airavata Registry ..... SUCCESS [0.084s] [INFO] Airavata Registry API ..... SUCCESS [5.419s] [INFO] Airavata Workflow Model Parent ..... SUCCESS [0.065s] [INFO] Airavata Workflow Model Core ..... SUCCESS [35.816s] [INFO] Airavata WS Messenger ..... SUCCESS [0.098s] [INFO] Airavata Messenger Commons ..... SUCCESS [7.771s] [INFO] Airavata Messenger Client ..... SUCCESS [13.343s] [INFO] Airavata WS Monitor ..... SUCCESS [51.453s] [INFO] Airavata Workflow Execution Context ..... SUCCESS [9.836s] [INFO] Airavata Client ..... SUCCESS [27.562s] [INFO] Airavata Workflow Tracking ..... SUCCESS [31.631s] [INFO] Airavata GFac Core ..... SUCCESS [8:07.161s] [INFO] Airavata GFac Axis2 Interface ..... SUCCESS [39.509s] [INFO] Airavata Message Box ..... SUCCESS [11.099s] [INFO] Airavata Message Broker ..... SUCCESS [43.387s] [INFO] Airavata Workflow Component ..... SUCCESS [2.730s] [INFO] Airavata Workflow Component Node ..... SUCCESS [0.923s] [INFO] Airavata JPA Registry ..... SUCCESS [2:12.381s] [INFO] Airavata REST ..... SUCCESS [0.096s] [INFO] airavata-rest-mappings ..... SUCCESS [7.903s] [INFO] airavata-rest-client ..... SUCCESS [40.164s] [INFO] Airavata XBaya ..... SUCCESS [3:25.629s] [INFO] Airavata Registry Service ..... SUCCESS [6.301s] [INFO] Airavata Registry Test ..... SUCCESS [26.741s] [INFO] Airavata Security Implementation ..... FAILURE [18.395s] [INFO] airavata-rest-services ..... SKIPPED [INFO] airavata-rest-service-webapp ..... SKIPPED [INFO] Airavata Tools ..... SKIPPED [INFO] Airavata Phoebus Integration ..... SKIPPED [INFO] registry-tool ..... SKIPPED [INFO] Simple Math Service ..... SKIPPED [INFO] Complex Math Service ..... SKIPPED [INFO] Levenshtein Distance Service ..... SKIPPED [INFO] Airavata Embedded Tomcat Implementation ..... SKIPPED [INFO] Airavata Test Suite ..... SKIPPED [INFO] distribution ..... SKIPPED [INFO] Samples - Airavata Client Samples ..... SKIPPED [INFO] Samples - Workflow Create Application ..... SKIPPED [INFO] Samples - Workflow Execution ..... SKIPPED [INFO] Airavata server distribution ..... SKIPPED [INFO] Airavata client distribution ..... SKIPPED	1	0	https://issues.apache.org/jira/browse/AIRAVATA-345

#	Requirement	Questions	Answered	Successfully Answered	URL
19088	AIRAVATA-345	Hi Sujan, I am not a windows users so unfortunately cannot help much with this issues. Can you search in mail archives, there were issues reported and solutions on windows build. Also, can you please post any non-gsoc related debugging like this one to users mailing lists and raise jira if they are bugs?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-345
19089	AIRAVATA-345	1:48:12 AM org.apache.coyote.AbstractProtocol init INFO: Initializing ProtocolHandler ["http-bio-8080"] Mar 26, 2013 1:48:12 AM org.apache.catalina.core.StandardService startInternal INFO: Starting service Tomcat Mar 26, 2013 1:48:12 AM org.apache.catalina.core.StandardEngine startInternal INFO: Starting Servlet Engine: Apache Tomcat/7.0.22 [WARN] Password hash method is not configured. Setting default to plaintext. [INFO] No cache or cacheManager properties have been set. Authorization cache cannot be obtained. Mar 26, 2013 1:48:13 AM com.sun.jersey.api.core.PackagesResourceConfig init INFO: Scanning for root resource and provider classes in the packages: org.apache.airavata.services.registry.rest org.codehaus.jackson.jaxrs Mar 26, 2013 1:48:13 AM com.sun.jersey.api.core.ScanningResourceConfig logClasses INFO: Root resource classes found: class org.apache.airavata.services.registry.rest.resources.ProjectRegistryResource class org.apache.airavata.services.registry.rest.resources.DescriptorRegistryResource class org.apache.airavata.services.registry.rest.resources.ExperimentRegistryResource class org.apache.airavata.services.registry.rest.resources.BasicRegistryResouce class org.apache.airavata.services.registry.rest.resources.PublishWorkflowRegistryResource class org.apache.airavata.services.registry.rest.resources.ProvenanceRegistryResource class org.apache.airavata.services.registry.rest.resources.UserWorkflowRegistryResource class org.apache.airavata.services.registry.rest.resources.ConfigurationRegistryResource Mar 26, 2013 1:48:13 AM com.sun.jersey.api.core.ScanningResourceConfig logClasses INFO: Provider classes found: class org.codehaus.jackson.jaxrs.JacksonJsonProvider class org.codehaus.jackson.jaxrs.JsonMappingExceptionMapper class org.codehaus.jackson.jaxrs.JacksonJaxbJsonProvider class org.codehaus.jackson.jaxrs.JsonParseExceptionMapper Mar 26, 2013 1:48:13 AM com.sun.jersey.server.impl.application.WebApplicationImpl _initiate INFO: Initiating Jersey application, version 'Jersey: 1.13 06/29/2012 05:14 PM' [INFO] Clustering has been disabled [INFO] Deploying module: addressing-1.5.1 - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/modules/addressing-1.5.1.mar [INFO] Deploying module: mtompolicy-1.5.1 - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/modules/mtompolicy-1.5.1.mar [INFO] Deploying module: ping-1.5.1 - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/modules/ping-1.5.1.mar [INFO] Deploying module: metadataExchange-1.5.1 - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/modules/mex-1.5.1.mar [INFO] Deploying module: jaxws-1.5.1 - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/modules/axis2-jaxws-mar-1.5.1.mar [INFO] Deploying module: script-1.5.1 - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/modules/axis2-scripting-1.5.1.mar [INFO] Deploying module: soapmonitor-1.5.1 - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/modules/soapmonitor-1.5.1.mar [INFO] Deploying module: metadataExchange-1.5.1-impl - file:/home/learner/Downloads/apache-airavata-server-0.6/lib/mex-1.5.1-impl.jar [INFO] Deploying Web service: airavata-xbaya-gui-0.6.jar - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/services/airavata-xbaya-gui-0.6.jar Tue Mar 26 01:48:15 IST 2013 : Could not listen on port 1527 on host 0.0.0.0: java.net.BindException: Address already in use [INFO] Database already created for Registry! 37 airavata_data INFO [Thread-3] openjpa.Runtime - Starting OpenJPA 2.2.0 97 airavata_data INFO [Thread-3] openjpa.jdbc.JDBC - Using dictionary class "org.apache.openjpa.jdbc.sql.DerbyDictionary". [INFO] Deploying Web service: airavata-registry-service-0.6.jar - file:/home/learner/Downloads/apache-airavata-server-0.6/repository/services/airavata-registry-service-0.6.jar [INFO] Deploying	1	0	https://issues.apache.org/jira/browse/AIRAVATA-345
19090	AIRAVATA-345	Hi Sanchit, As you can see there was a port in use error - "Could not listen on port 1527 on host 0.0.0.0: java.net.BindException: Address already in use", so the registry did not get started. Can you free up the port and restart? Please take this to the mailing list so we do not clutter this JIRA. -- Suresh	1	1	https://issues.apache.org/jira/browse/AIRAVATA-345
19091	AIRAVATA-345	Hi Sanchit, It seems port 1527 is not available when the airavata server is starting. Can you make the default derby port (1527) available and start airavata server again. Also if you change the port information in airavata-server.properties for the database configuration, you will be able to run derby in a different port. Once the database is ready, you should be able to log in using XBaya without any issues. Let us know if you get more issues. Regards, Chathuri	1	0	https://issues.apache.org/jira/browse/AIRAVATA-345
19092	AIRAVATA-343	I went through your description and things mentioned by Surresh in mailing list. I have couple of questions regarding the scope of this feature. 1) is this feature targeting on a debugging module that integrated in backend of Airavata (in Workflow Interpreter) or debugging module in XBaya GUI ? According to my understanding it should contains both of them to provide better debugging framework. 2) This feature covers lots of debugging functionalities and I feel, time constraints will be a problem to address all these things. So what would be the initial requirements that fit for a GSOC project? So we can set a milestone plan. Any suggestions are very much appreciated.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-343
19093	AIRAVATA-343	Hasitha, Until my request for becoming a mentor for ASF was approved through GSoC Melange site please consider the following as few comments on the proposal [1] you've submitted. i. There are a few spelling & grammar mistakes ii. While I know what you meant, I feel some of the following areas are a little unclear. Perhaps consider revising those a little? a. it is not clear why current XBaya debugging cannot be directly associated with a backend execution of the workflow with a simple RMI mechanism. b. Might be a good idea to mention what you plan do to if you do not have enough time to implement the all the tasks you've mentioned. (Note that it's not crucial for all the tasks you've mentioned to be implemented) c. How do you plan to carry out the testing phase. A brief description would do. d. When reading the proposal I see several paragraphs that could have a subheading (for clarity purposes). iii. If you have time, a simple box diagram would be helpful to others to understand the architecture you are proposing. Overall I think the proposal looks good. Saminda 1. http://www.google-melange.com/gsoc/proposal/review/google/gsoc2012/hasithaaravinda/1	1	1	https://issues.apache.org/jira/browse/AIRAVATA-343

#	Requirement	Questions	Answered	Successfully Answered	URL
19094	AIRAVATA-343	Suresh April 6, 2012, 12:23 p.m. Hi Hasitha, I echo with Chris, great proposal but a bit ambitious. If you are motivated enough and have good time, you will be able to do all of it, but as you can see from the OODT jira, these can easliy result in half a dozen significant sub-tasks. I do not want to discourage you, but caution you to scope your project well enough. Even you can finish the phase 1, that by itself will be very helpful for the airavata user community. Airavata workflows are typically very long running in nature and tapping into existing executions not only assit the workflow debugging but opens up venue for other exhilerating opportunities like providing deferred (or just in time binding) of workflow inputs, provide workflow adaptability and so on. Very encouraging to see you are thinking beyond the current project. Few questions and comments: Why do you propose new debugging server as opposed to enhancing the workflow interpreter itself? You have a very elaborative development plan but a week testing plan, please try and elaborate on your test cases as well. You have a good granular details in the implementation section, but your schedule section is very course grained. Can you submit a more fine grained schedule? Over all, good proposal and an important topic. These capabilities not only improve airavata, but also saves valuable computational time on lage super computers and should have a very broad impact. Me and others in airavata community will be happy to guide you through out the project. Good luck with the proposal, Suresh	1	0	https://issues.apache.org/jira/browse/AIRAVATA-343
19095	AIRAVATA-340	Hello. I am interested in this topic as my GSoC project and currently I study Amazon SWF to understand how it works. As I see it now, the main goal of this project is to create set of tools allowing transportation of DAG Model of the workflow between XBaya and Amazon SWF (may include development of new pluggable compiler module or use one of already existing) and access to the resources. Plus it will need message exchange to monitor workflow execution. Is it correct, or I missed something? Best regards	1	0	https://issues.apache.org/jira/browse/AIRAVATA-340
19096	AIRAVATA-340	Hi Yaroslav, Sorry for the late reply. You got the requirement very correctly. As you already figured it out, there are two tasks for this project, first one is a Airavata XWF format to SWF compiler. Airavata has already XWF to BPEL, XWF to python (jython) compilers. So you can look into implementing XSF compilers. Second challenge is the monitoring progress. Airavata workflow monitoring schema is reasonably well defined and the WS Messenger client publish and subscribe API's have good examples bundled with. You should be able to easily figure out Amazon SWF monitoring information and translate into to equivalent Airavata messages. If it helps, we can provide a definition for the Airavata workflow format XWF. Also, this project has elastic goals which could be easily extended if you accomplish your goals early. Are you planning to submit a proposal on this task? Do you have any further questions?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-340
19097	AIRAVATA-339	For tracking purposes, please keep this jira updated. Suresh April 6, 2012, 5:36 a.m. Hi Pulasthi, A very well written proposal. Glad to see you put so much thought into implementation. You have nicely uncovered the layers in ws messenger where you need to make changes. Your time line is very comprehensive and reasonable. Few suggestions to further improve: WS Messenger is heavily used for topic based subscriptions right now and it was seamless to support it in both eventing and notification specs. Also, we had XPath based subscriptions which we are not actviely supporting. Can you write some details on how you would like to integrate AMQP in terms of topic vs content based subscriptions? Based on your initial exploration, which AMQP implementations you would like to evaluate this against? You have good number of testing time allocated which is very nice. Can you discuss some rough test plans and how you will validate existing functionality works the same. Can you squeeze into your time frames tasks to evaluate performance over head (if any)caused by adding an additional protocol support? Over all good proposal, keep it going, Suresh pulasthi April 6, 2012, 6:35 a.m. Hi Suresh I updated the proposal with your feedback. I am not able to provide a detailed answer to topic based and content based subscription because of the time limit to the deadline. i will have to do further reading to answer that accurately. thanks for the feedback. Thanks Pulasthi. Suresh April 6, 2012, 11:23 a.m. Pulashi, Your changes looks good. Glad to see you have the right insights into the problems, you should be able to figure out the implementation details as work on the project. The airavata dev community is very knowledgable on this topic and will provide guidance. Good luck with your proposal, Suresh	1	1	https://issues.apache.org/jira/browse/AIRAVATA-339
19098	AIRAVATA-339	Harsha: to my knowledge Airavata does not have a stand alone server in debugging mode for either services or for xbaya. That may be a topic for a project, post an email after you think a little more deeper on how you can accomplish it. Danushaka: Can you please look up a bit on if other apache projects have used RabbitMQ? MPL in general seem to be acceptable but you may want to verify emailing to legal-discuss@a.o. Also, can you propose your thoughts on which layers within Airavata or at API level should employ AMQP?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-339
19099	AIRAVATA-338	Hi Bhathiya, Can you attach the patches you mentioned in [1] to this jira. You should give your consent to be released those patches under Apache. You will agree to this when you attach the patch to the jira. [1] - http://summerwithairavata.blogspot.com/2012/08/how-to-run-airavata-with-new-sling.html	1	1	https://issues.apache.org/jira/browse/AIRAVATA-338

#	Requirement	Questions	Answered	Successfully Answered	URL
19100	AIRAVATA-338	[INFO] --- maven-compiler-plugin:2.3.2:compile (default-compile) @ airavata-common-registry-api --- [INFO] Compiling 10 source files to /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/target/classes [INFO] --- [ERROR] COMPILATION ERROR : [INFO] --- [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[184,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[186,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[187,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[188,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[189,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[190,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[191,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[192,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[193,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[194,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[196,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-api/src/main/java/org/apache/airavata/common/registry/api/impl/SlingRegistry.java:[197,0] class, interface, or enum expected [ERROR] /home/heshan/Dev/gsoc/airavata-sling/modules/commons/common-registry-	1	1	https://issues.apache.org/jira/browse/AIRAVATA-338
19101	AIRAVATA-338	Hi Heshan, I attached 2 patches. Patch-1 contains only new files and new dependencies. Patch-2 contains all content of Patch-1 and, all the modifications to references of registry class. So, for now I think adding Patch-1 to the trunk will be enough, because we have 2 more methods to implement (as mentioned in a previous comment). What's your idea of that? Can you please let me know what instructions are missing in the blog? Thank you.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-338
19102	AIRAVATA-338	Hi Bhathiya, There were some changes in the trunk during last Friday and I'm getting a compilation error when building with Patch-1. Can you resubmit the patch again with the required changes. If the build does not break with the "not implemented" methods, I'd say you go ahead and include it to the patch too. Let's move those changes to the trunk and then try to move on from there. So add everything that you done so far to a single patch and attach it to the jira. Furthermore, I noticed that you changed the XRegistryMigrator as well. You don't have to do the changes to the XRegistryMigrator now. Just focus on the basic functionality first. We'll update the tool later for Sling migration. In the "Step 3 (Configure External Repository)" of your post, you have not mentioned where to find such a file. I had to run the jackrabbit.sh to startup the jackrabbit instance and then use the generated configuration file for it (I did the modifications after shutting down the jackrabbit server started by the above mentioned script.)	1	1	https://issues.apache.org/jira/browse/AIRAVATA-338
19103	AIRAVATA-329	Hi Raman, When I'm migrating the applications, I'm saving an application to the Airavata registry if there aren't any applications with the same name in the Airavata registry. So, if there are multiple applications with the same name, I am only migrating the just the first application. In case of migrating applications with same name (with different namespaces) to Airavata registry, do you have any idea on how we should do this? Should we rename them? WDYT?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-329
19104	AIRAVATA-321	Hi Ye, Did you do an online build (mvn clean install) ? I do not see this error in my machine. Anyway I'm running a build with a clean repo to verify this. I will give you an update as soon as my build finishes.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-321
19105	AIRAVATA-315	Is there any authorization process to restrict / access on the basis of client type (user or admin)?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-315
19106	AIRAVATA-309	Can anybody please provide some documents regarding how to port entries in the old XRegistry to the new registry? (Where is the tool and how to use it)	1	0	https://issues.apache.org/jira/browse/AIRAVATA-309

#	Requirement	Questions	Answered	Successfully Answered	URL
19107	AIRAVATA-260	I ran the build with tests enabled and all the tests passed. I was able to successfully run all 3 tutorials but had a minor issue with the 10 minute tutorial. When the jce jar is removed I am getting the following exception. One thing we could do is to update the 10 minute tutorial to have the Temporary Directory location as a directory path (eg: /home/heshan/temp) instead of 'tmp' directory. Then the 10 minute tutorial will work fine. WDYT? org.apache.airavata.core.gfac.exception.ProviderException: Cannot make directory at org.apache.airavata.core.gfac.provider.impl.LocalProvider.makeFileSystemDir(LocalProvider.java:101) at org.apache.airavata.core.gfac.provider.impl.LocalProvider.makeDirectory(LocalProvider.java:111) at org.apache.airavata.core.gfac.provider.AbstractProvider.initialize(AbstractProvider.java:48) at org.apache.airavata.core.gfac.services.impl.AbstractSimpleService.execute(AbstractSimpleService.java:103) at org.apache.airavata.services.gfac.axis2.reciever.GFacMessageReciever.invokeApplication(GFacMessageReciever.java:300) at org.apache.airavata.services.gfac.axis2.reciever.GFacMessageReciever.access$000(GFacMessageReciever.java:92) at org.apache.airavata.services.gfac.axis2.reciever.GFacMessageReciever$1.run(GFacMessageReciever.java:160) 598602 [Thread-15] ERROR org.apache.airavata.services.gfac.axis2.reciever.GFacMessageReciever - Error invoking GFac Service org.apache.airavata.core.gfac.exception.ProviderException: Cannot make directory at org.apache.airavata.core.gfac.provider.impl.LocalProvider.makeFileSystemDir(LocalProvider.java:101) at org.apache.airavata.core.gfac.provider.impl.LocalProvider.makeDirectory(LocalProvider.java:111) at org.apache.airavata.core.gfac.provider.AbstractProvider.initialize(AbstractProvider.java:48) at org.apache.airavata.core.gfac.services.impl.AbstractSimpleService.execute(AbstractSimpleService.java:103) at org.apache.airavata.services.gfac.axis2.reciever.GFacMessageReciever.invokeApplication(GFacMessageReciever.java:300) at org.apache.airavata.services.gfac.axis2.reciever.GFacMessageReciever.access$000(GFacMessageReciever.java:92) at org.apache.airavata.services.gfac.axis2.reciever.GFacMessageReciever$1.run(GFacMessageReciever.java:160)	1	1	https://issues.apache.org/jira/browse/AIRAVATA-260
19108	AIRAVATA-173	Patch applied .. thanks saminda for the contribution.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-173
19109	AIRAVATA-166	Patch applied .. thanks heshan for the contribution. Regards Lahiru	1	0	https://issues.apache.org/jira/browse/AIRAVATA-166
19110	AIRAVATA-148	add the image files in AIRAVATA-148-patch-images-2.zip to "trunk/modules/xbaya-gui/src/main/resources/images" add the patch file "AIRAVATA-148-2.patch" to "trunk/modules". this patch adds the workflows to the browser & capability of refreshing a node in the tree.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-148
19111	AIRAVATA-143	Thanks Heshan for the contribution. The patch seems to have problem with spaces in front, can you retry creating the patch? I get the following: patch: **** malformed patch at line 17: Index: modules/xbaya-gui/src/main/java/org/apache/airavata/xbaya/globus/FileTransferConstants.java	1	0	https://issues.apache.org/jira/browse/AIRAVATA-143
19112	AIRAVATA-143	Hi Suresh, I am reattaching the patch herewith. This patch contains the fix for AIRAVATA-146 as well. Since, there was no publicly hosted jar for Transfer API, I moved the code of TransferAPI during testing. Therefore, this patch contains source for TransferAPI as well. If we manage to host the jar somewhere, I might be able to add it as a dependency in the build and use it. WDYT?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-143
19113	AIRAVATA-143	I uploaded the jar to maven repo: <dependency> <groupId>globus</groupId> <artifactId>transfer-api-client</artifactId> <version>0.10</version> </dependency> can you please submit one more patch without the source of the transferapi?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-143
19114	AIRAVATA-138	Hi Patanachai, Are you still working on this? I am trying to make a RC, if you want to add more test cases, can you please change the FixRelease to 0.2? if not resolve and close this issue? Suresh	1	1	https://issues.apache.org/jira/browse/AIRAVATA-138
19115	AIRAVATA-123	Thanks Saminda for contribution.. Patch applied ! Lahiru	1	0	https://issues.apache.org/jira/browse/AIRAVATA-123
19116	AIRAVATA-105	Heshan, Sorry we have not applied your patch in a while. Can you please sync the code and resubmit the patch?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-105
19117	AIRAVATA-104	Hi Saminda, I have applied AIRAVATA-104-UM-API-JackRabbit-Impl.patch but when I try to apply AIRAVATA-104-UM-API.patch I can see its not updated with latest code. Can you please take an svn update and provide a patch wiht latest code. Thank you saminda for the contribution. Lahiru	1	1	https://issues.apache.org/jira/browse/AIRAVATA-104
19118	AIRAVATA-104	Hi Lahiru, Does this mean I dont have to re-submit the patches?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-104
19119	AIRAVATA-102	This is really awesome .. why don't you send this to dev list ? Lahiru	1	0	https://issues.apache.org/jira/browse/AIRAVATA-102
19120	AIRAVATA-56	Thanks patanachai for reporting the problem. From your description, it looks like you already fixed the build, can you please contribute a patch?	1	0	https://issues.apache.org/jira/browse/AIRAVATA-56
19121	AIRAVATA-52	Hi, I am University Of Moratuwa student who is currently involved in research of creating Geo Science gateway for the Apache AIRAVATA. Since this is inline with our research area I am interested in this project to do as my GSOC 2013 project. AFAIK apache OODT is consisted of Product server/client, Query server/client and Catalog and Archive Server. Is this project directly involve with Gfac core? Or this extension is relevant to Xbaya too?. I would like to discuss more about this project. Any previous implementations of OODT CAS will be really helpful. Thank You !	1	0	https://issues.apache.org/jira/browse/AIRAVATA-52
19122	AIRAVATA-38	Can we add a parent pom at trunk? Advantage is to import all the project as single step using Eclipse M2 plugin.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-38
19123	AIRAVATA-38	What exactly needs to be done? (I do not know enough about Eclipse M2 plugin.)	0	0	https://issues.apache.org/jira/browse/AIRAVATA-38
19124	AIRAVATA-38	Could you update build instructions in website? Does it work for all versions of Eclipse, not just 3.7? I would like to test it.	0	0	https://issues.apache.org/jira/browse/AIRAVATA-38
19125	AIRAVATA-38	If the eclipse support and documentation is good enough for now, can we resolve this issue? Otherwise, can we explicitly state what more needs to be done?	0	0	https://issues.apache.org/jira/browse/AIRAVATA-38
19126	AIRAVATA-27	Also it looks like your patch adds/move/delete some files. Can you please give a list of the files that are deleting and a list of files that needs to be added.. thanks, Thilina	1	1	https://issues.apache.org/jira/browse/AIRAVATA-27
19127	AIRAVATA-27	Yes I had a look, please remove that, it's my IDE issue. Since this patch is kind of big I prefer committing it at this point, since thing are getting complex if we keep lengthy patches. Please remove messagebox/resources directory and commit. I have moved it some how my patch doesn't reflect that change it seems. WDYT ? Regards Lahiru	1	1	https://issues.apache.org/jira/browse/AIRAVATA-27
19128	AIRAVATA-11	Thans Patanachai for the patch. Can you also specify with the patches which directory that has to be applied? the main trunk or gfac. In general it will help us to review the patches if you can specify the location.	1	1	https://issues.apache.org/jira/browse/AIRAVATA-11
19129	AIRAVATA-6	The mockup looks good to me too. Can we shrink the size of the banner height without loosing clarity of the image?	1	1	https://issues.apache.org/jira/browse/AIRAVATA-6

#	Requirement	Questions	Answered	Successfully Answered	URL
19130	AMBARI-21437	Reopening. URL for this is wrong since missing cluster id http://c6404.ambari.apache.org:8080/api/v1/requests {noformat} 11 Jul 2017 23:24:58,659 WARN [ambari-action-scheduler] ExecutionCommandWrapper:230 - Unable to lookup the cluster by ID; assuming that there is no cluster and therefore no configs for this execution command: Cluster not found, clusterName=clusterID=-1 11 Jul 2017 23:24:58,659 WARN [ambari-action-scheduler] ExecutionCommandWrapper:230 - Unable to lookup the cluster by ID; assuming that there is no cluster and therefore no configs for this execution command: Cluster not found, clusterName=clusterID=-1 {noformat}	1	1	https://issues.apache.org/jira/browse/AMBARI-21437
19131	AMBARI-21360	[~sangeetar], looks like the patch was accidentally [added to git|https://git1-us-west.apache.org/repos/asf?p=ambari.git;a=blob;f=AMBARI-21360.patch;h=c26f3a031f8851d25b47c53d7dd2108ec7e5ce08;hb=a3681c0199f31511170770d454697206cbeeeda3]. This causes RAT check failure. Can you please remove? Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-21360
19132	AMBARI-21326	Reviewing the batch, what should happen to existing clusters that are missing {noformat} -XX:ReservedCodeCacheSize=256m{noformat} in hbase-env content? The current patch will only affect new clusters.	1	1	https://issues.apache.org/jira/browse/AMBARI-21326
19133	AMBARI-21325	Hey [~cmirashi], could you please add a review board link?	1	0	https://issues.apache.org/jira/browse/AMBARI-21325
19134	AMBARI-21206	server.log} 26 Jun 2017 21:36:34,523 ERROR [main] ModuleFileUnmarshaller:141 - Cannot parse /var/lib/ambari-server/resources/common-services/YARN/3.0.0.3.0/metainfo.xml 26 Jun 2017 21:36:34,524 ERROR [main] ModuleFileUnmarshaller:143 - The prefix "xsi" for attribute "xsi:type" associated with an element type "condition" is not bound. org.xml.sax.SAXParseException; systemId: file:/var/lib/ambari-server/resources/common-services/YARN/3.0.0.3.0/metainfo.xml; lineNumber: 84; columnNumber: 54; The prefix "xsi" for attribute "xsi:type" associated with an element type "condition" is not bound. at com.sun.org.apache.xerces.internal.util.ErrorHandlerWrapper.createSAXParseException(ErrorHandlerWrapper.java:203) at com.sun.org.apache.xerces.internal.util.ErrorHandlerWrapper.fatalError(ErrorHandlerWrapper.java:177) at com.sun.org.apache.xerces.internal.impl.XMLErrorReporter.reportError(XMLErrorReporter.java:400) at com.sun.org.apache.xerces.internal.impl.XMLErrorReporter.reportError(XMLErrorReporter.java:327) at com.sun.org.apache.xerces.internal.impl.XMLErrorReporter.reportError(XMLErrorReporter.java:284) at com.sun.org.apache.xerces.internal.impl.XMLNSDocumentScannerImpl.scanStartElement(XMLNSDocumentScannerImpl.java:308) at com.sun.org.apache.xerces.internal.impl.XMLDocumentFragmentScannerImpl$FragmentContentDriver.next(XMLDocumentFragmentScannerImpl.java:2784) at com.sun.org.apache.xerces.internal.impl.XMLDocumentScannerImpl.next(XMLDocumentScannerImpl.java:602) at com.sun.org.apache.xerces.internal.impl.XMLNSDocumentScannerImpl.next(XMLNSDocumentScannerImpl.java:112) at com.sun.org.apache.xerces.internal.impl.XMLDocumentFragmentScannerImpl.scanDocument(XMLDocumentFragmentScannerImpl.java:505) at com.sun.org.apache.xerces.internal.parsers.XML11Configuration.parse(XML11Configuration.java:841) at com.sun.org.apache.xerces.internal.parsers.XML11Configuration.parse(XML11Configuration.java:770) at com.sun.org.apache.xerces.internal.parsers.XMLParser.parse(XMLParser.java:141) at com.sun.org.apache.xerces.internal.parsers.AbstractSAXParser.parse(AbstractSAXParser.java:1213) at com.sun.org.apache.xerces.internal.jaxp.SAXParserImpl$JAXPSAXParser.parse(SAXParserImpl.java:643) at com.sun.xml.bind.v2.runtime.unmarshaller.UnmarshallerImpl.unmarshal0(UnmarshallerImpl.java:216) at com.sun.xml.bind.v2.runtime.unmarshaller.UnmarshallerImpl.unmarshal(UnmarshallerImpl.java:189) at javax.xml.bind.helpers.AbstractUnmarshallerImpl.unmarshal(AbstractUnmarshallerImpl.java:157) at javax.xml.bind.helpers.AbstractUnmarshallerImpl.unmarshal(AbstractUnmarshallerImpl.java:162) at javax.xml.bind.helpers.AbstractUnmarshallerImpl.unmarshal(AbstractUnmarshallerImpl.java:171) at javax.xml.bind.helpers.AbstractUnmarshallerImpl.unmarshal(AbstractUnmarshallerImpl.java:189) at org.apache.ambari.server.stack.ModuleFileUnmarshaller.unmarshal(ModuleFileUnmarshaller.java:136) at org.apache.ambari.server.stack.ModuleFileUnmarshaller.unmarshal(ModuleFileUnmarshaller.java:81) at	1	1	https://issues.apache.org/jira/browse/AMBARI-21206
19135	AMBARI-21206	[~ssathish@hortonworks.com] - can you explain to me when exactly this breaks? I have applied this patch again to a fresh trunk workspace, it builds fine, I managed to install ambari-server and the agents and got to the stage where I was about to deploy a cluster with HDFS and ZOOKEEPER when I hit this unrelated error: 29 Jun 2017 08:57:37,000 ERROR [ambari-client-thread-29] BaseManagementHandler:63 - Bad request received: The properties [ServiceInfo/desired_repository_version_id] specified in the request or predicate are not supported for the resource type Service.	1	0	https://issues.apache.org/jira/browse/AMBARI-21206

#	Requirement	Questions	Answered	Successfully Answered	URL
19136	AMBARI-21206	following error {code} 26 Jun 2017 21:36:34,523 ERROR [main] ModuleFileUnmarshaller:141 - Cannot parse /var/lib/ambari-server/resources/common-services/YARN/3.0.0.3.0/metainfo.xml 26 Jun 2017 21:36:34,524 ERROR [main] ModuleFileUnmarshaller:143 - The prefix "xsi" for attribute "xsi:type" associated with an element type "condition" is not bound. org.xml.sax.SAXParseException; systemId: file:/var/lib/ambari-server/resources/common-services/YARN/3.0.0.3.0/metainfo.xml; lineNumber: 84; columnNumber: 54; The prefix "xsi" for attribute "xsi:type" associated with an element type "condition" is not bound. at com.sun.org.apache.xerces.internal.util.ErrorHandlerWrapper.createSAXParseException(ErrorHandlerWrapper.java:203) at com.sun.org.apache.xerces.internal.util.ErrorHandlerWrapper.fatalError(ErrorHandlerWrapper.java:177) at com.sun.org.apache.xerces.internal.impl.XMLErrorReporter.reportError(XMLErrorReporter.java:400) at com.sun.org.apache.xerces.internal.impl.XMLErrorReporter.reportError(XMLErrorReporter.java:327) at com.sun.org.apache.xerces.internal.impl.XMLErrorReporter.reportError(XMLErrorReporter.java:284) at com.sun.org.apache.xerces.internal.impl.XMLNSDocumentScannerImpl.scanStartElement(XMLNSDocumentScannerImpl.java:308) at com.sun.org.apache.xerces.internal.impl.XMLDocumentFragmentScannerImpl$FragmentContentDriver.next(XMLDocumentFragmentScannerImpl.java:2784) at com.sun.org.apache.xerces.internal.impl.XMLDocumentScannerImpl.next(XMLDocumentScannerImpl.java:602) at com.sun.org.apache.xerces.internal.impl.XMLNSDocumentScannerImpl.next(XMLNSDocumentScannerImpl.java:112) at com.sun.org.apache.xerces.internal.impl.XMLDocumentFragmentScannerImpl.scanDocument(XMLDocumentFragmentScannerImpl.java:505) at com.sun.org.apache.xerces.internal.parsers.XML11Configuration.parse(XML11Configuration.java:841) at com.sun.org.apache.xerces.internal.parsers.XML11Configuration.parse(XML11Configuration.java:770) at com.sun.org.apache.xerces.internal.parsers.XMLParser.parse(XMLParser.java:141) at com.sun.org.apache.xerces.internal.parsers.AbstractSAXParser.parse(AbstractSAXParser.java:1213) at com.sun.org.apache.xerces.internal.jaxp.SAXParserImpl$JAXPSAXParser.parse(SAXParserImpl.java:643) at com.sun.xml.bind.v2.runtime.unmarshaller.UnmarshallerImpl.unmarshal0(UnmarshallerImpl.java:216) at com.sun.xml.bind.v2.runtime.unmarshaller.UnmarshallerImpl.unmarshal(UnmarshallerImpl.java:189) at javax.xml.bind.helpers.AbstractUnmarshallerImpl.unmarshal(AbstractUnmarshallerImpl.java:157) at javax.xml.bind.helpers.AbstractUnmarshallerImpl.unmarshal(AbstractUnmarshallerImpl.java:162) at javax.xml.bind.helpers.AbstractUnmarshallerImpl.unmarshal(AbstractUnmarshallerImpl.java:171) at javax.xml.bind.helpers.AbstractUnmarshallerImpl.unmarshal(AbstractUnmarshallerImpl.java:189) at org.apache.ambari.server.stack.ModuleFileUnmarshaller.unmarshal(ModuleFileUnmarshaller.java:136) at org.apache.ambari.server.stack.ModuleFileUnmarshaller.unmarshal(ModuleFileUnmarshaller.java:81) at	1	1	https://issues.apache.org/jira/browse/AMBARI-21206
19137	AMBARI-21206	[~ssathish@hortonworks.com] - I've been able to get the stage where I was about to deploy an HDP 3.0 cluster with HDFS and ZOOKEEPER. That would be past the point where you hit the error. Did you want to try a fresh build using my commit? I can't really see what else I can do because I don't see that error.	1	0	https://issues.apache.org/jira/browse/AMBARI-21206
19138	AMBARI-21205	[~t-dule], Can you add details to this JIRA? Enabling and disabling Kerberos as well as adding and deleting services are not _experimental_ features, they have been first class features of Ambari through many releases.	1	1	https://issues.apache.org/jira/browse/AMBARI-21205
19139	AMBARI-21191	[~afernandez] Thank you so much for reviewing my patch on reviewboard, and could you be so kind as to commit the patch for me please? Thanks in advance.	1	1	https://issues.apache.org/jira/browse/AMBARI-21191
19140	AMBARI-21191	[~wangjianfei], can you link the code review from here?	1	1	https://issues.apache.org/jira/browse/AMBARI-21191
19141	AMBARI-21191	[~wangjianfei], could you please fix the following unit test error? {noformat:title=https://builds.apache.org/job/Ambari-trunk-Commit/7652/} Tests run: 1, Failures: 0, Errors: 1, Skipped: 0, Time elapsed: 0.394 sec <<< FAILURE! - in org.apache.ambari.server.state.ServicePropertiesTest validatePropertySchemaOfServiceXMLs(org.apache.ambari.server.state.ServicePropertiesTest) Time elapsed: 0.387 sec <<< ERROR! org.apache.ambari.server.AmbariException: File /home/jenkins/jenkins-slave/workspace/Ambari-trunk-Commit/ambari-server/target/test-classes/TestAmbaryServer.samples/../../../src/main/resources/common-services/FLUME/1.4.0.2.0/configuration/flume-log4j.xml didn't pass the validation. Error message is : cvc-complex-type.2.4.b: The content of element 'property' is not complete. One of '{filename, deleted, final, on-ambari-upgrade, on-stack-upgrade, property-type, value-attributes, depends-on, property_depended_by, used-by}' is expected. at org.apache.ambari.server.state.ServicePropertiesTest.validatePropertySchemaOfServiceXMLs(ServicePropertiesTest.java:49) {noformat} CC [~afernandez]	1	1	https://issues.apache.org/jira/browse/AMBARI-21191
19142	AMBARI-21191	[~wangjianfei], thanks, {{ServicePropertiesTest}} is fine with the new patch. Now there are some Python test failures in {{ambari-server/src/test/python/stacks/2.0.6/FLUME/test_flume.py}}: {noformat} Total run:1168 Total errors:0 Total failures:7 {noformat} You can run Python tests only with: {noformat} mvn -am -pl ambari-server -DskipSurefireTests clean test {noformat} Can you please check?	1	1	https://issues.apache.org/jira/browse/AMBARI-21191
19143	AMBARI-21145	Hello Oliver, While processing the *_"path"_* of _*"input"*_ parameter from a <service>-logsearch-conf.xml file, is it possible to select multiple files matching the regex provided in the "path" or just the first match from the many options matching the regex gets selected for Logfeeder input? In my analysis only the 1st match from the "path" regex gets picked up. There is existing logic to check for wildcards, but this code is executed only for those files in the "notReadyList" of the InputManager class. Files in this list are put when no matching file is found for the "path" regex during config loading stage. However, for a regex that matches with at least one of the files in the location, the 1st match gets added to the "inputList" of InputManager class. Please share your input on my understanding of how this logic works. Thank you, Keta	1	0	https://issues.apache.org/jira/browse/AMBARI-21145

#	Requirement	Questions	Answered	Successfully Answered	URL
19144	AMBARI-21145	About the problem in the description: In fact Log Feeder doesn't support wild card characters in the log folder names. Our assumption was that if there are log files in separate folders the users would surely don't want to load them under the same log id (called type in the input configuration). So you can create a separate input elements for each one of them, and they will be loaded each under their own log id. Please confirm if you really want to load input files from different folders under the same log id, and then we can take a look at the issue, but consider the suggestions above. About the problem in comment #1: if the "tail" parameter is "true", then Log Feeder loads only the first file found for a path with wildcard characters in it. If you set the "tail" parameter to "false", then it will load all of them, but will also stop after it loaded them i.e. it won't wait for new entries to come. We are planning to modify this behavior, and make it possible to load all the files and follow the first one. This is good for rolled over log files, where the user wants to load all the previous parts, but only follow the actual one. However if you still want to load log files from different directories, or specify a path with wild card in it in order to track multiple active log files, and follow all of them then it is not good.	0	0	https://issues.apache.org/jira/browse/AMBARI-21145
19145	AMBARI-21123	Hi [~rzang], Would you please commit this patch to trunk(3.0.0) or whether need more +1 on review board ? Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-21123
19146	AMBARI-21098	Hello Hui, Do you have design doc on what you are trying to accomplish ?	1	0	https://issues.apache.org/jira/browse/AMBARI-21098
19147	AMBARI-21065	How about upgrade from 2.5 to 2.6 stack ? [~sseth]	0	0	https://issues.apache.org/jira/browse/AMBARI-21065
19148	AMBARI-21065	New property in HDP-2.6 should get auto added when upgrading from an earlier version. Is that so [~jonathan.hurley]/[~ncole@hortonworks.com]?	1	1	https://issues.apache.org/jira/browse/AMBARI-21065
19149	AMBARI-21057	[~dgrinenko], the patch introduced checkstyle errors, can you please fix? {noformat} [INFO] --- maven-checkstyle-plugin:2.17:check (checkstyle) @ ambari-server --- [INFO] Starting audit... [ERROR] /home/jenkins/jenkins-slave/workspace/Ambari-trunk-Commit/ambari-server/src/main/java/org/apache/ambari/server/controller/internal/StageResourceProvider.java:50:8: Unused import - org.apache.ambari.server.orm.dao.RequestDAO. [UnusedImports] [ERROR] /home/jenkins/jenkins-slave/workspace/Ambari-trunk-Commit/ambari-server/src/main/java/org/apache/ambari/server/upgrade/UpgradeCatalog251.java:36: Wrong order for 'org.slf4j.Logger' import. [ImportOrder] [ERROR] /home/jenkins/jenkins-slave/workspace/Ambari-trunk-Commit/ambari-server/src/test/java/org/apache/ambari/server/controller/internal/RequestStageContainerTest.java:21:15: Unused import - org.easymock.EasyMock.anyString. [UnusedImports] Audit done. {noformat}	1	1	https://issues.apache.org/jira/browse/AMBARI-21057
19150	AMBARI-21035	Minor tweaks: * wrap request object in {{"Clusters"}} * document {{"health_report"}} structure [~bsari], can you please review?	1	0	https://issues.apache.org/jira/browse/AMBARI-21035
19151	AMBARI-21016	[~yaolei], I do not think your solution covers all cases. Did you try the case where you give the user more permissions or try using direct REST API calls using cookies. The changes in [^AMBARI-21016.patch] only seem to cover the case in an Ambari UI scenario where the user is given less permissions. This causes the UI to hide operations that the user is no longer supposed to have. However I believe due to the stored session information, the backend still thinks the user has the removed permissions. This can be tested by giving the user more permissions, allowing your patch to show more options for the user. But I believe that the backend will return errors since the user is not authorized to perform the operations because the _old_ permission set is stored in session. I don't necessarily think that this is an issue since changing roles/groups/permissions for a user when they are logged in usually requires the user to log out and the log back in to realize the changes. I believe that UNIX and Linux works this way. So, I was not concerted about this _issue_. I wonder what other team members feel about this. [~u39kun], [~ncole@hortonworks.com], [~jonathan.hurley], [~stoader]	1	1	https://issues.apache.org/jira/browse/AMBARI-21016
19152	AMBARI-21016	Is there a way to invalidate where the session information is bound to the permissions so that you don't need to logout to see the changes take effect?	1	1	https://issues.apache.org/jira/browse/AMBARI-21016
19153	AMBARI-21016	I don't think you want to invalidate the entire session, however I am not entirely sure what else is stored in session. [~u39kun], do you know what is stored in session? i think the frequency at which this will occur is small enough that we can require the user to log out and then back in to get the changes. Else we will use a bunch of cycles trying to determine whether to refresh the user's authentication information or not for no reason.	1	1	https://issues.apache.org/jira/browse/AMBARI-21016
19154	AMBARI-21016	If the frequency is not very often, then why not just invalidate it and force a logout? Isn't that better than having a bunch of HTTP exceptions for unauthorized actions from the backend?	1	1	https://issues.apache.org/jira/browse/AMBARI-21016
19155	AMBARI-21016	Hi [~rlevas] Sorry to trouble you again. I happen to find the subclasses of BaseManagementHandler will return 403 if catching AuthorizationException, i wonder whether we can return 401. Are there some reasonse to do so? Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-21016
19156	AMBARI-21016	[~yaolei] {quote} I happen to find the subclasses of BaseManagementHandler will return 403 if catching AuthorizationException, i wonder whether we can return 401. Are there some reasonse to do so? {quote} I am not sure why we would need to change this? I believe that technically, if no user is authenticated we should return a 401 status with an appropriate {{WWW-Authenticate}} header; however, I believe the accepted standard on the web is to return a 403 status.	1	1	https://issues.apache.org/jira/browse/AMBARI-21016
19157	AMBARI-20994	[~bsari], the addendum patch addresses NPE when processing URI: {{/api/v1/services/AMBARI/components/AMBARI_AGENT}}. Can you please review?	1	0	https://issues.apache.org/jira/browse/AMBARI-20994
19158	AMBARI-20977	Hi [~afernandez] , Thank you so much for reviewing my patch on reviewboard, and could you be so kind as to commit the patch for me please? Thanks in advance	1	0	https://issues.apache.org/jira/browse/AMBARI-20977
19159	AMBARI-20951	[~bsari], the addendum patch is needed to prevent an NPE due to an accidentally removed line. Can you please review?	1	0	https://issues.apache.org/jira/browse/AMBARI-20951
19160	AMBARI-20827	Hi [~jonathan.hurley] and [~adoroszlai], would you push this to appropriate branches?	1	1	https://issues.apache.org/jira/browse/AMBARI-20827
19161	AMBARI-20813	SUCCESS: Integrated in Jenkins build Ambari-branch-2.5 #1423 (See [https://builds.apache.org/job/Ambari-branch-2.5/1423/]) AMBARI-20813. Convert ambari-web build process to use yarn package (yusaku: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=c80dc4db10e50b3b668188601d183226ca95f91d]) * (add) ambari-web/yarn.lock	1	0	https://issues.apache.org/jira/browse/AMBARI-20813
19162	AMBARI-20813	FAILURE: Integrated in Jenkins build Ambari-trunk-Commit #7328 (See [https://builds.apache.org/job/Ambari-trunk-Commit/7328/]) AMBARI-20813. Convert ambari-web build process to use yarn package (yusaku: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=2075f52e327e6bc624ec043dc87ea74672bcc7f5]) * (add) ambari-web/yarn.lock	1	0	https://issues.apache.org/jira/browse/AMBARI-20813
19163	AMBARI-20777	[~swagle] , [~dsen] Can you please take a look at addendum patch?	1	1	https://issues.apache.org/jira/browse/AMBARI-20777
19164	AMBARI-20739	Hi [~jonathan.hurley] Sorry to bother you again. If we can configure the script path for script-based dispatchers on web client, whether there are some security factors need to be take into accout. ex.the location of script,etc. Do you think we should limit the location of dispatch script?	1	1	https://issues.apache.org/jira/browse/AMBARI-20739

#	Requirement	Questions	Answered	Successfully Answered	URL
19165	AMBARI-20739	I don't think that limiting the location is a bad idea. It might be overkill, but I'm OK with it. Why don't we do something like this: - Add a new configuration property to ambari.properties which indicates the directory where alert scripts can live. Default it to /var/lib/ambari-server/resources/scripts (or something like that). - Any new script alert being created would need to exist in this location. - As an admin, I can change this if I really want to something like /net/share/ambari-server/scripts.	1	1	https://issues.apache.org/jira/browse/AMBARI-20739
19166	AMBARI-20739	Besides, if we can configure the dispatch script filename(location) in alert target, i think we will don't need ambari.dispatch-property.script anymore. Do you think so?	1	1	https://issues.apache.org/jira/browse/AMBARI-20739
19167	AMBARI-20727	[~mgergely], could you please fix this checkstyle error? {noformat} [ERROR] /home/jenkins/jenkins-slave/workspace/Ambari-trunk-Commit/ambari-server/src/main/java/org/apache/ambari/server/upgrade/UpgradeCatalog300.java:27: Wrong order for 'java.util.Collections' import. [ImportOrder] {noformat}	1	1	https://issues.apache.org/jira/browse/AMBARI-20727
19168	AMBARI-20549	frontend-maven-plugin 1.3 requires at least Maven 3.1. Could you please update the recommended version in the [wiki|https://cwiki.apache.org/confluence/display/AMBARI/Ambari+Development#AmbariDevelopment-ToolsneededtobuildAmbari]? Thanks.	1	0	https://issues.apache.org/jira/browse/AMBARI-20549
19169	AMBARI-20534	+1 for the patch. Can you retrigger the Hadoop QA run to see if it passes.	1	0	https://issues.apache.org/jira/browse/AMBARI-20534
19170	AMBARI-20534	[~jluniya] sure. would you mind if I combine this with the patch for AMBARI-20686 ? Technical you need both to build in a clean environment for ARM64 and PPC	1	0	https://issues.apache.org/jira/browse/AMBARI-20534
19171	AMBARI-20534	This doesn't appear to be committed. Can someone please help?	1	0	https://issues.apache.org/jira/browse/AMBARI-20534
19172	AMBARI-20356	Hi [~afernandez] Thank you so much for reviewing my patch on reviewboard, and could you commit the patch ? Thanks	1	0	https://issues.apache.org/jira/browse/AMBARI-20356
19173	AMBARI-20290	[~ncole@hortonworks.com]Thank you so much for reviewing my patch on reviewboard, and could you commit the patch ? Thanks a lot.	0	0	https://issues.apache.org/jira/browse/AMBARI-20290
19174	AMBARI-20220	Hello @Jonathan Hurley, can you tell me the exact location of the files to patch? I have a hard time finding them on my server.	1	0	https://issues.apache.org/jira/browse/AMBARI-20220
19175	AMBARI-20220	Hello, Can you give me the links of all the file to patch? I sent you a screenshot to show you where I am stuck.	1	0	https://issues.apache.org/jira/browse/AMBARI-20220
19176	AMBARI-20220	I don't think you can attach screenshots here... But aside from the one file I pointed out, the other file is a Java file which is build into the JAR, so you wouldn't be able to replace it. Why are you trying to take this as a patch on top of an install? If that's what you're doing then you only need to change ambari_server_main.py until you can do an official ambari-server upgrade.	1	0	https://issues.apache.org/jira/browse/AMBARI-20220
19177	AMBARI-20220	Thanks for the time you give me .. I'm new to HDP, I downloaded the virtual machine on VMware Fusion. I am a computer security student, and I am French. I try to set up KnoxSSO and when I restart the server with the command "ambari-server restart" I encounter the famous error. So I did some research that led me to you and your patch. I will try to update ambari-server. Otherwise can you tell me what I need to change on the file "ambari_server_main.py" or I have to apply the patch on it? Sorry for my English…	0	0	https://issues.apache.org/jira/browse/AMBARI-20220
19178	AMBARI-19996	{color:red}-1 overall{color}. Here are the results of testing the latest attachment http://issues.apache.org/jira/secure/attachment/12852624/AMBARI-19996.PATCH against trunk revision . {color:green}+1 @author{color}. The patch does not contain any @author tags. {color:red}-1 tests included{color}. The patch doesn't appear to include any new or modified tests. Please justify why no new tests are needed for this patch. Also please list what manual steps were performed to verify this patch. {color:green}+1 javac{color}. The applied patch does not increase the total number of javac compiler warnings. {color:green}+1 release audit{color}. The applied patch does not increase the total number of release audit warnings. {color:green}+1 core tests{color}. The patch passed unit tests in ambari-server. Test results: https://builds.apache.org/job/Ambari-trunk-test-patch/10576//testReport/ Console output: https://builds.apache.org/job/Ambari-trunk-test-patch/10576//console This message is automatically generated.	0	0	https://issues.apache.org/jira/browse/AMBARI-19996
19179	AMBARI-19928	[~vivekratnavel] / [~avijayan], can you guys review this?	1	1	https://issues.apache.org/jira/browse/AMBARI-19928
19180	AMBARI-19906	[~jmarron], this is a very useful feature. Will the tar file also be compressed?	1	0	https://issues.apache.org/jira/browse/AMBARI-19906
19181	AMBARI-19906	Sorry for the late reply. Thanks [~afernandez],[~mgergely][~oleewere] for the comments and suggestions. [~oleewere] I love the idea of having a separate component (infra-solr-manager) to schedule jobs and keep track of files in HDFS. Is there a jira regarding this. If so I would love to be part of this and collaborate to create the component. Meanwhile, would it be a good idea to use my proposed system as a temporary solution to download multiple logs?I understand Log Search is stateless. So when the log download is initiated, the UI can wait for the download to get completed (Show the progress of the task). Like Mike suggested, if the number of log messages is limited, can this be used?	1	1	https://issues.apache.org/jira/browse/AMBARI-19906
19182	AMBARI-19886	Should this be reverted when "llap" is disabled? {{putCapSchedProperty("yarn.scheduler.capacity.root.ordering-policy", "priority-utilization"}}	0	0	https://issues.apache.org/jira/browse/AMBARI-19886
19183	AMBARI-19805	[~odiachenko] -- (is / does it need to be) the addition of extra parameter backward-compatible to previous PXF releases ?	1	0	https://issues.apache.org/jira/browse/AMBARI-19805
19184	AMBARI-19719	[~sumitmohanty] - upgradeCatalog250 has been modified. Do I need to change something else?	1	1	https://issues.apache.org/jira/browse/AMBARI-19719
19185	AMBARI-19702	[~daijy], [~afernandez], this prevents Ambari Server start, and also causes unit test failures: {noformat} ERROR [main] ModuleFileUnmarshaller:141 - Cannot parse /var/lib/ambari-server/resources/stacks/HDP/2.5/upgrades/upgrade-2.6.xml ERROR [main] ModuleFileUnmarshaller:143 - Upgrade pack must contain pre-downgrade elements if pre-upgrade exists for processing component PIG/PIG {noformat} Could you please check?	1	0	https://issues.apache.org/jira/browse/AMBARI-19702
19186	AMBARI-19702	Sorry about this. Attach AMBARI-19702.3.patch. Can you check if this fixes the issue? Or if you want me to create a separate Jira, let me know.	1	0	https://issues.apache.org/jira/browse/AMBARI-19702
19187	AMBARI-19702	[~afernandez], could you please commit [^AMBARI-19702.3.patch], too, to resolve this issue? Thanks.	1	0	https://issues.apache.org/jira/browse/AMBARI-19702
19188	AMBARI-19702	Thanks - [~afernandez] Can you commit this on behalf of Daniel?	0	0	https://issues.apache.org/jira/browse/AMBARI-19702
19189	AMBARI-19702	Ping [~afernandez] Can we commit this on Daniel's behalf?	1	1	https://issues.apache.org/jira/browse/AMBARI-19702
19190	AMBARI-19702	[~avijayan], [~afernandez], can you help commit to 2.5.0 branch? Otherwise, we leave it in a half done status.	0	0	https://issues.apache.org/jira/browse/AMBARI-19702
19191	AMBARI-19557	Hi, [~afernandez] [~jluniya]. Thank you so much for reviewing my patch on reviewboard, and could you be so kind as to commit the patch for me please? Thanks in advance	1	1	https://issues.apache.org/jira/browse/AMBARI-19557
19192	AMBARI-19557	[~alicezhangxiao], is this also intended for branches 2.5 and trunk? I don't believe we're planning additional minor releases of Ambari 2.4	1	1	https://issues.apache.org/jira/browse/AMBARI-19557
19193	AMBARI-19449	[~sumitmohanty] How can I assign this jira to myself? Also would you mind giving it a review? Thanks!	1	1	https://issues.apache.org/jira/browse/AMBARI-19449

#	Requirement	Questions	Answered	Successfully Answered	URL
19194	AMBARI-19433	{color:red}-1 overall{color}. Here are the results of testing the latest attachment http://issues.apache.org/jira/secure/attachment/12846462/AMBARI-19433.patch against trunk revision . {color:green}+1 @author{color}. The patch does not contain any @author tags. {color:red}-1 tests included{color}. The patch doesn't appear to include any new or modified tests. Please justify why no new tests are needed for this patch. Also please list what manual steps were performed to verify this patch. {color:green}+1 javac{color}. The applied patch does not increase the total number of javac compiler warnings. {color:green}+1 release audit{color}. The applied patch does not increase the total number of release audit warnings. {color:green}+1 core tests{color}. The patch passed unit tests in ambari-server. Test results: https://builds.apache.org/job/Ambari-trunk-test-patch/9974//testReport/ Console output: https://builds.apache.org/job/Ambari-trunk-test-patch/9974//console This message is automatically generated.	0	0	https://issues.apache.org/jira/browse/AMBARI-19433
19195	AMBARI-19429	Is it too late to push this into 2.5 as well? Not sure how close we are to a 2.5 release but this shouldn't have any negative impacts.	1	0	https://issues.apache.org/jira/browse/AMBARI-19429
19196	AMBARI-19429	[~jluniya] I really like your idea of turning this into a management pack so we can have independent release cycles from Ambari. Can you, please, take a first pass at turning the stack into the pack (since I'm sure it'll take less time for you than it would for me). I can most definitely help with testing it, etc.	1	0	https://issues.apache.org/jira/browse/AMBARI-19429
19197	AMBARI-19429	[~rshaposhnik] Please can you apply the patch https://issues.apache.org/jira/secure/attachment/12858833/AMBARI-19429-mpack_trunk.patch, build the ODPi mpack and try to install ODPi with the built ODPi stack Instructions to build and deploy ODPi mpack: *Building ODPi Management Pack* {code} 1. cd contrib/management-packs/odpi-ambari-mpack 2. mvn clean package 3. ls -lh target/odpi-ambari-mpack-1.0.0.0-SNAPSHOT.tar.gz {code} The ODPi Mpack is built at target/odpi-ambari-mpack-1.0.0.0-SNAPSHOT.tar.gz *Installing ODPi Managment Pack* {code} 1. yum install ambari-server 2. ambari-server install-mpack --mpack=/path/to/odpi-ambari-mpack-1.0.0.0-SNAPSHOT.tar.gz --purge --verbose 3. ambari-server setup 4. ambari-server start {code} Installing the ODPi mpack with --purge flag will remove the HDP stack definition and slipstream in the ODPi stack definition. When you log into Ambari Web you will only see ODPi 2.0 stack as an option.	1	1	https://issues.apache.org/jira/browse/AMBARI-19429
19198	AMBARI-19429	[~jluniya] thanks for providing the patch! It seems to be working for the service that rely on common services definitions (HDFS and ZOOKEEPER) but fails for the custom services (like YARN and HIVE). The failure seems to have to do with agents not being synchronized with the stack definition. E.g. here's how the YARN fails: {noformat} Caught an exception while executing custom service command: <class 'ambari_agent.AgentException.AgentException'>: 'Script /var/lib/ambari-agent/cache/stacks/ODPi/2.0/services/YARN/package/scripts/nodemanager.py does not exist'; 'Script /var/lib/ambari-agent/cache/stacks/ODPi/2.0/services/YARN/package/scripts/nodemanager.py does not exist' Command failed after 1 tries {noformat} and indeed that script is missing. In fact, on agent side all of services are missing. E.g. {noformat} $ ls -l /var/lib/ambari-agent/cache/stacks/ODPi/2.0 total 4 drwxr-xr-x. 7 root root 4096 Mar 17 01:02 hooks {noformat} As you can see there are no services subdirectory there -- just hooks. Can you help figure out what's going on here?	1	0	https://issues.apache.org/jira/browse/AMBARI-19429
19199	AMBARI-19429	[~rshaposhnik] Did you start/restart ambari-server after installing the mpack? Can you delete the following and restart ambari-server and see if the issue goes away? In my patch, I accidentally had a .hash file, I am wondering if it could have caused an issue. {code} /var/lib/ambari-server/resources/stacks/ODPi/2.0/services/YARN/package/.hash /var/lib/ambari-server/resources/stacks/ODPi/2.0/services/YARN/package/archive.zip /var/lib/ambari-server/resources/stacks/ODPi/2.0/services/HIVE/package/.hash /var/lib/ambari-server/resources/stacks/ODPi/2.0/services/HIVE/package/archive.zip {code} If you still see this problem, can you share the ambari agent logs?	1	1	https://issues.apache.org/jira/browse/AMBARI-19429
19200	AMBARI-19429	[~jluniya] I followed the steps in your previous comments (which means I didn't have to restart the ambari server -- it was never started to begin with). But I restarted it just in case and it didn't help. I also though that the .hash and archive.zip would be a problem, but since I used latest patch from [~mradhakrishnan] I don't see them in the localtions you provided. E.g.: {noformat} $ rm /var/lib/ambari-server/resources/stacks/ODPi/2.0/services/YARN/package/.hash rm: cannot remove `/var/lib/ambari-server/resources/stacks/ODPi/2.0/services/YARN/package/.hash': No such file or directory $ rm /var/lib/ambari-server/resources/stacks/ODPi/2.0/services/YARN/package/archive.zip rm: cannot remove `/var/lib/ambari-server/resources/stacks/ODPi/2.0/services/YARN/package/archive.zip': No such file or directory $ rm /var/lib/ambari-server/resources/stacks/ODPi/2.0/services/HIVE/package/.hash rm: cannot remove `/var/lib/ambari-server/resources/stacks/ODPi/2.0/services/HIVE/package/.hash': No such file or directory $ rm /var/lib/ambari-server/resources/stacks/ODPi/2.0/services/HIVE/package/archive.zip rm: cannot remove `/var/lib/ambari-server/resources/stacks/ODPi/2.0/services/HIVE/package/archive.zip': No such file or directory {noformat} Attaching log files from an agent. Look for 'Script /var/lib/ambari-agent/cache/stacks/ODPi/2.0/services/YARN/package/scripts/application_timeline_server.py does not exist'	1	0	https://issues.apache.org/jira/browse/AMBARI-19429

#	Requirement	Questions	Answered	Successfully Answered	URL
19201	AMBARI-19429	definition (Roman Shaposhnik via (jluniya: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=075cecbff95842b86cb2420318742a7345da6ad1]) * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/files/hcatSmoke.sh * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/package/alerts/alert_nodemanagers_summary.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/hooks/before-INSTALL/scripts/hook.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/hooks/after-INSTALL/scripts/shared_initialization.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/hooks/before-ANY/scripts/hook.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/hooks/after-INSTALL/scripts/params.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/package/scripts/nodemanager_upgrade.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/kerberos.json * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/hive_server.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/hive_service_interactive.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/ZOOKEEPER/metainfo.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/webhcat_service_check.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/package/scripts/application_timeline_server.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/files/hiveTezSetup.cmd * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/hooks/before-START/files/fast-hdfs-resource.jar * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/configuration/hivemetastore-site.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/blueprints/singlenode-default.json * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/etc/hive-schema-0.12.0.mysql.sql * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/role_command_order.json * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/etc/hive-schema-0.12.0.oracle.sql * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/etc/hive-schema-0.13.0.mysql.sql * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/package/scripts/params_windows.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/mysql_users.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/hooks/before-START/scripts/rack_awareness.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/configuration/webhcat-site.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/package/files/validateYarnComponentStatusWindows.py * (add) contrib/management-packs/odpi-ambari-	0	0	https://issues.apache.org/jira/browse/AMBARI-19429
19202	AMBARI-19429	definition (Roman Shaposhnik via (jluniya: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=dd193e155f38ea6e079c4e157fa441c0a331ec5c]) * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/webhcat_service.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/configuration/yarn-site.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/assemblies/odpi-ambari-mpack.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/kerberos.json * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/hooks/before-START/scripts/shared_initialization.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/kerberos.json * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/alerts/alert_llap_app_status.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/configuration/yarn-env.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/etc/hive-schema-0.12.0.mysql.sql * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/etc/hive-schema-0.13.0.mysql.sql * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/metainfo.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/params.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/hooks/before-START/files/checkForFormat.sh * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/metainfo.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/configuration/webhcat-log4j.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/hive_service.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/templates/startHiveserver2.sh.j2 * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/MAPREDUCE2_metrics.json * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/configuration/capacity-scheduler.xml * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/hive_server.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/hive_interactive.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/files/hiveserver2Smoke.sh * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/package/scripts/yarn.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/package/templates/taskcontroller.cfg.j2 * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/HIVE/package/scripts/hive_server_interactive.py * (add) contrib/management-packs/odpi-ambari-mpack/src/main/resources/stacks/ODPi/2.0/services/YARN/package/alerts/alert_nodemanager_health.py * (add) contrib/management-packs/odpi-ambari-	0	0	https://issues.apache.org/jira/browse/AMBARI-19429
19203	AMBARI-19369	[~qinliu], what version of Ambari should this go in? It is marked for trunk, but I believe that is for Ambari 3.0.0, should this go in Ambari 2.5.2 (branch-2.5) as well?	1	1	https://issues.apache.org/jira/browse/AMBARI-19369

#	Requirement	Questions	Answered	Successfully Answered	URL
19204	AMBARI-19369	[~qinliu] I committed the patch to trunk (for Ambari 3.0.0): {noformat} commit 4aaf259e191344076a88391f5853da4bf85b8a80 Author: Qin Liu <qinliu5678@gmail.com> Date: Thu Jun 8 16:23:34 2017 -0400 {noformat} However, the patch fails to apply to branch-2.5: {noformat} error: ambari-server/src/main/resources/common-services/HBASE/2.0.0.3.0/configuration/hbase-env.xml: No such file or directory error: ambari-server/src/main/resources/common-services/HBASE/2.0.0.3.0/package/templates/hbase_master_jaas.conf.j2: No such file or directory error: ambari-server/src/main/resources/common-services/HBASE/2.0.0.3.0/package/templates/hbase_regionserver_jaas.conf.j2: No such file or directory error: ambari-server/src/main/resources/common-services/HDFS/3.0.0.3.0/package/scripts/hdfs.py: No such file or directory error: patch failed: ambari-server/src/main/resources/common-services/YARN/2.1.0.2.0/package/scripts/yarn.py:192 error: ambari-server/src/main/resources/common-services/YARN/2.1.0.2.0/package/scripts/yarn.py: patch does not apply error: ambari-server/src/main/resources/common-services/YARN/3.0.0.3.0/configuration-mapred/mapred-env.xml: No such file or directory error: ambari-server/src/main/resources/common-services/YARN/3.0.0.3.0/configuration/yarn-env.xml: No such file or directory error: ambari-server/src/main/resources/common-services/YARN/3.0.0.3.0/package/scripts/params_linux.py: No such file or directory error: ambari-server/src/main/resources/common-services/YARN/3.0.0.3.0/package/scripts/yarn.py: No such file or directory error: ambari-server/src/main/resources/common-services/YARN/3.0.0.3.0/package/templates/yarn_jaas.conf.j2: No such file or directory error: ambari-server/src/main/resources/stacks/HDP/3.0/services/HDFS/configuration/hadoop-env.xml: No such file or directory error: ambari-server/src/main/resources/stacks/HDP/3.0/services/YARN/configuration-mapred/mapred-env.xml: No such file or directory error: patch failed: ambari-server/src/test/python/stacks/2.0.6/HDFS/test_zkfc.py:381 error: ambari-server/src/test/python/stacks/2.0.6/HDFS/test_zkfc.py: patch does not apply {noformat} Can you update the patch for branch-2.5?	1	1	https://issues.apache.org/jira/browse/AMBARI-19369
19205	AMBARI-19249	Hi akovalenko, can you please provide me the information for handling the below scenario: I have added the below code at "/opt/apache-ambari-2.2.2-src/ambari-web/public/javascripts/app.js" under: "var hdp2properties = " section, like below: { "name": "ldap_ad_authentication", "value": "", "recommendedValue": "system", "options": [{ displayName: 'system', hidden: false }, { displayName: 'LDAP', hidden: false }, { displayName: 'AD', hidden: false },], "displayType": "radio button", "radioName": "ldap-ad-authentication", "serviceName": "AUTHSERVER", "filename": "ldap-ad-auth.xml", "category": "Advanced ldap-ad-auth", "index": 2 }, { "name": "content", "serviceName": "AUTHSERVER", "filename": "ldap-ad-auth.xml", "category": "Advanced ldap-ad-auth", "index": 3 }, I want to select the "system" radio option as default. How to set the default value? Also, if i try to select the radio options manually, i see the below error in the console: app.js:180958 Uncaught TypeError: Cannot read property '0' of null(…) Which is pointing to the "handleDBConnectionProperty" Also i need the test connection button to get enabled when we select the LDAP or AD radio options. The test connection should ideally take values from he config fields and should validate the connection to ldap server. Your inputs are highly appreciated.	1	0	https://issues.apache.org/jira/browse/AMBARI-19249
19206	AMBARI-19149	I think it's easier to review if the cleanup is done in steps, grouped by the type of change. [^AMBARI-19149-1.patch] is the first patch, which cleans up imports: * Remove unused imports * Get rid of {{*}} imports * Rearrange imports according to specified order * Add Checkstyle check for build-time validation Testing done: * ran Java unit tests {noformat} Tests run: 4797, Failures: 0, Errors: 0, Skipped: 37 {noformat} * ran checkstyle check: {noformat} [INFO] --- maven-checkstyle-plugin:2.17:check (checkstyle) @ ambari-server --- [INFO] Starting audit... Audit done. {noformat} (Note: RAT check fails on 2 files added for AMBARI-19137.) [~jonathan.hurley], [~ncole], [~lpuskas] This patch is too big for ReviewBoard (got HTTP 502). Could you please review here?	1	1	https://issues.apache.org/jira/browse/AMBARI-19149
19207	AMBARI-19149	[~jonathan.hurley], [~ncole]: * {{utility}} was introduced for AMBARI-14439 in non-{{org.apache.ambari}} group. It was also surprising to me, but didn't want to change that now. * The two separate {{dependency}} sections for {{com.puppycrawl.tools}} have different {{type}} and {{scope}}. We need the {{test-jar}} to be able to use {{BaseCheckTestSupport}} from the Checkstyle project. This is not distributed with their regular {{jar}}, since it's needed only for compiling and running custom checks (eg. {{AvoidTransactionalOnPrivateMethodsCheckTest}} in our case). * {{$checkstyle.version}} for the {{checkstyle/test-jar}} dependency, and {{<version>}} for the regular {{checkstyle}} dependency are inherited from the parent pom, ie. {{ambari-project/pom.xml}}. * There is no need to always compile the whole root project. However, when working with a single module (eg. {{ambari-server}}), one needs to get the dependencies in one of two ways: ** Always make dependencies, too, with the {{-am}} Maven flag. Note that Ambari Web, Logsearch, etc. are skipped, since those are not required for the building the server. ** Get the dependencies from some repository. In our case this means having to install {{utility}}, its parents, and some other Ambari projects in a local repository once. One needs to re-install them if the dependencies themselves are updated. * Now I've noticed that running {{mvn checkstyle:check}} performs the default checks instead of our custom list, because the configuration was done only for the execution tied to the test phase, not for the plugin itself. This is being fixed in AMBARI-19158. {noformat:title=Always making dependencies} $ mvn -am -pl ambari-server clean compile ... [INFO] Reactor Build Order: [INFO] [INFO] Ambari Main [INFO] Apache Ambari Project POM [INFO] Ambari Views [INFO] utility [INFO] ambari-metrics [INFO] Ambari Metrics Common [INFO] Ambari Server {noformat} {noformat:title=Installing once} $ mvn -am -pl ambari-server -DskipTests clean install $ cd ambari-server $ mvn -DskipTest clean test $ mvn checkstyle:check {noformat}	1	1	https://issues.apache.org/jira/browse/AMBARI-19149
19208	AMBARI-19149	So, where are we at with this? I think we wanted to: - Have the {{test}} phase fail if these style checks fail. - Allow developers to very easily run the checks directly (ie {{mvn checkstyle:check}} from {{ambari-server}}) and just have it work.	1	1	https://issues.apache.org/jira/browse/AMBARI-19149
19209	AMBARI-19149	Yep, I confirmed that AMBARI-19158 fixes the issues we were seeing running checkstyle. Thanks Attila! What are your thoughts about creating a new Jira to track moving the {{utility}} project under {{org.apache.ambari}} ?	1	1	https://issues.apache.org/jira/browse/AMBARI-19149
19210	AMBARI-19126	Hi [~rkamath], the patch causes build failure due to missing Apache license in {{ambari-server/src/main/resources/common-services/ZEPPELIN/0.6.0.2.5/package/scripts/spark2_config_template.py}}. Can you please check?	1	0	https://issues.apache.org/jira/browse/AMBARI-19126
19211	AMBARI-19079	[~forestsissi]... {quote} if we use the ibm jdk,it should use the jdk's kinit {quote} I disagree with this. Ambari will install the MIT krb5 client libs (ie, krb5-workstation) on the hosts, so it will be available for use when needed. Why do you think Ambari should ignore the {{kinit}} provided by this package and use one provided by a JVM that may or may not be supported?	1	0	https://issues.apache.org/jira/browse/AMBARI-19079
19212	AMBARI-19079	Has it ever be considered that supporting Ambari deployment on IBM PPC?	1	1	https://issues.apache.org/jira/browse/AMBARI-19079

#	Requirement	Questions	Answered	Successfully Answered	URL
19213	AMBARI-19079	[~forestsissi]... I am not sure if the IBM PPC platform is supported by Ambari. As far as I know, the following platforms are supported by Ambari 2.4.2: * Red Hat Enterprise Linux (RHEL) v7.x * Red Hat Enterprise Linux (RHEL) v6.x * CentOS v7.x * CentOS v6.x * Debian v7.x * Oracle Linux v7.x * Oracle Linux v6.x * SUSE Linux Enterprise Server (SLES) v12 SP1 (HDP 2.5) * SUSE Linux Enterprise Server (SLES) v11 SP4 (HDP 2.2 and later) * SUSE Linux Enterprise Server (SLES) v11 SP3 (HDP 2.2 and later) * SUSE Linux Enterprise Server (SLES) v11 SP1 (HDP 2.2) * Ubuntu Precise v12.04 * Ubuntu Trusty v14.04 In any case, if the MIT Kerberos client packages are available on the IBM PPC platform all should be good. Else there may be issues other than just the arguments passed to kinit. What kerberos packages are available for this platform?	1	0	https://issues.apache.org/jira/browse/AMBARI-19079
19214	AMBARI-19079	[~forestsissi]... What is the OS being run on the IBM PPC platform? Is it IBM LinuxONE? bq. HDP must provide versions that support PPC machine, in other words, HDP needs repack within JDK environment on PPC machine; HDP does not package a JVM. It is a set of binaries that get installed expecting certain requirement to be met - one if which is a compatible JVM being installed. bq. Several configurations of service components need to be edited, such as HBase, Metrics. For example, kind of com.sun needs to be edited as com.ibm. Wow.. this seems to be an issue. Can you use OpenJDK instead? bq. Kinit's path in Ambari needs to be set as configurable, because IBM JDK contains self-matching Kinit,applying Kinit of Kerberos will cause conflict. However, for x86 PC, it is fine to keep using Kinit under Kerberos path. The path to kinit and the other Kerberos command line utilities are configurable. Just the options passed to them are not. I think your issue is becoming a bit more clear to me and maybe we should provide a way to customize calls to the Kerberos utilities - kinit, klist, kadmin. However it will take some work to determine the best way to do this. Can you provide _help_/_usage_ info for kinit, klist, and kadmin? For example: {noformat} Usage: kinit [-V] [-l lifetime] [-s start_time] [-r renewable_life] [-f | -F] [-p | -P] -n [-a | -A] [-C] [-E] [-v] [-R] [-k [-t keytab_file]] [-c cachename] [-S service_name] [-T ticket_armor_cache] [-X <attribute>[=<value>]] [principal] options: -V verbose -l lifetime -s start time -r renewable lifetime -f forwardable -F not forwardable -p proxiable -P not proxiable -n anonymous -a include addresses -A do not include addresses -v validate -R renew -C canonicalize -E client is enterprise principal name -k use keytab -t filename of keytab to use -c Kerberos 5 cache name -S service -T armor credential cache -X <attribute>[=<value>] {noformat}	1	0	https://issues.apache.org/jira/browse/AMBARI-19079
19215	AMBARI-19038	[~ghugare], I believe that this JIRA has been orphaned. However the topic has been restarted and I have a patch that contains much of the code from the original patch, plus changes for Ambari 2.5. Do you mind if I take this JIRA over?	1	1	https://issues.apache.org/jira/browse/AMBARI-19038
19216	AMBARI-19019	[~jeffreyr97] What is the status of this JIRA? Can this be moved to 3.0.0 if it still needs some work?	0	0	https://issues.apache.org/jira/browse/AMBARI-19019
19217	AMBARI-18944	Hi [~oleewere] ,[~rlevas] [~adoroszlai] [~jonathan.hurley], Thank you so much for reviewing my patch on reviewboard, and could you be so kind as to commit the patch for me please? Thanks in advance.	1	0	https://issues.apache.org/jira/browse/AMBARI-18944
19218	AMBARI-18892	Hi [~rlevas] Thank you so much for reviewing my patch on reviewboard, and could you commit the patch ? Thanks	1	0	https://issues.apache.org/jira/browse/AMBARI-18892
19219	AMBARI-18888	Reverting this from trunk/branch-2.5 because it's causing agent heartbeats to fail on commands executed before a cluster is created: {code} 02 Dec 2016 15:30:55,079 WARN [qtp-ambari-agent-42] AgentResource:136 - Error in HeartBeat org.apache.ambari.server.ClusterNotFoundException: Cluster not found, clusterName=null at org.apache.ambari.server.state.cluster.ClustersImpl.getCluster(ClustersImpl.java:277) at org.apache.ambari.server.agent.HeartBeatHandler.sendCommands(HeartBeatHandler.java:335) at org.apache.ambari.server.agent.HeartBeatHandler.handleHeartBeat(HeartBeatHandler.java:298) at org.apache.ambari.server.agent.rest.AgentResource.heartbeat(AgentResource.java:130) at sun.reflect.GeneratedMethodAccessor114.invoke(Unknown Source) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at com.sun.jersey.spi.container.JavaMethodInvokerFactory$1.invoke(JavaMethodInvokerFactory.java:60) at com.sun.jersey.server.impl.model.method.dispatch.AbstractResourceMethodDispatchProvider$TypeOutInvoker._dispatch(AbstractResourceMethodDispatchProvider.java:185) at com.sun.jersey.server.impl.model.method.dispatch.ResourceJavaMethodDispatcher.dispatch(ResourceJavaMethodDispatcher.java:75) at com.sun.jersey.server.impl.uri.rules.HttpMethodRule.accept(HttpMethodRule.java:302) at com.sun.jersey.server.impl.uri.rules.RightHandPathRule.accept(RightHandPathRule.java:147) at com.sun.jersey.server.impl.uri.rules.ResourceClassRule.accept(ResourceClassRule.java:108) at com.sun.jersey.server.impl.uri.rules.RightHandPathRule.accept(RightHandPathRule.java:147) at com.sun.jersey.server.impl.uri.rules.RootResourceClassesRule.accept(RootResourceClassesRule.java:84) {code}	0	0	https://issues.apache.org/jira/browse/AMBARI-18888
19220	AMBARI-18846	[~Tim Thorpe] {{branch-2.5}} still has this problem: {code} == FAIL: test_update_directory_archive (TestResourceFilesKeeper.TestResourceFilesKeeper) ---------- Traceback (most recent call last): File "/Users/jhurley/src/apache/ambari/ambari-common/src/test/python/mock/mock.py", line 1199, in patched return func(*args, **keywargs) File "/Users/jhurley/src/apache/ambari/ambari-server/src/test/python/TestResourceFilesKeeper.py", line 191, in test_update_directory_archive self.assertTrue(write_hash_sum_mock.called) AssertionError: False is not true ---------- Ran 270 tests in 7.216s FAILED (failures=1) ---------- Failed tests: FAIL: test_update_directory_archive (TestResourceFilesKeeper.TestResourceFilesKeeper) ---------- Traceback (most recent call last): File "/Users/jhurley/src/apache/ambari/ambari-common/src/test/python/mock/mock.py", line 1199, in patched return func(*args, **keywargs) File "/Users/jhurley/src/apache/ambari/ambari-server/src/test/python/TestResourceFilesKeeper.py", line 191, in test_update_directory_archive self.assertTrue(write_hash_sum_mock.called) AssertionError: False is not true {code} Should this patch be reverted from there as well?	1	0	https://issues.apache.org/jira/browse/AMBARI-18846
19221	AMBARI-18792	Need input on what the correct value for the configuration file is. Should this be changed in hiveserver-site explicitly?	1	1	https://issues.apache.org/jira/browse/AMBARI-18792
19222	AMBARI-18792	Is this good to commit?	1	1	https://issues.apache.org/jira/browse/AMBARI-18792

#	Requirement	Questions	Answered	Successfully Answered	URL
19223	AMBARI-18792	org.apache.ambari.server.upgrade.UpgradeCatalog250.updateHiveLlapConfigs(UpgradeCatalog250.java:212) at org.apache.ambari.server.upgrade.UpgradeCatalog250$$EnhancerByCGLIB$$48cf24f9.CGLIB$updateHiveLlapConfigs$6(<generated>) at org.apache.ambari.server.upgrade.UpgradeCatalog250$$EnhancerByCGLIB$$48cf24f9$$FastClassByCGLIB$$8832139.invoke(<generated>) at org.easymock.cglib.proxy.MethodProxy.invokeSuper(MethodProxy.java:228) at org.easymock.internal.ClassProxyFactory$MockMethodInterceptor.intercept(ClassProxyFactory.java:94) at org.apache.ambari.server.upgrade.UpgradeCatalog250$$EnhancerByCGLIB$$48cf24f9.updateHiveLlapConfigs(<generated>) at org.apache.ambari.server.upgrade.UpgradeCatalog250.executeDMLUpdates(UpgradeCatalog250.java:138) at org.apache.ambari.server.upgrade.UpgradeCatalog250$$EnhancerByCGLIB$$48cf24f9.CGLIB$executeDMLUpdates$11(<generated>) at org.apache.ambari.server.upgrade.UpgradeCatalog250$$EnhancerByCGLIB$$48cf24f9$$FastClassByCGLIB$$8832139.invoke(<generated>) at org.easymock.cglib.proxy.MethodProxy.invokeSuper(MethodProxy.java:228) at org.easymock.internal.ClassProxyFactory$MockMethodInterceptor.intercept(ClassProxyFactory.java:94) at org.apache.ambari.server.upgrade.UpgradeCatalog250$$EnhancerByCGLIB$$48cf24f9.executeDMLUpdates(<generated>) at org.apache.ambari.server.upgrade.UpgradeCatalog250Test.testExecuteDMLUpdates(UpgradeCatalog250Test.java:239) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:47) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:50) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:222) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.junit.runner.JUnitCore.run(JUnitCore.java:157) at com.intellij.junit4.JUnit4IdeaTestRunner.startRunnerWithArgs(JUnit4IdeaTestRunner.java:69) at com.intellij.rt.execution.junit.JUnitStarter.prepareStreamsAndStart(JUnitStarter.java:234) at com.intellij.rt.execution.junit.JUnitStarter.main(JUnitStarter.java:74) at	0	0	https://issues.apache.org/jira/browse/AMBARI-18792
19224	AMBARI-18736	Hi [~vbrodetskyi], this commit causes error in unit tests: https://builds.apache.org/job/Ambari-trunk-Commit/6127/testReport/org.apache.ambari.server.state/ServicePropertiesTest/validatePropertySchemaOfServiceXMLs/ Could you please fix?	1	0	https://issues.apache.org/jira/browse/AMBARI-18736
19225	AMBARI-18694	[~xiaobingo] can you please check if the fix missed file ambari-server/src/main/resources/stacks/HDP/2.3/services/HDFS/configuration/hadoop-env.xml ? Also can you please upload the patch at https://reviews.apache.org/dashboard/ for ambari project ?	1	1	https://issues.apache.org/jira/browse/AMBARI-18694
19226	AMBARI-18650	[~sumitmohanty] What is the status of this JIRA? Can this be moved out?	1	0	https://issues.apache.org/jira/browse/AMBARI-18650
19227	AMBARI-18650	Thanks @sumitmohanty. When was that API introduced? 2.5?	0	0	https://issues.apache.org/jira/browse/AMBARI-18650
19228	AMBARI-18638	Hi [~aantonenko], can the patch be pushed to 2.5 as well?	1	1	https://issues.apache.org/jira/browse/AMBARI-18638

#	Requirement	Questions	Answered	Successfully Answered	URL
19229	AMBARI-18573	Bootstrap3 framework (akovalenko) (akovalenko: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=2a77e39ca78d157c3286a317325ac6cba03670cb]) * (edit) ambari-web/test/views/main/admin/highAvailability/progress_view_test.js * (edit) ambari-web/app/utils/hosts.js * (edit) ambari-web/app/templates/main/admin/highAvailability/hawq/addStandby/wizard.hbs * (edit) ambari-web/app/templates/main/admin/stack_upgrade/stack_upgrade_wizard.hbs * (edit) ambari-web/app/templates/main/service/widgets/create/step2_template.hbs * (edit) ambari-web/app/views/main/admin/highAvailability/resourceManager/step4_view.js * (edit) ambari-web/app/templates/common/time_range.hbs * (edit) ambari-web/app/templates/main/admin/highAvailability/resourceManager/step3.hbs * (edit) ambari-web/app/templates/main/service/reassign/step3.hbs * (edit) ambari-web/app/mixins/wizard/wizardProgressPageView.js * (edit) ambari-web/app/routes/rm_high_availability_routes.js * (edit) ambari-web/app/templates/main/admin/kerberos/notify_security_off_popup.hbs * (edit) ambari-web/app/templates/installer.hbs * (edit) ambari-web/app/templates/common/log_tail.hbs * (edit) ambari-web/app/templates/main/admin/kerberos/step1.hbs * (edit) ambari-web/app/templates/common/host_progress_popup_footer.hbs * (edit) ambari-web/app/templates/main/admin/stack_upgrade/upgrade_options.hbs * (edit) ambari-web/app/views/main/admin/highAvailability/hawq/removeStandby/step3_view.js * (edit) ambari-web/app/templates/common/modal_popups/hosts_table_list_popup.hbs * (edit) ambari-web/test/views/common/configs/widgets/config_widget_view_test.js * (edit) ambari-web/vendor/scripts/ember-latest.js * (edit) ambari-web/app/views/wizard/step1_view.js * (edit) ambari-web/app/controllers/main/alerts/alert_definitions_actions_controller.js * (edit) ambari-web/app/utils/load_timer.js * (edit) ambari-web/app/templates/main/alerts/configs/alert_config_threshold.hbs * (edit) ambari-web/app/views/common/configs/widgets/label_config_widget_view.js * (edit) ambari-web/app/templates/common/selectable_popup.hbs * (edit) ambari-web/app/templates/main/admin/stack_upgrade/upgrade_version_column.hbs * (edit) ambari-web/app/templates/main.hbs * (edit) ambari-web/app/templates/main/admin/highAvailability/resourceManager/wizard.hbs * (edit) ambari-web/app/templates/common/custom_date_popup.hbs * (edit) ambari-web/app/templates/main/admin/highAvailability/nameNode/rollbackHA/rollback_wizard.hbs * (edit) ambari-web/vendor/scripts/bootstrap-datepicker.js * (edit) ambari-web/app/templates/main/admin/highAvailability/hawq/removeStandby/step2.hbs * (edit) ambari-web/app/templates/main/dashboard/widgets/cluster_metrics.hbs * (edit) ambari-web/app/templates/main/service/services/hdfs.hbs * (edit) ambari-web/test/controllers/main/admin/stack_and_upgrade_controller_test.js * (edit) ambari-web/app/data/host/categories.js * (edit) ambari-web/app/templates/wizard/step3/step3_host_warnings_popup.hbs * (add) ambari-web/app/styles/bootstrap_overrides.less * (edit) ambari-web/app/templates/main/dashboard/config_history.hbs * (edit) ambari-web/app/styles/stack_versions.less * (edit) ambari-web/app/templates/main/service/services/oozie.hbs * (edit) ambari-web/app/routes/add_host_routes.js * (edit) ambari-web/app/templates/common/configs/widgets/controls.hbs * (delete) ambari-web/vendor/scripts/jquery-1.7.2.min.js * (edit) ambari-web/app/templates/common/rolling_restart_view.hbs * (edit) ambari-web/app/templates/main/host/combo_search_box.hbs * (edit) ambari-web/app/templates/common/configs/service_config_wizard.hbs * (edit) ambari-web/app/styles/log_file_search.less * (edit) ambari-web/vendor/scripts/bootstrap.js *	0	0	https://issues.apache.org/jira/browse/AMBARI-18573
19230	AMBARI-18476	This seems not enough, or am I missing something?	1	0	https://issues.apache.org/jira/browse/AMBARI-18476
19231	AMBARI-18476	Can you clarify what you mean by that?	1	1	https://issues.apache.org/jira/browse/AMBARI-18476
19232	AMBARI-18456	If this is resolved, can the feature branch be removed?	1	0	https://issues.apache.org/jira/browse/AMBARI-18456
19233	AMBARI-18431	[~jluniya] any reason for changing the fix version. The patch is tested and reviewed. Can you merge it in?	1	0	https://issues.apache.org/jira/browse/AMBARI-18431
19234	AMBARI-18425	[~gautamborad], Can you take a look at this? The Review Board is at https://reviews.apache.org/r/53722/. Also, can you recommend anyone else that might be able to review this?	1	0	https://issues.apache.org/jira/browse/AMBARI-18425
19235	AMBARI-18259	[~alicezhangxiaolu], what branch is this for? Also, please combine these small code reviews that fix typos or text descriptions into a single patch 2.5 is already in a blocker-only mode	1	1	https://issues.apache.org/jira/browse/AMBARI-18259
19236	AMBARI-18245	Can the feature branch created for this be removed?	1	0	https://issues.apache.org/jira/browse/AMBARI-18245
19237	AMBARI-18180	The failure above is not related to this patch: {code} ls: cannot access /home/hudson/tools/java/latest1.7: No such file or directory {code}	0	0	https://issues.apache.org/jira/browse/AMBARI-18180
19238	AMBARI-18051	Hi [~Tim Thorpe], This change causes build failure when running unit tests with Maven without {{clean}} if any tests need recompilation. This makes running specific unit tests after changes more expensive in terms of time, since we need to clean, too. {noformat} $ mvn -pl ambari-server -DskipPythonTests -Dtest=ServiceModuleTest clean test ... [INFO] BUILD SUCCESS [INFO] Total time: 33.124 s ... $ mvn -pl ambari-server -DskipPythonTests -Dtest=ServiceModuleTest test ... [INFO] BUILD SUCCESS [INFO] Total time: 16.026 s ... $ touch ambari-server/src/test/java/org/apache/ambari/server/stack/ServiceModuleTest.java $ mvn -pl ambari-server -DskipPythonTests -Dtest=ServiceModuleTest test ... [INFO] BUILD FAILURE [ERROR] Failed to execute goal org.apache.maven.plugins:maven-jar-plugin:3.0.2:test-jar (create-sample-upgrade-check-jar) on project ambari-server: Error assembling JAR: A zip file cannot include itself ... $ touch ambari-server/src/test/java/org/apache/ambari/server/stack/ServiceModuleTest.java $ mvn -pl ambari-server clean process-test-classes ... [INFO] BUILD SUCCESS ... {noformat} Could you please check? Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-18051
19239	AMBARI-17994	[~cnormile] Can you add more details for this feature?	0	0	https://issues.apache.org/jira/browse/AMBARI-17994
19240	AMBARI-17981	Hi [~swapanshridhar], could you please check the unit test that started giving error with this commit? https://builds.apache.org/job/Ambari-trunk-Commit/5836/testReport/	1	1	https://issues.apache.org/jira/browse/AMBARI-17981
19241	AMBARI-17898	Thanks you so much Robert! Will correct it.	0	0	https://issues.apache.org/jira/browse/AMBARI-17898
19242	AMBARI-17898	[~dsen]hello Dmytro, sorry to bother you, would you be pushing the patch back to trunk and branch 2.5 since python-kerberos issue is fixed now ?	1	1	https://issues.apache.org/jira/browse/AMBARI-17898
19243	AMBARI-17888	[~swagle] I believe this is already committed to 2.4, so resolving this JIRA. Can you confirm?	1	0	https://issues.apache.org/jira/browse/AMBARI-17888
19244	AMBARI-17717	[~lavjain] Do we still need this JIRA? Also moving this out of 2.4.0	1	0	https://issues.apache.org/jira/browse/AMBARI-17717
19245	AMBARI-17645	[~afernandez], could you review the latest patch? Thank you.	1	0	https://issues.apache.org/jira/browse/AMBARI-17645
19246	AMBARI-17645	Hi [~afernandez], could you review the patch? Thank you.	1	0	https://issues.apache.org/jira/browse/AMBARI-17645
19247	AMBARI-17645	Hi [~afernandez], could you review the patch? Thank you.	1	0	https://issues.apache.org/jira/browse/AMBARI-17645
19248	AMBARI-17457	Pushed to trunk. {code} commit 875f1efb63cbbe81259a79c0b1530eba7ef2a3ec Author: Aravindan Vijayan <avijayan@hortonworks.com> Date: Tue Sep 20 15:54:32 2016 -0700 AMBARI-17457 : Modify the AMS stack scripts to support distributed collector (dsen via avijayan) {code}	1	0	https://issues.apache.org/jira/browse/AMBARI-17457

#	Requirement	Questions	Answered	Successfully Answered	URL
19249	AMBARI-17457	FAILURE: Integrated in Jenkins build Ambari-trunk-Commit #5700 (See [https://builds.apache.org/job/Ambari-trunk-Commit/5700/]) AMBARI-17457 : Modify the AMS stack scripts to support distributed (avijayan: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=875f1efb63cbbe81259a79c0b1530eba7ef2a3ec]) * (edit) ambari-server/src/main/resources/common-services/AMBARI_METRICS/0.1.0/package/scripts/service_check.py * (edit) ambari-server/src/main/resources/common-services/AMBARI_METRICS/0.1.0/package/scripts/metrics_grafana_util.py * (edit) ambari-server/src/main/resources/common-services/AMBARI_METRICS/0.1.0/package/templates/metrics_grafana_datasource.json.j2 * (edit) ambari-server/src/main/resources/common-services/AMBARI_METRICS/0.1.0/package/scripts/params.py * (add) ambari-common/src/main/python/ambari_commons/parallel_processing.py * (edit) ambari-common/src/main/python/ambari_commons/ambari_metrics_helper.py	0	0	https://issues.apache.org/jira/browse/AMBARI-17457
19250	AMBARI-17241	Thanks for your suggestions [~tctruong213] From your suggestion I infer that you would only show a button to the user saying 'Apply Configuration Changes'. (i) if all configs that were changed does not require restart ... show no warning, apply changes (ii) if configs that were changed contains configs requiring restart and configs not requiring restart ... show warning, restart service (iii) if all configs that were changed requires restart ... show warning, restart service Please confirm if my understanding described above is correct. It would need only some minor UI changes to the prototype I presented, to make it work that way. Just putting my thoughts on scenario (ii) above, where user changes some configs, which need restart and some which don't. Would there be a situation where user wants to apply the configurations, which do not require restart, immediately, without affecting running jobs? And later restart the service to apply other configs which were not applied before, after the running jobs have been completed? This action would apply the configs 'partially' - apply changes which do not require restart first. From implementation perspective, this would be challenging to implement - tracking if each config was applied or not is more complex. Any thoughts on this [~tctruong213] whether users/customers require this partial way of applying configs? This can also be an enhancement for this JIRA, if required.	1	1	https://issues.apache.org/jira/browse/AMBARI-17241
19251	AMBARI-17241	Hi [~mithmatt] I like the idea of Ambari supporting restart-less configuration changes. This may be a naive question but how will Ambari communicate the configuration change to the service when restart is not required? I missed that from a quick look at the doc. Also is the supports-reload or restart-required parameter expected to be configured by the administrator? That could be error prone.	1	0	https://issues.apache.org/jira/browse/AMBARI-17241
19252	AMBARI-17241	{color:blue}*This may be a naive question but how will Ambari communicate the configuration change to the service when restart is not required?*{color} In the prototype I implemented, I added a new method *reload* in Script.py. The service components has its own .py file, which implements the *reload* method overriding the implementation from Script.py. When the reload button was clicked on the UI, a RELOAD command was POSTed on the REST API (similar to how START, STOP of individual host components are done) for all the host components of the service. The versioning of configs were handled by existing code. Code snippet I introduced in ajax.js: {code} 'reload.hostComponents': { 'real': '/clusters/{clusterName}/requests', 'mock': '', 'format': function (data) { return { type: 'POST', data: JSON.stringify({{ "RequestInfo": { "command": "RELOAD", "context": data.context, "operation_level": data.operation_level }, "Requests/resource_filters": data.resource_filters }) } } } {code} {color:blue}*Also is the supports-reload or restart-required parameter expected to be configured by the administrator?*{color} I do not expect admins to go change the config files. Once this feature is in, any service can use the feature by specifying that additional property <supports-reload> or <require-restart>false<require-restart> for the config, in addition to the command that has to be run on the component host for reload. I've added an example on the last page of the design proposal, in case you missed it. Think of it as something similar as supporting final properties, or editable-only-at-install properties.	1	1	https://issues.apache.org/jira/browse/AMBARI-17241
19253	AMBARI-17241	[~swagle] [~arpitagarwal] Considering that each service component might have its own way to refresh its configs (similar to how each component has start and stop), in my prototype I made RELOAD a command similar to CONFIGURE, START or STOP that is recognized by the REST API and runs the reload method on each host component. Code snippet from ambari-server/src/main/java/org/apache/ambari/server/metadata/ActionMetadata.java: {code} private void fillHostComponentCommands() { //Standart commands for any host component defaultHostComponentCommands.add("RESTART"); defaultHostComponentCommands.add("START"); ... defaultHostComponentCommands.add("RELOAD"); defaultHostComponentCommands.add("CONFIGURE"); ... } {code} Code snippet from ambari-common/src/main/python/resource_management/libraries/script/script.py: {code} def reload(self, env): """ To be overridden by subclasses """ pass {code} Code snippet from ambari-server/src/main/java/org/apache/ambari/server/agent/HeartbeatProcessor.java: {code} if (report.getCustomCommand() != null && "RELOAD".equals(report.getCustomCommand())){ ... scHost.updateActualConfigs(report.getConfigurationTags()); ... } {code} [~swagle] Having an endpoint which would support querying configs which do not require restart, is nice :) - something that I never thought about. Is there a specific use case that you have in mind, where this would be beneficial?	1	1	https://issues.apache.org/jira/browse/AMBARI-17241
19254	AMBARI-17129	[~elserj], Can you fix the merge issues? {noformat} error: patch failed: ambari-server/src/main/java/org/apache/ambari/server/upgrade/UpgradeCatalog240.java:361 error: ambari-server/src/main/java/org/apache/ambari/server/upgrade/UpgradeCatalog240.java: patch does not apply error: patch failed: ambari-server/src/test/java/org/apache/ambari/server/upgrade/UpgradeCatalog240Test.java:498 error: ambari-server/src/test/java/org/apache/ambari/server/upgrade/UpgradeCatalog240Test.java: patch does not apply {noformat} Please fix for the trunk and branch-2.4 since they may have diverged significantly by now.	1	1	https://issues.apache.org/jira/browse/AMBARI-17129
19255	AMBARI-17129	bq. Can you fix the merge issues? Sure thing. Will rebase and check both branches.	1	1	https://issues.apache.org/jira/browse/AMBARI-17129
19256	AMBARI-17126	[~bergenholtz] Can you explain, how this is already possible today? Do you refer to an active-active setup with both Ambari's configured exactly the same way and accessible via a load-balancer?	0	0	https://issues.apache.org/jira/browse/AMBARI-17126
19257	AMBARI-17126	I have two questions: 1.) Is this actively being looked at for 2.5? (Is there an estimated release date for 2.5?) 2.) Also, why exactly was AMBARI-7896 closed and being tracked here? Active-Passive fault tolerance/redundancy isn't necessarily the same concept as high availability.	1	0	https://issues.apache.org/jira/browse/AMBARI-17126

#	Requirement	Questions	Answered	Successfully Answered	URL
19258	AMBARI-17126	[~allmusic76] {quote} 1.) Is this actively being looked at for 2.5? (Is there an estimated release date for 2.5?) {quote} - It is not for Ambari-2.5. {quote} 2.) Also, why exactly was AMBARI-7896 closed and being tracked here? Active-Passive fault tolerance/redundancy isn't necessarily the same concept as high availability. {quote} - Here also, we are working on having one Ambari server active at one point of time and other passive/non-active, and we are not doing the load balancing part as part of this JIRA.	0	0	https://issues.apache.org/jira/browse/AMBARI-17126
19259	AMBARI-17074	[~srimanth.gunturi] The patch has been updated based on the review comments. Could you help to review it again? Thanks.	1	0	https://issues.apache.org/jira/browse/AMBARI-17074
19260	AMBARI-17074	[~mahadev] [~srimanth.gunturi] Patch has been updated. Could you help to review it again? Thanks.	1	0	https://issues.apache.org/jira/browse/AMBARI-17074
19261	AMBARI-17041	[~tctruong213] According to this JIRA, it is assumed that all property-types are mutually exclusive. Would there be a scenario where a property is of multiple types?	1	0	https://issues.apache.org/jira/browse/AMBARI-17041
19262	AMBARI-17041	[~patelket@us.ibm.com] saw this just now : "if (Em.get(property,'propertyType') != null) {" why we are checking Em.get(property,'propertyType') against null ? not simply if(Em.get(property,'propertyType'))	1	1	https://issues.apache.org/jira/browse/AMBARI-17041
19263	AMBARI-17006	Hi [~Sandeep Nemuri] Can I write a patch to resolve this issue?	0	0	https://issues.apache.org/jira/browse/AMBARI-17006
19264	AMBARI-16864	Hi [~jluniya], two questions about unit test: 1. How do I verify the unit test locally? 2. How to get the config file like "configs/default.json" which has Spark2 definitions? Thanks a lot.	1	1	https://issues.apache.org/jira/browse/AMBARI-16864
19265	AMBARI-16864	[~jerryshao] Sorry dropped the ball on this. Can you please resubmit a new patch with the latest?	1	0	https://issues.apache.org/jira/browse/AMBARI-16864
19266	AMBARI-16864	[~jerryshao] I still see the reviewboard has older patch. Can you cross check?	1	0	https://issues.apache.org/jira/browse/AMBARI-16864
19267	AMBARI-16755	Hi, [~zjffdu] I want to submit the patch. Could you assign this jira to me? Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-16755
19268	AMBARI-16630	[~dgergely] can we add taskId as there are parallel execution so logs may intersperse. Rest of the changes look good. e.g. {code} logger.info("Command execution metadata - taskId = {taskId}, retry enabled = {retryAble}, max retry duration (sec) = {retryDuration}". format(taskId=taskId, retryAble=retryAble, retryDuration=retryDuration)) {code}	1	1	https://issues.apache.org/jira/browse/AMBARI-16630
19269	AMBARI-16362	Hi Alejandro, Sorry for that. It was not captured on the UTs, fresh installation, and service_check tests I am missing this import : from resource_management.libraries.functions.check_process_status import check_process_status in /ambari/ambari-server/src/main/resources/common-services/HDFS/2.1.0.2.0/package/scripts/hdfs_snamenode.py Should I create a new patch just with this fix or recreate the original patch adding this change?	1	1	https://issues.apache.org/jira/browse/AMBARI-16362
19270	AMBARI-16272	[~masatana], what is the branch and mvn/JDK versions you are using? I ran tests on branch-2.4 and trunk, and they were all green.	1	1	https://issues.apache.org/jira/browse/AMBARI-16272
19271	AMBARI-16171	{code} /** * Update Phoenix Query Server Kerberos configurations. 2.4 will move the Phoenix Query Server to * supporting SPNEGO authentication which requires that the "HTTP/_HOST" principal and corresponding * keytab file instead of the generic HBase service keytab it previously had. */ protected void updatePhoenixConfigs() throws AmbariException { AmbariManagementController controller = injector.getInstance(AmbariManagementController.class); Clusters clusters = controller.getClusters(); if (null != clusters) { Map<String, Cluster> clusterMap = clusters.getClusters(); if (null != clusterMap && !clusterMap.isEmpty()) { for (final Cluster cluster : clusterMap.values()) { Set<String> installedServices = cluster.getServices().keySet(); // HBase and Kerberos are both installed/enabled if (installedServices.contains("HBASE") && installedServices.contains("KERBEROS")) { Config hbaseSite = cluster.getDesiredConfigByType(HBASE_SITE_CONFIG); if (null != hbaseSite) { Map<String, String> hbaseSiteProperties = hbaseSite.getProperties(); // Get Phoenix Query Server kerberos config properties String pqsKrbPrincipal = hbaseSiteProperties.get(PHOENIX_QUERY_SERVER_PRINCIPAL_KEY); String pqsKrbKeytab = hbaseSiteProperties.get(PHOENIX_QUERY_SERVER_KEYTAB_KEY); // Principal and Keytab are set if (null != pqsKrbPrincipal && null != pqsKrbKeytab) { Map<String, String> updatedKerberosProperties = new HashMap<>(); // TODO Values? updatedKerberosProperties.put(PHOENIX_QUERY_SERVER_PRINCIPAL_KEY, "HTTP/_HOST@REALM"); updatedKerberosProperties.put(PHOENIX_QUERY_SERVER_KEYTAB_KEY, "/etc/security/keytabs/spnego.service.keytab"); updateConfigurationProperties(HBASE_SITE_CONFIG, updatedKerberosProperties, true, false); } } } } } } } {code} [~rlevas] the above is what I presently have worked out. The only confusion I have is how to determine what values should be placed in {{updatedKerberosProperties}} (by the TODO). Can you point me at some examples of how to get the location of the spnego keytab and how to get the proper principal name (I could probably just modify the existing principal, but I'm not sure if that would work as I'd hope).	1	1	https://issues.apache.org/jira/browse/AMBARI-16171
19272	AMBARI-16171	.003 is just rebased on top of origin/trunk. Nothing else changed. I couldn't get a real test to work -- the express upgrade dialog wouldn't let me proceed, claiming that the last status check was run days prior. I figured out where that was being stored in postgres, tried to update them, but even that wasn't working (so obviously I was missing something). So, apologies in advance if something isn't 100% correct in the upgrade path, but I think we should get this in now and deal with any upgrade problems later :). Can you commit for me [~rlevas]?	1	1	https://issues.apache.org/jira/browse/AMBARI-16171
19273	AMBARI-16171	.003 is just rebased on top of origin/trunk. Nothing else changed. I couldn't get a real test to work – the express upgrade dialog wouldn't let me proceed, claiming that the last status check was run days prior. I figured out where that was being stored in postgres, tried to update them, but even that wasn't working (so obviously I was missing something). So, apologies in advance if something isn't 100% correct in the upgrade path, but I think we should get this in now and deal with any upgrade problems later . Can you commit for me [~rlevas]?	1	1	https://issues.apache.org/jira/browse/AMBARI-16171

#	Requirement	Questions	Answered	Successfully Answered	URL
19274	AMBARI-16171	Reviewboard is linked (if you didn't see the email about it already). Also, trying again with the upgrade testing locally: {noformat} Exception in thread "main" org.apache.ambari.server.AmbariException at org.apache.ambari.server.upgrade.SchemaUpgradeHelper.executeDMLUpdates(SchemaUpgradeHelper.java:234) at org.apache.ambari.server.upgrade.SchemaUpgradeHelper.main(SchemaUpgradeHelper.java:308) Caused by: java.lang.NullPointerException at org.apache.ambari.server.controller.internal.ArtifactResourceProvider$ClusterTypeRegistration.instanceExists(ArtifactResourceProvider.java:711) at org.apache.ambari.server.controller.internal.ArtifactResourceProvider.validateParent(ArtifactResourceProvider.java:403) at org.apache.ambari.server.controller.internal.ArtifactResourceProvider.access$200(ArtifactResourceProvider.java:62) at org.apache.ambari.server.controller.internal.ArtifactResourceProvider$2.invoke(ArtifactResourceProvider.java:301) at org.apache.ambari.server.controller.internal.ArtifactResourceProvider$2.invoke(ArtifactResourceProvider.java:297) at org.apache.ambari.server.controller.internal.AbstractResourceProvider.getResources(AbstractResourceProvider.java:307) at org.apache.ambari.server.controller.internal.ArtifactResourceProvider.getResources(ArtifactResourceProvider.java:208) at org.apache.ambari.server.controller.KerberosHelperImpl.getKerberosDescriptor(KerberosHelperImpl.java:1068) at org.apache.ambari.server.upgrade.UpgradeCatalog240.updatePhoenixConfigs(UpgradeCatalog240.java:2081) at org.apache.ambari.server.upgrade.UpgradeCatalog240.executeDMLUpdates(UpgradeCatalog240.java:337) at org.apache.ambari.server.upgrade.AbstractUpgradeCatalog.upgradeData(AbstractUpgradeCatalog.java:666) at org.apache.ambari.server.upgrade.SchemaUpgradeHelper.executeDMLUpdates(SchemaUpgradeHelper.java:231) ... 1 more {noformat} I got past one exception already -- had to add the following to the {{updatePhoenixConfigs()}} method: {code} final ViewRegistry viewRegistry = injector.getInstance(ViewRegistry.class); ViewRegistry.initInstance(viewRegistry); {code} But, to be quite honest, I'm really not sure where to start with this exception. Some direction would be great :)	1	1	https://issues.apache.org/jira/browse/AMBARI-16171
19275	AMBARI-16171	[~jonathan.hurley], [~ncole@hortonworks.com], if you get a chance, can you take a look at the review for this - https://reviews.apache.org/r/47428/. We could use a few more eyes on it.	1	1	https://issues.apache.org/jira/browse/AMBARI-16171
19276	AMBARI-16171	[~elserj], can you add them to the review... that might be more effective. :)	1	1	https://issues.apache.org/jira/browse/AMBARI-16171
19277	AMBARI-15679	[~oleewere] There were a bunch of new files added for which there are different licenses in play. However, the top level license file has not been updated. Likewise does any of these new additions need an update to the NOTICE file? Additionally, ambari-logsearch/ambari-logsearch-portal/src/main/webapp/libs/bower/globalize/generator/HijriCalendar.js ambari-logsearch/ambari-logsearch-portal/src/main/webapp/libs/bower/globalize/generator/UmAlQuraCalendar.js Could you provide the original source for these? What license was the original source under? \cc [~sumitmohanty]	0	0	https://issues.apache.org/jira/browse/AMBARI-15679
19278	AMBARI-15636	[~ncole] Can we close this EPIC?	0	0	https://issues.apache.org/jira/browse/AMBARI-15636
19279	AMBARI-15577	[~jmarron] By when do you think you would be able to do HIVE featurization work? Can you provide an ETA?	1	1	https://issues.apache.org/jira/browse/AMBARI-15577
19280	AMBARI-15443	[~dili], can you post this on ReviewBoard?	1	0	https://issues.apache.org/jira/browse/AMBARI-15443
19281	AMBARI-15443	FAILURE: Integrated in Ambari-trunk-Commit #4575 (See [https://builds.apache.org/job/Ambari-trunk-Commit/4575/]) AMBARI-15443:Make Host bulk command menu item list stack driven instead (dili: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=2410929485dc8d8b3baac28528d8762322d748a4]) * ambari-server/src/main/resources/common-services/STORM/0.9.1.2.1/metainfo.xml * ambari-web/app/models/stack_service_component.js * ambari-server/src/test/resources/stacks/HDP/2.0.5/services/HDFS/metainfo.xml * ambari-server/src/test/resources/stacks/HDP/2.0.6/services/HBASE/metainfo.xml * ambari-server/src/test/resources/stacks/HDP/2.0.5/services/HBASE/metainfo.xml * ambari-server/src/test/java/org/apache/ambari/server/controller/AmbariManagementControllerTest.java * ambari-server/src/main/java/org/apache/ambari/server/controller/internal/StackServiceComponentResourceProvider.java * ambari-server/src/main/resources/common-services/PXF/3.0.0/metainfo.xml * ambari-server/src/main/resources/common-services/HAWQ/2.0.0/metainfo.xml * ambari-web/app/mappers/stack_service_mapper.js * ambari-server/src/main/resources/common-services/HBASE/0.96.0.2.0/metainfo.xml * ambari-web/app/views/main/host/hosts_table_menu_view.js * ambari-server/src/main/java/org/apache/ambari/server/state/ComponentInfo.java * ambari-server/src/main/resources/properties.json * ambari-server/src/main/resources/common-services/YARN/2.1.0.2.0/metainfo.xml * ambari-server/src/main/java/org/apache/ambari/server/state/BulkCommandDefinition.java * ambari-server/src/test/java/org/apache/ambari/server/stack/ComponentModuleTest.java * ambari-server/src/main/java/org/apache/ambari/server/stack/ComponentModule.java * ambari-server/src/main/resources/common-services/HDFS/2.1.0.2.0/metainfo.xml * ambari-web/test/mappers/stack_service_mapper_test.js * ambari-server/src/main/java/org/apache/ambari/server/controller/StackServiceComponentResponse.java	0	0	https://issues.apache.org/jira/browse/AMBARI-15443
19282	AMBARI-15436	Hi [~jwang], the patch does not apply on branch-2.2. Can you upload a branch-2.2 patch?	1	1	https://issues.apache.org/jira/browse/AMBARI-15436
19283	AMBARI-15412	Hello Robert, Could you please apply the trunk patch(AMBARI-15412-trunk.patch) on branch-2.2? It is working on branch-2.2. I was mistakenly looking at different branch when I developed patch for branch-2.2. Thank you,	1	0	https://issues.apache.org/jira/browse/AMBARI-15412
19284	AMBARI-15412	[~arborkar] I am sorry but I guess I missed you last message. Do you still need to patch applied to branch-2.2? If so, unfortunately we missed the boat for the Ambari 2.2.2 release. If I apply it to branch-2.2 now, it may get into future Ambari 2.2.x releases. Please make sure that the patch will apply to the branch.	0	0	https://issues.apache.org/jira/browse/AMBARI-15412

#	Requirement	Questions	Answered	Successfully Answered	URL
19285	AMBARI-15388	How is <order> being maintained? The <order> element is key to indicate dependencies for upgrading, and is not random. Oozie must go before Falcon (at the service level). Datanodes must be upgraded before RegionServers before NodeManagers within the same host. The CLIENTS group spans multiple services. What about pre-requisite checks? Order there also may be important. ClusterGrouping definitely has order - you can't execute Save Database State before HDFS Finalization occurs. <priority> within a service check group is also important. It's difficult to see all the changes based on just a plain text diff, can you add Review Board review or create a feature branch? This is a really large commit. We also have to balance stability here, especially when trying to maintain upgrade packs. It's already hard enough maintaining the source->target upgrade packs, but if we now need one per service as well, that increases the number of files as a multiple of the number of services.	0	0	https://issues.apache.org/jira/browse/AMBARI-15388
19286	AMBARI-15388	I agree about the priority/ordering. If you're going to break up the files, then you still need some sort of master XML to define ordering. The idea of adding new services on your own to be a part of upgrade is an interesting proposal. But it comes with a lot of questions: - Are other stack services affected by this service's upgrade? - When does the new service upgrade? - Does the service upgrade all together, or broken apart (things like clients/slaves/masters) I suppose you could always argue that the "master" XML file defines ordering for the stack as it ships. And then you can tack on the to end of it services which define their own upgrade. But now you hit issues with services that are already part of the main upgrade overriding their parent and so on. I think that: - At the very least, this needs a design sent out the community - especially before any code is written - A feature branch is a great starting point too; allows iterative development and smaller, more understandable reviews. The sheer size of this review makes it very difficult to find problems - ReviewBoard use is a must	0	0	https://issues.apache.org/jira/browse/AMBARI-15388
19287	AMBARI-15348	This commit breaks {{ambari-server setup}} {code} [root@c6401 yum.repos.d]# ambari-server setup Using python /usr/bin/python Setup ambari-server Checking SELinux... SELinux status is 'disabled' Customize user account for ambari-server daemon [y/n] (n)? Adjusting ambari-server permissions and ownership... Checking firewall status... Checking JDK... [1] Oracle JDK 1.8 + Java Cryptography Extension (JCE) Policy Files 8 [2] Oracle JDK 1.7 + Java Cryptography Extension (JCE) Policy Files 7 [3] Custom JDK == Enter choice (1): To download the Oracle JDK and the Java Cryptography Extension (JCE) Policy Files you must accept the license terms found at http://www.oracle.com/technetwork/java/javase/terms/license/index.html and not accepting will cancel the Ambari Server setup and you must install the JDK and JCE files manually. Do you accept the Oracle Binary Code License Agreement [y/n] (y)? Downloading JDK from http://public-repo-1.hortonworks.com/ARTIFACTS/jdk-8u60-linux-x64.tar.gz to //var/lib/ambari-server/resources/jdk-8u60-linux-x64.tar.gz jdk-8u60-linux-x64.tar.gz... 100% (172.8 MB of 172.8 MB) Successfully downloaded JDK distribution to //var/lib/ambari-server/resources/jdk-8u60-linux-x64.tar.gz Installing JDK to //usr/jdk64/ Successfully installed JDK to //usr/jdk64/ Downloading JCE Policy archive from http://public-repo-1.hortonworks.com/ARTIFACTS/jce_policy-8.zip to //var/lib/ambari-server/resources/jce_policy-8.zip Successfully downloaded JCE Policy archive to //var/lib/ambari-server/resources/jce_policy-8.zip Installing JCE policy... Completing setup... Configuring database... Enter advanced database configuration [y/n] (n)? Configuring database... Default properties detected. Using built-in database. WARNING: Command chown None //etc/ambari-server/conf/password.dat returned exit code //etc/ambari-server/conf/password.dat with message: chown: invalid user: `None' Configuring ambari database... Checking PostgreSQL... Running initdb: This may take up to a minute. Initializing database: [OK] About to start PostgreSQL Configuring local database... Connecting to local database...done. Configuring PostgreSQL... Restarting PostgreSQL Extracting system views... ERROR: Exiting with exit code 1. REASON: Multiple versions of javax.servlet.jsp.jstl.jar found in java class path (/usr/lib/ambari-server/javax.servlet.jsp.jstl-1.2.2.jar and /usr/lib/ambari-server/javax.servlet.jsp.jstl-1.2.0.v201105211821.jar). Make sure that you include only one javax.servlet.jsp.jstl.jar in the java class path '//etc/ambari-server/conf:/usr/lib/ambari-server/*:/usr/share/java/postgresql-jdbc.jar'. {code} Reverted from trunk and branch-2.2 {code} Author: Jonathan Hurley <jhurley@hortonworks.com> Date: Mon Mar 14 15:51:58 2016 -0400 Revert "AMBARI-15348 Ambari upgrade/start commands should detect any ambari-server hotfix and fail upgrade/start from proceeding (dsen)" This reverts commit 2b596be9b93682ed436837f4f2b75197ae4349c7. {code}	1	0	https://issues.apache.org/jira/browse/AMBARI-15348
19288	AMBARI-15256	Tested on trunk as well and it works fine. [~aantonenko] -- can you please provide details on how to reproduce the issue you saw (branch-2.2 or trunk, cluster topology, stack version, machine in topology that is missing components). I could not reproduce the issue neither on branch-2.2 nor or trunk in my environments. All hosts have their components listed properly -- see the attached screenshots.	1	1	https://issues.apache.org/jira/browse/AMBARI-15256
19289	AMBARI-15256	[~aantonenko] I looked into this issue myself out of curiosity. I built Ambari branch-2.2 from this [commit|https://github.com/apache/ambari/commit/7abb79cc6b390db189345339a196a766a977a817] and trunk from this [commit|https://github.com/apache/ambari/commit/f6ecddf69f615d800a0146b730a44ba075013760] (both before you reverted the commit of AMBARI-15256) With these Ambari builds, I deployed HDP 2.4 stack with a bunch of services: HDFS, YARN, MAPREDUCE2, HBASE, HIVE, PIG, TEZ, AMBARI METRICS and ZOOKEEPER. The host details page was showing all the host components that I had installed (through manual wizard installation). I was checking the console log at the same time, and there was no instance where commandMap was undefined. Do you remember which service's host component caused the issue? Possibly, a service component which I hadn't deployed caused the issue. The *if* condition introduced by [~adenissov] in the new set of patches should solve the issue you might have seen. I've seen that the new patches work on HDP cluster without any host components disappearing. I'm committing the new patches to trunk and branch-2.2. If the commit breaks something again, it should be because of some nasty edge case (which doesn't happen with the main/core services) - you may log a bug and I'd be happy to drill down the root cause and fix it.	1	1	https://issues.apache.org/jira/browse/AMBARI-15256
19290	AMBARI-15226	[~Tim Thorpe] Can you update the ReviewBoard with the latest patch? It would be easier to review and comment. Also can you add me to the reviewers list so that it gets on my reviews dashboard?	1	1	https://issues.apache.org/jira/browse/AMBARI-15226
19291	AMBARI-15226	[~Tim Thorpe] Doesnt look like the review https://reviews.apache.org/r/44210/ got updated. Did you hit the Publish button?	1	1	https://issues.apache.org/jira/browse/AMBARI-15226
19292	AMBARI-15172	[~adenisso], could you please add a reviewboard for especially the trunk patch. It is bigger because I believe it includes some unit tests as trunk has some unit tests for HAWQ.	1	1	https://issues.apache.org/jira/browse/AMBARI-15172
19293	AMBARI-15133	Looks like this was committed to trunk and branch-2.2. Is there a ReviewBoard for this? If so, please link it from this JIRA. Also, I reverted the changes from branch-2.2, because it was causing a rat failure (mvn clean apache-rat:check failed).	1	1	https://issues.apache.org/jira/browse/AMBARI-15133
19294	AMBARI-15083	Is there a default recommendation for hawq_global_rm_type when YARN is available, besides the validation that is covered here?	0	0	https://issues.apache.org/jira/browse/AMBARI-15083

#	Requirement	Questions	Answered	Successfully Answered	URL
19295	AMBARI-15039	Some Design Issues: 1. Enable PAM mode 2. Bootstrap group for admin authority 3. Importing of groups into Ambari? Is it needed?	0	0	https://issues.apache.org/jira/browse/AMBARI-15039
19296	AMBARI-14958	[~onechiporenko] Is this JIRA still valid? Moving it out of 2.4.0	0	0	https://issues.apache.org/jira/browse/AMBARI-14958
19297	AMBARI-14914	what about hdp 2.4?	1	1	https://issues.apache.org/jira/browse/AMBARI-14914
19298	AMBARI-14914	Also would update to 1.6.1 in HDP 2.4	1	1	https://issues.apache.org/jira/browse/AMBARI-14914
19299	AMBARI-14902	hi [~dsen], with your change, ambari-metrics-collector and ambari-metrics-hadoop-sink are not longer produced. Is it a expected behavior? Could you let us know why they are no longer required after adding grafana?	1	0	https://issues.apache.org/jira/browse/AMBARI-14902
19300	AMBARI-14902	hi [~dsen], the [assembly pom expects|https://github.com/apache/ambari/blob/trunk/ambari-metrics/ambari-metrics-assembly/pom.xml#L287] a folder ambari-metrics/ambari-metrics-assembly/target//ambari-metrics-collector-$\{project.version\} to be existed in rpm maven plugin. However our build fails as it does not exists. Folder structure on the failure {code} ambari / ambari-metrics / ambari-metrics-assembly / target / antrun embedded resources/rpm/sink rpm/ambari-metrics-collector (all files in zip) {code} Error log {code} [ERROR] Failed to execute goal org.codehaus.mojo:rpm-maven-plugin:2.0.1:rpm (ambari-metrics-collector) on project ambari-metrics-assembly: Source location /home/pivotal/ambari/ambari-metrics/ambari-metrics-assembly/target/ambari-metrics-collector-2.2.2.1-SNAPSHOT/ambari-metrics-collector-2.2.2.1-SNAPSHOT/hbase/lib/hadoop-native does not exist -> [Help 1] org.apache.maven.lifecycle.LifecycleExecutionException: Failed to execute goal org.codehaus.mojo:rpm-maven-plugin:2.0.1:rpm (ambari-metrics-collector) on project ambari-metrics-assembly: Source location /home/pivotal/ambari/ambari-metrics/ambari-metrics-assembly/target/ambari-metrics-collector-2.2.2.1-SNAPSHOT/ambari-metrics-collector-2.2.2.1-SNAPSHOT/hbase/lib/hadoop-native does not exist at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:217) at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:153) at org.apache.maven.lifecycle.internal.MojoExecutor.execute(MojoExecutor.java:145) at org.apache.maven.lifecycle.internal.LifecycleModuleBuilder.buildProject(LifecycleModuleBuilder.java:84) at org.apache.maven.lifecycle.internal.LifecycleModuleBuilder.buildProject(LifecycleModuleBuilder.java:59) at org.apache.maven.lifecycle.internal.LifecycleStarter.singleThreadedBuild(LifecycleStarter.java:183) {code} Do you know if we are missing something?	1	0	https://issues.apache.org/jira/browse/AMBARI-14902
19301	AMBARI-14892	Hi [~avijayan], After your patch added exclude_packages to install_packages() here https://github.com/apache/ambari/commit/1eb846c153b156f87a7e55b65decb1f2118f90dc#diff-fe82ec590b6de80b4ced0eed8473f33eL34 our CI starts failing. {code} Traceback (most recent call last): File "/var/lib/ambari-agent/cache/common-services/AMBARI_METRICS/0.1.0/package/scripts/metrics_collector.py", line 133, in <module> AmsCollector().execute() File "/usr/lib/python2.6/site-packages/resource_management/libraries/script/script.py", line 238, in execute method(env) File "/var/lib/ambari-agent/cache/common-services/AMBARI_METRICS/0.1.0/package/scripts/metrics_collector.py", line 34, in install self.install_packages(env, exclude_packages = ['ambari-metrics-grafana']) TypeError: install_packages() got an unexpected keyword argument 'exclude_packages' {code} The definition of install_packages() does not seem to have "exclude_packages" either. https://github.com/apache/ambari/blob/trunk/ambari-common/src/main/python/resource_management/libraries/script/script.py#L424 Do you know if we are missing something here?	1	1	https://issues.apache.org/jira/browse/AMBARI-14892
19302	AMBARI-14854	Management Packs also need to consider other "pluggable" assets of Ambari beyond full-stack definitions. Custom Services and Views for example.	1	1	https://issues.apache.org/jira/browse/AMBARI-14854
19303	AMBARI-14822	ABORTED: Integrated in Ambari-trunk-Commit #4246 (See [https://builds.apache.org/job/Ambari-trunk-Commit/4246/]) AMBARI-14822 Combo Search: Create auto suggest for Host attribute (rzang: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=5e09053deb0700b09ced4f8f12902ff2ea60154a]) * ambari-web/app/utils/ajax/ajax.js * ambari-web/app/templates/main/admin/stack_upgrade/stack_upgrade_wizard.hbs * ambari-web/app/messages.js * ambari-web/app/controllers/main/admin/stack_and_upgrade_controller.js Revert "AMBARI-14822 Combo Search: Create auto suggest for Host (rzang: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=cad69e6132b554386e82d45950fc1b302f689af7]) * ambari-web/app/utils/ajax/ajax.js * ambari-web/app/templates/main/admin/stack_upgrade/stack_upgrade_wizard.hbs * ambari-web/app/messages.js * ambari-web/app/controllers/main/admin/stack_and_upgrade_controller.js AMBARI-14822 Combo Search: Create auto suggest for Host attribute (rzang: [http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=b037ef5cfc9cfd72126818cf8c5c5e7c3525dc52]) * ambari-web/app/utils/ajax/ajax.js * ambari-web/vendor/scripts/visualsearch.js * ambari-web/app/controllers/main/host/combo_search_box.js * ambari-web/app/views/main/host/combo_search_box.js	0	0	https://issues.apache.org/jira/browse/AMBARI-14822
19304	AMBARI-14820	[~jaimin] [~u39kun] Can you review this patch?	1	0	https://issues.apache.org/jira/browse/AMBARI-14820
19305	AMBARI-14820	The new patches provided by [~bhuvnesh2703] is due to missing modifications in his original patch, which has been committed already. What it means; 1. Some HAWQ functionality is broken currently (only HAWQ HA activation part, but nothing else) 2. The new patch only has what were missing from the original. I wait a few more hours to see if anybody has something to say, but will commit as it is a broker to our use case and won't affect any other part of Ambari.	1	1	https://issues.apache.org/jira/browse/AMBARI-14820
19306	AMBARI-14776	WDYT, [~billie.rinaldi]? Should this also go to some of the older maintenance release branches?	1	0	https://issues.apache.org/jira/browse/AMBARI-14776
19307	AMBARI-14776	Can anyone look at this, please? Let me know if I need to rebase it by now.	0	0	https://issues.apache.org/jira/browse/AMBARI-14776
19308	AMBARI-14746	cc: [~bikassaha] I believe this is a stale JIRA. Can you confirm?	1	0	https://issues.apache.org/jira/browse/AMBARI-14746
19309	AMBARI-14731	Joe, did you mean the failure is NOT related to the patch?	1	0	https://issues.apache.org/jira/browse/AMBARI-14731

#	Requirement	Questions	Answered	Successfully Answered	URL
19310	AMBARI-14690	Hi [~oleewere] - Vinod, Sid and I had a discussion today. We came up with a solution which needs to be provided in 2 phases - one short term and the second long term. \\ Now let me provide a list of 3 different solutions (we have already discussed 1 and 2 above) - *Solution 1*: system resource properties in ambari-agent.ini (current patch you provided) *Solution 2*: system resource properties in an external file, whose path is hard coded in ambari-agent.ini *Solution 3*: system resource properties spread across multiple files located in an external directory (say D1), whose path is hard coded in ambari-agent.ini. A property say _*system_resource_dir*_ could be added to the _*agent*_ section of ambari-agent.ini with value set to D1. \\ A little bit of background: the resource-limits file that Yarn will generate, will be created when the container comes up. Since ambari-agent.ini file will not exist at this point, the current patch with solution (1) will not work. Now solution (2) will work, but we do not want to restrict it to just 1 file. Hence we are proposing solution (3). \\ In solution (3) all json files under D1 will be read by ambari agent and the properties will override the default system values. If a property (say K1) is specified across 2 or more json files under D1, then ambari agent should not bother about which value should prevail. Ambari agent will read the files in any random order. The value of K1 will be set to whichever file is read last. The onus is on the creator of files under D1 to ensure that property K1 is specified in one and only one json file. In this case the onus will be on yarn. \\ Phases of solution: *Phase A - short term*: solution (3) above where the external directory (D1) is hard coded to this path - {{/yarn-private/ambari}} *Phase B - long term*: solution (3) above where the external directory (D1) can be set to a configurable value from the Ambari Server side \\ Do you think you can provide Solution (3) as per Phase (A) in a reasonable time? We can have a chat to clear any doubts on confusion on Friday Jan 22.	1	0	https://issues.apache.org/jira/browse/AMBARI-14690
19311	AMBARI-14690	[~oleewere], [~sumitmohanty] if we have +1 from reviewers can we merge this patch to the 2.2 line?	0	0	https://issues.apache.org/jira/browse/AMBARI-14690
19312	AMBARI-14529	Committed to trunk. 215c21cf08185908f82139477d2e12a66b2d8203 [~mithmatt], could you please attach a patch for branch-2.2.	1	1	https://issues.apache.org/jira/browse/AMBARI-14529
19313	AMBARI-14512	I check the service check script for other component and would like to add similar script for spark. But don't know details about the python api. Could anyone guide me on that ? Or I can collaborate with someone else who's familiar with the python api, I can help on what needs to be checked for spark	1	0	https://issues.apache.org/jira/browse/AMBARI-14512
19314	AMBARI-14512	[~u39kun] Why is it reverted ?	1	0	https://issues.apache.org/jira/browse/AMBARI-14512
19315	AMBARI-14459	[~sumitmohanty] I'll make the changes in the recommendation to reflect dfs.replication=[No. of DATANODEs] if (No. of hosts < 3). Is the validation required for dfs.replication property? In this case, check if dfs.replication < 3 if (No. of hosts < 3)?	1	1	https://issues.apache.org/jira/browse/AMBARI-14459
19316	AMBARI-14459	[~zjffdu] https://reviews.apache.org/r/41795/ has the discussion. In summary, the problem is that if its not a test cluster then at least three DNs are expected. Also, if someone started with one DN and then added more but forgot to modify the replication value then HDFS will not create replicas. So the risk out-weighs the advantages. Also, did you run into any problem with replication being 3 and number of DNs being 1. There are a few scenarios where I can see it to be an issue but largely its a non-issue.	1	0	https://issues.apache.org/jira/browse/AMBARI-14459
19317	AMBARI-14459	I have not encountered this issue ever - as in HDFS restart failing because number of DNs are less than number of replica. Could it be because of any other config settings?	0	0	https://issues.apache.org/jira/browse/AMBARI-14459
19318	AMBARI-14427	Hi [~atkach], could you please fix this to branch 2.2 as well?	1	1	https://issues.apache.org/jira/browse/AMBARI-14427
19319	AMBARI-14384	Hi Dmytro, Are you actively working on this jira? Do you have an ETA on it? Thanks!	1	1	https://issues.apache.org/jira/browse/AMBARI-14384
19320	AMBARI-14375	Hi [~mpapirkovskyy] We would like to fix the user name conflict issues. Can you please let us know if any work is going on for this jira?	0	0	https://issues.apache.org/jira/browse/AMBARI-14375
19321	AMBARI-14301	Hi [~jaimin] Could you please take a look at the new implementation https://reviews.apache.org/r/41835/ (uses recommendation from stack_advisor)? I reverted the javascript changes made in https://reviews.apache.org/r/41151/.	1	1	https://issues.apache.org/jira/browse/AMBARI-14301
19322	AMBARI-14301	Committed to trunk fe8a8d1e1626a5a2f294622174f34cb0f65543fe [~lavjain], are you going to make a patch for branch-2.2 as well?	1	1	https://issues.apache.org/jira/browse/AMBARI-14301
19323	AMBARI-14132	This has been bugging me for a while too. Why not just set it to 9092 like the Apache Kafka docs? I can't see anything else using that port - what was the reasoning behind changing it from the default set in Apache?	0	0	https://issues.apache.org/jira/browse/AMBARI-14132
19324	AMBARI-14029	Hi [~mithmatt], for some reasons I can not apply AMBARI-14029-branch-2-2.patch to branch 2.2 I have 2338b7564b97d83c1d15fa31a8c83076bd29c8a3 on my work copy that seems to me the right branch 2.2 HEAD at this point. Am I mistaken something? {code} jaoki@junubuntu:~/coding/ref/ambari-oss$ git apply AMBARI-14029-branch-2-2.patch error: patch failed: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/params.py:34 error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/params.py: patch does not apply jaoki@junubuntu:~/coding/ref/ambari-oss$ git log -1 commit 2338b7564b97d83c1d15fa31a8c83076bd29c8a3 Author: Andrew Onishuk <aonishuk@hortonworks.com> Date: Fri Dec 25 15:57:52 2015 +0200 AMBARI-14454. File Permission is not applied properly while using fast-hdfs-resource.jar (aonishuk) {code}	1	1	https://issues.apache.org/jira/browse/AMBARI-14029
19325	AMBARI-14023	[~ndimiduk] Has this patch been committed? Can we have the patch file attached on this JIRA?	1	0	https://issues.apache.org/jira/browse/AMBARI-14023
19326	AMBARI-13903	[~natefalkibm] Hello Nathan, For example, if you look under Advanced core-site section, the Ambari web UI allows property "hadoop.security.key.provider.path" to be left blank (empty value). This is controlled on the core-site.xml level as shown below.... <property> <name>hadoop.security.key.provider.path</name> <value></value> <value-attributes> <empty-value-valid>true</empty-value-valid> </value-attributes> Have you tried specifying properties to allow empty values to meet your requirement? It's of course up to component's Python script to handle what it means for a prop to have an empty value... Di	1	0	https://issues.apache.org/jira/browse/AMBARI-13903
19327	AMBARI-13894	Exception on downgrade: {code} java.lang.IllegalArgumentException: Cannot merge an entity that has been removed: org.apache.ambari.server.orm.entities.ClusterConfigEntity@85f5a9eb {code}	1	1	https://issues.apache.org/jira/browse/AMBARI-13894
19328	AMBARI-13721	[~smnaha], it will be a good addition to the test suite. Can you provide some general idea of how these can be run - assuming these tests are different than the unit tests.	0	0	https://issues.apache.org/jira/browse/AMBARI-13721
19329	AMBARI-13689	So [~PankajSingh], what's the proposal here? Are you suggesting to validate, check and prevent the instance creation if the same named one already exists? I would personally prefer to check and prevent duplicate instances with same display name (or whatever the identifier we use). Let me know if you have any other proposal here. Thanks, Rajendra	1	0	https://issues.apache.org/jira/browse/AMBARI-13689
19330	AMBARI-13665	Hi Ambari community, Could you please take a look into the patch? Hadoop QA failures are due to its attempt to apply it to trunk but the patch is meant to branch-2.1. as [~mithmatt] [mentioned|https://issues.apache.org/jira/browse/AMBARI-13665?focusedCommentId=14983679&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-14983679], this is a patch that are consolidated from some patches that have been already reviewed and committed to trunk. We internally tested and observed the patch does not affect HDP package.	1	1	https://issues.apache.org/jira/browse/AMBARI-13665

#	Requirement	Questions	Answered	Successfully Answered	URL
19331	AMBARI-13665	[~mithmatt] [~adenisso] What was the testing done to ensure the changes (esp to stackadvisor, ambari-web) do not regress HDP deployments?	1	1	https://issues.apache.org/jira/browse/AMBARI-13665
19332	AMBARI-13665	Hi [~jluniya], do you really insist on adding those checks ? These lines are based on the patch committed to trunk 3 weeks ago as part of https://issues.apache.org/jira/browse/AMBARI-13442 and the changes you propose now were not proposed then. The way I see it, these lines add a few metadata entries, do not cause any performance impact and should not cause any issues if the services are not part of the stack. In fact, it might not even be possible to learn which services are actually available in this place in code and making the relevant data available might cause more widespread changes.	1	0	https://issues.apache.org/jira/browse/AMBARI-13665
19333	AMBARI-13665	[~jaoki] Disabling failing unit tests makes me much more nervous. What is the reason behind having to disable the unit test? [~jaimin] can you review the UI changes.	1	1	https://issues.apache.org/jira/browse/AMBARI-13665
19334	AMBARI-13665	error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/configuration/gpcheck-env.xml: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/configuration/hawq-limits-env.xml: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/configuration/hawq-site.xml: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/configuration/hawq-sysctl-env.xml: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/metainfo.xml: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/common.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/constants.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/hawqmaster.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/hawqsegment.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/hawqstandby.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/hawqstatus.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/master_helper.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/params.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/service_check.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/scripts/utils.py: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/templates/hawq-hosts.j2: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/templates/hawq-profile.sh.j2: No such file or directory error: ambari-server/src/main/resources/common-services/HAWQ/2.0.0/package/templates/slaves.j2: No such file or directory error: ambari-server/src/main/resources/common-services/PXF/3.0.0/configuration/pxf-site.xml: No such file or directory error: ambari-server/src/main/resources/common-services/PXF/3.0.0/metainfo.xml: No such file or directory error: ambari-server/src/main/resources/common-services/PXF/3.0.0/package/scripts/params.py: No such file or directory error: ambari-server/src/main/resources/common-services/PXF/3.0.0/package/scripts/pxf.py: No such file or directory error: ambari-server/src/main/resources/common-services/PXF/3.0.0/package/templates/pxf-env.j2: No such file or directory error: patch failed: ambari-server/src/main/resources/stacks/HDP/2.3/role_command_order.json:8 error: ambari-server/src/main/resources/stacks/HDP/2.3/role_command_order.json: patch does not apply error: patch failed: ambari-server/src/main/resources/stacks/HDP/2.3/services/stack_advisor.py:22 error: ambari-server/src/main/resources/stacks/HDP/2.3/services/stack_advisor.py: patch does not apply error: ambari-server/src/test/python/stacks/2.3/common/hosts-1-host.json: No such file or directory error: ambari-server/src/test/python/stacks/2.3/common/hosts-3-hosts.json: No such file or directory error: ambari-server/src/test/python/stacks/2.3/common/services-hawq-1-host.json: No such file or directory error: ambari-server/src/test/python/stacks/2.3/common/services-hawq-3-hosts.json: No such file or directory error: ambari-server/src/test/python/stacks/2.3/common/services-master_ambari_colo-3-hosts.json: No such file or directory error: ambari-	1	1	https://issues.apache.org/jira/browse/AMBARI-13665

#	Requirement	Questions	Answered	Successfully Answered	URL
19335	AMBARI-13450	===== --- ambari-web/app/controllers/wizard/step3_controller.js (date 1444764262000) +++ ambari-web/app/controllers/wizard/step3_controller.js (revision) @@ -270,6 +270,7 @@ 'sshKey': this.get('content.installOptions.sshKey'), 'hosts': this.getBootstrapHosts(), 'user': this.get('content.installOptions.sshUser'), + 'sshPort': this.get('content.installOptions.sshPort'), 'userRunAs': App.get('supports.customizeAgentUserAccount') ? this.get('content.installOptions.agentUser') : 'root' }); App.router.get(this.get('content.controllerName')).launchBootstrap(bootStrapData, function (requestId) { @@ -444,6 +445,7 @@ 'sshKey': this.get('content.installOptions.sshKey'), 'hosts': hosts.mapProperty('name'), 'user': this.get('content.installOptions.sshUser'), + 'sshPort': this.get('content.installOptions.sshPort'), 'userRunAs': App.get('supports.customizeAgentUserAccount') ? this.get('content.installOptions.agentUser') : 'root' }); this.set('numPolls', 0); Index: ambari-web/app/messages.js IDA additional info: Subsystem: com.intellij.openapi.diff.impl.patch.CharsetEP <+>UTF-8 ===== --- ambari-web/app/messages.js (date 1444764262000) +++ ambari-web/app/messages.js (revision) @@ -571,6 +571,9 @@ 'installer.step2.sshUser.toolTip':'The user account used to install the Ambari Agent on the target host(s) via SSH. This user must be set up with passwordless SSH and sudo access on all the target host(s)', 'installer.step2.sshUser.placeholder':'Enter user name', 'installer.step2.sshUser.required':'User name is required', + 'installer.step2.sshPort':'SSH Port Number', + 'installer.step2.sshPort.toolTip':'SSH Port Number', + 'installer.step2.sshPort.required':'SSH Port Number is required.', 'installer.step2.agentUser':'Ambari Agent User Account', 'installer.step2.agentUser.toolTip':'The user account used to run the Ambari Agent daemon on the target host(s). This user must be set up with passwordless sudo access on all the target host(s)', 'installer.step2.bootStrap.error':'Errors were encountered while setting up Ambari Agents on the hosts.', Index: ambari-web/app/templates/wizard/step2.hbs IDA additional info: Subsystem: com.intellij.openapi.diff.impl.patch.CharsetEP <+>UTF-8 ===== --- ambari-web/app/templates/wizard/step2.hbs (date 1444764262000) +++ ambari-web/app/templates/wizard/step2.hbs (revision) @@ -90,6 +90,18 @@ {{/if}} + + <label rel="tooltip" {{translateAttr title="installer.step2.sshPort.toolTip"}} class="ssh-port pull-left span4"> + {{t installer.step2.sshPort}} + </label> + + <div {{bindAttr class="sshPortError:error :control-group"}}> + {{view view.textFieldView valueBinding="content.installOptions.sshPort" isEnabledBinding="content.installOptions.useSsh" }} + {{#if sshPortError}} + {{sshPortError}} + {{/if}} + + {{#if App.supports.customizeAgentUserAccount}} <label rel="tooltip" {{translateAttr title="installer.step2.agentUser.toolTip"}} class="ssh-user pull-left span4"> Index: ambari-server/src/main/java/org/apache/ambari/server/bootstrap/SshHostInfo.java IDA additional info: Subsystem: com.intellij.openapi.diff.impl.patch.CharsetEP <+>UTF-8 ===== --- ambari-server/src/main/java/org/apache/ambari/server/bootstrap/SshHostInfo.java (date 1444764262000) +++ ambari-server/src/main/java/org/apache/ambari/server/bootstrap/SshHostInfo.java (revision) @@ -49,6 +49,9 @@ private String user; @XmlElement + private String sshPort; + + @XmlElement private String password; @XmlElement @@ -85,6 +88,10 @@ public void setUser(String user) { this.user = user; } + + public	1	1	https://issues.apache.org/jira/browse/AMBARI-13450
19336	AMBARI-13442	Got 2 +1s on reviewboard. [~adenisso], could you please rebase the patch? I can not apply it.	1	1	https://issues.apache.org/jira/browse/AMBARI-13442
19337	AMBARI-13439	Committed to trunk 2f8636e5eafa4d9d26097a6e0d8f8dcde59d2e75 [~nalex], is this only for trunk?	1	0	https://issues.apache.org/jira/browse/AMBARI-13439
19338	AMBARI-13364	Can we make root_dir default to /opt/apache?	1	1	https://issues.apache.org/jira/browse/AMBARI-13364
19339	AMBARI-13364	[~afernandez] [~jluniya], can you help creating a branch to start on this work?	1	0	https://issues.apache.org/jira/browse/AMBARI-13364
19340	AMBARI-13364	Hi [~afernandez], [~jluniya], [~eyang], hope you all have a good Thanksgiving break. Can one of you help create the branch so we can proceed with this work?	1	0	https://issues.apache.org/jira/browse/AMBARI-13364
19341	AMBARI-13364	Hi [~jluniya], Juanjo and I just quickly scanned thru the patch you submitted, and we both agree that this is a good improvement on our original design. We are very happy that you are able to spend bandwidth on this. Juanjo is actively working on this feature. If Alex and other committers support your proposal, Juanjo can proceed refactoring the other services based on this new design. Thanks again for your help. We are looking forward for you leadership on this effort. [~afernandez], do you have any comment on Jayush's approach before Juanjo start refactoring?	1	0	https://issues.apache.org/jira/browse/AMBARI-13364
19342	AMBARI-13364	[~tctruong] Sure, I would be happy to guide on this. Can we have a quick chat on this and spec this out a bit more? I am free this afternoon.	1	0	https://issues.apache.org/jira/browse/AMBARI-13364
19343	AMBARI-13364	Attached patch (AMBARI-14435.patch) implementing the new parametrization approach for Zookeeper: https://issues.apache.org/jira/browse/AMBARI-14435 [~jluniya] [~afernandez] Per our last conversation, this patch removes the hardcoded stack versions in ZK common-services code (centralized version and parameterized stack configurations will be faced in other JIRAS) stack_features is now a property in .json format in configuration/cluster-env.xml for each stack version. With the structure: { "stack_features": [{ "name": "feature1", "description" : "Feature1 support", "min_version" : "Y.Y.Y.Y" , "max_version" : "X.X.X.X"} , ...] } where min_version/max_version are optional constraints. New resource_management/libraries/functions/stack_fetaures.py has been introduced to parse the json file and called from service code to check if the stack supports the required feature. Zookeeper fresh installation has been tested using HDP stack 2.3 and AMBARI-13364 branch code. Please, can you review when you have the chance and validate if it is the expected behavior? Would it fit in the other set of changes you are planning to do? Thanks	1	1	https://issues.apache.org/jira/browse/AMBARI-13364
19344	AMBARI-13364	[~jmarron] Can you submit the patch to the review board?	1	0	https://issues.apache.org/jira/browse/AMBARI-13364
19345	AMBARI-13205	This scripts run snmp command for v2 . Does this command works for SNMPv3?	0	0	https://issues.apache.org/jira/browse/AMBARI-13205
19346	AMBARI-12916	[~nitiraj.rathore], can you comment on this one?	0	0	https://issues.apache.org/jira/browse/AMBARI-12916
19347	AMBARI-12916	[~jimbobhickville]. I am concerned with _faking_ the Ambari server's hostname. There is lots of code that relies on the FQDN of the host that the Ambari server executes on and if we allow for an additional name to be used, than we will wind up with some interesting scenarios. My main concern is with Kerberos-related tasks where the _discovered_ FQDN of the host is used for many things. However the effects on other parts of Ambari may take a while to uncover. For example each view. Who will be tasked with testing each view to make sure nothing breaks when an alternate host name is used to access the API? Can you describe a valid use case for this?	1	0	https://issues.apache.org/jira/browse/AMBARI-12916
19348	AMBARI-12916	[~jimbobhickville]... I now have a better understanding of your issue and the need for your change. You are correct in that the name of the variable ({{masterHostname}}) is confusing. One thing that is still unclear to me is why this change is only relevant for Views. Doesn't SSL certificate verification come into play for Ambari agents as well? Shouldn't they use the a specific hostname when communicating with the Ambari server? If so, we need to make sure that when Ambari is bootstrapping the agents, it properly configures them with the _correct_ Ambari sever hostname.	1	0	https://issues.apache.org/jira/browse/AMBARI-12916

#	Requirement	Questions	Answered	Successfully Answered	URL
19349	AMBARI-12915	Hi Greg. For the record, can you are more details? What JVM were you using when you were trying suggestion? Not all version of JVM have fixed this issue. and the JDK workaround adding only apply to sun/openjdk releases that added this. You can still check DH length using openssl are suggested.	0	0	https://issues.apache.org/jira/browse/AMBARI-12915
19350	AMBARI-12885	[~Tim Thorpe] Can you please add some steps on how to install a custom service as an extension? I would like to try this out from a build of your patch and see if I can install a custom service on it seamlessly.	1	0	https://issues.apache.org/jira/browse/AMBARI-12885
19351	AMBARI-12885	Brand new install ambari-server. Couldn't start ambari-server. In /var/log/ambari-server/ambari-server-check-database.log, it says {code} 2016-06-22 20:39:40,252 ERROR - Unexpected error, database check failed com.google.inject.ProvisionException: Guice provision errors: 1) Error injecting method, javax.persistence.PersistenceException: Exception [EclipseLink-4002] (Eclipse Persistence Services - 2.6.2.v20151217-774c696): org.eclipse.persistence.exceptions.DatabaseException Internal Exception: org.postgresql.util.PSQLException: ERROR: relation "extensionlink" does not exist Position: 64 Error Code: 0 Call: SELECT t1.link_id, t1.extension_id, t1.stack_id FROM stack t0, extensionlink t1 WHERE (((t0.stack_name = ?) AND (t0.stack_version = ?)) AND (t0.stack_id = t1.stack_id)) bind => [2 parameters bound] Query: ReadAllQuery(name="ExtensionLinkEntity.findByStack" referenceClass=ExtensionLinkEntity sql="SELECT t1.link_id, t1.extension_id, t1.stack_id FROM stack t0, extensionlink t1 WHERE (((t0.stack_name = ?) AND (t0.stack_version = ?)) AND (t0.stack_id = t1.stack_id))") at org.apache.ambari.server.api.services.AmbariMetaInfo.init(AmbariMetaInfo.java:261) at org.apache.ambari.server.api.services.AmbariMetaInfo.class(AmbariMetaInfo.java:135) while locating org.apache.ambari.server.api.services.AmbariMetaInfo 1 error at com.google.inject.internal.InjectorImpl$4.get(InjectorImpl.java:987) at com.google.inject.internal.InjectorImpl.getInstance(InjectorImpl.java:1013) at org.apache.ambari.server.checks.DatabaseConsistencyCheckHelper.checkServiceConfigs(DatabaseConsistencyCheckHelper.java:401) at org.apache.ambari.server.checks.DatabaseConsistencyChecker.main(DatabaseConsistencyChecker.java:115) Caused by: javax.persistence.PersistenceException: Exception [EclipseLink-4002] (Eclipse Persistence Services - 2.6.2.v20151217-774c696): org.eclipse.persistence.exceptions.DatabaseException {code}	0	0	https://issues.apache.org/jira/browse/AMBARI-12885
19352	AMBARI-12860	Does OpenStack set resource constraint on the docker containers or is it default to use host resource limit?	0	0	https://issues.apache.org/jira/browse/AMBARI-12860
19353	AMBARI-12767	Hi [~pgunasekaran], are you still working on this issue? I have some ideas about implement this. Here is my concern: when a component die, ambari-agent should know it first through its monitoring behavior and it restart it with given number of retries which can be configurable.	0	0	https://issues.apache.org/jira/browse/AMBARI-12767
19354	AMBARI-12701	Can someone review [^AMBARI-12701.additional.fix.patch] ?	1	1	https://issues.apache.org/jira/browse/AMBARI-12701
19355	AMBARI-12556	Hi [~ncole], is there a design doc for this work that you can attach to the JIRA? I am interested in the details. Thank you..	0	0	https://issues.apache.org/jira/browse/AMBARI-12556
19356	AMBARI-12488	[~ncole], [~jonathan.hurley], can you review [^AMBARI-12488.v2.patch]? I tested it on an RU from HDP 2.2.x -> 2.2.y and forgot to add this.	1	1	https://issues.apache.org/jira/browse/AMBARI-12488
19357	AMBARI-12263	[~swagle], [~u39kun], can you comment on the more integrated LDAP/AD support for Ambari user management? I recalled Yusaku mentioned at the HackFest what there is a plan to synch group with LDAP then grant permission to group. When users log in, use LDAP/AD to get the user's group, and enable privileges base on the group. I am not seeing any other JIRA for that work. This JIRA is sort of related to the more integrated LDAP feature.	0	0	https://issues.apache.org/jira/browse/AMBARI-12263
19358	AMBARI-12263	Patch looks good. Thanks! We were able to successfully port it to current Ambari 2.4.0.1 Something we noticed is that in a secured cluster we have issues with the views, getting the following exception for the Hive view as an example: {code} Struct:TOpenSessionResp(status:TStatus(statusCode:ERROR_STATUS, infoMessages:[*org.apache.hive.service.cli.HiveSQLException:Failed to validate proxy privilege of ambari for org.apache.ambari.server.security.authorization.AmbariPamAuthenticationProvider$1@34511119:33:32, ... sqlState:08S01, errorCode:0, errorMessage:Failed to validate proxy privilege of ambari for org.apache.ambari.server.security.authorization.AmbariPamAuthenticationProvider$1@34511119), serverProtocolVersion:null) {code} As you can see it tries to impersonte {{"org.apache.ambari.server.security.authorization.AmbariPamAuthenticationProvider$1@34511119:33:32"}}. Changing the {{UsernamePasswordAuthenticationToken}} from {{Principal}} to username fixes this. So instead of : {code} UsernamePasswordAuthenticationToken token = new UsernamePasswordAuthenticationToken(principal, null, userAuthorities); {code} We use: {code} UsernamePasswordAuthenticationToken token = new UsernamePasswordAuthenticationToken(user.getUserName(), null, userAuthorities); {code} What could potential also work is, overriding {{toString}} of the principal like: {code} Principal principal = new Principal() { @Override public String getName() { return user.getUserName(); } @Override public String toString(){ return user.getUserName().toString(); } }; {code} We did not test this! As a little side note, I notices you are using String concatenation in your error logging like this: {{LOG.error("Message"+ ex.getMessage())}} I think the {{public void error(String msg, Throwable t);}} interface would be preferable in such scenarios, so: {{LOG.error("Message", ex)}}	1	0	https://issues.apache.org/jira/browse/AMBARI-12263
19359	AMBARI-12263	[~ghugare] I had issues merging the patch into trunk. Can you fix? {noformat} git apply ~/Downloads/AMBARI-12263.patch ... error: patch failed: ambari-server/pom.xml:1436 error: ambari-server/pom.xml: patch does not apply error: patch failed: ambari-server/src/main/java/org/apache/ambari/server/configuration/Configuration.java:2355 error: ambari-server/src/main/java/org/apache/ambari/server/configuration/Configuration.java: patch does not apply error: patch failed: ambari-server/src/main/java/org/apache/ambari/server/controller/internal/UserPrivilegeResourceProvider.java:196 error: ambari-server/src/main/java/org/apache/ambari/server/controller/internal/UserPrivilegeResourceProvider.java: patch does not apply error: patch failed: ambari-server/src/main/java/org/apache/ambari/server/upgrade/UpgradeCatalog250.java:27 error: ambari-server/src/main/java/org/apache/ambari/server/upgrade/UpgradeCatalog250.java: patch does not apply error: patch failed: ambari-server/src/test/java/org/apache/ambari/server/security/authorization/TestUsers.java:216 error: ambari-server/src/test/java/org/apache/ambari/server/security/authorization/TestUsers.java: patch does not apply error: patch failed: ambari-server/src/test/java/org/apache/ambari/server/upgrade/UpgradeCatalog250Test.java:30 error: ambari-server/src/test/java/org/apache/ambari/server/upgrade/UpgradeCatalog250Test.java: patch does not apply {noformat}	1	1	https://issues.apache.org/jira/browse/AMBARI-12263

#	Requirement	Questions	Answered	Successfully Answered	URL
19360	AMBARI-12263	[~ghugare], The latest patch looks good, but fails to apply to the trunk {noformat} error: patch failed: ambari-server/src/main/java/org/apache/ambari/server/configuration/Configuration.java:2459 error: ambari-server/src/main/java/org/apache/ambari/server/configuration/Configuration.java: patch does not apply {noformat} Can you rebase and resubmit the patch to this JIRA and the review?	1	1	https://issues.apache.org/jira/browse/AMBARI-12263
19361	AMBARI-12263	Hello Henning, Can you please provide me steps/details to reproduce the hive view issue?	0	0	https://issues.apache.org/jira/browse/AMBARI-12263
19362	AMBARI-11601	[~tedyu], do you know if "/apps/hbase/staging" will be created by HBase and available?	1	1	https://issues.apache.org/jira/browse/AMBARI-11601
19363	AMBARI-11573	[~tedyu], is this change applicable to all versions of HBase (e.g. ones used by HDP-2.1 to 2.3)?	1	1	https://issues.apache.org/jira/browse/AMBARI-11573
19364	AMBARI-11480	FAILURE: Integrated in Ambari-trunk-Commit #2739 (See [https://builds.apache.org/job/Ambari-trunk-Commit/2739/]) AMBARI-11480 - Do not use ranger related configs until Ranger is installed (abaranchuk: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=869709914977e5671cf7207dc5da18fe65bf2ad7) * ambari-server/src/main/resources/common-services/HBASE/0.96.0.2.0/package/scripts/params_linux.py * ambari-server/src/main/resources/common-services/HIVE/0.12.0.2.0/package/scripts/params_linux.py * ambari-server/src/main/resources/common-services/YARN/2.1.0.2.0/package/scripts/resourcemanager.py * ambari-server/src/main/resources/common-services/YARN/2.1.0.2.0/package/scripts/params_linux.py	0	0	https://issues.apache.org/jira/browse/AMBARI-11480
19365	AMBARI-11350	[~sposetti], [~rlevas], with the latest update, the admin and operator roles are still reversed. Is this intentional? This seems to be inconsistent with IT standard.	1	0	https://issues.apache.org/jira/browse/AMBARI-11350
19366	AMBARI-11350	[~rlevas] In the description, the table has no mention of the "View User" role. This role is there in the attached PDF file. Can you clarify on the use-case for this role. Also, it seems that non-"View User" roles do not have access to use view per the table info. Can you clarify this? I would think that cluster level users should be able to use the view if the admin grant them access to the view.	1	0	https://issues.apache.org/jira/browse/AMBARI-11350
19367	AMBARI-11314	A little late to comment myself. Patch 2 looks good. Should address make sure that the tabletserver will have sufficient memory after index+data block cache and sort buffers during log recovery. Thanks for making the change, Billie!	0	0	https://issues.apache.org/jira/browse/AMBARI-11314
19368	AMBARI-11290	In that case, should we bump up the multiplier in (5 * # of RS) ? Threshold for Warning is 150 and Critical is currently set to 250, so for 1 node HBase it starts up with Critical.	0	0	https://issues.apache.org/jira/browse/AMBARI-11290
19369	AMBARI-11268	Di Li, Thanks for doing this work! I've tried to review the design and the patch to see if it meets all of the needs of a custom service. I don't know all of the internals of Ambari well enough to claim that I understand all of the changes, but I would like to ask one question... One of the changes I've had to make to add a quick link for a custom service was to update the requiredSiteNames list in ambari-web/app/views/common/quick_view_link_view.js. If the custom service specifies "site": "custom-site" in its quicklinks.json, will the Ambari GUI know to load custom-site?	0	0	https://issues.apache.org/jira/browse/AMBARI-11268
19370	AMBARI-11246	[~mahadev], [~u39kun], [~afernandez] Can you review the attached patch? This is to update pom.xml version in 2.0.maint branch.	1	1	https://issues.apache.org/jira/browse/AMBARI-11246
19371	AMBARI-11049	[~ncole@hortonworks.com], this change removed the hadoop_prefix property from Accumulo's params.py (previously set to /usr/hdp/current/hadoop-client). Accumulo needs this property. The conf_select utility does not seem to have a method that returns the hadoop home directory. Do you have any suggestions for the best way to obtain this directory?	1	1	https://issues.apache.org/jira/browse/AMBARI-11049
19372	AMBARI-11003	Tested on live cluster. Tests pass. Failing tests are not related to this JIRA's work {noformat} Results : Tests in error: testCreateResources(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest) testCreateResources(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest) testGetResources(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest): Could not initialize class org.apache.ambari.server.controller.internal.ActiveWidgetLayoutResourceProvider testGetResources(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest) testValidateRepositoryVersion(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest): Could not initialize class org.apache.ambari.server.controller.internal.ActiveWidgetLayoutResourceProvider testValidateRepositoryVersion(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest) testDeleteResources(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest): Could not initialize class org.apache.ambari.server.controller.internal.ActiveWidgetLayoutResourceProvider testDeleteResources(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest) testUpdateResources(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest): Could not initialize class org.apache.ambari.server.controller.internal.ActiveWidgetLayoutResourceProvider testUpdateResources(org.apache.ambari.server.controller.internal.RepositoryVersionResourceProviderTest) Tests run: 2966, Failures: 0, Errors: 10, Skipped: 20 [INFO] -- [INFO] BUILD FAILURE [INFO] -- [INFO] Total time: 42:34.351s [INFO] Finished at: Thu May 07 14:13:27 PDT 2015 [INFO] Final Memory: 29M/332M [INFO] -- [ERROR] Failed to execute goal org.apache.maven.plugins:maven-surefire-plugin:2.12:test (default-test) on project ambari-server: There are test failures. [ERROR] [ERROR] Please refer to /Users/sgunturi/dev/git/ambari/ambari-server/target/surefire-reports for the individual test results. [ERROR] -> [Help 1] [ERROR] [ERROR] To see the full stack trace of the errors, re-run Maven with the -e switch. [ERROR] Re-run Maven using the -X switch to enable full debug logging. [ERROR] [ERROR] For more information about the errors and possible solutions, please read the following articles: [ERROR] [Help 1] http://cwiki.apache.org/confluence/display/MAVEN/MojoFailureException {noformat}	1	1	https://issues.apache.org/jira/browse/AMBARI-11003
19373	AMBARI-10992	Resolving this. Everything planned for 2.1 is done now (overall it included 10 jiras). There are some other things we can do, but they are bit risky, I plan them for 2.2. The time of *start of all services reduced for me from ~27minutes to ~13minutes* Also install tasks reduced time much as well (don't the benchmarks for that).	0	0	https://issues.apache.org/jira/browse/AMBARI-10992
19374	AMBARI-10873	[~fbarca], [~afernandez] [~sumitmohanty] Can one of you review the attached patch? Its a one line change?	1	1	https://issues.apache.org/jira/browse/AMBARI-10873
19375	AMBARI-10873	{{hadoop-distcp}} was taken out. Is that OK?	0	0	https://issues.apache.org/jira/browse/AMBARI-10873
19376	AMBARI-10848	[~lazurinke], Why do you think that Ambari does not support integration with Active Directory? Have you seen the featured added via AMBARI-7204?	1	1	https://issues.apache.org/jira/browse/AMBARI-10848

#	Requirement	Questions	Answered	Successfully Answered	URL
19377	AMBARI-10848	Hi Robert. Actually I have not seen this JIRA beforehand. and therefore i believed that active directory was not supported . Looks good. Can this jira create the required scripts and other artificats required in order set up kerberos correctly "offline" , based on the client's topology ? If the answer is no. i think this JIRA (AMBARI-10709), could be a good match to support this effort . The resources and scripts were tested with various customers and seem to be quite mature .	1	1	https://issues.apache.org/jira/browse/AMBARI-10848
19378	AMBARI-10777	Since this patch is shipping with IOP 4.1, when do you think you can share it?	0	0	https://issues.apache.org/jira/browse/AMBARI-10777
19379	AMBARI-10709	[~lazurinke], Does AMBARI-9783 address any of your issues with enabled Kerberos via Ambari?	1	0	https://issues.apache.org/jira/browse/AMBARI-10709
19380	AMBARI-10675	[~srimanth.gunturi], can you take a look?	1	0	https://issues.apache.org/jira/browse/AMBARI-10675
19381	AMBARI-10675	[~tedyu], {{hbase.regionserver.wal.codec}} has already been added to HDP-2.2 stack - [https://github.com/apache/ambari/blob/trunk/ambari-server/src/main/resources/stacks/HDP/2.2/services/HBASE/configuration/hbase-site.xml#L110]. Its default value is {{org.apache.hadoop.hbase.regionserver.wal.WALCellCodec}}, which will change to {{org.apache.hadoop.hbase.regionserver.wal.IndexedWALEditCodec}} when Phoenix is enabled. Can this config be removed from the patch?	1	1	https://issues.apache.org/jira/browse/AMBARI-10675
19382	AMBARI-10645	Could any expert please help assign this jira to me? I already made a patch and did necessary tests to ensure it works as expectation. So I attach the patch here and also opened a review request: #33428 . Thanks!	1	0	https://issues.apache.org/jira/browse/AMBARI-10645
19383	AMBARI-10645	Hi Experts, It's really weird that this patch could be applied to trunk in my local env, but always failed on Apache jira system which said(https://builds.apache.org/job/Ambari-trunk-test-patch/2758//console): error: src/main/resources/stacks/HDP/2.2/services/stack_advisor.py: No such file or directory PATCH APPLICATION FAILED However, current patch plans to change file 'ambari-server/src/main/resources/stacks/HDP/2.2/services/stack_advisor.py', not 'src/main/resources/stacks/HDP/2.2/services/stack_advisor.py', as it includes following sentences: --- ambari-server/src/main/resources/stacks/HDP/2.2/services/stack_advisor.py +++ ambari-server/src/main/resources/stacks/HDP/2.2/services/stack_advisor.py This patch can be applied to trunk branch in my local env and works well. Could any expert please help take a review on this issue? Thanks in advance!	1	0	https://issues.apache.org/jira/browse/AMBARI-10645
19384	AMBARI-10490	[~afernandez], can you push to dist as well?	1	0	https://issues.apache.org/jira/browse/AMBARI-10490
19385	AMBARI-10472	Hi Nate, Ambari repositories are distinct versions, such as HDP/2.0, and HDP/2.1. This has certain user experience issues associated with updating ambari.repo to point to the correct repository to install Ambari and stack of software. Wouldn't it be better if we manage RPMs more like Redhat, where there is a Major release repository, and all minor releases are just symlinks to the major version repositories? Any compatible changes are actually put in "updates" directory for incremental updates. This will reduce the tight coupling of stack definition and compatible stacks for upgrade purposes. In CentOS, they can simply replace the URL to migrate from Redhat owned repository to their own CentOS stack of software. This removes the need for Ambari to manage yum repository metadata. Ambari only tracks dependencies to be satisified cross node dependencies, which yum repository doesn't track. Stack compatibility can be tracked using RPM signing key package metadata. This will ensure that future version can track what version that they are backward compatible with without having software dependency logic in Ambari.	1	1	https://issues.apache.org/jira/browse/AMBARI-10472
19386	AMBARI-10457	[~u39kun] [~jaimin] It looks like the builds for components went fine. [INFO] Ambari Main ..... SUCCESS [4.166 s] [INFO] Apache Ambari Project POM ..... SUCCESS [0.069 s] [INFO] Ambari Web ..... SUCCESS [11.442 s] [INFO] Ambari Views ..... SUCCESS [1.330 s] [INFO] Ambari Admin View ..... SUCCESS [15.069 s] [INFO] Ambari Metrics Common ..... SUCCESS [1.258 s] [INFO] Ambari Server ..... SUCCESS [02:36 min] [INFO] Ambari Agent ..... SUCCESS [2.506 s] [INFO] Ambari Client ..... SUCCESS [0.048 s] [INFO] Ambari Python Client ..... SUCCESS [0.470 s] [INFO] Ambari Groovy Client ..... SUCCESS [2.669 s] [INFO] Ambari Shell ..... SUCCESS [0.035 s] [INFO] Ambari Python Shell ..... SUCCESS [0.390 s] [INFO] Ambari Groovy Shell ..... SUCCESS [1.069 s] [INFO] -- [INFO] BUILD SUCCESS [INFO] -- [INFO] Total time: 03:17 min [INFO] Finished at: 2015-04-13T23:13:50+00:00 [INFO] Final Memory: 81M/419M [INFO] -- before checkReleaseAuditWarnings 1 Can you please take a look why the tests failed.	1	1	https://issues.apache.org/jira/browse/AMBARI-10457
19387	AMBARI-10457	[~sriharsha] The failure is not related to this patch. However, can you post the patch in ReviewBoard so that we can post feedback?	1	1	https://issues.apache.org/jira/browse/AMBARI-10457
19388	AMBARI-10385	Is this feature usable ?What's more to be done here ? Why is this still "Unresolved"?	0	0	https://issues.apache.org/jira/browse/AMBARI-10385
19389	AMBARI-10253	Are there any more details on this JIRA? What is the content of this JIRA and what has been done? I am asking to find out what and how Ambari supports SSL in many of the stack components.	0	0	https://issues.apache.org/jira/browse/AMBARI-10253
19390	AMBARI-10196	+1 LGTM What does the endpoint and sample output look like?	0	0	https://issues.apache.org/jira/browse/AMBARI-10196
19391	AMBARI-10106	Is the title on this JIRA correct? The patch looks like it addresses HBase master and not YARN RM HA.	0	0	https://issues.apache.org/jira/browse/AMBARI-10106
19392	AMBARI-10029	Is any documentation going to be added around this? I didn't see any in the Ambari Confluence or on the Ambari website. The only documentation that seems to exist is the PDF that is attached to this JIRA.	0	0	https://issues.apache.org/jira/browse/AMBARI-10029
19393	AMBARI-10023	[~vijikarthi] I've reviewed the proposal and it looks good to me, how far along are you with the development of this? Are you targeting for 2.1 release of Ambari? When you have a working patch available, let me know and I can assist in testing and flushing out any potential issues.	1	0	https://issues.apache.org/jira/browse/AMBARI-10023
19394	AMBARI-9998	Hi, any progress on this activity? What is the timeline? Thank you	0	0	https://issues.apache.org/jira/browse/AMBARI-9998
19395	AMBARI-9993	Here's a patch that introduces basic service info for Phoenix query server. How do I test it? Are there build instructions for Ambari someplace? I see nothing obvious in the repository root (ie, BUILD.txt).	1	1	https://issues.apache.org/jira/browse/AMBARI-9993
19396	AMBARI-9993	[~ndimiduk], thank you for your contribution. Can you upload the patch to reviews.apache.org so others can comment? A couple of comments, daemon_script = format('/usr/hdp/current/phoenix/bin/queryserver.py') doesn't seen format since it's just a string Are the packages already available on the various OS'es? Your service will likely also need start/stop/status to be functional. Can you confirm these work. Lastly, we typically also add unit tests for new services, take a look at ambari-server/src/test/python/stacks/2.3/MAHOUT/test_mahout_service_check.py as an example.	1	1	https://issues.apache.org/jira/browse/AMBARI-9993
19397	AMBARI-9993	I noticed in my debugging that the PQS is now automatically selected by default on the "slaves and clients select" step when HBase is installed. I think the intention is that this is not selected by default. How can this behavior be modified?	1	1	https://issues.apache.org/jira/browse/AMBARI-9993

#	Requirement	Questions	Answered	Successfully Answered	URL
19398	AMBARI-9993	Good discussion happening over on the RB [~afernandez], [~jluniya]. Can you please confirm what additional changes I need to make to the addendum posted here? It seems I also need to merge the contents of stacks/HDP/2.2/services/HBASE/configuration into stacks/HDP/2.3/services/HBASE/configuration -- is this correct? Please confirm. Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-9993
19399	AMBARI-9993	[~ndimiduk] Can you create a review board for the addendum patch?	1	0	https://issues.apache.org/jira/browse/AMBARI-9993
19400	AMBARI-9993	[~jluniya] have you looked at the addendum patch? Did I handle the config files properly? The file hbase-site.xml has diverged between 2.2 and 2.3; how should these changes be resolved? Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-9993
19401	AMBARI-9993	[~aonishuk] I see you committed AMBARI-10626 which undoes much of this patch. Can you ensure my addendum changes are integrated correctly?	1	1	https://issues.apache.org/jira/browse/AMBARI-9993
19402	AMBARI-9908	[~aantonenko], could you take a look at AMBARI-9908-test-fix.1.patch and see if that is an acceptable fix for the test?	1	1	https://issues.apache.org/jira/browse/AMBARI-9908
19403	AMBARI-9687	{code} org.apache.ambari.server.controller.AmbariManagementControllerTest.testUpdateRepoUrlController Failing for the past 1 build (Since Failed#1799) Took 15 sec. Error Message Could not access base url . http://public-repo-1.hortonworks.com/HDP-1.1.1.16/repos/centos5 . java.net.SocketTimeoutException: Read timed out Stacktrace java.lang.IllegalArgumentException: Could not access base url . http://public-repo-1.hortonworks.com/HDP-1.1.1.16/repos/centos5 . java.net.SocketTimeoutException: Read timed out at org.apache.ambari.server.controller.AmbariManagementControllerImpl.verifyRepository(AmbariManagementControllerImpl.java:3117) at org.apache.ambari.server.controller.AmbariManagementControllerImpl.updateRepositories(AmbariManagementControllerImpl.java:3040) at org.apache.ambari.server.controller.AmbariManagementControllerTest.testUpdateRepoUrlController(AmbariManagementControllerTest.java:8161) {code} This was happening before this commit (see https://builds.apache.org/job/Ambari-trunk-Commit/1797/testReport/)	1	0	https://issues.apache.org/jira/browse/AMBARI-9687
19404	AMBARI-9612	Maybe I'm missing something, but I don't see how can this patch be responsible for the UpgradeResourceProviderTest failure, since this issue was spotted in trunk and tracked by AMBARI-9624.	1	1	https://issues.apache.org/jira/browse/AMBARI-9612
19405	AMBARI-9609	[~tbeerbower], can you upload the patch to ReviewBoard so it's easier to see?	1	1	https://issues.apache.org/jira/browse/AMBARI-9609
19406	AMBARI-9527	Actually, Tom pointed out that setting it on an Item will also probably work. Can you double check an upgrade from scratch and verify it works?	0	0	https://issues.apache.org/jira/browse/AMBARI-9527
19407	AMBARI-9527	[~ncole], I tried it earlier today, set UpgradeItem in ABORTED state, and Upgrade transitioned into HOLDING state. Is it right, which fix is right?	1	1	https://issues.apache.org/jira/browse/AMBARI-9527
19408	AMBARI-9527	[~tbeerbower], can you comment on this?	1	1	https://issues.apache.org/jira/browse/AMBARI-9527
19409	AMBARI-9468	Reverted since it seems like this is causing Add Services to fail with NPE. {code} java.lang.NullPointerException at org.apache.ambari.server.controller.AmbariManagementControllerImpl.updateCluster(AmbariManagementControllerImpl.java:1219) at org.apache.ambari.server.controller.AmbariManagementControllerImpl.updateClusters(AmbariManagementControllerImpl.java:1154) at org.apache.ambari.server.controller.internal.ClusterResourceProvider$2.invoke(ClusterResourceProvider.java:245) at org.apache.ambari.server.controller.internal.ClusterResourceProvider$2.invoke(ClusterResourceProvider.java:242) at org.apache.ambari.server.controller.internal.AbstractResourceProvider.modifyResources(AbstractResourceProvider.java:331) at org.apache.ambari.server.controller.internal.ClusterResourceProvider.updateResources(ClusterResourceProvider.java:242) at org.apache.ambari.server.controller.internal.ClusterControllerImpl.updateResources(ClusterControllerImpl.java:310) at org.apache.ambari.server.api.services.persistence.PersistenceManagerImpl.update(PersistenceManagerImpl.java:100) at org.apache.ambari.server.api.handlers.UpdateHandler.persist(UpdateHandler.java:42) at org.apache.ambari.server.api.handlers.BaseManagementHandler.handleRequest(BaseManagementHandler.java:72) at org.apache.ambari.server.api.services.BaseRequest.process(BaseRequest.java:135) at org.apache.ambari.server.api.services.BaseService.handleRequest(BaseService.java:103) at org.apache.ambari.server.api.services.BaseService.handleRequest(BaseService.java:72) at org.apache.ambari.server.api.services.ClusterService.updateCluster(ClusterService.java:151) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) {code} Also, the property names in ams-hbase-security-site.xml are malformed. Was this tested E2E?	1	0	https://issues.apache.org/jira/browse/AMBARI-9468

#	Requirement	Questions	Answered	Successfully Answered	URL
19410	AMBARI-9439	SUCCESS: Integrated in Ambari-trunk-Commit #1727 (See [https://builds.apache.org/job/Ambari-trunk-Commit/1727/]) AMBARI-9439. Kerberos: Do not validate host health or maintenance state when enabling Kerberos (rlevas) (rlevas: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=c775c60710095259f9c089378d7cd960393ea49f) * ambari-server/src/main/java/org/apache/ambari/server/orm/entities/KerberosPrincipalHostEntityPK.java * ambari-server/src/main/resources/common-services/KERBEROS/1.10.3-10/package/scripts/params.py * ambari-server/src/main/java/org/apache/ambari/server/serveraction/kerberos/DestroyPrincipalsServerAction.java * ambari-server/src/main/resources/common-services/KERBEROS/1.10.3-10/package/scripts/kerberos_common.py * ambari-server/src/main/java/org/apache/ambari/server/upgrade/UpgradeCatalog200.java * ambari-server/src/test/java/org/apache/ambari/server/controller/KerberosHelperTest.java * ambari-server/src/main/resources/Ambari-DDL-Postgres-EMBEDDED-CREATE.sql * ambari-server/src/main/resources/Ambari-DDL-Postgres-CREATE.sql * ambari-server/src/main/resources/Ambari-DDL-Oracle-CREATE.sql * ambari-server/src/test/java/org/apache/ambari/server/agent/HeartBeatHandlerInjectKeytabTest.java * ambari-server/src/main/java/org/apache/ambari/server/orm/entities/KerberosPrincipalHostEntity.java * ambari-server/src/main/java/org/apache/ambari/server/controller/KerberosHelper.java * ambari-server/src/main/resources/Ambari-DDL-MySQL-CREATE.sql * ambari-server/src/main/java/org/apache/ambari/server/serveraction/kerberos/CreatePrincipalsServerAction.java * ambari-server/src/main/resources/META-INF/persistence.xml * ambari-server/src/main/java/org/apache/ambari/server/orm/dao/KerberosPrincipalDAO.java * ambari-server/src/test/java/org/apache/ambari/server/upgrade/UpgradeCatalog200Test.java * ambari-server/src/main/java/org/apache/ambari/server/serveraction/kerberos/KerberosServerAction.java * ambari-server/src/test/java/org/apache/ambari/server/agent/TestHeartbeatHandler.java * ambari-server/src/main/resources/Ambari-DDL-SQLServer-CREATE.sql * ambari-server/src/main/java/org/apache/ambari/server/orm/dao/KerberosPrincipalHostDAO.java * ambari-server/src/main/resources/common-services/KERBEROS/1.10.3-10/package/scripts/status_params.py * ambari-server/src/main/java/org/apache/ambari/server/agent/HeartBeatHandler.java * ambari-server/src/main/java/org/apache/ambari/server/orm/entities/KerberosPrincipalEntity.java * ambari-server/src/main/java/org/apache/ambari/server/serveraction/kerberos/CreateKeytabFilesServerAction.java * ambari-server/src/main/resources/common-services/KERBEROS/1.10.3-10/package/scripts/kerberos_client.py	0	0	https://issues.apache.org/jira/browse/AMBARI-9439
19411	AMBARI-9363	Thanks for your contribution [~Tim Thorpe]. Will you be doing further refactoring of the RCO?	1	0	https://issues.apache.org/jira/browse/AMBARI-9363
19412	AMBARI-9255	Hi, Can you post a screen shot? Any errors in your browser JavaScript console?	0	0	https://issues.apache.org/jira/browse/AMBARI-9255
19413	AMBARI-9224	[~brandonli], can you please make sure that the tests pass. Doing a {{mvn clean test}} in {{ambari/}} should succeed.	1	1	https://issues.apache.org/jira/browse/AMBARI-9224
19414	AMBARI-9224	[~brandonli] Patch looks good, can you add unit tests for your changes? Unit tests will guarantee that unwarranted future changes will be caught in dev cycle.	1	0	https://issues.apache.org/jira/browse/AMBARI-9224
19415	AMBARI-9224	[~swagle], thanks for the review! Is there any example I use as a reference to add the unit test?	1	1	https://issues.apache.org/jira/browse/AMBARI-9224
19416	AMBARI-9100	This is still a problem - did any work get done, or was this abandoned?	0	0	https://issues.apache.org/jira/browse/AMBARI-9100
19417	AMBARI-9029	[~jaimin] Can we move this to 2.1.0? cc: [~rlevas]	1	1	https://issues.apache.org/jira/browse/AMBARI-9029
19418	AMBARI-9021	[~atkach], you have been working on this feature. Can you link appropriate JIRAs and mark this as a dup?	0	0	https://issues.apache.org/jira/browse/AMBARI-9021
19419	AMBARI-9014	[~rlevas] Is there a ReviewBoard for this JIRA? If so, please create a link to it.	1	1	https://issues.apache.org/jira/browse/AMBARI-9014
19420	AMBARI-8997	[~ncole], Review Board will not allow me to upload the attached patch for some reason. If you have some time, could you review the patch attached to this Jira? Thanks.	1	0	https://issues.apache.org/jira/browse/AMBARI-8997
19421	AMBARI-8995	ReviewBoard is still not working for me, [~ncole], [~yshylov], [~jonathan.hurley], can you review [http://reviews.llvm.org/D6860]?	1	1	https://issues.apache.org/jira/browse/AMBARI-8995
19422	AMBARI-8878	[~fbarca], [~mahadev], [~jspeidel], [~u39kun] Can you +1 on the attached patch? This change is refactoring of YARN service. Testing Done: 1. Installed HDP stacks on Linux 2. Installed HDPWIN stacks on Windows mvn clean test	1	1	https://issues.apache.org/jira/browse/AMBARI-8878
19423	AMBARI-8804	Did an HDP 2.2 install with all services except AMS after applying patches for AMBARI-8804, AMBARI-8779, AMBARI-8782 and verified that all services install without issues and all alert definitions come up green.	1	1	https://issues.apache.org/jira/browse/AMBARI-8804
19424	AMBARI-8775	[~jspeidel] Can you provide ETA for resolving this JIRA and AMBARI-9143 and whether it will be done as part of AMBARI 2.0.0 release?	1	0	https://issues.apache.org/jira/browse/AMBARI-8775
19425	AMBARI-8724	+1 - I have comments b/c I'm confused and don't fully understand the system yet and want to make sure what I'm seeing isn't a boog. {noformat} - <task xsi:type="execute"> + <task xsi:type="execute" hosts="master"> {noformat} - Can this be hosts="active" maybe? It's a little confusing; makes it seem like you're saying "Run it on all masters" when I think we're saying "Run it only on the main master (ignoring others in standby". {noformat} // !!! revisit if needed if (hostsType.master != null && hostsType.secondary != null) { {noformat} - Should this be secondary? I would think that a secondary master (like Secondary NN) would not affect this task, but another standby master would {noformat} if (execwrappers.size() > 0) { results.addAll(execwrappers); {noformat} - Change to !isEmpty() to appease the anti-size() gods?	1	0	https://issues.apache.org/jira/browse/AMBARI-8724
19426	AMBARI-8695	[~jspeidel], [~ncole], [~mahadev], [~rlevas], [~swagle], [~u39kun] I have attached the patch for the JIRA as reviewboard is having issues. Can you review the latest patch? I would like to get this committed today as this blocks refactoring of other services in the HDP and HDPWIN stacks.	1	1	https://issues.apache.org/jira/browse/AMBARI-8695
19427	AMBARI-8695	Hi [~jluniya] Excited to see the progress on Common Service! Are we taking 2 steps (in 2 tickets), just like 1) Taking services out from resources/stacks/STACK_A/n.n.n/services/SERVICE_B to to resources/common-services/SERVICE_B/n.n.n (which still has the vendor specific hard code values [~adenisso] mentioned previously) 2) Then get rid of the vendor specific hard code from resources/common-services/SERVICE_B/n.n.n to make it more abstract, less vendor-specific. Is this ticket tacking care of only 1) and do we have another ticket to take care of 2) ?	1	0	https://issues.apache.org/jira/browse/AMBARI-8695
19428	AMBARI-8695	[~jluniya], if I can break down your previous comment to [2 steps that I referred earlier|https://issues.apache.org/jira/browse/AMBARI-8695?focusedCommentId=14250434&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-14250434], You are focusing 1) and moving some services configured specific to HDP under common-service. Then you (or anyone from the community) will focus on 2) and removing HDP specific logic from common-services to make it abstract. Is that right?	0	0	https://issues.apache.org/jira/browse/AMBARI-8695

#	Requirement	Questions	Answered	Successfully Answered	URL
19429	AMBARI-8648	[~mpapirkovskyy] [~swagle] Was the additional patch for phoenix tests committed? Can you resolve the JIRA if it was?	0	0	https://issues.apache.org/jira/browse/AMBARI-8648
19430	AMBARI-8610	Robert instead of requiring the user to use a REST API to retrieve the CSV of principals - can this instead be a download button in the add hosts/services wizard menu as is done when initially kerberizing a new cluster? That seems like it would be much more practical and user friendly.	1	0	https://issues.apache.org/jira/browse/AMBARI-8610
19431	AMBARI-8452	[~rlevas] [~bergenholtz] Is this JIRA still valid and required for 2.0.0? Can you update the fix version accordingly?	1	0	https://issues.apache.org/jira/browse/AMBARI-8452
19432	AMBARI-8447	[~rlevas], I'm trying to understand the requirements for this better. Could you provide a more concrete use case that I could step through? Thanks.	1	0	https://issues.apache.org/jira/browse/AMBARI-8447
19433	AMBARI-8447	Add API support to allow for session attributes to be set for a cluster. Expose the {{session_attributes}} property on a cluster resource. To add session attributes for a cluster the {{session_attributes}} property can be set on a cluster resource. The value of the {{session_attributes}} should be a collection of name value pairs, as follows ... {code} PUT api/v1/clusters/c1 [{ "session_attributes" : { "attr1" : "v1", "attr2" : "v2", "attr3" : {"sub1" : "v31", "sub2" : "v32"} } }] {code} The above example would add the following attributes to the session for cluster c1: ||attribute||value |attr1 | v1 |attr2 | v2 |attr3/sub1 | v31 |attr3/sub2 | v32 The cluster session attribute property is not available to read through the REST API. The map of cluster session attribute properties is available to resource providers on the Ambari server through the following method on the {{org.apache.ambari.server.state.Clusters}} singleton... {code} /** * Get the map of session attributes for the cluster identified by the given name. * * @param name the cluster name * * @return the map of session attributes for the cluster */ public Map<String, Object> getSessionAttributes(String name); {code} All session attributes set through the above API are scoped by cluster name and session. When the session dies, so do the attribute values for that session. [~rlevas], [~jspeidel], Could you review the above the above?	1	0	https://issues.apache.org/jira/browse/AMBARI-8447
19434	AMBARI-8435	[~xiwang0309] [~u39kun] Is this JIRA still valid and take action accordingly?	0	0	https://issues.apache.org/jira/browse/AMBARI-8435
19435	AMBARI-8426	[~rlevas], If we expose the following on AmbariSessionManager ... * setAttribute(String key, Object attribute) * getAttribute(String key) * removeAttribute(String key) Does that meet your requirements? The AmbariSessionManager could be injected into any ResourceProvider instance.	1	0	https://issues.apache.org/jira/browse/AMBARI-8426
19436	AMBARI-8420	SUCCESS: Integrated in Ambari-trunk-Commit #1165 (See [https://builds.apache.org/job/Ambari-trunk-Commit/1165/]) AMBARI-8420. Add Amabri Metric Provider to get metrics from AMS. Check liveliness. (swagle) (swagle: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=d2bc0ac8193e13988879bde4078a7406bacd0935) * ambari-server/src/main/java/org/apache/ambari/server/controller/metrics/timeline/AMSPropertyProvider.java	0	0	https://issues.apache.org/jira/browse/AMBARI-8420
19437	AMBARI-8357	[~rlevas] [~darumugam] [~bergenholtz] Is this JIRA still relevant?	1	0	https://issues.apache.org/jira/browse/AMBARI-8357
19438	AMBARI-8279	SUCCESS: Integrated in Ambari-trunk-Commit #910 (See [https://builds.apache.org/job/Ambari-trunk-Commit/910/]) AMBARI-8279. Slider View: "+ Create App" button effect is disabled, but is clickable (alexantonenko) (hiveww: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=6b98cfc084e9194ff4df578796953d1f8e93d642) * contrib/views/slider/src/main/resources/ui/app/styles/app.less	0	0	https://issues.apache.org/jira/browse/AMBARI-8279
19439	AMBARI-8279	FAILURE: Integrated in Ambari-trunk-Commit-docker #201 (See [https://builds.apache.org/job/Ambari-trunk-Commit-docker/201/]) AMBARI-8279. Slider View: "+ Create App" button effect is disabled, but is clickable (alexantonenko) (hiveww: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=6b98cfc084e9194ff4df578796953d1f8e93d642) * contrib/views/slider/src/main/resources/ui/app/styles/app.less	0	0	https://issues.apache.org/jira/browse/AMBARI-8279
19440	AMBARI-8269	[~fbarca], [~u39kun], [~mahadev] Can you review the patch on the JIRA itself as ReviewBoard fails to load a large patch?	1	0	https://issues.apache.org/jira/browse/AMBARI-8269
19441	AMBARI-8269	Thanks [~jluniya] Can you break up this big patch into several pieces, like by subproject or other chunks so that it's easier for reviewers to review and make comments?	1	1	https://issues.apache.org/jira/browse/AMBARI-8269
19442	AMBARI-8269	[~fbarca] Can you look at these unit test failures? They don't fail on Mac.	0	0	https://issues.apache.org/jira/browse/AMBARI-8269
19443	AMBARI-8250	[~rnettleton] [~bergenholtz] Can we close this JIRA if all Blueprint HA functionality has been committed?	1	0	https://issues.apache.org/jira/browse/AMBARI-8250
19444	AMBARI-8247	[~jspeidel], {quote} The descriptor is quite lengthy and it seems that it should have it's own resource. {quote} Good point, but where would the descriptor resource live? Are you suggesting that the descriptors be sub-resources of root or cluster (or something else) and just referenced from the stack service?	1	0	https://issues.apache.org/jira/browse/AMBARI-8247
19445	AMBARI-8247	[~rlevas] This ticket does not add necessary kerberos.json to all stack services. Stacks API won't expose kerberos related information until this patch along with necessary kerberos.json files are committed. Any particular reason for not committing kerberos.json to all services along with this work? Is that being tracked as different ticket for some reason ?	1	0	https://issues.apache.org/jira/browse/AMBARI-8247
19446	AMBARI-8225	Thanks [~nalex]. This is a great start! A few things: 1) When showing sync status on the services summary area, if a subset of hosts are out of sync, what will it display? Think a red icon with "X of Y hosts out of sync" 2) I think that text would link to hosts page so users can find the out of sync hosts 3) Would need to think about how this is shown and filtered per host on the Hosts page so a user can find host(s) that have out of sync config status 4) And on the Host details page, show sync icon info with each component (i.e. my DataNode component might be out of sync but not the NodeManager on a given host) 5) Also will need a phase 3: on save config, if configs are out of sync, how will the flow go. Should prompt user to proceed + overwrite or merge?	1	0	https://issues.apache.org/jira/browse/AMBARI-8225

#	Requirement	Questions	Answered	Successfully Answered	URL
19447	AMBARI-8225	Thanks [~sposetti] 1) When showing sync status on the services summary area, if a subset of hosts are out of sync, what will it display? Think a red icon with "X of Y hosts out of sync" Newton> Sounds like a good idea. Was initially thinking about just a yes/no like solution. 2) I think that text would link to hosts page so users can find the out of sync hosts Newton> Thats correct. I briefly mentioned about it in the Web UI section. The new page will have the details provided in the "JSON structures" section of the doc 3) Would need to think about how this is shown and filtered per host on the Hosts page so a user can find host(s) that have out of sync config status Newton> Thanks, I missed this one. I'll add a section for the Hosts screen too. I'm thinking of this one to be similar to the Service screen summary. When the user clicks on the link, it could take them to a details page that only shows the differences for that host. 4) And on the Host details page, show sync icon info with each component (i.e. my DataNode component might be out of sync but not the NodeManager on a given host) Newton> In the current design proposed, the JSON structure deals at the level of files. So in order to implement this requirement, we will have to map the config files back to components. I guess this can be done as an improvement? 5) Also will need a phase 3: on save config, if configs are out of sync, how will the flow go. Should prompt user to proceed + overwrite or merge? Newton> Sure, will add a phase 3 to the doc. Prompting user to proceed with overwrite should be easy and straight forward but auto-merging might be a bit complicated. For example, if the user has done similar changes on a bunch of nodes, we will have to find all nodes with similar changes, group them into the same config groups and deploy them. Unless, we do some more intelligent diff'ing at the agent side, it might be difficult.	1	0	https://issues.apache.org/jira/browse/AMBARI-8225
19448	AMBARI-8211	[~dmitriusan], can you please commit this? The patch has two +1 in code review.	1	0	https://issues.apache.org/jira/browse/AMBARI-8211
19449	AMBARI-8189	Hi [~aonishuk], parallel execution is disabled by default, so this patch shouldn't have changed anything, it must be something else. When was the last time you measured?	1	0	https://issues.apache.org/jira/browse/AMBARI-8189
19450	AMBARI-8189	Thanks Dmitry. Hm... I am not sure I understand, if parallel execution is disabled, all commands are executed serially, same way as before. There are no threads being created based on what I can tell... can you please check the patch and let us know which piece causes parallel execution? {code} + if self.parallel_execution == 0: + command = self.commandQueue.get(True, self.EXECUTION_COMMAND_WAIT_TIME) + self.process_command(command) {code}	1	0	https://issues.apache.org/jira/browse/AMBARI-8189
19451	AMBARI-8189	[~ivanmi] Which version of AMbari has this been merged to ? Is this available trough some Ambari properties? A quick answer will be appreciated.	0	0	https://issues.apache.org/jira/browse/AMBARI-8189
19452	AMBARI-8189	[~ivanmi] I also see that this patch is already in Ambari 2.2.2. I could set parallel_execution=1 in ambari. I am yet to see the test results. Can you please confirm/share your test results and any pitfall of enabling this? Another parallel question: The default concurrency is set to 5 in Ambari code. Can this be exposed through ambari-agent.ini property?	0	0	https://issues.apache.org/jira/browse/AMBARI-8189
19453	AMBARI-8166	There are 2 checks mentioned here. - One from Ambari Server where you will do connectivity check (and potentially credential check). - Second from Cluster hosts (in Ambari agent) where you will do connectivity check. My understanding is Ambari server validates that Server can reach out to KDC and do any administrative actions it needs to perform on KDC/AD with the provided credentials. The agent checks are to ensure that there are no connection issues between cluster host and KDC server as it has to reach out to KDC server for cluster operation. Are we saying check in agent is not necessary and we will assume (pretty safe, in my opinion) that if Ambari Server can reach KDC (IP/hostname + port), then so can all the cluster nodes?	1	1	https://issues.apache.org/jira/browse/AMBARI-8166
19454	AMBARI-8166	[~rpidva], I think we are saying that a check on the agent side _is_ necessary. It is may be possible (though probably rare) for the Ambari server to reach the KDC, but one or more hosts in the cluster to be blocked - firewall, host configuration, etc... Regarding authenticating with the KDC, that will (according to [~jaimin]) be a step later in the process - after the Kerberos client is installed.	1	1	https://issues.apache.org/jira/browse/AMBARI-8166
19455	AMBARI-8166	[~jaimin], [~rlevas] -- a follow-up question. In terms of workflow, is our assumption that this point KDC server has already been setup and it's up and running (i.e. listening on KDC port)? If yes, why do we have external AD/KDC details on this screen? Are we going to modify the configuration of that KDC? If no, then cluster hosts will not be able to connect to specified KDC host on specified port.	1	1	https://issues.apache.org/jira/browse/AMBARI-8166
19456	AMBARI-8166	[~rlevas], external KDC is fine but I am asking about internal KDC or Ambari-managed KDC. Do you mean it will also be set at this point?	1	1	https://issues.apache.org/jira/browse/AMBARI-8166

#	Requirement	Questions	Answered	Successfully Answered	URL
19457	AMBARI-8166	clean test -Dtest=KdcServerConnectionVerificationTest} ---------- Test set: org.apache.ambari.server.api.rest.KdcServerConnectionVerificationTest ---------- Tests run: 5, Failures: 1, Errors: 0, Skipped: 0, Time elapsed: 0.189 sec <<< FAILURE! testWithoutPortFailure(org.apache.ambari.server.api.rest.KdcServerConnectionVerificationTest) Time elapsed: 0.004 sec <<< FAILURE! java.lang.AssertionError at org.junit.Assert.fail(Assert.java:92) at org.junit.Assert.assertTrue(Assert.java:43) at org.junit.Assert.assertFalse(Assert.java:68) at org.junit.Assert.assertFalse(Assert.java:79) at org.apache.ambari.server.api.rest.KdcServerConnectionVerificationTest.testWithoutPortFailure(KdcServerConnectionVerificationTest.java:86) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:45) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:15) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:42) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:20) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.runners.ParentRunner.runLeaf(ParentRunner.java:263) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:68) at org.junit.runners.BlockJUnit4ClassRunner.runChild(BlockJUnit4ClassRunner.java:47) at org.junit.runners.ParentRunner$3.run(ParentRunner.java:231) at org.junit.runners.ParentRunner$1.schedule(ParentRunner.java:60) at org.junit.runners.ParentRunner.runChildren(ParentRunner.java:229) at org.junit.runners.ParentRunner.access$000(ParentRunner.java:50) at org.junit.runners.ParentRunner$2.evaluate(ParentRunner.java:222) at org.junit.internal.runners.statements.RunBefores.evaluate(RunBefores.java:28) at org.junit.internal.runners.statements.RunAfters.evaluate(RunAfters.java:30) at org.junit.runners.ParentRunner.run(ParentRunner.java:300) at org.apache.maven.surefire.junit4.JUnit4Provider.execute(JUnit4Provider.java:236) at org.apache.maven.surefire.junit4.JUnit4Provider.executeTestSet(JUnit4Provider.java:134) at org.apache.maven.surefire.junit4.JUnit4Provider.invoke(JUnit4Provider.java:113) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:39) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:25) at java.lang.reflect.Method.invoke(Method.java:597) at org.apache.maven.surefire.util.ReflectionUtils.invokeMethodWithArray(ReflectionUtils.java:189) at org.apache.maven.surefire.booter.ProviderFactory$ProviderProxy.invoke(ProviderFactory.java:165) at org.apache.maven.surefire.booter.ProviderFactory.invokeProvider(ProviderFactory.java:85) at org.apache.maven.surefire.booter.ForkedBooter.runSuitesInProcess(ForkedBooter.java:103) at	1	1	https://issues.apache.org/jira/browse/AMBARI-8166
19458	AMBARI-8166	[~jaimin], do you have a local KDC running on your machine? If yes, that would explain it. The test which failed in your local environment makes an assumption that nothing is listening on port 88 (stated in comments). We can change the port to something which is typically not used by any service like 11111.	1	1	https://issues.apache.org/jira/browse/AMBARI-8166
19459	AMBARI-8156	SUCCESS: Integrated in Ambari-trunk-Commit #863 (See [https://builds.apache.org/job/Ambari-trunk-Commit/863/]) AMBARI-8156. Usability: Do not force a mysql restart to apply Hive changes (via ncole) (ncole: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=af31f3ec343043bb39e54175e7deb3f6950fa481) * ambari-server/src/main/resources/stacks/HDP/2.0.6/services/HIVE/metainfo.xml * ambari-server/src/test/resources/stacks/HDP/2.0.5/services/HIVE/metainfo.xml * ambari-server/src/main/resources/stacks/HDP/2.1/services/HIVE/metainfo.xml	0	0	https://issues.apache.org/jira/browse/AMBARI-8156
19460	AMBARI-8156	SUCCESS: Integrated in Ambari-trunk-Commit-docker #157 (See [https://builds.apache.org/job/Ambari-trunk-Commit-docker/157/]) AMBARI-8156. Usability: Do not force a mysql restart to apply Hive changes (via ncole) (ncole: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=af31f3ec343043bb39e54175e7deb3f6950fa481) * ambari-server/src/test/resources/stacks/HDP/2.0.5/services/HIVE/metainfo.xml * ambari-server/src/main/resources/stacks/HDP/2.1/services/HIVE/metainfo.xml * ambari-server/src/main/resources/stacks/HDP/2.0.6/services/HIVE/metainfo.xml	0	0	https://issues.apache.org/jira/browse/AMBARI-8156
19461	AMBARI-8150	[~afernandez] Is this JIRA going to fixed as part of 2.0.0, if not can you update the fix version?	1	0	https://issues.apache.org/jira/browse/AMBARI-8150
19462	AMBARI-8147	[~afernandez] Is AMBARI-8150 going to be part of 2.0.0? If so can you provide ETA for resolving AMBARI-8150 and AMBARI-8147. Else can we update the JIRAs to 2.1.0?	1	0	https://issues.apache.org/jira/browse/AMBARI-8147
19463	AMBARI-8120	[~abaranchuk] Can you update the description to summarize the issues?	0	0	https://issues.apache.org/jira/browse/AMBARI-8120
19464	AMBARI-8022	[~tedyu@apache.org] / [~devaraj] There's v1 and v2 posted. Which one do we want? v1 says to remove "hbase.hstore.flush.retries.number", while v2 does not have this entry. From the description, we do want to remove it, correct? Also, if we want 4GB as the default heapsize, shouldn't the value be 4096? Is 4196 intentional?	1	1	https://issues.apache.org/jira/browse/AMBARI-8022
19465	AMBARI-7998	[~ncole@hortonworks.com], how does one trigger this event via the API? I am interested in this because it is more an an action call versus a call to update a resource and I may need to do something similar in the near future.	0	0	https://issues.apache.org/jira/browse/AMBARI-7998
19466	AMBARI-7985	- _Server-side commands should be specified as a task in a stage_ Does this mean a server action is a part of Role Command order, and if it is, since all server actions have the same role AMBARI_SERVER_ACTION, how can it be specified in the ordering. - _each server-side task should be executed in its own thread_ A stage in Ambari is a way of grouping tasks that can be executed in parallel (one command per host), if server action follows max one action per stage, it would be executed serially by the Action Scheduler. I am not sure I understand, why the need for separate thread to execute server actions, could you explain the need for this ?	1	0	https://issues.apache.org/jira/browse/AMBARI-7985

#	Requirement	Questions	Answered	Successfully Answered	URL
19467	AMBARI-7985	[~swagle], {{quote}} Server-side commands should be specified as a task in a stage Does this mean a server action is a part of Role Command order, and if it is, since all server actions have the same role AMBARI_SERVER_ACTION, how can it be specified in the ordering {{quote} I guess this does mean that a server action is part of a Role Command order and thus will have the role {{AMBARI_SERVER_ACTION}}. Thinking about this more, I believe that a role can be specified only once within a stage, so multiple stages will need to be created to handle multiple server-side actions. With this, aren't stages ordered and guaranteed not to start unless the previous stage completes successfully? {{quote}} {{quote}} each server-side task should be executed in its own thread A stage in Ambari is a way of grouping tasks that can be executed in parallel (one command per host), if server action follows max one action per stage, it would be executed serially by the Action Scheduler. I am not sure I understand, why the need for separate thread to execute server actions, could you explain the need for this ? {{quote}} It appeared to me that the Action Scheduler pulled tasks from the queue and processed them by simply updating their states so that another process could come by and pick them up - somewhere in the Heartbeat monitor code. This is a quick task and could be done inline without causing a perceivable delay. However if we allow the server-side task to be handled in this thread, it could take some time and delay the queue processing. It may be that case that this is ok, especially since each server-side task may be handled in its own stage. However, are stages created by different request ordered? If not, then we may not want to hold the Action Scheduler hostage while we process our server-side task.	1	0	https://issues.apache.org/jira/browse/AMBARI-7985
19468	AMBARI-7985	_With this, aren't stages ordered and guaranteed not to start unless the previous stage completes successfully?_ Yes that will be guaranteed by the Action Execution framework. - There is a ServiceActionManager which is invoked from ActionScehduler, I don't think we should create separate ActionScheduler for server actions. - Another thing that is unclear to me is, if the server action is a Host Role Command (HRC), then this is basically something queued up to be executed by the agent running on Ambari server? Or is it a java class that is a part of a jar file and in the class path of AmabriServer? - If it is a conventional HRC, it should be a part of the Action queue for the server host. (There is 1 queue of actions per host). *Note*: In certain Ambari deployments the Amabri server might not have an agent running on the same host as the server. - The ActionScehduler, is non-blocking it reads stages from DB and queues up actions, the delay is that heartbeats are resposible for dequeue and heartbeat happen every 6 seconds. Currently there is no delay for server actions, I see the point of spawning a thread because this is a synchronous op.	1	0	https://issues.apache.org/jira/browse/AMBARI-7985
19469	AMBARI-7985	Would there ever be a case where a server action and a host action would be grouped together and need to be run in serial? If not, then I would that we'd want a different scheduler handling the server-side tasks. Reimplementing the ServerActionManager to handle this might be the cleanest way to go.	1	0	https://issues.apache.org/jira/browse/AMBARI-7985
19470	AMBARI-7985	[~rlevas] - It looks like your patch increased the test time from 35 minutes to over 47 minutes causing the builds to timeout. How many tests did you add? Are they spinning up a database - that's usually what adds time to the test suite. I've actually never seen a patch increase the test time this much.	1	0	https://issues.apache.org/jira/browse/AMBARI-7985
19471	AMBARI-7985	Can this feature be Stack driven? I can define a service level customized action that can be invoked in Ambari Server.	1	1	https://issues.apache.org/jira/browse/AMBARI-7985
19472	AMBARI-7985	[~cloneman], Currenly this is not stack driven and required server-side code to setup and invoked the command. Ideally this can be made to be more dynamic. What sort of stack-driven commands do you have in mind?	1	0	https://issues.apache.org/jira/browse/AMBARI-7985
19473	AMBARI-7890	[~kminder] [~sumit.gupta] The description of this ticket is taken from the link: http://knox.apache.org/books/knox-0-5-0/knox-0-5-0.html#Secure+Clusters Most of the properties that needs to be changed has *knox* in it's name. For example {{hadoop.proxyuser.knox.groups}} *What does knox part in the property name signify over here ?* I want to make sure that when user customizes the service name for knox service or primary name of the knox principal (knox@EXAMPLE.COM) stated in the Jaas file, the new configs that are injected in other service's site.xml holds true	0	0	https://issues.apache.org/jira/browse/AMBARI-7890
19474	AMBARI-7811	All 3 integration as independent services would be better idea I guess. With HDP 2.4 log search preview is already made with ambari infra and it uses solr so definitely ELK is not coming for default log analysis server. But elastic search integration definitely make sense. My team is already working on it since we are integrating predictionio which needs elastic search for meta store so it make sense for us to integrate.	1	1	https://issues.apache.org/jira/browse/AMBARI-7811
19475	AMBARI-7811	May I ask how did you integrate elasticsearch to Ambari? Did you use any public Ambari's elasticsearch services on Github?	1	1	https://issues.apache.org/jira/browse/AMBARI-7811
19476	AMBARI-7807	[~jaimin] can you please assign this to me	1	0	https://issues.apache.org/jira/browse/AMBARI-7807
19477	AMBARI-7804	What about an upgrade using a custom rpm? Would this be supported. I was just thinking for a patch delivered for a customer, it might be nice to allow the user to select the rpms instead of only repos.	0	0	https://issues.apache.org/jira/browse/AMBARI-7804
19478	AMBARI-7804	[~ncole@hortonworks.com] The mockups and designs mostly revolve around the availability of versioned rpms. Not sure if you have thought through screens for non-versioned distributions like BIGTOP ? Features like Downgrade might not work for non-versioned rpms. Will there be provisions to disable such features based on the stack?	1	0	https://issues.apache.org/jira/browse/AMBARI-7804
19479	AMBARI-7799	[~sumit.gupta] The [doc link|http://knox.apache.org/books/knox-0-5-0/knox-0-5-0.html#Secure+Clusters] asks to update krb5.conf at */etc/knox/conf/krb5.conf* on Knox host. The default location for krb5.conf is /etc/krb5.conf. Does knox overrides KRB5_CONFIG environment variable to */etc/knox/conf/krb5.conf* ?	1	1	https://issues.apache.org/jira/browse/AMBARI-7799
19480	AMBARI-7767	[~chliu] Thanks for submitting the patch. Could you write up concrete use cases that you have in mind that will be addressed by this new feature?	1	0	https://issues.apache.org/jira/browse/AMBARI-7767
19481	AMBARI-7767	[~chuanliu], can you take a look at the updated patch.	1	0	https://issues.apache.org/jira/browse/AMBARI-7767
19482	AMBARI-7681	Shouldn't the title should be add metrics service to HDP stack instead of common services stack?	0	0	https://issues.apache.org/jira/browse/AMBARI-7681
19483	AMBARI-7679	SUCCESS: Integrated in Ambari-trunk-Commit #1184 (See [https://builds.apache.org/job/Ambari-trunk-Commit/1184/]) AMBARI-7679. Add psutil based resource monitoring to collect host metrics. Unitests.py python version. (swagle) (swagle: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=8235e7d63fbb0fb4f30ba209e3ac7f95bf3bf76f) * ambari-metrics/ambari-metrics-host-monitoring/src/test/python/unitTests.py	0	0	https://issues.apache.org/jira/browse/AMBARI-7679
19484	AMBARI-7656	Hi, [~267ada9d] Do you mean to let Ambari manage the start and stop of HBase Rest Gateway servers? Do you have any update on this? Thanks,	1	1	https://issues.apache.org/jira/browse/AMBARI-7656
19485	AMBARI-7656	Hi, I have upload the patch for this feature. could you please help me to take a look at ? Because REST api is a native api for hbase, so I add this changes in common-services.	1	1	https://issues.apache.org/jira/browse/AMBARI-7656

#	Requirement	Questions	Answered	Successfully Answered	URL
19486	AMBARI-7656	Hi, [~xieshiju] Can you add your hbase_restgatewayserver.py to the patch too?	1	1	https://issues.apache.org/jira/browse/AMBARI-7656
19487	AMBARI-7656	Hi Jerry, the failed test case is for HDP 2.2(ambari-server/src/test/python/stacks/2.2/common/test_stack_advisor.py), but my patch is on common-service. I think the failure is not caused by the patch, does it?	1	1	https://issues.apache.org/jira/browse/AMBARI-7656
19488	AMBARI-7621	[Yusaku Sako] I have attached the patch as a zip file as the patch itself is ~20MB and JIRA has a limit of 10MB. Can you take a look? hw11039:merge jluniya$ diff -qr --exclude=".git" win_hw win_apache hw11039:merge jluniya$	1	1	https://issues.apache.org/jira/browse/AMBARI-7621
19489	AMBARI-7617	FAILURE: Integrated in Ambari-trunk-Commit #461 (See [https://builds.apache.org/job/Ambari-trunk-Commit/461/]) AMBARI-7617. Cannot run build as root. (jaimin) (jaimin: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=354bdbac9291ba1ac6a335572ad8b521911d6f95) * ambari-admin/src/main/resources/ui/admin-web/package.json	1	1	https://issues.apache.org/jira/browse/AMBARI-7617
19490	AMBARI-7605	[~xiwang0309] Can you add unit tests for this?	1	0	https://issues.apache.org/jira/browse/AMBARI-7605
19491	AMBARI-7602	Stale patch. Can you resubmit the patch. cc: [~u39kun] [~jaimin] hw11039:ambari jluniya$ git apply ~/Downloads/AMBARI-7602.patch error: patch failed: ambari-web/app/styles/application.less:2448 error: ambari-web/app/styles/application.less: patch does not apply	1	0	https://issues.apache.org/jira/browse/AMBARI-7602
19492	AMBARI-7574	[~xiwang0309] Can you submit the patch on ReviewBoard. cc: [~u39kun] [~jaimin]	1	0	https://issues.apache.org/jira/browse/AMBARI-7574
19493	AMBARI-7574	{color:red}-1 overall{color}. Here are the results of testing the latest attachment http://issues.apache.org/jira/secure/attachment/12672184/AMBARI-7574.patch against trunk revision . {color:green}+1 @author{color}. The patch does not contain any @author tags. {color:red}-1 tests included{color}. The patch doesn't appear to include any new or modified tests. Please justify why no new tests are needed for this patch. Also please list what manual steps were performed to verify this patch. {color:green}+1 javac{color}. The applied patch does not increase the total number of javac compiler warnings. {color:green}+1 release audit{color}. The applied patch does not increase the total number of release audit warnings. {color:green}+1 core tests{color}. The patch passed unit tests in ambari-admin. Test results: https://builds.apache.org/job/Ambari-trunk-test-patch/4083//testReport/ Console output: https://builds.apache.org/job/Ambari-trunk-test-patch/4083//console This message is automatically generated.	1	1	https://issues.apache.org/jira/browse/AMBARI-7574
19494	AMBARI-7497	I got your point. displayOrder array is only to determine the preferable display order of the services and should not be iterated as it is being done. The attached patch does not apply cleanly on the current trunk bits. Can you create a new patch on the top of the latest trunk bits? Thanks [~wangy6]	1	1	https://issues.apache.org/jira/browse/AMBARI-7497
19495	AMBARI-7469	[~rlevas] Could you add several concrete use cases to the design doc?	1	1	https://issues.apache.org/jira/browse/AMBARI-7469
19496	AMBARI-7450	[~jaimin] Is this JIRA still valid? cc: [~bergenholtz] [~rlevas]	1	0	https://issues.apache.org/jira/browse/AMBARI-7450
19497	AMBARI-7449	[~rlevas] Is this JIRA planned for 2.0.0 or can we move it to 2.1.0	1	0	https://issues.apache.org/jira/browse/AMBARI-7449
19498	AMBARI-7417	Why was LDAP Sync removed from UI? If I recall correctly Hue does LDAP Sync from the UI and I don't see a problem - I'm not sure why you'd want to remove functionality like that from the admin UI interface?	0	0	https://issues.apache.org/jira/browse/AMBARI-7417
19499	AMBARI-7357	[~keyki] nice- did you get a chacne to run through unit tests? Also can you add some unit tests around here?	1	1	https://issues.apache.org/jira/browse/AMBARI-7357
19500	AMBARI-7357	Hi, [~u39kun] Indeed, I noticed that too. Should I do something about it? [~mahadev] Sorry, I forgot the tests, literally this was the first time I ever touched python code, but I'll look into this, and submit a new patch.	1	1	https://issues.apache.org/jira/browse/AMBARI-7357
19501	AMBARI-7357	Stale patch. Can you check if the issue is still valid and resubmit the patch. hw11039:ambari jluniya$ git apply ~/Downloads/AMBARI-7357-2.patch error: patch failed: ambari-agent/src/main/python/ambari_agent/hostname.py:75 error: ambari-agent/src/main/python/ambari_agent/hostname.py: patch does not apply	1	0	https://issues.apache.org/jira/browse/AMBARI-7357
19502	AMBARI-7325	[~dfj] Do you want to work on the patch for this?	0	0	https://issues.apache.org/jira/browse/AMBARI-7325
19503	AMBARI-7323	[~dfj] Do you want to work on the patch for this?	1	0	https://issues.apache.org/jira/browse/AMBARI-7323
19504	AMBARI-7302	Would this information include data relavant to the service itself (for example, SQL queries) or relevant to its relationship with Ambari (for example, can be managed, can be monitored, etc...)?	1	0	https://issues.apache.org/jira/browse/AMBARI-7302
19505	AMBARI-7292	Newton, should I commit it? Have we got maven profiles changes checked in?	1	1	https://issues.apache.org/jira/browse/AMBARI-7292
19506	AMBARI-7292	[~nalex] Following files are missing the License notice. Can you pl. fix them. Not sure if we need to checkin these files. /root/ambari-commits/ambari/ambari-server/src/main/resources/stacks/BIGTOP/0.8/services/HDFS/package/scripts/balancer-emulator/balancer.log /root/ambari-commits/ambari/ambari-server/src/main/resources/stacks/BIGTOP/0.8/services/HDFS/package/scripts/balancer-emulator/balancer-err.log	1	1	https://issues.apache.org/jira/browse/AMBARI-7292
19507	AMBARI-7246	[~keyki], can you please add use case(s) to the Jira?	1	1	https://issues.apache.org/jira/browse/AMBARI-7246
19508	AMBARI-7209	Patch available! I'm drawing the whole solution right here. There is a job to detect all Ambari Patch Available tickets. https://builds.apache.org/view/A-D/view/Ambari/job/Ambari-trunk-find-patches-available/ This currently points to one of files in the patch on my github account, and it is based upon oozie project's https://github.com/apache/oozie/blob/master/bin/test-patch-find-new-patch-available-jiras Once detected, each detected Jira number will be sent to the Ambari test-patch job. https://builds.apache.org/view/A-D/view/Ambari/job/Ambari-trunk-test-patch/ This will run a set of tests on the latest trunk with the patch applied. The script is based upon hadoop-common https://github.com/apache/hadoop-common/blob/trunk/dev-support/test-patch.sh What these jobs generate is Hadoop QA comments here https://issues.apache.org/jira/browse/AMBARI-7397 Once this whole solution is set in place, the committers can only focus on reviewing jiras that get +1 from Hadoop QA.	1	1	https://issues.apache.org/jira/browse/AMBARI-7209
19509	AMBARI-7209	[~jaoki] So we commit these scripts under "dev-support" and then update the repo pointer from your personal one to that of Ambari's in the Jenkins jobs? How do we activate it so that this runs automatically when "Submit Patch" is clicked?	1	0	https://issues.apache.org/jira/browse/AMBARI-7209
19510	AMBARI-7209	[~u39kun], thank you for looking into it. {quote} So we commit these scripts under "dev-support" and then update the repo pointer from your personal one to that of Ambari's in the Jenkins jobs? {quote} That's correct. {quote} How do we activate it so that this runs automatically when "Submit Patch" is clicked? {quote} This is not a hook on Jira, but rather we can let [the find patch available job|https://builds.apache.org/view/A-D/view/Ambari/job/Ambari-trunk-find-patches-available/] run periodically. That's what [oozie team does |https://builds.apache.org/job/oozie-trunk-find-patches-available/], I believe for every 10 mins	1	1	https://issues.apache.org/jira/browse/AMBARI-7209

#	Requirement	Questions	Answered	Successfully Answered	URL
19511	AMBARI-7209	[~jaoki] Once this feature is enabled, we need to formalize the dev process around this. Do you want to come up with a proposal so we can discuss on the dev mailing list?	1	0	https://issues.apache.org/jira/browse/AMBARI-7209
19512	AMBARI-7209	{quote} Do you want to come up with a proposal so we can discuss on the dev mailing list? {quote} Sure thing [~u39kun]. I will make a Google document to draw a proposal and will post it for discussion. {quote} Please update the Jenkins build jobs as needed. {quote} Will do. Thank you for taking care of the patch.	1	1	https://issues.apache.org/jira/browse/AMBARI-7209
19513	AMBARI-7204	Some review comments: 1. Support numerous cross-realm trusts. The cluster's home realm may require trust relationships with numerous other realms, both northward (trusted) and southward (trusting) of the cluster realm. For example, users from numerous AD realms may communicate with the cluster. The cluster may in turn communicate with numerous backend storage systems (each in an independent realm). 2. Improve the organization of the KDC information within the plan. The terms 'managed' and 'unmanaged' appear to serve numerous purposes and are overloaded. For example, in the current schema, how would one define an unmanaged KDC for the cluster plus a trust relationship with another (umanaged) realm? The term 'unmanaged' cannot be used twice as a key. My suggestion is to provide three sections: a. "cluster", defining the KDC for the Hadoop cluster itself, with either a 'managed' or 'unmanaged' section within. b. "trusted", defining a +list+ of realms that are to be trusted, each with a 'trust_password' (if the cluster KDC is managed). c. "trusting", defining a +list+ of realms that are to be trusting, each with a 'trust_password' (if the cluster KDC is managed). 3. Forwardable tickets. Please ensure that any managed KDC is configured to generate forwardable tickets. 4. Clarify whether 'trust' principals are automatically created in the 'unmanaged' cluster KDC scenario. I think it is clear that keytabs are created in every scenario. But the cross-realm trust principal for a given trust relationship, in what cases is it created by Ambari? My vote is, only in the managed scenario. 5. Provide a user extension point. For settings not contemplated in the proposal, perhaps the user may provide snippets to be included in krb5.conf.	1	0	https://issues.apache.org/jira/browse/AMBARI-7204
19514	AMBARI-7204	[~rlevas], well documented! Not knowing the Kerberos's best practice, I have a comment on the 2.1 Use Cases. To simplify, how about we implement only 2.1.1, 2.1.2 and 2.1.3? Clusters managed by Ambari always talk to the internal KDC, (basically 2.1.1 only as far as clusters and nodes can see) and if an organization has an existing KDC or active directory, we also provide a plugin or middleware to support 2.1.2 and 2.1.3 (one way trust to external source) This may not be needed if one way trust is a part of KDC. For use case 2.1.4 and 2.1.5 Ambari still set up a mandatory local KDC, we can simply apply an idea of 2.1.2 and 2.1.3 to let it connect to the existing ones.	1	0	https://issues.apache.org/jira/browse/AMBARI-7204
19515	AMBARI-7204	[~eronwright], thanks for your comments.... I will take them inconsideration while building this. Some comments on your comments.... {quote} 1. Support numerous cross-realm trusts. The cluster's home realm may require trust relationships with numerous other realms, both northward (trusted) and southward (trusting) of the cluster realm. For example, users from numerous AD realms may communicate with the cluster. The cluster may in turn communicate with numerous backend storage systems (each in an independent realm). {quote} I thought about this, however I could not seem to find a valid use case for any of these. I could imagine in a large enterprise that there may be multiple AD realms, and setting up the (one-way) cross-domain trusts for these will not be that difficult. My guess is that initially Ambari will not support such a case but it should not be that hard to add it later, if deemed necessary. {quote} 2. Improve the organization of the KDC information within the plan. The terms 'managed' and 'unmanaged' appear to serve numerous purposes and are overloaded. For example, in the current schema, how would one define an unmanaged KDC for the cluster plus a trust relationship with another (umanaged) realm? The term 'unmanaged' cannot be used twice as a key. My suggestion is to provide three sections: {quote} I will need to re-read the document. I didn't notice this issue. Essentially, there are two main cases with some sub-cases: # Managed KDC ## Stand-alone ## Cross-domain trust with unmanaged KDC # Unmanaged KDC ## KDC ## Active Directory I hope this clears it up a bit. {quote} 3. Forwardable tickets. Please ensure that any managed KDC is configured to generate forwardable tickets. {quote} This is the plan, thanks for pointing out that it wasn't clear in the design. {quote} 4. Clarify whether 'trust' principals are automatically created in the 'unmanaged' cluster KDC scenario. I think it is clear that keytabs are created in every scenario. But the cross-realm trust principal for a given trust relationship, in what cases is it created by Ambari? My vote is, only in the managed scenario. {quote} It should be assumed that whether the KDC is managed or unmanaged, principals for the Hadoop services will be created by Ambari. It think that it would be much easier if Ambari didn't have to create the principals in the unmanaged scenario; however that would defeat one of the purposes of this exercise which is to simplify he process of setting up Kerberos. {quote} 5. Provide a user extension point. For settings not contemplated in the proposal, perhaps the user may provide snippets to be included in krb5.conf. {quote} Let me think on this a while. I am not sure how "dangerous" it would be to let the user add their own configuration data to the krb5.conf file. However maybe the user is more educated on Kerberos than I assume. :)	1	1	https://issues.apache.org/jira/browse/AMBARI-7204
19516	AMBARI-7204	[~eronwright], I am in the process of updating the technical document. In it, I have an updated description of the API and will add a section on configuring the KDCs (single KDC, multiple cross-realm trusts, etc...). I know that JCE is needed the larger bit encryption algorithms like AES-256, however I am not aware of any JCE configuration requirements. Can you elaborate? In any case, I believe that the initial take on this feature will require JCE to be installed manually if needed.	1	1	https://issues.apache.org/jira/browse/AMBARI-7204
19517	AMBARI-7204	Thanks for incorporating much of the feedback. The introduction of the "kdc-<unique identifier>" configuration block seems to allow a multitude of realms to be defined in krb5.conf. Accordingly a use case should be defined. Seems we are now able to satisfy the use case I mentioned earlier, where the managed KDC is the 'trusted' party to some external KDC. But we must clarify how the cross-realm trust principals are specified. Would it be correct to define 'identity-' blocks for the trust principals? At minimum a password must be specified. Here is an example to illustrate the use case. A user in USER realm submits a job to a cluster in HADOOP realm, which reads data from a data storage system in DATABASE realm. The HADOOP realm is managed. Keep in mind that trust principals are defined as "krbtgt/TRUSTING_REALM@TRUSTED_REALM". {code} "configurations": [{ "cluster-env": { "kerberos_domain": "HADOOP.EXAMPLE.COM" } }, { "kdc--users": { "realm": "USERS.EXAMPLE.COM", "kdc_server": "kdc.users.example.com", "admin_server": "kdc.users.example.com" } }, { "kdc-hadoop": { "realm": "HADOOP.EXAMPLE.COM" } }, { "kdc--database": { "realm": "DATABASE.EXAMPLE.COM", "kdc_server": "kdc.database.example.com", "admin_server": "kdc.database.example.com" } }, { "identity--user-trust": { "realm": "HADOOP.EXAMPLE.COM", "principal": "krbtgt/HADOOP.EXAMPLE.COM@USERS.EXAMPLE.COM", "password": "hadoop" } }, { "identity--database-trust": { "realm": "HADOOP.EXAMPLE.COM", "principal": "krbtgt/DATABASE.EXAMPLE.COM@HADOOP.EXAMPLE.COM", "password": "hadoop" } }] {code} Thanks	1	0	https://issues.apache.org/jira/browse/AMBARI-7204

#	Requirement	Questions	Answered	Successfully Answered	URL
19518	AMBARI-7204	Since it is likely that any user-supplied KDC/AD administrator principals will not be properly secured due to the way Ambari handles (request) data, would it be unacceptable to declare that any KDC/AD administrative credentials are to be temporary and should be reset after any Ambari-related activity? What I mean by "declare" is via documentation and the UI.	0	0	https://issues.apache.org/jira/browse/AMBARI-7204
19519	AMBARI-7201	Nicely written document! Here are some comments from my side: 1) What would be the common service definitions tested against? Are we looking at Bigtop rpms? 2) Minor fix: On Page 3, in the directory tree: | | | | |- {Base: common-services/1.0.0/SERVICEA} should be | | | | |- {Base: common-services/SERVICEA/1.0.0} 3) I have some comments on the versioning of the service definitions. We are dealing with two variables here: i. The versions of the services. e.g. HDFS-2.4, HDFS-2.6 ii. The versions of the service definitions themselves e.g. 1.0.0, 1.0.1 Consider a scenario where we release the service definitions for HDFS-2.4 and call it HDFS-1.0.0. After some time we release the service definition for HDFS-2.5 and call it HDFS-1.1.0. Now if for some reason we need to patch the service definition for the older HDFS-2.4, what would the version number for such a fix look like? Do we call it 1.0.0.1, 1.0.1 or 1.1.1? Basically we will have to establish some guidelines on how we version the service definitions. Versioning the common services will help avoid vendor stack breakages. I can think of two approaches here: I> Establish a policy on how to increment the version numbers for the service definitions We can probably bump up the first digit for newer versions of services. E.g. HDFS-2.4 can be in common-services/HDFS/1.0.0 HDFS-2.5 can be in common-services/HDFS/2.0.0 And bump up the second digit for all patches or revisions to the service definition. E.g.: Version 2 of the service definition for HDFS-2.4 would be common-services/HDFS/1.1.0 Version 2 of the service definition for HDFS-2.5 would be common-services/HDFS/2.1.0 Pros: a. Compact representation Cons: a. The version numbers might not linearly increase. What if one vendor chooses to ship HDFS 2.5 but an another doesn't? How do we ensure that the versioning doesn't get messed up? II> Second option to have explicit versioning like the following: common-services/HDFS/{service_version}/version/{service_definition_version} E.g.: common-services/HDFS/2.4.1/versions/1.0 common-services/HDFS/2.4.1/versions/1.1 common-services/HDFS/2.6/versions/1.0 Pros: a. Natural/Easy to read service version hierarchy b. Avoid the necessity for internal version mapping between service and service-definition versions (inside the metainfo.xml files I guess?) c. If we don't end up supporting HDFS 2.5, then we can simply leave it empty. I personally lean towards the second approach. Let me know if any of the above approaches will work for us or if there is anything better? Thanks	1	1	https://issues.apache.org/jira/browse/AMBARI-7201
19520	AMBARI-7201	Sorry about the auto linking. Rewording to make it slightly better: "Consider a scenario where we release the service definitions for HDFS-2.4 and call it common-services/hdfs/1.0.0. After some time we release the service definition for HDFS-2.5 and call it common-services/hdfs/. Now if for some reason we need to patch the service definition for the older HDFS-2.4, what would the version number for such a fix look like? Do we call it common-services/hdfs/1.0.0.1, common-services/hdfs/1.0.1 or common-services/hdfs/1.1.1? "	1	1	https://issues.apache.org/jira/browse/AMBARI-7201
19521	AMBARI-7201	[~nalex], Thanks! bq. 1) What would be the common service definitions tested against? Are we looking at Bigtop rpms? I am not too sure. I was hoping the community will join together and ensure that we have a good place for that. bq. 2) Minor fix: I appreciate the fix for my documentation issue. I will correct it and upload a new PDF. bq. 3) I have some comments on the versioning of the service definitions. We are dealing with two variables here: I really like your ideas on versioning... thanks for pointing out the lack of details on this in the specification. You obviously put a lot of thought into this and I think that both schemas are great. My need to be explicit favors flavor II... even if it is a bit wordy. I will add one of these conventions to the document, but I want to see if anyone chimes in against or for either of those scenarios.	1	1	https://issues.apache.org/jira/browse/AMBARI-7201
19522	AMBARI-7201	[~nalex], Sorry for the delay on this... after discussing a bit with [~jspeidel], it seems like the 2 levels of version numbers may be overkill. I was trying to come up with a use case for multiple versions of a service declaration for a particular service version and failed to do so. Could you help with that. For example, why would we have {quote} common-services/HDFS/2.4.1/versions/1.0 common-services/HDFS/2.4.1/versions/1.1 common-services/HDFS/2.6/versions/1.0 {quote} Maybe we only really need to worry about one version of each {quote} common-services/HDFS/2.4.1 common-services/HDFS/2.6 {quote} My inclination is to be the most flexible, but maybe I was trying to be too flexible. Thoughts?	1	1	https://issues.apache.org/jira/browse/AMBARI-7201
19523	AMBARI-7201	[~rlevas] Here is what I had in mind while proposing the second set of versions. Imagine a scenario, where we release the service definitions for HDFS/2.4.1 and may be HDP and PHD inherit from it. At a later point of time we identify a bug in the HDFS service definition and we amend it. HDP might decide to skip the newer version but PHD might want to adopt the fix. If we don't have the version numbers for the service definitions itself, how would we identify the two different service definitions for HDFS/2.4.1 available now? I don't think retro-fixing any bugs on a already published version will be a good idea as it might break dependent stacks. Let me know if this sounds reasonable. Thanks	1	1	https://issues.apache.org/jira/browse/AMBARI-7201
19524	AMBARI-7201	[~nalex] Regarding your proposal for second set of versioning in common services, it would lead to more complications and versioning confusion as if common-services/HDFS/2.6 needs to inherit from common-services/HDFS/2.4.1, if 2.4.1 has 2 versions which version will the new 2.6 version inherit from or will we create new versions in the child service for each version of the parent service. Also this kinda defeats the purpose of common services if different stacks inherit from different versions of the same service version. Further, for the example you mentioned in which PHD stack needs a bug fix, we can always fix that bug in the PHD stack by overriding for example the default config property in the PHD stack. To keep things simple we will go without support for the second set of versioning. Let us know if you have any concerns.	1	1	https://issues.apache.org/jira/browse/AMBARI-7201

#	Requirement	Questions	Answered	Successfully Answered	URL
19525	AMBARI-7192	Retrieving which host components are installed to a specific host is already provided by the `clusters/<clustername>/hosts/<hostname>/host_components` resource. Retrieving the other way around is "sort of" provided with one small modification. The `clusters/<clustername>/host_components/<hostcomponent>` would be able to retrieve, but it is not allowed not to specify the <hostname> [here](https://github.com/apache/ambari/blob/trunk/ambari-server/src/main/java/org/apache/ambari/server/api/services/HostComponentService.java#L78) probably, because you can get the same result by the `/clusters/<clustername>/hosts/<hostname>/host_components/<hostcomponent>`. If I remove the check whether the <hostname> is provided or not I'll get back all the host responses where the components is installed to, but the json serializer messes up the response, e.g invalid response (response of 3 hosts): { "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/hosts/amb1.mycorp.kom/host_components/DATANODE", "HostRoles" : { "cluster_name" : "multi-node-hdfs-yarn", "component_name" : "DATANODE", "desired_admin_state" : "INSERVICE", "desired_stack_id" : "HDP-2.1", "desired_state" : "STARTED", "host_name" : "amb1.mycorp.kom", "maintenance_state" : "OFF", "service_name" : "HDFS", "stack_id" : "HDP-2.1", "stale_configs" : false, "state" : "INSTALL_FAILED", "actual_configs" : { } }, "host" : { "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/hosts/amb1.mycorp.kom" }, "processes" : [ ], "component" : [{ "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/services/HDFS/components/DATANODE", "ServiceComponentInfo" : { "cluster_name" : "multi-node-hdfs-yarn", "component_name" : "DATANODE", "service_name" : "HDFS" } }] }{ "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/hosts/amb2.mycorp.kom/host_components/DATANODE", "HostRoles" : { "cluster_name" : "multi-node-hdfs-yarn", "component_name" : "DATANODE", "desired_admin_state" : "INSERVICE", "desired_stack_id" : "HDP-2.1", "desired_state" : "STARTED", "host_name" : "amb2.mycorp.kom", "maintenance_state" : "OFF", "service_name" : "HDFS", "stack_id" : "HDP-2.1", "stale_configs" : false, "state" : "INSTALL_FAILED", "actual_configs" : { } }, "host" : { "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/hosts/amb2.mycorp.kom" }, "processes" : [ ], "component" : [{ "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/services/HDFS/components/DATANODE", "ServiceComponentInfo" : { "cluster_name" : "multi-node-hdfs-yarn", "component_name" : "DATANODE", "service_name" : "HDFS" } }] }{ "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/hosts/amb3.mycorp.kom/host_components/DATANODE", "HostRoles" : { "cluster_name" : "multi-node-hdfs-yarn", "component_name" : "DATANODE", "desired_admin_state" : "INSERVICE", "desired_stack_id" : "HDP-2.1", "desired_state" : "STARTED", "host_name" : "amb3.mycorp.kom", "maintenance_state" : "OFF", "service_name" : "HDFS", "stack_id" : "HDP-2.1", "stale_configs" : false, "state" : "INSTALL_FAILED", "actual_configs" : { } }, "host" : { "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/hosts/amb3.mycorp.kom" }, "processes" : [ ], "component" : [{ "href" : "http://172.24.0.21:8080/api/v1/clusters/multi-node-hdfs-yarn/services/HDFS/components/DATANODE", "ServiceComponentInfo" : { "cluster_name" : "multi-node-hdfs-yarn", "component_name" : "DATANODE", "service_name" : "HDFS" } }] } Is it a bug in the serializer or it os absolutely not recommended to remove that <hostname> check I mentioned above for some reason (in below) layers it is handled perfectly. Could any of you advise how should I proceed, please? [~u39kun][~swagle][~jspeidel] Thank you.	0	0	https://issues.apache.org/jira/browse/AMBARI-7192
19526	AMBARI-7190	Caught a runtime exception executing a query java.lang.NullPointerException at org.apache.ambari.server.view.ViewRegistry.getDefinition(ViewRegistry.java:261) at org.apache.ambari.server.controller.internal.ViewPermissionResourceProvider.getResources(ViewPermissionResourceProvider.java:139) at org.apache.ambari.server.controller.internal.ClusterControllerImpl.getResources(ClusterControllerImpl.java:134) at org.apache.ambari.server.api.query.QueryImpl.doQuery(QueryImpl.java:438) at org.apache.ambari.server.api.query.QueryImpl.queryForSubResources(QueryImpl.java:408) at org.apache.ambari.server.api.query.QueryImpl.queryForSubResources(QueryImpl.java:421) at org.apache.ambari.server.api.query.QueryImpl.queryForResources(QueryImpl.java:381) at org.apache.ambari.server.api.query.QueryImpl.execute(QueryImpl.java:196) at org.apache.ambari.server.api.handlers.ReadHandler.handleRequest(ReadHandler.java:68) at org.apache.ambari.server.api.services.BaseRequest.process(BaseRequest.java:135) at org.apache.ambari.server.api.services.BaseService.handleRequest(BaseService.java:103) at org.apache.ambari.server.api.services.BaseService.handleRequest(BaseService.java:72) at org.apache.ambari.server.api.services.ViewService.getViews(ViewService.java:75) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:57) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:606) at com.sun.jersey.spi.container.JavaMethodInvokerFactory$1.invoke(JavaMethodInvokerFactory.java:60) at com.sun.jersey.server.impl.model.method.dispatch.AbstractResourceMethodDispatchProvider$ResponseOutInvoker._dispatch(AbstractResourceMethodDispatchProvider.java:205) at com.sun.jersey.server.impl.model.method.dispatch.ResourceJavaMethodDispatcher.dispatch(ResourceJavaMethodDispatcher.java:75) at com.sun.jersey.server.impl.uri.rules.HttpMethodRule.accept(HttpMethodRule.java:288) at com.sun.jersey.server.impl.uri.rules.ResourceClassRule.accept(ResourceClassRule.java:108) at com.sun.jersey.server.impl.uri.rules.RightHandPathRule.accept(RightHandPathRule.java:147) at com.sun.jersey.server.impl.uri.rules.RootResourceClassesRule.accept(RootResourceClassesRule.java:84) at com.sun.jersey.server.impl.application.WebApplicationImpl._handleRequest(WebApplicationImpl.java:1469) at com.sun.jersey.server.impl.application.WebApplicationImpl._handleRequest(WebApplicationImpl.java:1400) at com.sun.jersey.server.impl.application.WebApplicationImpl.handleRequest(WebApplicationImpl.java:1349) at com.sun.jersey.server.impl.application.WebApplicationImpl.handleRequest(WebApplicationImpl.java:1339) at com.sun.jersey.spi.container.servlet.WebComponent.service(WebComponent.java:416) at com.sun.jersey.spi.container.servlet.ServletContainer.service(ServletContainer.java:537) at	1	0	https://issues.apache.org/jira/browse/AMBARI-7190
19527	AMBARI-7187	Stale patch [~keyki] can you check if the issue is still valid and resubmit the patch?	0	0	https://issues.apache.org/jira/browse/AMBARI-7187
19528	AMBARI-7175	Sorry I might not be aware of the complete existing stack inheritance implementation. From the current proposal, I'm guessing it will not be possible to have information distributed between the base and the derived files. If some one plans to inherit, will they have to copy all the information from the base file into the derived one as the entire file will get overwritten? Please correct me if I'm wrong.	1	0	https://issues.apache.org/jira/browse/AMBARI-7175
19529	AMBARI-7150	[~rlevas] That sounds perfect, is there a JIRA for this so I can keep an eye on it? thanks,	1	1	https://issues.apache.org/jira/browse/AMBARI-7150
19530	AMBARI-7121	[~sposetti] Hi Jeff, can you go ahead and commit this patch. Thanks in advance.	1	0	https://issues.apache.org/jira/browse/AMBARI-7121
19531	AMBARI-7072	[~keyki] [~matyix] Is this JIRA still and its sub-tasks still relevant?	0	0	https://issues.apache.org/jira/browse/AMBARI-7072

#	Requirement	Questions	Answered	Successfully Answered	URL
19532	AMBARI-7063	FAILURE: Integrated in Ambari-trunk-Commit #79 (See [https://builds.apache.org/job/Ambari-trunk-Commit/79/]) AMBARI-7063. Cannot query serviceConfigVersion with a config group name NULL.(mpapirkovskyy) (mpapyrkovskyy: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=9c281087d096a48695d4382a4e863c24a50b2699) * ambari-server/src/main/java/org/apache/ambari/server/state/ConfigImpl.java * ambari-server/src/test/resources/stacks/HDP/0.1/services/MAPREDUCE/metainfo.xml * ambari-server/src/test/java/org/apache/ambari/server/state/cluster/ClusterTest.java * ambari-server/src/main/java/org/apache/ambari/server/state/configgroup/ConfigGroupImpl.java * ambari-server/src/main/java/org/apache/ambari/server/orm/dao/ConfigGroupHostMappingDAO.java * ambari-server/src/main/java/org/apache/ambari/server/state/cluster/ClusterImpl.java	0	0	https://issues.apache.org/jira/browse/AMBARI-7063
19533	AMBARI-7063	FAILURE: Integrated in Ambari-trunk-Commit #80 (See [https://builds.apache.org/job/Ambari-trunk-Commit/80/]) AMBARI-7063. Cannot query serviceConfigVersion with a config group name NULL. Additional fix.(mpapirkovskyy) (mpapyrkovskyy: http://git-wip-us.apache.org/repos/asf?p=ambari.git&a=commit&h=99319de44248e59463560179db608e798ddc360d) * ambari-server/src/test/java/org/apache/ambari/server/state/host/HostTest.java	0	0	https://issues.apache.org/jira/browse/AMBARI-7063
19534	AMBARI-6926	[~afernandez] Is this JIRA fixed and part of 2.0.0 release?	1	0	https://issues.apache.org/jira/browse/AMBARI-6926
19535	AMBARI-6646	I'm interested in exploiting Rack awareness/rack ID in laying down Spectrum Scale(GPFS) file system. How can I get next level of details? Of course code can be downloaded but is there a detailed documentation of this feature. Thanks!	1	0	https://issues.apache.org/jira/browse/AMBARI-6646
19536	AMBARI-6554	Thanks [~jimbobhickville] Just a doubt from my side.When you create a cluster from blueprint do you already install all the rpms on all the nodes. Blueprints work well if you have deployed from image which contain preinstalled rpms for all services. Can you share some insights in to the setup which your using. Else the patch looks good. Thanks	1	0	https://issues.apache.org/jira/browse/AMBARI-6554
19537	AMBARI-6553	[~jimbobhickville] [~subin11] Is this JIRA still relevant? If not can you resolve it?	1	0	https://issues.apache.org/jira/browse/AMBARI-6553
19538	AMBARI-6527	What is the fix version for this?	0	0	https://issues.apache.org/jira/browse/AMBARI-6527
19539	AMBARI-6461	[~dballar1] Could you upload a patch?	1	1	https://issues.apache.org/jira/browse/AMBARI-6461
19540	AMBARI-6432	I have implemented IPA support in my fork at https://github.com/bolkedebruin/ambari. It implements it as a separate kerberos service (ie like MIT, AD). On a couple of things I need some feedback or help: * Where do I need to place (ie. which stack) the additional parameters "admin_keytab" and "group"? * It seems that not everything is picked up for the ui (like description or location) for these additional parameters. What am I doing wrong? Caveats: * Some idiosyncracies of IPA and Ambari are difficult to match. For example Ambari likes to create its key tabs programmatically, but this is incompatible with IPA as additional encryption types are generated only after a key tab is requested (and generated in IPA 4) on the server. This means that, in my implementation, the Ambari generated password is ignored for user principals or service principals that get a key tab generated. * Also, due to the fact IPA requires a password change on first use for user principals. Also if a key tab is generated for this user (which is a bit weird and I will need to ask to the IPA guys). Therefore, the user administrator (the principal in admin_keytab) requires additional permissions to be able to write to krbPasswordExpiration (ipa permission-add “Set Password Expiry” –permissions=write –type=user –attrs=krbPasswordExpiration, ipa privilege-add-permission “User Administrators” –permission=”Set Password Expiry”). Question: is this the right way to go or should an "expect" utility be implemented so that the password can be set? * Tests still need to be written.	1	1	https://issues.apache.org/jira/browse/AMBARI-6432
19541	AMBARI-6432	Unfortunately I'm neither an Ambari nor an IPA expert . Can someone on the security/identity side of things reveiew this with thoughts. Possibly some red hatters might be interested also, [~subin11] [~erinaboyd] [~screeley] , since it is FreeIPA after all :)	0	0	https://issues.apache.org/jira/browse/AMBARI-6432
19542	AMBARI-6432	[~u39kun], [~jaimin]... Can you take a look at this patch and review the UI updates? Or suggest someone that could do it?	1	1	https://issues.apache.org/jira/browse/AMBARI-6432
19543	AMBARI-6432	Can someone help out to get builds running? I am not sure what is off and I'd like to get this merged.	1	1	https://issues.apache.org/jira/browse/AMBARI-6432
19544	AMBARI-6432	[~bolke]... I think we wanted to ask you to add a _feature flag_ to the UI portion of your code. [~u39kun] should have the details on what this means. After that and if the review board has enough +1's I will be happy to commit the patch. BTW: I just too a look at the review and it appear to be closed as submitted. Was the patch already committed?	1	1	https://issues.apache.org/jira/browse/AMBARI-6432
19545	AMBARI-6432	[~u39kun] Thanks that will help. In the meantime I also hit a bug on which I need some guidance how to fix it. FreeIPA does not support uppercase user principal names. If the cluster name is in uppercase a test identity will be generated with "myname@REALM" . For the tests to pass I need it to be "myname@REALM". What would be the best way to fix this? 1) Generate the test identity in lowercase (where is this generated? I could not find it yet) 2) Ask the user at step 1 to verify the cluster name is in lower case 3) adjust service_check.py to convert to lowercase if a user principal is encountered 4) use auth_to_local rules and apply these to the test identity In my opinion option 1 seems the best option. It would affect other Kerberos providers as well, but as it is only the test identity I would say it would not matter. Please advice	1	1	https://issues.apache.org/jira/browse/AMBARI-6432
19546	AMBARI-6432	[~rlevas] [~u39kun] v6 patch should address all issues mentioned. I would like to consider it Final if it passes the tests (can someone kicks those off, it does not seem to run on uploading a patch?). * FreeIPA 3 should work * Marked experimental by 'enableIpa' at the experimental page * Tested on 10+2 node system * toLower Function implemented in VariableReplacementHelper and applied where needed Note: * There seems to be a bug in how Ambari handles adding and removing principals: it gets called multiple times for these operations. By turning on debugging output for the Kerberos classes (and use IPA as the other classes dont log enough) you can see this. I haven't been able to look at this.	1	1	https://issues.apache.org/jira/browse/AMBARI-6432
19547	AMBARI-6432	How is the testing going? I almost off to holidays and it would be really nice to have this in before that ;-)	1	1	https://issues.apache.org/jira/browse/AMBARI-6432
19548	AMBARI-6416	I know there is also a python client that [~subin11] is working on. do the groovy and python clients overlap functionality?	1	0	https://issues.apache.org/jira/browse/AMBARI-6416
19549	AMBARI-6363	[~subin11] Is this JIRA still valid.	0	0	https://issues.apache.org/jira/browse/AMBARI-6363

#	Requirement	Questions	Answered	Successfully Answered	URL
19550	AMBARI-6354	Nate, thanks for sharing the design of alert monitoring in Ambari. Here are some review comments and questions about the design, -- For various reasons, It is very important for Ambari to support Alert monitoring through an inbuilt mechanism via ambari agents and not solely relying on the external monitoring services like Nagios. So thanks for initiating the efforts! -- Although having a pluggable interface for Ambari to work with many existing monitoring systems like Nagios/Zabbix etc. would be really important for some users who has already invested in those monitoring systems for their other infrastructure components. Looking at current design, the approach is to have Ambari do the alert monitoring via its agents and then dispatch these alerts to other monitoring systems. This may not be ideal for some users(see some examples below). ** Users familiar with existing Monitoring systems may want to configure all the alerts including Hadoop services, all in one place. ** Advance alerts management features such as alert aggregation, correlation, mitigation etc. can be done uniformly across various alerts with their existing monitoring system as opposed to partly in Ambari. ** Etc. -- So I was wondering, if we make Ambari server internally communicate with existing monitoring systems with standard pluggable interface? This design point is orthogonal to current Ambari alert monitoring design. Ambari's alert monitoring would be one of the implementations of the pluggable monitoring interface and serve as out-of-box default implementation for Ambari managed services. Although advantage is that, it would also facilitate other users to add the implementation for their existing monitoring system. I can propose the initial draft of the interface for alert monitoring and see how we can get the existing Ambari alert monitoring to plug in under that interface.. Let me know what you think?	0	0	https://issues.apache.org/jira/browse/AMBARI-6354
19551	AMBARI-6354	I think by design it would be good to have Ambari collection of alerts be default but optional i.e. if I choose to go with Nagios, I may not want Ambari's default alert monitoring. i.e. Ambari's inbuilt alert monitoring can be treated as optional pluggable service. Is this achievable after refactor work? Again current refactor work to remove Nagios dependency is good but at the same time I wanted to see how can we make it pluggable design and wanted to help out along those lines. I am reviewing the design and code and will propose some specific changes w/ minimal disruption to exiting work as possible..	0	0	https://issues.apache.org/jira/browse/AMBARI-6354
19552	AMBARI-6354	Thanks for the quick reply. I see. Is there a JIRA for web notifications or it is not planned at all? Our use case is to scale the cluster based on different type of metrics.	1	1	https://issues.apache.org/jira/browse/AMBARI-6354
19553	AMBARI-6354	I'm not clear on exactly what you're asking about. Can you explain a bit more what you mean when you say "web notifications"?	1	1	https://issues.apache.org/jira/browse/AMBARI-6354
19554	AMBARI-6275	[~jspeidel] Thanks for posting the proposal. Are these new API endpoints meant to be idempotent so that we can easily recover from failures/partial failures? That's one issue with our current blueprint cluster creation API (you can't simply retry by issuing the same call so it's hard to recover when failures occur).	1	1	https://issues.apache.org/jira/browse/AMBARI-6275
19555	AMBARI-6275	[~jspeidel] Ok, thanks for clarifying. > Failure tolerance and retry behavior are really important to any long running asynchronous request so any solution would likely augment the current lower level request model. Do we have a JIRA for this already? If not, can you open one?	1	1	https://issues.apache.org/jira/browse/AMBARI-6275
19556	AMBARI-6275	I like John's approach. A REST end point should ideally do only one thing. So doing the post-add tasks should not be tied to the add-hosts REST end point. It should be the UI's job to make it easier for user. The REST APIs are targeted mostly for machines or developers. Some questions: 1. In case of reference node approach, what would be the response if there are master roles and slave roles on the node. E.g. Zookeeper nodes (master) and Datanodes (slaves)? 2. Will agents have to be pre-installed on the newly added nodes? 3. Should this API be aware of cluster security? How does Ambari add security to newly added nodes today? Are there separate APIs for doing that? BTW, is there a plan for removing nodes too (may be in a different ticket)? Thanks	1	0	https://issues.apache.org/jira/browse/AMBARI-6275
19557	AMBARI-6275	[~jspeidel] Questions: 1. Is it possible to use the inline host_group format to add hosts to a cluster that was not created via a blueprint? 2. Is it possible to have this API perform in a transactional manner? Say I'm adding a host and 5 host components are to be added. If 2 host components fail to install, then I want to roll back the whole thing. Or say I'm adding 5 hosts and 2 hosts failed to install, etc. 3. Once you add a host by specifying a host group, is the meta info about the added host belonging to the specified host group made available via the API?	1	1	https://issues.apache.org/jira/browse/AMBARI-6275
19558	AMBARI-6275	[~jspeidel] Based on issues in EPIC looks like only documentation work is remaining in this EPIC. Can you address the documentation work and resolve the EPIC since this is part of 2.0.0 release?	1	0	https://issues.apache.org/jira/browse/AMBARI-6275
19559	AMBARI-6148	Hi [~erinaboyd], I wonder if you are still working on this. I'd like to add a set of service definition for Spark so that spark can be installed through Ambari. Please let me know if you can give me this ticket. I have started investigating a little bit and found a challenge; Which stack (PHD, BIGTOP and PHD) to add the spark definition first? HDP went ahead on trunk with versioned rpm so it should be a preferable option but their yum repo does not provide spark. BIGTOP does but is not versioned rpm yet. HDP {code} bash-4.1# repoquery --repofrompath=AA,http://s3.amazonaws.com/dev.hortonworks.com/HDP/centos6/2.x/BUILDS/2.2.1.0-2048/ -qa | grep spark python-sparklines-0:0.9-2.el6.noarch # spark-core is not included {code} BIGTOP http://bigtop01.cloudera.org:8080/job/Bigtop-trunk-Repository/label=centos6/lastSuccessfulBuild/artifact/repo/RPMS/	0	0	https://issues.apache.org/jira/browse/AMBARI-6148
19560	AMBARI-6125	new: addHost - adds a registered host to the cluster removeHost - gracefully decommission and remove the host from the cluster (synchronous) waitForRequestsToFinish - waits for any given request to finish, used to synchronize rest calls decommissionDataNode - decommissions the datanode on a given host decommissionNodeManager - decommissions the nodemanager on a given host decommission - can decommission any given components which supports decommission deleteHostComponents - deletes the given components from a host deleteHost - deletes the host from the cluster installComponentsToHost - can install components from either blueprint or from the specified component list restartServiceComponents - restarts the given services components on all hosts startComponentsOnHost - starts the provided components on a host stopComponentsOnHost - stops the provided components on a host validateBlueprint - can validate the blueprint for recommendAssignments modifyConfiguration - can modify any services configuration (the whole config must be provided) getHostNamesByState - retrieves the hosts by state getPublicHostNames - returns the public host names which the provided host components are installed to getInternalHostNames - returns the internal host names which the provided host components are installed to getClusterHosts - returns only the hosts which are part of the cluster resolveInternalHostName - can resolve a host's internal name to the public one modified: recommendAssignments - host group and host recommendation changed: - at least host group number of host is required - if there are more hosts than host groups in the blueprint all host group get a single node and the leftover hosts are distributed among all the slave_X host groups getInstallProgress - renamed to getRequestProgress startAllServices - return type changed stopAllServices - return type changed blueprints: HDP single/multi node blueprints added to defaults dependencies: commons-io added	1	0	https://issues.apache.org/jira/browse/AMBARI-6125
19561	AMBARI-6125	Hi [~u39kun], Could you please take a look at the patch? Thank you.	1	0	https://issues.apache.org/jira/browse/AMBARI-6125

#	Requirement	Questions	Answered	Successfully Answered	URL
19562	AMBARI-6086	[~u39kun] in AMBARI-4447 [~jaimin] enabled the ability for globals to be defined for multiple services (as we did with the hadoop_heapsize for HDFS & GlusterFS) So, why can we not do the same with the namenode_heapsize? If it's defined in the global props with the service being GLUSTERFS, I don't understand how it breaks HDFS. We are having multiple issues with various services not starting up (Nagios, HBASE, because of these global paratemeters not being defined). The error looking like this (configuration parameters vary) File "/var/lib/ambari-agent/cache/stacks/HDP/2.0.6/services/NAGIOS/package/scripts/params.py", line 137, in <module> nn_ha_host_port_map[config['clusterHostInfo']['namenode_host'][0]] = namenode_metadata_port File "/usr/lib/python2.6/site-packages/resource_management/libraries/script/config_dictionary.py", line 75, in __getattr__ raise Fail("Configuration parameter '"+self.name+"' was not found in configurations dictionary!") Fail: Configuration parameter 'namenode_host' was not found in configurations dictionary! Please advise...	1	1	https://issues.apache.org/jira/browse/AMBARI-6086
19563	AMBARI-6086	Hi [~jaimin] & [~u39kun]The property we are looking at is namenode_host which I assume will be mutually exclusive as we would only use this prop in gluster when HDFS isn't installed. Can we add this back into the config then as well for trunk and 1.6.1? Also, we are struggling with HBASE not starting because of slave_hosts is not in the config dictionary. (Same error as above). However, this value doesn't appear to be set as a configuration value. I see it checking it on serveral services, Nagios, HDFS, Ganglia, but perhaps it's set in the params.py on the before-INSTALL script in the hooks directory?? If so, how can we default this to something else as this is done before install you don't have a list of what you are installing, therefore, we can't do a conditional for isGluster....Thoughts?	0	0	https://issues.apache.org/jira/browse/AMBARI-6086
19564	AMBARI-5980	[~xiwang0309] 1. Please convert the code from making a synchronous AJAX call to asynchronous (we should not be using "async: false"). 2. self.set('javaHome', property.split("")[3]); <- This code is mysterious and also possibly error prone. Why 3? What's the expected API response? What if after splitting, the length of the array is less than 4? We should document the actual expected response from the server (ideally, we should put the mock file as we do for many other AJAX calls)	1	1	https://issues.apache.org/jira/browse/AMBARI-5980
19565	AMBARI-5980	[~xiwang0309] It looks like AMBARI-5998 changed the format for structured_out. This is now a stringified JSON, so we don't need to do split magic anymore. We can simply use the JSON parser. Can you go ahead and make the change?	1	1	https://issues.apache.org/jira/browse/AMBARI-5980
19566	AMBARI-5811	[~rlevas] [~vbrodetskyi] Is this JIRA still valid and take necessary steps accordingly?	1	0	https://issues.apache.org/jira/browse/AMBARI-5811
19567	AMBARI-5707	+1 for this ,fantastic proposal .Something I had always wanted. Considering the conflicting ganglia/nagios versions on existing customer setups.This would help a lot. Will this daemon be coded in python? I am suggesting something in C/scala/erlang (C best) to capture data and dump it out as json. And agent itself can also have a webui(disabled by default) which can show the metrics. But OpenTSDB is already has python tcollector and a LGPL .So would we be inspired from it ,and rewrite it ? One last question I have is regarding whether you are considering to get the visualization as good/similar as in ganglia Ganglia uses RRDTool(i guess) which is again GPL. So does the plan also include improving the server side to render the data in a different way.	0	0	https://issues.apache.org/jira/browse/AMBARI-5707
19568	AMBARI-5707	It would be a good result but --- at first thought: Im a little wary of this: If Ambari re-builds alot of components in the stack, it might wind up being to hard to maintain as a open source project. One of the better ways maybe is reach out to ganglia community about the problems we face. Is a combination of custom JMX hooks + Ganglia not enough to accomplish the above bullets.	0	0	https://issues.apache.org/jira/browse/AMBARI-5707
19569	AMBARI-5613	[~simon.zhangsm@hotmail.com] Is this JIRA still valid?	0	0	https://issues.apache.org/jira/browse/AMBARI-5613
19570	AMBARI-5482	Thanks Yusaku, How should we proceed further? We already have some awesome new features and improvements. Should we create a new JIRA issue for each and submit patches or just create one issue and put everything there?	1	0	https://issues.apache.org/jira/browse/AMBARI-5482
19571	AMBARI-5465	[~swagle] Is this JIRA still valid?	0	0	https://issues.apache.org/jira/browse/AMBARI-5465
19572	AMBARI-5463	Hi [~simon.zhangsm@hotmail.com]. Can you submit a patch for this? Some info here on the review process: https://cwiki.apache.org/confluence/display/AMBARI/How+to+Contribute	1	0	https://issues.apache.org/jira/browse/AMBARI-5463
19573	AMBARI-5445	With some input from [~mahadev] how about a more generic solution: * Introduce the concept of a "monitoring" service ** Nagios and Ganglia are two such services * A monitoring service, by definition, needs to be restarted when host components are added/removed ** Restart is indicated by setting a flag (if we can user stale_config = true that will be great) on all service host components * The indicator itself is reset when a host component starts [~dsen], can you check if the above steps fit and provide a detailed proposal	1	0	https://issues.apache.org/jira/browse/AMBARI-5445
19574	AMBARI-5194	HostInfo.py path with mvn clean test results: Results : Failed tests: testNoNagiosServerCompoonent(org.apache.ambari.server.controller.nagios.NagiosPropertyProviderTest): Expected no alerts Tests in error: testDeleteUsers(org.apache.ambari.server.controller.AmbariManagementControllerTest): Could not remove user user1. System should have at least one user with administrator role. Tests run: 1404, Failures: 1, Errors: 1, Skipped: 7 [INFO] -- [INFO] Reactor Summary: [INFO] [INFO] Ambari Main ..... SUCCESS [0.063s] [INFO] Apache Ambari Project POM ..... SUCCESS [0.024s] [INFO] Ambari Web ..... SUCCESS [5.396s] [INFO] Ambari Views ..... SUCCESS [1.291s] [INFO] Ambari Server ..... FAILURE [8:43.111s] [INFO] Ambari Agent ..... SKIPPED [INFO] Ambari Client ..... SKIPPED [INFO] -- [INFO] BUILD FAILURE [INFO] -- [INFO] Total time: 8:50.247s [INFO] Finished at: Tue Mar 25 14:48:26 GMT-05:00 2014 [INFO] Final Memory: 34M/812M [INFO] --	1	1	https://issues.apache.org/jira/browse/AMBARI-5194
19575	AMBARI-5194	[~screeley] thats part of agent unit tests. What I meant was would you be able to add more unit tests to TestHostInfo.py? You can run the ambari-agent unit tests by doing: cd ambari-agent mvn clean test. That should run all he unit testso n the ambari agent side.	1	1	https://issues.apache.org/jira/browse/AMBARI-5194

#	Requirement	Questions	Answered	Successfully Answered	URL
19576	AMBARI-5194	[~mahadev] - thanks for the tips... So, here is the successful unit test: ============= specific line: ============= test_checkIptables (TestHostInfo.TestHostInfo) ... ok ============= All HostInfo tests ============= test_analyze_yum_output (TestHostInfo.TestHostInfo) ... ok test_analyze_yum_output_err (TestHostInfo.TestHostInfo) ... ok test_analyze_zypper_out (TestHostInfo.TestHostInfo) ... ok test_checkFolders (TestHostInfo.TestHostInfo) ... ok test_checkIptables (TestHostInfo.TestHostInfo) ... ok test_checkLiveServices (TestHostInfo.TestHostInfo) ... ok test_checkUsers (TestHostInfo.TestHostInfo) ... ok test_dirType (TestHostInfo.TestHostInfo) ... ok test_etcAlternativesConf (TestHostInfo.TestHostInfo) ... ok test_getReposToRemove (TestHostInfo.TestHostInfo) ... ok test_hadoopVarLogCount (TestHostInfo.TestHostInfo) ... ok test_hadoopVarRunCount (TestHostInfo.TestHostInfo) ... ok test_hostinfo_register (TestHostInfo.TestHostInfo) ... ok test_hostinfo_register_suse (TestHostInfo.TestHostInfo) ... ok test_javaProcs (TestHostInfo.TestHostInfo) ... ok test_osdiskAvailableSpace (TestHostInfo.TestHostInfo) ... ok test_perform_package_analysis (TestHostInfo.TestHostInfo) ... ok so it passed, but I had to modify the TestHostInfo.py for that particular test, as it was still expecting the values from the old method. So what is the process for that, do I now need a separate JIRA so I can update the TestHostInfo.py as well? Below is what I changed (old routine used faulty returncode, new routine uses boolean values) result = hostInfo.checkIptables() self.assertTrue(result == True) result = hostInfo.checkIptables() self.assertFalse(result == False)	1	1	https://issues.apache.org/jira/browse/AMBARI-5194
19577	AMBARI-5181	It appears at some point in time there was a generic HCFS integration that got ported back to GlusterFS specific one (AMBARI-3167, AMBARI-2721). Do we know why this got moved (sorry JIRA comments are not really helpful) It looks like the GlusterFS integration is not just handled only at the stack level (inheritence model) but also touches many core components (ambari-server, ambari-web, ambari-client modules etc.,). It appears there is no clean way of extending the framework to add new services (HCFS). If we need to introduce new HCFS support to Ambari, should we need to follow same approach as adopted for GlusterFS (extending stack and also modifying core Ambari components - read hardcoded) or is there a generic way of introducing new HCFS stack support. I am planning to post my findings and propose new JIRA task but wanted to find if there is something on the radar already.	1	1	https://issues.apache.org/jira/browse/AMBARI-5181
19578	AMBARI-5181	[~vijikarthi] You are right, hcfs JIRAs may be related to various components, eg. AMBARI-8658 is clearly a stack related, while there are issues related to multiple components as well: eg. AMBARI-8468. On the other hand, there are issues without component (eg. AMBARI-8505) as it's not immediatelly clear which is the right one (from the reporter perspective). And since I needed some easy way to track hcfs related issues, I created [hcfs label|https://issues.apache.org/jira/browse/AMBARI-8658?jql=labels%20%3D%20hcfs] some time ago. What do you think about this approach? The downside is that me and [~dahorak] are the only ones currently using it so far.	1	0	https://issues.apache.org/jira/browse/AMBARI-5181
19579	AMBARI-5116	[~xiwang0309], can we take care of the CRIT and WARN scenarios with time 0?	1	1	https://issues.apache.org/jira/browse/AMBARI-5116
19580	AMBARI-5116	[~xiwang0309], we need to externalize all static strings, including "for ", "Last Checked:", etc. Can you go ahead and do that?	1	1	https://issues.apache.org/jira/browse/AMBARI-5116
19581	AMBARI-5044	Thans [~xiwang0309]. Can we change it so that "OK" button says "Confirm Add" and "Confirm Delete" respectively?	1	1	https://issues.apache.org/jira/browse/AMBARI-5044
19582	AMBARI-4908	Thank you for your answer Jeff. I cant wait to see this new version 1.5.0 of Ambari released ;-). Any idea when it will be ? Best regards. Gwenael Le Barzic	1	1	https://issues.apache.org/jira/browse/AMBARI-4908
19583	AMBARI-4881	Unrelated - if we have a choice can we rename STORM_REST_API to something else such as STORM_REST_WEB_SERVICE. {noformat} "properties": { "default_port": "8745" # Storm port is always hardcoded, so we will just use the default value "relevant_config": "storm-site" # Just example, not needed for Storm "port_property": "storm.port" # Just example, not needed for Storm } {noformat} Can the above be consumed by AbstractProviderModule (modify JMXHostProvider into Generic HostProvider that provided host, port, protocol) and we can slowly move port related details for the jmx providers to the metrics.json file? APM can then make it available to JMXPropertyProvider as well as RESTPropertyProvider. Instead of "relevant_config" and "port_property" I would suggest "port_config_type" and "port_property_name". I am thinking the eventually we can have "some_other_default", "some_other_config_type" and "some_other_property_name". Little confused by {noformat} "metrics": { "metrics/api/cluster/summary/nimbus.uptime": { "metric": "nimbus.uptime", "pointInTime": true, "temporal": true } {noformat} Should we try to maintain similarity with the jmx/ganglia metrics? {noformat} "metrics": { "metrics/nimbus/nimbus.uptime": { <== URL part from Ambari Resource "metric": "api/cluster/summary/nimbus.uptime", <== selector/filter for the JSON received from Storm "pointInTime": true, "temporal": true } {noformat}	1	0	https://issues.apache.org/jira/browse/AMBARI-4881
19584	AMBARI-4798	In service_configs.js: {noformat} case 'yarn.ahs.webapp.address': var hsHost = masterComponentHostsInDB.findProperty('component', 'HISTORYSERVER').hostName; this.setDefaultValue("(0.0.0.0)(?=:)", hsHost); break; case 'yarn.timeline-service.webapp.https.address': var hsHost = masterComponentHostsInDB.findProperty('component', 'HISTORYSERVER').hostName; this.setDefaultValue("(0.0.0.0)(?=:)", hsHost); break; {noformat} 1. Should yarn.ahs.webapp.address be renamed to yarn.timeline-service.webapp.address? 2. Should HISTORYSERVER should be APP_TIMELINE_SERVER?	1	1	https://issues.apache.org/jira/browse/AMBARI-4798
19585	AMBARI-4786	Hi [~jspeidel], I can see that the output consist of "name" : "AMBARI_SERVER" Is there a special meaning to that? I wonder if it was because we exported the blueprint via a HTTP URL fired to the Ambari Server Host? Thanks!	0	0	https://issues.apache.org/jira/browse/AMBARI-4786
19586	AMBARI-4782	This change will impact the API spec so we need to follow the needed procedure. This is an additive change anyway. Some clarifications: * Will we see "state", "desired_state", "current_state" as three items in the response? * Query will then support filtering based on any of the three.	1	1	https://issues.apache.org/jira/browse/AMBARI-4782

#	Requirement	Questions	Answered	Successfully Answered	URL
19587	AMBARI-4782	[~mahadev] / [~sumitmohanty] / [~u39kun] I think it would be good to change the desired solution. Currently we have an inconsistency in our API: "state" parameter has a different meaning in request and in responce. Talked to UI folks regarding that, they agree that it is an issue, but are not sure how much the proposed fix will affect the UI. h2. What is wrong: HostComponentResourceProvider has two properties defined: "state" and "desired_state". When we issue a read request like http://vm-3.vm:8080/api/v1/clusters/cc/host_components?HostRoles/state=INSTALL_FAILED , the "state" parameter is mapped to "desired_state". At the same time, "state" property meaning at responce json is a live service state (so-called current state). {code} { href: "http://vm-3.vm:8080/api/v1/clusters/cc/hosts/vm-3.vm/host_components/HDFS_CLIENT", HostRoles: { cluster_name: "cc", component_name: "HDFS_CLIENT", host_name: "vm-3.vm", state: "INSTALLED" }, host: { href: "http://vm-3.vm:8080/api/v1/clusters/cc/hosts/vm-3.vm" } }, {code} This inconsistency is not obvious and may lead to various issues like state=INSTALL_FAILED issue mentioned above. So I believe that the right thing would be to remap "state" parameter to service live state both in requests and responses (and adjust appropriately web-UI and any third-party tools if needed). I assume that most party tools will not even notice a difference because they use "state" parameter as "live state". "current_state" property should not be added as it will only introduce more confusion.	1	0	https://issues.apache.org/jira/browse/AMBARI-4782
19588	AMBARI-4782	Published a new patch. The updated patch seems to implement all requested things. If the approach is appropriate, I'll add unit test coverage. Also, I'm considering fixing also Service resource (we have the same issue here, state is mapped to desired_state), but I'm not sure if it is needed. [~sumitmohanty], can you please your feedback regarding modifying Service resource as well?	1	0	https://issues.apache.org/jira/browse/AMBARI-4782
19589	AMBARI-4782	[~dmitriusan], I'm looking at another issue related to the host state, https://hortonworks.jira.com/browse/BUG-20676 Could these 2 Jiras be related?	0	0	https://issues.apache.org/jira/browse/AMBARI-4782
19590	AMBARI-4782	Is this JIRA still valid?	0	0	https://issues.apache.org/jira/browse/AMBARI-4782
19591	AMBARI-4770	[~xiwang0309], the patch is installing TEZ_CLIENT on NAGIOS_SERVER. Shouldn't that be HIVE_CLIENT?	1	1	https://issues.apache.org/jira/browse/AMBARI-4770
19592	AMBARI-4504	[~mahadev]/[~sumitmohanty], can you please take a look on implementation proposal?	0	0	https://issues.apache.org/jira/browse/AMBARI-4504
19593	AMBARI-4481	[~mahadev]/[~swagle], can you please take a look on proposal?	1	1	https://issues.apache.org/jira/browse/AMBARI-4481
19594	AMBARI-4481	[~swagle], I've reverted your change of Jetty servlet handler for resources/* introduced at revision 60f3e6e. Now servlet handler at /org/apache/ambari/server/controller/AmbariServer.java:294 looks like: {code} ServletHolder resources = new ServletHolder(ServletContainer.class); resources.setInitParameter("com.sun.jersey.config.property.resourceConfigClass", "com.sun.jersey.api.core.PackagesResourceConfig"); resources.setInitParameter("com.sun.jersey.config.property.packages", "org.apache.ambari.server.resources.api.rest;"); root.addServlet(resources, "/resources/*"); resources.setInitOrder(6); {code} Is it ok? Otherwise, resources requests any url starting with http://vm-0.vm:8080/resources/stacks/ return JSON responces (the same responces as API at "/api/v1/stacks" returns)	0	0	https://issues.apache.org/jira/browse/AMBARI-4481
19595	AMBARI-4358	[~dmitriusan] Some questions/comments: - When would the scripts that create symlinks run ? - How would upgrade scenario be handled with existing symlinks? - Is there an alternative to creating symlinks which could be fragile: meaning an accidental symlink delete leads to broken agent? - Is it possible to consider an approach where scripts, hooks and templates are written for a service say HDFS without tie up to a specific stack version and then in the definition of a stack say 2.1.1, we declaratively pin point with script to use for that version. This way it becomes easy to develop a inheritance chain of python classes and unit test changes as stack is added.	1	0	https://issues.apache.org/jira/browse/AMBARI-4358
19596	AMBARI-4324	h1. Implementation proposal: 1. Add new command type CANCEL_COMMAND to agent-server protocol. CANCEL_COMMAND contains identifier (task_id + stage_id) of an exact command for cancellation. 2. At the server side, commands of this type are issued when tasks are considered timed out. I'm going to do that here: org.apache.ambari.server.actionmanager.ActionScheduler#processInProgressStage. 3. At the agent side, CANCEL_COMMANDs are executed inside Controller.py right after arrival (they are not put into ActionQueue). If command mentioned by the CANCEL_COMMAND is not present in ActionQueue (it is already in progress or completed), CANCEL_COMMAND is silently ignored. 4. Also, agent clears entire action queue when it can not continue exchanging heartbeats with the server (disconnect or registration requested). I'm going to add an appropriate logic to src.main.python.ambari_agent.Controller.Controller#registerAndHeartbeat. The motivation is to make recovery from network/server fail more reliable and fast (agent will have an empty ActionQueue and can start executing new EXECUTION_COMMANDS and STATUS_COMMANDS right after registration). 5. In both cases described above (executing a single CANCEL_COMMAND or clearing entire ActionQueue) EXECUTION_COMMANDS are considered transactional-like. I mean that EXECUTION_COMMANDs that are already IN_PROGRESS are never interrupted. Thus we decrease chanses of leaving system in misconfigured/unpredictable state. Also, I'm going to fix a bug at org.apache.ambari.server.actionmanager.ActionScheduler#processInProgressStage . Here, we pass stage timeout instead of task timeout as a parameter of org.apache.ambari.server.actionmanager.ActionScheduler#timeOutActionNeeded . After bugfix, task timeout + some small time will be passed as a parameter value. Additional smal time (10-30 seconds) is needed to avoid sending CANCEL_COMMAND without absolute necessary (task will timeout at agent automatically in most cases). This implementation should also solve another related jira AMBARI-4324 [~mahadev] and/or [~sumitmohanty], can you please take a look on this proposal?	1	0	https://issues.apache.org/jira/browse/AMBARI-4324
19597	AMBARI-4324	[~dmitriusan], this is a good feature. We will also need to support cancel for request or request/task through API. I am working on a proposal for that. Can you ensure there are hooks to add Cancel commands when it is initiated through a API call. Typically, agents do not have a long queue of pending tasks. Just wondering if Controller.py should process it cancel synchronously. Perhaps it is OK. We should have the ability to interrupt INPROGRESS tasks as well. I am thinking run-away tasks or misconfigured timeouts. What would it take to have this ability? What is the other JIRA ID? The current link is to this JIRA. {quote} This implementation should also solve another related jira AMBARI-4324 {quote}	1	1	https://issues.apache.org/jira/browse/AMBARI-4324

#	Requirement	Questions	Answered	Successfully Answered	URL
19598	AMBARI-4324	[~sumitmohanty], {quote} We will also need to support cancel for request or request/task through API. I am working on a proposal for that. Can you ensure there are hooks to add Cancel commands when it is initiated through a API call. {quote} I'll add a relevant hook addCancelCommandAction() to /org/apache/ambari/server/controller/AmbariCustomCommandExecutionHelper.java {quote} Typically, agents do not have a long queue of pending tasks. Just wondering if Controller.py should process it cancel synchronously. Perhaps it is OK. {quote} during cluster install, queue may contain 5-10 pending tasks (all tasks for stage) + 20 status commands, so it takes a long time to execute. Also, if we process "cancel" request synchronously, it will be placed to queue after task we are trying to cancel {quote} We should have the ability to interrupt INPROGRESS tasks as well. I am thinking run-away tasks or misconfigured timeouts. What would it take to have this ability? {quote} Cancelling command in progress is possible (we may invoke kill-on-timeout callback method manually, effectively killing all subprocesses), but may occasionally leave system configuration in a broken state. {quote} What is the other JIRA ID? The current link is to this JIRA. This implementation should also solve another related jira AMBARI-4324 {quote} Implementation proposal was intended for AMBARI-4323, I've posted it to current (related) jira by mistake.	1	0	https://issues.apache.org/jira/browse/AMBARI-4324
19599	AMBARI-4323	Looks good. Few points to think about: * Agree that we do not need a new status, FAILED is OK. Somehow using STDOUT/STDERR does not seem right but it is better than adding yet another column. * I assume when a CANCEL is issued for a task the state of the task in the server side is IN_PROGRESS. Only, when the agent returns the command status that the server will update the state to FAILED. * Are we also implementing canceling of IN_PROGRESS commands? Just wanted to confirm. * _clears entire action queue on every registration_: When expected and received heartbeat ids are not the same the server asks agents to re-register. I do not know if in this case server assumes that in progress commands are aborted. If server does not assumes that commands are aborted then agent should not clear the action queue. Can you verify the behavior here? * _we pass a stage timeout instead of task timeout as a parameter_: That part of the ActionScheduler always seems more complex than needed. Anyway, I have a question. Today stage timeout is the largest "sum of task timeouts on a single host". One a one node, during install the last_attempt_time is basically the same for all install tasks. So if we pass "task timeout + some small time" while checking timeOutActionNeeded() for the later tasks - ones in the end of the queue in Agent - will we not have a problem?	1	0	https://issues.apache.org/jira/browse/AMBARI-4323
19600	AMBARI-4323	Is the purpose of the cancel command only to remove from the Queue? What about those items that are IN_PROGRESS - how do you actually stop them? For example, pretend you issue a START command, and the backing OS command (could be, say, some hdfs command) blocks. Will you issue a kill -9 on the command? Or are we just removing the "monitoring" for that request/task combination?	1	0	https://issues.apache.org/jira/browse/AMBARI-4323
19601	AMBARI-4323	[~ncole@hortonworks.com], yes, we perform the same action as for commands that timed out: we kill the process and all child subprocesses. [~sumitmohanty], {quote} I assume when a CANCEL is issued for a task the state of the task in the server side is IN_PROGRESS. Only, when the agent returns the command status that the server will update the state to FAILED. {quote} CANCEL command may be issued for any task id (both IN_PROGRESS or just scheduled) {quote} Are we also implementing canceling of IN_PROGRESS commands? Just wanted to confirm. {quote} yes, we may cancel even in_progress commands. If command has been already successfully finished by the moment when agent receives relevant CANCEL command, we display command as COMPLETED (cancellation has not effect) {quote} clears entire action queue on every registration: When expected and received heartbeat ids are not the same the server asks agents to re-register. I do not know if in this case server assumes that in progress commands are aborted. If server does not assumes that commands are aborted then agent should not clear the action queue. Can you verify the behavior here? {quote} According to code /org/apache/ambari/server/agent/HeartbeatMonitor.java:144 , we clear queue server-side queue on heartbeat lost {code} //Purge action queue actionQueue.dequeueAll(host); //notify action manager actionManager.handleLostHost(host); {code} But I'll recheck that when implementing the jira {quote} we pass a stage timeout instead of task timeout as a parameter: That part of the ActionScheduler always seems more complex than needed. Anyway, I have a question. Today stage timeout is the largest "sum of task timeouts on a single host". One a one node, during install the last_attempt_time is basically the same for all install tasks. So if we pass "task timeout + some small time" while checking timeOutActionNeeded() for the later tasks - ones in the end of the queue in Agent - will we not have a problem? {quote} Good point. Actually, we currently are passing stage timeout value, but it is named taskTimeout. So the fix is just to rename a variable.	0	0	https://issues.apache.org/jira/browse/AMBARI-4323
19602	AMBARI-4320	Discussed this question with [~dsen]. Actually, we already have support of issuing CONFIGURE command via api. Exact call looks like {code} curl -X POST 'http://localhost:8080/api/v1/clusters/cc/requests' -u admin:admin -H "X-Requested-By:anything" --data '{"RequestInfo":{"context":"Zookeeper server CONFIGURE", "command":"CONFIGURE","service_name":"ZOOKEEPER", "component_name":"ZOOKEEPER_SERVER", "hosts" : "vm-0.vm"}}' {code} Command is actually issued and sent to an agent. There are few minor troubles (at few service scripts, relevant method is called "config" instead of "configure"), but everything else seems to work well. I see two possible ways of issuing this command: 1) issue command directly from the UI during cluster deployment and also when needed afterwards (prefered). 2) during cluster deployment, issue this command authomatically at ambari-server during stage creation (like we do with service checks). That means that we have some internal flag like "unconfigured" for every service component host. This flag is stored at DB and defaults to "true". If this flag value is true for some STARTING component during stage creation process, additional CONFIGURE command is issued. After cluster deployment is finished, CONFIGURE command is issued manually by UI when configuration is updated. This way seems to be counter-intuitive. [~sumitmohanty], are we going to go the 1 way? In this case, I'll open a ticket for UI and close this jira.	1	1	https://issues.apache.org/jira/browse/AMBARI-4320
19603	AMBARI-4320	[~dmitriusan], how about wiring up START command to always do a configure and then start. This will also fix the RESTART issue where RESTART was not applying new configuration. That way we have CONFIGURE as a command available for scenarios that do not require START - e.g. client components. [~mahadev], what's your recommendation?	1	1	https://issues.apache.org/jira/browse/AMBARI-4320
19604	AMBARI-4320	Is this JIRA still valid. Please resolve if not.	0	0	https://issues.apache.org/jira/browse/AMBARI-4320
19605	AMBARI-4320	[~sumitmohanty] If start calls configure explicitly: * how do we deal with namenode starts which will be a low slower if they have to re-run configuration first * why bother having a separate configure command?	1	1	https://issues.apache.org/jira/browse/AMBARI-4320
19606	AMBARI-4283	[~mahadev], how are we planning to forward PUT and POST http requests (that have body): using http proxy protocol or as POST requests of some custom format? How will we deal with URL length limitation, or we ignore it?	1	1	https://issues.apache.org/jira/browse/AMBARI-4283

#	Requirement	Questions	Answered	Successfully Answered	URL
19607	AMBARI-4283	types, that we got from API call. And then, return back result, that we got after url execution. We can implement this functionality by API. But i think it's not correct. Because, to my mind, we can't map this functionality to our API architecture. In our API architecture, one of the main places, takes resources. But in our proxy, there are no resources, as i understand. At all, we need only service and provider, which will forward requests and responces. But if we need to implement it in API we can do the next. I'm not sure that names were choosen correctly(reason described before). We should create new resource type and resourcedefinition: {code} Resource.Type.Proxy {code} {code} public class ProxyResourceDefinition extends BaseResourceDefinition { public ProxyResourceDefinition () { super(Resource.Type.Proxy); } @Override public String getPluralName() { return "urls"; } @Override public String getSingularName() { return "url"; } @Override public Set<SubResourceDefinition> getSubResourceDefinitions() { return Collections.emptySet(); } } {code} After resource type and resourcedefinition were created, we need to add them to ResourceInstanceFactoryImpl.getResourceDefinition(...) It will look like: {code} case Proxy: resourceDefinition = new ProxyResourceDefinition(); break; {code} Then we can create new service class, something like ProxyService. It will look like: {code} @Path("/proxy/") public class ProxyService extends BaseService { @GET @Path("{url}") @Produces("text/plain") public Response proceedGetUrl(@Context HttpHeaders headers, @Context UriInfo ui, @PathParam("url") String url) { return handleRequest(headers, null, ui, Request.Type.GET, createProxyResource(url)); } @POST @Path("{url}") @Produces("text/plain") public Response proceedPostURL(String body, @Context HttpHeaders headers, @Context UriInfo ui, @PathParam("url") String url) { return handleRequest(headers, body, ui, Request.Type.POST, createProxyResource(url)); } @PUT @Path("{url}") @Produces("text/plain") public Response proceedPutURL(String body, @Context HttpHeaders headers, @Context UriInfo ui, @PathParam("url") String url) { return handleRequest(headers, body, ui, Request.Type.PUT, createProxyResource(url)); } @DELETE @Path("{url}") @Produces("text/plain") public Response proceedDeleteURL(@Context HttpHeaders headers, @Context UriInfo ui, @PathParam("url") String url) { return handleRequest(headers, null, ui, Request.Type.DELETE, createProxyResource(url)); } private ResourceInstance createProxyResource(String url) { return createResource(Resource.Type.Proxy, Collections.singletonMap(Resource.Type.Proxy, url)); } } {code} And now we can create provider class for our new resource. For example ProxyResourceProvider. Then we should add this provider to AbstractControllerResourceProvider.getResourceProvider(...) {code} case Proxy: return new ProxyResourceProvider(propertyIds, keyPropertyIds, managementController); {code} In ProxyPropertyProvider we should implement main functionality. Our main task is to execute URL request, which was sent for us by API, and return response for user. In our project we have URLStreamProvider class, which can help us to execute URL request. This class is used by Ganglia and Nagios to get some properties/info. Here are some code of the main method: {code} @Override public InputStream readFrom(String spec, String requestMethod, String params) throws IOException { if (LOG.isDebugEnabled()) { LOG.debug("readFrom spec:" + spec); } HttpURLConnection connection = spec.startsWith("https") ? (HttpURLConnection)getSSLConnection(spec) : (HttpURLConnection)getConnection(spec); String appCookie = appCookieManager.getCachedAppCookie(spec); if (appCookie != null) { LOG.debug("Using cached app cookie for URL:" + spec); connection.setRequestProperty(COOKIE, appCookie); } connection.setConnectTimeout(connTimeout); connection.setReadTimeout(readTimeout); connection.setDoOutput(true); connection.setRequestMethod(requestMethod); if (params != null) connection.getOutputStream().write(params.getBytes());	1	1	https://issues.apache.org/jira/browse/AMBARI-4283
19608	AMBARI-4205	Hi Michael, what branch is this patch for? I'd like to apply it and give it a try. I understand it is extremely hard to write tests for this, but if you have any please include them. Also, when you submit a patch, please use the "Submit Patch" button above - that will trigger a notice to be sent. Thanks!	0	0	https://issues.apache.org/jira/browse/AMBARI-4205
19609	AMBARI-4205	[~miharp], what is the status of the patch/JIRA on ReviewBoard?	1	0	https://issues.apache.org/jira/browse/AMBARI-4205
19610	AMBARI-4150	Overall looks good. Few comments and questions ... {quote} if (!isSchedulerAvailable()) { throw new AmbariException("Scheduler unavailable."); } {quote} The above check is there in few places. The check itself is good but what are the condition for scheduler not being available. If it is not available then can one still submit requests the do not have a schedule. * ambari-server/src/main/java/org/apache/ambari/server/scheduler/ExecutionScheduleManager.java ** If any exception is thrown in the method (scheduleBatch) then how does RequestExecution recover? As I understand, triggers are created when jobs are persisted and when ambari-server starts. So if an error is encountered while triggering a job then does it mean that till ambari-server restarted no new jobs for that RequestExecution will not be created ** Is there a validation somewhere to ensure that batch orderIds are unique within the batch? ** synchronized Long executeBatchRequest() - seems incomplete * ambari-server/src/main/java/org/apache/ambari/server/state/scheduler/BatchRequestResponse.java ** Is it intentionally empty? * ambari-server/src/main/resources/Ambari-DDL-MySQL-CREATE.sql ** Lets move the quartz related tables closer to rest of the tables in "ambari" db * ambari-server/src/test/java/org/apache/ambari/server/orm/InMemoryDefaultTestModule.java ** Do we need these changes? * ambari-server/src/test/java/org/apache/ambari/server/state/ConfigGroupTest.java ** Revert these changes	1	1	https://issues.apache.org/jira/browse/AMBARI-4150
19611	AMBARI-4150	Thanks for the feedback. Few concerns addressed below: - "Scheduler unavailable." This is a preemptive abort, since a scheduler instances that is shutdown cannot be started again, so server needs to be restarted. This ensures that we do not start creating jobs / triggers and then fail. Agree that exception message can be improved. - If any exception is thrown in the method (scheduleBatch) then how does RequestExecution recover? A chain of request will execute based on schedule even if the previous request failed or was aborted midway, the last_execution_status is meant to capture the status of the previous chain of requests. - Is there a validation somewhere to ensure that batch orderIds are unique within the batch? org.apache.ambari.server.controller.internal.RequestScheduleResourceProvider#validateRequest - synchronized Long executeBatchRequest() - seems incomplete This will be addressed in the followup Jira which enables request execution. - ambari-server/src/main/java/org/apache/ambari/server/state/scheduler/BatchRequestResponse.java Yes, same as above - ambari-server/src/main/resources/Ambari-DDL-MySQL-CREATE.sql Done - ambari-server/src/test/java/org/apache/ambari/server/orm/InMemoryDefaultTestModule.java Unnecessary imports added by previous patch reverted. - ambari-server/src/test/java/org/apache/ambari/server/state/ConfigGroupTest.java Formatting change.	1	1	https://issues.apache.org/jira/browse/AMBARI-4150
19612	AMBARI-4034	Looks good. Few comments: Can we rename classes with *RequestExecution* in name to *RequestSchedule*? * ambari-server/src/main/java/org/apache/ambari/server/state/cluster/ClusterImpl.java ** If we delete a RequestSchedule resource, what happens to the underlying Request that got spawned because of the one being deleted? For now, we could not delete RS instances as we also do not allow delete of Requests either. ** Can a RS be disabled? So that no new request execution instance get fired. * ambari-server/src/test/java/org/apache/ambari/server/orm/InMemoryDefaultTestModule.java ** No change in this file. * We need some unit test for the RequestSchedule API calls, and the changes to the JSON body parser	1	1	https://issues.apache.org/jira/browse/AMBARI-4034
19613	AMBARI-3909	Is the MySQL DDL missing from the patch? Ambari-DML-MySQL-INSERT_METAINFO.sql	1	1	https://issues.apache.org/jira/browse/AMBARI-3909

#	Requirement	Questions	Answered	Successfully Answered	URL
19614	AMBARI-3835	Reviewing the patch, and am thinking that instead of {code} {{bindAttr data-original-title="view.removeHostTooltip"...}} {code} it might be better to do this if it works {code} data-original-title={{t .....}} {code}	1	1	https://issues.apache.org/jira/browse/AMBARI-3835
19615	AMBARI-3786	Thanks [~BuzhorDenys]. +1 for the patch.	0	0	https://issues.apache.org/jira/browse/AMBARI-3786
19616	AMBARI-3762	+1 on the patch. [~odiachenko], could you add a comment to this Jira describing the fix?	1	0	https://issues.apache.org/jira/browse/AMBARI-3762
19617	AMBARI-3753	Hi Erin, The patch looks good - do you want to take a look at the new stack definition that add the new service pluggability to the stack definition? Though we can do that later. DO note that everythign will be moving to a new stack definition extensivbility. [~ncole@hortonworks.com] should have more info and can help - again that can be done later as well but we should keep track of that. Why is the patch for AMBARI-3167 needed for ambari server chages?	1	1	https://issues.apache.org/jira/browse/AMBARI-3753
19618	AMBARI-3753	THat sounds good Erin. The patch looks good to me. Can you please update the jira with mvn test output with htis patch? We can go ahead and commit thsi patch if the mvn test passes all fine.	1	1	https://issues.apache.org/jira/browse/AMBARI-3753
19619	AMBARI-3731	Comments: - public Stage getAction(String actionId); This should be renamed to getStage or getStageByActionId - if (!requestEntity.getIsCommand()) Is there a possibility for neither action nor command (in future), can this be simplified to: request.isAction() - validateCustomAction Is there a case for checking if a Service belongs to the cluster and not just the stack? - AmbariManagementControllerImpl.CLUSTER_DEFAULT_TAG Duplicate definition in ConfigHelper.java - Changes need to be replicated to other DDL scripts at ambari-server/src/main/resources For MySQL, Oracle and Pgsql remote.	0	0	https://issues.apache.org/jira/browse/AMBARI-3731
19620	AMBARI-3731	{quote} public Stage getAction(String actionId); This should be renamed to getStage or getStageByActionId {quote} Renamed. {quote} if (!requestEntity.getIsCommand()) Is there a possibility for neither action nor command (in future), can this be simplified to: request.isAction() {quote} Line refactored out. {quote} validateCustomAction Is there a case for checking if a Service belongs to the cluster and not just the stack? {quote} This is done when execution is scheduled. If the service does not belong to the stack then no hosts are identified as a target which is an error. {quote} AmbariManagementControllerImpl.CLUSTER_DEFAULT_TAG Duplicate definition in ConfigHelper.java {quote} Removed one. {quote} Changes need to be replicated to other DDL scripts at ambari-server/src/main/resources For MySQL, Oracle and Pgsql remote. {quote} Modified the upgrade script. Opened a JIRA for Oracle changes. MySQL is quite outdated and so is PgSQL remote. {quote} at ExecutionCommand.java, we already have 'String role' field that stores component name. Maybe it's better to rename it instead of adding new 'String componentName' field {quote} Role is being used as the name of the custom action or commands. We may need to revisit the structure after agent side changes are finalized for service pluggability.	1	0	https://issues.apache.org/jira/browse/AMBARI-3731
19621	AMBARI-3730	[~xiwang0309], thanks for the updated patch. There are AJAX calls with "async: false". Can we change the code to not use this option?	1	1	https://issues.apache.org/jira/browse/AMBARI-3730
19622	AMBARI-3698	Do we need the quotes around "users" option? If not, let's drop the quotes so that it's consistent with the description (and a little cleaner). Also, let's change "cleanup" to "clean up" in the description, as "cleanup" is not a verb.	1	1	https://issues.apache.org/jira/browse/AMBARI-3698
19623	AMBARI-3698	Uploaded patch does the fixes as suggested by [~u39kun]. Moves the note text as suggested by [~sposetti]. > 2) anyway to make the font one smaller so the whole code instruction line fits on a single line? host check.png image is just the pop-up box image. host checks-2.png is entire screen image with updated fixes. The font size for the code is already sufficiently small. Also doing copy of multiple lines of a command on the ui screen and pasting in the shell to run single command works as expected. Also apart from these suggested changes, replaced options -s with --silent and -k users with --skip=users as it's more narrative then short hand notations.	1	1	https://issues.apache.org/jira/browse/AMBARI-3698
19624	AMBARI-3672	Patch is looking good. One question though. Let's pretend you have a cluster, that has some config groups defined. But then, you have just one host, already part of a config group where you want to tweak, say, some heap settings. Do you create a completely new ConfigGroup for that one host, or do/can you override that one host directly?	1	1	https://issues.apache.org/jira/browse/AMBARI-3672
19625	AMBARI-3672	Few comments: == Not part of this patch but can you rename the input parameter? ambari-server/src/main/java/org/apache/ambari/server/agent/HeartbeatMonitor.java HeartbeatMonitor(Clusters fsm, <-- == ambari-server/src/main/java/org/apache/ambari/server/state/ConfigHelper.java getEffectiveDesiredTags() - Comments "override (1) with service-specific overrides" and "override (2) with host-specific overrides" are no longer applicable. And, the code now collects all tags and does not do any override. Line 280: if (serviceInfo.hasPropertyFor(type, keys) || !hasPropertyFor(stackId, type, keys)) For types that are not global why don't we just declare the SCH to be "stale" when actual does not contain the "type". I think the per property check sounds risky in two ways - too many properties to check and assumes that the stack definition is maintained correctly wrt. property mapping. Line 273: if (!serviceInfo.hasConfigType(type)) We should have a JIRA to make sure that the mapping of type to service is verified for 1.3.2, 1.3.3 and 2.0.6. Can you open one? Line 384: desiredSet.removeAll(actualSet); Isn't this only true if the size of the sets are the same? Line 294: Collection<String> changed = findChangedKeys(cluster, type, tags.values(), actualTags.values()); Should we do it only for "global"? == ambari-server/src/main/java/org/apache/ambari/server/controller/internal/ConfigGroupResourceProvider.java No change in this file. == asf-ambari/ambari-server/src/main/java/org/apache/ambari/server/state/host/HostImpl.java public Map<String, DesiredConfig> getDesiredConfigs(long clusterId) <-- this should be removed, is it being done through another JIRA? Do we have a unit test for agent returning actual config in the old form and it being handled by server? I am thinking upgrade.	1	1	https://issues.apache.org/jira/browse/AMBARI-3672

#	Requirement	Questions	Answered	Successfully Answered	URL
19626	AMBARI-3606	_Sorry for the delay. Let's open a separate JIRA for the ones needed to be incorporated._ * What happens when a host gets deleted? Does the ref from CG to host also get deleted? Is there a unit test for that? * (Minor nit) Javadoc missing for new methods on AmbariManagementController * ambari-server/src/main/java/org/apache/ambari/server/controller/internal/ConfigGroupResourceProvider.java ** Set<ConfigGroupResponse> getConfigGroups () - is it possible to have a request with more than one predicate, e.g. groupName and hostName? * These files can be reverted (unused imports or fields) ** ambari-server/src/main/java/org/apache/ambari/server/controller/internal/RequestStatusImpl.java ** ambari-server/src/main/java/org/apache/ambari/server/controller/internal/ResourceImpl.java ** ambari-server/src/main/java/org/apache/ambari/server/state/Clusters.java ** ambari-server/src/main/java/org/apache/ambari/server/state/cluster/ClustersImpl.java * ambari-server/src/main/resources/Ambari-DDL-Postgres-REMOTE-CREATE.sql' ** Should not there be a corresponding change to Ambari-DDL-Postgres-CREATE.sql Following unit tests fail: Tests in error: * agentRegistration(org.apache.ambari.server.agent.AgentResourceTest): Guice creation errors * agentHeartBeat(org.apache.ambari.server.agent.AgentResourceTest): Guice creation errors * agentHeartBeatWithEnv(org.apache.ambari.server.agent.AgentResourceTest): Guice creation errors * deserializeClasses(org.apache.ambari.server.agent.AgentResourceTest): Guice creation errors	0	0	https://issues.apache.org/jira/browse/AMBARI-3606
19627	AMBARI-3600	Overall changes look good. The actual staleness determination logic seem complex (probably necessary) and hence needs some more comments. Some comments and questions: * ConfigHelper.java ** isStaleConfigs() - I think we can add more details to this comment "Configs are considered stale when:". _It, as it stands today, is a little cryptic to me_ Also when we add support for ConfigGroup we need to update this comment. ** hasPropertyFor() - why is this required to determine if config is stale for a SCH? ** findChangedKeys() - We do a putAll() in this method. Does it assume that there are no over-lapping key names? Also, did we not need this only for "global" types? _In that case, we are OK as there are no over-lapping names in global_ ** It seems we need a ConfigHelperTest that is dedicated for this type. There are some *relatively-complex* comparison logic in this class. The tests can also have comments in them to layout the scenarios and why it should (or should not) result in stale config * StackServiceResponse.java ** How do we capture this relationship, TASKTRACKER/JOBTRACKER can be related to core-site and core-site change needs to result in these guys restart? * ambari-server/src/main/java/org/apache/ambari/server/state/ServiceInfo.java ** buildConfigLayout() - its probably not a concern but do we need a lock around populating "configLayout"? Or, at least create the layout on a temp variable and then assign it to "configLayout". * ambari-server/src/test/java/org/apache/ambari/server/state/svccomphost/ServiceComponentHostTest.java ** testStaleConfigs(). Can we use HDP-2.0.6 for test? This test probably covers my ask for dedicated test class for ConfigHelper.	1	1	https://issues.apache.org/jira/browse/AMBARI-3600
19628	AMBARI-3600	* ConfigHelper ** Will update comment to be more clear ** hasPropertyFor() - This is required because we need to check to see if a property is stale for ANY OTHER service than the SCH that we are checking, because the OTHER one will be considered stale. That check is designed to look forcustom properties that are assigned to a type but are not defined as part of the stack. For example, if you assign a value to core-site that is undefined in any stack, any service relying on core-site will be considered stale. ** findChangedKeys() - it's ok for overlap in this case, because even overlapping (two properties defined in two different services) means it's stale. ** The method is (more or less) directly tested in ServiceComponentHostTest.testStaleConfigs() as you mentioned. * StackServiceResponse ** TT/JT - the stack service response only indicates what configs a service defines, not the relationship of components to the config. Any stale config value will make both TT and JT stale. * ServiceInfo ** Will add synchronization to the method, but note that ServiceInfo is effectively singleton-per-stack, and no other synch activity happens around stacks as it is today. * testStaleConfigs() addresses the SCHTest you mentioned. Will switch to 2.0.6 and verify.	1	1	https://issues.apache.org/jira/browse/AMBARI-3600
19629	AMBARI-3600	LGTM, +1 One question: Is it expected that when there is a tag mismatch for global, then SCH restart must be done only if changed properties affect the SCH. E.g. if ZK properties change it should not restart other SCH	1	1	https://issues.apache.org/jira/browse/AMBARI-3600
19630	AMBARI-3585	Hmm, I am not convinced that the sqlite stuff is even used. Default seems to be derby and postgres. Do we need it at all? Should we remove JDBCProviderModule, DBHelper, and SQLiteConnectionFactory?	1	1	https://issues.apache.org/jira/browse/AMBARI-3585
19631	AMBARI-3567	committed to trunk. Do we need this on any existing branch?	1	0	https://issues.apache.org/jira/browse/AMBARI-3567
19632	AMBARI-3536	[~odiachenko] [~ncole] Is this JIRA still valid?	1	0	https://issues.apache.org/jira/browse/AMBARI-3536
19633	AMBARI-3524	why did this issue not assigned for so long a time?? Is there no one had concern on this issue?	0	0	https://issues.apache.org/jira/browse/AMBARI-3524
19634	AMBARI-3482	Chris, Want to upload it as a patch?	0	0	https://issues.apache.org/jira/browse/AMBARI-3482
19635	AMBARI-3482	Will this make it into the next version of Ambari?	1	0	https://issues.apache.org/jira/browse/AMBARI-3482
19636	AMBARI-3353	I think the best way here, is to convert the odd bootstrap response, to the usual format, so user and our postproccessing functions can work with that w/o troubles. [~subin11], could you review this please?	1	0	https://issues.apache.org/jira/browse/AMBARI-3353
19637	AMBARI-3317	Patch looks good, minor comments below: 1. Is this comment valid? + // two classes define the top-level element "metainfo", so we need 2 contexts. 2. Could we add a negative unit test case for the testing a JAXBException? Overall, LGTM +1	1	1	https://issues.apache.org/jira/browse/AMBARI-3317
19638	AMBARI-3266	[~mahadev], So the LICENSE file can be removed because the one at the Ambari project root applies? Are you saying that the text of MSPL-LICENSE file should be included in the NOTICE.txt file at the Ambari project root? Should the MSPL-LICENSE file stay under contrib/ambari-scom/ or be removed also? I thought that including the MSPL-LICENSE file in the project was sufficient. If this is true, should we add it to NOTICE.txt? From the Apache legal reference ... {quote} It is important to keep NOTICE as brief and simple as possible, as each addition places a burden on downstream consumers. Do not add anything to NOTICE which is not legally required. {quote}	1	1	https://issues.apache.org/jira/browse/AMBARI-3266
19639	AMBARI-3264	Will this we worked on?	1	0	https://issues.apache.org/jira/browse/AMBARI-3264
19640	AMBARI-3252	* setup_security(args). Could you use get_validated_string_input() to cover the prompt and choice selection. Something similar to: {code} idType = "0" idType = get_validated_string_input("Choose one of the following options:\n1 - " + "Enable HTTPS for Ambari server" + "\n2 - " + "Enable HTTPS for Ganglia service" + "\n3..... ", idType, "^[12345]$", "Invalid choice.", False) {code} Its a good change and will make the prompt clearer. LGTM, +1.	1	1	https://issues.apache.org/jira/browse/AMBARI-3252
19641	AMBARI-3211	LGTM, +1. Is there an unit test dependency on this change?	1	0	https://issues.apache.org/jira/browse/AMBARI-3211
19642	AMBARI-3201	subin, can you please review this?	1	0	https://issues.apache.org/jira/browse/AMBARI-3201
19643	AMBARI-3192	[~xiwang0309], actually this does not apply on the latest trunk code. Can you update the patch?	0	0	https://issues.apache.org/jira/browse/AMBARI-3192

#	Requirement	Questions	Answered	Successfully Answered	URL
19644	AMBARI-3167	Erin A Boyd, I am trying to understand why generic HCFS option was changed to GlusterFS? Are there any issues that you have noticed in supporting generic HCFS implementation?	0	0	https://issues.apache.org/jira/browse/AMBARI-3167
19645	AMBARI-3140	the patch looks good - can we add a unit test for this?	1	1	https://issues.apache.org/jira/browse/AMBARI-3140
19646	AMBARI-3130	There is a method in this patch that loads properties from new api. It is used in next patch that also uses new api. Should i make a new patch with this metod or this one can be commited?	1	1	https://issues.apache.org/jira/browse/AMBARI-3130
19647	AMBARI-3073	[~subin11] unfortunately [~aonishuk] has quite a few patches on client at: AMBARI-3078, AMBARI-3079, AMBARI-3080, AMBARI-3081, AMBARI-3082. It would be counter productive to have zip attachments for the client package. Can we do patches and diffs because it would not be possible to collaborate on the client library if the updates are as re writes?	1	0	https://issues.apache.org/jira/browse/AMBARI-3073
19648	AMBARI-3073	Hi Mahadev, unfortunately ever since Ambari code has moved away from svn ,I am unable to create proper patches with git. I had given the above code on 30/8/2013 to you in mail.As of now I have merged code from my branches into my trunk code and I cannot find it feasible to go back and recreate a git patch which matches with AUG-30 code. Here is what I propose after I looked at AMBARI-3078, AMBARI-3079, AMBARI-3080, AMBARI-3081, AMBARI-3082. 1)accept the zip code and checkin 2)I will update the architecture of client code, so that Andrew or anyone else can add the code at appropriate files. This will result in: 1)AMBARI-3080 issue becoming fixed. My next patch contains the code for create cluster.This will fix or help Andrew re-create his patches (1)AMBARI-3079 .AMBARI-3081 ,AMBARI-3082 , The above 3 issues are related to create hosts .here are few things which will come in the next patch. (1.a) /api/v1/hosts ...is done by ...myclient.create_hosts(host_ip_rack_tuple_list) /api/v1/clusters/c1/hosts..is done by...mycluster.create_hosts(host_ip_rack_tuple_list) (1.b) all the functions in a module are private. (1.c) all responses have to wrapped in a StatusModel object.(so 3082 has to re-written) 2)AMBARI-3078...this is available on my latest trunk.Anyways It is better if Andrew re-creates the patch once he plays with AUG_30 code In short .Of the 5 issues - 1 will become fixed.2-have to be recreated because they need further improvements.2 others can be postponed by few days. I dont see too much impact,Kindly accept the code before more issues are raised and this becomes unmanageable. The client code has lot of methods which I have not been merged properly due to my lack of time.But I have taken care by exposing only a certain number of methods in the ambari_api.py file.this file contains the entry point for accessing all the client code. Hope this clears the confusion Many thanks Subin	1	1	https://issues.apache.org/jira/browse/AMBARI-3073
19649	AMBARI-3073	Committed to trunk. Thanks a lot Subin!	0	0	https://issues.apache.org/jira/browse/AMBARI-3073
19650	AMBARI-3045	Looks like the patches were never committed properly. We added this because the new version of Ganglia was filling up disk and we didn't really know why. Since then the RRD config files have been fixed so this is not really an issue anymore. Marking this as Won't Fix.	0	0	https://issues.apache.org/jira/browse/AMBARI-3045
19651	AMBARI-2989	[~atkach], what was the issue with earlier patch?	0	0	https://issues.apache.org/jira/browse/AMBARI-2989
19652	AMBARI-2984	Dax, Good point. Want to upload a patch?	0	0	https://issues.apache.org/jira/browse/AMBARI-2984
19653	AMBARI-2954	[~ababiichuk], I see the following properties on the server side for HDP2 stack - {{mapreduce.cluster.local.dir}}, {{mapreduce.cluster.reducememory.mb}} in the file {{ambari-server/resources/stacks/HDP/2.0.5/services/MAPREDUCE2/configuration/mapred-site.xml}}. Did you test that these 2 properties did not end up after install? If they do end up in there, we might need cleaning up on server stack definition also.	1	1	https://issues.apache.org/jira/browse/AMBARI-2954
19654	AMBARI-2942	Hey guys, can anyone please tell why this was done?	0	0	https://issues.apache.org/jira/browse/AMBARI-2942
19655	AMBARI-2852	Any tentative date to get this fix? which version will have this feature?	0	0	https://issues.apache.org/jira/browse/AMBARI-2852
19656	AMBARI-2852	Does somebody has any feedback or updates on this issue? Thanks	0	0	https://issues.apache.org/jira/browse/AMBARI-2852
19657	AMBARI-2819	Nice. The code looks pretty good. Few minor comments. We need to add more unit tests - see below for suggested ones. Comments: *File*: ambari-server/src/main/java/org/apache/ambari/server/controller/StackVersionResponse.java * private String parentVersion; ** Why do we need to expose the parent version. Is this relevant to the API consumer? *File*: ambari-server/src/main/java/org/apache/ambari/server/controller/internal/StackVersionResourceProvider.java * public Set<Resource> getResources(Request request, Predicate predicate) ** Add back the @Override tag *Test related* * We should add tests for ** New configuration property being added to a service ** Similarly, removing a configuration and also throw in a check for configs that did not change ** Adding a component ** Also pick one test from AmbariManagementController unit tests and modify it to use 2.0.6 rather than 2.0.5. This will validate that unit tests can run with the inheritance model * Add some unit tests for key helper methods such as ** getParentStacksInOrder ** getAllApplicableServices	1	1	https://issues.apache.org/jira/browse/AMBARI-2819
19658	AMBARI-2819	*private String parentVersion;* Why do we need to expose the parent version. Is this relevant to the API consumer? - Since the only way to look at a deployed cluster and find out how the service definition is computed, is by following the parent version chain. This is for someone with access to API but not the server host with XML definitions. - Added \@Override tag - Unit tests added: # New config property # Property that did not change # AmbariManagementController.testReInstallClientComponent - 1.2.0 to 2.0.6 # Added test for getParentStacksInOrder # Added test for getAllApplicableServices	1	1	https://issues.apache.org/jira/browse/AMBARI-2819
19659	AMBARI-2738	{code} - if (newState == oldSchState) { + // Client component reinstall allowed + if (newState == oldSchState && !sc.isClientComponent()) { sch.setDesiredState(newState); {code} What happens we the new desired state is INIT, INSTALL_FAILED, MAINTENANCE for the client component?	1	1	https://issues.apache.org/jira/browse/AMBARI-2738
19660	AMBARI-2714	1) Would like to see concept of stack inheritance. A user can define a stack that extends another stack. The resulting stack can add a new service (Stack A is extended by Stack B, which is exactly the same except Stack B has N-new service) and/or turn off a service in the parent stack. 2) In the service/configurations/*, there should be a way to define how config files + properties show in the UI (including a way to group). For example, how would you describe the Configs tab on HDFS service via meta data (instead of having to code the JS page)? If you can define the UI in meta data, then the UI can read that info and be dynamic.	1	1	https://issues.apache.org/jira/browse/AMBARI-2714
19661	AMBARI-2714	Currently using Supervisor (http://supervisord.org/) to manage these types of processes. Perhaps a simple start to this activity would be to automatically configure Supervisor on each node, and deploy the necessary java and configuration files, etc. What do you think?	1	0	https://issues.apache.org/jira/browse/AMBARI-2714
19662	AMBARI-2714	Is the idea that the user configure all the xml/json files prior to loading the new service or will the UI provide a way to just load the new service and deploy it accordingly? Also, we need to consider what resources need to be loaded with the new service (rpms, tarball, etc..) to enable the new service. Might want to also consider the 'control' from Ambari to the new service. If it's just a 'viewer' or controller of the service itself.	1	0	https://issues.apache.org/jira/browse/AMBARI-2714

#	Requirement	Questions	Answered	Successfully Answered	URL
19663	AMBARI-2714	@Erin - initially it will be configuring XML and have the UI recognize it as an addable service. The capability exists to reload stack definitions already, so that's no big deal. Long term I agree - we want to be supply new services using the UI itself. You are also correct in that we want to be able to supply all the artifacts surrounding a new service including tarballs, scripts etc. To start with, at least, we'll probably focus on defining a service without modifying code. Then we can start adding real capability. @Abhishek - I think a great use-case will be to try to add Storm without modifying Ambari code. @K. Eric - I'm not too familiar with Supervisor - how would you envision that fitting with Ambari (it handles monitoring and control already)?	1	1	https://issues.apache.org/jira/browse/AMBARI-2714
19664	AMBARI-2714	@Nate - Supervisor is a daemon implemented in Python (https://pypi.python.org/pypi/supervisor) that we have been using to automatically start and manage services on Linux. If Ambari provides those capabilities so much the better. For the Supervisor configuration you specify the command line that invokes the service in a manner that will send the log to standard output. Supervisor confirms that records are being written to standard out and terminates the process if it hangs, then restarts using the command line. Primitive, but that is what we are using for example to have the Storm UI daemon running when the machine boots. Anyone who is using Supervisor should be able to migrate to Ambari if a similar configuration is expected. What do you think?	1	0	https://issues.apache.org/jira/browse/AMBARI-2714
19665	AMBARI-2714	These design docs seem to focus quite a bit on backend components and reporting - is there any thought going into modularizing the installation wizard GUI's at this point and what that would look like? Also, is there any prototyping going on for the designs yet? It seems like many aspects have been fairly well thought out. Would it be beneficial to start looking at the Java classes mentioned in the design docs and prototyping implementations of the spec?	1	0	https://issues.apache.org/jira/browse/AMBARI-2714
19666	AMBARI-2699	Does Ambari Server supoorts only one or multiple version of API at the same time? If multiple versions are supported how user will know to which version he sends requests? I assume for each version of Amabri we should have different version segment in the url. For instance: if one of supported version is "version3", url for sending requests should bi http://aaa:8080/api/v3/...., if version is "version5" url should be http://aaa:8080/api/v5/...	1	0	https://issues.apache.org/jira/browse/AMBARI-2699
19667	AMBARI-2699	[~brian_swan], I like the idea of getting the supported API versions through /api. I'm not sure about having the server resource at that level since it would be un-versioned, meaning we would never be able to change it. What do you think about adding an "api" resource as well as a "server" resource? Getting the "api" resource might look like ... {code} GET api/ { "href" : "http://aaa:8080/api", "API" : { "supported_versions" : "1,2,3" } }{code} The "server" resource would look something like what is described above and would be used to get the Ambari server/agent version and configuration.	1	1	https://issues.apache.org/jira/browse/AMBARI-2699
19668	AMBARI-2691	[~tbeerbower] the patch looks good but we will have to enable encrypt passwords for any password entries in ambari.properties so that the truststore is only readable by ambari-server [~swagle] can help with how to do (should be simple). Also what kind of trusstore can it be? Can we import multiple certs into the trusstore? Also can we run the setup nagios https multiple times? Also does removing the 2 properties disable the certs? Also what happens when ambari server is running as a different user. Will the trustore perms need ot be changed?	1	1	https://issues.apache.org/jira/browse/AMBARI-2691
19669	AMBARI-2691	Thanks [~mahadev], {quote} we will have to enable encrypt passwords {quote} Okay, I'll discuss it with [~swagle]. {quote} what kind of trusstore can it be? {quote} You can set the truststore type in ambari.properties. Should I prompt for that in the setup? {quote} Can we import multiple certs into the trusstore? {quote} Yes, I use the aliases 'ganglia_cert' and 'nagios_cert'. Others can be added as long as they have different aliases. {quote} Also can we run the setup nagios https multiple times? Also does removing the 2 properties disable the certs? {quote} It works like setup-https. The first time you run the setup it will prompt to make sure that you want to enable https, prompt for the truststore path and passord (if needed) and then prompt for the cert to be added. If you run setup when https for the component (ganglia or nagios) is already enabled, it will prompt to see if you want to disable https. If yes then the cert will be removed and https marked disabled (set ganglia.https=false or nagios.https=false in ambari.properties). You can set the properties directly in ambari.properties to disable https but the cert will not be removed from the truststore. {quote} Also what happens when ambari server is running as a different user. Will the trustore perms need ot be changed? {quote} I think ambari server will need read access. I don't know any way around that.	1	1	https://issues.apache.org/jira/browse/AMBARI-2691
19670	AMBARI-2691	Has the additions for encrypted password and truststore types made any changes in the set up procedure?	0	0	https://issues.apache.org/jira/browse/AMBARI-2691
19671	AMBARI-2562	Looks like there are some unnnecessary formatting changes. Can we upload a cleaner patch?	1	1	https://issues.apache.org/jira/browse/AMBARI-2562
19672	AMBARI-2555	Hi [~atkach], I applied the patch on top of r1499654 in trunk. This does not work in App.testMode. Can you make sure that this works in testMode?	1	1	https://issues.apache.org/jira/browse/AMBARI-2555
19673	AMBARI-2555	You are right, the patch failed to apply for some files. Can you upload a patch that works directly on top of the current latest trunk r1499654?	1	1	https://issues.apache.org/jira/browse/AMBARI-2555
19674	AMBARI-2555	[~atkach], can you do a clean checkout of trunk and apply AMBARI-2555.patch posted here? When I do that, it fails to apply on two files.	1	1	https://issues.apache.org/jira/browse/AMBARI-2555
19675	AMBARI-2541	[~atkach], the patch does not apply on trunk. Can you rebase it to the latest trunk (r1500106)?	1	1	https://issues.apache.org/jira/browse/AMBARI-2541
19676	AMBARI-2529	Tom, Can you please go aehad and commit the patch to trunk and 1.2.5 branch?	1	1	https://issues.apache.org/jira/browse/AMBARI-2529
19677	AMBARI-2487	[~jaimin], can you describe which dependencies are being removed? "Resolution of this ticket includes reducing agent's dependency on global" is a bit vague.	1	0	https://issues.apache.org/jira/browse/AMBARI-2487
19678	AMBARI-2487	"There are other globals that are sent by web-ui but not used by Agent. Also some globals that are sent to globals can be retrieved from site properties but agent uses globals like kerberos_domain, hbase_primary_name and hbase_regionserver_principal." Are you referring to what the web UI and agent do currently, and that upon resolution of this ticket, we will not be doing so?	0	0	https://issues.apache.org/jira/browse/AMBARI-2487
19679	AMBARI-2487	[~jaimin] Is this JIRA still valid? Can you update the JIRA accordingly?	0	0	https://issues.apache.org/jira/browse/AMBARI-2487

#	Requirement	Questions	Answered	Successfully Answered	URL
19680	AMBARI-2453	Most of the comments are minor - some are optional to incorporate. Let's add a black-list of items to avoid when erasing. Based on our experiments, we can figure out what needs to be in the list. ========== +HOST_CHECK_FILEPATH_DEFAULT = "/var/lib/ambari-agent/data/hostcheck.result" +OUTPUT_FILEPATH_DEFAULT = "/var/lib/ambari-agent/data/hostcleanup.result" ===== Let's read the "prefix" value from ambari-agent.ini and use that as the base path. ========== + if options.skip: + global SKIP_LIST + SKIP_LIST = options.skip.splt(',') ===== Let's by default skip the users. ========== + raise RuntimeError('HostCleanp needs to be run as root.') ===== Spelling - ...Cleanp ========== + if propMap: + h.do_cleanup(propMap) ===== Lets add some conclusion message. ========== + def do_erase_dir_silent(self, pathList): ... + else: + logger.info(path + " is a file and not a directory, deleting file") + try: + os.remove(path) ===== Let's call the do_erase_file_silent. That way it will be easier to maintain. ========== + def do_erase_packages(self, packageList): + packageStr = None + if packageList: + packageStr = ' '.join(packageList) ===== If erasing one package fails will this method remove remaining packages. Otherwise, we can erase then one by one. ========== + else: + logger.warn("Unsupported OS type, cannot get repository location.") + return [] ===== There are one or two places where we also check for OS type. We can report the same error on unsupported OS type. ========== + def find_repo_files_for_repos(self, repoNameList): ... + with open(filePath, 'r') as file: + contents = file.read() + for repoName in repoNameList: + if repoName and repoName in contents and filePath not in repoFiles: + repoFiles.append(filePath) ====== This I am a bit concerned about. We need to tighten the search to make sure that we are not accidentally removing repos that we need. Let's talk. ========== Can we add an unit test that fakes a host-check report and then makes the necessary calls?	1	1	https://issues.apache.org/jira/browse/AMBARI-2453
19681	AMBARI-2442	[~subin11] the patch looks good to me. Have you done any end to end testing with delete services?	0	0	https://issues.apache.org/jira/browse/AMBARI-2442
19682	AMBARI-2420	[~ncole@hortonworks.com], can you give us an example of how to batch update calls? For example, HDP-1.3.0 has 8 different operating system subresources. I want to be able to specify the base url for all 8 operating systems in one call.	1	1	https://issues.apache.org/jira/browse/AMBARI-2420
19683	AMBARI-2420	Also, how can we retrieve all the repo base URLs for a given stack-version in one call? Or better yet, how we can retrieve all base URLs for all stack versions in one call (preferred).	1	1	https://issues.apache.org/jira/browse/AMBARI-2420
19684	AMBARI-2420	To clarify: The following is currently set: "http://<server>/HDP/centos5/1.x/GA/1.3.0.0" Let's pretend someone wants to customize their repo. What would the base be? Suppose they use "http://<internal>/companyRepo/hadoop-repo" I think we can do this with a replace-token so if you specify a URL like "http://<internal>/companyRepo/{os_type}/myversion" then we can replace the OS when we save it. Would that be acceptable?	1	0	https://issues.apache.org/jira/browse/AMBARI-2420
19685	AMBARI-2420	But, that doesn't resolve the fact that we have multiple repos per OS. Do you want to set ALL base urls, including HDP-epel?	1	0	https://issues.apache.org/jira/browse/AMBARI-2420
19686	AMBARI-2373	Thanks [~xiwang0309]. A couple of things... 1. Can we externalize the Handlebar templates to a separate file, rather than embedding in the Views? For very simple templates, embedding is OK, but when it gets to a moderate size it's better to put on separate files so that the IDE can apply syntax highlighting, etc. 2. Instead of commenting out lines, we should just remove them.	1	1	https://issues.apache.org/jira/browse/AMBARI-2373
19687	AMBARI-2330	This was resolved in branch of tickets. Like AMBARI-14983, AMBARI-15194, AMBARI-15592 and other	1	1	https://issues.apache.org/jira/browse/AMBARI-2330
19688	AMBARI-2294	Is there an enumerated list of "rules" to follow for what each progress meter should be displaying? I think my original swing at this problem may have been off target because I wasn't aware of the overall expectations. Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-2294
19689	AMBARI-2294	[~croberts], have you had a chance to take a look at this?	1	1	https://issues.apache.org/jira/browse/AMBARI-2294
19690	AMBARI-2256	+1, LGTM. The title says 2.0.5, is that right?	1	0	https://issues.apache.org/jira/browse/AMBARI-2256
19691	AMBARI-2174	*ambari-server/conf/unix/ambari.properties* * ...isPersisted --> ...persisted (other true/false properties do not seem to use "is" *ambari-server/src/main/java/org/apache/ambari/server/configuration/Configuration.java* * sort the import statements * MASTER_KEY_FILENAME_DEFAULT - should we use a more descriptive name such as "master.key" or "ambari.master.key"? * credentialProvider (it will be initialized to null by default) * should loadCredentialProvider() use a lock around instantiating the provider? * getDatabasePassword() - use { } within the if-else - similar to readPasswordFromStore * isMasterKeyPersisted() - properties.getProperty(MASTER_KEY_PERSISTED, "true"); - this will return true if the property does not exist. Is this by design? * getMasterKeyLocation - Can you check the default value for MASTER_KEY_LOCATION - I was not sure if {code}properties.getProperty(SRVR_KSTR_DIR_KEY + File.separator + MASTER_KEY_FILENAME_DEFAULT, SRVR_KSTR_DIR_DEFAULT + File.separator + MASTER_KEY_FILENAME_DEFAULT){code} produces the desired default. *ambari-server/src/main/java/org/apache/ambari/server/controller/AmbariServer.java* * does not seem to have any changes. *ambari-server/src/main/java/org/apache/ambari/server/security/CertificateManager.java* * no change here either *ambari-server/src/main/python/ambari-server.py* * Could we have used "update_ambari_properties" instead of defining "update_properties" TBD: ambari-server/src/main/java/org/apache/ambari/server/security/encryption	0	0	https://issues.apache.org/jira/browse/AMBARI-2174
19692	AMBARI-2174	isPersisted --> more of a convention. http://stackoverflow.com/questions/3874350/naming-conventions-for-java-methods-that-return-booleanno-question-mark We should follow it other places. MASTER_KEY_FILENAME_DEFAULT - Path is /var/lib/ambari-server/keys/master (Maybe we should be able to take full path as param) should loadCredentialProvider() use a lock around instantiating the provider - Unnecessary since only read operation isMasterKeyPersisted() - Yes, this makes sure we look for the file and load it. If credential not in store we use what is in the file. Could we have used "update_ambari_properties" instead of defining "update_properties"? Cannot reuse, "update_ambari_properties" updates from old file to new file, "update_properties" is from live map to file.	1	1	https://issues.apache.org/jira/browse/AMBARI-2174
19693	AMBARI-2130	The patch looks good on a cursory look. Does mvn test pass on trunk with this patch?	1	1	https://issues.apache.org/jira/browse/AMBARI-2130
19694	AMBARI-2130	[~tmckay] was just reviewing the patch. Sorry about the delay. There are some changes in the patch for hadoop 2.0 stack: eq: {code} src/test/resources/stacks/HDP/2.0.1/services/MAPREDUCE2/configuration/container-executor.cfg {code} Is this accidental?	1	1	https://issues.apache.org/jira/browse/AMBARI-2130
19695	AMBARI-2130	[~erinaboyd] The patch looks good. Can you please post the result of mvn clean test on trunk?	1	1	https://issues.apache.org/jira/browse/AMBARI-2130
19696	AMBARI-2130	Erin, Looks like some issues with the test case. Can you please open a jira for this and if possible provide a patch so that it works in your environment? Since the patch doesnt change python code Ill go aehad and commit thsi.	1	1	https://issues.apache.org/jira/browse/AMBARI-2130

#	Requirement	Questions	Answered	Successfully Answered	URL
19697	AMBARI-2130	[~erinaboyd] just ran mvn test on the latest patch. {code} Results : Failed tests: testGetStackVersions(org.apache.ambari.server.controller.AmbariManagementControllerTest): expected:<3> but was:<5> Tests run: 1084, Failures: 1, Errors: 0, Skipped: 2 [INFO] ---------- [INFO] Reactor Summary: [INFO] {code} Can you please look into that?	1	1	https://issues.apache.org/jira/browse/AMBARI-2130
19698	AMBARI-2128	Can this be closed as 'fixed'? Any known gaps?	0	0	https://issues.apache.org/jira/browse/AMBARI-2128
19699	AMBARI-2118	Hello [~croberts], here's my feedback so far. Sorry for the delay: * /app/assets/data/wizard/deploy/5_hosts/poll_6.json: Please don't include the changes to this file, as this is used for testing scenarios where the server returns a response for the start API. It might be a good idea to introduce another set of mock data for which the step for start API is empty. * /app/controllers/wizard/step8_controller.js: {code} + case 'HCFS': + return {config: {'global': 'version1', 'core-site': 'version1'}}; {code} The above does not apply because getConfigForService function has been removed (this function is no longer used). * /app/controllers/wizard/step9_controller.js: This needs further review. This step needs to handle various edge cases (retry upon installation failure, page refreshes, etc). * /app/data/config_mapping.js: fs.default.name definition is duplicated. We should remove the redundant entry and manage the value in code (e.g., hdfs://namenode_host:8020 for HDFS and whatever is appropriate for glusterfs). * /app/data/config_properties.js: the newly added glusterfs properties have the "filename" attribute. These should be removed. * /app/data/service_configs.js: For the newly added HCFS object, "configs" should be an empty array, and "sites" should be added (see others) * /app/config.js: Skip this file from the patch * /app/messages/js: installer.step4.multileDFS.popup.header and installer.step4.multileDFS.popup.body are misspelled.	1	1	https://issues.apache.org/jira/browse/AMBARI-2118
19700	AMBARI-2115	Ximo, Can you add more details on the implmenetation? Sorry about the delay on the review. Also, lets say we have a delete hosts api would that satisfy the requirement? Removing a host from ambari server would mean getting all the host componenents in MAINTENANCE state and then remove all the host components and the host. Would that be an acceptable soln?	1	0	https://issues.apache.org/jira/browse/AMBARI-2115
19701	AMBARI-1930	[~andres_danter] did you test this patch out? Where did you get the accumulo RPM?	1	0	https://issues.apache.org/jira/browse/AMBARI-1930
19702	AMBARI-1930	The patch did not work. It is posted here as reference. The RPM used was created from Accumulo 1.4.3, but I had to slightly modify the POM to create an x86_64 RPM instead of an amd64 RPM. Version 1.4.4 of Accumulo already includes that fix.	1	1	https://issues.apache.org/jira/browse/AMBARI-1930
19703	AMBARI-1865	[~ncole@hortonworks.com], is it OK to leave the data in service config mapping as is? Any harm if we delete those? One question related to the feature itself. Earlier, we used to keep the mapping between service_name and config_type (in serviceconfigmapping). It seems, that in Condor we do not keep this information any more. How do we know the mapping that MAPREDUCE is associated with mapred-site, global, core-site?	1	1	https://issues.apache.org/jira/browse/AMBARI-1865
19704	AMBARI-1865	Will this be safe? {code} SELECT MAX(config_tag) {code} Is tag supposed to be ordered such that max is the correct one. I thought one can put arbitrary tags when using API directly. Is MAX of the timestamp more reliable? Rest looks good.	1	1	https://issues.apache.org/jira/browse/AMBARI-1865
19705	AMBARI-1843	Hi, Can you pls commit this, My other changes are pending because of this . Many thanks	1	0	https://issues.apache.org/jira/browse/AMBARI-1843
19706	AMBARI-1843	+1 for the patch [~subin11]. Sorry abt the delay in getting to the patch. Can you please make sure mvn clean apache-rat:check passes with the patch?	0	0	https://issues.apache.org/jira/browse/AMBARI-1843
19707	AMBARI-1843	Committed this to trunk. Thanks a lot Subin! Ill fix the apache rat check in another jira since we have a couple of more issues on that.	0	0	https://issues.apache.org/jira/browse/AMBARI-1843
19708	AMBARI-1790	Nice one Trevor. Can you also provide a patch for trunk?	0	0	https://issues.apache.org/jira/browse/AMBARI-1790
19709	AMBARI-1783	OpenJDK is not a service. It is a package instead of a service. Firewall could be an example of isManageable service. How about simplify the cluster blue print, software stack definition and host manifest? A cleaner version would be: Cluster blueprint: {code}{ "version": "1.3.0", "cluster": { "name": "example", "services": [{ "name": "filesystem", "components": ["hdfs"], }, { "name": "firewall", "components": ["iptables"], "isManageable": "false" }], "applications": [{ "name": "my hadoop killer app", "components": ["text-analytics"], "isManageable": "false" }], "environment": { "fs.default.name": "hdfs://localhost:50070/", "HADOOP_NAMENODE_OPTS": { "value": "-Xmx5G", "file": "hadoop-env.sh" } } } {code} Software stack example: {code}{ "software": { "repository": [{ "name": "local", "url": "file:///image" }], "platform": "x86_64", "osType": "el6", "components": [{ "name":"glusterfs", "conflicts": ["hdfs"], "provides" : "filesystem", "packages": [{ "name":"glusterfs", "version":"4.3", "filename":"gluster-4.3.el6.x86_64.rpm", "hostGroup": ["node"] }] }, { "name":"hdfs", "conflicts": ["glusterfs"], "provides" : "filesystem", "packages": [{ "name":"java", "version":"1.7", "filename":"java-1.7.el6.x86_64.rpm" }, { "name":"hadoop", "version":"1.2.0", "filename":"hadoop-1.2.0.el6.x86_64.rpm", "dependency":"java" "hostGroup": ["namenode", "secondarynamenode", "datanode"] }] }, { "name":"text-analytics", "packages": [{ "name":"hadoop", "version":"1.2.0", "text-analytics-1.2.0.el6.x86_64.rpm" "dependency":"filesystem" }] }] } {code} Host manifest: {code}{ "namenode": ["nn1"], "secondarynamenode": ["snn1"], "datanode": ["dn1", "dn2", "dn3"] } {code} Notice in cluster blueprint, the new addition section of applications is designed to deploy application in the hadoop environment (HDFS) or virtual environment, there is no local copy stored in local file system. The running environment is a flatten list of key value pairs, this is designed to make it easier for dev ops to manage the running environment more easily. The additional overwrite to environment allows to set a specific value for a particular config file. Host manifest is a flatten list of hostGroups for user friendliness for dev ops. Conflict, and provide syntax can support swappable component for Hadoop software stack.	1	1	https://issues.apache.org/jira/browse/AMBARI-1783
19710	AMBARI-1783	I saw Blueprints on the Hortonworks site, but see the ticket hasn't been updated since last year. Is there a newer ticket that should be followed?	1	0	https://issues.apache.org/jira/browse/AMBARI-1783
19711	AMBARI-1779	Trevor I am not sure I understand, why would a default of python2 work?	1	0	https://issues.apache.org/jira/browse/AMBARI-1779
19712	AMBARI-1779	Mahadev (is that right?) are you referring to my comment on AMBARI-1721 above? my understanding: that issue reported by Jeff occurs on Fedora 18 (for example) where the python2.6 executable does not exist. Instead, you've got /usr/bin/python2 linked to /usr/bin/python2.7. This results in the 'Cannot run program "python2.6"' If maven scripts invoke python2 instead of explicitly invoking python2.x to run tests, etc, it should work regardless of whether a platform has pyhon2.6 or python2.7 installed (or even some other python2.x but that's unlikely, I think). The /usr/bin/python2 link should be standard for python installs. I'll spin up an F18 box and confirm. Are we on the same page? Have I missed something? Thanks.	1	1	https://issues.apache.org/jira/browse/AMBARI-1779

#	Requirement	Questions	Answered	Successfully Answered	URL
19713	AMBARI-1779	[~tmckay] the issue is that a pyton2 executable is not standard. It might work on fedora 18. Instead of python2 why dont we just use python and just make sure its > 2.0. The patch attached will cause issues on different platforms (lets say on a mac, there is no python2). Does that make sense? Also for changing the structure, we will also have to do some post install changes wherein the prior configs that have some paths in there that reference old 2.6/site packages are upgraded to new ones when you upgrade the agent.	1	0	https://issues.apache.org/jira/browse/AMBARI-1779
19714	AMBARI-1778	Nice Suggestions Amit. Want to take some of these up? We can open specific subtasks for some of them if you want to work on these. Thanks.	0	0	https://issues.apache.org/jira/browse/AMBARI-1778
19715	AMBARI-1777	Mostly nits, LGTM, +1. Remove line - +// throw new AmbariException("Delete hosts not supported"); What will happen if a DELETE is invoked for a host that does not exist under api/v1/hosts? Looks like it may get added into the okToRemove list and clusters.deleteHost will get called. It might be OK as the method can handle this scenario but wanted to bring it up. + if (!component.isClientComponent() && So we now allow removal of client components. In that case, should we modify the error text as well? Something like - _To remove master or slave components they must be ..._. We can also modify the comment in the code. + public void removeAllHost(long clusterId, String hostName) { Should we rename it to removeHost (no All)	1	1	https://issues.apache.org/jira/browse/AMBARI-1777
19716	AMBARI-1769	Subin, Can you attach it as a patch? I think best would be to create a separate module for this, something like ambari-client and put the code in src/main/python. This way we can also add some unit tests for this and include it in the builds. Does that make sense?	1	0	https://issues.apache.org/jira/browse/AMBARI-1769
19717	AMBARI-1769	Subin, One quick thing, can you please remove the @author tags? Since the code gets owned/rewritten by lot of other folks, we like to avoid the @author tags in the code. Is that ok?	1	0	https://issues.apache.org/jira/browse/AMBARI-1769
19718	AMBARI-1769	Sure not a problem. That was added automatically by the IDE i was using. Can you pls give me all the changes you need so that i can do it all at once and submit the new patch? Thanks Subin	1	1	https://issues.apache.org/jira/browse/AMBARI-1769
19719	AMBARI-1769	Subin, Still not done with the review. Will probably be done by EOD today. Would that work?	1	0	https://issues.apache.org/jira/browse/AMBARI-1769
19720	AMBARI-1769	Subin, Just got the time to take a look at your patch. Nice work (specifically on the py docs)! The patch looks good. Some minor nits: - as metioned earlier the @ author tags and the created tags should be removed - http_utils.py has ^M characters in the file, not sure why (might want to do dos2unix on the files you have) - http_utils does not have an apache header and also says method copied from Django. If you copied the method, do you know what license/copyright it is? - main.py has ^M characters as well, same with init.py, cluster.py, paths.py, utils.py - the main.py does ganglia.stop() - i think that might be an oversight on having that in the main.py. Also why do we need a main.py? - the most important one is that to make it part of the main codebase (am sure lot of people will want to use this), we really need to make sure its maintained, which means unit tests (atleast testing main features of the client library). Can you please add some? You can also use the mock libary in ambari-common for python. - Also what about adding some help/examples into the codebase as well which will help folks who want to use the library?	1	0	https://issues.apache.org/jira/browse/AMBARI-1769
19721	AMBARI-1769	*??as metioned earlier the @ author tags and the created tags should be removed?? bq. I will remove this in the next patch *??http_utils.py has ^M characters in the file, not sure why (might want to do dos2unix on the files you have)?? bq. I will remove this in the next patch.Will do a dos2unix before i ship the code. *??http_utils does not have an apache header and also says method copied from Django. If you copied the method, do you know what license/copyright it is?? bq. I plan to remove this code in the next patch. I have found another way to do this. *??main.py has ^M characters as well, same with init.py, cluster.py, paths.py, utils.py?? bq. I will remove this in the next patch.Will do a dos2unix before i ship the code. *??the main.py does ganglia.stop() - i think that might be an oversight on having that in the main.py. Also why do we need a main.py?? bq. main.py was meant as a tutorial for python library. I will rename this as an example . *??the most important one is that to make it part of the main codebase (am sure lot of people will want to use this), we really need to make sure its maintained, which means unit tests (atleast testing main features of the client library). Can you please add some? You can also use the mock libary in ambari-common for python.?? bq. I have put about 2 months into this. Currently I support a lot more features than the features in the code i have given to you.I wanted to give the code in stages so that it can be reviewed and checked in and also hope it would be useful for future committers to understand the code.My plan is to give skeleton code in Phase1. Phase2 will contain some modifications for phase 1 and addition of new features ans also meet the ambari coding standards. Phase 3 onwards i would add the junits and some refactoring.Phase4 will contain all the features currently supported by me.Each phase would be 1 week long. Are you ok with my plan? I have few junits in place But i plan add them in phase 2. are you ok with this? *??Also what about adding some help/examples into the codebase as well which will help folks who want to use the library?? bq. main.py was meant as a tutorial for python library. I will rename this as a example . Please let me know if you have any concerns? Can you also tell me how to checkin this code into main codebase TRUNk ? Thanks Subin	1	1	https://issues.apache.org/jira/browse/AMBARI-1769
19722	AMBARI-1769	bq. I have put about 2 months into this. Currently I support a lot more features than the features in the code i have given to you.I wanted to give the code in stages so that it can be reviewed and checked in and also hope it would be useful for future committers to understand the code.My plan is to give skeleton code in Phase1. Phase2 will contain some modifications for phase 1 and addition of new features ans also meet the ambari coding standards. Phase 3 onwards i would add the junits and some refactoring.Phase4 will contain all the features currently supported by me.Each phase would be 1 week long. Subin, I can understand what you want to do here. I dont mind you creating another jira where units tests would be coming later. Can you create one related jira for unit tests and mention what/how you plan to add the unit tests? As long as we know when/how/what, I think we hsould be ok. Also, for committing to trunk, only committers can commit. Usually contributors become committers after working sometime with the community and help reviewing patches and getting other folks onboard. I am assuming this might be your first time on ASF. Take a look at: http://www.apache.org/foundation/how-it-works.html on how it works.	1	0	https://issues.apache.org/jira/browse/AMBARI-1769
19723	AMBARI-1769	Mahadev What i would like is to make my code meet your standards. I understand the meritocracy of the ASF. All i wanted to know was whether i would be allowed to do it or whether i would pass the code to you. both the options are fine for me.All izz well. So, I will add the junits and also take care of ^M ,main.py and the http_utils issue . I will give the phase1 patch to you.And you can do the further review and checkin the code. ok?	1	1	https://issues.apache.org/jira/browse/AMBARI-1769
19724	AMBARI-1769	Subin, Looks like we might need to rethink where to add the rpm or add install instructions to be able to install it as a python module so that it can be imported without changing python path. What do you think?	0	0	https://issues.apache.org/jira/browse/AMBARI-1769

#	Requirement	Questions	Answered	Successfully Answered	URL
19725	AMBARI-1769	+1 for the patch. I just committed this to trunk. Thanks Subin. Since this is just beginning of the work we will be doing for the client I just committed this as it is. Here are some improvements we need to make for this: - Add unit tests (whcih you have a jira for). - Make the client component part of the top level pom.xml (this will help in running unit tests/builds from a single place) - add some documentation on how/what of ambari-client that can be published with a release Subin, Can you please open a jira for the last 2. Again thanks a lot for your patience and contribution.	0	0	https://issues.apache.org/jira/browse/AMBARI-1769
19726	AMBARI-1760	Configurable on the command line during server setup? Modifying ambari.properties is less obvious and harder to automate.	1	0	https://issues.apache.org/jira/browse/AMBARI-1760
19727	AMBARI-1760	Good point Michael. Want to provide a patch? Should be simple to add a property to the config file while doing setup. Take a look at ambari-server.py.	1	1	https://issues.apache.org/jira/browse/AMBARI-1760
19728	AMBARI-1760	Michael, Thanks for the patch. The patch looks good. A couple of comments: - the indentation is 2 space (no tabs) for python code as well - can you provide a unit test for this as well? Take a look at the python unit tests in the ambari-server (src/test/python/TestAmbaryServer.py) to see how to add one. Other than that this looks good. Thanks a lot!	1	1	https://issues.apache.org/jira/browse/AMBARI-1760
19729	AMBARI-1760	mock is new to me but want to learn how to use it, any good resources to look at? Should I add the updated patch and unit test at the same time?	1	0	https://issues.apache.org/jira/browse/AMBARI-1760
19730	AMBARI-1703	+1 for the patch. One minor nit, can we write a unit test for this so that this does not break again?	1	0	https://issues.apache.org/jira/browse/AMBARI-1703
19731	AMBARI-1703	+1 on the patch. Thanks a lot Nate!	1	0	https://issues.apache.org/jira/browse/AMBARI-1703
19732	AMBARI-1639	My feedback: * This is a great feature to add and it was done really well! * app/controllers/global/update_controller.js's updateServiceMetric method used to ask for specific metrics (/services?fields=components/ServiceComponentInfo,components/host_components,components/host_components/HostRoles,components/host_components/metrics/jvm/memHeapUsedM,components/host_components/metrics/jvm/memHeapCommittedM,components/host_components/metrics/mapred/jobtracker/trackers_decommissioned). The patch is querying for /services?fields=components/ServiceComponentInfo,components/host_components,components/host_components/HostRoles,components/host_components/metrics. It seems that doing so would significantly increase the payload since it will grab all metrics available... We should explicitly list out all the metrics that we are asking for, like we were doing before. * app/models/service/hdfs.js and app/models/service/mapreduce.js have unnecessary comment tags. Please remove. * app/templates/main/service.hbs has an empty id attribute, like so: {code} {code} Please remove this id attribute. * ambari/ambari-web/app/messages.js - Change dashboard.customerWidgets.* to dashboard.widgets.*. Prefixing with "customer" does not make much sense. Change 'dashboard.customerWidgets': 'Customer Widgets' to 'dashboard.widgets': 'Cluster Status and Metrics'. * Replace "views/main/dashboard/customer_widgets/*" with "views/main/dashboard/widgets/*". Also, rename the directory views/main/dashboard/customer_widgets to widgets. * vendor/scripts/jquery.ui.sortable.js {{code}} this._delay(function() { if(counter == this.counter) this.refreshPositions(!hardRefresh); //Precompute after each DOM insertion, NOT on mousemove }); {{code}} The above code was commented out. Can you add a comment in this file as to why you did that? * General comment about code format. It does not conform to the guideline here: https://cwiki.apache.org/confluence/display/AMBARI/Coding+Guidelines+for+Ambari. Please correct it (but only for the lines you changed). * Certain features do not work in testMode (AMBARI-2306 was filed to track this).	0	0	https://issues.apache.org/jira/browse/AMBARI-1639
19733	AMBARI-1639	[~xiwang0309], can you address the comment regarding the change in vendor/scripts/jquery.ui.sortable.js?	0	0	https://issues.apache.org/jira/browse/AMBARI-1639
19734	AMBARI-1574	Hi Gunnar, thanks for your suggestion. What do you specifically have in mind? Creation of HDFS users and their home directories, quota management, etc?	1	1	https://issues.apache.org/jira/browse/AMBARI-1574
19735	AMBARI-1566	Good one Gunnar. Do you want to take a stab at parameterizing the JAVA_OPTS in the configurations for these services?	1	1	https://issues.apache.org/jira/browse/AMBARI-1566
19736	AMBARI-1562	Can you tell me where you saw this error message? Maybe in /var/log/ambari-server/ambari-server.log or /var/log/ambari-agent/ambari-agent.log? I suspect that I have met this issue but I'm not sure. Sometimes my install process goes wrong but I didn't found any exception in this log files.	1	0	https://issues.apache.org/jira/browse/AMBARI-1562
19737	AMBARI-1551	Jos, Here are a couple of pointers on how the monitoring of hadoop processes are done on the Ambari Agent side. Take a look at ambari-agent/src/main/python/ambari_agent/LiveStatus.py. What we can do is make LiveStatus.py pluggable and also as of now the server send commands to check for status of each component of the host. We can use that command to just send a list of components to check for on that host and send it just once in an hour or so. We can then do active live status checks on the agents itself using a pluggable module which could be a wrapper around daemon tools. Does that sound ok?	1	0	https://issues.apache.org/jira/browse/AMBARI-1551
19738	AMBARI-1551	[~mahadev] [~josb] Is this JIRA still relevant? If not can you resolve the JIRA?	0	0	https://issues.apache.org/jira/browse/AMBARI-1551
19739	AMBARI-1522	http://opentsdb.net/ Does it make sense to consider these features we'd like to build in conjunction with something like opentsb? It's not quite dumping things in HDFS, but opentsb is essentially built on hbase. There's a strange meta-loop going here.....where the logs of the platform itself should really be run through said platform in the same flow as one of the use cases the platform services.	0	0	https://issues.apache.org/jira/browse/AMBARI-1522
19740	AMBARI-1522	[~sposetti] Hello Jeff, I posted the first design doc draft for review. Could you please help take a look and provide feedback? Thank you. Di	1	0	https://issues.apache.org/jira/browse/AMBARI-1522
19741	AMBARI-1518	Does ambari server support this feature? As far as I know Ambari UI do not support it, but maybe the API yes. Do anybody tested it?	1	1	https://issues.apache.org/jira/browse/AMBARI-1518
19742	AMBARI-1518	Halacs, The server does have provision to support it. Its been a while since I tested it. You might find some issues when trying to do it via the api's. Want to give it a try and file issues based on your trials?	0	0	https://issues.apache.org/jira/browse/AMBARI-1518
19743	AMBARI-1518	Is that patch in mainline now?	0	0	https://issues.apache.org/jira/browse/AMBARI-1518
19744	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19745	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19746	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19747	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19748	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19749	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19750	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19751	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518

#	Requirement	Questions	Answered	Successfully Answered	URL
19752	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19753	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19754	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19755	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19756	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19757	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19758	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19759	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19760	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19761	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19762	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19763	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19764	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19765	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19766	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19767	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19768	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19769	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19770	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19771	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19772	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19773	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19774	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19775	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19776	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19777	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19778	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19779	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19780	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19781	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19782	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19783	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19784	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19785	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19786	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19787	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19788	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19789	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19790	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19791	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19792	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19793	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19794	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19795	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19796	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19797	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19798	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19799	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19800	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19801	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19802	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19803	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19804	AMBARI-1518	Where can we download the patch ? thanks ~~	1	0	https://issues.apache.org/jira/browse/AMBARI-1518
19805	AMBARI-1518	Has there been any progress on this issue?	0	0	https://issues.apache.org/jira/browse/AMBARI-1518
19806	AMBARI-1502	Nate, when I do a GET on /api/v1/clusters/clustername, I see 'desired_configs' as {}. Should it not list the latest desired type-tag pairs, just like the 'configurations' in that response.	1	0	https://issues.apache.org/jira/browse/AMBARI-1502
19807	AMBARI-1427	Hello, Is this feature available only at individual services' level? Is it possible to increase the time range for the graphs shown on the Ambari dashboard also?	0	0	https://issues.apache.org/jira/browse/AMBARI-1427
19808	AMBARI-1406	[~jspeidel], what is the syntax for specifying a query string in the body?	1	1	https://issues.apache.org/jira/browse/AMBARI-1406

#	Requirement	Questions	Answered	Successfully Answered	URL
19809	AMBARI-1369	Arpit, I am seeing the following tests fail. Can you please take a look? {noformat} == FAIL: test_check_iptables (TestAmbaryServer.TestAmbariServer) ---------- Traceback (most recent call last): File "/Users/mahadev/workspace/ambari-workspace/ambari-git/ambari-common/src/test/python/mock/mock.py", line 1199, in patched return func(*args, **keywargs) File "/Users/mahadev/workspace/ambari-workspace/ambari-git/ambari-server/src/test/python/TestAmbaryServer.py", line 660, in test_check_iptables self.assertEqual(0, rcode) AssertionError: 0 != (0, '') == FAIL: test_get_choice_string_input (TestAmbaryServer.TestAmbariServer) ---------- Traceback (most recent call last): File "/Users/mahadev/workspace/ambari-workspace/ambari-git/ambari-common/src/test/python/mock/mock.py", line 1199, in patched return func(*args, **keywargs) File "/Users/mahadev/workspace/ambari-workspace/ambari-git/ambari-server/src/test/python/TestAmbaryServer.py", line 97, in test_get_choice_string_input self.assertEquals(result, True) AssertionError: 'default blablabla' != True == FAIL: test_get_validated_string_input (TestAmbaryServer.TestAmbariServer) ---------- Traceback (most recent call last): File "/Users/mahadev/workspace/ambari-workspace/ambari-git/ambari-common/src/test/python/mock/mock.py", line 1199, in patched return func(*args, **keywargs) File "/Users/mahadev/workspace/ambari-workspace/ambari-git/ambari-server/src/test/python/TestAmbaryServer.py", line 155, in test_get_validated_string_input self.assertEquals(get_pass_method.return_value, result) AssertionError: 'dfdsfdsfds' != 'default_pass' {noformat}	1	1	https://issues.apache.org/jira/browse/AMBARI-1369
19810	AMBARI-1260	LGTM. Few nits. AbstractProviderModule.java Line 320 - Should we move the port numbers to their own constants - e.g. DEFAULT_NAME_NODE_PORT - if you take this change there are 4 other values in this section. Line 381 - Should the log contain the clusterName, serviceName, etc. for easy diagnosis? Line 385 - Should we define a constant string for "version1" and use that here. Line 387 - The loop will continue to overwrite versionTag for every non-null configs. Did you intend to break after the first non-null value? Line 396 - We can move the var declaration to just before the if (configResources != null) check at 412 Line 409 - Should the log contain the clusterName, configType, configTag, etc. for easy diagnosis? LIne 413 - It also seems that the loop will overwrite values for keys if there are multiple config resources. If we expect only one config resource should we assert that? JMXPropertyProviderTest.java Are there unused imports? Do we need additional unit tests that actually read from mock config files to test the service and configResource readers used by AbstractProviderModule? Any need for negative test cases? (unless they already exist)	1	1	https://issues.apache.org/jira/browse/AMBARI-1260
19811	AMBARI-1260	Line 387 - The loop will continue to overwrite versionTag for every non-null configs. Did you intend to break after the first non-null value? Expecting to get the last version tag id	0	0	https://issues.apache.org/jira/browse/AMBARI-1260
19812	AMBARI-1260	LGTM. JMXHostProvider Line 42 - Nit, Get the port for the given cluster name and component name Line 44 - Nit, add clusterName to the javadoc. AbstractProviderModule Line 191 - Nit, !currVersion.equals(oldVersion) will remove the need for null check Line 197 - Hashmap is not synchronized. Aren't we allowing multiple threads to do put which may result in inconsistency Line 417 - We return "version1" as the default value. Is there always a config with version name "version1"?	1	1	https://issues.apache.org/jira/browse/AMBARI-1260
19813	AMBARI-1255	The patch looks good Nate. Can we just verify that the agents register with the configured hostname at the server (some end to end testing), wherein we use a configurable hostname for an agent and then make sure it registers with the same hostname at the server by looking at <ambari-server>:8080/api/v1/hosts?	1	1	https://issues.apache.org/jira/browse/AMBARI-1255
19814	AMBARI-1242	Is there something I can do to help with this ? Would love to see this happen soon.	1	0	https://issues.apache.org/jira/browse/AMBARI-1242
19815	AMBARI-1241	Has this been put on hold for now? It looks like there hasn't been any activity for over a year.	0	0	https://issues.apache.org/jira/browse/AMBARI-1241
19816	AMBARI-1201	Added new patch that covers missing rpm (thanks Mahadev!). When merging, use AMBARI-1201.patch.1	1	1	https://issues.apache.org/jira/browse/AMBARI-1201
19817	AMBARI-1065	Like to hear how to apply the patch. Been looking for awhile for this feature to come to ambari. Does only apply to only version 1.2?	1	0	https://issues.apache.org/jira/browse/AMBARI-1065
19818	AMBARI-1065	Thank you Gunturi, Does this include the Hortonworks 1.2 release ? -Charles ---------- Sent using BlackBerry	1	0	https://issues.apache.org/jira/browse/AMBARI-1065
19819	AMBARI-1065	Thank you Gunturi, Does this include the Hortonworks 1.2 release ? -Charles ---------- Sent using BlackBerry	1	0	https://issues.apache.org/jira/browse/AMBARI-1065
19820	AMBARI-1065	Thank you Gunturi, Does this include the Hortonworks 1.2 release ? -Charles ---------- Sent using BlackBerry	1	0	https://issues.apache.org/jira/browse/AMBARI-1065
19821	AMBARI-1065	Thank you Gunturi, Does this include the Hortonworks 1.2 release ? -Charles ---------- Sent using BlackBerry	1	0	https://issues.apache.org/jira/browse/AMBARI-1065
19822	AMBARI-1044	I don't see any related exceptions in the server log which means that either its not attempting to get the metrics for this host or they are just not being set on the host resource. I think that I see what is happening. One of the arguments that can be specified for the rrd query is the Ganglia cluster (HDPHBaseMaster, HDPJobTracker, HDPNameNode or HDPSlaves). The question is, for a host level query which Ganglia cluster should we specify? Its hard to say since a host isn't necessarily with any of the services related to those clusters... or maybe more than one. It turns out it doesn't really matter. In this case I can see the system level rrd files that we use for host level metrics for ip-10-224-42-108.ec2.internal under any of the Ganglia cluster folders. For example ... {code} [root@ip-10-40-91-121 rrds]# ls ./HDPHBaseMaster/ip-10-224-42-108.ec2.internal boottime.rrd bytes_out.rrd cpu_idle.rrd cpu_num.rrd cpu_system.rrd cpu_wio.rrd disk_total.rrd load_five.rrd mem_buffers.rrd mem_free.rrd mem_total.rrd pkts_in.rrd proc_run.rrd swap_free.rrd bytes_in.rrd cpu_aidle.rrd cpu_nice.rrd cpu_speed.rrd cpu_user.rrd disk_free.rrd load_fifteen.rrd load_one.rrd mem_cached.rrd mem_shared.rrd part_max_used.rrd pkts_out.rrd proc_total.rrd swap_total.rrd ... [root@ip-10-40-91-121 rrds]# ls HDPNameNode/ip-10-224-42-108.ec2.internal boottime.rrd bytes_out.rrd cpu_idle.rrd cpu_num.rrd cpu_system.rrd cpu_wio.rrd disk_total.rrd load_five.rrd mem_buffers.rrd mem_free.rrd mem_total.rrd pkts_in.rrd proc_run.rrd swap_free.rrd bytes_in.rrd cpu_aidle.rrd cpu_nice.rrd cpu_speed.rrd cpu_user.rrd disk_free.rrd load_fifteen.rrd load_one.rrd mem_cached.rrd mem_shared.rrd part_max_used.rrd pkts_out.rrd proc_total.rrd swap_total.rrd {code} The approach that I've been using is to look through the host components for the host that we are interested in and try to map one of its component names back to a Ganglia cluster. In this case it looks like the host with the missing metrics is not associated with any component that would map back given the mapping method that I am using. Given what I am currently seeing with the system level metrics, I think that it would be safe to simply use HDPSlaves as the Ganglia cluster for host level queries.	1	0	https://issues.apache.org/jira/browse/AMBARI-1044
19823	AMBARI-147	Good stuff. Could you please update the patch so that it applies in the current trunk.	1	0	https://issues.apache.org/jira/browse/AMBARI-147

#	Requirement	Questions	Answered	Successfully Answered	URL
19824	AMBARI-25	Should "Property" be "Properties" and "ConfigurationCategory" be "ConfigurationCategories". I would prefer plural. In Blueprint "Components" element it is referred as "Properties". Also what is a significance of "@" prefix? Should it be only be applied to keys in Configuration element or other places too, like in "Components" or "Roles" etc. ?	1	0	https://issues.apache.org/jira/browse/AMBARI-25
19825	AMBARI-25	In xml, it now looks like: {noformat} <Configuration> <ConfigurationCategory Name="mapred-site"> <Property name="mapred.system.dir" value="/mapred/mapredsystem" force="false"/> <Property name="mapred.local.dir" value="/tmp/mapred" force="false"/> </ConfigurationCategory> <ConfigurationCategory Name="hdfs-site"> <Property name="dfs.name.dir" value="/tmp/namenode" force="false"/> <Property name="dfs.data.dir" value="/tmp/datanode" force="false"/> </ConfigurationCategory> </Configuration> {noformat} in JSON, it looks like: {noformat} "Configuration": { "ConfigurationCategory": [{"@Name": "mapred-site", "Property": [{"@name": "mapred.system.dir", "@value": "/mapred/mapredsystem", "@force": "false"}, {"@name": "mapred.local.dir", "@value": "/tmp/mapred", "@force": "false"}]}, {"@Name": "hdfs-site", "Property": [{"@name": "dfs.name.dir", "@value": "/tmp/namenode", "@force": "false"}, {"@name": "dfs.data.dir", "@value": "/tmp/datanode", "@force": "false"}]}]} {noformat} In answer to the questions, the '@' is saying that it represents an attribute in XML. Is there a way to make the JSON use 'Properties' while the XML uses 'Property'?	0	0	https://issues.apache.org/jira/browse/AMBARI-25
19826	ACE-592	Your first few questions I don't quite understand what you are asking for (I'm not a native speaker, so that could well be my fault). I am guessing you are interested to learn what exactly the rules are for naming features and distributions? They should be given unique names. Can you try explaining your question again? To answer your question about the deployment not working, what set of bundles did you try to deploy? Are you sure those bundles all resolve (and start)? One thing to diagnose this issue is to add the "-v" option to the ace-launcher, which will make it a lot more verbose. Usually then you will be able to figure out what is going wrong. If not, paste the output here so we can have a look. If possible, then also share the bundles you were trying to deploy. Your other questions: (1) Again, not sure what you are asking here. In the ACE UI you can click on a target to see what distributions, features and artifacts it is linked to (or on any other entity). You cannot start/stop/uninstall a bundle on the target, but you can remove a bundle from a feature and store that new configuration to uninstall it from all targets that had that feature installed. (2) The ace-launcher basically is an instance of Apache Felix with a management agent. So it already is integrated. What are you looking for? Using the management agent with a specific version of Felix, or?	1	0	https://issues.apache.org/jira/browse/ACE-592
19827	ACE-592	Thank you very much for your reply! I installed ACE （apache-ace-2.1.0） ON C:/ace I installed felix ON C:/felix My goal is to copy the file: test2-1.0.0.201605281422.jar from C:\ace\server-allinone\store to C:\felix\plugins Is this possible? The realization of the process is? Thank you very much for the reply.	0	0	https://issues.apache.org/jira/browse/ACE-592
19828	ACE-532	As the bot already noticed, I've added a somewhat bigger PR that not only replaces the org.apache.ace.http.listener but also binds all Servlets registered by ACE to a custom ServletContext. This makes it possible to prefix all ACE servlets (and resources) with a single configuration option. This is a lot easier (and less error prone) than the per Servlet endpoint configuration. For example the RepositoryServlet, by default this is registered at {{http://my.host/repository}} by changing the context root for the {{AceServletContextHelper}} to {{/ace}} the repository servlet will be available at {{http://my.host/ace/repository}} (all other servlets will move to {{http://my.host/ace/.....}} aswel. Authentication is also handled in the {{AceServletContextHelper}} so enabling authentication is also in a single place, and the duplicated authentication code is removed from the Servlets. What isn't fixed (yet) in this PR is are the "where to find" configurations, there are a lot of configurations for example to configure the url of the repository servlet. Now you can move the complete application to another we can change this we just need a place to configure the location of the server and obr. What would be the best way to get this url to all components that need it? Should / can we use the {{Discovery}} interface for this?	1	0	https://issues.apache.org/jira/browse/ACE-532
19829	ACE-530	Can you elaborate a bit on what the issue exactly is? Are you referring to the web UI of Ace, or some other project? I'm not aware of any concept named "service desk", so I'm not sure where I could help you with.	1	1	https://issues.apache.org/jira/browse/ACE-530
19830	ACE-522	Applied in rev #1726789. Thanks for the patch!	0	0	https://issues.apache.org/jira/browse/ACE-522
19831	ACE-513	Given that gogo has alias support and it is just one client.. would it not be best to remove all gogo registration and command short-hands from the ace.client.workspace bundle altogether? I think we should prefer meaningful names on the core interfaces and move gogo support into the ace.gogo bundle.	1	1	https://issues.apache.org/jira/browse/ACE-513
19832	ACE-503	Committed the patch. Thanks for your contribution [~brampouwelse]!	0	0	https://issues.apache.org/jira/browse/ACE-503
19833	ACE-483	This patch was generated with, cd trunk svn up svn diff >ace-1.patch On 24 Sept 2014 at ~10:00 UTC.	1	1	https://issues.apache.org/jira/browse/ACE-483
19834	ACE-483	I modified the patch a bit: I left out the generics changes for Java 8/9 because we currently build for Java 7 and I feel we need to fix that in the whole project, and as a separate "patch". Changed the shutdown hook to explicitly wait for the framework to shut down. Not doing that made the JVM shutdown way before bundles were actually stopped. It will wait forever: if the framework hangs we depend on the script "killing the JVM after a while". Some typographical changes to the script. I like the way you split up the target into a "run-develop" and "run-target". I am even thinking about not providing the "run-develop" option in our binary distribution anymore, but I'd like to get some community feedback about that.	1	1	https://issues.apache.org/jira/browse/ACE-483
19835	ACE-459	I agree that it makes sense to display that information in the LogViewer. As you already mention, the audit log does log the BSN and version of the bundle, but only when it is installed. That was done on purpose, because after that event, BSN and version cannot change anymore, so we did not want to keep wasting bytes and just identify the installed bundle by its ID. That being said, I would rather see a solution where the BSN and version are looked up in the viewer, instead of adding them to each event. Does that make sense?	1	1	https://issues.apache.org/jira/browse/ACE-459
19836	ACE-459	So why not make a component that interprets the log data and provides the information you need? That could be a component that can be used anywhere (not bound to client or server) and provides a richer view.	0	0	https://issues.apache.org/jira/browse/ACE-459
19837	ACE-456	I'm all for promoting reuse, but have my doubts about creating a separate project just for this functionality. Isn't there another project that can expose this functionality as additional bundle? Why not create sub-bundles in o.a.a.client.rest?	0	0	https://issues.apache.org/jira/browse/ACE-456

#	Requirement	Questions	Answered	Successfully Answered	URL
19838	ACE-456	Jan Willem, if there's a better place to put it, that's perfectly fine with me, I can move it and resubmit a patch if you want. However, I don't think it should be a sub-bundle of o.a.a.client.rest specifically. I don't think it has anything to do with the REST-part of the ACE client. The way I look at it, the Workspace is generic client functionality that can be used (as a service) by different "manifestations" of the client (i.e. REST, script-based, even the Vaadin client could potentially re-use it to maintain its "dirty" state on the server side). Perhaps it could be combined with for example the o.a.a.client.automation services into o.a.a.client.services or something similar? What do you think?	0	0	https://issues.apache.org/jira/browse/ACE-456
19839	AMQ-6701	Thanks for the reply. Attaching the jstack. Do you think this is due to GC?	1	0	https://issues.apache.org/jira/browse/AMQ-6701
19840	AMQ-6695	thanks for the patch. Changes like this can be tricky (as in it will change behaviour for folks) - that code has not changed since 2007. I think there should always be a copy of the message because sending does change the metadata and of course the usage. From the comments on setCopyMessage it seems the intent is to set the forward destination on the copy, maybe that flags should just do that and there is always a copy. Can you add some unit tests to show the problem and that can protect this changes into the future. Also verify the few unit tests that exercise the mirror queue There is a related issue https://issues.apache.org/jira/browse/AMQ-3669 that seems to be in the same area and has a relevant test that can give inspiration.	1	1	https://issues.apache.org/jira/browse/AMQ-6695
19841	AMQ-6630	Commit 4bf3152fa99bd97f04845e69b8d442eeb2a39308 in activemq's branch refs/heads/master from [~tabish121] [https://git-wip-us.apache.org/repos/asf?p=activemq.git;h=4bf3152] https://issues.apache.org/jira/browse/AMQ-6630 Update to QpidJMS 0.22.0 and Netty 4.1.9.Final	1	1	https://issues.apache.org/jira/browse/AMQ-6630
19842	AMQNET-557	OK thanks. I tried to contact “sbohlen” …but was not successful yet. Btw. The STOMP library is also not released as normal download. The latest available version is 1.5.4. See: http://activemq.apache.org/nms/stomp-downloads.html Can you release the library as “normal” download? (Maybe this download is the source of the nugget package, and “sbohlen” cannot create a new version of the nugget package without this download.)	1	1	https://issues.apache.org/jira/browse/AMQNET-557
19843	AMQNET-517	Commit 1724526 from [~jgomes] in branch 'MSMQ/trunk' [https://svn.apache.org/r1724526] Applied patch from Stephane Ramet to implement QueueBrowser and DeleteDestination. Thanks Stephane! Fixes [AMQNET-517]. (See https://issues.apache.org/jira/browse/AMQNET-517)	1	1	https://issues.apache.org/jira/browse/AMQNET-517
19844	AMQNET-490	Thanks for the quick reply. I did some searching around but don't see anything on changing the destination name passing. Do you have an example you can point me to on how to do this for the clients?	1	0	https://issues.apache.org/jira/browse/AMQNET-490
19845	AMQNET-490	Is there a later released version of the Stomp Client that I could be using where it is configurable via the wire format. I am not married to version 1.5.4, that just seemed like the latest official version so that is what I went with.	1	0	https://issues.apache.org/jira/browse/AMQNET-490
19846	AMQNET-470	This doesn't really seem like the best idea. Can you elaborate on the use case a bit. The send method for a persistent message returns only after receiving an ack from the broker, where a send for a non-persistent message returns immediately. You can force sync sends always via configuration or async sends always as well.	1	1	https://issues.apache.org/jira/browse/AMQNET-470
19847	AMQNET-470	I use the send method in IMessageProducer, Will it have the delegate to invoke to? B.T.W - The exception type is NotSupportedException :)	1	0	https://issues.apache.org/jira/browse/AMQNET-470
19848	AMQNET-470	Timothy - Should I add the delegate (optional) parameter to the IMessageProducer interface too?	1	0	https://issues.apache.org/jira/browse/AMQNET-470
19849	AMQNET-470	Hi. This is too much changes for a little modification, I wonder if I can do something like this : * Add to the "ITransport" (From ActiveMQ dll) an OnewaySent event (That will raise every time that a Command is sent from the TcpTransport) * Implement the OnewaySent event for all the decorator ITransports easily. * In my project : Decorate the IMessage and use it from the that is raised. What do you think about it?	1	0	https://issues.apache.org/jira/browse/AMQNET-470
19850	AMQNET-464	isn't a big task for one developer ?	1	0	https://issues.apache.org/jira/browse/AMQNET-464
19851	AMQNET-464	Would it make more sense to create a new REST provider?	0	0	https://issues.apache.org/jira/browse/AMQNET-464
19852	AMQNET-454	Patch doesn't seem to want to apply for me, is your repo up to date with the latest changes?	1	1	https://issues.apache.org/jira/browse/AMQNET-454
19853	AMQNET-454	Hi Jim, Thanks for your attention to my checkins. Maybe you could advise me on how best to proceed with my unmanaged content. h4. Populating \lib with Apache Qpid vendor files There is an issue with copying the Apache Qpid vendor files from the local .nant repo into lib\net-2.0 or lib\net-4.0. For net-2.0 or net-4.0 the Qpid vendor files are delivered in debug and release sets. h5. Distinguishing debug from release * qpid library file names have a "d" suffix. * boost library file names have an extra "-gd" embedded in them. * The library binding org.apache.qpid.messaging.dll files have the *identical name* for debug and release. h5. Distingushing net-2.0 from net-4.0 * qpid library file names are the same net-2.0 and net-4.0 * boost library file names have the vc compiler version embedded in them: ** net-2.0 identified by "-vc90" ** net-4.0 identified by "-vc100" * The library binding org.apache.qpid.messaging.dll files have the *identical name* for net-2.0 and net-4.0 It would be best if the net-2.0/net-4.0 debug/release files could be kept in four directories in \lib. Without debug and release folders this isn't possible and the debug and release files wind up on top of each other. This isn't a huge issue until the release version of org.apache.qpid.messaging.dll overwrites the debug version. As soon as the debug version of code is rebuilt in the \build folder the release version of the dll is copied into the build debug folder and nothing works after that. The .NET runtime error messages are not helpful in diagnosing the issue: it says it can't load a dll but it does not explain why. h4. Populating \build with libraries from \lib Scripts must account for the file naming variations so that all the desired files wind up in the proper places in build. It does not hurt to have the unmanaged debug and release libraries both installed in the build folder but care must be taken not to install the wrong build variant of org.apache.qpid.messaging.dll. h4. How to proceed I notice that products like NUnit dodge this issue by having the same file in the \debug and \release in the .nant local repo. I believe that org.apache.qpid could do the same thing. If the same file set were in two places in the local repo then there would never be a conflict in the \lib folders nor in the \build folders. Release versions of products built on top of debug versions of the qpid unmanged code would still work correctly. There is still a conflict between the net-2.0 and the net-4.0 versions of boost file names but that can be solved by adding the other names to the file copy list added in http://svn.apache.org/r1617620 I'll give this a try before my next commit and see how it goes. Please reply if you have ideas on how to proceed.	1	1	https://issues.apache.org/jira/browse/AMQNET-454
19854	APLO-329	Can you post the stack trace? Can you try running on latest version of Apollo?	1	0	https://issues.apache.org/jira/browse/APLO-329
19855	APLO-327	Should be good to go on my private branch. Hiram, if you wanna review it, please go ahead. Since this requires Java 7, which Apollo release should we plan to include this jira?	0	0	https://issues.apache.org/jira/browse/APLO-327

#	Requirement	Questions	Answered	Successfully Answered	URL
19856	APLO-325	So I have a patch to fix this. It relies on setting the underlying websocket connection max text message size: WebSocket.Connection#setMaxTextMessageSize There is also a method for #setMaxBinaryMessageSize Think it's the same in ActiveMQ 5.x (see the WSTransportServer.java#doStart) Just getting the last of the unit tests to pass (having some issues with the WSS test of course, unrelated to this change... will hopefully have it ready by tomorrow and attach the patch to this JIRA if you want to review). This will all change when the Jetty 9 API is used, but will be similar (there is a websocket client factory on which you can make the same text/binary message size settings). Should we also make the change in the client JS libs?	1	0	https://issues.apache.org/jira/browse/APLO-325
19857	APLO-325	Cool! Oliver, can you give the latest snapshot a try? I've also updated the docs: http://activemq.apache.org/apollo/versions/99-trunk-SNAPSHOT/website/documentation/user-manual.html#WebSocket_Transports	0	0	https://issues.apache.org/jira/browse/APLO-325
19858	APLO-322	Apollo supports master/slave fault tolerance, but if you're looking for "network of brokers" the way activemq 5.x does it, that doesn't exist in Apollo *yet*. What sort of clustering are you looking for? What's your usecase?	1	1	https://issues.apache.org/jira/browse/APLO-322
19859	APLO-322	The usecase is that we use the apollo broker to realize the interaction in class.In addition ,is the total number of the apollo broker no limit ?	1	0	https://issues.apache.org/jira/browse/APLO-322
19860	ARTEMIS-1302	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1416 Is this just logging ?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1302
19861	ARTEMIS-1276	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1407 it's probably an issue on the test. should I merge this already?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1276
19862	ARTEMIS-1276	Github user jdanekrh commented on the issue: https://github.com/apache/activemq-artemis/pull/1407 > messages only expire on client on activemq5 while it's on server & client on artemis Did you mean expire on _server_ on activemq5? If it expired on client, how would it then get to DLQ queue on server? --- I think get it now. When the test receives for the first time (first `for` loop), the client gets first 10 messages sent, which are those with short expiration time. One of them is delivered ok, the 9 remaining expire on the client while it is in sleep(). That means these 9 never get into the DLQ, because for that they would have to expire on the server. Which explains the `actual: 90, expected: 99` result when it is then receiving from DLQ. I did not know about this difference between 5 and Artemis. From user point of view, the fact that prefetch changes behavior (and is not just performance optimization) and that it can cause what can be thought of as a message loss (these 9 messages have irrevocably disappeared, right?) could be considered a bug, I think. I guess there is nothing to be done about it here, so removing [wip].	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1276
19863	ARTEMIS-1276	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1407 >>> If it expired on client, how would it then get to DLQ queue on server? There's a special protocol on OpenWire to tell the server the message expired. On Artemis we would just reject the message and let the server deal with it again. Having said that, perhaps there's something wrong on treating that message..so it's worth to check if it's a bug on that area.. or just a test issue. I didn't look into the details specifically to this test BTW: I was just mentioning that I had cases like this where I had to tweak the test because of that semantic change.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1276
19864	ARTEMIS-1270	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1385 just one request here.. don't leave master on an intermediate state this week.. I want to release before end of this week / early next week as I sent the heads up earlier. We can merge this now if we're sure it's ready.. otherwise lets wait till next week.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1270
19865	ARTEMIS-1270	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1385 @clebertsuconic i think I've address'ed now 1-9 on the list Dan gave. Only bit left is adding/updating artemis LICENSE/NOTICE. Whats the process on this. Is it manual or is it automated, who needs to verify it? After this the only last bit is aligning up the login (so don't have to bypass)	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1270
19866	ARTEMIS-1270	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1385 From here: http://www.apache.org/legal/resolved.html#category-a We just have to list the components and the license and point to the location of the dependency in LICENSE It's seems there is no need to touch NOTICE file for these licenses. Can this be confirmed?	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1270
19867	ARTEMIS-1270	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1385 You think it would be good on the release ? How much work to get the authentication working ? Or ready to merge but not release. We can talk on Monday. I'm not doing much coding this weekend.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1270
19868	ARTEMIS-1270	I find that it is a great feature that the Artemis server distribution doesn't include a advanced web console today. Not having a web console bundled reduces the distribution size, configuration and build complexity, limits the dependencies, reduces vulnerability exposure and limits the number of security updates required. I understand that a web console like Hawtio brings convenience and functionality to Artemis, and I understand that it can be attractive to new users. Would it be possible to offer two separate artifacts for the binary server distribution: One for just the Artemis Server binary as it is shipped today, and one that bundles both Artemis Server and the customized HawtIO console? I apologize if I've missed any relevant previous discussion about this on the mailing lists, I haven't had time to follow Artemis development closely for some time.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1270
19869	ARTEMIS-1270	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1385 @michaelandrepearce small tweak, can you add a link to the console on the main index? When you go to localhost you see this list: - API - User Manual - Hacking Guide - Examples - Apache ActiveMQ Artemis Website it would be useful having the webconsole as the first link here.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1270
19870	ARTEMIS-1265	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1380#discussion_r125054051 --- Diff: docs/hacking-guide/en/code-coverage-report.md --- @@ -0,0 +1,37 @@ +# Code coverage report + +## Getting JaCoCo exec files + ++Before you can generate code coverage report by JaCoCo tool, +you need to get data about what lines of code were executed --- End diff -- You're not telling how to generate this information ^^ how to generate the data about what lines where executed?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1265
19871	ARTEMIS-1263	What kind of back-end are you using for your user & role information? Can you not simply put all the required users into a single role and set the permission for that single role? Even if you're using a properties file it should be pretty simple.	0	0	https://issues.apache.org/jira/browse/ARTEMIS-1263

#	Requirement	Questions	Answered	Successfully Answered	URL
19872	ARTEMIS-1252	Github user roelens8 commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1366#discussion_r124116668 --- Diff: artemis-commons/src/main/java/org/apache/activemq/artemis/utils/PasswordMaskingUtil.java --- @@ -108,22 +110,29 @@ public static HashProcessor getHashProcessor() { // semi colons String[] parts = codecDesc.split(";"); - if (parts.length < 1) throw new ActiveMQException(ActiveMQExceptionType.ILLEGAL_STATE, "Invalid PasswordCodec value: " + codecDesc); final String codecClassName = parts[0]; // load class - codecInstance = AccessController.doPrivileged(new PrivilegedAction<SensitiveDataCodec<String>>() { - @Override - public SensitiveDataCodec<String> run() { - ClassLoader loader = Thread.currentThread().getContextClassLoader(); + codecInstance = AccessController.doPrivileged((PrivilegedAction<SensitiveDataCodec<String>>) () -> { + ServiceLoader<SensitiveDataCodec> serviceLoader = ServiceLoader.load(SensitiveDataCodec.class, PasswordMaskingUtil.class.getClassLoader()); --- End diff -- Added a couple unit tests. Are they sufficient enough? Also, should the documentation be updated to mention this new feature?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1252
19873	ARTEMIS-1252	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1366#discussion_r124569686 --- Diff: artemis-commons/src/main/java/org/apache/activemq/artemis/utils/PasswordMaskingUtil.java --- @@ -108,22 +110,29 @@ public static HashProcessor getHashProcessor() { // semi colons String[] parts = codecDesc.split(";"); - if (parts.length < 1) throw new ActiveMQException(ActiveMQExceptionType.ILLEGAL_STATE, "Invalid PasswordCodec value: " + codecDesc); final String codecClassName = parts[0]; // load class - codecInstance = AccessController.doPrivileged(new PrivilegedAction<SensitiveDataCodec<String>>() { - @Override - public SensitiveDataCodec<String> run() { - ClassLoader loader = Thread.currentThread().getContextClassLoader(); + codecInstance = AccessController.doPrivileged((PrivilegedAction<SensitiveDataCodec<String>>) () -> { + ServiceLoader<SensitiveDataCodec> serviceLoader = ServiceLoader.load(SensitiveDataCodec.class, PasswordMaskingUtil.class.getClassLoader()); --- End diff -- I think so... This is an area where @jbertram has done more work than me.. I was going to wait on him.. I work with him and I believe he's out of the office today. is that ok?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1252
19874	ARTEMIS-1252	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1366#discussion_r124801911 --- Diff: artemis-commons/src/main/java/org/apache/activemq/artemis/utils/PasswordMaskingUtil.java --- @@ -108,22 +109,34 @@ public static HashProcessor getHashProcessor() { // semi colons String[] parts = codecDesc.split(";"); - if (parts.length < 1) throw new ActiveMQException(ActiveMQExceptionType.ILLEGAL_STATE, "Invalid PasswordCodec value: " + codecDesc); final String codecClassName = parts[0]; // load class - codecInstance = AccessController.doPrivileged(new PrivilegedAction<SensitiveDataCodec<String>>() { - @Override - public SensitiveDataCodec<String> run() { - ClassLoader loader = Thread.currentThread().getContextClassLoader(); + codecInstance = AccessController.doPrivileged((PrivilegedAction<SensitiveDataCodec<String>>) () -> { + ServiceLoader<SensitiveDataCodec> serviceLoader = ServiceLoader.load(SensitiveDataCodec.class, PasswordMaskingUtil.class.getClassLoader()); + try { + // Service load the codec, if a service is available + for (SensitiveDataCodec<String> codec : serviceLoader) { + if ((codec.getClass().getCanonicalName()).equals(codecClassName)) { + return codec.getClass().newInstance(); + } + } + // If a service is not available, load the codec class using this class's class loader + return (SensitiveDataCodec<String>) PasswordMaskingUtil.class.getClassLoader().loadClass(codecClassName).newInstance(); + } catch (Exception e) { + try { + Class<?> clazz = loader.loadClass(codecClassName); + return (SensitiveDataCodec<String>) clazz.newInstance(); + } catch (Exception e) { + throw ActiveMQUtilBundle.BUNDLE.errorCreatingCodec(e, codecClassName); + // If service loading failed, load the codec class using this class's class loader + return (SensitiveDataCodec<String>) PasswordMaskingUtil.class.getClassLoader().loadClass(codecClassName).newInstance(); --- End diff -- I was just about to hit enter on merging when I realized something... isn't this a mistake? Shouldn't this be using Thread.currentThread().classLoader()? The previous call (inside the try... catch) was using this class classLoader.. if you catch an exception and try, I would expect to do somethng different.. like using the classLoader, not doing the same operation. it seems a typo to me... can you check?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1252
19875	ARTEMIS-1248	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1361 @franz1981 I see.. you could still use a ThreadLocal there.. right? also: Why don't you use a PooledBuffer on the write?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1248
19876	ARTEMIS-1248	Github user franz1981 commented on the issue: https://github.com/apache/activemq-artemis/pull/1361 @michaelandrepearce @clebertsuconic You're right: these are the results of a test writing 100 Pages of 10MB with messages of 100 bytes. Just 2 notes: - I've used the Artemis JVM default settings: -XX:+PrintClassHistogram -XX:+UseG1GC -XX:+AggressiveOpts -XX:+UseFastAccessorMethods -Xms512M -Xmx2G. With smaller HEAP the effect is even more pronounced. - I've forced a GC just at the beginning to have the HEAP clean. Old version: ![unpooled](https://user-images.githubusercontent.com/13125299/27471968-ee3e8036-57f9-11e7-962f-c8c5ae1f1276.png) ![unpooled_gc](https://user-images.githubusercontent.com/13125299/27471969-ee3f12a8-57f9-11e7-96f1-88103375c08d.png) New version: ![pooled](https://user-images.githubusercontent.com/13125299/27472080-59ff4fda-57fa-11e7-9645-34aac2446281.png) ![pooled_gc](https://user-images.githubusercontent.com/13125299/27472079-59fcd516-57fa-11e7-8d81-0d566e041c58.png) wdyt?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1248
19877	ARTEMIS-1235	Github user mtaylor commented on the issue: https://github.com/apache/activemq-artemis/pull/1345 @michaelandrepearce Sure. I get the configuration management tool use case. Two comments: 1. Does it make sense to enable this at the address-settings level vs just a global wide param? 2. If this is switched on, do we want to disable other means of managing addresses? We have discussed in the past, making the XML configuration the single source of truth for queues and addresses. I think the original reason for not supporting address/queue deletes via configuration was to avoid message loss. In this scenario, I think dropping a queue with messages is appropriate. We just need to make that clear in the documentation. This topic has cropped up in the past, do you guys have any input? @andytaylor @clebertsuconic @jbertram. On the PR. It looks good, I did a few extra manual tests with wildcard addresses, it works as expected Thanks	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1235

#	Requirement	Questions	Answered	Successfully Answered	URL
19878	ARTEMIS-1235	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1345 @andytaylor re: > My opinion removing queues should never be done automatically even if that is a configured flag. We actually would want to automatically remove the queue still when using tools like puppet/chef. Use case we need this for is that a durable subscription is being decom'd the app would need to be deployed (or underplayed) with the subscription gone, after which we remove the queue for that subscription. In this case there maybe build up of messages as message produce flow continues and other subscriptions (queues) would still exist. As such we'd want the queue removed still even if outstanding messages to deliver. Re: Removing bindings journal where would you put auto-created queues in this scenario, as we would only control in puppet persistent/durable queues typically. as such couldn't have these written to this file else puppet/chef would re-push the config file and remove the auto-created queues created for non-durable topic consumers.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1235
19879	ARTEMIS-1235	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1345 @michaelandrepearce I'm just back from a 2 weeks vacations.. and I'm still slow... So, help me to understand.. you added a DeletionPolicy on Address Settings. if you set it to FORCE, you will delete the queue if it disappeared from the configuration? my only concern about this is on subscriptions... that is queues created from client APIs. This solves the case for the CLI creating queues.. but what about the client calling .createQueue from a client session or something equivalent from AMQP Clients, OpenWire.. etc? you wouldn't set it to FORCE on such addresses?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1235
19880	ARTEMIS-1235	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1345#discussion_r123343028 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/ActiveMQServerImpl.java --- @@ -2313,6 +2318,53 @@ private void deploySecurityFromConfiguration() { } } + private void undeployAddressesAndQueueNotInConfiguration() throws Exception { + undeployAddressesAndQueueNotInConfiguration(configuration); + } + + private void undeployAddressesAndQueueNotInConfiguration(Configuration configuration) throws Exception { + Set<String> addressesInConfig = configuration.getAddressConfigurations().stream() + .map(CoreAddressConfiguration::getName) + .collect(Collectors.toSet()); + + Set<String> queuesInConfig = configuration.getAddressConfigurations().stream() + .map(CoreAddressConfiguration::getQueueConfigurations) + .flatMap(List::stream).map(CoreQueueConfiguration::getName) + .collect(Collectors.toSet()); + + for (SimpleString addressName : listAddressNames()) { + AddressSettings addressSettings = getAddressSettingsRepository().getMatch(addressName.toString()); + + if (!addressesInConfig.contains(addressName.toString()) && addressSettings.getConfigDeleteAddresses() == DeletionPolicy.FORCE) { + for (Queue queue : listQueues(addressName)) { + ActiveMQServerLogger.LOGGER.undeployQueue(queue.getName()); + queue.deleteQueue(true); + } + ActiveMQServerLogger.LOGGER.undeployAddress(addressName); + removeAddressInfo(addressName, null); + } else if (addressSettings.getConfigDeleteQueues() == DeletionPolicy.FORCE) { + for (Queue queue : listConfiguredQueues(addressName)) { + if (!queuesInConfig.contains(queue.getName().toString())) { --- End diff -- a simplification wouldn't work for you? if auto-created.. it means.. it was meant to be created.. no delete... if you really need the semantic differentiation.. we would need to find a better name.. NO_QUEUES doesn't translate IMO.	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1235
19881	ARTEMIS-1235	Github user clebertsuconic commented on a diff in the pull request: https://github.com/apache/activemq-artemis/pull/1345#discussion_r123536227 --- Diff: artemis-server/src/main/java/org/apache/activemq/artemis/core/server/impl/ActiveMQServerImpl.java --- @@ -2313,6 +2318,53 @@ private void deploySecurityFromConfiguration() { } } + private void undeployAddressesAndQueueNotInConfiguration() throws Exception { + undeployAddressesAndQueueNotInConfiguration(configuration); + } + + private void undeployAddressesAndQueueNotInConfiguration(Configuration configuration) throws Exception { + Set<String> addressesInConfig = configuration.getAddressConfigurations().stream() + .map(CoreAddressConfiguration::getName) + .collect(Collectors.toSet()); + + Set<String> queuesInConfig = configuration.getAddressConfigurations().stream() + .map(CoreAddressConfiguration::getQueueConfigurations) + .flatMap(List::stream).map(CoreQueueConfiguration::getName) + .collect(Collectors.toSet()); + + for (SimpleString addressName : listAddressNames()) { + AddressSettings addressSettings = getAddressSettingsRepository().getMatch(addressName.toString()); + + if (!addressesInConfig.contains(addressName.toString()) && addressSettings.getConfigDeleteAddresses() == DeletionPolicy.FORCE) { + for (Queue queue : listQueues(addressName)) { + ActiveMQServerLogger.LOGGER.undeployQueue(queue.getName()); + queue.deleteQueue(true); + } + ActiveMQServerLogger.LOGGER.undeployAddress(addressName); + removeAddressInfo(addressName, null); + } else if (addressSettings.getConfigDeleteQueues() == DeletionPolicy.FORCE) { + for (Queue queue : listConfiguredQueues(addressName)) { + if (!queuesInConfig.contains(queue.getName().toString())) { --- End diff -- On that case this may be good as is to me then... I missed that.. let me read the code again... @michaelandrepearce sorry.. still slow after a Caribiean vacation :)	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1235
19882	ARTEMIS-1230	Can you elaborate a bit more about what exactly you're looking for? I've used a Maven BOM, but never created one.	1	0	https://issues.apache.org/jira/browse/ARTEMIS-1230
19883	ARTEMIS-1214	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1324 I'm a bit concerned with isolation of concern on Page.. Why don't you bring your change to SequentialFile and use a method on SequentialFile to read it... if reading files through mmap is that faster than NIO, lets always read using mmap... even the load on libaio would speed up, as it's using (as far as I remember). is this the same on any env? lets just use that way always?	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1214
19884	ARTEMIS-1214	Github user clebertsuconic commented on the issue: https://github.com/apache/activemq-artemis/pull/1324 @franz1981 there are only two parts... paging/reading.. journal/loading.. journal/printdata (which is using the load.. so I'm not counting this one).. only two	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1214

#	Requirement	Questions	Answered	Successfully Answered	URL
19885	ARTEMIS-1214	Github user michaelandrepearce commented on the issue: https://github.com/apache/activemq-artemis/pull/1324 @franz1981 @clebertsuconic been going through this the last two days, one bit I'm unclear on is where or what logic flow are we intending this to improve things for? Journal Paging I ask as if its for the latter (paging), and we're doing it to use page caching as another layer of keeping as much as possible in memory (though not on heap), then i don't think AIO is an option there, as such should probably just scope the work to NIO and MAPPED. If its for Journal, i thought the journal was write heavy, as the read is only on broker restart, where do we really care, for the extra complexity it will cause if we had to deal with AIO? (i may have miss understood the read path for Journal btw) Cheers Mike	1	1	https://issues.apache.org/jira/browse/ARTEMIS-1214
19886	AMQCPP-606	Yep, aware of Qpid's library. However, that doesn't provide a standard programming API. How would adding AMQP to CMS be different than the current protocol support for openwire and stomp in CMS?	1	0	https://issues.apache.org/jira/browse/AMQCPP-606
19887	AMQCPP-564	I implemented following methods in the Threading namespace: // void Threading::releaseCurrentThreadHandle() { ThreadHandle* self = (ThreadHandle*)PlatformThread::getTlsValue(library->selfKey); if (self != NULL) { detachFromCurrentThread(self); } } // void Threading::detachFromCurrentThread(ThreadHandle* self) { PlatformThread::lockMutex(library->globalLock); // Destroy given Foreign Thread Facade that was created during runtime. std::vector<Thread*>::iterator iter = library->osThreads.begin(); bool isFound = false; for (; iter != library->osThreads.end(); ++iter) { if (self->parent == *iter) { isFound = true; break; } } if (isFound) { PlatformThread::setTlsValue(library->threadKey, NULL); PlatformThread::setTlsValue(library->selfKey, NULL); // Ensure all of this thread's local values are purged. threadExitTlsCleanup(self); // Destroy OS thread including self thread handle. delete *iter; // Remove thread form the global list. library->osThreads.erase(iter); } PlatformThread::unlockMutex(library->globalLock); } Both methods are static, the method void Threading::releaseCurrentThreadHandle() is public and I am using it in the destructor of my thread object. It seems that the code runs correctly at least on Windows. I will continue testing it on Linux and Solaris. I encountered no memory leaks, but it would be nice, if someone could check the correctness of the implementation. If it's correct, could you add these or similar methods to a future release of the library?	1	1	https://issues.apache.org/jira/browse/AMQCPP-564
19888	AMQCPP-543	Hi, I compare with the java client, and uploaded a patch that I think does pretty much what the java do. The java assume the latest version of the protocol by default and when it receives a wireFormatInfo command from the broker, saves the protocol version. this is done in an AtomicInteger. ActiveMQMessageProducer.java {code} // Enable producer window flow control if protocol > 3 and the window // size > 0 if (session.connection.getProtocolVersion() >= 3 && this.info.getWindowSize() > 0) { producerWindow = new MemoryUsage("Producer Window: " + producerId); producerWindow.setExecutor(session.getConnectionExecutor()); producerWindow.setLimit(this.info.getWindowSize()); producerWindow.start(); } {code} Can you explain why the comment says protocol > 3, but the code actually does >= 3? Perhaps you can also review the patch?	1	0	https://issues.apache.org/jira/browse/AMQCPP-543
19889	AMQCPP-517	Hi. Timothy Bish How are you? Could you help me?? I need to building ActiveMQ-CPP with C++. Do you have a clue or tip??? Tks. VIlmar Resende from Brazil 2013/10/9 Timothy Bish (JIRA) <jira@apache.org> -- *ViLMeR_**ReSeNDe* Vida é rsrsrsrsrs (risadas mesmo, não RS)	1	0	https://issues.apache.org/jira/browse/AMQCPP-517
19890	OPENWIRE-5	Can you update this issue, at least with ETA please?	1	0	https://issues.apache.org/jira/browse/OPENWIRE-5
19891	ADDR-38	Dasarath, which repo do you want me to put this in? please specify the full patch of the directory. -- dims	1	1	https://issues.apache.org/jira/browse/ADDR-38
19892	ADDR-38	Quoting "Davanum Srinivas (JIRA)" <jira@apache.org>: How about-- http://mirrors.ibiblio.org/maven/addressing/jars thanks, --dasarath page=com.atlassian.jira.plugin.system.issuetabpanels:comment- tabpanel#action_12474583	1	1	https://issues.apache.org/jira/browse/ADDR-38
19893	ADDR-21	Sorry I meant to create this issue in the Tuscany project instead, for some reason JIRA didn't point me to Tuscany as usual this time... Could you please cancel this issue. Thanks.	0	0	https://issues.apache.org/jira/browse/ADDR-21
19894	ADDR-20	Sorry I meant to create this issue in the Tuscany project instead, for some reason JIRA didn't point me to Tuscany as usual this time... Could you please cancel this issue. Thanks.	0	0	https://issues.apache.org/jira/browse/ADDR-20
19895	ADDR-10	Hi, the content of the patch looks good, but it's not in the correct format. Can you please regenerate it using "svn diff" (as described at http://logging.apache.org/log4net/contributing.html#patch)? Thanks, Ian	1	0	https://issues.apache.org/jira/browse/ADDR-10
19896	ADDR-10	Thomas, Do you really want us to check in the code with your copyright headers? Is it a bid deal to remove them? i think i'd prefer a clean patch w/o the copyright headers. thanks, dims	1	1	https://issues.apache.org/jira/browse/ADDR-10
19897	ADDR-10	Our policy is to require the copyright headers, as otherwise people won't know that we have copyright in the files. Without the extra header, it would say that Apache owns the complete copyright, whereas we are only granting you a license, as I understand it. I'm not a legal expert, and I'll have to check with the people here who are authorised to make these decisions. What is the normal way of handling this situation? Would it be acceptable if we just used a shorter header?	1	0	https://issues.apache.org/jira/browse/ADDR-10
19898	ADDR-10	I've had a talk with the managers here. I think we could shrink the header down a lot, if that's the issue. The main question they have is: who owns the copyright to code we submit? Presumably, we still own it (and are free to relicense it and use the code in other projects). If we don't have a header at all, the file would say that Apache owns it. Of course, this particular patch is rather trivial. However, we want to sort out a policy for future contributions, since we use several Apache libraries in our software (see http://www.gria.org/ for an LGPL example) and we have other patches we could submit too. Thanks,	1	0	https://issues.apache.org/jira/browse/ADDR-10
19899	AIRFLOW-1414	[~bolke] Any though on that this (the corresponding PR is https://github.com/apache/incubator-airflow/pull/2444) We are trying to switch from Oozie to Airflow at Booking.com and would really need this feature (or maybe there's already a way to achieve this) for us to go forward with the migration.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-1414
19900	AIRFLOW-1414	[~bolke] thanks for your feedback, would you have time to have a quick look at this ticket too: https://issues.apache.org/jira/browse/AIRFLOW-1398?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-1414
19901	AIRFLOW-1401	I'm particularly interested in thoughts from [~alexvanboxel], [~wwlian], and [~fenglu]	1	1	https://issues.apache.org/jira/browse/AIRFLOW-1401
19902	AIRFLOW-1401	SGTM, maybe we can include other standardization efforts on GCP operators as well (see https://issues.apache.org/jira/browse/AIRFLOW-820)? Let's also make sure to do it in a backwards-compatible way and not break any existing DAGs that depend on an older version of these operators.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-1401

#	Requirement	Questions	Answered	Successfully Answered	URL
19903	AIRFLOW-1398	[~bolke] Any though on that this (the corresponding PR is https://github.com/apache/incubator-airflow/pull/2431) We are trying to switch from Oozie to Airflow at Booking.com and would really need this feature (or maybe there's already a way to achieve this) for us to go forward with the migration.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-1398
19904	AIRFLOW-1398	[~bolke] I implemented your suggestion, would you mind reviewing?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-1398
19905	AIRFLOW-1325	To be clear, are you planning to REMOVE support for S3/GCS? We are heavy users of GCS for logging, and this would be a show stopper for us.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-1325
19906	AIRFLOW-1310	It was exactly the same sort of thinking on my side. Good point about the logs. It was sort of design decision as we capture all logs centrally in ELK so we had no need to stream the logs to airflow as such. It is not hard to add though. Also regarding your comment that it would be more powerful if you can supply a pod definition directly - thats another decision to hide the pod definition as such to make deployment of simple jobs simpler. I have a version where you can give a pod definition instead in case you need to do something complex, but for most jobs we have we just have an image (often with default command) and we dont want to force people to learn yet another library just to run a simple command. Do you mind collaborating to get out something nice and useful to everybody?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-1310
19907	AIRFLOW-1019	[~aoen], changing this to blocker for 1.8.1. LMK if you don't think we should block the release on it.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-1019
19908	AIRFLOW-983	Was this confirmed yet?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-983
19909	AIRFLOW-983	It seems like the commit was reverted but I'm not sure if testing was done: https://github.com/apache/incubator-airflow/pull/2195? If not I can try to take a look soon but just wanted to confirm first [~bolke]	0	0	https://issues.apache.org/jira/browse/AIRFLOW-983
19910	AIRFLOW-979	Please review my Pull Request #2149: https://github.com/apache/incubator-airflow/pull/2149	1	1	https://issues.apache.org/jira/browse/AIRFLOW-979
19911	AIRFLOW-932	verify_integrity in models.py is where you would like to look at. As mentioned it sets a state to REMOVED if a task instance does not have a corresponding task in a DAG. It is there for auditing and versioning reasons. Does it affect running anything? A changed DAG cannot have a dependency on a non existent task. Ie. please explain why this is a blocker.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-932
19912	AIRFLOW-932	Command: {quote}airflow backfill teztg -t create_and_transfer_roster -s 2017-06-14 -e 2017-06-14{quote} DAG: {code} import airflow from airflow.executors import SequentialExecutor from airflow.operators.dummy_operator import DummyOperator from airflow.operators.python_operator import PythonOperator from airflow.operators.subdag_operator import SubDagOperator from airflow.models import DAG from datetime import datetime, timedelta default_args = { 'owner': 'airflow', 'start_date': datetime(2016, 3, 15), } dag = DAG('teztg', default_args=default_args) subdag = DAG('teztg.create_and_transfer_roster', default_args=default_args) subdag_operator = SubDagOperator(task_id='create_and_transfer_roster', subdag=subdag, executor=SequentialExecutor(), dag=dag) def create_roster(ds, macros, **kwargs): pass PythonOperator(task_id='create_roster', python_callable=create_roster, dag=subdag, provide_context=True) first_task = DummyOperator(task_id="first_task", dag=dag) {code} Expected: first_task TI does not change (stays as none) create_and_transfer_roster TI is run Actual: first_task TI gets its state set to removed create_and_transfer_roster TI is not run, backfill errors with airflow.exceptions.AirflowException: Task first_task not found [~bolke] if you think these are related to your changes (addition of task.REMOVED state/backfill changes) do you mind taking a look? Otherwise let me know and I can do so.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-932
19913	AIRFLOW-932	select * from task_instance; +----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+ | task_id | dag_id | execution_date | start_date | end_date | duration | state | try_number | hostname | unixname | job_id | pool | queue | priority_weight | operator | queued_dttm | pid | +----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+ | create_and_transfer_roster | teztg | 2016-06-14 00:00:00.000000 | 2017-03-02 08:58:35.867219 | 2017-03-02 08:58:46.868470 | 11.0013 | success | 1 | bolkes-macbook-pro.local | bolke | 2 | NULL | default | 1 | SubDagOperator | NULL | 46856 | | create_roster | teztg.create_and_transfer_roster | 2016-06-14 00:00:00.000000 | 2017-03-02 08:58:42.681159 | 2017-03-02 08:58:42.702009 | 0.02085 | success | 1 | bolkes-macbook-pro.local | bolke | 4 | NULL | default | 1 | PythonOperator | NULL | 46862 | +----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+----------+ {code} {code} [2017-03-02 08:58:29,039] {jobs.py:1756} INFO - Checking run <DagRun teztg @ 2016-06-14 00:00:00: backfill_2016-06-14T00:00:00, externally triggered: False> [2017-03-02 08:58:29,044] {models.py:1125} INFO - Dependencies all met for <TaskInstance: teztg.create_and_transfer_roster 2016-06-14 00:00:00 [None]> [2017-03-02 08:58:29,047] {base_executor.py:50} INFO - Adding to queue: airflow run teztg create_and_transfer_roster 2016-06-14T00:00:00 --local -sd DAGS_FOLDER/task_removed_bug.py [2017-03-02 08:58:33,985] {jobs.py:1949} INFO - [backfill progress] | dag run 1 of 1 | tasks waiting: 1 | succeeded: 0 | kicked_off: 1 | failed: 0 | skipped: 0 | deadlocked: 0 [2017-03-02 08:58:33,987] {local_executor.py:45} INFO - LocalWorker running airflow run teztg create_and_transfer_roster 2016-06-14T00:00:00 --local -sd DAGS_FOLDER/task_removed_bug.py [2017-03-02 08:58:33,989] {models.py:1125} INFO - Dependencies all met for <TaskInstance: teztg.create_and_transfer_roster 2016-06-14 00:00:00 [scheduled]> /Users/bolke/Documents/dev/airflow_env/lib/python2.7/site-packages/airflow-1.9.0.dev0+apache.incubating-py2.7.egg/airflow/configuration.py:548: DeprecationWarning: This method will be removed in future versions. Use 'parser.read_file()' instead. self.readfp(StringIO.StringIO(string)) [2017-03-02 08:58:34,561] {__init__.py:56} INFO - Using executor LocalExecutor Logging into: /Users/bolke/airflow/logs/teztg/create_and_transfer_roster/2016-06-14T00:00:00 [2017-03-02 08:58:38,987] {jobs.py:1949} INFO - [backfill progress] | dag run 1 of 1 | tasks waiting: 1 | succeeded: 0 | kicked_off: 1 | failed: 0 | skipped: 0 | deadlocked: 0 [2017-03-02 08:58:38,990] {models.py:1119} INFO - Dependencies not met for <TaskInstance: teztg.create_and_transfer_roster 2016-06-14 00:00:00 [running]>, dependency 'Task Instance State' FAILED: Task is in the 'running' state which is not a valid state for execution. The task must be cleared in order to be run. [2017-03-02 08:58:38,990] {models.py:1119} INFO - Dependencies not met for <TaskInstance: teztg.create_and_transfer_roster 2016-06-14 00:00:00 [running]>, dependency 'Task Instance Not Already Running' FAILED: Task is already running, it started on 2017-03-02 08:58:35.867219. [2017-03-02 08:58:43,994] {jobs.py:1949} INFO - [backfill progress] | dag run 1 of 1 | tasks waiting: 1 | succeeded: 0 | kicked_off: 1 | failed: 0 | skipped: 0 | deadlocked: 0 [2017-03-02 08:58:43,997] {models.py:1119} INFO - Dependencies not met for <TaskInstance: teztg.create_and_transfer_roster 2016-06-14 00:00:00 [running]>, dependency 'Task Instance State' FAILED: Task is in the 'running' state which is not a valid state for execution. The task must be cleared in order to be run. [2017-03-02 08:58:43,997] {models.py:1119} INFO - Dependencies not met for <TaskInstance:	1	1	https://issues.apache.org/jira/browse/AIRFLOW-932
19914	AIRFLOW-932	Ah I think I know why it's not reproducing for you. Can you try letting the scheduler run a dagrun, and then backfill the dagrun once it completes (same DAG)? so e.g. let it run for one dagrun via scheduler (might need to unpause it if it's paused) then airflow backfill teztg -t create_and_transfer_roster -s 2016-03-15 -e 2016-03-15	1	1	https://issues.apache.org/jira/browse/AIRFLOW-932

#	Requirement	Questions	Answered	Successfully Answered	URL
19915	AIRFLOW-932	Issue resides in cli.py and will only happen when a specific task is used: {code} if args.task_regex: dag = dag.sub_dag(task_regex=args.task_regex, include_upstream=not args.ignore_dependencies) {code} This creates a subset of the tasks from a dag_run, with the same name as the original. Hence it will set a task to removed if you verify the integrity of a dag run, cause that will consider the tree of the original (whole). get_task_instances picks up all instances, including "removed' ones from the original (whole) dag and this is not filtered in the backfill. Hence the lists mismatch and an AirflowException is thrown. The quick and dirty fix is to not mark removed if run from backfill and filter the list of "get_task_instances". However the functionality if sub_dag is awkward imho - and seems not to be tested, and might need a real fix. What do you think?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-932
19916	AIRFLOW-932	PR is available with unit tests [~aoen]. Please review	1	1	https://issues.apache.org/jira/browse/AIRFLOW-932
19917	AIRFLOW-931	[~bolke] I added a bunch of manual logging to verify this. What seems to happen (in my simple gist example) is DAG starts with 4 tasks. The multiple scheduler processes kick them off (the jobs.py:1095) because at that time they are all runnable (since none are yet running). That models.py:1291 does another check to make sure it is actually runnable (2 of them are, but then due to the dag concurrency of 2, the third and fourth one aren't). The 2 tasks stay marked as queued but never actually get picked up by the scheduler again (unless you restart it). Is there some other state that that models.py:1291 check could move the tasks to if that last minute check found they aren't actually runnable? (I found by making them State.NONE it worked, but that seems hackish as it keeps bouncing back and forth between QUEUED and NONE until it can actually run)	1	0	https://issues.apache.org/jira/browse/AIRFLOW-931
19918	AIRFLOW-920	[~aoen] It does propagate in its current form: if you mark success a "subdag" operator from its parent dag it will mark success all tasks part of that subdag. My PR just verifies the tree of the dag to make sure dag runs are created for any missing subdag dag runs. I consider this to be something of a hack as it is in my opinion the wrong place to do it (the webserver). BTW: Also the unit tests confirm propagation. So are we talking about the same thing?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-920
19919	AIRFLOW-908	Can i take this feature ?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-908
19920	AIRFLOW-908	[~aoen] Would you give me an instruction ?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-908
19921	AIRFLOW-891	Any reason why you are suggesting only adding the day like that? It isn't very intuitive to read and could be misunderstood as some time between 21:29 and 23:29. I would suggest keeping the same time format as we use everywhere else YYYY-MM-DD HH:MM:SS	0	0	https://issues.apache.org/jira/browse/AIRFLOW-891
19922	AIRFLOW-890	Are the example endpoints just the two in the docs? Specifically: /api/experimental/dags/<DAG_ID>/tasks/<TASK_ID> returns info for a task (GET). /api/experimental/dags/<DAG_ID>/dag_runs creates a dag_run for a given dag id (POST). Or is there a roadmap somewhere?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-890
19923	AIRFLOW-846	Hi, I see that 1.8 has finally been tagged, is there any plan to release it to pypi ? How is that automated/processed ? Cheers	1	0	https://issues.apache.org/jira/browse/AIRFLOW-846
19924	AIRFLOW-846	BTW [~artwr], would you please be kind enough to point me where this vote / release tracking is taking place ? Looks like I'm lost on JIRA...	1	0	https://issues.apache.org/jira/browse/AIRFLOW-846
19925	AIRFLOW-679	I see this has a closing PR. Is this ticket still relevant or can it be closed?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-679
19926	AIRFLOW-539	Not sure if you saw the note on the related issue, but perhaps this could be implemented via a kwargs passed through to the API, so that one-off feature changes don't require end-to-end plumbing in the future?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-539
19927	AIRFLOW-539	I agree re: kwargs, but I'd rather keep it as a separate thing from this just for simplicity's sake. Could you open a second JIRA to use kwargs for Google hook/operator params?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-539
19928	AIRFLOW-539	[~criccomini]: athough this ticket is closed (can you add a "gcp" label on this ticket), you're probably an admin. Thanks.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-539
19929	AIRFLOW-513	What's wrong with workaround 2? The point of parallelism is to control how many processes are running. A sensor that's waiting is in a while loop eating processing time, so you don't really want to ignore them for parallelism's sake. Short timeout, with tuned retry delay, and retry counts, will work just fine.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-513
19930	AIRFLOW-513	Porting convo from gitter to ticket: {code} [kyuen] option 2 is a valid way of preventing the problem. Just created the issue to see what people think the sensors should do. Didn't want to make a pr before seeing how people feel about how the sensors should behave [George Leslie-Waksman] Sensor doesn't really give the best sense for how they work. I think @mistercrunch was talking about more complex operations in the scheduler as an alternative way of doing things, but I think that's still a work in progress. {code} Would love to hear more from others on how they'd like sensors to behave	1	1	https://issues.apache.org/jira/browse/AIRFLOW-513
19931	AIRFLOW-511	[~criccomini] any thoughts on this?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-511
19932	AIRFLOW-479	Can I take this issue?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-479
19933	AIRFLOW-461	This is for GCS to BQ operator, right?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-461
19934	AIRFLOW-454	Where in the docs would you place this? Submit a PR if you think this is a necessary feature. For example, if someone were to require this link in the natural process of reading through a getting-started guide, I'd say it makes sense. Otherwise, i'd expect folks who are working on PRs to use Google to look into Flask-Admin, SQLAlchemy, or someone other 3rd party framework that Airflow leverages.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-454
19935	AIRFLOW-445	[~jlowin][~criccomini] What do you think?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-445
19936	AIRFLOW-409	[~mtdewulf] I'm considering taking ownership of this issue since you mention on your PR: "Since I've posted this PR, my team has decided that Airflow is not appropriate for our use case, so I probably won't be able to spend any cycles updating it." I assume you're happy for me to assign this to myself and take it over?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-409
19937	AIRFLOW-382	+1 [~maxime.beauchemin@apache.org], {noformat} While we're [decreasingly] comfortable installing all sorts of packages (R packages, Python libs, related apt packages, ...) on Airflow workers, we clearly don't want any of that on the Yarn boxes directly.{noformat} It could be easier to implement once Docker on YARN infrastructure is complete? YARN-3611, YARN-4793 https://hadoop.apache.org/docs/r2.7.2/hadoop-yarn/hadoop-yarn-site/DockerContainerExecutor.html	0	0	https://issues.apache.org/jira/browse/AIRFLOW-382
19938	AIRFLOW-374	> The task can monitor the time since last heartbeat and kill itself to prevent such cases. In your example, if the task can't get access to the DB how does it know? Are you planning on keeping some local in-mem variable that tracks the last successful heartbeat write?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-374
19939	AIRFLOW-350	Could you add some more description to this ticket? I don't understand it.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-350
19940	AIRFLOW-340	Also, should I track other unreferenced dependencies? I'm sure there's a couple of others lying around. If so, should I open an another ticket/PR to cover this extended work?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-340
19941	AIRFLOW-331	Ah, nice catch. Can you send a pull request, please?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-331

#	Requirement	Questions	Answered	Successfully Answered	URL
19942	AIRFLOW-329	I'm a little concerned about performance. The dag runs page, for example, barely loads for us, since we have 100s of DAGs at a 15m interval. Perhaps recent DAG runs is better?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-329
19943	AIRFLOW-325	Isn't this possible with an ExternalTaskSensor(same_dag_id, previous_task_id, allowed_states=[State.SUCCESS, State.FAILED])?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-325
19944	AIRFLOW-312	Can you try running off of master? I'm wondering if my PR in AIRFLOW-200 fixes this.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-312
19945	AIRFLOW-288	Thoughts on support for this? It would be nice to see this at the DAG level so individual DAGs could be triggered in a timezone aware way by the the scheduler.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-288
19946	AIRFLOW-276	Considering https://issues.apache.org/jira/browse/AIRFLOW-758, would it not be more straight forward to just revert back to execpv and run gunicorn with the --reload argument? Granted, this functionality is listed under the "Debugging" section of gunicorn's settings documentation, but it might be still be appropriate for this use case. Link to gunicorn docs: http://docs.gunicorn.org/en/19.0/settings.html#reload	0	0	https://issues.apache.org/jira/browse/AIRFLOW-276
19947	AIRFLOW-275	[~bolke], you committed the PR for this, right?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-275
19948	AIRFLOW-258	Makes sense to me. Want to send a PR?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-258
19949	AIRFLOW-258	Hey guys - this seems like the closest thing to an issue I'm seeing. I have a custom executor and a custom operator as a plugin. If I specify my custom executor in the config, airflow takes a different path than if I don't specify using the custom executor. If I'm configured with a built in executor like LocalExecutor, the plugin loads fine. If I'm configured to use my custom executor, then it loads the plugins which also loads the modules my custom operator is defined in, which is normal. I'm subclassing BaseSensorOperator and as its importing modules it throws an error about not being able to import BaseOperator from airflow.models from airflow/operators/__init__.py. Long story, any ideas on how to tackle this? Would be glad to work on a PR with a bit of guidance. Thanks!	0	0	https://issues.apache.org/jira/browse/AIRFLOW-258
19950	AIRFLOW-249	This is really awesome analysis. Can you send a PR? (unless I missed it)	0	0	https://issues.apache.org/jira/browse/AIRFLOW-249
19951	AIRFLOW-247	[~sanand], not sure what the protocol is for creating issues related to something I'd like to implement. Is what I've done here acceptable?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-247
19952	AIRFLOW-247	[~rfroetscher], are you still planning to work on this? If not, we can re-assign to [~gjreda].	1	0	https://issues.apache.org/jira/browse/AIRFLOW-247
19953	AIRFLOW-247	[~gjreda], want to have a look at the PR?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-247
19954	AIRFLOW-247	I'm searching for documentation related to how Airflow works with EMR. I'm struggling to find anything here: https://airflow.incubator.apache.org/integration.html#aws My main question is, can Airflow create an EMR cluster and bring it back down like AWS Data Pipeline? Thanks! EDIT: Found some information here: Spark, EMR: - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/example_dags/example_emr_job_flow_automatic_steps.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/example_dags/example_emr_job_flow_manual_steps.py - (uses emr hooks, operators) https://docs.google.com/presentation/d/1NG1P86HRlX43qTVucCTOsFqIbCvYdOhq_np90VlbVRc/edit#slide=id.gd40eeee67_1_0 - (uses shells scripts to launch and terminate emr clusters) https://www.agari.com/automated-model-building-emr-spark-airflow/ - (use shell script to spark-submit on a local spark installation) https://blog.insightdatascience.com/scheduling-spark-jobs-with-airflow-4c66f3144660 - (installs spark on each airflow worker node and runs local spark jobs without use of spark submit) https://medium.com/@calvertmg/airflow-integrating-with-apache-spark-50a7704dcebd - (alternative mozilla implementation for emr spark job) https://github.com/mozilla/telemetry-airflow/blob/master/dags/operators/emr_spark_operator.py EMR: - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/hooks/emr_hook.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/operators/emr_create_job_flow_operator.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/operators/emr_add_steps_operator.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/operators/emr_terminate_job_flow_operator.py Spark: - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/hooks/spark_submit_hook.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/operators/spark_submit_operator.py	1	0	https://issues.apache.org/jira/browse/AIRFLOW-247
19955	AIRFLOW-245	[~jlowin] Thanks for the feedback. But I don't quite get it. There's `BaseExecutor` class that all executors derived from, it has a field `queued_tasks` that's a dictionary of tasks where the key is the task key and the value is a tuple of `(command, priority, queue, task_instance)`. During heartbeat we already [retrieve the task from the db.	https://github.com/apache/incubator-airflow/blob/f4c44146cd38afe21c517230bff71d33ab041246/airflow/executors/base_executor.py#L110] We can pass to the `execute_async` method. I have a question then: I write a custom executor. How can I reuse available Operators and their fields in the executor. Say I write executor for Mesos, I want to have access to Docker operator fields from there and transform them to an HTTP call. What would be the preferable way to do that in Airflow?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-245
19956	AIRFLOW-245	The Scheduler is an object that decides what tasks to run; the Executor is an object that decides where to run those tasks. So the Scheduler hands a TaskInstance to the Executor and says "I want to run this task". That's why the Executor has access to the TaskInstance when tasks are being queued. The Executor then extracts the TI's "command" (for example: airflow run my_task) and its sole job is to decide how to run that command. execute_async() is where that decision is made. The important thing is that the thing being run is the string command, not the TI object, and the signature of execute_async should reflect that. The larger question here is why would an operator ever influence the executor? An Operator is a series of instructions that are designed to be run in a self-contained way, no matter what the executor is. An executor is a way of kicking off an Operator. There should be no dependency between them. To be explicit: an Operator must be able to run under any Executor. However, I am not familiar with Mesos so unfortunately I don't know exactly what you're trying to do -- but have you looked at the existing MesosExecutor in airflow.contrib.executors? Is it insufficient for your objectives? If your operator requires special handling, then it sounds like maybe you want a MesosOperator that runs your docker command?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-245
19957	AIRFLOW-199	This can be done via the tree view by clicking on the failed task(s), and clicking "clear" with "downstream". Is this what you're looking for? It will allow the DAG to resume.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-199
19958	AIRFLOW-177	[~criccomini] I was thinking that would it be helpful to have a cli support for that.. say {{ resume_dag }} which takes a execution_date as input and rerun only failed tasks?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-177
19959	AIRFLOW-169	What do you mean by expired schedules? You mean DagRuns that have finished? Which page in the UI are you referring to?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-169
19960	AIRFLOW-167	This sounds reasonable to me. Want to send a PR?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-167

#	Requirement	Questions	Answered	Successfully Answered	URL
19961	AIRFLOW-167	Specifically, the CLI solution sounds reasonable. The REST solution is better, but we haven't yet set up a REST API for Airflow yet. In the mean time, want to send a PR for the CLI {{dag_state}} command?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-167
19962	AIRFLOW-165	I think this is already done, right? https://pythonhosted.org/airflow/concepts.html#task-documentation-notes You can do: {code} t = BashOperator("foo", dag=dag) t.doc_md = """\ #Title" Here's a [url](www.airbnb.com) """ {code}	0	0	https://issues.apache.org/jira/browse/AIRFLOW-165
19963	AIRFLOW-161	What context is this under? You mean like you want to use the EmailOperator or SlackOperator to notify people to download a file that's been created as part of the DAG? We do this using XCom+EmailOperator. XCom variables can be accessed via templates. We store the file in a blob store (like S3). The file location is stored in XCom, which the EmailOperator references when it sends the email.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-161
19964	AIRFLOW-161	I don't think that we want to embed Quoble logic directly in Airflow. I'm a bit out of my element on the UI-front, though. Perhaps there's a way to achieve this through plugins, or by simply putting the link in the logs?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-161
19965	AIRFLOW-161	Sure, can you please cc top contributors to PR or here.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-161
19966	AIRFLOW-135	I'm +1 on cutting branches for releases (instead of tags). I'm also +1 on deleting dead branches, but could you post a list of the proposed branches to delete so we can allow people time to notify if they have a branch that they don't want deleted?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-135
19967	AIRFLOW-129	[~maxime.beauchemin@apache.org], [~aoen], [~artwr], could one of you take a look at this? I think you guys are the big Celery users.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-129
19968	AIRFLOW-128	[~bolke], OK to close this?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-128
19969	AIRFLOW-123	What exactly do you define as a timeout? # Task is killed because it's a zombie # A connection to some infrastructure times out # A task runs for longer than it should (1) seems to be the one that Airflow could do right now. (2) is hard--it requires cooperation from the operator to throw some kind of known TimeoutException (unless you want to use Python built-in TimeoutError). (3) I supposed could be handled as a task-level SLA.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-123
19970	AIRFLOW-121	Hello. I think you should file a new bug issue with an example DAG that triggers this bug. Please also specify which version you are running.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-121
19971	AIRFLOW-118	Was there a PR for this?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-118
19972	AIRFLOW-115	[~artwr] Are you working on this issue? I'm facing many errors when using the current `S3Hook` mostly caused by it using `boto` instead of `boto3` (e.g. https://github.com/boto/boto/issues/2836)	0	0	https://issues.apache.org/jira/browse/AIRFLOW-115
19973	AIRFLOW-113	[~Shenghu], can you: # Test this against master? # Post your DAG (or some equivalent) Python code so we can try and replicate	1	1	https://issues.apache.org/jira/browse/AIRFLOW-113
19974	AIRFLOW-113	Also, I wonder if perhaps you want max_active_runs instead of concurrency? From [FAQ|https://pythonhosted.org/airflow/faq.html]: {quote} concurency defines how many running task instances a DAG is allowed to have, beyond which point things get queued {quote} Whereas: {quote} Is the max_active_runs parameter of your DAG reached? max_active_runs defines how many running concurrent instances of a DAG there are allowed to be. {quote}	1	1	https://issues.apache.org/jira/browse/AIRFLOW-113
19975	AIRFLOW-111	Thanks [~criccomini]. I fixed the description by changing 'jobs' to 'task instances'. Can you please help on review the related PR: https://github.com/apache/incubator-airflow/pull/1579 ?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-111
19976	AIRFLOW-101	[~bolke], which PR was this part of?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-101
19977	AIRFLOW-93	[~underyx] I'm rapidly learning about ExternalTaskSensor but could you not create 6 ETS in your hourly DAG, one for each of the 10 minute tasks (so one with delta of 10 minutes, one with 20 minutes, one with 30 minutes, etc.) and then have your hourly task depend on all of those sensors completing successfully? See for example AIRFLOW-47.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-93
19978	AIRFLOW-91	ok i pushed new PR https://github.com/apache/incubator-airflow/pull/1497 but i really misunderstand (if i replace 'ssl_certfile': None, section) i got : {color:red} Traceback (most recent call last): File "/usr/local/bin/airflow", line 13, in <module> parser = get_parser() File "/usr/local/lib/python2.7/dist-packages/airflow/bin/cli.py", line 689, in get_parser default=configuration.get('webserver', 'ssl_keyfile'), File "/usr/local/lib/python2.7/dist-packages/airflow/configuration.py", line 520, in get return conf.get(section, key, **kwargs) File "/usr/local/lib/python2.7/dist-packages/airflow/configuration.py", line 428, in get "in config".format(**locals())) airflow.configuration.AirflowConfigException: section/key [webserver/ssl_keyfile] not found in config {color} so please check ones more my changes ? thanks !	0	0	https://issues.apache.org/jira/browse/AIRFLOW-91
19979	AIRFLOW-91	i update related issue, pls take a look ones more ?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-91
19980	AIRFLOW-91	hi for all, i already update PR, plz take a look on it? thanks !	1	0	https://issues.apache.org/jira/browse/AIRFLOW-91
19981	AIRFLOW-81	What do you think about writing a new sensor based off of the TimeSensor that blocks during the blackout period?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-81
19982	AIRFLOW-79	Is this dead in the water? I've been reading about it in the pull request and on the mailing list and on the Gitter channel...	1	0	https://issues.apache.org/jira/browse/AIRFLOW-79
19983	AIRFLOW-78	[~abridgett] Can you tell me more about your expectations of clear?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-78
19984	AIRFLOW-78	Assigning this to Norman. We can clear dag runs issued by clear in both the UI and CLI. [~bolke] [~jlowin][~maxime.beauchemin@apache.org] Do you see any issue with this idea?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-78
19985	AIRFLOW-78	[~abridgett]: Just to clarify, the issue here is still that the cleared dagruns get rerun without respecting the max_active_runs amount?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-78
19986	AIRFLOW-78	[~abridgett]Is max_active_runs in this scenario set to 1? Did you set it to depend on upstream? And nothing else is running?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-78
19987	AIRFLOW-71	[~asnir], can you describe what this is a bit more?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-71
19988	AIRFLOW-71	[~asnir], do you want add config to the DockerOperator, such that it can execute a `docker login` before it runs the Docker command?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-71
19989	AIRFLOW-71	Hi [~criccomini], I'm really needing it, and I think it's simple to do that... I was thinking in: Expose the Client (docker-py.client.Client) that is instantiated inside of DockerOperator.execute() method, changing it to instantiate in __init__(), so we will be able to execute the login passing username, password, and the specific registry too.... eg: d = DockerOperator(task_id='example_task_id', image='example/image-name', force_pull=True, owner='felipetancredo', command="./bi_jobs/etl/example.py {{ execution_date }} {{ execution_date }}") user = 'test' pwd='pwd' if d.cli: d.cli.login(username=user, password=pwd, registry='https://registry.example.com', reauth=True) Can I contribute to docker_operator? Do u see any problem? (sorry, I'm beginner in the open source projects)	1	0	https://issues.apache.org/jira/browse/AIRFLOW-71
19990	AIRFLOW-71	When will [~asnir]'s PR be approved? It's been 8 months.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-71
19991	AIRFLOW-58	[~maxime.beauchemin@apache.org], any reason why this was closed?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-58
19992	AIRFLOW-58	[~sanand], did you mean to put a real external URL in? You just put the JIRA in. ;)	1	0	https://issues.apache.org/jira/browse/AIRFLOW-58
19993	AIRFLOW-55	Overall, this looks good to me. Given that it's in contrib, and basically can't be supported by the committers (unless one of us wants to use it?), I'm OK without tests. [~sanand], what's your take?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-55

#	Requirement	Questions	Answered	Successfully Answered	URL
19994	AIRFLOW-54	+1 for this feature as well. cc @jlowin (the original creator of xcom) [~amustafa] Would you like to take this on? The one case you need to think about is subdags. I think the relationships above translate across subdags, but please do test with subdags as well. Also, please change '@upstream' to '@parent'	1	0	https://issues.apache.org/jira/browse/AIRFLOW-54
19995	AIRFLOW-54	[~amustafa] Any PRs that we can look at?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-54
19996	AIRFLOW-52	[~criccomini] Can we do an 'airbnb' release for now then? Ie. not use the apache brand? Cause this release is long overdue and 1.7.0 is a bit buggy. Chasing licenses will cost at least another week or even more (with broken PRs to add)	1	1	https://issues.apache.org/jira/browse/AIRFLOW-52
19997	AIRFLOW-52	[~jghoman]/[~hitesh]/[~cnauroth] what's the best way to handle this? We basically had an RC in flight for like a month before Apache.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-52
19998	AIRFLOW-52	So, the release will be coordinated (and built from) github and github resources? I think Chris is right, but I'd be concerned about running the vote or whatever process was being used for the release on ASF resources. I'd want to avoid the appearance of trying to make an end-run around asf procedure (which I know isn't the case). I'm honestly not sure though; may be worth pinging general@incubator.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-52
19999	AIRFLOW-52	[~aoen], are you planning to burn in [~jlowin]'s fix before release?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-52
20000	AIRFLOW-52	This one right https://github.com/apache/incubator-airflow/pull/1473 ? In this case yes.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-52
20001	AIRFLOW-52	Should this be closed?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-52
20002	AIRFLOW-51	Shouldn't we support both, rather than deleting the `cpus` field? As I read it, it sounds like Docker supports both CPU count and cpu_shares percentage.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-51
20003	AIRFLOW-48	It's likely not implemented because no one is using it. Can you provide an example of your uri? Is it something like *postgresql://localhost:2345/your_db?sslmode=require*? What version of postgresql are you running? I'm not sure how well this parameter is supported for versions of postgresql before 9.2. Another alternative is to take control of this by implementing this logic in a PythonOperator's python callable. That could read this param and use psycopg2 to connect to your db. Of course, then you would not be using the Connection or Hook, but, you would be able to achieve your end goal. In the meantime, feel free to submit a PR for review -- we'd be happy to accept it.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-48
20004	AIRFLOW-45	What do you think about defining a heuristic that variables would be hidden from the UI if they contain the text: hidden, password, secret	1	1	https://issues.apache.org/jira/browse/AIRFLOW-45
20005	AIRFLOW-45	{quote} I think we should have a configuration variable to hide/expose the encrypted values in the UI for both Variables and Connections. It should likely affect all Variables and/or Connections in a given Airflow installation and not be DAG specific. {quote} This seems OK to me. So, current proposal is: # Add a config param to airflow.cfg called hide_encrypted_ui_fields # Default hide_encrypted_ui_fields to true The hide_encrypted_ui_fields param would hide `passwords` and `extras` in the hooks view, as well as the `value` of variables in the variables view. Does that sound OK?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-45
20006	AIRFLOW-45	Sounds good. [~cheny258], would you like to try this?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-45
20007	AIRFLOW-45	[~cheny258] As part of this Variables work, there are currently 2 views. One is the edit view, the other a list view of all variables. Currently, the list view only shows the variables names, not the values. i'd like to display the values if we the configuration param proposed above is set to False. Can you add this feature? Currently, one must click edit to view a variable value.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-45
20008	AIRFLOW-45	[~sanand] will do that for the list view of variables. Also, for conn edit page, {{extra}} and {{password}} are already in a hidden_fields list. Do we want to make any changes in edit or list view for that?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-45
20009	AIRFLOW-45	Could we not handle this with a regexp of fields that should be hidden rather than a boolean?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-45
20010	AIRFLOW-45	{quote} I'd have to go change the config whenever I do this, which is kind of annoying {quote} (Yes, quoting myself :)). Perhaps not. As long as the regex was something like '.*secret.*', and I stuck to that as a naming convention, I'd be fine. Still, to limit burden on [~cheny258], I'd like to keep the regex to a separate ticket. I'll open it up after this one is closed off. [~withnale]/[~sanand], does that sounds OK?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-45
20011	AIRFLOW-39	Can you provide sample code?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-39
20012	AIRFLOW-39	Oh, so not version 1.7.0? What does 'git log | head' look like for you?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-39
20013	AIRFLOW-35	[~sanand], any particular reason you closed this?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-35
20014	AIRFLOW-31	I'm not sure where the best place to comment is... Wherever you're comfortable? Let me know if you discover anything -- in my experience getting objects to behave like modules is tricky (supporting full {{import}} syntax etc). Maybe some of what I did for plugins would be reusable?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-31
20015	AIRFLOW-30	Initially, this seemed fine to me, but it actually sounds a little tricky to me. The concern I have is what happens if the DAG fails half way through and the transaction is reverted? In such a case, Airflow will show prior operators as having successfully run, but their state mutation will not have taken place. This effectively means that the entire DAG needs to be re-run, even if some of the operators show a successfully executed view.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-30
20016	AIRFLOW-30	In your use cases, are the source and destination DBs part of the same DB cluster?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-30
20017	AIRFLOW-29	One thing to consider: what happens if the scheduler machine is overloaded, and imports are taking a long time because of that? (similar to a GC pause in Java).	1	1	https://issues.apache.org/jira/browse/AIRFLOW-29
20018	AIRFLOW-29	[~underyx], re (1), I wasn't suggesting GC would be an issue--was just saying it's a similar problem to timeouts in Java when GC occurs. Anyway, I agree, probably not worth worrying about. What default did you have in mind? 5s?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-29
20019	AIRFLOW-29	[~underyx] I'm not connecting the configuration suggestion to how such a change would resolve your problem. Do you mind providing a little more context on impact of such a change? If the import is actually only taking 1 second on your box and if you are not hitting a timeout, then why does the larger configured value have an impact on you?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-29
20020	AIRFLOW-28	Okay, but what's preventing someone else from implementing this? How come the issue was closed as won't fix?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-28
20021	AIRFLOW-23	Hey [~alexvanboxel], checking in to see how this and AIRFLOW-24 are coming.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-23
20022	AIRFLOW-22	Can public buckets/BQ datasets (thinking about GDELT) be used without a project?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-22
20023	AIRFLOW-22	What's GDELT? In the case of GCS, yes I believe we can use public buckets. I'm uncomfortable baking in someone's own buckets, though. It'd have to be public buckets that Google, themselves, provide. For BQ, unfortunately not. The problem with BQ is that it requires a project to run its queries on. Even if the datasets are public, you still need to have access to a whole project, and be a viewer for it, in order to execute any queries against the data.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-22
20024	AIRFLOW-21	Do you need some help? I could do it.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-21

#	Requirement	Questions	Answered	Successfully Answered	URL
20025	AIRFLOW-21	[~alexvanboxel], that would be awesome. Can you wait until AIRFLOW-16 goes in, though, to avoid merge conflicts? Should be today.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-21
20026	AIRFLOW-20	The DagRun.previous handling is used to connect previous related DagRuns in order to make depends_on_past work with a moving interval or an automatically aligned start_date. Imagine the following Use Case 1 Dag1 has an assigned start_date "2016-04-24 00:00:00" with a cron schedule of “10 1 * * *”, ie run every day at 1.10am. Depends_on_past is true. The current scheduler kicks off a run “2016-04-24 00:00:00” and then stops as the previous interval of “2016-04-24 01:10:00” is “2016-04-23 01:10:00”. This cannot be found as it has not taken place and thus the whole thing grinds to a halt and the TaskInstances refuse to run. What the user expects here is that the first run is “2016-04-24 01:10:00”, ie start_date + interval, unless the start_date is on the interval, ie. start_date is first interval. This is what I address by start_date normalization in the PR. However, the second issue then kicks in as the “previous” run can still not be found. Use Case 2 Dag2 has an assigned start_date "2016-04-24 01:10:00" with a cron schedule of “10 1 * * *”. Depends_on_past is true. The scheduler happily goes through a few runs, but then the dag is updated and the schedule adjusted. Because the previous interval cannot be found by the TaskInstance (emphasis), tasks get stuck again requiring an manual override. What the user expects here is that the scheduler is smart enough to figure out that we are still running the same dag and that it needs to look up the previous run for that dag and make sure dependencies are met with that previous dagrun in mind. I don’t think those two use cases are edge cases, considering the amount of questions we get on these subjects. To resolve the remaining issues (aside from start_date normalization) I first made DagRun aware of its predecessors. Then I strengthened the relationship between TaskInstances and DagRuns slightly, by optionally including a dagrun_id in the TaskInstance. Now a TaskInstance can lookup its predecessors in the previous DagRun and know for sure that it is either the first run or it has a predecessor somewhere in time instead of guessing. What I am unsure of is what to consider is the unit of work: 1) is a TaskInstance dependent on its past predecessor ignoring the outcome of the DagRun? (previous TaskInstance can be successful, but previous DagRun as a whole can fail) or 2) is it dependent on the outcome of the DagRun? 3) Can it be both? In case of 1 and 3 my logic needs to be updated slightly, but that should not be too much of a big deal. However I have difficulty imagining why you want to do that.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-20
20027	AIRFLOW-20	Max I understand your concerns, I'm working with Jeremiah to see if I can remove the need for the dag run id in taskinstances. I would be happy to as it would simplify the change. Rigorous testing needs to take place before making it part of a release, current unit tests do not cover enough. Documentation is forthcoming but I will only assemble it when all assumptions have been worked out and discussions finished. Please also note that we (Jeremiah and I) are seeing this as the first step towards his bigger pr. Sort of paving the way. If you can regard it in that context and not just "making it easier for schedules that move". In that regard the more fundamental questions I have also asked on the list have not been answered yet: what do we consider the unit of work? Is it a DagRun or is it a taskinstance? We talk about dags, but we can run tasks without a DagRun. I consider taskinstances part of a DagRun. They should not be able to look beyond the borders of the containing DagRun. If they want to do so they should query the DagRun for this. If you do this backfills and so on become much easier to handle.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-20
20028	AIRFLOW-20	I took the plunge and made the "dag_run_id" field for TaskInstances NOT NULL. This allowed me to catch places where TaskInstances are manipulated and to fix these. So indeed there are some strange places like the PythonOperator and views.py that create TaskInstances. That is really a refactor task towards the future (eg. dao principle would be smart maybe) The good news is these have all been fixed. All tests pass, including some additional ones I added. Now I can dive in to some other cases. Like what to do when a backfill has an execution_date that is equal to a previous scheduled run or backfill? 1. Do we update the tasks and run new ones and skip ones that were successful? If so do we connect old ones to the backfill run or do we leave them as is (ie connected to the old run). Think lineage here. How do we handle dependencies if the old tasks stay connected to their previous run? 2. Do we refuse to run? That would disallow you to "add tasks" in the past. I tend to lean towards no 1 and connect all tasks to the backfill and setting the state of the "old" dagrun to "SUPERSEDED" or "OVERRIDDEN". The same flexibility as today would exist but an audit trail is preserved.	0	0	https://issues.apache.org/jira/browse/AIRFLOW-20
20029	AIRFLOW-17	Got it. Is there any way to fix it?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-17
20030	AIRFLOW-16	Pull request up at: https://github.com/airbnb/airflow/pull/1452 This PR should also fix: https://github.com/airbnb/airflow/issues/1354 [~jlowin]/[~alexvanboxel], I could use a review from you.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-16
20031	AIRFLOW-16	> Got it working with the new hooks Did you find any bugs, or started up without incident? > GCP operations and hooks be prefixed with gcp ... Could be a ticket, I'll pick it up (after this is merged). Sounds good. Please do! Go ahead and open the JIRA.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-16
20032	AIRFLOW-15	Pull request: https://github.com/airbnb/airflow/pull/1448 [~jlowin], could use a review.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-15
20033	AIRFLOW-15	[~criccomini] is this issue resolved?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-15
20034	AIRFLOW-14	Overall, major +1 on this. One thing that I think is glossed over is what happens when a DagRunJob dies ({{kill -9}}, for example)? You mention heartbeats, but I don't see a mention of who is going to clean up/take over the execution of the wedged DagRuns when their job dies. I see this problem a lot with Airflow right now (DagRuns stuck in 'running' forever), so it'd be good to think through.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-14
20035	AIRFLOW-14	If I'm not mistaken, I think this work should solidify running multiple schedulers as a first-class citizen, right? We are not going to run Celery for distributed execution, so having a trustworthy way to run more than one scheduler is useful. /cc [~sanand]	1	0	https://issues.apache.org/jira/browse/AIRFLOW-14
20036	AIRFLOW-14	How does a DRJ grab a Dag to execute? It seems to me like there's a race condition where two DRJs could simultaneously see that a DAG needs to be run, and both could insert entries into the dag_run table simultaneously. Is this being enforced by a primary key/unique constraint in the DB?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-14
20037	AIRFLOW-14	> DRJ needs to make sure to refresh from the db to check for a lock immediately before running it Isn't there a race condition where two DRJs check the lock in the DB before running, both see it's unset, and both set it to themselves? Even if they then check to verify that they are the owners of the lock, it's not trustworthy. The only way I can see this working is if a real read/write lock is obtained. Is SQLAlchemy doing this?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-14
20038	AIRFLOW-14	The good news is that lock_id is already set to the job's id, so I just need to add the additional check that it matches (right now it just assumes if it was able to lock, then it owns the lock, which is not as strict). But I'm still curious about your point: > Even if they then check to verify that they are the owners of the lock, it's not trustworthy. Is that not still an issue?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-14
20039	AIRFLOW-14	To clarify, would different DagRuns be evaluated in the same process, or in different processes? In the current work-in-progress branch, I noticed that task scheduling is done by each dag run. However, with the current approach, pool limits may not be correctly respected?	0	0	https://issues.apache.org/jira/browse/AIRFLOW-14

#	Requirement	Questions	Answered	Successfully Answered	URL
20040	AIRFLOW-14	While I understand that for being truly distributed it would be great to have a distributed scheduler, are not not looking for a distributed executor instead [~criccomini]? Now the scheduler and the executor are coupled, but that does not need to be and does not even make sense imho.	1	1	https://issues.apache.org/jira/browse/AIRFLOW-14
20041	AIRFLOW-14	{quote} Now the scheduler and the executor are coupled {quote} Can you expand on this a bit? I know [~sanand] is running multiple LocalExecutor schedulers. We were planning to do this as well.	1	0	https://issues.apache.org/jira/browse/AIRFLOW-14
20042	AIRFLOW-14	Now the scheduler is, in case of using the LocalExecutor, tied to the executor: ie. both are same process. This is nor really required and maybe not even the best way forward. A scheduler should, imho, only schedule and an executor should execute. With celery you already get a kind of decoupling. The celery executor puts tasks on a queue and then celery workers pick these up. It is not so difficult to imagine that the local executor can do the same and pick up tasks for a queue that is either something like rabbitmq or just the db. In other words the scheduler should queue tasks in a certain order and the executor should pick these up from anuwhere. Ie. producer/consumer like. Does this make sense?	1	1	https://issues.apache.org/jira/browse/AIRFLOW-14
20043	AIRFLOW-14	Yea, I agree with this. I guess the two things that I'm worried about are: # Will this mean we can't safely run more than one LocalExecutor? # Will this break the CeleryExecutor?	1	0	https://issues.apache.org/jira/browse/AIRFLOW-14
20044	ANY23-297	GitHub user lewismc opened a pull request: https://github.com/apache/any23/pull/30 ANY23-297 Any23 doesn't build under JDK1.8 This issue addresses https://issues.apache.org/jira/browse/ANY23-297 @ansell can you please check how this affects your open pull request? Do you have any objections to merging this into master branch? Thanks You can merge this pull request into a Git repository by running: $ git pull https://github.com/lewismc/any23 ANY23-297 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/any23/pull/30.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #30 ---- commit bb55685859a26b36ca4a9893bb93aa9eb7687b8c Author: Lewis John McGibbney <lewis.mcgibbney@gmail.com> Date: 2016-11-25T21:33:57Z ANY23-297 Any23 doesn't build under JDK1.8 ----	1	1	https://issues.apache.org/jira/browse/ANY23-297
20045	ANY23-297	Github user lewismc commented on the issue: https://github.com/apache/any23/pull/30 Hi @ansell any objections to merging into master branch? Seems the like the semargl issues are progressing which is great, however I don't think that should block us here in Any23. I'll commit by CoB today unless there are objections. Thanks	1	1	https://issues.apache.org/jira/browse/ANY23-297
20046	ANY23-291	I agree [~thomas.francart], are you working off of master branch? If so, are you able to provide us with a patch and a test case if possible?	1	1	https://issues.apache.org/jira/browse/ANY23-291
20047	ANY23-280	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/24#issuecomment-212761670 Anyone get a chance to take a look at this? This is THE critical issue for Any23 to address right now IMHO.	0	0	https://issues.apache.org/jira/browse/ANY23-280
20048	ANY23-280	Github user ansell commented on the issue: https://github.com/apache/any23/pull/24 There are a large number of whitespace modifications to change from tabs to 2-space indentation. Is 2-space indentation what Any23 is aiming for, given that most java code is either tab or 4-space indentation.	0	0	https://issues.apache.org/jira/browse/ANY23-280
20049	ANY23-276	Github user ansell commented on the issue: https://github.com/apache/any23/pull/27 I rebased this to fix the conflicts. There are some test failures, some of which I @Ignore'd while other I left so far. Could you have a look in to see if you can figure out why the tests that have been ignored or still failing are broken?	1	1	https://issues.apache.org/jira/browse/ANY23-276
20050	ANY23-276	Github user lewismc commented on the issue: https://github.com/apache/any23/pull/27 Hi @ansell whats the current state of this patch? Thanks for merging in the PR.	1	1	https://issues.apache.org/jira/browse/ANY23-276
20051	ANY23-276	Github user lewismc commented on the issue: https://github.com/apache/any23/pull/27 Hi @ansell this is starting to look REALLY good. I've pulled your suggestions. I did notice that the appassembler is still generated in the core/target/appassembler. I think that this can be removed based upon the suggestion for CLI to be modularized as you've implemented. Is there anything else you want to add to this patch? If not then I suggest we merge into master.	1	1	https://issues.apache.org/jira/browse/ANY23-276
20052	ANY23-259	[~p_ansell] is this the issue which is leading to the unstable builds?	1	1	https://issues.apache.org/jira/browse/ANY23-259
20053	ANY23-259	bq. There also appear to be a changed (lower, unless the junit assertEquals was applied backwards) number of extractors than expected. Hi [~p_ansell] what do you mean by this :(	1	1	https://issues.apache.org/jira/browse/ANY23-259
20054	ANY23-257	I added implementations of Sesame RDFParser's for common OWL formats, including Manchester Syntax, Functional Syntax (using the obo to owl algorithm that is in owlapi), and OWL/XML to owlapi-rio:4. See the list at: https://github.com/owlcs/owlapi/blob/version4/rio/src/main/resources/META-INF/services/org.openrdf.rio.RDFParserFactory As OWL isn't an input format, per se, which particular format were you wanting to target?	1	1	https://issues.apache.org/jira/browse/ANY23-257
20055	ANY23-250	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/13#issuecomment-76448755 all tests pass locally. Would be appreciated if someone could confirm.	1	1	https://issues.apache.org/jira/browse/ANY23-250

#	Requirement	Questions	Answered	Successfully Answered	URL
20056	ANY23-247	Github user ansell commented on a diff in the pull request: https://github.com/apache/any23/pull/17#discussion_r27437088 --- Diff: core/src/main/java/org/apache/any23/validator/rule/MissingItemscopeAttributeValueRule.java --- @@ -0,0 +1,52 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ + +package org.apache.any23.validator.rule; + +import org.apache.any23.validator.DOMDocument; +import org.apache.any23.validator.Fix; +import org.apache.any23.validator.Rule; +import org.apache.any23.validator.RuleContext; + +/** + * This fixes missing attribute values for the 'itemscope' attribute, + * which was be associated with nodes. + * Typically when such a snippet of XHTML is fed through the + * {@link org.apache.any23.extractor.rdfa.RDFa11Extractor}, and + * subsequently to Sesame's {@link org.semarglproject.sesame.rdf.rdfa.SesameRDFaParser}, + * it will result in the following behavior. + * <pre> + * {@code + * [Fatal Error] :23:15: Attribute name "itemscope" associated with an element type "div" must be followed by the ' = ' character. + * } + * </pre> + * This Fix is an effort to mitigate against that happening. + */ +public class MissingItemscopeAttributeValueRule implements Fix { --- End diff -- How is this class recognised or instantiated? META-INF/services/ or another method?	1	1	https://issues.apache.org/jira/browse/ANY23-247
20057	ANY23-247	Any ideas about this [~p_ansell]. Did you manage to debug where and when rules and/or fixes are invoked and when they are applied?	1	1	https://issues.apache.org/jira/browse/ANY23-247
20058	ANY23-247	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/17#issuecomment-201537723 hi @ansell OK I've added in the correct rule and fix as well as a test to verify that empty itemscope values are identified and fixed. Whilst debugging this however the core issue persists. Reasoning for this is that ```RDFa11Extractor extends BaseRDFExtractor``` which inherits the [parser function inputstream parameter](https://github.com/apache/any23/blob/63ba2fc82966cc056a2e475af849154d0dfdcf93/core/src/main/java/org/apache/any23/extractor/rdf/BaseRDFExtractor.java#L105). This input stream is not the 'fixed' steam but the raw document. The only way I can think around this is for us to * refactor the [RDFa1.1Extractor](https://github.com/apache/any23/blob/master/core/src/main/java/org/apache/any23/extractor/rdfa/RDFa11Extractor.java) such that it extends [TagSoupDomExtractor](https://github.com/apache/any23/blob/63ba2fc82966cc056a2e475af849154d0dfdcf93/api/src/main/java/org/apache/any23/extractor/Extractor.java#L60) as oppose to (eventually) the [ContentExtractor](https://github.com/apache/any23/blob/63ba2fc82966cc056a2e475af849154d0dfdcf93/api/src/main/java/org/apache/any23/extractor/Extractor.java#L44), or * undertake a mass refactoring which essentially removes the [ContentExtractor](https://github.com/apache/any23/blob/63ba2fc82966cc056a2e475af849154d0dfdcf93/api/src/main/java/org/apache/any23/extractor/Extractor.java#L44) altogether... this would provide us with a much more flexible and adaptable extraction framework IMHO. What do you think?	1	1	https://issues.apache.org/jira/browse/ANY23-247
20059	ANY23-247	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/17#issuecomment-201550510 I agree. Jumping through this in the debugged made me think the same. I think it is different if Any23 is to be a PURE implementation... But that is clearly not the case. Any23 fits in best when it can be used to extract semantics from any old crap input that it is fed. Parsers and extractors *should not* fail when there is a piece of crap input HTML. Currently, that's exactly what happens and it is extremely limiting. I would like to propose that this PR is committed to master as is, we then open a brand new issue which acts exactly your comments refactoring out content extractor and reusing the input stream which has been fixed, etc. Any thoughts Peter? Thanks fr quick response. On Friday, March 25, 2016, Peter Ansell <notifications@github.com> wrote: > The system does seem a little too complex for our purposes and isn't > usable because of that. > > > Removing generics would be the first step IMO as there are too many > rawtypes definitions which indicate generics are being used badly. > > > ContentExtractor may be able to be completely removed instead of being > refitted into the process after that and the parser should always be set to > parse as far as practical for our purposes. > > > It is a little strange that there isn't a buffered, markable, InputStream > provided for all of the steps to reuse as necessary rather than pushing a > raw InputStream or other source into different extractors. > > > — > You are receiving this because you authored the thread. > Reply to this email directly or view it on GitHub > <https://github.com/apache/any23/pull/17#issuecomment-201545776> > -- *Lewis*	1	1	https://issues.apache.org/jira/browse/ANY23-247
20060	ANY23-247	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/17#issuecomment-202702530 @ansell any further comments here? I will try to get to work on the larger issue this week.	1	1	https://issues.apache.org/jira/browse/ANY23-247
20061	ANY23-238	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/9#issuecomment-55443359 Hi @Timpy thank you so much fo this PR. I absolutely agree with you that this is a problem as I do not like the look of the BNode hash either, it is useless for many purposes. For the record, can you please post here what the new outout will look like. Thank you p.s. I posted https://issues.apache.org/jira/browse/ANY23-238 so that we can add this to our release report once we release Any23 1.0	1	0	https://issues.apache.org/jira/browse/ANY23-238
20062	ANY23-226	[~p_ansell], can you detail the current status of this over on jsonld-java? I feel that it would be best to bake in support in jsonld-java library which we wrap in Any23. Thanks in advance.	1	1	https://issues.apache.org/jira/browse/ANY23-226

#	Requirement	Questions	Answered	Successfully Answered	URL
20063	ANY23-226	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/16#issuecomment-84384643 Ok Peter thank you for looking. This is great. I have not seen the test failures. Can you please tell me if it is in Any23 or in jsonld-Java? We could upgrade the Jsonld-Java implementation as well. To the 0.5.1 release On Saturday, March 21, 2015, Peter Ansell <notifications@github.com> wrote: > The main bug was that the entire script node was being sent to > JSONLD-Java, and not just its content. > > However, I also made a few other changes while doing that testing. > > It turned out that the jsonld was invalid, but somehow the exception when > parses fail was changed to be silently swallowed, so the only indication > was that the count was 0. I turned on the exception propagation again (no > reason it should be swallowed outside of temporary testing). > > However, in addition to the 4 tests currently failing on the core tests, > there are now other tests failing due to an inability to parse " > " > > — > Reply to this email directly or view it on GitHub > <https://github.com/apache/any23/pull/16#issuecomment-84257478>. > -- *Lewis*	1	1	https://issues.apache.org/jira/browse/ANY23-226
20064	ANY23-207	Hi [~cuibonobo] I am VERY sorry for getting back to you so late. Right now we have a Microformats2 GSoC project underway so we can definitely sync. Are you up for that?	1	1	https://issues.apache.org/jira/browse/ANY23-207
20065	ANY23-207	Hi Folks, there is a PR for this now available at https://github.com/apache/any23/pull/18 Would be great if you guys could help us peer review the patch. Thanks.	1	1	https://issues.apache.org/jira/browse/ANY23-207
20066	ANY23-202	Do you want to do Tomcat monitoring? It would mean less code for us and more reuse of existing monitoring. We can automate this as well along with the nightly pull script. wdyt? I've classified this for 1.1 as it is certainly non-blocking as as the WebService at any23.org will be updated nightly there is no need for us to hold a 1.0 release for this.	1	0	https://issues.apache.org/jira/browse/ANY23-202
20067	ANY23-202	Does Apache have a privacy policy that details this collection of information by Google Analytics?	1	1	https://issues.apache.org/jira/browse/ANY23-202
20068	ANY23-202	[~p_ansell], I honestly do not know! I am not really too happy about Google Analytics being part of Any23 source or binary release. It means that everyone who deploys an Any23 server with the Google Analytics code embedded can be tracked! I am not really too happy with that. [~michele.mostarda], did you consider adding Tomcat logging metrics instead of Google ones? This would mean that it is down to whoever consumes Any23 to do their own log analysis, not for us to do it for them.	1	0	https://issues.apache.org/jira/browse/ANY23-202
20069	ANY23-202	Hi guys, I used Google Analytics because as you can easily verify other Apache projects already embed it within their landing pages (see for example http://hadoop.apache.org/). Any metrics provided by Tomcat is not comparable with the informations obtained by Google Analytics like geographic distribution and disambiguation of accesses. The tracker is embedded only in the demo front-end service page and anybody can customize (change the tracking ID) or remove it. Suggestions to make it less invasive: we could perhaps add the tracking ID within the POM so that it is more visible and left it empty in the distributions? we could add a disclaimer in the README? Suggestions? In general I suggest to keep a way to fine track user activity because we need to quantify the usage, the community and the distribution to better improve support.	1	1	https://issues.apache.org/jira/browse/ANY23-202
20070	ANY23-202	The fact that some other projects have it doesn't absolve the privacy issue. There are many questions that should have been asked before the tracker was added. Who is using the information other than yourself and Lewis now or in the future? Was this a request from your employer or just from personal interest? Why didn't you add a privacy policy when you added the tracker? Why doesn't the entire project team have access to the information? I am not sure what you mean by adding information to the POM, but this is certainly something that is never done by any other projects.	1	1	https://issues.apache.org/jira/browse/ANY23-202
20071	ANY23-198	Hi Norbert. Are you using the Any23 source code? Are you able to contribute a patch for this feature? It would be a nice contribution.	1	0	https://issues.apache.org/jira/browse/ANY23-198
20072	ANY23-189	what is the plan in this ticket? to integrate MO in any23 in some way?	1	0	https://issues.apache.org/jira/browse/ANY23-189
20073	ANY23-185	hi [~nisala12], nice patch :) Does it comply with the following ruleset "...Note that the global name has a specific semantic in the <meta> element and that the itemprop must not be set when one of the name, http-equiv or charset is already used."?	1	1	https://issues.apache.org/jira/browse/ANY23-185
20074	ANY23-185	Hi Lewis John McGibbney, One thing I have notice in the spec[1] itemprop is stated as "unimplemented", and I cant really find a reason behind this labeling. WDYT ? ---------- if the itemprop is set, a user-defined metadata, transparent for the user-agent as the semantics of the metadata is user-specific. Living Standard Unimplemented ---------- Also for the cases other cases name and http-quiv and charset (attribute type), this labeling has not been used and more information is given for these case where as for the itemprop case no any other information or sample case is provided. (see examples in the bottom of the the spec) So my patch is written considering these 3 cases except itemprop. I am more than happy to include this case of itemprop if more information is available. :) [1]https://developer.mozilla.org/en-US/docs/Web/HTML/Element/meta Regards Nisala	1	0	https://issues.apache.org/jira/browse/ANY23-185
20075	ANY23-185	Hi [~nisala12] bq. One thing I have notice in the spec[1] itemprop is stated as "unimplemented", and I cant really find a reason behind this labeling. WDYT ? Right now I also don't have an idea however id the spec states that it is unimplemented then we can ignore this attribute from the <meta> node for the time being. I have tested this patch against master branch and will commit.	1	1	https://issues.apache.org/jira/browse/ANY23-185
20076	ANY23-177	Hey [~lanthaler], do you have any interest in working on this issue? If so I can assign it to you.	1	1	https://issues.apache.org/jira/browse/ANY23-177
20077	ANY23-177	[~lewismc], it's been a couple of years since I programmed in Java the last time and I'm not familiar with the code of Any23 either. [~p_ansell], could you look into this please? I think you are quite familiar with both code bases and already have a quite good idea of what would be needed. I don't how much work it would be, but it would perhaps also be interesting to extract JSON-LD embedded in HTML. See http://www.w3.org/TR/json-ld/#embedding-json-ld-in-html-documents Thanks!	1	0	https://issues.apache.org/jira/browse/ANY23-177
20078	ANY23-177	[~p_ansell] integrating this into the extractor/rdf package was pretty straight forward. Is there anything else you would want to see implemented here?	1	1	https://issues.apache.org/jira/browse/ANY23-177
20079	ANY23-155	Hi Debmalya, can you submit a patch we can apply against trunk? Thanks for opening this issue.	1	1	https://issues.apache.org/jira/browse/ANY23-155
20080	ANY23-155	Hi, Thanks for your mail. I will feel happy to apply the patch and work for Any23. I am new to this project. Let me know how to apply patch. Is there any link to follow? Please give me some time. Is there any date line for this? Best Regards, Debmalya Jash On Tue, Apr 2, 2013 at 3:47 AM, Lewis John McGibbney (JIRA) <jira@apache.org	1	1	https://issues.apache.org/jira/browse/ANY23-155
20081	ANY23-155	The only test that is failing for me right now in the core module is testBlankNodesViaURL, so the RDFSchemaUtilsTest is working for me. Could you post your java and maven versions (mvn -version and java -version) here	1	1	https://issues.apache.org/jira/browse/ANY23-155

#	Requirement	Questions	Answered	Successfully Answered	URL
20082	ANY23-155	Could you also post the output of: core/target/surefire-reports/org.apache.any23.util.DiscoveryUtilsTest.txt After you revert your changes to RDFSchemaUtilsTest and post the output of: core/target/surefire-reports/org.apache.any23.vocab.RDFSchemaUtilsTest.txt	1	1	https://issues.apache.org/jira/browse/ANY23-155
20083	ANY23-155	So RDFSchemaUtilsTest is working now for you after you reverted your changes to it?	1	1	https://issues.apache.org/jira/browse/ANY23-155
20084	ANY23-147	http://purl.org/dc/terms/ doesn't look to be deprecated. Are you thinking of http://purl.org/dc/elements/1.1/ ?	1	1	https://issues.apache.org/jira/browse/ANY23-147
20085	ANY23-147	Well its a redirect right? Which URI so we use as the NS?	1	1	https://issues.apache.org/jira/browse/ANY23-147
20086	ANY23-147	I am not sure what you mean. Everyone I know is still using http://purl.org/dc/terms/ for current applications, although older applications may still support http://purl.org/dc/elements/1.1/. The DCTERMS vocabulary in Any23 has a small subset of the http://purl.org/dc/terms/ properties, but it could definitely be improved, particularly to include the classes. Most people refer to the dc/elements/1.1 namespace using the prefix "dc", with the newer dc/terms either being "dcterms" or "dct". Is that part of the confusion?	1	1	https://issues.apache.org/jira/browse/ANY23-147
20087	ANY23-144	What do you suggest then Peter? I am happy to go with a positive consensus on this one.	1	1	https://issues.apache.org/jira/browse/ANY23-144
20088	ANY23-144	{bq}My preference is for all uppercase vocabulary class names if the commonly used prefix is an acronym, and camel-case where it is not, however, in some cases that distinction is not necessarily clear.{bq} OK so in an attempt to move towards a more appropriate (partial) solution... the DOAP vocab would remain named as DOAP and LIKF-Core would be changed to CamelCase e.g. LKIFCoreAction, etc... Can you confirm? It appears that this was not so much of a pressing issue before Simo pointed it out and you commented. This being said I appreciate that and I think it is important to get it right :) {bq}To clear up the naming I would prefer if we avoided the getInstance() pattern.... like access to the properties without storing them in a local private field to get around the getInstance() call.{bq} Excuse me if I am wrong, but I see this as a fundamental different issue from renaming of classes. Can we address this more directly in a related but separate issue?	1	1	https://issues.apache.org/jira/browse/ANY23-144
20089	ANY23-144	patch for trunk which addresses comments and concerns. If someone can review we can sort this one out. Thank you	1	1	https://issues.apache.org/jira/browse/ANY23-144
20090	ANY23-144	I can't seem to apply either of the patches using "patch -p0 < ANY23-144.patch" or similar for ANY23-144-v2.patch. Do you need to apply both of the patches to get it to work, because that wasn't working for me either but I may just be doing something wrong? What command have you been using the apply the patches?	0	0	https://issues.apache.org/jira/browse/ANY23-144
20091	ANY23-141	Expected release date of Sesame 2.7.0 has been set for about 3 weeks time. Is it worth waiting until then to release 0.8.0 or should we still focus on 0.9.0?	1	1	https://issues.apache.org/jira/browse/ANY23-141
20092	ANY23-141	Sesame-2.7.0 is now available, and the Git repository is now writable. What should we do about this issue?	1	1	https://issues.apache.org/jira/browse/ANY23-141
20093	ANY23-138	Hi [~simone.tripodi], can you provide some guidance as to this one? Do you have any 'typical' steps or guidance which we could work on? Thanks	0	0	https://issues.apache.org/jira/browse/ANY23-138
20094	ANY23-137	Hi Lev. You've marked this issue as being related to a number of other issues. Can you please provide info on how many of these have been addressed within the scope of your proposal? Thank you.	1	1	https://issues.apache.org/jira/browse/ANY23-137
20095	ANY23-137	I think it would be great to push on with this RDFa proposal then make an interim release for Any23. Is anyone interested in working to get this done? This issue (with a few others) would IMO justify pushing an RC for 0.9.0. Any thought please?	1	1	https://issues.apache.org/jira/browse/ANY23-137
20096	ANY23-137	Is there still any glue code that needs to be written to achieve this? or is ready for testing? Are there any instructions somewhere on how to build any23 with this patch so I can test it?	1	0	https://issues.apache.org/jira/browse/ANY23-137
20097	ANY23-137	What's the status of this issue to merge semargl into Any23? Is there anything I can test to make progress?	1	0	https://issues.apache.org/jira/browse/ANY23-137
20098	ANY23-137	Hi! I've prepared pull request to fix RDFa related tests. See https://github.com/apache/any23/pull/2 . Are there any other things to do with this branch?	1	0	https://issues.apache.org/jira/browse/ANY23-137
20099	ANY23-137	Success. I tested https://github.com/apache/any23/pull/2 on ubuntu 13.04 and it worked with java 1.7.0_51 and mvn 3.0.4! Not sure why it's not working on Mac OS X, do I need java 7?	1	0	https://issues.apache.org/jira/browse/ANY23-137
20100	ANY23-137	Thanks, Stephane! Completely missed that RDFa was used as a part of extraction process in other tests. I've added related fixes. Brief description. *ServletTest* Old RDFa implementation produces {{<issue level="Warning" row="14" col="5">Error while processing node /HTML(1)/HEAD(1)/META(9) : 'Cannot map prefix 'fb''</issue>}} while {{<fb:app_id>}} is completely valid predicate which shouldn't be resolved against fb: prefix. *Any23Test* *RoverTest* Changed RDFXMLWriter to NTriplesWriter in some tests to improve precision (they basically check line count). Changed expected triples count. It was reduced in most cases, because old RDFa parsed produced a lot of invalid triples like: {quote} <http://host.com/service> <http://host.com/serviceexternal> <http://host.com/service/ambiente/> . <http://host.com/service> <http://host.com/serviceexternal> <http://host.com/service/salute/> . <http://host.com/service> <http://host.com/serviceexternal> <http://host.com/service/legalita/> . <http://host.com/service> <http://host.com/serviceexternal> <http://www.ansamed.info/> . <http://host.com/service> <http://host.com/serviceexternal> <http://host.com/service/web/notizie/regioni/lazio/provinciadiroma/> . {quote} Fixed markup in {{test-resources/src/test/resources/html/rdfa/ansa_2010-02-26_12645863.html}} to conform declared XHTML 1.0 Strict. Fixed RDFa markup in {{test-resources/src/test/resources/html/encoding-test.html}} otherwise it shouldn't produce any triples. Disabled second part of {{Any23Test.testExtractionParameters}}. Should it do anything after RDFa parser replacement? Also, ExtractionException thrown from BaseRDFExtractor is escalated in test suite. It leads to some failed tests in Any23Test. What's the correct behaviour for ANY23 parser in case it gets SAXException?	1	0	https://issues.apache.org/jira/browse/ANY23-137
20101	ANY23-134	Will you have time to work on this before the 0.8.0 release or is it okay to postpone it to 0.9.0?	1	0	https://issues.apache.org/jira/browse/ANY23-134
20102	ANY23-127	[~lewismc] I'm failing to understand the point of this story. The purpose of the Semantic Web is that parsing tools like any23 should be vocabulary agnostic. Baking support for a whole load of vocab is by design out dated as soon as it ships, given that vocabularies tend to evolve (e.g. schema.org). Maybe if you explaine what you are trying to achieve I will be better able to help...	1	1	https://issues.apache.org/jira/browse/ANY23-127

#	Requirement	Questions	Answered	Successfully Answered	URL
20103	ANY23-127	.bq The purpose of the Semantic Web is that parsing tools like any23 should be vocabulary agnostic. Yes of course. I take your point. However if we are not namespace (or more appropriately vocabulary) aware then do we just _expect_ to recognize some namespace declarations? .bq ...is by design out dated as soon as it ships, given that vocabularies tend to evolve (e.g. schema.org). I need to disagree here. I would agree if you had added "given that _some_ vocabularies..." Some vocab's do evolve you are right. You've picked probably the most dynamic of all as your example however this ignores the point that _some_ are not as dynamic. For example, I added some Legal vocab's for a specific project I was working on. These may or may not be useful to other people. I also know that the vocabularly has NOT changed, therefore they are as useful now as they were when I wrote the implementation. If you navigate to the BBC vocabularies you will also see that these are well established and are not as dynamic as you may indicate. Furthermore, by implementing other developers' vocabularies we are directly inviting those same developers to use Any23 as a tool for extracting triples from their markup. Personally I've reached out the DOAP and LKIF communities to state that we are now namespace aware. This is the wider goal. I hope this provides some justification behind my reasoning on this one.	0	0	https://issues.apache.org/jira/browse/ANY23-127
20104	ANY23-121	Has this got anything to do with Any23?	1	1	https://issues.apache.org/jira/browse/ANY23-121
20105	ANY23-110	I honestly didn't realise that this was committed. Have you verified it works Peter? I will try tomorrow	1	0	https://issues.apache.org/jira/browse/ANY23-110
20106	ANY23-105	One fix for this may be to create a Sesame Sail implementation that can talk to TDB. Has anyone looked into this before?	1	1	https://issues.apache.org/jira/browse/ANY23-105
20107	ANY23-105	[~andy.seaborne] [~rvesse] [~p_ansell] what do you think is the best implementation here? I am _thinking_ towards a more abstract interface which wouldn't just restrict us to TDB. Is there anything out there which I could use?	1	1	https://issues.apache.org/jira/browse/ANY23-105
20108	ANY23-103	So this has been addressed right? It doesn't quite look like what you implemented in your commit @ r1371537 is what you initially proposed, but all seems well anyway.	1	1	https://issues.apache.org/jira/browse/ANY23-103
20109	ANY23-83	Having had a browse of both o.a.a.vocab.RDFSchemaUtils, and o.a.a.rdf.RDFUtils respectively, I understand where you are coming from. I also notice that you propose a more generic factory type implementation similar to Tika which (based on some criteria) guesses the specific parser which should be used per situation. What other mimetypes would we encounter other than the ones mentioned in both of the above classes as specified within the hardcoded RDFParsers? I am interested to find out what you think. Thank you	1	1	https://issues.apache.org/jira/browse/ANY23-83
20110	ANY23-83	Hi Peter, {bq} I am not familiar enough with the rest of the codebase to know why the shortlist of supported parsers is in place [bq} Although I have tried to get to grips with as much of the codebase as I have a use-case for, I am sorry as I cannot provide insight into the original design considerations which you refer to... maybe one of the other guys could chime in here and clear the picture for us? I must admit though, I am in agreement with your overall analysis and think that the more extensible we can make the Any23 libraries the better. There was some discussion however over in ANY23-19 where the vision was to have an abstraction layer between Any23 and the underlying RDF API's. In my opinion, this would provide far greater potential for us to make Any23 a more flexible library... wdyt?	1	1	https://issues.apache.org/jira/browse/ANY23-83
20111	ANY23-83	I like these two patches and based on the conversation on any23-dev give the +1. Marking for 0.7.0 and happy to commit on behalf of Peter unless there are objections?	0	0	https://issues.apache.org/jira/browse/ANY23-83
20112	ANY23-83	I rebased my git repository onto a copy of the current trunk to check how many patches I still had in this area and noticed that RDFSchemaUtils still hardcodes the supported formats. I may not be seeing the significance of VocabularyFormat currently, but I assume that it could be generalised to any RDFFormat. Is it related to the ability to type in a human readable alias for the format in VocabPrinter? If so, have you tried using RDFFormat.valueOf(String) in the converter inside VocabPrinter which matches against registered formats using: format.getName().equalsIgnoreCase(formatName) These names may not exactly match the historical names for each format that were previously used, but they are extensible at runtime so it makes RDFSchemaUtils easier to maintain. This may also help with ANY23-3 as N3 would be available, although it would just be the Turtle subset in practice as we only hold RDF triples in memory.	1	0	https://issues.apache.org/jira/browse/ANY23-83
20113	ANY23-75	Setting for 0.7.0-incubating. Is it possible for you to explain a bit about the patch and the underlying reason as to why the existing parser implementation seems to clog up? Also this is really trivial but can you please have a look at the coding format if it differs or not. My initial thoughts of the patch are great, it's a nice one to have caught, but some additional explanation would really help us out. Thank you very much. Lewis	1	1	https://issues.apache.org/jira/browse/ANY23-75
20114	ANY23-70	Yep, unfortunately, as Lewis put it they are "horrible for searching". I usually end up doing searches elsewhere. Is there a way to search on the Apache mail archive?	0	0	https://issues.apache.org/jira/browse/ANY23-70
20115	ANY23-69	Does this duplicate ANY23-7 ?	1	0	https://issues.apache.org/jira/browse/ANY23-69
20116	ANY23-69	Yeah OK thnaks for clarification. The W3C guys that chimed in on the thread on any23-dev, stated that they would like us to produce this by the end of the month so we have plenty of time to cook something up. I had a look into the RDFa 1.1. spec, and subsequently into the EARL spec... unsurprisingly it's left me a bit puzzled as to how to pull this together so I was going to start a thread on our dev list. However this is as good a place as any I suppose. Have you any initial thoughts Michele?	1	0	https://issues.apache.org/jira/browse/ANY23-69

#	Requirement	Questions	Answered	Successfully Answered	URL
20117	ANY23-65	Well, last night I found that no, it isn't working flawlessly. I was just looking at the whole thing again to trace what is going wrong. I was about to send my XSLT file to the Sindice-dev mailing list, to bridge the period between now and the moment Sindice starts using Any23 0.7.0, when I thought of an edge case to test against. My stylesheet doesn't handle it well, and neither does my build from the SVN code. The case: when a prefix attribute contains "rnews:http://www.iptc.org/std/rNews/1.0: foaf:http://xmlns.com/foaf/0.1/ dbpedia:http://dbpedia.org/resource/" (no space between prefix name and prefix URI, but at least one prefix URI ending with a colon), my stylesheet wrongly assumes there are spaces between prefix name and prefix URI, because it tests whether the attribute contains ": ". The RDFa11Parser outputs warnings that it cannot map the foaf prefix when I test it on my 0.7.0 build. It is an edge test case because it is invalid content for a prefix attribute, but since I saw that Any23 is accepting / testing against no-space prefix definitions and I think namespace URIs with a colon are valid (e.g. <http://dbpedia.org/resource/Category:>), I figured it made an interesting test case. I sent the file anyway, with this test result too :) Another issue I have is that the RDFa11Parser doesn't infer the right triples from <link> elements with a @rel in the <head> section. I believe the @rel values "icon", "stylesheet", "bookmark" etc are to be treated specially. Sindice (Any23 0.6.1) produces URIs for these like <http://www.w3.org/1999/xhtml/vocab#icon> from my blog post. When I extract from my blog post locally, I get properties like <http://ben.companjen.name/2011/08/het-gezin-timmer-de-bruijn-in-amsterdam/icon>. The parser (also) complains my HTML doesn't declare 'xmlns="http://www.w3.org/1999/xhtml"'. I couldn't find the part of the HTML or RDFa specifications that says it should do so - as far as I can tell it's not necessary in HTML5. Looking at the code, I see some references to RDFa 1.0 in the comments in the processDocument method. This method seems to be the source for the complaint. Maybe the problem and complaint (well, warning actually) are linked? And maybe the incorrect handling of the @rel values is also linked to "// TODO: introduce support for RDFa profiles. (http://www.w3.org/TR/rdfa-core/#s_profiles)"? BTW, these profiles don't exist (anymore) in RDFa Core. I hope someone can shed more light on this, I'm too confused from reading all the RDFa related docs and drafts right now ;)	1	1	https://issues.apache.org/jira/browse/ANY23-65
20118	ANY23-61	The repo directory is a maven repository. Are the artifacts in this directory in another maven repository? We are currently deploying snapshots to the Apache Snapshots repository and new releases to the Apache staging server, based on the definitions in the apache-10.pom parent. http://svn.apache.org/repos/asf/maven/pom/tags/apache-10/pom.xml	1	0	https://issues.apache.org/jira/browse/ANY23-61
20119	ANY23-50	i had a look to the patch and it looks very nice. If the META-INF data is built at compile time and it is not necessary to package the plugins to annotate them I don't have any other concern about applying this simplification. How does it work /does it work the hot-loading with ServiceLoader?	0	0	https://issues.apache.org/jira/browse/ANY23-50
20120	ANY23-50	Hola! META-INF/services metadata are always generated at compile time, indeed plugins are annotated with [@MetaInfService|http://weblogs.java.net/blog/kohsuke/archive/2009/03/my_project_of_t.html] that is part of a provided-scope Annotation Processor wrote by Kohsuke (the Hudson/Jenkin's father) In Any23 the ServiceLoader is used inside the the Any23PluginManager, relying on the already managed dynamic ClassLoader, so when adding new jars/resources at runtime to the ClassLoader, SL queries it. I think that the same strategy can be applied to Extractors' factories as well, WDYT? Anyway, please double-ack if you agree I can commit the modification! TIA!	1	1	https://issues.apache.org/jira/browse/ANY23-50
20121	ANY23-50	I think this should definitely be commited? Once question, is there anything specific about the Any23 plugin system which make this issue only applicable to Any23? What I mean, is would it be possible to aplly this issue in a project elsewhere which relies heavily on plugins? (e.g. Nutch)	1	1	https://issues.apache.org/jira/browse/ANY23-50
20122	ANY23-48	OK so Committed @ revision 1303963 in site I also committed some changes to trunk but I can't remember which revision. If someone could please review where we are with this one, unless someone can point me to further improvement (which obviously I'm very open to) then I'd like to close this one off so we can focus on some of the other issues which lie between us and the 0.7.0-incubating release. Thanks	1	1	https://issues.apache.org/jira/browse/ANY23-48
20123	ANY23-46	Thanks Andy. Taking this forward, obviously we can copy in Any23-dev@ anyway, but do we wish to combine these resources requests to infra@, e.g. is there any way that we can share resources to limit the infra@ involvement if you were wanting to bring the SPARQL stuff to ASF as well?	1	1	https://issues.apache.org/jira/browse/ANY23-46
20124	ANY23-38	@Lewis, you mean org.apache.any23.util.LogUtils or what? That one is used to activate / deactivate the verbose mode of CLI tools by changing the log level programmatically.	1	1	https://issues.apache.org/jira/browse/ANY23-38
20125	ANY23-38	Upon closer inspection looks like your correct Michele. What about we re-factor this issue to address Simone's opinions? I'm a +1 for Slf4j	1	0	https://issues.apache.org/jira/browse/ANY23-38
20126	ANY23-36	I genuinely have not got a clue what has to be done to resolve this one. I've tried using the maven changes plugin but getting nowhere. Anyone else got a scooby?	0	0	https://issues.apache.org/jira/browse/ANY23-36
20127	ANY23-36	Isn't this issue related to Any23 history munging?	0	0	https://issues.apache.org/jira/browse/ANY23-36
20128	ANY23-33	I attempted to replace the classes with public implementations, but the once found within the classpath are marked as @deprecated. Ideas? {code} -import com.sun.org.apache.xml.internal.serialize.OutputFormat; -import com.sun.org.apache.xml.internal.serialize.XMLSerializer; +import org.apache.xml.serialize.OutputFormat; +import org.apache.xml.serialize.XMLSerializer; {code}	1	1	https://issues.apache.org/jira/browse/ANY23-33
20129	ANY23-33	Are you getting this from your compiler output? I've not tried to implement yet. We don't explicitly depend upon any xerces related stuff in the parent pom.xml, so I am curious where the imports are being pulled from for you to get this output.	0	0	https://issues.apache.org/jira/browse/ANY23-33
20130	ANY23-32	Thanks Mic for reporting! I need indeed doing more work on the service, my bad that I didn't notice! there is a specialized README in the produces distribution package that explains how to install&run, can you double chek please? TIA!	0	0	https://issues.apache.org/jira/browse/ANY23-32
20131	ANY23-26	Hi Peter. Regarding your edited comment above. Have you started/finished work on the Extractor factory patterns? I would like to get this issue integrated into trunk... therefore I need to test it. If the answer to the above is yes, are you able to produce a patch which I can apply to the SVN trunk HEAD? Thank you	1	1	https://issues.apache.org/jira/browse/ANY23-26
20132	ANY23-26	Has anyone else had the opportunity to test Peters patches? There are some trivial issues which we have discussed off list but generally speaking the spi and tika upgrades look dynamite to me. I also confirm that all tests pass locally for these patches.	1	1	https://issues.apache.org/jira/browse/ANY23-26
20133	ANY23-21	Is there an automated tool for doing this?	0	0	https://issues.apache.org/jira/browse/ANY23-21
20134	ANY23-19	Hey Paolo, with regards to Chris' recent comments on the dev list and the stub report he filed, how did you get on wrt moving away from the packages you've highlighted above and also with progressing with more collaboration between ASF RDF-type projects? I'm using this issue as a basis to kick the converation off and hopefully others will join in.	0	0	https://issues.apache.org/jira/browse/ANY23-19

#	Requirement	Questions	Answered	Successfully Answered	URL
20135	ANY23-19	Hi Max, the thing is: adding yet another dependency is going to make Any23 bigger and more complex. Can something simple be done without adding yet another dependency? As I said, an example of a small RDF distiller which does this is java-rdfa (https://github.com/shellac/java-rdfa). Third party RDF APIs become only an optional dependency and people are free to choose which one they want/need simply adding that to their classpath (in a very similar way to what SLF4j does for logging frameworks). The only thing tricky is "validation", but Any23 could chose not to validate URLs/IRIs, etc. and simply relay on the third party RDF APIs to do that.	1	1	https://issues.apache.org/jira/browse/ANY23-19
20136	ANY23-19	Hi Paolo, The example library that is continually being referred to here, java-rdfa, abstracts away from clerezza, sesame and jena interfaces by representing everything as strings. See https://github.com/shellac/java-rdfa/blob/master/core/src/main/java/net/rootdev/javardfa/StatementSink.java The java-rdfa library is also not being developed or used actively, so it isn't the best example of a successful library that doesn't use a type-safe RDF Statement/Value API. I worked on java-rdfa a little but there is no push behind it. Even its author refers to it as "The cruftiest RDFa parser in the world" Two of the three goals of Any23, command line utility, and web service, are completely ambivalent to the technology being used internally, as long as it is high quality, and the Sesame libraries are very high quality in my opinion and experience. The only goal that would be affected would be the use of Any23 as a library. How often is Any23 currently being used as a library? Could another project easily implement the same functionality using Jena with less effort than it would take to either create a custom string based solution, or move to another framework that may have just as many or more dependencies? In terms of the use of Any23 as a library, is there anything about the Sesame Model hierarchy (Value/Resource/Literal/BNode/URI) that would be better represented using a custom solution? As one example, I have been working with OWLAPI recently and its RDF handling is shocking, it merges URIs with Blank Nodes to form what it refers to as IRIs. It has a custom internal solution that only recognises two types of triples, those with an IRI in the object position (where IRI is not type-safely defined between URI and BlankNode) and those with Literals in the object position. I can't imagine Any23 going down this route, but it is the worst case scenario if the API is converted without a reason. In the simplest scenario, it may be possible to reuse the Sesame Model hierarchy to produce Values that work across all three libraries, using custom ValueImpl etc., implementations that actually implement the relevant interfaces from other libraries, along a custom ValueFactory to produce these multi-library-compatible Values (custom ValueFactories can be plugged into any Rio Parser using Rio.getParser(RDFFormat, ValueFactory), a functionality which I haven't seen in other libraries). In terms of the actual packages that are currently used, there are four basic packages sesame-model, sesame-rio-api, sesame-repository-api, sesame-sail-api, sesame-sail-memory. These base libraries are small dependencies. One other dependency is some small utilities, sesame-util that are used by sesame-model and other sesame libraries. 82K - sesame-model-2.6.5.jar 36K - sesame-repository-api-2.6.5.jar 22K - sesame-rio-api-2.6.5.jar 56K - sesame-sail-api-2.6.5.jar 54K - sesame-sail-memory-2.6.5.jar 53K - sesame-util-2.6.5.jar The value Impl classes should not be directly referenced. They should be accessed using a ValueFactory and used as their Interfaces. This doesn't change any of the libraries that are used, but it is better practice. The other libraries that pull in the Rio parsers can be linked in dynamically without compiling in the dependency, so the use of Any23 as a library would enable people to pull them in as needed. See Rio.getParser(RDFFormat, ValueFactory) and Rio.getWriter(RDFFormat) methods. It would be valuable if Any23 could dynamically pull in all of its parsers and writers using the Rio.* static methods. Then it could be used with the absolute minimum number of parsers and writers for the current user. 4.6K - sesame-rio-n3-2.6.5.jar 14K - sesame-rio-ntriples-2.6.5.jar 33K - sesame-rio-rdfxml-2.6.5.jar 17K - sesame-rio-turtle-2.6.5.jar Switching to another library may cause the bloat that you say you do not want. For example, Jena and its immediate dependencies is quite large compared to the modular sesame jar files, and that doesn't include the SPARQL parsing libraries from ARQ, as indeed the sesame libraries quote above do not include the sparql libraries. 1.7M - jena-core-2.7.0-incubating.jar 151K - jena-iri-0.9.0-incubating.jar 3.1M - icu4j-3.4.4.jar 1.4M - xercesImpl-2.10.0.jar	1	1	https://issues.apache.org/jira/browse/ANY23-19
20137	ANY23-18	Ciao Michele, what are the problems you found with java-rdfa? There is an open issue for java-rdfa to add RDFa 1.1 support (https://github.com/shellac/java-rdfa/issues/24). In any case, good to see the XSLT replaced by a new programmatic RDFa extractor.	1	0	https://issues.apache.org/jira/browse/ANY23-18
20138	ANY23-11	I've updated the w3c page above, but can't seem to create an account for the RDFizers!!! Is there anything else we wish to add to this issue? I think it goes to say that at any given opportunity make people aware of Any23.	1	1	https://issues.apache.org/jira/browse/ANY23-11
20139	ANY23-4	I think this certainly duplicates ANY23-69 ?	1	1	https://issues.apache.org/jira/browse/ANY23-4
20140	ANY23-3	Is this not addressed?	1	1	https://issues.apache.org/jira/browse/ANY23-3
20141	ANY23-3	Are there any full N3 implementations in Java? Sesame and Jena both parse the RDF subset, but not the full N3 grammar. In practice, the RDF subset amounts to Turtle. Is Any23 interested in the non-RDF components of N3?	1	1	https://issues.apache.org/jira/browse/ANY23-3
20142	ANY23-2	@Peter, Are you interested in committing this in line with thr new SPI Factory implementation you wrote?	1	1	https://issues.apache.org/jira/browse/ANY23-2
20143	APEXCORE-764	Could you please add some more description as to what the refactoring entails?	1	0	https://issues.apache.org/jira/browse/APEXCORE-764
20144	APEXCORE-735	Could you add the reason for the change.	0	0	https://issues.apache.org/jira/browse/APEXCORE-735

#	Requirement	Questions	Answered	Successfully Answered	URL
20145	APEXCORE-734	core/pull/532#discussion_r127563982 ########## File path: engine/src/main/java/com/datatorrent/stram/StramLocalCluster.java ########## @@ -463,103 +463,123 @@ public void run() @SuppressWarnings({"SleepWhileInLoop", "ResultOfObjectAllocationIgnored"}) public void run(long runMillis) { - if (!perContainerBufferServer) { - StreamingContainer.eventloop.start(); - bufferServer = new Server(StreamingContainer.eventloop, 0, 1024 * 1024, 8); - try { - bufferServer.setSpoolStorage(new DiskStorage()); - } catch (IOException e) { - throw new RuntimeException(e); + Thread eventLoopThread = null; + List<Thread> containerThreads = new LinkedList<>(); try { + if (!perContainerBufferServer) { + eventLoopThread = StreamingContainer.eventloop.start(); + bufferServer = new Server(StreamingContainer.eventloop, 0, 1024 * 1024, 8); + try { + bufferServer.setSpoolStorage(new DiskStorage()); + } catch (IOException e) { + throw new RuntimeException(e); + } + bufferServerAddress = InetSocketAddress.createUnresolved(LOCALHOST, bufferServer.run().getPort()); + LOG.info("Buffer server started: {}", bufferServerAddress); } - bufferServerAddress = InetSocketAddress.createUnresolved(LOCALHOST, bufferServer.run().getPort()); - LOG.info("Buffer server started: {}", bufferServerAddress); - } - long endMillis = System.currentTimeMillis() + runMillis; - while (!appDone) { + long endMillis = runMillis > 0 ? System.currentTimeMillis() + runMillis : Long.MAX_VALUE; + while (!appDone && System.currentTimeMillis() < endMillis) { - for (String containerIdStr: dnmgr.containerStopRequests.values()) { - // teardown child thread - StreamingContainer c = childContainers.get(containerIdStr); - if (c != null) { - ContainerHeartbeatResponse r = new ContainerHeartbeatResponse(); - r.shutdown = StreamingContainerUmbilicalProtocol.ShutdownType.ABORT; - c.processHeartbeatResponse(r); + for (String containerIdStr : dnmgr.containerStopRequests.values()) { + // teardown child thread + StreamingContainer c = childContainers.get(containerIdStr); + if (c != null) { + ContainerHeartbeatResponse r = new ContainerHeartbeatResponse(); + r.shutdown = StreamingContainerUmbilicalProtocol.ShutdownType.ABORT; + c.processHeartbeatResponse(r); + } + dnmgr.containerStopRequests.remove(containerIdStr); + LOG.info("Container {} restart.", containerIdStr); + dnmgr.scheduleContainerRestart(containerIdStr); + //dnmgr.removeContainerAgent(containerIdStr); } - dnmgr.containerStopRequests.remove(containerIdStr); - LOG.info("Container {} restart.", containerIdStr); - dnmgr.scheduleContainerRestart(containerIdStr); - //dnmgr.removeContainerAgent(containerIdStr); - } - // start containers - while (!dnmgr.containerStartRequests.isEmpty()) { - ContainerStartRequest cdr = dnmgr.containerStartRequests.poll(); - if (cdr != null) { - new LocalStreamingContainerLauncher(cdr, containerThreads); + // start containers + while (!dnmgr.containerStartRequests.isEmpty()) { + ContainerStartRequest cdr = dnmgr.containerStartRequests.poll(); + if (cdr != null) { + new LocalStreamingContainerLauncher(cdr, containerThreads); + } } - } - if (heartbeatMonitoringEnabled) { - // monitor child containers - dnmgr.monitorHeartbeat(false); - } - if (childContainers.isEmpty() && dnmgr.containerStartRequests.isEmpty()) { - appDone = true; - } + if (heartbeatMonitoringEnabled) { + // monitor child containers + dnmgr.monitorHeartbeat(false); + } - if (runMillis > 0 && System.currentTimeMillis() > endMillis) { - appDone = true; - } + if (childContainers.isEmpty() && dnmgr.containerStartRequests.isEmpty()) { + appDone = true; + } - try { - if (exitCondition != null && exitCondition.call()) { - if (runMillis > 0 && System.currentTimeMillis() > endMillis) { - appDone = true; - } - } catch (Exception ex) { - break; - } - if (Thread.interrupted()) { - break; + try { + if (exitCondition != null && exitCondition.call()) { + appDone = true; + } + } catch (Exception ex) { + break; + } + if (Thread.interrupted()) { + break; + } } + } finally { + if (eventLoopThread != null) { + StreamingContainer.eventloop.stop(eventLoopThread); + } + containerThreads.forEach(Thread::interrupt); }	0	0	https://issues.apache.org/jira/browse/APEXCORE-734
20146	APEXCORE-725	PR 624 is actually unrelated to this ticket; since PR 529 is now merged, shouldn't this ticket be resolved ?	0	0	https://issues.apache.org/jira/browse/APEXCORE-725
20147	APEXCORE-714	Thomas, My responses are inline.. On Tue, Apr 25, 2017 at 8:42 AM, Thomas Weise <thw@apache.org> wrote: Pramod, Sounds like some sort of alternative "processing mode" that from engine perspective allows potentially inconsistent state when there is a pipeline failure. This is of course only something the user can decide. >> Calling it an alternate processing mode is a good idea. Does the proposal assume that the operator state is immutable (or what is sometimes tagged with the stateless annotation)? For example an operator that has to load a large amount of state from another source before it can process the first tuple? >> Operator state can change, not necessarily stateless. Stateless may not automatically fall into this category as our current definition of stateless denotes window level stateless and not necessarily tuple level. Also, it would be an optimization but not something that will help with SLA if the operator still needs to be recovered when its own container fails. It might help to clarify that and also why there is a need to recover in the batch use case (vs. reprocess). >> Correct, as I mentioned in the last statement, if the container where the operator is running itself goes down then it is recovery from checkpoint and business as usual. I may have misspoken about batch, meant to say apps where operators have large state not necessarily batch, the use case we are dealing with is batch and restart is not practical as the run takes a long time. Thanks, Pramod	1	0	https://issues.apache.org/jira/browse/APEXCORE-714
20148	APEXCORE-712	I think you mean to have SSL_CONFIG as a more complex object like a POJO instead of a filename that refers to a file that contains all ssl related properties. That works for me in terms of more direct way to specifying these properties. Why would we need the long naming, using hadoop specific names, for the properties of that object. Why can't they be specified with regular field names like keystoreLocation, keystorePassord, keyPassword etc. This way the user does not need to be aware of the actual property names being used by hadoop and is more portable in case we support another runtime in the future. For convenience, for the user, the entire object with the field values can be specified as a single String property as defined by the StringCodec.Object2String class. I would prefer a more targeted name like MANAGEMENT_SSL_CONFIG as I am not sure if this would be enough to support any SSL configuration needed in future (like if this is used for encryption data flow itself). But, if you want to go with SSL_CONFIG I would be ok.	1	0	https://issues.apache.org/jira/browse/APEXCORE-712
20149	APEXCORE-711	[~sanjaypujare] There seems to be a strong dependency on Hadoop version where YARN-6457 will be fixed. As YARN-6457 by itself seems to depend on YARN-4562 that is only fixed in 2.9, will the fix require Hadoop 2.9 as well?	1	0	https://issues.apache.org/jira/browse/APEXCORE-711
20150	APEXCORE-711	[~sanjaypujare] While fix does not introduce a regression, it introduces an attribute that has no meaning unless a user upgrades to Hadoop 2.9. I'd like to see if there is a solution that does not depend on Hadoop upgrade. Additionally, how was it prototyped without YARN-6457 being fixed?	1	0	https://issues.apache.org/jira/browse/APEXCORE-711
20151	APEXCORE-711	I'd prefer not to rely on YARN-6457 and YARN-4562 be ported to older versions of hadoop-yarn-common. YARN-4562 is marked as resolved on 2.9 and 3.0, so it does not look that it will eventually make into 2.7.1. There is still no agreement on how YARN-6457 will be fixed. Instead of relying on YARN-6457, can we implement a class that extends org.apache.hadoop.yarn.webapp.WebApps or provides similar functionality and that does not have the same limitation as YARN-6457?	1	0	https://issues.apache.org/jira/browse/APEXCORE-711

#	Requirement	Questions	Answered	Successfully Answered	URL
20152	APEXCORE-711	Recapping my earlier analysis about eliminating dependencies on Hadoop/Yarn: - In org.apache.hadoop.yarn.webapp.WebApps.Builder<T> the start(WebApp) method does bulk of the work in creating the WebApp container object. The problematic line in this method currently is WebAppUtils.loadSslConfiguration(builder); which does not pass (honor) the config object passed to Builder.with(Configuration). The above line should be changed to WebAppUtils.loadSslConfiguration(builder, conf); which passes the required config object. Since a class static method (WebAppUtils.loadSslConfiguration) is being called there is no way to get the new behavior by subclassing a class and overriding a simple method. All of the code in org.apache.hadoop.yarn.webapp.WebApps.Builder.start(WebApp) will need to be copied in a new class on the Apex core side. BTW, precisely because of this, https://issues.apache.org/jira/browse/YARN-4562 was created and fixed (albeit in 2.9.x and not in any earlier version) - with this we still have an unresolved issue in the current Yarn. The current code in org.apache.hadoop.yarn.webapp.util.WebAppUtils.loadSslConfiguration(Builder, Configuration) is as follows (trimmed): public static HttpServer2.Builder loadSslConfiguration(HttpServer2.Builder builder, Configuration sslConf) { if (sslConf == null) { sslConf = new Configuration(false); } boolean needsClientAuth = YarnConfiguration.YARN_SSL_CLIENT_HTTPS_NEED_AUTH_DEFAULT; sslConf.addResource(YarnConfiguration.YARN_SSL_SERVER_RESOURCE_DEFAULT); return ... } which has 2 major issues: a) The caller can set SSL properties in the passed sslConf object but they are overridden by the default ssl-server.xml file because of the sslConf.addResource(...) call. There is no way to point to a different XML file (because it uses the hardcoded YARN_SSL_SERVER_RESOURCE_DEFAULT value and not the value of a property) which defeats the very purpose. b) the only way to solve (a) above is to mark the value final in the passed sslConf object in which case the default ssl-server.xml values won't be applied even if they are final there. Both of these are addressed by the changes in this JIRA.	1	0	https://issues.apache.org/jira/browse/APEXCORE-711
20153	APEXCORE-699	I guess that the proper way to support compatible Plugin interface modifications and versioning is to extend an existing interface. For example, if a new version of a Plugin introduces new method foo(), the new version should be defined as {code} interface PluginV2 extends Plugin { void foo(); } {code} On the other side, plugins should be able to check engine version and prevent the platform from loading plugin if the platform does not meet version requirements. Such information can be passed using PluginContext and does not require Plugin interface changes. With the proposed design, will it be OK to exclude Plugin versioning changes from 3.6.0 release?	1	0	https://issues.apache.org/jira/browse/APEXCORE-699
20154	APEXCORE-694	jsr305 is dormant, I see musings on the web that it never became a standard. Can you verify?	1	0	https://issues.apache.org/jira/browse/APEXCORE-694
20155	APEXCORE-673	Would this be to a public or some internal location? Don't you want to push to public location only on releases?	1	0	https://issues.apache.org/jira/browse/APEXCORE-673
20156	APEXCORE-670	Ah, of course the _target_ needs to be included in there as well. In my understanding, your cURL request would change the log level of the classes used by the stram? This would differ from setting the log levels in the operators, right? Example: {code} curl -H "Content-Type: application/json" -X POST \ -d '{"loggers":[{"target":"org.apache.apex.examples.pi.*","logLevel":"TRACE"}]}' \ http://apex-sandbox:9090/ws/v2/applications/application_1499808956620_0803/loggers {code} In this case, should we use two separate commands for setting the log levels in stram / operators?	1	1	https://issues.apache.org/jira/browse/APEXCORE-670
20157	APEXCORE-670	So an example for the new command in the CLI could be: {code}apex>set-log-level application_1499808956620_0873 org.apache.apex.examples.pi.PiCalculateOperator TRACE{code} OR {code}apex (application_1499808956620_0873) >set-log-level org.apache.apex.examples.pi.PiCalculateOperator TRACE{code} I will open a pull request with an exemplary implementation of this. Should the validation of the logLevel happen on the APEX CLI side? Currently it will just ignore any logLevel that is not valid	1	1	https://issues.apache.org/jira/browse/APEXCORE-670
20158	APEXCORE-670	Do other commands that interact with applications accept an application parameter? If not, then we should not introduce it here either.	1	0	https://issues.apache.org/jira/browse/APEXCORE-670
20159	APEXCORE-635	The end goal of the override BufferServerPublisher.write() in is not clear. Why is the override necessary? It is also not clear why Kryo (that is a library designed to do efficient serialization and de-serialization) will not match the proposed mechanism. If the goal is to serialize Collections (list in the example) using Kryo, Kryo has CollectionsSerializer and will (should?) do something similar to what is proposed. IMO, it will be much better not to even use Kryo when Tuples are strongly typed and implement StreamCodec without using Kryo. Default StreamCodec is provided to handle cases when Tuples are not strongly typed, in which case it is required to use reflection. The same will apply to unchecked Collections. (n) till all those issues are resolved.	1	0	https://issues.apache.org/jira/browse/APEXCORE-635
20160	APEXCORE-635	- Memory copy is required as it moves data from JVM heap to direct (off-heap) buffer. It is done prior to writing the buffer to a network channel. - Please benchmark copying one large block of memory in a single step and in multiple steps. What is the difference (in %)? Do the same for a direct buffer. - How will control tuples be injected into a continuous block of memory? - The current approach is fire and forget, the proposed requires some kind of an acknowledgment. I am not convinced that the design that requires an acknowledgment will be faster. - There is a plan to move to netty from netlet and netty provides buffer pool allocator that may better fit the proposed design. - The current format uses variable length (1, 2, 3 or 4 bytes) prefix and it is not possible to know offset where data needs to be written prior to the lenght of the serialized data is known.	1	0	https://issues.apache.org/jira/browse/APEXCORE-635

#	Requirement	Questions	Answered	Successfully Answered	URL
20161	APEXCORE-635	- If we use direct buffer, it seems ByteBuffer has no way to avoid memcopy. What I meant to avoid memcopy and garbage collection is explained as following by comparing details of current and new procedure(assume each tuple serialized to N bytes, and write M tuples in one time): current procedure(use kryo serde as example): 1) kryo serde: write data to buffer of kryo's output, assume the length of bytes is N 2) call Output.toBytes() to get bytes from kryo: need memcopy, copy N bytes 3) put message type and partition: need memcopy, copy N bytes, the total length become N+5; N bytes need to be garbage collected 4) repeat 1 to 3 M times, and generate M slices 5) copy the M slices to ByteBuffer: M*(N+5) bytes need to be garbage collected 6) write ByterBuffer to socket new procedure(use kryo serde as example): 1) write message type and partition to SerializationBuffer 2) kryo serde: serde to SerializationBuffer, assume the length of bytes is N 3) call SerializationBuffer.toSlice() to get the slice: return the wrapper instead copy the data 4) repeat 1 to 3 M times, and generate M slices 5) merge the M slices into one in most of case(if these slices in same block). If in different block, probably handle one block each time. 6) copy or wrapper the to ByteBuffer. Probably need copy if want to use direct mode. 7) write ByterBuffer to socket. reset the SerializationBuffer So, by comparing current and new procedure, current procedure need to copy around 3*N*M bytes and garbage collection of 2*N*M bytes, but new procedure only need copy N*M bytes and no need garbage collection. So new approach save 2*N*M memory copy and garbage collection. - The memory of different blocks is not continuous as we want dynamic allocate block, but the slices in same block are continuous. As the size of block is much large than slices. - The current fire and forget approach means no way to reuse the memory. That's why need to override write() method by walk around way. The new implementation of write() will reset the SerializationBuffer after wrote data to socket. The new approach don't need acknowledgement neither. - Yes, for variable length, it is not possible to reserve the space for length. But I doubled why we have to use variable length. It save at most 3 bytes for each message, but made parser harder. It's not common for protocols to use variable length. Even IP the fixed length. In worst case, even if we don't use the reserve feature of SerializationBuffer, we still can save 2*N*M memory copy and garbage collection.	1	1	https://issues.apache.org/jira/browse/APEXCORE-635
20162	APEXCORE-635	[~brightchen] What is the overhead of garbage collection? You can take the benchmark application in the Malhar for measurements.	1	0	https://issues.apache.org/jira/browse/APEXCORE-635
20163	APEXCORE-635	Hi Vald, - As I understand, array is a kind of object in java. Does java implemented differently? And how can JVM reuse allocated bytes without collection? Do you have any reference I can take a look? - I copied the implementation of kryo Output.toBytes() here (kryo-2.24.0.jar). From this implementation toBytes() do allocate new memory and copy. public byte[] toBytes () { byte[] newBuffer = new byte[position]; System.arraycopy(buffer, 0, newBuffer, 0, position); return newBuffer; } - I can get number, but that means need at least simply prototype. And based on my previous test for serialization string of 1000 ascii, The performance can gain 30% - Yes, it's pretty complex to implement all the proposal. But fortunately, most of memory manage stuff already implemented in Malhar. The things need to do is add some feature and integrate to BufferServer	1	1	https://issues.apache.org/jira/browse/APEXCORE-635
20164	APEXCORE-635	Hi Vlad, I think what you meant the issues I not addressed is a following: - Please benchmark copying one large block of memory in a single step and in multiple steps. What is the difference (in %)? Do the same for a direct buffer. - How will control tuples be injected into a continuous block of memory? my comments as following: - From my previous comments, we can the problem of current mechanism is waste 2 times memory copy and garbage collection. It's not the problem of copy one big bytes in one time or divided into multiple copies. - The control tuples also write to the same SerializationBuffer.	1	1	https://issues.apache.org/jira/browse/APEXCORE-635
20165	APEXCORE-635	bq. As I understand, array is a kind of object in java. Does java implemented differently? And how can JVM reuse allocated bytes without collection? Do you have any reference I can take a look? Yes, an array is an object in java that may hold either primitive types or other objects. When an array is GC, it is GC as a single object. If an array holds primitive types, the array's memory is reclaimed in a single step, there is no GC of individual bytes/int/long in the array. bq. I copied the implementation of kryo Output.toBytes() here (kryo-2.24.0.jar). From this implementation toBytes() do allocate new memory and copy. I was referring to Output.getBuffer() and DefaultKryoStreamCodec that uses Output.getBuffer(). You are right that Output.toBytes() allocates a new array and copies bytes from Output internal buffer to the newly allocated buffer. This is a bug in DefaultStatefulStreamCodec, it should use Output.getBuffer(). bq. I can get number, but that means need at least simply prototype. And based on my previous test for serialization string of 1000 ascii, The performance can gain 30% As I pointed to you, the prototype has a significant bug (reusing byte arrays that are not reusable due to fire and forget netlet behavior). What I would like to see is a series of small prototypes that prove that the approach provides significant performance gains. I don't see how without modifying the current definition of StreamCodec and StatefulStreamCodec the solution can be generalized. It may be possible to improve DefaultStatefulStreamCodec (concrete implementation of StatefulStreamCodec), but it seems to require a large amount of modifications to netlet and buffer server that other StreamCodec implementations will not benefit from. Additionally, the proposal seems to require netlet client to be aware of codec implementation and any changes to tuple format will require changes to serialization and de-serialization code.	1	0	https://issues.apache.org/jira/browse/APEXCORE-635

#	Requirement	Questions	Answered	Successfully Answered	URL
20166	APEXCORE-635	bq. Yes, Agree, One array of byte only treat as one object. But as the memory is relative big, it probably will not allocated in TLA(Thread Local Area), which means it need to get lock for each allocation, so the allocation is not a very light operation. Is your concern GC or allocating new byte array? From what I understand, it was GC. bq. Output.getBuffer() probably not work in our case. The getBuffer() returns the reference of Output's internal buffer, which was used for putting the serialized bytes. As we need to share one instance of Output for multiple tuples, the buffer was reused when serialize next tuple. Which cause problem no matter we set the output's position to zero or not. If set position to zero, which means clean the previous serialized data. If not changed position, then the serialized data of next tuple concat to the previous one, but we don't know the boundary of each tuple, and the buffer keeping on growing which means output have to keep on allocate new big memory and copy the old data(if the initial buffer size is not very big). Output.getBuffer() works. Tuples are serialized one at a time and after each serialization, bytes from the shared Output buffer are copied to the SerializedTuple that allocates a new buffer each time. It avoids extra copy that you refer to. bq. Yes, the original prototype reset the buffer when it should not. But as I pointed in this proposal, the buffer will be reset after the data sent to socket. That why need to override the write() method. The original prototype has a serious bug, so the results of benchmarking are not reliable. Prior to implementing the proposed approach, I suggest to implement a prototype in stages and have a prove that the approach will provide a significant performance improvement. Without such prove, the complexity of the proposal is not justifiable. bq. Yes, agree. But I think it was worthy to change StreamCode, as suggested in the proposal, it's not easy to implemented the customized reusable serialization without memory copy with the previous interface. And it doesn't compatible with kryo's interface. The current implementation(DefaultStatefulStreamCodec) of StatefulStreamCodec seems not natural for me. The pairs are cleared after first tuple. And for each other tuples, the pairs need to checked, and an instance of DataStatePair need to created to wrapper state(which is null except first tuple) and data. I agree that the interface is not optimal and needs to be changed, but it is not possible to change it without breaking backward compatibility. It means that new interface needs to be introduced and BufferServer will need to deal with 3 different codec interfaces. Before this is done, I'd like to have a prove that there is significant performance benefit that justifies the complexity. bq. The netlet client need to aware of the memory management mechanism in order to reset the memory. But the code is maintained in same class (current BufferServerPublisher; an extended class in proposal), so this should be ok. The serialization formation don't have to change. it can compatible with the previous one. But if we want to take advantage of reserve() function of SerializationBuffer to avoid another extra copy, the format need to be changed to avoid use variable length. Again, I'd like to see a staged prototype that proves that there is a performance benefit and that it is possible to avoid tight code coupling between Codec, Publisher and Tuple serialization format. Overall, I am still (n) on the approach and I think that it needs to be revisited when netlet->netty migration is done. Netty has it's own memory allocator and it will be necessary to see how it can be integrated with Kryo and other serializers.	1	0	https://issues.apache.org/jira/browse/APEXCORE-635
20167	APEXCORE-620	What is the proposed mechanism to resolve dependencies? Apex application developer should provide .apa file with all dependencies (except the one that will be provided at run-time) resolved.	1	0	https://issues.apache.org/jira/browse/APEXCORE-620
20168	APEXCORE-612	Can you please add more details to this? Are you saying you see Out of Sequence for endWindow Tuple?	1	0	https://issues.apache.org/jira/browse/APEXCORE-612
20169	APEXCORE-592	QUEUE_CAPACITY is an attribute and not a property of the port. Attributes are defined by the Apex platform and do not change across ports. What is the significance of providing the description for attributes?	1	1	https://issues.apache.org/jira/browse/APEXCORE-592
20170	APEXCORE-592	[~vrozov] The description in properties is required for displaying user the description of properties configurable by user. Would you suggest me filtering out this list so that we return description only for the operator properties?	1	1	https://issues.apache.org/jira/browse/APEXCORE-592
20171	APEXCORE-592	Have you looked at get-jar-operator-properties & get-jar-operator-classes commands in the CLI? They use OperatorDiscoverer (TypeGraph), which reads from the JavaDoc. It is better to explore that option.	1	0	https://issues.apache.org/jira/browse/APEXCORE-592
20172	APEXCORE-592	Ok, I am fine with keeping this simple. Why limit it to defaultProperties?	0	0	https://issues.apache.org/jira/browse/APEXCORE-592
20173	APEXCORE-587	Agree with both of you. But how about adding methods prepareDAG to launcher API and launchDAG that takes a prepared DAG and launches it, launchApp will be a combination of both. Also yes for making getDAG() part of AppHandler. Also thomas, the issue is current assigned, do you want to work on it? If so please go ahead otherwise I can work on it as well.	1	0	https://issues.apache.org/jira/browse/APEXCORE-587
20174	APEXCORE-581	Do the streams need to be control tuple aware? It is still a single stream isn't it and the nodes do the de-multiplexing isn't it?	1	0	https://issues.apache.org/jira/browse/APEXCORE-581
20175	APEXCORE-580	UserDefinedControlTuple is a long name for users to use. How about just calling it ControlTuple?	1	0	https://issues.apache.org/jira/browse/APEXCORE-580
20176	APEXCORE-577	Why specifically limit it to client side preparation of the DAG before the application is launched? Why not make it possible to have general hooks that can apply even when the application is running in the different distributed components of the application such as the containers and stram. The hooks could be registered to be asynchronous or synchronous. The asynchronous ones could be handled via a bus, we already have a bus called mbassador that we use today and it could potentially be used for this. The entire functionality is akin to something like dtrace without the dynamic part.	1	1	https://issues.apache.org/jira/browse/APEXCORE-577
20177	APEXCORE-576	As far as I understand then the main enhancement here is the "registry" and a process to pick items from the registry to merge into the main repository and we will be using the current git process of pull-requests and merging across forks. Are there existing models we can look at for the registry? A few things that need to be hammered out: - registry access control, and format - criteria for merging external contributions into the main repo (quality, unit tests, demand, licensing) - level of automation available, or otherwise the manual process used to merge contributions - dependency matrix management (e.g. an item depends on malhar 3.5.0 minimum) - guidelines for adding items to the registry to be followed by contributors	1	0	https://issues.apache.org/jira/browse/APEXCORE-576
20178	APEXCORE-576	I had the same concerns as [~ananthg.apex] and I agree that having this registry can dilute the quality and the Apex story as a rich ecosystem. Also the stated goal is to help/facilitate external contributions to Apex but how will the registry help with that goal?	0	0	https://issues.apache.org/jira/browse/APEXCORE-576
20179	APEXCORE-570	From the JIRA description, this sounds like a combination of backpressure with configurable limit on the queue. For the intra-process case the queue capacity can be set. For buffer server, it will spill over to disk, hence the upstream operator will produce at the speed of spooling. Are you thinking of a capacity limit here?	1	1	https://issues.apache.org/jira/browse/APEXCORE-570
20180	APEXCORE-570	If the spooling capacity limit is reached, would the operator just get blocked? If so, wouldn't that solve the back pressure problem at least partially?	1	1	https://issues.apache.org/jira/browse/APEXCORE-570

#	Requirement	Questions	Answered	Successfully Answered	URL
20181	APEXCORE-570	Let's first address the process issue (it may warrant a separate discussion and additions to the contributor guidelines also). If you think there was a conclusion then this may indicate that there was offline discussion that isn't captured here or anywhere else. Just by looking at this ticket it is everything but clear what lead to your PR. This is not how the community can work, discussion has to be in the open. Now back to the topic: I think this is a good approach. It will avoid the fast producing operator running ahead at the speed of writing to disk indefinitely. With this we will not need to limit the spooling at all? What about the case where you have two subscribers (and those could be different operators) where one can keep up with the rate at which data is published and the other one may be slow, albeit maybe temporarily? This will slow down the fast subscriber and introduce latency.	1	1	https://issues.apache.org/jira/browse/APEXCORE-570
20182	APEXCORE-570	I will go into the topic specific comments first and then address your concern about my approach with the process. "Now back to the topic: I think this is a good approach. It will avoid the fast producing operator running ahead at the speed of writing to disk indefinitely. With this we will not need to limit the spooling at all?" Correct, spooling limit is not needed. Spooling size limit as a mechanism for back pressure will not work because we do not know how much data will be generated between two commits (committed windows). Also since the checkpoint length is configurable we cannot set it to some "reasonable" high value. Hence, spooling size limit is something that would not be practical. "What about the case where you have two subscribers (and those could be different operators) where one can keep up with the rate at which data is published and the other one may be slow, albeit maybe temporarily? This will slow down the fast subscriber and introduce latency." Yes, you are correct, the slowest subscriber will slow down the publisher (unless it is parallel partition all the way through). But, this is expected isn't it with back pressure. "Let's first address the process issue (it may warrant a separate discussion and additions to the contributor guidelines also). If you think there was a conclusion then this may indicate that there was offline discussion that isn't captured here or anywhere else. Just by looking at this ticket it is everything but clear what lead to your PR. This is not how the community can work, discussion has to be in the open." I think you have misunderstood my approach. When I created the JIRA, that started the discussion, I was facing a problem with a production application and had proposed using window difference as the way to create the back pressure and block publisher. Limiting spooling was suggested as an approach by both you and David and in my comment on 02nd Nov 16th at 22:53 I mentioned that it won't work because it will cause a deadlock. David's had another comment on this approach about suspending publisher on spool limit till committed which is effectively the same deadlock problem as the commit will not happen till publisher moves forward. There were no other approaches suggested so I proceeded with attempting to solve the problem via the proposed window difference approach. As I got into the weeds of the implementation and figured out all the details of how the current implementation works, I figured that instead of a window difference using the block difference was a better way to accomplish this. To me, the fundamental approach I originally suggested of blocking publisher till subscriber caught up hadn't changed rather an implementation detail. Second, the majority of the time during the implementation was spent in how to accomplish the task with the original assumption of window difference and coming to the conclusion to use blocks instead of windows and the actual coding a day or two so what you see in the PR is a relatively new discovery. In your comments, you have made a couple of statements, first that there may have been offline discussions on the implementation. This has not happened, I assure you. You are all seeing the implementation at the same time including the reviewers. The second statement is stronger about this being detrimental to community and discussions have to be open. I take personal offense to this statement. I know you want the best for the community but suggest you ascertain the truth before making such strong statements.	1	1	https://issues.apache.org/jira/browse/APEXCORE-570
20183	APEXCORE-565	When will inactive operators be activated? Are there automation use cases where operators (and streams) will be activated based on a condition/event that happens in the application itself?	0	0	https://issues.apache.org/jira/browse/APEXCORE-565
20184	APEXCORE-536	So can we add anything that is new after 2.2 or we should stay backward compatible?	0	0	https://issues.apache.org/jira/browse/APEXCORE-536
20185	APEXCORE-536	When are we planning to upgrade? We can right away use, application tags feature. I think we don't have to worry about backward compatibility. If the users are still using Hadoop 2.2, then they should be happy with Apex 3.4 (or 3.5).	0	0	https://issues.apache.org/jira/browse/APEXCORE-536
20186	APEXCORE-534	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-site/pull/50#discussion_r78984813 --- Diff: src/md/contributing.md --- @@ -14,7 +14,38 @@ This project welcomes new contributors and invites everyone to participate. Our People that help with the project in any of the above categories or other ways are contributors. See the [roles](http://www.apache.org/foundation/how-it-works.html#roles) as defined by the ASF. Community members that make sustained, welcome contributions to the project may be invited to become a [committer](/people.html). -## Code Style +## Before coding: One time Setup + +### JIRA + +Apache JIRA is used for issue tracking. If you do not already have an Apache JIRA account, sign up [here](https://issues.apache.org/jira/). Note that the user name should have no white spaces or other special characters that complicate auto-completion within JIRA comments etc. + +Please use a single JIRA account only (don't create multiple with different email addresses) to retain the issue history. Please use a permanent email address, for an existing account it can be changed in the profile. If you absolutely have to change your user name, contact INFRA. + +Apex has 2 JIRA projects: + +1. [APEXCORE](https://issues.apache.org/jira/browse/APEXCORE/) for [apex-core](https://github.com/apache/apex-core) and [apex-site](https://github.com/apache/apex-site) +2. [APEXMALHAR](https://issues.apache.org/jira/browse/APEXMALHAR/) for [apex-malhar](https://github.com/apache/apex-malhar) + +Before working on changes for any of the repositories, please locate an existing JIRA ticket or submit a new one. In order to assign an issue to yourself, you need to be listed as contributor in the JIRA project. PMC members have access to add new contributors, please request to be added through a comment on the candidate ticket. --- End diff -- Rather than adding a comment on tkt, should we ask new contributor to send a mail to dev@apex for contributor access. That way, the PMC/Committers can follow up with new contributor and help the person if required. This way, he/she gets introduced to the community at the same time, existing community members gets chance to give new contributers next step.	1	1	https://issues.apache.org/jira/browse/APEXCORE-534
20187	APEXCORE-526	Can you please provide information on how to maintain this. Where is the configuration, how do we add new branches as part of the release process (may need to add relevant info to http://apex.apache.org/release.html)? Also, will commits to respective branches trigger the build?	1	1	https://issues.apache.org/jira/browse/APEXCORE-526
20188	APEXCORE-526	It would be good to document what we have for Malhar so that others can add the branches when needed. I would like to make this part of the upcoming 3.6 release. Would it be (later) possible to generate the user docs with this process also?	1	1	https://issues.apache.org/jira/browse/APEXCORE-526
20189	APEXCORE-526	Can it be setup to build whenever there are commits pushed to branches? That would be better than scheduling.	1	1	https://issues.apache.org/jira/browse/APEXCORE-526
20190	APEXCORE-526	There seems to be part of the documentation missing. Are you building it with -Pall-modules ?	1	1	https://issues.apache.org/jira/browse/APEXCORE-526

#	Requirement	Questions	Answered	Successfully Answered	URL
20191	APEXCORE-526	What's missing ? After adding that option what should I look for ?	1	1	https://issues.apache.org/jira/browse/APEXCORE-526
20192	APEXCORE-518	Do you mean to change parent pom to use the newer version of Apache parent pom? IMO, upgrading parent pom to the newer Apache version should be done in a separate commit. Changes to the codestyle-config pom.xml are completely independent. Changes to shaded-ning will require publishing new shaded-ning version.	0	0	https://issues.apache.org/jira/browse/APEXCORE-518
20193	APEXCORE-518	I don't see why. Changing the code and releasing are separate. The whole reason I asked for a separate JIRA as that this Apex repository should not use different ASF parent poms for no obvious reason. It is the parent pom change (not the version) that has nothing to do with checkstyle and is not even needed to publish a new version of the code style artifact.	1	1	https://issues.apache.org/jira/browse/APEXCORE-518
20194	APEXCORE-514	[~thw] I agree with you.. Some more change as suggested by you are required. I also have some suggestions about website updates. Would it make sense to start a thread on dev@apex about this? Otherwise, I can provide my suggestions here too.	1	0	https://issues.apache.org/jira/browse/APEXCORE-514
20195	APEXCORE-510	What classes overwrite DefaultOutputPort.setup()?	1	1	https://issues.apache.org/jira/browse/APEXCORE-510
20196	APEXCORE-498	[~sandesh] this is a good feature to have, it enables reprocessing. Can you outline how it will work, without getting into code details for now. Specifically, there are different areas where state is tracked (master, buffer servers, operator state, external systems). How will this play with all of those? Please consider end-to-end exactly once and see if and how the proposed names checkpoints (also referred to as safepoints?) will fit in?	1	0	https://issues.apache.org/jira/browse/APEXCORE-498
20197	APEXCORE-498	Discussed this with [~davidyan] & [~PramodSSImmaneni]. There was a concern that to use this feature properly a lot of effort is required beyond this feature. Existing Operators that use WindowDataManager won't work properly with this feature, they have to be made SavePoint compatible. What is your take on this [~thomasWeise] ?	1	0	https://issues.apache.org/jira/browse/APEXCORE-498
20198	APEXCORE-498	I think this is an important feature that should be supported. It needs to be designed properly, so I was actually expecting a design level discussion that is visible to everyone. There will of course be changes needed to existing components and there will also be limitations that pertain to external systems that need to be understood and discussed. Why should operators that use WindowDataManager not work?	0	0	https://issues.apache.org/jira/browse/APEXCORE-498
20199	APEXCORE-496	We support shared stats listener, But user do not have any way to identify the operator for which the stat listener is called. The BatchedOperatorStats only provides operatorId, and there is no api available to use to get operator information from operatorId. Can we include operator name as part of BatchedOperatorStats? This will allow users to take decision based on which operator stats are being processed.	1	0	https://issues.apache.org/jira/browse/APEXCORE-496
20200	APEXCORE-496	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/361#discussion_r71943358 --- Diff: api/src/main/java/com/datatorrent/api/StatsListener.java --- @@ -113,6 +113,8 @@ long getLatencyMA(); List<OperatorResponse> getOperatorResponse(); + + String getOperatorName(); --- End diff -- Why operator name is part of BatchedOperatorStats when operator name does not change?	0	0	https://issues.apache.org/jira/browse/APEXCORE-496
20201	APEXCORE-496	Github user tushargosavi commented on a diff in the pull request: https://github.com/apache/apex-core/pull/361#discussion_r74229007 --- Diff: api/src/main/java/com/datatorrent/api/StatsListener.java --- @@ -115,6 +115,28 @@ List<OperatorResponse> getOperatorResponse(); } + /** + * An interface to the DAG. Stats listener can get information about + * operator or other elements in the DAG through this interface. currerntly + * we only provide method to extract the operator name based on the physical + * id of the operator. In future more methods can be added. + */ --- End diff -- exposing all information may expose internal implementation of the DAG. we can only have dependency on apex-api. Can you suggest which information can be provided to stat listener without adding dependency on the apex-engine. I can think of following - name - number of partitions for self. - aggregated stats at logical operator level. This list could be expanded in future. This jira puts the required infrastructure without breaking backward compatibility.	0	0	https://issues.apache.org/jira/browse/APEXCORE-496
20202	APEXCORE-495	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/360#discussion_r71920123 --- Diff: engine/src/main/java/com/datatorrent/stram/cli/ApexCli.java --- @@ -3715,6 +3718,50 @@ DTConfiguration getLaunchAppPackageProperties(AppPackage ap, ConfigPackage cp, L return launchProperties; } + private List<AppInfo> getAppsFromPackageAndConfig(AppPackage ap, ConfigPackage cp, boolean onlyConfigApps) + { + if (cp == null) { + return ap.getApplications(); + } + + File src = new File(cp.tempDirectory(), "app"); + File dest = new File(ap.tempDirectory(), "app"); + if (!src.exists()) { + return ap.getApplications(); + } + if (onlyConfigApps) { + try { + FileUtils.cleanDirectory(dest); + } catch (IOException e) { + e.printStackTrace(); + } + } else { + for (File app: src.listFiles()) { + try { + FileUtils.deleteQuietly(new File(dest, app.getName())); + } + } + for (File app: src.listFiles()) { + try { + FileUtils.moveFileToDirectory(app, dest, true); + } catch (IOException e) { + LOG.error("Application from the config file {} failed while processing", app.getName()); --- End diff -- Should it be a warn since it's swallowing the exception and moving on?	1	1	https://issues.apache.org/jira/browse/APEXCORE-495
20203	APEXCORE-495	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/360#discussion_r72669919 --- Diff: engine/src/main/java/com/datatorrent/stram/cli/ApexCli.java --- @@ -3850,6 +3908,9 @@ private Option add(Option opt) final Option exactMatch = add(new Option("exactMatch", "Only consider applications with exact app name")); final Option queue = add(OptionBuilder.withArgName("queue name").hasArg().withDescription("Specify the queue to launch the application").create("queue")); final Option force = add(new Option("force", "Force launch the application. Do not check for compatibility")); + final Option onlyConfigApps = add(new Option("onlyConfigApps", "Only use the apps present in the config package.")); + final Option skipConfigApps = add(new Option("skipConfigApps", "Skip the apps present in the config package.")); + --- End diff -- Instead of having two options that conflict each other, why don't we have one option that takes a value? for example --useConfigPackageApps={exclusive|inclusive|none}	1	1	https://issues.apache.org/jira/browse/APEXCORE-495

#	Requirement	Questions	Answered	Successfully Answered	URL
20204	APEXCORE-495	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/360#discussion_r73581539 --- Diff: engine/src/test/java/com/datatorrent/stram/cli/ApexCliTest.java --- @@ -180,4 +188,40 @@ public void testLaunchAppPackagePrecedenceWithConfigPackage() throws Exception Assert.assertEquals("app-default", props.get("dt.test.5")); Assert.assertEquals("package-default", props.get("dt.test.6")); } + + @Test + public void testAppFromOnlyConfigPackage() throws Exception + { + ApexCli.LaunchCommandLineInfo commandLineInfo = ApexCli.getLaunchCommandLineInfo(new String[]{"-conf", configFile.getAbsolutePath(), appFile.getAbsolutePath(), "-useConfigApps", "exclusive"}); + ApexCli apexCli = new ApexCli(); + apexCli.init(); + Assert.assertEquals("configApps", "exclusive", commandLineInfo.useConfigApps); + apexCli.getLaunchAppPackageArgs(ap, cp, commandLineInfo, new ConsoleReader()); + Assert.assertEquals(ap.getApplications().size(), 1); + } + + @Test + public void testMergeAppFromConfigAndAppPackage() throws Exception + { + ApexCli.LaunchCommandLineInfo commandLineInfo = ApexCli.getLaunchCommandLineInfo(new String[]{"-conf", configFile.getAbsolutePath(), appFile.getAbsolutePath(), "-useConfigApps", "inclusive"}); + Assert.assertEquals("configApps", "inclusive", commandLineInfo.useConfigApps); + ApexCli apexCli = new ApexCli(); + apexCli.init(); + try { + apexCli.getLaunchAppPackageArgs(ap, cp, commandLineInfo, new ConsoleReader()); + } catch (Throwable throwable) { --- End diff -- why are we catching throwable instead of the exception you expect it to throw? Also, there is an extra space before the closing paren.	1	1	https://issues.apache.org/jira/browse/APEXCORE-495
20205	APEXCORE-495	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/360#discussion_r73581887 --- Diff: engine/src/test/java/com/datatorrent/stram/cli/ApexCliTest.java --- @@ -180,4 +188,40 @@ public void testLaunchAppPackagePrecedenceWithConfigPackage() throws Exception Assert.assertEquals("app-default", props.get("dt.test.5")); Assert.assertEquals("package-default", props.get("dt.test.6")); } + + @Test + public void testAppFromOnlyConfigPackage() throws Exception + { + ApexCli.LaunchCommandLineInfo commandLineInfo = ApexCli.getLaunchCommandLineInfo(new String[]{"-conf", configFile.getAbsolutePath(), appFile.getAbsolutePath(), "-useConfigApps", "exclusive"}); + ApexCli apexCli = new ApexCli(); + apexCli.init(); + Assert.assertEquals("configApps", "exclusive", commandLineInfo.useConfigApps); + apexCli.getLaunchAppPackageArgs(ap, cp, commandLineInfo, new ConsoleReader()); + Assert.assertEquals(ap.getApplications().size(), 1); + } + + @Test + public void testMergeAppFromConfigAndAppPackage() throws Exception + { + ApexCli.LaunchCommandLineInfo commandLineInfo = ApexCli.getLaunchCommandLineInfo(new String[]{"-conf", configFile.getAbsolutePath(), appFile.getAbsolutePath(), "-useConfigApps", "inclusive"}); + Assert.assertEquals("configApps", "inclusive", commandLineInfo.useConfigApps); + ApexCli apexCli = new ApexCli(); + apexCli.init(); + try { + apexCli.getLaunchAppPackageArgs(ap, cp, commandLineInfo, new ConsoleReader()); + } catch (Throwable throwable) { --- End diff -- Also, why should this throw an exception? It should be okay to have the conf package to include multiple apps and in which case, let the user choose by showing a menu, just like app package can contain multiple apps.	1	1	https://issues.apache.org/jira/browse/APEXCORE-495
20206	APEXCORE-495	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/360#discussion_r73582122 --- Diff: engine/src/test/java/com/datatorrent/stram/cli/ApexCliTest.java --- @@ -103,8 +107,8 @@ public static void setup() public void testLaunchAppPackagePropertyPrecedence() throws Exception { // set launch command options - ApexCli.LaunchCommandLineInfo commandLineInfo = ApexCli --- End diff -- Why are we making such changes and have an extra import static? Please restrict your changes to **only** those that are relevant to this feature.	1	1	https://issues.apache.org/jira/browse/APEXCORE-495
20207	APEXCORE-495	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/360#discussion_r73582758 --- Diff: engine/src/test/resources/testAppPackage/mydtapp/src/main/resources/com/example/mydtapp/Application/ExtRes.html --- @@ -0,0 +1,21 @@ +<!-- + Licensed to the Apache Software Foundation (ASF) under one + or more contributor license agreements. See the NOTICE file + distributed with this work for additional information + regarding copyright ownership. The ASF licenses this file + to you under the Apache License, Version 2.0 (the + "License"); you may not use this file except in compliance + with the License. You may obtain a copy of the License at + + http://www.apache.org/licenses/LICENSE-2.0 + + Unless required by applicable law or agreed to in writing, + software distributed under the License is distributed on an + "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + KIND, either express or implied. See the License for the + specific language governing permissions and limitations + under the License. + --> + <html>example</html> --- End diff -- What is this file for?	1	1	https://issues.apache.org/jira/browse/APEXCORE-495
20208	APEXCORE-495	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/360#discussion_r73583092 --- Diff: engine/src/test/java/com/datatorrent/stram/cli/ApexCliTest.java --- @@ -53,16 +56,16 @@ private static File configFile; - private static AppPackage ap; - private static ConfigPackage cp; - private static TemporaryFolder testFolder = new TemporaryFolder(); - private static ApexCli cli = new ApexCli(); + private AppPackage ap; + private ConfigPackage cp; + TemporaryFolder testFolder = new TemporaryFolder(); + ApexCli cli = new ApexCli(); - private static Map<String, String> env = new HashMap<String, String>(); - private static String userHome; + Map<String, String> env = new HashMap<String, String>(); + String userHome; - @BeforeClass - public static void starting() + @Before + public void starting() --- End diff -- Can you explain why this change is needed?	1	1	https://issues.apache.org/jira/browse/APEXCORE-495
20209	APEXCORE-489	Could you provide additional details such as: What does your DAG look like ? What operators (custom as well as from Malhar) are you using ? Does this exception happen immediately upon launch or after some time ?	1	0	https://issues.apache.org/jira/browse/APEXCORE-489
20210	APEXCORE-485	Github user vrozov commented on the issue: https://github.com/apache/apex-core/pull/355 @davidyan74 David, can you merge this one as well.	1	1	https://issues.apache.org/jira/browse/APEXCORE-485
20211	APEXCORE-484	Github user vrozov commented on the issue: https://github.com/apache/apex-core/pull/354 @davidyan74 This should help with other open PRs build failures. Can you review and merge.	1	1	https://issues.apache.org/jira/browse/APEXCORE-484
20212	APEXCORE-482	Why is Mx1 special ?	0	0	https://issues.apache.org/jira/browse/APEXCORE-482
20213	APEXCORE-482	Github user vrozov commented on the issue: https://github.com/apache/apex-core/pull/358 What is the reason unifier does not have it's own StreamMeta? IMO, it will be better to introduce separate StreamMeta for unifiers instead of the locality field that is used only for unifiers.	0	0	https://issues.apache.org/jira/browse/APEXCORE-482

#	Requirement	Questions	Answered	Successfully Answered	URL
20214	APEXCORE-482	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/358#discussion_r72490248 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/physical/PTOperator.java --- @@ -85,6 +85,7 @@ private static final long serialVersionUID = 201312112033L; public final LogicalPlan.StreamMeta logicalStream; + public Locality locality; --- End diff -- Why can't locality be set on streammeta object?	0	0	https://issues.apache.org/jira/browse/APEXCORE-482
20215	APEXCORE-475	We can implement part of your suggestion now, that is to include DataTorrent & ApacheApex while retrieving. Then in the 4.0 we can rename the category to ApacheApex. By doing this we can avoid the problem during development that can happen at a later point as well. What say?	1	1	https://issues.apache.org/jira/browse/APEXCORE-475
20216	APEXCORE-475	Please see discussion on the PR. Do you have examples for other application types?	0	0	https://issues.apache.org/jira/browse/APEXCORE-475
20217	APEXCORE-467	What's the problem with having the buffer server in each container? Does it allocate resources that are not used?	0	0	https://issues.apache.org/jira/browse/APEXCORE-467
20218	APEXCORE-466	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/339#discussion_r65788854 --- Diff: engine/src/main/java/com/datatorrent/stram/client/RecordingsAgent.java --- @@ -640,7 +640,7 @@ public JSONObject process(WebResource.Builder webResource, Class<JSONObject> cla }); } catch (Exception ex) { - LOG.error("Exception caught", ex); + LOG.error("Cannot start recording: ", ex); --- End diff -- @davidyan74 : Should logging be not done at same levels? I see here it is "error" and earlier it is done at "warn" and also this will print stacktrace but log messages at 546 line won't... Any reason for inconsistent log levels?	1	1	https://issues.apache.org/jira/browse/APEXCORE-466
20219	APEXCORE-457	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/333#discussion_r63443468 --- Diff: docs/security.md --- @@ -142,7 +168,15 @@ When operators are running there will be effective processing rate differences b Like STRAM, streaming containers also need to communicate with NameNode to use HDFS persistence for reasons such as saving the state of the operators. In secure mode they also use NameNode delegation tokens for authentication. These tokens are also seeded by STRAM for the streaming containers. +#### Stram Webservices + +Clients connects to STRAM and make web service requests to obtain operational information about a runtime application. When security is enabled we want this connection to also be authenticated. In this mode the client passes a web service token in the request and the STRAM checks this token. If the token is valid, then the request is processed else it is denied. + +How does the client get the web service token in the first place The client will first have to first connect to STRAM via the Resource Manager Web Services Proxy which is a service run by Hadoop to proxy requests to application web services. This connection is authenticated by the proxy service using a protocol called SPNEGO when secure mode is enabled. SPNEGO is Kerberos over HTTP and the client also needs to support it. If the authentication is successful the proxy forwards the request to STRAM. STRAM in processing the request generates and send back a web service token similar to a delegation token. This token is then used by client in subsequent requests it makes directly to STRAM and STRAM is able to validate it since it generated the token in the first place. --- End diff -- ... in the first place (add a question mark) The client ...	1	1	https://issues.apache.org/jira/browse/APEXCORE-457
20220	APEXCORE-457	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/333#discussion_r63443528 --- Diff: docs/security.md --- @@ -142,7 +168,15 @@ When operators are running there will be effective processing rate differences b Like STRAM, streaming containers also need to communicate with NameNode to use HDFS persistence for reasons such as saving the state of the operators. In secure mode they also use NameNode delegation tokens for authentication. These tokens are also seeded by STRAM for the streaming containers. +#### Stram Webservices + +Clients connects to STRAM and make web service requests to obtain operational information about a runtime application. When security is enabled we want this connection to also be authenticated. In this mode the client passes a web service token in the request and the STRAM checks this token. If the token is valid, then the request is processed else it is denied. + +How does the client get the web service token in the first place The client will first have to first connect to STRAM via the Resource Manager Web Services Proxy which is a service run by Hadoop to proxy requests to application web services. This connection is authenticated by the proxy service using a protocol called SPNEGO when secure mode is enabled. SPNEGO is Kerberos over HTTP and the client also needs to support it. If the authentication is successful the proxy forwards the request to STRAM. STRAM in processing the request generates and send back a web service token similar to a delegation token. This token is then used by client in subsequent requests it makes directly to STRAM and STRAM is able to validate it since it generated the token in the first place. --- End diff -- "The client will first have to first connect ..." Double "first"	1	1	https://issues.apache.org/jira/browse/APEXCORE-457
20221	APEXCORE-457	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/333#discussion_r63443609 --- Diff: docs/security.md --- @@ -142,7 +168,15 @@ When operators are running there will be effective processing rate differences b Like STRAM, streaming containers also need to communicate with NameNode to use HDFS persistence for reasons such as saving the state of the operators. In secure mode they also use NameNode delegation tokens for authentication. These tokens are also seeded by STRAM for the streaming containers. +#### Stram Webservices + +Clients connects to STRAM and make web service requests to obtain operational information about a runtime application. When security is enabled we want this connection to also be authenticated. In this mode the client passes a web service token in the request and the STRAM checks this token. If the token is valid, then the request is processed else it is denied. + +How does the client get the web service token in the first place The client will first have to first connect to STRAM via the Resource Manager Web Services Proxy which is a service run by Hadoop to proxy requests to application web services. This connection is authenticated by the proxy service using a protocol called SPNEGO when secure mode is enabled. SPNEGO is Kerberos over HTTP and the client also needs to support it. If the authentication is successful the proxy forwards the request to STRAM. STRAM in processing the request generates and send back a web service token similar to a delegation token. This token is then used by client in subsequent requests it makes directly to STRAM and STRAM is able to validate it since it generated the token in the first place. --- End diff -- STRAM in processing the request generates and send**s** back	1	1	https://issues.apache.org/jira/browse/APEXCORE-457

#	Requirement	Questions	Answered	Successfully Answered	URL
20222	APEXCORE-457	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/333#discussion_r63443653 --- Diff: docs/security.md --- @@ -142,7 +168,15 @@ When operators are running there will be effective processing rate differences b Like STRAM, streaming containers also need to communicate with NameNode to use HDFS persistence for reasons such as saving the state of the operators. In secure mode they also use NameNode delegation tokens for authentication. These tokens are also seeded by STRAM for the streaming containers. +#### Stram Webservices + +Clients connects to STRAM and make web service requests to obtain operational information about a runtime application. When security is enabled we want this connection to also be authenticated. In this mode the client passes a web service token in the request and the STRAM checks this token. If the token is valid, then the request is processed else it is denied. +How does the client get the web service token in the first place The client will first have to first connect to STRAM via the Resource Manager Web Services Proxy which is a service run by Hadoop to proxy requests to application web services. This connection is authenticated by the proxy service using a protocol called SPNEGO when secure mode is enabled. SPNEGO is Kerberos over HTTP and the client also needs to support it. If the authentication is successful the proxy forwards the request to STRAM. STRAM in processing the request generates and send back a web service token similar to a delegation token. This token is then used by client in subsequent requests it makes directly to STRAM and STRAM is able to validate it since it generated the token in the first place. --- End diff -- This token is then used by (add **the**) client	1	1	https://issues.apache.org/jira/browse/APEXCORE-457
20223	APEXCORE-451	Github user shubham-pathak22 commented on the pull request: https://github.com/apache/incubator-apex-core/pull/326#issuecomment-217246009 Hi, @tushargosavi @PramodSSImmaneni @tweise @davidyan74 Could you please review the changes.	1	1	https://issues.apache.org/jira/browse/APEXCORE-451
20224	APEXCORE-451	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/326#discussion_r62241719 --- Diff: engine/src/main/java/com/datatorrent/stram/webapp/TypeGraph.java --- @@ -348,13 +348,13 @@ public int size() public Set<String> getAllDTInstantiableOperators() { - TypeGraphVertex tgv = typeGraph.get(Operator.class.getName()); + TypeGraphVertex tgv = typeGraph.get(DAG.GenericOperator.class.getName()); if (tgv == null) { return null; } Set<String> result = new TreeSet<>(); for (TypeGraphVertex node : tgv.allInstantiableDescendants) { - if ((isAncestor(InputOperator.class.getName(), node.typeName) || !getAllInputPorts(node).isEmpty())) { + if ((isAncestor(DAG.GenericOperator.class.getName(), node.typeName) || !getAllInputPorts(node).isEmpty())) { --- End diff -- Can you explain why this change? It looks like it's changing the intention of the original code.	1	1	https://issues.apache.org/jira/browse/APEXCORE-451
20225	APEXCORE-448	There is another attribute of such read-only ilk, IS_OUTPUT_UNIFIED on PortContext. The read-only behavior is not enforced by the platform but user changing it has no effect. However, in this case, it might be confusing for user to see a method in DAG to add operator with a name and an operator attribute NAME. How about calling the method getOperatorName() to be more specific. Do you want to propose it on dev or would you like me to?	1	0	https://issues.apache.org/jira/browse/APEXCORE-448
20226	APEXCORE-448	Github user ilooner commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r72900235 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,98 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import java.util.Collection; +import java.util.Map; +import java.util.concurrent.CountDownLatch; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; + +public class OperatorContextTest +{ + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + final CountDownLatch latch = new CountDownLatch(1); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + input.countDownLatch = latch; + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); + dag.addStream("stream", input.output, output.ip1); + } + }; + lma.prepareDAG(testApp, new Configuration()); + LocalMode.Controller lc = lma.getController(); + lc.runAsync(); + latch.await(); + lc.shutdown(); + } + + private static class MockInputOperator extends BaseOperator implements InputOperator, AutoMetric.Aggregator + { + private transient CountDownLatch countDownLatch; + + @AutoMetric + String operatorName; + + public final transient DefaultOutputPort<String> output = new DefaultOutputPort<>(); + + @Override + public void setup(Context.OperatorContext context) + { + operatorName = Preconditions.checkNotNull(context.getOperatorName(), "operator name"); + } + + @Override + public void emitTuples() + { + } + + @Override + public Map<String, Object> aggregate(long windowId, Collection<AutoMetric.PhysicalMetricsContext> physicalMetrics) + { + String name = (String)physicalMetrics.iterator().next().getMetrics().get("operatorName"); + Assert.assertEquals("operator name", "input", name); --- End diff -- Could we just set a static variable and check it in the test instead of using a latch and the assert?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20227	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73065873 --- Diff: api/src/main/java/com/datatorrent/api/Context.java --- @@ -328,6 +328,8 @@ */ int getWindowsFromCheckpoint(); + String getOperatorName(); --- End diff -- To be consistent with `getId()` method, should it be `getName()`?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448

#	Requirement	Questions	Answered	Successfully Answered	URL
20228	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73066955 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/OperatorContext.java --- @@ -84,12 +87,13 @@ public void setIdleTimeout(long idleTimeout) * @param attributes the value of attributes * @param parentContext */ - public OperatorContext(int id, AttributeMap attributes, Context parentContext) + public OperatorContext(int id, String operatorName, AttributeMap attributes, Context parentContext) { super(attributes, parentContext); this.lastProcessedWindowId = Stateless.WINDOW_ID; this.id = id; this.stateless = super.getValue(OperatorContext.STATELESS); + this.operatorName = Preconditions.checkNotNull(operatorName, "operator name"); --- End diff -- Should the condition be extended to check for empty operator name as well?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20229	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73068052 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/OperatorContext.java --- @@ -84,12 +87,13 @@ public void setIdleTimeout(long idleTimeout) * @param attributes the value of attributes * @param parentContext */ - public OperatorContext(int id, AttributeMap attributes, Context parentContext) + public OperatorContext(int id, String operatorName, AttributeMap attributes, Context parentContext) --- End diff -- id, name and attributes seem to always come from OperatorDeployInfo. If this is the case, should OperatorContext take OperatorDeployInfo as input parameter instead?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20230	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73365188 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,99 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import java.util.Collection; +import java.util.Map; +import java.util.concurrent.CountDownLatch; +import java.util.concurrent.TimeUnit; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; + +public class OperatorContextTest +{ + public class MockInputOperator extends BaseOperator implements InputOperator, AutoMetric.Aggregator + { + private transient CountDownLatch countDownLatch; + + @AutoMetric + String operatorName; + + public final transient DefaultOutputPort<String> output = new DefaultOutputPort<>(); + + @Override + public void setup(Context.OperatorContext context) + { + operatorName = Preconditions.checkNotNull(context.getName(), "operator name"); --- End diff -- Why is it necessary to use metrics? Will it be sufficient to assertNotNull in setup() method itself?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448

#	Requirement	Questions	Answered	Successfully Answered	URL
20231	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73423271 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,99 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import java.util.Collection; +import java.util.Map; +import java.util.concurrent.CountDownLatch; +import java.util.concurrent.TimeUnit; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; + +public class OperatorContextTest +{ + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + final CountDownLatch latch = new CountDownLatch(1); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + input.countDownLatch = latch; + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); + dag.addStream("stream", input.output, output.ip1); + } + }; + lma.prepareDAG(testApp, new Configuration()); + LocalMode.Controller lc = lma.getController(); + lc.runAsync(); + latch.await(2, TimeUnit.SECONDS); + lc.shutdown(); + } + + private static class MockInputOperator extends BaseOperator implements InputOperator, AutoMetric.Aggregator + { + private transient CountDownLatch countDownLatch; + + @AutoMetric + String operatorName; + + public final transient DefaultOutputPort<String> output = new DefaultOutputPort<>(); + + @Override + public void setup(Context.OperatorContext context) + { + operatorName = Preconditions.checkNotNull(context.getName(), "operator name"); --- End diff -- Will not calling `countDownLatch.countDown()` in setup() solve the problem?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20232	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73425701 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/OperatorContext.java --- @@ -84,12 +87,13 @@ public void setIdleTimeout(long idleTimeout) * @param attributes the value of attributes * @param parentContext */ - public OperatorContext(int id, AttributeMap attributes, Context parentContext) + public OperatorContext(int id, String operatorName, AttributeMap attributes, Context parentContext) { super(attributes, parentContext); this.lastProcessedWindowId = Stateless.WINDOW_ID; this.id = id; this.stateless = super.getValue(OperatorContext.STATELESS); + this.operatorName = Preconditions.checkNotNull(operatorName, "operator name"); --- End diff -- Do you mean that LogicalPlan.addOperator() allows empty name? Does not it also allows `null`? Is there any reason to allow `null` or empty string?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20233	APEXCORE-448	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73429227 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/OperatorContext.java --- @@ -84,12 +87,13 @@ public void setIdleTimeout(long idleTimeout) * @param attributes the value of attributes * @param parentContext */ - public OperatorContext(int id, AttributeMap attributes, Context parentContext) + public OperatorContext(int id, String name, AttributeMap attributes, Context parentContext) { super(attributes, parentContext); this.lastProcessedWindowId = Stateless.WINDOW_ID; this.id = id; this.stateless = super.getValue(OperatorContext.STATELESS); + this.name = Preconditions.checkNotNull(name, "operator name"); --- End diff -- When can it be null?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448

#	Requirement	Questions	Answered	Successfully Answered	URL
20234	APEXCORE-448	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73430406 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,98 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import java.util.Collection; +import java.util.Map; +import java.util.concurrent.CountDownLatch; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; + +public class OperatorContextTest +{ + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + final CountDownLatch latch = new CountDownLatch(1); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + input.countDownLatch = latch; + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); + dag.addStream("stream", input.output, output.ip1); + } + }; + lma.prepareDAG(testApp, new Configuration()); + LocalMode.Controller lc = lma.getController(); + lc.runAsync(); + latch.await(); + lc.shutdown(); + } + + private static class MockInputOperator extends BaseOperator implements InputOperator, AutoMetric.Aggregator + { + private transient CountDownLatch countDownLatch; + + @AutoMetric + String operatorName; + + public final transient DefaultOutputPort<String> output = new DefaultOutputPort<>(); + + @Override + public void setup(Context.OperatorContext context) + { + operatorName = Preconditions.checkNotNull(context.getOperatorName(), "operator name"); + } + + @Override + public void emitTuples() + { + } + + @Override + public Map<String, Object> aggregate(long windowId, Collection<AutoMetric.PhysicalMetricsContext> physicalMetrics) + { + String name = (String)physicalMetrics.iterator().next().getMetrics().get("operatorName"); + Assert.assertEquals("operator name", "input", name); + } +} --- End diff -- Wouldn't it rely on the current local mode implementation specific that the aggregator is not getting serialized and deserialized back (like the operator). That is not a specified contract isn't it? I think it would be safer to rely on the in-process execution nature and use static volatiles.	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20235	APEXCORE-448	Github user chandnisingh commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73433050 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/OperatorContext.java --- @@ -84,12 +87,13 @@ public void setIdleTimeout(long idleTimeout) * @param attributes the value of attributes * @param parentContext */ - public OperatorContext(int id, AttributeMap attributes, Context parentContext) + public OperatorContext(int id, String operatorName, AttributeMap attributes, Context parentContext) { super(attributes, parentContext); this.lastProcessedWindowId = Stateless.WINDOW_ID; this.id = id; this.stateless = super.getValue(OperatorContext.STATELESS); + this.operatorName = Preconditions.checkNotNull(operatorName, "operator name"); --- End diff -- @vrozov I am running this app with a ```null``` operator name and NullPointerException is thrown: ``` @ApplicationAnnotation(name = "TestApp") class SplitterApp implements StreamingApplication { MockReceiver receiver; @Override public void populateDAG(DAG dag, Configuration configuration) { dag.setAttribute(DAG.APPLICATION_PATH, baseTestMeta.dataDirectory); MockFileInput fileInput = dag.addOperator("Input", new MockFileInput()); fileInput.filePaths = baseTestMeta.filePaths; FileSplitterBase splitter = dag.addOperator(null, new FileSplitterBase()); splitter.setFile(baseTestMeta.dataDirectory); receiver = dag.addOperator("Receiver", new MockReceiver()); dag.addStream("files", fileInput.files, splitter.input); dag.addStream("file-metadata", splitter.filesMetadataOutput, receiver.fileMetadata); } } java.lang.NullPointerException at com.datatorrent.stram.plan.logical.LogicalPlan$OperatorMeta.equals(LogicalPlan.java:1187) at java.util.ArrayList.remove(ArrayList.java:528) at com.datatorrent.stram.plan.logical.LogicalPlan$StreamMeta.addSink(LogicalPlan.java:531) at com.datatorrent.stram.plan.logical.LogicalPlan.addStream(LogicalPlan.java:1391) at com.datatorrent.stram.plan.logical.LogicalPlan.addStream(LogicalPlan.java:1442) at com.datatorrent.stram.plan.logical.LogicalPlan.addStream(LogicalPlan.java:119) at com.datatorrent.lib.io.fs.FileSplitterBaseTest$SplitterApp.populateDAG(FileSplitterBaseTest.java:231) at com.datatorrent.stram.plan.logical.LogicalPlanConfiguration.prepareDAG(LogicalPlanConfiguration.java:2226) at com.datatorrent.stram.LocalModeImpl.prepareDAG(LocalModeImpl.java:67) at com.datatorrent.lib.io.fs.FileSplitterBaseTest.testSplitterInApp(FileSplitterBaseTest.java:205) at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method) at sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:62) at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:43) at java.lang.reflect.Method.invoke(Method.java:498) at org.junit.runners.model.FrameworkMethod$1.runReflectiveCall(FrameworkMethod.java:47) at org.junit.internal.runners.model.ReflectiveCallable.run(ReflectiveCallable.java:12) at org.junit.runners.model.FrameworkMethod.invokeExplosively(FrameworkMethod.java:44) at org.junit.internal.runners.statements.InvokeMethod.evaluate(InvokeMethod.java:17) at org.junit.rules.TestWatcher$1.evaluate(TestWatcher.java:55)	1	1	https://issues.apache.org/jira/browse/APEXCORE-448

#	Requirement	Questions	Answered	Successfully Answered	URL
20236	APEXCORE-448	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73444164 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,112 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import javax.validation.constraints.NotNull; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.Attribute; +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; +import com.datatorrent.stram.api.OperatorDeployInfo; + +public class OperatorContextTest +{ + private static boolean foundOperatorName; + + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); + dag.addStream("stream", input.output, output.ip1); + } + }; + lma.prepareDAG(testApp, new Configuration()); + LocalMode.Controller lc = lma.getController(); + lc.runAsync(); + long startTime = System.currentTimeMillis(); + while (!foundOperatorName && System.currentTimeMillis() - startTime < 2000) { + Thread.sleep(100); + } + Assert.assertTrue("operator name found", foundOperatorName); + lc.shutdown(); + } + + private static class MockInputOperator extends BaseOperator implements InputOperator + { + @AutoMetric --- End diff -- Why is this annotation needed?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20237	APEXCORE-448	Github user chandnisingh commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73447483 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,112 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import javax.validation.constraints.NotNull; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.Attribute; +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; +import com.datatorrent.stram.api.OperatorDeployInfo; + +public class OperatorContextTest +{ + private static boolean foundOperatorName; + + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); + --- End diff -- Does it hurt to have an output operator here in the test local application?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448

#	Requirement	Questions	Answered	Successfully Answered	URL
20238	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73448436 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,99 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import java.util.Collection; +import java.util.Map; +import java.util.concurrent.CountDownLatch; +import java.util.concurrent.TimeUnit; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; + +public class OperatorContextTest +{ + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + final CountDownLatch latch = new CountDownLatch(1); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + input.countDownLatch = latch; + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); + dag.addStream("stream", input.output, output.ip1); + } + }; + lma.prepareDAG(testApp, new Configuration()); + LocalMode.Controller lc = lma.getController(); + lc.runAsync(); + latch.await(2, TimeUnit.SECONDS); + lc.shutdown(); + } + + private static class MockInputOperator extends BaseOperator implements InputOperator, AutoMetric.Aggregator + { + private transient CountDownLatch countDownLatch; + + @AutoMetric + String operatorName; + + public final transient DefaultOutputPort<String> output = new DefaultOutputPort<>(); + + @Override + public void setup(Context.OperatorContext context) + { + operatorName = Preconditions.checkNotNull(context.getName(), "operator name"); --- End diff -- Please double check. Exception thrown in setup() should terminate the container/JVM, so the test will fail. In case it does not fail, should not Preconditions.checkNotNull() be removed?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20239	APEXCORE-448	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73448985 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,112 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import javax.validation.constraints.NotNull; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.Attribute; +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; +import com.datatorrent.stram.api.OperatorDeployInfo; + +public class OperatorContextTest +{ + private static boolean foundOperatorName; + + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); --- End diff -- Is this code needed to test the feature being test.. No.. So adding redundant code to test case is definitely is not a good idea...	1	1	https://issues.apache.org/jira/browse/APEXCORE-448

#	Requirement	Questions	Answered	Successfully Answered	URL
20240	APEXCORE-448	Github user chandnisingh commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73452693 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,112 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import javax.validation.constraints.NotNull; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.Attribute; +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; +import com.datatorrent.stram.api.OperatorDeployInfo; + +public class OperatorContextTest +{ + private static boolean foundOperatorName; + + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); --- End diff -- @vrozov How is adding another output operator here makes it difficult to understand the test logic? I believe if the GenericOperator fails, the test will pass. IMO nitpicking to this level results in reduction of contribution.	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20241	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73453695 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,112 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; + +import javax.validation.constraints.NotNull; + +import org.junit.Assert; +import org.junit.Test; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.Attribute; +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; +import com.datatorrent.stram.api.OperatorDeployInfo; + +public class OperatorContextTest +{ + private static boolean foundOperatorName; + + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); --- End diff -- It is hard to say what is better - to have a test that may fail on unrelated changes or not to have a test at all. How long will it take to remove an extra operator? How long will it take to modify test not to rely on aggregate? It is a small test and it is not necessary to make it more complex than it needs to be. I suggest that we stop discussion on "nitpicking" and "hard time" on the PR and use it only to discuss code changes.	1	1	https://issues.apache.org/jira/browse/APEXCORE-448

#	Requirement	Questions	Answered	Successfully Answered	URL
20242	APEXCORE-448	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/364#discussion_r73454152 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/OperatorContextTest.java --- @@ -0,0 +1,99 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.engine; +import java.util.Collection; +import java.util.Map; +import java.util.concurrent.CountDownLatch; +import java.util.concurrent.TimeUnit; +import org.junit.Assert; +import org.junit.Test; +import org.apache.hadoop.conf.Configuration; +import com.google.common.base.Preconditions; +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; +public class OperatorContextTest +{ + @Test + public void testInjectionOfOperatorName() throws Exception + { + final LocalMode lma = LocalMode.newInstance(); + final CountDownLatch latch = new CountDownLatch(1); + StreamingApplication testApp = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + MockInputOperator input = dag.addOperator("input", new MockInputOperator()); + input.countDownLatch = latch; + GenericNodeTest.GenericOperator output = dag.addOperator("output", new GenericNodeTest.GenericOperator()); + dag.addStream("stream", input.output, output.ip1); + } + }; + lma.prepareDAG(testApp, new Configuration()); + LocalMode.Controller lc = lma.getController(); + lc.runAsync(); + latch.await(2, TimeUnit.SECONDS); + lc.shutdown(); + } + private static class MockInputOperator extends BaseOperator implements InputOperator, AutoMetric.Aggregator + { + private transient CountDownLatch countDownLatch; + @AutoMetric + String operatorName; + public final transient DefaultOutputPort<String> output = new DefaultOutputPort<>(); + @Override + public void setup(Context.OperatorContext context) + { + operatorName = Preconditions.checkNotNull(context.getName(), "operator name"); --- End diff -- How `Preconditions.checkNotNull` is different from `Assert.assertEquals("operator name", "input", operatorName)`? At the end both will throw an exception. If that exception does not lead to test failure, why `Preconditions.checkNotNull` is necessary? Why modifications to operatorName on operator thread will be immediately visible on the main thread?	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20243	APEXCORE-448	Github user PramodSSImmaneni commented on the issue: https://github.com/apache/apex-core/pull/364 I second it, please reopen this PR and I will merge it. I think we have spent enough time on the tests and everyone's inputs have been valuable. There will be other opportunities in future to add more tests and refactor them. My only small concern is adding a dependency to OperatorDeployInfo in OperatorContext, can we completely justify that? I liked the earlier way where @chandnisingh was passing the extra name as well, null check or without works for me.	1	1	https://issues.apache.org/jira/browse/APEXCORE-448
20244	APEXCORE-426	Adding an answer to the question asked by [~PramodSSImmaneni] Q: Just for completeness can you explain on the JIRA what happens if there were plan changes that were saved but before stram could affect them it got shutdown. Will the new instance make those changes. A: When a Stram recovers, It will use the last checkpointed plan, if the running containers are any different from the expected plan, following things will happen 1. Containers unknown to Stram are killed/rejected 2. Containers which are not responding will be rescheduled.	1	1	https://issues.apache.org/jira/browse/APEXCORE-426
20245	APEXCORE-426	How long will it take for 2. to happen? HEARTBEAT_TIMEOUT?	0	0	https://issues.apache.org/jira/browse/APEXCORE-426
20246	APEXCORE-419	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/293#issuecomment-205469590 Please check what happens after this message is logged. If the affected operators are reset, then why do we need this message?	1	1	https://issues.apache.org/jira/browse/APEXCORE-419
20247	APEXCORE-411	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/294#discussion_r58482768 --- Diff: engine/src/main/java/com/datatorrent/stram/cli/DTCli.java --- @@ -1861,49 +1862,60 @@ public void execute(String[] args, ConsoleReader reader) throws Exception } catch (Exception ex) { throw new CliException("Error opening the config XML file: " + configFile, ex); } - String fileName = expandFileName(commandLineInfo.args[0], true); StramAppLauncher submitApp; AppFactory appFactory = null; - String matchString = commandLineInfo.args.length >= 2 ? commandLineInfo.args[1] : null; - if (fileName.endsWith(".json")) { - File file = new File(fileName); - submitApp = new StramAppLauncher(file.getName(), config); - appFactory = new StramAppLauncher.JsonFileAppFactory(file); - if (matchString != null) { - LOG.warn("Match string \"{}\" is ignored for launching applications specified in JSON", matchString); - } - } else if (fileName.endsWith(".properties")) { - File file = new File(fileName); - submitApp = new StramAppLauncher(file.getName(), config); - appFactory = new StramAppLauncher.PropertyFileAppFactory(file); - if (matchString != null) { - LOG.warn("Match string \"{}\" is ignored for launching applications specified in properties file", matchString); + String matchString = null; + if (commandLineInfo.args.length == 0) { + if (commandLineInfo.origAppId == null) { + throw new CliException("APA or JAR file is required for launching an application when not recovering from a terminated application"); --- End diff -- "when not resuming terminated application" ?	1	1	https://issues.apache.org/jira/browse/APEXCORE-411

#	Requirement	Questions	Answered	Successfully Answered	URL
20248	APEXCORE-411	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/294#discussion_r58482984 --- Diff: engine/src/main/java/com/datatorrent/stram/client/StramAppLauncher.java --- @@ -210,14 +239,67 @@ public StramAppLauncher(FileSystem fs, Path path, Configuration conf) throws Exc public StramAppLauncher(String name, Configuration conf) throws Exception { this.propertiesBuilder = new LogicalPlanConfiguration(conf); + this.conf = conf; init(name); } + public StramAppLauncher(FileSystem fs, Configuration conf) throws Exception + { + this.propertiesBuilder = new LogicalPlanConfiguration(conf); + this.fs = fs; + this.conf = conf; + init(); + } + public String getMvnBuildClasspathOutput() { return mvnBuildClasspathOutput.toString(); } + /** + * This is for recovering an app without specifying apa or appjar file + * @throws Exception + */ + private void init() throws Exception + { + String originalAppId = propertiesBuilder.conf.get(ORIGINAL_APP_ID); + if (originalAppId == null) { + throw new AssertionError("Need original app id if launching without apa or appjar"); + } + Path appsBasePath = new Path(StramClientUtils.getDTDFSRootDir(fs, conf), StramClientUtils.SUBDIR_APPS); + Path origAppPath = new Path(appsBasePath, originalAppId); + StringWriter writer = new StringWriter(); + try (FSDataInputStream in = fs.open(new Path(origAppPath, "meta.json"))) { + IOUtils.copy(in, writer); + } + JSONObject metaJson = new JSONObject(writer.toString()); + String originalLibJars = null; + try { + JSONObject attributes = metaJson.getJSONObject("attributes"); + originalLibJars = attributes.getString(LogicalPlan.LIBRARY_JARS.getSimpleName()); + recoveryAppName = attributes.getString(Context.DAGContext.APPLICATION_NAME.getSimpleName()); + } catch (JSONException ex) { + // ignore + } + + LinkedHashSet<URL> clUrls = new LinkedHashSet<>(); --- End diff -- What's happening here? Weren't all the dependencies already copied to HDFS?	1	1	https://issues.apache.org/jira/browse/APEXCORE-411
20249	APEXCORE-411	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/294#discussion_r59095579 --- Diff: engine/src/main/java/com/datatorrent/stram/client/StramAppLauncher.java --- @@ -203,21 +231,75 @@ public StramAppLauncher(FileSystem fs, Path path, Configuration conf) throws Exc this.fs = fs; fs.copyToLocalFile(path, new Path(localJarFile.getAbsolutePath())); this.jarFile = localJarFile; + this.conf = conf; this.propertiesBuilder = new LogicalPlanConfiguration(conf); init(this.jarFile.getName()); } public StramAppLauncher(String name, Configuration conf) throws Exception { this.propertiesBuilder = new LogicalPlanConfiguration(conf); + this.conf = conf; init(name); } + public StramAppLauncher(FileSystem fs, Configuration conf) throws Exception + { + this.propertiesBuilder = new LogicalPlanConfiguration(conf); + this.fs = fs; + this.conf = conf; + init(); + } + public String getMvnBuildClasspathOutput() { return mvnBuildClasspathOutput.toString(); } + /** + * This is for recovering an app without specifying apa or appjar file --- End diff -- Might be better to move the comment on the constructor?	1	1	https://issues.apache.org/jira/browse/APEXCORE-411
20250	APEXCORE-408	removed word document, and added link to shared google drive document. Can we make use of https://cwiki.apache.org/confluence/display/APEX/Apache+Apex+Home for enhancement proposals and design documents.	1	1	https://issues.apache.org/jira/browse/APEXCORE-408
20251	APEXCORE-408	How about widening the scope to also trigger the changes when major events happen in the execution such as operators getting killed or choosing to shutdown? In general whenever a stram event gets created we could potentially trigger the changes.	1	0	https://issues.apache.org/jira/browse/APEXCORE-408
20252	APEXCORE-405	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/351#discussion_r69830044 --- Diff: api/src/main/java/com/datatorrent/api/ClusterMode.java --- @@ -0,0 +1,67 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.api; + +import com.datatorrent.api.Launcher.AppHandle; + +/** + * + */ +public abstract class ClusterMode<H extends AppHandle> extends Launcher<H> { + public static final Attribute<String> LIB_JARS = new Attribute<String>(new StringCodec.String2String()); --- End diff -- Maybe use the diamond operator in Java 7?	1	1	https://issues.apache.org/jira/browse/APEXCORE-405
20253	APEXCORE-405	Github user tweise commented on the issue: https://github.com/apache/apex-core/pull/351 Before getting into details, how will packaging and configuration related aspects be handled when using the proposed API? In embedded mode all of that is implied, but for launching an app on the YARN cluster the configuration is critical. The CLI expects an app package and sources configuration from different locations.	1	1	https://issues.apache.org/jira/browse/APEXCORE-405
20254	APEXCORE-405	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/351#discussion_r69962835 --- Diff: api/src/main/java/com/datatorrent/api/Attribute.java --- @@ -146,6 +146,8 @@ public String toString() Set<Map.Entry<Attribute<?>, Object>> entrySet(); + <T> T getValue(Attribute<T> key); --- End diff -- why do you need this? There is get(Attribute<T> key) already..	1	1	https://issues.apache.org/jira/browse/APEXCORE-405
20255	APEXCORE-405	Github user tweise commented on the issue: https://github.com/apache/apex-core/pull/351 @PramodSSImmaneni Parity with embedded mode is not sufficient. Launching an application on a YARN cluster is very different and needs to be considered accordingly. How does the user figure out what the dependencies are, how will the configuration object be put together?	1	1	https://issues.apache.org/jira/browse/APEXCORE-405
20256	APEXCORE-405	Github user gauravgopi123 commented on the issue: https://github.com/apache/apex-core/pull/351 @PramodSSImmaneni : Can you please point me to the discussion on dev for this? I don't see anything on Jira for this..	1	1	https://issues.apache.org/jira/browse/APEXCORE-405
20257	APEXCORE-402	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/286#discussion_r61825724 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/AbstractReservoir.java --- @@ -47,7 +49,7 @@ { private static final Logger logger = LoggerFactory.getLogger(AbstractReservoir.class); static final String reservoirClassNameProperty = "com.datatorrent.stram.engine.Reservoir"; - private static final int USE_SPSC_CAPACITY = 8 * 1024; + private static final int USE_SPSC_CAPACITY = 64 * 1024; --- End diff -- Name, the constant now is the threshold for blocking queue, it's spsc in either case. Why did it increase?	1	1	https://issues.apache.org/jira/browse/APEXCORE-402

#	Requirement	Questions	Answered	Successfully Answered	URL
20258	APEXCORE-401	Github user vrozov commented on the pull request: https://github.com/apache/incubator-apex-core/pull/290#issuecomment-204434631 Is it necessary to have a separate artifactId for codestyle and publish it to apache maven? Can the same goal be achieved by uploading codestyle definition file to a central location? For item 1, should checkstyle module be moved into a separate project?	1	1	https://issues.apache.org/jira/browse/APEXCORE-401
20259	APEXCORE-401	Github user chandnisingh commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/290#discussion_r58271890 --- Diff: codestyle-config/pom.xml --- @@ -0,0 +1,35 @@ +<?xml version="1.0" encoding="UTF-8"?> +<!-- + + Licensed to the Apache Software Foundation (ASF) under one + or more contributor license agreements. See the NOTICE file + distributed with this work for additional information + regarding copyright ownership. The ASF licenses this file + to you under the Apache License, Version 2.0 (the + "License"); you may not use this file except in compliance + with the License. You may obtain a copy of the License at + + http://www.apache.org/licenses/LICENSE-2.0 + + Unless required by applicable law or agreed to in writing, + software distributed under the License is distributed on an + "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + KIND, either express or implied. See the License for the + specific language governing permissions and limitations + under the License. + +--> +<project xmlns="http://maven.apache.org/POM/4.0.0" + xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" + xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd"> + <parent> + <groupId>org.apache.apex</groupId> + <artifactId>apex</artifactId> + <version>3.4.0-incubating-SNAPSHOT</version> + </parent> + <modelVersion>4.0.0</modelVersion> + + <artifactId>apex-codestyle-config</artifactId> + <name>Apache Apex (incubating) Codestyle Config</name> + --- End diff -- Add a version here? <version>1.0.0-incubating</version>	1	1	https://issues.apache.org/jira/browse/APEXCORE-401
20260	APEXCORE-401	Github user chandnisingh commented on the pull request: https://github.com/apache/incubator-apex-core/pull/290#issuecomment-204581927 @tweise @vrozov ok. does that mean that I right now I can't change the main apex pom.xml to use the new artifact as it is not present at the maven central repo. Only once the new artifact is available there, then I can modify the apex pom?	1	1	https://issues.apache.org/jira/browse/APEXCORE-401
20261	APEXCORE-401	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/290#issuecomment-204616808 checkstyle-suppressions.xml contains configuration that is only valid for -core. Can it be kept local?	1	1	https://issues.apache.org/jira/browse/APEXCORE-401
20262	APEXCORE-399	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/278#discussion_r56826832 --- Diff: engine/src/main/java/com/datatorrent/stram/security/StramWSFilterInitializer.java --- @@ -59,9 +59,11 @@ public void initFilter(FilterContainer container, Configuration conf) for (String rmId : ConfigUtils.getRMHAIds(conf)) { proxies.add(getResolvedRMWebAppURLWithoutScheme(conf, rmId)); } + logger.info("HA proxy addresses {}", proxies); --- End diff -- Should it be logged at the debug level?	1	1	https://issues.apache.org/jira/browse/APEXCORE-399
20263	APEXCORE-386	[~cowtowncoder] That should be a small change, given it only moves the patch version. Can you submit a pull request?	1	0	https://issues.apache.org/jira/browse/APEXCORE-386
20264	APEXCORE-386	Were any sanity tests specific to this change run before merging the PR?	1	0	https://issues.apache.org/jira/browse/APEXCORE-386
20265	APEXCORE-380	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/271#issuecomment-196593216 @vrozov how much does this change increase the throughput?	1	1	https://issues.apache.org/jira/browse/APEXCORE-380
20266	APEXCORE-380	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/271#discussion_r56440964 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/GenericNode.java --- @@ -608,9 +610,9 @@ else if (tracker.ports[trackerIndex] == null) { if (need2sleep) { if (handleIdleTime && insideWindow) { ((IdleTimeHandler) operator).handleIdleTime(); - } - else { + } else { Thread.sleep(spinMillis); + spinMillis = Math.min(maxSpinMillis, spinMillis + 1); --- End diff -- I had looked at this before in a difference scenario. Can we do a better heuristic. We can have another timeout setting (which is less than streaming window size) where we continue to do busy way till that timeout and if there is no data till then, then initiate the stepwise sleep.	1	1	https://issues.apache.org/jira/browse/APEXCORE-380
20267	APEXCORE-378	[Yogi] Hi, I studied source code, docs related to windowing for my presentation on "Windowing in Apex". After going through the phase of learning, teaching apex windowing concepts to others; I realized that attribute APPLICATION_WINDOW_COUNT from OperatorContext should have better name. Because: 1. From the name it looks like a attribute of the APPLICATION; but actually it is an Operator attribute. 2. It is very confusing to say that APPLICATION_WINDOW_COUNT can be different for different Operators within an Application. Can we have some better name for this attribute? Any suggestions for the new name? I propose following names: 1. OPERATOR_WINDOW_COUNT 2. OPERATOR_WINDOW_SPAN I understand that, introducing new name to the existing attribute might break backward compatibility. But, we can deal with that by deprecating the older name and keeping it alive till next major release. What is your opinion on this?	1	1	https://issues.apache.org/jira/browse/APEXCORE-378

#	Requirement	Questions	Answered	Successfully Answered	URL
20268	APEXCORE-354	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-site/pull/14#discussion_r54684986 --- Diff: src/md/community.md --- @@ -4,9 +4,9 @@ The following mailing lists are available for Apex. -- [dev@apex.incubator.apache.org](http://mail-archives.apache.org/mod_mbox/incubator-apex-dev/) is for people who want to contribute code to Apex. [subscribe](mailto:dev-subscribe@apex.incubator.apache.org?subject=send%20this%20email%20to%20subscribe), [unsubscribe](mailto:dev-unsubscribe@apex.incubator.apache.org?subject=send%20this%20email%20to%20unsubscribe), [archives](http://mail-archives.apache.org/mod_mbox/incubator-apex-dev/) -- [users@apex.incubator.apache.org](http://mail-archives.apache.org/mod_mbox/incubator-apex-users/) is a discussion forum for people who want to build apps and/or solutions using Apex. [subscribe](mailto:users-subscribe@apex.incubator.apache.org?subject=send%20this%20email%20to%20subscribe), [unsubscribe](mailto:users-unsubscribe@apex.incubator.apache.org?subject=send%20this%20email%20to%20unsubscribe), [archives](http://mail-archives.apache.org/mod_mbox/incubator-apex-users/) -- [commits@apex.incubator.apache.org](http://mail-archives.apache.org/mod_mbox/incubator-apex-commits/) is for commit messages and patches to Apex. [subscribe](mailto:commits-subscribe@apex.incubator.apache.org?subject=send%20this%20email%20to%20subscribe), [unsubscribe](mailto:commits-unsubscribe@apex.incubator.apache.org?subject=send%20this%20email%20to%20unsubscribe), [archives](http://mail-archives.apache.org/mod_mbox/incubator-apex-commits/) +- [dev@apex.incubator.apache.org](http://mail-archives.apache.org/mod_mbox/incubator-apex-dev/) is for people who want to contribute code to Apex. [subscribe](mailto:dev-subscribe@apex.incubator.apache.org?subject=send%20this%20email%20to%20subscribe), [unsubscribe](mailto:dev-unsubscribe@apex.incubator.apache.org?subject=send%20this%20email%20to%20unsubscribe), [forum](https://s.apache.org/apex-dev), [archives](http://mail-archives.apache.org/mod_mbox/incubator-apex-dev/) --- End diff -- See comment below. How does that work? Communication needs to be on the mailing list.	1	1	https://issues.apache.org/jira/browse/APEXCORE-354
20269	APEXCORE-354	Github user sashadt commented on the pull request: https://github.com/apache/incubator-apex-site/pull/14#issuecomment-194001421 @chinmaykolhatkar What makes something a forum is ability to discuss the topic. However, if this discussion never makes it back to dev@ or user@ mailing lists, then we have a problem. Are you saying Nabble takes care of sending the communication back via shared mailing lists? That's what I'm trying to confirm.	1	1	https://issues.apache.org/jira/browse/APEXCORE-354
20270	APEXCORE-348	Conversation from email Pramod Immaneni <pramod@datatorrent.com> Inline On Thu, Feb 11, 2016 at 4:32 PM, Timothy Farkas <tim@datatorrent.com> wrote: Comments inline +1 Overall as well provided Apex-339 is implemented first and it is documented that the mechanism should not be used with some stateful operators. On Thu, Feb 11, 2016 at 4:20 PM, Pramod Immaneni <pramod@datatorrent.com> wrote: > Comments inline > > > On Thu, Feb 11, 2016 at 4:13 PM, Timothy Farkas <tim@datatorrent.com> > wrote: > > > > Hey Pramod, > > > > I agree if APEX-339 is in place then it would work without redeploying > > containers for operators that are Stateless, or a subset of Stateful > > operators. > > > > Addressing your previous questions. > > > > - The StatsListener can be used to see how far behind operators are. You > > could determine what window the operator is on, or the number of tuples > > it's processed so far, or how long > > it takes it to complete a window. > > > > What if tuples are different sizes and number of tuples processed doesn't > reflect how far ahead or behind a downstream partition is? How is the > information from StatsListener made available to the upstream partition > codecs. > What is the information Buffer Server can provide that the StatsListener cannot? The stats information would have to be relayed down to the upstream operators. It's possible. The StatsListener can trigger a repartition. The information in the StatsListener can be shared with the partitioner by setting the same object for both in populate Dag. The partitioner can then compute the new Stream Codec. The mechanism by which the upstream would be updated with the new Stream Codec would have to be implemented as it's currently not there. > > > > > > - Some examples of Stateful operators that require repartitioning of > state > > are the following: > > - Deduper > > In this case after updating the stream codec the operator may > > allow a previously seen value to pass because the partition didn't > receive > > that value with the previous stream codec. > > - A key value store that holds aggregations for each key. > > In this case multiple partitions would hold partial > aggregations > > for a key, when they are expecting to hold the complete aggregation. > > > > Agreed for deduper. For the second case a unifier is a better approach so > that you are not affected by key skew in general. > This is not always possible. We can discuss this offline, since it won't add much to the discussion here to go into the details. Yes not always.	1	0	https://issues.apache.org/jira/browse/APEXCORE-348
20271	APEXCORE-347	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/226#discussion_r53106756 --- Diff: pom.xml --- @@ -63,6 +63,7 @@ <jersey.version>1.9</jersey.version> <!-- do not change jetty version as later versions have problems with DefaultServlet --> <jetty.version>8.1.10.v20130312</jetty.version> + <oldVersion.version>3.3.0-incubating</oldVersion.version> --- End diff -- Would japicmp.oldVersion be more appropriate?	1	1	https://issues.apache.org/jira/browse/APEXCORE-347
20272	APEXCORE-347	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/226#discussion_r53107901 --- Diff: pom.xml --- @@ -63,6 +63,7 @@ <jersey.version>1.9</jersey.version> <!-- do not change jetty version as later versions have problems with DefaultServlet --> <jetty.version>8.1.10.v20130312</jetty.version> + <oldVersion.version>3.3.0-incubating</oldVersion.version> --- End diff -- project.old.version? japicmp.oldVersion sounds related to japicmp version.	1	1	https://issues.apache.org/jira/browse/APEXCORE-347
20273	APEXCORE-347	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/226#issuecomment-184961199 Looks good. How about waiting with this PR until semver option in the plugin is available?	1	1	https://issues.apache.org/jira/browse/APEXCORE-347
20274	APEXCORE-344	Shutdown takes a while to finish. Can you please file a jira with your observation?	0	0	https://issues.apache.org/jira/browse/APEXCORE-344
20275	APEXCORE-340	Besides the symlink, the proper name change would involve: - Change the name of DTCli class to ApexCli - Change the prompt from dt> to apex> - Change the initial greeting when running the CLI from something like "DT CLI 3.3.0-incubating ..." to "Apex CLI ...". Although the above will break existing 3rd party scripts that rely on the the string "DT", the existence of such scripts is unlikely. What do people think if we actually make those changes without waiting for the major version bump?	1	1	https://issues.apache.org/jira/browse/APEXCORE-340
20276	APEXCORE-340	Hmm. How will it break the yahoo benchmark if we retain dtcli as a symlink?	1	1	https://issues.apache.org/jira/browse/APEXCORE-340

#	Requirement	Questions	Answered	Successfully Answered	URL
20277	APEXCORE-340	Made following changes in my fork: https://github.com/chinmaykolhatkar/incubator-apex-core/tree/APEXCORE-340 1. Script name changed to apex. Created dtcli as symlink to apex 2. Class name changed to ApexCli from DTCli. Similar for test class. 3. Changed console prompt to "apex>" from "dt>" 4. Changed Welcome message to "Apex CLI ..." from "DT CLI ..." Commit link is: https://github.com/chinmaykolhatkar/incubator-apex-core/commit/5ef35e83323504027aef115dd01f7278c69cb1c1 Verified with unit test and launching pi demo. Before I create a PR, I have some questions: 1. [~sandesh] : Even though there is dtcli symlink to apex, the prompt of the cli will change to "apex>" from "dtcli>"... Would that cause any problem in yahoo streaming benchmark? 2. [~davidyan] [~thw] : We There are still traces of "dt" all over in this. For eg. ".dt" folder, DT_CLASSPATH and other such variables in script. Should those also be changed as a part of this?	1	1	https://issues.apache.org/jira/browse/APEXCORE-340
20278	APEXCORE-340	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/303#issuecomment-209099592 docs/images/security/image02_apex.png Why are we not replacing the existing image?	1	1	https://issues.apache.org/jira/browse/APEXCORE-340
20279	APEXCORE-337	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/224#discussion_r52406007 --- Diff: pom.xml --- @@ -328,7 +328,7 @@ <dependency> <groupId>com.puppycrawl.tools</groupId> <artifactId>checkstyle</artifactId> - <version>6.14</version> + <version>6.15</version> --- End diff -- @vrozov : Why is this upgrade required? I don't see anything in Jira	1	1	https://issues.apache.org/jira/browse/APEXCORE-337
20280	APEXCORE-330	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/321#discussion_r61779139 --- Diff: engine/src/main/java/com/datatorrent/stram/util/StackTrace.java --- @@ -0,0 +1,77 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.util; + +import java.util.Map; + +import org.codehaus.jettison.json.JSONArray; +import org.codehaus.jettison.json.JSONException; +import org.codehaus.jettison.json.JSONObject; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +public class StackTrace +{ + private static final Logger LOG = LoggerFactory.getLogger(StackTrace.class); + + public static String getJsonFormat() --- End diff -- How about just returning JSONObject instead of String?	1	1	https://issues.apache.org/jira/browse/APEXCORE-330
20281	APEXCORE-330	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/321#discussion_r61779943 --- Diff: engine/src/main/java/com/datatorrent/stram/util/StackTrace.java --- @@ -0,0 +1,77 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.util; + +import java.util.Map; + +import org.codehaus.jettison.json.JSONArray; +import org.codehaus.jettison.json.JSONException; +import org.codehaus.jettison.json.JSONObject; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +public class StackTrace +{ + private static final Logger LOG = LoggerFactory.getLogger(StackTrace.class); + + public static String getJsonFormat() --- End diff -- How about moving this utility function to StramUtils.java?	1	1	https://issues.apache.org/jira/browse/APEXCORE-330
20282	APEXCORE-330	Github user PramodSSImmaneni commented on the pull request: https://github.com/apache/incubator-apex-core/pull/321#issuecomment-216398640 Can we provide the user a way to distinguish between threads directly related to the operators and user code vs system threads that are apex specific. It could be a flag in the json.	1	1	https://issues.apache.org/jira/browse/APEXCORE-330
20283	APEXCORE-330	I noticed that the individual containers are sending a JSON serialized string of the stack trace information to STRAM. The serialization is typically handled by RPC and the containers just send objects. Could you convert the Map<Thread, StackTraceElement[]> into an object that can be serialized, maybe into a Map<ThreadInfo, StackTraceElement[]>. Secondly the web service call results in a request flag to be set in the container agent which gets reset immediately after initiating a request to the container to send the stack trace. I believe this flag should be reset after receiving the stack trace from the container as there is a chance the request may not make it to the container if there are any intermittent network errors and the web service call is anyway waiting for a timeout period of 10s for the stack trace which is several heartbeats. I apologize I was not able to give these comments during the pull request review.	0	0	https://issues.apache.org/jira/browse/APEXCORE-330
20284	APEXCORE-328	Found the following message in the `screen.txt` that may explain why maven invocation from the unit test fails: {noformat} Error: Could not find or load main class org.codehaus.plexus.classworlds.launcher.Launcher {noformat}	1	0	https://issues.apache.org/jira/browse/APEXCORE-328
20285	APEXCORE-327	Can you please make the equivalent change for the release-3.3 branch? Also, is there a way to automate this, i.e. have it find out automatically based on the version number that modifications are not allowed (if the version is not X.Y.0, then enable the constraint)? Or is this an enhancement for the plugin?	0	0	https://issues.apache.org/jira/browse/APEXCORE-327
20286	APEXCORE-321	Github user chinmaykolhatkar commented on the pull request: https://github.com/apache/incubator-apex-site/pull/22#issuecomment-197170552 What is the parent page for bylaws page to be redirected from?	1	1	https://issues.apache.org/jira/browse/APEXCORE-321

#	Requirement	Questions	Answered	Successfully Answered	URL
20287	APEXCORE-321	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/incubator-apex-site/pull/22#discussion_r56286607 --- Diff: src/md/bylaws.md --- @@ -0,0 +1,381 @@ +#*** draft subject to top level project approval *** +#Apache Apex Project Bylaws + +This document defines the bylaws under which the Apache Apex project operates. It +defines the roles and responsibilities of the project, who may vote, how voting works, +how conflicts are resolved, etc. + +Apex is a project of the Apache Software Foundation. The Foundation +holds the copyright on Apache code including the code in the Apex codebase. The +Foundation +FAQ explains the operation and background of the foundation. + +Apex is typical of Apache projects in that it operates under a set of principles, +known collectively as the "Apache Way". If you are new to Apache +development, please refer to the Incubator project for more information on how Apache +projects operate. + +##Roles and Responsibilities + +Apache projects define a set of roles with +associated rights and responsibilities. These roles govern what tasks an individual +may perform within the project. The roles are defined in the following sections: + +###Users + +The most important participants in the project are people who use our software. +The majority of our developers start out as users and guide their development +efforts from the user’s perspective. + +Users contribute to the Apache projects by providing feedback to developers in +the form of bug reports and feature suggestions. As well, users participate in the +Apache community by helping other users on mailing lists and user support +forums. + +###Contributors + +All of the volunteers who are contributing time, code, documentation, or +resources to the Apex Project. A contributor that makes sustained, welcome +contributions to the project may be invited to become a Committer, though the exact +timing of such invitations depends on many factors. + +###Committers + +The project’s Committers are responsible for the project’s technical +management. All committers have write access to the project’s source +repositories. Committers may cast binding votes on any technical discussion --- End diff -- I was under impression that PPMC members have binding votes. Is this changed?	1	1	https://issues.apache.org/jira/browse/APEXCORE-321
20288	APEXCORE-321	src/md/bylaws.md --- @@ -0,0 +1,381 @@ +#*** draft subject to top level project approval *** +#Apache Apex Project Bylaws + +This document defines the bylaws under which the Apache Apex project operates. It +defines the roles and responsibilities of the project, who may vote, how voting works, +how conflicts are resolved, etc. + +Apex is a project of the Apache Software Foundation. The Foundation +holds the copyright on Apache code including the code in the Apex codebase. The +Foundation +FAQ explains the operation and background of the foundation. + +Apex is typical of Apache projects in that it operates under a set of principles, +known collectively as the "Apache Way". If you are new to Apache +development, please refer to the Incubator project for more information on how Apache +projects operate. + +##Roles and Responsibilities + +Apache projects define a set of roles with +associated rights and responsibilities. These roles govern what tasks an individual +may perform within the project. The roles are defined in the following sections: + +###Users + +The most important participants in the project are people who use our software. +The majority of our developers start out as users and guide their development +efforts from the user’s perspective. + +Users contribute to the Apache projects by providing feedback to developers in +the form of bug reports and feature suggestions. As well, users participate in the +Apache community by helping other users on mailing lists and user support +forums. + +###Contributors + +All of the volunteers who are contributing time, code, documentation, or +resources to the Apex Project. A contributor that makes sustained, welcome +contributions to the project may be invited to become a Committer, though the exact +timing of such invitations depends on many factors. + +###Committers + +The project’s Committers are responsible for the project’s technical +management. All committers have write access to the project’s source +repositories. Committers may cast binding votes on any technical discussion +regarding the project. + +Committer access is by invitation only and must be approved by lazy consensus of +the active PMC members. A Committer may request removal of their commit privileges +by their own declaration. A committer will be considered +“emeritus/inactive” by not contributing in any form to the project for +over 1 year. An emeritus committer may request reinstatement of commit access from +the PMC. Such reinstatement is subject to lazy consensus of active PMC members + +Commit access can be revoked by a unanimous vote of all the active PMC members +(except the committer in question if they are also a PMC member). + +All Apache committers are required to have a signed Contributor License +Agreement (CLA) on file with the Apache Software +Foundation. There is a Committer FAQ which provides more +details on the requirements for Committers + +A committer who makes a sustained contribution to the project may be invited to +become a member of the PMC. The form of contribution is not limited to code. It can +also include code review, helping out users on the mailing lists, documentation, +etc. + +###Project Management Committee + +The Project Management Committee (PMC) for Apache Apex was created by a +resolution of the board of the Apache Software Foundation on ***TBD***. The +PMC is responsible to the board and the ASF for the management and oversight of the +Apache Apex codebase. The responsibilities of the PMC include: + +* Deciding what is distributed as products of the project. In particular all releases must be approved by the PMC *	1	1	https://issues.apache.org/jira/browse/APEXCORE-321

#	Requirement	Questions	Answered	Successfully Answered	URL
20289	APEXCORE-321	src/md/bylaws.md --- @@ -0,0 +1,381 @@ +#*** draft subject to top level project approval *** +#Apache Apex Project Bylaws + +This document defines the bylaws under which the Apache Apex project operates. It +defines the roles and responsibilities of the project, who may vote, how voting works, +how conflicts are resolved, etc. + +Apex is a project of the Apache Software Foundation. The Foundation +holds the copyright on Apache code including the code in the Apex codebase. The +Foundation +FAQ explains the operation and background of the foundation. + +Apex is typical of Apache projects in that it operates under a set of principles, +known collectively as the "Apache Way". If you are new to Apache +development, please refer to the Incubator project for more information on how Apache +projects operate. + +##Roles and Responsibilities + +Apache projects define a set of roles with +associated rights and responsibilities. These roles govern what tasks an individual +may perform within the project. The roles are defined in the following sections: + +###Users + +The most important participants in the project are people who use our software. +The majority of our developers start out as users and guide their development +efforts from the user’s perspective. + +Users contribute to the Apache projects by providing feedback to developers in +the form of bug reports and feature suggestions. As well, users participate in the +Apache community by helping other users on mailing lists and user support +forums. + +###Contributors + +All of the volunteers who are contributing time, code, documentation, or +resources to the Apex Project. A contributor that makes sustained, welcome +contributions to the project may be invited to become a Committer, though the exact +timing of such invitations depends on many factors. + +###Committers + +The project’s Committers are responsible for the project’s technical +management. All committers have write access to the project’s source +repositories. Committers may cast binding votes on any technical discussion +regarding the project. + +Committer access is by invitation only and must be approved by lazy consensus of +the active PMC members. A Committer may request removal of their commit privileges +by their own declaration. A committer will be considered +“emeritus/inactive” by not contributing in any form to the project for +over 1 year. An emeritus committer may request reinstatement of commit access from +the PMC. Such reinstatement is subject to lazy consensus of active PMC members + +Commit access can be revoked by a unanimous vote of all the active PMC members +(except the committer in question if they are also a PMC member). + +All Apache committers are required to have a signed Contributor License +Agreement (CLA) on file with the Apache Software +Foundation. There is a Committer FAQ which provides more +details on the requirements for Committers + +A committer who makes a sustained contribution to the project may be invited to +become a member of the PMC. The form of contribution is not limited to code. It can +also include code review, helping out users on the mailing lists, documentation, +etc. + +###Project Management Committee + +The Project Management Committee (PMC) for Apache Apex was created by a +resolution of the board of the Apache Software Foundation on ***TBD***. The +PMC is responsible to the board and the ASF for the management and oversight of the +Apache Apex codebase. The responsibilities of the PMC include: + +* Deciding what is distributed as products of the project. In particular all releases must be approved by the PMC *	1	1	https://issues.apache.org/jira/browse/APEXCORE-321
20290	APEXCORE-321	src/md/bylaws.md --- @@ -0,0 +1,381 @@ +#*** draft subject to top level project approval *** +#Apache Apex Project Bylaws + +This document defines the bylaws under which the Apache Apex project operates. It +defines the roles and responsibilities of the project, who may vote, how voting works, +how conflicts are resolved, etc. + +Apex is a project of the Apache Software Foundation. The Foundation +holds the copyright on Apache code including the code in the Apex codebase. The +Foundation +FAQ explains the operation and background of the foundation. + +Apex is typical of Apache projects in that it operates under a set of principles, +known collectively as the "Apache Way". If you are new to Apache +development, please refer to the Incubator project for more information on how Apache +projects operate. + +##Roles and Responsibilities + +Apache projects define a set of roles with +associated rights and responsibilities. These roles govern what tasks an individual +may perform within the project. The roles are defined in the following sections: + +###Users + +The most important participants in the project are people who use our software. +The majority of our developers start out as users and guide their development +efforts from the user’s perspective. + +Users contribute to the Apache projects by providing feedback to developers in +the form of bug reports and feature suggestions. As well, users participate in the +Apache community by helping other users on mailing lists and user support +forums. + +###Contributors + +All of the volunteers who are contributing time, code, documentation, or +resources to the Apex Project. A contributor that makes sustained, welcome +contributions to the project may be invited to become a Committer, though the exact +timing of such invitations depends on many factors. + +###Committers + +The project’s Committers are responsible for the project’s technical +management. All committers have write access to the project’s source +repositories. Committers may cast binding votes on any technical discussion +regarding the project. + +Committer access is by invitation only and must be approved by lazy consensus of +the active PMC members. A Committer may request removal of their commit privileges +by their own declaration. A committer will be considered +“emeritus/inactive” by not contributing in any form to the project for +over 1 year. An emeritus committer may request reinstatement of commit access from +the PMC. Such reinstatement is subject to lazy consensus of active PMC members + +Commit access can be revoked by a unanimous vote of all the active PMC members +(except the committer in question if they are also a PMC member). + +All Apache committers are required to have a signed Contributor License +Agreement (CLA) on file with the Apache Software +Foundation. There is a Committer FAQ which provides more +details on the requirements for Committers + +A committer who makes a sustained contribution to the project may be invited to +become a member of the PMC. The form of contribution is not limited to code. It can +also include code review, helping out users on the mailing lists, documentation, +etc. + +###Project Management Committee + +The Project Management Committee (PMC) for Apache Apex was created by a +resolution of the board of the Apache Software Foundation on ***TBD***. The +PMC is responsible to the board and the ASF for the management and oversight of the +Apache Apex codebase. The responsibilities of the PMC include: + +* Deciding what is distributed as products of the project. In particular all releases must be approved by the PMC *	1	1	https://issues.apache.org/jira/browse/APEXCORE-321

#	Requirement	Questions	Answered	Successfully Answered	URL
20291	APEXCORE-321	src/md/bylaws.md --- @@ -0,0 +1,381 @@ +#*** draft subject to top level project approval *** +#Apache Apex Project Bylaws + +This document defines the bylaws under which the Apache Apex project operates. It +defines the roles and responsibilities of the project, who may vote, how voting works, +how conflicts are resolved, etc. + +Apex is a project of the Apache Software Foundation. The Foundation +holds the copyright on Apache code including the code in the Apex codebase. The +Foundation +FAQ explains the operation and background of the foundation. + +Apex is typical of Apache projects in that it operates under a set of principles, +known collectively as the "Apache Way". If you are new to Apache +development, please refer to the Incubator project for more information on how Apache +projects operate. + +##Roles and Responsibilities + +Apache projects define a set of roles with +associated rights and responsibilities. These roles govern what tasks an individual +may perform within the project. The roles are defined in the following sections: + +###Users + +The most important participants in the project are people who use our software. +The majority of our developers start out as users and guide their development +efforts from the user’s perspective. + +Users contribute to the Apache projects by providing feedback to developers in +the form of bug reports and feature suggestions. As well, users participate in the +Apache community by helping other users on mailing lists and user support +forums. + +###Contributors + +All of the volunteers who are contributing time, code, documentation, or +resources to the Apex Project. A contributor that makes sustained, welcome +contributions to the project may be invited to become a Committer, though the exact +timing of such invitations depends on many factors. + +###Committers + +The project’s Committers are responsible for the project’s technical +management. All committers have write access to the project’s source +repositories. Committers may cast binding votes on any technical discussion +regarding the project. + +Committer access is by invitation only and must be approved by lazy consensus of +the active PMC members. A Committer may request removal of their commit privileges +by their own declaration. A committer will be considered +“emeritus/inactive” by not contributing in any form to the project for +over 1 year. An emeritus committer may request reinstatement of commit access from +the PMC. Such reinstatement is subject to lazy consensus of active PMC members + +Commit access can be revoked by a unanimous vote of all the active PMC members +(except the committer in question if they are also a PMC member). + +All Apache committers are required to have a signed Contributor License +Agreement (CLA) on file with the Apache Software +Foundation. There is a Committer FAQ which provides more +details on the requirements for Committers + +A committer who makes a sustained contribution to the project may be invited to +become a member of the PMC. The form of contribution is not limited to code. It can +also include code review, helping out users on the mailing lists, documentation, +etc. + + +###Project Management Committee + +The Project Management Committee (PMC) for Apache Apex was created by a +resolution of the board of the Apache Software Foundation on ***TBD***. The +PMC is responsible to the board and the ASF for the management and oversight of the +Apache Apex codebase. The responsibilities of the PMC include: + +* Deciding what is distributed as products of the project. In particular all releases must be approved by the PMC *	1	1	https://issues.apache.org/jira/browse/APEXCORE-321
20292	APEXCORE-321	src/md/bylaws.md --- @@ -0,0 +1,381 @@ +#*** draft subject to top level project approval *** +#Apache Apex Project Bylaws + +This document defines the bylaws under which the Apache Apex project operates. It +defines the roles and responsibilities of the project, who may vote, how voting works, +how conflicts are resolved, etc. + +Apex is a project of the Apache Software Foundation. The Foundation +holds the copyright on Apache code including the code in the Apex codebase. The +Foundation +FAQ explains the operation and background of the foundation. + +Apex is typical of Apache projects in that it operates under a set of principles, +known collectively as the "Apache Way". If you are new to Apache +development, please refer to the Incubator project for more information on how Apache +projects operate. + +##Roles and Responsibilities + +Apache projects define a set of roles with +associated rights and responsibilities. These roles govern what tasks an individual +may perform within the project. The roles are defined in the following sections: + +###Users + +The most important participants in the project are people who use our software. +The majority of our developers start out as users and guide their development +efforts from the user’s perspective. + +Users contribute to the Apache projects by providing feedback to developers in +the form of bug reports and feature suggestions. As well, users participate in the +Apache community by helping other users on mailing lists and user support +forums. + +###Contributors + +All of the volunteers who are contributing time, code, documentation, or +resources to the Apex Project. A contributor that makes sustained, welcome +contributions to the project may be invited to become a Committer, though the exact +timing of such invitations depends on many factors. + +###Committers + +The project’s Committers are responsible for the project’s technical +management. All committers have write access to the project’s source +repositories. Committers may cast binding votes on any technical discussion +regarding the project. + +Committer access is by invitation only and must be approved by lazy consensus of +the active PMC members. A Committer may request removal of their commit privileges +by their own declaration. A committer will be considered +“emeritus/inactive” by not contributing in any form to the project for +over 1 year. An emeritus committer may request reinstatement of commit access from +the PMC. Such reinstatement is subject to lazy consensus of active PMC members + +Commit access can be revoked by a unanimous vote of all the active PMC members +(except the committer in question if they are also a PMC member). + +All Apache committers are required to have a signed Contributor License +Agreement (CLA) on file with the Apache Software +Foundation. There is a Committer FAQ which provides more +details on the requirements for Committers + +A committer who makes a sustained contribution to the project may be invited to +become a member of the PMC. The form of contribution is not limited to code. It can +also include code review, helping out users on the mailing lists, documentation, +etc. + + +###Project Management Committee + +The Project Management Committee (PMC) for Apache Apex was created by a +resolution of the board of the Apache Software Foundation on ***TBD***. The +PMC is responsible to the board and the ASF for the management and oversight of the +Apache Apex codebase. The responsibilities of the PMC include: + +* Deciding what is distributed as products of the project. In particular all releases must be approved by the PMC *	1	1	https://issues.apache.org/jira/browse/APEXCORE-321

#	Requirement	Questions	Answered	Successfully Answered	URL
20293	APEXCORE-321	src/md/bylaws.md --- @@ -0,0 +1,381 @@ +#*** draft subject to top level project approval *** +##Apache Apex Project Bylaws + +This document defines the bylaws under which the Apache Apex project operates. It +defines the roles and responsibilities of the project, who may vote, how voting works, +how conflicts are resolved, etc. + +Apex is a project of the Apache Software Foundation. The Foundation +holds the copyright on Apache code including the code in the Apex codebase. The +Foundation +FAQ explains the operation and background of the foundation. + +Apex is typical of Apache projects in that it operates under a set of principles, +known collectively as the "Apache Way". If you are new to Apache +development, please refer to the Incubator project for more information on how Apache +projects operate. + +##Roles and Responsibilities + +Apache projects define a set of roles with +associated rights and responsibilities. These roles govern what tasks an individual +may perform within the project. The roles are defined in the following sections: + +###Users + +The most important participants in the project are people who use our software. +The majority of our developers start out as users and guide their development +efforts from the user’s perspective. + +Users contribute to the Apache projects by providing feedback to developers in +the form of bug reports and feature suggestions. As well, users participate in the +Apache community by helping other users on mailing lists and user support +forums. + +###Contributors + +All of the volunteers who are contributing time, code, documentation, or +resources to the Apex Project. A contributor that makes sustained, welcome +contributions to the project may be invited to become a Committer, though the exact +timing of such invitations depends on many factors. + +###Committers + +The project’s Committers are responsible for the project’s technical +management. All committers have write access to the project’s source +repositories. Committers may cast binding votes on any technical discussion +regarding the project. + +Committer access is by invitation only and must be approved by lazy consensus of +the active PMC members. A Committer may request removal of their commit privileges +by their own declaration. A committer will be considered +“emeritus/inactive” by not contributing in any form to the project for +over 1 year. An emeritus committer may request reinstatement of commit access from +the PMC. Such reinstatement is subject to lazy consensus of active PMC members + +Commit access can be revoked by a unanimous vote of all the active PMC members +(except the committer in question if they are also a PMC member). + +All Apache committers are required to have a signed Contributor License +Agreement (CLA) on file with the Apache Software +Foundation. There is a Committer FAQ which provides more +details on the requirements for Committers + +A committer who makes a sustained contribution to the project may be invited to +become a member of the PMC. The form of contribution is not limited to code. It can +also include code review, helping out users on the mailing lists, documentation, +etc. + + +###Project Management Committee + +The Project Management Committee (PMC) for Apache Apex was created by a +resolution of the board of the Apache Software Foundation on ***TBD***. The +PMC is responsible to the board and the ASF for the management and oversight of the +Apache Apex codebase. The responsibilities of the PMC include: + +* Deciding what is distributed as products of the project. In particular all releases must be approved by the PMC +*	1	1	https://issues.apache.org/jira/browse/APEXCORE-321
20294	APEXCORE-319	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/incubator-apex-site/pull/19#discussion_r56011075 --- Diff: src/md/compatibility.md --- @@ -0,0 +1,112 @@ +#Apache Apex Compatibility + +##Purpose + +This document captures the compatibility goals of the Apache Apex project. The different types of compatibility between Apex releases that affect contributors, downstream projects, and end-users are enumerated. For each type of compatibility we: + +* describe the impact on downstream projects or end-users +* where applicable, call out the policy adopted when incompatible changes are permitted. + +Apache Apex follows [semantic versioning](http://semver.org/). Depending in the compatibility type, there may be different tools or mechanisms to ensure compatibility, for example by comparing artifacts during the build process. + +The type of change will inform the required target version number. Given a version number MAJOR.MINOR.PATCH, increment the: + +* MAJOR version when you make incompatible API changes, +* MINOR version when you add functionality in a backwards-compatible manner, and +* PATCH version when you make backwards-compatible bug fixes. + +Additional labels for pre-release and build metadata are available as extensions to the MAJOR.MINOR.PATCH format. + +The overall goal is to avoid backward incompatible changes and major release upgrades. Accordingly we attempt to release new features with minor versions that are incremental to the prior release and offer our users a frictionless upgrade path. When planning contributions, please consider compatibility and release road map upfront. Specifically, certain changes that conflict with the versioning may need to be documented in JIRA and deferred until a future major release. + +##Compatibility types + +###Java API + +Public API compatibility is required to ensure end-user programs and downstream projects continue to work without modification. + +The public API consists of: + +* apex-core: all interfaces and classes in `api` and `common` modules +* apex-malhar: all interfaces and classes in all modules except `demos`, `samples`, `benchmark` + +Interfaces and classes that are part of the public API and are annotated with [interface stability](https://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-common/InterfaceClassification.html) are treated according to the rules defined by the annotation. + +Policy + +Changes to the public API must follow semantic versioning. Public APIs must be deprecated for at least one major release prior to their removal in a major release. The japicmp plugin is used to enforce compatibility as part of the Travis pre-commit builds. + +###Semantic compatibility + +The behavior of APIs needs to remain consistent over versions, though changes for correctness may result in changes in behavior. Tests and javadocs specify the behavior. Over time, test suites should be expanded to verify compliance with the specification, effectively creating a formal specification for the subset of behaviors that can be easily tested. + +Policy + +The behavior of API may be changed to fix incorrect behavior, changes to be accompanied by tests coverage for the exact behavior. + +###REST API + +REST API compatibility corresponds to both the URLs and request/response content over the wire. REST APIs are specifically meant for stable use by clients across releases, even major releases. + +Policy + +The REST API is separately versioned. This is to allow for co-existence of old and new API should there be a need for backward incompatible changes in the future. + +###Command Line Interface (CLI) + +The CLI may be used either directly via the system shell or via shell scripts. Changing the path, removing or renaming command line options, the order of arguments, or the command return code and output break compatibility and may adversely affect users. + +Policy + +CLI commands are to be deprecated (warning when used) in a prior minor release before they are removed or incompatibly modified in a subsequent major release. --- End diff -- What are CLI commands not preserved for one major version in deprecated capacity like the APIs	1	1	https://issues.apache.org/jira/browse/APEXCORE-319

#	Requirement	Questions	Answered	Successfully Answered	URL
20295	APEXCORE-319	src/md/compatibility.md --- @@ -0,0 +1,112 @@ +#Apache Apex Compatibility + +##Purpose + +This document captures the compatibility goals of the Apache Apex project. The different types of compatibility between Apex releases that affect contributors, downstream projects, and end-users are enumerated. For each type of compatibility we: + +* describe the impact on downstream projects or end-users +* where applicable, call out the policy adopted when incompatible changes are permitted. + +Apache Apex follows [semantic versioning](http://semver.org/). Depending in the compatibility type, there may be different tools or mechanisms to ensure compatibility, for example by comparing artifacts during the build process. + +The type of change will inform the required target version number. Given a version number MAJOR.MINOR.PATCH, increment the: + +* MAJOR version when you make incompatible API changes, +* MINOR version when you add functionality in a backwards-compatible manner, and +* PATCH version when you make backwards-compatible bug fixes. + +Additional labels for pre-release and build metadata are available as extensions to the MAJOR.MINOR.PATCH format. + +The overall goal is to avoid backward incompatible changes and major release upgrades. Accordingly we attempt to release new features with minor versions that are incremental to the prior release and offer our users a frictionless upgrade path. When planning contributions, please consider compatibility and release road map upfront. Specifically, certain changes that conflict with the versioning may need to be documented in JIRA and deferred until a future major release. + +##Compatibility types + +###Java API + +Public API compatibility is required to ensure end-user programs and downstream projects continue to work without modification. + +The public API consists of: + +* apex-core: all interfaces and classes in `api` and `common` modules +* apex-malhar: all interfaces and classes in all modules except `demos`, `samples`, `benchmark` + +Interfaces and classes that are part of the public API and are annotated with [interface stability](https://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-common/InterfaceClassification.html) are treated according to the rules defined by the annotation. + +Policy + +Changes to the public API must follow semantic versioning. + +Public APIs must be deprecated for at least one major release prior to their removal in a major release. + +The japicmp plugin is used to enforce compatibility as part of the Travis pre-commit builds. + +###Semantic compatibility + +The behavior of APIs needs to remain consistent over versions, though changes for correctness may result in changes in behavior. Tests and javadocs specify the behavior. Over time, test suites should be expanded to verify compliance with the specification, effectively creating a formal specification for the subset of behaviors that can be easily tested. + +Policy + +The behavior of API may be changed to fix incorrect behavior, changes to be accompanied by tests coverage for the exact behavior. + +###REST API + +REST API compatibility corresponds to both the URLs and request/response content over the wire. REST APIs are specifically meant for stable use by clients across releases, even major releases. + +Policy + +The REST API is separately versioned. This is to allow for co-existence of old and new API should there be a need for backward incompatible changes in the future. + +###Command Line Interface (CLI) + +The CLI may be used either directly via the system shell or via shell scripts. Changing the path, removing or renaming command line options, the order of arguments, or the command return code and output break compatibility and may adversely affect users. + +Policy + +CLI commands are to be deprecated (warning when used) in a prior minor release before they are removed or incompatibly modified in a subsequent major release. + +###Configuration Files + +Configuration files are used for engine or application settings. Changes to keys and default values affect users. + +Policy + +Name, location, format, keys of configuration files should be deprecated in a prior minor release and can only be changed in major release. Best effort should be made to	1	1	https://issues.apache.org/jira/browse/APEXCORE-319
20296	APEXCORE-319	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-site/pull/19#discussion_r56069972 --- Diff: src/md/compatibility.md --- @@ -0,0 +1,112 @@ +#Apache Apex Compatibility + +##Purpose + +This document captures the compatibility goals of the Apache Apex project. The different types of compatibility between Apex releases that affect contributors, downstream projects, and end-users are enumerated. For each type of compatibility we: + +* describe the impact on downstream projects or end-users +* where applicable, call out the policy adopted when incompatible changes are permitted. + +Apache Apex follows [semantic versioning](http://semver.org/). Depending in the compatibility type, there may be different tools or mechanisms to ensure compatibility, for example by comparing artifacts during the build process. --- End diff -- Depending "on"?	1	1	https://issues.apache.org/jira/browse/APEXCORE-319

#	Requirement	Questions	Answered	Successfully Answered	URL
20297	APEXCORE-319	src/md/compatibility.md --- @@ -0,0 +1,112 @@ +#Apache Apex Compatibility + +##Purpose + +This document captures the compatibility goals of the Apache Apex project. The different types of compatibility between Apex releases that affect contributors, downstream projects, and end-users are enumerated. For each type of compatibility we: + +* describe the impact on downstream projects or end-users +* where applicable, call out the policy adopted when incompatible changes are permitted. + +Apache Apex follows [semantic versioning](http://semver.org/). Depending in the compatibility type, there may be different tools or mechanisms to ensure compatibility, for example by comparing artifacts during the build process. + +The type of change will inform the required target version number. Given a version number MAJOR.MINOR.PATCH, increment the: + +* MAJOR version when you make incompatible API changes, +* MINOR version when you add functionality in a backwards-compatible manner, and +* PATCH version when you make backwards-compatible bug fixes. + +Additional labels for pre-release and build metadata are available as extensions to the MAJOR.MINOR.PATCH format. + +The overall goal is to avoid backward incompatible changes and major release upgrades. Accordingly we attempt to release new features with minor versions that are incremental to the prior release and offer our users a frictionless upgrade path. When planning contributions, please consider compatibility and release road map upfront. Specifically, certain changes that conflict with the versioning may need to be documented in JIRA and deferred until a future major release. + +##Compatibility types + +###Java API + +Public API compatibility is required to ensure end-user programs and downstream projects continue to work without modification. + +The public API consists of: + +* apex-core: all interfaces and classes in `api` and `common` modules +* apex-malhar: all interfaces and classes in all modules except `demos`, `samples`, `benchmark` + +Interfaces and classes that are part of the public API and are annotated with [interface stability](https://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-common/InterfaceClassification.html) are treated according to the rules defined by the annotation. + +Policy + +Changes to the public API must follow semantic versioning. Public APIs must be deprecated for at least one major release prior to their removal in a major release. The japicmp plugin is used to enforce compatibility as part of the Travis pre-commit builds. + +###Semantic compatibility + +The behavior of APIs needs to remain consistent over versions, though changes for correctness may result in changes in behavior. Tests and javadocs specify the behavior. Over time, test suites should be expanded to verify compliance with the specification, effectively creating a formal specification for the subset of behaviors that can be easily tested. + +Policy + +The behavior of API may be changed to fix incorrect behavior, changes to be accompanied by tests coverage for the exact behavior. + +###REST API + +REST API compatibility corresponds to both the URLs and request/response content over the wire. REST APIs are specifically meant for stable use by clients across releases, even major releases. + +Policy + +The REST API is separately versioned. This is to allow for co-existence of old and new API should there be a need for backward incompatible changes in the future. + +###Command Line Interface (CLI) + +The CLI may be used either directly via the system shell or via shell scripts. Changing the path, removing or renaming command line options, the order of arguments, or the command return code and output break compatibility and may adversely affect users. + +Policy + +CLI commands are to be deprecated (warning when used) in a prior minor release before they are removed or incompatibly modified in a subsequent major release. + +###Configuration Files + +Configuration files are used for engine or application settings. Changes to keys and default values affect users. + +Policy + +Name, location, format, keys of configuration files should be deprecated in a prior minor release and can only be changed in major release. Best effort should be made to	1	1	https://issues.apache.org/jira/browse/APEXCORE-319
20298	APEXCORE-319	src/md/compatibility.md --- @@ -0,0 +1,112 @@ +#Apache Apex Compatibility + +##Purpose + +This document captures the compatibility goals of the Apache Apex project. The different types of compatibility between Apex releases that affect contributors, downstream projects, and end-users are enumerated. For each type of compatibility we: + +* describe the impact on downstream projects or end-users +* where applicable, call out the policy adopted when incompatible changes are permitted. + +Apache Apex follows [semantic versioning](http://semver.org/). Depending in the compatibility type, there may be different tools or mechanisms to ensure compatibility, for example by comparing artifacts during the build process. + +The type of change will inform the required target version number. Given a version number MAJOR.MINOR.PATCH, increment the: + +* MAJOR version when you make incompatible API changes, +* MINOR version when you add functionality in a backwards-compatible manner, and +* PATCH version when you make backwards-compatible bug fixes. + +Additional labels for pre-release and build metadata are available as extensions to the MAJOR.MINOR.PATCH format. + +The overall goal is to avoid backward incompatible changes and major release upgrades. Accordingly we attempt to release new features with minor versions that are incremental to the prior release and offer our users a frictionless upgrade path. When planning contributions, please consider compatibility and release road map upfront. Specifically, certain changes that conflict with the versioning may need to be documented in JIRA and deferred until a future major release. + +##Compatibility types + +###Java API + +Public API compatibility is required to ensure end-user programs and downstream projects continue to work without modification. + +The public API consists of: + +* apex-core: all interfaces and classes in `api` and `common` modules +* apex-malhar: all interfaces and classes in all modules except `demos`, `samples`, `benchmark` + +Interfaces and classes that are part of the public API and are annotated with [interface stability](https://hadoop.apache.org/docs/stable/hadoop-project-dist/hadoop-common/InterfaceClassification.html) are treated according to the rules defined by the annotation. + +Policy + +Changes to the public API must follow semantic versioning. Public APIs must be deprecated for at least one major release prior to their removal in a major release. The japicmp plugin is used to enforce compatibility as part of the Travis pre-commit builds. + +###Semantic compatibility + +The behavior of APIs needs to remain consistent over versions, though changes for correctness may result in changes in behavior. Tests and javadocs specify the behavior. Over time, test suites should be expanded to verify compliance with the specification, effectively creating a formal specification for the subset of behaviors that can be easily tested. + +Policy + +The behavior of API may be changed to fix incorrect behavior, changes to be accompanied by tests coverage for the exact behavior. + +###REST API + +REST API compatibility corresponds to both the URLs and request/response content over the wire. REST APIs are specifically meant for stable use by clients across releases, even major releases. + +Policy + +The REST API is separately versioned. This is to allow for co-existence of old and new API should there be a need for backward incompatible changes in the future. + +###Command Line Interface (CLI) + +The CLI may be used either directly via the system shell or via shell scripts. Changing the path, removing or renaming command line options, the order of arguments, or the command return code and output break compatibility and may adversely affect users. + +Policy + +CLI commands are to be deprecated (warning when used) in a prior minor release before they are removed or incompatibly modified in a subsequent major release. + +###Configuration Files + +Configuration files are used for engine or application settings. Changes to keys and default values affect users. + +Policy + +Name, location, format, keys of configuration files should be deprecated in a prior minor release and can only be changed in major release. Best effort should be made to	1	1	https://issues.apache.org/jira/browse/APEXCORE-319
20299	APEXCORE-318	Chris if we follow a similar model like hadoop of using a separate email list is it ok to ask infra for an email list like security@apex.apache.org (without using incubator) or should we wait till the last possible moment before becoming top level to request this. Secondly since the nature of the groups like these would be to not announce issues to the public until the vulnerabilities are fixed, is it right to assume that the membership to this group is selective and may not be every committer. Also how do groups like these track issues before the vulnerabilities are fixed, can JIRA still be used without making the information public for these till the fix.	1	1	https://issues.apache.org/jira/browse/APEXCORE-318
20300	APEXCORE-318	Chris, Received a message from INFRA "Please have someone with the proper karma fill out the mailing list request form: https://infra.apache.org/officers/mlreq/incubator" Who has the required karma? Thanks	1	1	https://issues.apache.org/jira/browse/APEXCORE-318

#	Requirement	Questions	Answered	Successfully Answered	URL
20301	APEXCORE-306	Would this be different from APEXCORE-279 or complementary.	1	0	https://issues.apache.org/jira/browse/APEXCORE-306
20302	APEXCORE-304	Github user DT-Priyanka commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/208#discussion_r49971245 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -156,6 +156,7 @@ private transient int nodeIndex = 0; // used for cycle validation private transient Stack<OperatorMeta> stack = new Stack<OperatorMeta>(); // used for cycle validation private transient Map<String, ArrayListMultimap<OutputPort<?>, InputPort<?>>> streamLinks = new HashMap<>(); + private transient Map<String, Boolean> extraResources = new LinkedHashMap<>(); --- End diff -- why map? Can we instead use a list/set of absolute path of resources?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20303	APEXCORE-304	Github user DT-Priyanka commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/208#discussion_r49976059 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -156,6 +156,7 @@ private transient int nodeIndex = 0; // used for cycle validation private transient Stack<OperatorMeta> stack = new Stack<OperatorMeta>(); // used for cycle validation private transient Map<String, ArrayListMultimap<OutputPort<?>, InputPort<?>>> streamLinks = new HashMap<>(); + private transient Map<String, Boolean> extraResources = new LinkedHashMap<>(); --- End diff -- What is the scenario where user will want to use different flag for resources, i.e. keep few and delete others?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20304	APEXCORE-304	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/208#discussion_r50059334 --- Diff: engine/src/main/java/com/datatorrent/stram/StramClient.java --- @@ -455,6 +462,13 @@ public void startApplication() throws YarnException, IOException } String libJarsCsv = copyFromLocal(fs, appPath, localJarFiles.toArray(new String[]{})); + // Now that the extra libs are copied, delete if it is required + for (Map.Entry<String, Boolean> jars : extraResources.entrySet()) { + if (jars.getValue()) { + FileUtils.deleteQuietly(new File(jars.getKey())); --- End diff -- Why quietly? If a user added a temporary jar and specified that it needs to be deleted, should not warning be at minimum logged if the file can not be deleted for any reason?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20305	APEXCORE-304	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/208#discussion_r50161978 --- Diff: engine/src/main/java/com/datatorrent/stram/StramLocalCluster.java --- @@ -308,6 +313,24 @@ public StramLocalCluster(LogicalPlan dag) throws IOException, ClassNotFoundExcep this.dnmgr = new StreamingContainerManager(dag); this.umbilical = new UmbilicalProtocolLocalImpl(); + Method addURL; + URLClassLoader loader; + try { + loader = (URLClassLoader)ClassLoader.getSystemClassLoader(); --- End diff -- What is the purpose of casting loader to URLCLassLoader? Can't this be coded as ``` ClassLoader loader; try { loader = ClassLoader.getSystemClassLoader(); addURL = loader.getClass().getDeclaredMethod("addURL", URL.class); ```	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20306	APEXCORE-304	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/208#discussion_r50163000 --- Diff: engine/src/main/java/com/datatorrent/stram/StramLocalCluster.java --- @@ -308,6 +313,24 @@ public StramLocalCluster(LogicalPlan dag) throws IOException, ClassNotFoundExcep this.dnmgr = new StreamingContainerManager(dag); this.umbilical = new UmbilicalProtocolLocalImpl(); + Method addURL; + URLClassLoader loader; + try { + loader = (URLClassLoader)ClassLoader.getSystemClassLoader(); + addURL = URLClassLoader.class.getDeclaredMethod("addURL", new Class[] { URL.class }); + addURL.setAccessible(true); + } catch (NoSuchMethodException e) { + throw new RuntimeException("No method named addURL found.", e); + } + + for (String jarPath : dag.getJarResources().keySet()) { + try { + addToClasspath(jarPath, loader, addURL); + } catch (InvocationTargetException | IllegalAccessException e) { --- End diff -- Is it necessary to catch exception inside the loop? Both InvocationTargetException and IllegalAccessException don't depend on the URL. In case StramLocalCluster delegates IOException handling to the consumer of the constructor, can other exception also be propagated?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20307	APEXCORE-304	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/208#discussion_r50308131 --- Diff: engine/src/test/java/com/datatorrent/stram/StramMiniClusterTest.java --- @@ -227,6 +242,7 @@ public void testSetupShutdown() throws Exception private LogicalPlan createDAG(LogicalPlanConfiguration lpc) throws Exception { LogicalPlan dag = new LogicalPlan(); + addExtraTestJar(dag); --- End diff -- Why this function is called for every createDAG? What functionality does it add? Should not this be just one of test cases?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20308	APEXCORE-304	Github user chinmaykolhatkar commented on the pull request: https://github.com/apache/incubator-apex-core/pull/208#issuecomment-174406481 All comments are addressed here. If there are no further comment, Can this pull request be merged? Thanks.	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20309	APEXCORE-304	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/208#discussion_r51064718 --- Diff: engine/src/main/java/com/datatorrent/stram/StramLocalCluster.java --- @@ -318,6 +333,21 @@ public StramLocalCluster(LogicalPlan dag) throws IOException, ClassNotFoundExcep } } + private void addExtraJarsToClasspath(Set<String> jars) + throws InvocationTargetException, IllegalAccessException, MalformedURLException, NoSuchMethodException + { + List<URL> jarUrls = new LinkedList<>(); + for (String jarPath : jars) { + File file = new File(jarPath); + URL url = file.toURI().toURL(); + jarUrls.add(url); + } + + ClassLoader prevCl = Thread.currentThread().getContextClassLoader(); + ClassLoader urlCl = URLClassLoader.newInstance(jarUrls.toArray(new URL[jarUrls.size()]), prevCl); + Thread.currentThread().setContextClassLoader(urlCl); --- End diff -- @tweise Is it expected that StramAppLauncher.runLocal() exposes different behavior compared to LocalMode.runApp()? I would expect StramAppLauchar.runLocal to use LocalMode.runApp() so functionality is the same when application is launched in application test and using dtcli local mode.	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20310	APEXCORE-304	Github user chinmaykolhatkar commented on the pull request: https://github.com/apache/incubator-apex-core/pull/216#issuecomment-180388829 Hello @vrozov @tweise Can this be merged if there are no further comments? Thanks.	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20311	APEXCORE-304	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/311#issuecomment-211509289 The ticket seems to suggest a broader change than just local mode. What am I missing?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304

#	Requirement	Questions	Answered	Successfully Answered	URL
20312	APEXCORE-304	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/311#discussion_r60113016 --- Diff: api/src/main/java/com/datatorrent/api/StreamingApplication.java --- @@ -47,6 +47,12 @@ String ENVIRONMENT = DT_PREFIX + "environment"; /** + * Use this config variable to to add any external library to classpath for the DAG. + * Value of this variable is same as StramAppLauncher.LIBJARS_CONF_KEY_NAME. + */ + String LIBJARS_CONF_KEY_NAME = "tmplibjars"; --- End diff -- Just looking at the string literal suggests something isn't right here. We configure applications through attributes, not random properties in the config file. Can you please check if there is a better way to express this dependency.	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20313	APEXCORE-304	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/311#discussion_r60265837 --- Diff: api/src/main/java/com/datatorrent/api/StreamingApplication.java --- @@ -47,6 +47,12 @@ String ENVIRONMENT = DT_PREFIX + "environment"; /** + * Use this config variable to to add any external library to classpath for the DAG. + * Value of this variable is same as StramAppLauncher.LIBJARS_CONF_KEY_NAME. + */ + String LIBJARS_CONF_KEY_NAME = "tmplibjars"; --- End diff -- Hmmm.. got your point and agree with you.. I find attribute as a better place than properties for such configuration. Here is my suggestion about exposing this as an attribute: 1. Have a attribute in DAGContext: Attribute<Map<String, Boolean>> APP_EXTERNAL_JARS = new Attribute<Map<, >(new Map2String<String, Boolean>(",", "=", new String2String<>(), new String2String<>())); 2. We pick up this attribute in following places: a. For cluster mode: StramAppLauncher.launchApp b. For local mode (apex cli) : StramAppLauncher.runLocal c. For local mode (ApplicationTest): LocalModeImpl.getController 3. For individual approaches, jar becomes available from that point onwards. Can you share though on above approach?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20314	APEXCORE-304	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/311#discussion_r60265869 --- Diff: api/src/main/java/com/datatorrent/api/StreamingApplication.java --- @@ -47,6 +47,12 @@ String ENVIRONMENT = DT_PREFIX + "environment"; /** + * Use this config variable to to add any external library to classpath for the DAG. + * Value of this variable is same as StramAppLauncher.LIBJARS_CONF_KEY_NAME. + */ + String LIBJARS_CONF_KEY_NAME = "tmplibjars"; --- End diff -- Hmmm.. got your point and agree with you.. I find attribute as a better place than properties for such configuration. Here is my suggestion about exposing this as an attribute: 1. Have a attribute in DAGContext: Attribute<Map<String, Boolean>> APP_EXTERNAL_JARS = new Attribute<Map<, >(new Map2String<String, Boolean>(",", "=", new String2String<>(), new String2String<>())); 2. We pick up this attribute in following places: a. For cluster mode: StramAppLauncher.launchApp b. For local mode (apex cli) : StramAppLauncher.runLocal c. For local mode (ApplicationTest): LocalModeImpl.getController 3. For individual approaches, jar becomes available from that point onwards. Can you share though on above approach?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20315	APEXCORE-304	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/311#discussion_r60498298 --- Diff: engine/src/test/java/com/datatorrent/stram/StramLocalClusterTest.java --- @@ -274,4 +286,103 @@ public WindowGenerator setupWindowGenerator() localCluster.shutdown(); } + @Test + public void testDynamicLoading() throws IOException, ClassNotFoundException + { + final String generatedJar = generatejar(); + File file = new File(generatedJar); + final Class<?> pojo = URLClassLoader.newInstance(new URL[]{file.toURI().toURL()}).loadClass("POJO"); + + StreamingApplication app = new StreamingApplication() + { + @Override + public void populateDAG(DAG dag, Configuration conf) + { + TestGeneratorInputOperator genNode = dag.addOperator("genNode", TestGeneratorInputOperator.class); + genNode.setMaxTuples(2); + + DynamicLoader dynamicLoader = dag.addOperator("DynamicLoader", new DynamicLoader()); + dynamicLoader.setClass("POJO"); + + dag.addStream("fromNode1", genNode.outport, dynamicLoader.in); + dag.setInputPortAttribute(dynamicLoader.in, Context.PortContext.TUPLE_CLASS, pojo); + conf.set(LIBJARS_CONF_KEY_NAME, generatedJar); + } + }; + + LocalMode.runApp(app, 10000); + FileUtils.forceDelete(new File("src/test/resources/dynamicJar/POJO.class")); --- End diff -- Why is the unit test (creating? and) deleting files in the source tree?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20316	APEXCORE-304	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/311#discussion_r61008036 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -2386,7 +2381,7 @@ public static void write(DAG dag, OutputStream os) throws IOException public static LogicalPlan read(InputStream is) throws IOException, ClassNotFoundException { - return (LogicalPlan)new ObjectInputStream(is).readObject(); + return (LogicalPlan)new ClassLoaderObjectInputStream(Thread.currentThread().getContextClassLoader(), is).readObject(); --- End diff -- Why is this required (or how was it working prior to this change)?	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20317	APEXCORE-304	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/311#discussion_r61535209 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -2386,7 +2381,7 @@ public static void write(DAG dag, OutputStream os) throws IOException public static LogicalPlan read(InputStream is) throws IOException, ClassNotFoundException { - return (LogicalPlan)new ObjectInputStream(is).readObject(); + return (LogicalPlan)new ClassLoaderObjectInputStream(Thread.currentThread().getContextClassLoader(), is).readObject(); --- End diff -- This is required for Local Mode case. Lets say a jar path is set to LIBRARY_JARS and this jar contains a POJO that needs to be set to TUPLE_CLASS attr. Now, one can do that in populateDAG method, that would cause a problem in LocalMode as follows: 1. populateDAG sets TUPLE_CLASS attr with new POJO class - Works. 2. Local mode serializes the LogicalPlan - Works. 3. Local mode deserialized the LogicalPlan - Fails. 3rd step fails becuase The ObjectInputStream deserializer uses System ClassLoader which does that have that POJO. ClassLoaderObjectInputStream is an overridden version of ObjectInputStream which allows to set the custom classloader. By this time of callstack, Thread.currentThread().getClassLoader() will have a classloader which is super of all. As previously, there was no functionality of LIBRARY_JARS, I guess this was not tested and so did not surface as an issue. For cluster mode, there is no difference. i.e. There is only 1 classloader i.e. System.Classloader and that has new jars as well. i.e. For cluster mode, Thread.currentThread().getClassLoader() = System classloader i.e. ClassLoaderObjectInputStream will behave as if ObjectInputStream.	1	1	https://issues.apache.org/jira/browse/APEXCORE-304
20318	APEXCORE-294	What is the advantage of end stream tuple over ShutdownException?	0	0	https://issues.apache.org/jira/browse/APEXCORE-294
20319	APEXCORE-294	Is it that the shutdownException is controlled more by STRAM and the control tuple will flow naturally in the DAG? As I understand, when a ShutdownException is thrown by an operator, it propagates to its downstream operators via STRAM as shutdown requests for each container.	0	0	https://issues.apache.org/jira/browse/APEXCORE-294
20320	APEXCORE-294	What will be the trigger for this? Could you please provide some details?	0	0	https://issues.apache.org/jira/browse/APEXCORE-294

#	Requirement	Questions	Answered	Successfully Answered	URL
20321	APEXCORE-294	would sending END_STREAM tuple from all input operators partition will work. Is there any other complication with this approach? If this approach is fine I would start working on it. The flow would be - Master receives shutdown command from cli/rest api - Master set shutdown on each container and send it as part of heartbeat response to all containers - On receiving shutdown command StreamingContainer would mark alive false for all InputNode causing them to terminate with sending END_STREAM tuple downstream. same could be achieve using operatorRequest for all input operator partition and this request will be handled by StreamingContainer to shutdown the input operator thread.	1	1	https://issues.apache.org/jira/browse/APEXCORE-294
20322	APEXCORE-293	Documentation source code will be hosted in apex-core repository on ASF. Generated documentation can be hosted on apex.incubator.apache.org, but it involves more steps to process, generate, and deal with document deployments, when compared with solutions like readthedocs.org, readme.io, etc. I notice several other Apache projects host their docs outside ASF, including Apache Ignite (https://apacheignite.readme.io/docs), Apache Libcloud (https://libcloud.readthedocs.org/en/latest/), etc. Does ASF mandate that documentation be hosted on apache.org? And if ASF infrastructure is required, is there a simple solution similar to readthedocs for markdown docs, which supports multiple release versions, etc?	1	1	https://issues.apache.org/jira/browse/APEXCORE-293
20323	APEXCORE-293	Is the intention to completely remove any references to external documentation, or to add rel="nofollow" to external links?	1	1	https://issues.apache.org/jira/browse/APEXCORE-293
20324	APEXCORE-293	[~sasha] We should review the need for each of these links. It is not clear, for example, why we would point to RTS docs for configuration. [~ptgoetz] In those cases where external links are appropriate, is it a matter of adding rel="nofollow"?	1	1	https://issues.apache.org/jira/browse/APEXCORE-293
20325	APEXCORE-291	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/190#discussion_r48397509 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/MetricAggregatorMeta.java --- @@ -0,0 +1,82 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram.plan.logical; + +import java.io.Serializable; +import java.util.Collection; +import java.util.Map; + +import javax.validation.constraints.NotNull; + +import com.google.common.base.Preconditions; + +import com.datatorrent.api.AutoMetric; + +public final class MetricAggregatorMeta implements Serializable --- End diff -- Can you add javadoc which describes the class.	1	1	https://issues.apache.org/jira/browse/APEXCORE-291
20326	APEXCORE-283	GitHub user ashishtadose opened a pull request: https://github.com/apache/incubator-apex-core/pull/184 APEXCORE-283 - checkpointing in distributed in-memory store @tushargosavi Can you please review the changes. You can merge this pull request into a Git repository by running: $ git pull https://github.com/ashishtadose/incubator-apex-core APEXCORE-283.checkpoint Alternatively you can review and apply these changes as the patch at: https://github.com/apache/incubator-apex-core/pull/184.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #184 ---- commit 1988d460e921dc82249d60c0bbe5a1f1b5913910 Author: Ashish <ashish@ampool.io> Date: 2015-12-09T14:45:27Z APEXCORE-283 #comment Interface for storage agent to retrieve application id commit f23c6814bf2964dc46d50f60ad93e4faa94b3880 Author: Ashish <ashish@ampool.io> Date: 2015-12-09T14:48:44Z APEXCORE-283 #comment Abstract implementation of KeyValue storage agent which can be configured with implementation of KeyValue store for checkpointing. commit f092f046810acfaeb00955ffcde6a3e6c27cff4d Author: Ashish <ashish@ampool.io> Date: 2015-12-09T14:53:02Z APEXCORE-283 #comment Concrete implementation of In memory storage agent for Apache Geode commit ab110cb0e77bfd50991fd16abc01fada7201e781 Author: Ashish <ashish@ampool.io> Date: 2015-12-09T15:11:00Z APEXCORE-283 #comment Updated AppIdAwareStorage agent changes in StramLocalCluster commit aae555660022051a367f7be822833098f8dcc8fb Author: Ashish <ashish@ampool.io> Date: 2015-12-09T15:14:10Z APEXCORE-283 #comment Excluded Geode storage related test cases as it requires some manual setup commit 142f0212eccda92dca90569f79d9c8a5a9874439 Author: Ashish <ashish@ampool.io> Date: 2015-12-09T15:15:42Z APEXCORE-283 #comment fixed checkstyle fixes causing build failures ----	1	1	https://issues.apache.org/jira/browse/APEXCORE-283
20327	APEXCORE-283	Github user tushargosavi commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/184#discussion_r47122916 --- Diff: api/src/main/java/com/datatorrent/api/AppIdAwareStorageAgent.java --- @@ -0,0 +1,36 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.api; + +/** + * Interface to pass application id to storage agent + * + * @since 3.2.0 + */ +public interface AppIdAwareStorageAgent extends StorageAgent +{ + /** + * Save the application id so it can be used by StorageAgent for checkpointing + * + * @param applicationId + */ + public void saveAppId(String applicationId); --- End diff -- Instead of just applicationId, can you pass Attributes. The DAG attributes has following members which may be useful for a storage agent. - Application Name - Application Id - Application Path	1	1	https://issues.apache.org/jira/browse/APEXCORE-283
20328	APEXCORE-283	Github user tushargosavi commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/184#discussion_r47125023 --- Diff: common/src/main/java/com/datatorrent/common/util/FSStorageAgent.java --- @@ -71,12 +71,10 @@ public FSStorageAgent(String path, Configuration conf) -- - End diff -- This file (FSStorageAgent.java) only contains formatting changes. Can you revert it back.	1	1	https://issues.apache.org/jira/browse/APEXCORE-283

#	Requirement	Questions	Answered	Successfully Answered	URL
20329	APEXCORE-283	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/184#discussion_r47592052 --- Diff: api/src/main/java/com/datatorrent/api/ApplicationAwareStorageAgent.java --- @@ -0,0 +1,38 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.api; +import com.datatorrent.api.Attribute.AttributeMap; +/** + * Interface to pass application attributes to storage agent + * + * + */ +public interface ApplicationAwareStorageAgent extends StorageAgent { + /** + * Passes attributes of application to storage agent + * + * + * @param map attributes of application + */ + public void saveAppAttributes(AttributeMap map); --- End diff -- Rename to setApplicationAttributes ?	1	1	https://issues.apache.org/jira/browse/APEXCORE-283
20330	APEXCORE-283	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/184#discussion_r47592207 --- Diff: api/src/main/java/com/datatorrent/api/ApplicationAwareStorageAgent.java --- @@ -0,0 +1,38 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.api; +import com.datatorrent.api.Attribute.AttributeMap; +/** + * Interface to pass application attributes to storage agent + * + * + */ +public interface ApplicationAwareStorageAgent extends StorageAgent { --- End diff -- Rename to StorageAgent.ApplicationAttributesAware ?	1	1	https://issues.apache.org/jira/browse/APEXCORE-283
20331	APEXCORE-276	Github user chandnisingh commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/193#discussion_r48619961 --- Diff: api/src/main/java/com/datatorrent/api/Context.java --- @@ -335,10 +335,16 @@ Attribute<String> APPLICATION_DATA_LINK = new Attribute<String>(new String2String()); /** * Transport to push the stats and the metrics, "builtin:{topic}" if STRAM should push the data directly - * using websocket with the given topic. + * using websocket with the given topic. + * If using a custom transport, please use Object2String codec to specify the transport object. The object must + * be from a class that implements the AutoMetric.Transport interface */ Attribute<String> METRICS_TRANSPORT = new Attribute<String>(new String2String()); --- End diff -- So what it means here is that we are asking the user to convert the object to string using Object2String codec before setting this attribute? Why don't we deprecate this and add another attribute which is of type Attribute<AutoMetric.Transport> which is consistent with rest of the complex attributes?	1	1	https://issues.apache.org/jira/browse/APEXCORE-276
20332	APEXCORE-276	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/193#issuecomment-168046724 This is why is is important to not assume things as stable right of the bat. Is this attribute in its present form used anywhere? I don't think an end user would even be able to understand what to do with it? If so, we should consider it a bug and fix it.	1	1	https://issues.apache.org/jira/browse/APEXCORE-276
20333	APEXCORE-276	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/193#discussion_r48622829 --- Diff: api/src/main/java/com/datatorrent/api/Context.java --- @@ -335,10 +335,16 @@ Attribute<String> APPLICATION_DATA_LINK = new Attribute<String>(new String2String()); /** * Transport to push the stats and the metrics, "builtin:{topic}" if STRAM should push the data directly - * using websocket with the given topic. + * using websocket with the given topic. + * If using a custom transport, please use Object2String codec to specify the transport object. The object must + * be from a class that implements the AutoMetric.Transport interface */ Attribute<String> METRICS_TRANSPORT = new Attribute<String>(new String2String()); /** + * The interval in milliseconds to resend metrics schema. Default value is 10,000 (10 seconds) + */ + Attribute<Long> METRICS_SCHEMA_RESEND_INTERVAL = new Attribute<Long>(10000L); --- End diff -- Shouldn't the resend interval be part of the transport as the need to do that is driven by the properties of the transport, i.e. need to resend if messages could be lost?	1	1	https://issues.apache.org/jira/browse/APEXCORE-276
20334	APEXCORE-272	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/191#discussion_r48492262 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1359,7 +1409,8 @@ private void addDAGToCurrentDAG(ModuleMeta moduleMeta) String name; for (OperatorMeta operatorMeta : subDag.getAllOperators()) { name = subDAGName + MODULE_NAMESPACE_SEPARATOR + operatorMeta.getName(); - this.addOperator(name, operatorMeta.getOperator()); + Operator op = this.addOperator(name, operatorMeta.getOperator()); --- End diff -- Can LogicalPlan class provide a new method that allows to add Operator with existing OperatorMeta? Something like ``` private <T extends Operator> T addOperator(final String name, final T operator, final OperatorMeta meta) { } ```	1	1	https://issues.apache.org/jira/browse/APEXCORE-272
20335	APEXCORE-272	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/191#discussion_r48571279 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1359,7 +1409,8 @@ private void addDAGToCurrentDAG(ModuleMeta moduleMeta) String name; for (OperatorMeta operatorMeta : subDag.getAllOperators()) { name = subDAGName + MODULE_NAMESPACE_SEPARATOR + operatorMeta.getName(); - this.addOperator(name, operatorMeta.getOperator()); + Operator op = this.addOperator(name, operatorMeta.getOperator()); + this.getMeta(op).copyFrom(operatorMeta); OperatorMeta operatorMetaNew = this.getOperatorMeta(name); --- End diff -- Is operatormeta returned by `this.getMeta(op)` different from `this.getOperatorMeta(name)`? Can you not use `operatorMetaNew.copyFrom(operatorMeta)`?	1	1	https://issues.apache.org/jira/browse/APEXCORE-272

#	Requirement	Questions	Answered	Successfully Answered	URL
20336	APEXCORE-272	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/191#discussion_r48589840 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -907,6 +912,51 @@ else if (field.getType() == float.class || field.getType() == Float.class || getValue(OperatorContext.METRICS_DIMENSIONS_SCHEME)); } + /** + * Copy attribute from source attributeMap to destination attributeMap. + * + * @param dest destination attribute map. + * @param source source attribute map. + */ + private void copyAttributes(AttributeMap dest, AttributeMap source) --- End diff -- Shouldn't this be putAll on AttributeMap?	1	1	https://issues.apache.org/jira/browse/APEXCORE-272
20337	APEXCORE-272	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/191#discussion_r48646622 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1359,7 +1409,8 @@ private void addDAGToCurrentDAG(ModuleMeta moduleMeta) String name; for (OperatorMeta operatorMeta : subDag.getAllOperators()) { name = subDAGName + MODULE_NAMESPACE_SEPARATOR + operatorMeta.getName(); - this.addOperator(name, operatorMeta.getOperator()); + Operator op = this.addOperator(name, operatorMeta.getOperator()); --- End diff -- Are we keeping module DAG after adding it's operators to the final DAG? If not, there is no meta object sharing between DAGs. Introducing copy constructor will allow to share some of final collection references. As operator reference is shared between meta objects, should other references be also shared?	1	1	https://issues.apache.org/jira/browse/APEXCORE-272
20338	APEXCORE-272	Github user tushargosavi commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/191#discussion_r48647943 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -907,6 +912,51 @@ else if (field.getType() == float.class || field.getType() == Float.class || getValue(OperatorContext.METRICS_DIMENSIONS_SCHEME)); } + /** + * Copy attribute from source attributeMap to destination attributeMap. + * + * @param dest destination attribute map. + * @param source source attribute map. + */ + private void copyAttributes(AttributeMap dest, AttributeMap source) --- End diff -- AttributeMap does not have pullAll method. Do you want me to add this in AttributeMap?	1	1	https://issues.apache.org/jira/browse/APEXCORE-272
20339	APEXCORE-268	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/199#discussion_r49207712 --- Diff: common/src/main/java/com/datatorrent/common/metric/MetricsAggregator.java --- @@ -129,13 +128,13 @@ public void addAggregators(@NotNull String metric, @NotNull SingleMetricAggregat * be used for the result of aggregators[i]. */ public void addAggregators(@NotNull String metric, @NotNull SingleMetricAggregator[] aggregators, - @NotNull String[] logicalMetricNames) + @NotNull String[] logicalMetricNames) { Preconditions.checkNotNull(metric, "metric"); Preconditions.checkNotNull(aggregators, "aggregators"); Preconditions.checkNotNull(logicalMetricNames, "logicalMetricNames"); Preconditions.checkArgument(aggregators.length == logicalMetricNames.length, "different length aggregators and" + - " logical names"); + " logical names"); --- End diff -- Avoid split on the string literal. Can the previous line be split after the comma?	1	1	https://issues.apache.org/jira/browse/APEXCORE-268
20340	APEXCORE-268	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/199#discussion_r49236401 --- Diff: common/pom.xml --- @@ -56,7 +56,7 @@ <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-checkstyle-plugin</artifactId> <configuration> - <maxAllowedViolations>114</maxAllowedViolations> + <logViolationsToConsole>true</logViolationsToConsole> --- End diff -- Which mvn goal do you use? I only see errors displayed when I explicitly use checkstyle:check. For the "verify" goal I don't see individual errors printed, only the total number of errors.	1	1	https://issues.apache.org/jira/browse/APEXCORE-268
20341	APEXCORE-268	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/199#issuecomment-170126536 Can it be configured to only print errors?	1	1	https://issues.apache.org/jira/browse/APEXCORE-268
20342	APEXCORE-268	Github user vrozov commented on the pull request: https://github.com/apache/incubator-apex-core/pull/199#issuecomment-170134683 What is the reason to suppress warnings? Do we want them to be fixed over time? If not, should warning rule be disabled completely?	1	1	https://issues.apache.org/jira/browse/APEXCORE-268
20343	APEXCORE-268	Github user vrozov commented on the pull request: https://github.com/apache/incubator-apex-core/pull/199#issuecomment-170188868 @chandnisingh I don't think it is the question of importance. It is the question of convenience for developers. If compiler warnings don't clutter build log, checkstyle warnings also should not clutter. Will suppression of LineLength warnings be a workable solution? In this case you may add ``` <suppress checks="LineLength" files="" /> ``` to checkstyle-suppresions.xml	1	1	https://issues.apache.org/jira/browse/APEXCORE-268
20344	APEXCORE-268	Github user chandnisingh commented on the pull request: https://github.com/apache/incubator-apex-core/pull/199#issuecomment-170189818 There should be a separate discussion on that since now it is a debate. I don't think it is a problem to run just another command to see the exact violation. Don't the users already do that for license, etc?	1	1	https://issues.apache.org/jira/browse/APEXCORE-268
20345	APEXCORE-268	Github user chandnisingh commented on the pull request: https://github.com/apache/incubator-apex-core/pull/199#issuecomment-170200390 I don't agree with it. A lot of folks in the community wanted fixed line length which is why I added it as a warning. That pull request was merged without any objections. So why suppress the warnings on an unrelated change. With this change it's not that I am increasing the steps. They are same as what they were. This is a separate topic and the whole community should participate in making that decision	1	1	https://issues.apache.org/jira/browse/APEXCORE-268
20346	APEXCORE-259	Did you mean Buffer Server resides in the same node?	1	0	https://issues.apache.org/jira/browse/APEXCORE-259
20347	APEXCORE-254	why is this needed?	1	0	https://issues.apache.org/jira/browse/APEXCORE-254
20348	APEXCORE-252	Would it be better to keep a layer of abstraction with CircularBuffer proxying SpscArrayQueue. If there is a different queue implementation in future all code using it doesn't have to change.	1	1	https://issues.apache.org/jira/browse/APEXCORE-252
20349	APEXCORE-247	Hi Sandesh, do you have a status update on this item?	0	0	https://issues.apache.org/jira/browse/APEXCORE-247
20350	APEXCORE-246	Was there a discussion about this new feature?	0	0	https://issues.apache.org/jira/browse/APEXCORE-246
20351	APEXCORE-225	Why is this needed ?	1	1	https://issues.apache.org/jira/browse/APEXCORE-225
20352	APEXCORE-225	What was the reason for not providing it in the input port.	0	0	https://issues.apache.org/jira/browse/APEXCORE-225
20353	APEXCORE-222	Why is this improvement needed?	0	0	https://issues.apache.org/jira/browse/APEXCORE-222

#	Requirement	Questions	Answered	Successfully Answered	URL
20354	APEXCORE-222	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-core/pull/350#discussion_r67060191 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/StreamingContainer.java --- @@ -769,7 +769,11 @@ public void processHeartbeatResponse(ContainerHeartbeatResponse rsp) } if (rsp.committedWindowId != lastCommittedWindowId) { + lastCommittedWindowId = rsp.committedWindowId; + + bufferServer.purge(lastCommittedWindowId); --- End diff -- Shouldn't this be a request that is sent to bufferserver from the streaming container otherwise this introduces a dependency which we didn't rely on before that bufferserver is in the same process.	1	1	https://issues.apache.org/jira/browse/APEXCORE-222
20355	APEXCORE-222	Shouldn't this be an option because bufferserver can be shared between multiple containers and can also exist outside the streaming container.	1	0	https://issues.apache.org/jira/browse/APEXCORE-222
20356	APEXCORE-222	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-core/pull/350#discussion_r68319480 --- Diff: engine/src/main/java/com/datatorrent/stram/StramLocalCluster.java --- @@ -72,7 +72,7 @@ protected final StreamingContainerManager dnmgr; private final UmbilicalProtocolLocalImpl umbilical; private InetSocketAddress bufferServerAddress; - private boolean perContainerBufferServer; + private boolean perContainerBufferServer = true; --- End diff -- Why is it necessary to change the default value for perContainerBufferServer?	1	1	https://issues.apache.org/jira/browse/APEXCORE-222
20357	APEXCORE-222	Github user tweise commented on a diff in the pull request: https://github.com/apache/apex-core/pull/350#discussion_r68493398 --- Diff: bufferserver/src/main/java/com/datatorrent/bufferserver/internal/DataList.java --- @@ -176,21 +176,20 @@ public void reset() numberOfInMemBlockPermits.set(MAX_COUNT_OF_INMEM_BLOCKS - 1); } - public void purge(final int baseSeconds, final int windowId) --- End diff -- Are these changes necessary as part of this PR?	1	1	https://issues.apache.org/jira/browse/APEXCORE-222
20358	APEXCORE-222	Github user tweise commented on a diff in the pull request: https://github.com/apache/apex-core/pull/350#discussion_r68496167 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/StreamingContainer.java --- @@ -769,7 +769,11 @@ public void processHeartbeatResponse(ContainerHeartbeatResponse rsp) } if (rsp.committedWindowId != lastCommittedWindowId) { + lastCommittedWindowId = rsp.committedWindowId; + + bufferServer.purge(lastCommittedWindowId); --- End diff -- Right, since that window was already materialized in the checkpointed state it shouldn't be needed. Have you checked that by looking where the subscriber starts consuming from?	1	1	https://issues.apache.org/jira/browse/APEXCORE-222
20359	APEXCORE-222	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/apex-core/pull/350#discussion_r71005395 --- Diff: bufferserver/src/main/java/com/datatorrent/bufferserver/internal/DataList.java --- @@ -176,21 +176,20 @@ public void reset() numberOfInMemBlockPermits.set(MAX_COUNT_OF_INMEM_BLOCKS - 1); } - public void purge(final int baseSeconds, final int windowId) + public void purge(long windowId) --- End diff -- should windowId be made final?	1	1	https://issues.apache.org/jira/browse/APEXCORE-222
20360	APEXCORE-215	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/incubator-apex-site/pull/11#discussion_r51406961 --- Diff: src/md/verification.md --- @@ -0,0 +1,82 @@ +# How to verify Apache Apex (incubating) release candidate builds + +## Prerequisites +1. *gpg* program should be installed on your system. +2. Download the KEYS file and import it. This is one time activity. +```bash +wget https://dist.apache.org/repos/dist/release/incubator/apex/KEYS +gpg --import KEYS +``` +3. You can also created your own key which is required if you would like to sign the build. This step is optional. +```bash +gpg --gen-key +``` +Please provide 4096 as keysize while generating the key. + +## File integrity check +Download all files present in staging directory of the RC build. Staging directory link is shared in VOTE thread of the release candidate. +```bash +wget -r -np -nd <staging-area-link>/ +``` +Note the link should end with "/". + +Define the apex release candidate variable. We will set it up *apex-3.3.0-incubating* as an example. +```bash +APEX_RELEASE_CANDIDATE=apex-3.3.0-incubating +``` + +Verify integrity of tar.gz file: +```bash +gpg --verify $APEX_RELEASE_CANDIDATE-source-release.tar.gz.asc --- End diff -- http://www.apache.org/info/verification.html mentions additional step for validating authenticity of a key. Should we add this step?	1	1	https://issues.apache.org/jira/browse/APEXCORE-215
20361	APEXCORE-215	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-site/pull/11#discussion_r51590627 --- Diff: src/md/verification.md --- @@ -0,0 +1,82 @@ +# How to verify Apache Apex (incubating) release candidate builds + +## Prerequisites +1. *gpg* program should be installed on your system. +2. Download the KEYS file and import it. This is one time activity. +```bash +wget https://dist.apache.org/repos/dist/release/incubator/apex/KEYS +gpg --import KEYS +``` +3. You can also created your own key which is required if you would like to sign the build. This step is optional. +```bash +gpg --gen-key +``` +Please provide 4096 as keysize while generating the key. + +## File integrity check +Download all files present in staging directory of the RC build. Staging directory link is shared in VOTE thread of the release candidate. +```bash +wget -r -np -nd <staging-area-link>/ +``` +Note the link should end with "/". + +Define the apex release candidate variable. We will set it up *apex-3.3.0-incubating* as an example. +```bash +APEX_RELEASE_CANDIDATE=apex-3.3.0-incubating +``` + +Verify integrity of tar.gz file: +```bash +gpg --verify $APEX_RELEASE_CANDIDATE-source-release.tar.gz.asc +md5sum --check $APEX_RELEASE_CANDIDATE-source-release.tar.gz.md5 +sha512sum --check $APEX_RELEASE_CANDIDATE-source-release.tar.gz.sha +``` + +Verify integrity of .zip file: +```bash +gpg --verify $APEX_RELEASE_CANDIDATE-source-release.zip.asc +md5sum --check $APEX_RELEASE_CANDIDATE-source-release.zip.md5 +sha512sum --check $APEX_RELEASE_CANDIDATE-source-release.zip.sha +``` + +## Source code verification +You can extract source either using .tar.gz file or .zip file. + +### Using .tar.gz source +Extract source using .tar.gz: +```bash +tar -zxvf $APEX_RELEASE_CANDIDATE-source-release.tar.gz +``` + +### Using .tar.gz source +```bash +unzip $APEX_RELEASE_CANDIDATE-source-release.zip +``` + +Any of the two commands above will create a directory named after *apex release candidate*. + +### Check for compilation, license headers, etc. +The last step is optional and needs prerequisite 3 given above. +```bash +cd $APEX_RELEASE_CANDIDATE +mvn clean package +mvn apache-rat:check verify -Dlicense.skip=false -Pall-modules -DskipTests --- End diff -- all-module profile is not defined for apex-core. Why is it necessary to run 3 distinct mvn?	1	1	https://issues.apache.org/jira/browse/APEXCORE-215

#	Requirement	Questions	Answered	Successfully Answered	URL
20362	APEXCORE-215	Github user vrozov commented on a diff in the pull request: https://github.com/apache/incubator-apex-site/pull/11#discussion_r53062507 --- Diff: src/md/verification.md --- @@ -0,0 +1,88 @@ +# How to verify Apache Apex (incubating) release candidate builds + +## Prerequisites +1. *gpg* program should be installed on your system. +2. Download the KEYS file and import it. This is one time activity. +```bash +wget https://dist.apache.org/repos/dist/release/incubator/apex/KEYS +gpg --import KEYS +``` +3. You can also created your own key which is required if you would like to sign the build. This step is optional. +```bash +gpg --gen-key +``` +Please provide 4096 as keysize while generating the key. + +## File integrity check +Download all files present in staging directory of the RC build. Staging directory link is shared in VOTE thread of the release candidate. +```bash +wget -r -np -nd <staging-area-link>/ +``` +Note the link should end with "/". +Define the apex release candidate variable. We will set it up *apex-3.3.0-incubating* as an example. +```bash +APEX_RELEASE_CANDIDATE=apex-3.3.0-incubating +``` + +Verify integrity of tar.gz file: +```bash +gpg --verify $APEX_RELEASE_CANDIDATE-source-release.tar.gz.asc +md5sum --check $APEX_RELEASE_CANDIDATE-source-release.tar.gz.md5 +sha512sum --check $APEX_RELEASE_CANDIDATE-source-release.tar.gz.sha +``` + +Verify integrity of .zip file: +```bash +gpg --verify $APEX_RELEASE_CANDIDATE-source-release.zip.asc +md5sum --check $APEX_RELEASE_CANDIDATE-source-release.zip.md5 +sha512sum --check $APEX_RELEASE_CANDIDATE-source-release.zip.sha +``` + +## Source code verification +You can extract source either using .tar.gz file or .zip file. + +### Using .tar.gz source +Extract source using .tar.gz: +```bash +tar -zxvf $APEX_RELEASE_CANDIDATE-source-release.tar.gz +``` + +### Using .zip source +```bash +unzip $APEX_RELEASE_CANDIDATE-source-release.zip +``` + +Any of the two commands above will create a directory named after *apex release candidate*. + +### Check for compilation, license headers, etc. +The last step is optional and needs prerequisite 3 given above. +```bash +cd $APEX_RELEASE_CANDIDATE +mvn clean apache-rat:check verify -Dlicense.skip=false -Pall-modules verify -Papache-release package +``` +--- End diff -- Is it necessary to have multiple verify and package goal? Will it be sufficient to have only one of them?	1	1	https://issues.apache.org/jira/browse/APEXCORE-215
20363	APEXCORE-215	@tweise and @vrozov : Could you please comment on Chinmay's comment on compiling Malhar with new Apex to check backward compatibility?	1	0	https://issues.apache.org/jira/browse/APEXCORE-215
20364	APEXCORE-215	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-site/pull/11#discussion_r53426395 --- Diff: src/md/verification.md --- @@ -0,0 +1,88 @@ +# How to verify Apache Apex (incubating) release candidate builds + +## Prerequisites +1. *gpg* program should be installed on your system. +2. Download the KEYS file and import it. This is one time activity. +```bash +wget https://dist.apache.org/repos/dist/release/incubator/apex/KEYS +gpg --import KEYS +``` +3. You can also created your own key which is required if you would like to sign the build. This step is optional. +```bash +gpg --gen-key +``` +Please provide 4096 as keysize while generating the key. + +## File integrity check +Download all files present in staging directory of the RC build. Staging directory link is shared in VOTE thread of the release candidate. +```bash +wget -r -np -nd <staging-area-link>/ +``` +Note the link should end with "/". +Define the apex release candidate variable. We will set it up *apex-3.3.0-incubating* as an example. +```bash +APEX_RELEASE_CANDIDATE=apex-3.3.0-incubating +``` + +Verify integrity of tar.gz file: +```bash +gpg --verify $APEX_RELEASE_CANDIDATE-source-release.tar.gz.asc +md5sum --check $APEX_RELEASE_CANDIDATE-source-release.tar.gz.md5 +sha512sum --check $APEX_RELEASE_CANDIDATE-source-release.tar.gz.sha +``` + +Verify integrity of .zip file: +```bash +gpg --verify $APEX_RELEASE_CANDIDATE-source-release.zip.asc +md5sum --check $APEX_RELEASE_CANDIDATE-source-release.zip.md5 +sha512sum --check $APEX_RELEASE_CANDIDATE-source-release.zip.sha +``` + +## Source code verification +You can extract source either using .tar.gz file or .zip file. + +### Using .tar.gz source +Extract source using .tar.gz: +```bash +tar -zxvf $APEX_RELEASE_CANDIDATE-source-release.tar.gz +``` + +### Using .zip source +```bash +unzip $APEX_RELEASE_CANDIDATE-source-release.zip +``` + +Any of the two commands above will create a directory named after *apex release candidate*. + +### Check for compilation, license headers, etc. +The last step is optional and needs prerequisite 3 given above. +```bash +cd $APEX_RELEASE_CANDIDATE +mvn clean apache-rat:check verify -Dlicense.skip=false -Pall-modules verify -Papache-release package +``` + +In case of Apex release, you can skip -Papache-release option as there are no modules defined in Apex as of now. + +## Launch demos +Launch few demos to make sure everything is working fine using dtcli utitlity. The dtcli script is present at <apex-core-folder>/engine/src/main/scripts/dtcli . If apex-core is being verified, then simply use engine/src/main/scripts/dtcli to launch the script. Otherwise, prefix it with apex-core folder path. +```bash +engine/src/main/scripts/dtcli +``` +You will get dtcli prompt, where demos can be launched. +```bash +dt> launch <demo-apk-file> +``` +--- End diff -- Do you mean launch existing version of demo with new Apex version? If so, then yes. In general, these are recommendations what folks can do the verify a release and not everyone will run through all of them.	1	1	https://issues.apache.org/jira/browse/APEXCORE-215
20365	APEXCORE-202	Shouldn't this be in Malhar?	0	0	https://issues.apache.org/jira/browse/APEXCORE-202
20366	APEXCORE-199	How about if it is compatible with the CLI version but incompatible with the RTS installation ?	0	0	https://issues.apache.org/jira/browse/APEXCORE-199
20367	APEXCORE-199	CLI does not know anything about RTS. But can you give an example of such situation? What kind of app would be compatible with the platform but not with the distribution that includes it?	0	0	https://issues.apache.org/jira/browse/APEXCORE-199
20368	APEXCORE-199	Why was this reopened ?	1	0	https://issues.apache.org/jira/browse/APEXCORE-199
20369	APEXCORE-168	Is $\{package} replaced during project creation? @davidyan The log4j properties discussed here are used for testing in the IDE only. They are not part of the .apa and therefore don't affect execution on the cluster.	1	0	https://issues.apache.org/jira/browse/APEXCORE-168
20370	APEXCORE-150	Can you add an example of where the MxN partitioning with multiple unifier partitions in parallel with the downstream operator partitions would not be sufficient to illustrate the point better.	1	0	https://issues.apache.org/jira/browse/APEXCORE-150
20371	APEXCORE-150	Why are cascading unifiers not sufficient?	0	0	https://issues.apache.org/jira/browse/APEXCORE-150
20372	APEXCORE-143	Could you please add a test case for this.	0	0	https://issues.apache.org/jira/browse/APEXCORE-143
20373	APEXCORE-143	What exactly are you trying to test? There is no guarantee that teardown will be executed on the cluster either, when the container process terminates due to a failure or because it was killed by the node manager.	0	0	https://issues.apache.org/jira/browse/APEXCORE-143
20374	APEXCORE-143	If an application is shutdown gracefully, teardown of all the operators is invoked. In case of local cluster, which is used to write test cases, Priyanka observed that teardown is not invoked on shutdown call on the cluster. We are shutting down the cluster and hence the application running within the local cluster is getting killed? No graceful shutdown here. Is there a way to shutdown the application in local mode ?	0	0	https://issues.apache.org/jira/browse/APEXCORE-143

#	Requirement	Questions	Answered	Successfully Answered	URL
20375	APEXCORE-143	teardown is not the place to effect changes to data. It is outside of the data flow and as we discussed, you cannot rely on it being called. Instead of looking for a workaround for the unit test, I would suggest you fix the real problem. Why can you not do this by implementing CheckpointListener or as part of rolling the output file, for example?	1	0	https://issues.apache.org/jira/browse/APEXCORE-143
20376	APEXCORE-139	Can I take this one?	1	0	https://issues.apache.org/jira/browse/APEXCORE-139
20377	APEXCORE-131	Can someone close this? This has been completed and backported to the various branches.	0	0	https://issues.apache.org/jira/browse/APEXCORE-131
20378	APEXCORE-128	why is this needed? Please open a discussion on dev@.	0	0	https://issues.apache.org/jira/browse/APEXCORE-128
20379	APEXCORE-119	Sounds like some sort of IMDG powered by partitions of an operator. Can you elaborate on performance and scalability characteristics? Can this be done without changes to the engine? Can this be done within the single thread per operator paradigm?	0	0	https://issues.apache.org/jira/browse/APEXCORE-119
20380	APEXCORE-119	Can delay operator be used for synchronizing windows across partitions?	0	0	https://issues.apache.org/jira/browse/APEXCORE-119
20381	APEXCORE-107	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/206#discussion_r50075888 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -2059,4 +2106,64 @@ public int hashCode() return result; } + /** + * This class hold Proxy Ports for a stream, and resolves them later after modules + * are expanded. + * + * If any source and sink added to a stream is ProxyInputPort or ProxyOutputPort + * we create a StreamLinkInfo with the stream reference and keep ProxyPorts In this class. + * Later we go over each Proxy port and find out the actual port connected to the ProxyPort + * and update StreamMeta. + */ + private static final class StreamLinkInfo { + OutputPort<?> source; + List<InputPort<?>> sinks = new ArrayList<>(); + final StreamMeta smeta; + + private StreamLinkInfo(StreamMeta smeta) { + this.smeta = smeta; + } + + private void setSource(OutputPort<?> source) { + if (!(source instanceof ProxyOutputPort)) { + throw new IllegalArgumentException("Invalid port type expected ProxyOutputPort"); --- End diff -- We are making proxy ports compulsory here. Is it not possible for the port to be a normal Default output port?	1	1	https://issues.apache.org/jira/browse/APEXCORE-107
20382	APEXCORE-107	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/313#discussion_r61536916 --- Diff: api/src/main/java/com/datatorrent/api/DagNode.java --- @@ -0,0 +1,24 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.api; + +/** + * Marker interface for a node in logical DAG + */ +public interface DagNode --- End diff -- How about naming it Vertex? Other suggestions?	1	1	https://issues.apache.org/jira/browse/APEXCORE-107
20383	APEXCORE-107	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/313#discussion_r61787475 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlanConfiguration.java --- @@ -2187,6 +2190,27 @@ public void populateDAG(LogicalPlan dag) } + private Vertex addNode(LogicalPlan dag, String name, Vertex vertex) --- End diff -- addVertex?	1	1	https://issues.apache.org/jira/browse/APEXCORE-107
20384	APEXCORE-103	Here are some comments I have about the design. These are repeated from the pull request. 1. DAG scoping currently is global for modules, each module's populateDAG sees the entire DAG. 2. The module has a populateDAG method with exact same syntax as in StreamingApplication. Is StreamingApplication also a module, should it extend that interface. 3. Setting properties for modules is too verbose. Module developer needs to repeat every property they want exposed with a setter and getter in JAVA. I don't disagree that module developer should be able to choose which properties from which operators need to be exposed but the current way seems to duplicate code. Here is a suggestion. a. Allow modules to specify which operators and properties can be accessible from outside. One way is in the "populateDAG" method of the module when adding the operator have the ability to specify if this operator can be accessible from outside and which or all properties can be accessible. b. Provide methods in ModuleMeta or elsewhere to set property values by specifying the operator name (friendly name) inside the module and property name. If this is allowed by a. above it is successful else it should fail. c. Allow a syntax in property files to specify the property in b. Example syntax dt.module.<modulename>.operator.<operatorname>.prop.<propname> 4. For attributes same mechanism as in 3 should apply for the operators that are exposed by the module. For property file, example syntax dt.module.<modulename>.operator.<operatorname>.attr.<attrname> 5. Module developers in addition to 3. and 4. above may choose to support module level properties and attributes. These should not be the default when 3. and 4. are possible but complementary in addition to them. In this case for properties they can implement setters and getters in the module. For attributes the user should still be able to set the attributes using the dag setAttribute method. You could introduce a method in the module to process attributes that can get called by the engine once everything is set. 6. For 5. above setting global properties and attributes for module is akin to ideas that have been proposed for the application as well. A consistent way must be possible for applications as well even if it is not implemented now. 7. For 5. or 6. above there should be a property file way of specifying the global module properties and attributes. Example syntax dt.module.<modulename>.prop.<propname>, dt.module.<modulename>.attr.<attrname>. Notice the difference with 3. c. and 4 above that there is no operator keyword here. 8. Partitioning needs to be consistent with what the user will expect when they see module as an entity. I will send an image of possible examples of how the user will expect the physical plan to look in certain cases.	0	0	https://issues.apache.org/jira/browse/APEXCORE-103
20385	APEXCORE-90	I've SSH'ed into people.apache.org and see where all the project websites are located, however the user:group ownership is all svnwc. I'm assuming this is the user that performs actions on behalf of the CMS, however since I can't get the "bookmarklet" thing to work, I'm not sure how to proceed. The relevant part of the guide actually says *THIS IS VERY OUTDATED AND NEEDS FIXING* and so I am not sure what to trust in that regard. I suppose I could simply create the necessary folder (/www/apex.apache.org) and add the files I want to that, then simply chown it to svnwc:svnwc, but something tells me I'm going down the wrong path with this tactic. [~cnauroth] or [~amolkekre], can you provide any guidance on this? I'll also send this out via dev@ update: just tried to create a folder in /www and it looks like I don't have permission to do that anyway so there you go...	0	0	https://issues.apache.org/jira/browse/APEXCORE-90

#	Requirement	Questions	Answered	Successfully Answered	URL
20386	APEXCORE-78	Why should this be changed?	1	1	https://issues.apache.org/jira/browse/APEXCORE-78
20387	APEXCORE-73	Hi Atri, That disk spooling should only happen when the memory used is getting close to configured limit. Are you seeing a different behavior. Thanks	1	0	https://issues.apache.org/jira/browse/APEXCORE-73
20388	APEXCORE-71	Given that this is really a workaround, can you please explain why we need to do this in first place? Should mention how it should work ideally (i.e. ability to pull the actual config values from the services).	1	0	https://issues.apache.org/jira/browse/APEXCORE-71
20389	APEXCORE-69	did you send an example on what the problem is? I cannot locate it... can you post it here? Thanks.	0	0	https://issues.apache.org/jira/browse/APEXCORE-69
20390	APEXCORE-69	my understanding is that Slowing down of the upstream should happen automatically as tcp is used for communication and tcp does the congestion control. So that is not working ?	0	0	https://issues.apache.org/jira/browse/APEXCORE-69
20391	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47722258 --- Diff: api/src/main/java/com/datatorrent/api/Operator.java --- @@ -99,6 +99,24 @@ public String toString() } /** + * DelayOperator is an operator that increments the outgoing streaming window + * id by one and allows loops in the "DAG". It is meant for iterative algorithms --- End diff -- can delay not be more than 1 window?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20392	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47722554 --- Diff: common/src/main/java/com/datatorrent/common/util/SimpleDelayOperator.java --- @@ -0,0 +1,71 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.common.util; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; + +/** + * SimpleDelayOperator + */ +public class SimpleDelayOperator<T> implements Operator.DelayOperator --- End diff -- What is the purpose of this class? I don't see any delay being implemented here...	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20393	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47722655 --- Diff: common/src/main/java/com/datatorrent/common/util/SimpleDelayOperator.java --- @@ -0,0 +1,71 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.common.util; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; + +/** + * SimpleDelayOperator + */ +public class SimpleDelayOperator<T> implements Operator.DelayOperator --- End diff -- Can this be done as follows to make it cleaner? `public class SimpleDelayOperator<T> extends BaseOperator implements Operator.DelayOperator`	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20394	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47723181 --- Diff: engine/src/main/java/com/datatorrent/stram/StreamingContainerManager.java --- @@ -1907,6 +1920,12 @@ public void updateRecoveryCheckpoints(PTOperator operator, UpdateCheckpointsCont if (operator.getState() == PTOperator.State.ACTIVE && (ctx.currentTms - operator.stats.lastWindowIdChangeTms) > operator.stats.windowProcessingTimeoutMillis) { // if the checkpoint is ahead, then it is not blocked but waiting for activation (state-less recovery, at-most-once) if (ctx.committedWindowId.longValue() >= operator.getRecoveryCheckpoint().windowId) { + LOG.info("Marking operator {} blocked: time since last windowId change={}, timeout={}, committed windowId={}, recovery windowId={}", --- End diff -- What is the impact of an operator being blocked? Should log level be warn if it impacts performance?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20395	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47723719 --- Diff: api/src/main/java/com/datatorrent/api/Context.java --- @@ -161,6 +161,13 @@ */ Attribute<Class<?>> TUPLE_CLASS = new Attribute<>(new Class2String<>()); + /** + * Attribute of input port. + * This is a read-only attribute to query whether the input port is connected to a DelayOperator + * This is for iterative processing. + */ + Attribute<Boolean> IS_CONNECTED_TO_DELAY_OPERATOR = new Attribute<>(false); --- End diff -- Currently all the attributes for both input and output ports are under same umbrella which some times is confusing about which attribute applies to which type of port. Should we separate them out?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20396	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47724731 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1613,6 +1636,37 @@ else if (stack.contains(successor)) { } } + public void findInvalidDelays(OperatorMeta om, List<List<String>> invalidDelays) { + stack.push(om); + // depth first successors traversal + boolean isDelayOperator = om.getOperator() instanceof Operator.DelayOperator; + if (isDelayOperator) { + if (om.getValue(OperatorContext.APPLICATION_WINDOW_COUNT) != 1) { + LOG.warn("detected DelayOperator having APPLICATION_WINDOW_COUNT not equal to 1"); --- End diff -- Should it not be that APPLICATION_WINDOW_COUNT of delay operator can't be greater than the down stream operator??	1	1	https://issues.apache.org/jira/browse/APEXCORE-60

#	Requirement	Questions	Answered	Successfully Answered	URL
20397	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47724883 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1613,6 +1636,37 @@ else if (stack.contains(successor)) { } } + public void findInvalidDelays(OperatorMeta om, List<List<String>> invalidDelays) + { + stack.push(om); + // depth first successors traversal + boolean isDelayOperator = om.getOperator() instanceof Operator.DelayOperator; + if (isDelayOperator) { + if (om.getValue(OperatorContext.APPLICATION_WINDOW_COUNT) != 1) { + LOG.warn("detected DelayOperator having APPLICATION_WINDOW_COUNT not equal to 1"); + invalidDelays.add(Collections.singletonList(om.getName())); + } + } + for (StreamMeta downStream: om.outputStreams.values()) { + for (InputPortMeta sink : downStream.sinks) { + OperatorMeta successor = sink.getOperatorWrapper(); + if (isDelayOperator) { + // Check whether all downstream operators are already visited in the path + if (successor != null && !stack.contains(successor)) { + LOG.warn("detected DelayOperator does not immediately output to a visited operator {}.{}->{}.{}", --- End diff -- should the log level be error as if there is invalid delay operator, validation throws ValidationException?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20398	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47725228 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/physical/PhysicalPlan.java --- @@ -327,8 +327,11 @@ public PhysicalPlan(LogicalPlan dag, PlanContext ctx) { boolean upstreamDeployed = true; - for (StreamMeta s : n.getInputStreams().values()) { - if (s.getSource() != null && !this.logicalToPTOperator.containsKey(s.getSource().getOperatorMeta())) { + for (Map.Entry<InputPortMeta, StreamMeta> entry : n.getInputStreams().entrySet()) { + InputPortMeta port = entry.getKey(); + StreamMeta s = entry.getValue(); + boolean delay = port.getValue(PortContext.IS_CONNECTED_TO_DELAY_OPERATOR); --- End diff -- I don't see `port` getting used anywhere else. Can we do `entry.getKey().getValue()`?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20399	APEXCORE-60	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47725499 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1613,6 +1636,37 @@ else if (stack.contains(successor)) { } } + public void findInvalidDelays(OperatorMeta om, List<List<String>> invalidDelays) + { + stack.push(om); + // depth first successors traversal + boolean isDelayOperator = om.getOperator() instanceof Operator.DelayOperator; + if (isDelayOperator) { + if (om.getValue(OperatorContext.APPLICATION_WINDOW_COUNT) != 1) { + LOG.warn("detected DelayOperator having APPLICATION_WINDOW_COUNT not equal to 1"); --- End diff -- My understanding is that the APPLICATION_WINDOW_COUNT of delay operator should be independent of the downstream. Can you shed some lights on this since you have implemented sliding window and probably have better understanding than me in terms of windowing.	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20400	APEXCORE-60	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47725662 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1613,6 +1636,37 @@ else if (stack.contains(successor)) { } } + public void findInvalidDelays(OperatorMeta om, List<List<String>> invalidDelays) + { + stack.push(om); + // depth first successors traversal + boolean isDelayOperator = om.getOperator() instanceof Operator.DelayOperator; + if (isDelayOperator) { + if (om.getValue(OperatorContext.APPLICATION_WINDOW_COUNT) != 1) { + LOG.warn("detected DelayOperator having APPLICATION_WINDOW_COUNT not equal to 1"); + invalidDelays.add(Collections.singletonList(om.getName())); + } + } + for (StreamMeta downStream: om.outputStreams.values()) { + for (InputPortMeta sink : downStream.sinks) { + OperatorMeta successor = sink.getOperatorWrapper(); + if (isDelayOperator) { + // Check whether all downstream operators are already visited in the path + if (successor != null && !stack.contains(successor)) { + LOG.warn("detected DelayOperator does not immediately output to a visited operator {}.{}->{}.{}", --- End diff -- It was actually a debug when detected cycles in existing code. should we change both to error()?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20401	APEXCORE-60	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47725699 --- Diff: engine/src/main/java/com/datatorrent/stram/StreamingContainerManager.java --- @@ -1907,6 +1920,12 @@ public void updateRecoveryCheckpoints(PTOperator operator, UpdateCheckpointsCont if (operator.getState() == PTOperator.State.ACTIVE && (ctx.currentTms - operator.stats.lastWindowIdChangeTms) > operator.stats.windowProcessingTimeoutMillis) { // if the checkpoint is ahead, then it is not blocked but waiting for activation (state-less recovery, at-most-once) if (ctx.committedWindowId.longValue() >= operator.getRecoveryCheckpoint().windowId) { + LOG.info("Marking operator {} blocked: time since last windowId change={}, timeout={}, committed windowId={}, recovery windowId={}", --- End diff -- How is this related to iteration support? Looks like some code was moved around.	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20402	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47729009 --- Diff: api/src/main/java/com/datatorrent/api/Context.java --- @@ -161,6 +161,13 @@ */ Attribute<Class<?>> TUPLE_CLASS = new Attribute<>(new Class2String<>()); + /** + * Attribute of input port. + * This is a read-only attribute to query whether the input port is connected to a DelayOperator + * This is for iterative processing. + */ + Attribute<Boolean> IS_CONNECTED_TO_DELAY_OPERATOR = new Attribute<>(false); --- End diff -- Should this be even an attribute? This set internally by platform so why expose it as an attribute..	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20403	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47782134 --- Diff: api/src/main/java/com/datatorrent/api/Context.java --- @@ -161,6 +161,13 @@ */ Attribute<Class<?>> TUPLE_CLASS = new Attribute<>(new Class2String<>()); + /** + * Attribute of input port. + * This is a read-only attribute to query whether the input port is connected to a DelayOperator + * This is for iterative processing. + */ + Attribute<Boolean> IS_CONNECTED_TO_DELAY_OPERATOR = new Attribute<>(false); --- End diff -- @davidyan74 : I checked the implementation of InputPortMeta.getValue() and looks like you don't need to check for the null.. I made changes in my local branch (https://github.com/gauravgopi123/incubator-apex-core/commit/9f4cf5cd3e4895c268c1d76f436e4db2223bfd5d) and ran tests and it worked fine Does that work for you?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60

#	Requirement	Questions	Answered	Successfully Answered	URL
20404	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47783467 --- Diff: common/src/main/java/com/datatorrent/common/util/SimpleDelayOperator.java --- @@ -0,0 +1,59 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.common.util; + +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; + +/** + * SimpleDelayOperator. This is an implementation of the DelayOperator that has one input port and one output + * port, and does a simple pass-through from the input port to the output port. Subclass of this operator can + * override this behavior by overriding processTuple(T tuple). + * + * Since the firstWindow method does not do anything, using this operator as-is means data loss during recovery. In + * order to achieve zero data loss during recovery, implementations must persist relevant tuples before the recovery + * checkpoint for emitting during the first window after recovery. + * + * Note that the engine will automatically does a +1 on the output window ID since it is a DelayOperator. --- End diff -- Should `engine will automatically does` be rephrased to `engine automatically does` or `engine will automatically do`? @amberarrow: any thoughts?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20405	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47783618 --- Diff: common/src/main/java/com/datatorrent/common/util/SimpleDelayOperator.java --- @@ -0,0 +1,59 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.common.util; + +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; + +/** + * SimpleDelayOperator. This is an implementation of the DelayOperator that has one input port and one output + * port, and does a simple pass-through from the input port to the output port. Subclass of this operator can + * override this behavior by overriding processTuple(T tuple). + * + * Since the firstWindow method does not do anything, using this operator as-is means data loss during recovery. In + * order to achieve zero data loss during recovery, implementations must persist relevant tuples before the recovery + * checkpoint for emitting during the first window after recovery. + * + * Note that the engine will automatically does a +1 on the output window ID since it is a DelayOperator. + */ +public class SimpleDelayOperator<T> extends BaseOperator implements Operator.DelayOperator --- End diff -- Does it make sense to name it SimpleNoOpDeplayOperator?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20406	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47822097 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1503,6 +1517,16 @@ private void validateThreadLocal(OperatorMeta om) { throw new ValidationException(msg); } + if (sm.source.operatorMeta.getOperator() instanceof Operator.DelayOperator) { + String msg = String.format("Locality %s invalid for delay operator %s", Locality.THREAD_LOCAL, sm.source.operatorMeta); + throw new ValidationException(msg); + } + } + if (om.getOperator() instanceof Operator.DelayOperator) { --- End diff -- Should this be outside the `for` loop?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20407	APEXCORE-60	Github user sandeshh commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r47980510 --- Diff: common/src/main/java/com/datatorrent/common/util/DefaultDelayOperator.java --- @@ -0,0 +1,57 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.common.util; + +import java.util.ArrayList; +import java.util.List; + +/** + * DefaultDelayOperator. This is the version of BaseDelayOperator that provides no data loss during recovery. It + * incurs a run-time cost per tuple, and all tuples of the checkpoint window will be part of the checkpoint state. + * Therefore if your application can tolerate data loss at recovery, BaseDelayOperator should be used instead. + */ +public class DefaultDelayOperator<T> extends BaseDelayOperator<T> --- End diff -- Do we need an implementation of the Delay Operator with windowDataManager ? Performance should be same for both the implementation. Only advantage I can think of for windowDataManager implementation is for delay of more than 1 window, it doesn't have to keep tuples of multiple windows in memory.	1	1	https://issues.apache.org/jira/browse/APEXCORE-60

#	Requirement	Questions	Answered	Successfully Answered	URL
20408	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r48219663 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/physical/PhysicalPlan.java --- @@ -940,17 +942,17 @@ private void updateStreamMappings(PMapping m) PTOperator slidingUnifier = StreamMapping.createSlidingUnifier(sourceOut.logicalStream, this, sourceOM.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT), slidingWindowCount); StreamMapping.addInput(slidingUnifier, sourceOut, null); - input = new PTInput(ipm.getKey().getPortName(), ipm.getValue(), oper, null, slidingUnifier.outputs.get(0)); + input = new PTInput(ipm.getKey().getPortName(), ipm.getValue(), oper, null, slidingUnifier.outputs.get(0), ipm.getKey().getValue(LogicalPlan.IS_CONNECTED_TO_DELAY_OPERATOR)); sourceMapping.outputStreams.get(ipm.getValue().getSource()).slidingUnifiers.add(slidingUnifier); } else { - input = new PTInput(ipm.getKey().getPortName(), ipm.getValue(), oper, null, sourceOut); + input = new PTInput(ipm.getKey().getPortName(), ipm.getValue(), oper, null, sourceOut, ipm.getKey().getValue(LogicalPlan.IS_CONNECTED_TO_DELAY_OPERATOR)); } oper.inputs.add(input); } } } - } else { + } else if (sourceMapping != null) { --- End diff -- Why is this additional required? sourceMapping is used in `if` condition too where this check is not being made	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20409	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r48221148 --- Diff: engine/src/test/java/com/datatorrent/stram/engine/GenericNodeTest.java --- @@ -277,6 +277,8 @@ public int getCount(boolean reset) gn.connectInputPort("ip1", reservoir1); gn.connectInputPort("ip2", reservoir2); gn.connectOutputPort("op", output); + gn.firstWindowMillis = 0L; --- End diff -- 387 line has firstWindowMillis initialized as 0 and here firstWindowMillis is initialized as 0L. Can we use same format?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20410	APEXCORE-60	Github user sandeshh commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r48757578 --- Diff: api/src/main/java/com/datatorrent/api/Operator.java --- @@ -99,6 +99,25 @@ public String toString() } /** + * DelayOperator is an operator of which the outgoing streaming window id is incremented by *one* by the + * engine, thus allowing loops in the "DAG". The output ports of a DelayOperator, if connected, *must* + * immediately connect to an upstream operator in the data flow path. Note that at least one output port of + * DelayOperator should be connected in order for the DelayOperator to serve its purpose. + * + * This is meant for iterative algorithms in the topology. A larger window increment can be simulated by an + * implementation of this interface. */ interface DelayOperator extends Operator { /** + * This method gets called at the first window of the execution. + * The implementation is expected to emit tuples for initialization and/or + * recovery. */ void firstWindow(); --- End diff -- Should this function have "Window ID" as an argument?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20411	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r48768935 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/GenericNode.java --- @@ -212,30 +225,60 @@ public final void run() long spinMillis = context.getValue(OperatorContext.SPIN_MILLIS); final boolean handleIdleTime = operator instanceof IdleTimeHandler; int totalQueues = inputs.size(); + int regularQueues = totalQueues; + // regularQueues is the number of queues that are not connected to a DelayOperator + for (String portName : inputs.keySet()) { + if (isInputPortConnectedToDelayOperator(portName)) { + regularQueues--; + } + } - ArrayList<SweepableReservoir> activeQueues = new ArrayList<SweepableReservoir>(); - activeQueues.addAll(inputs.values()); + ArrayList<Map.Entry<String, SweepableReservoir>> activeQueues = new ArrayList<>(); + activeQueues.addAll(inputs.entrySet()); int expectingBeginWindow = activeQueues.size(); int receivedEndWindow = 0; + long firstWindowId = -1; TupleTracker tracker; LinkedList<TupleTracker> resetTupleTracker = new LinkedList<TupleTracker>(); - try { do { - Iterator<SweepableReservoir> buffers = activeQueues.iterator(); + Iterator<Map.Entry<String, SweepableReservoir>> buffers = activeQueues.iterator(); activequeue: while (buffers.hasNext()) { - SweepableReservoir activePort = buffers.next(); + Map.Entry<String, SweepableReservoir> activePortEntry = buffers.next(); + SweepableReservoir activePort = activePortEntry.getValue(); Tuple t = activePort.sweep(); + boolean needResetWindow = false; if (t != null) { + boolean delay = (operator instanceof Operator.DelayOperator); + long windowAhead = 0; + if (delay) { + windowAhead = WindowGenerator.getAheadWindowId(t.getWindowId(), firstWindowMillis, windowWidthMillis, 1); + if (WindowGenerator.getBaseSecondsFromWindowId(windowAhead) > t.getBaseSeconds()) { --- End diff -- Why is this `if` condition needed here? needResetWindow is only used in BEGIN_WINDOW tuple type (line 271). Have this `if` condition there itself	1	1	https://issues.apache.org/jira/browse/APEXCORE-60

#	Requirement	Questions	Answered	Successfully Answered	URL
20412	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r48773375 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/GenericNode.java --- @@ -364,29 +413,49 @@ else if (!doCheckpoint) { if (tracker.ports[trackerIndex] == null) { tracker.ports[trackerIndex++] = activePort; break; - } - else if (tracker.ports[trackerIndex] == activePort) { + } else if (tracker.ports[trackerIndex] == activePort) { break; } trackerIndex++; } - if (trackerIndex == totalQueues) { - trackerIterator = resetTupleTracker.iterator(); + if (trackerIndex == regularQueues) { Iterator<TupleTracker> trackerIterator = resetTupleTracker.iterator(); while (trackerIterator.hasNext()) { if (trackerIterator.next().tuple.getBaseSeconds() <= baseSeconds) { trackerIterator.remove(); } } - for (int s = sinks.length; s-- > 0;) { - sinks[s].put(t); + if (!delay) { for (int s = sinks.length; s-- > 0;) { sinks[s].put(t); } controlTupleCount++; } - controlTupleCount++; - assert (activeQueues.isEmpty()); - activeQueues.addAll(inputs.values()); + if (!activeQueues.isEmpty()) { // make sure they are all queues from DelayOperator for (Map.Entry<String, SweepableReservoir> entry : activeQueues) { if (!isInputPortConnectedToDelayOperator(entry.getKey())) { assert (false); } } activeQueues.clear(); } activeQueues.addAll(inputs.entrySet()); expectingBeginWindow = activeQueues.size(); if (firstWindowId == -1) { if (delay) { for (int s = sinks.length; s-- > 0;) { sinks[s].put(t); } // if it's a DelayOperator and this is the first RESET_WINDOW (start) or END_STREAM --- End diff -- @davidyan74 : Do we need to send reset_tuple downstream? I don't we need it. If yes then do we not need to do controlTupleCount++;?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20413	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r48773751 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/GenericNode.java --- @@ -364,29 +413,49 @@ else if (!doCheckpoint) { if (tracker.ports[trackerIndex] == null) { tracker.ports[trackerIndex++] = activePort; break; - } - else if (tracker.ports[trackerIndex] == activePort) { + } else if (tracker.ports[trackerIndex] == activePort) { break; } trackerIndex++; } - if (trackerIndex == totalQueues) { - trackerIterator = resetTupleTracker.iterator(); + if (trackerIndex == regularQueues) { Iterator<TupleTracker> trackerIterator = resetTupleTracker.iterator(); while (trackerIterator.hasNext()) { if (trackerIterator.next().tuple.getBaseSeconds() <= baseSeconds) { trackerIterator.remove(); } } - for (int s = sinks.length; s-- > 0;) { - sinks[s].put(t); + if (!delay) { for (int s = sinks.length; s-- > 0;) { sinks[s].put(t); } controlTupleCount++; } - controlTupleCount++; - assert (activeQueues.isEmpty()); - activeQueues.addAll(inputs.values()); + if (!activeQueues.isEmpty()) { // make sure they are all queues from DelayOperator for (Map.Entry<String, SweepableReservoir> entry : activeQueues) { if (!isInputPortConnectedToDelayOperator(entry.getKey())) { assert (false); } } activeQueues.clear(); } activeQueues.addAll(inputs.entrySet()); expectingBeginWindow = activeQueues.size(); if (firstWindowId == -1) { if (delay) { for (int s = sinks.length; s-- > 0;) { sinks[s].put(t); } // if it's a DelayOperator and this is the first RESET_WINDOW (start) or END_STREAM // (recovery), fabricate the first window fabricateFirstWindow((Operator.DelayOperator)operator, windowAhead); } firstWindowId = t.getWindowId(); --- End diff -- Should this `firstWindowId` not be assigned when `delay` is `true` and `firstWindowId == -1`?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20414	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r48777786 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/GenericNode.java --- @@ -343,12 +389,15 @@ else if (!doCheckpoint) { * we will receive tuples which are equal to the number of input streams. */ activePort.remove(); - buffers.remove(); + if (isInputPortConnectedToDelayOperator(activePortEntry.getKey())) { break; // breaking out of the switch/case } buffers.remove(); --- End diff -- Should this not be before `if` condition above?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20415	APEXCORE-60	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r48784895 --- Diff: engine/src/main/java/com/datatorrent/stram/engine/GenericNode.java --- @@ -364,29 +413,49 @@ else if (!doCheckpoint) { if (tracker.ports[trackerIndex] == null) { tracker.ports[trackerIndex++] = activePort; break; - } - else if (tracker.ports[trackerIndex] == activePort) { + } else if (tracker.ports[trackerIndex] == activePort) { break; } trackerIndex++; } - if (trackerIndex == totalQueues) { - trackerIterator = resetTupleTracker.iterator(); + if (trackerIndex == regularQueues) { Iterator<TupleTracker> trackerIterator = resetTupleTracker.iterator(); while (trackerIterator.hasNext()) { if (trackerIterator.next().tuple.getBaseSeconds() <= baseSeconds) { trackerIterator.remove(); } } - for (int s = sinks.length; s-- > 0;) { - sinks[s].put(t); + if (!delay) { for (int s = sinks.length; s-- > 0;) { sinks[s].put(t); } controlTupleCount++; } - controlTupleCount++; - assert (activeQueues.isEmpty()); - activeQueues.addAll(inputs.values()); + if (!activeQueues.isEmpty()) { // make sure they are all queues from DelayOperator for (Map.Entry<String, SweepableReservoir> entry : activeQueues) { if (!isInputPortConnectedToDelayOperator(entry.getKey())) { assert (false); } } activeQueues.clear(); } activeQueues.addAll(inputs.entrySet()); expectingBeginWindow = activeQueues.size(); if (firstWindowId == -1) { if (delay) { for (int s = sinks.length; s-- > 0;) { sinks[s].put(t); } // if it's a DelayOperator and this is the first RESET_WINDOW (start) or END_STREAM --- End diff -- do we not need to do controlTupleCount++ as happening in `if` loop at line 430?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60

#	Requirement	Questions	Answered	Successfully Answered	URL
20416	APEXCORE-60	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/185#discussion_r49259156 --- Diff: engine/src/main/java/com/datatorrent/stram/StreamingContainerManager.java --- @@ -1917,25 +1930,30 @@ public void updateRecoveryCheckpoints(PTOperator operator, UpdateCheckpointsCont long currentWindowId = WindowGenerator.getWindowId(ctx.currentTms, this.vars.windowStartMillis, this.getLogicalPlan().getValue(LogicalPlan.STREAMING_WINDOW_SIZE_MILLIS)); maxCheckpoint = currentWindowId; } + ctx.visited.add(operator); // DFS downstream operators - for (PTOperator.PTOutput out : operator.getOutputs()) { - for (PTOperator.PTInput sink : out.sinks) { - PTOperator sinkOperator = sink.target; - if (!ctx.visited.contains(sinkOperator)) { - // downstream traversal - updateRecoveryCheckpoints(sinkOperator, ctx); - } - // recovery window id cannot move backwards - // when dynamically adding new operators - if (sinkOperator.getRecoveryCheckpoint().windowId >= operator.getRecoveryCheckpoint().windowId) { - maxCheckpoint = Math.min(maxCheckpoint, sinkOperator.getRecoveryCheckpoint().windowId); - } + if (operator.getOperatorMeta().getOperator() instanceof Operator.DelayOperator) { + addVisited(operator, ctx); + } else { --- End diff -- @tweise I think I'm not doing this correctly and hence the out-of-sequence tuple in the unit test. My debugging indicates that the recovery checkpoints are not in sync for the operators that try to recover. Can you please review this and see what I'm doing wrong?	1	1	https://issues.apache.org/jira/browse/APEXCORE-60
20417	APEXCORE-59	why is this minor priority?	1	0	https://issues.apache.org/jira/browse/APEXCORE-59
20418	APEXCORE-55	h1. Module support. h2. Introduction A module framework provides a way to use a group of operator as a single entity in a application. This will allow group of operator to implement a higher level functionality, which can be embedded in another streaming application. For example a module M1 is defined as following DAG. where M2 is also another module. {code} O3 | O1 -> O2 | M2 {code} and Module M2 as {code} O4 -> O5 {code} And application is defined as {code} A -> M1 -> B {code} Then during logical plan expansion module M1 will be replaced by its operator and resulting DAG will be exapnded again till no modules remain in the DAG. For example given application will expand as below. {code} O3 | A -> O1 -> O2 -> B | O4 -> O5 {code} h2. Goals Logically a Module will contain. * A List of input ports. * A list of output ports. * Configuration properties. ** Name of the property. ** Type of the property. * Attributes? (not decided which attribute make sense for a module). For most of the modules all internal component and their connection is already know. The other type of module will be populate the internal components based on properties set on the module. We will be targeting first type of modules with static sub-dag. h3. Properties A module will have its own properties. Module will hide properties of internal components, means user will only be able to set properties defined on module, and need not know the internal structure of module. Module will populate internal components properties as required based on its properties. Does it needs to have support for setting module property at runtime? This will require us to maintain runtime object for module. i.e module definition can not go away after logical plan expansion. Also any complex execution done my module inside setter needs to be tracked and sent to container running operators in the module. h3. Attributes partitions?? h3. Module Definition Languages User should be able to write modules in java. And other JVM compatible languages. We plan to provide compact api for defining module in javascript which are loaded in java using Javascript/Java bridge (JSR 292). A similar interface will be provided for python using Jython. The syntax will look like mostly like configuration description. h3. Application Builder Compatibility Another goal is that user should be able to define modules using application builder, where user will build a DAG for module and save it as a module. To support this functionality we will have support for defining module in json. As json does not support defining complex expressions, this will be applicable for simpler modules. h3. Port Type Checking We need to have type checking for the module ports. If Module is defined in java then we can have type check enforced by the compiler. In other cases we need to have a way to enforce type check during port connection time or logical dag validation time. h2. Implementation h3. API {code} public interface Module extends Component<ModuleContext> { /** * get list of input port exported by the module. */ public Map<String, Operator.InputPort> getInputPorts(); /** * get list of output ports exported by the module. */ public Map<String, Operator.OutputPort> getOutputPorts(); /** * This function will be called by engine to add operators/modules in the application DAG and connect them as per modules requirement. */ public void flattenDAG(DAG dag, Configuration conf); /** * return property of the module */ public Object getProperty(String key); /** * set property on the module */ public Object setProperty(String key, Object value); /** * return list of properties supported by the module. */ public Collection<String> getProperties(); } {code}	0	0	https://issues.apache.org/jira/browse/APEXCORE-55
20419	APEXCORE-55	What is the rationale to make ports on modules dynamic? Do you have an example use case? What is the runtime representation of a module? What drives the dynamic structure of a module? Can setting a property change ports and operators in it? If so, what happens when such properties are set at runtime? Full partitioning support for modules will require significant changes to the current physical plan translation. Has the impact been assessed? Or are you planning to have corresponding restrictions in the initial module implementation? How does this work converge with APEXCORE-3? This JIRA needs to be broken down into subtasks and the relationship to other parts of the engine more clearly called out.	1	0	https://issues.apache.org/jira/browse/APEXCORE-55
20420	APEXCORE-34	Did you confirm that after running the command all version strings have been replaced?	0	0	https://issues.apache.org/jira/browse/APEXCORE-34
20421	APEXCORE-34	Github user chinmaykolhatkar commented on the pull request: https://github.com/apache/incubator-apex-site/pull/10#issuecomment-181725435 Closing this PR.. Will reopen when changes in core gets in.	0	0	https://issues.apache.org/jira/browse/APEXCORE-34
20422	APEXCORE-31	Can we also evaluate akka, few projects are using it. http://www.slideshare.net/fhueske/flink-akka20150128	0	0	https://issues.apache.org/jira/browse/APEXCORE-31
20423	APEXCORE-22	Does this fix handle the below scenario? I have two different DAGs to run, I shall use A in one DAG and B in another. A extends B, I have declared output ports in A and B as well. When I use A in my DAG, it shows output port overridden error in 3.1.0 release. This requirement can be fulfilled with different implementation for the use case but I am just curious.	1	1	https://issues.apache.org/jira/browse/APEXCORE-22
20424	APEXCORE-19	can you list the a couple of use cases which need this?	0	0	https://issues.apache.org/jira/browse/APEXCORE-19
20425	APEXCORE-17	[~davidyan] in ClusterAppStats you have used @AppDataPushAgent.Metric on methods. Is that really needed? Do you extract the names of these metrics for the getterMethod. If yes, where is that being done?	1	0	https://issues.apache.org/jira/browse/APEXCORE-17
20426	APEXCORE-17	1. If this change is not included, then we know for sure that it needs to be changed. If there is a constant added to spell out that it can support phase, then may be we will need to change it may be we will not need to change it. which one is better? 2. I realize that there is not a lot of time left hence I recommended putting the stubs in and having the implementation later. You can defer even specifying various phases. The one which you have implemented, have it specified (or even defaulted to). -- Chetan	1	1	https://issues.apache.org/jira/browse/APEXCORE-17

#	Requirement	Questions	Answered	Successfully Answered	URL
20427	APEXCORE-17	I have created another ticket for phases. Can I resolve this?	1	0	https://issues.apache.org/jira/browse/APEXCORE-17
20428	APEXCORE-16	What is the agreement on white spaces for type cast? Should APEX use white space after the cast? Please see [http://www.oracle.com/technetwork/java/javase/documentation/codeconventions-141388.html]: Casts should be followed by a blank space. Examples: {noformat} myMethod((byte) aNum, (Object) x); myMethod((int) (cp + 5), ((int) (i + 3)) + 1); {noformat}	0	0	https://issues.apache.org/jira/browse/APEXCORE-16
20429	APEXCORE-13	Why cause unnecessarily additional HDFS operations in the master when the same is available for free in the container? The container should propagate the checkpoint only after it has been written (approach 1).	1	1	https://issues.apache.org/jira/browse/APEXCORE-13
20430	APEXCORE-10	this is a very important requirement, when can we have it	0	0	https://issues.apache.org/jira/browse/APEXCORE-10
20431	APEXCORE-10	There are two pieces to this: 1) how you express the constraint and 2) how you implement it using the RM API. 1) I understand you are proposing to define the operators that should not be on the same host as an attribute. How will this translate to the physical plan? How do we express that 2 partitions of the same operator are not on the same host? 2) Yes, we may run out of resources and launch time will increase of we sequence the sets. All containers that don't have the affinity constraint will still be allocated first? Why does the rest need to be sequential?	0	0	https://issues.apache.org/jira/browse/APEXCORE-10
20432	APEXCORE-10	As an additional note, is there any way that you could see specifying anti-affinity between partitions of an operator? At the moment we can specify affinity/anti-affinity between logical operators but it would be very useful to specify anti-affinity between partitions of a single logical operator (for the sake of durability, for example).	0	0	https://issues.apache.org/jira/browse/APEXCORE-10
20433	APEXCORE-10	A couple of questions - Would anti-affinity for partitions take precedence over the anti-affinity of operators? Can this be configured? Wouldn't there be a cascading effect of partitioning on the anti-affinity of operators? In case of clashes between anti-affinity and locality, what takes precedence? (For example, would the partitions be created across racks and different operators remain in the nodes on the same rack, or would all the partitions be within the same rack and the locality for operators will not be honored?) What would happen if a node crashes and YARN allocates a node on which another operator/partition is already active. YARN would not have any knowledge about the anti-affinity. So what would be our approach in this case? Ideally, if we get resources from YARN on a certain node, we shoul dbe able to run the instance/operator on the same node. However, is there a guarantee that YARN will allocate the exact resources we have asked for? Further even if it does, would we honour the "strict" policy and prevent the instance/operator from coming up or we bypass the same.	0	0	https://issues.apache.org/jira/browse/APEXCORE-10
20434	APEXCORE-10	In case of anti-affinity and host-locality set among partitions of operator, yes, we need to set precedence order. Another option is: should we check for such contradicting settings this in dag.validate phase itself?	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20435	APEXCORE-10	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/250#issuecomment-201640472 @ishark there are merge conflicts, can you please rebase.	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20436	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r57506464 --- Diff: api/src/main/java/com/datatorrent/api/AffinityRule.java --- @@ -0,0 +1,211 @@ + +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.api; + +import java.io.Serializable; +import java.util.List; + +import com.datatorrent.api.DAG.Locality; + +/** + * Affinity rule specifies constraints for physical deployment of operator + * containers. There are two types of rules that can be specified: Affinity and + * Anti-affinity. Each rule contains list of operators or pair of 2 operators or + * a regex that should match at least 2 operators. Based on the type of rule, + * affinity or anti-affinity, the operators will be deployed together or away + * from each other. The locality indicates the level at which the rule should be + * applied. E.g. CONTAINER_LOCAL affinity would indicate operators Should be + * allocated within same container NODE_LOCAL anti-affinity indicates that the + * operators should not be allocated on the same node. The rule can be either + * strict or relaxed. + * + */ +public class AffinityRule implements Serializable +{ + @Override + public String toString() + { + return "AffinityRule {operatorsList=" + operatorsList + ", operatorRegex=" + operatorRegex + ", operators=" + operators + ", locality=" + locality + ", type=" + type + ", relaxLocality=" + relaxLocality + "}"; + } + + private static final long serialVersionUID = 1071315049298753 86L; + + /** + * Pair of operator names to specify affinity rule + * The order of operators is not considered in this class + * i.e. OperatorPair("O1", "O2") is equal to OperatorPair("O2", "O1") + */ + public static class OperatorPair implements Serializable --- End diff -- Why is it a pair and not a collection, looks like the regex can match more operators?	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20437	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r57507058 --- Diff: api/src/main/java/com/datatorrent/api/StringCodec.java --- @@ -377,4 +383,43 @@ public String toString(Class<? extends T> clazz) private static final long serialVersionUID = 201312082053L; } + public class JsonStringCodec<T> implements StringCodec<T>, Serializable + { + private static final long serialVersionUID = 2513932518264776006L; + Class<?> clazz; + + public JsonStringCodec(Class<T> clazz) + { + this.clazz = clazz; + } + + @Override + public T fromString(String string) + { + try { + ObjectMapper mapper = new ObjectMapper(); + mapper.configure(DeserializationConfig.Feature.FAIL_ON_UNKNOWN_PROPERTIES, false); + ObjectReader reader = mapper.reader(clazz); + return reader.readValue(string); + } catch (IOException e) { + DTThrowable.wrapIfChecked(e); + } + return null; + } + + @Override + public String toString(T pojo) + { + try { + ObjectMapper mapper = new ObjectMapper(); + mapper.configure(DeserializationConfig.Feature.FAIL_ON_UNKNOWN_PROPERTIES, false); + ObjectWriter writer = mapper.writer(); + return writer.writeValueAsString(pojo); + } catch (IOException e) { + DTThrowable.wrapIfChecked(e); --- End diff -- This is not specific to this PR but in general, should we consider using Throwables.propagate() instead?	1	1	https://issues.apache.org/jira/browse/APEXCORE-10

#	Requirement	Questions	Answered	Successfully Answered	URL
20438	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r57507255 --- Diff: engine/src/main/java/com/datatorrent/stram/ContainerRequestHandler.java --- @@ -0,0 +1,59 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.stram; + +import java.util.List; +import java.util.Map; + +import org.apache.commons.lang3.tuple.MutablePair; +import org.apache.hadoop.yarn.client.api.AMRMClient; +import org.apache.hadoop.yarn.client.api.AMRMClient.ContainerRequest; + +/** + * Handles creating container requests and reissuing them on timeout + * + */ +public class ContainerRequestHandler +{ + protected static final int NUMBER_MISSED_HEARTBEATS = 30; + + public void reissueContainerRequests(AMRMClient<ContainerRequest> amRmClient, Map<StreamingContainerAgent.ContainerStartRequest, MutablePair<Integer, ContainerRequest>> requestedResources, int loopCounter, ResourceRequestHandler resourceRequestor, List<ContainerRequest> containerRequests, List<ContainerRequest> removedContainerRequests) + { + if (!requestedResources.isEmpty()) { + // resourceRequestor.clearNodeMapping(); + for (Map.Entry<StreamingContainerAgent.ContainerStartRequest, MutablePair<Integer, ContainerRequest>> entry : requestedResources.entrySet()) { + if ((loopCounter - entry.getValue().getKey()) > NUMBER_MISSED_HEARTBEATS) { --- End diff -- What happens otherwise?	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20439	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r57508355 --- Diff: engine/src/main/java/com/datatorrent/stram/ResourceRequestHandler.java --- @@ -108,6 +113,23 @@ public void updateNodeReports(List<NodeReport> nodeReports) } } + public List<String> getNodesExceptHost(List<String> hostNames) --- End diff -- Does hostNames always contain one element? The contains operation is expensive...	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20440	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r57508918 --- Diff: engine/src/main/java/com/datatorrent/stram/StreamingAppMasterService.java --- @@ -668,7 +671,10 @@ private void execute() throws YarnException, IOException int numReleasedContainers = 0; int nextRequestPriority = 0; ResourceRequestHandler resourceRequestor = new ResourceRequestHandler(); + // Use override for container requestor in case of cloudera distribution,to handle host specific requests + ContainerRequestHandler containerRequestor = System.getenv().containsKey("CDH_HADOOP_BIN") ? new ContainerRequestHandlerCloudera() : new ContainerRequestHandler(); --- End diff -- Would it makes sense to combine ResourceRequestHandler and ContainerRequestHandler?	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20441	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r58620829 --- Diff: api/src/main/java/com/datatorrent/api/AffinityRule.java --- @@ -0,0 +1,153 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.api; + +import java.io.Serializable; +import java.util.Arrays; +import java.util.List; + +import com.datatorrent.api.DAG.Locality; + +/** + * Affinity rule specifies constraints for physical deployment of operator + * containers. There are two types of rules that can be specified: Affinity and + * Anti-affinity. Each rule contains list of operators or pair of 2 operators or + * a regex that should match at least 2 operators. Based on the type of rule, + * affinity or anti-affinity, the operators will be deployed together or away + * from each other. The locality indicates the level at which the rule should be + * applied. E.g. CONTAINER_LOCAL affinity would indicate operators Should be + * allocated within same container NODE_LOCAL anti-affinity indicates that the + * operators should not be allocated on the same node. The rule can be either + * strict or relaxed. + */ +public class AffinityRule implements Serializable +{ + @Override + public String toString() + { + return "AffinityRule {operatorsList=" + operatorsList + ", operatorRegex=" + operatorRegex + ", locality=" + locality + ", type=" + type + ", relaxLocality=" + relaxLocality + "}"; + } + + private static final long serialVersionUID = 107131504929875386L; + + /** + * Type of affinity rule setting affects how operators are scheduled for + * deployment by the platform. + */ + public static enum Type + { + /** + * AFFINITY indicates that operators in the rule should be deployed within + * locality specified in the rule + */ + AFFINITY, + /** + * ANTI_AFFINITY indicates that operators in the rule should NOT deployed + * within same locality as specified in rule + */ + ANTI_AFFINITY + } + + private List<String> operatorsList; + private String operatorRegex; + private Locality locality; + private Type type; + private boolean relaxLocality; + + public AffinityRule() + { + } + + public AffinityRule(Type type, Locality locality, boolean relaxLocality) + { + this.type = type; + this.locality = locality; + this.setRelaxLocality(relaxLocality); + } + + public AffinityRule(Type type, Locality locality, boolean relaxLocality, String... operators) --- End diff -- Does it require at least 2 operators? Is so, maybe make it String, String...	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20442	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r58926771 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/physical/PhysicalPlan.java --- @@ -357,25 +417,51 @@ public PhysicalPlan(LogicalPlan dag, PlanContext ctx) { if (!container.operators.isEmpty()) { LOG.warn("Operator {} shares container without locality contraint due to insufficient resources.", oper); } + // TODO: Check if PT Operators conflict in anti-affinity, Pick the first container that does not conflict --- End diff -- Why would this happen here and not during logical plan validation?	1	1	https://issues.apache.org/jira/browse/APEXCORE-10

#	Requirement	Questions	Answered	Successfully Answered	URL
20443	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r58927033 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/physical/PhysicalPlan.java --- @@ -386,6 +472,38 @@ public PhysicalPlan(LogicalPlan dag, PlanContext ctx) { this.undeployOpers.clear(); } + public void setAntiAffinityForContainers(LogicalPlan dag, Collection<AffinityRule> affinityRules, Map<PTOperator, PTContainer> operatorContainerMap) --- End diff -- Would it make sense to maintain the relationship in PhysicalPlan instead of adding to each container?	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20444	APEXCORE-10	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/250#discussion_r58927359 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/physical/PhysicalPlan.java --- @@ -344,9 +346,67 @@ public PhysicalPlan(LogicalPlan dag, PlanContext ctx) { addLogicalOperator(n); } } + + inlinePrefs.prefs.clear(); + localityPrefs.prefs.clear(); + + // Add inlinePrefs and localityPreds for affinity rules --- End diff -- How do we handle this with dynamic plan changes?	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20445	APEXCORE-10	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/250#issuecomment-207202664 Are the affinity sets updated when containers are added or removed due to dynamic partitioning? If they are not, please create a follow up JIRA so we can address it later.	1	1	https://issues.apache.org/jira/browse/APEXCORE-10
20446	APEXCORE-9	this is important when can we get this	1	0	https://issues.apache.org/jira/browse/APEXCORE-9
20447	APEXCORE-4	We forgot to discuss this on Friday. I think Thomas should probably be in this discussion and he is coming back next weekend. Do you want to take this up after Thomas comes back?	0	0	https://issues.apache.org/jira/browse/APEXCORE-4
20448	APEXCORE-3	Which version of DTRTS this feature will be available? We may need a patch for 2.0.1/	0	0	https://issues.apache.org/jira/browse/APEXCORE-3
20449	APEXCORE-3	In a DAG say A -> B -> C Can a single physical operator say B, be spawned as B1 -> B2 -> B3 and append itself to the DAG at run time as part of this feature? Here B1, B2, B3 are instances of B. The values of B1, B2, B3 and so on are populated based on the output of Operator A.	0	0	https://issues.apache.org/jira/browse/APEXCORE-3
20450	APEXCORE-3	Can you explain a bit more as to how this will benefit APEXCORE-3? From what I understand, Module will be basically a replacable operator in DAG. It might be statically generated or runtime expanded. Is there any benefit that you see apart from code clarity (which might be a good enough reason)?	0	0	https://issues.apache.org/jira/browse/APEXCORE-3
20451	APEXCORE-3	Github user tushargosavi commented on the pull request: https://github.com/apache/incubator-apex-core/pull/189#issuecomment-166283164 @tweise following commits fixes checkstyle violations. https://github.com/tushargosavi/incubator-apex-core/commit/7e211ce69a433fdc9bef72b3b0cba0a4fb7111b7 I have these changes in APEXCORE-3 branch, with rebase on top of current apache/devel-3. https://github.com/tushargosavi/incubator-apex-core/tree/APEXCORE-3 How to include this patch, Should I open a pull request against your devel-3 branch?	1	1	https://issues.apache.org/jira/browse/APEXCORE-3
20452	APEXCORE-3	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/189#discussion_r48163483 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1116,13 +1320,78 @@ public StreamMeta addStream(String id) public <T> StreamMeta addStream(String id, Operator.OutputPort<? extends T> source, Operator.InputPort<? super T>... sinks) { StreamMeta s = addStream(id); - s.setSource(source); - for (Operator.InputPort<?> sink: sinks) { - s.addSink(sink); + id = s.id; + ArrayListMultimap<Operator.OutputPort<?>, Operator.InputPort<?>> streamMap = ArrayListMultimap.create(); + if (!(source instanceof ProxyOutputPort)) { + s.setSource(source); + } + for (Operator.InputPort<?> sink : sinks) { + if (source instanceof ProxyOutputPort || sink instanceof ProxyInputPort) { + streamMap.put(source, sink); + streamLinks.put(id, streamMap); + } else { + if (s.getSource() == null) { + s.setSource(source); + } + s.addSink(sink); + } } return s; } + /** + * This will be called once the Logical Dag is expanded, and the proxy input and proxy output ports are populated with the actual ports that they refer to + * This method adds sources and sinks for the StreamMeta objects which were left empty in the addStream call. */ public void applyStreamLinks() { for (String id : streamLinks.keySet()) { StreamMeta s = getStream(id); for (Map.Entry<Operator.OutputPort<?>, Operator.InputPort<?>> pair : streamLinks.get(id).entries()) { if (s.getSource() == null) { Operator.OutputPort<?> outputPort = pair.getKey(); while (outputPort instanceof ProxyOutputPort) { outputPort = ((ProxyOutputPort<?>)outputPort).get(); } s.setSource(outputPort); } Operator.InputPort<?> inputPort = pair.getValue(); while (inputPort instanceof ProxyInputPort) { inputPort = ((ProxyInputPort<?>)inputPort).get(); } s.addSink(inputPort); } } } @SuppressWarnings({ "unchecked", "rawtypes" }) private void addDAGToCurrentDAG(ModuleMeta moduleMeta) { LogicalPlan subDag = moduleMeta.getDag(); String subDAGName = moduleMeta.getName(); String name; for (OperatorMeta operatorMeta : subDag.getAllOperators()) { name = subDAGName + MODULE_NAMESPACE_SEPARATOR + operatorMeta.getName(); this.addOperator(name, operatorMeta.getOperator()); OperatorMeta operatorMetaNew = this.getOperatorMeta(name); operatorMetaNew.setModuleName(operatorMeta.getModuleName() == null ? subDAGName : subDAGName + MODULE_NAMESPACE_SEPARATOR + operatorMeta.getModuleName()); } for (StreamMeta streamMeta : subDag.getAllStreams()) { --- End diff -- Why loop over all streams after having visited all operators already?	1	1	https://issues.apache.org/jira/browse/APEXCORE-3

#	Requirement	Questions	Answered	Successfully Answered	URL
20453	APEXCORE-3	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-core/pull/189#discussion_r48165286 --- Diff: engine/src/main/java/com/datatorrent/stram/plan/logical/LogicalPlan.java --- @@ -1116,13 +1320,78 @@ public StreamMeta addStream(String id) public <T> StreamMeta addStream(String id, Operator.OutputPort<? extends T> source, Operator.InputPort<? super T>... sinks) { StreamMeta s = addStream(id); - s.setSource(source); - for (Operator.InputPort<?> sink: sinks) { - s.addSink(sink); + id = s.id; + ArrayListMultimap<Operator.OutputPort<?>, Operator.InputPort<?>> streamMap = ArrayListMultimap.create(); + if (!(source instanceof ProxyOutputPort)) { + s.setSource(source); + } + for (Operator.InputPort<?> sink : sinks) { + if (source instanceof ProxyOutputPort || sink instanceof ProxyInputPort) { + streamMap.put(source, sink); + streamLinks.put(id, streamMap); + } else { + if (s.getSource() == null) { + s.setSource(source); + } + s.addSink(sink); + } } return s; } + /** + * This will be called once the Logical Dag is expanded, and the proxy input and proxy output ports are populated with the actual ports that they refer to + * This method adds sources and sinks for the StreamMeta objects which were left empty in the addStream call. + */ + public void applyStreamLinks() { + for (String id : streamLinks.keySet()) { + StreamMeta s = getStream(id); + for (Map.Entry<Operator.OutputPort<?>, Operator.InputPort<?>> pair : streamLinks.get(id).entries()) { + if (s.getSource() == null) { + Operator.OutputPort<?> outputPort = pair.getKey(); + while (outputPort instanceof ProxyOutputPort) { + outputPort = ((ProxyOutputPort<?>)outputPort).get(); + } + s.setSource(outputPort); + } + Operator.InputPort<?> inputPort = pair.getValue(); + while (inputPort instanceof ProxyInputPort) { + inputPort = ((ProxyInputPort<?>)inputPort).get(); + } + s.addSink(inputPort); + } + } + } + @SuppressWarnings({ "unchecked", "rawtypes" }) + private void addDAGToCurrentDAG(ModuleMeta moduleMeta) { + LogicalPlan subDag = moduleMeta.getDag(); + String subDAGName = moduleMeta.getName(); + String name; + for (OperatorMeta operatorMeta : subDag.getAllOperators()) { + name = subDAGName + MODULE_NAMESPACE_SEPARATOR + operatorMeta.getName(); + this.addOperator(name, operatorMeta.getOperator()); + OperatorMeta operatorMetaNew = this.getOperatorMeta(name); + operatorMetaNew.setModuleName(operatorMeta.getModuleName() == null ? subDAGName : subDAGName + MODULE_NAMESPACE_SEPARATOR + operatorMeta.getModuleName()); + } + for (StreamMeta streamMeta : subDag.getAllStreams()) { + OutputPortMeta sourceMeta = streamMeta.getSource(); + List<InputPort<?>> ports = new LinkedList<>(); + for (InputPortMeta inputPortMeta : streamMeta.getSinks()) { + ports.add(inputPortMeta.getPortObject()); + } + InputPort[] inputPorts = ports.toArray(new InputPort[]{}); + name = subDAGName + MODULE_NAMESPACE_SEPARATOR + streamMeta.getName(); + StreamMeta streamMetaNew = this.addStream(name, sourceMeta.getPortObject(), inputPorts); + streamMetaNew.setModuleName(streamMeta.getModuleName() == null ? subDAGName : subDAGName + "_" + streamMeta.getModuleName()); --- End diff -- Shouldn't the module name be same for all operators and streams? Why construct it in the loops?	1	1	https://issues.apache.org/jira/browse/APEXCORE-3
20454	APEXCORE-3	Github user tushargosavi commented on the pull request: https://github.com/apache/incubator-apex-core/pull/189#issuecomment-166447656 I have made changes in my https://github.com/tushargosavi/incubator-apex-core/tree/pull189 branch. I will have to do force push in your branch because checkstyle changes are worked into original commits. Is it fine?	1	1	https://issues.apache.org/jira/browse/APEXCORE-3
20455	APEXMALHAR-2507	Hi Bhupesh, It looks like we already have an example doing pojo inner join using windowed merge operator ([https://github.com/apache/apex-malhar/blob/master/library/src/test/java/org/apache/apex/malhar/lib/window/impl/PojoInnerJoinTestApplication.java]). Should we move this app to the example folder? -Shunxin	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2507
20456	APEXMALHAR-2507	Hi all, Hitesh, could you elaborate on what you meant by "lost its way"? I will take a closer look at the app. Thomas, I agree [APEXMALHAR-2300] should be closed. I think the current example in innjoin folder will be removed in [APEXMALHAR-2488]	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2507
20457	APEXMALHAR-2489	[~vrozov] what do you mean by _with modification_? I implemented the algorithm as described and it worked as expected, and also it removed the overflow issue.	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2489
20458	APEXMALHAR-2459	IMO there is no need to touch 0.8 connectors, the decision to leave them as they are was made already and the can remain (and expire) in the contrib module. I wonder if we really need such elaborate structure for 0.9 and 0.10 or just have separate modules. How much of difference/overlap do you estimate to be there?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2459
20459	APEXMALHAR-2458	[~thw] are you suggesting the code of the application use the 0.9 operators or leave the application code and the change the dependency to 0.9. Is Kafka 0.8 not much in use? If not the example with 0.8 code may be useful to folks.	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2458
20460	APEXMALHAR-2438	I've been working with the above stores for a while now, could this be assigned to me ?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2438
20461	APEXMALHAR-2386	Could I work on this?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2386
20462	APEXMALHAR-2376	Hi, I am unable to assign this card to myself. Can someone please add my name to the contributor list?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2376
20463	APEXMALHAR-2376	[~thw] : Could you please make Akshay contributor and assign this JIRA to him.	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2376
20464	APEXMALHAR-2366	As the number of elements increases, the probability of false positives increases; are we planning to recreate the filter with a larger bit vector or is this automatically done by the Hadoop implementation ?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2366
20465	APEXMALHAR-2366	Hi [~brightchen] Sorry for commenting so late. A bloom filter implementation is already implemented by [~chaithu] in the Megh library. You can see it here: https://github.com/DataTorrent/Megh/blob/master/library/src/main/java/com/datatorrent/lib/bucket/bloomFilter/BloomFilter.java Can you please see if this implementation can be reused? I am asking this because the one in Megh is well tested as part of an earlier Deduper implementation.	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2366
20466	APEXMALHAR-2365	Hi, I am unable to assign this card to myself. Can someone please add my name to the contributor list?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2365
20467	APEXMALHAR-2365	Can someone from PMC list assign this JIRA to Shraddha - jogshraddha . [~thw]: Could you please assign this and make Shraddha contributor.	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2365
20468	APEXMALHAR-2365	Couple of thoughts: 1. Some use cases why this is needed would be helpful. 2. Is the plan to convert the byte array to a String as a first step ? If so, the encoding for that conversion should be a parameter.	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2365
20469	APEXMALHAR-2352	- The insert to the map with item of ImmutablePair cost 2 to 3 times than item of com.datatorrent.common.util.Pair - Kyro serializer serialize cost 6 to 8 times than customized serializer. The tuple which to serialize is generated with code "new ImmutablePair(new Window.TimeWindow<>(now + random.nextInt(100), random.nextInt(100)), "" + random.nextInt(1000));"	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2352
20470	APEXMALHAR-2339	when use default AccumulationMode DISCARDING, it will have big impact for performance due to AbstractWindowedOperator.clearWindowData(Window)	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2339

#	Requirement	Questions	Answered	Successfully Answered	URL
20471	APEXMALHAR-2322	Shouldn't we fix the issues instead? I would prefer the CI fails when code with invalid javadoc syntax is submitted.	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2322
20472	APEXMALHAR-2321	I changed to use keyStream and valueStream to manage the memory. But as we can't enforce the caller to use keyStream and valueStream to manage the memory, So it is also need to cleanup the maps. [~csingh] How do you think? and I saw you also clear the fileCache and close readers in freeMemory(). Is it better to move to another function or rename this method into something like releaseResource?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2321
20473	APEXMALHAR-2283	The exactly-once output logic is suspect. Why is it using the same key for all messages (appId+operatorId), why does it track extra window state in the operator and why does it rely on the hashcode of the object. In cases where the application can provide a unique message id, it should also be possible to use it for the key. It should be possible with the state stored in Kafka alone to do the dedup. The operator is also not easy to extend, we tried to implement output to topic depending on the tuple and found ourselves stuck with some private methods and unfriendly hooks. There is a need for redesign and good example.	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2283
20474	APEXMALHAR-2283	The hashCode() assumption isn't documented. It is also not sufficient, there should be control in the operator itself how the hash code is computed. I think we should revisit the options for tracking the state. Is there an example someone can start from for this operator? Documentation?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2283
20475	APEXMALHAR-2283	There are couple of solutions to "exactly-once". To me they are all different and they all have different assumptions. First of all, if we assume there is unique message id, we can definitely use that for dedup, but that is not always the case, then we need appid and operatorid to do dedup. Then this information can be store in either key or extra topic, I wouldn't say either of them is better than the other, really depends on how user's requirement. And no matter what we do, as long as there are number of operators writes to same kafka partition, I'm afraid there is no way to do perfect dedup because we don't know the safe place to do dedup from or too late to do dedup(too much noise from the safe place if other operator instances are fully loaded) I don't remember hashcode, but I think hashcode will include some false positive? And there is other solution like, create topic and number of partitions automatically based on kafka operator instances, in this case, it is much easier, we are always use what messages needs to be dedup because one operator only write to one kafka partition. This solution, my understanding, is most reliable and some user might want it. But the metadata of that kafka topic is kind of automatic created and it's very hard to support dynamic partition in this case. That's the reason why I say there is no general solution for exactly-once kafka output operator. We may need to provide different solutions in "examples" for people to choose from. Anyways, Sandesh, can you wrap up the current solution, post and discuss it in mailing list?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2283
20476	APEXMALHAR-2278	GitHub user ananthc opened a pull request: https://github.com/apache/apex-malhar/pull/486 APEXMALHAR-2278 Implement KuduNonTransactional Output Operator @tweise / @PramodSSImmaneni : Could you please review or suggest the right person who can complete this review for me? You can merge this pull request into a Git repository by running: $ git pull https://github.com/ananthc/apex-malhar APEXMALHAR-2278.KuduNonTransactionalOutputOperator Alternatively you can review and apply these changes as the patch at: https://github.com/apache/apex-malhar/pull/486.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #486 ---- commit af6944f53c86d0fb355798b9c48d0156964648c9 Author: ananthc <ananthg.apex@gmail.com> Date: 2016-11-08T13:28:28Z APEXMALHAR-2278 Implement KuduNonTransactional Output Operator ----	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2278
20477	APEXMALHAR-2255	Is the request to upgrade JDK or upgrade couchbase client library? It seems that it is later. Please update title accordingly and specify upgrade version. Is it backward compatible in a sense that it can be used to connect to older version of couchbase server?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2255
20478	APEXMALHAR-2251	Do you mean Direct or Heap? Majority of the code should work with ByteBuffer and not worry whether it is Direct or Heap allocated.	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2251
20479	APEXMALHAR-2244	How will the spillable DS know which time bucket to look for a given key in (get)? That's so that we don't look into multiple time buckets since it is not necessary in this case and the number of time buckets can also be high.	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2244
20480	APEXMALHAR-2223	Why delay writing to the WAL to endWindow? You get the best throughput by writing the data immediately (asynchronous), unless the same key is updated multiple times in a window. The other question is when the WAL should be flushed (blocking). Does it need to be in endWindow for idempotency or otherwise beforeCheckpoint()?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2223
20481	APEXMALHAR-2220	[~siyuan]I'm not sure where we should merge them into. Should it be a new package? Or should we put them under somewhere under "com.datatorrent.lib.stream"? Please clarify. Thanks!	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2220
20482	APEXMALHAR-2220	[~hsy541] There seems to be a circular dependency problem if we move them to org.apache.apex.malhar.lib it (malhar-library) depends on malhar-stream, however at the same time malhar-stream depends on malhar-library How about we merge them under org.apache.apex.malhar.lib in malhar-stream?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2220
20483	APEXMALHAR-2205	Is it just writing to the WAL or also writing the data files? Any information about key cardinality?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2205
20484	APEXMALHAR-2194	Not clear what the purpose is and why should it be in core? CheckpointNotificationListener is an interface which an operator can implement, begin/endWindow are already there.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2194
20485	APEXMALHAR-2193	@timothyfarkas can you please comment on this issue? Removing entries from ArrayList is expensive, should we use a different data structure?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2193
20486	APEXMALHAR-2192	GitHub user davidyan74 opened a pull request: https://github.com/apache/apex-malhar/pull/377 *Review Only* APEXMALHAR-2192 Added some spillable method implementations This is not ready for merge yet. Please do not comment on missing unit tests, code style, debug statements, etc. @brightchen @ilooner I added removeAll method to SpillableByteArrayListMultimapImpl but the testRemoveAll test method fails. Can you please take a look at what I did wrong? You can merge this pull request into a Git repository by running: $ git pull https://github.com/davidyan74/apex-malhar APEXMALHAR-2192 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/apex-malhar/pull/377.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #377 ---- commit f8168445e1f78ca48924a90a2ee0140eb85bde23 Author: David Yan <david@datatorrent.com> Date: 2016-08-22T20:15:36Z added some spillable method implementations ----	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2192

#	Requirement	Questions	Answered	Successfully Answered	URL
20487	APEXMALHAR-2184	Github user amberarrow commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/368#discussion_r74607706 --- Diff: docs/operators/fsInputOperator.md --- @@ -0,0 +1,101 @@ +File Input Operator +============== + +## Operator Objective +This operator scans a directory for files. Files are then read and split into tuples, which are emitted. The default implementation scans a single directory. The operator is fault tolerant. It tracks previously read files and current offset as part of checkpoint state. In case of failure the operator will skip files that were already processed and fast forward to the offset of the current file. Supports partitioning and changes to number of partitions. The directory scanner is responsible to only accept the files that belong to a partition. + +File Input Operator is **idempotent**, **fault-tolerant** and **partitionable**. + +## Operator Usecase +1. Read all files of a directory and then keep scanning it for newly added files. + +## Operator Information +1. Operator location: ***malhar-library*** +2. Available since: ***1.0.2*** +3. Operator state: ***Stable*** +3. Java Packages: + * Operator: ***[com.datatorrent.lib.io.fs.AbstractFileInputOperator](https://www.datatorrent.com/docs/apidocs/com/datatorrent/lib/io/fs/AbstractFileInputOperator.html)*** + +### AbstractFileInputOperator +This is the abstract implementation that serves as base class for scanning a directory for files and read the files one by one. This class doesn't have any ports. + ![AbstractFileInputOperator.png](images/fsInput/operatorsClassDiagram.png) + +## Properties, Attributes and Ports +### Properties of AbstractFileInputOperator +	**Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+	--------	----------	----	----------	----------	+	*directory*	absolute path of directory to be scanned	String	Yes	N/A	+	*scanIntervalMillis*	Interval in milliseconds after which directory should be scanned for new files	int	No	5000	+	*emitBatchSize*	Number of tuples to emit in a batch	int	No	1000	+	*partitionCount*	Desired number of partitions count	int	No	1	+	*maxRetryCount*	Maximum number of times the operator will attempt to process a file	No	5	+	*scanner*	Scanner to scan new files in directory	[DirectoryScanner](#DirectoryScanner)	No	DirectoryScanner	+ +#### Properties of DirectoryScanner +	**Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+	--------	----------	----	----------	----------	+	*filePatternRegexp*	regex to select files from input directory	String	No	N/A	+ +## Abstract Methods +void emit(T tuple): Abstract method that emits tuple read from file. +T readEntity(): Abstract method to read file entity (can be a line or block). + +## Partitioning of File Input Operator +Configure parameter "partitionCount" to define desired number of partitions. + +```xml + <property> + <name>dt.operator.{OperatorName}.prop.partitionCount</name> + <value>4</value> + </property> +``` + +where {OperatorName} is the name of the input operator. +Above lines will partition operator statically 4 times. Above value can be changed accordingly to change the number of static partitions. + +#### Dynamic Partitioning of JDBC Poller +Implement ```getNewPartitionCount()``` method to do dymanic partitioning. + +## Derived Classes +### 1. AbstractFTPInputOperator +The class is used to read files from FTP file system. +#### Properties, Attributes and Ports +--- End diff -- I don't see any ports mentioned; if there are none, please add a note saying something like: "This operator has no ports" in each case. Also, is there any reason to mention attributes since they are not operator specific ?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2184
20488	APEXMALHAR-2184	docs/operators/fsInputOperator.md --- @@ -0,0 +1,101 @@ +File Input Operator +============== + +## Operator Objective +This operator scans a directory for files. Files are then read and split into tuples, which are emitted. The default implementation scans a single directory. The operator is fault tolerant. It tracks previously read files and current offset as part of checkpoint state. In case of failure the operator will skip files that were already processed and fast forward to the offset of the current file. Supports partitioning and changes to number of partitions. The directory scanner is responsible to only accept the files that belong to a partition. + +File Input Operator is **idempotent**, **fault-tolerant** and **partitionable**. + +## Operator Usecase +1. Read all files of a directory and then keep scanning it for newly added files. + +## Operator Information +1. Operator location: ***malhar-library*** +2. Available since: ***1.0.2*** +3. Operator state: ***Stable*** +3. Java Packages: + * Operator: ***[com.datatorrent.lib.io.fs.AbstractFileInputOperator](https://www.datatorrent.com/docs/apidocs/com/datatorrent/lib/io/fs/AbstractFileInputOperator.html)*** + +### AbstractFileInputOperator +This is the abstract implementation that serves as base class for scanning a directory for files and read the files one by one. This class doesn't have any ports. + ![AbstractFileInputOperator.png](images/fsInput/operatorsClassDiagram.png) + +## Properties, Attributes and Ports +### Properties of AbstractFileInputOperator +| **Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+| --------	----------	----	----------	----------	+| *directory*	absolute path of directory to be scanned	String	Yes	N/A	+| *scanIntervalMillis*	Interval in milliseconds after which directory should be scanned for new files	int	No	5000	+| *emitBatchSize*	Number of tuples to emit in a batch	int	No	1000	+| *partitionCount*	Desired number of partitions count	int	No	1	+| *maxRetryCount*	Maximum number of times the operator will attempt to process a file	No	5	+| *scanner*	Scanner to scan new files in directory	[DirectoryScanner](#DirectoryScanner)	No	DirectoryScanner	+ +#### Properties of DirectoryScanner +| **Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+| --------	----------	----	----------	----------	+| *filePatternRegexp*	regex to select files from input directory	String	No	N/A	+ +## Abstract Methods +void emit(T tuple): Abstract method that emits tuple read from file. +T readEntity(): Abstract method to read file entity (can be a line or block). + +## Partitioning of File Input Operator +Configure parameter "partitionCount" to define desired number of partitions. + +```xml + <property> + <name>dt.operator.{OperatorName}.prop.partitionCount</name> + <value>4</value> + </property> +``` + +where {OperatorName} is the name of the input operator. +Above lines will partition operator statically 4 times. Above value can be changed accordingly to change the number of static partitions. + +#### Dynamic Partitioning of JDBC Poller +Implement ```getNewPartitionCount()``` method to do dymanic partitioning. + +## Derived Classes +### 1. AbstractFTPInputOperator +The class is used to read files from FTP file system. +#### Properties, Attributes and Ports +#### Properties of AbstractFTPInputOperator +| **Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+| --------	----------	----	----------	----------	+| *host*	Hostname of ftp server.	String	Yes	N/A	+| *port*	Port of ftp server.	int	No	21 (default ftp port)	+| *userName*	Username which is used for login to the server.	String	No	anonymous	+| *password*	Password which is used for login to the server.	String	No	gues	+ +### 2. AbstractParquetFileReader +Reads Parquet files from input directory using GroupReadSupport. Derived classes need to implement	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2184

#	Requirement	Questions	Answered	Successfully Answered	URL
20489	APEXMALHAR-2184	Github user amberarrow commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/368#discussion_r74611442 --- Diff: docs/operators/fsInputOperator.md --- @@ -0,0 +1,101 @@ +File Input Operator +============== + +## Operator Objective +This operator scans a directory for files. Files are then read and split into tuples, which are emitted. The default implementation scans a single directory. The operator is fault tolerant. It tracks previously read files and current offset as part of checkpoint state. In case of failure the operator will skip files that were already processed and fast forward to the offset of the current file. Supports partitioning and changes to number of partitions. The directory scanner is responsible to only accept the files that belong to a partition. + +File Input Operator is **idempotent**, **fault-tolerant** and **partitionable**. + +## Operator Usecase +1. Read all files of a directory and then keep scanning it for newly added files. + +## Operator Information +1. Operator location: ***malhar-library*** +2. Available since: ***1.0.2*** +3. Operator state: ***Stable*** +3. Java Packages: + * Operator: ***[com.datatorrent.lib.io.fs.AbstractFileInputOperator](https://www.datatorrent.com/docs/apidocs/com/datatorrent/lib/io/fs/AbstractFileInputOperator.html)*** + +### AbstractFileInputOperator +This is the abstract implementation that serves as base class for scanning a directory for files and read the files one by one. This class doesn't have any ports. --- End diff -- Add an overview here of what the operator does here, what parts need to be implemented by concrete subclasses and show code fragments for example from the LineByLine operator. Describe what the DirectoryScanner does. Some obvious questions that will come up in the reader's mind should be answered, for example: 1. What happens if a file that has already been processed has new data appended to it ? 2. What happens if a processed file is deleted ? 3. What happens if a separate actor is in the process of writing a new file in the monitored directory ? How will this operator know when the file is ready to be read ? 4. What if the number of files in the scanned directory grows over time to be very large -- what impact, if any, does it have on this operator ? 5. Can this operator be used to monitor multiple directories ? (Point people to the fileIO-multidir example) 6. This class already implements the Partitioner interface; can a custom partitioner be set on this operator ? More generally, since we are saying this operator supports dynamic partitioning, it is useful to describe (with a code fragment if possible) how to trigger it.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2184
20490	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r72958642 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -80,201 +93,48 @@ * @tags database, sql, jdbc, partitionable,exactlyOnce */ @Evolving -public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> - implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> implements ActivationListener<OperatorContext>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> { + private static int DEFAULT_QUEUE_CAPACITY = 4 * 1024 * 1024; + private static int DEFAULT_POLL_INTERVAL = 10 * 1000; + private static int DEFAULT_FETCH_SIZE = 20000; /** * poll interval in milliseconds */ - private static int pollInterval = 10000; + private int pollInterval = DEFAULT_POLL_INTERVAL; + private int queueCapacity = DEFAULT_QUEUE_CAPACITY; + private int fetchSize = DEFAULT_FETCH_SIZE; @Min(1) private int partitionCount = 1; protected transient int operatorId; - protected transient boolean isReplayed; - protected transient boolean isPollable; + protected boolean isPollerPartition; protected int batchSize; - protected static int fetchSize = 20000; + protected int fetchSize = 20000; /** * Map of windowId to <lower bound,upper bound> of the range key */ - protected transient MutablePair<String, String> currentWindowRecoveryState; + protected transient MutablePair<Integer, Integer> currentWindowRecoveryState; + private transient DSLContext create; - /** - * size of the emit queue used to hold polled records before emit - */ - private static int queueCapacity = 4 * 1024 * 1024; private transient volatile boolean execute; private transient AtomicReference<Throwable> cause; - private String tableName; - private String columnsExpression; - private String whereCondition = null; - private String key; - protected int spinMillis; - private transient OperatorContext context; - protected String tableName; - protected String key; - protected long currentWindowId; - protected KeyValPair<String, String> rangeQueryPair; - protected String lower; - protected String upper; - protected boolean recovered; - protected boolean isPolled; - protected String whereCondition = null; - protected String previousUpperBound; - protected String highestPolled; - private static final String user = ""; - private static final String password = ""; - /** - * thread to poll database - */ private transient Thread dbPoller; - protected transient ArrayBlockingQueue<List<T>> emitQueue; + protected KeyValPair<Integer, Integer> rangeQueryPair; + protected int lower; + protected transient boolean isPolled; + protected transient Integer lastPolledBound; + protected transient Integer lastEmittedRecord; --- End diff -- Can lower, lastPolledBound and lastEmittedRecord be int instead of Integer?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20491	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73319997 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -80,201 +93,48 @@ * @tags database, sql, jdbc, partitionable,exactlyOnce */ @Evolving -public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> - implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> implements ActivationListener<OperatorContext>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> { + private static int DEFAULT_QUEUE_CAPACITY = 4 * 1024 * 1024; + private static int DEFAULT_POLL_INTERVAL = 10 * 1000; + private static int DEFAULT_FETCH_SIZE = 20000; /** * poll interval in milliseconds */ - private static int pollInterval = 10000; + private int pollInterval = DEFAULT_POLL_INTERVAL; + private int queueCapacity = DEFAULT_QUEUE_CAPACITY; + private int fetchSize = DEFAULT_FETCH_SIZE; @Min(1) private int partitionCount = 1; protected transient int operatorId; - protected transient boolean isReplayed; - protected transient boolean isPollable; + protected boolean isPollerPartition; protected int batchSize; - protected static int fetchSize = 20000; + protected int fetchSize = 20000; /** * Map of windowId to <lower bound,upper bound> of the range key */ - protected transient MutablePair<String, String> currentWindowRecoveryState; + protected transient MutablePair<Integer, Integer> currentWindowRecoveryState; + private transient DSLContext create; - /** - * size of the emit queue used to hold polled records before emit - */ - private static int queueCapacity = 4 * 1024 * 1024; private transient volatile boolean execute; private transient AtomicReference<Throwable> cause; - private String tableName; - private String columnsExpression; - private String whereCondition = null; - private String key; - protected int spinMillis; - private transient OperatorContext context; - protected String tableName; - protected String key; - protected long currentWindowId; - protected KeyValPair<String, String> rangeQueryPair; - protected String lower; - protected String upper; - protected boolean recovered; - protected boolean isPolled; - protected String whereCondition = null; - protected String previousUpperBound; - protected String highestPolled; - private static final String user = ""; - private static final String password = ""; - /** - * thread to poll database - */ private transient Thread dbPoller; - protected transient ArrayBlockingQueue<List<T>> emitQueue; + protected KeyValPair<Integer, Integer> rangeQueryPair; + protected int lower; + protected transient boolean isPolled; + protected transient Integer lastPolledBound; + protected transient Integer lastEmittedRecord; --- End diff -- Can it be renamed to something like ```columnsListStr``` ?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172

#	Requirement	Questions	Answered	Successfully Answered	URL
20492	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73320816 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -80,201 +93,48 @@ * @tags database, sql, jdbc, partitionable,exactlyOnce */ @Evolving -public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> - implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> implements + ActivationListener<OperatorContext>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> { + private static int DEFAULT_QUEUE_CAPACITY = 4 * 1024 * 1024; + private static int DEFAULT_POLL_INTERVAL = 10 * 1000; + private static int DEFAULT_FETCH_SIZE = 20000; /** * poll interval in milliseconds */ - private static int pollInterval = 10000; + private int pollInterval = DEFAULT_POLL_INTERVAL; + private int queueCapacity = DEFAULT_QUEUE_CAPACITY; + private int fetchSize = DEFAULT_FETCH_SIZE; @Min(1) private int partitionCount = 1; - protected transient int operatorId; - protected transient boolean isReplayed; - protected transient boolean isPollable; + protected boolean isPollerPartition; + protected int batchSize; + protected static int fetchSize = 20000; + /** * Map of windowId to <lower bound,upper bound> of the range key */ - protected transient MutablePair<String, String> currentWindowRecoveryState; + protected transient MutablePair<Integer, Integer> currentWindowRecoveryState; + private transient DSLContext create; - /** - * size of the emit queue used to hold polled records before emit - */ - private static int queueCapacity = 4 * 1024 * 1024; + private transient volatile boolean execute; + private transient AtomicReference<Throwable> cause; + private String tableName; + private String columnsExpression; + private String whereCondition = null; + private String key; + protected transient int spinMillis; + private transient OperatorContext context; + protected String tableName; + protected String key; protected long currentWindowId; - protected KeyValPair<String, String> rangeQueryPair; + protected KeyValPair<Integer, Integer> rangeQueryPair; protected String lower; protected String upper; protected boolean recovered; protected boolean isPolled; - protected String whereCondition = null; protected String previousUpperBound; protected String highestPolled; - private static final String user = ""; - private static final String password = ""; - /** - * thread to poll database - */ - private transient Thread dbPoller; - protected transient ArrayBlockingQueue<List<T>> emitQueue; + protected KeyValPair<Integer, Integer> rangeQueryPair; protected Integer lower; protected transient boolean isPolled; protected transient Integer lastPolledBound; --- End diff -- can this actually be called ```upperBound```?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20493	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73321098 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -80,201 +93,48 @@ * @tags database, sql, jdbc, partitionable,exactlyOnce */ @Evolving -public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> - implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> implements + ActivationListener<OperatorContext>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> { + private static int DEFAULT_QUEUE_CAPACITY = 4 * 1024 * 1024; + private static int DEFAULT_POLL_INTERVAL = 10 * 1000; + private static int DEFAULT_FETCH_SIZE = 20000; /** * poll interval in milliseconds */ - private static int pollInterval = 10000; + private int pollInterval = DEFAULT_POLL_INTERVAL; + private int queueCapacity = DEFAULT_QUEUE_CAPACITY; + private int fetchSize = DEFAULT_FETCH_SIZE; @Min(1) private int partitionCount = 1; - protected transient int operatorId; - protected transient boolean isReplayed; - protected transient boolean isPollable; + protected boolean isPollerPartition; + protected int batchSize; + protected static int fetchSize = 20000; + /** * Map of windowId to <lower bound,upper bound> of the range key */ - protected transient MutablePair<String, String> currentWindowRecoveryState; + protected transient MutablePair<Integer, Integer> currentWindowRecoveryState; + private transient DSLContext create; - /** - * size of the emit queue used to hold polled records before emit - */ - private static int queueCapacity = 4 * 1024 * 1024; + private transient volatile boolean execute; + private transient AtomicReference<Throwable> cause; + private String tableName; + private String columnsExpression; + private String whereCondition = null; + private String key; + protected transient int spinMillis; + private transient OperatorContext context; + protected String tableName; + protected String key; protected long currentWindowId; - protected KeyValPair<String, String> rangeQueryPair; + protected KeyValPair<Integer, Integer> rangeQueryPair; protected String lower; protected String upper; protected boolean recovered; protected boolean isPolled; - protected String whereCondition = null; protected String previousUpperBound; protected String highestPolled; - private static final String user = ""; - private static final String password = ""; - /** - * thread to poll database - */ - private transient Thread dbPoller; - protected transient ArrayBlockingQueue<List<T>> emitQueue; + protected KeyValPair<Integer, Integer> rangeQueryPair; protected Integer lower; protected transient boolean isPolled; protected transient Integer lastPolledBound; + protected transient Integer lastEmittedRecord; + protected transient ScheduledExecutorService scanService; + protected transient LinkedBlockingDeque<T> emitQueue; + protected transient PreparedStatement ps; protected WindowDataManager windowManager; --- End diff -- Perhaps a lot of the above variables can be private. Can you check which ones can be made private?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172

#	Requirement	Questions	Answered	Successfully Answered	URL
20494	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73326772 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -286,149 +146,119 @@ public AbstractJdbcPollInputOperator() public void setup(OperatorContext context) { super.setup(context); + intializeDSLContext(); + if (scanService == null) { + scanService = Executors.newScheduledThreadPool(partitionCount); + } spinMillis = context.getValue(Context.OperatorContext.SPIN_MILLIS); execute = true; cause = new AtomicReference<Throwable>(); - emitQueue = new ArrayBlockingQueue<List<T>>(queueCapacity); - this.context = context; + emitQueue = new LinkedBlockingDeque<>(queueCapacity); operatorId = context.getId(); + windowManager.setup(context); + } - try { + private void intializeDSLContext() + { + create = DSL.using(store.getConnection(), JDBCUtils.dialect(store.getDatabaseUrl())); + } - //If its a range query pass upper and lower bounds - //If its a polling query pass only the lower bound - if (getRangeQueryPair().getValue() != null) { - ps = store.getConnection() - .prepareStatement(- JdbcMetaDataUtility.buildRangeQuery(getTableName(), getKey(), rangeQueryPair.getKey(), - rangeQueryPair.getValue()), java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + @Override + public void activate(OperatorContext context) + { + initializePreparedStatement(); + long largestRecoveryWindow = windowManager.getLargestRecoveryWindow(); + if (largestRecoveryWindow == Stateless.WINDOW_ID + || context.getValue(Context.OperatorContext.ACTIVATION_WINDOW_ID) > largestRecoveryWindow) { + scanService.scheduleAtFixedRate(new DBPoller(), 0, pollInterval, TimeUnit.MILLISECONDS); + } + } + protected void initializePreparedStatement() + { + try { + // If its a range query pass upper and lower bounds, If its a polling query pass only the lower bound + if (isPollerPartition) { + ps = store.getConnection().prepareStatement(buildRangeQuery(rangeQueryPair.getKey(), Integer.MAX_VALUE), --- End diff -- Are you expecting rangeQueryPair form the user? If so, mark it ```@NotNull```. This is not being set anywhere..	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20495	APEXMALHAR-2172	Github user DT-Priyanka commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73647661 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -80,201 +93,48 @@ * @tags database, sql, jdbc, partitionable,exactlyOnce */ @Evolving -public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> - implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> implements ActivationListener<OperatorContext>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> { - private static final int DEFAULT_QUEUE_CAPACITY = 4 * 1024 * 1024; - private static final int DEFAULT_POLL_INTERVAL = 10 * 1000; - private static final int DEFAULT_FETCH_SIZE = 20000; - /** - * poll interval in milliseconds - */ - private int pollInterval = 10000; - private int queueCapacity = DEFAULT_QUEUE_CAPACITY; - private int fetchSize = DEFAULT_FETCH_SIZE; - @Min(1) - private int partitionCount = 1; - protected transient int operatorId; - protected transient boolean isReplayed; - protected transient boolean isPollable; - protected boolean isPollerPartition; - protected int batchSize; --- End diff -- batchSize is optional, let me add some default value to it. What's the suggestion?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20496	APEXMALHAR-2172	Github user DT-Priyanka commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73647799 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -80,201 +93,48 @@ * @tags database, sql, jdbc, partitionable,exactlyOnce */ @Evolving -public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> - implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> implements ActivationListener<OperatorContext>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> { - private static final int DEFAULT_QUEUE_CAPACITY = 4 * 1024 * 1024; - private static final int DEFAULT_POLL_INTERVAL = 10 * 1000; - private static final int DEFAULT_FETCH_SIZE = 20000; - /** - * poll interval in milliseconds - */ - private int pollInterval = 10000; - private int queueCapacity = DEFAULT_QUEUE_CAPACITY; - private int fetchSize = DEFAULT_FETCH_SIZE; - @Min(1) - private int partitionCount = 1; - protected transient int operatorId; - protected transient boolean isReplayed; - protected transient boolean isPollable; - protected boolean isPollerPartition; - protected int batchSize; - protected int fetchSize = 20000; - /** - * Map of windowId to <lower bound,upper bound> of the range key - */ - protected transient MutablePair<String, String> currentWindowRecoveryState; + protected transient DSLContext create; - private transient DSLContext create; - /** - * size of the emit queue used to hold polled records before emit - */ - private int queueCapacity = 4 * 1024 * 1024; - private transient volatile boolean execute; - private transient AtomicReference<Throwable> cause; + protected String tableName; + protected String columnsExpression; --- End diff -- should we be consistent with already present JDBCPojoInputOperator? People can use same set of properties with both operators.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172

#	Requirement	Questions	Answered	Successfully Answered	URL
20497	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73826686 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -565,67 +379,110 @@ public void emitReplayedTuples(PreparedStatement ps) */ @Override public Collection<com.datatorrent.api.Partitioner.Partition<AbstractJdbcPollInputOperator<T>>> definePartitions(- Collection<com.datatorrent.api.Partitioner.Partition<AbstractJdbcPollInputOperator<T>>> partitions, - com.datatorrent.api.Partitioner.PartitioningContext context) + Collection<Partition<AbstractJdbcPollInputOperator<T>>> partitions, PartitioningContext context) { List<Partition<AbstractJdbcPollInputOperator<T>>> newPartitions = new ArrayList<Partition<AbstractJdbcPollInputOperator<T>>>(getPartitionCount()); - JdbcStore jdbcStore = new JdbcStore(); - jdbcStore.setDatabaseDriver(store.getDatabaseDriver()); - jdbcStore.setDatabaseUrl(store.getDatabaseUrl()); - jdbcStore.setConnectionProperties(store.getConnectionProperties()); - jdbcStore.connect(); - - HashMap<Integer, KeyValPair<String, String>> partitionToRangeMap = null; + HashMap<Integer, KeyValPair<Integer, Integer>> partitionToRangeMap = null; try { - partitionToRangeMap = JdbcMetaDataUtility.getPartitionedQueryMap(getPartitionCount(), - jdbcStore.getDatabaseDriver(), jdbcStore.getDatabaseUrl(), getTableName(), getKey(), - store.getConnectionProperties().getProperty(user), store.getConnectionProperties().getProperty(password), - whereCondition, emitColumnList); + store.connect(); + intializeDSLContext(); + partitionToRangeMap = getPartitionedQueryRangeMap(getPartitionCount()); } catch (SQLException e) { LOG.error("Exception in initializing the partition range", e); + throw new RuntimeException(e); + } finally { + store.disconnect(); } KryoCloneUtils<AbstractJdbcPollInputOperator<T>> cloneUtils = KryoCloneUtils.createCloneUtils(this); + // The n given partitions are for range queries and n + 1 partition is for polling query for (int i = 0; i <= getPartitionCount(); i++) { - AbstractJdbcPollInputOperator<T> jdbcPoller = null; - - jdbcPoller = cloneUtils.getClone(); - - jdbcPoller.setStore(store); - jdbcPoller.setKey(getKey()); - jdbcPoller.setPartitionCount(getPartitionCount()); - jdbcPoller.setPollInterval(getPollInterval()); - jdbcPoller.setTableName(getTableName()); - jdbcPoller.setBatchSize(getBatchSize()); - jdbcPoller.setEmitColumnList(getEmitColumnList()); - - store.connect(); - //The n given partitions are for range queries and n + 1 partition is for polling query - //The upper bound for the n+1 partition is set to null since its a pollable partition + AbstractJdbcPollInputOperator<T> jdbcPoller = cloneUtils.getClone(); if (i < getPartitionCount()) { - jdbcPoller.setRangeQueryPair(partitionToRangeMap.get(i)); - isPollable = false; + jdbcPoller.rangeQueryPair = partitionToRangeMap.get(i); + jdbcPoller.lastEmittedRecord = partitionToRangeMap.get(i).getKey(); --- End diff -- Set ```jdbcPoller.isPollerPartition = false``` here?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20498	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73830019 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -438,119 +261,110 @@ public void endWindow() currentWindowRecoveryState = new MutablePair<>(); } - public int getPartitionCount() { - return partitionCount; - } - - public void setPartitionCount(int partitionCount) { + public void deactivate() { - this.partitionCount = partitionCount; + scanService.shutdownNow(); + store.disconnect(); } - @Override - public void activate(Context cntxt) { + protected void pollRecords() { - if (context.getValue(OperatorContext.ACTIVATION_WINDOW_ID) != Stateless.WINDOW_ID - && context.getValue(OperatorContext.ACTIVATION_WINDOW_ID) < windowManager.getLargestRecoveryWindow()) { - // If it is a replay state, don't start any threads here + if (isPolled) { return; } - } - - @Override - public void deactivate() { - try { - if (dbPoller != null && dbPoller.isAlive()) { - dbPoller.interrupt(); - dbPoller.join(); + ps.setFetchSize(getFetchSize()); + ResultSet result = ps.executeQuery(); + if (result.next()) { + do { + emitQueue.add(getTuple(result)); + } while (result.next()); } - } catch (InterruptedException ex) { - // log and ignore, ending execution anyway - LOG.error("exception in poller thread: ", ex); + isPolled = true; + } catch (SQLException ex) { + throw new RuntimeException(String.format("Error while running query"), ex); + } finally { + store.disconnect(); } - } - - @Override - public void handleIdleTime() { - if (execute) { - try { - Thread.sleep(spinMillis); - } catch (InterruptedException ie) { - throw new RuntimeException(ie); - } - } else { - LOG.error("Exception: ", cause); - DTThrowable.rethrow(cause.get()); - } + public abstract T getTuple(ResultSet result); + protected void replay(long windowId) throws SQLException { - isReplayed = true; - MutablePair<String, String> recoveredData = new MutablePair<String, String>(); try { - recoveredData = (MutablePair<String, String>)windowManager.load(operatorId, windowId); + MutablePair<Integer, Integer> recoveredData = (MutablePair<Integer, Integer>)windowManager.load(operatorId, windowId); - if (recoveredData != null) { - //skip the window and return if there was no incoming data in the window - if (recoveredData.left == null || recoveredData.right == null) { - return; - } - if (recoveredData.right.equals(rangeQueryPair.getValue()) || recoveredData.right.equals(previousUpperBound)) { - LOG.info("Matched so returning"); - return; - } + if (recoveredData != null && shouldReplayWindow(recoveredData)) { + LOG.debug("[Recovering Window ID - {} for record range: {}, {}]", windowId, recoveredData.left, recoveredData.right); - JdbcPollInputOperator jdbcPoller = new JdbcPollInputOperator(); - jdbcPoller.setStore(store); - jdbcPoller.setKey(getKey()); - jdbcPoller.setPartitionCount(getPartitionCount()); - jdbcPoller.setPollInterval(getPollInterval()); - jdbcPoller.setTableName(getTableName()); - jdbcPoller.setBatchSize(getBatchSize()); - isPollable = false; - LOG.debug("[Window ID -" + windowId + "," + recoveredData.left + "," + recoveredData.right + "]"); + ps = store.getConnection().prepareStatement(buildRangeQuery(rangeQueryPair.getKey(), (rangeQueryPair.getValue() - rangeQueryPair.getKey())), --- End diff -- Should this fetch data from ```recoveredData.left``` to ```recoveredData.right```?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172

#	Requirement	Questions	Answered	Successfully Answered	URL
20499	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r73840151 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -438,119 +261,110 @@ public void endWindow() currentWindowRecoveryState = new MutablePair<>(); } - public int getPartitionCount() { - return partitionCount; } - public void setPartitionCount(int partitionCount) + @Override + public void deactivate() { - this.partitionCount = partitionCount; + scanService.shutdownNow(); + store.disconnect(); } - @Override - public void activate(Context cntxt) + protected void pollRecords() { - if (context.getValue(OperatorContext.ACTIVATION_WINDOW_ID) != Stateless.WINDOW_ID - && context.getValue(OperatorContext.ACTIVATION_WINDOW_ID) < windowManager.getLargestRecoveryWindow()) { - // If it is a replay state, don't start any threads here + if (isPolled) { return; } - } - - @Override - public void deactivate() - { try { - if (dbPoller != null && dbPoller.isAlive()) { - dbPoller.interrupt(); - dbPoller.join(); + ps.setFetchSize(getFetchSize()); + ResultSet result = ps.executeQuery(); + if (result.next()) { + do { + emitQueue.add(getTuple(result)); + } while (result.next()); } - } catch (InterruptedException ex) { - // log and ignore, ending execution anyway - LOG.error("exception in poller thread: ", ex); + isPolled = true; --- End diff -- Does this mean that we want to poll records just once in a window? Is this because the poll interval is sufficiently long so as to prevent multiple calls to ```pollRecords()``` within a single window? Even in that case, if the poll interval is reduced, there might be multiple calls where only one call will happen in this case.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20500	APEXMALHAR-2172	library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -286,149 +141,117 @@ public AbstractJdbcPollInputOperator() public void setup(OperatorContext context) { super.setup(context); - spinMillis = context.getValue(Context.OperatorContext.SPIN_MILLIS); + intializeDSLContext(); + if (scanService == null) { + scanService = Executors.newScheduledThreadPool(1); + } execute = true; - cause = new AtomicReference<Throwable>(); - emitQueue = new ArrayBlockingQueue<List<T>>(queueCapacity); - this.context = context; + emitQueue = new LinkedBlockingDeque<>(queueCapacity); operatorId = context.getId(); + windowManager.setup(context); + } - try { + private void intializeDSLContext() { + create = DSL.using(store.getConnection(), JDBCUtils.dialect(store.getDatabaseUrl())); + } - //If its a range query pass upper and lower bounds - //If its a polling query pass only the lower bound - if (getRangeQueryPair().getValue() != null) { - ps = store.getConnection() - .prepareStatement(- JdbcMetaDataUtility.buildRangeQuery(getTableName(), getKey(), rangeQueryPair.getKey(), - rangeQueryPair.getValue()), - java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + @Override + public void activate(OperatorContext context) + { + intializePreparedStatement(); + long largestRecoveryWindow = windowManager.getLargestRecoveryWindow(); + if (largestRecoveryWindow == Stateless.WINDOW_ID + || context.getValue(Context.OperatorContext.ACTIVATION_WINDOW_ID) > largestRecoveryWindow) { + scanService.scheduleAtFixedRate(new DBPoller(), 0, pollInterval, TimeUnit.MILLISECONDS); + } + } + + protected void intializePreparedStatement() + { + try { + // If its a range query pass upper and lower bounds, If its a polling query pass only the lower bound + if (isPollerPartition) { + ps = store.getConnection().prepareStatement(buildRangeQuery(rangeQueryPair.getKey(), Integer.MAX_VALUE), + TYPE_FORWARD_ONLY, CONCUR_READ_ONLY); + } else { + ps = store.getConnection().prepareStatement(+ JdbcMetaDataUtility.buildPollableQuery(getTableName(), getKey(), rangeQueryPair.getKey()), + java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + isPollable = true; + buildRangeQuery(rangeQueryPair.getKey(), + (rangeQueryPair.getValue() - rangeQueryPair.getKey())), + TYPE_FORWARD_ONLY, CONCUR_READ_ONLY); } - } catch (SQLException e) { + } catch (SQLException e) { LOG.error("Exception in initializing the range query for a given partition", e); throw new RuntimeException(e); } - windowManager.setup(context); - LOG.debug("super setup done..."); } @Override public void beginWindow(long windowId) { currentWindowId = windowId; - isReplayed = false; - if (currentWindowId <= windowManager.getLargestRecoveryWindow()) { try { replay(currentWindowId); return; } catch (SQLException e) { LOG.error("Exception in replayed windows", e); throw new RuntimeException(e); } } - if (isReplayed && currentWindowId == windowManager.getLargestRecoveryWindow()) { try { if (!isPollable && rangeQueryPair.getValue() != null) { ps = store.getConnection().prepareStatement(JdbcMetaDataUtility.buildGTRangeQuery(getTableName(), getKey(), previousUpperBound, rangeQueryPair.getValue()), java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); } else { String bound = null; if (previousUpperBound == null) { bound = getRangeQueryPair().getKey(); } else { bound = previousUpperBound; } ps = store.getConnection().prepareStatement(JdbcMetaDataUtility.buildPollableQuery(getTableName(), getKey(), bound), java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); isPollable = true; } isReplayed = false;	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20501	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r74004081 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -438,119 +261,110 @@ public void endWindow() currentWindowRecoveryState = new MutablePair<>(); } - public int getPartitionCount() { - return partitionCount; } - public void setPartitionCount(int partitionCount) + @Override + public void deactivate() { - this.partitionCount = partitionCount; + scanService.shutdownNow(); + store.disconnect(); } - @Override - public void activate(Context cntxt) + protected void pollRecords() { - if (context.getValue(OperatorContext.ACTIVATION_WINDOW_ID) != Stateless.WINDOW_ID - && context.getValue(OperatorContext.ACTIVATION_WINDOW_ID) < windowManager.getLargestRecoveryWindow()) { - // If it is a replay state, don't start any threads here + if (isPolled) { return; } - } - - @Override - public void deactivate() - { try { - if (dbPoller != null && dbPoller.isAlive()) { - dbPoller.interrupt(); - dbPoller.join(); + ps.setFetchSize(getFetchSize()); + ResultSet result = ps.executeQuery(); + if (result.next()) { + do { + emitQueue.add(getTuple(result)); + } while (result.next()); } - } catch (InterruptedException ex) { - // log and ignore, ending execution anyway - LOG.error("exception in poller thread: ", ex); + isPolled = true; --- End diff -- What about non-poller partitions? ```pollRecords()``` is just fetching ```fetchSize``` records and making ```isPolled = true```. This way non-poller partitions only fetch ```fetchSize``` records.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172

#	Requirement	Questions	Answered	Successfully Answered	URL
20502	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r74005794 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/JdbcPollInputOperator.java --- @@ -18,214 +18,75 @@ */ package com.datatorrent.lib.db.jdbc; - import java.sql.PreparedStatement; import java.sql.ResultSet; import java.sql.SQLException; import java.util.ArrayList; -import java.util.List; - -import org.slf4j.Logger; -import org.slf4j.LoggerFactory; import org.apache.hadoop.classification.InterfaceStability.Evolving; - import com.google.common.collect.Lists; - import com.datatorrent.api.Context.OperatorContext; -import com.datatorrent.api.annotation.OperatorAnnotation; +import com.datatorrent.api.DefaultOutputPort; import com.datatorrent.api.annotation.OutputPortFieldAnnotation; - /** - * A concrete implementation for {@link AbstractJdbcPollInputOperator}} for - * consuming data from MySQL using JDBC interface --- End diff -- Can we remove explicit mention for MySQL? I think it is generic as of now..	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20503	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r74005996 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/JdbcPOJOPollInputOperator.java --- @@ -0,0 +1,321 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.db.jdbc; + +import java.math.BigDecimal; +import java.sql.Date; +import java.sql.PreparedStatement; +import java.sql.ResultSet; +import java.sql.ResultSetMetaData; +import java.sql.SQLException; +import java.sql.Time; +import java.sql.Timestamp; +import java.sql.Types; +import java.util.List; +import java.util.Map; + +import javax.validation.constraints.NotNull; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.hadoop.classification.InterfaceStability.Evolving; + +import com.google.common.collect.Lists; +import com.google.common.collect.Maps; + +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.db.jdbc.JdbcPOJOInputOperator.ActiveFieldInfo; +import com.datatorrent.lib.util.FieldInfo; +import com.datatorrent.lib.util.PojoUtils; + +/** + * A concrete implementation for {@link AbstractJdbcPollInputOperator} for --- End diff -- Remove reference to MySQL?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20504	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r74013531 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/JdbcPOJOPollInputOperator.java --- @@ -0,0 +1,321 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.db.jdbc; + +import java.math.BigDecimal; +import java.sql.Date; +import java.sql.PreparedStatement; +import java.sql.ResultSet; +import java.sql.ResultSetMetaData; +import java.sql.SQLException; +import java.sql.Time; +import java.sql.Timestamp; +import java.sql.Types; +import java.util.List; +import java.util.Map; + +import javax.validation.constraints.NotNull; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.hadoop.classification.InterfaceStability.Evolving; + +import com.google.common.collect.Lists; +import com.google.common.collect.Maps; + +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.db.jdbc.JdbcPOJOInputOperator.ActiveFieldInfo; +import com.datatorrent.lib.util.FieldInfo; +import com.datatorrent.lib.util.PojoUtils; + +/** + * A concrete implementation for {@link AbstractJdbcPollInputOperator} for + * consuming data from MySQL using JDBC interface + * User needs to provide tableName,dbConnection,setEmitColumnList,look-up key + * Optionally batchSize,pollInterval,Look-up key and a where clause can be given + * This operator uses static partitioning to arrive at range queries for exactly + * once reads + * Assumption is that there is an ordered column using which range queries can + * be formed + * If an emitColumnList is provided, please ensure that the keyColumn is the + * first column in the list + * Range queries are formed using the {@link JdbcMetaDataUtility} Output - comma + * separated list of the emit columns eg columnA,columnB,columnC + * + * @displayName Jdbc Polling Input Operator + * @category Input + * @tags database, sql, jdbc + */ +@Evolving +public class JdbcPOJOPollInputOperator extends AbstractJdbcPollInputOperator<Object> +{ + private final transient List<ActiveFieldInfo> columnFieldSetters = Lists.newArrayList(); + protected List<Integer> columnDataTypes; + protected transient Class<?> pojoClass; + @NotNull + private List<FieldInfo> fieldInfos; --- End diff -- For consistency with Jdbc output operator, can you please use ```JdbcFieldInfo``` instead?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172

#	Requirement	Questions	Answered	Successfully Answered	URL
20505	APEXMALHAR-2172	library/src/test/java/com/datatorrent/lib/db/jdbc/JdbcPojoPollableOpeartorTest.java --- @@ -0,0 +1,251 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.db.jdbc; + +import java.io.File; +import java.io.IOException; +import java.sql.Date; +import java.util.Collection; +import java.util.Iterator; +import java.util.List; +import java.util.concurrent.ScheduledExecutorService; +import java.util.concurrent.ScheduledFuture; +import java.util.concurrent.TimeUnit; +import org.junit.After; +import org.junit.Assert; +import org.junit.Before; +import org.junit.Test; +import org.mockito.Mock; +import org.mockito.MockitoAnnotations; +import org.apache.apex.malhar.lib.wal.WindowDataManager; +import org.apache.commons.io.FileUtils; +import org.apache.commons.lang3.tuple.MutablePair; +import com.google.common.collect.Lists; +import com.datatorrent.api.Attribute; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultPartition; +import com.datatorrent.api.Partitioner; +import com.datatorrent.lib.helper.OperatorContextTestHelper; +import com.datatorrent.lib.helper.TestPortContext; +import com.datatorrent.lib.testbench.CollectorTestSink; +import com.datatorrent.lib.util.FieldInfo; +import com.datatorrent.lib.util.KeyValPair; +import com.datatorrent.lib.util.TestUtils; +import static org.mockito.Matchers.any; +import static org.mockito.Matchers.anyLong; +import static org.mockito.Mockito.times; +import static org.mockito.Mockito.verify; +import static org.mockito.Mockito.when; +public class JdbcPojoPollableOpeartorTest extends JdbcOperatorTest +{ + public String dir = null; + @Mock + private ScheduledExecutorService mockschedular; + @Mock + private ScheduledFuture futureTaskMock; + @Mock + private WindowDataManager windowDataManagerMock; + + @Before + public void beforeTest() + { + dir = "target/" + APP_ID + "/"; + MockitoAnnotations.initMocks(this); + when(mockschedular.scheduleWithFixedDelay(any(Runnable.class), anyLong(), anyLong(), any(TimeUnit.class))) + .thenReturn(futureTaskMock); + } + @After + public void afterTest() throws IOException + { + cleanTable(); + FileUtils.deleteDirectory(new File(dir)); + } + @Test + public void testDBPoller() throws InterruptedException + { + insertEvents(10, true, 0); + JdbcStore store = new JdbcStore(); + store.setDatabaseDriver(DB_DRIVER); + store.setDatabaseUrl(URL); + List<FieldInfo> fieldInfos = getFieldInfos(); + Attribute.AttributeMap.DefaultAttributeMap portAttributes = new Attribute.AttributeMap.DefaultAttributeMap(); + portAttributes.put(Context.PortContext.TUPLE_CLASS, TestPOJOEvent.class); + TestPortContext tpc = new TestPortContext(portAttributes); + JdbcPOJOPollInputOperator inputOperator = new JdbcPOJOPollInputOperator(); + inputOperator.setStore(store); + inputOperator.setTableName(TABLE_POJO_NAME); + inputOperator.setKey("id"); + inputOperator.setFieldInfos(fieldInfos); + inputOperator.setFetchSize(100); + inputOperator.setBatchSize(100); + inputOperator.setPartitionCount(2); + TestUtils.MockBatchedOperatorStats	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20506	APEXMALHAR-2172	library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -286,271 +141,238 @@ public AbstractJdbcPollInputOperator() public void setup(OperatorContext context) { super.setup(context); - spinMillis = context.getValue(Context.OperatorContext.SPIN_MILLIS); + intializeDSLContext(); + if (scanService == null) { + scanService = Executors.newScheduledThreadPool(1); + } execute = true; - cause = new AtomicReference<Throwable>(); - emitQueue = new ArrayBlockingQueue<List<T>>(queueCapacity); - this.context = context; + emitQueue = new LinkedBlockingDeque<>(queueCapacity); operatorId = context.getId(); + windowManager.setup(context); + } - try { + private void intializeDSLContext() + { + create = DSL.using(store.getConnection(), JDBCUtils.dialect(store.getDatabaseUrl())); + } + @Override + public void activate(OperatorContext context) + { + initializePreparedStatement(); + long largestRecoveryWindow = windowManager.getLargestRecoveryWindow(); + if (largestRecoveryWindow == Stateless.WINDOW_ID || context.getValue(Context.OperatorContext.ACTIVATION_WINDOW_ID) > largestRecoveryWindow) { + scanService.scheduleAtFixedRate(new DBPoller(), 0, pollInterval, TimeUnit.MILLISECONDS); + } + } - //If its a range query pass upper and lower bounds - //If its a polling query pass only the lower bound - if (getRangeQueryPair().getValue() != null) { - ps = store.getConnection() - .prepareStatement(- JdbcMetaDataUtility.buildRangeQuery(getTableName(), getKey(), rangeQueryPair.getKey(), - rangeQueryPair.getValue()), - java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + protected void initializePreparedStatement() + { + try { + // If its a range query pass upper and lower bounds, If its a polling query pass only the lower bound + if (isPollerPartition) { + ps = store.getConnection().prepareStatement(buildRangeQuery(rangeQueryPair.getKey(), Integer.MAX_VALUE), + TYPE_FORWARD_ONLY, CONCUR_READ_ONLY); + } else { + ps = store.getConnection().prepareStatement(+ JdbcMetaDataUtility.buildPollableQuery(getTableName(), getKey(), rangeQueryPair.getKey()), + java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + isPollable = true; + buildRangeQuery(rangeQueryPair.getKey(), + (rangeQueryPair.getValue() - rangeQueryPair.getKey())), + TYPE_FORWARD_ONLY, CONCUR_READ_ONLY); + } + } catch (SQLException e) { + LOG.error("Exception in initializing the range query for a given partition", e); + throw new RuntimeException(e); + } + windowManager.setup(context); + LOG.debug("super setup done..."); + } + @Override + public void beginWindow(long windowId) + { + currentWindowId = windowId; + isReplayed = false; + if (currentWindowId <= windowManager.getLargestRecoveryWindow()) { + try { + replay(currentWindowId); + return; + } catch (SQLException e) { + LOG.error("Exception in replayed windows", e); + throw new RuntimeException(e); + } + } + if (isReplayed && + currentWindowId == windowManager.getLargestRecoveryWindow()) { + try { + if (!isPollable && rangeQueryPair.getValue() != null) { + ps = store.getConnection().prepareStatement(+ JdbcMetaDataUtility.buildGTRangeQuery(getTableName(), getKey(), previousUpperBound, + rangeQueryPair.getValue()), + java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + } else { + String bound = null; + if (previousUpperBound == null) { + bound = getRangeQueryPair().getKey(); + } else { + bound = previousUpperBound; + } + ps = store.getConnection().prepareStatement(+ JdbcMetaDataUtility.buildPollableQuery(getTableName(), getKey(), bound), + java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + isPollable = true; + } + isReplayed = false;	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172

#	Requirement	Questions	Answered	Successfully Answered	URL
20507	APEXMALHAR-2172	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/358#discussion_r74187632 --- Diff: library/src/main/java/com/datatorrent/lib/db/jdbc/JdbcPOJOPollInputOperator.java --- @@ -0,0 +1,321 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.db.jdbc; + +import java.math.BigDecimal; +import java.sql.Date; +import java.sql.PreparedStatement; +import java.sql.ResultSet; +import java.sql.ResultSetMetaData; +import java.sql.SQLException; +import java.sql.Time; +import java.sql.Timestamp; +import java.sql.Types; +import java.util.List; +import java.util.Map; + +import javax.validation.constraints.NotNull; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.hadoop.classification.InterfaceStability.Evolving; + +import com.google.common.collect.Lists; +import com.google.common.collect.Maps; + +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.db.jdbc.JdbcPOJOInputOperator.ActiveFieldInfo; +import com.datatorrent.lib.util.FieldInfo; +import com.datatorrent.lib.util.PojoUtils; + +/** + * A concrete implementation for {@link AbstractJdbcPollInputOperator} for + * consuming data from MySQL using JDBC interface + * User needs to provide tableName,dbConnection,setEmitColumnList,look-up key + * Optionally batchSize,pollInterval,Look-up key and a where clause can be given + * This operator uses static partitioning to arrive at range queries for exactly + * once reads + * Assumption is that there is an ordered column using which range queries can + * be formed + * If an emitColumnList is provided, please ensure that the keyColumn is the + * first column in the list + * Range queries are formed using the {@link JdbcMetaDataUtility} Output - comma + * separated list of the emit columns eg columnA,columnB,columnC + * + * @displayName Jdbc Polling Input Operator + * @category Input + * @tags database, sql, jdbc + */ +@Evolving +public class JdbcPOJOPollInputOperator extends AbstractJdbcPollInputOperator<Object> +{ + private final transient List<ActiveFieldInfo> columnFieldSetters = Lists.newArrayList(); + protected List<Integer> columnDataTypes; + protected transient Class<?> pojoClass; + @NotNull + private List<FieldInfo> fieldInfos; --- End diff -- Can you please create a JIRA for this?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20508	APEXMALHAR-2172	library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -286,271 +141,238 @@ public AbstractJdbcPollInputOperator() public void setup(OperatorContext context) { super.setup(context); - spinMillis = context.getValue(Context.OperatorContext.SPIN_MILLIS); + intializeDSLContext(); + if (scanService == null) { + scanService = Executors.newScheduledThreadPool(1); + } execute = true; - cause = new AtomicReference<Throwable>(); - emitQueue = new ArrayBlockingQueue<List<T>>(queueCapacity); - this.context = context; + emitQueue = new LinkedBlockingDeque<>(queueCapacity); operatorId = context.getId(); + windowManager.setup(context); + } - try { + private void intializeDSLContext() { + create = DSL.using(store.getConnection(), JDBCUtils.dialect(store.getDatabaseUrl())); + } + + @Override + public void activate(OperatorContext context) { + initializePreparedStatement(); + long largestRecoveryWindow = windowManager.getLargestRecoveryWindow(); + if (largestRecoveryWindow == Stateless.WINDOW_ID || context.getValue(Context.OperatorContext.ACTIVATION_WINDOW_ID) > largestRecoveryWindow) { + scanService.scheduleAtFixedRate(new DBPoller(), 0, pollInterval, TimeUnit.MILLISECONDS); + } + } - //If its a range query pass upper and lower bounds - //If its a polling query pass only the lower bound - if (getRangeQueryPair().getValue() != null) { - ps = store.getConnection() - .prepareStatement(- JdbcMetaDataUtility.buildRangeQuery(getTableName(), getKey(), rangeQueryPair.getKey(), - rangeQueryPair.getValue()), - java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + protected void initializePreparedStatement() { + try { + // If its a range query pass upper and lower bounds, If its a polling query pass only the lower bound + if (isPollerPartition) { + ps = store.getConnection().prepareStatement(buildRangeQuery(rangeQueryPair.getKey(), Integer.MAX_VALUE), TYPE_FORWARD_ONLY, CONCUR_READ_ONLY); + } else { + ps = store.getConnection().prepareStatement(+ JdbcMetaDataUtility.buildPollableQuery(getTableName(), getKey(), rangeQueryPair.getKey()), + java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + isPollable = true; + buildRangeQuery(rangeQueryPair.getKey(), + (rangeQueryPair.getValue() - rangeQueryPair.getKey())), TYPE_FORWARD_ONLY, CONCUR_READ_ONLY); + } + } catch (SQLException e) { + LOG.error("Exception in initializing the range query for a given partition", e); + throw new RuntimeException(e); + } + windowManager.setup(context); + LOG.debug("super setup done..."); + } + @Override + public void beginWindow(long windowId) { + currentWindowId = windowId; + isReplayed = false; + if (currentWindowId <= windowManager.getLargestRecoveryWindow()) { + try { + replay(currentWindowId); + return; + } catch (SQLException e) { + LOG.error("Exception in replayed windows", e); + throw new RuntimeException(e); + } + } + if (isReplayed && currentWindowId == windowManager.getLargestRecoveryWindow()) { + try { + if (!isPollable && rangeQueryPair.getValue() != null) { + ps = store.getConnection().prepareStatement(+ JdbcMetaDataUtility.buildGTRangeQuery(getTableName(), getKey(), previousUpperBound, + rangeQueryPair.getValue()), + java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + } else { + String bound = null; + if (previousUpperBound == null) { + bound = getRangeQueryPair().getKey(); + } else { + bound = previousUpperBound; + } + ps = store.getConnection().prepareStatement(+ JdbcMetaDataUtility.buildPollableQuery(getTableName(), getKey(), bound), + java.sql.ResultSet.TYPE_FORWARD_ONLY, java.sql.ResultSet.CONCUR_READ_ONLY); + isPollable = true; + } + isReplayed = false;	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2172
20509	APEXMALHAR-2162	[~thomasWeise] We can use this in KafkaOutputOperatorTest but not in KafkaInputOperatorTest. Because, EmbeddedZookeeper supports only for single cluster. It doesn't work for multiple clusters. So, this can't be used it for multi-cluster test cases which are present in KafkaInputOperatorTest. Is it fine to add it only in KafkaOutputOperatorTest ?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2162
20510	APEXMALHAR-2162	Why can we not use multiple embedded zookeepers with different ports when the goal is to simulate multiple clusters?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2162

#	Requirement	Questions	Answered	Successfully Answered	URL
20511	APEXMALHAR-2161	Github user yogidevendra commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/359#discussion_r73849553 --- Diff: library/src/test/java/com/datatorrent/lib/io/fs/AbstractThroughputFileInputOperatorTest.java --- @@ -0,0 +1,129 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.io.fs; +import java.io.IOException; +import java.util.Collections; +import org.junit.Assert; +import org.junit.Before; +import org.junit.Test; +import org.mockito.Mock; +import org.mockito.MockitoAnnotations; +import org.apache.commons.lang.mutable.MutableLong; +import com.datatorrent.api.Partitioner.Partition; +import com.datatorrent.api.Stats.OperatorStats; +import com.datatorrent.api.StatsListener.BatchedOperatorStats; +import com.datatorrent.api.StatsListener.Response; +import com.datatorrent.lib.counters.BasicCounters; +import com.datatorrent.lib.io.fs.AbstractFileInputOperator.FileCounters; +import static org.mockito.Matchers.any; +import static org.mockito.Mockito.when; +public class AbstractThroughputFileInputOperatorTest {+ private AbstractThroughputFileInputOperator<String> underTest; + @Mock + private Partition<AbstractFileInputOperator<String>> mockPartition; + @Mock + private BatchedOperatorStats mockBatchStats; + @Mock + private OperatorStats mockOperatorStats; + @Mock + private BasicCounters<MutableLong> fileCountersMock; + @Mock + private MutableLong fileCounterMock; + + @Before + public void setup() + { + underTest = new ThroughputFileInputOperator(); + MockitoAnnotations.initMocks(this); + + when(mockPartition.getPartitionedInstance()).thenReturn(underTest); + } + + @Test + public void testInitialPartitioning() + { + underTest.setPartitionCount(4); + underTest.setPreferredMaxPendingFilesPerOperator(6); + + for (int i = 0; i < 74; i++) { + underTest.pendingFiles.add("file-" + i); + } + + int partitioncount = underTest.getNewPartitionCount(Collections.singleton(mockPartition), null); + Assert.assertEquals(4, partitioncount); + } + @Test + public void testProcessStats() throws Exception + { + underTest.setPartitionCount(4); + underTest.setPreferredMaxPendingFilesPerOperator(10); + + for (int i = 0; i < 21; i++) { + underTest.pendingFiles.add("file-" + i); + } + + mockOperatorStats.counters = fileCountersMock; + when(mockPartition.getStats()).thenReturn(mockBatchStats); + when(mockBatchStats.getLastWindowedStats()).thenReturn(Collections.singletonList(mockOperatorStats)); + when(fileCountersMock.getCounter(any(FileCounters.class))).thenReturn(fileCounterMock); + when(fileCounterMock.getValue()).thenReturn(20L); + + Response response = underTest.processStats(mockBatchStats); + Assert.assertTrue(response.repartitionRequired); --- End diff -- Could you please add test scenario for repartitionRequired=false.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2161
20512	APEXMALHAR-2157	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/348#discussion_r71688855 --- Diff: library/src/main/java/com/datatorrent/lib/formatter/Formatter.java --- @@ -48,6 +49,13 @@ { protected transient Class<?> clazz; + @AutoMetric + protected long errorTupleCount; --- End diff -- Why does the AutoMetric variables needs to be protected?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2157
20513	APEXMALHAR-2157	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/348#discussion_r71710711 --- Diff: library/src/test/java/com/datatorrent/lib/formatter/JsonFormatterTest.java --- @@ -182,14 +173,19 @@ public void testJSONToPOJONoFieldPOJO() Assert.assertEquals(0, validDataSink.collectedTuples.size()); Assert.assertEquals(1, invalidDataSink.collectedTuples.size()); Assert.assertEquals(o, invalidDataSink.collectedTuples.get(0)); + Assert.assertEquals(1, operator.getIncomingTuplesCount()); + Assert.assertEquals(0, operator.getEmittedObjectCount()); + Assert.assertEquals(1, operator.getErrorTupleCount()); } public static class Test1Pojo { public int a; public long b; + @JsonProperty("c2") public String c; public List<String> d; + @JsonFormat(shape = JsonFormat.Shape.STRING, pattern = "yyyy/MM/dd") --- End diff -- This looks good... Can you also add few other fields and see how the values interact for them: 1. java.util.Date without annotation.. What is the expected format and does the test pass with default format? 2. java.util.Time... Is there any annotation for that? If yes, can you test out Time field for both annotation and without annotation.... If there is no annotation, please add just a new Time field and see default format. 3. Please add tests for different patterns as well.. Just a single pattern is not enough.. Please include other time and timezone etc fields also in pattern. 4. What happens if the shape is change? What are the possible values for that?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2157
20514	APEXMALHAR-2156	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/347#discussion_r72279473 --- Diff: library/src/main/java/com/datatorrent/lib/io/jms/JMSBase.java --- @@ -76,18 +76,24 @@ private transient Session session; private transient Destination destination; + private ConnectionFactoryBuilder connectionFactoryBuilder; private String connectionFactoryClass; private Map<String, String> connectionFactoryProperties = Maps.newHashMap(); private String ackMode = "CLIENT_ACKNOWLEDGE"; - private String clientId = "TestClient"; + private String clientId; private String subject = "TEST.FOO"; + private String subject; --- End diff -- Is subject a parameter that user has to set? In that case add a constraint validation.	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2156
20515	APEXMALHAR-2156	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/347#discussion_r72343723 --- Diff: library/src/main/java/com/datatorrent/lib/io/jms/JMSBase.java --- @@ -355,10 +366,13 @@ public void createConnection() throws JMSException */ protected ConnectionFactory getConnectionFactory() { - logger.debug("class {} properties {}", connectionFactoryClass, connectionFactoryProperties); + logger.debug("connectionFactoryBuilder {} connectionFactoryClass {} properties {}", "" + connectionFactoryBuilder, + connectionFactoryClass, connectionFactoryProperties); ConnectionFactory cf; try { - if (connectionFactoryClass != null) { + if (connectionFactoryBuilder != null) { + cf = connectionFactoryBuilder.buildConnectionFactory(); + } else if (connectionFactoryClass != null) { --- End diff -- Also shouldn't the connection factory properties be actually handled by the builder?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2156

#	Requirement	Questions	Answered	Successfully Answered	URL
20516	APEXMALHAR-2156	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/347#discussion_r72991276 --- Diff: contrib/src/main/java/com/datatorrent/contrib/kafka/AbstractKafkaInputOperator.java --- @@ -250,12 +250,7 @@ public void setup(OperatorContext context) consumer.create(); // reset the offsets to checkpointed one if (consumer instanceof SimpleKafkaConsumer && !offsetStats.isEmpty()) { - Map<KafkaPartition, Long> currentOffsets = new HashMap<>(); - // Increment the offsets and set it to consumer - for (Map.Entry<KafkaPartition, Long> e: offsetStats.entrySet()) { - currentOffsets.put(e.getKey(), e.getValue() + 1); - } - ((SimpleKafkaConsumer)consumer).resetOffset(currentOffsets); + ((SimpleKafkaConsumer)consumer).resetOffset(offsetStats); --- End diff -- Why are kafka changes in this pull request?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2156
20517	APEXMALHAR-2156	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/347#discussion_r72993901 --- Diff: library/src/main/java/com/datatorrent/lib/io/jms/JMSBase.java --- @@ -76,18 +76,52 @@ private transient Session session; private transient Destination destination; - private String connectionFactoryClass; - private Map<String, String> connectionFactoryProperties = Maps.newHashMap(); + private ConnectionFactoryBuilder connectionFactoryBuilder = new DefaultConnectionFactoryBuilder(); private String ackMode = "CLIENT_ACKNOWLEDGE"; - private String clientId = "TestClient"; - private String subject = "TEST.FOO"; + private String clientId; + private String subject; private int batch = 10; private int messageSize = 255; private boolean durable = false; private boolean topic = false; private boolean verbose = false; + protected boolean transacted = true; + public abstract static class ConnectionFactoryBuilder { + protected Map<String, String> connectionFactoryProperties = Maps.newHashMap(); + public Map<String, String> getConnectionFactoryProperties() { return connectionFactoryProperties; } + public void setConnectionFactoryProperties(Map<String, String> connectionFactoryProperties) { this.connectionFactoryProperties = connectionFactoryProperties; } + public abstract ConnectionFactory buildConnectionFactory(); } + public static class DefaultConnectionFactoryBuilder extends ConnectionFactoryBuilder { @Override public ConnectionFactory buildConnectionFactory() { ConnectionFactory cf; try { --- End diff -- Where is the connectionFactoryClass option?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2156
20518	APEXMALHAR-2156	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/347#discussion_r72995064 --- Diff: stream/src/test/java/org/apache/apex/malhar/stream/FunctionOperator/FunctionOperatorTest.java --- @@ -1,332 +0,0 @@ -/** --- End diff -- Why are changes to this operator included?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2156
20519	APEXMALHAR-2156	Github user PramodSSImmaneni commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/347#discussion_r73784474 --- Diff: library/src/main/java/com/datatorrent/lib/io/jms/JMSBase.java --- @@ -76,18 +76,52 @@ private transient Session session; private transient Destination destination; - private String connectionFactoryClass; - private Map<String, String> connectionFactoryProperties = Maps.newHashMap(); + private ConnectionFactoryBuilder connectionFactoryBuilder = new DefaultConnectionFactoryBuilder(); private String ackMode = "CLIENT_ACKNOWLEDGE"; - private String clientId = "TestClient"; - private String subject = "TEST.FOO"; + private String clientId; + private String subject; private int batch = 10; private int messageSize = 255; private boolean durable = false; private boolean topic = false; private boolean verbose = false; + protected boolean transacted = true; + public abstract static class ConnectionFactoryBuilder { + protected Map<String, String> connectionFactoryProperties = Maps.newHashMap(); + public Map<String, String> getConnectionFactoryProperties() { return connectionFactoryProperties; } + public void setConnectionFactoryProperties(Map<String, String> connectionFactoryProperties) { this.connectionFactoryProperties = connectionFactoryProperties; } + public abstract ConnectionFactory buildConnectionFactory(); } + public static class DefaultConnectionFactoryBuilder extends ConnectionFactoryBuilder { @Override public ConnectionFactory buildConnectionFactory() { ConnectionFactory cf; try { --- End diff -- That means they would have to implement a builder that has similar functionality as the default builder isn't it. Why not add the connectionFactoryClass logic to the default builder if (connectionFactoryClass != null) { @SuppressWarnings("unchecked") Class<ConnectionFactory> clazz = (Class<ConnectionFactory>)Class.forName(connectionFactoryClass); cf = clazz.newInstance(); } else { cf = new org.apache.activemq.ActiveMQConnectionFactory(); }	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2156
20520	APEXMALHAR-2155	Is this related to join operators?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2155
20521	APEXMALHAR-2154	Is there no behavior in this operator that can benefit from beforeCheckpoint?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2154
20522	APEXMALHAR-2153	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/353#discussion_r72930458 --- Diff: docs/operators/enricher.md --- @@ -0,0 +1,169 @@ +POJO Enricher +========== + +## Operator Objective +This operator receives an POJO ([Plain Old Java Object](https://en.wikipedia.org/wiki/Plain_Old_Java_Object)) as an incoming tuple and uses an external source to enrich the data in the incoming tuple and finally emits the enriched data as a new enriched POJO. + +POJOEnricher supports enrichment from following external sources: +1. **JSON File Based** - Reads the file in memory having content stored in JSON format and use that to enrich the data. This can be done using FSLoader implementation. --- End diff -- Should this be small enough to fit in memory. If so, please mention.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2153

#	Requirement	Questions	Answered	Successfully Answered	URL
20523	APEXMALHAR-2153	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/353#discussion_r72930659 --- Diff: docs/operators/enricher.md --- @@ -0,0 +1,169 @@ +POJO Enricher +========== + +## Operator Objective +This operator receives an POJO ([Plain Old Java Object](https://en.wikipedia.org/wiki/Plain_Old_Java_Object)) as an incoming tuple and uses an external source to enrich the data in the incoming tuple and finally emits the enriched data as a new enriched POJO. + +POJOEnricher supports enrichment from following external sources: + 1. **JSON File Based** - Reads the file in memory having content stored in JSON format and use that to enrich the data. This can be done using FSLoader implementation. 2. **JDBC Based** - Any JDBC store can act as an external entity to which enricher can request data for enriching incoming tuples. This can be done using JDBCLoader implementation. + +POJO Enricher does not hold any state and is **idempotent**, **fault-tolerance** and **statically/dynamically partitionable**. + +## Operator Usecase +1. Bank ***transaction records*** usually contains customerId. For further analysis of transaction one wants the customer name and other customer related information. +Such information is present in another database. One could enrich the transaction's record with customer information using POJOEnricher. 2. ***Call Data Record (CDR)*** contains only mobile/telephone numbers of the customer. Customer information is missing in CDR. POJO Enricher can be used to enricher --- End diff -- nit: can be used to enrich*	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2153
20524	APEXMALHAR-2153	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/353#discussion_r72930798 --- Diff: docs/operators/enricher.md --- @@ -0,0 +1,169 @@ +POJO Enricher +========== + +## Operator Objective +This operator receives an POJO ([Plain Old Java Object](https://en.wikipedia.org/wiki/Plain_Old_Java_Object)) as an incoming tuple and uses an external source to enrich the data in the incoming tuple and finally emits the enriched data as a new enriched POJO. + +POJOEnricher supports enrichment from following external sources: + 1. **JSON File Based** - Reads the file in memory having content stored in JSON format and use that to enrich the data. This can be done using FSLoader implementation. 2. **JDBC Based** - Any JDBC store can act as an external entity to which enricher can request data for enriching incoming tuples. This can be done using JDBCLoader implementation. + +POJO Enricher does not hold any state and is **idempotent**, **fault-tolerance** and **statically/dynamically partitionable**. + +## Operator Usecase +1. Bank ***transaction records*** usually contains customerId. For further analysis of transaction one wants the customer name and other customer related information. +Such information is present in another database. One could enrich the transaction's record with customer information using POJOEnricher. 2. ***Call Data Record (CDR)*** contains only mobile/telephone numbers of the customer. Customer information is missing in CDR. POJO Enricher can be used to enricher +CDR with customer data for further analysis. + +## Operator Information +1. Operator location: ***malhar-contrib*** +2. Available since: ***3.4.0*** +3. Operator state: ***Evolving*** +3. Java Packages: + * Operator: ***[com.datatorrent.contrib.enrich.POJOEnricher](https://www.datatorrent.com/docs/apidocs/com/datatorrent/contrib/enrich/POJOEnricher.html)*** + * FSLoader: ***[com.datatorrent.contrib.enrich.FSLoader](https://www.datatorrent.com/docs/apidocs/com/datatorrent/contrib/enrich/FSLoader.html)*** + * JDBCLoader: ***[com.datatorrent.contrib.enrich.JDBCLoader](https://www.datatorrent.com/docs/apidocs/com/datatorrent/contrib/enrich/JDBCLoader.html)*** + +## Properties, Attributes and Ports +### Properties of POJOEnricher +| **Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+| --------	-----------	----	----------	-------------	+| *includeFields*	List of fields from database that needs to be added to output POJO.	List<String\>	Yes	N/A	+| *lookupFields*	List of fields from input POJO which will form a *unique composite* key for querying to database	List<String\>	Yes	N/A	--- End diff -- Will composite key work for JSON?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2153

#	Requirement	Questions	Answered	Successfully Answered	URL
20525	APEXMALHAR-2153	docs/operators/enricher.md --- @@ -0,0 +1,169 @@ +POJO Enricher +========== + +## Operator Objective +This operator receives an POJO ([Plain Old Java Object](https://en.wikipedia.org/wiki/Plain_Old_Java_Object)) as an incoming tuple and uses an external source to enrich the data in +the incoming tuple and finally emits the enriched data as a new enriched POJO. + +POJOEnricher supports enrichment from following external sources: + +1. **JSON File Based** - Reads the file in memory having content stored in JSON format and use that to enrich the data. This can be done using FSLoader implementation. +2. **JDBC Based** - Any JDBC store can act as an external entity to which enricher can request data for enriching incoming tuples. This can be done using JDBCLoader implementation. + +POJO Enricher does not hold any state and is **idempotent**, **fault-tolerance** and **statically/dynamically partitionable**. + +## Operator Usecase +1. Bank ***transaction records*** usually contains customerId. For further analysis of transaction one wants the customer name and other customer related information. +Such information is present in another database. One could enrich the transaction's record with customer information using POJOEnricher. +2. ***Call Data Record (CDR)*** contains only mobile/telephone numbers of the customer. Customer information is missing in CDR. POJO Enricher can be used to enricher +CDR with customer data for further analysis. + +## Operator Information +1. Operator location: ***malhar-contrib*** +2. Available since: ***3.4.0*** +3. Operator state: ***Evolving*** +3. Java Packages: + * Operator: ***[com.datatorrent.contrib.enrich.POJOEnricher](https://www.datatorrent.com/docs/apidocs/com/datatorrent/contrib/enrich/POJOEnricher.html)*** + * FSLoader: ***[com.datatorrent.contrib.enrich.FSLoader](https://www.datatorrent.com/docs/apidocs/com/datatorrent/contrib/enrich/FSLoader.html)*** + * JDBCLoader: ***[com.datatorrent.contrib.enrich.JDBCLoader](https://www.datatorrent.com/docs/apidocs/com/datatorrent/contrib/enrich/JDBCLoader.html)*** + +## Properties, Attributes and Ports +### Properties of POJOEnricher +| **Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+| --------	----------	----	----------	----------	+| *includeFields*	List of fields from database that needs to be added to output POJO.	List<String\>	Yes	N/A	+| *lookupFields*	List of fields from input POJO which will form a *unique composite* key for querying to database	List<String\>	Yes	N/A	+| *store*	Backend Store from which data should be queried for enrichment	[BackendStore](#backendStore)	Yes	N/A	+| *cacheExpirationInterval*	Cache entry expiry in ms. After this time, the lookup to store will be done again for given key	int	No	1 * 60 * 60 * 1000 (1 hour)	+| *cacheCleanupInterval*	Interval in ms after which cache will be removed for any stale entries.	int	No	1 * 60 * 60 * 1000 (1 hour)	+| *cacheSize*	Number of entry in cache after which eviction will start on each addition based on LRU	int	No	1000	+ +#### Properties of FSLoader (BackendStore) +| **Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+| --------	----------	----	----------	----------	+| *fileName*	Path of the file, the data from which will be used for enrichment. See [here](#JSONFileFormat) for JSON File format.	String	Yes	N/A	+ +#### Properties of JDBCLoader (BackendStore) +| **Property**	**Description**	**Type**	**Mandatory**	**Default Value**	+| --------	----------	----	----------	----------	+| *databaseUrl*	Connection string for connecting to JDBC	String	Yes	N/A	+| *databaseDriver*	JDBC Driver class for connection to JDBC Store. This driver should be there in classpath	String	Yes	N/A	+| *tableName*	Name of the table from which data needs to be retrieved	String	Yes	N/A	+| *connectionProperties*	Command seperated list of advanced connection properties that need to be passed to JDBC Driver. For eg. *prop1:val1,prop2:val2*	String				1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2153
20526	APEXMALHAR-2130	What's the use case? I think this needs a clearer explanation as to what the purpose of the proposed feature is. We do have managed state and on top of it abstractions to plug it into operators: https://github.com/apache/apex-malhar/tree/master/library/src/main/java/org/apache/apex/malhar/lib/state/	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2130
20527	APEXMALHAR-2130	Here are some thoughts to implement it based on state manager - Use one bucket only, use the Event Window and Key to generate key for the bucket, and use Stream Window as timeBucket. The problem of this approach is get the whole map of one Event window as it lack of key. Is possible the client code get value by Event Window and key instead of get whole map? - One Event window map to one bucket, and Stream window id as timeBucket. But the problem is the size of Event Window are variable and could be huge. and also the Event window size are vary. So this approach probably not doable. - One Streaming window map to a bucket. It will have problem to map Event window. And also have problem get whole map by Event Window, as we don't know which bucket to read the data.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2130

#	Requirement	Questions	Answered	Successfully Answered	URL
20528	APEXMALHAR-2130	library/src/main/java/org/apache/apex/malhar/lib/window/impl/SpillableWindowedPlainStorage.java --- @@ -0,0 +1,188 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window.impl; + +import java.util.Iterator; +import java.util.Map; + +import org.apache.apex.malhar.lib.state.spillable.Spillable; +import org.apache.apex.malhar.lib.state.spillable.SpillableComponentImpl; +import org.apache.apex.malhar.lib.state.spillable.SpillableStateStore; +import org.apache.apex.malhar.lib.state.spillable.managed.ManagedStateSpillableStateStore; +import org.apache.apex.malhar.lib.utils.serde.Serde; +import org.apache.apex.malhar.lib.utils.serde.SerdeKryoSlice; +import org.apache.apex.malhar.lib.window.Window; +import org.apache.apex.malhar.lib.window.WindowedStorage; + +import com.datatorrent.api.Context; +import com.datatorrent.netlet.util.Slice; + +/** + * This is an implementation of WindowedPlainStorage that makes use of {@link Spillable} data structures + * + * @param <T> Type of the value per window + */ +public class SpillableWindowedPlainStorage<T> implements WindowedStorage.WindowedPlainStorage<T> +{ + private SpillableStateStore store; + private transient SpillableComponentImpl sccImpl; + private long bucket; + private Serde<Window, Slice> windowSerde; + private Serde<T, Slice> valueSerde; + + protected transient Spillable.SpillableByteMap<Window, T> internMap; + + public SpillableWindowedPlainStorage() + { + } + + public SpillableWindowedPlainStorage(long bucket, Serde<Window, Slice> windowSerde, Serde<T, Slice> valueSerde) + { + this.bucket = bucket; + this.windowSerde = windowSerde; + this.valueSerde = valueSerde; + } + + public void setStore(SpillableStateStore store) + { + this.store = store; + } + + public void setBucket(long bucket) + { + this.bucket = bucket; + } + + public void setWindowSerde(Serde<Window, Slice> windowSerde) + { + this.windowSerde = windowSerde; + } + + public void setValueSerde(Serde<T, Slice> valueSerde) + { + this.valueSerde = valueSerde; + } + + @Override + public void put(Window window, T value) + { + internMap.put(window, value); + } + + @Override + public T get(Window window) + { + return internMap.get(window); + } + + @Override + public Iterable<Map.Entry<Window, T>> entrySet() + { + return internMap.entrySet(); + } + + @Override + public Iterator<Map.Entry<Window, T>> iterator() + { + return internMap.entrySet().iterator(); + } + + @Override + public boolean containsWindow(Window window) + { + return internMap.containsKey(window); + } + + @Override + public long size() + { + return internMap.size(); + } + + @Override + public void remove(Window window) + { + internMap.remove(window); + } + + @Override + public void migrateWindow(Window fromWindow, Window toWindow) + { + internMap.put(toWindow, internMap.remove(fromWindow)); + } + + @Override + public void setup(Context.OperatorContext context) + { + if (store == null) { + // provide a default store + store = new ManagedStateSpillableStateStore(); + } + if (bucket == 0) { + //	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2130
20529	APEXMALHAR-2130	[~davidyan] As far I can tell what you have implemented ontop of spillable datastructures is 95% of the way there. I'm not sure what has changed. 1. SpillableByteMap already supports deleting a key by setting the value for a key to be an empty byte array. So that's there. This can easily be extended to delete all the values in SpillableArrayListMultimap. 2. The only missing piece is the ability to iterate over the set of keys in a map. This can be done with another SpillableByteMap let's call this map the linkedListMap. The key of the linkedListMap represents the current node, the value represents the next node. You then keep track of your "head" key and iterate by taking your current node, and getting the value. The value then becomes the current node and so on. When the value for the current node is null you are done traversing the list. You can take this and wrap it in an iterator.	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2130
20530	APEXMALHAR-2130	Which keys belong to which window can be seen as derived information when using a key in managed state that is <window><key>. We can then retrieve all keys for a given window by doing a prefix scan for <window>. This becomes difficult however when including the state that wasn't flushed to the data files yet (they are compacted asynchronously). Using another spillable datastructure to store the keys for each window is possible, but also comes with drawbacks, since it duplicates information. Once of them is obviously performance, since it is increasing the HDFS usage. It also needs to use the same store, as otherwise there is the situation of two resources that need to be updated atomically. Assuming that both collections can use the same store (key prefix) and share the same WAL and data files, then we need to see whether time based purging will still work as intended. David?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2130
20531	APEXMALHAR-2129	Please explain the requirement and why is it needed? This helps with purging of data. How will you purge data?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2129
20532	APEXMALHAR-2126	Why is it necessary to have Slice here? Slice allows referring to a portion of byte array without allocating individual byte arrays and creating many byte array objects. The example tries to solve a completely different problem - concatenate existing byte arrays into a single large byte array. There is no need for Slice here at all and Slice is misused in SpillableByteMapImpl. Instead of using Slice, it will be better to use Collection (for example ArrayList) of identifiers and serialized keys pairs and create byte array after all pairs are added to the collection.	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2126
20533	APEXMALHAR-2121	Github user siyuanh commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/323#discussion_r67717267 --- Diff: contrib/src/main/java/com/datatorrent/contrib/kafka/AbstractKafkaInputOperator.java --- @@ -203,6 +204,14 @@ public AbstractKafkaInputOperator() */ protected abstract void emitTuple(Message message); --- End diff -- It is good to change this method signature, but can we do that?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2121

#	Requirement	Questions	Answered	Successfully Answered	URL
20534	APEXMALHAR-2116	Github user amberarrow commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/326#discussion_r67969125 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/fs/FSRecordReader.java --- @@ -0,0 +1,175 @@ + +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.fs; +import java.io.IOException; +import org.apache.commons.beanutils.ConvertUtils; +import org.apache.commons.beanutils.converters.AbstractConverter; +import org.apache.hadoop.fs.FSDataInputStream; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.lib.io.block.BlockMetadata; +import com.datatorrent.lib.io.block.FSSliceReader; +import com.datatorrent.lib.io.block.ReaderContext; +/** + * This operator can be used for reading records/tuples from Filesystem + * in parallel (without ordering guarantees between tuples). + * Records can be delimited (e.g. newline) or fixed with records. --- End diff -- fixed with => fixed width	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2116
20535	APEXMALHAR-2116	Github user amberarrow commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/326#discussion_r67970457 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/fs/FSRecordReader.java --- @@ -0,0 +1,175 @@ + +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.fs; +import java.io.IOException; +import org.apache.commons.beanutils.ConvertUtils; +import org.apache.commons.beanutils.converters.AbstractConverter; +import org.apache.hadoop.fs.FSDataInputStream; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.lib.io.block.BlockMetadata; +import com.datatorrent.lib.io.block.FSSliceReader; +import com.datatorrent.lib.io.block.ReaderContext; +/** + * This operator can be used for reading records/tuples from Filesystem + * in parallel (without ordering guarantees between tuples). + * Records can be delimited (e.g. newline) or fixed with records. + * Output tuples are byte[]. + * + * Typically, this operator will be connected to output of FileSplitterInput + * to read records in parallel. + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class FSRecordReader extends FSSliceReader +{ + /** + * Record reader mode decides how to split the records. + */ + public static enum RECORD_READER_MODE + { + DELIMITED_RECORD, FIXED_WIDTH_RECORD; + } + + /** + * Criteria for record split + */ + private RECORD_READER_MODE mode = RECORD_READER_MODE.DELIMITED_RECORD; + + /** + * Length for fixed width record + */ + private int recordLength; + + /** + * Port to emit individual records/tuples as byte[] + */ + public final transient DefaultOutputPort<byte[]> records = new DefaultOutputPort<byte[]>(); + + /** + * Initialize appropriate reader context based on mode selection + */ + @Override + public void setup(OperatorContext context) + { + super.setup(context); + if (mode == RECORD_READER_MODE.FIXED_WIDTH_RECORD) { + ReaderContext.FixedBytesReaderContext<FSDataInputStream> fixedBytesReaderContext = new ReaderContext.FixedBytesReaderContext<FSDataInputStream>(); + fixedBytesReaderContext.setLength(recordLength); + readerContext = fixedBytesReaderContext; + } else { + readerContext = new ReaderContext.ReadAheadLineReaderContext<FSDataInputStream>(); + } + } + + /** + * Read the block data and emit records based on reader context + * + * @param blockMetadata block + * @throws IOException + */ + protected void readBlock(BlockMetadata blockMetadata) throws IOException + { + readerContext.initialize(stream, blockMetadata, consecutiveBlock); + ReaderContext.Entity entity; + while ((entity = readerContext.next()) != null) { + counters.getCounter(ReaderCounterKeys.BYTES).add(entity.getUsedBytes()); + byte[] record = entity.getRecord(); + //If the record is partial then ignore the record. --- End diff -- why 'partial' rather than 'null' ?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2116

#	Requirement	Questions	Answered	Successfully Answered	URL
20536	APEXMALHAR-2116	library/src/main/java/org/apache/apex/malhar/lib/fs/FSRecordReader.java --- @@ -0,0 +1,180 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.fs; +import java.io.IOException; +import org.apache.commons.beanutils.ConvertUtils; +import org.apache.commons.beanutils.converters.AbstractConverter; +import org.apache.hadoop.fs.FSDataInputStream; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.lib.io.block.BlockMetadata; +import com.datatorrent.lib.io.block.FSSliceReader; +import com.datatorrent.lib.io.block.ReaderContext; +/** + * This operator can be used for reading records/tuples from Filesystem in + * parallel (without ordering guarantees between tuples). Records can be + * delimited (e.g. newline) or fixed width records. Output tuples are byte[]. + * + * Typically, this operator will be connected to output of FileSplitterInput to + * read records in parallel. + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class FSRecordReader extends FSSliceReader +{ + /** + * Record reader mode decides how to split the records. + */ + public static enum RECORD_READER_MODE + { + DELIMITED_RECORD, FIXED_WIDTH_RECORD; + } + + /** + * Criteria for record split + */ + private RECORD_READER_MODE mode = RECORD_READER_MODE.DELIMITED_RECORD; + + /** + * Length for fixed width record + */ + private int recordLength; + + /** + * Port to emit individual records/tuples as byte[] + */ + public final transient DefaultOutputPort<byte[]> records = new DefaultOutputPort<byte[]>(); + + /** + * Initialize appropriate reader context based on mode selection + */ + @Override + public void setup(OperatorContext context) + { + super.setup(context); + if (mode == RECORD_READER_MODE.FIXED_WIDTH_RECORD) { + ReaderContext.FixedBytesReaderContext<FSDataInputStream> fixedBytesReaderContext = new ReaderContext.FixedBytesReaderContext<FSDataInputStream>(); + fixedBytesReaderContext.setLength(recordLength); + readerContext = fixedBytesReaderContext; + } else { + readerContext = new ReaderContext.ReadAheadLineReaderContext<FSDataInputStream>(); + } + } + + /** + * Read the block data and emit records based on reader context + * + * @param blockMetadata + * block + * @throws IOException + */ + protected void readBlock(BlockMetadata blockMetadata) throws IOException + { + readerContext.initialize(stream, blockMetadata, consecutiveBlock); + ReaderContext.Entity entity; + while ((entity = readerContext.next()) != null) { + counters.getCounter(ReaderCounterKeys.BYTES).add(entity.getUsedBytes()); + byte[] record = entity.getRecord(); + if (record != null) { + counters.getCounter(ReaderCounterKeys.RECORDS).increment(); + records.emit(record); + } + } + } + + /** + * Criteria for record split + * + * @param mode + * Mode + */ + public void setMode(RECORD_READER_MODE mode) + { + this.mode = mode; + } + + /** + * Criteria for record split + * + * @return mode	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2116
20537	APEXMALHAR-2116	library/src/main/java/org/apache/apex/malhar/lib/fs/FSRecordReaderModule.java --- @@ -0,0 +1,332 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.fs; +import javax.validation.constraints.Min; +import javax.validation.constraints.NotNull; +import javax.validation.constraints.Size; +import org.apache.apex.malhar.lib.fs.FSRecordReader.RECORD_READER_MODE; +import org.apache.hadoop.conf.Configuration; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.Module; +import com.datatorrent.common.partitioner.StatelessPartitioner; +import com.datatorrent.lib.codec.KryoSerializableStreamCodec; +import com.datatorrent.lib.io.block.BlockMetadata; +import com.datatorrent.lib.io.block.FSSliceReader; +import com.datatorrent.lib.io.fs.FileSplitterInput; +/** + * This module is used for reading records/tuples from FileSystem. Records can + * be read in parallel using multiple partitions of record reader operator. + * (Ordering is not guaranteed when records are read in parallel) + * + * Input directory is scanned at specified interval to poll for new data. + * + * The module reads data in parallel, following parameters can be configured + * + * 1. files: list of file(s)/directories to read + * 2. filePatternRegularExp: Files names matching given regex will be read + * 3. scanIntervalMillis: interval between two scans to discover new files in + * input directory + * 4. recursive: if scan recursively input directories + * 5. blockSize: block size used to read input blocks of file + * 6. readersCount: count of readers to read input file + * 7. sequentialFileRead: If emit file blocks in sequence? + * 8. blocksThreshold: number of blocks emitted per window + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class FSRecordReaderModule implements Module +{ + @NotNull + @Size(min = 1) + private String files; + private String filePatternRegularExp; + @Min(0) + private long scanIntervalMillis; + private boolean recursive = true; + private boolean sequentialFileRead = false; + private int readersCount; + @Min(1) + protected int blocksThreshold; + public final transient ProxyOutputPort<byte[]> records = new ProxyOutputPort<byte[]>(); + /** + * Criteria for record split + */ + private RECORD_READER_MODE mode; + /** + * Length for fixed width record + */ + private int recordLength; + public FileSplitterInput createFileSplitter() + { + return new FileSplitterInput(); + } + public FSRecordReader createBlockReader() + { + FSRecordReader recordReader = new FSRecordReader(); + recordReader.setMode(mode); + recordReader.setRecordLength(recordLength); + return recordReader; + } + @Override + public void populateDAG(DAG dag, Configuration configuration) + { + FileSplitterInput fileSplitter = dag.addOperator("FileSplitter", createFileSplitter()); + FSRecordReader recordReader = dag.addOperator("BlockReader", createBlockReader()); + dag.addStream("BlockMetadata", fileSplitter.blocksMetadataOutput, recordReader.blocksMetadataInput); + if (sequentialFileRead) { + dag.setInputPortAttribute(recordReader.blocksMetadataInput,	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2116

#	Requirement	Questions	Answered	Successfully Answered	URL
20538	APEXMALHAR-2116	library/src/main/java/org/apache/apex/malhar/lib/fs/FSRecordReaderModule.java --- @@ -0,0 +1,334 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.fs; +import javax.validation.constraints.Min; +import javax.validation.constraints.NotNull; import javax.validation.constraints.Size; +import org.apache.apex.malhar.lib.fs.FSRecordReader.RECORD_READER_MODE; import org.apache.hadoop.conf.Configuration; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.Module; +import com.datatorrent.common.partitioner.StatelessPartitioner; +import com.datatorrent.lib.codec.KryoSerializableStreamCodec; +import com.datatorrent.lib.io.block.BlockMetadata; +import com.datatorrent.lib.io.block.FSSliceReader; +import com.datatorrent.lib.io.fs.FileSplitterInput; +/** * This module is used for reading records/tuples from FileSystem. Records can be read in parallel using multiple partitions of record reader operator. * (Ordering is not guaranteed when records are read in parallel) * Input directory is scanned at specified interval to poll for new data. * The module reads data in parallel, following parameters can be configured * * 1. files: list of file(s)/directories to read * 2. filePatternRegularExp: Files with names matching given regex will be read * 3. scanIntervalMillis: interval between two scans to discover new files in * input directory * 4. recursive: if true, scan input directories recursively * 5. blockSize: block size used to read input blocks of file * 6. readersCount: count of readers to read input file * 7. sequentialFileRead: if true, then each reader partition will read different file. * instead of reading different offsets of the same file. * (File level parallelism instead of block level parallelism) * 8. blocksThreshold: number of blocks emitted per window */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class FSRecordReaderModule implements Module +{ @NotNull @Size(min = 1) private String files; private String filePatternRegularExp; @Min(0) private long scanIntervalMillis; private boolean recursive = true; private boolean sequentialFileRead = false; private int readersCount; @Min(1) protected int blocksThreshold; public final transient ProxyOutputPort<byte[]> records = new ProxyOutputPort<byte[]>(); /** * Criteria for record split */ private RECORD_READER_MODE mode; /** * Length for fixed width record */ private int recordLength; public FileSplitterInput createFileSplitter() { return new FileSplitterInput(); } public FSRecordReader createBlockReader() { FSRecordReader recordReader = new FSRecordReader(); recordReader.setMode(mode); recordReader.setRecordLength(recordLength); return recordReader; } @Override public void populateDAG(DAG dag, Configuration configuration) { FileSplitterInput fileSplitter = dag.addOperator("FileSplitter", createFileSplitter()); FSRecordReader recordReader = dag.addOperator("BlockReader", createBlockReader()); dag.addStream("BlockMetadata", fileSplitter.blocksMetadataOutput,	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2116
20539	APEXMALHAR-2116	Github user yogidevendra commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/326#discussion_r70930388 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/fs/FSRecordReaderModule.java --- @@ -0,0 +1,344 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.fs; +import javax.validation.constraints.Min; +import javax.validation.constraints.NotNull; import javax.validation.constraints.Size; +import org.apache.apex.malhar.lib.fs.FSRecordReader.RECORD_READER_MODE; import org.apache.hadoop.conf.Configuration; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.Module; +import com.datatorrent.common.partitioner.StatelessPartitioner; +import com.datatorrent.lib.codec.KryoSerializableStreamCodec; +import com.datatorrent.lib.io.block.BlockMetadata; +import com.datatorrent.lib.io.block.FSSliceReader; +import com.datatorrent.lib.io.fs.FileSplitterInput; +/** * This module is used for reading records/tuples from FileSystem. Records can be read in parallel using multiple partitions of record reader operator. * (Ordering is not guaranteed when records are read in parallel) * Input directory is scanned at specified interval to poll for new data. * The module reads data in parallel, following parameters can be configured * * 1. files: list of file(s)/directories to read * 2. filePatternRegularExp: Files with names matching given regex will be read * 3. scanIntervalMillis: interval between two scans to discover new files in * input directory * 4. recursive: if true, scan input directories recursively * 5. blockSize: block size used to read input blocks of file * 6. readersCount: count of readers to read input file * 7. sequentialFileRead: if true, then each reader partition will read different file. * instead of reading different offsets of the same file. * (File level parallelism instead of block level parallelism) * 8. blocksThreshold: number of blocks emitted per window */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class FSRecordReaderModule implements Module +{ @NotNull @Size(min = 1) private String files; private String filePatternRegularExp; @Min(1) private long scanIntervalMillis = 5000; private boolean recursive = true; private boolean sequentialFileRead = false; @Min(1) private int readersCount = 1; @Min(1) protected int blocksThreshold = 1; public final transient ProxyOutputPort<byte[]> records = new ProxyOutputPort<byte[]>(); /** * Criteria for record split */ private RECORD_READER_MODE mode = RECORD_READER_MODE.DELIMITED_RECORD; /** * Length for fixed width record */ @Min(1) private int recordLength; --- End diff -- Oh. Good point. Should we separate this into two different classes. FixedWidthRecordReader, DelimitedRecordReader? That will make configuration clean. @amberarrow Any thoughts?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2116

#	Requirement	Questions	Answered	Successfully Answered	URL
20540	APEXMALHAR-2116	Github user yogidevendra commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/326#discussion_r71698302 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/fs/FSRecordReaderModule.java --- @@ -0,0 +1,344 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.fs; +import javax.validation.constraints.Min; +import javax.validation.constraints.NotNull; +import javax.validation.constraints.Size; +import org.apache.apex.malhar.lib.fs.FSRecordReader.RECORD_READER_MODE; +import org.apache.hadoop.conf.Configuration; +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.Module; +import com.datatorrent.common.partitioner.StatelessPartitioner; +import com.datatorrent.lib.codec.KryoSerializableStreamCodec; +import com.datatorrent.lib.io.block.BlockMetadata; +import com.datatorrent.lib.io.block.FSSliceReader; +import com.datatorrent.lib.io.fs.FileSplitterInput; +/** * This module is used for reading records/tuples from FileSystem. Records can * be read in parallel using multiple partitions of record reader operator. * (Ordering is not guaranteed when records are read in parallel) * * Input directory is scanned at specified interval to poll for new data. * * The module reads data in parallel, following parameters can be configured * * 1. files: list of file(s)/directories to read * 2. filePatternRegularExp: Files with names matching given regex will be read * 3. scanIntervalMillis: interval between two scans to discover new files in * input directory * 4. recursive: if true, scan input directories recursively * 5. blockSize: block size used to read input blocks of file * 6. readersCount: count of readers to read input file * 7. sequentialFileRead: if true, then each reader partition will read different file. * instead of reading different offsets of the same file. * (File level parallelism instead of block level parallelism) * 8. blocksThreshold: number of blocks emitted per window * */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class FSRecordReaderModule implements Module { + @NotNull + @Size(min = 1) + private String files; + private String filePatternRegularExp; + @Min(1) + private long scanIntervalMillis = 5000; + private boolean recursive = true; + private boolean sequentialFileRead = false; + @Min(1) + private int readersCount = 1; + @Min(1) + protected int blocksThreshold = 1; + + public final transient ProxyOutputPort<byte[]> records = new ProxyOutputPort<byte[]>(); + + /** + * Criteria for record split + */ + private RECORD_READER_MODE mode = RECORD_READER_MODE.DELIMITED_RECORD; + + /** + * Length for fixed width record + */ + @Min(1) + private int recordLength; --- End diff -- @amberarrow As per your suggestion, removed @ min annotation for record length. Added validation in the code to ensure that the value is positive if FIXED_WIDTH_RECORD. Also, added test case for validating this. Could you have a look and merge this if it looks OK.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2116
20541	APEXMALHAR-2111	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/311#discussion_r66675231 --- Diff: library/src/main/java/com/datatorrent/lib/projection/ProjectionOperator.java --- @@ -308,5 +303,45 @@ private void handleProjection(Object t) } } + /** + * set selectFields, a list of fields to be selected from incoming POJO + * + * @param selectFields List of fields from POJO to be selected + */ + public void setSelectFields(List<String> selectFields) --- End diff -- Now that we have proper configuration for selectFields, can you please add @useSchema and @description javadoc metatag. You can refer to POJOEnricher setIncludeFields setter method for this.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2111
20542	APEXMALHAR-2106	Should this of type be improvement?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2106
20543	APEXMALHAR-2106	GitHub user ilganeli opened a pull request: https://github.com/apache/incubator-apex-malhar/pull/307 [APEXMALHAR-2106] Support multiple streams in StreamMerger * Added ability to support more than one stream in the operator Note: I am unclear whether this approach will work if the operator or the DAG restarts, will the ports be reconnected properly? If not, what mechanism prevents this? You can merge this pull request into a Git repository by running: $ git pull https://github.com/ilganeli/incubator-apex-malhar APEXMALHAR_2106 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/incubator-apex-malhar/pull/307.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #307 ---- commit 5a486ff777c3b5c41fa2cbbff7a196bc9da18a44 Author: Ilya Ganelin <ilya.ganelin@capitalone.com> Date: 2016-06-01T23:36:49Z Initial implementation of a Stream merger operator that can handle more than two input streams. Updated test cases ----	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2106
20544	APEXMALHAR-2106	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/307#discussion_r65472797 --- Diff: library/src/main/java/com/datatorrent/lib/stream/StreamMerger.java --- @@ -41,35 +47,57 @@ @Stateless public class StreamMerger<K> extends BaseOperator { - /** - * Data input port 1. - */ - public final transient DefaultInputPort<K> data1 = new DefaultInputPort<K>() + private int portCount = 0; + private final transient ArrayList<DefaultInputPort<K>> ports = new ArrayList<>(portCount); --- End diff -- Did you use this operator in a DAG and see if it works?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2106

#	Requirement	Questions	Answered	Successfully Answered	URL
20545	APEXMALHAR-2106	library/src/main/java/com/datatorrent/lib/stream/MultipleStreamMerger.java --- @@ -0,0 +1,277 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.stream; + +import java.util.ArrayList; +import java.util.Iterator; + +import org.apache.hadoop.conf.Configuration; + +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Module; + +/** + * Module that adds functionality to bypass the platform limitations of combining more than two streams at a time with + * Stream Merger. + * + * Usage: + * + * dag.addOperator("Stream_1", op1); + * dag.addOperator("Stream_2", op2); + * dag.addOperator("Stream_3", op3); + * + * MultipleStreamMerger merger = new MultipleStreamMerger(); + * merger.merge(op1.out) + * .merge(op2.out) + * .merge(op3.out) + * .insertInto(dag, conf); + * + * dag.addModule("Merger", merger); + * + * @param <K> + */ +public class MultipleStreamMerger<K> implements Module +{ + public class Stream + { + DefaultInputPort destPort; + DefaultOutputPort sourcePort; + String name; + + public Stream(String name, DefaultOutputPort sourcePort, DefaultInputPort destPort) + { + this.destPort = destPort; + this.sourcePort = sourcePort; + this.name = name; + } + } + + public class NamedMerger + { + StreamMerger<K> merger; + String name; + + public NamedMerger(String name, StreamMerger<K> merger) + { + this.merger = merger; + this.name = name; + } + } + + private int streamCount = 0; + ArrayList<DefaultOutputPort<K>> streamsToMerge = new ArrayList<>(); + public transient ProxyOutputPort<K> streamOutput = new ProxyOutputPort<>(); + + /** + * Used to define all the sources to be merged into a single stream. + * + * @param sourcePort - The output port from the upstream operator that provides data + * @return The updated MultipleStreamMerger object that tracks which streams should be unified. + */ + public MultipleStreamMerger<K> merge(DefaultOutputPort<K> sourcePort) + { + streamsToMerge.add(sourcePort); + return this; + } + + /** + * To merge more than two streams at a time, we construct a binary tree of thread-local StreamMerger operators + * E.g. + * Tier 0 Tier 1 Tier 2 + * Stream 1 -> StreamMerger_1 -> Stream 2 -> StreamMerger_Final -> Out + * Stream 3 -> StreamMerger_2 -> Stream 4 -> + * + * This function updates the provided DAG with the relevant streams. + */ + public void mergeStreams(DAG dag, Configuration conf) + { + if (streamsToMerge.size() < 2) { + throw new IllegalArgumentException("Not enough streams to merge, at least two streams must be selected for " + "merging with `.merge()`."); + } + ArrayList<Stream> streamsToAddToDag = new ArrayList<>(); + ArrayList<NamedMerger> operatorsToAdd = new ArrayList<>(); + // Determine operators and streams to add to the DAG + constructMergeTree(streamsToAddToDag, operatorsToAdd); + for (NamedMerger m : operatorsToAdd) { + dag.addOperator(m.name, m.merger); + } + for (Stream s : streamsToAddToDag) { + dag.addStream(s.name, s.sourcePort,	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2106
20546	APEXMALHAR-2106	library/src/main/java/com/datatorrent/lib/stream/MultipleStreamMerger.java --- @@ -0,0 +1,277 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.stream; + +import java.util.ArrayList; +import java.util.Iterator; + +import org.apache.hadoop.conf.Configuration; + +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Module; + +/** + * Module that adds functionality to bypass the platform limitations of combining more than two streams at a time with + * Stream Merger. + * + * Usage: + * + * dag.addOperator("Stream_1", op1); + * dag.addOperator("Stream_2", op2); + * dag.addOperator("Stream_3", op3); + * + * MultipleStreamMerger merger = new MultipleStreamMerger(); + * merger.merge(op1.out) + * .merge(op2.out) + * .merge(op3.out) + * .insertInto(dag, conf); + * + * dag.addModule("Merger", merger); + * + * @param <K> + */ +public class MultipleStreamMerger<K> implements Module +{ + public class Stream + { + DefaultInputPort destPort; + DefaultOutputPort sourcePort; + String name; + + public Stream(String name, DefaultOutputPort sourcePort, DefaultInputPort destPort) + { + this.destPort = destPort; + this.sourcePort = sourcePort; + this.name = name; + } + } + + public class NamedMerger + { + StreamMerger<K> merger; + String name; + + public NamedMerger(String name, StreamMerger<K> merger) + { + this.merger = merger; + this.name = name; + } + } + + private int streamCount = 0; + ArrayList<DefaultOutputPort<K>> streamsToMerge = new ArrayList<>(); + public transient ProxyOutputPort<K> streamOutput = new ProxyOutputPort<>(); + + /** + * Used to define all the sources to be merged into a single stream. + * + * @param sourcePort - The output port from the upstream operator that provides data + * @return The updated MultipleStreamMerger object that tracks which streams should be unified. + */ + public MultipleStreamMerger<K> merge(DefaultOutputPort<K> sourcePort) + { + streamsToMerge.add(sourcePort); + return this; + } + + /** + * To merge more than two streams at a time, we construct a binary tree of thread-local StreamMerger operators + * E.g. + * Tier 0 Tier 1 Tier 2 + * Stream 1 -> StreamMerger_1 -> Stream 2 -> StreamMerger_Final -> Out + * Stream 3 -> StreamMerger_2 -> Stream 4 -> + * + * This function updates the provided DAG with the relevant streams. + */ + public void mergeStreams(DAG dag, Configuration conf) + { + if (streamsToMerge.size() < 2) { + throw new IllegalArgumentException("Not enough streams to merge, at least two streams must be selected for " + "merging with `.merge()`."); + } + ArrayList<Stream> streamsToAddToDag = new ArrayList<>(); + ArrayList<NamedMerger> operatorsToAdd = new ArrayList<>(); + // Determine operators and streams to add to the DAG + constructMergeTree(streamsToAddToDag, operatorsToAdd); + for (NamedMerger m : operatorsToAdd) { + dag.addOperator(m.name, m.merger); + } + for (Stream s : streamsToAddToDag) { + dag.addStream(s.name, s.sourcePort,	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2106

#	Requirement	Questions	Answered	Successfully Answered	URL
20547	APEXMALHAR-2106	library/src/main/java/com/datatorrent/lib/stream/MultipleStreamMerger.java --- @@ -0,0 +1,277 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.stream; + +import java.util.ArrayList; +import java.util.Iterator; + +import org.apache.hadoop.conf.Configuration; + +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Module; + +/** + * Module that adds functionality to bypass the platform limitations of combining more than two streams at a time with + * Stream Merger. + * + * Usage: + * + * dag.addOperator("Stream_1", op1); + * dag.addOperator("Stream_2", op2); + * dag.addOperator("Stream_3", op3); + * + * MultipleStreamMerger merger = new MultipleStreamMerger(); + * merger.merge(op1.out) + * .merge(op2.out) + * .merge(op3.out) + * .insertInto(dag, conf); + * + * dag.addModule("Merger", merger); + * + * @param <K> + */ +public class MultipleStreamMerger<K> implements Module +{ + public class Stream + { + DefaultInputPort destPort; + DefaultOutputPort sourcePort; + String name; + + public Stream(String name, DefaultOutputPort sourcePort, DefaultInputPort destPort) + { + this.destPort = destPort; + this.sourcePort = sourcePort; + this.name = name; + } + } + + public class NamedMerger + { + StreamMerger<K> merger; + String name; + + public NamedMerger(String name, StreamMerger<K> merger) + { + this.merger = merger; + this.name = name; + } + } + + private int streamCount = 0; + ArrayList<DefaultOutputPort<K>> streamsToMerge = new ArrayList<>(); + public transient ProxyOutputPort<K> streamOutput = new ProxyOutputPort<>(); + + /** + * Used to define all the sources to be merged into a single stream. + * + * @param sourcePort - The output port from the upstream operator that provides data + * @return The updated MultipleStreamMerger object that tracks which streams should be unified. + */ + public MultipleStreamMerger<K> merge(DefaultOutputPort<K> sourcePort) + { + streamsToMerge.add(sourcePort); + return this; + } + + /** + * To merge more than two streams at a time, we construct a binary tree of thread-local StreamMerger operators + * E.g. + * Tier 0 Tier 1 Tier 2 + * Stream 1 -> StreamMerger_1 -> Stream 2 -> StreamMerger_Final -> Out + * Stream 3 -> StreamMerger_2 -> Stream 4 -> + * + * This function updates the provided DAG with the relevant streams. + */ + public void mergeStreams(DAG dag, Configuration conf) + { + if (streamsToMerge.size() < 2) { + throw new IllegalArgumentException("Not enough streams to merge, at least two streams must be selected for " + + "merging with `.merge()`."); + } + + ArrayList<Stream> streamsToAddToDag = new ArrayList<>(); + ArrayList<NamedMerger> operatorsToAdd = new ArrayList<>(); + + // Determine operators and streams to add to the DAG + constructMergeTree(streamsToAddToDag, operatorsToAdd); + + for (NamedMerger m : operatorsToAdd) { + dag.addOperator(m.name, m.merger); + } + + for (Stream s : streamsToAddToDag) { + dag.addStream(s.name, s.sourcePort,	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2106
20548	APEXMALHAR-2106	library/src/main/java/com/datatorrent/lib/stream/MultipleStreamMerger.java --- @@ -0,0 +1,277 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.stream; + +import java.util.ArrayList; +import java.util.Iterator; + +import org.apache.hadoop.conf.Configuration; + +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Module; + +/** + * Module that adds functionality to bypass the platform limitations of combining more than two streams at a time with + * Stream Merger. + * + * Usage: + * + * dag.addOperator("Stream_1", op1); + * dag.addOperator("Stream_2", op2); + * dag.addOperator("Stream_3", op3); + * + * MultipleStreamMerger merger = new MultipleStreamMerger(); + * merger.merge(op1.out) + * .merge(op2.out) + * .merge(op3.out) + * .insertInto(dag, conf); + * + * dag.addModule("Merger", merger); + * + * @param <K> + */ +public class MultipleStreamMerger<K> implements Module +{ + public class Stream + { + DefaultInputPort destPort; + DefaultOutputPort sourcePort; + String name; + + public Stream(String name, DefaultOutputPort sourcePort, DefaultInputPort destPort) + { + this.destPort = destPort; + this.sourcePort = sourcePort; + this.name = name; + } + } + + public class NamedMerger + { + StreamMerger<K> merger; + String name; + + public NamedMerger(String name, StreamMerger<K> merger) + { + this.merger = merger; + this.name = name; + } + } + + private int streamCount = 0; + ArrayList<DefaultOutputPort<K>> streamsToMerge = new ArrayList<>(); + public transient ProxyOutputPort<K> streamOutput = new ProxyOutputPort<>(); + + /** + * Used to define all the sources to be merged into a single stream. + * + * @param sourcePort - The output port from the upstream operator that provides data + * @return The updated MultipleStreamMerger object that tracks which streams should be unified. + */ + public MultipleStreamMerger<K> merge(DefaultOutputPort<K> sourcePort) + { + streamsToMerge.add(sourcePort); + return this; + } + + /** + * To merge more than two streams at a time, we construct a binary tree of thread-local StreamMerger operators + * E.g. + * Tier 0 Tier 1 Tier 2 + * Stream 1 -> StreamMerger_1 -> Stream 2 -> StreamMerger_Final -> Out + * Stream 3 -> StreamMerger_2 -> Stream 4 -> + * + * This function updates the provided DAG with the relevant streams. + */ + public void mergeStreams(DAG dag, Configuration conf) + { + if (streamsToMerge.size() < 2) { + throw new IllegalArgumentException("Not enough streams to merge, at least two streams must be selected for " + + "merging with `.merge()`."); + } + + ArrayList<Stream> streamsToAddToDag = new ArrayList<>(); + ArrayList<NamedMerger> operatorsToAdd = new ArrayList<>(); + + // Determine operators and streams to add to the DAG + constructMergeTree(streamsToAddToDag, operatorsToAdd); + + for (NamedMerger m : operatorsToAdd) { + dag.addOperator(m.name, m.merger); + } + + for (Stream s : streamsToAddToDag) { + dag.addStream(s.name, s.sourcePort,	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2106

#	Requirement	Questions	Answered	Successfully Answered	URL
20549	APEXMALHAR-2106	library/src/test/java/com/datatorrent/lib/stream/MultipleStreamMergerTest.java --- @@ -0,0 +1,145 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.stream; + +import java.io.IOException; +import java.util.ArrayList; + +import javax.validation.ConstraintViolationException; + +import org.junit.Assert; +import org.junit.Before; +import org.junit.Test; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.hadoop.conf.Configuration; + +import com.datatorrent.api.DAG; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.lib.io.ConsoleOutputOperator; +import com.datatorrent.lib.testbench.RandomWordGenerator; + +import static org.junit.Assert.assertEquals; + +public class MultipleStreamMergerTest { + private static Logger LOG = LoggerFactory.getLogger(MultipleStreamMergerTest.class); + + StreamMerger<byte[]> finalMerger; + ArrayList<MultipleStreamMerger<byte[]>.Stream> streamsToAddToDag; + ArrayList<MultipleStreamMerger<byte[]>.NamedMerger> operatorsToAdd; + + @Before + public void setUp() throws Exception { + finalMerger = new StreamMerger<>(); + streamsToAddToDag = new ArrayList<>(); + operatorsToAdd = new ArrayList<>(); + } + + @Test + public void mergeTwoStreams() { + RandomWordGenerator randomWordGenerator = new RandomWordGenerator(); + RandomWordGenerator randomWordGenerator2 = new RandomWordGenerator(); + randomWordGenerator.setTuplesPerWindow(1); + randomWordGenerator2.setTuplesPerWindow(1); + MultipleStreamMerger<byte[]> merger = new MultipleStreamMerger<>(); + merger.merge(randomWordGenerator.output) + .merge(randomWordGenerator2.output); + merger.constructMergeTree(streamsToAddToDag, operatorsToAdd); + assertEquals("Count of created streams", 2, streamsToAddToDag.size()); + assertEquals("Count of created operators", 1, operatorsToAdd.size()); + // Next check actual connections + assertEquals("Generator 1 stream", randomWordGenerator.output, streamsToAddToDag.get(0).sourcePort); + assertEquals("Generator 2 stream", randomWordGenerator2.output, streamsToAddToDag.get(1).sourcePort); + assertEquals("Final operator input_1", operatorsToAdd.get(0).merger.data1, streamsToAddToDag.get(0).destPort); + assertEquals("Final operator input_2", operatorsToAdd.get(0).merger.data2, streamsToAddToDag.get(1).destPort); + } + + @Test(expected = IllegalArgumentException.class) + public void mergeOneStream() { + RandomWordGenerator randomWordGenerator = new RandomWordGenerator(); + MultipleStreamMerger<byte[]> merger = new MultipleStreamMerger<>(); + merger.merge(randomWordGenerator.output); + merger.constructMergeTree(streamsToAddToDag, operatorsToAdd); + } + + @Test(expected = IllegalArgumentException.class) + public void mergeZeroStream() { + RandomWordGenerator randomWordGenerator = new RandomWordGenerator(); + MultipleStreamMerger<byte[]> merger = new MultipleStreamMerger<>(); + merger.constructMergeTree(streamsToAddToDag,	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2106
20550	APEXMALHAR-2106	library/src/main/java/com/datatorrent/lib/stream/MultipleStreamMerger.java --- @@ -0,0 +1,215 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.stream; + +import java.util.ArrayList; +import java.util.Queue; + +import org.eclipse.jetty.util.ArrayQueue; + +import org.apache.hadoop.conf.Configuration; + +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; + +/** + * A helper class that adds functionality to bypass the platform limitations of combining more than two + * streams at a time with Stream Merger. + * + * Usage: + * + * dag.addOperator("Stream_1", op1); + * dag.addOperator("Stream_2", op2); + * dag.addOperator("Stream_3", op3); + * + * MultipleStreamMerger merger = new MultipleStreamMerger(); + * DefaultOutputPort streamOut = merger.merge(op1.out) + * .merge(op2.out) + * .merge(op3.out) + * .mergeStreams(dag, conf); + * + * dag.addStream("merger-counter", streamOut, counterOp.counter); + * + * @param <K> + */ +public class MultipleStreamMerger<K> +{ + public class Stream + { + DefaultInputPort<K> destPort; + SourcedOutputPort sourcePort; + String name; + + public Stream(String name, SourcedOutputPort sourcePort, DefaultInputPort<K> destPort) + { + this.destPort = destPort; + this.sourcePort = sourcePort; + this.name = name; + } + } + + public class NamedMerger + { + StreamMerger<K> merger; + String name; + + public NamedMerger(String name, StreamMerger<K> merger) + { + this.merger = merger; + this.name = name; + } + } + + /** + * A simple class to allow us to track whether the port to be merged is internal (allowing it to be thread local) or external + */ + public class SourcedOutputPort + { + boolean internal; + DefaultOutputPort<K> port; + + public SourcedOutputPort(DefaultOutputPort<K> port) + { + this.internal = false; + this.port = port; + } + + public SourcedOutputPort(boolean internal, DefaultOutputPort<K> port) + { + this.internal = internal; + this.port = port; + } + } + + ArrayList<DefaultOutputPort<K>> streamsToMerge = new ArrayList<>(); + private DefaultOutputPort<K> streamOutput = new DefaultOutputPort<>(); + + /** + * Used to define all the sources to be merged into a single stream. + * + * @param sourcePort - The output port from the upstream operator that provides data + * @return The updated MultipleStreamMerger object that tracks which streams should be unified. + */ + public MultipleStreamMerger<K> merge(DefaultOutputPort<K> sourcePort) + { + streamsToMerge.add(sourcePort); + return this; + } + + /** + * Given the streams to merge have been selected with {@link #merge(DefaultOutputPort)}, create a subDAG and add it + * to an existing DAG. + * + * To merge more than two streams at a time, we construct a tiered hierarchy of thread-local StreamMerger operators + * E.g. + * Stream 0 -> StreamMerger_1 -> + * Stream 1 -> StreamMerger_Final -> Out + * Stream 2 -> StreamMerger_2 -> + * Stream 3 -> + * Note that we don't use the populateDAG function because	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2106

#	Requirement	Questions	Answered	Successfully Answered	URL
20551	APEXMALHAR-2105	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/306#discussion_r65411560 --- Diff: contrib/src/main/java/com/datatorrent/contrib/formatter/CsvFormatter.java --- @@ -58,215 +68,156 @@ public class CsvFormatter extends Formatter<String> { - private ArrayList<Field> fields; - @NotNull - protected String classname; - @NotNull - protected int fieldDelimiter; - protected String lineDelimiter; + /** + * Names of all the fields in the same order that would appear in output + * records + */ + private transient String[] nameMapping; + /** + * Cell processors are an integral part of reading and writing with Super CSV + * they automate the data type conversions, and enforce constraints. */ + private transient CellProcessor[] processors; + /** + * Writing preferences that are passed through schema */ + private transient CsvPreference preference; + /** + * Contents of the schema.Schema is specified in a json format as per + * {@link DelimitedSchema} */ @NotNull - protected String fieldInfo; + private String schema; + /** + * Schema is read into this object to access fields */ + private transient DelimitedSchema delimitedParserSchema; - public enum FIELD_TYPE - { - BOOLEAN, DOUBLE, INTEGER, FLOAT, LONG, SHORT, CHARACTER, STRING, DATE - }; + /** + * metric to keep count of number of tuples emitted on error port port */ + @AutoMetric + private long errorTupleCount; + /** + * metric to keep count of number of tuples emitted on out port */ + @AutoMetric + private long emittedObjectCount; --- End diff -- Can you please make this private?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2105
20552	APEXMALHAR-2105	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/306#discussion_r65411586 --- Diff: contrib/src/main/java/com/datatorrent/contrib/formatter/CsvFormatter.java --- @@ -58,215 +68,156 @@ public class CsvFormatter extends Formatter<String> { - private ArrayList<Field> fields; - @NotNull - protected String classname; - @NotNull - protected int fieldDelimiter; - protected String lineDelimiter; + /** + * Names of all the fields in the same order that would appear in output + * records + */ + private transient String[] nameMapping; + /** + * Cell processors are an integral part of reading and writing with Super CSV + * they automate the data type conversions, and enforce constraints. */ + private transient CellProcessor[] processors; + /** + * Writing preferences that are passed through schema */ + private transient CsvPreference preference; + /** + * Contents of the schema.Schema is specified in a json format as per + * {@link DelimitedSchema} */ @NotNull - protected String fieldInfo; + private String schema; + /** + * Schema is read into this object to access fields */ + private transient DelimitedSchema delimitedParserSchema; - public enum FIELD_TYPE - { - BOOLEAN, DOUBLE, INTEGER, FLOAT, LONG, SHORT, CHARACTER, STRING, DATE - }; + /** + * metric to keep count of number of tuples emitted on error port port */ + @AutoMetric + private long errorTupleCount; + /** + * metric to keep count of number of tuples emitted on out port */ + @AutoMetric + private long emittedObjectCount; - protected transient String[] nameMapping; - protected transient CellProcessor[] processors; - protected transient CsvPreference preference; + /** + * metric to keep count of number of tuples emitted on input port */ + @AutoMetric + protected long incomingTuplesCount; --- End diff -- Can you make this private?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2105
20553	APEXMALHAR-2105	Diff: contrib/src/main/java/com/datatorrent/contrib/formatter/CsvFormatter.java --- @@ -58,215 +68,156 @@ public class CsvFormatter extends Formatter<String> { - private ArrayList<Field> fields; - @NotNull - protected String classname; - @NotNull - protected int fieldDelimiter; - protected String lineDelimiter; + /** + * Names of all the fields in the same order that would appear in output + * records + */ + private transient String[] nameMapping; + /** + * Cell processors are an integral part of reading and writing with Super CSV + * they automate the data type conversions, and enforce constraints. */ + private transient CellProcessor[] processors; + /** + * Writing preferences that are passed through schema */ + private transient CsvPreference preference; + /** + * Contents of the schema.Schema is specified in a json format as per + * {@link DelimitedSchema} */ @NotNull - protected String fieldInfo; + private String schema; + /** + * Schema is read into this object to access fields */ + private transient DelimitedSchema delimitedParserSchema; - public enum FIELD_TYPE - { - BOOLEAN, DOUBLE, INTEGER, FLOAT, LONG, SHORT, CHARACTER, STRING, DATE - }; + /** + * metric to keep count of number of tuples emitted on error port port */ + @AutoMetric + private long errorTupleCount; + /** + * metric to keep count of number of tuples emitted on out port */ + @AutoMetric + private long emittedObjectCount; - protected transient String[] nameMapping; - protected transient CellProcessor[] processors; - protected transient CsvPreference preference; + /** + * metric to keep count of number of tuples emitted on input port */ + @AutoMetric + protected long incomingTuplesCount; + public CsvFormatter() { - fields = new ArrayList<Field>(); - fieldDelimiter = ','; - lineDelimiter = "\r\n"; } + @Override + public void beginWindow(long windowId) { + super.beginWindow(windowId); + errorTupleCount = 0; + emittedObjectCount = 0; + incomingTuplesCount = 0; } + @Override public void setup(Context.OperatorContext context) { super.setup(context); - - //fieldInfo information - fields = new ArrayList<Field>(); - String[] fieldInfoTuple = fieldInfo.split(","); - for (int i = 0; i < fieldInfoTuple.length; i++) { - String[] fieldTuple = fieldInfoTuple[i].split(":"); - Field field = new Field(); - field.setName(fieldTuple[0]); - String[] typeFormat = fieldTuple[1].split("\\|"); - field.setType(typeFormat[0].toUpperCase()); - if (typeFormat.length > 1) { - field.setFormat(typeFormat[1]); - } - getFields().add(field); - } - preference = new CsvPreference.Builder('"', fieldDelimiter, lineDelimiter).build(); - int countKeyValue = getFields().size(); - nameMapping = new String[countKeyValue]; - processors = new CellProcessor[countKeyValue]; - initialise(nameMapping, processors); + delimitedParserSchema = new DelimitedSchema(schema); + preference = new CsvPreference.Builder(delimitedParserSchema.getQuoteChar(), delimitedParserSchema.getDelimiterChar(), delimitedParserSchema.getLineDelimiter()).build(); + nameMapping = delimitedParserSchema.getFieldNames().toArray(new String[delimitedParserSchema.getFieldNames().size()]); + processors = getProcessor(delimitedParserSchema.getFields()); } - private void initialise(String[] nameMapping, CellProcessor[] processors) + /** + * Returns array of cellprocessors, one for each field */ + private CellProcessor[] getProcessor(List<Field> fields) { - for (int i = 0; i < getFields().size(); i++) { - FIELD_TYPE type = getFields().get(i).type; - nameMapping[i] = getFields().get(i).name; - if (type == FIELD_TYPE.DATE) { - String dateFormat = getFields().get(i).format; - processors[i] = new Optional(new FmtDate(dateFormat == null ? "dd/MM/yyyy" : dateFormat)); - } - CellProcessor[] processor = new CellProcessor[fields.size()]; - int fieldCount = 0; - for (Field field : fields) { - if (field.getType() == FieldType.DATE) { - String format =	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2105

#	Requirement	Questions	Answered	Successfully Answered	URL
20554	APEXMALHAR-2105	Diff: contrib/src/main/java/com/datatorrent/contrib/formatter/CsvFormatter.java --- @@ -58,215 +68,156 @@ public class CsvFormatter extends Formatter<String> { - private ArrayList<Field> fields; - @NotNull - protected String classname; - @NotNull - protected int fieldDelimiter; - protected String lineDelimiter; + /** + * Names of all the fields in the same order that would appear in output + * records + */ + private transient String[] nameMapping; + /** + * Cell processors are an integral part of reading and writing with Super CSV + * they automate the data type conversions, and enforce constraints. + */ + private transient CellProcessor[] processors; + /** + * Writing preferences that are passed through schema + */ + private transient CsvPreference preference; + /** + * Contents of the schema.Schema is specified in a json format as per + * {@link DelimitedSchema} + */ @NotNull - protected String fieldInfo; + private String schema; + /** + * Schema is read into this object to access fields + */ + private transient DelimitedSchema delimitedParserSchema; - public enum FIELD_TYPE - { - BOOLEAN, DOUBLE, INTEGER, FLOAT, LONG, SHORT, CHARACTER, STRING, DATE - }; + /** + * metric to keep count of number of tuples emitted on error port port + */ + @AutoMetric + protected long errorTupleCount; + /** + * metric to keep count of number of tuples emitted on out port + */ + @AutoMetric + protected long emittedObjectCount; - protected transient String[] nameMapping; - protected transient CellProcessor[] processors; - protected transient CsvPreference preference; + /** + * metric to keep count of number of tuples emitted on input port + */ + @AutoMetric + protected long incomingTuplesCount; - public CsvFormatter() - { - fields = new ArrayList<Field>(); - fieldDelimiter = ','; - lineDelimiter = "\r\n"; + @Override + public void beginWindow(long windowId) + { + super.beginWindow(windowId); + errorTupleCount = 0; + emittedObjectCount = 0; + incomingTuplesCount = 0; } @Override public void setup(Context.OperatorContext context) { super.setup(context); - - - //fieldInfo information - fields = new ArrayList<Field>(); - String[] fieldInfoTuple = fieldInfo.split(","); - for (int i = 0; i < fieldInfoTuple.length; i++) { - String[] fieldTuple = fieldInfoTuple[i].split(":"); - Field field = new Field(); - field.setName(fieldTuple[0]); - String[] typeFormat = fieldTuple[1].split("\\|"); - field.setType(typeFormat[0].toUpperCase()); - if (typeFormat.length > 1) { - field.setFormat(typeFormat[1]); - } - getFields().add(field); - } - preference = new CsvPreference.Builder('"', fieldDelimiter, lineDelimiter).build(); - int countKeyValue = getFields().size(); - nameMapping = new String[countKeyValue]; - processors = new CellProcessor[countKeyValue]; - initialise(nameMapping, processors); + delimitedParserSchema = new DelimitedSchema(schema); + preference = new CsvPreference.Builder(delimitedParserSchema.getQuoteChar(), delimitedParserSchema.getDelimiterChar(), delimitedParserSchema.getLineDelimiter()).build(); + nameMapping = delimitedParserSchema.getFieldNames().toArray(new String[delimitedParserSchema.getFieldNames().size()]); + processors = getProcessor(delimitedParserSchema.getFields()); } - private void initialise(String[] nameMapping, CellProcessor[] processors) + /** + * Returns array of cellprocessors, one for each field + */ + private CellProcessor[] getProcessor(List<Field> fields) { - for (int i = 0; i < getFields().size(); i++) { - FIELD_TYPE type = getFields().get(i).type; - nameMapping[i] = getFields().get(i).name; - if (type == FIELD_TYPE.DATE) { - String dateFormat = getFields().get(i).format; - processors[i] = new Optional(new FmtDate(dateFormat == null ? "dd/MM/yyyy" : dateFormat)); - CellProcessor[] processor = new CellProcessor[fields.size()]; - int fieldCount = 0; - for (Field field : fields) { - if (field.getType() == FieldType.DATE) { - String format =	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2105
20555	APEXMALHAR-2102	I've been looking into the code. It seems that a lot of the logic within the StatelessPartitioner would go away in a ClonePartitioner. Is there any issue with assigning an identical set of PartitionKeys to each input port / partition combination? Does there even need to be more than one key for other logic to work? Repartitioning logic becomes somewhat undefined. It doesn't make sense to partition based on load, I think the only function needed is the definePartitions() function which simply assigns an arbitrary key assignment to each partition, e.g. a 0 key. Is my understanding correct?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2102
20556	APEXMALHAR-2102	GitHub user ilganeli opened a pull request: https://github.com/apache/incubator-apex-core/pull/344 [APEXCORE-146] Add ClonePartitioner * Created a Clone partitioner similar to the StatelessPartitioner that assigns all data to all partitions * Added a simple unit test suite to test scale up and scale down **Questions** Is there a good place/way to add a test to verify that data is partitioned appropriately? For example, the stateless partitioner assigns keys based on the nearest power of two - is that validated anywhere? How is the serialVersionUUID generated? Does the manual assignment of ports to PartitionKeys need to happen for each Partition? There is a comment in the Partition class that by default all data is sent to all partitions. However, digging into the implementation in depth doesn't show that behavior. Instead, it seems data is assigned based on the mask associated with each Partition. This references: https://issues.apache.org/jira/browse/APEXCORE-146 You can merge this pull request into a Git repository by running: $ git pull https://github.com/ilganeli/incubator-apex-core APEXCORE-146 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/incubator-apex-core/pull/344.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #344 ---- commit e2f50274ccf4e3d1cc515d9770ee9e46a0e2e57b Author: Ilya Ganelin <ilya.ganelin@capitalone.com> Date: 2016-05-25T16:59:55Z Started updates on clone partitioner. commit 92c1eb31521deff1e9314751cf41892e3ed6793c Author: Ilya Ganelin <ilya.ganelin@capitalone.com> Date: 2016-05-26T15:13:58Z Added a new partitioner that replicates data across all partitions by default. ----	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102
20557	APEXMALHAR-2102	Github user gauravgopi123 commented on the pull request: https://github.com/apache/incubator-apex-core/pull/344#issuecomment-221932938 Should this be part of malhar under https://github.com/apache/incubator-apex-malhar/tree/master/library/src/main/java/com/datatorrent/lib/partitioner?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102
20558	APEXMALHAR-2102	Github user ilganeli commented on the pull request: https://github.com/apache/incubator-apex-core/pull/344#issuecomment-221936008 @gauravgopi123 Don't really know - the StatelessPartitioner lives in Apex Common. What is the intended structure?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102
20559	APEXMALHAR-2102	Github user tweise commented on the pull request: https://github.com/apache/incubator-apex-core/pull/344#issuecomment-221943349 @ilganeli I also think this should go into Malhar. The stateless partitioner is directly used by the engine, hence it is in common. Should I move the JIRA to APEXMALHAR?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102
20560	APEXMALHAR-2102	Github user ilganeli commented on the pull request: https://github.com/apache/incubator-apex-core/pull/344#issuecomment-221969005 @tweise Sure, let's move the JIRA to Malhar and I'll move the class to Malhar. Is the approach sound?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102

#	Requirement	Questions	Answered	Successfully Answered	URL
20561	APEXMALHAR-2102	GitHub user ilganeli opened a pull request: https://github.com/apache/incubator-apex-malhar/pull/297 [APEXCORE-146] Add ClonePartitioner **Note:** This is still marked as an Apex issue even though it should live in Malhar. Once the JIRA is updated I'll update title and reference. * Created a Clone partitioner similar to the StatelessPartitioner that assigns all data to all partitions * Added a simple unit test suite to test scale up and scale down **Questions** Is there a good place/way to add a test to verify that data is partitioned appropriately? For example, the stateless partitioner assigns keys based on the nearest power of two - is that validated anywhere? How is the serialVersionUUID generated? Does the manual assignment of ports to PartitionKeys need to happen for each Partition? There is a comment in the Partition class that by default all data is sent to all partitions. However, digging into the implementation in depth doesn't show that behavior. Instead, it seems data is assigned based on the mask associated with each Partition. This references: https://issues.apache.org/jira/browse/APEXCORE-146 You can merge this pull request into a Git repository by running: $ git pull https://github.com/ilganeli/incubator-apex-malhar APEXCORE-146 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/incubator-apex-malhar/pull/297.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #297 ---- commit 74f6fd5113a3d488015a9817790018dd63b5e8cb Author: Ilya Ganelin <ilya.ganelin@capitalone.com> Date: 2016-05-26T19:39:24Z Added a new partitioner that replicates data across all partitions by default. ----	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102
20562	APEXMALHAR-2102	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/315#discussion_r66321909 --- Diff: library/src/test/java/com/datatorrent/lib/partitioner/ClonePartitionerTest.java --- @@ -0,0 +1,314 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.partitioner; + +import java.io.IOException; +import java.util.ArrayList; +import java.util.Collection; +import java.util.List; + +import javax.validation.ConstraintViolationException; + +import org.junit.Assert; +import org.junit.Test; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.collect.Lists; + +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.DefaultPartition; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.Operator; +import com.datatorrent.api.Operator.InputPort; +import com.datatorrent.api.Partitioner; +import com.datatorrent.api.Partitioner.Partition; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.api.StringCodec.Object2String; + +import static junit.framework.TestCase.assertFalse; +import static org.junit.Assert.assertEquals; + +public class ClonePartitionerTest +{ + private static Logger LOG = LoggerFactory.getLogger(ClonePartitionerTest.class); + + private static int numPartitions = 20; + private static Integer countPerWindow = 100; + private static Integer totalReceived = 0; --- End diff -- Is this variable required?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102
20563	APEXMALHAR-2102	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/315#discussion_r66322765 --- Diff: library/src/test/java/com/datatorrent/lib/partitioner/ClonePartitionerTest.java --- @@ -0,0 +1,314 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.partitioner; + +import java.io.IOException; +import java.util.ArrayList; +import java.util.Collection; +import java.util.List; + +import javax.validation.ConstraintViolationException; + +import org.junit.Assert; +import org.junit.Test; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.hadoop.conf.Configuration; + +import com.google.common.collect.Lists; + +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.DefaultPartition; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.Operator; +import com.datatorrent.api.Operator.InputPort; +import com.datatorrent.api.Partitioner; +import com.datatorrent.api.Partitioner.Partition; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.api.StringCodec.Object2String; + +import static junit.framework.TestCase.assertFalse; +import static org.junit.Assert.assertEquals; + +public class ClonePartitionerTest +{ + private static Logger LOG = LoggerFactory.getLogger(ClonePartitionerTest.class); + + private static int numPartitions = 20; + private static Integer countPerWindow = 100; + private static Integer totalReceived = 0; + private static Integer totalSent = 0; --- End diff -- Is this variable required?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102

#	Requirement	Questions	Answered	Successfully Answered	URL
20564	APEXMALHAR-2102	Github user brennonyork commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/315#discussion_r66347754 --- Diff: library/src/main/java/com/datatorrent/lib/partitioner/ClonePartitioner.java --- @@ -0,0 +1,143 @@ + +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.partitioner; + +import java.io.Serializable; +import java.util.ArrayList; +import java.util.Collection; +import java.util.List; +import java.util.Map; +import java.util.Set; + +import javax.validation.constraints.Min; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import com.google.common.collect.Lists; +import com.google.common.collect.Sets; + +import com.datatorrent.api.DefaultPartition; +import com.datatorrent.api.Operator; +import com.datatorrent.api.Operator.InputPort; +import com.datatorrent.api.Partitioner; + +/** + * This is a simple partitioner, which replicates data across all partitions of an operator. + * + * @param <T> The type of the operator + * @since 2.0.0 + */ +public class ClonePartitioner<T extends Operator> implements Partitioner<T>, Serializable +{ + private static final Logger logger = LoggerFactory.getLogger(ClonePartitioner.class); + + private static final long serialVersionUID = 201462371710L; + + /** + * The number of partitions for the default partitioner to create. + */ + @Min(1) + private int partitionCount = 1; + + /** + * This creates a partitioner which creates only one partition. + */ + public ClonePartitioner() + { + } + + /** + * This constructor is used to create the partitioner from a property. + * @param value A string which is an integer of the number of partitions to create + */ + public ClonePartitioner(String value) + { + this(Integer.parseInt(value)); + } + + /** + * This creates a partitioner which creates partitonCount partitions. + * @param partitionCount The number of partitions to create. + */ + public ClonePartitioner(int partitionCount) + { + this.partitionCount = partitionCount; + } + + /** + * This method sets the number of partitions for the ClonePartitioner to create. + * @param partitionCount The number of partitions to create. + */ + public void setPartitionCount(int partitionCount) + { + this.partitionCount = partitionCount; + } + + /** + * This method gets the number of partitions for the ClonePartitioner to create. + * @return The number of partitions to create. + */ + public int getPartitionCount() + { + return partitionCount; + } + + @Override + public Collection<Partition<T>> definePartitions(Collection<Partition<T>> partitions, PartitioningContext context) + { + final int newPartitionCount = DefaultPartition.getRequiredPartitionCount(context, this.partitionCount); + logger.debug("define partitions, partitionCount current {} requested {}", partitions.size(), newPartitionCount); + + //Get a partition + DefaultPartition<T> partition = (DefaultPartition<T>)partitions.iterator().next(); + + ArrayList<Partition<T>> newPartitions; --- End diff -- why not `ArrayList<Partition<T>> newPartitions = Lists.newArrayList();` and remove line 112?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2102
20565	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,247 @@ + +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * + context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100

#	Requirement	Questions	Answered	Successfully Answered	URL
20566	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,247 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100
20567	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,248 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected transient Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100

#	Requirement	Questions	Answered	Successfully Answered	URL
20568	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,248 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected transient Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100
20569	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/AbstractManagedStateInnerJoinOperator.java --- @@ -0,0 +1,354 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.util.Iterator; +import java.util.List; +import java.util.Map; +import java.util.concurrent.ExecutionException; +import java.util.concurrent.Future; + +import org.joda.time.Duration; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.apex.malhar.lib.state.managed.ManagedTimeStateImpl; +import org.apache.apex.malhar.lib.state.managed.ManagedTimeStateMultiValue; +import org.apache.apex.malhar.lib.state.spillable.Spillable; +import org.apache.apex.malhar.lib.state.spillable.SpillableComplexComponent; +import org.apache.apex.malhar.lib.utils.serde.Serde; +import org.apache.hadoop.fs.Path; + +import com.google.common.collect.Maps; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DAG; +import com.datatorrent.api.Operator; +import com.datatorrent.lib.fileaccess.FileAccessFSImpl; +import com.datatorrent.netlet.util.Slice; + +/** + * An abstract implementation of inner join operator over Managed state which extends from + * AbstractInnerJoinOperator. + * + * Properties: + * noOfBuckets: Number of buckets required for Managed state. + * bucketSpanTime: Indicates the length of the time bucket. + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public abstract class AbstractManagedStateInnerJoinOperator<K,T> extends AbstractInnerJoinOperator<K,T> implements Operator.CheckpointNotificationListener, Operator.CheckpointListener,Operator.IdleTimeHandler +{ + private static final transient Logger LOG = LoggerFactory.getLogger(AbstractManagedStateInnerJoinOperator.class); + public static final String stateDir = "managedState"; + public static final String stream1State = "stream1Data"; + public static final String stream2State = "stream2Data"; + private transient Map<JoinEvent<K,T>, Future<List>> waitingEvents = Maps.newLinkedHashMap(); + private int noOfBuckets = 1; + private Long bucketSpanTime; + protected ManagedTimeStateImpl stream1Store; + protected ManagedTimeStateImpl stream2Store; + + /** + * Create Managed states and stores for both the streams. + */ + @Override + public void createStores() + { + stream1Store = new ManagedTimeStateImpl(); + stream2Store = new ManagedTimeStateImpl(); + stream1Store.setNumBuckets(noOfBuckets); + stream2Store.setNumBuckets(noOfBuckets); + if (bucketSpanTime != null) { + stream1Store.getTimeBucketAssigner().setBucketSpan(Duration.millis(bucketSpanTime)); + stream2Store.getTimeBucketAssigner().setBucketSpan(Duration.millis(bucketSpanTime)); + } + stream1Store.getTimeBucketAssigner().setExpireBefore(Duration.millis(getExpiryTime())); + stream2Store.getTimeBucketAssigner().setExpireBefore(Duration.millis(getExpiryTime())); + component = new ManagedSpillableComplexComponent(); +	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100

#	Requirement	Questions	Answered	Successfully Answered	URL
20570	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,247 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * + context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100
20571	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,248 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected transient Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * + context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100

#	Requirement	Questions	Answered	Successfully Answered	URL
20572	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,248 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected transient Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * + context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100
20573	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,248 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected transient Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * + context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100

#	Requirement	Questions	Answered	Successfully Answered	URL
20574	APEXMALHAR-2100	library/src/main/java/com/datatorrent/lib/join/POJOInnerJoinOperator.java --- @@ -0,0 +1,247 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.join; + +import java.lang.reflect.Array; +import java.util.HashMap; +import java.util.Map; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.commons.lang3.ClassUtils; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.StreamCodec; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.api.annotation.OutputPortFieldAnnotation; +import com.datatorrent.lib.util.PojoUtils; + +/** + * Concrete implementation of AbstractManagedStateInnerJoinOperator and receives objects from both streams. + * + * @displayName POJO Inner Join Operator + * @tags join + */ +public class POJOInnerJoinOperator extends AbstractManagedStateInnerJoinOperator<Object,Object> implements Operator.ActivationListener<Context> +{ + private static final transient Logger LOG = LoggerFactory.getLogger(POJOInnerJoinOperator.class); + private transient long timeIncrement; + private transient FieldObjectMap[] inputFieldObjects = (FieldObjectMap[])Array.newInstance(FieldObjectMap.class, 2); + protected Class<?> outputClass; + private long time = System.currentTimeMillis(); + + @OutputPortFieldAnnotation(schemaRequired = true) + public final transient DefaultOutputPort<Object> outputPort = new DefaultOutputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + outputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input1 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[0].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + @Override + public void process(Object tuple) + { + processTuple(tuple,true); + } + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(true); + } + }; + + @InputPortFieldAnnotation(schemaRequired = true) + public transient DefaultInputPort<Object> input2 = new DefaultInputPort<Object>() + { + @Override + public void setup(Context.PortContext context) + { + inputFieldObjects[1].inputClass = context.getValue(Context.PortContext.TUPLE_CLASS); + } + @Override + public void process(Object tuple) + { + processTuple(tuple,false); + } + @Override + public StreamCodec<Object> getStreamCodec() + { + return getInnerJoinStreamCodec(false); + } + }; + + @Override + public void setup(Context.OperatorContext context) + { + timeIncrement = context.getValue(Context.OperatorContext.APPLICATION_WINDOW_COUNT) * context.getValue(Context.DAGContext.STREAMING_WINDOW_SIZE_MILLIS); + super.setup(context); + for (int i = 0; i < 2; i++) { + inputFieldObjects[i] = new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2100
20575	APEXMALHAR-2099	The biggest lack we presently have seems to be GroupByKey. Implementing GroupByKey will need either a new operator (that emits a map of {Key -> List[Value]} on window boundaries) or a pairing with the ApexStream implementation (that an incoming stream is split up into multiple new streams) (Preferred). Implementing the flatten operator should be straightforward. In terms of windowing and triggering, I think Siyuan's present API which provides the "with" construct, should be more than sufficient for specifying windowing or trigger semantics that can then be translated by the ApexStream implementation. [~thw][~siyuan] Thoughts? Which route would you recommend?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2099
20576	APEXMALHAR-2099	[~ilganeli] When we started working on high level API, we asked the question. We already have Operator API which supports a much greater variety of transformations, why do we need high-level API which would definitely sacrifice some sort of flexibility. Beam API is kind of in the middle of low level API like our Operator API and some high level API like Flink/Spark. A main reason is we want to make another API that is easy to start with and also easy to use in most cases. Flink, Spark and even JDK 8 pick same flavour for some reason. Think about some developer who's new to parallel stream processing. How can they find out which existing transformation library fit into their needs? By search for all sub classes of PTransform and use apply, apply, apply to chain them together or look at the first-order self-descriptive methods(map, filter, flatmap...) in one class? If I'm a beginner I definitely prefer the facade pattern. I'm not saying Beam API is ugly, if they can have only one API. They have no choice but the full flexibility to support Google Data Flow Model.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2099
20577	APEXMALHAR-2096	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/283#discussion_r63816258 --- Diff: library/src/main/java/com/datatorrent/lib/io/fs/FSInputModule.java --- @@ -107,6 +110,9 @@ public void populateDAG(DAG dag, Configuration conf) dag.setAttribute(blockReader, Context.OperatorContext.PARTITIONER, new StatelessPartitioner<FSSliceReader>(readersCount)); fileSplitter.setBlocksThreshold(readersCount); } + if (blocksThreshold != 0) { --- End diff -- Do you need this check as blocksThreshold is always >=1 according to constraint?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2096
20578	APEXMALHAR-2096	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/283#discussion_r63816277 --- Diff: library/src/main/java/com/datatorrent/lib/io/fs/FSInputModule.java --- @@ -57,6 +58,8 @@ private long blockSize; private boolean sequencialFileRead = false; private int readersCount; + @Min(1) + protected int blocksThreshold; --- End diff -- why not initialize it to 1 so that user don't have to set anything if user goes with default value	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2096
20579	APEXMALHAR-2096	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/287#discussion_r64828460 --- Diff: library/src/main/java/com/datatorrent/lib/io/fs/AbstractFileSplitter.java --- @@ -53,6 +53,7 @@ * This is a threshold on the no. of blocks emitted per window. A lot of blocks emitted * per window can overwhelm the downstream operators. This setting helps to control that. */ + @NotNull --- End diff -- Is this required?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2096
20580	APEXMALHAR-2096	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/287#discussion_r64828487 --- Diff: library/src/main/java/com/datatorrent/lib/io/fs/FSInputModule.java --- @@ -59,7 +59,7 @@ private boolean sequencialFileRead = false; private int readersCount; @Min(1) - protected int blocksThreshold; + protected Integer blocksThreshold; --- End diff -- Is this required?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2096
20581	APEXMALHAR-2096	Github user gauravgopi123 commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/287#discussion_r64840077 --- Diff: library/src/main/java/com/datatorrent/lib/io/fs/FSInputModule.java --- @@ -59,7 +59,7 @@ private boolean sequencialFileRead = false; private int readersCount; @NotNull @Min(1) - protected int blocksThreshold; + protected Integer blocksThreshold; --- End diff -- why Integer type and not primitive int type?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2096

#	Requirement	Questions	Answered	Successfully Answered	URL
20582	APEXMALHAR-2094	Sandeep, are you planning to add more functionality based on the datasketches library? If so, would it make sense to create a separate module for it?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2094
20583	APEXMALHAR-2094	Thomas Weise Thomas, I do intend to add more functionality. DataSketches includes algorithms like hyperloglog, and heavy hitters / frequent items and a "theta sketch" framework that will allow set operations on multiple streams, and I want to build Apex operators for each of these. By "module", do you mean a package like {{statistics}} or {{algo}} in the Malhar library? If so, yes, it might make sense to move this operator from {{statistics}} into a {{sketches}} package that will encapsulate all stream sketching algorithms.	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2094
20584	APEXMALHAR-2094	[~thw] Thomas, I created a new “sketches” module. However, working in this module, IntelliJ is unable to import com.datatorrent.lib.util.TestUtils even though the sketches pom.xml contains malhar-library as a dependency. Where should this be imported from?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2094
20585	APEXMALHAR-2094	sketches/src/main/java/org/apache/apex/malhar/sketches/QuantilesEstimator.java --- @@ -0,0 +1,182 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.sketches; + +import com.yahoo.sketches.quantiles.QuantilesSketch; + +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.OperatorAnnotation; +import com.datatorrent.common.util.BaseOperator; + +/** + * An implementation of BaseOperator that computes a "sketch" (a representation of the probability distribution using + * a low memory footprint) of the incoming numeric data, and evaluates/outputs the cumulative distribution function and + * quantiles of the probability distribution. Leverages the quantiles sketch implementation from the Yahoo Datasketches + * Library. + * <p/> + * Input Port(s) : + * data : Data values input port. + * + * Output Port(s) : + * cdfOutput : cumulative distribution function output port. + * quantilesOutput : quantiles output port. + * + * Partitions : No, no will yield wrong results. + * + */ +@OperatorAnnotation(partitionable = false) +public class QuantilesEstimator extends BaseOperator +{ + /** + * Output port that emits cdf estimated at the current data point + */ + public final transient DefaultOutputPort<Double> cdfOutput = new DefaultOutputPort<>(); + /** + * Emits quantiles of stream seen thus far + */ + public final transient DefaultOutputPort<double[]> quantilesOutput = new DefaultOutputPort<>(); + /** + * Emits probability masses on specified intervals + */ + public final transient DefaultOutputPort<double[]> pmfOutput = new DefaultOutputPort<>(); + private transient QuantilesSketch quantilesSketch = QuantilesSketch.builder().build(); + /** + * This field determines the specific quantiles to be calculated. + * Default is set to compute the standard quartiles. + */ + private double[] fractions = {0.0, 0.25, 0.50, 0.75, 1.00}; + /** + * This field determines the intervals on which the probability mass function is computed. + */ + private double[] pmfIntervals = {}; + /** + * This operator computes three different quantities which are output on separate output ports. If not using any of + * these quantities, these variables can be set to avoid unnecessary computation. + */ + private boolean computeCdf = true; + private boolean computeQuantiles = true; + private boolean computePmf = true; + public final transient DefaultInputPort<Double> data = new DefaultInputPort<Double>() + { + @Override + public void process(Double input) + { + quantilesSketch.update(input); + if (computeQuantiles) { + /** + * Computes and emits quantiles of the stream seen thus far + */ + quantilesOutput.emit(quantilesSketch.getQuantiles(fractions)); + } + if (computeCdf) { + /** + * Emits (estimate of the) cumulative distribution function evaluated at the input value, according to the + * sketched probability distribution of the stream seen thus far. + */ + cdfOutput.emit(quantilesSketch.getCDF(new double[]{input})[0]); --- End diff -- Why are you creating an array from the input here? If only passing a single value it	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2094

#	Requirement	Questions	Answered	Successfully Answered	URL
20586	APEXMALHAR-2094	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/301#discussion_r66151198 --- Diff: sketches/src/main/java/org/apache/apex/malhar/sketches/QuantilesEstimator.java --- @@ -0,0 +1,188 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.sketches; +import com.yahoo.sketches.quantiles.QuantilesSketch; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.OperatorAnnotation; +import com.datatorrent.common.util.BaseOperator; + +/** + * An implementation of BaseOperator that computes a "sketch" (a representation of the probability distribution using + * a low memory footprint) of the incoming numeric data, and evaluates/outputs the cumulative distribution function and + * quantiles of the probability distribution. Leverages the quantiles sketch implementation from the Yahoo Datasketches + * Library. + * <p/> + * Input Port(s) : + * data : Data values input port. + * + * Output Port(s) : + * cdfOutput : cumulative distribution function output port. + * quantilesOutput : quantiles output port. + * + * Partitions : No unifier. Merging these sketches is non-trivial. + * + + */ +@OperatorAnnotation(partitionable = false) +public class QuantilesEstimator extends BaseOperator +{ + /** + * Constructor that allows non-default initialization of the quantile sketch object + * + * @param k: Parameter that determines accuracy and memory usage of quantile sketch. See QuantilesSketch + * documentation for details + * @param seed: The quantile sketch algorithm is inherently random. Set seed to 0 for reproducibility in testing, but + * do not set otherwise. + */ + public QuantilesEstimator(int k, short seed) + { + quantilesSketch = QuantilesSketch.builder().setSeed(seed).build(k); + } + + private transient QuantilesSketch quantilesSketch = QuantilesSketch.builder().build(); + + /** + * This field determines the specific quantiles to be calculated. For a stream of numbers, the quantile at a value + * 0 <= p <= 1 is the number x such that a fraction p of the numbers in the sorted stream are less than x. E.g., the + * quantile at p = 0.5 is the median (half the numbers in the stream are less than the median). + * The default is set to compute the standard quartiles (4-quantiles). + */ + private double[] fractions = {0.0, 0.25, 0.50, 0.75, 1.00}; --- End diff -- Does 0.0 make sense here? This quantile would always be 0. Was the default originally intended to be a 4-quantile?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2094
20587	APEXMALHAR-2094	sketches/src/test/java/org/apache/apex/malhar/sketches/QuantilesEstimatorTest.java --- @@ -0,0 +1,204 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.sketches; +import java.io.IOException; +import java.util.ArrayList; +import java.util.Collections; +import java.util.List; +import java.util.Random; +import javax.validation.ConstraintViolationException; +import org.junit.Assert; +import org.junit.Test; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; +import org.apache.hadoop.conf.Configuration; +import com.datatorrent.api.DAG; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.InputOperator; +import com.datatorrent.api.LocalMode; +import com.datatorrent.api.StreamingApplication; +import com.datatorrent.common.util.BaseOperator; +import com.datatorrent.lib.io.ConsoleOutputOperator; +import com.datatorrent.lib.testbench.CollectorTestSink; +import com.datatorrent.lib.util.TestUtils; + +public class QuantilesEstimatorTest +{ + private static Logger LOG = LoggerFactory.getLogger(QuantilesEstimatorTest.class); + + public static class NumberSource extends BaseOperator implements InputOperator + { + public final DefaultOutputPort<Double> output = new DefaultOutputPort<>(); + private Random rand = new Random(1234L); + public NumberSource() {} + @Override + public void emitTuples() + { + output.emit(rand.nextGaussian()); + } + } + + public static class PmfSink extends BaseOperator + { + public final DefaultInputPort<double[]> input = new DefaultInputPort<double[]>() + { + @Override + public void process(double[] tuple) {} + }; + public PmfSink() {} + } + + public static class QuantileSink extends BaseOperator + { + public final DefaultInputPort<double[]> input = new DefaultInputPort<double[]>() + { + @Override + public void process(double[] tuple) {} + }; + public QuantileSink() {} + } + + public static class CdfSink extends BaseOperator + { + public final DefaultInputPort<Double> input = new DefaultInputPort<Double>() + { + @Override + public void process(Double tuple) {} + }; + public CdfSink() {} + } + + @Test + public void testQuantiles() + { + QuantilesEstimator quantilesOp = new QuantilesEstimator(128, (short)12345); + CollectorTestSink<double[]> sink = new CollectorTestSink<>(); + TestUtils.setSink(quantilesOp.quantilesOutput, sink); + Random rand = new Random(1234L); + ArrayList<Double> randArray = new ArrayList<>(); + quantilesOp.setup(null); + quantilesOp.beginWindow(0); + for (int i = 0; i < 101; i++) { + double r = rand.nextGaussian(); + quantilesOp.data.process(r); + randArray.add(r); + } + quantilesOp.endWindow(); + Collections.sort(randArray); + Assert.assertEquals("Captures all computed quantiles", sink.collectedTuples.size(), 101); + Assert.assertTrue("Computes median correctly", randArray.get(50) == sink.collectedTuples.get(100)[2]); + } + + @Test + public void testCDF() + { + QuantilesEstimator quantilesOp = new QuantilesEstimator(128, (short)12345); + CollectorTestSink<Double> sink = new CollectorTestSink<>(); + TestUtils.setSink(quantilesOp.cdfOutput, sink);	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2094
20588	APEXMALHAR-2087	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/289#discussion_r64593672 --- Diff: contrib/src/main/java/com/datatorrent/contrib/hive/FSPojoToHiveOperator.java --- @@ -40,14 +43,14 @@ * @tags fs, hive, database * @since 3.0.0 */ -public class FSPojoToHiveOperator extends AbstractFSRollingOutputOperator<Object> +public class FSPojoToHiveOperator<T> extends AbstractFSRollingOutputOperator<T> --- End diff -- What's the purpose of the type parameter?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2087

#	Requirement	Questions	Answered	Successfully Answered	URL
20589	APEXMALHAR-2087	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/289#discussion_r64996266 --- Diff: hive/pom.xml --- @@ -0,0 +1,265 @@ +<!-- + + Licensed to the Apache Software Foundation (ASF) under one + or more contributor license agreements. See the NOTICE file + distributed with this work for additional information + regarding copyright ownership. The ASF licenses this file + to you under the Apache License, Version 2.0 (the + "License"); you may not use this file except in compliance + with the License. You may obtain a copy of the License at + + http://www.apache.org/licenses/LICENSE-2.0 + + Unless required by applicable law or agreed to in writing, + software distributed under the License is distributed on an + "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + KIND, either express or implied. See the License for the + specific language governing permissions and limitations + under the License. + +--> +<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd"> <modelVersion>4.0.0</modelVersion> + <parent> <groupId>org.apache.apex</groupId> <artifactId>malhar</artifactId> <version>3.5.0-SNAPSHOT</version> </parent> + <artifactId>malhar-hive</artifactId> <name>Apache Apex Malhar Hive Support</name> <packaging>jar</packaging> + <properties> <checkstyle.console>true</checkstyle.console> </properties> + <build> <plugins> <!-- Publish tests jar --> <plugin> --- End diff -- Do we need this here?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2087
20590	APEXMALHAR-2087	hive/pom.xml --- @@ -0,0 +1,265 @@ +<!-- + + Licensed to the Apache Software Foundation (ASF) under one + or more contributor license agreements. See the NOTICE file + distributed with this work for additional information + regarding copyright ownership. The ASF licenses this file + to you under the Apache License, Version 2.0 (the + "License"); you may not use this file except in compliance + with the License. You may obtain a copy of the License at + + http://www.apache.org/licenses/LICENSE-2.0 + + Unless required by applicable law or agreed to in writing, + software distributed under the License is distributed on an + "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + KIND, either express or implied. See the License for the + specific language governing permissions and limitations + under the License. + +--> +<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/xsd/maven-4.0.0.xsd"> <modelVersion>4.0.0</modelVersion> + <parent> <groupId>org.apache.apex</groupId> <artifactId>malhar</artifactId> <version>3.5.0-SNAPSHOT</version> </parent> + <artifactId>malhar-hive</artifactId> <name>Apache Apex Malhar Hive Support</name> <packaging>jar</packaging> + <properties> <checkstyle.console>true</checkstyle.console> </properties> + <build> <plugins> <!-- Publish tests jar --> <plugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-jar-plugin</artifactId> <version>2.4</version> <executions> <execution> <goals> <goal>test-jar</goal> </goals> <phase>package</phase> </execution> </executions> </plugin> <!-- create resource directory for xml javadoc--> <plugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-antrun-plugin</artifactId> <executions> <execution> <id>createJavadocDirectory</id> <phase>generate-resources</phase> <configuration> <tasks> <delete dir="${project.build.directory}/generated-resources/xml-javadoc"/> <mkdir dir="${project.build.directory}/generated-resources/xml-javadoc"/> </tasks> </configuration> <goals> <goal>run</goal> </goals> </execution> </executions> </plugin> <!-- generate javdoc --> <plugin> <groupId>org.apache.maven.plugins</groupId> <artifactId>maven-javadoc-plugin</artifactId> <executions> <!-- generate xml javadoc --> <execution> <id>xml-doclet</id> <phase>generate-resources</phase> <goals> <goal>javadoc</goal> </goals> <configuration> <doclet>com.github.markusbernhardt.xmldoclet.XmlDoclet</doclet> <additionalparam>-d ${project.build.directory}/generated-resources/xml-javadoc -filename ${project.artifactId}-${project.version}-javadoc.xml</additionalparam> <useStandardDocletOptions>false</useStandardDocletOptions> <docletArtifact> <groupId>com.github.markusbernhardt</groupId> <artifactId>xml-doclet</artifactId> <version>1.0.4</version> </docletArtifact> </configuration> </execution> <!-- generate default javadoc jar with custom tags --> <execution> <id>attach-sources</id> <goals> <goal>jar</goal> </goals> <configuration> <skip>true</skip> <tags> <tag> <name>customTag1</name> <placement>a</placement> <head>Custom Tag One:</head> </tag> <tag> <name>customTag2</name> <placement>a</placement> <head>Custom Tag two:</head> </tag> <tag>	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2087
20591	APEXMALHAR-2086	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/298#discussion_r66122516 --- Diff: kafka/src/main/java/org/apache/apex/malhar/kafka/KafkaSinglePortOutputOperator.java --- @@ -0,0 +1,38 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; +import org.apache.kafka.clients.producer.ProducerRecord; +import com.datatorrent.api.DefaultInputPort; +public class KafkaSinglePortOutputOperator<K,V> extends AbstractKafkaOutputOperator --- End diff -- Can you make this evolving?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086

#	Requirement	Questions	Answered	Successfully Answered	URL
20592	APEXMALHAR-2086	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/298#discussion_r66130894 --- Diff: kafka/src/main/java/org/apache/apex/malhar/kafka/AbstractKafkaOutputOperator.java --- @@ -0,0 +1,141 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; + +import java.util.Properties; +import javax.validation.constraints.NotNull; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; +import org.apache.commons.lang3.StringUtils; +import org.apache.kafka.clients.producer.KafkaProducer; +import org.apache.kafka.clients.producer.Producer; +import com.datatorrent.api.Context; +import com.datatorrent.api.Operator; + +public abstract class AbstractKafkaOutputOperator<K, V> implements Operator { + @SuppressWarnings("unused") private static final Logger logger = LoggerFactory.getLogger(AbstractKafkaOutputOperator.class); + private transient Producer<K, V> producer; // K is key partitioner, V is value type + @NotNull private String topic = "topic1"; --- End diff -- Why is default value of "topic" property set to "topic1"... That invalidates the use of "@NotNull"...	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086
20593	APEXMALHAR-2086	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/298#discussion_r66130956 --- Diff: kafka/src/main/java/org/apache/apex/malhar/kafka/AbstractKafkaOutputOperator.java --- @@ -0,0 +1,141 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; + +import java.util.Properties; +import javax.validation.constraints.NotNull; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; +import org.apache.commons.lang3.StringUtils; +import org.apache.kafka.clients.producer.KafkaProducer; +import org.apache.kafka.clients.producer.Producer; +import com.datatorrent.api.Context; +import com.datatorrent.api.Operator; + +public abstract class AbstractKafkaOutputOperator<K, V> implements Operator --- End diff -- Can you add a javadoc for this class?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086
20594	APEXMALHAR-2086	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/298#discussion_r66132271 --- Diff: kafka/src/main/java/org/apache/apex/malhar/kafka/AbstractKafkaOutputOperator.java --- @@ -0,0 +1,141 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; + +import java.util.Properties; +import javax.validation.constraints.NotNull; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; +import org.apache.commons.lang3.StringUtils; +import org.apache.kafka.clients.producer.KafkaProducer; +import org.apache.kafka.clients.producer.Producer; +import com.datatorrent.api.Context; +import com.datatorrent.api.Operator; + +public abstract class AbstractKafkaOutputOperator<K, V> implements Operator { + @SuppressWarnings("unused") private static final Logger logger = LoggerFactory.getLogger(AbstractKafkaOutputOperator.class); --- End diff -- Can you remove the logger is unused?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086

#	Requirement	Questions	Answered	Successfully Answered	URL
20595	APEXMALHAR-2086	kafka/src/main/java/org/apache/apex/malhar/kafka/AbstractExactlyOnceKafkaOutputOperator.java --- @@ -0,0 +1,385 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; +import java.io.IOException; import java.util.ArrayList; import java.util.HashMap; import java.util.HashSet; import java.util.List; import java.util.Map; import java.util.Properties; import java.util.Set; import java.util.concurrent.ExecutionException; import org.slf4j.Logger; import org.slf4j.LoggerFactory; import org.apache.apex.malhar.lib.wal.FSWindowDataManager; import org.apache.apex.malhar.lib.wal.WindowDataManager; import org.apache.kafka.clients.consumer.ConsumerRecord; import org.apache.kafka.clients.consumer.ConsumerRecords; import org.apache.kafka.clients.consumer.KafkaConsumer; import org.apache.kafka.clients.producer.ProducerRecord; import org.apache.kafka.common.PartitionInfo; import org.apache.kafka.common.TopicPartition; import com.datatorrent.api.Context; import com.datatorrent.api.DefaultInputPort; import com.datatorrent.api.Operator; // 1. JavaDoc // 2. Unit Test // 3. Comparator -- not required // 4. Generic Type // 5. remove e.printStackTrace() // 6. Should the Consumer be kept open? +/** * This is a base implementation of a Kafka output operator, * which, in most cases, guarantees to send tuples to Kafka MQ only once. * Subclasses should implement the methods for converting tuples into a format appropriate for Kafka. * * Assuming messages kept in kafka are ordered by either key or value or keyvalue pair * (For example, use timestamps as key), this Kafka OutputOperator always retrieve the last message from MQ as initial offset. * So that replayed message wouldn't be sent to kafka again. * * This is not "perfect exact once" in 2 cases: * 1 Multiple producers produce messages to same kafka partition * 2 You have same message sent out and before kafka synchronized this message among all the brokers, the operator is * started again. * * * Ports: * Input: One input port * Output: No output port * * Properties: * configProperties * * Compile time checks: * Class derived from has to implement 2 methods: * tupleToKeyValue() to convert input tuples to kafka key value objects * compareToLastMsg() to compare incoming tuple with the last received msg in kafka so that the operator could skip the received ones * * Run time checks: * None * * Benchmarks: * TBD * * * @displayName Abstract Exactly Once Kafka Output(0.9.0) * @category Messaging * @tags output operator * * @since 3.5 */ public abstract class AbstractExactlyOnceKafkaOutputOperator<T> extends AbstractKafkaOutputOperator<String, T> implements Operator.CheckpointNotificationListener { private WindowDataManager windowDataManager = new FSWindowDataManager(); private String key; private Integer operatorId; private String appId; private transient Long windowId; private transient Set<T> recoveredTuples = new HashSet<>(); private int KAFKA_CONNECT_ATTEMPT = 10; +	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086
20596	APEXMALHAR-2086	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/298#discussion_r66693682 --- Diff: kafka/src/main/java/org/apache/apex/malhar/kafka/AbstractExactlyOnceKafkaOutputOperator.java --- @@ -0,0 +1,366 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; +import java.io.IOException; import java.util.ArrayList; import java.util.HashMap; import java.util.List; import java.util.Map; import java.util.Properties; import java.util.concurrent.ExecutionException; import org.slf4j.Logger; import org.slf4j.LoggerFactory; import org.apache.apex.malhar.lib.wal.FSWindowDataManager; import org.apache.apex.malhar.lib.wal.WindowDataManager; import org.apache.kafka.clients.consumer.ConsumerRecord; import org.apache.kafka.clients.consumer.ConsumerRecords; import org.apache.kafka.clients.consumer.KafkaConsumer; import org.apache.kafka.clients.producer.ProducerRecord; import org.apache.kafka.common.PartitionInfo; import org.apache.kafka.common.TopicPartition; import com.datatorrent.api.Context; import com.datatorrent.api.Operator; import static org.apache.kafka.clients.consumer.ConsumerConfig.BOOTSTRAP_SERVERS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG; import static org.apache.kafka.clients.producer.ProducerConfig.ACKS_CONFIG; +/** * Kafka output operator with exactly once processing semantics under certain conditions., * Requirement for Exactly Once: * Every message within the Window is unique * This operator uses *Key* to distinguish the messages written by particular instance of the Output operator. * Operator users can only use *value* for storing the data. * * @displayName Abstract Exactly Once Kafka Output(0.9.0) * @category Messaging * @tags output operator * * @since 3.5 --- End diff -- Can you remove @since? They'll get added during release time.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086

#	Requirement	Questions	Answered	Successfully Answered	URL
20597	APEXMALHAR-2086	Github user devtagare commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/298#discussion_r66696137 --- Diff: kafka/src/main/java/org/apache/apex/malhar/kafka/AbstractExactlyOnceKafkaOutputOperator.java --- @@ -0,0 +1,366 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; +import java.io.IOException; import java.util.ArrayList; import java.util.HashMap; import java.util.List; import java.util.Map; import java.util.Properties; import java.util.concurrent.ExecutionException; import org.slf4j.Logger; import org.slf4j.LoggerFactory; import org.apache.apex.malhar.lib.wal.FSWindowDataManager; import org.apache.apex.malhar.lib.wal.WindowDataManager; import org.apache.kafka.clients.consumer.ConsumerRecord; import org.apache.kafka.clients.consumer.ConsumerRecords; import org.apache.kafka.clients.consumer.KafkaConsumer; import org.apache.kafka.clients.producer.ProducerRecord; import org.apache.kafka.common.PartitionInfo; import org.apache.kafka.common.TopicPartition; import com.datatorrent.api.Context; import com.datatorrent.api.Operator; import static org.apache.kafka.clients.consumer.ConsumerConfig.BOOTSTRAP_SERVERS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG; import static org.apache.kafka.clients.producer.ProducerConfig.ACKS_CONFIG; +/** + * Kafka output operator with exactly once processing semantics under certain conditions., + * Requirement for Exactly Once: + * Every message within the Window is unique + * This operator uses *Key* to distinguish the messages written by particular instance of the Output operator. + * Operator users can only use *value* for storing the data. + * @displayName Abstract Exactly Once Kafka Output(0.9.0) + * @category Messaging + * @tags output operator + * @since 3.5 --- End diff -- Why unique records in a window ?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086
20598	APEXMALHAR-2086	kafka/src/main/java/org/apache/apex/malhar/kafka/AbstractExactlyOnceKafkaOutputOperator.java --- @@ -0,0 +1,366 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; +import java.io.IOException; import java.util.ArrayList; import java.util.HashMap; import java.util.List; import java.util.Map; import java.util.Properties; import java.util.concurrent.ExecutionException; import org.slf4j.Logger; import org.slf4j.LoggerFactory; import org.apache.apex.malhar.lib.wal.FSWindowDataManager; import org.apache.apex.malhar.lib.wal.WindowDataManager; import org.apache.kafka.clients.consumer.ConsumerRecord; import org.apache.kafka.clients.consumer.ConsumerRecords; import org.apache.kafka.clients.consumer.KafkaConsumer; import org.apache.kafka.clients.producer.ProducerRecord; import org.apache.kafka.common.PartitionInfo; import org.apache.kafka.common.TopicPartition; import com.datatorrent.api.Context; import com.datatorrent.api.Operator; import static org.apache.kafka.clients.consumer.ConsumerConfig.BOOTSTRAP_SERVERS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG; import static org.apache.kafka.clients.producer.ProducerConfig.ACKS_CONFIG; +/** + * Kafka output operator with exactly once processing semantics under certain conditions., + * Requirement for Exactly Once: + * Every message within the Window is unique + * This operator uses *Key* to distinguish the messages written by particular instance of the Output operator. + * Operator users can only use *value* for storing the data. + * @displayName Abstract Exactly Once Kafka Output(0.9.0) + * @category Messaging + * @tags output operator + * @since 3.5 + */ +@org.apache.hadoop.classification.InterfaceStability.Evolving +public abstract class AbstractExactlyOnceKafkaOutputOperator<T> extends AbstractKafkaOutputOperator<String, T> implements Operator.CheckpointNotificationListener { private transient String key; private transient String appId; private transient Integer operatorId; private transient Long windowId; private transient Map<T, Integer> partialWindowTuples = new HashMap<>(); private transient KafkaConsumer consumer; private WindowDataManager windowDataManager = new FSWindowDataManager(); private final int KAFKA_CONNECT_ATTEMPT = 10; private final String KEY_SEPARATOR = "#"; private final String KEY_SERIALIZER = "org.apache.kafka.common.serialization.StringDeserializer"; private final String VALUE_SERIALIZER = "org.apache.kafka.common.serialization.StringDeserializer"; @Override public void setup(Context.OperatorContext context) { super.setup(context); this.operatorId = context.getId(); this.windowDataManager.setup(context); this.appId = context.getValue(Context.DAGContext.APPLICATION_ID); this.key = appId + KEY_SEPARATOR + (new	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086

#	Requirement	Questions	Answered	Successfully Answered	URL
20599	APEXMALHAR-2086	kafka/src/main/java/org/apache/apex/malhar/kafka/KafkaSinglePortExactlyOnceOutputOperator.java --- @@ -0,0 +1,369 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; +import java.io.IOException; import java.util.ArrayList; import java.util.HashMap; import java.util.List; import java.util.Map; import java.util.Properties; import java.util.concurrent.ExecutionException; import org.slf4j.Logger; import org.slf4j.LoggerFactory; import org.apache.apex.malhar.lib.wal.FSWindowDataManager; import org.apache.apex.malhar.lib.wal.WindowDataManager; import org.apache.kafka.clients.consumer.ConsumerRecord; import org.apache.kafka.clients.consumer.ConsumerRecords; import org.apache.kafka.clients.consumer.KafkaConsumer; import org.apache.kafka.clients.producer.ProducerRecord; import org.apache.kafka.common.PartitionInfo; import org.apache.kafka.common.TopicPartition; import com.datatorrent.api.Context; import com.datatorrent.api.DefaultInputPort; import com.datatorrent.api.Operator; import static org.apache.kafka.clients.consumer.ConsumerConfig.BOOTSTRAP_SERVERS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG; +/** * Kafka output operator with exactly once processing semantics under certain conditions., * * This operator uses *Key* to distinguish the messages written by particular instance of the Output operator. * Operator users can only use *value* for storing the data. * * @displayName Single Port Exactly Once Kafka Output(0.9.0) * @category Messaging * @tags output operator * */ @org.apache.hadoop.classification.InterfaceStability.Evolving public class KafkaSinglePortExactlyOnceOutputOperator<T> extends AbstractKafkaOutputOperator<String, T> implements Operator.CheckpointNotificationListener { private transient String key; private transient String appName; private transient Integer operatorId; private transient Long windowId; private transient Map<T, Integer> partialWindowTuples = new HashMap<>(); private transient KafkaConsumer consumer; + private WindowDataManager windowDataManager = new FSWindowDataManager(); private final int KAFKA_CONNECT_ATTEMPT = 10; private final String KEY_SEPARATOR = "#"; private final String KEY_SERIALIZER = "org.apache.kafka.common.serialization.StringDeserializer"; private final String VALUE_SERIALIZER = "org.apache.kafka.common.serialization.StringDeserializer"; + public final transient DefaultInputPort<T> inputPort = new DefaultInputPort<T>() { @Override public void process(T tuple) { sendTuple(tuple); } }; + @Override public void setup(Context.OperatorContext context) { super.setup(context); + this.operatorId = context.getId(); this.windowDataManager.setup(context); this.appName = context.getValue(Context.DAGContext.APPLICATION_NAME); + this.key = appName + KEY_SEPARATOR + (new Integer(operatorId)); this.consumer = KafkaConsumerInit(); } + @Override public void	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086
20600	APEXMALHAR-2086	Isn't the producer API in 0.9 same as 0.8.x? If so, why do we need this operator. What is different between the existing operator and this? What problems are being addressed?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2086
20601	APEXMALHAR-2086	Github user tweise commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/298#discussion_r69250081 --- Diff: kafka/src/main/java/org/apache/apex/malhar/kafka/KafkaSinglePortExactlyOnceOutputOperator.java --- @@ -0,0 +1,374 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.kafka; +import java.io.IOException; import java.util.ArrayList; import java.util.HashMap; import java.util.List; import java.util.Map; import java.util.Properties; import java.util.concurrent.ExecutionException; import org.slf4j.Logger; import org.slf4j.LoggerFactory; import org.apache.apex.malhar.lib.wal.FSWindowDataManager; import org.apache.apex.malhar.lib.wal.WindowDataManager; import org.apache.kafka.clients.consumer.ConsumerRecord; import org.apache.kafka.clients.consumer.ConsumerRecords; import org.apache.kafka.clients.consumer.KafkaConsumer; import org.apache.kafka.clients.producer.Callback; import org.apache.kafka.clients.producer.ProducerRecord; import org.apache.kafka.clients.producer.RecordMetadata; import org.apache.kafka.common.PartitionInfo; import org.apache.kafka.common.TopicPartition; import com.datatorrent.api.Context; import com.datatorrent.api.DefaultInputPort; import com.datatorrent.api.Operator; import static org.apache.kafka.clients.consumer.ConsumerConfig.BOOTSTRAP_SERVERS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG; import static org.apache.kafka.clients.consumer.ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG; +/** * Kafka output operator with exactly once processing semantics under certain conditions., * * This operator uses *Key* to distinguish the messages written by particular instance of the Output operator. --- End diff -- I think this needs more explanation. What makes it exactly once, are there other restrictions or assumptions?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2086
20602	APEXMALHAR-2085	[~siyuan] Hi, Siyuan - I would like to get the ball rolling on this work. Since you've spent a fair amount of time digging into this, what would you recommend as a starting point? I'm ready to kick it off immediately :)	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2085
20603	APEXMALHAR-2085	The windowed operator base will consist of multiple pluggable components. We will need a PoC to better demarcate those and ensure they fit the use cases (aggregation, inner/outer join, topN/sort etc.) Based on our learning with dimension computation, we should be able to fork off the pluggable storage fairly soon. This should be based on the spillable data structures with mock/in-memory (for testing), Managed state or HDHT as backend: - Review spillable DS interfaces - PoC with in-memory impl - Spillable implementation for HDHT - Spillable impl for ManagedState Windowed operator base class - identify components to cover the various scenarios - accumulating vs discarding - session windows - usage for count, join Further discussion: - watermark handling - should we use the Beam API windowing classes directly	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2085

#	Requirement	Questions	Answered	Successfully Answered	URL
20604	APEXMALHAR-2085	I'm in the process of designing the API. I would like to start a quick discussion on whether we should use API from other libraries for our API. We already know that we need interfaces for things like the window policy, the triggering policy, the storage, the aggregation, functions that extract timestamp and key, etc. There are interfaces from Beam, Guava, etc., that let you specify some of those things. Some of them are complex (like windowing policy) and some of them are simple (like the function interface). Should we mirror (or duplicate) their interfaces in our API or just use their interfaces directly? Please provide your feedback.	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2085
20605	APEXMALHAR-2085	library/src/main/java/org/apache/apex/malhar/lib/window/impl/AbstractWindowedOperator.java --- @@ -0,0 +1,486 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window.impl; + +import java.util.ArrayList; +import java.util.Iterator; +import java.util.List; +import java.util.Map; + +import javax.validation.ValidationException; + +import org.joda.time.Duration; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.apex.malhar.lib.window.Accumulation; +import org.apache.apex.malhar.lib.window.ControlTuple; +import org.apache.apex.malhar.lib.window.TriggerOption; +import org.apache.apex.malhar.lib.window.Tuple; +import org.apache.apex.malhar.lib.window.Window; +import org.apache.apex.malhar.lib.window.WindowOption; +import org.apache.apex.malhar.lib.window.WindowState; +import org.apache.apex.malhar.lib.window.WindowedOperator; +import org.apache.apex.malhar.lib.window.WindowedStorage; + +import org.apache.hadoop.classification.InterfaceStability; + +import com.google.common.base.Function; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.common.util.BaseOperator; + +/** + * This is the abstract windowed operator class that implements most of the windowing, triggering, and accumulating + * concepts. The subclass of this abstract class is supposed to provide the implementation of how the accumulated + * values are stored in the storage. + * + */ +@InterfaceStability.Evolving +public abstract class AbstractWindowedOperator<InputT, OutputT, DataStorageT extends WindowedStorage, AccumulationT extends Accumulation> + extends BaseOperator implements WindowedOperator<InputT> +{ + protected WindowOption windowOption; + protected TriggerOption triggerOption; + protected long allowedLatenessMillis = -1; + protected WindowedStorage<WindowState> windowStateMap; + + private Function<InputT, Long> timestampExtractor; + + private long currentWatermark; + private boolean triggerAtWatermark; + private long earlyTriggerCount; + private long earlyTriggerMillis; + private long lateTriggerCount; + private long lateTriggerMillis; + private long currentDerivedTimestamp = -1; + private long windowWidthMillis; + protected DataStorageT dataStorage; + protected DataStorageT retractionStorage; + protected AccumulationT accumulation; + + private static final transient Logger LOG = LoggerFactory.getLogger(AbstractWindowedOperator.class); + + public final transient DefaultInputPort<Tuple<InputT>> input = new DefaultInputPort<Tuple<InputT>>() + { + @Override + public void process(Tuple<InputT> tuple) + { + processTuple(tuple); + } + }; + + // TODO: This port should be removed when Apex Core has native support for custom control tuples + @InputPortFieldAnnotation(optional = true) + public final transient DefaultInputPort<ControlTuple> controlInput = new DefaultInputPort<ControlTuple>() + { + @Override + public void process(ControlTuple tuple) + { + if (tuple instanceof ControlTuple.Watermark) {	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2085
20606	APEXMALHAR-2085	library/src/main/java/org/apache/apex/malhar/lib/window/impl/AbstractWindowedOperator.java --- @@ -0,0 +1,490 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window.impl; + +import java.util.ArrayList; +import java.util.Iterator; +import java.util.List; +import java.util.Map; + +import javax.validation.ValidationException; + +import org.joda.time.Duration; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.apex.malhar.lib.window.Accumulation; +import org.apache.apex.malhar.lib.window.ControlTuple; +import org.apache.apex.malhar.lib.window.TriggerOption; +import org.apache.apex.malhar.lib.window.Tuple; +import org.apache.apex.malhar.lib.window.Window; +import org.apache.apex.malhar.lib.window.WindowOption; +import org.apache.apex.malhar.lib.window.WindowState; +import org.apache.apex.malhar.lib.window.WindowedOperator; +import org.apache.apex.malhar.lib.window.WindowedStorage; + +import org.apache.hadoop.classification.InterfaceStability; + +import com.google.common.base.Function; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.common.util.BaseOperator; + +/** + * This is the abstract windowed operator class that implements most of the windowing, triggering, and accumulating + * concepts. The subclass of this abstract class is supposed to provide the implementation of how the accumulated + * values are stored in the storage. + * + * @param <InputT> The type of the input tuple + * @param <OutputT> The type of the output tuple + * @param <DataStorageT> The type of the data storage + * @param <AccumulationT> The type of the accumulation + */ +@InterfaceStability.Evolving +public abstract class AbstractWindowedOperator<InputT, OutputT, DataStorageT extends WindowedStorage, AccumulationT extends Accumulation> + extends BaseOperator implements WindowedOperator<InputT> +{ + protected WindowOption windowOption; + protected TriggerOption triggerOption; + protected long allowedLatenessMillis = -1; + protected WindowedStorage<WindowState> windowStateMap; + + private Function<InputT, Long> timestampExtractor; + + private long currentWatermark; + private boolean triggerAtWatermark; + private long earlyTriggerCount; + private long earlyTriggerMillis; + private long lateTriggerCount; + private long lateTriggerMillis; + private long currentDerivedTimestamp = -1; + private long windowWidthMillis; + protected DataStorageT dataStorage; + protected DataStorageT retractionStorage; + protected AccumulationT accumulation; + + private static final transient Logger LOG = LoggerFactory.getLogger(AbstractWindowedOperator.class); + + public final transient DefaultInputPort<Tuple<InputT>> input = new DefaultInputPort<Tuple<InputT>>() + { + @Override + public void process(Tuple<InputT> tuple) + { + processTuple(tuple); + } + }; + + // TODO: This port should be removed when Apex Core has native support for custom control tuples + @InputPortFieldAnnotation(optional = true) + public final	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2085

#	Requirement	Questions	Answered	Successfully Answered	URL
20607	APEXMALHAR-2085	library/src/main/java/org/apache/apex/malhar/lib/window/impl/AbstractWindowedOperator.java --- @@ -0,0 +1,490 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window.impl; + +import java.util.ArrayList; +import java.util.Iterator; +import java.util.List; +import java.util.Map; + +import javax.validation.ValidationException; + +import org.joda.time.Duration; +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.apex.malhar.lib.window.Accumulation; +import org.apache.apex.malhar.lib.window.ControlTuple; +import org.apache.apex.malhar.lib.window.TriggerOption; +import org.apache.apex.malhar.lib.window.Tuple; +import org.apache.apex.malhar.lib.window.Window; +import org.apache.apex.malhar.lib.window.WindowOption; +import org.apache.apex.malhar.lib.window.WindowState; +import org.apache.apex.malhar.lib.window.WindowedOperator; +import org.apache.apex.malhar.lib.window.WindowedStorage; +import org.apache.hadoop.classification.InterfaceStability; + +import com.google.common.base.Function; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.common.util.BaseOperator; + +/** + * This is the abstract windowed operator class that implements most of the windowing, triggering, and accumulating + * concepts. The subclass of this abstract class is supposed to provide the implementation of how the accumulated + * values are stored in the storage. + * + * @param <InputT> The type of the input tuple + * @param <OutputT> The type of the output tuple + * @param <DataStorageT> The type of the data storage + * @param <AccumulationT> The type of the accumulation + */ +@InterfaceStability.Evolving +public abstract class AbstractWindowedOperator<InputT, OutputT, DataStorageT extends WindowedStorage, AccumulationT extends Accumulation> + extends BaseOperator implements WindowedOperator<InputT> +{ + protected WindowOption windowOption; + protected TriggerOption triggerOption; + protected long allowedLatenessMillis = -1; + protected WindowedStorage<WindowState> windowStateMap; + + private Function<InputT, Long> timestampExtractor; + + private long currentWatermark; + private boolean triggerAtWatermark; + private long earlyTriggerCount; + private long earlyTriggerMillis; + private long lateTriggerCount; + private long lateTriggerMillis; + private long currentDerivedTimestamp = -1; + private long windowWidthMillis; + protected DataStorageT dataStorage; + protected DataStorageT retractionStorage; + protected AccumulationT accumulation; + + private static final transient Logger LOG = LoggerFactory.getLogger(AbstractWindowedOperator.class); + + public final transient DefaultInputPort<Tuple<InputT>> input = new DefaultInputPort<Tuple<InputT>>() + { + @Override + public void process(Tuple<InputT> tuple) + { + processTuple(tuple); + } + }; + + // TODO: This port should be removed when Apex Core has native support for custom control tuples + @InputPortFieldAnnotation(optional = true) + public final	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2085
20608	APEXMALHAR-2085	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/319#discussion_r68659163 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/window/TriggerOption.java --- @@ -0,0 +1,278 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window; + +import java.util.ArrayList; +import java.util.Collections; +import java.util.List; + +import org.joda.time.Duration; + +import org.apache.hadoop.classification.InterfaceStability; + +import com.esotericsoftware.kryo.serializers.FieldSerializer; +import com.esotericsoftware.kryo.serializers.JavaSerializer; + +/** + * This class describes how triggers should be fired for each window. + * For each window, a trigger can be fired before the watermark (EARLY), at the watermark (ON_TIME), or after the watermark (LATE). + * If a LATE trigger is specified and the accumulation mode is ACCUMULATING, it is important for the WindowOption to + * specify the allowed lateness because otherwise, all states must be kept in storage. + * + */ +@InterfaceStability.Evolving +public class TriggerOption +{ + public enum AccumulationMode + { + DISCARDING, + ACCUMULATING, + ACCUMULATING_AND_RETRACTING + } + private AccumulationMode accumulationMode = AccumulationMode.DISCARDING; + private boolean firingOnlyUpdatedPanes = false; + + /** + * Whether the trigger should be fired before the watermark, at the watermark, or after the watermark + */ + public enum WatermarkOpt + { + EARLY, --- End diff -- If I understand correctly from bream model, EARLY, ON-TIME & LATE are the attributes of data and/or of the results to be send to the downstream. Its not the attribute of a trigger. A trigger is an action which sends updated to downstream. The data containing that should be marked late, early or on-time. Shouldn't this enum be moved out as its not specific to trigger? Also the terminology of WatermarkOpt is confusion. As I said before these are not the attribute of watermark or trigger.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2085

#	Requirement	Questions	Answered	Successfully Answered	URL
20609	APEXMALHAR-2085	Github user davidyan74 commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/319#discussion_r68675229 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/window/TriggerOption.java --- @@ -0,0 +1,278 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window; + +import java.util.ArrayList; +import java.util.Collections; +import java.util.List; + +import org.joda.time.Duration; + +import org.apache.hadoop.classification.InterfaceStability; + +import com.esotericsoftware.kryo.serializers.JavaSerializer; + +/** + * This class describes how triggers should be fired for each window. + * For each window, a trigger can be fired before the watermark (EARLY), at the watermark (ON_TIME), or after the watermark (LATE). + * If a LATE trigger is specified and the accumulation mode is ACCUMULATING, it is important for the WindowOption to + * specify the allowed lateness because otherwise, all states must be kept in storage. + * + */ +@InterfaceStability.Evolving +public class TriggerOption +{ + public enum AccumulationMode + { + DISCARDING, + ACCUMULATING, + ACCUMULATING_AND_RETRACTING + } + + private AccumulationMode accumulationMode = AccumulationMode.DISCARDING; + private boolean firingOnlyUpdatedPanes = false; + + /** + * Whether the trigger should be fired before the watermark, at the watermark, or after the watermark + */ + public enum WatermarkOpt + { + EARLY, --- End diff -- We can. @siyuanh any problem with that?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2085
20610	APEXMALHAR-2085	library/src/main/java/org/apache/apex/malhar/lib/window/impl/AbstractWindowedOperator.java --- @@ -0,0 +1,490 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window.impl; + +import java.util.ArrayList; +import java.util.Iterator; +import java.util.List; +import java.util.Map; + +import javax.validation.ValidationException; + +import org.joda.time.Duration; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.apex.malhar.lib.window.Accumulation; +import org.apache.apex.malhar.lib.window.ControlTuple; +import org.apache.apex.malhar.lib.window.TriggerOption; +import org.apache.apex.malhar.lib.window.Tuple; +import org.apache.apex.malhar.lib.window.Window; +import org.apache.apex.malhar.lib.window.WindowOption; +import org.apache.apex.malhar.lib.window.WindowState; +import org.apache.apex.malhar.lib.window.WindowedOperator; +import org.apache.apex.malhar.lib.window.WindowedStorage; + +import org.apache.hadoop.classification.InterfaceStability; + +import com.google.common.base.Function; + +import com.datatorrent.api.Context; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.common.util.BaseOperator; + +/** + * This is the abstract windowed operator class that implements most of the windowing, triggering, and accumulating + * concepts. The subclass of this abstract class is supposed to provide the implementation of how the accumulated + * values are stored in the storage. + * + * @param <InputT> The type of the input tuple + * @param <OutputT> The type of the output tuple + * @param <DataStorageT> The type of the data storage + * @param <AccumulationT> The type of the accumulation + */ +@InterfaceStability.Evolving +public abstract class AbstractWindowedOperator<InputT, OutputT, DataStorageT extends WindowedStorage, AccumulationT extends Accumulation> + extends BaseOperator implements WindowedOperator<InputT> +{ + protected WindowOption windowOption; + protected TriggerOption triggerOption; + protected long allowedLatenessMillis = -1; + protected WindowedStorage<WindowState> windowStateMap; + + private Function<InputT, Long> timestampExtractor; + + private long currentWatermark; + private boolean triggerAtWatermark; + private long earlyTriggerCount; + private long earlyTriggerMillis; + private long lateTriggerCount; + private long lateTriggerMillis; + private long currentDerivedTimestamp = -1; + private long windowWidthMillis; + protected DataStorageT dataStorage; + protected DataStorageT retractionStorage; + protected AccumulationT accumulation; + + private static final transient Logger LOG = LoggerFactory.getLogger(AbstractWindowedOperator.class); + + public final transient DefaultInputPort<Tuple<InputT>> input = new DefaultInputPort<Tuple<InputT>>() + { + @Override + public void process(Tuple<InputT> tuple) + { + processTuple(tuple); + } + }; + + // TODO: This port should be removed when Apex Core has native support for custom control tuples + @InputPortFieldAnnotation(optional = true) + public final	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2085
20611	APEXMALHAR-2085	Github user bhupeshchawda commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/319#discussion_r69312802 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/window/WindowedStorage.java --- @@ -0,0 +1,88 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window; + +import java.util.Map; + +import org.apache.hadoop.classification.InterfaceStability; + +/** + * WindowedStorage is a key-value store with the key being the window. The implementation of this interface should + * make sure checkpointing and recovery will be done correctly. + * + * @param <T> The type of the data that is stored per window + * + * TODO: Look at the possibility of integrating spillable data structure: https://issues.apache.org/jira/browse/APEXMALHAR-2026 --- End diff -- This seems very similar to the managed state implementations. Can we also look at managed state directly before looking at Spillable data structures?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2085

#	Requirement	Questions	Answered	Successfully Answered	URL
20612	APEXMALHAR-2085	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/319#discussion_r69608324 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/window/Accumulation.java --- @@ -0,0 +1,76 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package org.apache.apex.malhar.lib.window; + +import org.apache.hadoop.classification.InterfaceStability; + +/** + * This interface is for the processing part of the WindowedOperator. + * We can assume that all stateful processing of the WindowedOperator is a form of accumulation. + * + * In most cases, AccumT is the same as OutputT. But in some cases, the accumulated type and the output type may be + * different. For example, if we are doing the AVERAGE of doubles, InputT will be double, and we need the SUM and the + * COUNT stored as type AccumT, and AccumT will be a pair of double and long, in which double is the sum of the inputs, + * and long is the number of inputs. OutputT will be double, because it represents the average of the inputs. + */ +@InterfaceStability.Evolving +public interface Accumulation<InputT, AccumT, OutputT> { + /** + * Returns the default accumulated value when nothing has been accumulated + * + * @return + */ + AccumT defaultAccumulatedValue(); + + /** + * Accumulates the input to the accumulated value + * + * @param accumulatedValue + * @param input + * @return + */ + AccumT accumulate(AccumT accumulatedValue, InputT input); --- End diff -- What is the reason to make the interface stateless? Should it be stateful? My concern is with the interface promoting boxing/unboxing pattern and leading to an excessive GC.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2085
20613	APEXMALHAR-2082	Github user chinmaykolhatkar commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/270#discussion_r64454657 --- Diff: library/src/main/java/com/datatorrent/lib/filter/FilterOperator.java --- @@ -0,0 +1,210 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.filter; + +import java.util.LinkedList; +import java.util.List; + +import org.slf4j.Logger; +import org.slf4j.LoggerFactory; + +import org.apache.hadoop.classification.InterfaceStability; + +import com.datatorrent.api.AutoMetric; +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.PortContext; +import com.datatorrent.api.DefaultInputPort; +import com.datatorrent.api.DefaultOutputPort; +import com.datatorrent.api.Operator; +import com.datatorrent.api.annotation.InputPortFieldAnnotation; +import com.datatorrent.common.util.BaseOperator; +import com.datatorrent.lib.expression.Expression; +import com.datatorrent.lib.util.PojoUtils; + +/** + * FilterOperator + * Filter Operator filter out tuples based on defined condition + * + * Parameters + * - condition: condition based on expression langugage + * + * Input Port takes POJOs as an input + * + * Output Ports + * - truePort emits POJOs meeting the given condition + * - falsePort emits POJOs not meeting the given condition + * - error port emits any error situation while evaluating expression + * + */ +@InterfaceStability.Evolving +public class FilterOperator extends BaseOperator implements Operator.ActivationListener<Context> { + private transient Class<?> inClazz = null; + private transient Expression<Boolean> expr = null; + private List<String> expressionFunctions = new LinkedList<>(); + + @AutoMetric + private long trueTuples; + + @AutoMetric + private long falseTuples; + + @AutoMetric + private long errorTuples; + + protected String condition; --- End diff -- Can this be made private? You already have setter and getter for this.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2082
20614	APEXMALHAR-2077	Can we not add a property to the base class to toggle this behavior ? That way, we avoid creating a whole new class whose only point is to add a property.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2077
20615	APEXMALHAR-2070	Github user tweise commented on a diff in the pull request: https://github.com/apache/incubator-apex-malhar/pull/262#discussion_r61682196 --- Diff: library/src/main/java/org/apache/apex/malhar/lib/state/spillable/SpillableComplexComponent.java --- @@ -105,6 +105,26 @@ Serde<V, byte[]> serdeValue); /** + * This is a method for creating a {@link SpillableByteMultiset}. This method + * auto-generates an identifier for the data structure. + * @param <T> The type of the elements. + * @param bucket The bucket that this {@link SpillableByteMultiset} will be spilled too. + * @param serde The Serializer/Deserializer to use for data stored in the {@link SpillableByteMultiset}. + * @return A {@link SpillableByteMultiset}. + */ + <T> SpillableByteMultiset<T> newSpillableMultiset(long bucket, Serde<T, byte[]> serde); + + /** + * This is a method for creating a {@link SpillableByteMultiset}. + * @param <T> The type of the elements. + * @param identifier The identifier for this {@link SpillableByteMultiset}. + * @param bucket The bucket that this {@link SpillableByteMultiset} will be spilled too. + * @param serde The Serializer/Deserializer to use for data stored in the {@link SpillableByteMultiset}. + * @return A {@link SpillableByteMultiset}. + */ + <T> SpillableByteMultiset<T> newSpillableMultiset(byte[] identifier, long bucket, Serde<T, byte[]> serde); --- End diff -- How will this work with ManagedState? Will the data be distributed over multiple data files, partitioned by time?	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2070
20616	APEXMALHAR-2069	These are the default values which are assigned in the constructor. It is not evident to me that why these should be done during operator setup(). Can you please explain?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2069

#	Requirement	Questions	Answered	Successfully Answered	URL
20617	APEXMALHAR-2069	Github user vrozov commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/329#discussion_r69063970 --- Diff: library/src/main/java/com/datatorrent/lib/io/fs/FileSplitterInput.java --- @@ -318,6 +317,9 @@ public TimeBasedDirectoryScanner() @Override public void setup(Context.OperatorContext context) { + if (scanService == null) { --- End diff -- why is it necessary to check for null? Under what condition it will be not null? As far as I know platform calls setup/terminate once per deployment. IMO, it is safer to create new execution service each time setup is called and possibly assert or throw RuntimeException if scanService is not null to avoid misusage in unit tests.	1	1	https://issues.apache.org/jira/browse/APEXMALHAR-2069
20618	APEXMALHAR-2069	Github user chandnisingh commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/333#discussion_r69828207 --- Diff: library/src/test/java/com/datatorrent/lib/io/fs/FileSplitterInputTest.java --- @@ -499,7 +499,7 @@ public void testRecoveryOfBlockMetadataIterator() throws InterruptedException testMeta.fileSplitterInput.getScanner().setScanIntervalMillis(500); - testMeta.fileSplitterInput.setup(testMeta.context); --- End diff -- The goal of writing tests is to test the operator not just setup of window data manager. For a trivial improvement, that is, to avoid instantiating objects in the logical instance, this change now doesn't execute setup of the operator which is a much bigger problem. In future if something is added to splitter setup, the test will never execute it and result in bugs. Also, there are multiple of operators in library which instantiate field that are transient when the operator is constructed. Are you going to change all of them?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2069
20619	APEXMALHAR-2069	Github user sanjaypujare commented on a diff in the pull request: https://github.com/apache/apex-malhar/pull/333#discussion_r69834329 --- Diff: library/src/test/java/com/datatorrent/lib/io/fs/FileSplitterInputTest.java --- @@ -499,7 +499,7 @@ public void testRecoveryOfBlockMetadataIterator() throws InterruptedException testMeta.fileSplitterInput.getScanner().setScanIntervalMillis(500); - testMeta.fileSplitterInput.setup(testMeta.context); --- End diff -- For the first point, we can (should) add another unit test just to test FileSplitterInput.setup() which in any case sounds like a better idea. The 2nd point is not very clear. Are you saying there are many other instances of non-transient fields which should be checked for null in setup() or are you saying there are many fields which should ideally be declared transient?	0	0	https://issues.apache.org/jira/browse/APEXMALHAR-2069
20620	APEXMALHAR-2066	library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -0,0 +1,592 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.db.jdbc; + +import java.io.IOException; +import java.sql.PreparedStatement; +import java.sql.ResultSet; +import java.sql.SQLException; +import java.util.ArrayList; +import java.util.Collection; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.concurrent.ArrayBlockingQueue; +import java.util.concurrent.atomic.AtomicReference; +import javax.validation.constraints.Min; +import org.slf4j.LoggerFactory; +import org.apache.apex.malhar.lib.wal.WindowDataManager; +import org.apache.commons.lang3.tuple.MutablePair; +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultPartition; +import com.datatorrent.api.Operator.ActivationListener; +import com.datatorrent.api.Operator.IdleTimeHandler; +import com.datatorrent.api.Partitioner; +import com.datatorrent.api.annotation.Stateless; +import com.datatorrent.lib.db.AbstractStoreInputOperator; +import com.datatorrent.lib.util.KeyValPair; +import com.datatorrent.lib.util.KryoCloneUtils; +import com.datatorrent.netlet.util.DTThrowable; + +/** + * Abstract operator for for consuming data using JDBC interface + * User needs + * User needs to provide tableName,dbConnection,setEmitColumnList,look-up key + * + * Optionally batchSize,pollInterval,Look-up key and a where clause can be given + * + * This operator uses static partitioning to arrive at range queries for exactly + * once reads + * Assumption is that there is an ordered column using which range queries can + * be formed + * If an emitColumnList is provided, please ensure that the keyColumn is the + * first column in the list + * Range queries are formed using the {@link JdbcMetaDataUtility}} Output - + * comma separated list of the emit columns eg columnA,columnB,columnC + * + * @displayName Jdbc Polling Input Operator + * @category Input + * @tags database, sql, jdbc, partitionable,exactlyOnce + */ +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> + implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +{ + /** + * poll interval in milliseconds + */ + private int pollInterval; + + @Min(1) + private int partitionCount = 1; + protected transient int operatorId; + protected transient boolean isReplayed; + protected transient boolean isPollable; + protected int batchSize; + protected int fetchSize; + /** + * Map of windowId to <lower bound,upper bound> of the range key + */ + protected transient MutablePair<String, String> currentWindowRecoveryState; + + /** + * size of the emit queue used to hold polled records before emit + */ + private int queueCapacity = 4 * 1024 * 1024; + private transient volatile boolean execute; + private transient AtomicReference<Throwable> cause; + protected transient int spinMillis; + private transient OperatorContext context = null; + protected String tableName; + protected String key; + protected transient long currentWindowId; + protected KeyValPair<String, String> rangeQueryPair; + protected String lower; +	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2066

#	Requirement	Questions	Answered	Successfully Answered	URL
20621	APEXMALHAR-2066	library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -0,0 +1,592 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.db.jdbc; + +import java.io.IOException; +import java.sql.PreparedStatement; +import java.sql.ResultSet; +import java.sql.SQLException; +import java.util.ArrayList; +import java.util.Collection; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.concurrent.ArrayBlockingQueue; +import java.util.concurrent.atomic.AtomicReference; +import javax.validation.constraints.Min; +import org.slf4j.LoggerFactory; +import org.apache.apex.malhar.lib.wal.WindowDataManager; +import org.apache.commons.lang3.tuple.MutablePair; +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultPartition; +import com.datatorrent.api.Operator.ActivationListener; +import com.datatorrent.api.Operator.IdleTimeHandler; +import com.datatorrent.api.Partitioner; +import com.datatorrent.api.annotation.Stateless; +import com.datatorrent.lib.db.AbstractStoreInputOperator; +import com.datatorrent.lib.util.KeyValPair; +import com.datatorrent.lib.util.KryoCloneUtils; +import com.datatorrent.netlet.util.DTThrowable; + +/** + * Abstract operator for for consuming data using JDBC interface + * User needs + * User needs to provide tableName,dbConnection,setEmitColumnList,look-up key + * + * Optionally batchSize,pollInterval,Look-up key and a where clause can be given + * + * This operator uses static partitioning to arrive at range queries for exactly + * once reads + * Assumption is that there is an ordered column using which range queries can + * be formed + * If an emitColumnList is provided, please ensure that the keyColumn is the + * first column in the list + * Range queries are formed using the {@link JdbcMetaDataUtility}} Output - + * comma separated list of the emit columns eg columnA,columnB,columnC + * + * @displayName Jdbc Polling Input Operator + * @category Input + * @tags database, sql, jdbc, partitionable,exactlyOnce + */ +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> + implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +{ + /** + * poll interval in milliseconds + */ + private int pollInterval; + + @Min(1) + private int partitionCount = 1; + protected transient int operatorId; + protected transient boolean isReplayed; + protected transient boolean isPollable; + protected int batchSize; + protected int fetchSize; + /** + * Map of windowId to <lower bound,upper bound> of the range key + */ + protected transient MutablePair<String, String> currentWindowRecoveryState; + /** + * size of the emit queue used to hold polled records before emit + */ + private int queueCapacity = 4 * 1024 * 1024; + private transient volatile boolean execute; + private transient AtomicReference<Throwable> cause; + protected transient int spinMillis; + private transient OperatorContext context = null; + protected String tableName; + protected String key; + protected transient long currentWindowId; --- End diff -- What happens when checkpoint happens between Application window?	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2066
20622	APEXMALHAR-2066	library/src/main/java/com/datatorrent/lib/db/jdbc/AbstractJdbcPollInputOperator.java --- @@ -0,0 +1,592 @@ +/** + * Licensed to the Apache Software Foundation (ASF) under one + * or more contributor license agreements. See the NOTICE file + * distributed with this work for additional information + * regarding copyright ownership. The ASF licenses this file + * to you under the Apache License, Version 2.0 (the + * "License"); you may not use this file except in compliance + * with the License. You may obtain a copy of the License at + * + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, + * software distributed under the License is distributed on an + * "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY + * KIND, either express or implied. See the License for the + * specific language governing permissions and limitations + * under the License. + */ +package com.datatorrent.lib.db.jdbc; + +import java.io.IOException; +import java.sql.PreparedStatement; +import java.sql.ResultSet; +import java.sql.SQLException; +import java.util.ArrayList; +import java.util.Collection; +import java.util.HashMap; +import java.util.List; +import java.util.Map; +import java.util.concurrent.ArrayBlockingQueue; +import java.util.concurrent.atomic.AtomicReference; +import javax.validation.constraints.Min; +import org.slf4j.LoggerFactory; +import org.apache.apex.malhar.lib.wal.WindowDataManager; +import org.apache.commons.lang3.tuple.MutablePair; +import com.datatorrent.api.Context; +import com.datatorrent.api.Context.OperatorContext; +import com.datatorrent.api.DefaultPartition; +import com.datatorrent.api.Operator.ActivationListener; +import com.datatorrent.api.Operator.IdleTimeHandler; +import com.datatorrent.api.Partitioner; +import com.datatorrent.api.annotation.Stateless; +import com.datatorrent.lib.db.AbstractStoreInputOperator; +import com.datatorrent.lib.util.KeyValPair; +import com.datatorrent.lib.util.KryoCloneUtils; +import com.datatorrent.netlet.util.DTThrowable; + +/** + * Abstract operator for for consuming data using JDBC interface + * User needs + * User needs to provide tableName,dbConnection,setEmitColumnList,look-up key + * + * Optionally batchSize,pollInterval,Look-up key and a where clause can be given + * + * This operator uses static partitioning to arrive at range queries for exactly + * once reads + * Assumption is that there is an ordered column using which range queries can + * be formed + * If an emitColumnList is provided, please ensure that the keyColumn is the + * first column in the list + * Range queries are formed using the {@link JdbcMetaDataUtility}} Output - + * comma separated list of the emit columns eg columnA,columnB,columnC + * + * @displayName Jdbc Polling Input Operator + * @category Input + * @tags database, sql, jdbc, partitionable,exactlyOnce + */ +public abstract class AbstractJdbcPollInputOperator<T> extends AbstractStoreInputOperator<T, JdbcStore> + implements ActivationListener<Context>, IdleTimeHandler, Partitioner<AbstractJdbcPollInputOperator<T>> +{ + /** + * poll interval in milliseconds + */ + private int pollInterval; + + @Min(1) + private int partitionCount = 1; + protected transient int operatorId; + protected transient boolean isReplayed; + protected transient boolean isPollable; + protected int batchSize; + protected int fetchSize; + /** + * Map of windowId to <lower bound,upper bound> of the range key + */ + protected transient MutablePair<String, String> currentWindowRecoveryState; + /** + * size of the emit queue used to hold polled records before emit + */ + private int queueCapacity = 4 * 1024 * 1024; + private transient volatile boolean execute; + private transient AtomicReference<Throwable> cause; + protected transient int spinMillis; + private transient OperatorContext context = null; + protected String tableName; + protected String key; + protected transient long currentWindowId; + protected KeyValPair<String, String> rangeQueryPair; + protected String lower; +	1	0	https://issues.apache.org/jira/browse/APEXMALHAR-2066